

Combining relatively hyperbolic groups over a complex of groups

Darius Alizadeh

October 6, 2025

Abstract

Given a complex of groups $G(\mathcal{Y}) = (G_\sigma, \psi_a, g_{a,b})$ where all G_σ are relatively hyperbolic, the ψ_a are inclusions of full relatively quasiconvex subgroups, and the universal cover X is CAT(0) and δ -hyperbolic, we show $\pi_1(G(\mathcal{Y}))$ is relatively hyperbolic. The proof extends the work of Dahmani and Martin by constructing a model for the Bowditch boundary of $\pi_1(G(\mathcal{Y}))$. We prove the model is a compact metrizable space on which G acts as a geometrically finite convergence group, and a theorem of Yaman then implies the result. More generally, this model shows how any suitable action of a relatively hyperbolic group on a simply connected cell complex encodes a decomposition of the Bowditch boundary into the boundary of the cell complex and the boundaries of cell stabilizers. We hope this decomposition will be helpful in answering topological questions about Bowditch boundaries.

Contents

1	Introduction	3
2	Background	5
2.1	Relatively Hyperbolic Groups	5
2.2	Quasiconvexity	8
2.3	M_κ -Complexes	15
2.4	Complexes of Groups	17
3	Constructing $\overline{Z}$	24
3.1	Domains and Their Geometry	29
4	Geometric Tools	35
4.1	Nesting and Families	35

4.2	The Crossing Lemma and Refined Families	40
4.3	The Refinement and Star Lemmas	45
5	The Topology of $\overline{Z}$	49
5.1	Definition of the Topology	49
5.2	A basis of neighborhoods	51
6	Properties of the Topology	57
6.1	Independence of basepoint and induced topologies	57
6.2	$\overline{Z}$ is T_0	60
6.3	$\overline{Z}$ is regular	62
6.4	$\overline{Z}$ is compact	67
7	Dynamics of the Action	70
7.1	G is a Convergence Group	70
7.2	G is geometrically finite	82

1 Introduction

This work is another chapter in the story of combination theorems which began with [BF92], where Bestvina and Feighn give conditions on a graph of hyperbolic spaces which guarantee the resulting space is itself hyperbolic. Their conditions deal directly with the metric on the resulting space, and this metric approach has been elaborated on in [Ali05], [MR08] with applications to limit groups and relatively hyperbolic groups. An alternative, dynamical approach was opened by a theorem of Bowditch in [Bow98], which equates geometrically finite convergence groups with relatively hyperbolic groups. Dahmani [Dah03] applied this theorem to prove a combination theorem for graphs of relatively hyperbolic groups and more. He combined Bowditch boundaries over the Bass–Serre tree to build a compact metrizable space on which the fundamental group acts as a geometrically finite convergence group, then applied Bowditch’s theorem to conclude the fundamental group was relatively hyperbolic. Martin [Mar14] pushed this dynamical strategy into arbitrary dimensions by combining boundaries of hyperbolic groups over the development of a complex of groups, and we extend Martin’s work to relatively hyperbolic groups.

Recall from [BH99][I.7] that for $\kappa \in \mathbb{R}$, an M_κ -complex is a cell complex where each n -cell is modeled on a simplex in the model space M_κ^n , which is the unique simply-connected Riemannian manifold of dimension n with constant curvature κ . For example, a cube complex is an M_0 -complex. See Definition 2.28 for details. Here is our main theorem.

Theorem 1.1. Let $G(\mathcal{Y}) = (G_\sigma, \psi_a, g_{a,b})$ be a nonpositively curved developable complex of groups over a scwol $\mathcal{Y}$, where each G_σ is a relatively hyperbolic group and each ψ_a is the inclusion of a full relatively quasiconvex subgroup. Let $\mathcal{X}$ be the universal cover of $G(\mathcal{Y})$, and let X be the geometric realization of $\mathcal{X}$ equipped with an M_κ structure. Suppose X is δ -hyperbolic and the action of $G = \pi_1(G(\mathcal{Y}))$ on X is acylindrical. Then G is relatively hyperbolic. The maximal parabolic subgroups of G are virtually parabolic subgroups of vertex stabilizers, and the stabilizer of each simplex in X is a full relatively quasiconvex subgroup of G .

The relevant terms are explained in Section 2. Just as a group action on a tree induces a graph of groups structure in Bass–Serre theory, a cocompact action on a simply connected cell complex induces a complex of groups over the quotient. These cocompact actions are usually easier to describe and arise more naturally, e.g., groups acting on CAT(0)-cube complexes. Here is a rephrasing of our main theorem from this point of view.

Theorem 1.2. Let X be a δ -hyperbolic, CAT(0) M_κ -complex with $\kappa \leq 0$. Let G be a group acting on X cocompactly and acylindrically. Write G_σ for the point-wise stabilizer of a simplex σ of X . Suppose that for every pair of simplices $\sigma' \subset \sigma$ in X , $G_{\sigma'}$ and G_σ are relatively hyperbolic and G_σ is a full RQC subgroup of $G_{\sigma'}$. Then G is relatively hyperbolic. The maximal parabolic subgroups of G are virtually the maximal parabolic subgroups of vertex stabilizers, and each G_σ is

a full RQC subgroup of G .

Like [Dah03] and [Mar14], we construct a model of the boundary for G by gluing the Bowditch boundaries of simplex stabilizers together and show G acts on our model as a geometrically finite convergence group. Unlike Martin however, we must deal with parabolic points in these boundaries. Much of the work in defining this model and its topology comes from [Mar14]. It is rather complex, and we hope to clarify many of the details involved. Further, while the conclusion of Theorem 1.2 is that a certain group is relatively hyperbolic, the real fruit of this labor is this model. Given a group acting on a tree, Bass–Serre theory explains how the tree encodes the combinatorial structure of the group. Given a relatively hyperbolic group acting on a cell complex, this model shows how the cell complex encodes the boundary of the group. We describe an illustrative example.

Let Σ be the surface of genus 2, c be a geodesic which separates Σ into two tori with boundary, and let $G = \pi_1(\Sigma)$. Since G acts geometrically on the universal cover $\mathbb{H}^2$, G is hyperbolic and $\partial G = S^1$. Let $\gamma \in G$ be the element corresponding to c . Then each conjugate of γ acts on $\mathbb{H}^2$ as a hyperbolic isometry fixing two points in S^1 , namely the endpoints of some lift of c . It turns out G is hyperbolic relative to $\langle \gamma \rangle$, and using a theorem of Tran [Tra13], the Bowditch boundary of $\partial_{\langle \gamma \rangle} G$ is given by contracting the boundary of each parabolic subgroup to a point. So each lift of c to $\mathbb{H}^2$ gives a pair of points in the boundary S^1 , and $\partial_{\langle \gamma \rangle} G$ is given by contracting each pair to a point. Each contraction takes the boundary circle and pinches it together at two points, so the resembles a tree of circles, see Figure 1. If we fix a basepoint $x_0 \in \mathbb{H}^2$, then some geodesics from x_0 to the boundary S^1 will cross infinitely many lifts of c . These points are not in one of the circles of the tree of circles, but correspond to points in the boundary of the underlying tree.

The two tori with boundary and c make Σ into a graph of spaces, which induces a graph of groups structure for G :

$$(F_2, [a_1, b_1]) \longleftarrow \langle \gamma \rangle \longrightarrow (F_2, [a_2, b_2]).$$

The edge group is simply γ , which is hyperbolic relative to itself and has Bowditch boundary a single point. Each vertex group is a free group, and γ maps to the commutator. A free group on two generators is hyperbolic relative to this commutator and its Bowditch boundary is S^1 , which can be seen by taking a hyperbolic structure on the cusped torus. Let X be the Bass–Serre tree of this splitting. Then G acts on X as in Theorem 1.2, and X is exactly the underlying tree of the tree of circles in the previous paragraph.

To construct our model, we take a copy of the Bowditch boundary for each cell of X and glue the boundaries using the inclusion of cell stabilizers and limit sets. In this example, this means a circle for each vertex, a point for each edge, and gluing these circles along those points. The resulting space is almost the tree of circles above, but is not compact – a sequence of points in circles which escape to infinity will not converge. Such a sequence represents a sequence of vertices in X , so we compactify by adding ∂X . The proof of Theorem 1.2 shows how to construct such a




Figure 1: On the left is $\mathbb{H}^2$ tiled by octagons, describing an action of G . The red curves represent lifts of c . In the center is the tree of circles we get by contracting the endpoints of each lift of c . These images come from [BZLW22]. On the right is part of the Bass–Serre tree of the splitting for G . Red and blue vertices are stabilized by conjugates of $[a_1, b_1]$ and $[a_2, b_2]$ respectively.

model in general.

We describe the layout. Section 2 provides the necessary background on relatively hyperbolic groups and M_κ -complexes. Section 3 defines our model $\overline{Z}$ as a set and proves some basic properties about it. Section 4 proves many lemmas which we use throughout the later sections. In Sections 5 and 6, we define a neighborhood basis for points in $\overline{Z}$ and show that it indeed forms the basis for a compact metrizable topology. In section 7, we show G acts as a convergence group on $\overline{Z}$, understand its limit set, and apply Yaman’s Theorem 2.5 to conclude.

2 Background

2.1 Relatively Hyperbolic Groups

For this subsection, let G be a group acting on a compact metrizable space M .

Definition 2.1 (Convergence Groups). Given a sequence $(g_n)_n$ in G and points $\xi_+, \xi_- \in M$, we call the triple (g_n, ξ_+, ξ_-) an *attracting repelling triple* (ART) if for any compact set $K \subset M \setminus \{\xi_-\}$, the sets $g_n K$ converge to ξ_+ uniformly. In this situation, $(g_n)_n$ is called a *convergence sequence* and the points ξ_+ and ξ_- are the *attracting* and *repelling* points of the sequence $(g_n)_n$. We say G acts as a *convergence group* on M if every infinite sequence in G has a subsequence which is a convergence sequence.

If (g_n, ξ_+, ξ_-) is an ART, it is possible that $\xi_+ = \xi_-$. It’s clear that being a convergence group passes to subgroups. Acting as a convergence group is equivalent to acting properly discontinuously

on the space of distinct triples of M [Bow99].

Definition 2.2 (Conical and Parabolic Limit Points). Let G be a convergence group on M . A point $\xi \in M$ is a *conical limit point* if there exists a sequence $(g_n)_n$ in G and points $\xi_+ \neq \xi_-$ in G such that $g_n \xi \rightarrow \xi_-$ and $g_n \xi' \rightarrow \xi_+$ for all $\xi' \neq \xi$ in M .

A subgroup $P < G$ is called *parabolic* if it is infinite, fixes a point $\xi \in M$, and contains no loxodromics (elements which fix exactly 2 points of M). This fixed point ξ is unique and is called a *parabolic point*. Further, ξ is a *bounded parabolic point* if $\text{Stab}_G(\xi)$ acts properly discontinuously and cocompactly on $M \setminus \{\xi\}$.

Definition 2.3 (Geometrically Finite). Let G act as a convergence group on M . The action of G is *geometrically finite* if every point of M is either a bounded parabolic point or conical limit point.

In [Hru10], Hruska lays out six equivalent definitions of relative hyperbolicity. We will only be concerned with two of them. We note that some of these definitions allow finite parabolic subgroups, which we do not allow. Recall from [BH99, III.H] that if Σ is a proper δ -hyperbolic metric space, then its boundary $\partial\Sigma$ is defined by equivalence classes of asymptotic rays. Then $\overline{\Sigma} := \Sigma \cup \partial\Sigma$ can be topologized as a compact metrizable space, and if G acts by isometries on Σ , this extends to an action on $\overline{\Sigma}$.

Definition 2.4 (Relatively Hyperbolic). Let G be a group acting properly discontinuously by isometries on a proper hyperbolic metric space Σ so that G acts on $\partial\Sigma$ as a geometrically finite convergence group. Let $\mathbb{P}$ be a collection of one representative from each conjugacy class of maximal parabolic subgroups, and assume each element of $\mathbb{P}$ is finitely generated. Then we say G is *hyperbolic relative to $\mathbb{P}$* , or $(G, \mathbb{P})$ is relatively hyperbolic. The collection $\mathbb{P}$ is the *peripheral structure*. The boundary $\partial\Sigma$ is the *Bowditch boundary* of $(G, \mathbb{P})$, denoted $\partial_{\mathbb{P}}G$.

The peripheral structure is fundamental to how we view G as a relatively hyperbolic group — different peripheral structures on the same group G lead to different Bowditch boundaries. However, when context makes the peripheral structure clear, we will be glib and say that G is relatively hyperbolic with Bowditch boundary ∂G . The following remarkable theorem shows that the only geometrically finite convergence actions on compact metrizable spaces are relatively hyperbolic groups acting on their Bowditch boundaries. This also implies the Bowditch boundary depends only on the group G and its peripheral structure.

Theorem 2.5. [Yam04] [Bow98] Let G be a geometrically finite convergence group on a perfect compact metrizable space M . Let $\mathbb{P}$ be a collection of representatives of conjugacy classes of maximal parabolic subgroups, and assume that each element of P is finitely generated. Then $(G, \mathbb{P})$ is relatively hyperbolic and M is equivariantly homeomorphic to $\partial_{\mathbb{P}}G$.

Our goal is to show a particular G is relatively hyperbolic. We will do it by constructing a candidate for M and applying Theorem 2.5. Tukia has shown the set $\mathbb{P}$ is always finite:

Proposition 2.6. [Tuk98] If G is a geometrically finite convergence group on a compact metrizable space M , then there are finitely many conjugacy classes of maximal parabolic subgroups.

The boundary of a hyperbolic group naturally compactifies the Cayley graph of the group, as described in [BH99, III.H]. In [GM08], Groves and Manning provide a space analogous to the Cayley graph for relatively hyperbolic groups by attaching combinatorial horoballs to the Cayley graph. Called the cusped space, this graph is compactified by the Bowditch boundary.

Definition 2.7. [GM08, Definition 3.1] Let Γ be any graph. The *combinatorial horoball based on* Γ has vertices $\Gamma^{(0)} \times (\mathbb{N} \sqcup \{0\})$ and three kinds of edges:

1. If e is an edge of Γ joining v to w , then there is a corresponding edge $\bar{e}$ joining $(v, 0)$ to $(w, 0)$.
2. If $k > 0$ and $0 < d_\Gamma(v, w) < 2^k$ then there is a single edge joining (v, k) to (w, k) .
3. If $k \geq 0$ and $v \in \Gamma^{(0)}$, then there is an edge joining (v, k) to $(v, k + 1)$.

Definition 2.8. [GM08, Definition 3.12] Let G be a finitely generated group with $\mathbb{P} = \{P_1, \dots, P_n\}$ a finite collection of finitely generated subgroups, and let S be a finite generating set for G so that $P_i \cap S$ generates P_i for each i . Let Γ be the Cayley graph of G with respect to S . The *cusped space* for $(G, \mathbb{P}, S)$ is defined by attaching a combinatorial horoball to each coset of P_i , viewed as a subset of Γ .

Theorem 2.9. [GM08, Theorem 3.25] Let G be a finitely generated group with $\mathbb{P}$ a finite collection of finitely generated subgroups, and let X be a cusped space for $(G, \mathbb{P})$. Then $(G, \mathbb{P})$ is relatively hyperbolic in a sense equivalent to Definition 2.4 if and only if X is δ -hyperbolic, and in this case $\partial X = \partial_{\mathbb{P}} G$.

[Hru10, Section 4] explains how to adapt the cusped space when the peripheral subgroups are not finitely generated, as well as showing the equivalence of the various definitions of relative hyperbolicity. The Bowditch boundary only contains information about infinite subgroups through their limit sets, and for us the cusped space is a convenient way to manage finite subgroups. Except for a few lemmas like the next one, we do not deal in the details of cusped spaces, so we direct the reader to [GM08] for further details.

Lemma 2.10. Let $(G, \mathbb{P})$ be relatively hyperbolic with cusped space C . Let P be a finite index subgroup of an element of $\mathbb{P}$ fixing $\xi \in \partial C = \partial_{\mathbb{P}} G$, and let $\mathbb{F} = \{F_1, \dots, F_m\}$ be a collection of finite subgroups of G . Then there exists a compact subset $K \subset \bar{C} = C \cup \partial C$ so that

1. $P(\partial C \cap K) = \partial C \setminus \{\xi\}$, that is, $K \cap \partial C$ is a coarse fundamental domain for P acting on $\partial G \setminus \{\xi\}$, and
2. for any coset gF_i , there is some $p \in P$ so that $pgF_i \cap K \neq \emptyset$.

Proof. Because G is relatively hyperbolic, ξ is a bounded parabolic point. Because P is finite index in the stabilizer of ξ , we can choose a compact set $K_1 \subset \partial C \setminus \{\xi\}$ so that $PK_1 = \partial C \setminus \{\xi\}$.

Consider the set $B = \{gF_i \mid d_C(gF_i, P) \text{ is realized by } d_C(gF_i, 1)\}$ and let $\overline{B}$ be the closure of $\bigcup_{gF_i \in B} gF_i$ in $\overline{C}$. Notice $\overline{B}$ is compact because $\overline{C}$ is compact. Any point $\eta \in \overline{B}$ is the limit of a sequence of cosets $(g_n F_{i_n})_n$, and by Arzela–Ascoli the geodesics $[1, g_n F_{i_n}]$ subconverge to a geodesic $[1, \eta)$. Because the geodesics $[1, g_n F_{i_n}]$ all travel immediately away from the horoball containing P for longer and longer distances, their limit cannot follow travel the vertical geodesic $[1, \xi)$ forever. Thus $\xi \notin \overline{B}$ and $\overline{B}$ is compact in $\overline{C} \setminus \{\xi\}$.

Let $K = K_1 \cup \overline{B}$ and we show K satisfies the conclusion. Clearly K is compact since $K_1, \overline{B}$ are and clearly (1) is satisfied because $K_1 \subset K \cap \partial C$. For any coset gF_i , suppose $d_C(gF_i, P)$ is realized by $d_C(gF_i, p)$ for some $p \in P$. Then $d_C(p^{-1}gF_i, P)$ is realized by $d_C(p^{-1}gF_i, 1)$, and $p^{-1}gF_i \in B$. So $p^{-1}gF_i \subset \overline{B}$, and (2) is satisfied. $\square$

2.2 Quasiconvexity

Quasiconvex subgroups play an important role in hyperbolic groups. Depending on context, relatively quasiconvex and full relatively quasiconvex subgroups play the analogous role in relatively hyperbolic groups.

Definition 2.11. [Limit Set, [Bow98][Tuk94]] If G is a convergence group on a compact metrizable space M and H is an infinite subgroup, the *limit set* ΛH has three equivalent characterizations:

1. the unique minimal nonempty closed H -invariant subset of M ,
2. the set of points in M at which H does not act properly discontinuously,
3. the set of attractive points of convergence sequences in H .

If H is finite, ΛH is empty.

Proposition 2.12. [Bow98][Prop 3.1, 3.2] Let G be a convergence group on a compact metrizable space M . Then conical limit points are in ΛG and conical limit points are not parabolic points.

[Hru10] lays out many definitions of relative quasiconvexity as well, including the first of the following definitions.

Definition 2.13. Let G act on a compact metrizable space M as a geometrically finite convergence group, let $\mathbb{P}$ be the set of maximal parabolic subgroups, and let H be a subgroup.

1. H is *relatively quasi-convex* in G , or *RQC*, if the action of H on ΛH is geometrically finite.
2. H is *fully relatively quasi-convex* in G , or *fully RQC*, if it is RQC and, for any infinite sequence of elements $(g_n)_n$ each in a distinct H coset, we have $\bigcap_n g_n \Lambda H = \emptyset$,

3. H is *full relatively quasi-convex* in G , or *full RQC*, if it is RQC and, for every maximal parabolic subgroup P of G , $H \cap P$ is either finite or finite index in P .

The reader should be disturbed by these names. Fortunately, fully RQC and full RQC are equivalent conditions, as we will see in Corollary 2.23. The fully implies full direction is due to Dahmani [Dah03, Lemma 1.7], and he also shows that a subgroup of a hyperbolic group is metrically quasiconvex as a subset of the Cayley graph if and only if it is fully RQC in the above sense.

The following proposition shows that an RQC subgroup inherits a peripheral structure from the larger group which makes it relatively hyperbolic.

Proposition 2.14. [Hru10, Theorem 9.1] Let H be an RQC subgroup of a relatively hyperbolic group $(G, \mathbb{P})$ and let

$$\overline{\mathbb{O}} = \{H \cap P^g \mid P \in \mathbb{P}, g \in G, H \cap P^g \text{ infinite}\}$$

Then $\overline{\mathbb{O}}$ consists of finitely many H conjugacy classes. If $\mathbb{O}$ is a set of representatives of H conjugacy classes from $\overline{\mathbb{O}}$, $(H, \mathbb{O})$ is relatively hyperbolic and $\partial_{\mathbb{O}} H$ is H -equivariantly homeomorphic to $\Lambda H \subset \partial_{\mathbb{P}} G$.

- Remark 2.15.**
1. If G is relatively hyperbolic and Σ is as in Definition 2.4, then ΛH can be seen as the points of $\partial \Sigma$ which are limit points of some H -orbit in Σ . Changing which base point defines the orbit translates the orbit a bounded amount, leaving the limit set unchanged.
 2. If $H_1 < H_2 < G$ are groups with G relatively hyperbolic, H_2 RQC in G , and H_1 RQC in H_2 , then H_1 is RQC in G . The limit set of H_1 in ∂G is the image of the inclusion $\partial H_1 \longrightarrow \partial H_2 \longrightarrow \partial G$.
 3. Parabolic subgroups are RQC since their limit set is a single point. Further, maximal parabolic subgroups are full RQC; If P, P' are maximal parabolic subgroups of a relatively hyperbolic group and $P \cap P'$ is infinite, we must have $P = P'$. This really shows maximal parabolic subgroups are almost malnormal, for if $P' = P^g$ and $P \cap P^g$ is infinite, then it must have a limit set, and it must be the unique fixed point of P . Hence g fixes this point and $g \in P$ by the maximality of P .

Proposition 2.16. [AGM09, Lemma 3.1] Let H be a RQC subgroup of a relatively hyperbolic group $(G, \mathbb{P})$ and let CH and CG be the corresponding cusped spaces. The inclusion $i : H \hookrightarrow G$ induces an H -equivariant Lipschitz map $\widehat{i} : CH \rightarrow CG$ with quasiconvex image.

Proof Sketch. As above, the peripheral structure of H can be expressed as $\mathbb{D} = \{D_i = H \cap P_i^{c_i} \mid P_i \in \mathbb{P}, c_1, \dots, c_\ell \in G, H \cap P_i^{c_i} \text{ infinite}\}$, where some of the P_i may be identical. To extend i to $\widehat{i}$, consider a point of a horoball in H is given by (hD_i, hd, n) , where $h \in H, d \in D_i$, and $n \geq 1$. Define

$$\widehat{i}(hD_i, hd, n) = (hc_i P_i, hdc_i, n).$$

Keeping track of the generating sets for H, G used to build CH, CG shows $\widehat{i}$ is Lipschitz.

In [GM08, Definition 3.11], a subgroup of $(G, \mathbb{P})$ is defined as C -relatively quasiconvex exactly when the image of $\widehat{i}$ is C -quasiconvex as a subset of the cusped space for G . [MMP10, Theorem A.10] shows H is C -relatively quasiconvex if and only if H is RQC in the sense of Definition 2.13. By assumption, H is RQC in the sense above, so [MMP10] implies this image is metrically quasiconvex as in the statement. $\square$

In [Yan12], Yang studies intersections of RQC subgroups and gives us the following useful properties.

Proposition 2.17. [Yan12, Thm 1.1] Let H_1, H_2 be RQC subgroups of a relatively hyperbolic group G . Then

$$\Lambda H_1 \cap \Lambda H_2 = \Lambda(H_1 \cap H_2) \sqcup E$$

where the exceptional set E consists of parabolic points $\xi \in \Lambda H_1 \cap \Lambda H_2$ so that $\text{Stab}_{H_1}(\xi) \cap \text{Stab}_{H_2}(\xi)$ is finite.

Proposition 2.18. [Yan12, Prop 1.3] If H_1, H_2 are RQC subgroups of a relatively hyperbolic group, then $H_1 \cap H_2$ is RQC.

Corollary 2.19 (Limit Set Property). If H_1, H_2 are full RQC subgroups of a relatively hyperbolic group G , then $H_1 \cap H_2$ is full RQC and

$$\Lambda H_1 \cap \Lambda H_2 = \Lambda(H_1 \cap H_2).$$

Proof. From the previous two propositions we know $H_1 \cap H_2$ is RQC and that the intersection of limit sets is the limit set of the intersection together with some exceptional points. If $\xi \in \Lambda H_1 \cap \Lambda H_2$ is the fixed point of a maximal parabolic subgroup P , then we must have $P \cap H_i$ finite index in P for each i since the H_i are full. Since the intersection of finite index subgroups is again finite index, we know $(P \cap H_1) \cap (P \cap H_2) = P \cap (H_1 \cap H_2)$ is finite index in P , and therefore infinite. So there are no exceptional points and we have the equality above. Similarly, if $P \cap (H_1 \cap H_2)$ is infinite, then $P \cap H_i$ is infinite for each i , and because each H_i is full, $P \cap (H_1 \cap H_2) = (P \cap H_1) \cap (P \cap H_2)$ is an intersection of finite index subgroups of P , hence is finite index in P . This shows $H_1 \cap H_2$ is full RQC. $\square$

To show that full RQC and fully RQC are equivalent, we need the following definitions.

Definition 2.20. Let G be a group and H a subgroup. The *height* of H in G is the minimal number n so that for any collection of distinct cosets $\{g_1 H, g_2 H, \dots, g_{n+1} H\}$, the intersection $\bigcap_i H^{g_i}$ is finite. If G is relatively hyperbolic, the *relative height* of H in G is the minimal number n so that for any collection of distinct cosets $\{g_1 H, g_2 H, \dots, g_{n+1} H\}$, the intersection $\bigcap_i H^{g_i}$ is either finite or parabolic.

In [HW09], Hruska and Wise use the term “height” for what we are calling “relative height” and prove the following.

Proposition 2.21. [HW09, Corollary 8.6] Let H be a RQC subgroup of a relatively hyperbolic group G . Then H has finite relative height in G .

If H is full RQC, we can upgrade this to truly finite height.

Proposition 2.22. If H is a full RQC subgroup of a relatively hyperbolic group $(G, \mathbb{P})$, then H has finite height.

Proof. As in Proposition 2.14, G induces a peripheral structure $\mathbb{O} = \{P_1, \dots, P_k\}$ on H . Explicitly, $P_i = H \cap P'_i$, where $P'_i < G$ is a maximal parabolic subgroup with $H \cap P'_i$ infinite, and $\mathbb{O}$ contains one representative of each H -conjugacy class of such subgroups. Let

$$m = \max_i [P'_i : P_i] = \max\{[P : H \cap P] \mid H \cap P \text{ infinite, } P \text{ maximal parabolic in } G\}.$$

From the previous proposition, H has finite relative height, say n . We claim H has height at most $N = n(m+1)(k+1)$.

For a contradiction, suppose $\{g_1 H, \dots, g_{N+1} H\}$ are distinct H -cosets in G with $\bigcap_i H^{g_i}$ infinite. After conjugating by g_1^{-1} , we may assume $g_1 = 1$ so that this intersection is contained in H . Because $N+1$ is larger than the relative height of H , we must have $\bigcap_i H^{g_i} < H \cap P$ for some maximal parabolic subgroup $P < G$. For each i , we have $\bigcap_i H^{g_i} \subset H^{g_i} \cap P$, hence $H \cap P^{g_i^{-1}}$ is infinite, hence $H \cap P^{g_i^{-1}}$ is a maximal parabolic subgroup of H . Apply the pigeonhole principle to $W := \{H \cap P^{g_i^{-1}} \mid i = 1, \dots, N+1\}$ where the pigeons are elements of W and the holes are H -conjugacy classes of maximal parabolic subgroups. Because $N+1 > (m+1)(k+1)$ and there are k different H -conjugacy classes of maximal parabolic subgroups, some H -conjugacy class has at least $m+1$ representatives in W . After reindexing, $H \cap P^{g_1^{-1}}, \dots, P^{g_{m+1}^{-1}}$ are conjugate in H , so there are $h_i \in H$ so that $H \cap P^{g_1^{-1}} = H \cap P^{h_i g_i^{-1}}$. Then

$$H \cap P^{g_1^{-1}} = H \cap P^{h_i g_i^{-1}} \subset P^{g_1^{-1}} \cap P^{h_i g_i^{-1}},$$

and the sets on the left are infinite. Because P is almost malnormal, this implies $P^{g_1^{-1}} = P^{h_i g_i^{-1}}$, equivalently $P = P^{g_1 h_i g_i^{-1}}$, and $g_1 h_i g_i^{-1} \in P$ for $i = 1, \dots, m+1$. By the definition of m , $D := H^{g_1} \cap P$ has index at most m in P , so we may write P as a disjoint union of cosets:

$$P = Dr_1 \sqcup Dr_2 \sqcup \dots \sqcup Dr_m.$$

Viewing these cosets as pigeonholes and the elements $g_1 h_i g_i^{-1}$ as pigeons, there is some coset which contains two of these elements. This means there is some $1 \leq i, j \leq m+1$ and $r \in P, p_i, p_j \in D$ so

that $p_i r = g_1 h_i g_i^{-1}$ and $p_j r = g_1 h_j g_j^{-1}$. Rearranging this, we have $p_i^{-1} g_1 h_i g_i^{-1} = r = p_j^{-1} g_1 h_j g_j^{-1}$. Since $p_i, p_j \in D := H^{g_1} \cap P$, we can write $p_i = g_1 q_i g_1^{-1}, p_j = g_1 q_j g_1^{-1}$ for some $q_i, q_j \in H$. Substituting this into the previous equality, we have $g_1 q_i^{-1} h_i g_i^{-1} = g_1 q_j^{-1} h_j g_j^{-1}$. Canceling the g_1 on each side and rearranging, we have

$$g_i = g_j h_j^{-1} q_j q_i h_i.$$

But $h_j^{-1} q_j q_i h_i$ is a product of elements of H , so this implies $g_i H = g_j H$, contradicting the assumption that g_i were distinct coset representatives. Thus N is a bound on the height of H . $\square$

We remark that there is a simpler proof that *fully* RQC subgroups have finite height, but finite height is necessary for proving the full implies fully direction in the following corollary.

Corollary 2.23 (Finite Height). Let H be a subgroup of a relatively hyperbolic group G . Then H is fully RQC if and only if H is full RQC, and either condition implies H has finite height in G .

Proof. The fully implies full direction is [Dah03, Lemma 1.7], and we repeat the proof here for completeness. Suppose H is fully RQC in G and P is a maximal parabolic subgroup fixing $\xi \in \partial G$. Suppose $H \cap P$ is infinite so that $\xi \in \Lambda(H \cap P)$. For a contradiction, suppose $(p_n)_n$ is an infinite sequence of distinct coset representatives for $H \cap P$. Each p_n represents a distinct H coset too, since $p_n H = p_m H$, implies $p_n p_m^{-1} \in H \cap P$, hence $p_n(H \cap P) = p_m(H \cap P)$ and $n = m$. Each p_n also fixes ξ , so

$$\xi \in \bigcap_n p_n \Lambda(H \cap P) \subset \bigcap_n p_n \Lambda H.$$

But the intersection on the right is empty because H is fully RQC. This contradiction shows H is also full RQC.

For the other direction, suppose H is full RQC. Then H has finite height in G by the previous proposition. If $(g_n)_n$ is an infinite sequence of elements in distinct H cosets, then $\bigcap_n H^{g_n}$ is finite because H has finite height. In fact, if H has height N , then $\bigcap_{n=1}^{N+1} H^{g_n}$ is finite. By the Limit Set Property 2.19 and because the limit set of a finite group is empty, we have

$$\bigcap_{n=1}^{\infty} g_n \Lambda H \subset \bigcap_{n=1}^{N+1} g_n \Lambda H = \bigcap_{n=1}^{N+1} \Lambda H^{g_n} = \Lambda \left(\bigcap_{n=1}^{N+1} H^{g_n} \right) = \emptyset.$$

$\square$

Proposition 2.24. (Convergence Property) Let $(G, \mathbb{P})$ be a relatively hyperbolic group, H a RQC subgroup, and $(g_n H)_n$ a sequence of distinct cosets. Let X be the cusped space for $(G, \mathbb{P})$. Let Y

be the image of the Lipschitz map provided by Proposition 2.16. After taking a subsequence, there is some $\xi \in \partial G$ so that $g_n(Y \sqcup \Lambda H) \rightarrow \xi$ uniformly.

Proof. Recall Y consists of H together with some points of positive depth and is quasiconvex for some constant K . Choose $x_n \in g_n Y$ so that $d(1, x_n) = d(1, g_n Y)$. The cosets $g_n H$ are all distinct, hence disjoint, and X is proper, so $d(1, x_n) \rightarrow \infty$. Applying Arzela–Ascoli, we can choose a subsequence so that the geodesics $[1, x_n]$ converge to some ray r from 1 to a point $\xi \in \partial X = \partial G$. The set $Y \sqcup \Lambda H$ is the closure of Y in the compact metrizable space $X \sqcup \partial X$, so to show the claim it suffices to show $g_n Y \rightarrow \xi$ uniformly. To show this, it suffices to show that for any M , there exists N so that if $n \geq N$ and $y \in Y$, the geodesic ray $[1, g_n y]$ stays within δ of r for time M .

We can assume $d(1, x_n) \geq K + \delta$ for all n . We claim that for all n and $y \in g_n Y$, $(1, y)_{x_n} \leq K + \delta$. Given $y \in g_n Y$, consider the geodesic triangle made by $1, x_n$ and y . Let $z \in [1, x_n]$ and $z' \in [x_n, y]$ be the corresponding tripod points. Then $d(z, z') \leq \delta$ because X is δ -hyperbolic, and there is some $z'' \in g_n Y$ so that $d(z', z'') \leq K$ because $g_n Y$ is K -quasiconvex. Because x_n achieves the distance $d(1, g_n Y)$, we must have $d(z, g_n Y) = d(z, x_n) = (1, y)_{x_n}$, but using z', z'' and the triangle inequality we have $d(z, g_n Y) \leq K + \delta$. Hence $(1, y)_{x_n} \leq K + \delta$ as claimed.

Fix M and choose N so that $n \geq N$ implies $[1, x_n]$ stays within δ of r for time at least $M + 2\delta + K$, which is possible because the x_n converge to ξ . For $n \geq N$ and $y' \in g_n Y$, let $x \in [1, x_n], y \in [1, y'], z \in r$ be the points at distance $M + 1\delta$ from 1. Considering the comparison tripod for the triangle on $1, y', x_n$ and noticing that $d(1, x_n) \geq M + \delta + K$ and $(1, y')_{x_n} \leq K + \delta$, we see that x, y are both closer to 1 than the tripod points on either leg, hence $d(x, y) \leq \delta$. Further, by the assumption on n , we have $d(x, z) \leq \delta$. Therefore $d(y, z) \leq 2\delta$ and

$$(y, z)_1 = M + 2\delta - \frac{1}{2}d(y, z) \geq M.$$

Since $y \in [1, y'], z \in r$, this implies $[1, y']$ and r stay within δ of each other for time at least M , and we are finished. $\square$

Dahmani proves an analogous result to the Convergence Property 2.24 in [Dah03, Prop 1.8] using only the Bowditch boundary, but we need the version above to manage finite groups which do not interact with the boundary.

Lemma 2.25. Suppose H_1, H_2 are full RQC subgroups of a relatively hyperbolic group $(G, \mathbb{P})$ and $(a_n)_n$ a sequence in G . Let X be the cusped space for $(G, \mathbb{P})$ and let Q_1, Q_2 be the images of cusped spaces for H_1, H_2 in X from Proposition 2.16. Then there is a sequence $(k_n)_n$ in H_1 and a subsequence of $(a_n)_n$ (still denoted $(a_n)_n$) so that the translates $k_n a_n Q_2$ are either constant or converge to a point in $\partial X \setminus \Lambda Q_1 = \partial G \setminus \Lambda H_1$.

Proof. By Proposition 2.16, Q_1, Q_2 are K -quasiconvex. Let $\pi : X \rightarrow Q_1$ be the projection of X onto Q_1 . Slightly abusing language, we will call an image of a horoball from the cusped space for

H_i a horoball of Q_i .

Suppose the peripheral structure on H_1 is given by $\mathbb{D} = \{H_1 \cap P_i^{c_i} \mid P_i \in \mathbb{P}, c_1, \dots, c_\ell \in G, H_1 \cap P_i^{c_i} \text{ infinite}\}$, where some of the P_i may be identical. A point of Q_1 with positive depth is in a horoball, so it has the form

$$(hc_i P_i, hdc_i, n)$$

where $h \in H_1$ and $d \in H_1 \cap c_i P_i c_i^{-1}$. We can translate such a point by $(hd)^{-1} \in H_1$ to the point $(c_i P_i, c_i, n)$. Points of depth 0 are simply elements of H_1 , which can be translated to the identity by an element of H_1 . Thus Q_1 is the H_1 orbit of the identity together with finitely many H_1 orbits of vertical paths starting at the $(c_i P_i, c_i, 1)$.

Because H_1 is full RQC, $H_1 \cap P_i^{c_i}$ is finite index in $P_i^{c_i}$ for every i , so there is a constant M so that if A' is a horoball of Q_1 contained in a horoball A of X , then A is contained in the M neighborhood of A' .

We claim that if $x \in X$ has depth 0, then $\pi(x)$ has depth at most M . If $\pi(x)$ has positive depth, let y be the first point of $[x, \pi(x)]$ in the horoball of X containing $\pi(x)$. Then y has depth 0 because it is the first point in this horoball, but also $d(y, Q_1) \leq M$, so $[y, \pi(x)]$ can travel at most M vertically. Thus $\pi(x)$ has depth at most M .

For each n , choose a point $x_n \in a_n Q_2$ achieving the distance $d(a_n Q_2, Q_1)$. If $a_n Q_2 \cap Q_1 \neq \emptyset$, we can assume x_n has depth 0. If $a_n Q_2 \cap Q_1 = \emptyset$, then $\pi(x_n)$ has depth at most M by the previous paragraph. We choose $k_n \in H_1$ so that $k_n \pi(x_n)$ is either the identity or some $(c_i P_i, c_i, n)$, where $n \leq M$. The definition of π is only coarse so π may not be equivariant, but this implies $\pi(k_n x_n)$ is some uniformly bounded distance from $k_n \pi(x_n)$. Since X is locally finite, this implies there are finitely many choices for $\pi(k_n x_n)$, and after a subsequence we can assume $\pi(k_n x_n)$ is constant, say $z \in Q_1$.

After a further subsequence, we can assume that $k_n x_n$ is either constant or the geodesics $[z, k_n x_n]$ converge to a geodesic $[z, \eta)$ for some $\eta \in \partial X$. If $k_n x_n$ is constant, then $k_n x_n \in k_n a_n Q_2$, and the translates of Q_2 are disjoint, so $k_n a_n Q_2$ is constant and we are finished. In the second case, the distances $d(k_n x_n, z) = d(k_n x_n, Q_1)$ must be unbounded and the geodesics $[z, k_n x_n]$ travel further and further from Q_1 , hence $\eta \notin \Lambda Q_1$. Applying the Convergence Property 2.24, we can assume the $k_n a_n Q_2$ converge to some point in ∂X , and since $k_n x_n \rightarrow \eta$, this point must be $\eta \notin \Lambda Q_1$, and we're finished. $\square$

Proposition 2.26. Let $(G, \mathbb{P})$ be a relatively hyperbolic group with cusped space C . There exists a cell complex $K(G)$ satisfying the following.

1. $K(G)$ is a $K(G, 1)$.
2. The universal cover $\widetilde{K(G)}$ has 1-skeleton equal to C , possibly with duplicated edges.

3. If H is a full RQC subgroup of G and $\psi : C_H \rightarrow C$ is a Lipschitz map as in Proposition 2.16, then there is a map $\varphi : K(H) \rightarrow K(G)$ realizing the inclusion $H \rightarrow G$ on fundamental groups and which lifts to a map $\tilde{\varphi} : K(H) \rightarrow K(G)$ restricting to ψ on 1-skeletons.

G admits a $K(G, 1)$ whose universal cover has 1-skeleton equal to C , with duplicated edges if G has 2-torsion. If H is a full RQC subgroup of G and C_H

Proof. Let C be a cusped space for $(G, \mathbb{P})$. [GM08] describes how to add 2-cells to C using a relative presentation for G to get a simply connected space. However, if G is not torsion free, the action on this space may not be free; An element of order 2 may invert an edge, and if an element fixes the boundary of an 2-cell, it will fix the center of the cell. Therefore if e is an edge of C which is inverted by some element of G , we replace e with two edges so that an element which inverts e swaps these two edges instead. For each 2-cell added to make C simply connected with a boundary of length n , we add n 2-cells along the same boundary so that an element which fixes this boundary permutes these added 2-cells instead. After these modifications, we have a simply connected 2-complex on which G acts freely, say C' . Then $\pi_1(C'/G) = G$, and it is a standard construction to add higher dimensional cells to C'/G so that it is aspherical, and these added cells lift to a higher dimensional cell structure on C' . This constructs $K(G)$.

With H, C_H, ψ from (3), ψ extends to a map $\widetilde{K(H)} \rightarrow \widetilde{K(G)}$ which projects to a map realizing the inclusion on fundamental groups. $\square$

2.3 M_κ -Complexes

Definition 2.27. [BH99, I.2] For $\kappa < 0$, the model space M_κ^n is simply $\mathbb{H}^n$ with the metric scaled so that the curvature is κ . For $\kappa = 0$, $M_\kappa^n = \mathbb{R}^n$, and for $\kappa > 0$, $M_\kappa^n = S^n$ with the metric scaled so that the curvature is κ .

Definition 2.28 (M_κ -complex). [BH99, I.7] Let $\kappa \leq 0$. A simplicial complex X is called an M_κ -complex if it satisfies the following conditions.

- Each simplex of X is modeled after a geodesic simplex in M_κ^n , which is the convex hull of finitely many points in general position.
- If σ, σ' are two simplices meeting in a face τ , then the identity map from $\tau \subset \sigma$ to $\tau \subset \sigma'$ is an isometry.

The path metric, also called the simplicial metric, on an M_κ -complex is natural and pleasant by the following theorem.

Theorem 2.29. [Bri91, Theorem 1.1] If X is an M_κ -complex with $\kappa \leq 0$ and finitely many isometry types of simplices, then the simplicial path metric is complete and geodesic.

M_κ -complexes are very general and Bridson's thesis [Bri91] developed much of the machinery we will use. For the remainder of this section, we fix a CAT(0) M_κ -complex X with finitely many isometry types of simplices.

Definition 2.30. Given two subsets K, K' of X , let $\text{Geod}(K, K')$ denote the set of points lying on a geodesic segment from a point of K to a point of K' .

Definition 2.31 (Simplicial Neighborhood). Let K be a subcomplex of X .

1. The *open simplicial neighborhood* of K , denoted $N(K)$, is the union of open simplices of X whose closure meets K . If $K = \sigma$ is a single simplex, we write $st(\sigma) = N(\sigma)$.
2. The *closed simplicial neighborhood* of K , denoted $\overline{N}(K)$, is subcomplex spanned by closed simplices that meet K . Equivalently, it is the closure of $N(K)$.
3. The *simplicial link* of K , denoted $Lk(K)$, is $N(K) \setminus K$.

