Algebraic & Geometric Topology 17 (2017) 2145–2176

Acylindrical group actions on quasi-trees

SAHANA H BALASUBRAMANYA

A group G is acylindrically hyperbolic if it admits a non-elementary acylindrical action on a hyperbolic space. We prove that every acylindrically hyperbolic group G has a generating set X such that the corresponding Cayley graph Γ is a (non-elementary) quasi-tree and the action of G on Γ is acylindrical. Our proof utilizes the notions of hyperbolically embedded subgroups and projection complexes. As an application, we obtain some new results about hyperbolically embedded subgroups and quasi-convex subgroups of acylindrically hyperbolic groups.

20F67; 20F65, 20E08

1 Introduction

Definition 1.1 An isometric action of a group G on a metric space (S, d) is *acylindrical* if for every $\epsilon > 0$ there exist $R, N > 0$ such that for every two points x, y with $\mathrm{d}(x, y) \geq R$, there are at most N elements $g \in G$ satisfying

$$\mathrm{d}(x, gx) \leq \epsilon \quad \text{and} \quad \mathrm{d}(y, gy) \leq \epsilon.$$

Obvious examples are provided by geometric (ie proper and cobounded) actions; note, however, that acylindricity is a much weaker condition.

In order to define an acylindrically hyperbolic group, we must define non-elementary actions, for which we will need the following definition and theorem.

Definition 1.2 Let G be a group acting on a hyperbolic metric space S. An element $g \in G$ is called *loxodromic* if the map $\mathbb{Z} \to S$ given by

$$n \mapsto g^n s$$

is a quasi-isometric embedding for some (equivalently any) $s \in S$. Every loxodromic element has exactly two limit points $\{g^{\pm\infty}\}$ on the Gromov boundary ∂S. Two loxodromic elements g, h are said to be *independent* if the sets $\{g^{\pm\infty}\}$ and $\{h^{\pm\infty}\}$ are disjoint.

Published: 3 August 2017 DOI: 10.2140/agt.2017.17.2145

Theorem 1.3 (Osin [12, Theorem 1.1]) *Let G be a group acting acylindrically on a hyperbolic space S. Then exactly one of the following holds:*

(a) *G has bounded orbits.*

(b) *G is virtually cyclic and contains a loxodromic element.*

(c) *G has infinitely many independent loxodromic elements.*

Definition 1.4 An acylindrical action of a group G is said to be *elementary* in cases (a) and (b) above, and *non-elementary* is case (c). Equivalently, a non-elementary acylindrical action of a group G on a hyperbolic space is an action with unbounded orbits, and where G is not virtually cyclic.

Definition 1.5 A group G is called *acylindrically hyperbolic* if it admits a non-elementary acylindrical action on a hyperbolic space.

Over the last few years, the class of acylindrically hyperbolic groups has received considerable attention. It is broad enough to include many examples of interest, eg non-elementary hyperbolic and relatively hyperbolic groups, all but finitely many mapping class groups of punctured closed surfaces, $\mathrm{Out}(F_n)$ for $n \geq 2$, most 3–manifold groups, and finitely presented groups of deficiency at least 2. On the other hand, the existence of a non-elementary acylindrical action on a hyperbolic space is a rather strong assumption, which allows one to prove non-trivial results. In particular, acylindrically hyperbolic groups share many interesting properties with non-elementary hyperbolic and relatively hyperbolic groups. For details we refer to Dahmani, Guirardel and Osin [5], Minasyan and Osin [10], Osin [12; 11] and references therein.

The main goal of this paper is to answer the following.

Question 1.6 Which groups admit non-elementary cobounded acylindrical actions on quasi-trees?

By a quasi-tree we mean a connected graph which is quasi-isometric to a tree. Quasi-trees form a very particular subclass of the class of all hyperbolic spaces. From the asymptotic point of view, quasi-trees are exactly "1–dimensional hyperbolic spaces".

The motivation behind our question comes from the following observation. If instead of cobounded acylindrical actions we consider cobounded proper (ie geometric) ones, then there is a crucial difference between the groups acting on hyperbolic spaces and quasi-trees. Indeed, a group G acts geometrically on a hyperbolic space if and only if G is a hyperbolic group. On the other hand, Stallings' theorem on groups with infinitely many ends and Dunwoody's accessibility theorem implies that groups admitting geometric actions on quasi-trees are exactly virtually free groups. Yet another related observation is that acylindrical actions on unbounded locally finite graphs are necessarily proper. Thus if we restrict to quasi-trees of bounded valence in Question 1.6, we again obtain the class of virtually free groups. Other known examples of groups having non-elementary, acylindrical and cobounded actions on quasi-trees include groups associated with special cube complexes and right-angled Artin groups (see Behrstock, Hagen and Sisto [1], Hagen [6] and Kim and Koberda [8]).

Thus one could expect that the answer to Question 1.6 would produce a proper subclass of the class of all acylindrically hyperbolic groups, which generalizes virtually free groups in the same sense as acylindrically hyperbolic groups generalize hyperbolic groups. Our main result shows that this does not happen.

Theorem 1.7 *Every acylindrically hyperbolic group admits a non-elementary cobounded acylindrical action on a quasi-tree.*

In other words, being acylindrically hyperbolic is equivalent to admitting a non-elementary acylindrical action on a quasi-tree. Although this result does not produce any new class of groups, it can be useful in the study of acylindrically hyperbolic groups and their subgroups. In this paper we concentrate on proving Theorem 1.7 and leave applications for future papers to explore (for some applications, see [10]).

It was known before that every acylindrically hyperbolic group admits a non-elementary cobounded action on a quasi-tree satisfying the so-called *weak proper discontinuity* property, which is weaker than acylindricity. Such a quasi-tree can be produced by using projection complexes introduced by Bestvina, Bromberg and Fujiwara [2]. To the best of our knowledge, whether the corresponding action is acylindrical is an open question. The main idea of the proof of Theorem 1.7 is to combine the Bestvina–Bromberg–Fujiwara approach with an "acylindrification" construction from Osin [12], in order to make the action acylindrical. An essential role in this process is played by the notion of a hyperbolically embedded subgroup, introduced by Dahmani, Guirardel and Osin [5]. This fact is of independent interest since it provides a new setting for the application of the Bestvina–Bromberg–Fujiwara construction.

The above-mentioned construction has been applied in the setting of geometrically separated subgroups (see [5, Section 4.5]). However, not every hyperbolically embedded subgroup $H \leq G$ arises from an action of G on a hyperbolic space in which H is geometrically separated. Nevertheless, it is possible to employ hyperbolically embedded subgroups in this construction, possibly with interesting applications. If fact, we prove much stronger results in terms of hyperbolically embedded subgroups (see Theorem 3.1) of which Theorem 1.7 is an easy consequence, and derive an application which is stated below (see Corollary 3.24).

Corollary 1.8 *Let G be a group. If $H \leq K \leq G$, H is countable and H is hyperbolically embedded in G, then H is hyperbolically embedded in K.*

This result continues to hold even when we have a finite collection $\{H_1, H_2, \dots, H_n\}$ of hyperbolically embedded subgroups in G such that $H_i \leq K$ for $i = 1, 2, \dots, n$. Interestingly, A Sisto obtains a similar result in [14, Corollary 6.10]. His result does not require H to be countable, but under the assumption that $H \cap K$ is a virtual retract of K, it states that $H \cap K \hookrightarrow_h K$. Although similar, these two theorems are independent in the sense that neither follows from the other.

Another application of Theorem 3.1 is to the case of finitely generated subgroups, as stated below (see Corollary 3.27).

Corollary 1.9 *Let H be a finitely generated subgroup of an acylindrically hyperbolic group G. Then there exists a subset $X \subset G$ such that*

(a) *$\Gamma(G, X)$ is hyperbolic, and the action of G on $\Gamma(G, X)$ is non-elementary and acylindrical, and*

(b) *H is quasi-convex in $\Gamma(G, X)$.*

This result indicates that in order to develop a theory of quasi-convex subgroups in acylindrically hyperbolic groups, the notion of quasi-convexity is not sufficient, ie a stronger set of conditions is necessary in order to prove results similar to those known for quasi-convex subgroups in hyperbolic groups. For example, using Rips' construction [13] and Corollary 1.9, one can easily construct an example of an infinite, infinite-index, normal subgroup in an acylindrically hyperbolic group, which is quasi-convex with respect to some non-elementary acylindrical action.

Acknowledgements My heartfelt gratitude to my advisor Denis Osin for his guidance and support, and to Jason Behrstock and Yago Antolin Pichel for their remarks. I am also grateful to the referee for remarks. My sincere thanks to Bryan Jacobson for his thorough proofreading and comments on this paper.

2 Preliminaries

We recall some definitions and theorems which we will need to refer to.

2.1 Relative metrics on subgroups

Definition 2.1 (relative metric) Let G be a group and $\{H_\lambda\}_{\lambda\in\Lambda}$ a fixed collection of subgroups of G. Let $X \subset G$ such that G is generated by X along with the union of all $\{H_\lambda\}_{\lambda\in\Lambda}$. Let $\mathcal{H} = \bigsqcup_{\lambda\in\Lambda} H_\lambda$. We denote the corresponding Cayley graph of G (whose edges are labeled by elements of $X \sqcup \mathcal{H}$) by $\Gamma(G, X \sqcup \mathcal{H})$.

Remark 2.2 It is important that the union in the definition above is disjoint. This disjoint union leads to the following observation: for every $h \in H_i \cap H_j$, the alphabet $\mathcal{H}$ will have two letters representing h in G, one from H_i and another from H_j. It may also be the case that a letter from $\mathcal{H}$ and a letter from X represent the same element of the group G. In this situation, the corresponding Cayley graph $\Gamma(G, X \sqcup \mathcal{H})$ has bigons (or multiple edges in general) between the identity and the element, one corresponding to each of these letters.

We think of $\Gamma(H_\lambda, H_\lambda)$ as a complete subgraph in $\Gamma(G, X \sqcup \mathcal{H})$. A path p in $\Gamma(G, X \sqcup \mathcal{H})$ is said to be *λ–admissible* if it contains no edges of the subgraph $\Gamma(H_\lambda, H_\lambda)$. In other words, the path p does not travel through H_λ in the Cayley graph. Using this notion, we can define a metric $\widehat{d}_\lambda\colon H_\lambda \times H_\lambda \to [0, \infty]$, known as the relative metric, by setting $\widehat{d}_\lambda(h, k)$ for $h, k \in H_\lambda$ to be the length of the shortest admissible path in $\Gamma(G, X \sqcup \mathcal{H})$ that connects h to k. If no such path exists, we define $\widehat{d}_\lambda(h, k) = \infty$. It is easy to check that $\widehat{d}_\lambda$ is a metric.

Definition 2.3 Let q be a path in the Cayley graph of $\Gamma(G, X \sqcup \mathcal{H})$. A non-trivial subpath p of q is said to be an *H_λ–subpath* if the label of p (denoted $\mathrm{Lab}(p)$) is a word in the alphabet H_λ. Such a subpath is further called an *H_λ–component* if it is not contained in a longer H_λ–subpath of q. If q is a loop, we must also have that p is not contained in a longer H_λ–subpath of any cyclic shift of q.

We refer to an H_λ–component of q (for some $\lambda \in \Lambda$) simply by calling it a *component* of q. We note that, on a geodesic, H_λ–components must be single H_λ–edges. In general, however, the subpath p of q may consist of more than one edge.

