

DRILLING HYPERBOLIC GROUPS

DANIEL GROVES, PETER HAÏSSINSKY, JASON F. MANNING, DAMIAN OSAJDA,
ALESSANDRO SISTO, AND GENEVIEVE S. WALSH

ABSTRACT. Given a hyperbolic group G and a maximal infinite cyclic subgroup $\langle g \rangle$, we define a *drilling of G along g* , which is a relatively hyperbolic group pair (\widehat{G}, P) . This is inspired by the well-studied procedure of drilling a hyperbolic 3-manifold along an embedded geodesic. We prove that, under suitable conditions, a hyperbolic group with 2-sphere boundary admits a drilling where the resulting relatively hyperbolic group pair (\widehat{G}, P) has relatively hyperbolic boundary S^2 . This allows us to reduce the Cannon Conjecture (in the residually finite case) to a relative version.

CONTENTS

1. Introduction	2
2. Outline of proof of Theorem A	7
3. Background	8
4. Coarse Topology	11
5. Linear connectivity	16
6. Shells and tube complements	26
7. Unwrapping the complement of a tube	48
8. Unwrapping a family of axes	65
9. Linearly connected limits of spheres	73
10. Proof of Theorem A	75
11. The Cohen–Lyndon property	77
12. Notation index	79
References	80

Date: June 20, 2024.

DG was partially supported by a Simons Fellowship, #1034629 to Daniel Groves and the National Science Foundation, DMS-1904913 and DMS-2203343. PH is partially supported by ANR-22-CE40-0004 GoFR. JM was partially supported by the Simons foundation, grants number 942496 and 524176 and by the National Science Foundation under grant DMS-1462263. DO was supported by the Carlsberg Foundation, grant CF23-1226. GW was partially supported by NSF Grants DMS-2005353 and DMS-1709964.

1. INTRODUCTION

An important tool in the study of hyperbolic 3-manifolds is the operation of drilling out an embedded geodesic k . It follows from Thurston's geometrization theorem that if M is a closed hyperbolic 3-manifold and k is an embedded geodesic loop then the complement $M \setminus k$ admits a complete hyperbolic metric of finite volume (see [37, 30]). The operation of drilling geodesics has been used in the proofs of some of the most important theorems about hyperbolic 3-manifolds. Some examples are Canary's proof that topologically tame hyperbolic 3-manifolds are geometrically tame [14], the proof of the Bers Density Theorem by Brock–Bromberg in [11], and Calegari–Gabai's proof of Tameness using Shrinkwrapping in [13], among many others. The purpose of this paper is to introduce a notion of 'drilling' in the group-theoretic setting. Here is what we mean by this.

Definition 1.1. Suppose that G is a hyperbolic group, and that g is an element so that $\langle g \rangle$ is a maximal cyclic subgroup. A *drilling of G along g* is a relatively hyperbolic pair (\widehat{G}, P) along with a normal subgroup $N \trianglelefteq P$ with an identification $P/N \cong \langle g \rangle$ so that $\widehat{G}/\langle\langle N \rangle\rangle \cong G$, with the quotient map inducing the identification of P/N with $\langle g \rangle$.

Here are some examples of (group-theoretic) drilling.

Example 1.2. Let M be a closed hyperbolic 3-manifold, and let k be an embedded geodesic loop. Let $M^\circ = M \setminus k$. If $G = \pi_1(M)$, g is the element represented by k , $\widehat{G} = \pi_1(M^\circ)$, and P is the subgroup of $\pi_1(M^\circ)$ corresponding to a regular neighborhood of k , then (\widehat{G}, P) is a drilling of G along g .

Example 1.3. Let H be a hyperbolic group and suppose that $h \in H$ generates an infinite maximal elementary subgroup. Let $T = \langle a, b \mid [a, b] = 1 \rangle$. Define $\widehat{H} = H *_{h=a} T$, and note that (\widehat{H}, T) is a drilling of H along h . We refer to this drilling as the *trivial drilling*.

Example 1.4. Let F be the fundamental group of a 3-dimensional handlebody, and $f \in F$ a diskbusting curve. We can realize f as an embedded geodesic in some negatively curved metric on the handlebody, and remove a neighborhood of this curve to obtain a drilling of a free group. It is possible to choose such a drilling to be nontrivial.

Similarly, one can perform a nontrivial drilling of a surface group $G = \pi_1(S)$ by realizing G as the fundamental group of some quasi-fuchsian $S \times I$, and drilling a simple geodesic γ out of $S \times I$. This can give a trivial drilling, but the drilling is nontrivial if γ is non-simple when realized as a geodesic in a hyperbolic metric on S itself.

Example 1.5. *Let (Γ, P) be a relatively hyperbolic group, and suppose that $N \trianglelefteq P$ is such that $P/N \cong \mathbb{Z}$. It follows from the relatively hyperbolic Dehn filling theorem [36] (see also [23]) that $\bar{\Gamma} = \Gamma / \langle\langle N \rangle\rangle$ is often hyperbolic, and that the image of P in $\bar{\Gamma}$ is maximal elementary. Thus (Γ, P) is a drilling of $\bar{\Gamma}$.*

It is notable that the remarkable 5–dimensional hyperbolic groups constructed by Italiano, Martelli, and Migliorini in [28] fall into the class described in Example 1.5, and so admit (nontrivial) drillings. Moreover, there are finite-volume hyperbolic manifolds of all dimensions at least 3 which have quotients as in Example 1.5, so it follows from Examples 1.4 and 1.5 that there are hyperbolic groups of all positive dimensions admitting nontrivial drillings.

Remark 1.6. Our condition that $P/N \cong \langle g \rangle$ is designed to mimic what happens on the level of the fundamental group when a geodesic is drilled. There are situations where it is interesting to remove other geometrically controlled sets (for example removing totally geodesic codimension–2 submanifolds from negatively curved manifolds of high dimension as in [2, 1]). Many of our basic tools work in such a setting, and we phrase our results to apply in these settings as much as possible. However, we do not pursue these ideas in this paper.

We are interested in drillings which look more like Example 1.2, in the following sense.

Let M be a closed hyperbolic 3–manifold and k an embedded geodesic loop. The fundamental group $\pi_1(M)$ is a (Gromov) hyperbolic group since it acts geometrically on \mathbb{H}^3 . The Gromov boundary of $\pi_1(M)$ is S^2 . The pair $(\pi_1(M \setminus k), \mathbb{Z}^2)$ is a relatively hyperbolic pair, and the relatively hyperbolic boundary $\partial(\pi_1(M \setminus k), \mathbb{Z}^2)$ is S^2 , since $\pi_1(M \setminus k)$ acts geometrically finitely on \mathbb{H}^3 . The parabolic subgroups in this action are the conjugates of $\mathbb{Z}^2 = \pi_1(T^2)$, where T^2 is the peripheral torus. Let $m \in \pi_1(M \setminus k)$ represent a meridian, i.e. a curve bounding a disk in M which is punctured exactly once by k . The quotient of $\pi_1(M \setminus k)$ by the normal closure of m is $\pi_1(M)$.

The main result of this paper is that given a hyperbolic group with boundary S^2 , there are often drillings in the sense of Definition 1.1 that are consistent with the above description of drilling a closed hyperbolic 3-manifold along an embedded geodesic. Precisely, here is our main result.

Theorem A. *Let X be a hyperbolic graph with boundary S^2 . For every hyperbolic isometry g of X , with $(g$ -invariant, quasi-geodesic) axis ℓ , there exists $\Sigma > 0$ so that the following holds:*

Suppose G is a subgroup of $\text{Isom}(X)$ so that

- (1) *G acts freely and cocompactly on X ; and*

(2) for any $h \in G \setminus \langle g \rangle$ the axes ℓ and $h \cdot \ell$ are at least Σ apart.

Then there exists a drilling (\widehat{G}, P) of G along g which satisfies the following properties:

- (1) (\widehat{G}, P) is relatively hyperbolic with P either free abelian of rank 2 or the fundamental group of a Klein bottle;
- (2) The Bowditch boundary of (\widehat{G}, P) is a 2-sphere.

The peripheral subgroup P is free abelian if and only if g preserves an orientation on ∂X .

Finally, \widehat{G} is torsion-free if and only if G is torsion-free.

Remark 1.7. In Theorem A we give a sufficient *geometric* condition for drilling. In [12, Section 5, p.79], Calegari and Fujiwara suggest that the condition of g having very small *stable commutator length* might be sufficient to perform some kind of drilling. We do not know if such a condition on the stable commutator length is sufficient, but believe it is an interesting question.

Our approach to proving Theorem A is to do “medium-scale geometric topology”. Our spaces are graphs, which of course do not have much useful local structure. However, using the Gromov hyperbolic geometry of X , we identify certain large (but fixed) scales where a neighborhood of the axis ℓ “looks like” a 3-dimensional solid tube $\mathbb{R} \times D^2$, and the “frontier” of this neighborhood “looks like” an open annulus $\mathbb{R} \times S^1$. Much of Section 6 is devoted to making this idea precise in a useful way.

Not only is Theorem A consistent with the relationship between the group-theoretical boundaries in the 3-manifold case, but we can use this to relate the understanding of hyperbolic groups with Gromov boundary S^2 to the understanding of relatively hyperbolic group pairs with Bowditch boundary S^2 . In that regard, we recall the statement of the Cannon Conjecture, one of the main open questions about hyperbolic groups, as well as its relatively hyperbolic counterpart.

Conjecture 1.8 (Cannon Conjecture). *Let G be a word hyperbolic group whose Gromov boundary is homeomorphic to S^2 . Then G is virtually Kleinian.*

Conjecture 1.9 (Toral Relative Cannon Conjecture). *Let (G, P) be a relatively hyperbolic group pair whose Bowditch boundary is homeomorphic to S^2 , $P \neq \emptyset$, and where every element of P is \mathbb{Z}^2 . Then G is Kleinian.*

In Conjecture 1.8, if Kleinian then G is virtually the fundamental group of a closed hyperbolic 3-manifold, whereas in Conjecture 1.9 G will be virtually the fundamental group of a finite-volume, non-compact hyperbolic 3-manifold.

The Cannon Conjecture was introduced by Cannon in [15], see also [34, Conjecture 6.2] for a broader conjecture. In the torsion-free case, the Cannon Conjecture is a special case of a question of Wall [40, Question G2, p.391]. In dimension 3, this question asks whether every finitely generated $PD(3)$ -group is the fundamental group of a closed aspherical 3-manifold. These conjectures are connected by work of Bestvina [3, Theorem 2.8, Remark 2.9], who proved (building on work of Bestvina–Mess [4]) that a torsion-free hyperbolic group is a $PD(3)$ group exactly when it has Gromov boundary homeomorphic to S^2 . Theorem A has the following consequence for $PD(3)$ groups and pairs.

Corollary B. *Suppose that G, X, g , and ℓ satisfy the hypotheses of Theorem A, and furthermore suppose that G is a $PD(3)$ -group (i.e. suppose that G is torsion-free). Let (\widehat{G}, P) be the drilling of G along g from the conclusion of Theorem A. Then (\widehat{G}, P) is a $PD(3)$ -pair.*

Proof. A $PD(3)$ -group has cohomological dimension 3, so is torsion-free. By the final conclusion of Theorem A, \widehat{G} is also torsion-free. That (\widehat{G}, P) is a $PD(3)$ -pair follows from [39, Corollary 1.3] (see also [33, Corollary 4.3]). \square

In [25, Corollary 1.4], the odd-numbered authors of this paper show that the Cannon Conjecture implies the Toral Relative Cannon Conjecture, using (group theoretic) Dehn Filling. In the setting of the Toral Relative Cannon Conjecture, the group pair is “Haken” in the sense of [29, p.282]. Many important breakthroughs, such as geometrization, and the cubulation of hyperbolic manifolds, were first solved in the Haken setting. Thus, it is reasonable to believe that Conjecture 1.9 may in fact be easier to prove than Conjecture 1.8. Theorem A implies the following converse in the residually finite case.

Theorem C. *Suppose that the Toral Relative Cannon Conjecture is true, and suppose that G is a residually finite hyperbolic group for which $\partial G \cong S^2$. Then G is virtually Kleinian.*

Proof. Let X be the Cayley graph of G with respect to some generating set. Since G is residually finite, there is a subgroup G_0 of finite index which is torsion-free and preserves an orientation on ∂G .

Let $g \in G_0$ be so that $\langle g \rangle$ is maximal cyclic. Since maximal abelian subgroups of residually finite groups are separable (see, for example, [31, Proposition, p.484]), there is a finite index subgroup G_1 of G_0 so that $g \in G_1$ and so that the axis of g in X satisfies the hypotheses of Theorem A with respect to the G_1 -action on X .

Therefore, there exists a drilling (\widehat{G}_1, P) of G_1 along g so that (\widehat{G}_1, P) is relatively hyperbolic with Bowditch boundary S^2 , P free abelian of rank 2, and an element $k \in P$ so that

$$G_1 \cong \widehat{G}_1 / \langle\langle k \rangle\rangle,$$

which identifies $\langle g \rangle$ with $P / \langle k \rangle$.

Since G_1 is torsion-free, \widehat{G}_1 is also torsion-free, by Theorem A. According to the Toral Relative Cannon Conjecture, \widehat{G}_1 is Kleinian. Suppose that M is a 1-cusped finite-volume hyperbolic 3-manifold such that $\pi_1(M) = \widehat{G}_1$. The group G_1 can be obtained as the fundamental group of a Dehn filling of M which implies that G_1 is the fundamental group of a closed 3-manifold. Since G_1 is word-hyperbolic, and $\partial G_1 \cong S^2$, it follows from the Geometrization Theorem that G_1 is Kleinian. This completes the proof. \square

The first three paragraphs of the proof of Theorem C prove the following unconditional statement.

Corollary D. *If G is a residually finite hyperbolic group with $\partial G \cong S^2$, then G virtually admits a drilling as in the statement of Theorem A.*

We expect many of the techniques used in this paper to be applicable to other situations. As an example, we highlight Theorem 6.35, which shows that, given an appropriate hyperbolic space X and quasi-convex subset Y , the coarse fundamental group of the ‘‘completed shell’’ around Y in X is isomorphic to $\pi_1(\partial X \setminus \Lambda Y)$.

Section 2 gives a general outline of the proof of Theorem A, while general background on relatively hyperbolic groups and their boundaries is contained in Section 3. In Section 4, we review and develop some notions of coarse topology, notably the coarse fundamental group, coarse Cartan-Hadamard, and coarse deformation retractions. Section 5 defines a condition we call ‘spherical connectivity’, which gives a local (in terms of bounded sized balls in a hyperbolic space) criterion for verifying that a boundary is linearly connected. We also relate linear connectivity of the boundary for visual metrics from varying basepoints to cut points on the boundary. In Section 6, we study the geometry of large tubes around geodesics, and maps from the boundary of a hyperbolic space to the ‘‘boundary’’ of a neighborhood of a convex subset. We show how to unwrap the complement of a tube about a geodesic (take a coarse infinite cyclic cover) in Section 7, and develop tools for studying the resulting space, and also the ‘‘unwrapped’’ space with a horoball attached to the pre-image of the boundary of the tube. In Section 8, we show how to unwrap a whole family of these tubes. In Section 9, we give a condition for a weak Gromov-Hausdorff

limit to be S^2 , which is used in Section 10 to prove that the boundary of the completely unwrapped space (the limit of the partially unwrapped spaces) is S^2 . In Section 10 we also prove the statement about torsion in Theorem A completing the proof. In Section 11 we prove that the quotient map from the drilled group to the original group is a “long filling” in a reasonable sense. In Section 12, we give a guide to some of our notation and constructions.

Acknowledgments. This project was started at the American Institute for Mathematics (AIM) workshop “Boundaries of groups”, and the bulk of the work was also carried out at AIM in the context of the SQuaRE program. We thank AIM for the very generous support. We thank C. G. Hruska and L. Paoluzzi in particular for early conversations about this work.

2. OUTLINE OF PROOF OF THEOREM A

Our proof of Theorem A is constructive. We start with a hyperbolic graph X_0 on which G acts geometrically and such that ∂X_0 is S^2 . We are given a g -invariant axis γ_0 whose G -translates are far apart. We want to build a space \widehat{Y} which is hyperbolic, $\partial \widehat{Y} = S^2$, and such that the drilled group (\widehat{G}, P) acts on \widehat{Y} , and this action is cusp-uniform with P as the stabilizer of a horoball. This implies that (\widehat{G}, P) is relatively hyperbolic, and that the Bowditch boundary is S^2 . Roughly speaking, we will build \widehat{Y} as a limit of (uniformly) hyperbolic spaces, and the boundary of \widehat{Y} will be a limit, in a suitable sense, of the boundaries of the approximating spaces, which will all be homeomorphic to S^2 .

Classically speaking “drilling” involves removing a regular neighborhood of a loop in a three-manifold. Since we are starting with a graph, we must take a more coarse approach. For example, for some scale D one can define the D -fundamental group of a graph to be the fundamental group of the space obtained by gluing on a disk to every loop of length at most D , see Section 4. A D -cover of a space is then a cover to which every loop of length at most D lifts. The precise scale D at which we work is fixed in Subsection 6.5.

We use this coarse theory to understand the connection between a shell around our quasi-geodesic axis γ_0 and the boundary $S^2 \setminus \Lambda \gamma_0$ (where $\Lambda \gamma_0$ denotes the limit set of γ_0 in $\partial X_0 = S^2$). For a large number K , let S denote the set of points at distance K from γ_0 . We prove that S is coarsely connected (Lemma 6.9), and that after S is “completed” to be a connected graph, $\pi_1^D(S)$ is isomorphic to \mathbb{Z} .

More precisely, there is a projection from $\partial X \setminus \Lambda \gamma_0$ to S which induces an isomorphism between $\pi_1(\partial X \setminus \Lambda \gamma_0)$ and $\pi_1^D(S)$. See Theorem 6.35 and Corollary 6.36 for precise statements.

Moreover, if T is the collection of points in X at distance at most K from γ_0 then $X \setminus T$ coarsely deformation retracts (see Definition 4.12) onto S (see Lemma 6.47). This implies their coarse fundamental groups are canonically isomorphic (see Lemma 4.13).

In Section 7, we fix a number R_0 large enough that the preceding results hold for $K = R_0$, and additionally so that every loop representing a nontrivial element of $\pi_1^D(S)$ is very long.

The D -universal cover \tilde{S} of S is a coarse plane, upon which a group E (either \mathbb{Z}^2 or the fundamental group of a Klein bottle) acts geometrically, see Lemma 7.21. The D -universal cover of $X \setminus T$ contains a copy of \tilde{S} . The space obtained by gluing a combinatorial horoball onto the copy of \tilde{S} is Gromov hyperbolic, and has boundary a 2-sphere (see Proposition 7.13 for hyperbolicity and Corollary 7.24 for the fact that the boundary is a 2-sphere).

We then proceed in Section 8 to inductively build spaces by unwrapping more and more lifts of complements of tubes around translates of γ_0 and gluing horoballs onto the resulting coarse planes. These spaces are all Gromov hyperbolic for a uniform hyperbolicity constant (this uses Lemma 7.12), and the visual metrics on their boundaries are uniformly linearly connected (this uses Corollary 7.30.) In Proposition 10.1 we show all these boundaries are homeomorphic to S^2 .

We thus obtain a sequence of based finite-valence graphs, and for any given N the N -ball about the basepoint stabilizes in this sequence. The limiting graph \hat{Y} is also Gromov hyperbolic (see Corollary 8.5), and admits a geometrically finite action of a group \hat{G} , which is hyperbolic relative to (a natural copy of) the group E , and is the required drilling of G along $\langle g \rangle$. This is shown in Theorem 8.9.

Realizing the boundary of \hat{Y} as a weak Gromov–Hausdorff limit (Definition 3.11) of the boundary two-spheres of the partially unwrapped and glued spaces, we use the characterization of the two-sphere from Theorem 9.1 to prove that $\partial\hat{Y}$ is a two-sphere. Finally, we prove the statement about torsion in G and \hat{G} in Proposition 10.2, thus completing the proof of Theorem A (see Section 10).

3. BACKGROUND

For background on hyperbolic spaces and groups and their Gromov boundaries we refer to [10, III.H].

3.1. Relative hyperbolicity. There are a number of equivalent definitions of relative hyperbolicity in the literature; see [27] for a survey. It is most

convenient for us to use Gromov’s original definition [22, 8.6]. In this definition (G, \mathcal{P}) is relatively hyperbolic if G has a *cuspidal uniform* action on a hyperbolic space with \mathcal{P} a choice of conjugacy classes of horoball stabilizers. We spell out what this means in the next few paragraphs.

Definition 3.1 (Busemann function, horofunction). [21, 8.1] Let Υ be a Gromov hyperbolic space, and let $p \in \partial\Upsilon$. For $x \in \Upsilon$ and $\gamma : [0, \infty) \rightarrow \Upsilon$ a geodesic ray tending to p , define

$$h_\gamma(x) = \limsup_{t \rightarrow \infty} (d_\Upsilon(x, \gamma(t)) - t).$$

The function $h_\gamma : \Upsilon \rightarrow \mathbb{R}$ is a *horofunction* based at p .

We obtain a function $\beta_p : \Upsilon \times \Upsilon \rightarrow \mathbb{R}$ by

$$\beta_p(x, y) = \sup \{h_\gamma(x) \mid \gamma(0) = y\}$$

This is called a *Busemann function based at p* .

Roughly speaking $\beta_p(x, y)$ measures how much “farther” x is from p as compared to y . Note that in some sources “Busemann function” refers to what we are here calling a horofunction (see for example [10, III.H.3.4]).

Definition 3.2 (Horoball). Let Υ be a Gromov hyperbolic space, and let $p \in \partial\Upsilon$. A *horoball centered at p* is any set \mathcal{H} , so that there is a horofunction h based at p and real numbers $a \leq b$, so that

$$h^{-1}((-\infty, a]) \subset \mathcal{H} \subset h^{-1}((-\infty, b]).$$

Definition 3.3. Let H be a group, and \mathcal{P} a finite family of subgroups of H . Suppose that H acts properly on a hyperbolic space Υ , preserving a family of disjoint open horoballs, and acting cocompactly on the complement of the union of that family. Suppose further that there is a collection of orbit representatives of horoballs in the family so that \mathcal{P} is the collection of their stabilizers. Then the pair (H, \mathcal{P}) is said to be *relatively hyperbolic*. If $\mathcal{P} = \{P\}$ consists of a single subgroup, we say that (H, P) is relatively hyperbolic.

There may be many spaces Υ which certify the relative hyperbolicity of a pair (H, \mathcal{P}) , but they all have H -equivariantly homeomorphic boundary at infinity, by a result of Bowditch [8, Theorem 9.4]. Accordingly, $\partial\Upsilon$ is referred to as the *Bowditch boundary* of the pair (H, \mathcal{P}) .

3.2. Visibility. Recall the following definition and result from [25].

Definition 3.4. [25, Definition 3.4] Let $\mu > 0$. A geodesic metric space Z is μ -*visible* if, for every $a, b \in Z$, there is a geodesic ray based at a passing within μ of b .

Proposition 3.5. [25, Proposition 6.3] *For every $v \geq 1$ the following holds. Let Υ be a proper v -hyperbolic space and suppose that for all $p, q \in \Upsilon$ with $d(p, q) \leq 100v$ there exists a geodesic $[p, q']$ of length at least $200v$ with $d(q, [p, q']) \leq v$. Then Υ is $5v$ -visible.*

Remark 3.6. Some sources say “almost geodesically complete” or “almost extendible” instead of “visible” (see for example [35]). The fact that hyperbolic groups are visible is due to Mihalik. See [4, Lemma 3.1] for a proof, which extends easily to any proper hyperbolic space Υ whose isometry group is co-compact. (In fact it is enough that the action on the space of pairs of distinct points in $\partial\Upsilon$ is co-compact.)

3.3. Visual metrics. Here we recall some notions and facts concerning metrics on the Gromov boundary. For details see e.g. [21, Sec. 7.3] or [10, Chap. III.H].

Definition 3.7. Let Υ be a Gromov hyperbolic space. Suppose that $x, y, w \in \Upsilon$. The *Gromov product* of x and y with respect to w is

$$(x | y)_w = \frac{1}{2} (d_{\Upsilon}(x, w) + d_{\Upsilon}(y, w) - d_{\Upsilon}(x, y)).$$

Suppose that α, β are geodesic rays with common basepoint w . The *Gromov product* of α and β with respect to w is

$$(\alpha | \beta)_w = \liminf_{s, t \rightarrow \infty} (\alpha(s) | \beta(t))_w.$$

Definition 3.8. Let Υ be Gromov hyperbolic space, with basepoint w . A *visual metric on $\partial\Upsilon$, based at w , with parameters ε, κ* is a metric $\rho(\cdot, \cdot)$ which is κ -bi-Lipschitz to $e^{-\varepsilon(\cdot)_w}$.

Proposition 3.9. [10, III.H.3.21] *Let $\delta > 0$. For any positive $\varepsilon \leq \frac{1}{6\delta}$, and $\kappa(\varepsilon, \delta) = \left(3 - 2e^{2\delta\varepsilon}\right)^{-1}$ we have the following:*

If Υ is a δ -hyperbolic space and $w \in \Upsilon$, then $\partial\Upsilon$ has a visual metric based at w with parameters ε, κ .

Definition 3.10. A δ -*adapted visual metric at w* is a visual metric as above where $\varepsilon(\delta) = \frac{1}{6\delta}$ and $\kappa = \left(3 - 2e^{1/3}\right)^{-1} = 4.78984\dots$ are determined as in Proposition 3.9. If the basepoint is understood or unimportant, we just say that the metric is δ -*adapted*.

3.4. Strong convergence and weak Gromov-Hausdorff convergence. The use of the following two types of convergence for metric spaces is important for this paper. For this we mostly follow [25].

Definition 3.11. Let $(Z_i)_{i \in \mathbb{N}}$ and Z be metric spaces. We say that Z is a *weak Gromov–Hausdorff limit* of the sequence (Z_i) if there exists $\lambda \geq 1$ and a sequence of (λ, ε_i) –quasi-isometries $Z \rightarrow Z_i$, with $\varepsilon_i \rightarrow 0$ as $i \rightarrow \infty$.

Definition 3.12. Let (Z, p) be a pointed metric space. We say that the sequence of pointed metric spaces (Z_i, p_i) *strongly converges* to (Z, p) if the following holds: For every $R > 0$, for all but finitely many i there are isometries $\phi_i: B_R(p) \rightarrow B_R(p_i)$ so that $\phi_i(p) = p_i$.

The connection between these two notions of convergence is given in the following.

Proposition 3.13. [25, Proposition 3.5] *Let $\delta > 0$. Let (Y_i, p_i) be a sequence of pointed δ –hyperbolic δ –visible metric spaces which strongly converge to the pointed δ –hyperbolic, δ –visible metric space (Y, p) . Let ρ be a δ –adapted visual metric at p on ∂Y and, for each i , let ρ_i be a δ –adapted visual metric at p_i on ∂Y_i . The space $(\partial Y, \rho)$ is a weak Gromov–Hausdorff limit of the spaces $(\partial Y_i, \rho_i)$.*

3.5. Guessing geodesics. In order to verify that a space $X_0^{\text{cusp}}(K)$ we build in Subsection 6.5 is hyperbolic, we need the following result of Bowditch. For a connected (simplicial) graph A , from now on we always consider a path metric d_A on its vertex set.

Proposition 3.14. [9, Proposition 3.1] *Given $h \geq 0$ there exists $k \geq 0$ with the following property. Suppose that Γ is a connected graph, and that for each $x, y \in V(\Gamma)$ there is an associated connected sub-graph $\mathbb{P}(x, y) \subseteq \Gamma$ so that*

(1) *For all $x, y, z \in V(\Gamma)$*

$$\mathbb{P}(x, y) \subseteq N_h(\mathbb{P}(x, z) \cup \mathbb{P}(z, y));$$

(2) *For any $x, y \in V(\Gamma)$ with $d_\Gamma(x, y) \leq 1$, the diameter of $\mathbb{P}(x, y)$ in Γ is at most h .*

Then Γ is k –hyperbolic. In fact, it suffices to take any $k \geq \frac{1}{2}(3m - 10h)$, where m is any positive real number so that

$$2h(6 + \log_2(m + 2)) \leq m.$$

Moreover, for all $x, y \in V(\Gamma)$, the Hausdorff distance between $\mathbb{P}(x, y)$ and any geodesic from x to y is at most $m - 4h$.

4. COARSE TOPOLOGY

In this section we introduce some tools used in Sections 6 and 7. We are working with graphs, so we assume that $D, Q \in \mathbb{N}$.

Definition 4.1. (Coarse fundamental group) Let Γ be a connected graph and let $D \geq 0$. Let Γ^D be the space obtained from Γ by gluing disks to all edge-loops of length $\leq D$.

Set

$$\pi_1^D(\Gamma, a_0) = \pi_1(\Gamma^D, a_0).$$

The D -universal cover of Γ is the natural preimage of Γ in the universal cover of Γ^D . More generally, a D -covering map of Γ is the restriction to the preimage of Γ of a covering of Γ^D .

Observe that a D -covering map of Γ is a covering for which all the loops of length at most D in Γ lift. The following lemma is a direct consequence of the definition.

Lemma 4.2. *Suppose that Γ is a connected graph, let $D > 0$, let $\tilde{\Gamma}$ be the D -universal cover of Γ , and let $p \in \tilde{\Gamma}$. Then $\pi_1^D(\tilde{\Gamma}, p) = \{1\}$.*

The coarse Cartan–Hadamard Theorem below uses the language of the following Definition 4.3, but our focus is on the length rather than the diameter of loops. We relate these two notions in Lemma 4.5 below.

Definition 4.3. [17, Appendix A] Let $D > 0$. A connected graph Γ is said to be D -simply connected if $\pi_1(\Gamma)$ is generated by loops freely homotopic to loops of diameter at most D . If Γ is D -simply-connected for some $D > 0$ we say that Γ is *coarsely simply-connected*.

Basic examples of coarsely simply-connected spaces are hyperbolic spaces. The following is an immediate consequence of [10, Lemma III.H.2.6].

Lemma 4.4. *Let Y be a δ -hyperbolic graph and, $D \geq 16\delta$. For all $y \in Y$ we have $\pi_1^D(Y, y) = \{1\}$.*

Lemma 4.5. *Let Γ be a graph, and $D > 0$. If Γ^{2D} is simply connected, then Γ is D -simply connected. Conversely, if Γ is D -simply connected, then Γ^{2D+1} is simply connected.*

Proof. Suppose Γ^{2D} is simply connected, and let α be an edge-loop in Γ . There exists a (singular) disk diagram $\Delta \rightarrow \Gamma^{2D}$ for α , that is, $\partial\Delta$ is mapped to α and 2-cells in Δ are mapped to 2-cells in Γ^{2D} . The boundary of each such 2-cell has diameter at most D , hence the class of α in $\pi_1(\Gamma)$ is generated by loops of such diameter, proving the first implication.

Now suppose that Γ is D -simply connected. Consider an edge-loop α in Γ^{2D+1} . Such a loop is in fact a loop in Γ . The class of α in $\pi_1(\Gamma)$ is generated by loops of diameter at most D hence there exists a map $\Delta^{(1)} \rightarrow \Gamma$ from the 1-skeleton of a combinatorial (singular) disk Δ such that $\partial\Delta$ is mapped onto α and the boundary of every 2-cell in Δ is mapped onto a loop of diameter at most D . Observe that any loop of such diameter can be

decomposed into finitely many loops of length at most $2D + 1$, hence the map $\Delta^{(1)} \rightarrow \Gamma$ induces a map $\Delta' \rightarrow \Gamma^{2D+1}$ from a subdivision Δ' of Δ , justifying triviality of the homotopy class of α in Γ^{2D+1} , and finishing the proof. \square

When calculating constants later, we implicitly use the following immediate consequence of Lemma 4.5. Note that constants are not optimal.

Corollary 4.6. *Let Γ be a connected graph, $a_0 \in \Gamma$, and $D > 0$. If $\pi_1^D(\Gamma, a_0) = \{1\}$ then Γ is D -simply-connected.*

As mentioned above, the reason we wish to control the scale of simple connectivity is so we can apply the following theorem.

Theorem 4.7 (Coarse Cartan–Hadamard). [17, Theorem A.1] *Let $v \geq 0$, and let $\sigma \geq 10^7 v$. Let Z be a geodesic space. If every ball of radius σ in Z is v -hyperbolic and if Z is $10^{-5}\sigma$ -simply-connected, then Z is $300v$ -hyperbolic.*

4.1. Combinatorial horoballs.

Definition 4.8. [23, Definition 3.1] Let Γ be a simplicial graph. The *combinatorial horoball based on Γ* , is the graph $\mathcal{H}(\Gamma)$ with vertex set $\Gamma \times \mathbb{Z}_{\geq 0}$ and two types of edges:

- vertical edges connecting (x, n) to $(x, n + 1)$
- horizontal edges connecting (x, n) to (y, n) if $d_\Gamma(x, y) \leq 2^n$.