For example, if X is a graph and $K = v$ is a vertex, then $N(K)$ is v together with the interior of each edge with v as an endpoint, $Lk(v)$ is the interior of each of these edges, and $\overline{N}(K)$ is $N(K)$ together with the other vertices at the end of those edges.

Definition 2.32 (Path of Simplices). A *path of simplices* is a sequence of open simplices $\sigma_1, \dots, \sigma_n$ so that either $\overline{\sigma_i} \subset \overline{\sigma_{i+1}}$ or $\overline{\sigma_{i+1}} \subset \overline{\sigma_i}$ for each $i = 1, \dots, n-1$. Equivalently, it is a finite path in the 1-skeleton of the first barycentric subdivision of X . The integer n is the *length* of the path of simplices.

Lemma 2.33. [Mar14, Lemma 3.5] For finite subcomplexes $K, K' \subset X$, $\text{Geod}(K, K')$ meets finitely many open simplices.

Lemma 2.34. [Bri91, Theorem 1.11] For every n , there exists a constant k so that every geodesic segment of length at most n meets at most k open simplices.

Lemma 2.35. [Bri91, Theorem 1.11][Containment] For every n there exists a constant k such that for every finite subcomplex K of X containing at most n simplices, any geodesic path in the open simplicial neighborhood of K meets at most k simplices.

See Figure 2 for an illustration of the set up for the following lemma.

Lemma 2.36. [Mar14, Lemma 3.7][Short Paths of Simplices] There is a function $F : \mathbb{N} \rightarrow \mathbb{N}$ so that the following holds: Let K be a convex subcomplex of X and K' a connected subcomplex of X both containing at most n simplices. Let $x, y \in K$ and $x', y' \in K'$, and assume there exists a path between x' and y' in K' that does not meet K . Let τ, τ' be two simplices of $Lk(K)$ so that $[x, x']$ (resp. $[y, y']$) meets τ (resp. τ'). Then there exists a path of simplices in $Lk(K)$ of length at most $F(n)$ between τ and τ' .




Figure 2: The situation of Lemma 2.36.

Definition 2.37 (Acylindrical). Let G be a group acting on an M_κ -complex X . The action is called *acylindrical* if there is a constant A so that for any set $K \subset X$ with $\text{diam}(K) \geq A$, K has finite pointwise stabilizer.

2.4 Complexes of Groups

Given a group acting cocompactly on a tree, Bass–Serre theory explains how to use vertex and edge stabilizers with the quotient graph to get a graph of groups structure for the original group. Conversely, given a graph of groups, we can build a tree on which the fundamental group acts with quotient our original graph of groups. In [BH99], Bridson and Haefliger extend this correspondence to groups acting on any simply connected category. Here we introduce the definitions and notation we need from their theory.

Definition 2.38 (Scwol). Let $\mathcal{Y}$ be a category with objects $V(\mathcal{Y})$ and morphisms (or arrows) $E(\mathcal{Y})$. For $a \in E(\mathcal{Y})$ we write $i(a), t(a)$ for the source and target of a . Then $\mathcal{Y}$ is a *small category without loops*, or a *scwol*, if $V(\mathcal{Y}), E(\mathcal{Y})$ are both sets and for any $a \in E(\mathcal{Y})$, we have

$$i(a) = t(a) \implies a \text{ is the identity morphism of } i(a) = t(a)$$

Further, $\mathcal{Y}$ is *simple* if there is at most one morphism between any two objects of $\mathcal{Y}$.

Definition 2.39. Let Y be a cell complex. The *scwolification* of Y , denoted $\mathcal{Y}$, is the scwol with objects corresponding to cell of Y and arrows corresponding to reverse inclusion.




Figure 3: From left to right, a simplex, the scwol constructed by reverse inclusion, and the geometric realization of this scwol.

Definition 2.40. Let $\mathcal{Y}$ be a scwol. The *geometric realization* of $\mathcal{Y}$, denoted $|\mathcal{Y}|$ is a the flag simplicial complex with vertices $V(\mathcal{Y})$ and edges $E(\mathcal{Y})$. An n -simplex of $|\mathcal{Y}|$ corresponds to an n -tuple of composable edges in $\mathcal{Y}$.

The only scwols we will be interested in are the scwolification of cell complexes. If Y is a cell complex and $\mathcal{Y}$ is its scwolification, then $|\mathcal{Y}|$ is the barycentric subdivision of Y . See 3 for an example where Y is a single triangle.

Note that if a group G acts on *simplicial* complex X and fixes some simplex setwise but not pointwise, then the quotient will not be simplicial, for example $\mathbb{Z}/3\mathbb{Z}$ acting on a triangle by rotation. To remedy this, one can consider the action on the barycentric subdivision of X , say X_b . If σ is an n -cell of X_b , then σ corresponds to a chain $\sigma_1 \subset \sigma_2 \subset \dots \subset \sigma_n$ where each σ_i is a cell of X . If an element of G fixes σ , it must fix each σ_i , so $\text{Stab}_G(\sigma)$ fixes σ pointwise.

Convention 2.41. Whenever G is a group acting on a simplicial complex X , we assume the quotient is simplicial and stabilizers of simplices fix simplices pointwise. As above, this can always be achieved by barycentrically subdividing if necessary.

Definition 2.42 ([BH99]). (Complex of Groups) A *complex of groups* over a scwol $\mathcal{Y}$, denoted $G(\mathcal{Y}) = (G_\sigma, \psi, g_{a,b})$, consists of the following data.

1. For each $\sigma \in V(\mathcal{Y})$, a *local group* G_σ ,
2. For each $a \in E(\mathcal{Y})$, an injective homomorphism $\psi_a : G_{i(a)} \longrightarrow G_{t(a)}$,
3. For each pair of composable arrows a, b , a *twisting element* $g_{a,b} \in G_{t(a)}$ satisfying
 - (a) $\text{Ad}(g_{a,b})\psi_{ab} = \psi_a\psi_b$, where $\text{Ad}(g_{a,b})$ is conjugation by $g_{a,b}$,
 - (b) for any composable edges a, b, c , we have $\psi_a(g_{b,c})g_{a,bc} = g_{a,b}g_{ab,c}$.

Definition 2.43. [BH99, III.C.3.7][Fundamental Group of a Complex of Groups] Let $G(\mathcal{Y}) = (G_\sigma, \varphi_a, g_{a,b})$ be a complex of groups with T a maximal tree in $\mathcal{Y}$, and let $E^\pm(\mathcal{Y})$ be the symbols $\{a^+, a^- \mid a \in E(\mathcal{Y})\}$. The *fundamental group* $\pi_1(G(\mathcal{Y}), T)$ has a presentation with the following generators

$$\left(\bigsqcup_{\sigma \in V(\mathcal{Y})} G_\sigma \right) \bigsqcup E^\pm(\mathcal{Y})$$

and the following relations

$$\left\{ \begin{array}{l} \text{the relations in the groups } G_\sigma, \\ (a^+)^{-1} = a^- \text{ and } (a^-)^{-1} = a^+, \\ a^+ b^+ = g_{a,b} (ab)^+ \text{ for all composable edges } a, b, \\ \varphi_a(g) = a^+ g a^- \text{ for all } a \in E(\mathcal{Y}), g \in G_{i(a)} \\ a^+ = 1 \text{ for all } a \in T \end{array} \right\}$$

Definition 2.44 (Universal Cover of a Complex of Groups). Given a complex of groups $G(\mathcal{Y})$ and a maximal tree T in $\mathcal{Y}$, let $G = \pi_1(G(\mathcal{Y}), T)$. The *universal cover* is a scwol $\mathcal{X}$ defined by

$$V(\mathcal{X}) = \left\{ (gG_\sigma, \sigma) \mid \sigma \in V(\mathcal{Y}), gG_\sigma \in G/G_\sigma \right\}$$

$$E(\mathcal{X}) = \left\{ (gG_{i(a)}, a) \mid a \in E(\mathcal{Y}), gG_{i(a)} \in G/G_{i(a)} \right\}$$

$$i(gG_{i(a)}, a) = (gG_{i(a)}, i(a)) \quad t(gG_{i(a)}, a) = (ga^- G_{t(a)}, t(a))$$

The group G acts on $\mathcal{X}$ by left multiplication with quotient $\mathcal{Y}$.

For the remainder of this section we fix a group G acting on a simply connected simplicial complex X with quotient Y , following Convention 2.41. Let $\mathcal{X}, \mathcal{Y}$ be their scwolifications and let $X_b = |\mathcal{X}|, Y_b = |\mathcal{Y}|$ be their geometric realizations.

We briefly recall how to induce a complex of groups structure on G . The interested reader can refer to [BH99, III.C] for more details. Begin by choosing for each object $\tau \in V(\mathcal{Y})$ a lift $\tilde{\tau} \in V(\mathcal{X})$. For each $a \in E(\mathcal{Y})$, this induces a unique choice of lift $\tilde{a} \in E(\mathcal{X})$ so that $i(\tilde{a}) = \widetilde{i(a)}$, and then we choose a (not unique) $h_a \in G$ so that $h_a t(\tilde{a}) = \widetilde{t(a)}$. For $\tau \in V(\mathcal{Y})$, we set $G_\tau = \text{Stab}_G(\tilde{\tau})$ and for each $a \in E(\mathcal{Y})$, the homomorphism $\psi_a : G_{i(a)} \rightarrow G_{t(a)}$ is simply conjugation by h_a inside G . For composable edges a, b , we set $g_{a,b} := h_a h_b h_{ab}^{-1}$. If X is a tree, there are no twisting elements and this is simply the graph of groups decomposition of G as in Bass–Serre theory, only with more categorical language.

Definition 2.45. A *complex of spaces compatible with the G action on X* consists of the following.

1. For each simplex $\sigma \subset X$, a space X_σ . For each face of $\sigma' \subset \sigma$, a map $\varphi_{\sigma', \sigma} : X_\sigma \rightarrow X_{\sigma'}$ such that if $\sigma'' \subset \sigma' \subset \sigma$, we have $\varphi_{\sigma'', \sigma'} \varphi_{\sigma', \sigma} = \varphi_{\sigma'', \sigma}$.

2. For each $g \in G$ and simplex $\sigma \subset X$, a homeomorphism $g : X_\sigma \rightarrow X_{g\sigma}$, so that whenever $\sigma' \subset \sigma$ or $h \in G$, the following diagrams commute.

$$\begin{array}{ccc} X_\sigma & \xrightarrow{g} & X_{g\sigma} \\ \varphi_{\sigma',\sigma} \downarrow & & \downarrow \varphi_{g\sigma',g\sigma} \\ X_{\sigma'} & \xrightarrow{g} & X_{g\sigma'} \end{array} \quad \begin{array}{ccc} X_\sigma & \xrightarrow{gh} & X_{gh\sigma} \\ & \searrow h & \nearrow g \\ & X_{h\sigma} & \end{array}$$

The following two definitions are the language needed to state Haefliger's Theorem 2.48. We use this theorem to prove Proposition 2.50, which we need to combine cusped spaces in our main theorem.

Definition 2.46. [Hae92, Section 1] For $\tau \in V(\mathcal{Y})$, let CD_τ be the scwol with $V(CD_\tau) = \{a \in E(\mathcal{Y}) \mid t(a) = \tau\}$ and $E(CD_\tau) = \{(a, b) \mid a, b \text{ composable edges of } \mathcal{Y} \text{ with } t(b) = i(a), t(a) = \tau\}$. For $(a, b) \in E(CD_\tau)$, we set $i(a, b) = ab$, $t(a, b) = a$. For composition, we set $(a, b)(a', b') = (a, bb')$. There is a functor $j_\tau : CD_\tau \rightarrow \mathcal{Y}$ which sends $(a, b) \in E(CD_\tau)$ to b , hence $j_\tau(a) = i(a)$. We write D_τ for the geometric realization of CD_τ . The functor j_τ induces a cellular map $D_\tau \rightarrow Y_b$, which we also denote j_τ . Given $a \in E(\mathcal{Y})$, there is a functor $j_a : CD_{i(a)} \rightarrow CD_{t(a)}$ which sends $(b, c) \in E(CD_{i(a)})$ to (ab, c) . This j_a induces a cellular map $D_{i(a)} \rightarrow D_{t(a)}$, which we also denote j_a .

Definition 2.47. [Hae92, Section 3.3] Let KY be a topological space with a continuous projection $\pi : KY \rightarrow Y_b$. Each object $\tau \in V(\mathcal{Y})$ corresponds to a vertex of Y_b , and we set $Y_\tau = \pi^{-1}(\tau)$. Let $Y(D_\tau)$ be the subset of $D_\tau \times KY$ of pairs (x, y) with $j_\tau(x) = \pi(y)$. Let $Y(j_\tau), \pi_\tau$ be the projections onto the first and second coordinates. We identify Y_τ with the fiber $\pi_\tau^{-1}(\tau) \subset Y(D_\tau)$. Any $a \in E(\mathcal{Y})$ induces a map $Y(j_a) : Y(D_{i(a)}) \rightarrow Y(D_{t(a)})$ sending (x, y) to $(j_a(x), y)$. If s is a section of π over the 1-skeleton of Y_b , then each fiber Y_τ has a basepoint $s(\tau)$. This induces a section s_σ of π_σ over the 1-skeleton of D_τ . With this notation, KY is a *complex of spaces associated to $G(\mathcal{Y})$* if the following hold.

1. For each $\tau \in V(\mathcal{Y})$, Y_τ is connected and there is a retraction $r_\tau : Y(D_\tau) \rightarrow Y_\tau$ which is homotopic to the identity relative to $Y(\tau)$ so that $r_\tau s_\tau(x) = s_\tau(\tau)$ for $x \in D_\tau^{(1)}$.
2. $\pi_1(Y_\tau, s(\tau)) = G_\tau$ and for any $a \in E(\mathcal{Y})$, restricting $Y(j_a)$ to $Y_{i(a)} \subset Y(D_{i(a)})$ gives a basepoint preserving map

$$r_{t(a)} Y(j_a) : Y_{i(a)} \rightarrow Y_{t(a)}$$

which induces ψ_a on fundamental groups.

3. For two composable edges $a, b \in E(\mathcal{Y})$ with $\tau = t(a)$, the edge (a, b) of D_τ maps to a loop in Y_τ under $r_\tau s_\tau$ representing the homotopy class of $g_{a,b}^{-1} \in \pi_1(Y_\tau, s(\tau)) = G_\tau$.

Further, if KY is a cell complex, π is a cellular map, and each Y_τ is a $K(G_\tau, 1)$, then we call KY an *aspherical cellular realization of $G(\mathcal{Y})$* .

Theorem 2.48. [Hae92, Theorem 3.4.1] For each $\tau \in V(\mathcal{Y})$, let Y_τ be a fixed choice of $K(G_\tau, 1)$ with a basepoint $s(\tau)$. For each $a \in E(\mathcal{Y})$, let $\varphi_a : Y_{i(a)} \rightarrow Y_{t(a)}$ be any map realizing ψ_a on fundamental groups. Then there is an aspherical cellular realization $\pi : KY \rightarrow Y$ where $\pi^{-1}(\tau)$ is the given complex Y_τ for each $\tau \in V(\mathcal{Y})$.

We collect some results from Haefliger's proof of the above that we will need.

Corollary 2.49. Suppose that for each $\tau \in V(\mathcal{Y})$, Y_τ is a cell-complex and for each $a \in E(\mathcal{Y})$, φ_a is a cellular map. Let $\pi : KY \rightarrow Y$ be an aspherical cellular realization obtained from Theorem 2.48 using these spaces and maps as input.

1. Each $Y(D_\tau)$ can be given a cell structure using the cell structure of the $Y_{\tau'}$ for $\tau' \in V(\mathcal{Y})$, together with some cubes of dimension at most the dimension of Y .
2. With this cell structure, the retraction r_τ is a cellular map for each $\tau \in V(\mathcal{Y})$. Further, if $\widetilde{r}_\tau$ is a lift of r_τ to universal covers, then $\widetilde{r}_\tau$ restricts to a quasi-isometry of k -skeletons for all $k \geq 0$.

Proof. We briefly sketch Haefliger's construction of KY . It proceeds inductively by building a space KY^k for $k = 0, 1, \dots$. For $k = 0$, KY^0 is the disjoint union of the chosen spaces Y_τ for each $\tau \in V(\mathcal{Y})$. For each $a \in E(\mathcal{Y})$, choose a map $\varphi_a : Y_{i(a)} \rightarrow Y_{t(a)}$ realizing ψ_a on fundamental groups. To construct KY^1 , take KY^0 together with a mapping cylinder for each φ_a , and identify the end of each mapping cylinder with the corresponding space in KY^0 . For a general k , KY^k is a quotient of KY^{k-1} and spaces $Y_{i(a_1)} \times [0, 1]^k$ for each k -tuple of composable arrows $(a_k, \dots, a_1)$ in $E(\mathcal{Y})$. The construction of KY^{k-1} tells us where to glue the faces of $Y_{i(a)} \times \partial[0, 1]^k$, and the extension properties of $K(G, 1)$ spaces allows us to extend across the interior of $Y_{i(a)} \times [0, 1]^k$. For $\tau \in V(\mathcal{Y})$, the retraction $r_\tau : Y(D_\tau) \rightarrow Y_\tau$ is a contraction of these added cubes. For example if Y is a graph and τ is a vertex, then $Y(D_\tau)$ is Y_τ with a mapping cylinder attached for each edge incident to τ , and the retraction is just the standard deformation retraction onto Y_τ .

If Y has dimension n , then $\mathcal{Y}$ does not have any $n+k$ -tuples of composable edges for any $k > 0$, meaning this iterated mapping cylinder construction ends after n steps. Further, if each φ_a for $a \in E(\mathcal{Y})$ is cellular, then each mapping cylinder can be given a cell structure so that each r_τ is the contraction of some cells, meaning r_τ is a cellular map. This also allows us to understand the skeletons of KY ; the k -skeleton $(KY)_k$ is the k -skeleton of each Y_τ for $\tau \in V(\mathcal{Y})$ together with cubes of dimension at most k . This explains 1.

Fix some $\tau \in V(\mathcal{Y})$ and consider the space $Y(D_\tau)$. Because r_τ is a homotopy equivalence and a retraction onto Y_τ , it follows that there is a unique lift of the inclusion $\widetilde{Y}_\tau \rightarrow \widetilde{Y(D_\tau)}$ to universal

covers, so we can identify $\widetilde{Y}_\tau$ with a subset of $\widetilde{Y(D_\tau)}$. Further, r_τ is homotopic to the identity, so if $h_t : Y(D_\tau) \rightarrow Y(D_\tau)$ is the homotopy from the identity to r_τ , then h_τ lifts to a deformation retraction of $\widetilde{Y(D_\tau)}$ onto this unique lift of $\widetilde{Y}_\tau$, say $\widetilde{h}_\tau$. From Haefliger's construction, h_τ is the contraction of the finitely many cubes added between the Y_τ . If $a : \tau' \rightarrow \tau$, we may view $Y_{\tau'}$ as a subset of $Y(D_\tau)$, and the restriction of r_τ is simply φ_a . Thus r_τ inherits the cellular properties of the φ_a , and in particular it can be restricted to any k -skeleton. Further, h_τ translates points a finite distance, so the lift of r_τ also translates points a finite distance. Thus $\widetilde{Y}_\tau$ is cobounded in $\widetilde{Y(D_\tau)}$ and it follows that r_τ is a quasi-isometry, and it can be restricted for a quasi-isometry between any k -skeletons. $\square$

Proposition 2.50. Let $\pi : KY \rightarrow Y$ be a complex of spaces associated to $G(\mathcal{Y})$. Then the universal cover $\widetilde{KY}$ induces a complex of spaces compatible with the G action on X . For $\sigma \in V(\mathcal{X})$, the space X_σ is the image of a lift $\pi_\tau : Y(D_\tau) \rightarrow KY$ to universal covers. For simplices $\sigma' \subset \sigma$ in X , the map $\varphi_{\sigma',\sigma} : X_\sigma \rightarrow X_{\sigma'}$ is a lift of the map $Y(j_a) : Y(D_{i(a)}) \rightarrow Y(D_{t(a)})$ to universal covers, where $a \in E(\mathcal{Y})$ is the image of the morphism $\sigma \rightarrow \sigma'$ of $\mathcal{X}$. If KY is an aspherical cellular realization, the $\varphi_{\sigma',\sigma}$ are cellular maps.

Proof. Lift the map $\pi : KY \rightarrow Y_b$ to a map p between universal covers which is $G = \pi_1(KY)$ -equivariant and makes the following diagram commute.

$$\begin{array}{ccc} \widetilde{KY} & \xrightarrow{p} & X_b \\ \downarrow & & \downarrow \\ KY & \xrightarrow{\pi} & Y_b \end{array}$$

Given a vertex σ of X_b , σ projects to a vertex τ of Y_b . The preimage of Y_τ in $\widetilde{KY}$ is the disjoint union of lifts of $Y_\tau \hookrightarrow KY$ to universal covers, hence $p^{-1}(\sigma)$ is exactly the image of some lift. This image is contained in the image of some lift of $\pi_\tau : Y(D_\tau) \rightarrow KY$ to universal covers, and we set X_σ to be the image of this lift. This is exactly $p^{-1}(D_\sigma)$, as illustrated in this diagram.

$$\begin{array}{ccccc} \widetilde{Y(D_\tau)} & \longrightarrow & X_\sigma & \xrightarrow{p} & D_\sigma \\ \downarrow & & \downarrow & & \downarrow \\ Y(D_\tau) & \xrightarrow{\pi_\tau} & KY & \xrightarrow{\pi} & D_\tau \end{array}$$

Since G acts on $\widetilde{KY}$ and permutes these lifts, we have maps $g : X_\sigma \rightarrow X_{g\sigma}$ for each $g \in G$, $\sigma \in V(\mathcal{X})$ and the triangular diagram in Definition 2.45 commutes.

Given $a \in E(\mathcal{Y})$, we can lift $Y(j_a)$ to a map $\widetilde{Y(D_\tau)} \rightarrow \widetilde{Y(D_{\tau'})}$. There are many choices for this lift, but each identifies $\widetilde{Y(D_\tau)}$ with a subspace of $\widetilde{Y(D_{\tau'})}$. If $\sigma \rightarrow \sigma'$ is a morphism of $\mathcal{X}$ covering $a \in E(\mathcal{Y})$, then we can use some such lift to get an inclusion $\varphi_{\sigma',\sigma} : X_\sigma \rightarrow X_{\sigma'}$. Again, since G permutes these lifts, the square diagram in Definition 2.45 commutes.

To check the first condition of Definition 2.45, suppose a, b are composable edges of $\mathcal{Y}$ with $\tau = i(b), \tau' = t(b) = i(a), \tau'' = t(a)$. Then ab is the unique morphism of $\mathcal{Y}$ with source τ and target τ'' because $\mathcal{Y}$ is the scwolification of Y and hence is a simple scwol. Therefore $j_a j_b = j_{ab}$, which implies $Y(j_a)Y(j_b) = Y(j_{ab})$. It follows that composing lifts of $Y(j_b), Y(j_a)$ gives a lift of $Y(j_{ab})$. If $\sigma'' \subset \sigma' \subset \sigma$ are simplices of X , the corresponding morphisms in $\mathcal{X}$ cover some arrows $a, b \in E(\mathcal{Y})$ as in the previous paragraph. The maps $\varphi_{\sigma'', \sigma'}, \varphi_{\sigma', \sigma}, \varphi_{\sigma'', \sigma}$ are lifts of $Y(j_a), Y(j_b), Y(j_{ab})$, hence $\varphi_{\sigma'', \sigma} = \varphi_{\sigma'', \sigma'} \varphi_{\sigma', \sigma}$ as needed.

Finally, if $\pi : KY \rightarrow Y$ is an aspherical cellular realization, then all these maps can be taken to be cellular. $\square$

The next definition is analogous to Definition 2.2 of [Mar14].

Definition 2.51. A complex of spaces compatible with $G(\mathcal{Y})$ consists of the following.

1. For each simplex τ of Y , a space Y_τ with a G_τ action.
2. For each arrow $a \in E(\mathcal{Y})$, an embedding $\varphi_a : Y_{i(a)} \rightarrow Y_{t(a)}$ which is ψ_a -equivariant, that is, for each $g \in G_{i(a)}$ and $x \in Y_{i(a)}$, we have

$$\varphi_a(g \cdot x) = \psi_a(g) \cdot \varphi_a(x),$$

and such that for every pair of composable edges $a, b \in E(\mathcal{Y})$, we have

$$g_{a,b} \circ \varphi_{ab} = \varphi_a \circ \varphi_b.$$

Martin does the following in section 9 of [Mar14]. Suppose $G(\mathcal{Y})$ is a complex of hyperbolic groups over a finite simplicial complex Y . Beginning with a finite generating set for local groups of cells with maximal dimension, we inductively define a generating set S_τ for each $\tau \in V(\mathcal{Y})$ so that if $a : \tau' \rightarrow \tau$ is a morphism of $\mathcal{Y}$, then $\psi_a(S_{\tau'}) \subset S_\tau$. Let Y_τ be the Rips complex $P_n(\Gamma_\tau)$ where Γ_τ is the Cayley graph of G_τ with respect to S_τ . Because there are finitely many hyperbolic groups here, we can choose n large enough so that each Y_τ is contractible. Whenever $\sigma \subset \sigma'$ in Y , we let $\varphi_{\sigma, \sigma'}$ be the induced map induced on these Rips complexes by $\psi_{\sigma, \sigma'}$.

With this notation, the Y_τ do *not* form a complex of spaces compatible with $G(\mathcal{Y})$ because the last condition fails. Explicitly, if a, b are composable edges of $\mathcal{Y}$, then $\varphi_a \circ \varphi_b$ will map $\Gamma_{i(b)}$ to the image of $\psi_a \psi_b(G_{i(b)})$, which is $Ad(g_{a,b})\psi_{ab}$ by the definition of a complex of groups. On the other hand, $g_{a,b} \varphi_{ab}(\Gamma_{i(b)})$ will be a translation of $\psi_{ab}(\Gamma_{i(b)})$, not a conjugation. In other words, there is a missing $g_{a,b}^{-1}$ on the right of $g_{a,b} \circ \varphi_{ab}$ which makes it different from $\varphi_a \circ \varphi_b$. Thus [Mar14, Proposition 9.4] is subtly flawed. This issue is why we use Theorem 2.48. We will not use it, but the following proposition shows how to construct structures as in Definition 2.51, thereby recovering some of Martin's formalism.

Proposition 2.52. A complex of spaces compatible with the G action on X induces a complex of spaces compatible with $G(\mathcal{Y})$.

Proof. This is essentially a diagram chase. To construct the complex of spaces compatible with $G(\mathcal{Y})$, we set $Y_\tau = X_{\widetilde{\tau}}$ and φ_a to be the composition $\varphi_{\widetilde{t(a)}, h_a \widetilde{i(a)}} h_a : X_{\widetilde{i(a)}} \rightarrow X_{\widetilde{t(a)}}$. It is immediate from (2) in Definition 2.45 that G_τ acts on Y_τ for each $\tau \in V(\mathcal{Y})$ and that the maps φ_a are ψ_a -equivariant. To check that $g_{a,b} \circ \varphi_{ab} = \varphi_a \circ \varphi_b$ for composable edges $a, b \in E(\mathcal{Y})$, consider the following diagram.

$$\begin{array}{ccccc}
 X_{\widetilde{i(b)}} & \xrightarrow{h_b} & X_{h_b \widetilde{i(b)}} & \xrightarrow{h_a} & X_{h_a h_b \widetilde{i(b)}} \\
 & \searrow & \downarrow & \nearrow g_{a,b} & \downarrow \\
 & & X_{\widetilde{t(b)}} = X_{\widetilde{i(a)}} & & X_{h_a \widetilde{i(a)}} \\
 & & \downarrow & & \downarrow \\
 & & X_{\widetilde{t(ab)}} & \nearrow g_{a,b} & X_{\widetilde{t(a)}}
 \end{array}$$

The vertical arrows are the embeddings from the assumed complex of spaces compatible with the G action on X and the horizontal arrows are the homeomorphisms from elements of G . The composition $\varphi_a \circ \varphi_b$ is the staircase path, and the composition $g_{a,b} \circ \varphi_{ab}$ goes along the long diagonal arrow, straight down, then along the short diagonal. The commutativity of each square and triangle comes from Definition 2.45, so the two paths are equivalent. $\square$

3 Constructing $\overline{Z}$

For clarity, we recall the main theorem.

Theorem 1.1. Let $G(\mathcal{Y}) = (G_\sigma, \psi_a, g_{a,b})$ be a nonpositively curved developable complex of groups over a scwol $\mathcal{Y}$, where each G_σ is a relatively hyperbolic group and each ψ_a is the inclusion of a full relatively quasiconvex subgroup. Let $\mathcal{X}$ be the universal cover of $G(\mathcal{Y})$, and let X be the geometric realization of $\mathcal{X}$ equipped with an M_κ structure. Suppose X is δ -hyperbolic and the action of $G = \pi_1(G(\mathcal{Y}))$ on X is acylindrical. Then G is relatively hyperbolic. The maximal parabolic subgroups of G are virtually parabolic subgroups of vertex stabilizers, and the stabilizer of each simplex in X is a full relatively quasiconvex subgroup of G .

Fix once and for all a complex X , a group G , an acylindricity constant A , a hyperbolicity constant δ_0 , and let $Y = X/G$. Recall from Convention 2.41 that stabilizers are considered pointwise.

Assumption 3.1. After rescaling the metric, we may assume that for every simplex $\sigma \subset X$, the distance from σ to the boundary of its closed simplicial neighborhood is at least 1. Equivalently, $d(\sigma, \overline{N}(\sigma) \setminus N(\sigma)) \geq 1$. As X is an M_κ -complex, this will scale κ, δ_0 , but it will not change that $\kappa \leq 0$, or hyperbolicity and CAT(0) properties.

Let $\mathcal{X}, \mathcal{Y}$ be the scwolifications of X and Y . By definition, the simplices of X correspond to objects in $V(\mathcal{X})$. For simplices $\tau' \subset \tau \subset Y$, we will write $[\tau'\tau]$ for the corresponding morphism in $E(\mathcal{Y})$. Choosing a maximal tree T in $\mathcal{Y}$, a lift of T to $\mathcal{X}$, and lifts for each element of $V(\mathcal{Y}) \setminus T$ induces a choice of twisting elements, hence a complex of groups $G(\mathcal{Y}) = (G_\sigma, \psi_a, g_{a,b})$. This identifies G with $\pi_1(G(\mathcal{Y}), T)$ and $\mathcal{X}$ with the universal cover of $G(\mathcal{Y})$ as described in Definition 2.43. We can recover X from this construction as

$$X = \bigsqcup_{g \in G, \tau \in V(\mathcal{Y})} (gG_\tau, \tau) \times \tau / \equiv$$

where $\equiv$ is defined as follows: if $\tau' \subset \tau$ are simplices of Y and $x \in \tau'$, then

$$(gG_\tau, \tau, x) \equiv (g[\tau'\tau]^{-1}G_{\tau'}, \tau', x).$$

In words, the disjoint union above is one simplex for each object of $\mathcal{X}$, and $\equiv$ glues these simplices along faces via the corresponding morphisms in $E(\mathcal{X})$. We will say a simplex $\sigma = (gG_\tau, \tau) \times \tau$ is *labeled by* (gG_τ, τ) .

Lemma 3.2. There is a complex of spaces compatible with the G action on X satisfying the following.

1. For each $\sigma \in V(\mathcal{X})$, there is a cellular retraction $r_\sigma : X_\sigma \longrightarrow C_\sigma$, where C_σ is a cusped space for G_σ , possibly with duplicated edges. This retraction is a quasi-isometry, making X_σ into a δ -hyperbolic space and identifying ∂X_σ with the Bowditch boundary ∂G_σ . We write $\overline{X}_\sigma := X_\sigma \sqcup \partial G_\sigma$.
2. For each pair of simplices $\sigma' \subset \sigma$ of X , the map $\varphi_{\sigma', \sigma} : X_\sigma \longrightarrow X_{\sigma'}$ is an inclusion. It is also a quasi-isometric embedding and extends to an embedding $\varphi_{\sigma', \sigma} : \overline{X}_\sigma \longrightarrow \overline{X}_{\sigma'}$. This extension identifies ∂G_σ with $\Lambda G_\sigma \subset \partial G_{\sigma'}$ and has closed image. Further, $r_\sigma(X_{\sigma'})$ is in a bounded neighborhood of some translate of a Lipschitz map between cusped spaces as in Proposition 2.16.
3. (Dichotomy Property) If two closed simplices σ_1, σ_2 of X intersect in a simplex σ , then exactly one of the following holds.

- (a) Some simplex σ' contains both $\sigma_1 \cup \sigma_2$, and $\varphi_{\sigma,\sigma_1}(X_{\sigma_1}) \cap \varphi_{\sigma,\sigma_2}(X_{\sigma_2}) = \varphi_{\sigma,\sigma'}(X_{\sigma'})$.
- (b) No simplex contains both σ_1, σ_2 , and $\varphi_{\sigma,\sigma_1}(\overline{X_{\sigma_1}}) \cap \varphi_{\sigma,\sigma_2}(\overline{X_{\sigma_2}}) \subset \partial X_\sigma$. In particular, $\varphi_{\sigma,\sigma_1}(X_{\sigma_1}) \cap \varphi_{\sigma,\sigma_2}(X_{\sigma_2}) = \emptyset$.

Proof. For each $\tau \in V(\mathcal{Y})$, G_τ is relatively hyperbolic, so we may apply Proposition 2.26 to construct a suitable $K(G_\tau, 1)$, say Y_τ . The universal cover of Y_τ has 1-skeleton equal to a cusped space for G_τ with possibly duplicated edges, so it is δ -hyperbolic for some δ . For the basepoint $s(\tau) \in Y_\tau$, we choose the image of 1 from the cusped space inside $\widetilde{Y}_\tau$.

Using 3 of Proposition 2.26 for each $a : \tau' \rightarrow \tau$ in $E(\mathcal{Y})$, we can choose a basepoint preserving map $Y_{\tau'} \rightarrow Y_\tau$ which induces ψ_a on fundamental groups and lifts to map between universal covers which restricts to a Lipschitz map on 1-skeletons as described in Proposition 2.16.

Applying Theorem 2.48 to $Y_\tau, s(\tau)$, and these specified maps, we receive an aspherical cellular realization of $G(\mathcal{Y})$, say $\pi : KY \rightarrow Y$. Applying Proposition 2.50 we receive a complex of spaces compatible with the G action on X where each X_σ is the image of a lift of $\pi_\tau : Y(D_\tau) \rightarrow KY$ to universal covers for some $\tau \in V(\mathcal{Y})$. The maps $\varphi_{\sigma,\sigma'}$ for $\sigma \subset \sigma'$ simplices of X are cellular, so they can be restricted to k -skeletons. Using this, we replace each X_σ with the 1-skeleton of X_σ and replace each $\varphi_{\sigma,\sigma'}$ with its restriction to 1-skeletons.

If σ is a simplex of X lying over a simplex τ of Y , then the retraction r_τ lifts to a retraction $\widetilde{r}_\tau$ of X_σ onto a copy of the 1-skeleton of $\widetilde{Y}_\tau$. By our application of Proposition 2.26, $\widetilde{Y}_\tau$ is a cusped space for G_τ with possibly duplicated edges, so it is δ -hyperbolic for some δ , and by Corollary 2.49, $\widetilde{r}_\tau$ is a quasi-isometry from X_σ onto this cusped space. This proves (1).

If $\sigma \subset \sigma'$ are simplices of X lying over simplices $\tau \subset \tau'$ in Y , then the map $\varphi_{\sigma,\sigma'}$ coming from our application of Proposition 2.50 is a lift of the inclusion $Y(D_{\tau'}) \rightarrow Y(D_\tau)$ to universal covers. Inside of $\widetilde{Y(D_{\tau'})}$, there is a copy $\widetilde{Y}_\tau$ which is a cusped space for $G_{\sigma'}$, and $\widetilde{r}_\tau$ identifies this copy with some translate of the image of a Lipschitz map as in Proposition 2.16 by our choices of the map $Y_{\tau'} \rightarrow Y_\tau$. Because $\widetilde{r}_\tau$ is a quasi-isometry, this establishes (2).

For (3), suppose $\sigma_1, \sigma_2, \sigma$ are simplices of X with $\sigma_1 \cap \sigma_2 = \sigma$. The objects in the categories $D_{\sigma_1}, D_{\sigma_2}$ are σ_1, σ_2 together with all higher dimensional simplices which contain σ_1, σ_2 respectively. These categories embed into D_σ , and the only objects in the intersection are simplices which contain both σ_1 and σ_2 . For example, σ_1, σ_2 might be edges meeting in a vertex. If σ_1, σ_2 are sides of a square σ' , then σ' is an object in $D_{\sigma_1} \cap D_{\sigma_2}$. On the other hand, if there is no higher dimensional cell containing both σ_1, σ_2 , then $D_{\sigma_1} \cap D_{\sigma_2} = \emptyset$.

In the case where there is a higher dimensional cell containing $\sigma_1 \cup \sigma_2$, then because X is a CAT(0) M_κ -complex, there is a unique cell of minimal dimension containing $\sigma_1 \cup \sigma_2$, which we call σ' . The uniqueness and minimal dimension implies that $D_{\sigma_1} \cap D_{\sigma_2} = D_{\sigma'}$ as subcategories of D_σ . Down in Y , there are corresponding simplices $\tau = \tau_1 \cap \tau_2$ and $\tau_1 \cup \tau_2 \subset \tau'$. The spaces $X_{\sigma_1}, X_{\sigma_2}$ are lifts to the universal cover of the spaces $Y(D_{\tau_1}), Y(D_{\tau_2})$. These two lifts intersect in a lift of

$Y(D_{\tau'})$, and 3a follows.

If there is no simplex containing both σ_1 and σ_2 , then the spaces $X_{\sigma_1}, X_{\sigma_2}$ do not intersect. This restricts any intersection of the images of $\overline{X_{\sigma_1}}$ and $\overline{X_{\sigma_2}}$ to ∂X_σ , which is exactly the statement of 3b. $\square$

The previous lemma gives us a suitable space for each object of $\mathcal{X}$ and maps between them. We add the geometry of the simplex itself back in and add a helpful label.

Definition 3.3. For $\sigma \in V(\mathcal{X})$, let $\widehat{X}_\sigma = \{\sigma\} \times \sigma \times \overline{X_\sigma} / \sim$, where $(\sigma, x, \xi) \sim (\sigma, x', \xi)$ for each $x, x' \in \sigma$ and $\xi \in \partial G_\sigma$. Write $\widehat{X}_\sigma^\circ := \{\sigma\} \times \sigma \times X_\sigma$ for the points of $\widehat{X}_\sigma$ not in the boundary. For simplices $\sigma \subset \sigma'$ of X , let $\widehat{X}_{\sigma'}|_\sigma$ be the points of $\widehat{X}_{\sigma'}$ with a second coordinate in σ , as in $\widehat{X}_{\sigma'}|_\sigma = \{(\sigma, x, y) \mid x \in \sigma\}$. We extend $\varphi_{\sigma, \sigma'}$ to a map $\widehat{X}_{\sigma'}|_\sigma \rightarrow \widehat{X}_\sigma$ by $\varphi_{\sigma, \sigma'}(\sigma', x, z) = (\sigma, x, \varphi_{\sigma, \sigma'}(z))$.

Since σ is compact, $\widehat{X}_\sigma^\circ$ is still a δ -hyperbolic metric space on which G_σ acts with $\partial \widehat{X}_\sigma^\circ = \partial X_\sigma = \partial G_\sigma$.

Definition 3.4. Let

$$Z \sqcup \partial_{\text{Stab}} G = \left(\bigsqcup_{\sigma \in V(\mathcal{X})} \widehat{X}_\sigma \right) / \simeq$$

where Z is the image of the $\widehat{X}_\sigma^\circ$ and $\partial_{\text{Stab}} G$ is the image of the ∂G_σ . To define $\simeq$, let $\sigma \subset \sigma'$ be simplices of X and $w \in \widehat{X}_{\sigma'}$, and set $w \simeq \varphi_{\sigma, \sigma'}(w)$. More explicitly, if $w = (\sigma', x, z)$, we set $(\sigma', x, z) \simeq (\sigma, x, \varphi_{\sigma, \sigma'}(z))$.

With this notation, G acts on $Z \sqcup \partial_{\text{Stab}} G$ diagonally – for $g \in G$, $w = (\sigma, x, z)$, we have $gw = (g\sigma, gx, gz)$. In particular, if $g \in G_\sigma$, then $gw = (\sigma, x, gz)$.

Definition 3.5. Let $\pi : \bigsqcup_{\sigma \in V(\mathcal{X})} \widehat{X}_\sigma \rightarrow Z \sqcup \partial_{\text{Stab}} G$ be the quotient map. For each simplex $\sigma \subset X$, we write $\pi_\sigma : \widehat{X}_\sigma \rightarrow Z \sqcup \partial_{\text{Stab}} G$ for the restriction to $\widehat{X}_\sigma$, and depending on context we may restrict the domain of π_σ to $\widehat{X}_\sigma^\circ$ or $\partial \widehat{X}_\sigma$.

Lemma 3.6. For any σ , the projection $\pi_\sigma : \widehat{X}_\sigma^\circ \rightarrow Z$ is injective.

Proof. If two points in the disjoint union defining Z are identified by $\simeq$, there must be a chain of relations, say

$$w_0 \simeq w_1 \simeq \cdots \simeq w_{n-1} \simeq w_n$$

where w_0, w_n are the points identified and each $w_i \simeq w_{i+1}$ is an equation of the form defining $\simeq$. Fix a simplex σ . Inducting on n , we show that for any such chain with $w_0, w_n \in \widehat{X}_\sigma^\circ$, we have $w_0 = w_n$.

The defining equation for $\simeq$ doesn't change the coordinate in the simplex, so for any chain, the second coordinate of each w_i is the same, say x . There is no chain of length $n = 1$ because then

the simplex labels could not match. If $n = 2$, the chain goes from σ to another simplex σ' , then back to σ . If $\sigma \subset \sigma'$, this is

$$w_0 = (\sigma, x, \varphi_{\sigma, \sigma'}(z)) \simeq (\sigma', x, z) \simeq (\sigma, x, \varphi_{\sigma, \sigma'}(z)) = w_2,$$

and if $\sigma' \subset \sigma$, this is

$$w_0 = (\sigma, x, z) \simeq (\sigma', x, \varphi_{\sigma', \sigma}(z)) \simeq (\sigma, x, z) = w_2.$$

In the above, we use that $\varphi_{\sigma', \sigma}$ is an embedding, which guarantees that $(\sigma', x, \varphi_{\sigma', \sigma}(z)) \neq (\sigma', x, \varphi_{\sigma', \sigma}(y))$ for $y \neq z$. This proves the base case.