Let p_1, p_2 be two H_λ–components of a path q for some $\lambda \in \Lambda$. These components are said to be *connected* if there exists a path p in $\Gamma(G, X \sqcup \mathcal{H})$ such that $\mathrm{Lab}(p)$ is a word consisting only of letters from H_λ, and p connects some vertex of p_1 to some vertex of p_2. In algebraic terms, this means that all vertices of p_1 and p_2 belong to the same (left) coset of H_λ. We refer to a component of a path q as *isolated* if it is not connected to any other component of q.

If p is a path, we denote its initial point by p_- and its terminating point by p_+.

Lemma 2.4 [5, Proposition 4.13] *Let G be a group and $\{H_\lambda\}_{\lambda\in\Lambda}$ a fixed collection of subgroups in G. Let $X \subset G$ such that G is generated by X together with the union of all $\{H_\lambda\}_{\lambda\in\Lambda}$. Then there exists a constant $C > 0$ such that for any n–gon p with geodesic sides in $\Gamma(G, X \sqcup \mathcal{H})$, any $\lambda \in \Lambda$ and any isolated H_λ–component a of p, $\widehat{d}_\lambda(a_-, a_+) \le Cn$.*

2.2 Hyperbolically embedded subgroups

Hyperbolically embedded subgroups will be our main tool in constructing the quasi-tree. The notion has been taken from Dahmani, Guirardel and Osin [5], where it was introduced. We recall the definition here.

Definition 2.5 (hyperbolically embedded subgroups) Let G be a group. Let X be a (not necessarily finite) subset of G and let $\{H_\lambda\}_{\lambda\in\Lambda}$ be a collection of subgroups of G. We say that $\{H_\lambda\}_{\lambda\in\Lambda}$ is *hyperbolically embedded* in G with respect to X (denoted by $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h (G, X)$) if the following conditions hold:

(a) The group G is generated by X together with the union of all $\{H_\lambda\}_{\lambda\in\Lambda}$.

(b) The Cayley graph $\Gamma(G, X \sqcup \mathcal{H})$ is hyperbolic, where $\mathcal{H} = \bigsqcup_{\lambda\in\Lambda} H_\lambda$.

(c) For every $\lambda \in \Lambda$, the metric space $(H_\lambda, \widehat{d}_\lambda)$ is proper, ie every ball of finite radius has finite cardinality.

Furthermore, we say that $\{H_\lambda\}_{\lambda\in\Lambda}$ is hyperbolically embedded in G (denoted by $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h G$) if $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h (G, X)$ for some $X \subseteq G$. The set X is called a *relative generating set*.

Since the notion of a hyperbolically embedded subgroup plays a crucial role in this paper, we include two examples borrowed from [5].

Example 2.6 Let $G = H \times \mathbb{Z}$ and $\mathbb{Z} = \langle x \rangle$. Let $X = \{x\}$. Then $\Gamma(G, X \sqcup H)$ is quasi-isometric to a line and is hence hyperbolic. The corresponding relative metric satisfies the inequality $\widehat{d}(h_1, h_2) \le 3$ for every $h_1, h_2 \in H$, which is easy to see from the Cayley graph (see Figure 1, left). Indeed, if Γ_H denotes the Cayley graph $\Gamma(H, H)$, then in its shifted copy $x\Gamma_H$, there is an edge e connecting xh_1 to xh_2 (labeled by $h_1^{-1}h_2 \in H$). There is thus an admissible path of length 3 connecting h_1 to h_2. We conclude that if H is infinite, then H is not hyperbolically embedded in (G, X), since the relative metric will not be proper. In this example, one can also note that the admissible path from h_1 to h_2 contains an H–subpath, namely the edge e, which is also an H–component of this path.

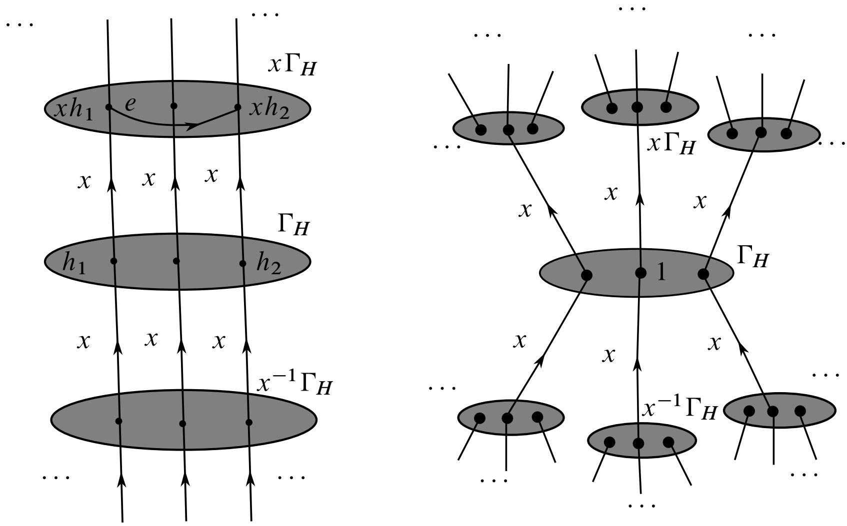

Figure 1: $H \times \mathbb{Z}$ (left) and $H * \mathbb{Z}$ (right)

Example 2.7 Let $G = H * \mathbb{Z}$ and $\mathbb{Z} = \langle x \rangle$. As in the previous example, let $X = \{x\}$. In this case $\Gamma(G, X \sqcup H)$ is quasi-isometric to a tree (see Figure 1, right) and it is easy to see that $\widehat{d}(h_1, h_2) = \infty$ unless $h_1 = h_2$. This means that every ball of finite radius in the relative metric has cardinality 1. We can thus conclude that $H \hookrightarrow_h (G, X)$.

2.3 A slight modification to the relative metric

The aim of this section is to modify the relative metric on countable subgroups that are hyperbolically embedded, so that the resulting metric takes values only in $\mathbb{R}$, ie is finite-valued. This will be of importance in Section 3. The main result of this section is the following.

Theorem 2.8 *Let G be a group. Let $H < G$ be countable and such that $H \hookrightarrow_h G$. Then there exists a left-invariant metric $\widetilde{d}\colon H \times H \to \mathbb{R}$ such that*

(a) $\widetilde{d} \le \widehat{d}$*, and*

(b) $\widetilde{d}$ *is proper, ie every ball of finite radius has finitely many elements.*

Proof There exists a collection of finite, symmetric (closed under inverses) subsets $\{F_i\}$ of H such that $H = \bigcup_{i=1}^{\infty} F_i$ and $1 \subseteq F_1 \subseteq F_2 \subseteq \cdots$.

Let $\widehat{d}$ be the relative metric on H. Let $H_0 = \{h \in H \mid \widehat{d}(1, h) < \infty\}$.

Define a function $w\colon H \to \mathbb{N}$ by

$$w(h) = \begin{cases} \widehat{d}(1, h) & \text{if } h \in H_0, \\ \min\{i \mid h \in F_i\} & \text{otherwise.} \end{cases}$$

Since the F_i are symmetric, $w(h) = w(h^{-1})$ for all $h \in H$. Define a function l on H as follows: for every word $u = x_1 x_2 \cdots x_k$ in the elements of H, set

$$l(u) = \sum_{i=1}^{k} w(x_i).$$

Define a length function on H by

$$|g|_w = \min\{l(u) \mid u \text{ is a word in the elements of } H \text{ that represents } g\},$$

for each g in H. We can now define a metric $d_w\colon H \times H \to \mathbb{N}$ by

$$d_w(g,h) = |g^{-1}h|_w.$$

It is easy to check that d_w is a (finite-valued) well defined metric. Since

$$d_w(ag,ah) = |(ag)^{-1}ah|_w = |g^{-1}a^{-1}ah|_w = |g^{-1}h|_w = d_w(g,h)$$

for all $a, g, h \in G$, the metric d_w is left-invariant. Further, it is easy to see that for all $h \in H$,

$$d_w(1,h) \le w(h).$$

It remains to show that d_w is proper. Let $N \in \mathbb{N}$. Suppose $h \in H$ such that $w(h) \le N$. If $h \in H_0$, then $\widehat{d}(1,h) \le N$, which implies that there are finitely many choices for h, since $\widehat{d}$ is proper. If $h \notin H_0$, then $h \in F_i$ for some minimal i. But each F_i is a finite set, so there are finitely many choices for h. Thus $|\{h \in H \mid w(h) \le N\}| < \infty$ for all $N \in \mathbb{N}$. This implies d_w is proper.

Indeed, if $y \ne 1$ is such that $|y|_w \le n$, then there exists a word u, written without the identity element (which has weight zero), representing y in the alphabet H such that $u = x_1 x_2 \cdots x_r$ and $\sum_{i=1}^{r} w(x_i) \le n$. Since $w(x_i) \ge 1$ for every $x_i \ne 1$, we have $r \le n$. Further, $w(x_i) \le n$ for all i. Thus $x_i \in \{x \in H \mid w(x) \le n\}$ for all i. So there are only finitely many choices for each x_i, which implies there are finitely many choices for y. By definition, $d_w \le \widehat{d}$. So we can set $\widetilde{d} = d_w$. □

2.4 Acylindrically hyperbolic groups

In the following theorem, ∂ represents the Gromov boundary.

Theorem 2.9 *For any group G, the following are equivalent:*

(AH_1) *There exists a generating set X of G such that the corresponding Cayley graph $\Gamma(G,X)$ is hyperbolic, $|\partial\Gamma(G,X)| \ge 2$ and the natural action of G on $\Gamma(G,X)$ is acylindrical.*

(AH_2) *G admits a non-elementary acylindrical action on a hyperbolic space.*

(AH_3) *G contains a proper infinite hyperbolically embedded subgroup.*

It follows from the definitions that $(\mathrm{AH}_1) \Longrightarrow (\mathrm{AH}_2)$. The implication $(\mathrm{AH}_2) \Longrightarrow (\mathrm{AH}_3)$ is non-trivial and was proved by Dahmani, Guirardel and Osin [5]. The implication $(\mathrm{AH}_3) \Longrightarrow (\mathrm{AH}_1)$ was proved by Osin [12].

Definition 2.10 We call a group G *acylindrically hyperbolic* if it satisfies any of the equivalent conditions (AH_1)–(AH_3) from Theorem 2.9.

Lemma 2.11 [5, Corollary 4.27] *Let G be a group, $\{H_\lambda\}_{\lambda\in\Lambda}$ a collection of subgroups of G, and X_1 and X_2 be relative generating sets. Suppose that $|X_1 \Delta X_2| < \infty$. Then $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h (G, X_1)$ if and only if $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h (G, X_2)$.*

Theorem 2.12 [12, Theorem 5.4] *Let G be a group, $\{H_\lambda\}_{\lambda\in\Lambda}$ a finite collection of subgroups of G, and X a subset of G. Suppose that $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h (G, X)$. Then there exists $Y \subset G$ such that the following conditions hold:*

(a) $X \subset Y$.