Vertices (x, n) are at depth n .

The graph Γ naturally includes as a subgraph of $\mathcal{H}(\Gamma)$ but its metric is distorted exponentially. More precisely we have the following.

Lemma 4.9. *Let Γ be a connected simplicial graph, and let $\mathcal{H} = \mathcal{H}(\Gamma)$ be the horoball based on Γ . For any distinct vertices $v, w \in \Gamma$, we have*

$$(1) \quad \frac{1}{2}d_{\mathcal{H}}((v, 0), (w, 0)) - 2 < \log_2(d_\Gamma(v, w)) < \frac{1}{2}d_{\mathcal{H}}((v, 0), (w, 0)) + 1.$$

Proof. As pointed out in [23, Lemma 3.10], any two vertices of \mathcal{H} are joined by a *regular* geodesic, which consists of an initial vertical subsegment, followed by a horizontal subsegment of length at most 3, followed by a final vertical subsegment. If both the vertical subsegments are nondegenerate, then the horizontal subsegment has length at least 2. If $d_\Gamma(v, w) < 6$, then $d_\Gamma(v, w) = d_{\mathcal{H}}((v, 0), (w, 0))$ and the inequalities (1) are easily verified.

Otherwise, the vertical segments are nondegenerate. Let $h \in \{2, 3\}$ be the length of the horizontal subsegment and let D be the length of each vertical subsegment, so $d_{\mathcal{H}}((v, 0), (w, 0)) = 2D + h$. We have

$$2^D < d_\Gamma(v, w) \leq h \cdot 2^D.$$

Taking logs and using the relationship between D and $d_{\mathcal{H}}((v,0), (w,0))$ yields the inequalities (1). \square

Lemma 4.10. *Let Γ be a connected simplicial graph, and let $\mathcal{H}(\Gamma)$ be the combinatorial horoball based on Γ . For any $D \geq 5$ and any $p \in \mathcal{H}(\Gamma)$ we have $\pi_1^D(\mathcal{H}(\Gamma), p) = \{1\}$.*

Proof. In [23, Definition 3.1] the combinatorial horoball is defined by including 2-cells which are triangles, squares, or pentagons. In [23, Proposition 3.7] it is proved that this complex is simply-connected. \square

4.2. Coarse topology of glued spaces.

Proposition 4.11. *Let Γ be a connected simplicial graph, and let $D \geq 5$. Suppose that Γ^D is simply connected. Let $\{\Gamma_i\}$ be disjoint connected sub-graphs of Γ . Let $\widehat{\Gamma}$ be obtained from Γ by gluing a combinatorial horoball \mathcal{H}_i based on Γ_i to each Γ_i . Then $\widehat{\Gamma}$ is connected and $\widehat{\Gamma}^D$ is simply connected.*

Proof. Since $\widehat{\Gamma}$ is obtained from gluing together connected graphs along connected sub-graphs, it is clear that $\widehat{\Gamma}$ is connected. Therefore, the basepoint is irrelevant, and we assume that the basepoint a is in Γ .

We first note since Γ_i is connected, any path p in \mathcal{H}_i with endpoints in Γ_i is D -homotopic rel endpoints to a path in Γ_i , by Lemma 4.10.

Consider now a loop ℓ based at $a \in \Gamma$. Up to homotopy in $\widehat{\Gamma}^D$, we can replace each maximal sub-path of ℓ contained in some \mathcal{H}_i with a path contained in Γ_i , thereby obtaining a loop ℓ' in Γ . Since $\pi_1^D(\Gamma, a) = \{1\}$, we have that ℓ' , hence ℓ , is D -homotopically trivial, as required. \square

4.3. Coarse deformation retractions.

Definition 4.12. Let \mathfrak{E} be a connected graph and $\Gamma \subset \mathfrak{E}$ a subset. A Q -deformation retraction of \mathfrak{E} onto Γ is a sequence of maps $f_0, \dots, f_n, \dots : \mathfrak{E} \rightarrow \mathfrak{E}$ satisfying the following properties:

- (1) f_0 is the identity,
- (2) $f_i|_{\Gamma}$ is the identity for all i ,
- (3) for all $b_1, b_2 \in \mathfrak{E}$ so that $d_{\mathfrak{E}}(b_1, b_2) \leq 1$ and all i we have

$$d_{\mathfrak{E}}(f_i(b_1), f_i(b_2)) \leq Q,$$

- (4) for all $b \in \mathfrak{E}$ and all i we have $d_{\mathfrak{E}}(f_i(b), f_{i+1}(b)) \leq 1$,
- (5) for each $b \in \mathfrak{E}$ there is some $n \geq 0$ and $a \in \Gamma$ so that $f_i(b) = a$ for all $i \geq n$.

Define $f(b)$ to be the point a as in item (5). We call f the *stable map*.

The Q -deformation retractions we require are built in Lemma 6.47 below. A natural example of a Q -deformation retraction is given by a projection of hyperbolic graphs onto a neighborhood of a quasi-convex sub-graph, see Lemma 6.45.

It is important to observe that in the following proposition the conditions required on the Q -deformation retraction are stronger than those implied by the definition. In particular, notice that in Item (1) the d_Γ -distance between $f(b_1)$ and $f(b_2)$ is bounded by Q .

Proposition 4.13. *Suppose that Ξ is a connected graph, and that $\Gamma \subset \Xi$ is a connected sub-graph. Let $a \in \Gamma$, and $Q > 0$. Suppose*

- (1) *there exists a Q -deformation retraction $f_0, \dots, f_n, \dots : \Xi \rightarrow \Xi$ of Ξ onto Γ with stable map f , so that for all $b_1, b_2 \in \Xi$ for which $d_\Xi(b_1, b_2) \leq 1$ we have $d_\Gamma(f(b_1), f(b_2)) \leq Q$; and*
- (2) *there exists $D > 2Q + 2$ so that all Γ -loops of length at most QD represent the (conjugacy class of the) identity element of $\pi_1^D(\Gamma, a)$.*

Then the inclusion $\iota : \Gamma \hookrightarrow \Xi$ induces an isomorphism:

$$\iota_* : \pi_1^D(\Gamma, a) \rightarrow \pi_1^D(\Xi, a).$$

Moreover, for any k , given any loop σ in Ξ of length k (possibly not based at a), with consecutive vertices $b_0, b_1, \dots, b_k = b_0$, let $f(\sigma)$ be a loop in Γ obtained by connecting $f(b_i)$ with $f(b_{i+1})$ by a path of length at most Q in Γ , for all i (this is possible by condition (2) above). Then $f(\sigma)$ represents the same conjugacy class as σ in $\pi_1^D(\Xi, a) \cong \pi_1^D(\Gamma, a)$.

Proof. Any inclusion $\iota : \Gamma \hookrightarrow \Xi$ between connected graphs induces a homomorphism $\iota_* : \pi_1^D(\Gamma, a) \rightarrow \pi_1^D(\Xi, a)$, for any $a \in \Gamma$.

Take a loop σ in Ξ , and construct a loop $f(\sigma)$ as in the last statement of the theorem. If σ was based at a , then so is $f(\sigma)$. Different choices of paths between b_i and b_{i+1} give loops differing by a collection of loops of length at most $2Q \leq QD$, so give the same element of $\pi_1^D(\Gamma, a)$. Moreover, a relation in $\pi_1^D(\Xi, a)$ is a loop of length at most D in Ξ , which is mapped to a loop of length at most QD in Γ , which implies (by Condition (2)) that there is a well-defined induced map

$$f_* : \pi_1^D(\Xi, a) \rightarrow \pi_1^D(\Gamma, a).$$

Since f is a Q -deformation retraction, we clearly have $f \circ \iota = \text{Id} |_\Gamma$, which means that $f_* \circ \iota_* = \text{Id} |_{\pi_1^D(\Gamma, a)}$, and thus ι_* is injective.

However, if σ is any Ξ -loop based at a , the fact that $D > 2Q + 2$ means that the maps (f_i) provide a homotopy in Ξ^D between σ and $f(\sigma)$, which means that ι_* is surjective.

The final statement of the result follows immediately from the above argument. \square

5. LINEAR CONNECTIVITY

In this section we investigate linear connectedness of boundaries. In particular, we link “uniform” linear connectedness to (weak) cut points (see Theorem 5.3 and also Proposition 5.6 and 5.7). We also define a condition on graphs called “spherical connectivity” which allows us to prove a local-to-global result about linearly connected boundaries (see Lemma 5.11). In later sections, we use this criterion to verify that the boundaries of our partially unwrapped-and-glued spaces are uniformly linearly connected, which is crucial in our calculation that the Bowditch boundary of the drilled group is a two-sphere (see Corollary 7.30).

Definition 5.1. A metric space M is L -linearly connected if every pair $\{x, y\}$ of points is contained in a connected set J of diameter at most $L \cdot d(x, y)$. A metric space is *linearly connected* if it is L -linearly connected for some L .

This condition is also called L - LLC_1 (see, for example [5, Section 2]), and has a companion property L - LLC_2 , which together make the property L -*locally linearly connected*. Connected boundaries of hyperbolic groups are locally linearly connected, but we only need the property from Definition 5.1 in this paper.

As noted by Mackay [32], when M is locally compact and locally connected, by increasing L by an arbitrarily small amount we may assume J is an arc.

We use the following notation several times in this section. If a, b are two positive quantities, we write $a \lesssim b$ or $b \gtrsim a$ if there is a constant $u = u(\delta) > 0$ such that $a \leq ub$. We write $a \asymp b$ if $a \lesssim b$ and $a \gtrsim b$. We say that $a \leq b + O(\delta)$ if there exists non-negative number $C(\delta)$ such that $a - b \leq C\delta$.

5.1. Linear connectivity versus cut points. A *continuum* is a compact, connected, Hausdorff topological space. A *cut point* of a continuum M is a point $x \in M$ such that $M \setminus x$ is not connected.

Definition 5.2. Suppose that M is a continuum. A point $x \in M$ is a *weak cut point* of M if there exist $p, q \in M \setminus \{x\}$ so that every sub-continuum of M that contains both p and q must also contain x .

An example of a weak cut point that is not a cut point is any point on the vertical line in a topologist’s sine curve.

The next result is one of the main results of this section. We use it in Lemma 7.28 for proving that our first unwrapped and glued space has boundary which is uniformly linearly connected.

Theorem 5.3. *Suppose Y_i are δ -hyperbolic and δ -visual metric spaces with finitely many isometry types of balls of radius R for every $R > 0$ (not*

depending on i). Endow each $M_i = \partial Y_i$ with a δ -adapted visual metric ρ_i with basepoint p_i , and suppose each M_i is connected. Further suppose that M_i is not L_i -linearly connected for $L_i \rightarrow \infty$.

Then there exists a sequence p'_i of basepoints such that a subsequence of (Y_i, p'_i) strongly converges to (Y_∞, p'_∞) with the property that $M_\infty = \partial Y_\infty$ has a weak cut point.

We need the following result, which is [25, Lemma 6.19].

Lemma 5.4. *Let M be a compact metric space. Suppose that there exists $L \geq 1$ so that each $p, q \in M$ can be joined by a chain of points $p = p_1, \dots, p_n = q$ so that $\text{diam}(\{p_1, \dots, p_n\}) \leq Ld(p, q)$ and $d(p_i, p_{i+1}) \leq d(p, q)/2$. Then M is $5L$ -linearly connected.*

We use the following lemma to find the new basepoints p'_i in the proof of Theorem 5.3:

Lemma 5.5. *For every $\delta > 0$, there exist constants $c, \lambda, t_0 > 1$ so that the following holds. Suppose that Y is δ -hyperbolic and δ -visual, and let $p \in Y$, and $x, y \in \partial Y$, with $x \neq y$. Let ρ be a δ -adapted visual metric on ∂Y at p . There exists $p' \in Y$ so that in any δ -adapted visual metric ρ' on ∂Y at p' we have:*

- (1) $\rho'(x, y) > 1/c$,
- (2) for every $t \geq t_0$, the set $\mathcal{A}_t = \{z \in \partial Y : \rho(z, x) \geq t\rho(x, y)\}$ has ρ' -diameter at most c/t and for any $z \in \mathcal{A}_t$, $\rho'(x, z), \rho'(y, z) \geq 1/c$,
- (3) on $\partial Y - \mathcal{A}_t$, $\frac{1}{\rho(x, y)}\rho$ and ρ' are λt^2 -bi-Lipschitz equivalent.

Proof. Let $a = e^{1/6\delta}$. By the definition of δ -adapted visual metric, we have $\rho(w, z) \asymp a^{-(w|z)_p}$ for all $w, z \in \partial Y$. We use the fact that if x, y are points on ∂Y , then as in [10, III.H.3.18.(3)],

$$(x|y)_p = d(p, (x, y)) + O(\delta),$$

where (x, y) is any bi-infinite geodesic from x to y .

By [21, Thm 2.12], there exists a constant $C = C(\delta) \geq 0$ such that, for any $\{p, x, y, z, w\} \in Y \cup \partial Y$, $p \in Y$, there is a $(1, C)$ -quasi-isometry

$$f: Y = [p, x) \cup [p, y) \cup [p, z) \cup [p, w) \rightarrow T$$

into a metric tree which preserves distances to p . Let us write $\bar{u} = f(u)$ for any $u \in Y$. For $u \in \partial Y = \{x, y, z, w\}$ we write \bar{u} for the corresponding point in ∂T . It follows that, for any $q \in Y$ and $u, v \in Y \cup \partial Y$, we have $(\bar{u}|\bar{v})_{\bar{q}} = d(\bar{q}, \bar{c}) = (u|v)_q + O(\delta)$ where \bar{c} is the center of the tripod $\{\bar{q}, \bar{u}, \bar{v}\}$.

Let us consider the point p' on a geodesic ray with origin p which is asymptotic to x such that $d(p, p') = (x|y)_p$. The approximating tree is a tripod

where \bar{p}' is $O(\delta)$ from the center. This implies that $\rho'(x, y)$ is comparable to 1, because $(x|y)_{\bar{p}'} = O(\delta)$. This proves (1), for $c \geq c_1(\delta)$

From now on, we assume that $\bar{p}' = f(p')$ is the center of the tripod $\{\bar{p}, \bar{x}, \bar{y}\}$.

We turn to the proof of (2). Let $t \geq 1$ and pick z, w such that $\rho(w, x), \rho(z, x) \geq t\rho(x, y)$, i.e., $\max\{(x|w)_p, (x|z)_p\} \leq (x|y)_p - \log_a t + O(\delta)$, see Figure 1(left). Thus,

$$(\bar{x}|\bar{z})_{\bar{p}} \leq (x|z)_p + O(\delta) \leq (x|y)_p - \log_a t + O(\delta) \leq d_T(\bar{p}, \bar{p}') - \log_a t + O(\delta)$$

and similarly

$$(\bar{y}|\bar{z})_{\bar{p}} \leq d_T(\bar{p}, \bar{p}') - \log_a t + O(\delta)$$

as

$$\rho(y, z) \geq (t-1)\rho(x, y) \text{ by the triangle inequality and } (y|z)_p \leq (x|y)_p - \log_a(t-1).$$

Thus, we may find some $t_0 > 1$ large enough such that, if $t \geq t_0$, then, in the approximate tree T , the point \bar{p}' belongs to $(\bar{z}, \bar{x}) \cap (\bar{z}, \bar{y})$.




FIGURE 1. The proof of Lemma 5.5: Conclusion (2) (left); conclusion (3), in case $(\bar{x}|\bar{w})_{\bar{p}} \geq d_T(\bar{p} - \bar{p}')$ (center); Conclusion (3), in case $(\bar{x}|\bar{w})_{\bar{p}} \leq d_T(\bar{p} - \bar{p}')$ (right).

We obtain, with the observation that $(\bar{x}|\bar{w})_{\bar{p}} = (\bar{p}'|\bar{w})_{\bar{p}}$ and $(\bar{x}|\bar{z})_{\bar{p}} = (\bar{p}'|\bar{z})_{\bar{p}}$,

$$\begin{cases} (\bar{x}|\bar{w})_{\bar{p}} + (\bar{p}|\bar{w})_{\bar{p}'} = (\bar{x}|\bar{y})_{\bar{p}} \\ (\bar{x}|\bar{z})_{\bar{p}} + (\bar{p}|\bar{z})_{\bar{p}'} = (\bar{x}|\bar{y})_{\bar{p}} \end{cases}$$

Therefore,

$$\begin{aligned} (\bar{w}|\bar{z})_{\bar{p}'} &\geq \min\{(\bar{p}|\bar{z})_{\bar{p}'}, (\bar{p}|\bar{w})_{\bar{p}'}\} \\ &\geq (\bar{x}|\bar{y})_{\bar{p}} - \max\{(\bar{x}|\bar{z})_{\bar{p}}, (\bar{x}|\bar{w})_{\bar{p}}\} \\ &\geq \log_a t + O(\delta). \end{aligned}$$

Hence $\rho'(z, w) \leq c_2/t$ for some $c_2 = c_2(\delta) > 0$. This implies that the diameter of \mathcal{A}_t is bounded by c_2/t . Note that \bar{p}' belongs to both geodesics (\bar{x}, \bar{z}) and (\bar{y}, \bar{z}) so that $\rho'(x, z), \rho'(y, z) \asymp 1$ thus $\rho'(x, z), \rho'(y, z) \gtrsim 1$, which ends the proof of (2), taking $c = \max\{c_1, c_2\}$.

Let us deal with the proof of (3). If $w, z \in \partial Y - \mathcal{A}_t$, then $\rho(w, x), \rho(z, x) \leq t\rho(x, y)$, i.e., $\min\{(x|w)_p, (x|z)_p\} \geq (x|y)_p - \log_a t + O(\delta)$. By the triangle inequality, we have

$$(2) \quad |(z|w)_{p'} - (z|w)_p| \leq d_Y(p - p') = (x|y)_p.$$

Therefore

$$\frac{\rho'(z, w)}{\rho(x, y)^2} \gtrsim \frac{\rho(z, w)}{\rho(x, y)} \gtrsim \rho'(z, w).$$

This is one of the coarse inequalities that we need. Now we find an upper bound independent of $\rho(x, y)$.

If $\log_a t \geq (x|y)_p$ then we have essentially no constraints: the inequality (2) implies that

$$\frac{\rho(z, w)}{\rho(x, y)} \lesssim \frac{\rho'(z, w)}{\rho(x, y)^2} \lesssim t^2 \rho'(z, w).$$

So we restrict our attention to the condition $\log_a t \leq (x|y)_p$. Let us work on the approximate tree T . Switching z and w if necessary, we may assume that $(\bar{x}|\bar{w})_{\bar{p}} \leq (\bar{x}|\bar{z})_{\bar{p}}$. Let us consider two cases. If $(\bar{x}|\bar{w})_{\bar{p}} \geq d_T(\bar{p} - \bar{p}')$ (see Figure 1(center)) then

$$(\bar{z}|\bar{w})_{\bar{p}} = d_T(\bar{p} - \bar{p}') + (\bar{z}|\bar{w})_{\bar{p}'} = (\bar{x}|\bar{y})_{\bar{p}} + (\bar{z}|\bar{w})_{\bar{p}'}$$

hence

$$\frac{\rho(z, w)}{\rho(x, y)} \asymp \rho'(z, w).$$

If, on the contrary, $(\bar{x}|\bar{w})_{\bar{p}} \leq d_T(\bar{p} - \bar{p}')$ (see Figure 1(right)) then we first observe that $w \in B(x, t\rho(x, y))$ implies that $(x|w)_p \geq (x|y)_p - \log_a t + O(\delta)$ holds. Together with $(\bar{z}|\bar{w})_{\bar{p}} = (\bar{x}|\bar{w})_{\bar{p}}$, we deduce that

$$\begin{aligned} (\bar{w}|\bar{z})_{\bar{p}'} \leq (\bar{w}|\bar{p})_{\bar{p}'} &= (\bar{x}|\bar{y})_{\bar{p}} - (\bar{p}'|\bar{w})_{\bar{p}} \\ &= (\bar{x}|\bar{y})_{\bar{p}} - (\bar{x}|\bar{w})_{\bar{p}} \\ &\leq \log_a t + O(\delta) \end{aligned}$$

and hence

$$\begin{aligned} (\bar{x}|\bar{y})_{\bar{p}} - (\bar{z}|\bar{w})_{\bar{p}} &\leq \log_a t + O(\delta) \\ &\leq 2\log_a t - (\bar{z}|\bar{w})_{\bar{p}'} + O(\delta). \end{aligned}$$

This translates into

$$\frac{\rho(z, w)}{\rho(x, y)} \lesssim t^2 \rho'(z, w)$$

and concludes the proof of (3) and of the lemma. \square

Before presenting the proof of Theorem 5.3, we introduce the following terminology. A *fine chain* joining two points x and y of a metric space M is a finite sequence of points m_0, \dots, m_n such that $m_0 = x$, $m_n = y$ and $d(m_j, m_{j+1}) \leq d(x, y)/2$ for all $0 \leq j < n$.

Proof of Theorem 5.3. We first pass to a subsequence so that $L_i \geq 5i$, so that the space M_i is not $5i$ -linearly connected. Lemma 5.4 implies that for each i we may find a pair of points (x_i, y_i) in M_i so that any fine chain in M_i joining x_i to y_i meets the (non-empty) set $\mathcal{A}_i = M_i \setminus B_{\rho_i}(x_i, i\rho_i(x_i, y_i))$.

Let ρ'_i be chosen according to Lemma 5.5, applied with $x = x_i, y = y_i$, and let ρ'_i be the corresponding re-centered δ -adapted visual metric from the conclusion of that lemma. Let (Y_∞, ρ'_∞) be a strong limit of a subsequence of (Y_i, ρ'_i) , which exists because there are finitely many isometry types of balls of any radius in the spaces Y_i . The spaces (M_i, ρ'_i) weakly Gromov–Hausdorff converge to (M_∞, ρ'_∞) where $M_\infty = \partial Y_\infty$ and ρ'_∞ is a δ -adapted visual metric, by Proposition 3.13.

Denote by $\phi_i: (M_i, \rho'_i) \rightarrow (M_\infty, \rho'_\infty)$ choices of (k, ε_i) -quasi-isometries that witness the weak Gromov–Hausdorff convergence. Choose z_i in \mathcal{A}_i . Let $\psi_i: M \rightarrow M_i$ be quasi-inverses for the ϕ_i 's. We may assume that, for each i , $\psi_i(\phi_i(z_i)) = z_i$ and that ψ_i is a (k, ε_i) -quasi-isometry, by replacing ε_i by $3k\varepsilon_i$ if necessary.

Passing to a subsequence and reindexing, we may also assume that $\varepsilon_i i^2$ tends to zero, and that $\phi_i(x_i)$ (resp. $\phi_i(y_i), \phi_i(z_i)$) converges to a point x_∞ (resp. y_∞, z_∞) in M_∞ . The points x_∞, y_∞ and z_∞ are at least $(1/kc)$ -apart by Lemma 5.5 (1) and (2).

Recall that every fine chain connecting x_i to y_i meets \mathcal{A}_i . The diameters with respect to ρ'_i of the \mathcal{A}_i go to 0 as $i \rightarrow \infty$ by Lemma 5.5(2), and $\phi_i(\mathcal{A}_i)$ converge to z_∞ in M_∞ . The intuitive idea is that the \mathcal{A}_i separate x_i from y_i , so in the limit z_∞ is a weak cut point of M_∞ . We now make this more precise.

In order to obtain a contradiction, suppose that z_∞ is not a weak cut point in M_∞ , so that there is a sub-continuum K of M_∞ which contains x_∞ and y_∞ but not z_∞ .

We first claim that, for i large enough, $\psi_i(K)$ does not meet \mathcal{A}_i . We show that there exists $\eta > 0$ such that, for any i large enough, $\rho'_i(\psi_i(w), w_i) \geq \eta$

for any $w \in K$ and $w_i \in \mathcal{A}_i$:

$$\begin{aligned}
\rho'_i(\psi_i(w), w_i) &\geq \rho'_i(\psi_i(w), \psi_i(\phi_i(z_i))) - \rho'_i(z_i, w_i) \\
&\geq \frac{1}{k} \rho'_\infty(w, \phi_i(z_i)) - \varepsilon_i - \text{diam}_{\rho'_i} \mathcal{A}_i \\
&\geq \frac{1}{k} \rho'_\infty(w, z_\infty) - \left(\frac{1}{k} \rho'_\infty(\phi_i(z_i), z_\infty) + \varepsilon_i + c/i \right) \\
&\geq \frac{1}{k} \rho'_\infty(z_\infty, K) - \left(\frac{1}{k} \rho'_\infty(\phi_i(z_i), z_\infty) + \varepsilon_i + c/i \right)
\end{aligned}$$

Thus, if i is sufficiently large, then $\eta = \rho'_\infty(z_\infty, K)/(2k)$ bounds $\rho'_i(\psi_i(w), w_i)$ from below.

For any $i \geq 0$, we may find a finite subset $F_i \subset K$ such that $K \subset \cup_{x \in F_i} B_{\rho'_\infty}(x, \varepsilon_i/k)$. Since K is a continuum, the nerve of this cover is connected. Applying ψ_i , we get that $\psi_i(K) \subset \cup_{x \in F_i} B_{\rho'_i}(\psi_i(x), 2\varepsilon_i)$ and the nerve of this cover is also connected. For i large enough, $\rho'_i(\psi_i(K), \mathcal{A}_i) \geq \eta > 2\varepsilon_i$, so Lemma 5.5 (3) implies that, for some $\lambda > 0$ and each $x \in F_i$,

$$B_{\rho'_i}(\psi_i(x), 2\varepsilon_i) \subset B_{\rho_i}(\psi_i(x), 2\lambda i^2 \varepsilon_i \rho_i(x_i, y_i))$$

so that

$$\psi_i(K) \subset \cup_{x \in F_i} B_{\rho_i}(\psi_i(x), 2\lambda i^2 \varepsilon_i \rho_i(x_i, y_i)).$$

The nerve of the cover $\{B_{\rho_i}(\psi_i(x), 2\lambda i^2 \varepsilon_i \rho_i(x_i, y_i))\}_{x \in F_i}$ of $\psi_i(K)$ is still connected.

Since we made sure that $\lim i^2 \varepsilon_i = 0$, we may pick i large enough so that $2\lambda i^2 \varepsilon_i \leq 1/4$. Thus, $\psi_i(F_i) \subset \psi_i(K)$ contains a fine chain connecting x_i and y_i . Since $\psi_i(K) \subset M_i \setminus \mathcal{A}_i$, we obtain a contradiction. \square

In the opposite direction we have the following statement about hyperbolic spaces with cut points in their boundary; the visual metrics on the boundary of such a space cannot be uniformly linearly connected. (A collection of metrics is *uniformly* linearly connected if there is some constant of linear connectedness which works for all of them.)

Proposition 5.6. *Let X be a proper geodesic δ -hyperbolic space whose boundary is connected and has a cut point. There is a sequence of points $\{x_i\}$ in X satisfying the following. If for each i , ρ_i is a δ -adapted visual metric based at x_i then the metrics $\{\rho_i\}$ are not uniformly linearly connected.*

Proof. Let $p \in \partial X$ be a cut point, and let a, b be in different components of $\partial X \setminus \{p\}$. Since every curve joining a and b has to go through the point p , it suffices to find points x_i so that the Gromov products $(a|p)_{x_i}$ and $(b|p)_{x_i}$ are bounded above, but the Gromov products $(a|b)_{x_i}$ tend to infinity.

Let γ be any geodesic ray in X tending to p . Then the points $x_i = \gamma(i)$ satisfy this property. \square

For the Bowditch boundary of a relatively hyperbolic pair, existence of a cut point prevents even non-uniform linear connectivity.

Proposition 5.7. *Let (H, \mathcal{P}) be a relatively hyperbolic pair, with each $P \in \mathcal{P}$ infinite. If the Bowditch boundary (when endowed with any visual metric) is linearly connected, then it does not have cut points.*

Proof. Suppose that the Bowditch boundary $\partial(H, \mathcal{P})$ has a cut point p . Then p is a parabolic point by [20, Theorem 1.1]. Let x and y lie in different components of $\partial(H, \mathcal{P}) \setminus \{p\}$. Let J be any arc connecting x and y , so containing p . It follows from Lemma 5.8 below that, for any visual distance ρ , we have for all $h \in \text{Stab}(p)$,

$$\frac{\rho(hx, hy)}{\text{diam}_\rho h(J)} \leq \frac{\rho(hx, hy)}{\rho(p, hx)} \lesssim \rho(p, hy).$$

Considering an infinite sequence (h_n) , the right-hand side tends to 0, so $\partial(H, \mathcal{P})$ cannot be linearly connected. \square

Lemma 5.8. *Let Υ be a proper geodesic δ -hyperbolic space and let P be a parabolic group of isometries of Υ fixing the parabolic point $p \in \partial\Upsilon$. Let ρ be a δ -adapted visual metric on $\partial\Upsilon$, and let $x, y \in \partial\Upsilon \setminus \{p\}$.*

There is a $C > 0$ so that, for all $g \in P$,

$$\rho(gx, gy) \leq C\rho(p, gx)\rho(p, gy).$$

Proof. It follows from [21, Prop. 8.11, Rem. 8.13.ii, Thm. 8.16] that horoballs based at p are quasi-preserved by P : there exists a constant $C_0 \geq 0$ depending only on δ such that, for any $x \in \Upsilon$ and any $g \in P$,

$$(3) \quad |\beta_p(x, g(x))| \leq C_0$$

where β_p denotes a Busemann function at p .

Let w be the basepoint for the visual metric ρ , and write $B = \max\{(x|p)_w, (y|p)_w\}$; fix $g \in P$. We claim that there is a constant $C_1 = C_1(\delta) \geq 0$ such that

$$(g(x)|g(y))_w \geq (g(x)|p)_w + (g(y)|p)_w - (C_1 + 2B).$$

From the claim, we may conclude:

$$\rho(gx, gy) \lesssim e^{-\varepsilon(g(x)|g(y))_w} \lesssim e^{-\varepsilon((g(x)|p)_w + (g(y)|p)_w)} \lesssim \rho(p, gx)\rho(p, gy).$$

To prove the claim, we first observe that, given any four points $x, y, w, p \in (\Upsilon \cup \partial\Upsilon)$ with $w \in \Upsilon$ and $p \in \partial\Upsilon$

$$(4) \quad \left| \left(\sup_{z \in (xy)} \beta_p(w, z) + (x|y)_w \right) - ((x|p)_w + (y|p)_w) \right| \leq C_2.$$

for some constant $C_2(\delta) \geq 0$. This may be checked by approximating these points by a tree [21, Thm 2.12].

Thus, we get

$$\sup_{z \in (xy)} \beta_p(w, z) \leq C_2 + 2B.$$

Since $g \in P$, we have by the quasi-cocycle property of Busemann functions [21, Prop. 8.2.iii] and by (3)

$$\left| \sup_{z \in (xy)} \beta_p(w, z) - \sup_{z \in (g(x)g(y))} \beta_p(w, z) \right| \leq C_0 + C_3$$

for some constant $C_3 = C_3(\delta)$. Hence, with a second application of (4) we obtain

$$(g(x)|g(y))_w \geq (g(x)|p)_w + (g(y)|p)_w - C_0 - C_3 - 2B - 2C_2$$

proving the claim. \square

5.2. Linear connectivity versus spherical connectivity. In this subsection we introduce our local-to-global criterion for linear connectedness of the boundary.

Definition 5.9. Let $\Delta, R \geq 0$, let Υ be a δ -hyperbolic and δ -visible metric space, and let $y \in \Upsilon$. We say that Υ is (Δ, R) -spherically connected at y if for any points $p, q \in \Upsilon$ with $d(p, y) = d(q, y) = R$, the following holds. There is a sequence $p = p_0, \dots, p_n = q$ so that $d(y, p_i) = R$, $(p_i|p_{i+1})_y \geq R - 5\delta$ and $(p|p_i)_y \geq (p|q)_y - \Delta$.

Lemma 5.10. For every δ, L there exists $\Delta = \Delta(\delta, L)$ such that the following holds. Let Υ be δ -hyperbolic and δ -visible and $y \in \Upsilon$ be such that any δ -adapted visual metric on $\partial\Upsilon$ with basepoint y is L -linearly connected. For every R we have that Υ is (Δ, R) -spherically connected at y .