Now suppose that $w_0 = w_n$ for every chain of length n and we have a chain of length $n + 1$ from w_0 to w_{n+1} . We shorten the chain by doing the first two steps in one step. Let σ', σ'' be the simplex labels for w_1, w_2 . If $\sigma'' \subset \sigma' \subset \sigma$, then the chain is

$$w_0 = (\sigma, x, z) \simeq (\sigma', x, \varphi_{\sigma', \sigma}(z)) \simeq (\sigma'', x, \varphi_{\sigma'', \sigma'} \varphi_{\sigma', \sigma}(z)) \simeq \cdots$$

But $\varphi_{\sigma'', \sigma'} \varphi_{\sigma', \sigma}(z) = \varphi_{\sigma'', \sigma}(z)$, so we can shorten this chain, using this instead

$$w_0 = (\sigma, x, z) \simeq (\sigma'', x, \varphi_{\sigma'', \sigma}(z)) \simeq \cdots$$

The case $\sigma \subset \sigma' \subset \sigma''$ is similar. If $\sigma \subset \sigma'' \subset \sigma'$, then the chain must start with

$$w_0 = (\sigma, x, z) \simeq (\sigma', x, z') \simeq (\sigma'', x, \varphi_{\sigma'', \sigma'}(z')) \simeq \cdots$$

where $z = \varphi_{\sigma, \sigma'}(z')$. But then $\varphi_{\sigma, \sigma''} \varphi_{\sigma'', \sigma'}(z') = \varphi_{\sigma, \sigma'}(z') = z$, so again we can shorten this chain in exactly the same way as before. All the other cases are similar. $\square$

Proposition 3.7. There is a G -equivariant projection $p : Z \longrightarrow X$ which projects to the simplex coordinate, $p(\sigma, x, z) = x$. This map can be extended to $p : Z \sqcup \partial X \longrightarrow \overline{X}$ by declaring $p(\eta) = \eta$ for all $\eta \in \partial X$. With the quotient topology on Z and the disjoint union topology on $Z \sqcup \partial X$, p is continuous.

That p is well defined, G -equivariant, and extends to ∂X is immediate from the definition of Z , so the only thing to prove is that p is continuous. We introduce some helpful notation which we will use throughout.

Notation 3.8. For any points $x, y \in \overline{X}$, let $[x, y]$ denote the (possibly infinite) geodesic from x to y . For $x \in X$, let σ_x be the unique simplex of X containing x in its interior. By Lemma 3.6, each $z \in Z$, has a unique representative of the form $(\sigma_{p(z)}, p(z), y)$. We call y the *cusped space coordinate* of z . For an arbitrary z , we write $c(z)$ for the cusped space coordinate of z .

Proof of Proposition 3.7. Suppose U is open in $\overline{X}$. Then $p^{-1}(U) \cap \partial X = U \cap \partial X$, which is clearly open in ∂X . To see $p^{-1}(U) \cap Z$ is open in Z , we must show that for each simplex σ of X , $\pi_\sigma^{-1}p^{-1}(U)$ is open in $\widehat{X}_\sigma^\circ$. For any such σ , $\pi_\sigma^{-1}p^{-1}(U) = \{\sigma\} \times (\sigma \cap U) \times X_\sigma^\circ$. This is a product of open sets in $\widehat{X}_\sigma^\circ$, hence open. $\square$

The following lets us understand a neighborhood of $z \in Z$ as a ball around $p(z)$ and a neighborhood of $c(z) \in \overline{X}_{\sigma_{p(z)}}$.

Lemma 3.9. Let $z \in Z$ and let $x = p(z)$. Suppose U is an open neighborhood of $c(z) \in X_{\sigma_x}$ and $0 < \delta$ is so that $B(x, \delta) \subset st(\sigma_x)$. For each simplex σ containing σ_x , let $W_\sigma = \varphi_{\sigma_x, \sigma}^{-1}(U) \subset \overline{X}_\sigma$. Then

$$W = W_z(U, \delta) = \{z' \in Z \mid p(z') \in B(x, \delta), c(z') \in W_{\sigma_{p(z')}}\}$$

is an open neighborhood of z in Z .

Proof. For any simplex σ of X , either σ is not a simplex of $st(\sigma_x)$ and $\pi_\sigma^{-1}(W) = \emptyset$, or $\sigma_x \subseteq \sigma$ and $\pi_\sigma^{-1}(W) = \{\sigma\} \times (B(x, \delta) \cap \sigma) \times W_\sigma$, which is a product of open sets from $\sigma, \overline{X}_\sigma$, hence is open in $\widehat{X}_\sigma^\circ$. Thus $\pi_\sigma^{-1}(W)$ is open in $\widehat{X}_\sigma^\circ$ for all $\sigma \subset X$, which by the definition of the quotient topology, shows W is open in Z . $\square$

Recall that X is δ_0 -hyperbolic so it has a boundary ∂X and the action of G on X extends to an action on this boundary.

Definition 3.10. Let

$$\overline{Z} = Z \sqcup \partial_{Stab} G \partial X \quad \partial G = \partial_{Stab} G \sqcup \partial X.$$

Our goal is to endow $\overline{Z}$ with a topology so that it is compact and metrizable, show G acts on it as a convergence group with limit set ∂G , and then use Yaman's Theorem 2.5 to conclude that G is relatively hyperbolic. See Figure 4 for a diagram of how the Bowditch boundaries are glued in $\partial_{Stab} G$. Note the inverse relationship between dimension of a cell and the 'size' of its limit set.

3.1 Domains and Their Geometry

The map $p : Z \rightarrow X$ in Proposition 3.7 projects a point in Z to X . The next definition is the analogous concept for $\xi \in \partial_{Stab} G$.

Definition 3.11. Let $\xi \in \partial_{Stab} G$. The *domain* of ξ , denoted $D(\xi)$, is the subcomplex of X spanned by simplices $\sigma \subset X$ so that $\xi \in \pi_\sigma(\partial G_\sigma)$. We use $V(\xi)$ to denote the vertices of $D(\xi)$. If $\eta \in \partial X$, we set $D(\eta) = \{\eta\}$.

Proposition 3.12. If σ is a simplex of X , then the projection $\pi_\sigma : \partial G_\sigma \rightarrow \partial_{Stab} G$ is injective.




Figure 4: On the left is a geometric simplex of X . On the right, each circle represents the Bowditch boundary of the corresponding cell. The arrows represent the maps $\varphi_{\tau, \tau'}$ between boundaries, and $\partial_{Stab} G$ is constructed by identifying points along these arrows. construct Z .

The content of Proposition 3.12 is that for a given ∂G_σ , no points are identified by the relation $\simeq$ in Definition 3.4. If σ is a simplex of X with vertex v , then $\pi_\sigma = \pi_v \varphi_{v, \sigma}$. Because $\varphi_{v, \sigma}$ is injective, to show π_σ is injective, it is enough to show π_v is injective. As in the proof of Lemma 3.6, identifications happen along chains and it suffices to consider identifications made along edges of X , motivating the following definition.

Definition 3.13. Let $\xi \in \partial_{Stab} G$. A ξ -path is the data $\{(v_i)_{0 \leq i \leq n}, (\xi_i)_{0 \leq i \leq n}, (x_i)_{1 \leq i \leq n}\}$ where

- the $v_0, \dots, v_n$ are vertices of X so that v_i is adjacent to v_{i+1} , and $e_i = [v_{i-1}, v_i]$ is the edge between v_{i-1}, v_i ,
- each $\xi_i \in \partial G_{v_i}$ and ξ_i represents the equivalence class of ξ in $\partial_{Stab} G$, so that $\xi = \pi_{v_i}(\xi_i)$,
- each $x_i \in \partial G_{[v_{i-1}, v_i]}$ with $\varphi_{v_{i-1}, e_i}(x_i) = \xi_{i-1}$ and $\varphi_{v_i, e_i}(x_i) = \xi_i$

This data will be denoted $[v_0, \dots, v_n]_\xi$, and the choices of x_i, ξ_i will be implicit. The *support* of a ξ -path is the path in the 1-skeleton of X given by the v_i . If $v_0 = v_n$, we call the ξ -path a ξ -loop.

Notation 3.14. If a group $H < G$ stabilizes a simplex $\sigma \subset X$ pointwise, then $H < G_\sigma$ and we can consider the action of H on $\widehat{X}_\sigma$. Since G_σ acts on $\widehat{X}_\sigma$ as a convergence group, we can consider the limit set of H in this action, denoted $\Lambda_\sigma H \subset \partial G_\sigma$, which is nonempty if H is infinite. If σ' is a face of σ , then H also stabilizes σ' and it follows from the definition of the extended $\varphi_{\sigma', \sigma}$ in Lemma 3.2 that $\varphi_{\sigma', \sigma}(\Lambda_\sigma H) = \Lambda_{\sigma'} H$.

Lemma 3.15. Suppose $H < G$ is infinite and stabilizes a connected subcomplex $K \subset X$ pointwise. If $v_0 \in K$ and $\xi \in \Lambda_{v_0} H \subset \partial G_{v_0}$, then $\pi_{v_0}(\xi) \in \pi_\sigma(\partial G_\sigma)$ for all $\sigma \subset K$, hence $K \subset D(\pi_{v_0}(\xi))$.

Proof. Because K is connected, v_0 can be connected to any $\sigma \subset K$ by a path of simplices $v_0 = \sigma_0, \sigma_1, \dots, \sigma_n = \sigma$, and we prove the claim by induction on n . If $n = 1$, then v_0 is a simplex of σ . By assumption, $\xi \in \Lambda_{v_0} H = \varphi_{v, \sigma}(\Lambda_\sigma H)$, so we can choose $\xi' \in \Lambda_\sigma H \subset \partial G_\sigma$ mapping to ξ . It follows that $\pi_\sigma(\xi') = \pi_{v_0}(\xi)$. This proves the base case.

Now suppose the claim holds for a path of simplices of length n , and $v_0 = \sigma_0, \sigma_1, \dots, \sigma_{n+1}$ is a path of simplices in K of length $n + 1$. By induction, there is some $\xi' \in \Lambda_{\sigma_n} H$ so that $\pi_{\sigma_n}(\xi') = \pi_{v_0}(\xi)$. If $\sigma_{n+1} \subset \sigma$, then

$$\pi_{v_0}(\xi) = \pi_{\sigma_n}(\xi') = \pi_{\sigma_{n+1}} \varphi_{\sigma_{n+1}, \sigma_n}(\xi')$$

and we are finished. If $\sigma_n \subset \sigma_{n+1}$, then as in the base case, $\xi' \in \Lambda_{\sigma_n} H = \varphi_{\sigma_n, \sigma_{n+1}}(\Lambda_{\sigma_{n+1}} H)$, so we can choose some $\xi'' \in \partial G_{\sigma_{n+1}}$ mapping to ξ' , and $\pi_{\sigma_{n+1}}(\xi'') = \pi_{\sigma_n}(\xi') = \pi_{v_0}(\xi)$. $\square$

Lemma 3.16. Let $\xi \in \partial_{Stab} G$ and $[v_0, \dots, v_n]_\xi$ be a ξ -path. Then $H := \bigcap_i G_{v_i}$ is an infinite, full RQC subgroup of G_{v_i} and the $\xi_i \in \partial G_{v_i}$ from the definition of a ξ -path have $\xi_i \in \Lambda_{v_i} H$ for each i .

Proof. It is clear from the definition of ξ -path that the support is contained in $D(\xi)$. Induct on n . If $n = 1$, then the ξ -path is a single edge $e_1 = [v_0, v_1]$. By definition, $G_{v_0} \cap G_{v_1} = G_{e_1}$, and G_{e_1} is infinite because $x_1 \in \partial G_{e_1} \neq \emptyset$. Further, G_{e_1} is a full RQC subgroup of the G_{v_i} by assumption. By definition of a ξ -path, $\xi_i = \varphi_{v_i, e_1}(x_1) \in \Lambda_{v_i} G_{e_1}$. This proves the base case.

Suppose the lemma holds for a ξ -path of length n , and $[v_0, \dots, v_{n+1}]_\xi$ is a ξ -path of length $n + 1$. By induction, $H = \bigcap_{i=1}^n G_{v_i}$ is an infinite, full RQC subgroup of G_{v_i} for each i . Then $H \cap G_{[v_n, v_{n+1}]}$ is the intersection of full RQC subgroups in G_{v_n} . By the Limit Set Property 2.19, this is also a full RQC subgroup in G_{v_n} , and by induction, we have

$$\varphi_{v_n, e_{n+1}}(x_{n+1}) = \xi_n \in \Lambda_{v_n} H \cap \Lambda_{v_n} G_{e_{n+1}} = \Lambda_{v_n}(H \cap G_{e_{n+1}})$$

Because this limit set is nonempty, $H' = H \cap G_{e_n} = \bigcap_{i=1}^{n+1} G_{v_i}$ is infinite. This also implies $x_{n+1} \in \Lambda_{e_{n+1}} H'$, hence $\xi_{n+1} \in \Lambda_{v_{n+1}} H'$. $\square$

Proposition 3.17. For any $\xi \in \partial_{Stab} G$, $D(\xi)$ is convex and $\text{diam}(D(\xi)) \leq A$ where A is the acylindricity constant of G acting on X . Further, there is a uniform bound on the number of simplices in $D(\xi)$.

Proof. Let x, x' be two points of $D(\xi)$, let $\sigma_x, \sigma_{x'}$ be the unique simplices of X with x, x' in their interior, and let v, v' be vertices of $\sigma_x, \sigma_{x'}$. Since $v, v' \in V(\xi)$, there must be a ξ -path joining them, say $[v_0, \dots, v_n]_\xi$ with $v = v_0, v' = v_n$. After making the ξ -path longer, we can assume every

vertex of $\sigma_x, \sigma_{x'}$ is in the support of this ξ -path. By Lemma 3.16, $H := \bigcap_i G_{v_i}$ is infinite, fixes $\sigma_x \cup \sigma_{x'}$ pointwise, and $\xi_i \in \Lambda_{v_i} H$ for each i . Since X is CAT(0), H also stabilizes the geodesic $[x, x']$, so $d(x, x') \leq A$ by the definition of acylindricity. Since x, x' were arbitrary, this shows $\text{diam}(D(\xi)) \leq A$.

Because H fixes $[x, x']$ pointwise, if $[x, x']$ meets the interior of a simplex σ , then $H\sigma = \sigma$ as a set, but H may not fix σ point wise (imagine the plane tiled by unit squares and geodesic between $x = (0, 0)$ and $x' = (1, 1)$ – reflecting across the line $x = y$ will fix the endpoints but will not fix the square that the endpoints are corners of). However because σ has finitely many vertices, a finite index subgroup of H fixes σ pointwise. Since $[x, x']$ meets finitely many simplices by Lemma 2.33, this implies a finite index subgroup of H fixes the closed subcomplex spanned by $[x, x']$ pointwise. Let K be this closed subcomplex spanned by $[x, x']$ and apply Lemma 3.15 with this K , v_0 as the start of our ξ -path, $\xi_0 \in \partial G_{v_0}$, and the finite index subgroup of H , we have that $K \subset D(\xi)$, hence $D(\xi)$ is convex.

Next we prove there is a uniform constant bounding the number of simplices in every domain. Recall that for every pair of simplices $\sigma \subset \sigma'$ of X , $G_{\sigma'}$ is a fully RQC subgroup of G_σ , and therefore has finite height by Corollary 2.23. Because Y is finite, we can choose N so that each such $G_{\sigma'}$ has height at most N in G_σ . Increasing N if necessary, we can also assume that any simplex τ of Y has at most N simplices in $st(\tau)$.

Let σ_0 be simplex of $D(\xi)$ and choose $\xi_0 \in \partial G_{\sigma_0}$ so that $\pi_{\sigma_0}(\xi_0) = \xi$. There are at most N orbits of simplices in $st(\sigma_0)$, and for any given orbit, at most N simplices can be in $D(\xi)$. Indeed, suppose $g_1\sigma, g_2\sigma, \dots, g_m\sigma$ are all simplices of $st(\sigma_0) \cap D(\xi)$ with $g_i \in G_{\sigma_0}$. Then by the Limit Set Property 2.19, we have the following in ∂G_{σ_0} :

$$\xi_0 \in \bigcap_i \Lambda_{\sigma_0} G_{g_i\sigma} = \bigcap_i \Lambda_{\sigma_0} G_\sigma^{g_i} = \Lambda_{\sigma_0} \left(\bigcap_i G_\sigma^{g_i} \right) \neq \emptyset.$$

Since the limit set of a finite group is empty, this implies $\bigcap_i G_\sigma^{g_i}$ is infinite. Because G_σ has height at most N in G_{σ_0} , this implies $m \leq N$. Thus each G orbit of simplices in $st(\sigma_0)$ contributes at most N simplices to $st(\sigma_0) \cap D(\xi)$, and there are at most N G -orbits, so $st(\sigma_0) \cap D(\xi)$ contains at most N^2 simplices.

Fix a vertex $v \in V(\xi)$ and inductively define a sequence of finite subcomplexes by $K_1 = v$ and $K_i := \overline{N}(K_{i-1}) \cap D(\xi)$. That is, K_i is K_{i-1} together with where its closed simplicial neighborhood meets $D(\xi)$. For example $K_2 = \overline{st}(v) \cap D(\xi)$. From the previous paragraph, it is clear that each K_i has finitely many simplices.

By Lemma 2.34, there is some constant M so that any geodesic in X of length at most A meets at most M simplices. Because $D(\xi)$ has diameter at most A , for any $x \in D(\xi)$, $[v, x]$ meets at most M simplices of $D(\xi)$. It follows that $[v, x] \subset K_M$, hence $K_M = D(\xi)$ and $D(\xi)$ has finitely many simplices. This bound depends only on N and M , which depend only on $G(\mathcal{Y})$, X and A , so it is

uniform for all domains. $\square$

Corollary 3.18. [Finiteness] There exists an integer $d_{\max}$ satisfying the following.

1. For every $\xi \in \partial_{\text{Stab}} G$, $D(\xi)$ contains at most $d_{\max}$ simplices and any geodesic in the open simplician neighborhood of $D(\xi)$ meets at most $d_{\max}$ simplices.
2. Any closed simplex of X is the union of at most $d_{\max}$ open simplices.

Proof. The first point follows immediately from the Containment Lemma 2.35 and the previous lemma. The second point follows from the finiteness of $Y = X/G$. $\square$

The integer $d_{\max}$ will be an essential tool going forward. The convexity of domains also allows us to upgrade the Dichotomy Property in Lemma 3.2.

Corollary 3.19. If two simplices σ_1, σ_2 of X intersect in a face σ , then exactly one of the following holds.

1. There is a simplex σ' containing $\sigma_1 \cup \sigma_2$ so that

$$\varphi_{\sigma, \sigma_1}(\overline{X_{\sigma_1}}) \cap \varphi_{\sigma, \sigma_2}(\overline{X_{\sigma_2}}) = \varphi_{\sigma, \sigma'}(\overline{X_{\sigma'}}).$$

2. There is no simplex containing both σ_1, σ_2 , and

$$\varphi_{\sigma, \sigma_1}(\overline{X_{\sigma_1}}) \cap \varphi_{\sigma, \sigma_2}(\overline{X_{\sigma_2}}) \subset \partial X_{\sigma}.$$

Proof. The only difference between this and the dichotomy property in Lemma 3.2 is in the first option, where the property is extended to the intersection in the boundary. Suppose we are in the first case and $\xi \in \varphi_{\sigma, \sigma_1}(\partial G_{\sigma_1}) \cap \varphi_{\sigma, \sigma_2}(\partial G_{\sigma_2})$. Simplices are convex, so if we choose points $x_1 \in \sigma_1$ and $x_2 \in \sigma_2$, then the geodesic $[x_1, x_2]$ is contained in σ' and connects points of $D(\xi)$, so $\sigma' \subset D(\xi)$. $\square$

Proof of Proposition 3.12. Fix a vertex v of X and suppose $\pi_v(\xi) = \pi_v(\xi') = \zeta \in \partial_{\text{Stab}} G$. Then there must be a ξ -loop $\{(v_i)_{0 \leq i \leq n}, (\xi_i)_{0 \leq i \leq n}, (x_i)_{1 \leq i \leq n}\}$ with $\xi = \xi_0$ and $\xi' = \xi_n$. Without loss of generality, we can assume the support of the ξ -path is injective and let P be the 1-complex given by this loop.

It is clear that $P \subset D(\zeta)$. Because X is CAT(0) and domains are convex, we can contract the loop P in $D(\zeta)$ by pulling along geodesics to v_0 . This gives a map of the disc into $D(\zeta)$, and after a homotopy, we can assume the image lies in the 2-skeleton of $D(\zeta)$. This yields a finite, contractible 2-complex with boundary P , say H . We will call such a subcomplex a *hull* of P . The result will follow immediately from the following claim.

Claim 3.19.1. For any such hull, we have $\xi = \xi'$ in ∂G_v .

Proof of Claim. Let d be the number of 2-cells in H and we prove the claim by induction on d . For simplicity, we assume all 2-cells are triangles, and explain at the end how to modify this to suit squares and more.

If $d = 1$, then P is the boundary of a single triangle, say σ , and $n = 3$. Because the hull is contained in $D(\xi)$, there is some ξ_σ so that $\varphi_{v_0, \sigma}(\xi_\sigma) = \xi \in \partial G_{v_0}$. The commutativity of the maps φ implies ξ_0 completely determines all the data of the ξ -path. Explicitly, because $\varphi_{v_0, \sigma} = \varphi_{v_0, e_1} \varphi_{e_1, \sigma}$, it follows that $\varphi_{e_1, \sigma}(\xi_\sigma) = x_1 \in \partial G_{e_1}$, and hence $\varphi_{v_1, \sigma}(\xi_\sigma) = \xi_1$. Continuing around the triangle, we have

$$\xi = \varphi_{v_0, \sigma}(\xi_\sigma) = \varphi_{v_0, e_3} \varphi_{e_3, \sigma}(\xi_\sigma) = \varphi_{v_0, e_3}(x_3) = \xi'.$$

This completes the base case. Now assume any hull with d 2-cells has $\xi = \xi'$ and let H be a hull with $d + 1$ 2-cells. Suppose σ is a 2-cell of H containing the edge $e_1 = [v_0, v_1]$. We show how to remove σ and get a smaller hull. Because $\sigma \subset D(\xi)$, we can choose $\xi_\sigma \in \partial G_\sigma$ so that $\varphi_{e_1, \sigma}(\xi_\sigma) = x_1 \in \partial G_{e_1}$. There are two cases.

If another side of σ is contained in P , say $e_2 = [v_1, v_2]$, then the third edge of σ is $e := [v_0, v_2]$. Set $x = \varphi_{e, \sigma}(\xi_\sigma)$. The commutativity of the inclusions yields the following, new ξ -loop:

$$(v_0, v_2, v_3, \dots, v_n), (\xi_0, \xi_2, \xi_3, \dots, \xi_n), (x_1, x, x_3, \dots, x_n).$$

The closure of $H \setminus \sigma$ gives a hull of this new loop containing d 2-cells.

If no other side of σ is contained in P , then let w be the remaining vertex of σ which is not in P . Let $e'_1 = [v_0, w]$ and $e'_2 = [w, v_1]$, and let x'_1, x'_2 be the image of ξ_σ in $\partial G_{e'_1}, \partial G_{e'_2}$ and $\xi_w = \varphi_{w, \sigma}(\xi_\sigma)$. Again the commutativity of the diagram yields the following new ξ -loop:

$$(v_0, w, v_1, \dots, v_n), (\xi_1, \xi_w, \xi_1, \xi_2, \dots, \xi_n), (x'_1, x'_2, x_3, \dots, x_n).$$

Again, the closure of $H \setminus \sigma$ gives a hull for the new loop containing at most d triangles, and we are done.

For cells which are not triangles, the base case $d = 1$ is the same, except that n will equal the number of sides of the single cell. For the inductive step, instead of cutting around a triangle, we can cut around an n -gon from the hull in a similar way. $\square$

With the claim proven, we have proven the lemma. $\square$

4 Geometric Tools

4.1 Nesting and Families

The following is key to defining open neighborhoods of points in $\partial_{Stab}G$.

Definition 4.1. (ξ -family) Let $\xi \in \partial_{Stab}G$ and for each vertex $v \in V(\xi)$, let $\xi_v \in \partial G_v$ be the representative of ξ . A ξ -family is choice of open neighborhoods $\{U_v \subset \overline{X_v} \mid \xi_v \in U_v, v \in V(\xi)\}$ so that for every edge $[v, v']$ of $D(\xi)$ and $x \in \overline{X_{[v, v']}}$, we have

$$\varphi_{v, [v, v']}(x) \in U_v \iff \varphi_{v', [v, v']}(x) \in U_{v'}.$$

Equivalently, $\varphi_{v, [v, v']}^{-1}(U_v) = \varphi_{v', [v, v']}^{-1}(U_{v'})$ for each edge $[v, v']$ of $D(\xi)$.

The following is a basic topological fact which we leave to the reader.

Lemma 4.2. Let B be a topological space and let $C_1, \dots, C_n$ be a finite collection of closed sets in B . Let $\sim$ be the relation on the abstract disjoint union $\bigsqcup_i C_i$ with $x \sim y$ whenever $x = y$ in B , and let $C = \bigsqcup_i C_i / \sim$. The natural map $p : C \rightarrow B$ is a homeomorphism onto $\bigcup_i C_i$. In particular, the quotient topology on C , the final topology induced by the projections $C_i \rightarrow C$, and the subspace topology on $\bigcup_i C_i$ are all equal.

Proposition 4.3. Let $\xi \in \partial_{Stab}G$, $v_0 \in V(\xi)$, and let U_{v_0} be an open neighborhood of ξ in $\overline{X_{v_0}}$. Then U_{v_0} can be extended to a ξ -family.

Proof. Suppose Σ is some set of closed simplices in $D(\xi)$ and $\{U_\sigma \subset \overline{X_\sigma}, \sigma \in \Sigma\}$ is a collection of open neighborhoods of ξ . We will call this collection of neighborhoods *aligned* if for each pair of simplices $\sigma \subset \sigma'$ in Σ , we have

$$U_\sigma \cap \varphi_{\sigma, \sigma'}(\overline{X_{\sigma'}}) = \varphi_{\sigma, \sigma'}(U_{\sigma'}).$$

Because $\varphi_{\sigma, \sigma'}$ is an embedding, the equation above is equivalent to $\varphi_{\sigma, \sigma'}^{-1}(U_\sigma) = U_{\sigma'}$. If Σ contains all of $D(\xi)$ and the collection of open neighborhoods is aligned, then the sets $\{U_v, v \in V(\xi)\}$ will be a ξ -family.

Consider the set $\{d(\sigma, v_0), \sigma \subset D(\xi)\}$. It is finite because $D(\xi)$ has finitely many simplices, and we enumerate it as $\{0 = t_0 < t_1 < \dots < t_m\}$. Let $\Sigma(s) = \{\sigma \subset D(\xi), d(\sigma, v_0) \leq t_s\}$ where elements of $\Sigma(s)$ are open simplices. Because the distance between sets is an infimum, we have $d(v_0, \sigma) = d(v_0, \bar{\sigma})$ for any simplex $\sigma \subset X$, but this distance may be realized on the boundary of a simplex. If $d(\sigma, v_0) = t_i$ so that $\sigma \in \Sigma(i)$ and $\sigma \subset \sigma'$, then $d(\sigma', v_0) \leq t_i$, hence $st(\sigma) \subset \Sigma(i)$. In particular, $\Sigma(0) = st(v_0)$.

For each $\sigma \subset D(\xi)$, we can choose a $y_\sigma \in \bar{\sigma}$ so that $d(y_\sigma, v_0) = d(\sigma, v_0)$. This choice is unique; if there were two choices for y_σ , say y and y' , then the geodesic $[y, y']$ is contained in $\bar{\sigma}$ because

simplices are convex, but the CAT(0) inequality implies $[y, y']$ comes strictly closer to v_0 than y or y' . This would imply y, y' are not actually the closest points in σ to v_0 , a contradiction.

There are finitely many simplices $\omega \in \Sigma(s) \setminus \Sigma(s-1)$ so that y_ω is in the interior of ω . For such an ω , any simplex $\sigma \subset st(\omega) \cap D(\xi) \setminus \Sigma(s-1)$ must have $y_\sigma = y_\omega$. For such a σ , ω must be the unique face of $\bar{\sigma}$ in $\Sigma(s)$, for if $\omega' \subset \sigma$ was another such face, we must have $y_{\omega'} = y_\sigma = y_\omega$, hence $\omega = \omega'$. We can enumerate these simplices where y_ω is in the interior of ω as $\omega_1, \dots, \omega_r$, and from this discussion we have

$$\Sigma(s) \setminus \Sigma(s-1) = D(\xi) \cap \bigsqcup_j (st(\omega_j) \setminus \Sigma(s-1)).$$

We imagine inflating a balloon centered at v_0 and recording the time and place where the balloon touches each closed simplex for the first time. These are the $t_0, \dots, t_m$ and y_σ . At time t_s , the balloon touches some *open* simplices for the first time – these are like the ω in the previous paragraph – and we add all of these open simplices and their disjoint stars to $\Sigma(s-1)$ to form $\Sigma(s)$. We inductively build an aligned collection of open neighborhoods for each $\Sigma(i)$.

For $\sigma \in \Sigma(0) = st(v_0)$, we must set $U_\sigma = \varphi_{v_0, \sigma}^{-1}(U_{v_0})$. To check these choices are aligned, suppose $v_0 \subset \sigma \subset \sigma'$ and $\varphi_{\sigma, \sigma'}(z) \in U_\sigma \cap \varphi_{\sigma, \sigma'}(\overline{X_{\sigma'}})$. Then $\varphi_{v_0, \sigma} \varphi_{\sigma, \sigma'}(z) = \varphi_{v_0, \sigma'}(z) \in U_{v_0}$ by definition of $U_\sigma = \varphi_{v_0, \sigma}^{-1}(U_{v_0})$, hence $z \in \varphi_{v_0, \sigma'}^{-1}(U_{v_0}) = U_{\sigma'}$. This shows the $\subseteq$ containment in the definition of aligned, and the $\supseteq$ is similar.

Suppose we have an aligned collection of sets $\{U_\sigma, \sigma \in \Sigma(s-1)\}$ and we extend this to an aligned collection for $\Sigma(s)$. Let $\omega_1, \dots, \omega_r$ be the simplices of $\Sigma(s) \setminus \Sigma(s-1)$ so that y_{ω_i} is in the interior of ω_i , let τ be any of the ω_i , and we focus on finding U_τ . Let $st(\tau) \cap \Sigma(s-1) = \{\sigma_1, \dots, \sigma_\ell\}$ be the simplices we have already made choices for, and set

$$C = \bigcup_i \varphi_{\tau, \sigma_i}(\overline{X_{\sigma_i}}) \quad U = \bigcup_i \varphi_{\tau, \sigma_i}(U_{\sigma_i}).$$

Claim 4.3.1. For each i , $U \cap \varphi_{\tau, \sigma_i}(\overline{X_{\sigma_i}}) = \varphi_{\tau, \sigma_i}(U_{\sigma_i})$.

Given the claim, we show how to finish the proof. Because each φ_{τ, σ_i} is injective, the claim implies $\varphi_{\tau, \sigma_i}^{-1}(U) = U_{\sigma_i}$ which is open in $\overline{X_{\sigma_i}}$, hence U is open in the final topology on C induced by the φ_{τ, σ_i} . By Lemma 4.2, this final topology on C coincides with the subspace topology on $C \subset \overline{X_\tau}$, so this claim allows us to choose U_τ open in $\overline{X_\tau}$ so that $U_\tau \cap C = U$, hence $U_\tau \cap \varphi_{\tau, \sigma_i}(\overline{X_{\sigma_i}}) = \varphi_{\tau, \sigma_i}(U_{\sigma_i})$ and U_τ is compatible with our choices for $\Sigma(s-1)$. For the simplices $\sigma \subset st(\tau) \cap D(\xi) \setminus \Sigma(s-1)$ that we have not made choices for yet, we must set $U_\sigma = \varphi_{\tau, \sigma}^{-1}(U_\tau)$. Because the $\Sigma(s) \setminus \Sigma(s-1)$ is a disjoint union of parts of the $st(\omega_i)$, we can repeat this process for each ω_i without making contradicting choices. To check this collection of sets is aligned, suppose $\sigma \subset \sigma'$ are simplices of $\Sigma(s)$. If $\sigma \in \Sigma(s-1)$, then so is σ' and by assumption our choices for $\Sigma(s-1)$

are aligned. If $\sigma \notin \Sigma(s-1)$, then $\sigma \subset st(\omega_i)$ for some i and $U_\sigma = \varphi_{\omega_i, \sigma}^{-1}(U_{\omega_i})$. If $\sigma' \in \Sigma(s-1)$, then $U_{\sigma'} = \varphi_{\omega_i, \sigma'}^{-1}(U_{\omega_i})$ by the claim and the choice of U_{ω_i} , and if $\sigma' \in \Sigma(s) \setminus \Sigma(s-1)$, then $U_{\sigma'} = \varphi_{\omega_i, \sigma'}^{-1}(U_{\omega_i})$ again, this time by definition. From here, we can repeat our reasoning for $\Sigma(0)$ and see that our choices are aligned. It remains to prove the claim.

Proof of Claim. Fix σ_i and suppose $z \in U \cap \varphi_{\tau, \sigma_i}(\overline{X_{\sigma_i}})$. Since $U = \bigcup_j \varphi_{\tau, \sigma_j}(U_{\sigma_j})$, we can fix j so that $z \in \varphi_{\tau, \sigma_j}(U_{\sigma_j})$. Because $\tau \subset \sigma_i \cap \sigma_j$, we can apply the Dichotomy Corollary 3.19.

If there is a simplex σ' containing $\sigma_i \cup \sigma_j$, then we must have $\sigma' = \sigma_k$ for some k and $\varphi_{\tau, \sigma_i}(\overline{X_{\sigma_i}}) \cap \varphi_{\tau, \sigma_j}(\overline{X_{\sigma_j}}) = \varphi_{\tau, \sigma_k}(\overline{X_{\sigma_k}})$, hence $z \in \varphi_{\tau, \sigma_k}(\overline{X_{\sigma_k}})$. We can pull z back along φ_{τ, σ_j} to some point in $U_{\sigma_j} \cap \varphi_{\sigma_j, \sigma_k}(\overline{X_{\sigma_k}})$ which equals $\varphi_{\sigma_j, \sigma_k}(U_{\sigma_k})$ because the choices for $\Sigma(s-1)$ are aligned. This means

$$z \in \varphi_{\tau, \sigma_k}(U_{\sigma_k}) = \varphi_{\tau, \sigma_i} \varphi_{\sigma_i, \sigma_k}(U_{\sigma_k}) \subset \varphi_{\tau, \sigma_i}(U_{\sigma_i}).$$

This last containment again uses that our choices were aligned, and is exactly what we need.

If we are in the second case of the dichotomy, there is no simplex containing σ_i, σ_j but their images only intersect in ∂G_τ , hence $z \in \varphi_{\tau, \sigma_i}(\partial G_{\sigma_i}) \cap \varphi_{\tau, \sigma_j}(U_{\sigma_j})$. Choose points x_i, x_j in the interior of σ_i, σ_j so that $d(x_i, y_\tau), d(x_j, y_\tau) < 1$ but also $d(x_i, v_0), d(x_j, v_0) < t_{s-1}$. The scaling assumption 3.1 and the definition of $\Sigma(s-1)$ implies $[x_j, x_i] \subset st(\tau) \cap \Sigma(s-1)$. Let $\sigma_j = \tau_0, \tau_1, \dots, \tau_q = \sigma_i$ be the path of simplices along $[x_j, x_i]$. Because $[x_j, x_i] \subset \Sigma(i-1)$, U_{τ_k} is already defined for each k , and because $[x_j, x_i] \subset st(\tau)$, we can compare the images of these sets in $\overline{X_\tau}$. Further, z represents a point of $\partial_{Stab} G$, $D(z)$ is convex, and $x_j, x_i \in D(z)$, so this path of simplices is contained in $D(z)$.

We show that $z \in \varphi_{\tau, \tau_k}(U_{\tau_k})$ implies $z \in \varphi_{\tau, \tau_{k+1}}(U_{\tau_{k+1}})$. Since $z \in \varphi_{\tau, \tau_0}(U_{\tau_0})$ by assumption, this will imply $z \in \varphi_{\tau, \tau_q}(U_{\tau_q}) = \varphi_{\tau, \sigma_i}(U_{\sigma_i})$, completing the proof of the claim. There are two cases. If $\tau_{k+1} \subset \tau_k$, then this is immediate since $z \in \varphi_{\tau, \tau_k}(U_{\tau_k}) \subset \varphi_{\tau, \tau_{k+1}}(U_{\tau_{k+1}})$. If $\tau_k \subset \tau_{k+1}$, then $z \in \varphi_{\tau, \tau_k}(U_{\tau_k})$ by assumption, and $z \in \varphi_{\tau, \tau_{k+1}}(\partial G_{\tau_{k+1}})$ because $\tau_{k+1} \subset D(z)$, so z is the image of a point $z' \in U_{\tau_k} \cap \varphi_{\tau_k, \tau_{k+1}}(\overline{X_{\tau_{k+1}}}) \subset \overline{X_{\tau_k}}$. Because our choices for $\Sigma(s-1)$ are aligned, this implies $z' \in \varphi_{\tau_k, \tau_{k+1}}(U_{\tau_{k+1}})$, so there is some $z'' \in U_{\tau_{k+1}}$ so that

$$z = \varphi_{\tau, \tau_k}(z') = \varphi_{\tau, \tau_k} \varphi_{\tau_k, \tau_{k+1}}(z'') = \varphi_{\tau, \tau_{k+1}}(z'').$$

Thus $z \in \varphi_{\tau, \tau_{k+1}}(U_{\tau_{k+1}})$, as desired. Now by assumption, $z \in \varphi_{\tau, \tau_0}(U_{\tau_0}) = \varphi_{\tau, \sigma_j}(U_{\sigma_j})$, so repeating this for $1, 2, \dots, q$, we have $z \in \varphi_{\tau, \tau_q}(U_{\tau_q}) = \varphi_{\tau, \sigma_i}(U_{\sigma_i})$, completing the proof of the claim. $\square$

With the claim proven, the lemma is complete. $\square$

Corollary 4.4. Let $\xi \in \partial_{Stab} G$ and for each $v \in V(\xi)$, suppose U_v is a neighborhood of ξ in $\overline{X_v}$. Then there exists a ξ -family $\mathcal{V}$ so that $V_v \subset U_v$ for each $v \in V(\xi)$.

Proof. If $\mathcal{V}, \mathcal{W}$ are ξ -families, then the sets $\{U_v \cap W_v, v \in V(\xi)\}$ are also a ξ -family. Indeed, if

$e = [v, v']$ is an edge of $D(\xi)$, then

$$\varphi_{v,e}^{-1}(U_v \cap W_v) = \varphi_{v,e}^{-1}(U_v) \cap \varphi_{v,e}^{-1}(W_v) = \varphi_{v',e}^{-1}(U_{v'}) \cap \varphi_{v',e}^{-1}(W_{v'}) = \varphi_{v',e}^{-1}(U_{v'} \cap W_{v'}),$$

where the middle equality follows from $\mathcal{V}, \mathcal{W}$ being ξ -families. For each $v \in V(\xi)$, we can use the previous+ proposition to get a ξ -family extending U_v , then take a finite intersection of these ξ -families as above to get the desired ξ -family. $\square$

Notation 4.5. To lighten notation, we will drop the maps $\varphi_{\sigma,\sigma'}$. Whenever $\sigma \subset \sigma'$, we will view $\overline{X_{\sigma'}}$ as a subset of $\overline{X_{\sigma}}$. Because $\varphi_{\sigma,\sigma''} = \varphi_{\sigma,\sigma'}\varphi_{\sigma',\sigma''}$ whenever $\sigma \subset \sigma' \subset \sigma''$, this is unambiguous. We will also identify an element of $\partial_{Stab}G$ with its representatives in each $\overline{X_{\sigma}}$ for $\sigma \subset D(\xi)$. With this convention, Definition 4.1 can be restated as the following: If $\xi \in \partial_{Stab}G$, then a ξ -family is a choice of open neighborhoods $U_v \subset \overline{X_v}$ of ξ for each vertex $v \in V(\xi)$ so that for every edge $[v, v']$ of $D(\xi)$ and $x \in \overline{X_{[v,v']}}$, we have

$$x \in U_v \iff x \in U_{v'}.$$

Equivalently, $U_v \cap \overline{X_{[v,v']}} = U_{v'} \cap \overline{X_{[v,v']}}$.

Definition 4.6. Let $\xi \in \partial_{Stab}G$, v a vertex of $D(\xi)$, and U a neighborhood of ξ in $\overline{X_v}$. We say that a subneighborhood $V \subset U$ is *nested* in U if its closure is contained in U and for every simplex σ of $st(v)$ not contained in $D(\xi)$ we have

$$\overline{X_{\sigma}} \cap V \neq \emptyset \implies \overline{X_{\sigma}} \subset U.$$

Lemma 4.7. Let $\sigma, (\sigma_n)_n$ be simplices of X so that $\sigma \subset \sigma_n$ for each n and each σ_n lies over the same simplex of Y . Then there is a subsequence, still denoted $(\sigma_n)_n$, and a point $\xi \in \partial G_{\sigma}$ so that $\overline{X_{\sigma_n}}$ converges uniformly to ξ in $\overline{X_{\sigma}}$.

Proof. This is the Convergence Property 2.24 combined with the r_{σ} of Lemma 3.2. The σ_n all correspond to cosets of the same subgroup of G_{σ} because they all lie over the same simplex of Y , so we can write $\sigma_n = g_n \sigma'$ for some simplex $\sigma' \subset st(\sigma)$. Then $r_{\sigma}(\overline{X_{\sigma'}})$ is the image of Lipschitz map between cusped spaces as in Proposition 2.16, and $r_{\sigma}(\overline{X_{\sigma_n}}) = g_n r_{\sigma}(\overline{X_{\sigma'}})$. Applying 3.8 4.7, we receive a subsequence and a point $\xi \in \partial G_{\sigma}$ so that $g_n r_{\sigma}(\overline{X_{\sigma'}}) \rightarrow \xi$ uniformly. Since r_{σ} moves points a finite distance in X_{σ}° , it follows that $g_n X_{\sigma_n} \rightarrow \xi$ in $\overline{X_{\sigma}}$. $\square$

The following lemmas will not be used until Section 7, but we include them here because of their similarity to the previous lemma.

Lemma 4.8. Suppose $\sigma \subset \sigma_1 \cap \sigma_2$ are simplices of X and $(a_n)_n$ is a sequence in G_{σ} . Then there is a sequence $(k_n)_n$ in G_{σ_1} and a subsequence of $(a_n)_n$ so that the sets $k_n a_n \overline{X_{\sigma_2}}$ are either constant or converge to a point in $\partial G_{\sigma} \setminus \partial G_{\sigma_1}$.

Proof. As in the proof the Convergence Property 4.7, this is simply Lemma 2.25 rephrased in our context using the maps r_σ of Lemma 3.2. $\square$

Lemma 4.9. Let v be a vertex of X , and let P be a finite index subgroup of a maximal parabolic subgroup of G_v fixing a point $\xi \in \partial G_v$. Then there exists a compact subset $K \subset \overline{X_\sigma} \setminus \{\xi\}$ so that

1. $P(\partial G_v \cap K) = \partial G_v \setminus \{\xi\}$, that is, $\partial G_v \cap K$ is a coarse fundamental domain for P acting on $\partial G_v \setminus \{\xi\}$, and
2. for any simplex $\sigma \subset st(v)$ with G_σ finite, there is some $p \in P$ with $p\overline{X_\sigma} \cap K = \emptyset$.

Proof. As with Lemma 4.8, this follows from Lemma 2.10 and the maps r_σ of Lemma 3.2. For $\mathbb{F}$ we choose a set of representatives of the G_v orbits of simplices in $st(v)$ with finite stabilizers. $\square$

Lemma 4.10. (Nesting) Let $\xi \in \partial_{Stab} G$, $v \in V(\xi)$, and U a neighborhood of ξ in $\overline{X_v}$. Then there exists a subneighborhood $\xi \in V \subset U$ so that V is nested in U .

