

Limit groups for relatively hyperbolic groups. I. The basic tools

DANIEL GROVES

We begin the investigation of Γ -limit groups, where Γ is a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. Using the results of Druţu and Sapir [17], we adapt the results from the author's paper [24]. Specifically, given a finitely generated group G and a sequence of pairwise nonconjugate homomorphisms $\{h_n: G \rightarrow \Gamma\}$, we extract an \mathbb{R} -tree with a nontrivial isometric G -action.

We then provide an analogue of Sela's shortening argument.

20F65; 20F67, 20E08, 57M07

In his remarkable series of papers [40; 42; 44; 43; 45; 46; 47], Z Sela has classified those finitely generated groups with the same *elementary theory* as the free group of rank 2 (see also Sela [41] for a summary). This class includes all nonabelian free groups, most surface groups and certain other hyperbolic groups. In particular, Sela answers in the positive some long-standing questions of Tarski (Kharlampovich and Myasnikov [31] have another approach to these problems).

In [40], Sela begins with a study of *limit groups*. Sela's definition of a limit group is geometric, though it turns out that a group is a limit group if and only if it is a finitely generated fully-residually free group. He then produces *Makanin–Razborov diagrams*, which give a parametrisation of $\text{Hom}(G, \mathbb{F})$, where G is an arbitrary finitely generated group and \mathbb{F} is a nonabelian free group (such a parametrisation is also given by Kharlampovich and Myasnikov [30]). Over the course of his six papers, two of the main tools Sela uses are the theory of isometric actions on \mathbb{R} -trees and the *shortening argument*.

Sela's work naturally raises the question of which other classes of groups can be understood using this approach. Many of Sela's methods (and, more strikingly, some of the answers) come from geometric group theory. Thus, when looking for classes of groups to apply these methods to, it seems natural to consider groups of interest in geometric group theory. In [48], Sela considers an arbitrary torsion-free hyperbolic group Γ and characterises those groups with the same elementary theory as Γ . Of particular note is Sela's result that any group which has the same elementary theory as

a torsion-free hyperbolic group is itself a torsion-free hyperbolic group. This result exhibits a deep connection between the logic of groups and geometric group theory. In [2], Alibegović constructs Makanin–Razborov diagrams for limit groups. In [24] the author began this study for certain torsion-free $\text{CAT}(0)$ groups.

This paper serves two purposes. First, we generalise the results of [24] to the context of torsion-free groups which are hyperbolic relative to a collection of abelian subgroups¹. We construct a space closely related to the Cayley graph (see Section 4) and use the results of Druţu and Sapir from [17] to analyse an asymptotic cone of this space. We then follow [24] to extract an \mathbb{R} -tree from this asymptotic cone. Armed with this \mathbb{R} -tree, we then develop an analogue of Sela’s shortening argument in this context.

The following result is a straightforward application of the shortening argument (see Section 7 for a definition of $\text{Mod}(\Gamma)$).

Theorem 0.1 *Suppose that Γ is a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. Then $\text{Mod}(\Gamma)$ has finite index in $\text{Aut}(\Gamma)$.*

The true context of this paper is as the beginning of the study of Γ -limit groups where Γ is torsion-free and hyperbolic relative to free abelian subgroups. In a continuation paper [23] we use the results of this paper and of [48] to construct *Makanin–Razborov diagrams* for such a group Γ . It is our hope that much, possibly all, of Sela’s program can be carried out for these groups.

The outline of this paper is as follows: In Section 1 we recall the concepts of *limit groups* and Γ -*limit groups*. In Section 2 we recall the definition of relatively hyperbolic groups and the basic results required for this paper. In Section 3 we recall the concept of asymptotic cones and some results of Druţu and Sapir from [17]. In Section 4 we define a space X , closely related to the Cayley graph of the relatively hyperbolic group Γ . The space X , equipped with a natural Γ -action, seems to be an appropriate space with which to study Γ -limit groups. In Section 5 a particular asymptotic cone X_ω is extracted from a sequence of nonconjugate homomorphisms $\{h_n: G \rightarrow \Gamma\}$, where G is an arbitrary finitely generated group and the limiting action of G on X_ω is studied. In Section 6 we extract an action of G on an \mathbb{R} -tree with no global fixed point. In Section 7 we state a version of Sela’s shortening argument in the context of this paper. In Section 8 we recall some of the theory of groups acting isometrically on \mathbb{R} -trees. In Sections 9–11 we present the proof of Theorem 7.5 (the shortening argument).