(b) $\{H_\lambda\}_{\lambda\in\Lambda} \hookrightarrow_h (G, Y)$. *In particular, the Cayley graph $\Gamma(G, Y \sqcup \mathcal{H})$ is hyperbolic.*

(c) *The action of G on $\Gamma(G, Y \sqcup \mathcal{H})$ is acylindrical.*

Definition 2.13 Let (X, d_X) and (Y, d_Y) be two metric spaces. A map $\phi\colon X \to Y$ is said to be a (λ, C)–*quasi-isometry* if there exist constants $\lambda > 1$, $C > 0$ such that

(a) $\frac{1}{\lambda} d_X(a,b) - C \le d_Y(\phi(a), \phi(b)) \le \lambda d_X(a,b) + C$, for all $a, b \in X$, and

(b) Y is contained in the C–neighborhood of $\phi(X)$.

The spaces X and Y are said to be *quasi-isometric* if such a map $\phi\colon X \to Y$ exists. It is easy to check that being quasi-isometric is an equivalence relation. If the map ϕ satisfies only condition (a), then it is said to be a (λ, C)–*quasi-isometric embedding*.

Definition 2.14 A graph Γ with the combinatorial metric d_Γ is said to be a *quasi-tree* if it is quasi-isometric to a tree T.

Definition 2.15 A *quasi-geodesic* is a quasi-isometric embedding of an interval $I \subseteq \mathbb{R}$ (bounded or unbounded) into a metric space X. Note that geodesics are $(1,0)$–quasi-geodesics. By slight abuse of notation, we may identify the map that defines a quasi-geodesic with its image in the space.

Theorem 2.16 [9, Theorem 4.6, bottleneck property] *Let Y be a geodesic metric space. The following are equivalent:*

(a) *Y is quasi-isometric to some simplicial tree Γ.*

(b) *There is some $\mu > 0$ such that for all $x, y \in Y$, there is a midpoint $m = m(x, y)$ with $d(x, m) = d(y, m) = \frac{1}{2}d(x, y)$ and the property that any path from x to y must pass within less than μ of the point m.*

We remark that if m is replaced with any point p on a geodesic between x and y, then the property that any path from x to y passes within less than μ of the point p still follows from (a), as proved below in Lemma 2.18. We will need the following lemma.

Lemma 2.17 [4, Proposition 3.1] *For all $\lambda \geq 1$, $C \geq 0$, $\delta \geq 0$, there exists an $R = R(\delta, \lambda, C)$ such that if X is a δ–hyperbolic space, γ is a (λ, C)–quasi-geodesic in X, and γ' is a geodesic segment with the same endpoints, then γ' and γ are Hausdorff distance less than R from each other.*

Lemma 2.18 *If Y is a quasi-tree, then there exists $\mu > 0$ such that for any point z on a geodesic connecting two points, any other path between the same endpoints passes within μ of z.*

Proof Let T be a tree and $q\colon Y \to T$ be the (λ, C)–quasi-isometry. Let d_Y and d_T denote the metrics in the spaces Y and T, respectively. Note that since T is 0–hyperbolic, Y is δ–hyperbolic for some δ.

Let x, y be two points in Y, joined by a geodesic γ. Let z be any point of γ and let α be another path from x to y. Let V denote the vertex set of α, ordered according to the geodesic γ. Take its image $q(V)$ and connect consecutive points by geodesics (of length at most $\lambda + C$) to get a path β in T from $q(x)$ to $q(y)$. Then the unique geodesic σ in T must be a subset of β. Since $q(V) \subset q \circ \alpha$, we get that any point of σ is at most $\lambda + C$ from $q \circ \alpha$. Also, $q \circ \gamma$ is a (λ, C)–quasi-isometric embedding of an interval, and hence a (λ, C)–quasi-geodesic. Thus, by Lemma 2.17 the distance from $q(z)$ to σ is less than $R = R(0, \lambda, C)$.

Let p be the point on σ closest to $q(z)$. There is a point $w \in Y$ on α such that $d(q(w), p) \leq \lambda + C$. Since $d(p, q(z)) < R$, we have $d(q(w), q(z)) \leq \lambda + C + R$. Thus
$$d(z, w) \leq \lambda^2 + 2\lambda C + R\lambda.$$
Thus α must pass within $\mu = \lambda^2 + 2\lambda C + R\lambda$ of the point z. □

2.5 A modified version of Bowditch's lemma

In this section, $\mathcal{N}_k(X)$ denotes the closed k–neighborhood of a set X in a metric space (S, d_S), ie
$$\mathcal{N}_k(X) = \{s \in S \mid \exists x \in X \text{ such that } d_S(s, x) \leq k\}.$$

In particular, $\mathcal{N}_k(x)$ denotes the closed k–neighborhood of a point x in a metric space.

The following theorem will be used in Section 3. Part (a) is a simplified form of a result taken from [7], which is in fact derived from a hyperbolicity criterion developed by Bowditch [3].

Theorem 2.19 *Let Σ be a hyperbolic graph and Δ be a graph obtained from Σ by adding edges.*

(a) [3] *Suppose there exists $M > 0$ such that for all vertices $x, y \in \Sigma$ joined by an edge in Δ and for all geodesics p in Σ between x and y, all vertices of p lie in an M–neighborhood of x, ie $p \subseteq \mathcal{N}_M(x)$ in Δ. Then Δ is also hyperbolic, and there exists a constant k such that for all vertices $x, y \in \Sigma$, every geodesic q between x and y in Σ lies in a k–neighborhood in Δ of every geodesic in Δ between x and y.*

(b) *If, under the assumptions of* (a)*, we additionally assume that Σ is a quasi-tree, then Δ is also a quasi-tree.*

Lemma 2.20 *Let p, q be two paths in a metric space S between points x and y, such that p is a geodesic and $q \subseteq \mathcal{N}_k(p)$. Then $p \subseteq \mathcal{N}_{2k}(q)$.*

Proof Let z be any point on p. Let p_1, p_2 denote the segments of the geodesic p with endpoints x, z and z, y, respectively.

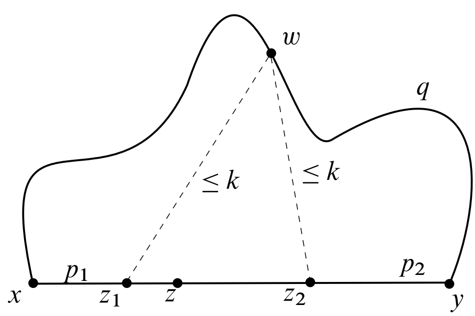

Figure 2: Lemma 2.20

Define a function $f\colon q \to \mathbb{R}$ by $f(s) = d(s, p_1) - d(s, p_2)$. Then f is a continuous function. Further, $f(x) < 0$ and $f(y) > 0$. By the intermediate value theorem, there exists a point w on q such that $f(w) = 0$. Thus $d(w, p_1) = d(w, p_2)$ (see Figure 2). For $i = 1, 2$, let z_i be a point of p_i such that $d(p_i, w) = d(z_i, w)$. Then $d(z_1, w) = d(z_2, w)$. By the hypothesis, $d(w, p) = \min\{d(w, p_1), d(w, p_2)\} \leq k$. So we get that $d(w, p_1) = d(w, p_2) \leq k$. Thus $d(z_1, z_2) \leq 2k$, which implies $d(z, w) \leq 2k$. $\square$

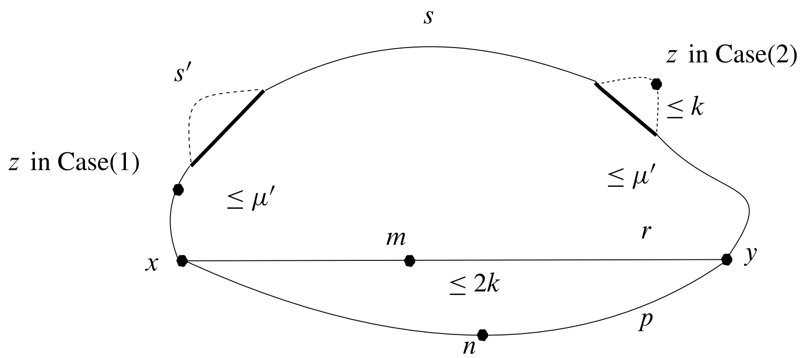

Figure 3: Theorem 2.19

Proof of Theorem 2.19 We proceed with the proof of part (b).

We prove that Δ is a quasi-tree by verifying the bottleneck property from Theorem 2.16. Let d_Σ and d_Δ denote the distances in the graphs Σ and Δ, respectively. Note that the vertex sets of the two graphs are equal.

Let x, y be two vertices. Let m be the midpoint of a geodesic r in Δ connecting them. Let s be any path from x to y in Δ. The path s consists of edges of two types:

(i) edges of the graph Σ;

(ii) edges added in transforming Σ to Δ (marked as bold edges on Figure 3).

Let p be a geodesic in Σ between x and y. By part (a), there exists k such that p is in the k–neighborhood of r in Δ. Applying Lemma 2.20 , we get a point n on p such that

$$d_\Delta(m,n) \leq 2k.$$

Let s' be the path in Σ between x and y, obtained from s by replacing every edge e of type (ii) by a geodesic path $t(e)$ in Σ between its endpoints (marked by dotted lines in Figure 3). Since Σ is a quasi-tree, by Lemma 2.18, there exists $\mu' > 0$ and a point z on s' such that

$$d_\Sigma(z,n) \leq \mu'.$$

Case 1 If z lies on an edge of s of type (i), then

$$d_\Delta(z,m) \leq d_\Delta(z,n) + d_\Delta(n,m) \leq d_\Sigma(z,n) + d_\Delta(n,m) \leq \mu' + 2k.$$

Case 2 If z lies on a path $t(e)$ that replaced an edge e of type (ii), then by part (a),

$$d_\Delta(e_-,m) \leq d_\Delta(e_-,z) + d_\Delta(z,n) + d_\Delta(n,m) \leq k + \mu' + 2k = \mu' + 3k.$$

Thus the bottleneck property holds for $\mu = \mu' + 3k > 0$. $\square$

3 Proof of the main result

Our main result is the following theorem, from which Theorem 1.7 and other corollaries stated in the introduction can be easily derived (see Section 3.5).

Theorem 3.1 *Let $\{H_1, H_2, \dots, H_n\}$ be a finite collection of countable subgroups of a group G such that $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h (G, Z)$ for some $Z \subset G$. Let K be a subgroup of G such that $H_i \leq K$ for all i. Then there exists a subset $Y \subset K$ such that:*

(a) $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h (K, Y)$.

(b) $\Gamma(K, Y \sqcup \mathcal{H})$ *is a quasi-tree, where* $\mathcal{H} = \bigsqcup_{i=1}^n H_i$.

(c) *The action of K on $\Gamma(K, Y \sqcup \mathcal{H})$ is acylindrical.*

(d) $Z \cap K \subset Y$.

3.1 Outline of the proof

Step 1 In order to prove Theorem 3.1, we first prove the following proposition. It is distinct from Theorem 3.1 since it does not require the action of K on the Cayley graph $\Gamma(K, X \sqcup \mathcal{H})$ to be acylindrical.

Proposition 3.2 *Let $\{H_1, H_2, \dots, H_n\}$ be a finite collection of countable subgroups of a group G such that $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h G$ with respect to a relative generating set Z. Let K be a subgroup of G such that $H_i \leq K$ for all i. Then there exists $X \subset K$ such that:*

(a) $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h (K, X)$.

(b) $\Gamma(K, X \sqcup \mathcal{H})$ *is a quasi-tree, where* $\mathcal{H} = \bigsqcup_{1=1}^n H_i$.