Proof. Since the only base point considered here is the point y , we write $(\cdot|\cdot)_y = (\cdot|\cdot)$. We first gather some estimates that are used in the proof.

Let $z \in \Upsilon$ and ξ be the endpoint of a ray $[y, \xi)$ that passes at distance δ from z . Then the δ -hyperbolicity implies

$$(5) \quad (z|\xi) \geq d(y, z) - 2\delta.$$

Let $\xi, \zeta \in \partial\Upsilon$, $R > 0$. Let us consider points p and q at distance R from the base point y that are on rays defining ξ and ζ respectively. Then

$$(p|q) \geq \min\{(p|\xi), (\xi|\zeta), (\zeta|q)\} - 2\delta \geq \min\{R, (\xi|\zeta)\} - 2\delta \geq (\xi|\zeta) - 2\delta$$

since $(p|q) \leq R$, and, by symmetry, we get

$$(6) \quad (p|q) - 2\delta \leq (\xi|\zeta) \leq (p|q) + 2\delta.$$

We now establish the spherical connectivity. Let p and q be on the sphere of radius R around y . If $(p|q) \geq R - 5\delta$, then we are done. Otherwise,

we approximate p and q by rays γ_p and γ_q that pass within δ of p and q , respectively. The associated points on ∂Y are ξ_p and ξ_q .

The hyperbolicity and (5) implies

$$\begin{aligned} (\xi_p|\xi_q) &\geq \min\{(\xi_p|p), (p|q), (q|\xi_q)\} - 2\delta \\ &\geq \min\{R - 2\delta, (p|q)\} - 2\delta \end{aligned}$$

Since $(p|q) \leq R - 5\delta$ we have

$$(7) \quad (\xi_p|\xi_q) \geq (p|q) - 2\delta.$$

Let J be an arc connecting ξ_p and ξ_q such that $\rho(\xi, \xi_p) \leq L\rho(\xi_p, \xi_q)$ holds for every point ξ in J and let $\eta = e^{-\varepsilon R}/\kappa$. Since J is an arc, there is a sequence of points $(\xi_j)_{1 \leq j \leq n-1}$ on J represented by rays $\gamma_p = \alpha_1, \dots, \alpha_{n-1} = \gamma_q$ with $\rho(\xi_j, \xi_{j+1}) \leq \eta$. Let p_j be the intersection of α_j with the sphere of radius R about y , and write $p = p_0$ and $q = p_n$. Let us check that this chain satisfies the conditions of spherical connectedness.

We first note that $(p_0|p_1), (p_{n-1}|p_n) \geq R - \delta$ by construction; for $1 \leq j \leq n-2$, the estimate (6) shows us that

$$\begin{aligned} (p_j|p_{j+1}) &\geq (\xi_j|\xi_{j+1}) - 2\delta \\ &\geq \frac{-1}{\varepsilon} \log(\kappa\eta) - 2\delta \\ &\geq R - 2\delta. \end{aligned}$$

If $j = 1, n$, then $(p|p_j) \geq \min\{(p|q), R - 2\delta\} \geq (p|q)$ by (5) and since $(p|q) \leq R - 5\delta$. If $2 \leq j \leq n-1$, then by (6) and (7),

$$\begin{aligned} (p_1|p_j) &\geq (\xi_p|\xi_j) - 2\delta \\ &\geq \frac{-1}{\varepsilon} \log \kappa \rho(\xi_p, \xi_j) - 2\delta \\ &\geq \frac{-1}{\varepsilon} \log \kappa L \rho(\xi_p, \xi_q) - 2\delta \\ &\geq \frac{-1}{\varepsilon} \log(\kappa^2 L) + (\xi_p|\xi_q) - 2\delta \\ &\geq (p|q) - \frac{1}{\varepsilon} \log(\kappa^2 L) - 4\delta \end{aligned}$$

so

$$(p|p_j) \geq \min\{(p|p_1), (p_1|p_j)\} - \delta \geq (p|q) - \frac{1}{\varepsilon} \log(\kappa^2 L) - 5\delta$$

since $(p|p_1) \geq R - \delta \geq R - 5\delta \geq (p|q)$. Therefore, we may set

$$\Delta = \frac{1}{\varepsilon} \log(\kappa^2 L) + 5\delta$$

that only depends on δ and L . □

We now prove a kind of converse to Lemma 5.10.

Lemma 5.11. *For every Δ, δ there exists $R_0 = R_0(\Delta, \delta)$ and $L = L(\Delta, \delta)$ so that the following holds. Suppose that Υ is a δ -hyperbolic and δ -visible metric space, and that there exists $R \geq R_0$ so that for every $y \in \Upsilon$ the space Υ is (Δ, R) -spherically connected at y . Then any δ -adapted visual metric on $\partial\Upsilon$ with any basepoint is L -linearly connected.*

Proof. We fix a base point $y \in \Upsilon$ and let ρ denote a δ -adapted visual distance. We check the criterion of Lemma 5.4. Let $\xi, \zeta \in \partial\Upsilon$ with $\xi \neq \zeta$ and assume that $R \geq R_0$ for some constant R_0 that will be fixed later on.

Let us apply Lemma 5.5 to the triple (y, ξ, ζ) and denote by y' the new basepoint with associated visual distance ρ' .

Let p, q be at distance R from the base point y' on rays $[y', \xi)$ and $[y', \zeta)$ respectively. Therefore, the argument of (6) leads us to

$$(8) \quad (p|q)_{y'} - 2\delta \leq (\xi|\zeta)_{y'} \leq (p|q)_{y'} + 2\delta.$$

Since Υ is (Δ, R) -spherically connected at y' , we may find $p = p_0, \dots, p_n = q$ so that $d(y', p_j) = R$, $(p_j|p_{j+1})_{y'} \geq R - 5\delta$ and $(p|p_j)_{y'} \geq (p|q)_{y'} - \Delta$.

For any $1 \leq j \leq n-1$, we may find a ray that passes at distance at most δ from p_j with end point $\xi_j \in \partial\Upsilon$.

On the one hand, we have

$$\begin{aligned} (\xi|\xi_j)_{y'} &\geq \min\{(\xi|p)_{y'}, (p|p_j)_{y'}, (p_j|\xi_j)_{y'}\} - 2\delta \\ &\geq \min\{R, (p|q)_{y'} - \Delta, R - 2\delta\} - 2\delta \\ &\geq (p|q)_{y'} - \Delta - 2\delta \\ &\geq (\xi|\zeta)_{y'} - \Delta - 4\delta \end{aligned}$$

so that

$$(9) \quad \rho'(\xi, \xi_j) \leq \kappa e^{-\varepsilon(\xi|\xi_j)_{y'}} \leq \kappa^2 e^{\varepsilon(\Delta+4\delta)} \rho'(\xi, \zeta).$$

This bound shows that there is a choice of t , depending only on δ and Δ , so that for each j we have $\xi_j \notin \mathcal{A}_t$, where \mathcal{A}_t is the set defined in Lemma 5.5.(2) with $x = \xi$, $y = \zeta$. Lemma 5.5.(3) tells us that there is a constant $\lambda_0 = \lambda t^2 \geq 1$, where λ_0 depends only on δ and Δ such that

$$\frac{1}{\lambda_0} \frac{\rho(\xi_i, \xi_j)}{\rho(\zeta, \xi)} \leq \rho'(\xi_i, \xi_j) \leq \lambda_0 \frac{\rho(\xi_i, \xi_j)}{\rho(\zeta, \xi)}$$

holds for all $0 \leq i, j \leq n$.

On the other hand, the same argument as for (7) leads to

$$(\xi_j|\xi_{j+1})_{y'} \geq (p_j|p_{j+1})_{y'} - 2\delta \geq R - 7\delta$$

so that, by definition of R , we have

$$\rho'(\xi_j, \xi_{j+1}) \leq \kappa e^{-\varepsilon(R-7\delta)}.$$

Hence, Lemma 5.5.(3) (with the same choice of t , and hence λ_0 , as above) implies that

$$\begin{aligned} \rho(\xi_j, \xi_{j+1}) &\leq \lambda_0 \rho'(\xi_j, \xi_{j+1}) \rho(\zeta, \xi) \\ &\leq \lambda_0 \kappa e^{-\varepsilon(R_0 - 7\delta)} \rho(\zeta, \xi) \\ &\leq \rho(\xi, \zeta)/2, \end{aligned}$$

if we pick R_0 large enough. This implies that the assumptions of Lemma 5.4 hold and we may deduce that ∂Y is L -linearly connected where L only depends on δ and Δ . \square

6. SHELLS AND TUBE COMPLEMENTS

In this section we study the geometry of large tubes and complements of large tubes around bi-infinite geodesics in hyperbolic spaces, in particular when the boundary of the hyperbolic space is S^2 . One of the main results of the section is Corollary 6.36, which shows that the coarse fundamental group of the boundary of a tube is isomorphic to the fundamental group of the boundary at infinity minus the limit points of the tube, which is isomorphic to \mathbb{Z} . The isomorphism arises from a closest point projection map, see Definition 6.3.

Another result important for the later parts of the paper is Theorem 6.24, which roughly speaking says that removing a large tube around bi-infinite geodesic and adding a combinatorial horoball on the corresponding shell results in a hyperbolic space; this space will serve as a “local” model for other hyperbolic spaces we construct later on.

Two key technical constructions in this section are completed shells and completed tube complements, Definitions 6.12 and 6.15, which allow us for example to make sense of the cusped space described above.

We now fix some notation and hypotheses which we use throughout this section. The reader might want to keep in mind a space X_0 acted upon geometrically by a hyperbolic group G , as will be the case in Section 8 below.

Assumption 6.1. *Let X_0 be a locally finite graph and suppose that X_0 is δ_0 -hyperbolic and δ_0 -visible for some δ_0 .*

Let $\lambda_0 \geq 0$ be a constant, and let Y_0 denote a λ_0 -quasi-convex subset of X_0 such that $\partial X_0 \setminus \Lambda Y_0$ is path connected, where ΛY_0 is the limit set of Y_0 in ∂X_0 . Suppose further that if $y \in Y_0$ and α is a geodesic in X_0 starting at y and limiting to a point in ΛY_0 then α stays entirely in the λ_0 -neighborhood of Y_0 . Fix a basepoint $w_0 \in Y_0$.

Suppose further that some δ_0 -adapted visual metric ρ_0 on ∂X_0 based at w_0 is L_0 -linearly connected and N_{doub} -doubling.

6.1. Shells and projections.

Definition 6.2. Suppose that Z is a metric space, $W \subset Z$, and let $K > 0$ be an integer.

The K -shell of W in Z is $S_K^Z(W) = \{z \in Z \mid d(z, W) = K\}$, and the (open) K -tube of W in Z is $T_K^Z(W) = \{z \in Z \mid d(z, W) < K\}$. The (closed) K -neighborhood of W in Z is $N_K^Z(W) = \{z \in Z \mid d(z, W) \leq K\}$.

Often, the space Z is implicit, and we write $S_K(W)$, $T_K(W)$, and $N_K(W)$.

We remark that if Z is a graph and W is a sub-graph, then $S_K^Z(W)$ is in the 0-skeleton of Z , since K is an integer. This will be important later.

Recall that the pair (X_0, Y_0) satisfies the conditions given in Assumption 6.1. Recall also that we fixed $w_0 \in Y_0$. Recall that the limit set of Y_0 in ∂X_0 is denoted by ΛY_0 .

Definition 6.3. For $K \geq 0$, define the map $\Pi_K = \Pi_{S_K(Y_0)}: \partial X_0 \setminus \Lambda Y_0 \rightarrow S_K(Y_0)$ as follows. For each $p \in \partial X_0 \setminus \Lambda Y_0$, choose a representative ray $r(p)$ from w_0 to p . Then $\Pi_K(p)$ is the last point on $r(p)$ at distance K from Y_0 .

Note that the map Π_K involves some choices, such as w_0 and $r(p)$. For large choices of K , Lemma 6.5 shows that the choices don't matter much. We use the chosen paths $r(p)$ later.

The next lemma shows that a geodesic ray starting on Y_0 which leaves $T_K(Y_0)$ does not stay near $S_K(Y_0)$ for very long.

Lemma 6.4. *Suppose $K \geq \lambda_0 + 3\delta_0$. Let α be a geodesic ray from $x \in Y_0$ tending to a point in $\partial X_0 \setminus \Lambda Y_0$. Let $y \in \alpha \cap S_K(Y_0)$. If z is any point on α with $d(z, S_K(Y_0)) \leq 2\delta_0$, then $d(z, y) \leq 6\delta_0$.*

Proof. Let u be a point of α so that z and y are both between u and $x \in Y_0$. Project u to Y_0 and call the projection u' . Let Δ be a geodesic triangle with one side that part of α from x to u , and third vertex u' . Since the side xu' lies within λ_0 of Y_0 , and $K \geq \lambda_0 + 3\delta_0$, z and y are at most δ_0 from points z' and y' on the uu' side of this triangle. However, the geodesic uu' travels as quickly to Y_0 as possible, so since $|d(z', Y_0) - K| \leq 3\delta_0$ and $|d(y', Y_0) - K| \leq \delta_0$, we have $d(z', y') \leq 4\delta_0$. Thus $d(z, y) < \delta_0 + \delta_0 + 4\delta_0$, as required. \square

Lemma 6.5. *For any $K \geq \lambda_0 + 3\delta_0$, and any $p \in \partial X_0 \setminus \Lambda Y_0$, changing the ray $r(p)$ or the basepoint w_0 in Definition 6.3 changes the location of $\Pi_K(p)$ by at most $8\delta_0$.*

Proof. Let w and w' be two arbitrary points on Y_0 . Let g and g' be two geodesics from w and w' limiting at p . There is a partially ideal triangle with vertices w, w', p , and sides g, g' and a geodesic $[w, w']$. Let y and y' be the points on g and g' at distance K from Y_0 . Since $K > \lambda_0 + 3\delta_0$, the point

y' on g' lies within $2\delta_0$ of some point y'' on g . By Lemma 6.4, this point y'' is within $6\delta_0$ of y . \square

6.2. The coarse topology of shells. Throughout this subsection, we let $S_K(Y_0)$ denote $S_K^{X_0}(Y_0)$, and similarly drop the superscript for tubes. For our projections, we continue to work with a fixed w_0 on Y_0 .

The following is well-known, we include an argument for completeness.

Lemma 6.6. *For any $K \geq \lambda_0$ the tube $T_K(Y_0)$ is $2\delta_0$ -quasi-convex.*

Proof. Let $x, y \in T_K(Y_0)$ and consider a geodesic $[x, y]$ between them. Also, let $x', y' \in Y_0$ be closest points to x, y respectively. Consider any $z \in [x, y]$, and we have to show that it lies within $2\delta_0$ of $T_K(Y_0)$. By hyperbolicity, z lies $2\delta_0$ close to a geodesic $[x, x']$, $[y, y']$, or $[x', y']$. The first two geodesics are contained in $T_K(Y_0)$, while the third one is contained in $T_{\lambda_0}(Y_0) \subseteq T_K(Y_0)$, so we are done. \square

Lemma 6.7. *For any $K \geq \lambda_0 + 3\delta_0$, the map Π_K is $8\delta_0$ -coarsely surjective. Furthermore, any $p \in \partial X_0 \setminus \Lambda Y_0$ has an open neighborhood U_p so that for every $q \in U_p$ we have $d(\Pi_K(p), \Pi_K(q)) \leq 8\delta_0$.*

Proof. By δ_0 -visibility, any point x on $S_K(Y_0)$ lies δ_0 -close to a geodesic ray r from w_0 to some $q \in \partial X_0$. It follows that x is within δ_0 of a point y on r . Let $r(q)$ be the geodesic ray used to define $\Pi_K(q)$ starting at w_0 . We can form a geodesic triangle between w_0 , and two points on r and $r(q)$ which are far away from $S_K(Y_0)$. Then y is distance at most δ_0 from a point z on $r(q)$. It follows from Lemma 6.4 that x is within $8\delta_0$ of $\Pi_K(q)$. This proves the first assertion.

Let p be in $\partial X_0 \setminus \Lambda Y_0$ and consider the geodesic $r(p)$ used to define $\Pi_K(p)$. Pick a neighborhood U_p of p in ∂X_0 such that any geodesic from w_0 to a point q in U_p fellow travels $r(p)$ until both geodesics are a distance at least $2K$ from Y_0 . Since (partially) ideal triangles are $2\delta_0$ -thin, $\Pi_K(q)$ is within $2\delta_0$ of a point on $r(p)$. Then using Lemma 6.4, $d(\Pi_K(p), \Pi_K(q)) \leq 8\delta_0$. \square

Definition 6.8. Let (A, d) be a metric space. An s -path in A is a function $f: \{0, 1, \dots, n_f\} \rightarrow A$ so that $d(f(i-1), f(i)) \leq s$ for each $i \in \{1, \dots, n_f\}$.

The number $\sum_{i=1}^{n_f} d(f(i-1), f(i))$ is called its *length*.

We call a space s -path-connected if every pair of points is joined by an s -path.

Notice that for any $x, y \in A$ and any $s > 0$ the length of any s -path between x and y is at least $d(x, y)$.

Lemma 6.9. *For any $K \geq \lambda_0 + 3\delta_0$ the shell $S_K(Y_0)$ is $8\delta_0$ -path connected.*

Proof. Let $x, y \in S_K(Y_0)$. By the first part of Lemma 6.7, there exist $p, q \in \partial X_0 \setminus \Lambda Y_0$ so that $d(x, \Pi_K(p)), d(y, \Pi_K(q)) \leq 8\delta_0$. Since $\partial X_0 \setminus \Lambda Y_0$ is path-connected, we can consider a path from p to q in it. We can then cover the path with finitely many open neighborhoods as described in the second part of Lemma 6.7, and apply Π_K to obtain a $8\delta_0$ -path between $\Pi_K(p)$ and $\Pi_K(q)$. \square

The following result is a more precise version of [25, Lemma 6.21], and we explain how to adapt the proof of [25]. Recall that $\varepsilon = \varepsilon(\delta_0)$ and κ are the constants from Definition 3.10, and that ∂X_0 is equipped with a δ_0 -adapted visual metric ρ_0 based at w_0 which is L_0 -linearly connected and N_{doub} -doubling.

Lemma 6.10. *For any $B > 0$ there exists $N = N(\delta_0, L_0, N_{\text{doub}}, B)$ such that for every pair of rays η_1, η_2 based at w_0 with different endpoints in ∂X_0 there exists $n \leq N$, and a sequence of rays $\eta_1 = \alpha_1, \dots, \alpha_n = \eta_2$ starting at w_0 with $(\alpha_i | \alpha_{i+1})_{w_0} \geq (\eta_1 | \eta_2)_{w_0} + B$ and $(\alpha_i | \eta_1)_{w_0} \geq (\eta_1 | \eta_2)_{w_0} - \Delta_0$, where $\Delta_0 = \frac{\log(2\kappa^2 L_0)}{\varepsilon} + 20\delta_0$.*

Proof. The statement of the lemma differs from that of [25, Lemma 6.21] in two ways. First, the constant R (which is our Δ_0) in the statement of [25, Lemma 6.21] is allowed to depend on the hyperbolic space rather than just on δ_0 and L_0 . Second, [25, Lemma 6.21] provides no bound on n .

Regarding the first difference, the explicit constant Δ_0 , which depends only on L_0 and δ_0 , is actually given in the proof in [25] (third line).

Regarding the bound on n , we first outline the proof of [25, Lemma 6.21]. First of all, in said proof the limit points p_1, p_2 of η_1, η_2 get connected by an arc I of diameter at most $2L_0\rho_0(p_1, p_2)$. The sequence of rays α_i is then obtained as any chain of rays with limit points a_i in I where the a_i form a chain that starts at p_1 , ends at p_2 , and any two consecutive points are sufficiently close. Explicitly, we need

$$\rho(a_i, a_{i+1}) \leq e^{-\varepsilon B} \rho(p_1, p_2) / \kappa^2,$$

which guarantees

$$e^{-\varepsilon(\alpha_i | \alpha_{i+1})_{w_0}} \leq \kappa \rho(a_i, a_{i+1}) \leq e^{-\varepsilon B} \rho(p_1, p_2) / \kappa \leq e^{-\varepsilon B} e^{-\varepsilon(\eta_1 | \eta_2)_{w_0}},$$

and in turn this yields the required $(\alpha_i | \alpha_{i+1})_{w_0} \geq (\eta_1 | \eta_2)_{w_0} + B$.

To get the desired bound on n , we observe that subspaces of an N_{doub} -doubling space are N_{doub}^2 -doubling (the intersection of a ball with a subspace is contained in a ball of twice the radius in the induced metric). From this, we see that I can be covered by at most $N_{\text{doub}}^{2\lceil \log_2(4L_0 e^{-\varepsilon B} / (2\kappa^2)) \rceil}$ balls of radius $e^{-\varepsilon B} \rho(p_1, p_2) / (2\kappa^2)$. Therefore, we can extract a chain of points with the required bound on the length from the centers of these balls. \square

Recall the hyperbolicity constant δ_0 and the linear connectedness constant L_0 were fixed in Assumption 6.1, as was the quasi-convexity constant λ_0 of Y_0 . In the following we also fix the constant Δ_0 from Lemma 6.10.

Lemma 6.11. *There is a function $\Phi: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ so that for any $C > 0$ and $K \geq C + \Delta_0 + \lambda_0 + 3\delta_0$, if $x, y \in S_K(Y_0)$ satisfy $d(x, y) \leq C$ then there exists an $8\delta_0$ -path in $S_K(Y_0)$ from x to y of length at most $\Phi(C)$.*

Proof. We use the same idea as in Lemma 6.9. Let $C > 0$, and suppose that $K > C + \Delta_0 + \lambda_0 + 3\delta_0$. Let $x, y \in S_K(Y_0)$ satisfy $d(x, y) \leq C$. Lemma 6.7 says that Π_K is $8\delta_0$ -coarsely surjective, so by moving x and y at most $8\delta_0$, we may assume that $x = \Pi_K(p)$ and $y = \Pi_K(q)$, for some $p, q \in \partial X_0 \setminus \Lambda Y_0$.

Let $r(p)$ and $r(q)$ denote the geodesic rays from w_0 defining $\Pi_K(p)$ and $\Pi_K(q)$, respectively. Note that because $x, y \in S_K(Y_0)$ and $d(x, y) \leq C$, we know that there are points $y_1 \in r(p)$ and $y_2 \in r(q)$ which lie at distance at least $K - (1/2C + \delta_0)$ from Y_0 and so that $d(y_1, y_2) \leq \delta_0$. As in the conclusion of Lemma 6.10, choose a sequence $r(p) = \alpha_1, \dots, \alpha_n = r(q)$ of geodesic rays so that $(\alpha_i \mid \alpha_{i+1})_{w_0} \geq K + 10\delta_0$ for each i , and $(\alpha_i \mid r(p))_{w_0} \geq (r(p) \mid r(q))_{w_0} - \Delta_0$. By the same lemma, we can also require that n is bounded in terms of $\delta_0, L_0, N_{\text{doub}}, C$ only (as $(r(p) \mid r(q))_{w_0}$ is at least K minus a constant depending on C and δ_0). It now follows that there exists a point z_i on each α_i at distance at least $K - ((1/2)C + \delta_0) - \Delta_0 \geq \lambda_0 + 2\delta_0$ from Y_0 . By the assumption on Y_0 in Assumption 6.1, any geodesic ray from w_0 which leaves the λ_0 -neighborhood of Y_0 cannot limit on ΛY_0 . Thus, the sequence of rays α_i each limit on a point p_i of $\partial X_0 \setminus \Lambda Y_0$.

As in the proof of Lemma 6.9, the points $x = \Pi_K(p_1), \dots, \Pi_K(p_n) = y$ form a $8\delta_0$ -path. \square

6.3. Completed shells.

Definition 6.12. Let Γ be a graph, and suppose $\Xi \subset \Gamma$ is a subgraph. Let K, s be positive integers, and let $S_K^\Gamma(\Xi)$ be the K -shell about Ξ in Γ (which consists of vertices). The *completed K -shell about Ξ in Γ at scale s* , denoted $\mathcal{S}_{K,s}^\Gamma(\Xi)$ is the graph obtained from $S_K^\Gamma(\Xi)$ by joining each pair of points $a, b \in S_K^\Gamma(\Xi)$ so that $d_\Gamma(a, b) \leq s$ with an edge-path of length $d_\Gamma(a, b)$.

Convention 6.13. *In applying Definition 6.12 above, we always use $s = 8\delta_0$. Thus, we use $\mathcal{S}_K^\Gamma(\Xi)$ to denote $\mathcal{S}_{K,8\delta_0}^\Gamma(\Xi)$.*

Our default space is X_0 , so we use $\mathcal{S}_K(X_0)$ to denote $\mathcal{S}_K^{X_0}(\Xi) = \mathcal{S}_{K,8\delta_0}^{X_0}(\Xi)$.

The following is an immediate corollary of Lemma 6.9.

Lemma 6.14. *If $K \geq \lambda_0 + 3\delta_0$ the completed shell $\mathcal{S}_K(Y_0)$ is connected.*

Definition 6.15. Let Γ be a graph and $\Xi \subset \Gamma$ a subgraph. Let K, s be positive integers. The *completed K -tube complement of Ξ in Γ at scale s* , denoted $\mathcal{C}_{K,s}^\Gamma(\Xi)$ is the space

$$\mathcal{C}_{K,s}^\Gamma(\Xi) = (\Gamma \setminus T_K^\Gamma(\Xi)) \sqcup_{S_K^\Gamma(\Xi)} \mathcal{S}_{K,s}^\Gamma(\Xi),$$

obtained from ‘completing’ the copy of $S_K^\Gamma(\Xi)$ in $\Gamma \setminus T_K^\Gamma(\Xi)$ at scale s .

Note that $\mathcal{C}_{K,s}^\Gamma(\Xi)$ and $\mathcal{S}_{K,s}^\Gamma(\Xi)$ are graphs.

Convention 6.16. As in Convention 6.13, when completing complements of tubes about Y_0 in X_0 , we always use $s = 8\delta_0$ and denote the space $\mathcal{C}_{K,s}^\Gamma(\Xi)$ by $\mathcal{C}_K^\Gamma(\Xi)$, and when $\Gamma = X_0$ we omit the superscript and simply write $\mathcal{C}_K(\Xi)$.

Lemma 6.17. Suppose that $K \geq \lambda_0 + 3\delta_0$. Then $\mathcal{C}_K(Y_0)$ is a connected graph.

Proof. By Lemma 6.14, $\mathcal{S}_K(Y_0)$ is connected, so we only need to show that every point of $\mathcal{C}_K(Y_0)$ is connected to the completed shell. Take a point in $X_0 \setminus T_K^Z(Y_0)$. It is connected to Y_0 by a path in X_0 (since X_0 is a connected graph). This path hits $S_K^{X_0}(Y_0) \subset \mathcal{S}_K(Y_0)$. \square

Convention 6.18. Whenever we talk about a completed tube complement, we ensure that Lemma 6.17 applies, and the implied metric is the path metric.

6.4. Geometry of cusped spaces. The main result in this subsection is Theorem 6.24, which gives us control over the geometry of a space obtained by attaching a horoball to a K -shell in a tube complement.

Definition 6.19. Suppose Z is a metric space, $W \subset Z$ and $z \in Z$. Then $P_W(z)$ denotes a point in W which is closest to z (amongst all points in W).

Remark 6.20. All of our spaces are locally compact and δ -hyperbolic, and our subsets W are closed, so such points $P_W(z)$ exist, though they typically are not unique. Nothing that we do depends on the particular choice of $P_W(z)$.

Lemma 6.21. Let Υ be a δ -hyperbolic metric space and A a λ -quasi-convex subset. Suppose $K \geq 4\delta + \lambda$ and let $y, z \in \Upsilon$, and suppose there is a geodesic $[z, y] \subset \Upsilon \setminus N_K(A)$. Suppose $y' \in [y, P_A(y)] \cap S_K(A)$ and $z' \in [z, P_A(z)] \cap S_K(A)$ (for some choices of geodesics $[y, P_A(y)]$, $[z, P_A(z)]$). Then $d(y', z') < 8\delta$.

Proof. Choose a geodesic $[P_A(y), P_A(z)]$ and consider the geodesic quadrilateral $[y, P_A(y)] \cup [P_A(y), P_A(z)] \cup [P_A(z), z] \cup [y, z]$. Since $K \geq 4\delta + \lambda$, every point of $[y, z]$ is distance at least 4δ from $[P_A(y), P_A(z)]$. Choose a geodesic

$[P_A(y), z]$ and consider the geodesic triangle $[y, P_A(y)] \cup [y, z] \cup [P_A(y), z]$. Since $y' \in [y, P_A(y)]$, y' is within δ of either $[y, z]$ or $[P_A(y), z]$. If y' is within δ of some point q on $[P_A(y), z]$ then q is not within δ of $[P_A(y), P_A(z)]$, so it is distance δ from a point q' on $[z, P_A(z)]$. We have $|d(q', A) - K| \leq 2\delta$, since $d(y', A) = K$. Hence the distance from q' to z' is at most 2δ . So $d(y', z') < 4\delta$ in this case.

Therefore, we assume that y' is within δ of some point $u \in [y, z]$. By a symmetric argument we can assume that z' is within δ of some $v \in [y, z]$. If $u = v$ we are done. We assume that the order along $[y, z]$ is y, u, v, z (else it follows easily that $d(y', z') < 4\delta$).

Using λ -quasi-convexity, $[y, z]$ lies outside of the 4δ -neighborhood of $[P_A(y), P_A(z)]$, so every point of the subsegment $[u, v] \subseteq [y, z]$ lies within 2δ of $[y, P_A(y)] \cup [z, P_A(z)]$. Let p be a point on $[u, v]$ which is distance at most 2δ from some point y_p on $[y, P_A(y)]$ and distance at most 2δ from a point z_p on $[z, P_A(z)]$.

If z_p is closer to A than z' , it is at most 2δ closer, so $d(z', z_p) < 2\delta$. Then $d(z', p) < 4\delta$. If z_p is further away from A than z' , then z_p is on $[z', z]$. Consider the triangle $[z', z], [z, v], [v, z']$, where the length of $[v, z']$ is at most δ . If z_p is δ -close to $[v, z']$, then $d(z_p, z') < 2\delta$ and $d(z', p) < 4\delta$. If z_p is δ -close to $[z, v]$, v is at most 3δ from p , so again $d(z', p) < 4\delta$. So in any case $d(z', p) < 4\delta$. A similar argument shows $d(y', p) < 4\delta$, so $d(y', z') < 8\delta$, as required. \square

Lemma 6.22. *Let Υ be a δ -hyperbolic space, let A be a λ -quasi-convex subset of Υ , and suppose α is a geodesic ray starting at $w \in A$ and not converging to a point in ΛA . For any $y, z \in \alpha$ which are distance greater than $6\delta + \lambda$ from A , we have*

$$d(y, z) \sim_{10\delta} |d(y, A) - d(z, A)|.$$

Proof. Without loss of generality we assume that $d(y, w) > d(z, w)$.

An easy argument shows that under the hypotheses, the subsegment of α joining z to y lies entirely outside the $(5\delta + \lambda)$ -neighborhood of A , so we may apply Lemma 6.21 with $K = 4\delta + \lambda$. Let $y'' = P_A(y)$, $z'' = P_A(z)$, and let $y' \in [y, y''] \cap S_K(A)$, $z' \in [z, z''] \cap S_K(A)$.

Another easy argument using quasi-convexity of A shows $(w | y)_{y''}$ and $(w | z)_{z''}$ are bounded above by $\lambda + \delta$. Since $K > \lambda + \delta$, there are points \hat{y}, \hat{z} on α within δ of y', z' , respectively, and satisfying $d(y, \hat{y}) = d(y, y')$, $d(z, \hat{z}) = d(z, z')$. Lemma 6.21 implies that $d(\hat{y}, \hat{z}) \leq 10\delta$. So in particular we have $d(z, \hat{z}) \sim_{10\delta} d(z, \hat{y})$, and so (since $d(y, A) = d(y, y') + d(y', y'') = d(y, \hat{y}) + K$, and similarly for $d(z, A)$)

$$d(y, A) - d(z, A) \sim_{10\delta} d(y, \hat{y}) - d(z, \hat{y}) = d(y, z).$$

The last equality holds because \widehat{y} lies between w and z on α . This is because \widehat{y} is at most $5\delta + \lambda$ from A , but, as noted in the first paragraph, the subsegment of α joining z to y lies entirely outside the $(5\delta + \lambda)$ -neighborhood of A . \square

Definition 6.23. Let $K > \lambda_0 + \delta_0$ and suppose $X_0^{\text{cusp}}(Y_0, K)$ is obtained by gluing a combinatorial horoball $\mathcal{D} = \mathcal{H}(\mathcal{S}_K(Y_0))$ based on $\mathcal{S}_K(Y_0)$ to $\mathcal{C}_K(Y_0)$.