Proof. Because $\overline{X_v}$ is metrizable, we can choose a countable basis $(V_n)_n$ of neighborhoods for $\xi \in \overline{X_v}$. For a contradiction, suppose that no V_n is nested in U , that is, for every n , there is a simplex $\sigma_n \subset st(v) \setminus D(\xi)$ so that $\overline{X_{\sigma_n}} \cap V_n \neq \emptyset$ but also $\overline{X_{\sigma_n}} \not\subset U$. If there were finitely many σ_n , after a subsequence we could assume the $\sigma_n = \sigma$ are constant. But then $\xi \notin \overline{X_\sigma}$ and $\overline{X_\sigma}$ is a closed subset of $\overline{X_v}$, so we could choose V_n small enough so that $\overline{X_\sigma} \cap V_n = \emptyset$ which is a contradiction. Thus there must be infinitely many σ_n . After a subsequence, we can assume the sequence $(\sigma_n)_n$ is injective.

Because there are finitely many G orbits of simplices, we can pass to a subsequence so all the σ_n lie in the same G_v orbit. The Convergence Property 2.24 then implies the $\overline{X_{\sigma_n}}$ converge uniformly to a point of ∂G_v , say ξ' . If $\xi' \neq \xi$, then we could choose disjoint neighborhoods V', V_n separating ξ' and ξ and take n large enough so that $\overline{X_{\sigma_n}} \subset V'$. But this contradicts that $\overline{X_{\sigma_n}} \cap V_n \neq \emptyset$, so $\xi' = \xi$. But then we could choose n large enough so that $\overline{X_{\sigma_n}} \subset U$, contradicting the choice of σ_n . Thus no such sequence can exist. $\square$

Definition 4.11. Let $\xi \in \partial_{Stab} G$ and let $\mathcal{U}, \mathcal{U}'$ be two ξ -families. We say $\mathcal{U}'$ is *nested* in $\mathcal{U}$ if for every vertex $v \in V(\xi)$, U'_v is nested in U_v . Further, we say $\mathcal{U}$ is *n-nested* in $\mathcal{U}$ if there exist ξ -families

$$\mathcal{U}' = \mathcal{U}^0 \subset \mathcal{U}^1 \subset \dots \subset \mathcal{U}^n = \mathcal{U}$$

so that $\mathcal{U}^i$ is nested in $\mathcal{U}^{i+1}$ for each $i = 0, \dots, n-1$.

Lemma 4.12. Given $\xi \in \partial_{Stab} G$, a ξ -family $\mathcal{U}$, and an integer $n > 0$, there exists a ξ -family $\mathcal{U}'$ which is n -nested in $\mathcal{U}$.




Figure 5: In this example, $D(\xi)$ is the single bold edge, $N(\xi)$ is the bold edge together with the interior of each of the thinner edges, and $Lk(\xi)$ is only the interior of these thinner edges. The neighborhood $D^\varepsilon(\xi)$ is everything contained in the dotted region.

Proof. To get a ξ -family which is nested in a given $\mathcal{U}$, apply the Nesting Lemma 4.10 to each $v \in V(\xi)$, $U_v \subset \overline{X_v}$ to get a family of sets $\{W'_v, v \in V(\xi)\}$ so that each W'_v is nested in U_v . Apply Corollary 4.4 to $\{W'_v, v \in V(\xi)\}$ to get a ξ -family $\mathcal{W} = \{W_v, v \in V(\xi)\}$, and it follows that $\mathcal{W}$ is nested in $\mathcal{U}$. To get a ξ -family which is n -nested in $\mathcal{U}$, repeat this process n times. $\square$

4.2 The Crossing Lemma and Refined Families

Notation 4.13. For a fixed basepoint $v_0 \in X$ and any point $x \in \overline{X}$, let γ_x denote the unit speed parameterization of $[v_0, x]$. Recall that for any $x \in X$, σ_x is the unique simplex containing x in its interior.

Definition 4.14. Let $\xi \in \partial_{Stab} G$ and $\varepsilon \in (0, 1)$. We use $D^\varepsilon(\xi)$ for the open ε -neighborhood of $D(\xi)$ in X , $N(\xi) = N(D(\xi))$ for the open simplicial neighborhood of $D(\xi)$, and

$$Lk(\xi) := Lk(D(\xi)) = N(D(\xi)) \setminus D(\xi)$$

for the simplicial link of $D(\xi)$. By Assumption 3.1, $D^\varepsilon(\xi) \subset N(\xi)$. See Figure 4.2 for an example.

Definition 4.15. We say a geodesic γ in X enters $D^\varepsilon(\xi)$ if $\gamma(t_0) \in D^\varepsilon(\xi)$ for some t_0 . We say γ goes through $D^\varepsilon(\xi)$ if γ enters $D^\varepsilon(\xi)$ and there is some $t_1 > t_0$ so that for all $t \geq t_1$, $\gamma(t) \notin D^\varepsilon(\xi)$.

Definition 4.16 (Exit Simplex). Fix a basepoint $v_0 \in X$ and let $x \in \overline{X}$, $\xi \in \partial_{Stab} G$, and $\varepsilon \in (0, 1)$. If γ_x enters $D^\varepsilon(\xi)$, then the *exit simplex* of x , denoted $\sigma_{\xi, \varepsilon}(x)$ is the last simplex met by γ_x in $D^\varepsilon(\xi)$. More explicitly, if $x \in D^\varepsilon(\xi)$, then $\sigma_{\xi, \varepsilon}(x) = \sigma_x$, and if $x \notin D^\varepsilon(\xi)$, then $\sigma_{\xi, \varepsilon}(x)$ is the last simplex of $Lk(\xi)$ met by γ_x .

Definition 4.17 (Cones). Let $\xi \in \partial_{Stab} G$, let $\mathcal{U}$ be a ξ -family, and let $\varepsilon \in (0, 1)$. Define $\widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ as the set of points $x \in \overline{X} \setminus D(\xi)$ so that




Figure 6: This figure shows various ways geodesics from a basepoint v_0 can interact with $D^\varepsilon(\xi)$ for some $\xi \in \partial_{Stab}G$. For example they can not meet it at all, they can enter $D^\varepsilon(\xi)$ but avoid $D(\xi)$, meet $D(\xi)$ and go through, etc. For the point x , we mark where γ_x exits $D^\varepsilon(\xi)$, and the simplex containing this point is $\sigma_{\xi,\varepsilon}(x)$.

1. the geodesic γ_x from the base point v_0 to x enters $D^\varepsilon(\xi)$,
2. for every vertex $v \in \sigma_{\xi,\varepsilon}(x) \cap V(\xi)$, $\overline{X_{\sigma_{\xi,\varepsilon}(x)}} \subset U_v$ in $\overline{X_v}$.

Further, let $\widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ be the points of $\widetilde{\text{Cone}}_{\xi,\varepsilon}(\xi)$ so that γ_x goes through $D^\varepsilon(\xi)$. Equivalently, $\text{Cone}_{\mathcal{U},\varepsilon}(\xi) = \{x \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi) \mid d(x, D(\xi)) \geq \varepsilon\}$.

Definition 4.18. For $\xi \in \partial_{Stab}G$ and a ξ -family $\mathcal{U}$, let $N_{\mathcal{U}}(\xi)$ be the union of simplices $\sigma \subset Lk(\xi)$ such that for some vertex v of $\sigma \cap D(\xi)$, we have $\overline{X_\sigma} \cap U_v \neq \emptyset$ in $\overline{X_v}$.

Proposition 4.19 (Consistency). Suppose $\xi \in \partial_{Stab}G$, $\mathcal{U}$ is a ξ -family, $\sigma \subset Lk(\xi)$ is a simplex. If there is some vertex $v \in \sigma \cap V(\xi)$ so that $\overline{X_\sigma} \subset U_v$, then for every vertex $w \in \sigma \cap V(\xi)$, we have $\overline{X_\sigma} \subset U_w$.

Proof. Suppose $v \in \sigma \cap V(\xi)$ is as in the assumption and $w \in \sigma \cap V(\xi)$ is some other vertex. Because closed simplices are convex, the geodesic $[v, w]$ is contained in $\overline{\sigma}$, hence there is some path in the 1-skeleton of $\sigma \cap D(\xi)$ from v to w , say $v = v_0, v_1, \dots, v_n = w$. We induct on n .

If $n = 1$, v and w are connected by an edge $e = [v, w]$. In $\overline{X_v}$, the assumption is that $\overline{X_\sigma} \subset \overline{X_e} \cap U_v$, and the definition of ξ -family implies $\overline{X_v} \subset \overline{X_e} \cap U_w$ in $\overline{X_w}$. If v, w are joined by a path of length $n + 1$, say $v = v_0, \dots, v_n, v_{n+1} = w$. Letting $e = [v_n, v_{n+1}]$ be the final edge of this path, we repeat the argument of the base case — $\overline{X_\sigma} \subset \overline{X_e} \cap U_{v_n}$ in $\overline{X_{v_n}}$ by induction, so the definition of a ξ -family implies $\overline{X_\sigma} \subset \overline{X_e} \cap U_{v_{n+1}}$ as desired. $\square$

Corollary 4.20. Let $\xi \in \partial_{Stab}G$, $\mathcal{U}$ be a ξ -family, and $\varepsilon \in (0, 1)$. If $x \in \overline{X}$, γ_x enters $D^\varepsilon(\xi)$, and for some vertex $v \in \sigma_{\xi,\varepsilon}(x)$, $\overline{X_{\sigma_{\xi,\varepsilon}(x)}} \subset U_v$ in $\overline{X_v}$, then $x \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$. In other words, the 'for all' in Definition 4.17 (2) can be replaced with 'for some' without changing the definition.

Proof. If x is as in the assumption, then Proposition 4.19 implies $\overline{X_{\sigma_{\xi,\varepsilon}(x)}} \subset U_w$ for every vertex $w \in \sigma_{\xi,\varepsilon}(x) \cap V(\xi)$, which is exactly the definition of $x \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$. $\square$

Fix a vertex v_0 of X as a base point for the remainder of this section.

Lemma 4.21 (Crossing Lemma). Let $\xi \in \partial_{\text{Stab}}G$, and let $\mathcal{U}, \mathcal{U}'$ be two ξ -families, with $\mathcal{U}'$ n -nested in $\mathcal{U}$. Let $\sigma_1, \dots, \sigma_n$ be a path of simplices in $Lk(\xi)$. Suppose $\sigma_1 \subset N_{\mathcal{U}'}(\xi)$. Then for any k and any vertex $v \in \sigma_k \cap D(\xi)$, we have $\overline{X_{\sigma_k}} \subset U_v$ in $\overline{X_v}$.

Proof. We induct on n . For $n = 1$, this is a combination of Definition 4.6 and Definition 4.18. Explicitly, if $\sigma_1 \subset N_{\mathcal{U}'}(\xi)$, then by Definition 4.18, $\overline{X_{\sigma_1}} \cap U'_v \neq \emptyset$ for some vertex v of $\sigma_1 \cap V(\xi)$. Because U'_v is nested in U_v in the sense of Definition 4.6, this implies $\overline{X_{\sigma_1}} \subset U_v$. By Proposition 4.19, we have $\overline{X_{\sigma}} \subset U_w$ for any other vertex w of $\sigma \cap D(\xi)$.

Suppose the result is true for k . Let $\sigma_1, \dots, \sigma_{k+1}$ be a path of simplices in $Lk(\xi)$ and $\mathcal{U}' = \mathcal{U}^0 \subset \dots \subset \mathcal{U}^{k+1} = \mathcal{U}$. By induction, the result is true for the path $\sigma_1, \dots, \sigma_k$ and the nesting families $\mathcal{U}^0 \subset \dots \subset \mathcal{U}^k$, so the only inclusion to prove is the last one.

If $\sigma_k \subset \sigma_{k+1}$, then every vertex v of $\sigma_k \cap D(\xi)$ is also a vertex of σ_{k+1} , and in $\overline{X_v}$ we have

$$\overline{X_{\sigma_{k+1}}} \subset \overline{X_{\sigma_k}} \subset U_v^k \subset U_v^{k+1}.$$

Since this holds for the vertices of $\sigma_k \cap D(\xi)$, it holds for all vertices of $\sigma_{k+1} \cap D(\xi)$ by Proposition 4.19.

Now suppose $\sigma_{k+1} \subset \sigma_k$. If v is a vertex of $\sigma_{k+1} \cap D(\xi)$ then it is also a vertex of σ_k , and by induction we have $\overline{X_{\sigma_k}} \subset U_v^k$, so

$$\overline{X_{\sigma_k}} \subset \overline{X_{\sigma_{k+1}}} \cap U_v^k \neq \emptyset.$$

The definition of nesting implies $\overline{X_{\sigma_{k+1}}} \subset U_v^{k+1}$. $\square$

Recall the definition of $d_{\max}$ from Corollary 3.18.

Definition 4.22 (Refined Families). Fix $n \geq 0$. Given two ξ -families $\mathcal{U}, \mathcal{U}'$, we say $\mathcal{U}'$ is n -refined in $\mathcal{U}$ if it is $F(\max(n, d_{\max}))$ -nested in $\mathcal{U}$, where F is the function described in Lemma 2.36.

Given any $\xi \in \partial_{\text{Stab}}G$, ξ -family $\mathcal{U}$, and $n \in \mathbb{N}$, there always exists a ξ -family $\mathcal{U}'$ which is n -refined in $\mathcal{U}$ by Lemma 4.12. When applying Lemma 2.36, we count open simplices. For example, a single triangle is the union of 3 vertices, 3 edges, and one 2-cell, for a total of 7 simplices. With this convention and point (2) of Corollary 3.18, when we apply Lemma 2.36 with $K = D(\xi)$ for some $\xi \in \partial_{\text{Stab}}G$ and K' a single closed simplex, we get a path of simplices of length at most $F(d_{\max})$.

We chose a vertex as a basepoint for convenience, but all the proofs in this subsection would work with any basepoint $x_0 \in X$ if we replace v_0 with σ_{x_0} when applying Lemma 2.36.

Definition 4.23. For $\xi \in \partial_{Stab}G$, $\mathcal{U}$ a ξ -family, define the *pseudocone* $\text{pCone}_{\mathcal{U}}(\xi)$ to be the set of points $x \in \overline{X} \setminus D(\xi)$ so that there is some simplex $\sigma \subset N_{\mathcal{U}}(\xi) \subset Lk(\xi)$ which meets $\text{Geod}(x, D(\xi))$. We call any such σ a *witness* for x .

Lemma 4.24. For any $\xi \in \partial_{Stab}G$ and basepoint v_0 , there is a ξ -family $\mathcal{V}$ so that $v_0 \notin \text{pCone}_{\mathcal{V}}(\xi)$. Further, if $\xi \notin \partial G_{v_0}$, then $st(v_0) \cap \text{pCone}_{\mathcal{V}}(\xi) = \emptyset$.

Proof. If $v_0 \in D(\xi)$, then v_0 cannot be in $\text{pCone}_{\mathcal{U}}(\xi)$ for any ξ -family $\mathcal{U}$ because $\text{pCone}_{\mathcal{U}}(\xi) \subset X \setminus D(\xi)$ by definition and $\xi \in \partial G_{v_0}$, so we don't need to consider the further statement. Therefore we can assume $v_0 \notin D(\xi)$ so that $\text{Geod}(v_0, D(\xi)) \cap Lk(\xi) \neq \emptyset$. Choose some $\sigma \subset Lk(\xi)$ which meets $\text{Geod}(v_0, D(\xi))$ and fix a vertex $v \in \sigma \cap V(\xi)$. Because $\overline{X_{\sigma}}$ is a closed subset of $\overline{X_v}$ missing ξ , we can choose a neighborhood $U_v \subset \overline{X_v}$ of ξ so that $U_v \cap \overline{X_{\sigma}} = \emptyset$. For every other vertex $w \in V(\xi)$, let $U_w = \overline{X_w}$. Apply Corollary 4.4 to the sets $\{U_w, w \in V(\xi)\}$ to get a ξ -family $\mathcal{V}'$, and let $\mathcal{V}$ be a ξ -family d_{max} -refined in $\mathcal{V}'$.

For a contradiction, suppose σ' is a witness for $v_0 \in \text{pCone}_{\mathcal{V}}(\xi)$. Applying Lemma 2.36 to the connected complex v_0 , the convex complex $D(\xi)$, and the simplices σ', σ , we receive a path of simplices of length at most $F(d_{max})$ from σ' to σ . Because $\sigma' \subset N_{\mathcal{V}}(\xi)$ and $\mathcal{V}$ is d_{max} -refined, hence $F(d_{max})$ -nested, in the family of sets $\{U_w, w \in V(\xi)\}$, the Crossing Lemma 4.21 applied to this path implies $\overline{X_{\sigma}} \subset U_v$. But U_v was chosen to avoid $\overline{X_{\sigma}}$, so this is impossible.

For the further statement, notice that $\xi \notin \partial G_{v_0}$ implies $D(\xi) \cap st(v_0) = \emptyset$ (here it matters that we are using the *open* star). Suppose $\tau \subset st(v_0)$ has $y \in \tau \cap \text{pCone}_{\mathcal{V}}(\xi)$ with σ' as a witness. We apply Lemma 2.36 to the connected complex τ with v_0, y as the points, the convex complex $D(\xi)$, and the simplices σ', σ . Since τ has at most d_{max} simplices too, the path is still of length at most $F(d_{max})$, and the Crossing Lemma 4.21 gives a contradiction as before. $\square$

Lemma 4.25. Let $\xi \in \partial_{Stab}G$ with $\mathcal{V}$ a ξ -family satisfying the conclusions of Lemma 4.24. Let $\mathcal{U} \subset \mathcal{U}' \subset \mathcal{V}$ be ξ -families, with $\mathcal{U}$ d_{max} -refined in $\mathcal{U}'$ and $\mathcal{U}'$ $[d_{max} + F(d_{max})]$ -nested in $\mathcal{V}$. If $x \in \text{pCone}_{\mathcal{U}}(\xi)$, then $[v_0, x]$ meets $D(\xi)$ and $x \in \widetilde{\text{Cone}_{\mathcal{V}, \varepsilon}(\xi)}$ for all $\varepsilon \in (0, 1)$. If $\xi \notin \partial G_{v_0}$, then $[y, x]$ meets $D(\xi)$ for each $y \in st(v_0)$.

Proof. Suppose $x \in \text{pCone}_{\mathcal{U}}(\xi)$ and let σ be a witness.

Claim 4.25.1. $\text{Geod}(v_0, D(\xi)) \cap \text{Geod}(x, D(\xi)) = D(\xi)$, and if $\xi \notin \partial G_{v_0}$, then $\text{Geod}(v_0, D(\xi)) \cap \text{Geod}(st(v_0), D(\xi)) = D(\xi)$.

Proof of Claim. Clearly $D(\xi)$ is contained in the intersection, so the claim is that $D(\xi)$ is the entire intersection. For a contradiction, suppose $y \in \text{Geod}(x, D(\xi)) \cap \text{Geod}(v_0, D(\xi)) \setminus D(\xi)$. Then there are points $z', z'' \in D(\xi)$ so that

$$y \in [x, z'] \cap [v_0, z''] \quad (\star)$$

and simplices σ', σ'' of $Lk(\xi)$ so that $[x, z']$ meets σ' and $[v_0, z'']$ meets σ'' . We apply Lemma 2.36 twice, first to $\sigma_x, D(\xi)$ to get a path of simplices from σ to σ' , then to $\sigma_y, D(\xi)$ to get a path of simplices from σ' to σ'' . Each of these paths has length at most $F(d_{max})$. By assumption, $\sigma \subset N_{\mathcal{U}}(\xi)$, so the Crossing Lemma 4.21 applied to each of these paths implies $\overline{X_{\sigma''}} \subset V_w$ for any vertex w of $\sigma'' \cap V(\xi)$. But then σ'' would be a witness for $v_0 \in \text{pCone}_{\mathcal{V}}(\xi)$, contradicting the choice of $\mathcal{V}$. Thus no such y can exist.

If $\xi \notin \partial G_{v_0}$ and $\tau \subset st(v_0)$, then we can replace v_0 with some $z \in \tau$ in $(\star)$, and find σ', σ'' as before. Again we get a path of simplices of length at most $2F(d_{max})$ from σ to σ'' , and again combining this with the Crossing Lemma 4.21 contradicts the definition of $\mathcal{V}$. This proves the claim. $\square$

Let $K := \text{Geod}(x, D(\xi)) \cup \text{Geod}(v_0, D(\xi))$ and let $[v_0, x]_K$ be the shortest path from v_0 to x in K .

Claim 4.25.2. $[v_0, x]_K = [v_0, x]$ and $[v_0, x]$ meets $D(\xi)$. If $\xi \notin \partial_{stab} G$, then $[y, x]_K = [y, x]$ for all $y \in st(v_0)$ and $[y, x]$ meets $D(\xi)$.

Proof of Claim. The path $[v_0, x]_K$ connects a point of $\text{Geod}(v_0, D(\xi))$ to a point of $\text{Geod}(x, D(\xi))$, so it must go through the intersection of these two sets. By the previous claim, that intersection is $D(\xi)$, so $[v_0, x]_K$ meets $D(\xi)$. Because $D(\xi)$ is convex, $[v_0, x]_K$ meets $D(\xi)$ in a connected segment, so we can write $[v_0, x]_K$ as a concatenation $[v_0, v_1] \cup [v_1, v_2] \cup [v_2, x]$ for some points $v_1, v_2 \in D(\xi)$. Since X is CAT(0), local geodesics are geodesics (see [BH99][II.1.4]), so if we can show $[v_0, x]_K$ is a local geodesic, we will be done. Since $[v_0, v_2]$ and $[v_1, x]$ are both geodesics in X , we only need to deal with the case $v_1 = v_2$.

If $[v_0, x]_K$ is not a local geodesic at $v_1 = v_2$, then for all $\varepsilon > 0$, there are points $a \in [v_0, v_1]$ and $b \in [v_1, x]$ so that $[a, b]$ is strictly shorter than the concatenation $[a, v_1] \cup [v_1, b]$. Then $[a, b]$ cannot meet $D(\xi)$ – if it did we could combine paths from v_0, x to this meeting point and get a path shorter than $[v_0, x]_K$. On the other hand, $D(\xi)$ is a convex subset of a CAT(0) space, so $[a, b]$ stays within ε of $D(\xi)$. The definition of d_{max} then gives us a path of simplices of length at most d_{max} from σ_b to σ_a along $[a, b]$.

Using Lemma 2.36 with $\sigma_x, D(\xi)$ and the simplices σ_b and the witness σ , we get a path of simplices of length at most $F(d_{max})$ from σ to σ_b . Combining this with the path from σ_b to σ_a , the Crossing Lemma 4.21 implies that $\overline{X_{\sigma_a}} \subset U_w \subset V_w$ for any vertex $w \in \sigma_a \cap D(\xi)$. Then σ_a is a witness for $v_0 \in \text{pCone}_{\mathcal{V}}(\xi)$, which again contradicts the choice of $\mathcal{V}$. This contradiction implies $[v_0, x]_K$ is a local geodesic at $v_1 = v_2$, so $[v_0, x]_K = [v_0, x]$.

To prove the second portion, we can replace v_0 with y and argue in exactly the same way, completing the proof of the claim. $\square$

For any $\varepsilon \in (0, 1)$, we apply Lemma 2.36 to $\sigma_x, D(\xi)$ with the witness σ and the exit simplex $\sigma_{\xi, \varepsilon}(x)$, which exists since $[v_0, x]$ goes through $D(\xi)$ by the previous claim. We get a path of simplices of length at most $F(d_{max})$ from σ to $\sigma_{\xi, \varepsilon}(x)$. Since σ is a witness and $\mathcal{U}$ is d_{max} -refined in $\mathcal{U}'$, the Crossing Lemma 4.21 implies $\overline{X_{\sigma_{\xi, \varepsilon}(x)}} \subset U'_v$ for any $v \in \sigma_{\xi, \varepsilon}(x) \cap D(\xi)$, which is the definition of $x \in \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$. This completes the proof of Lemma 4.25. $\square$

Assumption 4.26. When a basepoint v_0 is specified, we will assume that all ξ -families are contained a ξ -family $\mathcal{U}_\xi$ satisfying the assumptions of $\mathcal{U}$ in Lemma 4.25. Hence the conclusion of Lemma 4.25 applies to every $x \in \text{pCone}_{\mathcal{U}}(\xi)$. We will call such ξ -families *admissible with v_0 as basepoint*.

For a given $\xi \in \partial_{\text{Stab}} G$, ξ -family $\mathcal{U}$, and $\varepsilon \in (0, 1)$, we imagine shining a cone of light from v_0 onto $D(\xi)$. Roughly speaking, $\widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ is the shadow cast by $D(\xi)$ and $\mathcal{U}$ from this light. The previous assumption allows the following proposition to work, so that this shadow is a genuine shadow in the sense that everything in $\widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ truly lies behind $D(\xi)$ from the point of view of v_0 .

Proposition 4.27 (Genuine Shadow). Let $\xi \in \partial_{\text{Stab}} G$, $\mathcal{U}$ a ξ -family admissible with v_0 as a basepoint, and $\varepsilon \in (0, 1)$. If $x \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$, then $[v_0, x]$ meets $D(\xi)$, and if $\xi \notin \partial_{\text{Stab}} G$, then $[y, x]$ meets $D(\xi)$ for every $y \in \text{st}(v_0)$.

Proof. By the definition of cones 4.17, $[v_0, x]$ meets $D^\varepsilon(\xi)$, and for any vertex $v \in \sigma_{\xi, \varepsilon}(x) \cap V(\xi)$, we have $\overline{X_{\sigma_{\xi, \varepsilon}(x)}} \subset U_v$. This exit simplex may not be a witness for x because it may not meet $\text{Geod}(x, D(\xi))$. However, we can choose $y \in [v_0, x] \cap \sigma_{\xi, \varepsilon}(x)$, so that $\sigma_y = \sigma_{\xi, \varepsilon}(x)$. Clearly σ_y is a simplex of $Lk(\xi)$ which meets $\text{Geod}(y, D(\xi))$, so σ_y is a witness for $y \in \text{pCone}_{\mathcal{U}}(\xi)$. By Assumption 4.26, $\mathcal{U}$ is contained in a ξ -family $\mathcal{U}_\xi$ satisfying the assumptions of $\mathcal{U}$ in Lemma 4.25. Since $\mathcal{U}$ is contained in $\mathcal{U}_\xi$, σ_y is also a witness for $y \in \text{pCone}_{\mathcal{U}_\xi}(\xi)$, so the conclusion of Lemma 4.25 states that $[v_0, y]$, hence $[v_0, x]$, meets $D(\xi)$. The case where $\xi \notin \partial_{\text{Stab}} G$ is identical. $\square$

4.3 The Refinement and Star Lemmas

The next lemma tells us that if we have ξ families $\mathcal{U}' \subset \mathcal{U}$ with $\mathcal{U}'$ appropriately refined in $\mathcal{U}$, we can move away from a point in the refined cone while staying in the larger cone.

Lemma 4.28 (Refinement Lemma). Let $\xi \in \partial_{\text{Stab}} G$ and let $\mathcal{U}' \subset \mathcal{U}$ be ξ -families with $\mathcal{U}'$ n -refined in $\mathcal{U}$. Let $\sigma_1, \dots, \sigma_n$ be a path of simplices in $X \setminus D(\xi)$. For any $\varepsilon \in (0, 1)$, if there is some $x \in \sigma_1$ so that $[v_0, x_1]$ enters $D^\varepsilon(\xi)$ and $\sigma_{\xi, \varepsilon}(x_1) \subset N_{\mathcal{U}'}(\xi)$, then

$$\bigcup_{i=1}^n \sigma_i \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi).$$

Proof. Let $y \in \bigcup_{i=1}^n \sigma_i$. Because $\sigma_{\xi,\varepsilon}(x_1) \subset N_{\mathcal{U}'}(\xi)$ and $\mathcal{U}'$ is refined in $\mathcal{U}$, we have $x \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$, thus $[v_0, x_1]$ meets $D(\xi)$ by the Genuine Shadow Proposition 4.27. Let v be a vertex of $\sigma_{\xi,\varepsilon}(x_1) \cap D(\xi)$, let σ be a simplex of $Lk(\xi)$ met by $[v, y]$, and let w be a vertex of $\sigma \cap D(\xi)$.

Note that x_1 can be connected to y outside of $D(\xi)$ precisely because the path of simplices avoids $D(\xi)$. Therefore we can apply Lemma 2.36 to the connected subcomplex $\bigcup_{i=1}^n \sigma_i$ and the convex subcomplex $D(\xi)$, using the portion of $[v_0, x_1]$ between $D(\xi)$ and x_1 and $[w, y]$. We receive a path of simplices from $\sigma_{\xi,\varepsilon}(x_1)$ to σ of length at most $F(\max(n, d_{max}))$. Because $\mathcal{U}'$ is n -refined in $\mathcal{U}$ and $\sigma_{\xi,\varepsilon}(x_1) \subset N_{\mathcal{U}'}(\xi)$, the Crossing Lemma 4.21 implies $\overline{X_\sigma} \subset U_w$. Thus σ is a witness for $y \in \text{pCone}_{\mathcal{U}}(\xi)$. By Assumption 4.26, $\mathcal{U}$ is contained in ξ -family $\mathcal{U}_\xi$ satisfying the assumptions of $\mathcal{U}$ in Lemma 4.25. Because $\mathcal{U}$ is contained in $\mathcal{U}_\xi$, we have $\text{pCone}_{\mathcal{U}}(\xi) \subset \text{pCone}_{\mathcal{U}_\xi}(\xi)$, and in particular $y \in \text{pCone}_{\mathcal{U}_\xi}(\xi)$. The conclusion of Lemma 4.25 then applies to y , implying $[v_0, y]$ meets $D(\xi)$. This means $\sigma_{\xi,\varepsilon}(y)$ exists.

Apply Lemma 2.36 again, this time to the portions of $[v_0, x]$, $[v_0, y]$ in between $D(\xi)$ and x, y , and the simplices $\sigma_{\xi,\varepsilon}(x), \sigma_{\xi,\varepsilon}(y)$. We receive a path of simplices of length at most $F(\max(d_{max}, n))$, and the Crossing Lemma 4.21 applied to this path implies the result. $\square$

If x, y, z are points of a CAT(0) space and y, z are very close, then the geodesics $[x, y]$ and $[x, z]$ are very close. The following lemma uses this property to make the exit simplex of y close to the exit simplex of z , where closeness is measured by being in the open star.

Lemma 4.29 (Star Lemma). Let $\xi \in \partial_{\text{Stab}} G$, $\varepsilon \in (0, 1)$, and let $x \in X \setminus D^\varepsilon(\xi)$ be such that $[v_0, x]$ goes through $D^\varepsilon(\xi)$. Then there exists a δ such that for every $y \in B(x, \delta) \setminus D^\varepsilon(\xi)$, $[v_0, y]$ goes through $D^\varepsilon(\xi)$ and

$$\sigma_{\xi,\varepsilon}(y) \subset \text{st}(\sigma_{\xi,\varepsilon}(x)).$$

Proof. Let $T = d(v_0, x)$ and let $\gamma_x : [0, T] \rightarrow X$ parameterize the geodesic segment $[v_0, x]$. By assumption, γ_x goes through $D^\varepsilon(\xi)$, so there is some t_0 with $\gamma_x(t_0) \in \partial D^\varepsilon(\xi)$, as in $d(\gamma_x(t_0), D(\xi)) = \varepsilon$ and $\sigma_{\xi,\varepsilon}(x) = \sigma_{\gamma_x(t_0)}$.

Since X is CAT(0) and $D(\xi)$ is convex, the function $t \rightarrow d(\gamma_x(t), D(\xi))$ is convex. Because γ_x goes through $D^\varepsilon(\xi)$, this convexity implies there is some interval ending in t_0 which γ_x takes into $D^\varepsilon(\xi)$. Combining this with the fact that $\text{st}(\sigma_{\xi,\varepsilon}(x))$ is an open set, there exists some $r > 0$ so that

$$\gamma_x([t_0 - r, t_0]) \subset D^\varepsilon(\xi) \cap \text{st}(\sigma_{\xi,\varepsilon}(x)).$$

Again because stars are open, we can choose τ so that the τ neighborhood of $\gamma_x([t_0 - r, t_0])$ is contained in $\sigma_{\xi,\varepsilon}(x)$. Since $\gamma_x([t_0 - r, t_0])$ is also convex, the τ neighborhood of it is also convex. To see this, consider two points y_1, y_2 in this τ neighborhood and without loss of generality assume y_1 is further from $\gamma_x([t_0 - r, t_0])$ than y_2 is. Then convexity implies all the points along $[y_1, y_2]$ are

closer to $\gamma_x([t_0 - r, t_0])$ than y_1 , hence closer than τ . Let

$$k = \varepsilon - d(\gamma_x(t_0 - r), D(\xi)) > 0.$$

We separate into 2 cases for $d(x, D(\xi)) = \varepsilon$ and $d(x, D(\xi)) > \varepsilon$.

First, if $d(x, D(\xi)) = \varepsilon$, then using the notation above we have $\gamma(t_0) = x$ and $T = t_0$. Set $\delta = \frac{1}{10} \min(k, \tau, r)$. To see this δ works, choose $y \in B(x, \delta) \setminus D^\varepsilon(\xi)$ and let γ_y parameterize $[v_0, y]$. By the triangle inequality and $d(y, x) < \delta < r$, we see that $d(v_0, y) \geq t_0 - r$, so we can consider $\gamma_y(t_0 - r)$. By the CAT(0) inequality, we know $d(\gamma_x(t_0 - r), \gamma_y([t_0 - r, t_0])) < \delta$, thus

$$\begin{aligned} d(D(\xi), \gamma_y(t_0 - r)) &\leq d(D(\xi), \gamma_x(t_0 - r)) + d(\gamma_x(t_0 - r), \gamma_y(t_0 - r)) \quad (\star) \\ &\leq \varepsilon - \underbrace{k + \delta}_{< 0} \leq \varepsilon \end{aligned}$$

So $\gamma_y(t_0 - r) \in D^\varepsilon(\xi)$, and $[v_0, y]$ goes through $D^\varepsilon(\xi)$ since $y \notin D^\varepsilon(\xi)$. Further, $d(y, \gamma_x(t_0)) = d(y, x) < \delta < \tau$, so $\gamma_y(t_0 - r)$ and y are in the τ neighborhood of $\gamma_x([t_0 - r, t_0])$. The definition of τ then implies this terminal portion of $[v_0, y]$ is contained in $st(\sigma_{\xi, \varepsilon}(x))$, and hence so is the simplex $\sigma_{\xi, \varepsilon}(y)$. This completes the proof in the first case.

Now suppose $d(x, D(\xi)) > \varepsilon$. See Figure 7 for a diagram of what we will set up. Since $T > t_0$, we can choose some $0 < s < \min(\frac{1}{10}(T - t_0), \frac{1}{2}\tau)$ and set

$$\delta = \frac{1}{10} \min(k, r, s, d(\gamma_x(t_0 + s), D^\varepsilon(\xi))) > 0.$$

To see δ works, choose $y \in B(x, \delta) \setminus D^\varepsilon(\xi)$. By the triangle inequality and the definition of s , we have

$$d(v_0, y) \geq d(v_0, x) - d(x, y) > T - s > t_0 + \frac{9}{10}(T - t_0).$$

In particular $d(v_0, y) > t_0 + s > t_0 - r$, so the exact argument in $(\star)$ above implies $[v_0, y]$ goes through $D^\varepsilon(\xi)$ and $d(\gamma_y(t_0 - r), \gamma_x(t_0 - r)) < \delta < \tau$. By the CAT(0) inequality, $d(\gamma_y(t_0 + s), \gamma_y(t_0 + s)) < \delta$, but because $\delta < d(\gamma_x(t_0 + s), D^\varepsilon(\xi))$, the ball $B(\gamma_x(t_0 + s), \delta)$ is disjoint from $D^\varepsilon(\xi)$, so $\gamma_y(t_0 + s) \notin D^\varepsilon(\xi)$. Therefore $[v_0, y]$ leaves $D^\varepsilon(\xi)$ between times $t_0 - r$ and $t_0 + s$. We also know

$$B(\gamma_x(t_0 + s), \delta) \subset B(\gamma_x(t_0), \underbrace{s + \delta}_{< \tau}) \subset B(\gamma_x(t_0), \tau).$$

So the geodesic segment $\gamma_y([t_0 - r, t_0 + s])$ connects two points in the τ neighborhood of $\gamma_x([t_0 - r, t_0])$. The definition of τ implies this segment is contained in $st(\sigma_{\xi, \varepsilon}(x))$, and we have $\sigma_{\xi, \varepsilon}(y) \subset st(\sigma_{\xi, \varepsilon}(x))$. This completes the second case. $\square$




Figure 7: The second case of the Star Lemma 4.29.

Corollary 4.30. Let $\xi \in \partial_{\text{stab}} G$, $\mathcal{U}$ a ξ -family, and $\varepsilon \in (0, 1)$. Then $\widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ is open in $\overline{X}$.

Proof. Take $x \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$. By the Genuine Shadow Proposition 4.27, $[v_0, x]$ goes through $D(\xi)$. Suppose for now that $x \in X$.

If $x \in D^\varepsilon(\xi)$, it must be that $\sigma_{\xi, \varepsilon}(\xi) = \sigma_x \subset N_{\mathcal{U}}(\xi)$. We can take ε' so that $B(x, \varepsilon') \subset st(\sigma_x) \cap D^\varepsilon(\xi)$. If $y \in B(x, \varepsilon')$, then $\sigma_x \subset \sigma_y = \sigma_{\xi, \varepsilon}(y)$, hence $\overline{X_{\sigma_y}} \subseteq \overline{X_{\sigma_x}} \subset U_v$ for any vertex v of $\sigma_x \cap D(\xi)$. This is the definition of $y \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$.

If $x \notin D^\varepsilon(\xi)$, the Star Lemma 4.29 provides us with a suitable δ . If $x \notin \partial D^\varepsilon(\xi)$, equivalently $d(x, D(\xi)) > \varepsilon$, we can assume δ is small enough so that $D^\varepsilon(\xi) \cap B(x, \delta) = \emptyset$. As in the previous paragraph, the conclusion of the Star Lemma 4.29 implies $y \in \text{Cone}_{\mathcal{U}, \varepsilon}(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ for all $y \in B(x, \delta)$.

If $d(x, D(\xi)) = \varepsilon$, then $\sigma_{\xi, \varepsilon}(x) = \sigma_x \subset N_{\mathcal{U}}(\xi)$. Decreasing δ if necessary, we can assume $B(x, \delta) \subset st(\sigma_x)$. If $y \in B(x, \delta)$, then either $y \in D^\varepsilon(\xi)$ or not. If $y \in D^\varepsilon(\xi)$, then $\sigma_{\xi, \varepsilon}(y) = \sigma_y \subset st(\sigma_x)$, and as in the first paragraph this implies $y \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$. If $y \notin D^\varepsilon(\xi)$, then the conclusion of the Star Lemma 4.29 implies $\sigma_{\xi, \varepsilon}(y) \subset st(\sigma_x)$, and again this implies $y \in \text{Cone}_{\mathcal{U}, \varepsilon}(\xi)$.

If $x \in \partial X$, then we can choose some x' far along $[v_0, x]$ so that $\sigma_{\xi, \varepsilon}(x) = \sigma_{\xi, \varepsilon}(x')$. Using this x' and the argument above, we can get a δ so that $B(x', \delta) \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$. The set $\{y \in \overline{X} \mid \gamma_y \text{ goes through } B(x', \delta)\}$ defines an open neighborhood of x in $\overline{X}$ and is contained in $\widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$. $\square$

Notice that by definition $\text{Cone}_{\mathcal{U}, \varepsilon}(\xi) \cap \overline{D^\varepsilon(\xi)} \subset \partial \overline{D^\varepsilon(\xi)}$ and the Star Lemma 4.29 implies that points in $\text{Cone}_{\mathcal{U}, \varepsilon}(\xi) \setminus \overline{D^\varepsilon(\xi)}$ have an open neighborhood contained in $\text{Cone}_{\mathcal{U}, \varepsilon}(\xi)$. It is exactly the

points in this intersection which prevent $\text{Cone}_{\mathcal{U},\varepsilon}(\xi)$ from being open, since any ball around them would have points $x \in D^\varepsilon(\xi)$ so that γ_x does not leave $D^\varepsilon(\xi)$.

5 The Topology of $\overline{Z}$

5.1 Definition of the Topology

A point in $\overline{Z}$ is in either Z , $\partial_{\text{Stab}}G$ or ∂X . We define a topology on $\overline{Z}$ by defining a neighborhood basis for each type of point. Consider the example in the introduction. We know $\partial\pi_1(\Sigma) = S^1$, and see $\partial_{\langle\gamma\rangle}\pi_1(\Sigma)$ as a tree of circles by contracting the endpoints of lifts of c . We encourage the reader to reflect on the what topology the tree of circles has as a quotient of S^1 . This topology is what we are emulating in the definitions below.

Recall that each $\widehat{X}_\sigma$ is the product $\{\sigma\} \times \sigma \times \overline{X}_\sigma$, that $\widehat{X}_\sigma^\circ = \{\sigma\} \times \sigma \times X_\sigma$, and that Z is constructed by gluing the $\widehat{X}_\sigma^\circ$ to each other along faces with the $\varphi_{\sigma,\sigma'} : X_{\sigma'} \rightarrow X_\sigma$. We call the second coordinate the *simplex coordinate* and the third coordinate the *cusped coordinate*, see Notation 3.8. Also recall from Proposition 3.7 that $p : Z \sqcup \partial X \rightarrow \overline{X}$ is the G -equivariant map which projects to the simplex coordinate and is the identity on ∂X .

Definition 5.1. For $z \in Z$, let $\mathcal{O}_{\overline{Z}}(z)$ be the collection of open sets of Z containing z .

As a CAT(0) space, we have a natural topology on $\overline{X} = X \cup \partial X$, which we use to define neighborhoods of $\eta \in \partial X \subset \overline{X}$.

Definition 5.2. For a base point $v_0 \in X$ and $\eta \in \overline{X}$, let

$$V_{r,\delta}(\eta) = \{x \in \overline{X} \mid d(v_0, x) > r \text{ and } \gamma_x(r) \in B(\gamma_\eta(r), \delta)\}.$$

The $V_{r,\delta}(\eta)$ form a basis of (not necessarily open) neighborhoods for η in $\overline{X}$, denoted $\mathcal{O}_{\overline{X}}(\eta)$. For any open set $U \subset \overline{X}$ containing η , let

$$V_U(\eta) = p^{-1}(U) \sqcup \{\xi \in \partial_{\text{Stab}}G \mid D(\xi) \subset U\}.$$

To define a neighborhood basis, let $\mathcal{O}_{\overline{Z}}(\eta) = \{V_U(\eta) \mid U \subset \overline{X} \text{ the interior of some } V_{r,\delta}(\eta)\}$.

Lemma 5.3. [Mar14, 6.3] Let $\eta \in \partial X$ with U a neighborhood of η in $\overline{X}$, and $k \geq 0$. Then there exists a neighborhood $U' \subset U$ of η so that $d(X \setminus U, U') > k$.