(c) $Z \cap K \subset X$.

Step 2 Once we have proved Proposition 3.2, we will utilize an "acylindrification" construction from [12] to make the action acylindrical, which will prove Theorem 3.1. The details of this step are as follows.

Proof By Proposition 3.2, there exists $X \subseteq K$ such that:

(a) $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h (K, X)$.

(b) $\Gamma(K, X \sqcup \mathcal{H})$ is a quasi-tree.

(c) $Z \cap K \subset X$.

By applying Theorem 2.12 to the above, we get that there exists $Y \subset K$ such that:

(a) $X \subseteq Y$.

(b) $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h (K, Y)$. In particular, the Cayley graph $\Gamma(K, Y \sqcup \mathcal{H})$ is hyperbolic.

(c) The action of K on $\Gamma(K, Y \sqcup \mathcal{H})$ is acylindrical.

From the proof of Theorem 2.12 (see [12, Lemma 5.6] in particular), it is easy to see that the Cayley graph $\Gamma(G, Y \sqcup \mathcal{H})$ is obtained from $\Gamma(G, X \sqcup \mathcal{H})$ in a manner that satisfies the assumptions of Theorem 2.19, with $M = 1$ (see [12, Lemma 5.6]). Thus by Theorem 2.19, $\Gamma(K, Y \sqcup \mathcal{H})$ is also a quasi-tree. Further,

$$K \cap Z \subset X \subset Y.$$

Thus Y is the required relative generating set. □

We will thus now focus on proving Proposition 3.2. To do so, we will use a construction introduced by Bestvina, Bromberg and Fujiwara in [2]. We describe the construction below and will retain the same terminology as introduced by the authors in [2].

3.2 The projection complex

Definition 3.3 Let $\mathbb{Y}$ be a set and $\xi > 0$ be a constant. Suppose that for each $Y \in \mathbb{Y}$ we have a function

$$d_Y^\pi\colon (\mathbb{Y}\backslash\{Y\} \times \mathbb{Y}\backslash\{Y\}) \to [0, \infty)$$

that satisfies the following axioms:

(A1) $d_Y^\pi(A, B) = d_Y^\pi(B, A)$ for all $Y \in \mathbb{Y}$ and all $A, B \in \mathbb{Y}\backslash\{Y\}$.

(A2) $d_Y^\pi(A, B) + d_Y^\pi(B, C) \geq d_Y^\pi(A, C)$ for all $Y \in \mathbb{Y}$ and all $A, B, C \in \mathbb{Y}\backslash\{Y\}$.

(A3) $\min\{d_Y^\pi(A, B), d_B^\pi(A, Y)\} < \xi$ for all distinct $Y, A, B \in \mathbb{Y}$.

(A4) $\#\{Y \mid d_Y^\pi(A, B) \geq \xi\}$ is finite for all $A, B \in \mathbb{Y}$.

Let J be a positive constant. Then associated to this data we have the *projection complex* $P_J(\mathbb{Y})$, which is a graph constructed in the following manner: the set of vertices of $P_J(\mathbb{Y})$ is the set $\mathbb{Y}$. To specify the set of edges, one first defines a new function $d_Y\colon (\mathbb{Y}\backslash\{Y\} \times \mathbb{Y}\backslash\{Y\}) \to [0, \infty)$, which can be thought of as a small perturbation of d_Y^π. The exact definition of d_Y can be found in [2]. An essential property of the new function is the following inequality, which is an immediate corollary of [2, Proposition 3.2].

For every $Y \in \mathbb{Y}$ and every $A, B \in \mathbb{Y}\backslash\{Y\}$, we have

(1) $$|d_Y^\pi(A, B) - d_Y(A, B)| \leq 2\xi.$$

The set of edges of the graph $P_J(\mathbb{Y})$ can now be described as follows: two vertices $A, B \in \mathbb{Y}$ are connected by an edge if and only if $d_Y(A, B) \le J$ for every $Y \in \mathbb{Y} \backslash \{A, B\}$. This construction strongly depends on the constant J. Complexes corresponding to different J are not isometric in general.

We would like to mention that if $\mathbb{Y}$ is endowed with an action of a group G that preserves projections (ie $d^{\pi}_{g(Y)}(g(A), g(B)) = d^{\pi}_{Y}(A, B)$), then the action of G can be extended to an action on $P_J(\mathbb{Y})$. We also mention the following proposition, which has been proved under the assumptions of Definition 3.3.

Proposition 3.4 [2, Theorem 3.16] *For a sufficiently large $J > 0$, $P_J(\mathbb{Y})$ is connected and quasi-isometric to a tree.*

Definition 3.5 (nearest point projection) In a metric space (S, d), given a set Y and a point $a \in S$, we define the nearest point projection as

$$\mathrm{proj}_Y(a) = \{y \in Y \mid d(Y, a) = d(y, a)\}.$$

If A, Y are two sets in S, then

$$\mathrm{proj}_Y(A) = \bigcup_{a \in A} \mathrm{proj}_Y(a).$$

We note that in our case, since elements of $\mathbb{Y}$ will come from a Cayley graph, which is a combinatorial graph, the nearest point projection will exist. This is because distances on a combinatorial graph take discrete values in $\mathbb{N} \cup \{0\}$. Since this set is bounded below, we cannot have an infinite strictly decreasing sequence of distances.

We make all geometric considerations in the Cayley graph $\Gamma(G, Z \sqcup \mathcal{H})$. Let $\mathrm{d}_{Z \sqcup \mathcal{H}}$ denote the metric on this graph. Since $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h G$ under the assumptions of Proposition 3.2, from Remark 4.26 of [5] it follows that $H_i \hookrightarrow_h G$ for $i = 1, 2, \dots, n$. By Theorem 2.8, we can define a finite-valued, proper metric $\widetilde{d}_i$ on H_i, for $i = 1, 2, \dots, n$, satisfying

$$\widetilde{d}_i(x, y) \le \widehat{d}_i(x, y) \tag{2}$$

for all $x, y \in H_i$.

We can extend both $\widehat{d}_i$ and $\widetilde{d}_i$ to all cosets gH_i of H_i by setting $\widetilde{d}_i(gx, gy) = \widetilde{d}_i(x, y)$ and $\widehat{d}_i(gx, gy) = \widehat{d}_i(x, y)$ for all $x, y \in H_i$. Let $\widehat{\mathrm{diam}}$ (resp. $\widetilde{\mathrm{diam}}$) denote the diameter of a subset of H_i or a coset of H_i with respect to the $\widehat{d}_i$ (resp. $\widetilde{d}_i$) metric.

Let

$$\mathbb{Y} = \{kH_i \mid k \in K, i = 1, 2, \dots, n\}$$

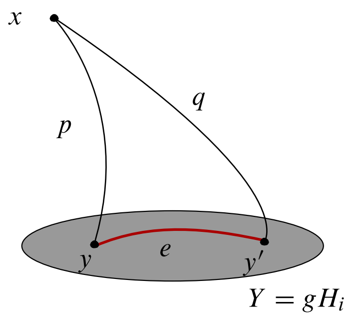

Figure 4: The bold red edge e denotes a single edge labeled by an element of $\mathcal{H}$

be the set of cosets of all H_i in K. We think of cosets of H_i as a subset of vertices of $\Gamma(G, Z \sqcup \mathcal{H})$.

For each $Y \in \mathbb{Y}$ and $A, B \in \mathbb{Y} \setminus \{Y\}$, define

$$d_Y^{\pi}(A, B) = \widetilde{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B)), \tag{3}$$

where $\operatorname{proj}_Y(A)$ is as in Definition 3.5. The fact that (3) is well-defined will follow from Lemma 3.6 and Lemma 3.8, which are proved below. We will also proceed to verify the axioms (A1)–(A4) of the Bestvina–Bromberg–Fujiwara construction in the above setting.

Lemma 3.6 *For any $Y \in \mathbb{Y}$ and any $x \in G$, $\widetilde{\operatorname{diam}}(\operatorname{proj}_Y(x)) \leq 3C$, where C is the constant as in Lemma 2.4. As a consequence, $\widetilde{\operatorname{diam}}(\operatorname{proj}_Y(x))$ is bounded.*

Proof By (2), it suffices to prove that $\widehat{\operatorname{diam}}(\operatorname{proj}_Y(x))$ is bounded. Let $y, y' \in \operatorname{proj}_Y(x)$. Then $\mathrm{d}_{Z \sqcup \mathcal{H}}(x, y) = \mathrm{d}_{Z \sqcup \mathcal{H}}(x, y') = \mathrm{d}_{Z \sqcup \mathcal{H}}(x, Y)$. Without loss of generality, $x \notin Y$, or else the diameter is zero.

Let $Y = gH_i$. Let e denote the edge connecting y and y', which is labeled by an element of H_i. Let p and q denote geodesics between x and y and between x and y', respectively (see Figure 4).

Consider the geodesic triangle T with sides e, p, q. Since p and q are geodesics between the point x and Y, e is an isolated component in T, ie e cannot be connected to either p or q. Indeed, if e is connected to, say, a component of p, then since e_+ and e_- are in Y, it would imply that the geodesic p passes through a point of Y before y. But then y is not the nearest point from Y to x, which is a contradiction. By Lemma 2.4, $\widehat{d}_i(y, y') \leq 3C$. Hence

$$\widehat{\operatorname{diam}}(\operatorname{proj}_Y(x)) \leq 3C.$$

□

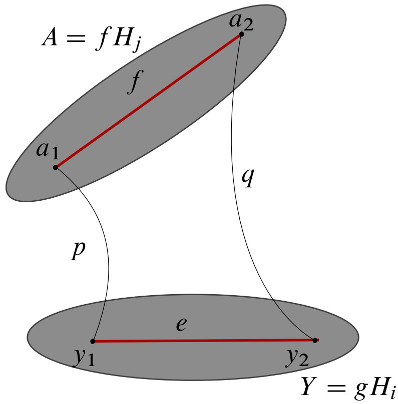

Figure 5: Lemma 3.8

Remark 3.7 Observe that in the previous lemma, we proved the following fact: If x is a point in G and $y \in \mathrm{proj}_Y(x)$, then every geodesic path p between x and y satisfies the property that no vertex of p, except for y, can belong to the coset Y. We will use this fact repeatedly in the following lemmas.

Lemma 3.8 *For every pair of distinct elements $A, Y \in \mathbb{Y}$, $\widehat{\mathrm{diam}}(\mathrm{proj}_Y(A)) \le 4C$, where C is the constant as in Lemma 2.4. As a consequence, $\widehat{\mathrm{diam}}(\mathrm{proj}_Y(A))$ is bounded.*

Proof Let $Y = gH_i$ and $A = fH_j$. Let $y_1, y_2 \in \mathrm{proj}_Y(A)$. Then there exist $a_1, a_2 \in A$ such that $\mathrm{d}_{Z\sqcup\mathcal{H}}(a_1, y_1) = \mathrm{d}_{Z\sqcup\mathcal{H}}(a_1, Y)$ and $\mathrm{d}_{Z\sqcup\mathcal{H}}(a_2, y_2) = \mathrm{d}_{Z\sqcup\mathcal{H}}(a_2, Y)$. Now y_1 and y_2 are connected by a single edge e, labeled by an element of H_i, and similarly, a_1 and a_2 are connected by an edge f, labeled by an element of H_j (see Figure 5). Let p and q be geodesics that connect y_1 to a_1 and y_2 to a_2, respectively. We note that p and/or q may be trivial paths (consisting of a single point), but this does not alter the proof.