We say that $X_0^{\text{cusp}}(Y_0, K)$ is obtained from X_0 by *cusping the K -shell around Y_0* . Note that this depends on the set Y_0 , which is fixed in Assumptions 6.1. Therefore, we drop the Y_0 from our notation and refer simply to $X_0^{\text{cusp}}(K)$.

The hyperbolicity constant of $\mathcal{D} = \mathcal{H}(\mathcal{S}_K(Y_0))$ does not depend on K .

For each pair (x, y) of points in $X_0^{\text{cusp}}(K)$, we define a path $\mathcal{P}(x, y)$ connecting points x, y . For each x outside the tube, choose once and for all some $x' \in \mathcal{S}_K(Y_0)$ closest to x and a geodesic $[x, x']$ in X_0 from x to x' . Two such choices of x' differ only by δ_0 . If x, y lie outside the horoball \mathcal{D} , then let $[x, y]$ be a geodesic in X_0 from x to y . If $[x, y]$ does not intersect $T_K(Y_0)$, then set $\mathcal{P}(x, y) = [x, y]$. Otherwise, let $\mathcal{P}(x, y)$ be the union of the two geodesics $[x, x']$ and $[y, y']$ and a geodesic in \mathcal{D} connecting x', y' . If x, y both lie in \mathcal{D} , then let $\mathcal{P}(x, y)$ just be a geodesic in \mathcal{D} connecting them. If one of them, say x , lies outside \mathcal{D} and y lies in \mathcal{D} , then again consider a shortest geodesic $[x, x']$ from x to (the closure of) $T_K(Y_0)$, and define $\mathcal{P}(x, y)$ to be the concatenation of $[x, x']$ and a geodesic in \mathcal{D} from x' to y .

Theorem 6.24. *Under Assumptions 6.1 there exists $\delta_1 = \delta_1(\delta_0)$ so that for any $K \geq \lambda_0 + 4\delta_0$ the space $X_0^{\text{cusp}}(K)$ is δ_1 -hyperbolic and δ_1 -visible.*

Proof. We prove that $X_0^{\text{cusp}}(K)$ is hyperbolic by applying Proposition 3.14 to the paths $\mathcal{P}(x, y)$.

Suppose $d_{X_0^{\text{cusp}}(K)}(x, y) = 1$. If x, y lie outside \mathcal{D} , then so does the edge connecting them, and $\mathcal{P}(x, y)$ is an edge. This is also the case if x, y both lie in \mathcal{D} . In the mixed case, the point lying in \mathcal{D} actually lies in the boundary horosphere, and again $\mathcal{P}(x, y)$ is a single edge.

To apply Proposition 3.14 we must show that the \mathcal{P} -triangles are uniformly slim. We consider the case of vertices x, y, z all lying outside \mathcal{D} , the other cases being similar. Suppose that none of $\mathcal{P}(x, y)$, $\mathcal{P}(x, z)$ and $\mathcal{P}(y, z)$ intersects \mathcal{D} . Slimness of the corresponding \mathcal{P} -triangle follows from hyperbolicity of X_0 and the following claim:

Claim 6.24.1. For any $a, b \in X_0 \setminus T_K(Y_0)$ if $d_{X_0}(a, b) \leq 8\delta_0$, then $d_{X_0^{\text{cusp}}(K)}(a, b) \leq 8\delta_0$.

Proof. There is a geodesic α in X_0 of length at most $8\delta_0$ joining a, b . Every maximal subpath of α that is contained in the closure of $T_K(Y_0)$ has endpoints at distance at most $8\delta_0$ from each other in X_0 . Since $s = 8\delta_0$ (see Convention 6.13) the distance between these endpoints in the completed shell (and in $X_0^{\text{cusp}}(K)$) is no more than the distance in X_0 . \square

Suppose that all sides $\mathcal{P}(x, y)$, $\mathcal{P}(x, z)$ and $\mathcal{P}(y, z)$ intersect \mathcal{D} , then we can decompose the triangle into

- 3 pairs of geodesics in X_0 that coincide, and
- a geodesic triangle in \mathcal{D} .

From hyperbolicity of \mathcal{D} , we get the desired conclusion (letting δ_1 be bigger than this hyperbolicity constant).

Let us now consider the case that $\mathcal{P}(x, y)$, $\mathcal{P}(x, z)$ intersect \mathcal{D} , but $\mathcal{P}(y, z)$ does not. Let x', y', z' be closest points to x, y, z in the closure of $T_K(Y_0)$ (chosen as in the definition of \mathcal{P}). In this case we decompose the triangle into

- a pair of coinciding geodesics starting at x ($[x, x']$),
- geodesics in \mathcal{D} from x' to y' and from x' to z' ,
- a geodesic $[y, z]$ and geodesics $[y, y']$, $[z, z']$.

Notice that, since $[y, z]$ does not intersect $T_K(Y_0)$, we have $d_{X_0}(y', z') \leq 8\delta_0$ by Lemma 6.21. Hence, hyperbolicity of \mathcal{D} says that the geodesics in the second item fellow-travel with uniform constant. By considering the quadrilateral with vertices y, y', z, z' , and using the fact that $y' = P_{S_K(Y_0)}(y)$ and $z' = P_{S_K(Y_0)}(z)$, we see any point on a geodesic from the third item is within X_0 -distance $6\delta_0$ of the union of the other two. So we are done in this case as well, since the $X_0^{\text{cusp}}(K)$ -distance is bounded in terms of this in Claim 6.24.1 The case that only one of the sides, say $\mathcal{P}(x, y)$ intersects \mathcal{D} is similar, except that in this case we have the bound $d_{X_0^{\text{cusp}}}(x', y') \leq d_{X_0^{\text{cusp}}}(x', z') + d_{X_0^{\text{cusp}}}(z', y') \leq 16\delta_0$.

So far we not only proved hyperbolicity of $X_0^{\text{cusp}}(K)$, but we also have a description of paths $\mathcal{P}(\cdot, \cdot)$ that stay close to geodesics.

It remains to show visibility. In order to show it, we prove that there exists a constant $v = v(\delta_0)$ such that every $p \in X_0^{\text{cusp}}(K)$ lies v -close to a bi-infinite $(v(\delta_0), v(\delta_0))$ -quasi-geodesic. This suffices since this implies that we can “prolong” any geodesic using one of the two rays from a geodesic line passing near its endpoint.

We first claim that there exists $v_0(\delta)$ such that given any $x \in X_0 \setminus T_K(Y_0)$ there is an X_0 geodesic ray $\alpha(x)$ in X_0 starting on $S_K(Y_0)$, passing v_0 -close to x , and intersecting $S_K(Y_0)$ only at its starting point. To see this, note that coarse surjectivity of $\Pi_{K'}$ for K' such that p lies in the K' -shell (Lemma 6.7) says that there exists an X_0 -ray starting on Y_0 and passing $8\delta_0$ -close to p . We claim that a subray intersecting $S_K(Y_0)$ only at its starting point q is

a uniform quasi-geodesic in $X_0^{\text{cusp}}(K)$. This is because, by Lemma 6.22, for q' on the ray, the length of the subpath of the ray from q to q' is coarsely equal to the distance in X_0 , hence coarsely equal to the distance in $X_0^{\text{cusp}}(K)$, from q' to $S_K(Y_0)$. (We apply Lemma 6.22 to a smaller tube to get this, and we apply it in X_0 .) This implies that there cannot be “shortcuts” from q to q' in $X_0^{\text{cusp}}(K)$.

Consider first the case that p lies in \mathcal{D} . The required geodesic line consists of a vertical ray in \mathcal{D} within distance 1 of p , and a geodesic ray $\alpha(x)$ where $x \in S_K(Y_0)$ is the starting point of the vertical ray. The fact that these two rays form a quasi-geodesic line is again an application of Lemma 6.22. If p does not lie in \mathcal{D} , we concatenate $\alpha(p)$ and a vertical ray in \mathcal{D} . \square

6.5. Fundamental groups of completed shells. The purpose of this subsection is to understand the (coarse) fundamental groups of completed shells.

We first define a map from the fundamental group of $\partial X_0 \setminus \Lambda Y_0$ to the coarse fundamental group of a completed shell, in the proof of Proposition 6.27 below. We do this by projecting paths using Π_K . Since Π_K is not continuous, it does not literally send paths to paths, but by sampling a path in $\partial X_0 \setminus \Lambda Y_0$ sufficiently densely, we obtain a coarse path in the shell. “Connecting the dots” in the completed shell gives an honest path, which is well defined up to homotopy. We formalize this process of projection in the following definition.

Definition 6.25. Let $\eta : [a, b] \rightarrow \partial X_0 \setminus \Lambda Y_0$ be a path, and let $K \geq \lambda_0 + 3\delta_0$, so we can apply Lemma 6.7. A *projection of η to $\mathcal{S}_K(Y_0)$* is any path α constructed by the following procedure: Let $a = t_0, \dots, t_n = b$ be a partition so that each $\eta([t_i, t_{i+1}])$ lies in an open neighborhood U_p as in Lemma 6.7. For each i let $z_i = \Pi_K(\eta(t_i))$, let γ_i be a geodesic segment in $\mathcal{S}_K(Y_0)$ joining z_i to z_{i+1} , and let α be the concatenation $\gamma_0 \cdots \gamma_{n-1}$.

The projection of η is not unique, but the next lemma says that any two projections are homotopic rel endpoints.

Lemma 6.26. *Let (X_0, Y_0) be as in Assumption 6.1. For any $K \geq \lambda_0 + 3\delta_0$ and any $D \geq 24\delta_0$, the following holds: Let η be a path in $\partial X_0 \setminus \Lambda Y_0$, and let α and α' be two projections of η to $\mathcal{S}_K(Y_0)$. Then α and α' are homotopic rel endpoints in $(\mathcal{S}_K(Y_0))^D$.*

Moreover, the paths α and α' are Hausdorff distance at most $24\delta_0$ from one another, with respect to the path metric on $\mathcal{S}_K(Y_0)$.

Proof. We use the notation of Definition 6.25, adding a prime to objects used in defining α' . Suppose first that α and α' arise from the same partition of the domain of η , so $t_i = t'_i$ for each i . For each i , the segments γ_i and γ'_i form a loop of length at most $16\delta_0 < D$ so they are homotopic rel endpoints in $(\mathcal{S}_K(Y_0))^D$. Thus α and α' are homotopic rel endpoints in $(\mathcal{S}_K(Y_0))^D$.

Now suppose that α' arises from a refinement of the partition for α , say by adding a single point $t'_{i+1} \in (t_i, t_{i+1})$. The geodesics γ_i , γ'_i and γ'_{i+1} form a geodesic triangle of length at most $24\delta_0 \leq D$, so again α and α' are homotopic rel endpoints in $(\mathcal{S}_K(Y_0))^D$.

The general case follows readily.

For the statement about Hausdorff distance, note that α and α' come from partitions with a common refinement. If α'' is a projection using this refinement, it is straightforward to see that α'' is Hausdorff distance at most $12\delta_0$ from either α or α' . \square

Proposition 6.27. *Let (X_0, Y_0) be as in Assumption 6.1. For any $K \geq \lambda_0 + 3\delta_0$ and any $D \geq 24\delta_0$, and any point $* \in \partial X_0 \setminus \Lambda Y_0$, the projection Π_K induces a well-defined homomorphism*

$$\Pi_* : \pi_1(\partial X_0 \setminus \Lambda Y_0, *) \rightarrow \pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*)).$$

Proof. Let $\eta : [0, 1] \rightarrow \partial X_0 \setminus \Lambda Y_0$ be a loop based at $*$. We define $\Pi_*(\eta) = [\alpha]$ where α is any projection of η to $\mathcal{S}_K(Y_0)$. Lemma 6.26 tells us this is a well-defined element of $\pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*))$.

To show that $\Pi_*([\eta])$ is independent of the representative η , we argue as follows: Suppose η, ζ represent the same element of $\pi_1(\partial X_0 \setminus \Lambda Y_0, *)$. The image of a homotopy from η to ζ is a compact subset of $\partial X_0 \setminus \Lambda Y_0$, so we can cover it with finitely many neighborhoods U_p as in Lemma 6.7. Now choose a sufficiently fine triangulation T of $I \times I$ so that the image of every triangle is contained in one of the U_p . Take projections to $\mathcal{S}_K(Y_0)$ of all the edges of this triangulation. Each edge projects to a path of length at most $8\delta_0$, so the boundary of each triangle projects to a loop of length at most $24\delta_0 \leq D$. Putting these together gives a homotopy (rel $\Pi_K(*)$) in $(\mathcal{S}_K(Y_0))^D$ from a projection of η to a projection of ζ . In particular these projections represent the same element of $\pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*))$.

Finally, we show that Π_* is a homomorphism. Let η_1 and η_2 be loops in $\partial X_0 \setminus \Lambda Y_0$ based at $*$. If we ensure $\frac{1}{2}$ is a vertex in the subdivision of $[0, 1]$ then we may obtain a projection of $\eta_1 \cdot \eta_2$ which is a concatenation $\alpha_1 \cdot \alpha_2$, where α_i is a projection of η_i for $i \in \{1, 2\}$. From this it follows that Π_* is a homomorphism. \square

Definition 6.28. Suppose that $w \in X_0$ and that α and β are rays in X_0 starting at w , ending at points $z, z' \in \partial X_0$, respectively. Let $T(\alpha, \beta)$ be the tripod obtained by gluing two rays together along an initial segment of length $(z | z')_w$. Let $c(w, \alpha, \beta)$ be the point of α in the preimage of the trivalent point of $T(\alpha, \beta)$ (or the point w if $(z | z')_w = 0$).

According to [25, Lemma 3.7], there is a $(1, 5\delta_0)$ -quasi-isometry from $\alpha \cup \beta$ to $T(\alpha, \beta)$ which is isometric on each of α and β . The utility of the

center $c(w, \alpha, \beta)$ is that it lies at distance $(z | z')_w$ from w , and therefore can be used to estimate the visual distance between z and z' as viewed from w .

We now fix some more constants: Let

$$C_0 = 2 \left(\frac{2 \log \kappa + \log L_0}{\varepsilon} + 29 \delta_0 \right) + 20 \delta_0,$$

and let

$$D_0 = 3\Phi(C_0) + 108\delta_0,$$

where Φ is the function from Lemma 6.11.

Lemma 6.29. *There exists $R_{6.29}$ depending only on δ_0 , λ_0 , and L_0 , so that if $K \geq R_{6.29}$, the following holds. Suppose that $x, y \in \partial X_0 \setminus \Lambda Y_0$ satisfy $d(\Pi_K(x), \Pi_K(y)) \leq 24\delta_0$. There exists a path η in $\partial X_0 \setminus \Lambda Y_0$ connecting x and y , and such that for every $z \in \eta$ we have $d_{X_0}(\Pi_K(x), \Pi_K(z)) \leq C_0$.*

Proof. Recall that ρ_0 is a δ_0 -adapted visual metric on ∂X_0 based at $w_0 \in Y_0$ with constants $\kappa, \varepsilon = \varepsilon(\delta_0) = \frac{1}{6\delta_0}$ (see Definition 3.10). That is, for all $x, y \in \partial X_0$ we have

$$(10) \quad \kappa^{-1} e^{-\varepsilon(x|y)_{w_0}} \leq \rho_0(x, y) \leq \kappa e^{-\varepsilon(x|y)_{w_0}}.$$

We prove the lemma with

$$R_{6.29} = \left(\frac{2 \log \kappa + \log L_0}{\varepsilon} \right) + 40\delta_0 + \lambda_0.$$

Accordingly, assume $K \geq R_{6.29}$. By linear connectivity, there exists a path η in ∂X_0 connecting x and y and such that, for every $z \in \eta$ we have

$$(11) \quad \rho_0(x, z) \leq L_0 \rho_0(x, y).$$

Fix some such z . We will see using the λ_0 -quasi-convexity of Y_0 that $z \notin \Lambda Y_0$. Let α be the geodesic ray defining $\Pi_K(x)$, and let β_y, β_z be the geodesic rays defining $\Pi_K(y), \Pi_K(z)$ respectively.

Let $c_y, c_z \in \alpha$ be the centers $c_y = c(w_0, \alpha, \beta_y)$ and $c_z = c(w_0, \alpha, \beta_z)$ from Definition 6.28. If $d_{X_0}(c_z, Y_0) \geq K$ then it is easy to bound $d_{X_0}(\Pi_K(x), \Pi_K(z))$, so we assume $d_{X_0}(c_z, Y_0) < K$.

By (10) we have (writing $|\cdot|$ for $d_{X_0}(w_0, \cdot)$)

$$\rho_0(x, z) \geq \kappa^{-1} e^{-\varepsilon|c_z|}, \text{ and}$$

$$\rho_0(x, y) \leq \kappa e^{-\varepsilon|c_y|}.$$

Taking logs and applying (11) we deduce

$$-\varepsilon|c_y| + \log \kappa + \log L_0 \geq -\varepsilon|c_z| - \log \kappa.$$

Rearranging gives

$$(12) \quad |c_z| \geq |c_y| - \frac{2 \log \kappa + \log L_0}{\varepsilon}.$$

Next we note that since $d_{X_0}(\Pi_K(x), \Pi_K(y)) \leq 24\delta_0$, the distance between the images of $\Pi_K(x)$ and $\Pi_K(y)$ in the comparison tripod $T(\alpha, \beta_y)$ is at most $29\delta_0$, and so

$$(13) \quad |c_y| \geq |\Pi_K(x)| - 29\delta_0.$$

Combining (12) with (13) gives

$$|c_z| \geq |\Pi_K(x)| - \left(\frac{2\log \kappa + \log L_0}{\varepsilon} + 29\delta_0 \right).$$

Since c_z and $\Pi_K(x)$ both lie on the geodesic ray α , and $d_{X_0}(\Pi_K(x), Y_0) = K$, we have

$$d_{X_0}(c_z, Y_0) \geq K - \left(\frac{2\log \kappa + \log L_0}{\varepsilon} + 29\delta_0 \right) \geq 11\delta_0 + \lambda_0.$$

Let $c'_z = c(w_0, \beta_z, \alpha) \in \beta_z$, so that $d_{X_0}(c_z, c'_z) \leq 5\delta_0$. We have $d_{X_0}(c'_z, Y_0) \geq 6\delta_0 + \lambda_0$. In particular, since β_z contains a point further than λ_0 (the quasi-convexity constant) from Y_0 , we deduce that z is not in the limit set of Y_0 . In particular we may apply Lemma 6.22 to estimate the distance between points on β_z in terms of the distance to Y_0 . We conclude $d_{X_0}(c'_z, \Pi_K(z)) \leq \left(\frac{2\log \kappa + \log L_0}{\varepsilon} + 29\delta_0 + 5\delta_0 \right) + 10\delta_0$, and finally

$$\begin{aligned} d_{X_0}(\Pi_K(x), \Pi_K(z)) &\leq d_{X_0}(\Pi_K(x), c_z) + d_{X_0}(c_z, c'_z) + d_{X_0}(c'_z, \Pi_K(z)) \\ &\leq \left(\frac{2\log \kappa + \log L_0}{\varepsilon} + 29\delta_0 \right) + 5\delta_0 \\ &\quad + \left(\frac{2\log \kappa + \log L_0}{\varepsilon} + 29\delta_0 + 5\delta_0 \right) + 10\delta_0 \\ &= 2 \left(\frac{2\log \kappa + \log L_0}{\varepsilon} + 29\delta_0 \right) + 20\delta_0 = C_0 \end{aligned}$$

as desired. \square

Proposition 6.30. *Let (X_0, Y_0) be as in Assumption 6.1. For any $D \geq D_0$, any $K \geq \max\{R_{6,29}, C_0 + \Delta_0 + \lambda_0 + 3\delta_0\}$, and any points $x, y \in \partial X_0 \setminus \Lambda Y_0$ the following holds: Let α be a path in $\mathcal{S}_K(Y_0)$ joining $\Pi_K(x)$ to $\Pi_K(y)$. Then there is a path η from x to y in $\partial X_0 \setminus \Lambda Y_0$ so that any projection of η to $\mathcal{S}_K(Y_0)$ is homotopic rel endpoints to α in $(\mathcal{S}_K(Y_0))^D$. Moreover such a projection lies in a $D_0/3$ -neighborhood of α .*

Proof. Let α be such a path, and divide it into subpaths α_i from x_i to x_{i+1} of length at most $8\delta_0$, where the points x_i lie in the shell $S_K(Y_0)$. For each i there is a point $p_i \in \partial X_0$ so that $d_{X_0}(\Pi_K(p_i), x_i) \leq 8\delta_0$. (Note that $p_i \notin \Lambda Y_0$.) Inserting geodesic segments back and forth between the points x_i and $\Pi_K(p_i)$ produces a path α' composed of subpaths $\alpha'_i := [\Pi_K(p_i), x_i] \cdot \alpha_i$.

$[x_{i+1}, \Pi_K(p_{i+1})]$. (Being sure to use the same geodesic segment $[x_{i+1}, \Pi_K(p_{i+1})]$ at the end of α'_i and, in reverse, at the beginning of α'_{i+1} .)

By Lemma 6.29 there is, for each i , a path η_i in $\partial X_0 \setminus \Lambda Y_0$ joining p_i to p_{i+1} so that $\Pi_K(\eta_i)$ lies in a C_0 -ball (with respect to d_{X_0}) around $\Pi_K(p_i)$. Since K is large enough to apply Lemma 6.11, we conclude that $d_{\mathcal{S}_K(Y_0)}(\Pi_K(p_i), \Pi_K(z)) \leq \Phi(C_0) < D/6$ for any $z \in \eta_i$. Let η be the path from p to q which is the concatenation of these paths η_i , and let α'' be a projection of η to $\mathcal{S}_K(Y_0)$ which passes through the points $\Pi_K(p_i)$. (Choose the partition to be a refinement of the one which includes the endpoints of the domains of the η_i .) Let α''_i be the subpath of α'' from $\Pi_K(p_i)$ to $\Pi_K(p_{i+1})$. This path is composed of segments $[\Pi_K(z), \Pi_K(z')]$ of length at most $8\delta_0$, so that both endpoints are X_0 -distance at most $8\delta_0$ from either $\Pi_K(p_i)$.

We claim the loop composed of α''_i and α'_i is contractible in $(\mathcal{S}_K(Y_0))^D$. We first note that each α''_i is homotopic to a $\mathcal{S}_K(Y_0)$ -geodesic γ_i joining its endpoints. This is because each projected η_i is composed of geodesic segments of length at most $8\delta_0$, each of whose endpoints can be connected to p_i (the initial point of α''_i) by a path in $\mathcal{S}_K(Y_0)$ of length at most $\Phi(C_0)$. This decomposes the loop consisting of α''_i and γ_i into triangular loops of length at most $2\Phi(C_0) + 8\delta_0 < D$. See Figure 2.




FIGURE 2. Homotoping a projection of η to α .

The paths γ_i and α'_i have length at most $24\delta_0$, so together they bound a loop of length at most $48\delta_0 < D$.

Putting the pieces together we obtain a homotopy rel endpoints from the projection α'' to α . Lemma 6.26 says any two projections are homotopic rel endpoints.

We also see from this picture that α'' is contained in the $\Phi(C_0) + 12\delta_0$ -neighborhood of α . Hence (by Lemma 6.26) any projection of η lies within $\Phi(C_0) + 36\delta_0 = D_0/3$ of α . \square

Corollary 6.31. *Let (X_0, Y_0) be as in Assumption 6.1. Then for $R_{6.31} = \max\{R_{6.29}, C + \Delta_0 + \lambda_0 + 3\delta_0\}$, for all $K \geq R_{6.31}$ and $D \geq D_0$, the following holds:*

For any point $$ in $\partial X_0 \setminus \Lambda Y_0$, the homomorphism $\Pi_* : \pi_1(\partial X_0 \setminus \Lambda Y_0, *) \rightarrow \pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*))$ from Proposition 6.27 is surjective.*

Recall the choice of rays η_p from Definition 6.3, used to define the projection function Π_K .

Lemma 6.32. *In the context of Proposition 6.30, there is a function $R_{6.32}(E, \iota)$ defined for positive E and ι so that for any $K > R_{6.32}(E, \iota)$ and any path α in $\mathcal{S}_K(Y_0)$ of diameter at most E , the path η produced by Proposition 6.30 has ρ_0 -diameter at most ι .*

Proof. Let z, z' be two points in η . There is a projection of η which passes through the points $a = \Pi_K(z)$ and $a' = \Pi_K(z')$, so the points a, a' lie in a $D_0/3$ -neighborhood of α . In particular $d_{X_0}(a, a') < E + D_0$. One can see from this that the center $c(w_0, r(z), r(z'))$ must be at least $K - (E + D_0 + 5\delta_0)$ from the basepoint w_0 . The distance from the basepoint to the center is the Gromov product $(z|z')_{w_0}$ so we deduce

$$\rho_0(z, z') \leq \kappa e^{-\varepsilon(K - (E + D_0 + 5\delta_0))}.$$

Fixing E , we may increase K to make the right-hand side less than ι , as desired. \square

Assumption 6.33. *In the setting of Assumption 6.1, assume further that $\partial X_0 \setminus \Lambda Y_0$ is locally simply connected, and that some $H_0 < \text{Isom}(X_0)$ preserves and acts cocompactly on Y_0 . Choose a constant A_0 so that for all $x \in Y_0$ there exists $h \in H_0$ so that $d_{X_0}(h \cdot x, w_0) \leq A_0$.*

Lemma 6.34. *In the setting of Assumption 6.33, for all $\zeta \in \partial X_0 \setminus \Lambda Y_0$ there exists $h \in H_0$ so that for any $\xi \in \Lambda Y_0$ and any choice of geodesic ray β_ξ from w_0 to ξ we have*

$$d(w_0, c(w_0, r(h \cdot \zeta), \beta_\xi)) \leq A_0 + \lambda_0 + 12\delta_0.$$

In particular, $h \cdot \zeta \in \partial X_0 \setminus \mathring{N}_{r_0}(\Lambda Y_0)$, for $r_0 = \kappa^{-1} e^{-\varepsilon(A_0 + \lambda_0 + 12\delta_0)}$.

Proof. Let $r(\zeta)$ be the geodesic ray from w_0 to ζ chosen in Definition 6.3, and let $z = \Pi_{\lambda_0 + 6\delta_0}(\zeta)$ be the projection from ζ to $S_{\lambda_0 + 6\delta_0}$ as in that definition. Then z lies at distance exactly $\lambda_0 + 6\delta_0$ from Y_0 , so there exists $h \in H_0$ so that $d(w_0, h \cdot z) \leq A_0 + \lambda_0 + 6\delta_0$.

Let $r(h \cdot \zeta)$ and $v = \Pi_{\lambda_0+6\delta_0}(h \cdot \zeta)$ be the analogous geodesic ray from w_0 to $h \cdot \zeta$ and the projection for $h \cdot \zeta$ to $S_{\lambda_0+6\delta_0}$. Lemma 6.5 implies that $d(h \cdot z, v) \leq 6\delta_0$, which implies that $d(w_0, v) \leq A_0 + \lambda_0 + 12\delta_0$. Now, v lies at distance $\lambda + 6\delta_0$ from Y_0 , and every point after v on $r(h \cdot \zeta)$ lies at least that distance from Y_0 .

Let $\xi \in \Lambda Y_0$ be arbitrary, and choose a geodesic ray β_ξ from w_0 to ξ . By the assumption on Y_0 in Assumption 6.1, β_ξ lies entirely in the λ_0 -neighborhood of Y_0 . It follows that the center $c(w_0, r(h \cdot \zeta), \beta_\xi)$ lies closer to w_0 on $r(h \cdot \zeta)$ than v does, which implies that

$$d(w_0, c(w_0, r(h \cdot \zeta), \beta_\xi)) \leq A_0 + \lambda_0 + 12\delta_0,$$

as required. \square

Theorem 6.35. *Let X_0, Y_0, H_0, A_0 be as in Assumption 6.33, and let $D_0 = 3\Phi(C_0) + 108\delta_0$ be as above and $R_{6.31}$ be the constant from Corollary 6.31 applied to the pair (X_0, Y_0) . For all $D \geq D_0$, and for all $\text{sys} > 0$ there exists*

$$R_{6.35} = R_{6.35}(X_0, Y_0, D, \text{sys}, A_0)$$

so that for all $K \geq R_{6.35}$ and all $* \in \partial X_0 \setminus \Lambda Y_0$ the following properties hold:

(1) *the homomorphism*

$$\Pi_*: \pi_1(\partial X_0 \setminus \Lambda Y_0, *) \rightarrow \pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*))$$

from Proposition 6.27 is an isomorphism;

(2) *any homotopically nontrivial loop in the one-skeleton of $(\mathcal{S}_K(Y_0))^D$ has length at least sys ;*

(3) *for any $h \in H_0$, there is a canonical isomorphism*

$$\pi_1^D(\mathcal{S}_K(Y_0, \Pi_K(h*))) \cong \pi_1^D(\mathcal{S}_K(Y_0, h\Pi_K(*)))$$

and moreover there is a commutative diagram as follows:

$$(\dagger) \quad \begin{array}{ccc} \pi_1(\partial X_0 \setminus \Lambda Y_0, *) & \xrightarrow{\Pi_*} & \pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*)) \\ \downarrow & & \downarrow \\ \pi_1(\partial X_0 \setminus \Lambda Y_0, h*) & \xrightarrow{\Pi_{h*}} & \pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(h*)) \end{array}$$

The left vertical arrow is the obvious map induced by the action of h on $\partial X_0 \setminus \Lambda Y_0$ and the right vertical arrow is the concatenation of the obvious map induced by the action of h on $\mathcal{S}_K(Y_0)$ with the above canonical isomorphism.

Proof. We claim first there is a function $\theta : (0, \infty) \rightarrow (0, 1]$, depending only on the pair (X_0, Y_0) and the visual metric ρ_0 on ∂X_0 chosen in Assumption 6.1, so that any loop in $\partial X_0 \setminus \mathring{N}_r(\Lambda Y_0)$ of diameter at most $\theta(r)$ is contractible in $\partial X_0 \setminus \Lambda Y_0$. (The proof of the existence of such a function θ uses the local simple connectivity of $\partial X_0 \setminus \Lambda Y_0$ and a compactness argument.) We may assume that θ is non-decreasing.

Let A_0 be the diameter of a fundamental domain for the action of H_0 on Y_0 as fixed in Assumption 6.33, and let $A = A_0 + \lambda_0 + 12\delta_0$. Fix $D \geq D_0$ and also sys , and let $\iota_0 = \min\{\frac{1}{2}\kappa^{-1}e^{-\varepsilon A}, \theta(\frac{1}{2}\kappa^{-1}e^{-\varepsilon A})\}$.

We claim that

$$R_{6.35} = \max\{R_{6.31}, R_{6.32}(\max\{D, \text{sys}\} + 24\delta_0, \iota_0)\}$$

satisfies the conclusions of the theorem.

We first establish Item (2). Let α be a loop in $\mathcal{S}_K(Y_0)$ of length less than sys . If α is not already contractible in $\mathcal{S}_K(Y_0)$, then it contains a point in $S_K(Y_0)$. Some point of α is at most $8\delta_0$ from some point $\Pi_K(x)$. In particular α is freely homotopic in $\mathcal{S}_K(Y_0)$ to a loop α' based at $\Pi_K(x)$, of length at most $\text{sys} + 16\delta_0$. Apply Lemma 6.34 to the point x to find an element h so that $h \cdot x \in \partial X_0 \setminus \mathring{N}_{r_0}(\Lambda Y_0)$, for $r_0 = \kappa^{-1}e^{-\varepsilon A}$. Let η be the path produced from $h \cdot \alpha'$ by Proposition 6.30.

By the choice of $\iota_0 = \min\{\frac{1}{2}\kappa^{-1}e^{-\varepsilon A}, \theta(\frac{1}{2}\kappa^{-1}e^{-\varepsilon A})\}$, and the fact that $K \geq R_{6.32}(\max\{D, \text{sys}\} + 24\delta_0, \iota_0)$, Lemma 6.32 implies that the ρ_0 -diameter of η is at most $\frac{1}{2}\kappa^{-1}e^{-\varepsilon A}$, so $\eta \subseteq \partial X_0 \setminus \mathring{N}_r(\Lambda Y_0)$, for $r = \frac{1}{2}\kappa^{-1}e^{-\varepsilon A}$. Since the ρ_0 -diameter of η is at most $\theta(\frac{1}{2}\kappa^{-1}e^{-\varepsilon A})$, the choice of the function θ implies that η is contractible in $\partial X_0 \setminus \Lambda Y_0$.