Definition 5.4. Let $\xi \in \partial_{\text{Stab}}G$ and suppose $\mathcal{U}$ is a ξ -family and $\varepsilon \in (0, 1)$. Let $W_{\mathcal{U},\varepsilon}(\xi)$ be the set of points $z \in Z$ such that $p(z) \in D^\varepsilon(\xi)$ and for every vertex $v \in \sigma_{p(z)} \cap D(\xi)$, the cusped coordinate $c(z)$ is in U_v . Let $\partial W_{\mathcal{U},\varepsilon}(\xi)$ be the points $\xi' \in \partial_{\text{Stab}}G$ such that $D(\xi') \setminus D(\xi) \subset \widehat{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ and for

every vertex v of $D(\xi) \cap D(\xi')$ we have $\xi' \in U_v$. Let

$$V_{\mathcal{U},\varepsilon}(\xi) := W_{\mathcal{U},\varepsilon}(\xi) \cup p^{-1}(\text{Cone}_{\mathcal{U},\varepsilon}(\xi)) \cup \partial W_{\mathcal{U},\varepsilon}(\xi)$$

and let $\mathcal{O}_{\overline{Z}}(\xi) = \{V_{\mathcal{U},\varepsilon}(\xi) \mid \mathcal{U} \text{ a } \xi\text{-family}, \varepsilon \in (0, 1)\}$.

Remark 5.5. Suppose $\xi \in \partial_{\text{Stab}} G$ and $V_{\mathcal{U},\varepsilon}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$. Suppose $z \in Z$ has $p(z) = x \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi) \setminus \text{Cone}_{\mathcal{U},\varepsilon}(\xi)$, then $x \in D^\varepsilon(\xi)$, $\sigma_{\xi,\varepsilon}(x) = \sigma_x$, and for any vertex $v \in \sigma_x \cap V(\xi)$, we have $\overline{X_{\sigma_x}} \subset U_v$ in $\overline{X_v}$. In particular, $c(z) \in U_v$, so $z \in W_{\mathcal{U},\varepsilon}(\xi)$. This shows

$$p^{-1}(\widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)) \subset W_{\mathcal{U},\varepsilon}(\xi) \cup p^{-1}(\text{Cone}_{\mathcal{U},\varepsilon}(\xi)).$$

This means the definition of $V_{\mathcal{U},\varepsilon}(\xi)$ does not change if we take the preimage of a $\widetilde{\text{Cone}}$ instead of a Cone , since we are only adding points which would be in $W_{\mathcal{U},\varepsilon}(\xi)$ anyway.

Note that by Assumption 4.26, the admissible ξ -families in this definition depend on the choice of base point. For unrelated ξ -families $\mathcal{U}, \mathcal{U}'$ and $\varepsilon, \varepsilon' \in (0, 1)$, there is no obvious relationship between $V_{\mathcal{U},\varepsilon}(\xi)$ and $V_{\mathcal{U}',\varepsilon'}(\xi)$, but we do have the following.

Lemma 5.6. Let $\xi \in \partial_{\text{Stab}} G$ with ξ -families $\mathcal{U}, \mathcal{U}'$ and $0 < \varepsilon' \leq \varepsilon \in (0, 1)$. If $\mathcal{U}'$ is $d_{\max}$ -nested in $\mathcal{U}$, then $V_{\mathcal{U}',\varepsilon'}(\xi) \subset V_{\mathcal{U},\varepsilon}(\xi)$.

Proof. Let $x \in \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon'}(\xi)$. Then $[v_0, x]$ meets $D(\xi)$ by the Genuine Shadow Proposition 4.27, and because a geodesic segment in $N(\xi)$ can meet at most $d_{\max}$ simplices, we get a path of simplices of length at most $d_{\max}$ from $\sigma_{\xi,\varepsilon'}(x)$ to $\sigma_{\xi,\varepsilon}(x)$. Because $x \in \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon'}(\xi)$, we know $\sigma_{\xi,\varepsilon'}(x) \subset N_{\mathcal{U}'}(\xi)$, so the Crossing Lemma 4.21 implies that $\overline{X_{\sigma_{\xi,\varepsilon'}(x)}} \subset U_v$ in $\overline{X_v}$ for any vertex v of $\sigma_{\xi,\varepsilon'}(x) \cap V(\xi)$. Thus $x \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ as well. This shows $\widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon'}(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$. If $\eta \in \partial X \cap V_{\mathcal{U}',\varepsilon'}(\xi)$ or $z \in p^{-1}(\widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon'}(\xi))$, this immediately implies $\eta, z \in V_{\mathcal{U},\varepsilon}(\xi)$ as well. If $\xi' \in \partial_{\text{Stab}} G \cap V_{\mathcal{U}',\varepsilon'}(\xi)$, then this implies $D(\xi') \setminus D(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$. For a vertex v of $D(\xi') \cap D(\xi)$, we have $\xi' \in U'_v \subset U_v$, so $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$ too.

If $z \in W_{\mathcal{U}',\varepsilon'}(\xi)$ then by definition $p(z) = x \in D^{\varepsilon'}(\xi)$ and $c(z) \in U'_v$ for every vertex $v \in \sigma_x \cap V(\xi)$. Since $\varepsilon' \leq \varepsilon$ and $U'_v \subset U_v$, we have $z \in W_{\mathcal{U},\varepsilon}(\xi)$. $\square$

The following lemma provides conditions where one half of the definition of $\xi' \in \partial W_{\mathcal{U},\varepsilon}(\xi)$ implies the other half.

Lemma 5.7. Let $\xi, \xi' \in \partial_{\text{Stab}} G$, $\varepsilon \in (0, 1)$, and let $\mathcal{U}'' \subset \mathcal{U}' \subset \mathcal{U}$ be ξ -families, each $d_{\max}$ -refined and $d_{\max}$ -nested in the next.

- (i) If $\xi' \in U'_v$ for some vertex v of $D(\xi) \cap D(\xi')$, then $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$.
- (ii) If $D(\xi') \cap \widetilde{\text{Cone}}_{\mathcal{U}'',\varepsilon}(\xi) \neq \emptyset$, then $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$.

Proof. We begin with (i). By the definition of a ξ -family and the assumption $\xi' \in U'_v \subset \overline{X_v}$, we have $\xi' \in U'_w$ for any vertex w of $D(\xi) \cap D(\xi')$. If $D(\xi') \subset D(\xi)$, this immediately implies $\xi' \in V_{\mathcal{U}',\varepsilon}(\xi) \subset V_{\mathcal{U},\varepsilon}(\xi)$ since the other condition is vacuously true. If not, let K be a connected component of $D(\xi') \setminus D(\xi)$ and choose a simplex $\sigma \subset K \cap Lk(\xi)$, which exists since $D(\xi')$ is connected. Any other simplex of K can be reached by a path of simplices of length at most d_{max} from σ since $D(\xi')$ only has at most d_{max} simplices. For any vertex w of $\sigma \cap D(\xi)$, we have

$$\xi' \in \partial G_\sigma \cap U'_w \implies \sigma \subset N_{\mathcal{U}'}(\xi)$$

so applying the Refinement Lemma 4.28 implies $K \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$. Repeating this for every such component tells us $D(\xi') \setminus D(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$, so $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$.

For (ii), choose $x \in D(\xi') \cap \widetilde{\text{Cone}}_{\mathcal{U}'',\varepsilon}(\xi)$. Because cones are defined to be disjoint from domains, $x \notin D(\xi)$. If $D(\xi) \cap D(\xi') = \emptyset$, then σ_x can be connected to any other simplex of $D(\xi')$ using a path of simplices of length at most d_{max} . The Refinement Lemma 4.28 then implies $D(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon}(\xi)$, hence $\xi \in V_{\mathcal{U}',\varepsilon}(\xi)$.

If $D(\xi') \cap D(\xi) \neq \emptyset$, by taking a geodesic from x to a point of $D(\xi') \cap D(\xi)$, we can connect σ_x to a simplex $\sigma \subset D(\xi') \cap Lk(\xi)$ with a path of simplices of length at most d_{max} . The Refinement Lemma 4.28 implies $\sigma \subset \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon}(\xi)$, hence $\xi' \in \partial G_\sigma \subset U'_v$ in ∂G_v for any vertex v of $\sigma \cap D(\xi)$. Now we can apply (i) and conclude that $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$. $\square$

5.2 A basis of neighborhoods

Let $\mathcal{O}_{\overline{Z}}$ be the collection of sets $\mathcal{O}_{\overline{Z}}(z), \mathcal{O}_{\overline{Z}}(\eta)$ as $\mathcal{O}_{\overline{Z}}(\xi)$ for all $z, \eta, \xi \in \overline{Z}$. In this subsection, we prove $\mathcal{O}_{\overline{Z}}$ forms a basis of neighborhoods for a topology. We begin with the following lemma.

Lemma 5.8 (Filtration). Given $z, z' \in \overline{Z}$, if $U \in \mathcal{O}_{\overline{Z}}(z)$ and $z' \in U$, then there is some $U' \in \mathcal{O}_{\overline{Z}}(z')$ so that $U' \subset U$.

There are 6 different cases depending on which part of $\overline{Z}$ that z, z' are in.

Lemma 5.9 (Case 1: $z, z' \in \partial X$). If $\eta, \eta' \in \partial X$ and $\eta' \in V_U(\eta) \in \mathcal{O}_{\overline{Z}}(\eta)$, then there is a neighborhood $U' \subset \overline{X}$ of η' so that $V_{U'}(\eta') \subset V_U(\eta)$.

Proof. If $\eta' \in V_U(\eta)$ then by definition $\eta' \in U \cap \partial X$, so $U' := U$ is an open neighborhood of η in $\overline{X}$, and it's clear that $V_{U'}(\eta') = V_U(\eta)$, so in particular we have containment. Any other neighborhood of η' contained in U would also work. $\square$

Lemma 5.10 (Case 2: $z \in \partial X, z' \in \partial_{Stab} G$). Let $\eta \in \partial X$, U be a neighborhood of η in $\overline{X}$, and $\xi \in \partial_{Stab} G$ with $\xi \in V_U(\eta)$. Then there is some $V_{\mathcal{U},\varepsilon}(\xi)$ so that $V_{\mathcal{U},\varepsilon}(\xi) \subset V_U(\eta)$.

Proof. By assumption, U is the interior of some $V_{r,\delta}(\eta)$ and $D(\xi) \subset U$. Domains are finite, hence compact, and U is open, so we can choose $\varepsilon \in (0, 1)$ so that $D^\varepsilon(\xi) \subset U$. Let $\mathcal{U}$ be any ξ -family. We claim that $\widetilde{V_{\mathcal{U},\varepsilon}(\xi)} \subset V_U(\eta)$.

For any $x \in \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$, the geodesic from v_0 to x meets $D(\xi)$ by the Genuine Shadow Proposition 4.27, say at a point y . Because $D(\xi) \subset U \subseteq V_{r,\delta}(\eta)$, we know $y \in U$ implies $x \in U$, so $D^\varepsilon(\xi) \cup \widetilde{V_{\mathcal{U},\varepsilon}(\xi)} \subset U$. For any $w \in (Z \cup \partial X) \cap V_{\mathcal{U},\varepsilon}(\xi)$, this implies $p(w) \in U$. For any $\xi' \in \partial_{\text{Stab}} G \cap V_{\mathcal{U},\varepsilon}(\xi)$, we have $D(\xi') \subset D^\varepsilon(\xi) \cup \widetilde{V_{\mathcal{U},\varepsilon}(\xi)} \subset U$. $\square$

Lemma 5.11 (Case 3: $z \in \partial_{\text{Stab}} G, z' \in \partial X$). Let $\eta \in \partial X$, $\xi \in \partial_{\text{Stab}} G$, and $V_{\mathcal{U},\varepsilon}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$. If $\eta \in V_{\mathcal{U},\varepsilon}(\xi)$, then there is some neighborhood $U \subset \overline{X}$ of η so that $V_U(\eta) \subset V_{\mathcal{U},\varepsilon}(\xi)$.

Proof. Since domains are bounded, we can choose N so that $D^\varepsilon(\xi) \subset B(v_0, N)$, and let $x = \gamma_\eta(N + 1)$. Because $\eta \in V_{\mathcal{U},\varepsilon}(\xi)$ and $d(x, D(\xi)) > \varepsilon$, we have $x \in \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$. By Corollary 4.30, there is some $\delta > 0$ so that $B(x, \delta) \subset \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$. Letting $U = V_{N+1,\delta}(\eta)$, we have $V_U(\eta) \subset V_{\mathcal{U},\varepsilon}(\xi)$. $\square$

Lemma 5.12 (Case 4: $z, z' \in \partial_{\text{Stab}} G$). Let $\xi, \xi' \in \partial_{\text{Stab}} G$ and suppose $\xi' \in V_{\mathcal{U},\varepsilon}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$. Then there exists a ξ' -family $\mathcal{V}$ and $\delta \in (0, 1)$ so that $V_{\mathcal{V},\delta}(\xi') \subset V_{\mathcal{U},\varepsilon}(\xi)$.

Proof. We define δ and a preliminary ξ' -family $\mathcal{U}'$ in two claims.

Claim 5.12.1. There exists δ so that if $y \in \overline{D^\delta(\xi')} \setminus D^\varepsilon(\xi)$, then $y \in \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$.

Proof of Claim. If $x \in D(\xi') \cap D^\varepsilon(\xi)$, we can choose δ_x so that $B(x, \delta_x) \subset D^\varepsilon(\xi)$. If $x \in D(\xi') \setminus D^\varepsilon(\xi)$, then $x \in \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$ because $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$. Therefore $[v_0, x]$ goes through $D(\xi)$ by the Genuine Shadow Proposition 4.27 and we can apply the Star Lemma 4.29 to receive a δ_x so that if $y \in B(x, \delta_x) \setminus D^\varepsilon(\xi)$, then $[v_0, y]$ goes through $D^\varepsilon(\xi)$ and $\sigma_{\xi,\varepsilon}(y) \subset st(\sigma_{\xi,\varepsilon}(x))$. In this case, for any vertex $v \in \sigma_x \cap V(\xi)$, we have $\overline{X_{\sigma_{\xi,\varepsilon}(y)}} \subset \overline{X_{\sigma_{\xi,\varepsilon}(x)}} \subset U_v$, hence $y \in \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$. Now $D(\xi') \setminus D^{\varepsilon/2}(\xi)$ is compact, so we can choose finitely many $x_1, \dots, x_m$ so that

$$D(\xi') \setminus D^{\varepsilon/2}(\xi) \subset \bigcup_i B(x_i, \frac{\delta_{x_i}}{2}).$$

Let $\delta' < \min(\frac{\delta_{x_i}}{2}, \frac{\varepsilon}{2})$ and suppose $y \in D^{\delta'}(\xi') \setminus D^\varepsilon(\xi)$. There is some $x \in D(\xi')$ so that $d(y, x) < \delta'$, and because $d(y, D(\xi)) \geq \varepsilon$, the triangle inequality implies $d(x, D(\xi)) \geq \varepsilon - \delta' \geq \frac{\varepsilon}{2}$. So $x \in D(\xi') \setminus D^{\varepsilon/2}(\xi)$, and our choice of open cover gives us some x_i so that $d(x, x_i) < \frac{\delta_{x_i}}{2}$, hence $d(y, x_i) < \delta + \delta_{x_i}/2 \leq \delta_{x_i}$. If $x_i \in D^\varepsilon(\xi)$, then our choice of δ_{x_i} would imply $y \in D^\varepsilon(\xi)$, which is a contradiction. Therefore $x_i \in D(\xi') \setminus D^\varepsilon(\xi)$, and as above this implies $y \in \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$. Thus $D^{\delta'}(\xi') \setminus D^\varepsilon(\xi) \subset \widetilde{V_{\mathcal{U},\varepsilon}(\xi)}$.

Letting $\delta = \frac{\delta'}{2}$, we have $\overline{D^\delta(\xi')} \subset D^{\delta'}(\xi')$. $\square$

Claim 5.12.2. There exists a ξ' -family $\mathcal{U}'$ satisfying the following.

A1. If $v \in V(\xi) \cap V(\xi')$, then $U'_v \subset U_v$.

A2. If $y \in D^\delta(\xi') \cap D^\varepsilon(\xi)$ and there is some $v' \in \sigma_y \cap V(\xi')$ with $\overline{X_{\sigma_y}} \subset U'_{v'}$, then $\sigma_y \cap D(\xi) \subset D(\xi')$ and for every $v \in \sigma_y \cap V(\xi)$, we have $\overline{X_{\sigma_y}} \subset U'_v \subset U_v$.

A3. If $x \in \text{Cone}_{\mathcal{U}', \delta}(\xi')$ leaves $D^\delta(\xi')$ at a point inside $D^\varepsilon(\xi)$, then $x \in \widetilde{\text{Cone}_{\mathcal{U}, \varepsilon}(\xi)}$.

Proof of Claim. For $v' \in V(\xi')$, if $v' \notin V(\xi)$, let $V_{v'} = \overline{X_{v'}}$. Otherwise, $v' \in V(\xi)$ and we let $V_{v'} = U_{v'}$, where $U_{v'}$ comes from the ξ -family $\mathcal{U}$. Note that $U_{v'}$ contains ξ' precisely because $\xi' \in V_{\mathcal{U}, \varepsilon}(\xi)$ and the definition of a ξ -family.

For each $v' \in V(\xi')$, let $st_\xi(v') = \{v \in (V(\xi) \cap \overline{st}(v')) \setminus V(\xi')\}$. In words, $st_\xi(v')$ is the vertices of the closed star $\overline{st}(v')$ which are in $D(\xi)$ but are *not* in $D(\xi')$. It's possible that $st_\xi(v')$ might be empty. For each $v \in st_\xi(v')$, the geodesic $[v', v]$ goes from v' , through a simplex we call $\sigma_{v', v} \subset Lk(v')$, then ends at v . For example if X is a graph, $\sigma_{v', v}$ is simply the edge connecting v, v' , but if X is a CAT(0) cube-complex, then v, v' might be opposite corners of a square, and $\sigma_{v', v}$ is that square. Note that $\sigma_{v', v}$ is not in $D(\xi')$ because $v \notin D(\xi')$. Therefore $\overline{X_{\sigma_{v', v}}}$ is a closed subset of $\overline{X_{v'}}$ not containing ξ' , and we can choose an open neighborhood $W_{v', v} = \overline{X_{v'}} \setminus \overline{X_{\sigma_{v', v}}}$ of ξ' . Set

$$W_{v'} = V_{v'} \cap \left(\bigcap_{v \in st_\xi(v')} W_{v', v} \right).$$

Repeating this process for each vertex $v' \in V(\xi')$, we let $\mathcal{U}'$ be a ξ' -family d_{max} -nested in the collection of sets $\{W_{v'}, v' \in V(\xi')\}$. Explicitly, for each vertex $v' \in V(\xi')$, we have open sets

$$U'_{v'} = U_{v'}^0 \subset U_{v'}^1 \subset \dots \subset U_{v'}^{d_{max}-1} \subset U_{v'}^{d_{max}} \subset W_{v'}$$

where each containment is a nesting. This is our $\mathcal{U}'$. Clearly A1 is immediate by the choice of $V_{v'}$.

Note that because a simplex $\sigma \subset X$ is isometric to the convex hull in some model space of some number of points in *general* position and domains are convex, any simplex of X which meets $D(\xi)$ does so in some maximum dimensional face. The dimension of that face is $|\sigma \cap V(\xi)| - 1$. For example if σ is an edge of $D(\xi)$, then σ meets $D(\xi)$ in σ , which contains 2 vertices, and an edge is a $2 - 1$ -simplex.

Let y, v' be as in the assumption of A2. Because $\varepsilon, \delta < 1$, we know $\sigma_y \subset N(\xi') \cap N(\xi)$, so σ_y meets each of $D(\xi'), D(\xi)$. Let v be any vertex of $\sigma_y \cap V(\xi)$. For a contradiction, suppose $v \notin V(\xi')$, then $v \in st_\xi(v')$ and σ_y (or one of its faces) is $\sigma_{v', v}$, hence $\overline{X_{\sigma_y}} \subseteq \overline{X_{\sigma_{v', v}}}$ in $\overline{X_{v'}}$. Putting this together with the assumption that $\overline{X_{\sigma_y}} \subset U'_{v'}$, we have

$$\overline{X_{\sigma_y}} \subseteq U'_{v'} \cap \overline{X_{\sigma_{v', v}}},$$

but this is impossible because $U'_{v'}$ and $\overline{X_{\sigma_{v', v}}}$ are disjoint by construction. This contradiction implies

$v \in V(\xi')$. By the definition of a ξ' -family and the assumption $\overline{X_{\sigma_y}} \subset U'_{v'}$, we know $\overline{X_{\sigma_y}} \subset U'_v$. Since v was an arbitrary vertex of $\sigma_y \cap D(\xi)$, this proves A2.

Suppose $x \in \text{Cone}_{\mathcal{U}', \delta}(\xi')$ leaves $D^\delta(\xi')$ at a point in $D^\varepsilon(\xi)$, and choose $y \in [v_0, x] \cap D^\varepsilon(\xi)$ just before $[v_0, x]$ leaves $D^\delta(\xi')$ so that $\sigma_y = \sigma_{\xi', \delta}(\xi)$. Then $[y, x]$ gives a path of simplices $\sigma_{\xi', \delta}(x) = \sigma_1, \sigma_2, \dots, \sigma_n = \sigma_{\xi, \varepsilon}(x)$ in $N(\xi)$ with $n \leq d_{\max}$. For $1 \leq k \leq n$, we prove the following.

1. $\sigma_k \subset Lk(\xi')$ and further, $\sigma_k \cap D(\xi) \subset D(\xi')$.
2. For any vertex $v' \in \sigma_k \cap V(\xi')$, $\overline{X_{\sigma_k}} \subset U_{v'}^k$.

Induct on k . For $k = 1$, $y \in D^\delta(\xi')$ by assumption so $\sigma_1 = \sigma_y = \sigma_{\xi', \delta}(x) \subset Lk(\xi')$ already. Because $x \in \text{Cone}_{\mathcal{U}', \delta}(\xi')$, we know A2 applies to y , which is exactly 1 and 2 above. Assuming the result for k , we prove 1 and 2 for $k + 1$ by breaking into two cases.

Suppose $\sigma_k \subset \sigma_{k+1}$. By induction, $\sigma_k \subset Lk(\xi')$, so $\sigma_{k+1} \subset Lk(\xi')$ too. Also by induction, $\overline{X_{\sigma_{k+1}}} \subset \overline{X_{\sigma_k}} \subset U_{v'}^k \subset U_{v'}^{k+1}$ for any vertex $v' \in \sigma_k \cap V(\xi')$, so the definition of ξ' -family implies 2 for $k + 1$. Fix a vertex $v' \in \sigma_{k+1} \cap V(\xi')$ and let v be any vertex of $\sigma_{k+1} \cap V(\xi)$. As in the proof of A2, if $v \notin V(\xi')$, then σ_{k+1} (or one of its faces) must be $\sigma_{v', v}$, and this contradicts $\overline{X_{v'}} \subset U_v^{k+1}$. This means $v \in V(\xi')$, and again because v was arbitrary, we have $\sigma_{k+1} \cap D(\xi) \subset D(\xi)$. This completes the inductive step for this case.

If $\sigma_{k+1} \subset \sigma_k$, then by induction $\sigma_k \subset Lk(\xi')$ so we can fix a vertex $v' \in \sigma_k \cap V(\xi')$. Because this path of simplices is in $N(\xi)$, we know there is some vertex $v \in \sigma_{k+1} \cap V(\xi)$. Because $\sigma_{k+1} \subset \sigma_k$, this v is also a vertex of σ_k , so by induction with 1, we know $v \in V(\xi')$. Because v was arbitrary, this implies 1 for σ_{k+1} . For 2, if $v' \in \sigma_{k+1} \cap V(\xi')$, then by induction we have $\overline{X_{\sigma_k}} \subset \overline{X_{\sigma_{k+1}}} \cap U_{v'}^k \neq \emptyset$, so the definition of nesting implies $\overline{X_{\sigma_{k+1}}} \subset U_{v'}^{k+1}$. This completes the induction.

Consider $k = n$ and recall that $\sigma_n = \sigma_{\xi, \varepsilon}(x)$. By 1, $\sigma_{\xi, \varepsilon}(x) \cap D(\xi) \subset D(\xi')$, so for any vertex $v \in \sigma_{\xi, \varepsilon}(x) \cap V(\xi)$, we have $v \in V(\xi')$ as well. For any such vertex v , 2 and the definition of our ξ' -family implies

$$\overline{X_{\sigma_{\xi, \varepsilon}(x)}} \subset U_v^n \subset V_v = U_v,$$

hence $x \in \widetilde{\text{Cone}_{\mathcal{U}, \varepsilon}(\xi)}$. This completes the proof of A3 and the claim. $\square$

Claim 5.12.3. $\widetilde{\text{Cone}_{\mathcal{U}', \delta}(\xi')} \subset \widetilde{\text{Cone}_{\mathcal{U}, \varepsilon}(\xi)}$.

Proof of Claim. Suppose $x \in \widetilde{\text{Cone}_{\mathcal{U}', \delta}(\xi')}$ and choose $y \in [v_0, x] \in D^\delta(\xi')$ just before $[v_0, x]$ leaves $D^\delta(\xi')$ so that $\sigma_y = \sigma_{\xi', \delta}(x)$. If $y \in D^\delta(\xi') \setminus D^\varepsilon(\xi)$, then the claim defining δ implies y and hence $x \in \widetilde{\text{Cone}_{\mathcal{U}, \varepsilon}(\xi)}$. If $y \in D^\delta(\xi') \cap D^\varepsilon(\xi)$, then $[v_0, x]$ either ends in σ_y or leaves $D^\delta(\xi)$ at a point inside $D^\varepsilon(\xi)$. Then A2 or A3 imply that $\overline{X_{\sigma_y}}$ interacts with the vertices of $D(\xi)$ in the right way so that $x \in \widetilde{\text{Cone}_{\mathcal{U}, \varepsilon}(\xi)}$. $\square$

Let $\mathcal{V}$ be a ξ' -family nested in $\mathcal{U}'$ and note that A1, A2, A3, and the previous claim all apply to $\mathcal{V}$. We show $V_{\mathcal{V}, \delta}(\xi') \subset V_{\mathcal{U}, \varepsilon}(\xi)$ by considering the three kinds of points.

If $z \in V_{\mathcal{V},\delta}(\xi') \cap Z$, then either $p(z) \in \widetilde{\text{Cone}}_{\mathcal{V},\delta}(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ and $z \in V_{\mathcal{U},\varepsilon}(\xi)$ immediately, or $z \in W_{\mathcal{V},\delta}(\xi')$. In this second case, let $y = p(z) \in D^\delta(\xi')$.

If $y \in D^\delta(\xi') \setminus D^\varepsilon(\xi)$, then again the definition of δ implies $y \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$, hence $z \in V_{\mathcal{U},\varepsilon}(\xi)$. If $y \in D^\delta(\xi') \cap D^\varepsilon(\xi)$, then by definition of $W_{\mathcal{V},\delta}(\xi')$, for every vertex $v' \in \sigma_y \cap V(\xi')$ we have $c(z) \in V_{v'} \subset U'_{v'}$. But then $c(z) \in \overline{X_{\sigma_y}} \cap V_{v'}$, so the definition of nesting implies $\overline{X_{\sigma_y}} \subset U'_{v'}$ (this is why we had to nest one more time after all the claims). Thus A2 applies to y and tells us that $\sigma_y \cap V(\xi) \subset D(\xi')$, so every vertex of $\sigma_y \cap V(\xi)$ is also a vertex of $V(\xi')$. Using the rest of the conclusion of A2 and A1, we have $\overline{X_{\sigma_y}} \subset V_v \subset U'_v \subset U_v$, hence $z \in W_{\mathcal{U},\varepsilon}(\xi)$.

If $\eta \in V_{\mathcal{V},\delta}(\xi') \cap \partial X$, then $\eta \in \widetilde{\text{Cone}}_{\mathcal{V},\delta}(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ immediately from the last claim, so $\eta \in V_{\mathcal{U},\varepsilon}(\xi)$.

If $\zeta \in V_{\mathcal{V},\delta}(\xi') \cap \partial_{\text{Stab}} G$, we need to show that $D(\zeta) \setminus D(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ and for all vertices $w \in V(\zeta) \cap V(\xi)$, we have $\zeta \in U_w$.

For the first condition, notice that a point in $D(\zeta) \setminus D(\xi)$ is either inside $D(\xi')$ or outside $D(\xi')$. The points inside $D(\xi')$ are controlled by $\xi' \in V_{\mathcal{U},\varepsilon}(\xi)$ and the points outside $D(\xi')$ are controlled by $\zeta \in V_{\mathcal{V},\delta}(\xi')$. Explicitly, if $x \in D(\zeta) \setminus D(\xi)$, then either $x \in D(\xi') \setminus D(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$, or $x \in D(\zeta) \setminus D(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{V},\delta}(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$.

For the second condition, suppose $w \in V(\zeta) \cap V(\xi)$. If w is also in $V(\xi')$, then because $\zeta \in V_{\mathcal{V},\delta}(\xi')$, we have $\zeta \in V_w \subset U_w$. If $w \notin V(\xi')$, then again using the definition of $\zeta \in V_{\mathcal{V},\delta}(\xi')$ and the previous claim, we have

$$w \in D(\zeta) \setminus D(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{V},\delta}(\xi') \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi),$$

but cones are disjoint from domains by definition, so this would imply $w \notin V(\xi)$, which is a contradiction. This completes the check on vertices of $V(\zeta) \cap V(\xi')$ finishes the lemma. $\square$

Lemma 5.13 (Case 5: $z \in Z, z' \in \partial X$). Let $z \in Z, \eta \in \partial X$ and let U be an open neighborhood of η in ∂X . If $z \in V_U(\eta)$, then there is some neighborhood $U' \in \mathcal{O}_{\overline{Z}}(z)$ so that $U' \subset V_U(\eta)$.

Proof. By definition, U is an open set containing $p(z)$ in $\overline{X}$, so it contains some ball say $B(p(z), \varepsilon)$. Because p is continuous by Proposition 3.7, $U' = p^{-1}(B(p(z), \varepsilon))$ is an open neighborhood of z in Z clearly contained in $V_U(\eta)$. $\square$

Lemma 5.14 (Case 6: $z \in Z, z' \in \partial_{\text{Stab}} G$). Let $z \in Z, \xi \in \partial_{\text{Stab}} G$ and $V_{\mathcal{U},\varepsilon}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$. If $z \in V_{\mathcal{U},\varepsilon}(\xi)$, then there is some neighborhood $U' \in \mathcal{O}_{\overline{Z}}(z)$ so that $U' \subset V_U(\eta)$.

Proof. There are two cases. Let $x = p(z)$.

If $z \in W_{\mathcal{U},\varepsilon}(\xi)$, then by definition $z \in D^\varepsilon(\xi)$ and $c(z) \in U_v$ for every vertex v of $\sigma_x \cap D(\xi)$. Fixing some choice of v , U_v contains a neighborhood U of $c(z)$ in X_v , which we can interpret as

an open neighborhood of $c(z)$ in X_{σ_x} by taking its preimage under φ_{v,σ_x} . We can also choose δ so that $B(x, \delta) \subset st(\sigma_x) \cap D^\varepsilon(\xi)$, and applying Lemma 3.9, we find an open neighborhood of $z \in Z$ contained in $W_{\mathcal{U},\varepsilon}(\xi)$.

If $x \in \text{Cone}_{\mathcal{U},\varepsilon}(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$, then there is some δ so that $B(x, \delta) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ by Corollary 4.30. Then $p^{-1}(B(x, \delta))$ is an open neighborhood of z in Z by Proposition 3.7 and is contained in $W_{\mathcal{U},\varepsilon}(\xi) \sqcup \text{Cone}_{\mathcal{U},\varepsilon}(\xi)$, hence contained in $V_{\mathcal{U},\varepsilon}(\xi)$. $\square$

Proof of Filtration Lemma 5.8. The previous lemmas cover all the cases except $z \in Z, z' \in \partial_{\text{Stab}}G$ and $z, z' \in Z$. This first case never happens because neighborhoods of $z \in Z$ are contained in Z . In the second case, we are given $z' \in U \in \mathcal{O}_{\overline{Z}}(z)$. By definition, U is open in Z , so we can take $U \in \mathcal{O}_{\overline{Z}}(z')$. $\square$

Theorem 5.15. $\mathcal{O}_{\overline{Z}}$ is the basis for a topology on $\overline{Z}$ making it into a second countable space. Further, Z embeds into $\overline{Z}$ as a dense subset.

Proof. We must prove that if $U_1, U_2 \in \mathcal{O}_{\overline{Z}}$ and $z \in U_1 \cap U_2$, then there is some $W \in \mathcal{O}_{\overline{Z}}(z)$ so that $z \in W \subset U_1 \cap U_2$. There are 3 cases.

1. If $z \in Z$, then by the Filtration Lemma 5.8 there are open neighborhoods $W_1, W_2 \subset Z$ so that $W_i \subset U_i$. Then $W = W_1 \cap W_2$ is open in Z and contained in $U_1 \cap U_2$ as desired.
2. If $z = \eta \in \partial X$, then by the Filtration Lemma 5.8 there are open neighborhoods W_1, W_2 of η in $\overline{X}$ so that $V_{W_i}(\eta) \subset U_i$. Then $W = W_1 \cap W_2$ is an open neighborhood of $\eta \in \overline{X}$ and it's clear that

$$V_W(\eta) \subset V_{W_1}(\eta) \cap V_{W_2}(\eta) \subset U_1 \cap U_2.$$

3. If $z = \xi \in \partial_{\text{Stab}}G$, then by the Filtration Lemma 5.8 there are $V_{\mathcal{U}_1,\varepsilon_1}(\xi), V_{\mathcal{U}_2,\varepsilon_2}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$ so that $V_{\mathcal{U}_i,\varepsilon_i}(\xi) \subset U_i$. Let $\mathcal{U}$ be a ξ -family which is $d_{\max}$ -nested in the family of sets $\{(U_1)_v \cap (U_2)_v, v \in D(\xi)\}$ and let $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$. Then Lemma 5.6 implies

$$V_{\mathcal{U},\varepsilon}(\xi) \subset V_{\mathcal{U}_1,\varepsilon_1}(\xi) \cap V_{\mathcal{U}_2,\varepsilon_2}(\xi) \subset U_1 \cap U_2.$$

To see ∂G is second countable, we need to provide a countable basis for the topology. Since Z is the quotient of countably many spaces $\widehat{X}_\sigma$ which are themselves second countable metric spaces, Z is second countable. We enumerate a basis $\{U_n, n \geq 0\}$.

Since X is a simplicial complex with countably many cells, it is a separable space, hence so is the set Λ of points on a geodesic from v_0 to some $\eta \in \partial X$ (here Λ may not equal X since a given geodesic segment may not extend to a ray). Let Λ' be a dense countable subset of Λ . Then the family of open sets $V_{N,\varepsilon}(\eta)$ with $\gamma_\eta(N) \in \Lambda', \varepsilon \in \mathbb{Q}$ forms a countable basis for the topology on ∂X . Enumerate them as $\{V_n, n \geq 0\}$.

A neighborhood of $\xi \in \partial_{Stab}G$ is defined by choosing a constant $\varepsilon \in (0, 1)$, a finite subcomplex of X , namely $D(\xi)$, and for every vertex v of that finite subcomplex, a neighborhood U_v of ξ in $\overline{X_v}$. There are countably many choices of domains since each domain is finite, and each $\overline{X_v}$ has a countable basis since it is a compact metric space. By only allowing open subsets from these countable bases, we receive countably many possible ξ -families $\mathcal{U}$ as ξ ranges over $\partial_{Stab}G$. By further only allowing $\varepsilon \in \mathbb{Q}$, we receive a countable collection of sets $V_{\mathcal{U}, \varepsilon}(\xi)$. Enumerate these as $\{W_n, n \geq 0\}$. It's clear that these form a neighborhood basis for every point $\xi \in \partial_{Stab}G$, and thus the collection $\{U_n, V_n, W_n, n \geq 0\}$ forms a countable basis of neighborhoods for $\overline{Z}$.

Finally, every element of our basis meets Z by construction, so Z is dense in $\overline{Z}$. $\square$

6 Properties of the Topology

6.1 Independence of basepoint and induced topologies

In the previous section we endowed $\overline{Z}$ with a topology. We begin by showing this topology is independent of our chosen basepoint and induces the topology we expect on some of the spaces we already understand. The following is another basic topological fact that we leave to the reader.

Lemma 6.1. Suppose X is a set with two topologies τ_1, τ_2 and for each $x \in X$, let $N_1(x), N_2(x)$ be a neighborhood basis for x . If for each $x \in X$ and $U \in N_1(x)$, there is some $U' \in N_2(x)$ so that $U' \subset U$, then any open set in τ_1 is also open in τ_2 .

Lemma 6.2. The topology of $\overline{Z}$ does not depend on the choice of basepoint.

Proof. Choose two points $x_1, x_2 \in X$, not necessarily vertices. For $z \in \overline{Z}$ and $i = 1, 2$, let $\mathcal{O}_{\overline{Z}}^i(z)$ be the neighborhood basis of z defined above using x_i as the basepoint.

Claim 6.2.1. For any $z \in \overline{Z}$, if $U \in \mathcal{O}_{\overline{Z}}^1(z)$, then there is some $U' \in \mathcal{O}_{\overline{Z}}^2(z)$ so that $U' \subset U$. Similarly, we indicate basic open sets and cones with superscripts to indicate which basepoint they are using.

Proof of Claim. There are 3 cases depending on if $z \in Z, \partial X$, or $\partial_{Stab}G$.

If $z \in Z$, then $\mathcal{O}_{\overline{Z}}^1(z) = \mathcal{O}_{\overline{Z}}^2(z)$ because these neighborhoods do not reference the basepoint, so we can take $U' = U$.

If $z \in \partial X$, then $U = V_W(\eta)$ where $W = V_{r, \delta}^1(\eta)$. From [BH99], the topology on $\overline{X}$ does not depend on the choice of basepoint and the $V_{r, \delta}(\eta)$ form a neighborhood basis with any basepoint. Therefore there must be r', δ' so that $V_{r', \delta'}^2(\eta) \subset V_{r, \delta}^1(\eta)$, where $V_{r', \delta'}^2(\eta)$ uses x_2 as a basepoint. Letting W' be the interior of $V_{r', \delta'}^2(\eta)$, it follows that $V_{W'}^2(\eta) \subset V_W^1(\eta)$, so $U' = V_{W'}^2(\eta)$ is our desired set.

If $\xi \in \partial_{Stab}G$, then $U = V_{\mathcal{U}_1, \varepsilon}^1(\xi)$ for some ξ -family $\mathcal{U}_1$ which is admissible using x_1 as a basepoint. Let $\mathcal{U}_2''$ be any ξ -family which is admissible with x_2 as a basepoint, let $\mathcal{U}_2'$ be a ξ -family $d_{max} + F(d_{max})$ -nested in the family of sets $\{(U_1)_v \cap (U_2'')_v \mid v \in V(\xi)\}$, and let $\mathcal{U}_2$ be d_{max} -refined in $\mathcal{U}_2'$. We claim $V_{\mathcal{U}_2, \varepsilon}^2(\xi) \subset V_{\mathcal{U}_1, \varepsilon}^1(\xi)$.

By definition, $\mathcal{U}_2$ is contained in the ξ -family $\mathcal{U}_1$ which is admissible with x_1 as basepoint, so $\mathcal{U}_2$ is also admissible with x_1 as basepoint. If $z \in W_{\mathcal{U}_2, \varepsilon}(\xi)$, then for every vertex $v \in \sigma_{p(z)} \cap V(\xi)$, we know $c(z) \in (U_2)_v \subset (U_1)_v$, so $z \in W_{\mathcal{U}_1, \varepsilon}(\xi)$.

If $y \in \widetilde{\text{Cone}}_{\mathcal{U}_2, \varepsilon}^2(\xi)$, then $\sigma_{\xi, \varepsilon}^2(y)$ is a witness for $y \in \text{pCone}_{\mathcal{U}_2}(\xi) \subset \text{pCone}_{\mathcal{U}_1}(\xi)$. We can apply Lemma 4.25 with $\mathcal{U} = \mathcal{U}_2$ and $\mathcal{V} = \mathcal{U}_1$ to conclude that $y \in \widetilde{\text{Cone}}_{\mathcal{U}_1, \varepsilon}(\xi)$ for any e . Since y was arbitrary, this implies $\widetilde{\text{Cone}}_{\mathcal{U}_2, \varepsilon}^2(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U}_1, \varepsilon}^1(\xi)$.

Finally, if $\xi' \in V_{\mathcal{U}_2, \varepsilon}^2(\xi)$, then we must check conditions on $v \in D(\xi') \cap D(\xi)$ and $D(\xi') \setminus D(\xi)$. If $v \in D(\xi') \cap D(\xi)$, then $\xi' \in (U_2)_v \subset (U_1)_v$, and by the previous paragraph, we have

$$D(\xi') \setminus D(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U}_2, \varepsilon}^2(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U}_1, \varepsilon}^1(\xi),$$

which checks all the necessary conditions. Hence $U' = V_{\mathcal{U}_2, \varepsilon}^2(\xi)$ satisfies the claim. $\square$

With the claim proven, we can apply the previous lemma to see that any set which is open with x_1 as a basepoint is also open with x_2 as a basepoint. Reversing the roles of x_1 and x_2 , we see that a subset of $\overline{Z}$ is open with respect to x_1 if and only if it is open with respect to x_2 , so the topologies are the same. $\square$

Proposition 6.3. The topology of $\overline{Z}$ induces the natural topology on $Z, \partial X$ and $\widehat{X}_v$ for every vertex v of X .

Proof. We must show that if U is an open subset of $Z, \partial X$, or any $\widehat{X}_v$, then U can be extended to an open set in $\overline{Z}$, and that if $U \in \mathcal{O}_{\overline{Z}}$, then $U \cap B$ is open in the topology on B .

Consider Z first. Open sets of Z are part of $\mathcal{O}_{\overline{Z}}$ already, so no extension is necessary. On the other hand, if $U \in \mathcal{O}_{\overline{Z}}$, then U is a basic open set around some point $z \in \overline{Z}$. If $z \in Z$, then U is already open in $\overline{Z}$, and if $z \in \partial G = \partial X \sqcup \partial_{Stab}G$, it follows from Lemma 5.13 and Lemma 5.14 that $U \cap Z$ is open in Z .

Now consider ∂X . If $U \subset \partial X$ is open, then U can be extended to an open set $U' \subset \overline{X}$. For any $\eta \in U$, $V_{U'}(\eta)$ is open in $\overline{Z}$ and $V_{U'}(\eta) \cap \partial X = U$, so $V_{U'}(\eta)$ extends U to an open set in $\overline{Z}$. On the other hand, if $U \in \mathcal{O}_{\overline{Z}}$, then again there are three cases to consider. If $U \subset Z$ then $U \cap \partial X = \emptyset$ and there is nothing to prove. If $U = V_{U'}(\eta)$ for some $\eta \in \partial X$, then $U \cap \partial X = U' \cap \partial X$, which is open in ∂X . If $U = V_{\mathcal{U}, \varepsilon}(\xi)$ for some $\xi \in \partial_{Stab}G$, then it follows from Lemma 5.11 that $U \cap \partial X$ is open in ∂X .