Consider e in the quadrilateral Q with sides p, f, q, e. By Remark 3.7, e cannot be connected to a component of p or q.

If $i = j$, then e cannot be connected to f since $A \neq Y$. If $i \neq j$, then obviously e and f cannot be connected. Thus e is isolated in this quadrilateral Q. By Lemma 2.4, $\hat{d}_i(y_1, y_2) \le 4C$. Thus

$$\widehat{\mathrm{diam}}(\mathrm{proj}_Y(A)) \le 4C. \qquad \square$$

Corollary 3.9 *The function d_Y^π defined by (3) is well-defined.*

Proof Since the $\widetilde{d}_i$ metric takes finite values for $i = 1, 2, \dots, n$, using Lemma 3.8, we have that d_Y^{π} also takes only finite values. $\square$

Lemma 3.10 *The function d_Y^{π} defined by (3) satisfies conditions* (A1) *and* (A2) *in Definition 3.3.*

Proof (A1) is obviously satisfied. For any $Y \in \mathbb{Y}$ and any $A, B, C \in \mathbb{Y} \setminus \{Y\}$, by the triangle inequality, we have that

$$\begin{aligned} d_Y^{\pi}(A, C) &= \widetilde{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(C)) \\ &\le \widetilde{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B)) + \widetilde{\operatorname{diam}}(\operatorname{proj}_Y(B) \cup \operatorname{proj}_Y(C)) \\ &= d_Y^{\pi}(A, B) + d_Y^{\pi}(B, C). \end{aligned}$$

Thus (A2) also holds. $\square$

Lemma 3.11 *The function d_Y^{π} from (3) satisfies condition* (A3) *in Definition 3.3 for any $\xi > 14C$, where C is the constant from Lemma 2.4*

Proof By (2), it suffices to prove that

$$\min\left\{\widehat{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B)), \widehat{\operatorname{diam}}(\operatorname{proj}_B(A) \cup \operatorname{proj}_B(Y))\right\} < \xi.$$

Let $A, B \in \mathbb{Y} \setminus \{Y\}$ be distinct. Let $Y = gH_i$, $A = fH_j$ and $B = tH_k$. If

$$\widehat{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B)) \le 14C,$$

then we are done. So suppose that

$$\widehat{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B))) > 14C. \tag{4}$$

Choose $a \in A$, $b \in B$ and $x, y \in Y$ such that $\mathrm{d}_{Z \sqcup \mathcal{H}}(A, Y) = \mathrm{d}_{Z \sqcup \mathcal{H}}(a, x)$ and $\mathrm{d}_{Z \sqcup \mathcal{H}}(B, Y) = \mathrm{d}_{Z \sqcup \mathcal{H}}(b, y)$. In particular,

$$x \in \operatorname{proj}_Y(A), \qquad y \in \operatorname{proj}_Y(B) \tag{5}$$

and $b \in \operatorname{proj}_B(Y)$. Let p, q denote geodesics connecting a to x and b to y, respectively. Let h_1 denote the edge connecting x and y, which is labeled by an element of H_i.

By (5), we have that

$$\widehat{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B)) \le \widehat{\operatorname{diam}}(\operatorname{proj}_Y(A)) + \widehat{\operatorname{diam}}(\operatorname{proj}_Y(B)) + \widehat{d}_i(x, y).$$

Combining this with (4) and Lemma 3.8, we get

$$\begin{aligned} \widehat{d}_i(x, y) &\ge \widehat{\operatorname{diam}}(\operatorname{proj}_Y(A) \cup \operatorname{proj}_Y(B)) - \widehat{\operatorname{diam}}(\operatorname{proj}_Y(A)) - \widehat{\operatorname{diam}}(\operatorname{proj}_Y(B)) \\ &> 14C - 8C = 6C. \end{aligned}$$

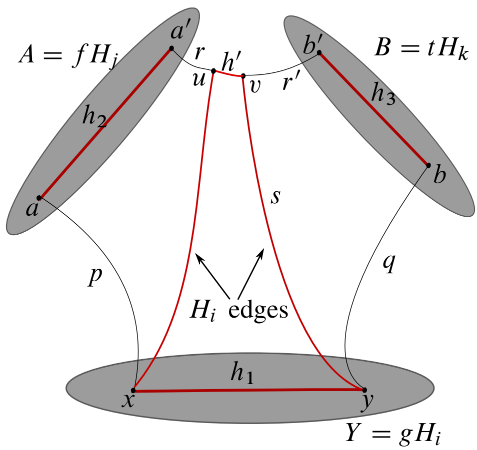

Figure 6: Condition (A3)

Choose any $a' \in A$ and $b' \in \mathrm{proj}_B(a')$. Then $\mathrm{d}_{Z \sqcup \mathcal{H}}(a', B) = \mathrm{d}_{Z \sqcup \mathcal{H}}(a', b')$ (see Figure 6). (Note that if $a' = a$, the following arguments still hold.) Let h_2 and h_3 denote the edges connecting a to a' and b to b', which are labeled by elements of H_j and H_k, respectively. Let r denote a geodesic connecting a' and b'. Consider the geodesic hexagon W with sides p, h_1, q, h_3, r, h_2. Then h_1 is not isolated in W; otherwise, by Lemma 2.4, $\widehat{d}_i(x, y) \leq 6C$, a contradiction.

Thus h_1 is connected to another H_i–component in W. By Remark 3.7, h_1 cannot be connected to a component of p or q. Since A, B, Y are all distinct, h_1 cannot be connected to h_2 or h_3. So h_1 must be connected to an H_i–component on the geodesic r. Let this edge be h', with endpoints u and v, as shown in Figure 6. Let s denote the edge (labeled by an element of H_i) that connects y and v. Let r' denote the segment of r that connects v to b'. Then r' is also a geodesic.

Consider the quadrilateral Q with sides r', h_3, q, s. By using arguments similar to those in the previous paragraph, h_3 cannot be connected to r', q or s. Thus h_3 is isolated in Q. By Lemma 2.4,

$$\widehat{d}_k(b, b') \leq 4C.$$

Since the above argument holds for any $a' \in A$ and for $b' \in \mathrm{proj}_B(A)$, we have that $\widehat{d}_k(b, b') \leq 4C$. Using Lemma 3.8 (see Figure 7), we get that

$$\widehat{\mathrm{diam}}(\mathrm{proj}_B(Y) \cup \mathrm{proj}_B(A)) \leq 4C + 4C = 8C < \xi. \qquad \square$$

Lemma 3.12 *The function d_Y^π defined by (3) satisfies condition* (A4) *in Definition 3.3, for $\xi > 14C$, where C is the constant from Lemma 2.4*

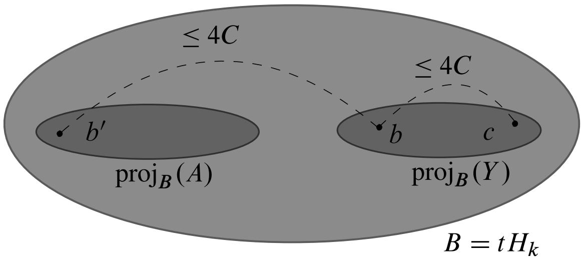

Figure 7: Estimating the distance between arbitrary points b' and c of $\mathrm{proj}_B(A)$ and $\mathrm{proj}_B(Y)$

Proof If $d^{\pi}_Y(A, B) \geq \xi$, then by (2),

$$\widehat{\mathrm{diam}}(\mathrm{proj}_Y(A) \cup \mathrm{proj}_Y(B)) \geq d^{\pi}_Y(A, B) \geq \xi.$$

Thus it suffices to prove that the number of elements $Y \in \mathbb{Y}$ satisfying

$$\widehat{\mathrm{diam}}(\mathrm{proj}_Y(A) \cup \mathrm{proj}_Y(B)) \geq \xi \tag{6}$$

is finite. Let $A, B \in \mathbb{Y}$, $A = fH_j$ and $B = tH_k$. Let $Y \in \mathbb{Y} \setminus \{A, B\}$, $Y = gH_i$. Let $a' \in A$, $b' \in \mathrm{proj}_B(a')$. By repeating the computations in Lemma 3.11, we can show that if Y is such that $\widehat{\mathrm{diam}}(\mathrm{proj}_Y(A) \cup \mathrm{proj}_Y(B)) \geq \xi$, then for any $a \in A$, $b \in B$, $x \in \mathrm{proj}_Y(a)$, $y \in \mathrm{proj}_Y(b)$, we have that $\widehat{d_i}(x, y) > 6C$.

Let h_1 denote the edge connecting x, y, which is labeled by an element of H_i (see Figure 8). Let h_2 denote the edge connecting a, a', which is labeled by an element of H_j, and let h_3 denote the edge connecting b, b', which is labeled by an element of H_k. Let p be a geodesic between a, x, let q be a geodesic between b, y and let r be a geodesic between a', b'. As argued in Lemma 3.11, we can show that h_1 cannot be isolated in the hexagon W with sides p, h_1, q, h_2, r, h_3 and must be connected to an H_i–component of r, say the edge h'.

We claim that the edge h' uniquely identifies Y. Indeed, let Y' be a member of $\mathbb{Y}$, with elements x', y' connected by an edge e (labeled by an element of the corresponding subgroup). Suppose that e is connected to h'. Then we must have that Y' is also a coset of H_i. But cosets of a subgroup are either disjoint or equal, so $Y = Y'$. Thus, the number of $Y \in \mathbb{Y}$ satisfying (6) is bounded by the number of distinct H_i–components of r, which is finite. □

3.3 Choosing a relative generating set

We now have the necessary details to choose a relative generating set X which will satisfy conditions (a) and (b) of Proposition 3.2. This set will later be altered slightly

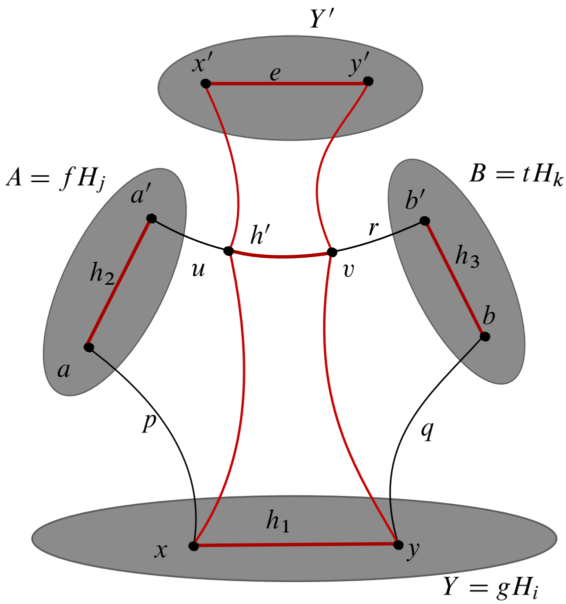

Figure 8: Condition (A4)

to obtain another relative generating set which will satisfy all three conditions of Proposition 3.2. We will repeat arguments similar to those made by Dahmani, Guirardel and Osin in [5, pages 60–63].