Applying Proposition 6.27 with $* = h \cdot x$ (the basepoint of η) implies that any projection of η to $\mathcal{S}_K(Y_0)$ is contractible in $(\mathcal{S}_K(Y_0))^D$. It follows from Proposition 6.30 that $h \cdot \alpha'$, and hence $h \cdot \alpha$ are contractible in $(\mathcal{S}_K(Y_0))^D$. Since $h \in \text{Isom}(X_0)$ preserves $\mathcal{S}_K(Y_0)$, we see that α' is contractible in $\mathcal{S}_K(Y_0)$, proving Item 2.

We now prove item (1). In view of Corollary 6.31, it is enough to prove that the homomorphism Π_* is injective. To that end, let α be a loop in $\mathcal{S}_K(Y_0)$ based at $\Pi_K(*)$ which represents the trivial element of $\pi_1^D(\mathcal{S}_K(Y_0), \Pi_K(*))$. By applying Proposition 6.30, we may assume that α is the projection to $\mathcal{S}_K(Y_0)$ of a path η in $\partial X_0 \setminus \Lambda Y_0$ based at $*$. Since α represents the identity, it can be filled with a van Kampen diagram whose 2-cells are labelled by loops in $\mathcal{S}_K(Y_0)$ of length at most D . Modifying this diagram, we may assume that all 2-cells are triangles of length at most $D + 24\delta_0$, all of whose corners are in the image of Π_K .

Applying Proposition 6.30 to the 1-cells of this van Kampen diagram Δ , we obtain a map from the 1-skeleton of Δ into $\partial X_0 \setminus \Lambda Y_0$ so that the map

restricted to the boundary of Δ is η , and the boundary of each 2–cell in Δ is mapped to a loop which projects to a loop in $\mathcal{S}_K(Y_0)$ at distance at most $D_0/3$ from the corresponding loop in the original van Kampen diagram. The argument above proving item (2) shows that the boundary of each 2–cell in Δ maps to a contractible loop in $\partial X_0 \setminus \Lambda Y_0$. This proves that η represents the trivial element of $\pi_1(\partial X_0 \setminus \Lambda Y_0, *)$, as required. This proves item (1).

Finally, we prove item 3. Observe that by Lemma 6.5, for any $\zeta \in \partial X_0 \setminus \Lambda Y_0$ and any $h \in H_0$, we have $d_{X_0}(h \cdot \Pi_K(\zeta), \Pi_K(h \cdot \zeta)) \leq 8\delta_0$. Since there is a unique D –homotopy class of path of length at most $8\delta_0$ between $\Pi_K(h \cdot *)$ and $h \cdot \Pi_K(*)$, there is a canonical change of basepoint isomorphism, which is the isomorphism from the statement.

We now establish the commutativity of (\dagger). Let η be a loop based at $*$ in $\partial X_0 \setminus \Lambda Y_0$, and let α be a projection of η to $\mathcal{S}_K(Y_0)$. Let $* = z_0, \dots, z_n = *$ be the points used in defining the projection α . Adding more points if necessary, we may use the points hz_i to define a projection α' of $h\eta$. The diagram

$$\begin{array}{ccc} \partial X_0 \setminus \Lambda Y_0 & \xrightarrow{\Pi_K} & \mathcal{S}_K(Y_0) \\ \downarrow h & & \downarrow h \\ \partial X_0 \setminus \Lambda Y_0 & \xrightarrow{\Pi_K} & \mathcal{S}_K(Y_0) \end{array}$$

commutes up to $8\delta_0$. Hence for each i , there is a path of length at most $8\delta_0$ joining $\Pi_K(hz_i)$ to $h\Pi_K(z_i)$ in $\mathcal{S}_K(Y_0)$. These combine with subsegments of $h\alpha$ and α' to form quadrilaterals of length at most $32\delta_0 \leq D$ determining a D –homotopy from $h\alpha$ to α' . In the diagram (\dagger) the loop α' represents the result of going down and then right; the loop $h\alpha$ represents the result of going right and then down. \square

The following is an immediate consequence of Item 1 of Theorem 6.35, and Lemma 6.14 which implies that the basepoint does not matter.

Corollary 6.36. *Let X_0, Y_0, H_0 , and A_0 be as in Assumption 6.33, and suppose furthermore that Y_0 is a bi-infinite quasi-geodesic, and that $\partial X_0 \cong S^2$. Let $D_0 = 3\Phi(C_0) + 108\delta_0$ be as above, and fix $D \geq D_0$ and sys . Let $R_{6.35} = R_{6.35}(X_0, Y_0, D_0, \text{sys}, A_0)$ be the constant from Theorem 6.35. Then for all $K \geq R_{6.35}$ and all $p \in \mathcal{S}_K(Y_0)$ we have*

$$\pi_1^D(\mathcal{S}_K(Y_0), p) \cong \mathbb{Z}.$$

Remark 6.37. In case $\partial X_0 \cong S^2$ and $\text{Isom}(X_0)$ acts cocompactly on X_0 , it follows from [5, Lemma 2.5] that ∂X_0 , equipped with any visual metric, is *linearly locally contractible*. This means that one can take the function θ

from the proof of Theorem 6.35 to be a linear function. It is then straightforward to prove that the conclusion of Corollary 6.36 remains true (for a (possibly larger) $R_{6.35}$) even without the assumption that Y_0 admits a cocompact action by H_0 , and without the reliance on the diameter of the fundamental domain. In fact, in this setting, the constant $R_{6.35}$ does not depend on the choice of Y_0 , only on its quasi-convexity constant (as well as X_0 , etc.).

Since this argument works only in a setting of considerably smaller generality than Theorem 6.35, we decided not to pursue it, since we believe that Theorem 6.35 should be applicable in many situations beyond the one considered in this paper.

However, we also believe that the uniformity of drilling radius sketched in this remark will be important in other contexts.

6.6. Coarse topology of tube complements. In this section we assume that we are in the setting of Assumption 6.33.

Recall the construction of the completed shell from Definition 6.12, the completed tube complement from Definition 6.15, and also Conventions 6.13 and 6.16.

In this subsection, we investigate the coarse topology of completed shells and completed tube complements.

Throughout this subsection, fix the following notation.

Notation 6.38. Let δ_1 be the hyperbolicity constant of $X_0^{\text{cusp}}(K)$ (for large enough K) from Theorem 6.24, and let $\delta_2 = 1500\delta_1$ (5 times the hyperbolicity constant from the conclusion of the Coarse Cartan-Hadamard 4.7). Without loss of generality, we assume that $\delta_1 \geq \delta_0$, where δ_0 is the hyperbolicity constant for X_0 from Assumption 6.1.

Definition 6.39. Suppose that Z is a metric space, that $Y \subset Z$ is a bi-infinite path, and that $\alpha > 0$. Then (Z, Y) is α -tube comparable to (X_0, Y_0) if there is an isomorphism of graphs from $N_\alpha^Z(Y)$ to $N_\alpha^{X_0}(Y_0)$ which takes Y (isomorphically) onto Y_0 .

Remark 6.40. In our context, all spaces are graphs. There are two natural metrics on $N_\alpha^{X_0}(Y_0)$ that one might consider. The first is the induced metric from d_{X_0} . The second is the induced graph metric on $N_\alpha^{X_0}(Y_0)$, and this is what is meant in Definition 6.39 above. We choose this second meaning since it is easier to verify.

For some of our applications, it is not enough that $N_\alpha^Z(Y)$ and $N_\alpha^{X_0}(Y_0)$ be isomorphic as graphs, and we need their induced metrics to be isometric. However, in our applications, each of these graphs are $2\delta_2$ -quasi-convex (by the obvious variant of Lemma 6.6). In this case, if $N_{\alpha+2\delta_2}^Z(Y)$ and

$N_{\alpha+2\delta_2}^{X_0}(Y_0)$ are isometric with the induced graph metrics then $N_{\alpha}^Z(Y)$ and $N_{\alpha}^{X_0}(Y_0)$ are isometric with their induced metrics from d_Z, d_{X_0} , respectively.

The issue raised in Remark 6.40 is the reason why many of the hypothesis about tube comparable have an extra $2\delta_2$ than might first be expected.

The following is immediate from Remark 6.40, from the definition of the completed shell, and from the isomorphism between shells induced by being $(K+4\delta_0)$ -tube comparable. We need $K+4\delta_0$ so that any path of length $8\delta_0$ between points on the K -shell is included in the tube comparison.

Lemma 6.41. *Suppose that (Z, Y) is $(K+4\delta_0)$ -tube comparable to (X_0, Y_0) . Then $\mathcal{S}_K^Z(Y)$ is isometric to $\mathcal{S}_K^{X_0}(Y_0)$.*

The proof of the next result is identical to that of Lemma 6.17, except that Lemma 6.41 above is used as well as Lemma 6.14.

Lemma 6.42. *Suppose that $K \geq \lambda_0 + 3\delta_0$, that Γ is a connected graph, and that (Γ, Y) is $(K+4\delta_0)$ -tube comparable to (X_0, Y_0) . Then $\mathcal{C}_K^{\Gamma}(Y)$ is a connected graph.*

Lemma 6.43. *Suppose that Υ is δ_2 -hyperbolic. If $K \geq \lambda_0 + \delta_2$ and (Υ, Y) is $(K+2\delta_2)$ -tube comparable to (X_0, Y_0) then Y is a λ_0 -quasi-convex subset in Υ .*

Proof. Suppose that σ is a geodesic with endpoints p, q on Y . By Remark 6.40, there is an isometry $\phi : N_K(Y_0) \rightarrow N_K(Y)$ taking Y_0 onto Y . Let σ' be a geodesic in X between $\phi^{-1}(p)$ and $\phi^{-1}(q)$. Since Y_0 is λ_0 -quasi-convex in X_0 , and $K > \lambda_0$, the geodesic σ' is in the domain of ϕ , and its image $\phi(\sigma')$ is an Υ -geodesic joining p to q , and lying in the λ_0 -neighborhood of Y . Since Υ is δ_2 -hyperbolic, σ lies in the $\lambda_0 + \delta_2$ -neighborhood of Y . Since $K \geq \lambda_0 + \delta_2$, the preimage $\phi^{-1}(\sigma)$ is a geodesic joining $\phi^{-1}(p)$ to $\phi^{-1}(q)$. Since Y_0 is λ_0 -quasi-convex, $\phi^{-1}(\sigma)$ actually lies in the λ_0 -neighborhood of Y_0 , and so σ lies in the λ_0 -neighborhood of Y , as required. \square

Recall the constant Δ_0 from Lemma 6.10 and the function Φ from Lemma 6.11. The following is an analogue of Lemma 6.11, and follows immediately from that result and Lemma 6.41 (and the definition of $\mathcal{S}_K^{\Upsilon}(Y)$, Definition 6.12).

Lemma 6.44. *Suppose that Υ is a δ_2 -hyperbolic space. Suppose that $C > 0$, that $K \geq C + \Delta_0 + \lambda_0 + 3\delta_0$, and that (Z, Y) is $(K+2\delta_2)$ -tube comparable to (X_0, Y_0) . For any $z_1, z_2 \in \mathcal{S}_K^Z(Y)$ so that $d_Z(z_1, z_2) \leq C$ we have $d_{\mathcal{S}_K^Z(Y)}(z_1, z_2) \leq \Phi(C)$.*

Recall the definition of Q -deformation retraction from Definition 4.12.

Lemma 6.45. *Suppose that Υ is a δ -hyperbolic graph, that $\Xi \subset \Upsilon$ is a λ -quasi-convex sub-graph, and that $K \geq 4\delta + \lambda$. There exists a $(2\delta + 1)$ -deformation retraction $\{r_i: \Upsilon \rightarrow \Upsilon\}$ of Υ onto $N_K(\Xi)$ so that for all $i \geq 1$ and all $z \in \Upsilon \setminus N_K(\Xi)$ we have $r_i(z) \in \Upsilon^{(0)}$.*

Moreover if $z \notin N_K(\Xi)$ then $d_\Upsilon(r_i(z), \Xi) \geq K$ for all i .

Proof. For $z \in \Upsilon$, let $P_\Xi(z)$ be a closest point of Ξ to z , and choose a unit speed geodesic $\alpha_z: [-d_\Upsilon(z, \Xi), 0] \rightarrow X$ joining z to $P_\Xi(z)$. (The domain is chosen so that whether or not z is a vertex, α_z sends integers to vertices.) Define $f_z(i) = \min\{-K, \lfloor i - d_\Upsilon(z, \Xi) \rfloor\}$, and $r_i(z) = \alpha_z(f_z(i))$. Thus $r_1(z)$ is the first vertex on α_z after z , $r_2(z)$ is the next, and so on, until for large i we have $r_i(z)$ a closest point to z in $N_K(\Xi)$.

Conditions (1), (2) and (4) from Definition 4.12 obviously hold. Condition (5) holds because every point in Υ is finite distance from Ξ . The condition that $r_i(z) \in \Upsilon^{(0)}$ for $i \geq 1$ and $z \in \Upsilon \setminus N_K(\Xi)$ is clear from the construction, as is the ‘‘Moreover’’ statement.

It remains to verify Condition (3) from Definition 4.12. To that end, suppose that $z_1, z_2 \in \Upsilon$ satisfy $d_\Upsilon(z_1, z_2) \leq 1$, and consider the geodesic quadrilateral with two sides α_{z_1} and α_{z_2} , a geodesic between z_1 and z_2 , and a geodesic between the projections $P_\Xi(z_1)$ and $P_\Xi(z_2)$. Since $K \geq \lambda + 4\delta$, it is straightforward to see from the fact that geodesic quadrilaterals are 2δ -thin, that for all i we have $d_\Upsilon(f_i(z_1), f_i(z_2)) \leq 2\delta + 1$, as required. \square

Recall the construction of the completed tube complement

$$\mathcal{C}_K^\Upsilon(Y) = \left(\Upsilon \setminus T_K^\Upsilon(Y) \right) \sqcup_{S_K^\Upsilon(Y)} \mathcal{S}_K^\Upsilon(Y)$$

from Definition 6.15, and let

$$\iota: \Upsilon \setminus T_K^\Upsilon(Y) \rightarrow \mathcal{C}_K^\Upsilon(Y)$$

denote the inclusion map.

Lemma 6.46. *Let $C > 0$ and suppose $K \geq C + \Delta_0 + \lambda_0 + 3\delta_0$. Suppose also that Υ is a δ_2 -hyperbolic metric space and that (Υ, Y) is $(K + \delta_2)$ -tube comparable to (X_0, Y_0) . For any $z_1, z_2 \in \Upsilon \setminus T_K^\Upsilon(Y)$ so that $d_\Upsilon(z_1, z_2) \leq C$ we have*

$$d_{\mathcal{C}_K^\Upsilon(Y)}(\iota(z_1), \iota(z_2)) \leq 2\Phi(C).$$

Proof. Let γ be a geodesic in Υ between z_1 and z_2 . If γ does not meet $S_K^\Upsilon(Y)$ there is nothing to show, since $\Phi(C) \geq C$.

Suppose that z'_1 is the first point (from z_1) on γ that intersects $S_K^\Upsilon(Y)$ and that z'_2 is the last point (closest to z_2 along γ) that intersects $S_K^\Upsilon(Y)$. Those parts of γ that lie between z_1 and z'_1 and between z'_2 and z_2 can be considered as paths in $\mathcal{C}_K^\Upsilon(Y)$. On the other hand, z'_1 and z'_2 are points in

$\mathcal{C}_K^\Upsilon(Y)$ which satisfy $d_\Upsilon(z'_1, z'_2) \leq d_\Upsilon(z_1, z_2) \leq C$, so by Lemma 6.44 we have $d_{\mathcal{C}_K^\Upsilon(Y)}(z'_1, z'_2) \leq \Phi(C)$. It follows that $d_{\mathcal{C}_K^\Upsilon(Y)}(z_1, z_2) \leq C + \Phi(C) \leq 2\Phi(C)$, as required. \square

We now apply Lemma 6.45 to a slightly different situation.

Lemma 6.47. *Suppose that $K \geq 5\delta_2 + \Delta_0 + \lambda_0 + 1$. Suppose further that Υ is a δ_2 -hyperbolic space and that (Υ, Y) is $(K + 2\delta_2)$ -tube comparable to (X_0, Y_0) . There exists a $2\Phi(2\delta_2 + 1)$ -deformation retraction*

$$\left\{ f_i: \mathcal{C}_K^\Upsilon(Y) \rightarrow \mathcal{S}_K^\Upsilon(Y) \right\}$$

of $\mathcal{C}_K^\Upsilon(Y)$ onto $\mathcal{S}_K^\Upsilon(Y)$.

Furthermore, if f is the stable map of the sequence $\{f_i\}$ then for any $z_1, z_2 \in \Upsilon$ we have

$$d_{\mathcal{S}_K^\Upsilon(Y)}(f(z_1), f(z_2)) \leq \Phi(2\delta_2 + 1).$$

Proof. Recall that

$$\mathcal{C}_K^\Upsilon(Y) = \left(\Upsilon \setminus T_K^\Upsilon(Y) \right) \sqcup_{\mathcal{S}_K^\Upsilon(Y)} \mathcal{S}_K^\Upsilon(Y).$$

For $i \geq 0$ we define maps $f_i: \mathcal{C}_K^\Upsilon(Y) \rightarrow \mathcal{C}_K^\Upsilon(Y)$ satisfying the conditions for Definition 4.12. If $y \in \mathcal{S}_K^\Upsilon(Y)$ then define $f_i(y) = y$ for all i , as required by Definition 4.12.(2).

On the other hand, suppose that $y \notin \mathcal{S}_K^\Upsilon(Y)$. Then $y \in \Upsilon \setminus T_K^\Upsilon(Y)$. Let $\{r_i: \Upsilon \rightarrow \Upsilon\}$ be the $(2\delta_2 + 1)$ -deformation retraction from Υ to $N_K(Y)$ afforded by Lemma 6.45. By the ‘‘Moreover’’ statement in that lemma, since $y \notin T_K^\Upsilon(Y)$ in Υ , for all i we have $r_i(y) \notin T_K^\Upsilon(Y)$. Therefore, since $r_i(y) \in \Upsilon \setminus T_K^\Upsilon(Y)$, we can define $f_i(y) = r_i(y) \in \mathcal{C}_K^\Upsilon(Y)$.

This defines the maps $\{f_i: \mathcal{C}_K^\Upsilon(Y) \rightarrow \mathcal{C}_K^\Upsilon(Y)\}$. We claim that this sequence of maps is a $2\Phi(2\delta_2 + 1)$ -deformation retraction of $\mathcal{C}_K^\Upsilon(Y)$ onto $\mathcal{S}_K^\Upsilon(Y)$, as required. Indeed, Conditions (1), (2), (4), and (5) all clearly hold by definition and Lemma 6.45. The only property left from Definition 4.12 is Condition 3, which follows quickly from Lemmas 6.46 and 6.45.

Finally, suppose that f is the stable map of the sequence $\{f_i\}$ and suppose that z_1, z_2 in Υ . The construction of $f(z_1)$ and $f(z_2)$ follow by applying the construction from Lemma 6.45, so we have $d_\Upsilon(f(z_1), f(z_2)) \leq 2\delta_2 + 1$ by that lemma. The definition of the function Φ from Lemma 6.11 now implies that

$$d_{\mathcal{S}_K^\Upsilon(Y)}(f(z_1), f(z_2)) \leq \Phi(2\delta_2 + 1),$$

as required. \square

Recall the constant A_0 was fixed in Assumption 6.33.

Proposition 6.48. *Let $Q_0 = \max\{2\Phi(2\delta_2 + 1), \delta_2\}$, let $D_0 = 3\Phi(C_0) + 108\delta_0$ be the constant from Subsection 6.5, let $D \geq \max\{D_0, 2Q_0 + 2\}$, let $\text{sys} = Q_0D_0$, and let $R_{6.35} = R_{6.35}(X_0, Y_0, D, \text{sys}, A_0)$ be the constant from that theorem applied with these constants. Let $K \geq \max\{R_{6.35}, 5\delta_2 + \Delta_0 + \lambda_0 + 1\}$ and let Υ be a connected δ_2 -hyperbolic graph with path γ so that (Υ, Y) is $(K + 2\delta_2)$ -tube comparable to (X_0, Y_0) . The inclusion*

$$\mathcal{S}_K^\Upsilon(Y) \hookrightarrow \mathcal{C}_K^\Upsilon(Y)$$

induces an isomorphism

$$\pi_1^D(\mathcal{S}_K^\Upsilon(Y)) \cong \pi_1^D(\mathcal{C}_K^\Upsilon(Y)).$$

Proof. According to Lemma 6.47 there exists a Q_0 -deformation retraction $f_i: \mathcal{C}_K^\Upsilon(Y) \rightarrow \mathcal{S}_K^\Upsilon(Y)$. Let $f: \mathcal{C}_K^\Upsilon(Y) \rightarrow \mathcal{S}_K^\Upsilon(Y)$ be the associated stable map. For any $b_1, b_2 \in \mathcal{C}_K^\Upsilon(Y)$ so that $d_{\mathcal{C}_K^\Upsilon(Y)}(b_1, b_2) \leq 1$ we have $d_{\mathcal{S}_K^\Upsilon(Y)}(f(b_1), f(b_2)) \leq Q_0$. By the choice of $R_{6.35}$ and Theorem 6.35.(2) all $\mathcal{S}_K^\Upsilon(Y)$ -loops of length at most Q_0D represent the (conjugacy class of the) identity element of $\pi_1(\mathcal{S}_K^\Upsilon(Y))$. The result now follows immediately from Proposition 4.13. \square

The space X_0 and the quasi-convex subset Y_0 satisfy the assumptions of Proposition 6.48. Therefore, we immediately have the following.

Corollary 6.49. *Let $Q_0 = \max\{2\Phi(2\delta_2 + 1), \delta_2\}$. Let $D_0 = 3\Phi(C_0) + 108\delta_0$ be the constant from Subsection 6.5, let $D \geq \max\{D_0, 2Q_0 + 2\}$, let $\text{sys} = QD_0$ and choose $R_{6.35} = R_{6.35}(X_0, Y_0, D, \text{sys}, A_0)$ to be the constant from Theorem 6.35. For any $K \geq \max\{R_{6.35}, 5\delta_2 + \Delta_0 + \lambda_0 + 1\}$, the inclusion induces an isomorphism*

$$\pi_1^D(\mathcal{S}_K(Y_0)) \cong \pi_1^D(\mathcal{C}_K(Y_0)).$$

7. UNWRAPPING THE COMPLEMENT OF A TUBE

In this section we start with the axis γ_0 of the hyperbolic isometry g of the hyperbolic space X_0 with boundary S^2 , and we study the space constructed (approximately) as follows: remove a large tube around γ_0 , take the universal cover of the resulting space, and glue a horoball on the lift of the boundary of the tube. We show that this space is hyperbolic (Lemma 7.12) and has boundary S^2 (Proposition 10.1), allowing us to repeat the procedure. However, we need careful control on all the constants involved to make sure that the procedure can be repeated, so we now fix once and for all a set of constants that allows us to do so.

Assumption 7.1. *For this section, and the remainder of the paper, assume we are in the setting of Assumption 6.1. Assume further that $\partial X_0 \cong S^2$, and*

that $Y_0 = \gamma_0$ is a (bi-infinite) continuous quasi-geodesic in X_0 which is λ_0 -quasi-convex, and let g be a loxodromic element g for which γ_0 is an axis, and let A_0 be a constant satisfying Assumption 6.33 (for the $\langle g \rangle$ -action on γ_0).

In particular, X_0 is δ_0 -hyperbolic and has boundary homeomorphic to S^2 .

We now fix a collection of constants needed later in the construction and the proof of Theorem A. Further constants are chosen at the beginning of Section 8.

Notation 7.2.

- (1) $s_0 = 8\delta_0$, the scale of coarse path connectedness of shells;
- (2) δ_1 is the hyperbolicity and visibility constant of $X_0^{\text{cusp}}(K)$ (see Definition 6.23), for sufficiently large values of K . We assume $\delta_1 \geq \max\{\delta_0, 100\}$;
- (3) $\delta_2 = 1500\delta_1$ (5 times the hyperbolicity constant from the conclusion of the Coarse Cartan-Hadamard Theorem 4.7);
- (4) L_0 is the constant of linear connectedness of ∂X_0 , from Assumption 6.1;
- (5) $Q_0 = \max\{2\Phi(2\delta_2 + 1), \delta_2\}$ is the constant from Lemma 6.47, where Φ is the function from Lemma 6.11;
- (6) $D_0 = 3\Phi(C_0) + 108\delta_0$ is the constant from Subsection 6.5 (just before the statement of Lemma 6.29);
- (7) $D_1 = \max\{D_0, 2Q_0 + 2, 16\delta_0 + 16\delta_0\Phi(16\delta_0)\}$, the scale of loops for defining the coarse fundamental group;
- (8) $\sigma_0 = \max\{10^7\delta_1, 10^5D_1\}$ is the size of balls needed for the Coarse Cartan-Hadamard to apply;
- (9) $\text{sys}_0 = 2^{25\sigma_0}Q_0$ is, because of Theorem 6.35.(2) and the choice of R_0 in the item below, a lower bound for the length of a path in the completed shell $\mathcal{S}_{R_0}(\gamma_0)$ representing a nontrivial conjugacy class in $\pi_1^{D_1}(\mathcal{S}_{R_0}(\gamma_0))$. This value of sys_0 is chosen to make the proof of Lemma 7.9 below work;
- (10) $R_0 = \max\{R_{6.35}, \lambda_0 + 2\delta_2, 6\sigma_0\}$ where $R_{6.35} = R_{6.35}(X_0, \gamma_0, D_1, \text{sys}_0, A_0)$ is the constant from Theorem 6.35. This is the radius of the tubes and shells that we henceforth consider.

Notation 7.3. Having fixed the constants above, for the remainder of the paper, we always take completed R_0 -shells, and completed R_0 -tube complements. Therefore, we henceforth drop the subscript R_0 from our notation.

Thus, for a space Υ , and subset γ , we write

$$\mathcal{S}_\Upsilon(\gamma) = \mathcal{S}_{R_0}^\Upsilon(\gamma), \text{ the completed shell}$$

and

$$\mathcal{C}_\Upsilon(\gamma) = \mathcal{C}_{R_0}^\Upsilon(\gamma), \text{ the completed tube complement.}$$

We are particularly interested in the completed shells and tube complements associated to γ_0 in X_0 , and for these, we further drop all of the decoration, and write

$$\mathcal{S} = \mathcal{S}_{R_0}^{X_0}(\gamma_0)$$

and

$$\mathcal{C} = \mathcal{C}_{R_0}^{X_0}(\gamma_0).$$

(The definitions of completed shell and completed tube complement are 6.12 and 6.15, and the definition of *tube comparable* is Definition 6.39.) The next result summarizes various results from the previous section.

Proposition 7.4. *Let X_0 be the space from Assumption 7.1. Suppose that Υ is δ_2 -hyperbolic and that (Υ, γ) is $(R_0 + 2\delta_2)$ -tube comparable with (X_0, γ_0) . With the choice of constants as in Notation 7.2, the following hold:*

- (1) $\pi_1^{D_1}(\mathcal{C}) \cong \pi_1^{D_1}(\mathcal{S}) \cong \mathbb{Z}$;
- (2) For any loop α in \mathcal{S} , if $[\alpha] \neq 1$ in $\pi_1^{D_1}(\mathcal{S})$ then the length of α is at least sys_0 ;
- (3) For any loop α' in $\mathcal{C}_\Upsilon(\gamma)$, if $[\alpha'] \neq 1$ in $\pi_1^{D_1}(\mathcal{C}_\Upsilon(\gamma))$ then the length of α' is at least sys_0/Q_0 ;
- (4) There are isomorphisms

$$\pi_1^{D_1}(\mathcal{C}_\Upsilon(\gamma)) \cong \pi_1^{D_1}(\mathcal{S}_\Upsilon(\gamma)) = \pi_1^{D_1}(\mathcal{S}) \cong \mathbb{Z}.$$

The first isomorphism is induced by inclusion.

Proof. The first isomorphism in (1) follows from Corollary 6.49 and the second from Corollary 6.36.

Item (2) is Theorem 6.35.(2). For Item (3), suppose that α is an edge-loop in \mathcal{C}_Υ of length less than sys_0/Q_0 . Lemma 6.47 (include the ‘‘Furthermore’’ statement) imply that the hypotheses of Proposition 4.13 hold (Item (2) from the statement holds by the assumption on sys_0). Therefore, by the ‘‘Moreover’’ part of the conclusion of Proposition 4.13, there is a loop $f(\alpha)$ on \mathcal{S}_Υ of length less than sys_0 which represents the same conjugacy class in $\pi_1^{D_1}(\mathcal{C}_\Upsilon)$. Since $\mathcal{S}_\Upsilon = \mathcal{S}$, we can interpret $f(\alpha)$ as a loop on \mathcal{S} , and by Item (2), $f(\alpha)$ represents the identity element of $\pi_1^{D_1}(\mathcal{S}) = \pi_1^{D_1}(\mathcal{S}_\Upsilon) \cong \pi_1^{D_1}(\mathcal{C}_\Upsilon)$. Therefore, α also represents the trivial element, proving Item (3). Item (4) follows immediately from Proposition 6.48 Item 1 and the fact that $\mathcal{S}_\Upsilon = \mathcal{S}$. \square

7.1. Unwrapping. In the rest of this section we study the spaces described below:

Notation 7.5. Let \mathcal{H} be the combinatorial horoball $\mathcal{H}(\widetilde{\mathcal{S}})$ based on the D_1 -universal cover $\widetilde{\mathcal{S}}$ of \mathcal{S} . Let $\mathcal{H}^\circ = \mathcal{H} \setminus \widetilde{\mathcal{S}}$ be the set of points in \mathcal{H} at positive depth. By $\widetilde{\mathcal{C}}$ we denote the D_1 -universal cover of \mathcal{C} .

By Proposition 7.4.(1), the D_1 -universal covers $\widetilde{\mathcal{S}} \rightarrow \mathcal{S}$, and $\widetilde{\mathcal{C}} \rightarrow \mathcal{C}$ are \mathbb{Z} -covers.

The following is an analogue of Definition 6.23. The construction is well-defined since by Proposition 7.4.(4) the preimage of the completed shell is its D_1 -universal cover.

Definition 7.6. Suppose that Υ is δ_2 -hyperbolic and that (Υ, γ) is $(R_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) . The space Υ^{cusp} is obtained by gluing a combinatorial horoball \mathcal{D} based on $\mathcal{S}^\Upsilon(\gamma) \cong \mathcal{S}$ to $\mathcal{C}^\Upsilon(\gamma)$.

In particular, X_0^{cusp} denotes $X_0^{\text{cusp}}(R_0)$, as defined in Definition 6.23.

Construction 7.7. (*Unwrap and glue for one tube*) Suppose that Υ is δ_2 -hyperbolic, and that (Υ, γ) is $(R_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) .

Let $\text{UG}(\Upsilon, \gamma)$ be obtained by taking the D_1 -universal cover of $\mathcal{C}^\Upsilon(\gamma)$ and gluing a copy of \mathcal{H} along the preimage of $\mathcal{S}^\Upsilon(\gamma)$.

The covering map

$$q_{R_0}: \widetilde{\mathcal{C}^\Upsilon(\gamma)} \rightarrow \mathcal{C}^\Upsilon(\gamma)$$

extends to a continuous depth-preserving map

$$q: \text{UG}(\Upsilon, \gamma) \rightarrow \Upsilon^{\text{cusp}}$$

Some edges in \mathcal{H} factor to loops under the natural quotient map, and these loops are not in the horoball based on $\mathcal{S}^\Upsilon(\gamma)$. These loops are mapped to their endpoint.

The map q turns out to be an isometry on balls of radius σ_0 with center not too deep in \mathcal{H} . We show this in Lemma 7.9 below, after stating and proving the following general lemma on, roughly, promoting local homeomorphisms to local isometries.

Lemma 7.8. Let Γ, Ξ be metric graphs, and let $B_R(p)$ be a ball of radius R in Γ for some integer $R > 3$. Suppose that we have an injective simplicial map $f: B_R(p) \rightarrow \Xi$ which is a local homeomorphism at every point of $B_{R-1}(p)$. Then f restricts to an isometry $f': B_{R/3}(p) \rightarrow B_{R/3}(f(p))$.

Proof. Note that the restriction f' of f to $B_{R/3}(p)$ is 1-Lipschitz, so that its image is indeed contained in $B_{R/3}(f(p))$. Therefore it suffices to construct a 1-Lipschitz map $g: B_{R/3}(f(p)) \rightarrow B_{R/3}(p)$ such that $g \circ f' = \text{id}$.