This paper relies heavily on the results of Druţu and Sapir [17], and also on the paper [24] of the author. We have stated the results from [17] which are required, but

¹See Section 2 below for a definition and discussion of relatively hyperbolic groups.

the reader who is not familiar with [17] may find it hard going. The construction of the space X in Section 4 is carefully designed so that many of the proofs from [24] go through without change. In particular, many of the results of the later part of Section 4 and also of Section 5 and Section 6 have proofs almost identical to results in [24], and we mostly refer the reader to that paper for details. Therefore, a full understanding of this paper will be difficult without some knowledge of [24].

Remark 0.2 This paper subsumes the results which were in the first version of this paper, and also in the author's preprint [22]. The exception to this is the proof that the groups under consideration are Hopfian. This result follows immediately from the author's paper [23, Theorem 5.2] (and the proof there does not depend on anything left out of this version of this paper). Since the proof of the Hopf property is technical, and no more enlightening than the proof of [23, Theorem 5.2], we chose to leave this result out of this paper. Leaving this proof out made it natural to take the results of [22] and the first version of this paper and merge them into this single paper.

Subsequent to the first version of this paper appearing on the ArXiv, the paper [18] of Druţu and Sapir appeared. In this, they take a finitely generated group G acting on an arbitrary tree-graded metric space and produce an isometric action of G on an \mathbb{R} -tree.

In this paper, we require that the tree-graded metric space has pieces isometric to Euclidean spaces and that the stabilizers of these pieces act by translations. However, there are substantial benefits to the construction in this paper over that of [18]. Namely, as we see in this paper and in [23], with our tree we are able to prove the appropriate analogue of Sela's shortening argument. There still does not exist a shortening argument for general actions on tree-graded spaces.

The construction in this paper has been used in [23] and by Dahmani and the author [14]. It will also be crucial in future work of the author on the elementary theory of relatively hyperbolic groups.

Acknowledgements The paper [24] was written before [17] appeared. Although I had already had a cursory glance at [17], I would like to thank both Mark Sapir and Jon McCammond for suggesting that the results in [17] may allow the results from [24] to be generalised as they have been in this paper. I would also like to thank Jason Manning and Indira Chatterji for fruitful discussions on the construction in Section 4.

The first version of this paper was written while the author was a Taussky-Todd Instructor at Caltech. The author was supported in part by the NSF. Both of these organisations are thanked for their support. Finally, thanks to the anonymous referee, whose careful reading and numerous comments have improved the exposition of this paper.

1 Limit groups and Γ -limit groups

Recall the following two definitions, due to Bestvina and Feighn [4].

Definition 1.1 [4, Definition 1.5] Let G and Ξ be finitely generated groups, and let $\{h_n: G \rightarrow \Xi\}$ be a sequence of homomorphisms. The *stable kernel* of $\{h_n\}$, denoted $\underline{\text{Ker}}(h_n)$, is the set of all elements $g \in G$ so that $g \in \ker(h_n)$ for all but finitely many n .

The sequence $\{h_n\}$ is *stable* if for all $g \in G$, either (i) $g \in \underline{\text{Ker}}(h_n)$; or (ii) $g \notin \ker(h_n)$ for all but finitely many n .

Definition 1.2 [4, Definition 1.5] A Ξ -limit group is a group of the form $G/\underline{\text{Ker}}(h_n)$ where G is a finitely generated group and $\{h_n: G \rightarrow \Xi\}$ is a stable sequence of homomorphisms.

Remark 1.3 If each of the h_n is equal to h , a single homomorphism, then the sequence $\{h_n\}$ is certainly stable and the associated Ξ -limit group is just $h(G)$. In particular, all finitely generated subgroups of Ξ are Ξ -limit groups.

A *limit group* is an \mathbb{F} -limit group, where \mathbb{F} is a finitely generated free group. This terminology is due to Sela [40], although the definition that Sela gave was in terms on an action of G on an \mathbb{R} -tree induced by the sequence $\{h_n: G \rightarrow \mathbb{F}\}$. Sela's geometric definition also makes sense for δ -hyperbolic groups (see Sela [48]). In this paper, we achieve a geometric definition of Γ -limit groups, where Γ is a torsion-free group which is hyperbolic relative to free abelian subgroups.

In case $\Xi = \mathbb{F}$, the geometric and algebraic definitions of Ξ -limit groups are the same [40, Lemma 1.3]. The two definitions are also the same when Ξ is a torsion-free δ -hyperbolic group [48, Lemma 1.3]. When Ξ is a torsion-free CAT(0) group with isolated flats whose flat stabilisers are abelian, a geometric definition of Ξ -limit group was given in [24, Definition 3.21] and it was proved [24, Theorem 5.1] that these two definitions are the same.