Recall that $\mathcal{H} = \bigsqcup_{i=1}^{n} H_i$ and that Z is the relative generating set for this collection of subgroups such that $\{H_1, H_2, \ldots, H_n\} \hookrightarrow_h (G, Z)$. Let $P_J(\mathbb{Y})$ be the projection complex corresponding to the vertex set $\mathbb{Y}$ as specified in Section 3.2, where the constant J is as in Proposition 3.4, ie $P_J(\mathbb{Y})$ is connected and a quasi-tree. Let d_P denote the combinatorial metric on $P_J(\mathbb{Y})$. Our definition of projections is K–equivariant and hence the action of K on $\mathbb{Y}$ extends to a cobounded action of K on $P_J(\mathbb{Y})$.

In what follows, by considering H_i to be vertices of the projection complex $P_J(\mathbb{Y})$, we denote by $\operatorname{star}(H_i)$ the set

$$\bigl\{e \text{ is an edge in } P_J(\mathbb{Y}) \mid e \text{ connects } H_i \text{ to } kH_j \text{ for some } k \in K \text{ and } 1 \le j \le n\bigr\}.$$

We choose the set X in the following manner. For all $i = 1, 2, \ldots, n$ and each edge e in $\operatorname{star}(H_i)$ in $P_J(\mathbb{Y})$ that connects H_i to kH_j, choose all elements $x_e \in H_i k H_j$ such that

$$\mathrm{d}_{Z \sqcup \mathcal{H}}(1, x_e) = \mathrm{d}_{Z \sqcup \mathcal{H}}(1, H_i k H_j).$$

We say that all such x_e have *type* (i, j). Since $H_i \le K$ for all i, $x_e \in K$. We observe the following:

(a) For each x_e of type (i,j) as above, there is an edge in $P_J(\mathbb{Y})$ connecting H_i and $x_e H_j$. Indeed if $x_e = h_1 k h_2$ for $h_1 \in H_i$ and $h_2 \in H_j$, then

$$\begin{aligned} d_P(H_i, x_e H_j) &= d_P(H_i, h_1 k h_2 H_j) = d_P(H_i, h_1 k H_j) \\ &= d_P(h_1^{-1} H_i, k H_j) = d_P(H_i, k H_j) \\ &= 1. \end{aligned}$$

(b) For each edge e connecting H_i and kH_j, there is a dual edge f connecting H_j and $k^{-1}H_i$. Thus for every element x_e of type (i,j), there is an element $x_f = (x_e)^{-1}$ of type (j,i). In particular, the set given by

$$X = \{x_e \neq 1 \mid e \in \mathrm{star}(H_i), i = 1, 2, \ldots, n\} \tag{7}$$

is symmetric, ie closed under taking inverses. Obviously, $X \subset K$.

(c) If $x_e \in X$ is of type (i,j), then x_e is not an element of H_i or H_j. Indeed, if $x_e = h_1 k h_2$ for some $h_1 \in H_i$ and some $h_2 \in H_j$ and x_e is an element of H_i or H_j, then $k = hf$ for some $h \in H_i$ and some $f \in H_j$. Consequently

$$\mathrm{d}_{Z \sqcup \mathcal{H}}(1, H_i k H_j) = \mathrm{d}_{Z \sqcup \mathcal{H}}(1, H_i H_j) = 0 = \mathrm{d}_{Z \sqcup \mathcal{H}}(1, x_e),$$

which implies $x_e = 1$, which is a contradiction to (7).

Lemma 3.13 (cf [5, Lemma 4.49]) *The subgroup K is generated by X together with the union of all the H_i. Further, the Cayley graph $\Gamma(K, X \sqcup \mathcal{H})$ is quasi-isometric to $P_J(\mathbb{Y})$, and hence a quasi-tree.*

Proof Let $\Sigma = \{H_1, H_2, \ldots, H_n\} \subseteq \mathbb{Y}$. Let $\mathrm{diam}(\Sigma)$ denote the diameter of the set Σ in the combinatorial metric d_P. Since Σ is a finite set, $\mathrm{diam}(\Sigma)$ is finite. Define

$$\phi\colon K \to \mathbb{Y}, \quad \phi(k) = kH_1.$$

By property (a) above, if $x_e \in X$ is of type (i,j),

$$\begin{aligned} d_P(x_e H_1, H_1) &\le d_P(x_e H_1, x_e H_j) + d_P(x_e H_j, H_i) + d_P(H_i, H_1) \\ &= d_P(H_1, H_j) + 1 + d_P(H_i, H_1) \\ &\le 2\,\mathrm{diam}(\Sigma) + 1. \end{aligned}$$

Further, for $h \in H_i$,

$$\begin{aligned} d_P(hH_1, H_1) &\le d_P(hH_1, hH_i) + d_P(hH_i, H_1) \\ &= d_P(H_1, H_i) + d_P(H_i, H_1) \\ &\le 2\,\mathrm{diam}(\Sigma). \end{aligned}$$

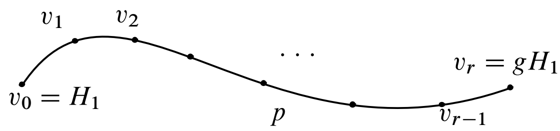

Figure 9: The geodesic p

Thus, for all $g \in \langle X \cup H_1 \cup H_2 \cup \cdots \cup H_n \rangle$, we have

$$d_P(\phi(1), \phi(g)) \leq (2\,\mathrm{diam}(\Sigma) + 1)|g|_{X \sqcup \mathcal{H}}, \tag{8}$$

where $|g|_{X \sqcup \mathcal{H}}$ denotes the length of g in the generating set $X \cup H_1 \cup H_2 \cup \cdots \cup H_n$. (We use this notation for the sake of uniformity).

Now let $g \in K$ and suppose $d_P(\phi(1), \phi(g)) = r$, ie $d_P(H_1, gH_1) = r$. If $r = 0$, then $H_1 = gH_1$, thus $g \in H_1$ and $|g|_{X \sqcup \mathcal{H}} \leq 1$. If $r > 0$, consider a geodesic p in $P_J(\mathbb{Y})$ connecting H_1 and gH_1. Let

$$\begin{aligned} v_0 &= H_1 = g_0 H_1 \quad (g_0 = 1), \\ v_1 &= g_1 H_{\lambda_1}, \\ v_2 &= g_2 H_{\lambda_2}, \\ &\;\;\vdots \\ v_{r-1} &= g_{r-1} H_{\lambda_{r-1}}, \\ v_r &= gH_1 \qquad\qquad (g_r = g) \end{aligned}$$

be the sequence of vertices of p, for some $\lambda_j \in \{1, 2, \ldots, n\}$ and some $g_i \in K$ (see Figure 9).

Now $g_i H_{\lambda_i}$ is connected by a single edge to $g_{i+1} H_{\lambda_{i+1}}$. Thus

$$d_P(g_i H_{\lambda_i}, g_{i+1} H_{\lambda_{i+1}}) = 1,$$

which implies

$$d_P(H_{\lambda_i}, g_i^{-1} g_{i+1} H_{\lambda_{i+1}}) = 1.$$

Then there exists $x \in X$ such that

$$x \in H_{\lambda_i} g_i^{-1} g_{i+1} H_{\lambda_{i+1}} \quad \text{and} \quad \mathrm{d}_{Z \sqcup \mathcal{H}}(1, x) = \mathrm{d}_{Z \sqcup \mathcal{H}}(1, H_{\lambda_i} g_i^{-1} g_{i+1} H_{\lambda_{i+1}}).$$

Thus $x = h g_i^{-1} g_{i+1} k$ for some $h \in H_{\lambda_i}$ and some $k \in H_{\lambda_{i+1}}$ which implies $g_i^{-1} g_{i+1} = h^{-1} x k^{-1}$. So $|g_i^{-1} g_{i+1}|_{X \sqcup \mathcal{H}} \leq 3$, which implies

$$|g|_{X \sqcup \mathcal{H}} = \left| \prod_{i=1}^{r} g_{i-1}^{-1} g_i \right|_{X \sqcup \mathcal{H}} \leq \sum_{i=1}^{r} |g_{i-1}^{-1} g_i|_{X \sqcup \mathcal{H}} \leq 3r = 3 d_P(\phi(1), \phi(g)). \tag{9}$$

The above argument also provides a representation for every element $g \in K$ as a product of elements from $X \cup H_1 \cup H_2 \cup \cdots \cup H_n$. Thus K is generated by the union of X and all the H_i. By (8) and (9), ϕ is a quasi-isometric embedding of $(K, |\cdot|_{X \sqcup \mathcal{H}})$ into $(P_J(\mathbb{Y}), d_P)$ satisfying

$$\tfrac{1}{3}|g|_{X\sqcup\mathcal{H}} \le d_P(\phi(1), \phi(g)) \le (2\operatorname{diam}(\Sigma)+1)|g|_{X\sqcup\mathcal{H}}.$$

Since $\mathbb{Y}$ is contained in the closed $\operatorname{diam}(\Sigma)$–neighborhood of $\phi(K)$, ϕ is a quasi-isometry. This implies that $\Gamma(K, X \sqcup \mathcal{H})$ is a quasi-tree. □

Let $\widetilde{d}_i$ denote the modified relative metric on H_i associated with the Cayley graph $\Gamma(G, Z \sqcup \mathcal{H})$ from Theorem 2.8. Let $\widehat{d_i^X}$ denote the relative metric on H_i associated with the Cayley graph $\Gamma(K, X \sqcup \mathcal{H})$. We will now show that $\widehat{d_i^X}$ is proper for all $i = 1, 2, \ldots, n$. We will use the fact that $\widetilde{d}_i$ is proper and derive a relation between $\widetilde{d}_i$ and $\widehat{d_i^X}$.

Lemma 3.14 (cf [5, Lemma 4.50]) *There exists a constant α such that for any $Y \in \mathbb{Y}$ and any $x \in X \sqcup \mathcal{H}$, if*

$$\widetilde{\operatorname{diam}}(\operatorname{proj}_Y\{1, x\}) > \alpha,$$

then $x \in H_j$ and $Y = H_j$ for some j.

Proof We prove the result for

$$\alpha = \max\{J + 2\xi, 6C\}.$$

Suppose that $\widetilde{\operatorname{diam}}(\operatorname{proj}_Y\{1, x\}) > \alpha$ and $x \in X$ has type (k, l), ie there exists an edge connecting H_k and gH_l in $P_J(\mathbb{Y})$, where $g \in K$. We consider three possible cases and arrive at a contradiction in each case.

Case 1 ($H_k \ne Y \ne xH_l$) Then

$$\widetilde{\operatorname{diam}}(\operatorname{proj}_Y\{1, x\}) \le d_Y^{\pi}(H_k, xH_l) \le d_Y(H_k, xH_l) + 2\xi \le J + 2\xi \le \alpha,$$

using (1) and the fact that H_k and xH_l are connected by an edge in $P_J(\mathbb{Y})$, which is a contradiction.