To construct the map, for $x \in B_{R/3}(f(p))$ we choose a geodesic from $f(p)$ to x , we lift it starting at p , and we define $g(x)$ to be the endpoint of the lift. Note that indeed $g(x) \in B_{R/3}(p)$. By construction we have $g(f'(x)) \in f^{-1}(f'(x))$ for each $x \in B_{R/3}(p)$, and since f is injective we have $f^{-1}(f'(x)) = \{x\}$. That is, $g(f'(x)) = f'(x)$, and $g \circ f' = id$. We are left to show that g is 1-Lipschitz. This is because any geodesic connecting points in $B_{R/3}(f(p))$ is contained in $B_{2R/3}(f(p))$, and hence can be lifted to a path in $B_{2R/3}(p)$. (Note that $R - 1 > 2R/3$ since $R > 3$.) \square

Lemma 7.9. *Suppose that Υ is δ_2 -hyperbolic, and that (Υ, γ) is $(R_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) . For any point $z \in \text{UG}(\Upsilon, \gamma)$ that either does not lie in \mathcal{H} or lies at depth at most $3\sigma_0$ in \mathcal{H} the map q as above is an isometric embedding when restricted to the ball $B(z, 3\sigma_0)$, and moreover*

$$q(B(z, 3\sigma_0)) = B(q(z), 3\sigma_0).$$

Proof. Let $z \in \text{UG}(\Upsilon, \gamma)$ be as in the statement of the lemma. We claim that the map q is injective on $B(z, 9\sigma_0)$. This and Lemma 7.8 immediately implies the desired conclusion, since q is then a covering on such balls.

In order to obtain a contradiction, suppose $z_1, z_2 \in B(z, 9\sigma_0)$ are distinct but satisfy $q(z_1) = q(z_2)$. Note that $d_{\text{UG}(\Upsilon, \gamma)}(z_1, z_2) \leq 18\sigma_0$. Since z has depth at most $3\sigma_0$, each of z_1 and z_2 have depth at most $12\sigma_0$. If the depth of z_1 or z_2 is positive, we replace it with a point on $\widetilde{\mathcal{F}}$ directly above it (this corresponds to replacing a point (x, n) with a point $(x, 0)$ in Definition 4.8). We may therefore suppose, at the cost of assuming that $d_{\text{UG}(\Upsilon, \gamma)}(z_1, z_2) \leq 42\sigma_0$, that z_1 and z_2 do not lie in the open horoball.

The points z_1 and z_2 may be joined by a path in $\widetilde{\mathcal{C}}_\Upsilon$ whose image represents a nontrivial element of $\pi_1^{D_1}(\mathcal{C}_\Upsilon)$. Therefore, by Proposition 7.4.(3), we have $d_{\widetilde{\mathcal{C}}_\Upsilon}(z_1, z_2) \geq \text{sys}_0 / Q_0 = 2^{25\sigma_0}$, by the choice of sys_0 in Notation 7.2.

On the other hand, $\text{UG}(\Upsilon, \gamma)$ is built from $\widetilde{\mathcal{C}}_\Upsilon$ by gluing on a copy of a horoball. Any path σ in $\text{UG}(\Upsilon, \gamma)$ may be pushed out the horoball at the cost of increasing the length by an exponential function. Precisely, we have (using Lemma 4.9 and the fact that any subsegment that is pushed out of the horoball has length at least 4):

$$d_{\widetilde{\mathcal{C}}_\Upsilon}(z_1, z_2) \leq 2 \cdot 2^{\frac{1}{2}d_{\text{UG}(\Upsilon, \gamma)}(z_1, z_2)} \leq 2^{21\sigma_0+1} < 2^{25\sigma_0}.$$

This is a contradiction. \square

From Lemma 7.9, we control the shape of balls in $\text{UG}(\Upsilon, \gamma)$ of reasonable size:

Corollary 7.10. *Let σ_0 and R_0 be as fixed in Notation 7.2. Suppose that Y is δ_2 -hyperbolic, and that (Y, γ) is $(R_0 + 6\sigma_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) .*

Any ball of radius σ_0 in $\text{UG}(Y, \gamma)$ is isometric to one of

- a ball in \mathcal{H}° ,
- a ball in Y , or
- a ball in X_0^{cusp} .

Proof. Consider a ball B as in the statement, and let B^+ be the ball with the same center and radius $3\sigma_0$. By Lemma 7.9 either B^+ is entirely contained in \mathcal{H}° (in which case Lemma 7.8 yields the required isometry), or q maps the ball B^+ isometrically to a ball in Y^{cusp} , whose center lies within $3\sigma_0$ of $\mathcal{C}^Y(\gamma)$. The second case has two sub-cases. If the ball $q(B^+)$ is disjoint from the horoball of Y^{cusp} , then $q(B^+)$ is isomorphic to a ball in Y , and as B is isometric to a ball in Y by Lemma 7.8. Otherwise, notice that the $6\sigma_0$ -neighborhood of the horoball of Y^{cusp} is isomorphic to the corresponding neighborhood in X_0^{cusp} , since (Y, γ) is $(R_0 + 6\sigma_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) . We can then again use Lemma 7.8 to conclude. \square

We now control the local geometry of an unwrapped and glued space.

Definition 7.11. Suppose that $\mathcal{A} = (Z_1, \dots, Z_n)$ is a finite collection of metric spaces and that $k > 0$. We say that a metric space Z is k -modeled on \mathcal{A} if for every $z \in Z$ there is an i so that the closed ball of radius k about z is isometric to a closed ball of radius k in Z_i .

Lemma 7.12. *If Y is connected, D_1 -simply-connected, and σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$ then Y is δ_2 -hyperbolic and δ_2 -visible.*

Proof. Hyperbolicity, with constant $\delta' = 300\delta_1$ is an immediate consequence of the Coarse Cartan–Hadamard Theorem 4.7, along with the choices of constants and the fact that both X_0^{cusp} and \mathcal{H}° are δ_1 -hyperbolic.

For visibility with constant $5\delta' = \delta_2$, by Proposition 3.5 it suffices to prove that if $p, q \in Y$ satisfy $d_Y(p, q) \leq 100\delta'$ there exists $q' \in Y$ with $d(p, q') \geq 200\delta'$ and a geodesic $[p, q']$ that passes within δ' of q .

Note that the closed $200\delta'$ -ball around p is compact, and therefore (since $\sigma_0 \geq 200\delta' = 60000\delta_1$) it is isometric to either a ball in X_0^{cusp} , or to a ball of \mathcal{H}° whose center is deeper than its radius.¹ In either case, we can find the required q' . Indeed, for \mathcal{H}° we can use the explicit description of geodesics (cf. [25, Lemma 2.23]), and for X_0^{cusp} we just apply δ_1 -visibility. \square

¹this is the reason why we model on \mathcal{H}° rather than \mathcal{H} .

Proposition 7.13. *Let Υ be a connected D_1 -simply connected graph, and let γ be a bi-infinite path in Υ . Assume that Υ is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$. Suppose further that (Υ, γ) is $(R_0 + 6\sigma_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) . Then $\text{UG}(\Upsilon, \gamma)$ is connected, D_1 -simply connected, and σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$. In particular, $\text{UG}(\Upsilon, \gamma)$ is δ_2 -hyperbolic and δ_2 -visible.*

Proof. Proposition 4.11 implies that $\text{UG}(\Upsilon, \gamma)$ is connected and D_1 -simply connected.

The σ_0 -modeled statement follows directly from Corollary 7.10 and the assumption that Υ is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$. The last assertion of the proposition now follows immediately from Lemma 7.12. \square

Definition 7.14. Let $\chi > 10R_0 + 60\sigma_0 + 20\delta_2$ be a constant. Suppose that Υ is a graph which is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$. A collection of subsets \mathcal{A} in Υ is a χ -separated collection of tubes and horoballs if the following hold:

- (1) Each component of \mathcal{A} is either isometric to $T_{R_0}(\gamma_0)$ ($= T_{R_0}^{X_0}(\gamma_0)$) and then called a *tube*, or to \mathcal{H}° , and called a *horoball*.
- (2) Each tube $A \in \mathcal{A}$ contains a specified path $\gamma(A)$ that we call the *core*, and the pair $(\Upsilon, \gamma(A))$ is $\frac{\chi}{10}$ -tube comparable to (X_0, γ_0) .
- (3) For each horoball $A \in \mathcal{A}$ the $\frac{\chi}{10}$ -neighborhood of A in Υ is isomorphic to the $\frac{\chi}{10}$ -neighborhood of the horoball in $\text{UG}(X_0, \gamma_0)$.
- (4) Any two components of \mathcal{A} are at distance at least χ apart.

Remark 7.15. Assuming that Υ is δ_2 -hyperbolic, each tube or horoball in a collection as described above is $2\delta_2$ -quasiconvex. In particular the metric on such a tube or horoball is determined by the isomorphism type of a $2\delta_2$ -neighborhood.

In Section 8 we need the following construction.

Construction 7.16. *Suppose that Υ is a graph which is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$ and that \mathcal{A} is a χ -separated collection of tubes and horoballs in Υ (for some $\chi > 10R_0 + 20\delta_2$).*

Let A be the union of all the cores of tubes of \mathcal{A} . Let Υ' be obtained from Υ by removing all open horoballs of \mathcal{A} . The hollowed-out space associated to Υ is the space

$$\text{HO}(\Upsilon, \mathcal{A}) = \mathcal{C}^{\Upsilon'}(A).$$

Lemma 7.17. *Let $\tau > R_0$ and suppose that Υ is a graph which is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$, and that γ is a path in Υ so that (Υ, γ) is τ -tube comparable to (X_0, γ_0) . Let $\text{Ann}_\tau = N_\tau^\Upsilon(\gamma) \setminus T_{R_0}^\Upsilon(\gamma)$.*

The preimage of Ann_τ in $\text{UG}(\Upsilon, \gamma)$ is isomorphic to the preimage of $N_\tau(\gamma_0) \setminus T_{R_0}(\gamma_0)$ in $\text{UG}(X_0, \gamma_0)$.

Proof. The deformation retraction $\mathcal{C}^\Upsilon(\gamma) \rightarrow \mathcal{S}^\Upsilon(\gamma)$ from Lemma 6.47 restricts to a deformation retraction $\text{Ann}_\tau \rightarrow \mathcal{S}^\Upsilon(\gamma)$, so the inclusion $\mathcal{S}^\Upsilon(\gamma) \hookrightarrow \text{Ann}_\tau$ induces an isomorphism on π_1^D by Proposition 4.13. Similarly, there is a deformation retraction from $\mathcal{C}^\Upsilon(\gamma)$ onto Ann_τ (built just as in the proof of Lemma 6.47), so the inclusion $\text{Ann}_\tau \hookrightarrow \mathcal{C}^\Upsilon(\gamma)$ induces an isomorphism on π_1^D .

It follows that the coarse universal cover of Ann_τ embeds in that of $\mathcal{C}^\Upsilon(\gamma)$, containing the distinguished copy of $\widetilde{\mathcal{S}^\Upsilon(\gamma)}$ onto which the horoball \mathcal{H} is glued.

The same considerations for the construction of $\text{UG}(X_0, \gamma_0)$ imply the result. \square

Lemma 7.18. *Let (Υ, γ) be as in Construction 7.7. Let $p: \widetilde{\mathcal{C}^\Upsilon(\gamma)} \rightarrow \widetilde{\mathcal{S}^\Upsilon(\gamma)}$ be the “projection” map defined by considering an (equivariantly chosen) ray r_x in $\text{UG}(\Upsilon, \gamma)$ starting at x and limiting to the point at infinity of \mathcal{H} and letting $p(x)$ be the intersection of r_x and the horosphere $\widetilde{\mathcal{S}^\Upsilon(\gamma)}$.*

Let $[x, y]$ be a geodesic in $\text{UG}(\Upsilon, \gamma)$ that does not intersect \mathcal{H} . Then $\text{diam}_{\mathcal{H}}(p([x, y])) \leq 8\delta_2 + 4$.

Proof. By Proposition 7.13 and Lemma 7.12 we have that $\text{UG}(\Upsilon, \gamma)$ is δ_2 -hyperbolic.

To prove the lemma it suffices to show that $d(p(x), p(y)) \leq 8\delta_2 + 4$ (since we can then apply this to sub-geodesics of a given geodesic).

We can form a geodesic quadrangle containing $[x, y]$, two segments of the rays r_x, r_y and a geodesic of length at most 1 and at depth larger than $4\delta_2 + 2$.

Consider x' in \mathcal{H} along r_x at distance exactly $2\delta_2 + 2$ from $p(x)$. Since $[x, y]$ does not intersect \mathcal{H} , x' is not $2\delta_2$ -close to $[x, y]$, and is also not $2\delta_2$ -close to the geodesic of length 1. Therefore, x' is $2\delta_2$ close to a point w on r_y , and in fact $6\delta_2 + 2$ -close to $p(y)$ (since w is at depth at most $4\delta_2 + 2$). Recalling that x is exactly $2\delta_2 + 2$ from x' , we obtain the desired bound on $d(p(x), p(y))$. Notice that all the distances in this argument occur in \mathcal{H} , completing the proof of the lemma. \square

Lemma 7.19. *With the same assumptions as in the previous lemma, if B is any ball in $\text{UG}(\Upsilon, \gamma)$ disjoint from \mathcal{H} , then the covering map $\widetilde{\mathcal{C}^\Upsilon(\gamma)} \rightarrow \widetilde{\mathcal{C}^\Upsilon(\gamma)}$ is injective on B .*

Proof. Lemma 7.18 implies that the diameter (measured in \mathcal{H}) of the projection of B to the horosphere is at most $16\delta_2 + 8$.

Suppose that there are two points x, y of B which project to the same point of $\mathcal{C}^Y(\gamma)$. Then there is an element of the deck group of $\widetilde{\mathcal{C}^Y(\gamma)} \rightarrow \mathcal{C}^Y(\gamma)$ which takes x to y , and hence $p(x)$ to $p(y)$.

In particular, this element has translation length in $\widetilde{\mathcal{S}^Y(\gamma)} \cong \mathcal{S}$ at most $2 \cdot 2^{\frac{1}{2}(16\delta_2+8)} < 2^{16\delta_2}$ (the exponential difference coming from the difference between $\mathcal{S}^Y(\gamma)$ and \mathcal{H} – see Lemma 4.9). But this contradicts the choice of sys_0 in Notation 7.2. \square

Proposition 7.20. *Suppose that Y is D_1 -simply-connected, δ_2 -hyperbolic graph. Suppose further that, for some $\chi > 10R_0 + 60\sigma_0 + 20\delta_2$, \mathcal{A} is a χ -separated collection of tubes and horoballs in Y .*

Let C be a tube of \mathcal{A} and let $\widehat{Y} = \text{UG}(Y, \gamma(C))$. Let $\widehat{\mathcal{A}}$ be the collection consisting of the horoball added to $\widetilde{\mathcal{C}^Y(\gamma(C))}$ and all of the lifts of the elements of $\mathcal{A} \setminus \{C\}$ to $\widetilde{\mathcal{C}^Y(\gamma(C))} \subset \widehat{Y}$. Then $\widehat{\mathcal{A}}$ forms a χ -separated collection of tubes and horoballs in \widehat{Y} . Moreover if $A \in \widehat{\mathcal{A}}$ is a lift of a tube (resp. horoball) of \mathcal{A} then it is a tube (resp. horoball).

Proof. First of all, note that all $\frac{\chi}{10}$ -neighborhoods of horoballs and tubes other than C do lift, to horoballs and tubes, respectively. This is because they are D_1 -simply-connected. Horoballs are D_1 -simply connected by Lemma 4.10. One can see that tubes are D_1 -simply connected by taking a loop in some tube, coning to the quasi-geodesic, and filling a triangulation of the resulting loop. The resulting triangulated disk can be cut into loops of length at most $3\delta_0 + 3 < D_1$.

The point (2) of Definition 7.14 follows from the fact that the cores $\gamma(A)$ of tubes all have $\frac{\chi}{10}$ -neighborhoods isometric to $T_{\frac{\chi}{10}}(\gamma_0)$ and so each component of the preimage of a tube which lifts also has such a neighborhood.

Similarly, a lifted horoball clearly still satisfies the point (3). Let \mathcal{H}' be the new horoball glued onto the unwrapped shell which is the coarse universal cover of the shell around C . By Lemma 7.17, the $\frac{\chi}{10}$ -neighborhood of \mathcal{H}' in $\widetilde{\mathcal{C}^Y(\gamma(C))}$ is isomorphic to the $\frac{\chi}{10}$ -neighborhood of the distinguished horoball in $\text{UG}(X_0, \gamma_0)$. Thus the point (3) from Definition 7.14 holds also.

Regarding the point (4), it is clear that lifts of distinct elements of \mathcal{A} are at least as far apart as the elements.

Denote by \mathcal{H} the horoball added to $\widetilde{\mathcal{C}^Y(\gamma(C))}$. Let $p : \widetilde{\mathcal{C}^Y(\gamma(C))} \rightarrow \mathcal{S}^Y(\gamma(C))$ be the “projection” map considered in Lemma 7.18, with $\gamma = \gamma(C)$.

By Lemma 7.18, for any $A \in \widehat{\mathcal{A}} \setminus \{\mathcal{H}\}$ we have that $\text{diam}(p(A)) \leq 24\delta_2 + 12$. Indeed, any two points of A can be joined by a concatenation of 3 geodesics that do not intersect \mathcal{H} .

What we are left to show is that given two lifts $\widehat{A}, \widehat{A}'$ of the same element A of $\mathcal{A} \setminus \{C\}$, the two lifts are χ -separated. Consider a shortest geodesic from A to C . We can lift this to paths starting at $\widehat{A}, \widehat{A}'$ and ending at \mathcal{H} , with endpoints $\widehat{c}, \widehat{c}'$. Since \widehat{c} and \widehat{c}' are translates of each other by a non-trivial deck transformation, by the choice of sys_0 from Notation 7.2 we have that they are further than $100\delta_2$ apart (this is not optimal). Moreover, $\widehat{c}, \widehat{c}'$ are (or more precisely, can be taken to be) the projections of their starting points to \mathcal{H} , any two points in $\widehat{A}, \widehat{A}'$ have projections more than $6\delta_2 + 2$ apart, and therefore by Lemma 7.18 any geodesic connecting them intersects \mathcal{H} . Since $\widehat{A}, \widehat{A}'$ are at least χ away from \mathcal{H} , they are at least χ (in fact, 2χ) apart. \square

Recall the element g fixed in Assumption 7.1. The isometry g naturally induces a fixed point free isometry of \mathcal{S}^{D_1} , which we continue to denote by g .

Lemma 7.21. *There is an isomorphism*

$$E = \pi_1 \left(\langle g \rangle \backslash \mathcal{S}^{D_1} \right) \cong \pi_1 \left(\langle g \rangle \backslash \partial X_0 \setminus \{\gamma_0^\pm\} \right).$$

Moreover, E is a semidirect product of infinite cyclic groups

$$E = \langle a \rangle \rtimes \langle b \rangle.$$

Realizing E as the deck group of the universal cover $\widetilde{\mathcal{S}^{D_1}} \rightarrow \langle g \rangle \backslash \mathcal{S}^{D_1}$, the element a is a generator of the deck group of the intermediate cover $\widetilde{\mathcal{S}^{D_1}} \rightarrow \mathcal{S}^{D_1}$ and b is a lift of g to $\widetilde{\mathcal{S}^{D_1}}$.

Finally, E acts geometrically on $\widetilde{\mathcal{S}}$, so that both E and $\widetilde{\mathcal{S}}$ are quasi-isometric to \mathbb{R}^2 .

Proof. The first statement follows immediately from Theorem 6.35 and Corollary 6.36. Indeed, $\pi_1 \left(\langle g \rangle \backslash \partial X_0 \setminus \{\gamma_0^\pm\} \right)$ is isomorphic to semi-direct product of $\pi_1(\partial X_0 \setminus \{\gamma_0^\pm\})$ by $\langle g \rangle$, where g acts on $\pi_1(\partial X_0 \setminus \{\gamma_0^\pm\})$ via the induced action on π_1 . The other π_1 admits a similar description, with the action on π_1 being compatible with the isomorphism given by item Theorem 6.35.(1) in view of item Theorem 6.35.(3).

The description of the deck group follows from the fact that we can describe the universal covering map as a composition of two covers $\widetilde{\mathcal{S}^{D_1}} \rightarrow \mathcal{S}^{D_1}$ and $\mathcal{S}^{D_1} \rightarrow \langle g \rangle \backslash \mathcal{S}^{D_1}$, the first cover having deck group $\pi_1^{D_1}(\mathcal{S})$ and the second cover having deck group $\langle g \rangle$. \square

Proposition 7.22. *The action of the group E from Lemma 7.21 on $\widetilde{\mathcal{S}}$ extends to an action on $\widetilde{\mathcal{C}}$.*

Proof. This follows from the fact that the inclusion of $\langle g \rangle \backslash \mathcal{S}^{D_1}$ into $\langle g \rangle \backslash \mathcal{C}^{D_1}$ induces an isomorphism on fundamental groups, by Proposition 6.48 and the fact that the action of g commutes with the inclusion of \mathcal{S}^{D_1} into \mathcal{C}^{D_1} . \square

7.2. Unwrapping: boundaries.

Proposition 7.23. *Let Υ be a connected D_1 -simply connected graph, and let γ be a bi-infinite path in Υ . Assume that Υ is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}\}$. Suppose further that (Υ, γ) is $(R_0 + 6\sigma_0 + 2\delta_2)$ -tube comparable to (X_0, γ_0) .*

Let p be the point at infinity of the horoball \mathcal{H} of $\partial\text{UG}(\Upsilon, \gamma)$. There exists a regular covering map $\Theta : \partial\text{UG}(\Upsilon, \gamma) \setminus \{p\} \rightarrow \partial\Upsilon \setminus \{\gamma^\pm\}$, and the deck transformation group is \mathbb{Z} .

Moreover, if $(\Upsilon, \gamma) = (X_0, \gamma_0)$ then Θ is the universal covering map.

Proof. Let us construct the covering map Θ . We have a (D_1) -universal covering map $q_{R_0} : \widetilde{\mathcal{C}^\Upsilon(\gamma)} \rightarrow \mathcal{C}^\Upsilon(\gamma)$ as in Construction 7.7.

Fix $w \in \partial\text{UG}(\Upsilon, \gamma) \setminus \{p\}$. We claim that we can choose a geodesic ray r_w beginning at $\widetilde{\mathcal{S}^\Upsilon(\gamma)}$ and ending at w so that $q(r_w)$ is a geodesic ray in $\Upsilon \setminus T_{R_0}(\gamma)$, and therefore determines a point in $\partial\Upsilon$ that we denote $\Theta(w)$.

Proof of Claim. Consider a sequence of points $w_i \in \text{UG}(\Upsilon, \gamma)$ tending to w . There are geodesic segments r_i determined by projecting from w_i to $\widetilde{\mathcal{S}^\Upsilon(\gamma)}$. Each of these have the property that $r_i(t)$ is distance t from $\widetilde{\mathcal{S}^\Upsilon(\gamma)}$. These sub-converge to an infinite ray r_w with this property representing w . Then $q(r_w)$ has the property that $q(r_w(t))$ is distance t from $\mathcal{S}^\Upsilon(\gamma)$ and is a geodesic. \square

Note that $\Theta(w) \neq \gamma^\pm$. Indeed, if $q(r_w)$ had point at infinity γ^\pm then $q(r_w)$ would intersect $T_{R_0}(\gamma)$ arbitrarily far from its starting point, since γ is λ_0 -quasi-convex by Lemma 6.43 and R_0 is large compared to λ_0 and δ_2 (see Notation 7.2), the latter being a hyperbolicity constant for Υ by Lemma 7.12. This is a contradiction.

The map Θ is easily seen to be continuous because, by taking $z \in \partial\text{UG}(\Upsilon, \gamma) \setminus \{p\}$ sufficiently close to w , it can be arranged for the corresponding rays $q(r_w), q(r_z)$ to fellow-travel for an arbitrarily long time. Note that by the same argument any choice of r_w gives the same $\Theta(w)$.

Let us prove that Θ is surjective. Consider $u \in \partial\Upsilon \setminus \{\gamma^\pm\}$, and choose a geodesic ray α_u that intersects $\mathcal{S}^\Upsilon(\gamma)$ only at its starting point and has point at infinity u . Then α_u can be lifted to $\text{UG}(\Upsilon, \gamma)$. We can in fact choose α_u that lifts to a geodesic ray in $\text{UG}(\Upsilon, \gamma)$; this uses the argument from the proof of the claim applied to a sequence of points along the lift, which

yields a ray in $\text{UG}(\Upsilon, \gamma)$ projecting to a ray in Υ at finite distance, in our case, from the original choice of α_u . For v the point at infinity of the lift, we can take said lift as r_v , and therefore we get $\Theta(v) = u$.

A similar argument shows that Θ is open (since in order to prove it one has to prove a surjectivity statement, meaning that the image of an open set contains a small neighborhood around each of its points). One can also see that, in view of Theorem 6.35.(2), which provides a lower bound on the translation distance of any non-trivial element of the deck group of q , different lifts of α_u yield different points in $\partial\text{UG}(\Upsilon, \gamma)$. More precisely, those lifts are rays that lie at distance at least sys_0 from each other in $\Upsilon \setminus T_{R_0}(\gamma)$, and it can be deduced from Lemma 4.9 that they lie at least $2(\log_2(\text{sys}_0) - 1)$ away from each other in $\text{UG}(\Upsilon, \gamma)$. This is much larger than the hyperbolicity constant δ_2 of $\text{UG}(\Upsilon, \gamma)$ (see Proposition 7.13), so that the lifts have distinct points at infinity. Moreover, again because of the lower bound on the translation distance, the deck group acts properly discontinuously on $\partial\text{UG}(\Upsilon, \gamma) \setminus \{p\}$ (roughly, as above, different orbit points project very far on the horosphere). Moreover, if $\Theta(w) = \Theta(w')$ then for any point on a ray r_w there is an element of the deck group mapping it within distance bounded in terms of δ_0 from a point on a ray $r_{w'}$. Picking the point on r_w sufficiently far from the horosphere, said element of the deck group must map w to w' ; if not, we would have orbit points of w accumulating to w' , contradicting proper discontinuity.

We now have that Θ factors through the quotient map

$$\Psi: \partial\text{UG}(\Upsilon, \gamma) \setminus \{p\} \rightarrow \text{Deck}(q) \backslash (\partial\text{UG}(\Upsilon, \gamma) \setminus \{p\}),$$

where $\text{Deck}(q)$ is the group of the extensions to $\partial\text{UG}(\Upsilon, \gamma) \setminus \{p\}$ of the elements of the deck transformation group of q , which acts properly discontinuously on $\partial\text{UG}(\Upsilon, \gamma) \setminus \{p\}$. Moreover, writing $\Theta = \iota \circ \Psi$, we have that ι is bijective and continuous, and it is open since Θ is. Hence, ι is a homeomorphism.

Since Ψ is a covering map with deck group isomorphic to \mathbb{Z} , so is Θ .

To prove the “moreover” statement we consider the following (not quite commutative) diagram:

$$(14) \quad \begin{array}{ccc} \partial\text{UG}(X_0, \gamma_0) \setminus \{p\} & \xrightarrow{\phi} & \widetilde{\mathcal{F}} \\ \Theta \downarrow & & \downarrow q \\ \partial X_0 \setminus \{\gamma_0^{\pm\infty}\} & \xrightarrow{\Pi_{R_0}} & \mathcal{S} \end{array}$$

The maps Θ, q were defined above: the map Π_{R_0} is from Definition 6.3. The map ϕ is defined via the choice of rays r_w above; namely $\phi(w) = r_w(0)$.

Claim. The diagram (14) commutes up to an error of $34\delta_0$.

Proof. The point $q(\phi(w))$ is the initial point of a geodesic ray (namely $q(r_w)$) in X_0 which tends to the point $\Theta(w)$, and only meets \mathcal{S} in one point. The point $\Pi_{R_0}(\Theta(w))$ is the initial point of another such ray, r' . Lemma 6.22 implies that for all $t \geq 0$,

$$d_{X_0}(r'(t), \gamma_0) \sim_{10\delta_0} R_0 + t$$

and similarly for r_w . Let $q \in \gamma_0$ be closest to $r_w(0)$ and let $q' \in \gamma_0$ be closest to $r'(0)$. Since r_w and r' eventually stay in $2\delta_0$ -neighborhoods of each other we may choose $z \in r_w$ and $z' \in r'$ which are distance at least $R_0 + 100\delta_0$ from γ_0 and so that $d_{X_0}(z, z') \leq 2\delta_0$. We consider a geodesic hexagon with vertices (cyclically ordered) $\{q, r_w(0), z, z', r'(0), q'\}$. Let $T = 15\delta_0$ and let $a = r_w(T)$. Since the hexagon is $4\delta_0$ -slim, there is a point a' on one of the other five sides so that $d_{X_0}(a, a') \leq 4\delta_0$.

By Lemma 6.22, the distance from a to γ_0 is at least $R_0 + T - 10\delta > R_0$, so a' cannot be on any of the three sides contained in $T_{R_0}(\gamma_0)$. It cannot be contained in the side z, z' , since this side is too far away. Thus the point a is equal to $r'(S)$ for some S . Again using Lemma 6.22, we have $S \leq R_0 + d_{X_0}(a', \gamma_0)$. But $d_{X_0}(a', \gamma_0) \leq d_{X_0}(a, \gamma_0) + 4\delta_0 \leq R_0 + T + 4\delta_0$. Putting these together we deduce $d_{X_0}(r'(0), a) \leq T + 4\delta_0$, and so $d_{X_0}(r_w(0), r'(0)) \leq 2T + 4\delta_0 = 34\delta_0$. \square

The universal covering map of the annulus $\partial X_0 \setminus \{\gamma_0^\pm\}$ can be characterized as the only \mathbb{Z} -cover with the property that the preimage of any given point is contained in a single connected component.

Consider any $h \in \text{Deck}(q)$ and any $x \in \partial \text{UG}(X_0, \gamma_0) \setminus \{p\}$. We want to show that $h \cdot x$ lies in the same component as x . Consider $\phi(x)$ and $\phi(h \cdot x) = h \cdot \phi(x)$. Since $\text{Deck}(q)$ is naturally isomorphic to $\pi_1^{D_1}(\mathcal{S})$, and Π_{R_0} induces an isomorphism from $\pi_1(\partial X_0 \setminus \{\gamma_0^\pm\})$ to $\pi_1^{D_1}(\mathcal{S})$, there exists a loop α in $\partial X_0 \setminus \{\gamma_0^\pm\}$ at $\Theta(x)$ and a loop α' in \mathcal{S} with the following properties. First, α' lifts to a path β' with endpoints $\phi(x)$ and $\phi(h \cdot x) = h \cdot \phi(x)$ in $\text{UG}(\Upsilon, \gamma)$. Secondly, for each t we have that $\Pi_{R_0}(\alpha(t))$ is $8\delta_0$ -close to $\alpha'(t)$ (see Lemma 6.7). We have that α lifts to a path β in $\partial \text{UG}(X_0, \gamma_0) \setminus \{p\}$ starting at x , and ending at point in the fiber of $\Theta(x)$, say $h' \cdot x$. We claim that $h' = h$, and this will conclude the proof.

Indeed, by the claim, for each t there exists an element g_t of $\text{Deck}(q) \cong \mathbb{Z}$ such that $\phi(\beta(t))$ is $42\delta_0$ -close to $g_t \beta'(t)$. We have $g_0 = id$ since $\beta(0) = x$ and $\beta'(0) = \phi(x)$.

By the same argument as in Lemma 6.7, each point in $\partial \text{UG}(X_0, \gamma_0) \setminus \{p\}$ has a neighborhood that has ϕ -image of diameter at most $8\delta_2$. Since any non-trivial element acts with translation distance larger than $84\delta_0 + 8\delta_2$ on

$\widetilde{\mathcal{F}}$ (by the choice of sys_0 in Notation 7.2), for any g in the deck group the set of t such that $g_t = g$ is open. Therefore, for all t we have $g_t = id$, and in particular for $t = 1$ we have that $\phi(h'(x))$ is $42\delta_0$ -close to $\phi(hx)$. Similarly to the previous argument, this implies that $h' = h$, as required. \square

Corollary 7.24. $\partial\text{UG}(X_0, \gamma_0)$ is homeomorphic to S^2 .

Proof. Let p be the point at infinity of the horoball. In view of Proposition 7.23 we have that $\partial\text{UG}(X_0, \gamma_0) \setminus \{p\}$ is homeomorphic to \mathbb{R}^2 , as it is the universal cover of an open annulus.