Now suppose v is a vertex of X and consider an open subset $U \subset \widehat{X}_v$. Choose $\delta > 0$ so that $B(v, \delta) \subset st(v)$, then

$$U' = \{z \in Z, p(z) \in B(v, \delta), c(z) \in U\}$$

is open in Z , and clearly $U' \cap \widehat{X}_v = U \cap Z$. For each $\xi \in U \cap \partial G_v$, we can apply Proposition 4.3 to U, ξ to receive a ξ -family $\mathcal{U}_\xi$, so that $(U_\xi)_v \subset U$. Then

$$\left(U' \cup \bigcup_{\xi \in U \cap \partial G_v} V_{\mathcal{U}_\xi, \frac{1}{2}}(\xi) \right) \cap \widehat{X}_v = U.$$

The large set on the left is a union of open sets in $\overline{Z}$ so it is open, and is the desired extension of U to $\overline{Z}$.

We fix some $U \in \mathcal{O}_{\overline{Z}}$ and we show $U \cap \widehat{X}_v$ is open. There are 3 cases, depending on what kind of basic open set U is.

If U is an open subset of Z , then $U \cap \widehat{X}_v$ is open for every simplex σ of X because that's the definition of the quotient topology on Z . In particular $U \cap \widehat{X}_v$ is open in $\widehat{X}_v$.

Suppose $U = V_{\mathcal{U}, \varepsilon}(\xi)$ and fix some $z \in V_{\mathcal{U}, \varepsilon}(\xi) \cap \widehat{X}_v$. We need to find an open neighborhood U' of z in $\widehat{X}_v$ so that $U' \subset U \cap \widehat{X}_v$, and consider the 3 possibilities for z .

If $z \in Z$, then $z \in U \cap \widehat{X}_v^\circ$ and we need find a neighborhood U' of z in $\widehat{X}_v$ so that $U' \subset U \cap \widehat{X}_v$. If $v \in V(\xi)$, then $c(z) \in U_v$ by definition of $W_{\mathcal{U}, \varepsilon}(\xi)$, so $U' = U_v \cap \widehat{X}_v^\circ$ is the desired neighborhood of z . If $v \notin V(\xi)$ then we must have $v \in \text{Cone}_{\mathcal{U}, \varepsilon}(\xi)$. In this case, $p(\widehat{X}_v^\circ) = v \subset \text{Cone}_{\mathcal{U}, \varepsilon}(\xi)$, so $U' = \widehat{X}_v^\circ$ is the desired neighborhood of z .

Suppose $z = \xi' \in V_{\mathcal{U}, \varepsilon}(\xi) \cap \partial G_v$. We know $v \in V(\xi')$, but v may or may not be in $V(\xi)$. If $v \in V(\xi)$, then set $U' = U_v$, and if $v \notin V(\xi)$, set $U' = \overline{X}_v$. Either way, U' is an open neighborhood of ξ' in $\overline{X}_v$. Using Proposition 4.3, extend U' to a ξ' -family $\mathcal{U}'$. Using the Filtration Lemma 5.8, there is a ξ' -family $\mathcal{V}$ and $\varepsilon' \in (0, 1)$ so that $V_{\mathcal{V}, \varepsilon'}(\xi') \subset V_{\mathcal{U}, \varepsilon}(\xi)$. Let $\mathcal{V}'$ be a d_{max} -refined and nested in $\mathcal{V}$. Recall that the intersection of ξ' -families is again ξ' -family (see Corollary 4.4), so we can combine $\mathcal{U}'$ and $\mathcal{V}'$ into a ξ' -family $\mathcal{W} = \{U'_v \cap V'_v \mid v \in V(\xi')\}$. Then W_v is an open neighborhood of $\xi' \in \widehat{X}_v$, and we claim $W_v \subset U \cap \widehat{X}_v$ so that W_v is the desired neighborhood of ξ' . Indeed, if $x \in W_v$, then either $x \in \widehat{X}_v^\circ$ or $x \in \partial G_v$. In the first case, we know $c(x) \in W_v \subset U_v$, so $x \in W_{\mathcal{W}, \varepsilon'}(\xi') \subset V_{\mathcal{U}, \varepsilon}(\xi)$. In the second case, $x \in W_v \subset V'_v$, and because $\mathcal{V}'$ is d_{max} -refined and nested in $\mathcal{V}$, Lemma 5.7 applied to x tells us $x \in V_{\mathcal{V}, \varepsilon'}(\xi') \subset V_{\mathcal{U}, \varepsilon}(\xi)$. This completes the proof that if $U = V_{\mathcal{U}, \varepsilon}(\xi)$, then $U \cap \widehat{X}_v$ is open.

Suppose $U = V_{U'}(\eta)$ for some $\eta \in \partial X, U' \subset \overline{X}$. Either $v \notin U'$ so $U \cap \widehat{X}_v = \emptyset$ and there is nothing to prove, or $v \in U'$ and

$$U \cap \widehat{X}_v = \widehat{X}_v^\circ \sqcup \{\xi \in \partial_{Stab} G, v \in D(\xi) \subset U'\}.$$

The first set on the right is clearly open in $\widehat{X}_v$. If ξ is in the second set, then by the Filtration Lemma 5.8, there is some $V_{\mathcal{U},\varepsilon}(\xi) \subset U$. Let $\mathcal{U}'$ be a ξ -family d_{max} -refined and nested in $\mathcal{U}$. By the previous case, $W = V_{\mathcal{U}',\varepsilon}(\xi) \cap \widehat{X}_v$ is open in $\widehat{X}_v$, and we claim $W \subset V_{U'}(\eta)$. For $x \in W \cap \widehat{X}_v^\circ$, there is nothing to check since $\widehat{X}_v^\circ \subset V_{U'}(\eta)$. For $\xi' \in W \cap \partial G_v$, we know that $\xi' \in U'_v$, so by the choice of $\mathcal{U}'$ and Lemma 5.7, $\xi' \in V_{\mathcal{U},\varepsilon}(\xi) \subset V_{U'}(\eta)$. This shows the second set on the right is open, hence $U \cap \widehat{X}_v$ is open. $\square$

The reader might wonder why we limit ourselves to vertices in the previous proposition and don't consider the induced topology on $\widehat{X}_\sigma$ for an arbitrary simplex σ of X . The answer is because the result would be false. For example, if σ was a higher dimensional simplex, we could choose an open neighborhood $U = \{\sigma\} \times U_1 \times U_2 \subset \widehat{X}_\sigma$, where $U_1 \subsetneq \sigma$ and $U_2 \subset \overline{X}_\sigma$ are open. If $\xi \in U_2 \cap \partial G_\sigma$, then any neighborhood of ξ in our topology on $\overline{Z}$ would contain some $W_{\mathcal{U},\varepsilon}(\xi)$, which would contain points z with $p(z) \in \sigma \setminus U_1$. Thus it will be impossible to find an open neighborhood of ξ which meets $\widehat{X}_\sigma$ in exactly U .

6.2 $\overline{Z}$ is T_0

By Urysohn's metrization theorem [Mun75, 4.4], to show $\overline{Z}$ is metrizable it is enough to show $\overline{Z}$ is Hausdorff, separable, and regular. We have already seen it is separable in Theorem 5.15, and we will prove it is regular and T_0 , a combination which implies Hausdorff.

Proposition 6.4. The space $\overline{Z}$ satisfies the T_0 separation condition, that is, for every pair of distinct points $z, z' \in \overline{Z}$ there is a neighborhood of z not containing z' .

The proof is broken into many cases. We continue to use v_0 as our notation for the basepoint of the topology.

Lemma 6.5 (Case 1: $z, z' \in \partial X$). If $\eta, \eta' \in \partial X$, then there is an open neighborhood U of η in $\overline{X}$ so that $\eta' \notin V_U(\eta)$.

Proof. The space $\overline{X}$ is metrizable, hence Hausdorff and we have a nice neighborhood basis for points in ∂X . We can choose r, δ, r', δ' so that $V_{r,\delta}(\eta)$ and $V_{r',\delta'}(\eta')$ are disjoint and let U, U' be their respective interiors. Then $V_U(\eta)$ and $V_{U'}(\eta')$ are disjoint. $\square$

Lemma 6.6 (Case 2: $z \in \partial X, z' \in \partial_{Stab}G$). If $\eta \in \partial X$ and $\xi \in \partial_{Stab}G$, then there is an open neighborhood U of η in $\overline{X}$ so that $\xi \notin V_U(\eta)$.

Proof. Pick R so that $D(\xi) \subset B(v_0, R)$ and let U be the interior of $V_{R+1,1}(\eta)$. Then $\xi \notin V_U(\eta)$. $\square$

Lemma 6.7. Suppose $\xi \in \partial_{Stab}G$ and $x \in \overline{X} \setminus D(\xi)$. Then there is some ξ -family $\mathcal{U}$ so that $x \notin \widehat{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ for any ε .

Proof. If $[v_0, x]$ does not meet $D(\xi)$, then $x \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ for any ξ -family $\mathcal{U}$ by the contrapositive of the Genuine Shadow Proposition 4.27.

If $[v_0, x]$ does meet $D(\xi)$, let σ be the first simplex met by $[v_0, x]$ after leaving $D(\xi)$. Choose a vertex v of $\sigma \cap D(\xi)$ and a neighborhood U_v of ξ in $\overline{X_v}$ avoiding $\overline{X_\sigma}$. Extend U_v to a ξ -family $\mathcal{U}$ using Proposition 4.3, and let $\mathcal{U}'$ be a ξ -family $d_{\max} + 1$ -nested in $\mathcal{U}$. Observe that the portion of $[v_0, x]$ in $D^\varepsilon(\xi)$ gives a path of simplices of length at most $d_{\max}$ from $\sigma_{\xi, \varepsilon}(x)$ to σ .

For a contradiction, suppose $\varepsilon \in (0, 1)$ and $x \in V_{\mathcal{U}', \varepsilon}(\xi)$. Then $\sigma_{\xi, \varepsilon}(x) \subset N_{\mathcal{U}'}(\xi)$, and the Crossing Lemma 4.21 applied to the path of simplices from $\sigma_{\xi, \varepsilon}(x)$ to σ implies $\overline{X_\sigma} \subset U_v$ in $\overline{X_v}$. But U_v was chosen to avoid $\overline{X_\sigma}$, so this is impossible. $\square$

Lemma 6.8 (Case 3: $z \in \partial_{\text{Stab}} G, z' \in \partial X$). If $\xi \in \partial_{\text{Stab}} G$ and $\eta \in \partial X$, then there is an open neighborhood $V_{\mathcal{U}, \varepsilon}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$ so that $\eta \notin V_{\mathcal{U}, \varepsilon}(\xi)$.

Proof. Apply the previous lemma with $x = \eta$ to receive a ξ -family $\mathcal{U}$ with $\eta \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ for any ε . Then $\eta \notin V_{\mathcal{U}, \varepsilon}(\xi)$ for any ε . $\square$

Lemma 6.9 (Case 4: $z, z' \in \partial_{\text{Stab}} G$). If $\xi, \xi' \in \partial_{\text{Stab}} G$, then there is an open neighborhood $V_{\mathcal{U}, \varepsilon}(\xi) \in \mathcal{O}_{\overline{Z}}(\xi)$ so that $\xi' \notin V_{\mathcal{U}, \varepsilon}(\xi)$.

Proof. If $D(\xi) \cap D(\xi') \neq \emptyset$, we can choose a vertex $v \in V(\xi) \cap V(\xi')$. Let U_v be a neighborhood of ξ in $\overline{X_v}$ not containing ξ' and extend U to a ξ -family $\mathcal{U}$. Then $\xi' \notin V_{\mathcal{U}, \varepsilon}(\xi)$ for any ε since it fails the first condition for $\xi' \in \partial W_{\mathcal{U}, \varepsilon}(\xi)$ in $\overline{X_v}$.

Otherwise, we can suppose $D(\xi) \cap D(\xi') = \emptyset$. Apply Lemma 6.7 to any $x \in D(\xi')$ to find a ξ -family $\mathcal{U}$ so that $x \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ for any ε . Then ξ' fails the second condition for $\xi' \in V_{\mathcal{U}, \varepsilon}(\xi)$. $\square$

Lemma 6.10 (Case 5: $z \in Z, z' \in \partial G$). If $z \in Z$ and $z' \in \partial G$, there is a neighborhood of z not containing z' .

Proof. This is immediate since by definition sets from $\mathcal{O}_{\overline{Z}}(z)$ do not meet ∂G . $\square$

Lemma 6.11 (Case 6: $z \in \partial G, z' \in Z$). If $z \in \partial G$ and $z' \in Z$, then there is a set $V \in \mathcal{O}_{\overline{Z}}(z)$ so that $z \notin Z$.

Proof. Let $x = p(z')$. If $z = \eta \in \partial X$, then any neighborhood U of η in $\overline{X}$ which doesn't contain x will have $z' \notin V_U(\eta)$. If $z = \xi \in \partial_{\text{Stab}} G$, then $[v_0, x]$ either meets $D(\xi)$ or doesn't. If it doesn't, then $x \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi) \cup D(\xi)$ for any ξ -family $\mathcal{U}$ or $\varepsilon \in (0, 1)$ by the contrapositive of the Genuine Shadow Proposition 4.27. If $[v_0, x]$ does meet $D(\xi)$, then it either goes through or $x \in D(\xi)$. If it goes through, we can apply Lemma 6.7 to x to find a suitable ξ -family. If $x \in D(\xi)$, then $\sigma_x \subset D(\xi)$. Choose a vertex v of σ_x and a neighborhood U_v of ξ in $\overline{X_v}$ so that $c(z) \notin U_v$. Extend U_v to a ξ -family $\mathcal{U}$ using Proposition 4.3. Then $z \notin W_{\mathcal{U}, \varepsilon}(\xi)$ and $z \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$, so $z \notin V_{\mathcal{U}, \varepsilon}(\xi)$. $\square$

6.3 $\overline{Z}$ is regular

Recall that a topological space B is *regular* if for every open set $U \subset B$ and $x \in U$, there exists another open neighborhood U' of x so that $\overline{U'} \subset U$, or equivalently, every point of $B \setminus U$ admits a neighborhood which doesn't meet U' . If B is regular and has the T_0 separation property, it's easy to see B is Hausdorff; given distinct points $x, y \in B$, the T_0 property implies we can find an open set U containing x but not y , and regularity implies we can find another open set $x \in U' \subset U$ so that $\overline{U'} \subset U$. Then U' and $B \setminus \overline{U'}$ are disjoint neighborhoods of x, y .

The goal of this subsection is the following proposition.

Proposition 6.12. The space $\overline{Z}$ is regular.

Since open sets are unions of basic open sets, it is enough to prove the regularity condition for basic open sets, leaving us with three cases to check.

Lemma 6.13 (Case 1: $\eta \in \partial X$). Let $\eta \in \partial X$ and $V_U(\eta) \in \mathcal{O}_{\overline{Z}}(\eta)$. Then there exists an open neighborhood U' of η so that every point of $\overline{Z} \setminus V_U(\eta)$ admits a neighborhood missing $V_{U'}(\eta)$.

Proof. Because there are finitely many orbits of simplices, there is some constant D so that all cells of X have diameter at most D . Recall that A is the acylindricity constant which bounds the diameter of domains.

By definition, U is the interior of $V_{r,\delta}(\eta)$ for some r, δ . Using Lemma 5.3, choose an open neighborhood $W_1 \subset U \subset \overline{X}$ of η so that $d(W_1, X \setminus U) \geq A + D + 1$. Since $\overline{X}$ is metrizable, it is regular and we can choose a neighborhood $W_2 \subset W_1$ so that $\overline{W_2} \subset W_1$. Because the $V_{r,\delta}(\eta)$ form a neighborhood basis for η in $\overline{X}$, we can choose r', δ' so that $V_{r',\delta'}(\eta) \subset W_2$, and let U' be the interior of $V_{r',\delta'}(\eta)$. Let $x = \gamma_\eta(r') \in U'$, and decreasing δ' if necessary, we can assume $B(x, \delta') \subset st(x)$. We show U' satisfies the conclusion.

If $\eta' \in \partial X \setminus V_U(\eta)$ then $\eta' \notin \overline{U'}$, so we can choose an open neighborhood W of η' so that $W \cap U' = \emptyset$, for example $W := \overline{X} \setminus \overline{U'}$. Then $V_W(\eta') \cap V_{U'}(\eta)$ contains no points of ∂X , since those points would have to be in both W and U' . It cannot contain any point $\xi \in \partial_{Stab} G$, since any such ξ would have $D(\xi) \subset W \cap U' = \emptyset$. It cannot contain any point $z \in Z$, since such a z would have $p(z) \in W \cap U' = \emptyset$. Thus $V_W(\eta') \cap V_{U'}(\eta) = \emptyset$ as needed.

If $z \in Z \setminus V_U(\eta)$, then $p(z) \notin U$. Since $\overline{U'} \subset U$, there is an open neighborhood W of $p(z)$ so that $W \cap U' = \emptyset$, for example $W := X \setminus \overline{U'}$. Then $p^{-1}(W)$ is an open neighborhood of z which doesn't meet $V_{U'}(\eta)$.

If $\xi \in \partial_{Stab} G \setminus V_U(\eta)$, then by definition $D(\xi) \not\subset U$, which means $D(\xi) \cap X \setminus U \neq \emptyset$. Because $\text{diam}(N(\xi)) \leq A + D$ and W_1 is at least $A + D + 1$ from $\overline{X} \setminus W_1$, this implies $N(\xi) \subset X \setminus W_1$.

Claim 6.13.1. There is a ξ -family $\mathcal{V}$ so that

$$U' \cap (D(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi)) = \emptyset.$$

Proof of Claim. By Lemma 2.33, $\text{Geod}(x, D(\xi))$ spans a finite subcomplex, say K . For every $v \in V(\xi)$ and simplex σ of $st(v) \cap K \setminus D(\xi)$, let $U_{v,\sigma}$ be an open neighborhood of ξ in $\overline{X_v}$ avoiding $\overline{X_\sigma}$. Then

$$V_v := \bigcap_{\sigma \subset st(v) \cap K \setminus D(\xi)} U_{v,\sigma}$$

is a finite intersection of open sets in $\overline{X_v}$, each containing ξ . If $v \in V(\xi)$ has $st(v) \cap K = \emptyset$, then this intersection above is empty and we set $V_v = \overline{X_v}$. Let $\mathcal{V}'$ be a ξ -family contained in the sets $\{V_v, v \in V(\xi)\}$, and let $\mathcal{V}$ be a ξ -family 1-refined in $\mathcal{V}'$. This is our desired $\mathcal{V}$.

For a contradiction, suppose $y \in U' \cap \widetilde{\text{Cone}}_{\mathcal{V}, \frac{1}{2}}(\xi)$. Since $y \in U' \subset V_{r',\delta'}(\eta)$, there is some $y' \in [v_0, y] \cap B(x, \delta') \setminus B(v_0, r')$. Because $B(x, \delta') \subset st(x)$, we know $\sigma_x \subseteq \sigma_{y'}$, so there is a path of simplices of length at most 1 from $\sigma_{y'}$ to σ_x which doesn't meet $D(\xi)$. Because $y \in \widetilde{\text{Cone}}_{\mathcal{V}, \frac{1}{2}}(\xi)$, $[v_0, y]$ meets $D^{\frac{1}{2}}(\xi)$, so $\sigma_{\xi, \frac{1}{2}}(y)$ exists. Because $y' \in U' \subset W_1$ and $N(\xi) \subset X \setminus W_1$, we know $[v_0, y]$ leaves $D^{\frac{1}{2}}(\xi)$ before reaching y' , hence $\sigma_{\xi, \frac{1}{2}}(y') = \sigma_{\xi, \frac{1}{2}}(y)$. Since $y \in \widetilde{\text{Cone}}_{\mathcal{V}, \frac{1}{2}}(\xi)$ and being in the $\widetilde{\text{Cone}}$ is a statement about exit simplices, we have $y' \in \widetilde{\text{Cone}}_{\mathcal{V}, \frac{1}{2}}(\xi)$. Combining this with the path of simplices from $\sigma_{y'}$ to σ_x and the Refinement Lemma 4.28, we see that $x \in \widetilde{\text{Cone}}_{\mathcal{V}', \frac{1}{2}}(\xi)$, which means that $\overline{X_{\sigma_{\xi, \frac{1}{2}}(x)}} \subset V'_v$ in $\overline{X_v}$ for every vertex $v \in \sigma_{\xi, \frac{1}{2}}(x) \cap V(\xi)$. But $\sigma_{\xi, \frac{1}{2}}(x)$ is a simplex of K and V'_v was chosen specifically to avoid simplices of K , so this is a contradiction. $\square$

With this claim, it's easy to see that $V_{\mathcal{V}, \frac{1}{2}}(\xi) \cap V_{U'}(\eta) = \emptyset$. Every point in $V_{\mathcal{V}, \frac{1}{2}}(\xi)$ has an associated point or set in $\overline{X}$, either $p(z), \eta$, or $D(\xi')$. This associated point or set is contained in $D(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{V}, \frac{1}{2}}(\xi)$, hence it cannot meet U' , so the point cannot be in $V_{U'}(\eta)$. $\square$

Lemma 6.14 (Case 2: $\xi \in \partial_{Stab} G$). Let $\xi \in \partial_{Stab} G$ and $V_{\mathcal{U}^0, \varepsilon}(\xi) \in \mathcal{O}_{\overline{X}}(\xi)$. Then there exists a ξ -family $\mathcal{U}$ and $\varepsilon' \in (0, 1)$ so that $V_{\mathcal{U}, \varepsilon'}(\xi) \subset V_{\mathcal{U}^0, \varepsilon}(\xi)$ and every point not in $V_{\mathcal{U}^0, \varepsilon}(\xi)$ admits a neighborhood avoiding $V_{\mathcal{U}, \varepsilon'}(\xi)$. Further, if $z' = \xi' \in \partial_{Stab} G$, then this neighborhood has the form $V_{\mathcal{V}, \varepsilon}(\xi')$ where $\widetilde{\text{Cone}}_{\mathcal{U}^0, \varepsilon}(\xi) \cap \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi') = \emptyset$.

Proof. For each $v \in V(\xi)$, $\overline{X_v}$ is metrizable, hence regular, so we can choose an open neighborhood V_v^1 of ξ so that $\overline{V_v^1} \subset U_v^0$, and then choose a ξ -family $\mathcal{U}^1$ which is d_{max} -refined and nested in the sets $\{V_v^1, v \in V(\xi)\}$. Repeat this process to construct ξ -families $\mathcal{U}^4 \subset \mathcal{U}^3 \subset \mathcal{U}^2 \subset \mathcal{U}^1 \subset \mathcal{U}^0$, so that for $i = 1, 2, 3, 4$, $\mathcal{U}^i$ is d_{max} -refined and nested in $\mathcal{U}^{i-1}$, and for every vertex $v \in V(\xi)$, we have $\overline{V_v^i} \subset U_v^{i-1}$. We show $V_{\mathcal{U}^4, \varepsilon/2}(\xi)$ is the desired neighborhood of ξ .

Claim 6.14.1. If $\eta \in \partial X \setminus V_{\mathcal{U}^0, \varepsilon}(\xi)$, then there is some open neighborhood $W \subset \overline{X}$ so that $V_W(\eta)$ avoids $V_{\mathcal{U}^3, \varepsilon}(\xi) = \emptyset$.

Proof of Claim. It is enough to choose R, δ so that

$$V_{R, \delta}(\eta) \cap \left(D(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi) \right) = \emptyset.$$

Since $N(\xi)$ is bounded, we can choose R so that $N(\xi) \subset B(v_0, R)$. Then $V_{R,\delta}(\eta) \cap D(\xi) = \emptyset$ for any δ , since any point in $V_{R,\delta}(\eta)$ is at least R from v_0 . Let $x = \gamma_\eta(R)$. There are two cases.

If $[v_0, \eta]$ does not meet $D(\xi)$, then we can choose $0 < \delta < d([v_0, \eta], D(\xi))$. If $y \in V_{R,\delta}(\eta)$, then the CAT(0) inequality implies $[v_0, y]$ stays within δ of $[v_0, \eta]$ for time R , hence $[v_0, y]$ avoids $D(\xi)$. Therefore $y \notin \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi)$ by the contrapositive of the Genuine Shadow Proposition 4.27.

If $[v_0, \eta]$ does meet $D(\xi)$, then because $x \notin \widetilde{\text{Cone}}_{\mathcal{U}^0, \varepsilon}(\xi)$ there is a vertex v of $\sigma_{\xi, \varepsilon}(x) \cap D(\xi)$ so that $\overline{X_{\sigma_{\xi, \varepsilon}(x)}} \not\subset U_v^0$ in $\overline{X_v}$. The Star Lemma 4.29 gives a constant $\delta > 0$ so that for every $y \in B(x, \delta) \setminus D^\varepsilon(\xi)$, $[v_0, y]$ goes through $D^\varepsilon(\xi)$ and $\sigma_{\xi, \varepsilon}(y) \subset st(\sigma_{\xi, \varepsilon}(x))$. Because $x \notin N(\xi)$, we can decrease δ if necessary to assume $B(x, \delta) \cap D^\varepsilon(\xi) = \emptyset$. For a contradiction, suppose $y \in V_{R,\delta}(\eta) \cap \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi)$. The star condition and the definition of the cone imply that

$$\overline{X_{\sigma_{\xi, \varepsilon}(y)}} \subset (\overline{X_{\sigma_{\xi, \varepsilon}(\eta)}} \cap U_v^3)$$

for any vertex $v \in \sigma_{\xi, \varepsilon}(x) \cap V(\xi)$. But $\mathcal{U}^3$ is nested in $\mathcal{U}^0$, so this implies $\overline{X_{\sigma_{\xi, \varepsilon}(\eta)}} \subset U_v^0$, hence $\eta \in V_{\mathcal{U}^0, \varepsilon}(\xi)$. This is a contradiction, so no such y can exist. $\square$

Claim 6.14.2. If $\xi' \in \partial_{\text{Stab}} G \setminus V_{\mathcal{U}^0, \varepsilon}(\xi)$, then there is some ξ' -family $\mathcal{V}$ so that $V_{\mathcal{V}, \varepsilon}(\xi')$ avoids $V_{\mathcal{U}^3, \varepsilon}(\xi)$, and $\widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi) \cap \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi') = \emptyset$.

Proof of Claim. We construct a ξ' -family very carefully. Consider a vertex $v \in V(\xi)$. If $v \notin V(\xi')$, let $V'_v = \overline{X_v}$. Otherwise, if $v \in V(\xi) \cap V(\xi')$, then $\xi' \notin U_v^1$ – if it was, then ξ' would be in $V_{\mathcal{U}^0, \varepsilon}(\xi)$ by the choice of $\mathcal{U}^1$ and Lemma 5.7, but it's not. Further, we chose U_v^2 so that $\overline{U_v^2} \subset U_v^1$, so we can choose a neighborhood $V'_v \subset \overline{X_v}$ of ξ' so that $V'_v \cap U_v^2 = \emptyset$.

By Lemma 2.33, $\text{Geod}(v_0, D(\xi))$ is contained in a finite subcomplex K . For every vertex v of $D(\xi')$ and simplex $\sigma \subset (st(v) \cap K) \setminus D(\xi')$, let $V'_{v, \sigma}$ be a neighborhood of ξ' in $\overline{X_v}$ disjoint from $\overline{X_\sigma}$. Let

$$W_v = V'_v \cap \left(\bigcap_{\sigma \subset (st(v) \cap K) \setminus D(\xi')} V'_{v, \sigma} \right)$$

and let $\mathcal{V}$ be a ξ' -family $d_{\max}$ -nested in a ξ' -family built from the sets $\{W_v, v \in V(\xi')\}$. We show $V_{\mathcal{U}^3, \varepsilon}(\xi) \cap V_{\mathcal{V}, \varepsilon}(\xi') = \emptyset$, and we begin by showing the cones are disjoint.

For a contradiction, suppose $x \in \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi) \cap \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi')$. By the Genuine Shadow Proposition 4.27, $[v_0, x]$ goes through both $D(\xi)$ and $D(\xi')$. Let σ, σ' be the last simplex met by $[v_0, x]$ in $D(\xi), D(\xi')$. There are 3 cases, depending on which of σ, σ' is met first and if they are distinct.

If $\sigma = \sigma'$, then $[v_0, x]$ leaves both domains at the same time. Let $\tau \subset Lk(\xi) \cap Lk(\xi')$ be the next simplex met by $[v_0, x]$ after $\sigma = \sigma'$. Going backwards along $[v_0, x]$ from the point it leaves $D^\varepsilon(\xi)$ to τ gives a path of simplices from $\sigma_{\xi, \varepsilon}(x)$ to τ of length at most $d_{\max}$, and the Crossing Lemma 4.21 combined with the assumption that $x \in \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi)$ implies $\overline{X_\tau} \subset U_v^2$ for any vertex v of $\sigma = \sigma'$.

Applying the same strategy to $\sigma_{\xi',\varepsilon}(x)$, we have $\overline{X_\tau} \subset W_v$. But then $\overline{X_\tau} \subset U_v^2 \cap W_v$, which is a contradiction because $U_v^2 \cap W_v = \emptyset$. Hence $\sigma \neq \sigma'$.

If $[v_0, x]$ meets σ first and then σ' , let $x' \in \sigma' \cap [v_0, x]$. If x, x' are both far from $D(\xi)$ then $\sigma_{\xi,\varepsilon}(x) = \sigma_{\xi,\varepsilon}(x')$, but if one of x, x' is close to $D(\xi)$, then these exit simplices might be different. Either way, $[v_0, x]$ gives a path of simplices in $Lk(\xi)$ from $\sigma_{\xi,\varepsilon}(x)$ to $\sigma_{\xi,\varepsilon}(x')$ of length at most d_{max} . Applying the Crossing Lemma 4.21 to this path with $x \in \widetilde{\text{Cone}_{\mathcal{U}^3,\varepsilon}(\xi)}$ and $\mathcal{U}^3$ d_{max} -nested in $\mathcal{U}^2$, we get $\overline{X_{\sigma_{\xi,\varepsilon}(x')}} \subset U_v^2$ for any vertex v of $\sigma_{\xi,\varepsilon}(x') \cap D(\xi)$. This means $x' \in D(\xi') \cap \widetilde{\text{Cone}_{\mathcal{U}^2,\varepsilon}(\xi)}$, so Lemma 5.7 implies $\xi' \in V_{\mathcal{U}^0,\varepsilon}(\xi)$, a contradiction.

Thus $[v_0, x]$ must meet σ' first and σ second. Then the first simplex after σ' along $[v_0, x]$, say τ , is contained in K . There is a path of simplices along $[v_0, x]$ from $\sigma_{\xi',\varepsilon}(x)$ to τ of length at most d_{max} . Applying the Crossing Lemma 4.21 to this path with $x \in \widetilde{\text{Cone}_{\mathcal{V},\varepsilon}(\xi')}$ and $\mathcal{V}$ d_{max} -nested in the sets $\{W_v, v \in V(\xi')\}$ implies $\overline{X_\tau} \subset W_v$ for any vertex v of $\tau \cap D(\xi')$. But W_v was explicitly chosen to avoid $\overline{X_\tau}$ because $\tau \subset K$ so this is again a contradiction.

This shows no such x can exist and the cones are disjoint. This immediately implies that $V_{\mathcal{U}^3,\varepsilon}(\xi)$ and $V_{\mathcal{V},\varepsilon}(\xi')$ cannot intersect in ∂X and we show they cannot intersect in Z or $\partial_{Stab}G$ either.

Suppose $z \in Z \cap V_{\mathcal{U}^3,\varepsilon}(\xi) \cap V_{\mathcal{V},\varepsilon}(\xi')$. Because the cones are disjoint, we must have $z \in W_{\mathcal{U}^3,\varepsilon}(\xi) \cap W_{\mathcal{V},\varepsilon}(\xi')$, hence $x := p(z) \in D^\varepsilon(\xi) \cap D^\varepsilon(\xi')$. If there is a vertex $v \in \sigma_x \cap V(\xi) \cap V(\xi')$, then in $\overline{X_v}$, we have $c(z) \in V_v \cap U_v^3$, which is impossible because these sets are disjoint. If there is no such vertex, then σ_x is not entirely contained in either domain and we can choose vertices $v \in \sigma_x \cap V(\xi)$ and $v' \in \sigma_x \cap V(\xi')$ so that $[v, v']$ is not contained in either domain. Thus σ_x or one of its faces is contained in K . and this again leads to a contradiction.

Suppose $\xi'' \in \partial_{Stab}G \cap V_{\mathcal{U}^3,\varepsilon}(\xi) \cap V_{\mathcal{V},\varepsilon}(\xi')$. If there is a vertex $v \in V(\xi') \cap V(\xi'') \setminus V(\xi)$, then $v \in \widetilde{\text{Cone}_{\mathcal{U}^3,\varepsilon}(\xi)} \cap D(\xi')$, and Lemma 5.7 implies $\xi' \in V_{\mathcal{U}^1,\varepsilon}(\xi)$, a contradiction. If there is a vertex $v \in V(\xi) \cap V(\xi'') \setminus V(\xi')$, then $v \in \widetilde{\text{Cone}_{\mathcal{V},\varepsilon}(\xi')}$, hence $[v_0, v]$ meets $D(\xi')$ by the Genuine Shadows Proposition 4.27. Hence $\sigma_{\xi',\varepsilon}(v) \subset K$ and we find a contradiction because $\widetilde{\text{Cone}_{\mathcal{V},\varepsilon}(\xi')}$ avoids the simplices of K . Thus $D(\xi'') \subset D(\xi) \cap D(\xi')$. But this is also impossible, for if $v \in V(\xi) \cap V(\xi') \cap V(\xi'')$, then we must have $\xi'' \in U_v^3 \cap V_v$, but this intersection is empty. Thus no such ξ'' can exist. $\square$

Claim 6.14.3. If $z \in Z \setminus V_{\mathcal{U}^0,\varepsilon}(\xi)$, then there is an open neighborhood $W \subset Z$ of z which avoids $V_{\mathcal{U}^2,\varepsilon/2}(\xi)$.

Proof. Let $x = p(z)$ and consider the two cases $d(x, D(\xi)) > \varepsilon/2$ and $d(x, D(\xi)) \leq \varepsilon/2$ separately.

If $d(x, D(\xi)) > \varepsilon/2$, then we can choose δ_1 so that $B(x, \delta_1) \cap D^{\varepsilon/2}(\xi) = \emptyset$. If $[v_0, x]$ does not meet $D(\xi)$, then using the CAT(0) inequality and the convexity of $D(\xi)$, we can choose a δ_2 so that $[v_0, y]$ does not meet $D(\xi)$ for any $y \in B(x, \delta_2)$, hence $B(x, \delta_2) \cap \widetilde{\text{Cone}_{\mathcal{W},\varepsilon'}(\xi)} = \emptyset$ for any ξ -family $\mathcal{W}$ and ε' by the contrapositive of Proposition 4.27. Taking $\delta = \min(\delta_1, \delta_2)$, $p^{-1}(B(x, \delta))$ is an open neighborhood of z avoiding $V_{\mathcal{U}^0,\varepsilon}(\xi)$. If $[v_0, x]$ does meet $D(\xi)$, we can apply the Star Lemma 4.29

to find a δ_2 so that for all $y \in B(x, \delta_2)$, $[v_0, y]$ goes through $D^{\varepsilon/2}(\xi)$ and $\sigma_{\xi, \varepsilon/2}(y) \subset st(\sigma_{\xi, \varepsilon/2}(x))$. Because $x \notin \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon/2}(\xi)$, we can let $\delta = \min(\delta_1, \delta_2)$ again and see that

$$B(x, \delta) \cap (D^{\varepsilon/2}(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon/2}(\xi)) = \emptyset.$$

Thus $p^{-1}(B(x, \delta))$ is the desired neighborhood of z .

If $d(x, D(\xi)) \leq \varepsilon/2$, then we can choose δ so that $B(x, \delta) \subset D^{\varepsilon}(\xi) \cap st(\sigma_x)$. Because $z \notin V_{\mathcal{U}^0, \varepsilon}(\xi)$, $c(z) \notin U_v^0$ for any vertex $v \in \sigma_x \cap V(\xi)$. By the choice of U_v^1 , we can choose a neighborhood U of $c(z)$ in $\overline{X_v}$ so that $U \cap U_v^1 = \emptyset$. Applying Lemma 3.9 to δ, U , we get the desired neighborhood of z . $\square$

By Lemma 5.6, $V_{\mathcal{U}^4, \varepsilon/2}(\xi)$ is contained in each of the neighborhoods of ξ used in the preceding claims. The preceding claims show that any point $z \in \overline{Z} \setminus V_{\mathcal{U}^0, \varepsilon}(\xi)$ has a neighborhood avoiding $V_{\mathcal{U}^4, \varepsilon/2}(\xi)$, proving the lemma. $\square$

Corollary 6.15. Let $\xi, \xi' \in \partial_{\text{stab}} G$ be distinct points and let $\varepsilon \in (0, 1)$. There exists a ξ -family $\mathcal{U}$ and ξ' -family $\mathcal{U}'$ so that $V_{\mathcal{U}, \varepsilon}(\xi) \cap V_{\mathcal{U}', \varepsilon}(\xi') = \emptyset$ and $\widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi) \cap \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi') = \emptyset$ for any $\varepsilon \in (0, 1)$.

Proof. Beginning with any neighborhood of ξ , we can apply the previous lemma to receive suitable neighborhoods of ξ, ξ' so that the cones are disjoint. $\square$

Lemma 6.16 (Case 3: $z \in Z$). If $z \in Z$ and $U \in \mathcal{O}_{\overline{Z}}(z)$, then there is some neighborhood of $W \subset U$ of z so that if $z' \in \overline{Z} \setminus U$, z' admits a neighborhood avoiding W .

Proof. Let $x = p(z)$. Because U is open in Z , U meets $\widehat{X_{\sigma_x}^\circ}$ in an open set, and because basic open sets are a basis for the product topology on $\widehat{X_{\sigma_x}^\circ}$, there are open sets $W_1 \subset \sigma_x$ and $W_2 \subset X_{\sigma_x}$ so that

$$\{\sigma_x\} \times W_1 \times W_2 \subset U \cap \widehat{X_{\sigma_x}^\circ}$$

and $x \in W_1 \subset \sigma_x$ and $c(z) \in W_2$. Choose $\delta > 0$ so that $B(x, \delta) \subset st(\sigma_x)$ and $\overline{B(x, \delta)} \cap \sigma_x$ is contained in the interior of σ_x . Decreasing δ further if necessary, we can assume $\overline{B(x, \delta)} \cap \sigma_x \subset W_1$. Because X_{σ_x} is metrizable and hence regular, we can choose a W'_2 so that $\overline{W'_2} \subset W_2$. Using Lemma 3.9 with δ and W'_2 , set

$$W = U \cap W_z(B(x, \delta), W'_2).$$

Note that for any $z' \in W$, $p(z') \in B(x, \delta)$ and $c(z')$ can be interpreted as a point of W'_2 . To show this W satisfies the lemma, fix some $z' \in \overline{Z} \setminus U$ and consider three cases.

If $z' \in Z$, set $x' = p(z')$. If $d(x, x') > \delta$, then there is a δ' so that $B(x', \delta') \cap B(x, \delta) = \emptyset$. Then $p^{-1}(B(x', \delta'))$ is an open neighborhood of z' which doesn't meet W . If $d(x, x') \leq \delta$, then because $\overline{B(x, \delta)} \cap \sigma_x \subset W_1$ we have $x' \in p(U)$. If $c(z')$ was in W_2 , then we would have $z' \in U$, which is

not true, so $c(z') \notin W_2$. Since $\overline{W'_2} \subset W_2$, we can choose a neighborhood $W_{z'} \subset X_{\sigma_x}$ of $c(z')$ which doesn't meet W'_2 . We can also choose δ' so that $B(x', \delta') \subset st(\sigma_x)$. Applying Lemma 3.9 to $\delta', W_{z'}$, we get a neighborhood of z' disjoint from W .

If $z' = \eta \in \partial X$, then we can choose any R, δ so that $B(x, \delta) \cap V_{R, \delta}(\eta) = \emptyset$. If U' is the interior of $V_{R, \delta}(\eta)$, then clearly $V_U(\eta) \cap W = \emptyset$.

If $z' = \xi \in \partial_{Stab} G$, then consider the cases $x \notin D(\xi)$ and $x \in D(\xi)$ separately. If $x \notin D(\xi)$, then apply Lemma 6.7 to x and receive a ξ -family $\mathcal{U}$ so that $x \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ for any ε . Choose ε so that $D^\varepsilon(\xi) \cap B(x, \delta) = \emptyset$ and let $\mathcal{U}'$ be a ξ -family 1-refined in $\mathcal{U}$. For a contradiction, suppose there is a point $x' \in B(x, \delta) \cap \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi)$. Then $\sigma_x \subseteq \sigma_{x'}$, so there is a path of simplices of length 1 from $\sigma_{x'}$ to σ_x . The Refinement Lemma 4.28 then implies $x \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$, contradicting the choice of $\mathcal{U}$. Thus

$$B(x, \delta) \cap (D^\varepsilon(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi)) = \emptyset.$$

It follows that $W \cap V_{\mathcal{U}, \varepsilon}(\xi) = \emptyset$.

If $x \in D(\xi)$, then because $\overline{X_{\sigma_x}}$ is metrizable, we can choose a neighborhood $U_{\sigma_x} \subset \overline{X_{\sigma_x}}$ of ξ avoiding W_2 . Using Proposition 4.3, we can extend U_{σ_x} to a ξ -family $\mathcal{U}$. For every $z' \in W$, because the cusped space coordinate $c(z') \in W'_2$ so $c(z') \notin U_{\sigma_x}$, z' cannot be in either $W_{\mathcal{U}, \delta}(\xi)$ or $\widetilde{\text{Cone}}_{\mathcal{U}, \delta}(\xi)$, thus $W \cap V_{\mathcal{U}, \delta}(\xi) = \emptyset$. $\square$

Corollary 6.17. $\overline{Z}$ is metrizable.

Proof. We have shown $\overline{Z}$ has the T_0 separation property and is regular, so $\overline{Z}$ is Hausdorff. Since $\overline{Z}$ is also second countable by Theorem 5.15, Urysohn's metrization theorem [Mun75, 4.4] implies $\overline{Z}$ is metrizable. $\square$

6.4 $\overline{Z}$ is compact

In this section we prove the following.

Theorem 6.18 (Compactness). $\overline{Z}$ is compact.

Lemma 6.19. Let $(x_n)_n$ be a sequence of points in X and let σ_k^n be the k^{th} simplex met by the geodesic $[v_0, x_n]$. Let $L(k) = \{\sigma_k^n, n \geq 0\}$.

1. If some subsequence of $(x_n)_n$ is bounded and all $L(k)$ are finite, then a further subsequence converges to a point $x \in X$.
2. If $x_n \rightarrow \infty$ and all $L(k)$ are finite, then some subsequence converges to a point $\eta \in \partial X$.
3. If some $L(k)$ is infinite, let k be minimal so that $L(k)$ is infinite. Then there is some $\sigma \in L(k-1)$, $\xi \in \partial G_\sigma$, and subsequence of $(x_n)_n$ so that for any ξ -family $\mathcal{U}$ and $\varepsilon \in (0, 1)$, we have $x_n \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$ for n large enough.