Case 2 ($H_k = Y$) Since $x \notin H_k$, let $y \in \operatorname{proj}_Y(x)$, ie

$$\mathrm{d}_{Z\sqcup\mathcal{H}}(x, y) = \mathrm{d}_{Z\sqcup\mathcal{H}}(x, H_k) = \mathrm{d}_{Z\sqcup\mathcal{H}}(x, Y).$$

By Lemma 3.6, if $\widehat{d}_k(1,y) \le 3C$, then

$$\begin{aligned}\widehat{\mathrm{diam}}(\mathrm{proj}_Y\{1,x\}) &\le \widehat{\mathrm{diam}}(\mathrm{proj}_Y(1)) + \widehat{\mathrm{diam}}(\mathrm{proj}_Y(x)) + \widehat{d}_k(\mathrm{proj}_Y(1),\mathrm{proj}_Y(x))\\ &\le 0 + 3C + \widehat{d}_k(1,y)\\ &\le 6C \le \alpha.\end{aligned}$$

Then by (2), we have

$$\widetilde{\mathrm{diam}}(\mathrm{proj}_Y\{1,x\}) \le \alpha,$$

which is a contradiction. Thus $\widehat{d}_k(1,y) > 3C$. This implies that $1 \notin \mathrm{proj}_Y(x)$ (see Figure 10). By definition of the nearest point projection, $\mathrm{d}_{Z\sqcup\mathcal{H}}(1,x) > \mathrm{d}_{Z\sqcup\mathcal{H}}(y,x)$, which implies $\mathrm{d}_{Z\sqcup\mathcal{H}}(1,x) > \mathrm{d}_{Z\sqcup\mathcal{H}}(1,y^{-1}x)$. Since $y^{-1}x \in H_k g H_l$, we obtain $\mathrm{d}_{Z\sqcup\mathcal{H}}(1,x) > \mathrm{d}_{Z\sqcup\mathcal{H}}(1,H_k g H_l)$, which is a contradiction to the choice of x.

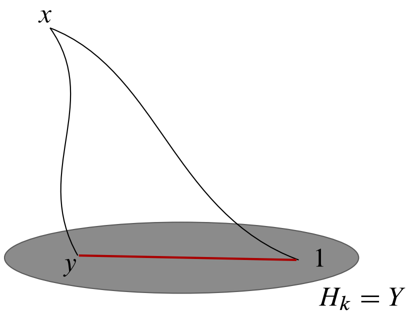

Figure 10: Case 2

Case 3 $(Y = xH_l,\ H_k \ne Y)$ This case reduces to Case 2, since we can translate everything by x^{-1}.

Thus we must have $x \in H_j$ for some j. Suppose that $H_j \ne Y$. But then

$$\widetilde{\mathrm{diam}}(\mathrm{proj}_Y\{1,x\}) \le \widetilde{\mathrm{diam}}(\mathrm{proj}_Y(H_j)) \le 4C \le \alpha,$$

by Lemma 3.8, which is a contradiction. □

Lemma 3.15 (cf [5, Lemma 4.45]) *If $H_i = fH_j$, then $H_i = H_j$ and $f \in H_i$. Consequently, if $gH_i = fH_j$, then $H_i = H_j$ and $g^{-1}f \in H_i$.*

Proof If $H_i = fH_j$, then $1 = fk$ for some $k \in H_j$. Then $f = k^{-1} \in H_j$, which implies $H_i = H_j$. □

Lemma 3.16 (cf [5, Theorem 4.42]) *For $i = 1,2,\dots,n$ and any $h \in H_i$, we have*

$$\alpha \widehat{d_i^X}(1,h) \ge \widetilde{d}_i(1,h),$$

where α is the constant from Lemma 3.14. Thus $\widehat{d_i^X}$ is proper.

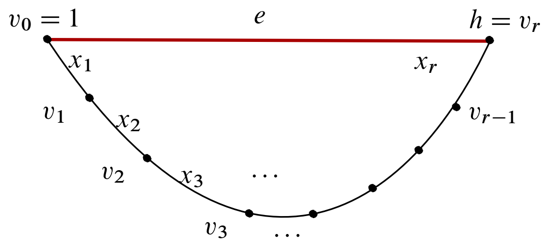

Figure 11: The cycle ep

Proof Let $h \in H_i$ be such that $\widehat{d}_i^X(1,h) = r$. Let e denote the H_i–edge in the Cayley graph $\Gamma(K, X \sqcup \mathcal{H})$ connecting h to 1, labeled by h^{-1}. Let p be an admissible (see Definition 2.1) path of length r in $\Gamma(K, X \sqcup \mathcal{H})$ connecting 1 and h. Then ep forms a cycle (see Figure 11). Since p is admissible, e is isolated in this cycle.

Let $\mathrm{Lab}(p) = x_1 x_2 \cdots x_r$ for some $x_1, x_2, \dots, x_r \in X \sqcup \mathcal{H}$. Let

$$v_0 = 1, \quad v_1 = x_1, \quad v_2 = x_1 x_2, \quad \dots, \quad v_r = x_1 x_2 \cdots x_r = h.$$

Since these are also elements of G, for $k = 1, 2, \dots, r$ we have

$$\begin{aligned}\widetilde{\mathrm{diam}}(\mathrm{proj}_{H_i}\{v_{k-1}, v_k\}) &= \widetilde{\mathrm{diam}}(\mathrm{proj}_{H_i}\{x_1 x_2 \cdots x_{k-1}, x_1 x_2 \cdots x_{k-1} x_k\}) \\ &= \widetilde{\mathrm{diam}}(\mathrm{proj}_Y\{1, x_k\}),\end{aligned}$$

where $Y = (x_1 x_2 \cdots x_{k-1})^{-1} H_i$.

If $\widetilde{\mathrm{diam}}(\mathrm{proj}_Y\{1, x_k\}) > \alpha$ for some k, then by Lemma 3.14, $x_k \in H_j$ and $Y = H_j$ for some j. By Lemma 3.15, $H_i = H_j$ and $x_1 x_2 \cdots x_{k-1} \in H_j$. But then e is not isolated in the cycle ep, which is a contradiction.

Hence

$$\widetilde{\mathrm{diam}}(\mathrm{proj}_{H_i}\{v_{k-1}, v_k\}) \le \alpha$$

for all $k = 1, 2, \dots, r$, which implies

$$\begin{aligned}\widetilde{d}_i(1,h) &\le \widetilde{\mathrm{diam}}(\mathrm{proj}_{H_i}\{v_0, v_r\}) \\ &\le \sum_{j=1}^{r} \widetilde{\mathrm{diam}}(\mathrm{proj}_{H_i}\{v_{j-1}, v_j\}) \\ &\le r\alpha = \alpha \widehat{d}_i^X(1,h).\end{aligned}$$ □

3.4 Proof of Proposition 3.2

The goal of this section is to alter our relative generating set X from Section 3.3, so that we obtain another relative generating set that satisfies all the conditions of

Proposition 3.2. To do so, we need to establish a relation between the set X and the set Z. We will need the following obvious lemma.

Lemma 3.17 *Let X and Y be generating sets of G, and suppose that*

$$\sup_{x\in X} |x|_Y < \infty \quad \text{and} \quad \sup_{y\in Y} |y|_X < \infty.$$

Then $\Gamma(G,X)$ is quasi-isometric to $\Gamma(G,Y)$. In particular, $\Gamma(G,X)$ is a quasi-tree if and only if $\Gamma(G,Y)$ is a quasi-tree.

Remark 3.18 This lemma implies that if we change a generating set by adding finitely many elements, then the property that the Cayley graph is a quasi-tree still holds.

We also need to note that from (1) in Definition 3.3, it easily follows that for each $Y \in \mathbb{Y}$ and every $A, B \in \mathbb{Y}\backslash\{Y\}$, we have

$$d_Y(A,B) \le d_Y^{\pi}(A,B) + 2\xi. \tag{10}$$

Lemma 3.19 *For a large enough J, the set X constructed in Section 3.3 satisfies the following property: if $z \in Z \cap K$ does not represent an element of H_i for any $i = 1, 2, \dots, n$, then $z \in X$.*

Proof Recall that $\mathrm{d}_{Z\sqcup\mathcal{H}}$ denotes the combinatorial metric on $\Gamma(G, Z \sqcup \mathcal{H})$. Let $z \in Z \cap K$ be as in the statement of the lemma. Then $z \in H_i z H_i$ for all i and $1 \notin H_i z H_i$. Thus

$$\mathrm{d}_{Z\sqcup\mathcal{H}}(1, H_i z H_i) \ge 1 = \mathrm{d}_{Z\sqcup\mathcal{H}}(1, z) \ge \mathrm{d}_{Z\sqcup\mathcal{H}}(1, H_i z H_i),$$

which implies

$$\mathrm{d}_{Z\sqcup\mathcal{H}}(1, H_i z H_i) = \mathrm{d}_{Z\sqcup\mathcal{H}}(1, z) \quad \text{for all } i.$$

In order to prove $z \in X$, we must show that H_i and zH_i are connected by an edge in $P_J(\mathbb{Y})$. By Definition 3.3, this is true if

$$d_Y(H_i, zH_i) \le J \quad \text{for all } Y \ne H_i, zH_i.$$

In view of (10), we will estimate $d_Y^{\pi}(H_i, zH_i)$.

Let $\mathrm{d}_{Z\sqcup\mathcal{H}}(h,x) = \mathrm{d}_{Z\sqcup\mathcal{H}}(H_i, Y)$ and $\mathrm{d}_{Z\sqcup\mathcal{H}}(f,y) = \mathrm{d}_{Z\sqcup\mathcal{H}}(zH_i, Y)$ for some $h \in H_i$, $f \in zH_i$ and for some $x, y \in Y = gH_j$. Let p be a geodesic connecting h and x, and let q be a geodesic connecting y and f. Let h_2 denote the edge connecting x and y, labeled by an element of H_j. Similarly, let s, t denote the edges connecting h to 1 and z to f respectively, which are labeled by elements of H_i. Let e denote the

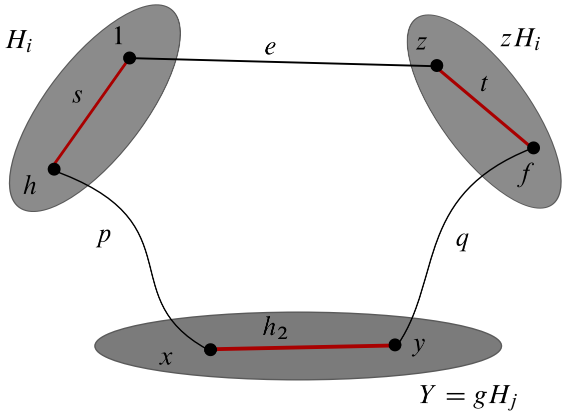

Figure 12: Dealing with elements of $Z \cap K$ that represent elements of $\mathcal{H}$

edge connecting 1 and z, labeled by z. Consider the geodesic hexagon W with sides p, h_2, q, t, e, s (see Figure 12).

By using Remark 3.7 and the fact that $Y \neq H_i, zH_i$, we can show that h_2 cannot be connected to q, p, s or t. Since z does not represent an element of H_i for any i, h_2 cannot be connected to e. Thus, h_2 is isolated in W. By Lemma 2.4, $\widehat{d}_j(x, y) \leq 6C$. By Lemma 3.8,

$$d_Y(H_i, zH_i) \leq d_Y^{\pi}(H_i, zH_i) + 2\xi \leq 14C + 2\xi.$$

We conclude that by taking the constant J to be sufficiently large so that Proposition 3.4 holds and J exceeds $14C + 2\xi$, we can ensure that $z \in X$ and the arguments of the previous section still hold. □

Lemma 3.20 *There are only finitely many elements of $Z \cap K$ that can represent an element of H_i for some $i \in \{1, 2, \ldots, n\}$.*

Proof Let $z \in Z \cap K$ represent an element of H_i for some $i = 1, 2, \ldots, n$. Then in the Cayley graph $\Gamma(G, Z \sqcup \mathcal{H})$, we have a bigon between the elements 1 and h, where one edge is labeled by z and the other edge is labeled by an element of H_i, say h_1 (see Remark 2.2 and Figure 13).