Since $\partial\text{UG}(X_0, \gamma_0)$ is a compact metric space, and removing a point makes it homeomorphic to \mathbb{R}^2 , it must be homeomorphic to S^2 . \square

We say that a combinatorial graph Γ has *uniform polynomial growth* if there exist constants $c < C$ and an integer d so that, for any ball B of radius r , we have

$$cr^d \leq \#B \cap \Gamma^{(0)} \leq Cr^d.$$

Note that a combinatorial graph with a co-compact group action and polynomial growth has uniform polynomial growth.

The following lemma is implicit in the work of Dahmani and Yaman [19].

Lemma 7.25. *Let Γ be a combinatorial graph of uniform polynomial growth. There is a bounded valence graph $\mathcal{H}_{bv}(\Gamma)$ containing Γ so that the identity map on Γ extends to a quasi-isometry from $\mathcal{H}_{bv}(\Gamma)$ to the combinatorial horoball based on Γ .*

Proof. Since Γ is of uniform polynomial growth, there are constants $c < C$ and an integer d so that, for any ball B of radius r , we have $cr^d \leq \#B \cap \Gamma^{(0)} \leq Cr^d$. For each n , choose a subset V_n of $\Gamma^{(0)}$ satisfying

- (1) For any $x \in \Gamma^{(0)}$, there is some $v \in V_n$ so that $d_\Gamma(x, v) < 2^{n-1}$.
- (2) For any $v, w \in V_n$, we have $d_\Gamma(v, w) \geq 2^{n-1}$ (so the radius 2^{n-2} balls around v, w are disjoint).

Now define a graph $\mathcal{H}_{bv} = \mathcal{H}_{bv}(\Gamma)$ with $V(\mathcal{H}_{bv}) = \bigsqcup_n V_n$ and edges defined as follows.

- (horizontal edges) Connect $v, w \in V_n$ with an edge if $d_\Gamma(v, w) < 2^{n+1}$.
- (vertical edges) Connect $v \in V_n$ to $w \in V_{n+1}$ if $d_\Gamma(v, w) < 2^{n+1}$.

Notice that $V_0 = \Gamma^{(0)}$ and the sub-graph induced by V_0 is canonically isomorphic to Γ .

Let $\mathcal{H} = \mathcal{H}(\Gamma)$ be the combinatorial horoball based on Γ . (Recall Definition 4.8 for the definition of a combinatorial horoball.) We will define coarsely Lipschitz maps

$$\phi : \mathcal{H}_{bv}^{(0)} \rightarrow \mathcal{H}^{(0)} \quad \text{and} \quad \psi : \mathcal{H}^{(0)} \rightarrow \mathcal{H}_{bv}^{(0)}$$

which are quasi-inverses of one another, which implies they are quasi-isometries.

For $v \in V_n$, we define $\phi(v) = (v, n)$. Suppose that $v, w \in V_n$ are connected by a horizontal edge. Then $d_\Gamma(v, w) < 2^{n+1}$, so there is a horizontal path in \mathcal{H} of length at most 2 connecting $\phi(v) = (v, n)$ to $\phi(w) = (w, n)$. If $v \in V_n, w \in V_{n+1}$ are connected by a vertical edge, then there is a path of length at most 2 connecting $\phi(v) = (v, n)$ to $\phi(w) = (w, n+1)$. Namely, connect (v, n) to $(v, n+1)$ by a vertical edge, and then $(v, n+1)$ to $(w, n+1)$ by a horizontal edge. Summarizing, we have shown that ϕ is 2-Lipschitz.

For $(v, n) \in \mathcal{H}^{(0)}$, we choose a point $v' \in V_n$ minimizing $d_\Gamma(v, v')$, and define $\psi(v, n)$ to be v' . We now show ψ is coarsely Lipschitz. Suppose first that (v, n) is connected to (w, n) by a horizontal edge, and that v', w' are the chosen points of V_n closest to v, w respectively. Then $d_\Gamma(v', w') < 2^{n+1} + 2^n = 6 \cdot 2^{n-1}$. There is therefore a sequence of vertices $v' = x_0, x_1, \dots, x_6 = w'$ in $\Gamma^{(0)}$ so that $d_\Gamma(x_i, x_{i+1}) \leq 2^{n-1}$. For each x_i there is a vertex $v_i \in V_n$ so that $d_\Gamma(v_i, x_i) < 2^{n-1}$, by (1). We thus have, for each i , $d_\Gamma(v_i, v_{i+1}) < 3 \cdot 2^{n-1} < 2^{n+1}$, so v_i is connected by a horizontal edge to v_{i+1} . In particular, we have $d_{\mathcal{H}_{bv}}(v', w') \leq 6$, so ψ sends vertices connected by a horizontal edge to vertices which are at most 6 apart.

Now consider a vertical edge, joining (v, n) to $(v, n+1)$ in \mathcal{H} . Let $v' \in V_n$ and $v'' \in V_{n+1}$ be the images under ψ of (v, n) and $(v, n+1)$, respectively. We have $d_\Gamma(v, v') < 2^{n-1}$ and $d_\Gamma(v, v'') < 2^n$. The vertex v' is connected by a vertical edge in \mathcal{H}_{bv} to some $w \in V_{n+1}$ so that $d_\Gamma(v', w) < 2^n$. We thus have $d_\Gamma(w, v'') < 2^{n+2}$, so v', w are connected by a horizontal edge. In particular $d_{\mathcal{H}_{bv}}(\psi(v, n), \psi(v, n+1)) \leq 2$. Together with the bound in the last paragraph this shows that ψ is 6-Lipschitz.

The composition $\psi \circ \phi$ is the identity. The composition $\phi \circ \psi$ maps a vertex (v, n) to some (v', n) so that $d_\Gamma(v, v') < 2^{n-1}$. In particular, (v, n) and (v', n) are connected by a horizontal edge in \mathcal{H} and so $d_\Gamma(\phi \circ \psi(v, n), (v, n)) \leq 1$.

It remains to show that \mathcal{H}_{bv} has bounded valence. Let $v \in V_n$ be some vertex of \mathcal{H}_{bv} , and let $w_1, \dots, w_k \in V_n$ be vertices connected by an edge to v . These points w_i all lie in the ball of radius 2^{n+1} around v , and the 2^{n-2} -balls around the w_i are disjoint, by (2). Each of these balls has volume at least $c(2^{n-2})^d$, and they are all contained in the 2^{n+2} -ball around v , which has volume at most $C(2^{n+2})^d$, so we have

$$kc(2^{n-2})^d \leq C(2^{n+2})^d,$$

implying that $k \leq 16^d \frac{C}{c}$.

If $u_1, \dots, u_l \in V_{n-1}$ are connected by vertical edges to v , their 2^{n-3} -balls in Γ are disjoint and contained in the 2^{n+2} -ball around v , so we similarly argue that $l \leq 32^d \frac{C}{c}$.

Finally, if $z_1, \dots, z_m \in V_{n+1}$ are connected by vertical edges to v , their (disjoint) 2^{n-1} -balls are contained in the 2^{n+3} -ball around v , so we have $m \leq 16^d \frac{C}{c}$. Put together we get a bound on the valence of v which is independent of v . \square

Lemma 7.26. *Suppose X_0 has bounded valence. Then the space $\text{UG}(X_0, \gamma_0)$ is quasi-isometric to a bounded valence graph. More precisely, there exists a bounded valence graph \mathcal{H}_{bv} which is quasi-isometric to a horoball in \mathbb{H}^3 and contains a copy of \mathcal{S} so that if \mathcal{B} is the space obtained by gluing \mathcal{H}_{bv} to \mathcal{C} along \mathcal{S} then there is a quasi-isometry between $\text{UG}(X_0, \gamma_0)$ and \mathcal{B} extending the identity map on \mathcal{C} .*

Proof. By Proposition 7.22, \mathcal{S} has a free and cocompact action by a group quasi-isometric to \mathbb{R}^2 . It follows from this that the combinatorial horoball \mathcal{H} based on \mathcal{S} is quasi-isometric to a horoball in \mathbb{H}^3 (see, for example [26, Corollary 9.2]).

It also follows that \mathcal{S} has uniform polynomial (quadratic) growth. Let \mathcal{H}_{bv} be the graph $\mathcal{H}_{bv}(\mathcal{S})$ from the conclusion of Lemma 7.25. Let \mathcal{B} be the space obtained from gluing \mathcal{C} to \mathcal{H}_{bv} along the natural copies of \mathcal{S} . Since we have assumed X_0 has bounded valence, and \mathcal{H}_{bv} does by Lemma 7.25, the space \mathcal{B} has bounded valence. The horoballs \mathcal{H} and \mathcal{H}_{bv} are quasi-isometric, so the identity on \mathcal{C} extends to coarsely inverse coarsely Lipschitz maps $\mathcal{B} \rightarrow \text{UG}(X_0, \gamma_0)$ and $\text{UG}(X_0, \gamma_0) \rightarrow \mathcal{B}$. This implies that the spaces are quasi-isometric. \square

7.3. Uniform linear connectivity. Up to this point, we have not required there to be a cocompact group of isometries of the space X_0 . However, the following result does require this assumption. Thus we make the following extra assumption.

Assumption 7.27. *Suppose that X_0 satisfies Assumption 7.1, and furthermore assume that X_0 admits a group of isometries containing g which is geometric.*

Notice that under Assumption 7.27, X_0 has bounded valence (rather than just being locally finite), so the assumptions of Lemma 7.26 hold.

In the proof of the following lemma we use a limiting argument, so the resulting constant is not explicit and depends on X_0, γ_0 , and R_0 .

Lemma 7.28. *There exists $L_1 = L_1(X_0, \gamma_0, R_0)$ so that any δ_2 -adapted visual metric on $\partial \text{UG}(X_0, \gamma_0)$ with any basepoint is L_1 -linearly connected.*

Proof. Let \mathcal{B} be the bounded valence graph from Lemma 7.26. This space is quasi-isometric to $\text{UG}(X_0, \gamma_0)$ and is δ' -hyperbolic for some δ' . Thus to prove the current lemma it suffices to show that the elements of any sequence of δ' -adapted visual metrics on $\partial\mathcal{B}$ are uniformly linearly connected.

We therefore assume (y_i) is a sequence of points in \mathcal{B} . Using bounded valence we may pass to a subsequence such that the pointed spaces (\mathcal{B}, y_i) strongly converge to a pointed space $(\mathcal{B}_\infty, y_\infty)$. By Proposition 3.13, the boundaries $\partial\mathcal{B}$ with adapted visual metrics at y_i weakly Gromov-Hausdorff converge to the boundary $\partial\mathcal{B}_\infty$ with adapted visual metric at y_∞ .

We claim that given any such sequence y_i and limit $(\mathcal{B}_\infty, y_\infty)$, the space $\partial\mathcal{B}_\infty$ does not have a weak cut point.

Given this claim, Theorem 5.3 implies that the δ' -adapted visual metrics on (\mathcal{B}, y_i) must have been uniformly linearly connected, establishing the lemma.

It remains to prove the claim. The group E from Proposition 7.22 acts on $\text{UG}(X_0, \gamma_0)$ and the quasi-isometry from Lemma 7.26 therefore induces a quasi-action of E on \mathcal{B} (by uniform quality quasi-isometries).

The possible limits depend on the position of y_i compared to the distinguished horoball in \mathcal{B} . If the distance from the horoball diverges, then \mathcal{B}_∞ is a Gromov-Hausdorff limit of X_0 , and hence its boundary is S^2 (by Assumption 7.27 there is a geometric action on X_0 , so all Gromov-Hausdorff limits of X_0 are quasi-isometric to X_0 , which by assumption has boundary S^2). If the y_i are deeper and deeper in the horoball, then the limit is a limit of \mathcal{H}_{bv} , which is quasi-isometric to a horoball in \mathbb{H}^3 , by Lemma 7.26. Hence, once again, the boundary of \mathcal{B}_∞ is S^2 since \mathcal{B}_∞ is quasi-isometric to \mathbb{H}^3 . Finally, if the y_i stay within bounded distance from \mathcal{I} , then we can move them to a bounded region of \mathcal{B} using the quasi-action on \mathcal{B} , which is cobounded on $\widetilde{\mathcal{I}}$. Therefore \mathcal{B}_∞ is quasi-isometric to \mathcal{B} , and hence to $\text{UG}(X_0, \gamma_0)$. Therefore, the boundary of \mathcal{B}_∞ is S^2 in this case too, by Corollary 7.24. Thus in every case $\partial\mathcal{B}_\infty$ is S^2 , which does not have a weak cut point, establishing the claim, and hence the lemma. \square

The following is an immediate consequence of Lemmas 5.10 and 7.28.

Corollary 7.29. *There exists $\Delta = \Delta(X_0, \gamma_0, R_0)$ so that for every $y \in \text{UG}(X_0, \gamma_0)$ and every $\Sigma \geq 0$ the space $\text{UG}(X_0, \gamma_0)$ is (Δ, Σ) -spherically connected at y .*

By Lemma 5.11, we therefore have the following.

Corollary 7.30. *There exist Σ_0 and L so that for any δ_2 -hyperbolic space Y which is Σ_0 -modeled on $\{\text{UG}(X_0, \gamma_0)\}$ and any $y \in Y$, any δ_2 -adapted visual metric on ∂Y with basepoint y is L -linearly connected.*

8. UNWRAPPING A FAMILY OF AXES

In the previous section we unwrapped the complement of a single tube. In this section we iterate this process, starting from a suitable family of axes in a hyperbolic space on which a group G acts. We construct a sequence of “partially” unwrapped spaces, and we show that they converge to a space where another group \widehat{G} acts. This \widehat{G} will be the drilled group from the main Theorem A, and the main goal of this section is to show that it is relatively hyperbolic, and a drilling of G along g .

8.1. Assumptions for this section. Fix a group G acting freely and co-compactly on a δ_0 -hyperbolic graph X_0 with boundary homeomorphic to S^2 , and up to increasing δ_0 we can and do assume that X_0 is δ_0 -visible, by [4, Lemma 3.1]. We fix a loxodromic element g and a continuous quasi-geodesic γ_0 stabilized by g . Fix also a linear connectedness constant L_0 and a doubling constant N_{doub} as in Assumption 6.1, which exist by [6, Proposition 4] and [7, Theorem 9.2] respectively. Fix λ_0 as in Assumptions 6.1 (for $Y_0 = \gamma_0$), and notice that at this point we have that Assumptions 6.1 and 7.1 hold for X_0 and $Y_0 = \gamma_0$, and the associated constants. Further, let $s_0, D_1, \delta_1, \delta_2, \sigma_0, \text{sys}_0$ and R_0 be chosen as in Notation 7.2.

We now make careful choices of some further constants:

- (1) Fix Σ_0 and L as in Corollary 7.30. Without loss of generality, we also assume that $\Sigma_0 \geq 10^9 \delta_1$.
- (2) Fix $\Sigma_1 = 20\Sigma_0 + 60\sigma_0 + 10R_0 + 50\delta_2$.
- (3) Fix $\Sigma = \Sigma_1 + 2R_0$. This is the constant Σ in our main theorem, Theorem A.

We furthermore assume that the collection of translates $\{h \cdot \gamma_0\}_{h \in G \setminus \langle g \rangle}$ forms a Σ -separated family (that is, each two distinct elements are at least Σ apart), so the tubes $\{h \cdot T_{R_0}(\gamma_0)\}_{h \in G \setminus \langle g \rangle}$ form a Σ_1 -separated family.

8.2. Construction of the intermediate spaces. We now describe a sequence of pointed spaces obtained by repeatedly applying the “unwrap and glue” Construction 7.7. The limit of these spaces will be the cusped space of the group obtained by drilling G along g .

Let x_0 be a base vertex of X_0 which lies further than R_0 from any $h \cdot \gamma_0$ with $h \in G$. We also regard x_0 as a basepoint for $\check{X}_0 = \mathcal{C}_{R_0}^{X_0}(\bigcup_{h \in G} h \cdot \gamma_0)$ (see Definition 6.15 and Convention 6.16), that is, the completed complement of all the tubes around the translates $h \cdot \gamma_0$.

The collection of edge-paths in \check{X}_0 beginning at x_0 and ending on some shell is countable. We choose an enumeration $\{p_i\}_{i \geq 1}$ of these paths. For each i we fix some h_i so that p_i ends on the shell $\mathcal{S}^{X_0}(h_i \cdot \gamma_0)$. It will happen for some (in fact many) $i \neq j$ that $h_i = h_j$.

Define the pair $(\widehat{Y}_1, y_1) = (\text{UG}(X_0, \gamma_0), x_1)$, where x_1 is any preimage of the basepoint $x_0 \in X_0$. By relabeling if necessary, we suppose that the path p_1 lifts to a path in \widehat{Y}_1 which starts at y_1 and finishes on the (single) horoball in \widehat{Y}_1 . Proposition 7.20 furnishes \widehat{Y}_1 with a Σ -separated collection of tubes and horoballs, \mathcal{A}_1 say (cf. Definition 7.14). Define $q_1: \text{HO}(\widehat{Y}_1, \mathcal{A}_1) \rightarrow \check{X}_0$ be the restriction of the covering map $\check{\mathcal{C}} \rightarrow \mathcal{C}$ from the construction of $\text{UG}(X_0, \gamma_0)$ to $\text{HO}(\widehat{Y}_1, \mathcal{A}_1)$. (The ‘‘hollowed out’’ space $\text{HO}(\cdot, \cdot)$ was defined in Construction 7.16.)

Let $j \geq 1$. Suppose, by induction, that for all l such that $1 \leq l \leq j$ we have constructed a quadruple $(\widehat{Y}_l, \mathcal{A}_l, y_l, q_l)$ where: \widehat{Y}_l is a space, \mathcal{A}_l is a collection of tubes and horoballs in \widehat{Y}_l , $y_l \in \widehat{Y}_l$ is a basepoint and

$$q_l: \text{HO}(\widehat{Y}_l, \mathcal{A}_l) \rightarrow \text{HO}(\widehat{Y}_{l-1}, \mathcal{A}_{l-1})$$

is a D_1 -covering map (see Definition 4.1) such that $q_l(y_l) = y_{l-1}$. Let $\check{Y}_l = \text{HO}(\widehat{Y}_l, \mathcal{A}_l)$. (Note that composing the D_1 -covers gives a D_1 -cover $q_1 \circ \cdots \circ q_l: \check{Y}_l \rightarrow \check{X}_0$.)

Furthermore, we suppose that the following properties hold, for each $1 \leq l \leq j$:

- (1) \widehat{Y}_l is connected and D_1 -simply-connected;
- (2) \widehat{Y}_l is Σ_0 -modeled on $\text{UG}(X_0, \gamma_0)$;
- (3) \widehat{Y}_l is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$ (see Notation 7.5 and Definition 7.6 for these spaces);
- (4) \widehat{Y}_l is δ_2 -hyperbolic;
- (5) \mathcal{A}_l is a Σ_1 -separated collection of tubes and horoballs, as in Definition 7.14;
- (6) Let $(p_i^l)_{i \geq 1}$ be the sequence of paths starting at y_l which are lifts of the p_i (with p_i^l being the lift of p_i). For each $k \geq 1$, the path p_k^l finishes on a shell or a horosphere associated to some element of \mathcal{A}_l . For each $k \leq l$ the path p_k^l finishes on a horosphere (not a shell).

Note that the third item follows immediately from the second, the fact that $\Sigma_0 > \sigma_0$ and Proposition 7.13 (observe that X_0 is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}^\circ\}$).

For the base case of the induction, let us prove that the properties hold for $j = 1$.

The property (1) is Proposition 7.13. A space is modeled on itself, so (2) holds. We have already noted that the property (3) follows from (2). The property (4) holds by Proposition 7.13 and Lemma 7.12. We have already noted that Proposition 7.20 implies (5). Finally, the choice of p_1 above implies the property (6).

Given these assumptions, we define a quadruple $(\widehat{Y}_{j+1}, \mathcal{A}_{j+1}, y_{j+1}, q_{j+1})$ inductively so that all of these assumptions are satisfied.

If p_{j+1}^j ends on a horosphere, then $\widehat{Y}_{j+1} = \widehat{Y}_j$, and all of the other data remains the same.

We now suppose that p_{j+1}^j ends on the (completed) shell S_j of a tube C_j from \mathcal{A}_j , with core c_j . We let $\widehat{Y}_{j+1} = \text{UG}(\widehat{Y}_j, c_j)$ (see Construction 7.7). That is, we do the following. We remove the tube C_j , we complete the space we obtain, we take the D_1 -universal cover, and finally we glue in a copy of \mathcal{H} .

Let y_{j+1} be any preimage of y_j . Let \mathcal{A}_{j+1} be the collection of all of the lifts of the elements of \mathcal{A}_j , along with the new horoball. It follows from Proposition 7.13 (and the assumptions on \widehat{Y}_j) that \widehat{Y}_{j+1} is connected and D_1 -simply-connected, so the property (1) holds. Let B be a ball of radius Σ_0 in \widehat{Y}_{j+1} . If B intersects the new horoball, then it is contained in the $2\Sigma_0$ -neighborhood of this horoball, and so by Lemma 7.17 it is isometric to a ball of radius Σ_0 in $\text{UG}(X_0, \gamma_0)$ (this uses the choice of Σ , Definition 7.14, and Remark 7.15). On the other hand, if B does not intersect the new horoball then, by Lemma 7.19, B injects into \widehat{Y}_j , and so is isometric to a ball in $\text{UG}(X_0, \gamma_0)$ by induction. Hence the property (2) holds, and so (3) does also. It follows from Lemma 7.12 that \widehat{Y}_{j+1} is δ_2 -hyperbolic so (4) holds. It follows from Proposition 7.20 that \mathcal{A}_{j+1} forms a Σ_1 -separated collection of tubes and horoballs, so (5) holds. The property (6) holds by the construction.

We define \check{Y}_{j+1} to be $\text{HO}(\widehat{Y}_{j+1}, \mathcal{A}_{j+1})$ and construct the map $q_{j+1}: \check{Y}_{j+1} \rightarrow \check{Y}_j$, which will be an infinite cyclic cover.

We claim that we have the following diagram where the vertical arrows are covers. The middle column is the same as the right column except that the shells around tubes have not been completed.

$$\begin{array}{ccccc} \widetilde{\mathcal{C}^{\widehat{Y}_j}(c_j)} & \longleftarrow & \widehat{Y}_{j+1} \setminus \bigcup \mathcal{A}_{j+1} & \longrightarrow & \check{Y}_{j+1} \\ \downarrow & & \downarrow & & \downarrow \\ \mathcal{C}^{\widehat{Y}_j}(c_j) & \longleftarrow & \widehat{Y}_j \setminus \bigcup \mathcal{A}_j & \longrightarrow & \check{Y}_j \end{array}$$

The left (D_1 -universal) cover indeed restricts to the central cover because of the definition of \mathcal{A}_{j+1} . Since tubes are Σ_1 -separated and Σ_1 is much larger than s_0 , and lifts of tubes are tubes by Proposition 7.20, the central cover extends to a cover on the right. Let q_{j+1} be the cover on the right. By construction, this is a D_1 -covering.

We can now consider all lifts p_i^{j+1} of the p_i via q_{j+1} starting at the basepoint y_{j+1} . We think of these lifts as obtained by first lifting to \check{Y}_j and then to \check{Y}_{j+1} . It is clear that all p_i^{j+1} finish on a shell or horosphere associated to an element of \mathcal{A}_{j+1} . Moreover, by construction we see that p_{j+1}^{j+1} finishes on a horosphere (since the full preimage of the shell we unwrap is a horosphere). For $1 \leq k \leq j$ the fact that p_k^j finishes on a horosphere implies that p_k^{j+1} also finishes on a horosphere.

This finishes the inductive construction of $(\widehat{Y}_{j+1}, \mathcal{A}_{j+1}, y_{j+1}, q_{j+1})$.

Since the spaces \check{Y}_j form a tower of D_1 -covers, the following result is immediate.

Lemma 8.1. *There is an inverse limit \check{Y} of the tower of covers \check{Y}_j , along with a D_1 -covering map $q : \check{Y} \rightarrow \check{X}_0$, and a basepoint \check{y} , as well as D_1 -covering maps $r_j : \check{Y} \rightarrow \check{Y}_j$, such that $r_j(\check{y}) = y_j$.*

Lemma 8.2. *Any component of the preimage of $\mathcal{S}^{X_0}(h_i \gamma_0)$ under $q : \check{Y} \rightarrow \check{X}_0$ is isometric to the D_1 -universal cover $\widetilde{\mathcal{S}^{X_0}(h_i \gamma_0)}$.*

Proof. Consider any preimage C as in the statement. An edge-path from the basepoint \check{y} to C is a lift of one of the paths p_j . By construction, the lift of p_j to \check{Y}_j starting at y_j ends on a horosphere. Therefore, the same holds in the further cover \check{Y} . \square

Definition 8.3. Denote by \widehat{Y} the space obtained from \check{Y} by gluing in copies of \mathcal{H} to each component in the pre-image of a completed shell as in Lemma 8.2.

We now want to show that the \widehat{Y}_j strongly converge to \widehat{Y} ; in order to do so we will use Lemma 7.8 (the lemma which, roughly, allows us to promote local homeomorphisms to local isometries).

Lemma 8.4. *The spaces (\widehat{Y}_j, y_j) strongly converge to (\widehat{Y}, \check{y}) .*

Proof. Consider the ball B_k of radius k in \widehat{Y} around \check{y} . By Lemma 7.8 it suffices to construct an injective simplicial map from a sub-graph of \widehat{Y} containing B_k to \widehat{Y}_j for any sufficiently large j , which is a local homeomorphism on a ball of slightly smaller radius. Note that B_k intersects finitely many horoballs, and we denote by C_k the union of B_k and said horoballs, and by C'_k the intersection of C_k with \check{Y} . We claim that the cover $r_j : \check{Y} \rightarrow \check{Y}_j$ restricted to C'_k is injective for all sufficiently large j . Indeed, because of convergence of the covers $\check{Y}_i \rightarrow \check{Y}$, the restriction to $B_k \cap \check{Y}$ of r_j is injective for all sufficiently large j . Moreover, as in the proof of Lemma 8.2 each of the horospheres of C'_k maps isometrically to \check{Y}_j for any sufficiently large j .

For any j such that r_j restricted to C'_k is injective, we claim that C_k embeds in \widehat{Y}_j . Indeed, C'_k does not intersect any completion of shells in \check{Y}_j (since C'_k contains all the horospheres it intersects), so that C'_k embeds in \widehat{Y}_j . Since this embedding of C'_k maps horospheres to horospheres, the embedding extends across the horoballs, giving the required embedding $C_k \rightarrow \widehat{Y}_j$. To see that the embedding is a local homeomorphism, notice that on C'_k it is a composition of the restriction of a covering map and a homeomorphism, and the extension to the horoballs preserves being a local homeomorphism. \square

Corollary 8.5. *The space \widehat{Y} is D_1 -simply connected and δ_2 -hyperbolic and δ_2 -visible.*

Proof. Note that \widehat{Y} is clearly connected and σ_0 -modeled on the $\{\widehat{Y}_i\}$, so by the construction and properties of the \widehat{Y}_i , the space \widehat{Y} is σ_0 -modeled on $\{X_0^{\text{cusp}}, \mathcal{H}\}$. Thus by Lemma 7.12, to prove that \widehat{Y} is δ_2 -hyperbolic and δ_2 -visible it suffices to prove that \widehat{Y} is D_1 -simply-connected. Therefore, suppose that σ_0 is a loop in \widehat{Y} . Then σ_0 lies in some ball $B_k(y)$. By strong convergence we can interpret σ_0 as a loop in \widehat{Y}_i for all sufficiently large i . These loops can be filled by a coarse disk which lies entirely in some $B_{k'}(y_i)$, where k' depends on the length of σ_0 , but not on i , since all of the \widehat{Y}_i are δ_2 -hyperbolic. For sufficiently large i , the ball $B_{k'}(y_i)$ is isometric to $B_{k'}(y)$, which means that this filling can be interpreted as a filling in \widehat{Y} . This proves that \widehat{Y} is D_1 -simply-connected, and completes the proof. \square

Recall from Definition 4.1 that for Γ any graph, and $D > 0$, we write Γ^D for the 2-complex obtained by attaching disks to all loops of length at most D .

Lemma 8.6. *The space \check{Y}^{D_1} is simply-connected. In particular, \check{Y}^{D_1} is the universal cover of $\check{X}_0^{D_1}$.*

Proof. This follows quickly from the fact that horospheres are D_1 -simply-connected. Take a loop in \check{Y} , fill it in \widehat{Y} , and then replace those parts of the fillings which lie in horoballs with fillings on the horospheres. \square

Note that since the construction of \check{X}_0 from X_0 is G -equivariant, G acts freely and cocompactly on \check{X}_0 , hence on $\check{X}_0^{D_1}$. Let $Z_{D_1} = G \backslash \check{X}_0^{D_1}$. We now have a sequence of regular covering maps:

$$\check{Y}^{D_1} \rightarrow \check{X}_0^{D_1} \rightarrow Z_{D_1},$$

where the first map extends q and the second map has deck group G .

Definition 8.7. Let \widehat{G} be the deck group of the cover $\check{Y}^{D_1} \rightarrow Z_{D_1}$. That is, let \widehat{G} be the fundamental group of Z_{D_1} .

The group \widehat{G} acts freely and cocompactly on the space \check{Y}^{D_1} , since it is a covering action with (compact) quotient Z_{D_1} . Thus \widehat{G} also acts freely and cocompactly on \check{Y} . Moreover, this action preserves the family of horospheres in \check{Y} , which means that \widehat{G} also acts freely by isometries on the δ_2 -hyperbolic space \widehat{Y} . The quotient $\widehat{Z} := \widehat{G} \backslash \widehat{Y}$ is clearly quasi-isometric to a ray.

Definition 8.8. Choose a horoball \mathcal{H} in \widehat{Y} and let $P = \text{Stab}_{\widehat{G}}(\mathcal{H})$.

Theorem 8.9. *The pair (\widehat{G}, P) is relatively hyperbolic, and $\partial(\widehat{G}, P) = \partial\widehat{Y}$.*

Proof. We claim that the action of \widehat{G} on \widehat{Y} satisfies Gromov's definition of relative hyperbolicity (see Definition 3.3). We know that \widehat{Y} is hyperbolic and the quotient by the group action is quasi-isometric to a ray, which is in fact the image of a vertical ray r in a copy of \mathcal{H} in \widehat{Y} (which we identify with \mathcal{H} below). What we are left to show is that \mathcal{H} is coarsely the sub-level set of a horofunction h based at the limit point of r . Notice that all vertical rays in \mathcal{H} are geodesic rays in \widehat{Y} with the same point at infinity. Therefore the value of h along any such ray r' coincides up to an additive constant with minus the depth plus some value $v_{r'}$. The values $v_{r'}$ can be chosen independently of r' since any two vertical rays come within bounded distance of each other and horofunctions are coarsely Lipschitz. This easily implies the claim that \mathcal{H} is coarsely the sub-level set of a horofunction, as required. \square

Since \widehat{G} is the deck group of the regular cover $\check{Y}^{D_1} \rightarrow Z_{D_1}$ and G is the deck group of the intermediate (regular) cover $\check{X}_0^{D_1} \rightarrow Z_{D_1}$, there is a natural quotient map $\Omega : \widehat{G} \rightarrow G$ whose kernel is canonically identified with $\pi_1(\check{X}_0^{D_1})$.

We have found a relatively hyperbolic group \widehat{G} mapping onto G . To show it is a drilling, we still need to show that the kernel of $\widehat{G} \rightarrow G$ is normally generated by a subgroup of P , and other related properties.

Note that each completed shell in X_0 gives a conjugacy class of subgroups of $\pi_1(\check{X}_0^{D_1})$. We refer to any such subgroup as a *shell group*. By Corollary 6.36 (and the definition of coarse fundamental group) each shell group is isomorphic to the integers.

Lemma 8.10. *The peripheral subgroup P is isomorphic to the \mathbb{Z} -by- \mathbb{Z} group E from Lemma 7.21. Moreover, the intersection of P with the kernel of the quotient $\widehat{G} \rightarrow G$ is a shell group N , and the image of P in G under the quotient map is a conjugate of $\langle g \rangle$.*

Proof. This follows from the proof of Lemma 7.21 and from restricting the covers $\check{Y}^{D_1} \rightarrow \check{X}_0^{D_1}$ and $\check{X}_0^{D_1} \rightarrow Z_{D_1}$ to the horosphere corresponding to \mathcal{H} , which is naturally identified with $\widetilde{\mathcal{S}}$. We have the following commutative diagram, where the vertical arrows are covers, and the horizontal arrows are inclusions. Furthermore the deck group actions on the left extend to the deck group actions on the right.

$$(15) \quad \begin{array}{ccc} \langle a \rangle \curvearrowright & \widetilde{\mathcal{S}}^{D_1} & \hookrightarrow \check{Y}^{D_1} \\ & \downarrow & \downarrow \\ \langle g \rangle \curvearrowright & \mathcal{S}^{D_1} & \hookrightarrow \check{X}_0^{D_1} \\ & \downarrow & \downarrow \\ & \langle g \rangle \backslash \mathcal{S}^{D_1} & \hookrightarrow Z_{D_1} \end{array}$$

By Lemma 7.21, the group E describes the deck group of the composition of the covers on the left. \square

We now show that $\pi_1(\check{X}_0^{D_1})$ is generated by the shell groups, and that these are all conjugate in \widehat{G} .