Proof. Suppose all $L(k)$ are finite. Then we can pass to a subsequence so that σ_k^n is eventually constant. For $r \in \mathbb{N}$, let π_r be the projection of X onto $\overline{B(v_0, r)}$. For $r = 1$, our choice of subsequence implies $\pi_r(x_n)$ is eventually a sequence in a single closed simplex, which is compact, so we can pass to a subsequence so that $\pi_r(x_n)$ converges. Iterating this for $r = 2, 3, \dots$ and taking a diagonal sequence, we get a subsequence so that $\pi_r(x_n)$ converges to a point in $\overline{B(v_0, r)}$ for each $r \in \mathbb{N}$. Because $\overline{X}$ can be topologized as a projective limit of these balls, this implies the $(x_n)_n$ converges to a point of $\overline{X}$. If this sequence is bounded, that point must be in X , and if the sequence is unbounded that point must be in ∂X . This proves 1 and 2.

Now suppose $L(k)$ is infinite for some minimal k . For any $k' > k$, $L(k')$ must also be infinite, for if $L(k')$ was finite, then $\text{Geod}(v_0, L(k'))$ would meet finitely many simplices by Lemma 2.33, but $L(k) \subset \text{Geod}(v_0, L(k'))$, which is a contradiction. After a subsequence, we can assume that

1. the geodesics $[v_0, x_n]$ all cross the same first $k - 1$ simplices, as in there are open simplices $\sigma_1, \dots, \sigma_{k-1}$ so that $\sigma_1^n = \sigma_1, \dots, \sigma_{k-1}^n = \sigma_{k-1}$ for all $n \geq 0$,
2. σ_k^n contains σ_{k-1} for all n , since σ_{k-1} has finitely many faces,
3. the sequence $(\sigma_k^n)_n$ is injective,
4. no σ_k^n is contained in $D(\xi)$, since $D(\xi)$ contains finitely many simplices,
5. all σ_k^n lie above the same simplex of Y , since there are finitely many G orbits of simplices in X .

The properties 3, 4 allow us to apply the Convergence Property 4.7 and receive a $\xi \in \partial G_{\sigma_{k-1}}$ so that $\overline{X_{\sigma_k^n}}$ converge to ξ uniformly in $\overline{X_{\sigma_{k-1}}}$. This is our candidate ξ and we fix a vertex $v \in \sigma_{k-1} \subset D(\xi)$.

Given some $V_{\mathcal{U}, \varepsilon}(\xi) \in \mathcal{O}_{\overline{X}}(\xi)$ and let $\mathcal{U}'$ be a ξ -family $d_{\max}$ -nested in $\mathcal{U}$. Each $[v_0, x_n]$ meets $D(\xi)$, passes through σ_{k-1} , and then meets $\sigma_k^n \not\subset D(\xi)$. Because domains are convex, this means $x_n \notin D(\xi)$ for each n and σ_k^n is the first simplex of $Lk(\xi)$ met by $[v_0, x_n]$. Following along $[v_0, x_n]$, we get a path of simplices of length at most $d_{\max}$ from σ_k^n to $\sigma_{\xi, \varepsilon}(x_n)$. For large enough n , we have $\overline{X_{\sigma_k^n}} \subset U'_v$ in $\overline{X_v}$, and for such n the Crossing Lemma 4.21 applied to this path of simplices implies $\overline{X_{\sigma_{\xi, \varepsilon}(x_n)}} \subset U_{v'}$ for any vertex $v' \in \sigma_{\xi, \varepsilon}(x_n) \cap V(\xi)$. Hence $x_n \in \widehat{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi)$, and we're done. $\square$

Corollary 6.20. Let v be a vertex of X and let $(g_n)_n$ be a sequence in G such that $d(v_0, g_nv) \rightarrow \infty$. Let σ_k^n be the k^{th} simplex met by $[v_0, g_nv]$ and let $L(k) = \{\sigma_k^n, n \geq 0\}$. After a subsequence, exactly one of the following holds.

1. If all $L(k)$ are finite, then there is some $\eta \in \partial X$ so that $g_n \widehat{X_v} \rightarrow \eta$ uniformly.
2. If k is minimal so that $L(k)$ is infinite, then there is some $\sigma \in L(k-1)$ and $\xi_0 \in \partial G_\sigma$ so that $g_n \widehat{X_v} \rightarrow \xi_0$ uniformly.

Proof. If all $L(k)$ are finite, then Lemma 6.19 implies there is some subsequence and some $\eta \in \partial X$ so that $g_nv \rightarrow \eta$. Given $V_U(\eta) \in \mathcal{O}_{\overline{Z}}(\eta)$, then U is an open neighborhood of η in $\overline{X}$ and we can find a neighborhood $U' \subset U$ of η so that $d(U', X \setminus U) > A + 1$ by Lemma 5.3, and then choose R, δ so that $V_{R,\delta}(\eta) \subset U'$. Let W be the interior of $V_{R,\delta}(\eta)$. Since $g_nv \rightarrow \eta$, $g_nv \in W$ for n large enough. For such n , any $z \in \widehat{X}_v$ has $p(z) = v \in W$, hence $z \in V_W(\eta)$. Further, if $g_nv \in W$, then because domains have diameter at most A ,

$$g_nv \in D(g_n\xi) \subset B(g_nv, A + 1) \subset U$$

for any $\xi \in \partial G_v$, so $g_n\xi \in V_U(\eta)$. This shows that for any $V_U(\eta)$, $g_n\widehat{X}_v \subset V_U(\eta)$ for n large enough, proving case 1.

If k is minimal so that $L(k)$ is infinite, then Lemma 6.19 implies that after a subsequence, there is some $\sigma \in L(k - 1)$ and $\xi_0 \in \partial G_\sigma$ so that for any ξ_0 -family $\mathcal{U}$ and $\varepsilon \in (0, 1)$, $g_nv \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ for n large enough. Because $d(D(\xi_0), g_nv) \rightarrow \infty$ and domains have bounded diameter, we can choose N so that $d(D(\xi_0), g_nv) \geq A + 1$ for all $n \geq N$. This implies that if $\xi \in \partial G_v$ and $n \geq N$, we have $D(g_n\xi) \cap D(\xi_0) = \emptyset$.

If $V_{\mathcal{U},\varepsilon}(\xi_0) \in \mathcal{O}_{\overline{Z}}(\xi_0)$, let $\mathcal{U}'$ be a ξ_0 -family $d_{\max}$ -refined in $\mathcal{U}$. Using the conclusion of Lemma 6.19 and increasing N if necessary, we can assume $g_nv \in \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon}(\xi_0)$ for all $n \geq N$, which immediately implies $g_nz \in V_{\mathcal{U},\varepsilon}(\xi_0)$ for all $z \in \widehat{X}_v$. If $\xi \in \partial G_v$ and $n \geq N$, then because $D(g_n\xi) \cap D(\xi_0) = \emptyset$, every point of $D(g_n\xi)$ can be connected to g_nv by a path of simplices avoiding $D(\xi_0)$. This path of simplices has length at most $d_{\max}$ because $D(\xi)$ has at most $d_{\max}$ simplices. Because $g_nv \in \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon}(\xi_0)$, the Refinement Lemma 4.28 implies $D(g_n\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi_0)$, hence $g_n\xi \in V_{\mathcal{U},\varepsilon}(\xi_0)$ as needed. $\square$

Lemma 6.21. If $(z_n)_n$ is a sequence in Z , then $(z_n)_n$ has a subsequence converging to a point in $\overline{Z}$.

Proof. Let $x_n = p(z_n)$ and apply Lemma 6.19.

If there is a bounded subsequence of x_n and all $L(k)$ are finite, then some subsequence converges to a point $x \in X$. For n large enough so that $x_n \in st(\sigma_x)$, so we can interpret $c(z_n)$ as points in $X_{\sigma_x}^\circ$. Because $\overline{X_{\sigma_x}^\circ}$ is compact, we can pass to a subsequence so that $c(z_n)$ converges to some point $y \in \overline{X_{\sigma_x}^\circ}$. If $y \in X_{\sigma_x}^\circ$, then it's clear that $z_n \rightarrow (\sigma_x, x, y) \in \widehat{X_{\sigma_x}^\circ}$, and if $y \in \partial G_{\sigma_x}$, it's clear that $z_n \rightarrow y$.

If there is an unbounded subsequence of $(x_n)_n$ with all $L(k)$ finite, then $x_n \rightarrow \eta \in \partial X$ and it is clear that $z_n \rightarrow \eta$ in $\overline{Z}$ as well.

Similarly, if some minimal k has $L(k)$ infinite, then we pass to a subsequence and receive some $\xi \in \partial_{\text{Stab}} G$ so that for any ξ -family $\mathcal{U}$ and $\varepsilon \in (0, 1)$, we have $x_n \in \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi)$ for n large enough. It follows that $z_n \rightarrow \xi$. $\square$

Proof of Theorem 6.18. Because $\overline{Z}$ is metrizable, it is enough to show $\overline{Z}$ is sequentially compact. By Theorem 5.15, Z is dense in $\overline{Z}$ and by Lemma 6.21, every sequence in Z has a convergent subsequence. It follows that $\overline{Z}$ is sequentially compact. $\square$

7 Dynamics of the Action

7.1 G is a Convergence Group

In this section, we show G acts as a convergence group on $\overline{Z}$ with limit set ∂G . Given an arbitrary sequence $(g_n)_n$, we must produce a ξ_+, ξ_- so that (g_n, ξ_+, ξ_-) is an ART (see Definition 2.1). In the course of these proofs we will translate $(g_n)_n$ on the left and right or use the sequence $(g_n^{-1})_n$ instead. The following lemma shows these changes can be undone at the end of the proof to make the original sequence a convergence sequence.

Lemma 7.1. Let H be a group acting on a compact metrizable space M , let $(g_n)_n$ be an infinite sequence in H , and let $\xi_+, \xi_- \in M$. The following are equivalent.

1. (g_n, ξ_+, ξ_-) is an ART.
2. (g_n^{-1}, ξ_-, ξ_+) is an ART.
3. $(hg_nk, h\xi_+, k^{-1}\xi_-)$ is an ART, for any $h, k \in H$.

Proof. Notice that (2) $\implies$ (1) follows from (1) $\implies$ (2) using (g_n^{-1}, ξ_-, ξ_+) . Similarly (3) $\implies$ (1) follows from (1) $\implies$ (3), so it suffices to show only two implications starting with (1). Suppose (g_n, ξ_+, ξ_-) is an ART.

Let U_+, U_- be arbitrary neighborhoods of ξ_+, ξ_- . For n large enough, we have $g_n(M \setminus U_-) \subset U_+$. For any such n , g_n can only use U_- to cover all of $M \setminus U_+$, as in $M \setminus U_+ \subset g_n U_-$. Hence $g_n^{-1}(M \setminus U_+) \subset U_-$. Thus the g_n^{-1} take complements of open neighborhoods of ξ_+ uniformly to ξ_- , and (g_n^{-1}, ξ_-, ξ_+) is an ART. This proves (1) $\implies$ (2).

Let $h, k \in H$ and let U, V be arbitrary neighborhoods of $h\xi_+, k^{-1}\xi_-$. Then $h^{-1}U, kV$ are neighborhoods of ξ_+, ξ_- , and for large enough n we have $g_n(M \setminus kV) \subset h^{-1}U$, hence $hg_nk(M \setminus V) \subset U$. This shows $(hg_nk, h\xi_+, k^{-1}\xi_-)$ is an ART and proves (1) $\implies$ (3). $\square$

Lemma 7.2. Let $(g_n)_n$ be an injective sequence in G . Suppose there are vertices v_0, v_1 of X so that $g_nv_0 = v_1$ for all n . Then $(g_n)_n$ is a convergence sequence.

Proof. By the previous lemma, it is enough to show $(g_1^{-1}g_n)_n$ is a convergence sequence, so we can assume $g_nv_0 = v_0$ for all n , hence $(g_n)_n$ is a sequence in G_{v_0} . Since G_{v_0} acts on $\overline{X_{v_0}}$ as a convergence group with limit set ∂G_{v_0} , we can pass to a subsequence and choose $\xi_+, \xi_- \in \partial G_{v_0}$ so that (g_n, ξ_+, ξ_-) is an ART for G_{v_0} acting on $\overline{X_{v_0}}$. This means for all compact subsets $K \subset$

$\overline{X_{v_0}} \setminus \{\xi_-\}$, $g_n K$ converges to ξ_+ uniformly. Our goal is to upgrade this to compact subsets $K \subset \overline{Z} \setminus \{\xi_-\}$ so that $(g_n)_n$ is a convergence sequence for G acting on $\overline{Z}$. We use v_0 as the base point for the topology.

Claim 7.2.1. After translating $(g_n)_n$, passing to a subsequence, and relabeling ξ_+, ξ_- , there is a finite subcomplex $F \subset D(\xi_-) \cap D(\xi_+)$ so that

- F1. all g_n fix F pointwise,
- F2. for all vertices v of F , (g_n, ξ_+, ξ_-) is an ART for G_v acting on $\overline{X_v}$, and
- F3. for all $\sigma \subset Lk(F)$, we have $g_n \overline{X_\sigma} \rightarrow \xi_+$ uniformly in $\overline{X_v}$ for any vertex v of $F \cap \sigma$.

Proof of Claim. First, assume that for every simplex $\sigma \subset st(v_0)$, the set $\{g_n \sigma, n \geq 0\}$ is infinite. We claim that after a subsequence, $F = v_0$ works. Indeed, F1 and F2 are immediate, and if $v_0 \in \sigma$ and $\xi_- \notin \overline{X_\sigma}$, then $\overline{X_\sigma}$ is a compact subset of $\overline{X_{v_0}} \setminus \{\xi_-\}$, so $g_n \overline{X_\sigma} \rightarrow \xi_+$ uniformly already. This proves F3 for all but the simplices $\sigma \subset Lk(v_0) \cap D(\xi_-)$. Enumerate the simplices of $D(\xi_-) \cap Lk(v_0)$ as $\sigma_1, \sigma_2, \dots, \sigma_m$, with $m \leq d_{max}$. By assumption, $\{g_n \sigma_1, n \geq 0\}$ is infinite, so using the Convergence Property 2.24 we can pass to a subsequence so that $g_n \overline{X_{\sigma_1}}$ converges uniformly to some $\xi_0 \in \partial G_{v_0}$. Picking a point $z \in X_{\sigma_1}^\circ \subset \overline{X_{v_0}}$, we have $z \neq \xi_-$, so $g_n z \rightarrow \xi_+$ because (g_n, ξ_+, ξ_-) is an ART, but also $g_n z \rightarrow \xi_0$, so $\xi_+ = \xi_0$. Repeating this subsequence maneuver for $\sigma_2, \sigma_3, \dots, \sigma_m$, we get that $g_n \overline{X_{\sigma_i}} \rightarrow \xi_+$ uniformly for each i . Since $g_n v_0 = v_0$ already, this shows $F = v_0$ would work.

If the assumption doesn't hold, there is some simplex $\sigma \subset Lk(v_0)$ so that $\{g_n \sigma, n \geq 0\}$ is finite. After a subsequence we can assume $g_n \sigma$ is constant. Then $g_1^{-1} g_n \sigma = \sigma$ for all n and we can replace the ART (g_n, ξ_+, ξ_-) for G_{v_0} acting on $\overline{X_{v_0}}$ with $(g_1^{-1} g_n, g_1^{-1} \xi_+, \xi_-)$. After this replacement, all the g_n fix σ , so they fix $\overline{X_\sigma}$ as a closed subset of $\overline{X_{v_0}}$. Fix some $z \in X_\sigma^\circ \subset \overline{X_v}$ and notice that $z \neq \xi_-, \xi_+$ because $\xi_-, \xi_+ \in \partial G_{v_0}$. For all n , $g_n z \in \overline{X_\sigma}$ and $g_n z \rightarrow \xi_+$, so $\xi_+ \in \partial G_\sigma$. Using g_n^{-1} in the same way, we have $\xi_- \in \partial G_\sigma$, hence $\sigma \subset D(\xi_-) \cap D(\xi_+)$.

Let v be another vertex of σ . Viewing $(g_n)_n$ as a sequence in $G_\sigma \subset G_v$ and applying the definition of a convergence group, we can pass to a subsequence and find points $\xi'_+, \xi'_- \in \partial G_v$ so that (g_n, ξ'_+, ξ'_-) is an ART for G_v acting on $\overline{X_v}$. Using the same $z \in \overline{X_\sigma}$ as before, we observe that $(g_n z)_n$ is a sequence in $X_\sigma^\circ \subset \overline{X_v}$ which converges to both ξ'_+ and ξ_+ , hence $\xi'_+ = \xi_+$. Using $g_n^{-1} z$ in the same way, we have $\xi'_- = \xi_-$. Repeating this for all the vertices of σ makes $(g_n)_n$ a convergence sequence for each vertex with attractive and repelling points ξ_+ and ξ_- . This extends the properties of v_0 to σ and all its vertices.

At this point, either $\{g_n \sigma', n \geq 0\}$ is infinite for all simplices $\sigma' \subset Lk(\sigma)$ or there is a simplex $\sigma' \subset Lk(\sigma)$ so that $\{g_n \sigma', n \geq 0\}$ is finite. In the first case, we can argue exactly as in the first paragraph, replacing v_0 with σ and see that $F = \sigma$ satisfies the claim. In the second case, we can argue as above with σ' in place of σ ; After possibly translating the sequence by g_1^{-1} and taking subsequences, each g_n fixes $\sigma \cup \sigma' \subset D(\xi_-) \cap D(\xi_+)$ and for each vertex v of $\sigma \cup \sigma'$, $(g_n)_n$ is a

convergence sequence in G_v with attractive and repelling points $\xi_+, \xi_- \in \partial G_v$. Notice that because g_1 fixes σ , translating the entire sequence by g_1^{-1} will give a new sequence which also fixes σ , so after translating and relabeling, the new sequence does indeed fix $\sigma \cup \sigma'$.

We iterate the previous paragraph with $\sigma \cup \sigma'$ instead of σ . Either $F = \sigma \cup \sigma'$ satisfies the claim, or we repeat the argument and add a third simplex. Because our growing subcomplex is contained in $D(\xi_-) \cap D(\xi_+)$, which has at most d_{max} simplices, we must repeat this argument at most d_{max} times before getting a suitable F . This proves the claim. $\square$

To show (g_n) is a convergence sequence for G acting on $\overline{Z}$, it is enough to fix arbitrary neighborhoods $V_{\mathcal{U},\varepsilon}(\xi_-), V_{\mathcal{V},\varepsilon}(\xi_+)$ and find some N so that $n \geq N$ implies

$$g_n(\overline{Z} \setminus V_{\mathcal{U},\varepsilon}(\xi_-)) \subset V_{\mathcal{V},\varepsilon}(\xi_+).$$

Let $\mathcal{U}^3 \subset \mathcal{U}^2 \subset \mathcal{U}^1 \subset \mathcal{U}$ be a sequence of ξ_- -families each d_{max} -refined and nested in the next. Let $\mathcal{V}'$ be a ξ_+ -family d_{max} -refined and nested in $\mathcal{V}$. If $v \in \sigma \subset D(\xi_+)$, we make V'_v smaller if necessary so that $\overline{X_\sigma} \not\subset V'_v$. We will use the contrapositive of this later; if $\sigma \subset N(\xi)$ is a simplex and $v \in \sigma \cap V(\xi_+)$ so that $\overline{X_\sigma} \subset V'_v$, then σ is *not* a simplex of $D(\xi_+)$.

Claim 7.2.2. With $(g_n)_n$ and F satisfying the conclusion of Claim 7.2.1, there is some N so that $n \geq N$ implies the following.

1. For all vertices v of F , $g_n(\overline{X_v} \setminus U_v^3) \subset V'_v$.
2. $g_n(\overline{X} \setminus (\widetilde{\text{Cone}}_{\mathcal{U}^2,\varepsilon}(\xi_-) \cup D(\xi_-))) \subset \widetilde{\text{Cone}}_{\mathcal{V},\varepsilon}(\xi_+)$.
3. $g_n(D(\xi_-) \setminus F) \subset \widetilde{\text{Cone}}_{\mathcal{V},\varepsilon}(\xi_+)$.

Proof of Claim. Applying F2 of Claim 7.2.1 to each of the finitely many vertices of F , we find N large enough to satisfy (1). There are also finitely many simplices $\sigma \subset Lk(F) \cap D(\xi_-)$, and applying F3 to each of these and possibly increasing N , we have $n \geq N$ implies $g_n \overline{X_\sigma} \subset V'_v$ for each vertex v of $\sigma \cap F$. This is our N .

To prove (2), consider a point $x \in \overline{X} \setminus (\widetilde{\text{Cone}}_{\mathcal{U}^2,\varepsilon}(\xi_-) \cup D(\xi_-))$. Then $x \notin F$ because $F \subset D(\xi_-)$, so let σ be the first simplex met by $[v_0, x]$ after leaving F (recall v_0 is the basepoint of our topology). Because the g_n fix F pointwise, the geodesics $g_n[v_0, x] = [v_0, g_n x]$ share the initial segment in F and only differ starting in $g_n \sigma$. For each n , the portion of $[v_0, g_n x]$ outside of F gives a path of simplices of length at most d_{max} through $N(\xi_+)$ from $g_n \sigma$ to $\sigma_{\xi_+,\varepsilon}(g_n x)$.

If $\sigma \subset Lk(F) \cap D(\xi_-)$, then we have chosen N large enough so that $n \geq N$ implies $g_n \overline{X_\sigma} \subset V'_v$, and by the choice of $\mathcal{V}'$ just before the claim, this implies $g_n \sigma$ is not a simplex of $D(\xi_+)$, hence $g_n \sigma$ is also the first simplex met by $[v_0, g_n x]$ outside of $D(\xi_+)$. The Crossing Lemma 4.21 applied to the path of simplices from $g_n \sigma$ to $\sigma_{\xi_+,\varepsilon}(g_n x)$ now implies $g_n x \in \widetilde{\text{Cone}}_{\mathcal{V},\varepsilon}(\xi_+)$.

If $\sigma \not\subset D(\xi_-)$, then $[v_0, x]$ gives a path of simplices of length at most d_{max} from σ to $\sigma_{\xi_-, \varepsilon}(x)$. Let v be a vertex of $\sigma \cap F$. If $\overline{X_\sigma} \cap U_v^3 \neq \emptyset$, then the Crossing Lemma 4.21 applied to this path of simplices would imply $\overline{X_{\sigma_{\xi_-, \varepsilon}(x)}} \subset U_v^2$ for any vertex v of $\sigma_{\xi_-, \varepsilon}(x) \cap D(\xi_-)$. But that is the definition of $x \in \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$, which is a contradiction. Thus $\overline{X_\sigma} \subset (\overline{X_v} \setminus U_v^3)$, and for $n \geq N$, we have

$$g_n \overline{X_\sigma} \subset g_n(\overline{X_v} \setminus U_v^3) \subset V'_v.$$

Exactly as in the previous paragraph, this implies $g_n x \in \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi_+)$. This proves (2).

For (3), let $\sigma' \subset D(\xi_-) \setminus F$. Then a geodesic from any point of σ' to v_0 gives a path of simplices of length at most d_{max} from σ' to some $\sigma \subset Lk(F) \cap D(\xi_-)$, since $v_0 \in D(\xi_-)$ and domains are convex. For $n \geq N$, we have $g_n \overline{X_\sigma} \subset V'_v$ for any vertex v of $\sigma \cap F$. The Refinement Lemma 4.28 then implies $g_n \sigma' \subset \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi_+)$. This proves (3). $\square$

Claim 7.2.3. With N from the previous claim, $n \geq N$ implies

$$g_n(\overline{Z} \setminus V_{\mathcal{U}, \varepsilon}(\xi_-)) \subset V_{\mathcal{V}, \varepsilon}(\xi_+).$$

Proof of Claim. Fix $z \in \overline{Z} \setminus V_{\mathcal{U}, \varepsilon}(\xi_-)$ and consider three cases.

If $z = \eta \in \partial X$, then by definition $\eta \notin \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi_-)$. In particular $\eta \notin \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$, so the claim implies that for $n \geq N$, we have $g_n \eta \in \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi_+)$, hence $g_n \eta \in V_{\mathcal{V}, \varepsilon}(\xi_+)$.

If $z = \xi \in \partial_{stab} G$, then $D(\xi) \cap \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-) = \emptyset$ and for any vertex v of $D(\xi) \cap D(\xi_-)$, we have $\xi \notin U_v^1$ in ∂G_v , since either of these conditions would imply $\xi \in V_{\mathcal{U}, \varepsilon}(\xi_-)$ by Lemma 5.7. To show $g_n \xi \in V_{\mathcal{V}, \varepsilon}(\xi_+)$, we have to show various conditions on how $D(\xi)$ interacts with $D(\xi_+)$. To check those conditions, we break $D(\xi)$ into three parts:

$$D(\xi) = (D(\xi) \cap F) \sqcup (D(\xi) \cap D(\xi_-) \setminus F) \sqcup (D(\xi) \setminus D(\xi_-)).$$

If x is a point of either the second or third parts, then either (3) or (2) of Claim 7.2.2 implies $g_n x \in \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi_+)$. Note that either because Cones are disjoint from domains by definition, or from following the proof Claim 7.2.2, we know $g_n x \notin D(\xi_+)$. This means that $n \geq N$ implies $D(g_n \xi) \setminus F \subset \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi)$, and $D(g_n \xi) \cap D(\xi_+) \subset F$, so only in F do we need to check the condition on vertices of $D(g_n \xi) \cap D(\xi)$. If $v = g_n v$ is a vertex of $D(g_n \xi) \cap F$, (1) of Claim 7.2.2 implies that $g_n \xi \in g_n(\overline{X_v} \setminus U_v^2) \subset V'_v$ as needed.

Finally, if $z \in Z$, then let $x = p(z)$. Suppose $x \in F$, so that $g_n x = x$ for all n . By definition of $z \notin V_{\mathcal{U}, \varepsilon}(\xi_-)$, we know that for any vertex v of σ_x , $c(z) \notin U_v$, so by (1) of the Claim 7.2.2, we have $g_n c(z) \in V'_v$ for all such v , hence $g_n z \in V_{\mathcal{V}, \varepsilon}(\xi_+)$. If $x \notin F$, then either $x \in D(\xi_-)$ or $x \in (X \setminus \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-) \cup D(\xi_-))$, and these cases are covered by (3) and (2) of Claim 7.2.2. $\square$

As explained just before Claim 7.2.2, because $V_{\mathcal{U}, \varepsilon}(\xi_-), V_{\mathcal{V}, \varepsilon}(\xi_+)$ were arbitrary neighborhoods

of ξ_+, ξ_- , the previous claim shows (g_n, ξ_+, ξ_-) is an ART, proving the lemma. $\square$

Lemma 7.3. Let $(g_n)_n$ be an injective sequence of elements. Suppose that for some (hence any) vertex v the sequence $(g_nv)_n$ is bounded but there are no vertices u, w so that $g_nu = w$ for infinitely many n . Then $(g_n)_n$ is a convergence sequence.

Proof. We begin with 2 claims which explore the assumption on how the $(g_n)_n$ act.

Claim 7.3.1.

- (R1) If K, K' are finite subcomplexes of X , then for all n large enough, $g_nK' \cap K = \emptyset$.
- (R2) For any vertex w of X and $\xi, \xi' \in \partial_{\text{Stab}}G$, the complex spanned by $g_n\text{Geod}(w, D(\xi))$ does not meet $D(\xi')$ for n large enough.
- (R3) For any vertex w of X , $x \in \overline{X}$, and $\xi \in \partial_{\text{Stab}}G$, the geodesic $g_n[w, x]$ does not meet $D(\xi)$ for n large enough.

Proof of Claim. For a contradiction, suppose $g_nK' \cap K \neq \emptyset$ for infinitely many n . Because K is finite, some vertex of K appears infinitely often in the sets $g_nK' \cap K$, say w . There are finitely many possibilities for $g_n^{-1}w \in K'$, so some vertex of K' appears infinitely many times, say u . But then $g_nu = w$ for infinitely many n , contradicting the assumption on $(g_n)_n$. This proves R1.

Both R2 and R3 are special cases of R1. For R2, take $K = D(\xi)$ and K' the subcomplex spanned by $\text{Geod}(w, D(\xi))$. Then K has at most $d_{\max}$ simplices and K' is finite by Lemma 2.33.

For R3, the set $\{g_nw, n \geq 0\}$ is bounded by assumption and $D(\xi)$ has diameter at most A , so there is some constant M so that $D(\xi) \subset B(g_nw, M)$ for all n , hence any intersection between $g_n[w, x]$ and $D(\xi)$ happens in $B(g_nw, M)$. By Lemma 2.34, the initial segment of $[w, x]$ of length M is contained in a finite subcomplex, say K' . Then R3 follows from R1 with $K = D(\xi)$. $\square$

Claim 7.3.2. Let v_0 be a vertex of X , $\varepsilon \in (0, 1)$, $W \subset \overline{X}$, and $\xi_1 \in \partial_{\text{Stab}}G$. Fix a vertex v_1 of $D(\xi_1)$ as the basepoint to consider cones and let $\mathcal{U}' \subset \mathcal{U}$ be two ξ_1 -families with $\mathcal{U}'$ $3d_{\max} + 1$ -nested in $\mathcal{U}$. Suppose that for all $n \geq N$,

1. $g_nv_0 \in \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi_1)$ and
2. for all $x \in W$, $g_n[v_0, x]$ does not meet $D(\xi_1)$.

Then $g_nW \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi_1)$ for all $n \geq N$.

Proof of Claim. Fix $x \in W$ and assume for now that $x \in X$. We will use Lemma 2.36 three times to get a path of simplices from $\sigma_{\xi_1, \varepsilon}(g_nv_0)$ to $\sigma_{\xi_1, \varepsilon}(g_nx)$ in $N(\xi_1)$, then apply the Crossing Lemma 4.21 to transfer assumption (1) from g_nv_0 to g_nx .




Figure 8: The situation of Claim 7.3.2. The blue line represents the path of simplices constructed. The left and right hand segments are the first two applications of Lemma 2.36 and the middle segment is the path of simplices along $[u, u']$.

Fix $n \geq N$. Because $D(\xi_1)$ is compact and disjoint from $g_n[v_0, x]$, we can choose $y \in D(\xi_1)$ achieving $d(g_n[v_0, x], D(\xi_1)) > 0$. Let τ be a simplex of $Lk(\xi_1)$ whose interior is met by $[y, g_nv_0]$ at a point u . Similarly, let τ' be a simplex of $Lk(\xi_1)$ whose interior is met by $[y, g_nx]$ at a point u' . By choosing u, u' closer to $D(\xi_1)$, we may assume $g_nv_0 \neq u, g_nx \neq u$. See Figure 8.

The first two applications of Lemma 2.36 are easy; Using $v_1, y \in D(\xi_1)$ as points in a convex subcomplex $D(\xi_1)$ and g_nv_0 as our finite subcomplex, we have a path of simplices of length at most d_{max} from $\sigma_{\xi_1, \varepsilon}(g_nv_0)$ to τ . Using g_nx instead of g_nv_0 , we have a path of simplices of length at most d_{max} from τ' to $\sigma_{\xi_1, \varepsilon}(g_nx)$.[†]

Now we need one from τ to τ' . Let w be the closest point of $D(\xi_1)$ to g_nv_0 and consider the geodesic triangle with corners g_nv_0, w, y . The leg $[w, y]$ is contained in $D(\xi_1)$ by convexity and $u \in [g_nv_0, y]$ by definition. The CAT(0) inequality implies $d(u, D(\xi_1)) < d(g_nv_0, D(\xi_1))$. Using g_nx instead of g_nv_0 , we get $d(u', D(\xi_1)) < d(g_nx, D(\xi_1))$. In any CAT(0) metric space, the distance to a convex set along a geodesic is a convex function, see [BH99, II.2]. Letting γ parameterize the geodesic $[u, u']$ and using the geodesic $g_n[v_0, x]$ as our convex set, we get the first of the following inequalities:

$$\begin{aligned} d(\gamma(t), g_n[v_0, x]) &\leq \max(d(u, g_n[v_0, x]), d(u', g_n[v_0, x])) \\ &< d(y, g_n[v_0, x]) \\ &= d(D(\xi_1), g_n[v_0, x]). \end{aligned}$$

This shows $[u, u']$ stays strictly closer to $g_n[v_0, x]$ than any point of $D(\xi_1)$, so $[u, u']$ does not meet $D(\xi_1)$. By the convexity of $D(\xi_1)$, we also have that $[u, u'] \subset D^\varepsilon(\xi)$ since $u, u' \in D^\varepsilon(\xi)$. Therefore $[u, u']$ gives a path of simplices in $Lk(\xi_1)$ from τ to τ' , which has length at most d_{max} since $[u, u']$ is a geodesic in $N(\xi_1)$. Combining these, we have a path of simplices

$$\sigma_{\xi_1, \varepsilon}(g_n v_0) \longrightarrow \tau \longrightarrow \tau' \longrightarrow \sigma_{\xi_1, \varepsilon}(g_n x)$$

Each subpath has length at most d_{max} , so the concatenation has length at most $3d_{max}$. Since $g_n v_0 \in \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi_1)$ and $\mathcal{U}'$ is $3d_{max} + 1$ nested in $\mathcal{U}$, the Crossing Lemma 4.21 applied to this path implies $g_n x \in \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi_1)$ as claimed.

If $x = \eta \in \partial X$, the only change is in the sentence marked $\dagger$. Since η doesn't have a finite subcomplex to apply Lemma 2.36 to, we use the CAT(0) inequality and choose x' far along $g_n[v_0, \eta]$ so that $\sigma_{\xi_1, \varepsilon}(g_n \eta) \subset st(\sigma_{\xi_1, \varepsilon}(x'))$, and use x' as we used x above. We get a path of simplices of length $3d_{max}$ from $\sigma_{\xi_1, \varepsilon}(g_n v_0)$ to $\sigma_{\xi_1, \varepsilon}(x')$, and the containment $\sigma_{\xi_1, \varepsilon}(g_n \eta) \subset \sigma_{\xi_1, \varepsilon}(x')$ adds possibly one step to the path of simplices, hence the $+1$ in $3d_{max} + 1$. This proves the claim. $\square$

When assumptions (1) and (2) of Claim 7.3.2 hold for some N , we say N is *suited to* $v_0, v_1, W, \varepsilon, \xi_1$ and $\mathcal{U}' \subset \mathcal{U}$, or suited to W if the rest is clear.

With these 2 claims, we can start on the real proof. If $\partial_{Stab} G = \emptyset$, G is hyperbolic because it acts geometrically on the δ -hyperbolic space X and our main theorem is trivial. Otherwise, choose $\xi_0 \in \partial_{Stab} G$. Because $\bar{Z}$ is compact and metrizable, after a subsequence we can find points $\xi_+, \xi_- \in \partial G$ so that $g_n \xi_0 \rightarrow \xi_+$ and $g_n^{-1} \xi_0 \rightarrow \xi_-$. Because the $g_n D(\xi_0)$ lie in a bounded set, $\xi_+, \xi_- \notin \partial X$. Fix a vertex $v_+ \in D(\xi_+)$ as a basepoint.

Claim 7.3.3. For any $\xi \in \partial_{Stab} G$, $g_n \xi \rightarrow \xi_+$ and $g_n^{-1} \xi \rightarrow \xi_-$.

Proof of Claim. With the terminology above, this is simple to argue. Fix an arbitrary open neighborhood of ξ_+ , say $V_{\mathcal{U}, \varepsilon}(\xi_+)$, and let $\mathcal{U}' \subset \mathcal{U}$ be a ξ_+ -family $3d_{max} + 1$ -nested in $\mathcal{U}$. Because $g_n \xi_0 \rightarrow \xi_+$, we can choose N so that $n \geq N$ implies $g_n \xi_0 \in V_{\mathcal{U}', \varepsilon}(\xi_+)$. Using R1 with $K = D(\xi_0)$ and $K' = D(\xi_+)$ and increasing N if necessary, we can assume $n \geq N$ implies $g_n D(\xi_0) \cap D(\xi_+) = \emptyset$, which implies $g_n v_0 \in \widetilde{\text{Cone}}_{\mathcal{U}', \varepsilon}(\xi_+)$. This satisfies (1) of Claim 7.3.2. Using R2 and possibly increasing N , we have $g_n \text{Geod}(v_0, D(\xi)) \cap D(\xi_+) = \emptyset$ for all $n \geq N$, so assumption (2) holds for all $x \in D(\xi)$. Then N is suited to $D(\xi)$, hence $g_n D(\xi) \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi_+)$, hence $g_n \xi \in V_{\mathcal{U}, \varepsilon}(\xi_+)$ for all $n \geq N$. Since $V_{\mathcal{U}, \varepsilon}(\xi_+)$ was arbitrary, this proves $g_n \xi \rightarrow \xi_+$.

Using g_n^{-1}, ξ_- instead of g_n, ξ_+ in the paragraph above gives the second half of the claim. $\square$

Claim 7.3.4. For any $z \in \bar{Z}$ with $z \neq \xi_-$ and open neighborhood $V_{\mathcal{U}, \varepsilon}(\xi_+)$ of ξ_+ , there is a neighborhood U of z and N so that $g_n U \subset V_{\mathcal{U}, \varepsilon}(\xi_+)$ for all $n \geq N$.

Proof of Claim. Fix $V_{\mathcal{U},\varepsilon}(\xi_+)$ and let $\mathcal{U}'$ be $3d_{max} + 1$ -nested in $\mathcal{U}$. Using Claim 7.3.3 with ξ_+ , we have $g_n \xi_+ \rightarrow \xi_+$ and using R1 with $K = K' = D(\xi_+)$, we can find N_0 large enough that $g_n v_+ \in \widetilde{\text{Cone}}_{\mathcal{U}',\varepsilon}(\xi_+)$ for all $n \geq N_0$. We assume all N below are larger than N_0 so that assumption (1) of Claim 7.3.2 is always satisfied. There are 3 cases.

Suppose $z = \eta \in \partial X$. We find a neighborhood $W \subset \overline{X}$ of η and N is suited to W so that $g_n W \subset \widetilde{\text{Cone}}_{\mathcal{U},\varepsilon}(\xi_+)$ for all $n \geq N$, hence $g_n V_W(\eta) \subset V_{\mathcal{U},\varepsilon}(\xi_+)$ as desired.

Because $(g_n v_+)_n$ and $N(\xi_+)$ are bounded, the number

$$M := \sup_{y \in N(\xi_+), n \geq 0} d(g_n v_+, y)$$

is finite. We claim there are finitely many n so that $g_n[v_+, \eta]$ meets $N(\xi_+)$. If there were infinitely many such n , then we could pass to a subsequence and find $x_n \in [v_+, \eta]$ so that $g_n x_n \in N(\xi)$ for all n . Each x_n has $d(v_+, x_n) \leq M$, and the initial segment of $[v_+, \eta]$ of length M meets finitely many simplices by Lemma 2.34, so after a subsequence we can assume all x_n are in the same simplex, say $\sigma = \sigma_{x_n}$ for all n . Then $g_n \sigma \cap D(\xi) \neq \emptyset$ for all n , and because $\sigma, D(\xi)$ each have finitely many vertices, this leads to a contradiction of the assumption on $(g_n)_n$. This shows that indeed there are only finitely many n so that $g_n[v_+, \eta]$ meets $N(\xi_+)$. Because $D^\varepsilon(\xi) \subset N(\xi)$ for all $\varepsilon \in (0, 1)$, we can fix some $\varepsilon \in (0, 1)$ and N so that $g_n[v_+, \eta]$ does not meet $D^\varepsilon(\xi)$ for all $n \geq N$.

We claim N is suited to $W = V_{M,\varepsilon}(\eta)$. Let $x \in V_{M,\varepsilon}(\eta)$ and for a contradiction suppose $g_n[v_+, x]$ meets $D(\xi_+)$ at a point y . By definition of M , $d(g_n v_+, y) \leq M$, so $g_n[v_+, x]$ meets $D(\xi_+)$ in time at most M . But during this time, $g_n[v_+, x]$ stays within ε of $g_n[v_+, \eta]$ because $x \in V_{M,\varepsilon}(\eta)$, hence avoids $D(\xi_+)$, so this contradicts the definition of ε and no such y can exist. Thus $V_W(\eta), N$ satisfy the claim.

Suppose $z = \xi \in \partial_{\text{Stab}} G$. We construct a neighborhood $V_{\mathcal{V}',\varepsilon'}(\xi)$ and N suited to $D(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{V},\varepsilon'}(\xi)$. It follows that $g_n V_{\mathcal{V},\varepsilon'}(\xi) \subset V_{\mathcal{U},\varepsilon}(\xi_+)$ for all $n \geq N$, satisfying the claim.

Because $\xi \neq \xi_-$, we can choose neighborhoods $V_{\mathcal{V},\varepsilon'}(\xi)$ and $V_{\mathcal{W},\varepsilon''}(\xi_-)$ with disjoint cones by Lemma 6.15. Let $\mathcal{V}'$ be d_{max} -nested in $\mathcal{V}$. Because $g_n^{-1} \xi_+ \rightarrow \xi_-$ by Claim 7.3.3 and using R1 with $K = D(\xi_-)$, $K' = D(\xi_+)$, we can choose N so that $g_n^{-1} D(\xi_+) \subset \widetilde{\text{Cone}}_{\mathcal{W},\varepsilon''}(\xi_-)$ for all $n \geq N$. Using R2 and possibly increasing N , we can assume $g_n \text{Geod}(v_+, D(\xi)) \cap D(\xi_+) = \emptyset$ for all $n \geq N$.

We claim that this N is suited to $D(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{V}',\varepsilon'}(\xi)$. It is clearly suited to $D(\xi)$ by the last sentence of the previous paragraph. Let $x \in \widetilde{\text{Cone}}_{\mathcal{V}',\varepsilon'}(\xi)$ and for a contradiction suppose $g_n[v_+, x]$ meets $D(\xi_+)$ for some $n \geq N$, say $y \in g_n[v_+, x] \cap D(\xi_+)$. See Figure 9. Now $[v_+, x]$ goes through $D(\xi)$ by the Genuine Shadow Proposition 4.27, so let x' be the last point of $D(\xi)$ along $[v_+, x]$ and decompose $g_n[v_+, x] = g_n[v_+, x'] \cup g_n[x', x]$. By the last sentence of the previous paragraph, y cannot be in the first half of this decomposition, so it must be in the second. Translating by g_n^{-1} and recalling that x' is the *last* point of $[v_+, x]$ in $D(\xi)$, we see that $[v_+, g_n^{-1}y]$ goes through $D(\xi)$.

If $g_n^{-1}y$ is outside of $N(\xi)$, then $\sigma_{\xi,\varepsilon'}(g_n^{-1}y) = \sigma_{\xi,\varepsilon'}(x)$, but maybe $g_n^{-1}y$ is close to $D(\xi)$ and these




Figure 9: A diagram of the situation in Claim 7.3.4 with $z = \xi$.

exit simplices are different. Either way, the portion of $[x', x]$ in $D^{\varepsilon'}(\xi)$ gives a path of simplices of length at most d_{max} in $Lk(\xi)$ between these exit simplices. Because $x \in \widetilde{\text{Cone}}_{\mathcal{V}', \varepsilon}(\xi)$ and $\mathcal{V}'$ is d_{max} -nested in $\mathcal{V}$, the Crossing Lemma 4.21 implies $g_n^{-1}y \in \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon}(\xi)$. On the other hand, $g_n^{-1}y \in g_n^{-1}D(\xi_+) \subset \widetilde{\text{Cone}}_{\mathcal{W}, \varepsilon''}(\xi_-)$ by the choice of N . Therefore

$$g_n^{-1}y \in \widetilde{\text{Cone}}_{\mathcal{W}, \varepsilon''}(\xi_-) \cap \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon'}(\xi).$$

But this set is empty because we chose the cones to be disjoint, so this is a contradiction. Hence N is indeed suited to $D(\xi) \cup \widetilde{\text{Cone}}_{\mathcal{V}, \varepsilon'}(\xi)$, and this case is done.