This implies that $\widehat{d}_i(1, z) \leq 1$, so $\widetilde{d}_i(1, z) \leq 1$. But then $z \in \widetilde{B}_i(1, 1)$, ie the ball of radius 1 in the subgroup H_i in the relative metric, centered at the identity. But this is a finite ball. Take

$$\rho = \left| \bigcup_{i=1}^{n} \widetilde{B}_i(1, 1) \right|.$$

Then z has at most ρ choices, which is finite. □

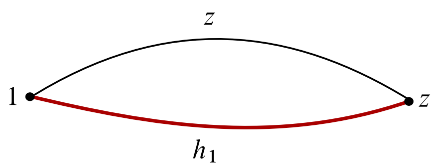

Figure 13: Bigons in the Cayley graph

By Lemma 3.20 and by selecting the constant J as specified in Lemma 3.19, we conclude that the set X from Section 3.3 does not contain at most finitely many elements of $Z \cap K$. By adding these finitely many remaining elements of $Z \cap K$ to X, we obtain a new relative generating set X' such that $|X' \Delta X| < \infty$. By Lemma 2.11, $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h (K, X')$ and $Z \cap K \subset X'$. By Remark 3.18, $\Gamma(K, X' \sqcup \mathcal{H})$ is also a quasi-tree. Thus X' is the required set in the statement of Proposition 3.2, which completes the proof.

3.5 Applications of Theorem 3.1

In order to prove Theorem 1.7, we first need to recall the following definitions.

Definition 3.21 (loxodromic element) Let G be a group acting on a hyperbolic space S. An element $g \in G$ is called *loxodromic* if the map $\mathbb{Z} \to S$ defined by $n \mapsto g^n s$ is a quasi-isometric embedding for some (equivalently, any) $s \in S$.

Definition 3.22 (elementary subgroup [5, Lemma 6.5]) Let G be a group acting acylindrically on a hyperbolic space S and $g \in G$ a loxodromic element. Then g is contained in a unique maximal *elementary subgroup* $E(g)$ of G given by

$$E(g) = \{h \in G \mid d_{\mathrm{Hau}}(l, h(l)) < \infty\},$$

where d_{Hau} denotes the Hausdorff distance and l is a quasi-geodesic axis of g in S.

Corollary 3.23 *A group G is acylindrically hyperbolic if and only if G has an acylindrical and non-elementary action on a quasi-tree.*

Proof If G has an acylindrical and non-elementary action on a quasi-tree, Theorem 2.9 implies that G is acylindrically hyperbolic. Conversely, let G be acylindrically hyperbolic, with an acylindrical non-elementary action on a hyperbolic space X. Let g be a loxodromic element for this action. By Lemma 6.5 of [5] the elementary subgroup $E(g)$ is virtually cyclic and thus countable. By Theorem 6.8 of [5], $E(g)$ is hyperbolically embedded in G. Taking $K = G$ and $E(g)$ to be the hyperbolically embedded subgroup in the statement of Theorem 3.1 now gives us the result. Since $E(g)$ is non-degenerate, by Lemma 5.12 of [12], the resulting action of G on the associated Cayley graph $\Gamma(G, X \sqcup E(g))$ is also non-elementary. □

The following corollary is an immediate consequence of Theorem 3.1.

Corollary 3.24 *Let $\{H_1, H_2, \dots, H_n\}$ be a finite collection of countable subgroups of a group G such that $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h G$. Let K be a subgroup of G. If $H_i \leq K$ for all $i = 1, 2, \dots, n$, then $\{H_1, H_2, \dots, H_n\} \hookrightarrow_h K$.*

Definition 3.25 Let (M, d) be a geodesic metric space, and $\epsilon > 0$ a fixed constant. A subset $S \subset M$ is said to be *ϵ–coarsely connected* if for any two points x, y in S, there exist points $x_0 = x, x_1, x_2, \dots, x_{n-1}, x_n = y$ in S such that for all $i = 0, \dots, n-1$,

$$d(x_i, x_{i+1}) \leq \epsilon.$$

We say that S is *coarsely connected* if it is ϵ–coarsely connected for some $\epsilon > 0$.

Recall that we denote the closed σ–neighborhood of S by $S^{+\sigma}$.

Definition 3.26 Let (M, d) be a geodesic metric space, and $\sigma > 0$ a fixed constant. A subset $S \subset M$ is said to be *σ–quasi-convex* if for any two points x, y in S, any geodesic connecting x and y is contained in $S^{+\sigma}$. Further, we say that S is *quasi-convex* if it is σ–quasi-convex for some $\sigma > 0$.

Corollary 3.27 *Let H be a finitely generated subgroup of an acylindrically hyperbolic group G. Then there exists a subset $X \subset G$ such that*

(a) *$\Gamma(G, X)$ is hyperbolic, and the action of G on $\Gamma(G, X)$ is non-elementary and acylindrical, and*

(b) *H is quasi-convex in $\Gamma(G, X)$.*

To prove this corollary, we need the following two lemmas.

Lemma 3.28 *Let T be a tree, and let $Q \subset T$ be ϵ–coarsely connected. Then Q is ϵ–quasi-convex.*

Proof Let $\epsilon > 0$ be the constant from Definition 3.25. Let x, y be two points in Q and p be any geodesic between them. Then there exist points $x_0 = x, x_1, x_2, \dots, x_{n-1}, x_n = y$ in Q such that $d(x_i, x_{i+1}) \leq \epsilon$ for all $i = 0, \dots, n-1$. Let p_i denote the geodesic segments between x_i and x_{i+1} for $i = 0, 1, \dots, n-1$. Since T is a tree, we must have that

$$p \subseteq \bigcup_{i=0}^{n-1} p_i.$$

By definition, $p_i \subseteq B(x_i, \epsilon)$, the ball of radius ϵ centered at x_i, for $i = 0, 1, \dots, n-1$. Since $x_i \in Q$ for $i = 0, 1, \dots, n-1$, we obtain

$$p_i \subseteq Q^{+\epsilon}.$$

This implies $p \subseteq Q^{+\epsilon}$. □

Lemma 3.29 *Let Γ be a quasi-tree, and $S \subset \Gamma$ be coarsely connected. Then S is quasi-convex.*

Proof Let T be a tree such that Γ is quasi-isometric to T. Let d_Γ and d_T denote distances in Γ and T, respectively. Let $\delta > 0$ be the hyperbolicity constant of Γ. Let $q\colon T \to \Gamma$ be a (λ, C)–quasi-isometry, ie

$$-C + \frac{1}{\lambda} d_T(a,b) \le d_\Gamma(q(a), q(b)) \le \lambda d_T(a,b) + C.$$

Let $\epsilon > 0$ be the constant from Definition 3.25 for S. Set $Q = q^{-1}(S)$. Then $Q \subset T$. It is easy to check that Q is ρ–coarsely connected with constant $\rho = \lambda(\epsilon + C)$. By Lemma 3.28, Q is ρ–quasi-convex.

Let x, y be two points in S and p be a geodesic between them. Choose points a, b in Q such that $q(a) = x$ and $q(b) = y$. Let r denote the (unique) geodesic in T between a and b. Since Q is ρ–quasi-convex, we have

$$r \subseteq Q^{+\rho}.$$

Set $\sigma = \lambda\rho + C$. Then

$$q(r) \subseteq S^{+\sigma}.$$

Further, $q \circ r$ is a quasi-geodesic between x and y. By Lemma 2.17, there exists a constant $R = R(\lambda, C, \delta)$ such that $q(r)$ and p are Hausdorff distance less than R from each other. This implies that $p \subseteq S^{+(R+\sigma)}$. Thus S is quasi-convex. □

Proof of Corollary 3.27 By Corollary 3.23, there exists a generating set X of G such that $\Gamma(G, X)$ is a quasi-tree (hence hyperbolic) and the action of G on $\Gamma(G, X)$ is acylindrical and non-elementary. Let d_X denote the metric on $\Gamma(G, X)$ induced by the generating set X. Let $H = \langle x_1, x_2, \dots, x_n \rangle$. Set

$$\epsilon = \max\{d_X(1, x_i^{\pm 1}) \mid i = 1, 2, \dots, n\}.$$

We claim that H is coarsely connected with constant ϵ. Indeed if u, v are elements of H, then $u^{-1}v = \prod_{j=1}^k w_j$, where $w_j \in \{x_1^{\pm 1}, \dots, x_n^{\pm 1}\}$. Set

$$z_0 = u, \quad z_1 = uw_1, \quad \dots, \quad z_{k-1} = uw_1w_2\cdots w_{k-1}, \quad z_k = v.$$

Clearly $z_i \in H$ for all $i = 0, 1, \ldots, k-1$. Further,

$$d_X(z_i, z_{i+1}) = d_X(1, w_{i+1}) \leq \epsilon$$

for all $i = 0, 1, \ldots, k-1$. By Lemma 3.29, H is quasi-convex in $\Gamma(G, X)$. $\square$

References

[1] **J Behrstock**, **M F Hagen**, **A Sisto**, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, preprint (2015) arXiv

[2] **M Bestvina**, **K Bromberg**, **K Fujiwara**, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publ. Math. Inst. Hautes Études Sci. 122 (2015) 1–64 MR

[3] **B H Bowditch**, *Intersection numbers and the hyperbolicity of the curve complex*, J. Reine Angew. Math. 598 (2006) 105–129 MR

[4] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999) MR

[5] **F Dahmani**, **V Guirardel**, **D Osin**, *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*, Mem. Amer. Math. Soc. 1156, Amer. Math. Soc., Providence, RI (2017) MR

[6] **M F Hagen**, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, J. Topol. 7 (2014) 385–418 MR

[7] **I Kapovich**, **K Rafi**, *On hyperbolicity of free splitting and free factor complexes*, Groups Geom. Dyn. 8 (2014) 391–414 MR

[8] **S-H Kim**, **T Koberda**, *The geometry of the curve graph of a right-angled Artin group*, Internat. J. Algebra Comput. 24 (2014) 121–169 MR

[9] **J F Manning**, *Geometry of pseudocharacters*, Geom. Topol. 9 (2005) 1147–1185 MR

[10] **A Minasyan**, **D Osin**, *Small subgroups of acylindrically hyperbolic groups*, preprint (2016)

[11] **D Osin**, *On acylindrical hyperbolicity of groups with positive first ℓ^2–Betti number*, Bull. Lond. Math. Soc. 47 (2015) 725–730 MR

[12] **D Osin**, *Acylindrically hyperbolic groups*, Trans. Amer. Math. Soc. 368 (2016) 851–888 MR

[13] **E Rips**, *Subgroups of small cancellation groups*, Bull. London Math. Soc. 14 (1982) 45–47 MR

[14] **A Sisto**, *On metric relative hyperbolicity*, preprint (2012) arXiv

Department of Mathematics, Vanderbilt University
1326 Stevenson Center Ln, Nashville, TN, 37240, United States

sahana.balasubramanya@vanderbilt.edu

Received: 4 March 2016 Revised: 19 October 2016