Lemma 8.11. *$\pi_1(\check{X}_0^{D_1})$ is generated by shell groups.*

Proof. Let

$$\text{CS} = \bigcup_{h \in G \setminus \langle g \rangle} \mathcal{S}^{X_0}(h \cdot \gamma_0),$$

and let

$$\text{CT} = \text{CS} \cup \left(\bigcup_{h \in G \setminus \langle g \rangle} T_{R_0}(h \cdot \gamma_0) \right).$$

We claim first that the space $X_0 \cup \text{CS}$ is $16\delta_0$ -simply connected. Indeed, let α be a loop in this space. By the definition of the completed shells, α may be homotoped to a loop in X_0 across disks of circumference at most $2s = 16\delta_0$. The space X_0 is δ_0 hyperbolic, so by Lemma 4.4 it is $16\delta_0$ -simply-connected. The claim follows.

Next we claim that the space

$$(16) \quad \check{X}_0^{D_1} \cup_{\text{CS}^{D_1}} \text{CT}^{D_1}$$

is simply connected. Let β be a loop in the 1–skeleton of this space. This loop is also a loop in the 1–skeleton of $X_0 \cup \text{CS}$. Since that space is $16\delta_0$ –simply connected, there is a $16\delta_0$ –filling of β . We now subdivide the 2–cells of this filling to obtain a D_1 –filling of β in the space from (16). Let α be one of the 2–cell boundaries of the filling. If it lies completely in \check{X}_0 or completely in CT we do not need to modify it. Otherwise, we consider a maximal sub-segment lying in some component of CT. The endpoints lie in a completed shell, and are distance $\leq 16\delta_0$ from one another, so they may be connected by a path in that completed shell of length at most $\Phi(16\delta_0)$, by Lemma 6.11. We can then replace our 2–cell boundary by one inside a component of CT and one which has one fewer component not in \check{X}_0 . This has length less than $16\delta_0 + \Phi(16\delta_0)$. Repeating this process, we obtain a subdivision into loops of length at most D_1 inside a component of CT and a single loop of length at most

$$16\delta_0 + 16\delta_0\Phi(16\delta_0)$$

in \check{X}_0 . This quantity is less than D_1 so the loop bounds a disk in $X_0^{D_1}$. Applying this subdivision process to all the loops in our original $16\delta_0$ –filling gives a contraction in the space described in (16).

By Seifert–van Kampen the fundamental group of $\check{X}_0^{D_1}$ is generated by the fundamental groups of the components of CS^{D_1} , in other words by the shell groups. \square

Lemma 8.12. *The shell groups are all conjugate in \widehat{G} .*

Proof. This follows from the fact that there is only one G –orbit of shells in \check{X}_0 . A specific conjugate of a shell group corresponds to a choice of path from a basepoint in \check{X}_0 to a shell, and given two such paths the corresponding paths give a loop in Z_{D_1} ; this loop represents the conjugator in \widehat{G} (thought of as $\pi_1(Z_{D_1})$) between the two shell groups. \square

From the previous two lemmas, $\pi_1(\check{X}_0^{D_1})$, considered as a subgroup of \widehat{G} , is normally generated by any single shell group. Choose the shell group N that lies in P . Then $\widehat{G}/\langle\langle N \rangle\rangle = G$.

Thus, we have proved that \widehat{G} is a drilling of G along g as required by Theorem A.

We must prove that P is free abelian if and only if the action of g on the boundary is orientation-preserving. This follows because by Lemma 7.21

P is the fundamental group of

$$\langle g \rangle \backslash (\partial X_0 \setminus \{\gamma_0^\pm\})$$

which is a torus, or a Klein bottle, according to whether the action of g on $\partial X_0 \setminus \{\gamma_0^\pm\}$ is, or is not, orientation-preserving.

To finish the proof of Theorem A it remains to prove that the Bowditch boundary of (\widehat{G}, P) is a 2-sphere, and also the final assertion about torsion (or lack thereof).

9. LINEARLY CONNECTED LIMITS OF SPHERES

In this section we prove the following criterion for a weak Gromov-Hausdorff limit to be S^2 :

Theorem 9.1. *Suppose that (M_i, d_i) are L -linearly connected metric spaces homeomorphic to S^2 , and that the metric space (M, d) is a weak Gromov-Hausdorff limit of the M_i . Suppose furthermore that:*

- (1) M is nondegenerate,
- (2) M is compact,
- (3) M has no cut-points,
- (4) (Self-similarity) Given any $p \in M$ and any open set $U \subseteq M$, there exists a neighborhood V of p and an open embedding $V \rightarrow U$.

Then M is homeomorphic to S^2 .

We start with a general lemma.

Lemma 9.2. *Let (M_i, d_i) be L -linearly connected spaces that weakly Gromov-Hausdorff converge to the compact metric space M . Then M is linearly connected, connected, and locally path-connected.*

Proof. A linearly connected space is both connected and locally path-connected. Indeed, if B is a ball of radius ε around a point x , then L -linear connectivity allows one to construct a path-connected $C \subset B$ containing the ε/L -ball B' around x . This set C is the union of the path-connected sets of controlled diameter containing the pairs $\{x, y\}$ for $y \in B'$. Therefore, it suffices to prove that M is linearly connected.

By assumption, there exist λ and $\varepsilon_i \rightarrow 0$ such that there exist (λ, ε_i) -quasi-isometries $\phi_i : M_i \rightarrow M$ and $\psi_i : M \rightarrow M_i$ that are ε_i -quasi-inverses of each other. Given any $p, q \in M$, consider connected sets $C_i \subseteq M_i$ of diameter at most $Ld_i(\phi_i(p), \phi_i(q)) \leq L\lambda d(p, q) + L\varepsilon_i$ containing $\phi_i(p), \phi_i(q)$. We can then consider the closures of $\psi_i(C_i)$, and consider a Hausdorff limit D of a subsequence. It is easy to see that D is connected, contains p and q , and has diameter bounded linearly in terms of $d(p, q)$. \square

The following two propositions are (essentially) proven in [25].

Proposition 9.3. *Suppose that (M_i, d_i) are L -linearly connected metric spaces homeomorphic to S^2 , and that the compact metric space (M, d) is a weak Gromov-Hausdorff limit of the M_i . Either M has a cut-point or it is planar.*

Proof. By [25, Lemma 3.9], non-planar graphs cannot topologically embed in M . Moreover, M is connected and locally connected by Lemma 9.2. Hence, M is a Peano continuum that does not contain any non-planar graphs, and therefore by [16] either it has a cut-point or it is planar. \square

Definition 9.4. [25, Definitions 8.1, 8.2] Let M be a metric space, and let $f: S^1 \rightarrow M$. For $\varepsilon > 0$, an ε -filling of f is a triangulation of the unit disk D^2 , along with a (not necessarily continuous) extension of f to $\bar{f}: D^2 \rightarrow M$ so that every simplex of the triangulation is mapped by \bar{f} to a set of diameter at most ε .

We say that M is *weakly simply connected* if, for every $\varepsilon > 0$, every continuous $f: S^1 \rightarrow M$ has an ε -filling.

The second result is an immediate consequence of [25, Theorem 8.6].

Proposition 9.5. *Suppose that (M_i, d_i) are simply connected L -linearly connected metric spaces, and that the compact metric space (M, d) is a weak Gromov-Hausdorff limit of the M_i . Then M is weakly simply connected.*

Observe that if a compact metric space M is weakly simply connected, and if M' is a metric space homeomorphic to M , then M' is also weakly simply connected. Thus we can speak about weakly simply connected compact spaces.

Lemma 9.6. *Let M be a subspace of S^2 (as topological spaces). Suppose that:*

- (1) M is nondegenerate,
- (2) M is compact,
- (3) M has no cut points,
- (4) M is connected and locally path connected,
- (5) M is weakly simply connected.

Then M contains an open subset homeomorphic to \mathbb{R}^2 .

Proof. We note first that (4) implies that M is path connected and that (4) and (3) together imply that M has no path cut point.

We consider the standard round metric d_{S^2} on S^2 (that is, of diameter π), and the induced metric on M .

Since M is nondegenerate, there are distinct points x and y in M . Since it is path connected, there is a path p in M between x and y . Take a third point z on the path. Since z is not a path cut point, there is another path p' in M between x and y . The union of p and p' contains an embedded circle C .

Jordan's theorem ensures that C cuts S^2 into two disks D_1 and D_2 . We claim that M contains at least one of the disks, which proves that M contains an open subset homeomorphic to \mathbb{R}^2 , as required.

Otherwise, there are points $p_1 \in D_1 \setminus M$, $p_2 \in D_2 \setminus M$. Since M is compact, the distance between $\{p_1, p_2\}$ and M is realized. Choose a positive ε so that

$$\varepsilon < \min \{ \pi/2, d_{S^2}(M, \{p_1, p_2\})/2 \}.$$

Let $f: S^1 \rightarrow M$ be so that $f(S^1) = C$. Since M is weakly simply connected, there is an ε -filling $\bar{f}: D^2 \rightarrow M$ of f . Since $\varepsilon < \pi/2$, the \bar{f} -images of every two adjacent vertices in the triangulation of D^2 can be connected by a geodesic segment in S^2 of length at most ε . The resulting triangles can be filled in without increasing their diameter. So there exists a continuous map $g: D^2 \rightarrow S^2$ such that $g|_{S^1} = f$ and $g(D^2)$ lies in the 2ε -neighbourhood of M . But then $g(D^2)$ must contain one of the two disks D_1 or D_2 , and hence either p_1 or p_2 , a contradiction. \square

We are ready to prove the theorem.

Proof of Theorem 9.1. By Lemma 9.2, M is connected and locally path-connected.

By Proposition 9.3, M embeds in S^2 , since it does not have cut-points. Moreover, it is weakly simply connected by Proposition 9.5. But then Lemma 9.6 applies, so M contains an open subset homeomorphic to \mathbb{R}^2 . By self-similarity, each point of M has an open neighborhood homeomorphic to \mathbb{R}^2 . Hence, any embedding $M \rightarrow S^2$ is open. Such an embedding is also closed (since M is compact and S^2 is Hausdorff), so it is a homeomorphism. \square

10. PROOF OF THEOREM A

As noted at the end of Section 8, it remains to show that the Bowditch boundary of (\widehat{G}, P) is a 2-sphere, and the statement about torsion. We retain the notation from the previous sections.

By Theorem 8.9 we have to show that $\partial\widehat{Y}$ is homeomorphic to the 2-sphere.

By Lemma 8.4, the spaces (\widehat{Y}_i, y_i) strongly converge to (\widehat{Y}, \tilde{y}) . By the construction of the \widehat{Y}_i , they are all δ_2 -hyperbolic and δ_2 -visible, so by Proposition 3.13 the spaces $\partial\widehat{Y}_i$ weakly Gromov-Hausdorff converge to $\partial\widehat{Y}$ (when all are equipped with δ_2 -adapted visual metrics at the given basepoints).

It is clear that $\partial\widehat{Y}$ is infinite (since the parabolic subgroup has infinite index, as can be seen in the quotient G of \widehat{G}) and compact. Moreover it is self-similar (as in Theorem 9.1), since it is the boundary of a relatively hyperbolic group with infinite index peripheral subgroups. Therefore, the

endpoints of loxodromic elements are dense in $\partial\widehat{Y}$. Since each of the \widehat{Y}_i is δ_2 -hyperbolic and Σ_0 -modeled on $\{\text{UG}(X_0, \gamma_0)\}$, by Corollary 7.30 there is an L such that the spaces $\partial\widehat{Y}_i$ are all L -linearly connected. We now show that each of the $\partial\widehat{Y}_i$ is homeomorphic to S^2 .

Proposition 10.1. *Each of $\partial\widehat{Y}_i$ is homeomorphic to S^2 .*

Proof. We prove this by induction on i . In the base case $i = 1$ we have $\widehat{Y}_1 = \text{UG}(X_0, \gamma_0)$, whose boundary is homeomorphic to S^2 by Corollary 7.24.

Suppose now that $i \geq 1$ and that $\partial\widehat{Y}_i$ is homeomorphic to S^2 . Then $\widehat{Y}_{i+1} = \widehat{Y}_i$ (and then $\partial\widehat{Y}_{i+1}$ is homeomorphic to S^2) or $\widehat{Y}_{i+1} = \text{UG}(\widehat{Y}_i, c_i)$, for a core c_i of a tube in \widehat{Y}_i .

Notice that $\partial\widehat{Y}_i \setminus \{c_i^\pm\}$ is an (open) annulus. By Corollary 7.30 we know that $\partial\text{UG}(\widehat{Y}_i, c_i)$ is uniformly linearly connected (and, in particular, connected). Hence, by Proposition 5.6, $\partial\text{UG}(\widehat{Y}_i, c_i)$ does not have cut-points, so $\partial\text{UG}(\widehat{Y}_i, c_i) \setminus \{p_i\}$ is connected, where p_i is the point at infinity of the new horoball in $\text{UG}(\widehat{Y}_i, c_i)$.

In view of Proposition 7.23 we have that $\partial\text{UG}(\widehat{Y}_i, c_i) \setminus \{p_i\}$ is a connected \mathbb{Z} -cover of an annulus, and therefore it is homeomorphic to \mathbb{R}^2 . As in the proof of Corollary 7.24, $\partial\text{UG}(\widehat{Y}_i, c_i)$ is a compact metric space, and removing a point makes it homeomorphic to \mathbb{R}^2 . So $\partial\text{UG}(\widehat{Y}_i, c_i)$ is homeomorphic to S^2 , and the induction is complete. \square

Proposition 5.7 now implies that $\partial\widehat{Y}$ has no cut-points. By Lemma 9.2, $\partial\widehat{Y}$ is locally connected. It now follows from Theorem 9.1 that $\partial\widehat{Y}$ is homeomorphic to S^2 , as required.

Finally, we consider the issue of torsion in G and \widehat{G} . The following statement immediately implies the final assertion of Theorem A (and in fact proves somewhat more).

Proposition 10.2. *Let $\Omega : \widehat{G} \rightarrow G$ be the quotient map described in Section 8. Then Ω is injective on any finite subgroup of \widehat{G} . Conversely, if F is a finite subgroup of G then there exists a finite subgroup $\widehat{F} \subset \Omega^{-1}(F)$ so that $\Omega|_{\widehat{F}} : \widehat{F} \rightarrow F$ is an isomorphism.*

Proof. Let \widehat{F} be a finite subgroup of \widehat{G} . Then \widehat{F} acts isometrically on the δ_2 -hyperbolic space \widehat{Y} . By [24, Lemma 4.2] there is an \widehat{F} -orbit $\widehat{F} \cdot \mathfrak{o}$ of diameter at most $4\delta_2 + 2$ in \widehat{Y} .

Suppose that \widehat{F} is nontrivial. Then it cannot be contained in a stabilizer of a horoball in \widehat{Y} , by Lemma 8.10. However, the horoballs in \widehat{Y} are Σ_1 -separated, so if $\widehat{F} \cdot \mathfrak{o}$ lies within $20\delta_2$ of a horoball, \widehat{F} must in fact stabilize

that horoball. It follows that $\widehat{F} \cdot o$ lie in a ball B in \widehat{Y} which does not lie within $15\delta_2$ of any horoball.

By the definition of \widehat{Y} as a strong limit of the \widehat{Y}_j , and by induction using Lemma 7.9, the ball B maps injectively to a ball \overline{B} in \check{X}_0 (since B is not close to any horoball).

Recall that $\widehat{G} = \pi_1(Z_{D_1})$, and that G corresponds to the deck group of the cover $\check{X}_0^{D_1} \rightarrow Z_{D_1}$ which is intermediate to the universal cover $\check{Y}^{D_1} \rightarrow Z_{D_1}$. The nontrivial elements of \widehat{F} are represented by loops in Z_{D_1} . These loops elevate to paths in \check{X}_0 which are not loops because B maps injectively to the ball \overline{B} in \check{X}_0 . Thus, Ω is injective on \widehat{F} .

Conversely, consider a finite subgroup F of G . Then F has an orbit $F \cdot o$ of diameter at most $4\delta_0 + 2$ in X_0 . If F is nontrivial the assumption that translates of γ_0 are Σ -separated forces a ball around $\widehat{F} \cdot o$ to lie in \check{X}_0 . This ball lifts to a ball B in \check{Y} , since this is a D_1 -cover, and $D_1 \gg 4\delta_0 + 2$. The lift of $\widehat{F} \cdot o$ in B defines a collection of deck transformations in \widehat{G} which forms a section of Ω on F and also forms a subgroup of \widehat{G} , completing the proof. \square

The proof of Theorem A is now complete.

11. THE COHEN–LYNDON PROPERTY

In this section we argue that the map from (\widehat{G}, P) to G is a “long filling” in a reasonable sense. Specifically, it satisfies the Cohen–Lyndon property, defined by Sun in [38, Definition 2.2]

Definition 11.1. A triple of groups (A, B, C) with $C \triangleleft B < A$ satisfies the *Cohen–Lyndon property* if there is a left transversal T of $B\langle\langle C \rangle\rangle$ in A so that

$$\langle\langle C \rangle\rangle = \ast_{t \in T} tCt^{-1}.$$

(To be more precise, the statement is that the natural map from the group on the right-hand side to $\langle\langle C \rangle\rangle$ is an isomorphism.)

Dahmani–Guirardel–Osin [18, Theorem 2.27] prove that the kernel of a long Dehn filling is a free product of conjugates of filling kernels, but do not give an explicit indexing set. Sun clarifies that the triple (A, B, C) satisfies the Cohen–Lyndon property whenever (A, B) is relatively hyperbolic and C is a subgroup of B which is “sufficiently long” in the sense it misses finitely many (short) elements. (Both Sun and Dahmani–Guirardel–Osin work in a more general setting, where B is (weakly) hyperbolically embedded in A .) In the relatively hyperbolic case this was proved in [25, Theorem 4.8]. In fact the proof there yields the following quantitative statement:

Theorem 11.2. *Suppose that (A, B) is a relatively hyperbolic group pair, with a geometrically finite action on a θ -hyperbolic space Υ , for some $\theta \geq 1$, containing an equivariant family of convex horoballs which are $10^3\theta$ -separated. Suppose also (Very Translating Condition) that $C \triangleleft B$ is chosen so that if H is the horoball stabilized by B , then $d_\Upsilon(x, n \cdot x) \geq 10^4\theta$ for every $n \in C \setminus \{1\}$ and $x \in \Upsilon \setminus H$.*

Then (A, B, C) satisfies the Cohen–Lyndon property.

Two remarks are in order. First, the Cohen-Lyndon property and Theorem 11.2 can be reformulated for triples $(A, \mathcal{B}, \mathcal{C})$ where \mathcal{B} is a family of subgroups and \mathcal{C} a family of subgroups of the form $C_B \triangleleft B$ for $B \in \mathcal{B}$ (see [38, Definition 3.13]). All the theorems mentioned above are known in this generality. But for the current purposes it is sufficient to consider a single subgroup.

Secondly, there is a misleading comment in [25] near the beginning of Section 4. It is there claimed [25, Lemma 4.1] that the Very Translating Condition is a *consequence* of being a sufficiently long filling, and the relatively hyperbolic Dehn filling theorem is referenced. In fact, the Very Translating Condition is only about the filling kernels, so it is satisfied whenever those filling kernels contain no short nontrivial elements of the parabolics. It is a consequence of the proof that the normal subgroup generated by the filling kernels also contains no such short elements.

We apply Theorem 11.2 to the setting of the current paper as follows.

Theorem 11.3. *Let (\widehat{G}, P) be the relatively hyperbolic pair from Theorem 8.9, and let $N \triangleleft P$ be the shell group described in Lemma 8.10. The triple (\widehat{G}, P, N) satisfies the Cohen-Lyndon property.*

Proof. We have built a δ_2 -hyperbolic space \widehat{Y} on which (\widehat{G}, P) acts geometrically finitely, so that P preserves a horoball \mathcal{H} . The family of translates of \mathcal{H} is Σ_1 -separated, and Σ_1 is much larger than $10^3\delta_2$, but the horoballs \mathcal{H} may not be convex. We replace them by nested $2\delta_2$ -horoballs $g\mathcal{H}' \subset g\mathcal{H}$, which *are* convex (see [23, Lemma 3.26] – the proof there applies verbatim). These nested horoballs are obviously still $10^3\delta_2$ -separated. (Here we are using that $\Sigma_1 \geq 10R_0 \geq 10 \cdot 2\sigma_0 \geq 20 \cdot 10^5 D_1 \geq 2 \cdot 10^6 \cdot 2Q_0 \geq 4 \cdot 10^6 \delta_2 > 10^3 \delta_2$.)

Next we verify the Very Translating Condition. Let x be a point outside \mathcal{H}' and let $n \in N \setminus \{1\}$. Let y be a closest point to x on ∂H .

Now we note that our systole requirement ($\text{sys}_0 = 2^{25\sigma_0}Q_0$) together with Lemma 4.9 gives

$$\begin{aligned} d_{\mathcal{H}}(y, n \cdot y) &\geq 2 \log_2 \text{sys}_0 - 2 \\ &\geq 20 \cdot 10^5 D_1 \\ &\geq 40 \cdot 10^5 Q_0 \\ &\geq 4 \cdot 10^6 \delta_2. \end{aligned}$$

It is not hard to see that if $d_{\mathcal{H}}(a, b) \geq 10\delta_2$, then the \widehat{Y} -distance is actually realized in \mathcal{H} . Moreover a simple argument gives $d_{\widehat{Y}}(x, n \cdot x) \geq d_{\widehat{Y}}(y, n \cdot y) - 8\delta_2$.

Thus the above estimate gives $d_{\widehat{Y}}(y, n \cdot y) \geq 4 \cdot 10^6 \delta_2$, so $d_{\widehat{Y}}(x, n \cdot x) \geq (4 \cdot 10^6 - 8)\delta_2 > 10^4 \delta_2$, thus verifying the Very Translating Condition. \square

12. NOTATION INDEX

In this section, we collect (some of) the notation from this paper, with links to where it is introduced, and some discussion about how it is used. Hopefully this will help the reader.

The reader may also want to consult Notation 7.2, where many of our constants are fixed.

- M is a general compact metric space.
- Γ is reserved for a general connected graph.
- Z denotes a general metric space.
- Y denotes a general Gromov hyperbolic metric space (with constant of hyperbolicity always locally specified, often δ).
- X_0 is the δ_0 -hyperbolic space we start with. The constant δ_0 is reserved for this specific value. See Assumption 6.1. From Section 8 on we consider a group G acting geometrically on X_0 .
- Inside X_0 is a λ_0 -quasi-convex subset Y_0 , with a fixed basepoint $w_0 \in Y_0$ (see also 6.1).

In Assumption 7.1, we further assume that $\partial X_0 \cong S^2$, and that Y_0 is a quasi-geodesic axis γ_0 for a loxodromic element g . It is $\langle g \rangle$ that we drill along in the proof of Theorem A. We make the assumption about the separation of the translates of γ_0 at the beginning of Section 8.

- For $K > 0$ an integer, the *completed K -shell* about Y_0 is obtained by taking all the points at distance K from Y_0 , and adding edges of length at most $8\delta_0$ to make a connected graph, denoted $\mathcal{S}_K(Y_0)$. See Definition 6.12 and Convention 6.13.

- The *completed tube complement*, $\mathcal{C}_K(Y_0)$ is obtained by removing the tube about Y_0 and adding the edges from the completed shell. See Definition 6.15 and Convention 6.16.
- $X_0^{\text{cusp}}(K)$ is obtained by gluing a combinatorial horoball \mathcal{D} onto the copy of $\mathcal{S}_K(Y_0)$ in $\mathcal{C}_K(Y_0)$. See Definition 6.23.
- In Notation 7.2, many constants are fixed, including the radius R_0 of the tubes that are considered from that point forwards. Further separation constants are fixed at the beginning of Section 8.
- \mathcal{S} is the completed R_0 -shell about γ_0 in X_0 , and \mathcal{C} is the associated completed tube complement. These spaces have canonical regular \mathbb{Z} -covers $\tilde{\mathcal{S}}$ and $\tilde{\mathcal{C}}$ which are D_1 -simply-connected (for a constant D_1 fixed in 7.2). See Notation 7.5.
- \mathcal{H} denotes the combinatorial horoball based on $\tilde{\mathcal{S}}$. See Notation 7.5.
- (Y, γ) is some space which is R -tube comparable (see Definition 6.39) to (X_0, γ_0) .
- $\text{UG}(Y, \gamma)$ means: Unwrap a completed tube complement $\mathcal{C}_{R_0}^Y(\gamma)$ and glue on a copy of \mathcal{H} . See Construction 7.7.
- \check{X}_0 is obtained by removing an equivariant family of tubes from X_0 and gluing on corresponding copies of $\mathcal{S}_{R_0}(\gamma_0)$. See Subsection 8.2.
- We form the space $\check{X}_0^{D_1}$ by attaching disks to edge-loops of length D_1 , see Definition 4.1. Then the quotient space is $Z_{D_1} = G \backslash \check{X}_0^{D_1}$. This is analogous to the manifold exterior in the 3-manifold case.
- The space \check{Y} is a limit of certain covers of \check{X}_0 . See Lemma 8.1.
- The space \hat{Y} is the space \check{Y} with an equivariant family of horoballs added. See Definition 8.3.

In Section 8 we construct the group \hat{G} in the conclusion of Theorem A and prove in that the space \hat{Y} admits a geometrically finite \hat{G} -action which exhibits the relative hyperbolicity of our drilled group.

REFERENCES

- [1] I. Belegradek. Complex hyperbolic hyperplane complements. *Math. Ann.*, 353(2):545–579, 2012.
- [2] I. Belegradek. Rigidity and relative hyperbolicity of real hyperbolic hyperplane complements. *Pure Appl. Math. Q.*, 8(1):15–51, 2012.
- [3] M. Bestvina. Local homology properties of boundaries of groups. *Michigan Math. J.*, 43(1):123–139, 1996.
- [4] M. Bestvina and G. Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.
- [5] M. Bonk and B. Kleiner. Quasisymmetric parametrizations of two-dimensional metric spheres. *Invent. Math.*, 150(1):127–183, 2002.

- [6] M. Bonk and B. Kleiner. Quasi-hyperbolic planes in hyperbolic groups. *Proc. Amer. Math. Soc.*, 133(9):2491–2494 (electronic), 2005.
- [7] M. Bonk and O. Schramm. Embeddings of Gromov hyperbolic spaces. *Geom. Funct. Anal.*, 10(2):266–306, 2000.
- [8] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.
- [9] B. H. Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific J. Math.*, 269(2):269–280, 2014.
- [10] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [11] J. F. Brock and K. W. Bromberg. On the density of geometrically finite Kleinian groups. *Acta Math.*, 192(1):33–93, 2004.
- [12] D. Calegari and K. Fujiwara. Stable commutator length in word-hyperbolic groups. *Groups Geom. Dyn.*, 4(1):59–90, 2010.
- [13] D. Calegari and D. Gabai. Shrinkwrapping and the taming of hyperbolic 3-manifolds. *J. Amer. Math. Soc.*, 19(2):385–446, 2006.
- [14] R. D. Canary. Ends of hyperbolic 3-manifolds. *J. Amer. Math. Soc.*, 6(1):1–35, 1993.
- [15] J. W. Cannon. The theory of negatively curved spaces and groups. In *Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989)*, Oxford Sci. Publ., pages 315–369. Oxford Univ. Press, New York, 1991.
- [16] S. Claytor. Topological immersion of Peanian continua in a spherical surface. *Ann. of Math. (2)*, 35(4):809–835, 1934.
- [17] R. Coulon. On the geometry of Burnside quotients of torsion free hyperbolic groups. *Internat. J. Algebra Comput.*, 24(3):251–345, 2014.
- [18] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.
- [19] F. Dahmani and A. Yaman. Bounded geometry in relatively hyperbolic groups. *New York J. Math.*, 11:89–95, 2005.
- [20] A. Dasgupta and G. C. Hruska. Local connectedness of boundaries for relatively hyperbolic groups. Preprint, arXiv:2204.02463, 2022.
- [21] É. Ghys and P. de la Harpe, editors. *Sur les groupes hyperboliques d’après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1990. Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988.
- [22] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [23] D. Groves and J. F. Manning. Dehn filling in relatively hyperbolic groups. *Israel Journal of Mathematics*, 168:317–429, 2008.
- [24] D. Groves and J. F. Manning. Dehn fillings and elementary splittings. *Trans. Amer. Math. Soc.*, 370(5):3017–3051, 2018.
- [25] D. Groves, J. F. Manning, and A. Sisto. Boundaries of Dehn fillings. *Geom. Topol.*, 23(6):2929–3002, 2019.
- [26] B. B. Healy and G. C. Hruska. Cusped spaces and quasi-isometries of relatively hyperbolic groups. Preprint, arXiv:2010.09876, 2020.
- [27] G. C. Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

- [28] G. Italiano, B. Martelli, and M. Migliorini. Hyperbolic 5-manifolds that fiber over S^1 . *Invent. Math.*, 231(1):1–38, 2023.
- [29] M. Kapovich and B. Kleiner. Coarse Alexander duality and duality groups. *J. Differential Geom.*, 69(2):279–352, 2005.
- [30] S. Kojima. Isometry transformations of hyperbolic 3-manifolds. *Topology Appl.*, 29(3):297–307, 1988.
- [31] D. D. Long. Immersions and embeddings of totally geodesic surfaces. *Bull. London Math. Soc.*, 19(5):481–484, 1987.
- [32] J. M. Mackay. Existence of quasi-arcs. *Proc. Amer. Math. Soc.*, 136(11):3975–3981, 2008.
- [33] J. F. Manning and O. H. Wang. Cohomology and the Bowditch boundary. *Michigan Math. J.*, 69(3):633–669, 2020.
- [34] G. J. Martin and R. K. Skora. Group actions of the 2-sphere. *Amer. J. Math.*, 111(3):387–402, 1989.
- [35] P. Ontaneda. Cocompact CAT(0) spaces are almost geodesically complete. *Topology*, 44(1):47–62, 2005.
- [36] D. V. Osin. Peripheral fillings of relatively hyperbolic groups. *Invent. Math.*, 167(2):295–326, 2007.
- [37] T. Sakai. Geodesic knots in a hyperbolic 3-manifold. *Kobe J. Math.*, 8(1):81–87, 1991.
- [38] B. Sun. Cohomology of group theoretic Dehn fillings I: Cohen-Lyndon type theorems. *J. Algebra*, 542:277–307, 2020.
- [39] B. Tshishiku and G. Walsh. On groups with S^2 Bowditch boundary. *Groups Geom. Dyn.*, 14(3):791–811, 2020.
- [40] C. T. C. Wall, editor. *Homological group theory*, volume 36 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge-New York, 1979.

DEPARTMENT OF MATHEMATICS, STATISTICS, AND COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT CHICAGO, 322 SCIENCE AND ENGINEERING OFFICES (M/C 249), 851 S. MORGAN ST., CHICAGO, IL 60607-7045

Email address: groves@math.uic.edu

AIX-MARSEILLE UNIV., CNRS, I2M, UMR 7373, MARSEILLE, FRANCE

Email address: phaissin@math.cnrs.fr

DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14853, USA

Email address: jfmanning@cornell.edu

DEPARTMENT OF MATHEMATICAL SCIENCES, UNIVERSITY OF COPENHAGEN, NØRREGADE 10, 1165 KØBENHAVN, DENMARK; MATHEMATICAL INSTITUTE, UNIVERSITY OF WROCLAW, PL. GRUNWALDZKI 2/4, 50-384 WROCLAW, POLAND

Email address: dosaj@math.uni.wroc.pl

DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY AND MAXWELL INSTITUTE FOR MATHEMATICAL SCIENCES, EDINBURGH, UK

Email address: a.sisto@hw.ac.uk

DEPARTMENT OF MATH, TUFTS UNIVERSITY, MEDFORD, MA 02155, USA

Email address: genevieve.walsh@tufts.edu