Finally, suppose $z \in Z$ and let $x = p(z)$. Choose δ so that $B(x, \delta) \subset st(\sigma_x)$, and let $U = p^{-1}(B(x, \delta))$. By Lemma 2.33, there is some constant M so that $[v_+, x]$ meets M simplices (M may be very large if x is far away, but it won't matter), and we let $\mathcal{U}''$ be a ξ_+ -family $M+1$ -refined in $\mathcal{U}$. Because $g_n \xi_+ \rightarrow \xi_+$, we can choose N large enough so that $g_n v_+ \in \widetilde{\text{Cone}}_{\mathcal{U}'', \varepsilon}(\xi_+)$ for all $n \geq N$. Using R3 and possibly increasing N , we can assume $g_n[v_+, x] \cap D(\xi_+) = \emptyset$ so that $[g_n v_+, x]$ gives a path of simplices in $X \setminus D(\xi_+)$ of length M from $g_n v_+$ to $g_n x$.

If $z' \in U$, let $x' = p(z')$. By the choice of U , $\sigma_x \subset \sigma_{x'}$, so we have a path of simplices of length at most $M+1$ from $g_n v_+$ to $g_n \sigma_{x'}$. The Refinement Lemma 4.28 applied to this path of simplices implies $g_n \sigma_{x'} \subset \widetilde{\text{Cone}}_{\mathcal{U}, \varepsilon}(\xi_+)$, hence $g_n z' \in V_{\mathcal{U}, \varepsilon}(\xi_+)$. Since z' was arbitrary, this implies $g_n U \subset V_{\mathcal{U}, \varepsilon}(\xi_+)$ and U, N satisfy the claim. $\square$

Now we show (g_n, ξ_+, ξ_-) is an ART. Fix a compact set $K \subset \overline{Z} \setminus \{\xi_-\}$ and neighborhood $V_{\mathcal{U}, \varepsilon}(\xi_+)$ of ξ_+ . By Claim 7.3.4, every $z \in K$ has a neighborhood U_z and N_z so that $g_n U_z \subset V_{\mathcal{U}, \varepsilon}(\xi_+)$ for $n \geq N_z$. By taking a finite cover of K and a maximum, we can find N so that $g_n K \subset V_{\mathcal{U}, \varepsilon}(\xi_+)$. $\square$

Lemma 7.4. Suppose $(g_n)_n$ is an injective sequence of elements and suppose for some vertex v_0 so that $d(v_0, g_nv_0) \rightarrow \infty$. Then $(g_n)_n$ is a convergence sequence.

Proof. As in the previous lemma, we may assume $\partial_{Stab}G \neq \emptyset$, fix $\xi_0 \in \partial_{Stab}G$ and $v_0 \in D(\xi_0)$, then use the compactness of $\bar{Z}$ to find a subsequence and points $\xi_+, \xi_- \in \partial G$ so that $g_n\xi_0 \rightarrow \xi_+$ and $g_n^{-1}\xi_0 \rightarrow \xi_-$. To show (g_n, ξ_+, ξ_-) is an ART it suffices to fix arbitrary neighborhoods W_- of ξ_- and W_+ of ξ_+ and show that for n large enough, $g_n(\bar{Z} \setminus W_-) \subset W_+$. Without loss of generality, W_- is a basic open neighborhood and we set $K = \bar{Z} \setminus W_-$. We will argue the cases of $\xi_- \in \partial X$ and $\xi_- \in \partial_{Stab}G$ in parallel.

If $\xi_- \in \partial X$, let $W_- = V_U(\xi_-)$ for some neighborhood $U \subset \bar{X}$ of ξ_- . Since $g_n^{-1}\xi_0 \rightarrow \xi_-$ in $\bar{Z}$, we know $g_n^{-1}v_0 \rightarrow \xi_-$ in $\bar{X}$. Because domains have diameter at most A and there are finitely many isometry types of simplices, we can choose a constant α so that $\text{diam}(N(\xi)) < \alpha$ for all $\xi \in \partial_{Stab}G$. By Lemma 5.3, we can find a subneighborhood, say $\xi_- \in U' \subset U$ so that $d(X \setminus U, U') > \alpha$. With these choices, if ξ is any point of $\partial_{Stab}G$, then either $D(\xi) \cap U' = \emptyset$, or $D(\xi) \cap U' \neq \emptyset$ hence $D(\xi) \subset U$.

If $\xi_- \in \partial_{Stab}G$, let $W_- = V_{\mathcal{U}, \varepsilon}(\xi_-)$ for some ξ_- -family $\mathcal{U}$ and $\varepsilon \in (0, 1)$. Recall that δ_0 is the hyperbolicity constant of X , and apply Lemma 2.34 to find a constant d so that any geodesic of length δ_0 meets at most d simplices. Choose ξ_- -families $\mathcal{U}^3 \subset \mathcal{U}^2 \subset \mathcal{U}^1 \subset \mathcal{U}$, each $\max(d_{\max}, d)$ -refined in the next.

For $x, y, z \in \bar{X}$, we use $(x, y)_z$ for the Gromov product of x, y based at z , and we let $W_k(x) = \{y \in \bar{X} \mid (y, x)_{v_0} \geq k\}$.

Claim 7.4.1. For any $k \geq 0$,

1. if $\xi_- \in \partial X$, then $g_n(\bar{X} \setminus U') \subset W_k(g_nv_0)$ for all n large enough, and
2. if $\xi_- \in \partial_{Stab}G$, then $g_n(\bar{X} \setminus \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)) \subset W_k(g_nv_0)$ for n large enough.

Proof of Claim. Fix k and consider the case $\xi_- \in \partial X$ first. We claim there is a constant C so that for any $x \in X \setminus U'$, $(g_n^{-1}v_0, x)_{v_0} \leq C$ for all n . If not, we could take a sequence of points $y_n \in X \setminus U'$ so that $(g_n^{-1}v_0, y_n)_{v_0} \rightarrow \infty$. But then by the definition of ∂X , the sequence $(y_n)_n$ represents ξ_- , which contradicts that $y_n \in X \setminus U'$ and $\xi_- \in U'$, so C exists. Two Gromov products add up to the side of a triangle, so for any $x \in X \setminus U'$ we have

$$d(v_0, g_nv_0) = (v_0, g_nx)_{g_nv_0} + (g_nv_0, g_nx)_{v_0}$$

The left side grows arbitrarily large by assumption and $(v_0, g_nx)_{g_nv_0} = (g_n^{-1}v_0, x)_{v_0} \leq C$ by the previous paragraph. Therefore we can choose N large enough that $d(v_0, g_nv_0) \geq k + C$ for all $n \geq N$, hence $(g_nv_0, g_nx)_{v_0} \geq k$, proving (1) for $x \in X \setminus U'$.

If $x \in \partial X \setminus U'$, x can be represented by a sequence $(x_m)_m$ in $X \setminus U'$. Our choice of N above implies $(g_n v_0, g_n x_m)_{v_0} \geq k$ for all m and $n \geq N$. In particular $\liminf_m (g_n v_0, g_n x_m) \geq k$ for all such n . Thus

$$(g_n v_0, g_n x)_{v_0} := \sup_{y_m \rightarrow g_n x} \liminf_m (g_n v_0, y_m)_{v_0} \geq \liminf_m (g_n v_0, g_n x_m) \geq k$$

Thus the same choice of N implies $g_n(\overline{X} \setminus U') \subset W_k(g_n v_0)$.

For (2), consider the case $\xi_- \in \partial_{Stab} G$. As above, we claim there is a constant C so that if $x \notin \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$, then $(g_n^{-1} v_0, x)_{v_0} \leq C$ for all n .

If not, we could take a sequence $y_n \notin \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$ so that $(g_n^{-1} v_0, y_n)_{v_0} \rightarrow \infty$. The tripod points of the triangles $v_0, y_n, g_n^{-1} v_0$ give points $a_n \in [v_0, y_n]$ and $b_n \in [v_0, g_n^{-1} v_0]$ so that $d(a_n, b_n) \leq \delta_0$ and $a_n, b_n \rightarrow \infty$. Because $g_n^{-1} \xi_0 \rightarrow \xi_-$, $v_0 \in V(\xi_0)$, and $g_n^{-1} v_0 \rightarrow \infty$, we can choose some N so that $n \geq N$ implies $g_n^{-1} v_0 \in \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi_-)$. Because $b_n \in [v_0, g_n^{-1} v_0]$ and $b_n \rightarrow \infty$, we can increase N and assume that b_n occurs after $[v_0, g_n^{-1} v_0]$ leaves $D^\varepsilon(\xi_-)$ for all $n \geq N$. Then $\sigma_{\xi_-, \varepsilon}(g_n^{-1} v_0) = \sigma_{\xi_-, \varepsilon}(b_n)$, and since $g_n^{-1} v_0 \in \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi_-)$, we have that $b_n \in \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi_-)$ as well. Increasing N even more, we can assume that $d(b_n, N(\xi_-)) \geq 1 + \delta_0$ so that $[a_n, b_n]$ does not meet $D^\varepsilon(\xi_-)$. Since $[a_n, b_n]$ has length at most δ_0 , it gives a path of simplices of length at most d outside $D(\xi_-)$ from b_n to a_n . Because $b_n \in \widetilde{\text{Cone}}_{\mathcal{U}^3, \varepsilon}(\xi_-)$ and $\mathcal{U}^3$ is at least d -refined in $\mathcal{U}^2$, the Refinement Lemma 4.28 implies $a_n \in \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$, hence $y_n \in \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$. But this contradicts the definition of the y_n , so C must exist.

With C in hand, we can proceed exactly as in the previous case with $\widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-)$ instead of U' to prove (2). $\square$

Recall there is a projection $p : Z \sqcup \partial X \rightarrow \overline{X}$ from Proposition 3.7. It is convenient to extend the definition of domains to points $\eta \in \partial X$ by declaring $D(\eta) = \{\eta\}$. For a subset $B \subset \overline{Z}$, we extend the definition of p by setting

$$p(B) = \{p(z), z \in Z \cap B\} \cup \bigcup_{z \in \partial G \cap B} D(z).$$

Claim 7.4.2. For any $k \geq 0$, we have $g_n p(K) \subset W_k(g_n v_0)$ for all n large enough.

Proof of Claim. Consider $\xi_- \in \partial X$ first. We claim that $p(K) \subset \overline{X} \setminus U'$, so that Claim 7.4.1 implies Claim 7.4.2 immediately. Indeed, if $z \in K \cap Z$ or $\eta \in K \cap \partial X$, then $p(z), p(\eta) \notin U$ by definition of $K = \overline{Z} \setminus V_U(\xi_-)$. If $\xi \in K \cap \partial_{Stab} G$, then $p(\xi) = D(\xi)$ and if $D(\xi)$ met U' , then the choice of U' would imply $D(\xi) \subset U$, contradicting $\xi \in K = \overline{Z} \setminus V_U(\xi_-)$. Hence $D(\xi) \subset X \setminus U'$, and we're done with this case.

Consider $\xi_- \in \partial_{Stab} G$ and fix $k \geq 0$. Note that $p(K) \subset N(\xi_-) \cup (\overline{X} \setminus \widetilde{\text{Cone}}_{\mathcal{U}^2, \varepsilon}(\xi_-))$. Indeed, if $z \in Z \setminus V_{\mathcal{U}, \varepsilon}(\xi_-)$, then either $p(z) \notin \text{Cone}_{\mathcal{U}, \varepsilon}(\xi_-)$ or $p(z) \in D^\varepsilon(\xi_-)$ but $c(z)$ prevents z from being

in $V_{\mathcal{U},\varepsilon}(\xi_-)$. If $\eta \in \partial X \setminus V_{\mathcal{U},\varepsilon}(\xi_-)$, then by definition $p(\eta) = \eta \notin \text{Cone}_{\mathcal{U},\varepsilon}(\xi)$. If $\xi \notin V_{\mathcal{U},\varepsilon}(\xi_-)$, then $D(\xi) \cap \widetilde{\text{Cone}_{\mathcal{U}^2,\varepsilon}(\xi_-)} = \emptyset$ by Lemma 5.7, but $D(\xi)$ may meet $D(\xi_-)$.

Because domains have bounded diameter, we can choose $M \geq 0$ so that $B(v_0, M)$ contains $N(\xi_-)$. Choose N so that $n \geq N$ implies $d(v_0, g_n v_0) > 2k + 2M$. If $y \in B(v_0, M)$ and $n \geq N$, then

$$(g_n v_0, g_n y)_{v_0} = \frac{1}{2} \left(\underbrace{d(g_n v_0, v_0)}_{\geq 2k+2M} + \underbrace{d(g_n y, v_0)}_{\geq 0} - \underbrace{d(g_n v_0, g_n y)}_{\leq 2M} \right) \geq \frac{1}{2}(2k + 2M - 2M) = k.$$

In other words, $n \geq N$ implies $g_n B(v_0, M) \subset W_k(g_n v_0)$. Using (2) of Claim 7.4.1 and possibly increasing N , we can assume that $n \geq N$ also implies $g_n \widetilde{\text{Cone}_{\mathcal{U}^2,\varepsilon}(\xi_-)} \subset W_k(g_n v_0)$. Putting all this together, $n \geq N$ implies

$$g_n K \subset g_n \left(N(\xi_-) \cup (\overline{X} \setminus \widetilde{\text{Cone}_{\mathcal{U}^2,\varepsilon}(\xi_-)}) \right) \subset W_k(g_n v_0).$$

□

Claim 7.4.3. For n large enough, $g_n K \subset W_+$.

Proof of Claim. First suppose $\xi_+ \in \partial X$ and let $W_+ = V_{U_+}(\xi_+)$. The $W_k(\xi_+)$ form a neighborhood basis for ξ_+ in $\overline{X}$, so we can fix some large k with $W_k(\xi_+) \subset U_+$. It is enough to show that for n large enough, $g_n p(K) \subset W_k(\xi_+) \subset U_+$.

Using Claim 7.4.2, choose N so that $n \geq N$ implies $g_n p(K) \subset W_{k+\delta_0}(g_n v_0)$. Since the $g_n \xi_0 \rightarrow \xi_+$ in $\overline{Z}$, the $g_n v_0$ must converge to ξ_+ in $\overline{X}$, so the Gromov products $(g_n v_0, g_m v_0)_{v_0} \rightarrow \infty$. Increasing N if necessary, we can assume $n, m \geq N$ implies $(g_n v_0, g_m v_0)_{v_0} > k + \delta_0$. If $x \in p(K)$ and $n, m \geq N$, then the four point condition of hyperbolicity implies

$$(g_n x, g_m v_0)_{v_0} \geq \min\{(g_n x, g_n v_0)_{v_0}, (g_n v_0, g_m v_0)_{v_0}\} - \delta_0 \geq k.$$

In particular, for any $n \geq N$, we have $\liminf_{m \rightarrow \infty} (g_n x, g_m v_0)_{v_0} \geq k$. Applying this to $(g_n x, \xi_+)_{v_0}$, $n \geq N$ implies

$$(g_n x, \xi_+) := \sup_{y_m \rightarrow \xi_+} \liminf_m (g_n x, y_m)_{v_0} \geq \liminf_m (g_n x, g_m v_0) \geq k,$$

hence $g_n x \in W_k(\xi_+)$ as needed.

Now suppose $\xi_+ \in \partial_{\text{Stab}} G$ and let $W_+ = V_{\mathcal{V},\varepsilon'}(\xi_+)$. Recall any geodesic of length at most δ_0 meets at most d simplices, and let $\mathcal{V}'$ be a ξ_+ -family d -refined in $\mathcal{V}$. Choose k large enough so that the δ_0 -neighborhood of $N(\xi_+)$ is contained in $B(v_0, k)$. Applying Claim 7.4.2 with this k , we find N so that $n \geq N$ implies $g_n p(K) \subset W_k(g_n v_0)$. Since $g_n \xi_0 \rightarrow \xi_+$ and $g_n v_0 \rightarrow \infty$, we can increase N if necessary so that $n \geq N$ also implies $g_n v_0 \in \widetilde{\text{Cone}_{\mathcal{V}',\varepsilon}(\xi_+)}$.

With this N , we claim $g_np(K) \subset \widetilde{\text{Cone}}_{V,\varepsilon'}(\xi_+)$. Indeed, if $x \in p(K)$, then we can consider the tripod points of the triangle g_nx, g_nv_0, v_0 to get $a \in [v_0, g_nv_0], b \in [v_0, g_nx]$ so that $d(a, b) \leq \delta_0$ and, because $g_np(K) \subset W_k(g_nv_0)$, $d(a, v_0) = d(b, v_0) \geq k$. Because $N(\xi_+) \subset B(v_0, k)$, we know $\sigma_{\xi_+, \varepsilon}(g_nv_0) = \sigma_{\xi_+, \varepsilon}(a)$, hence $a \in \widetilde{\text{Cone}}_{V', \varepsilon'}(\xi_+)$. Further, the geodesic $[a, b]$ gives a path of simplices of length at most d from σ_a to σ_b , and the choice of k implies this path of simplices doesn't meet $N(\xi)$. The Refinement Lemma 4.28 then implies $b \in \widetilde{\text{Cone}}_{V, \varepsilon'}(\xi_+)$. Just as with a, g_nv_0 , we know that $\sigma_{\xi_+, \varepsilon'}(g_nx) = \sigma_{\xi_+, \varepsilon'}(b)$, hence $g_nx \in \widetilde{\text{Cone}}_{V, \varepsilon'}(\xi_+)$. This proves that $g_np(K) \subset \widetilde{\text{Cone}}_{V, \varepsilon'}(\xi_+)$, and it follows that $g_nK \subset V_{V, \varepsilon'}(\xi_+)$ $\square$

Since $K = \overline{Z} \setminus W_-$ and W_+ were arbitrary, the previous claim completes the proof of the lemma. $\square$

Theorem 7.5. G acts on $\overline{Z}$ as a convergence group.

Proof. If $(g_n)_n$ is an infinite sequence in G , then one of Lemma 7.2, Lemma 7.3, or Lemma 7.4 implies $(g_n)_n$ is a convergence sequence, hence G acts as a convergence group on $\overline{Z}$. $\square$

7.2 G is geometrically finite

In the previous section, we saw that G acts as a convergence group on the compact metrizable space $\overline{Z}$ with limit set ∂G . To apply Yaman's Theorem 2.5, we must show G is geometrically finite, which means showing every point of ∂G is either a conical limit point or a bounded parabolic point.

Lemma 7.6. For each simplex $\sigma \subset X$, the limit set ΛG_σ is exactly ∂G_σ .

Proof. If G_σ is finite, then $\partial G_\sigma = \Lambda G_\sigma = \emptyset$. If G_σ is infinite, then ∂G_σ is nonempty, and G_σ -invariant. Further ∂G_σ is compact by Proposition 6.3, hence closed because $\overline{Z}$ is Hausdorff. $\square$

Lemma 7.7. If σ is a simplex of X and $\xi \in \partial G_\sigma$ is a conical limit point for G_σ acting on $\widehat{X}_\sigma$, then ξ is a conical limit point for G acting on $\overline{Z}$.

Proof. If $\xi \in \partial G_\sigma$ is a conical limit point, then there is a sequence $(g_n)_n$ in G_σ and distinct points $\xi_+, \xi_- \in \partial G_\sigma$ so that $g_n\xi \rightarrow \xi_-$ and $g_n\xi' \rightarrow \xi_+$ for all $\xi' \in \partial G_\sigma \setminus \{\xi\}$. It follows from Proposition 6.3 that $g_n\xi \rightarrow \xi_-$ as a sequence in $\overline{Z}$, so to show ξ is a conical limit point for G acting on $\overline{Z}$, we just need to upgrade the last part of the previous sentence to 'for all $\xi' \in \partial G \setminus \{\xi\}$ '.

Because G acts as a convergence group on $\overline{Z}$, we can pass to a subsequence and find $\xi'_+, \xi'_- \in \partial G$ so that (g_n, ξ'_+, ξ'_-) is an ART. Fix some $z \in \widehat{X}_\sigma^\circ$. Viewing $(g_nz)_n$ as a sequence in $\widehat{X}_\sigma$ and knowing that $z \neq \xi$, we know that $g_nz \rightarrow \xi_+$. On the other hand, $z \neq \xi_-$ so $g_nz \rightarrow \xi'_+$ too, so $\xi_+ = \xi'_+$. For a contradiction, suppose $\xi \neq \xi'_-$. Because (g_n, ξ'_+, ξ'_-) is an ART, we know $g_n\xi \rightarrow \xi'_+ = \xi_+$. But we also know that $g_n\xi \rightarrow \xi_-$ and that $\xi_- \neq \xi_+$, so this is a contradiction. Thus (g_n, ξ_+, ξ_-) is an ART, which completes the upgrade explained at the end of the previous paragraph. $\square$

Lemma 7.8 (Parabolic Points). Suppose v_0 is a vertex of X and $\xi \in \partial G_{v_0}$ is a bounded parabolic point for G_{v_0} acting on ∂G_{v_0} . Then ξ is a bounded parabolic point for G acting on $\partial_{Stab} G$ and the stabilizer of ξ in G is finitely generated.

Proof. Let P' be the stabilizer of ξ in G . Then P' fixes $D(\xi)$ set-wise and because domains have finitely many simplices, there is a finite index subgroup P which fixes $D(\xi)$ point-wise. If $v \in V(\xi)$, then $P < G_v$ and P is infinite, so ξ is also a parabolic point for G_v acting on ∂G_v . Let P_v be the stabilizer of ξ in G_v and recall that P_v is finitely generated by our definition of relatively hyperbolic groups. We have $P < P_v < P'$ and P finite index in P' , so P_v is finite index in P' , and P' must be finitely generated because it has a finitely generated finite index subgroup.

Applying Lemma 4.9 for each $v \in V(\xi)$, we receive a collection of compact sets $\{K_v \subset \overline{X_v} \setminus \{\xi\}, v \in V(\xi)\}$. Each K_v is closed, so we can choose a ξ -family $\mathcal{U}$ with $K_v \cap U_v = \emptyset$ in each $\overline{X_v}$. Using v_0 as our basepoint, $K := \partial_{Stab} G \setminus V_{\mathcal{U}, \frac{1}{2}}(\xi)$ is compact because it is closed in the compact space $\partial_{Stab} G$.

Claim 7.8.1. For any $x \in \overline{X} \setminus D(\xi)$, there is some $p \in P$ so that $px \notin \widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi)$.

Proof of Claim 7.8.1. Because $x \notin D(\xi)$ and $v_0 \in D(\xi)$ is our basepoint, $[v_0, x]$ leaves $D(\xi)$. Let y be the last point of $[v_0, x]$ in $D^{\frac{1}{2}}(\xi)$ so that $\sigma_y = \sigma_{\xi, \frac{1}{2}}(x)$ (it's possible that $y = x$ if x is close to $D(\xi)$). Because P fixes $D(\xi)$ point wise and acts by isometries, it follows that for any $p \in P$, py is the last point of $p[v_0, x] = [v_0, px]$ in $D^{\frac{1}{2}}(\xi)$, hence

$$p\sigma_y = p\sigma_{\xi, \frac{1}{2}}(x) = \sigma_{\xi, \frac{1}{2}}(px).$$

Fix $v \in \sigma_y \cap D(\xi)$ and consider G_{σ_y} . If G_{σ_y} is infinite, then $\partial G_{\sigma_y} \neq \emptyset$ in ∂G_v , and because K_v is a fundamental domain for P acting on $\partial G_v \setminus \{\xi\}$, there is some $p \in P$ so that $p\partial G_{\sigma_y} \cap K_v \neq \emptyset$. Since $K_v \cap U_v = \emptyset$, this means $p\partial G_{\sigma_y} = \partial G_{\sigma_{\mathcal{U}, \frac{1}{2}}(px)} \not\subseteq U_v$, hence $px \notin \widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi)$, as desired. On the other hand, if G_{σ_y} is finite, then by the choice of K_v , there is some $p \in P$ so that $p\overline{X_{\sigma_y}} \cap K_v \neq \emptyset$ in $\overline{X_v}$. Just as before, this implies $px \notin \widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi)$. $\square$

We show that for any $z \in \partial_{Stab} G \setminus \{\xi\}$, there is some $p \in P$ so that $pz \in K$, hence $K \cap \partial G$ is a compact coarse fundamental domain for the action of P on $\partial_{Stab} G \setminus \{\xi\}$.

If $z = \eta \in \partial X$, Claim 7.8.1 immediately implies there is some $p \in P$ so that $p\eta \notin \text{Cone}_{\mathcal{U}, \frac{1}{2}}(\xi)$, hence $p\eta \in K$.

If $z = \xi' \in \partial_{Stab} G \setminus \{\xi\}$, then we need to find $p \in P$ so that $p\xi'$ fails one of the two conditions necessary for $p\xi' \in V_{\mathcal{U}, \frac{1}{2}}(\xi)$. If $D(\xi') \cap D(\xi)$ is empty, then we can choose any $x \in D(\xi')$, apply Claim 7.8.1 to find $p \in P$ so that $px \notin \widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi)$, and conclude that $p\xi' \notin V_{\mathcal{U}, \frac{1}{2}}(\xi)$. If $D(\xi') \cap D(\xi)$ nonempty, then we can choose a vertex $v \in D(\xi') \cap D(\xi)$. Because K_v is a fundamental domain for P acting on K_v , there is some $p \in P$ so that $p\xi' \in K_v$, hence $p\xi' \notin V_{\mathcal{U}, \frac{1}{2}}(\xi)$. $\square$

Lemma 7.9. Every point $\eta \in \partial X$ is a conical limit point.

Proof. Fix a point $\eta \in \partial X$ and a vertex v_0 of X with $\partial G_{v_0} \neq \emptyset$. The geodesic $[v_0, \eta]$ meets infinitely many simplices and there are finitely many G orbits, so it meets some G orbit infinitely many times. This gives a simplex σ and a sequence $(g_n)_n$ so that $[v_0, \eta]$ meets the interior of $g_n^{-1}\sigma$ for each n (we take g_n^{-1} instead of g_n for notational convenience). Then $g_n[v_0, \eta]$ meets the interior of σ for each n . Fix a vertex v of σ .

Claim 7.9.1. There is a subsequence of $(g_n)_n$, elements $h_n \in G_v$, and a point $\xi_+ \in \partial G \setminus \partial G_v$ so that $h_n g_n \overline{X_v} \rightarrow \xi_+$ uniformly.

Proof of Claim. We use v as the basepoint of our topology. Let σ_n^1 be the first open simplex met by $[v, g_n v]$ after leaving v . After a subsequence, all σ_n^1 lie above the same simplex of Y , so they correspond to cosets of the same subgroup of G_v , and we can choose $h_n \in G_v$ so that all $h_n \sigma_n^1 = \sigma_1$ for some simplex $\sigma \subset st(v)$. Then each geodesic $[v, h_n g_n v]$ meets the simplices $v, \sigma_1, \dots$, and after a further subsequence, each $[v, h_n g_n v]$ leaves σ_1 along the same face, say τ_1 . Let $\sigma_n^2 \subset st(\tau_1)$ be the next simplex met by $[v, h_n g_n v]$ after τ_1 . We begin an iterative argument.

1. Take a subsequence so that all σ_n^2 lie over the same simplex of Y , which means there is a simplex $\sigma_2 \subset st(\tau_1)$ and a sequence $(a_n)_n$ in G_{τ_1} so that $\sigma_n^2 = a_n \sigma_2$. Applying Lemma 4.8 to $G_{\sigma_1}, G_{\sigma_2}$ in G_{τ_1} , we get a sequence $k_n \in G_{\sigma_1} < G_v$ so that the sets $k_n a_n \overline{X_{\sigma_2}}$ are either constant or converge to a point $\xi_+ \in \partial G_{\tau_1} \setminus \partial G_{\sigma_1}$.

Suppose we are in this second case where the $k_n a_n \overline{X_{\sigma_2}}$ converge to some $\xi_+ \in \partial G_{\tau_1} \setminus \partial G_{\sigma_1}$. Then the sequence $k_n a_n \sigma_2$ must be infinite, so we can apply the second case of Corollary 6.20 to the sequence $(k_n h_n g_n v)_n$. We receive a subsequence and a $\xi_0 \in \partial G_{\tau_1}$ so that $\overline{X_v} \rightarrow \xi_0$ uniformly. Since the $k_n a_n \overline{X_{\sigma_2}} = k_n \overline{X_{\sigma_2}}$ already converge to ξ_+ , we must have $\xi_0 = \xi_+$. If $\xi_+ \in \partial G_v$, then the initial portion of each $[v, k_n h_n g_n v]$ connects two simplices of $D(\xi_+)$, namely v, τ_1 . Because domains are convex, this initial portion must be contained in $D(\xi_+)$, but it goes through σ_1 and $\xi_+ \notin \partial G_{\sigma_1}$. This contradiction implies $\xi_+ = \xi_0 \notin \partial G_v$, and we have proven the claim.

So either we are finished as in the previous paragraph, or the simplices $k_n a_n \sigma_2 = k_n \sigma_n^2$ are all the same, and we relabel this simplex as σ_2 and all the geodesics $[v, k_n h_n g_n v]$ meet the simplices $v, \sigma_1, \tau_1, \sigma_2, \dots$. Let $g_n^1 = k_n h_n g_n$, and proceed to 2 or 3 depending on whether or not $G_{\sigma_1} \cap G_{\sigma_2}$ is finite or infinite.

2. If $G_{\sigma_1} \cap G_{\sigma_2}$ is finite, we apply Corollary 6.20 to $(g_n^1 v)_n$. If $g_n^1 \overline{X_v}$ converges uniformly to some $\eta' \in \partial X$, then we can take $\xi_+ = \eta'$ and prove the claim. Otherwise, as in the proof of Corollary 6.20, the path of simplices $\sigma_1, \tau_1, \sigma_2$ extends to a path of simplices $\sigma_1, \tau_1, \sigma_2, \dots, \sigma_m$ crossed by every geodesic $[v, g_n^1 v]$, and $g_n^1 \overline{X_v} \rightarrow \xi_+ \in \partial G_{\sigma_m}$. If $\xi_+ \in \partial G_v$, then the convexity

of domains implies this path of simplices is contained in $D(\xi_+)$. We can construct a ξ_+ -path going through each vertex of this path of simplices, and then Lemma 3.16 tells us the path of simplices is stabilized by an infinite subgroup H . But $H < G_{\sigma_1} \cap G_{\sigma_2}$, so this is a contradiction. Therefore $\xi_+ \notin \partial G_v$, and the claim is proven.

3. If $G_{\sigma_1} \cap G_{\sigma_2}$ is infinite, then we pass to a subsequence so that each $[v, g_n^1 v]$ leaves σ_2 in the same face τ_2 and let σ_n^3 be the next simplex met by $[v, g_n^1 v]$ after leaving τ_2 . We know $G_{\sigma_1} \cap G_{\sigma_2}$ is full RQC in G_{σ_2} by Corollary 2.19 and G_{σ_2} is full RQC in G_{τ_2} so $G_{\sigma_1} \cap G_{\sigma_2}$ is full RQC in G_{τ_2} as well. This means we can repeat the argument in step 1 replacing G_{σ_1} with $G_{\sigma_1} \cap G_{\sigma_2}$ and σ_n^2 with σ_n^3 . We either prove the claim, or find a simplex σ_3 and a sequence $k'_n \in G_{\sigma_1} \cap G_{\sigma_2} < G_v$ so that $[v, k'_n g_n^1 v]$ all cross the simplices $v, \sigma_1, \tau_1, \sigma_2, \tau_2, \sigma_3, \dots$. Letting $g_n^2 = k'_n g_n^1$, we again proceed to either step 2 or 3 depending on the cardinality of $G_{\sigma_1} \cap G_{\sigma_2} \cap G_{\sigma_3}$. Hopefully it is clear how to iterate this argument.

Each time we repeat step 3, we add a pair of simplices to a path of simplices $v, \sigma_1, \tau_1, \sigma_2, \dots$. By Lemma 2.34, there is a constant k so that any geodesic of length $A + 1$ meets at most k simplices. Contrapositively, a geodesic which meets more than k simplices must be longer than $A + 1$, hence have finite stabilizer by the definition of A as the acylindricity constant. This means we can repeat step 3 at most k times before getting to step 2, which finishes the proof of the claim. $\square$

We pass to a subsequence of $(g_n)_n$ and choose $h_n \in G_v, \xi_+ \in \partial G \setminus \partial G_v$ as in the claim and let $b_n = h_n g_n$. Recall that $[v_0, \eta]$ meets the interior of $g_n^{-1} \sigma$ for each n , so $b_n g_n \sigma = h_n \sigma \subset st(v)$ and $b_n[v_0, \eta]$ goes through the interior of some simplex of $st(v)$ for each n .

Because G is a convergence group, we can pass to a subsequence and find points $\xi'_+, \xi'_- \in \partial G$ so that (b_n, ξ'_+, ξ'_-) is an ART. Fix some $z \in X_v^\circ$ and note that $p(z) = v$. Because $z \neq \xi'_+, \xi'_-$, the sequences $b_n z, b_n^{-1} z$ converge to ξ'_+, ξ'_- respectively. From the claim, $b_n z \rightarrow \xi_+$, hence $\xi'_+ = \xi_+$. If $U \subset \bar{X}$ is any neighborhood of η , then for large enough n , we have $b_n^{-1} v = g_n^{-1} h_n^{-1} v = g_n^{-1} v \in U$. This shows $b_n^{-1} z \rightarrow \eta$, hence $\eta = \xi'_-$ and this ART is (b_n, ξ_+, η) . From the definition of an ART, we know that every point $\xi \in \partial G \setminus \{\eta\}$ has $b_n \xi \rightarrow \xi_+$, and after a subsequence we can assume $b_n \eta$ converges to some point of ∂G , say ξ_- . To show $(b_n)_n$ makes η into a conical limit point, it remains to show that $\xi_- \neq \xi_+$. For a contradiction, suppose that $b_n \eta \rightarrow \xi_+$.

If $\xi_+ \in \partial X$, then $b_n v_0 \rightarrow \xi_+$. Each geodesic ray $b_n[v_0, \eta]$ meets the bounded set $st(v)$, hence the Gromov products $(b_n v_0, b_n \eta)_v$ are bounded. But if both $b_n v_0, b_n \eta$ converge to ξ_+ , then these Gromov products should go to infinity. This contradiction implies $\xi_+ \notin \partial X$, hence $\xi_+ \in \partial_{Stab} G$.

Let $\mathcal{U}$ be any ξ_+ -family which is admissible with v as a basepoint and recall the last part of the Genuine Shadow Proposition 4.27, namely that because $\xi_+ \notin \partial G_v$, any geodesic from a point of $st(v)$ and a point of $\widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi_+)$ meets $D(\xi_+)$. For each n , the ray $b_n[v_0, \eta]$ can be written as a concatenation $[b_n v_0, x] \cup [x, b_n \eta]$ with $x \in st(v)$. For large enough n , $b_n v_0, b_n \eta \in \widetilde{\text{Cone}}_{\mathcal{U}, \frac{1}{2}}(\xi_+)$,

hence both pieces of the concatenation meet $D(\xi_+)$ and x is in a geodesic connecting two points of $D(\xi_+)$. Because domains are convex, this means $v \in \sigma_x \subset D(\xi_+)$, but this contradicts $\xi_+ \notin \partial G_v$.

Thus $b_n \eta$ cannot converge to ξ_+ , so ξ_+ and ξ_- are distinct and $(b_n)_n$ makes η into a conical limit point. $\square$

Proposition 7.10. G acts as a convergence group on $\overline{Z}$ with limit set ∂G . Further, G acts as a geometrically finite convergence group on ∂G , and the parabolic subgroups of G are finitely generated.

Proof. G acts as a convergence group on $\overline{Z}$ by Theorem 7.5. To understand its limit set, note that ∂G is clearly G -invariant, and is closed because every point of Z admits a neighborhood avoiding ∂G . Because ΛG is the smallest closed G -invariant subset of $\overline{Z}$, this implies $\Lambda G \subset \partial G$. On the other hand, every point of ∂G is either a conical limit point or a bounded parabolic point by Lemma 7.7, Lemma 7.8, and Lemma 7.9. Conical limit points are in ΛG by Proposition 2.12. Because ΛG can also be characterized as the points of $\overline{Z}$ on which G does *not* act properly discontinuously and parabolic points have infinite stabilizer, parabolic points are also in ΛG . Thus $\partial G \subset \Lambda G$ and we have equality.

Since every point of ∂G is either a conical limit point or a bounded parabolic point, G satisfies Definition 2.3 and is a geometrically finite convergence group on ∂G .

Every point in ∂X is a conical limit point by Lemma 7.9, hence not a parabolic point by Proposition 2.12. Thus all parabolic points are points of $\partial_{Stab} G$, and these have finitely generated stabilizer by Lemma 7.8. $\square$

We are finally ready to apply Theorem 2.5. We restate the main theorem and finish its proof.

Theorem 1.1. Let $G(\mathcal{Y}) = (G_\sigma, \psi_a, g_{a,b})$ be a nonpositively curved developable complex of groups over a scwol $\mathcal{Y}$, where each G_σ is a relatively hyperbolic group and each ψ_a is the inclusion of a full relatively quasiconvex subgroup. Let $\mathcal{X}$ be the universal cover of $G(\mathcal{Y})$, and let X be the geometric realization of $\mathcal{X}$ equipped with an M_κ structure. Suppose X is δ -hyperbolic and the action of $G = \pi_1(G(\mathcal{Y}))$ on X is acylindrical. Then G is relatively hyperbolic. The maximal parabolic subgroups of G are virtually parabolic subgroups of vertex stabilizers, and the stabilizer of each simplex in X is a full relatively quasiconvex subgroup of G .

Theorem 7.11. G is relatively hyperbolic. The maximal parabolic subgroups of G are virtually maximal parabolic subgroups of vertex stabilizers and each simplex stabilizer is a full RQC subgroup of G .

Proof. G acts on the compact metrizable space ∂G as a geometrically finite convergence group with finitely generated parabolic subgroups by Theorem 6.18 and Proposition 7.10. By Theorem 2S of [Tuk94], either $|\partial G| \leq 2$ or ∂G is an infinite perfect set. If $|\partial G| = 0, 1$ or 2 , then G is either finite,

hyperbolic relative to itself, or virtually cyclic, so our main theorem is trivial. Therefore we can assume ∂G is infinite and perfect and Theorem 2.5 implies G is relatively hyperbolic.

Towards understanding the maximal parabolic subgroups, all points of ∂X are conical limit points by Lemma 7.9, so if $\xi \in \partial_{Stab} G$ is a parabolic point with $P = Stab_G(\xi)$, then P permutes the simplices of $D(\xi)$. Because domains are finite, we can pass to a finite index subgroup, say P' , which fixes $D(\xi)$ pointwise, hence $P' < G_\sigma$ for every simplex $\sigma \subset D(\xi)$, and P' makes ξ a parabolic point for G_v acting on ∂G_v for any $v \in V(\xi)$.

It follows from Lemma 7.6 that simplex stabilizers are RQC subgroups of G . To see they are *full* RQC subgroups, take a simplex $\sigma \subset X$ and a maximal parabolic subgroup $P = Stab_G(\xi)$ with $\xi \in \partial_{Stab} G$ so that $G_\sigma \cap P$ is infinite.

We claim $\sigma \subset D(\xi)$. Take an infinite sequence $g_n \in G_\sigma \cap P$, then pass to a subsequence and find points $\xi_+, \xi_- \in \partial G$ so that (g_n, ξ_+, ξ_-) is an ART. Fixing any $z \in \widehat{X}_\sigma^\circ$, we know that $z \neq \xi_+, \xi_-$ so $g_n z \rightarrow \xi_+$ and $g_n^{-1} z \rightarrow \xi_-$. But $\widehat{X}_\sigma$ is closed in $\overline{Z}$, so it contains these limit points and $\xi_+, \xi_- \in \partial G_\sigma$. Now either $\xi = \xi_- \in \partial G_\sigma$, or $\xi \neq \xi_-$ so $g_n \xi = \xi \rightarrow \xi_+$, hence $\xi = \xi_+ \in \partial G_\sigma$. Either way, $\xi \in \partial G_\sigma$, which means $\sigma \subset D(\xi)$ and our claim is proven. As above, we use that domains are finite to pass to a finite index subgroup $P' < P$ which fixes $D(\xi)$ pointwise, hence fixes σ . Thus $P' < G_\sigma \cap P < P$ and $G_\sigma \cap P$ is finite index in P as needed. $\square$

References

- [AGM09] Ian Agol, Daniel Groves, and Jason Fox Manning. Residual finiteness, QCERF and fillings of hyperbolic groups. *Geom. Topol.*, 13(2):1043–1073, 2009.
- [Ali05] Emina Alibegović. A combination theorem for relatively hyperbolic groups. *Bull. London Math. Soc.*, 37(3):459–466, 2005.
- [BF92] Mladen Bestvina and Mark Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.
- [BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [Bow98] Brian H. Bowditch. A topological characterisation of hyperbolic groups. *J. Amer. Math. Soc.*, 11(3):643–667, 1998.
- [Bow99] Brian H. Bowditch. Convergence groups and configuration spaces. In *Geometric group theory down under (Canberra, 1996)*, pages 23–54. de Gruyter, Berlin, 1999.

- [Bri91] Martin R. Bridson. *Geodesics and curvature in metric simplicial complexes*. ProQuest LLC, Ann Arbor, MI, 1991. Thesis (Ph.D.)—Cornell University.
- [BZLW22] Michael Ben-Zvi, Jiayi Lou, and Genevieve S. Walsh. Hyperbolic boundaries vs. hyperbolic groups, 2022.
- [Dah03] Francois Dahmani. Combination of convergence groups. *Geom. Topol.*, 7:933–963, 2003.
- [GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [Hae92] André Haefliger. Extension of complexes of groups. *Ann. Inst. Fourier (Grenoble)*, 42(1-2):275–311, 1992.
- [Hru10] G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.
- [HW09] G. Christopher Hruska and Daniel T. Wise. Packing subgroups in relatively hyperbolic groups. *Geom. Topol.*, 13(4):1945–1988, 2009.
- [Mar14] Alexandre Martin. Non-positively curved complexes of groups and boundaries. *Geom. Topol.*, 18(1):31–102, 2014.
- [MMP10] Jason Fox Manning and Eduardo Martínez-Pedroza. Separation of relatively quasiconvex subgroups. *Pacific J. Math.*, 244(2):309–334, 2010.
- [MR08] Mahan Mj and Lawrence Reeves. A combination theorem for strong relative hyperbolicity. *Geom. Topol.*, 12(3):1777–1798, 2008.
- [Mun75] James R. Munkres. *Topology: a first course*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1975.
- [Tra13] Hung Cong Tran. Relations between various boundaries of relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 23(7):1551–1572, 2013.
- [Tuk94] Pekka Tukia. Convergence groups and Gromov’s metric hyperbolic spaces. *New Zealand J. Math.*, 23(2):157–187, 1994.
- [Tuk98] Pekka Tukia. Conical limit points and uniform convergence groups. *J. Reine Angew. Math.*, 501:71–98, 1998.
- [Yam04] Asli Yaman. A topological characterisation of relatively hyperbolic groups. *J. Reine Angew. Math.*, 566:41–89, 2004.
- [Yan12] Wenyan Yang. Limit sets of relatively hyperbolic groups. *Geom. Dedicata*, 156:1–12, 2012.