Suppose that Γ is a torsion-free group hyperbolic relative to free abelian subgroups. In this paper, we provide an appropriate geometric definition of Γ -limit group, in analogy with the definition from [24] (see Definition 5.8 below). It is proved in Theorem 6.7 that this definition is equivalent to Definition 1.2. As in Sela's definition, along with the geometric definition comes a faithful action of a (strict) Γ -limit group on an \mathbb{R} -tree.

The utility of the algebraic Definition 1.2 is that it has implications for the logic of Γ . In the case of Sela's limit groups, the nonabelian limit groups are exactly those that have

the same universal theory as a nonabelian free group. In general, if $T_{\forall}(H)$ denotes the *universal theory* of a group H then we have the following result. See Champetier and Guirardel [9] for a detailed discussion of this issue.

Lemma 1.4 *Let Ξ be a finitely presented group and suppose that L is a Ξ -limit group. Then $T_{\forall}(\Xi) \subseteq T_{\forall}(L)$.*

The utility of Sela’s geometric definition is that it allows the application of the (Rips) theory of isometric actions on \mathbb{R} -trees, and Sela uses this to make a very deep study of limit groups (and of Γ -limit groups, where Γ is a torsion-free hyperbolic group). It turns out that the class of limit groups is exactly the class of finitely generated fully residually free groups, which has been widely studied in the past.

2 Relatively hyperbolic groups

Recently there has been a large amount of interest in relatively hyperbolic groups. Relatively hyperbolic groups were originally defined by Gromov in his seminal paper [21], and an alternative definition was given by Farb [19]. Bowditch [6] gave two definitions, equivalent to Gromov’s and Farb’s, respectively (see Dahmani [11] for a proof of the equivalence of the definitions). Druţu and Sapir [17] gave a characterisation of relatively hyperbolic groups in terms of their asymptotic cones.²

The idea behind relatively hyperbolic groups is to generalise hyperbolic groups (those admitting a proper and cocompact action on a δ -hyperbolic space) to proper, “geometrically finite” actions of on δ -hyperbolic spaces. There are four basic approaches; the first is to mimic the action of a the fundamental group of a finite-volume hyperbolic manifold on \mathbb{H}^n ; the second to mimic its action on “electrified hyperbolic space”; the third to study the intrinsic (path metric) geometry of the Cayley graph; and the fourth to study (relative) isoperimetric functions. The first approach was Gromov’s, the second Farb’s. Bowditch used both the first and the second. Druţu and Sapir’s approach is the third, whereas the fourth has been very successfully followed by Osin [34].³

Examples of relatively hyperbolic groups include: (i) geometrically finite Kleinian groups (which are hyperbolic relative to their maximal parabolic subgroups); (ii) fundamental groups of hyperbolic manifolds of finite volume (hyperbolic relative to their cusp subgroups); (iii) hyperbolic groups (relative to the trivial group, or a finite collection of almost malnormal quasi-convex subgroups with pairwise finite

²As noted in the introduction, the results of this paper rely heavily on the results of [17].

³It is worth remarking that Osin’s theory allows infinitely generated groups. We will focus entirely on finitely generated groups in this paper.

intersections); (iv) free products (relative to the factors); and (v) limit groups (relative to their maximal noncyclic abelian subgroups). See Farb [19], Bowditch [6], Dahmani [12] and Szaeapański [49] for details.

For further recent work on relatively hyperbolic groups, see Alibegović [1], Yaman [50], Druţu and Sapir [17; 16; 18], Dahmani [13], Groves and Manning [25], Osin [34] and Chatterji and Ruane [10] (among others).

Definition 2.1 (Coned-off Cayley graph) Suppose that Γ is a finitely generated group, with finite generating set \mathcal{A} , and that $\{H_1, \dots, H_m\}$ is a collection of finitely generated subgroups of Γ . Let X be the Cayley graph of Γ with respect to \mathcal{A} . We form the *coned-off Cayley graph*, \tilde{X} , by adding to X a vertex c_{γ, H_i} for each coset γH_i of a parabolic subgroup, and for each coset γH_i , an edge from c_{γ, H_i} to γ' for each $\gamma' \in \gamma H_i$.

Definition 2.2 We say that Γ is *hyperbolic relative to* $\{H_1, \dots, H_m\}$ (a collection of finitely generated subgroups) if

- (1) the coned-off Cayley graph \tilde{X} is δ -hyperbolic for some δ ;
- (2) for each edge $e \in \tilde{X}$, and each $n \geq 1$, there are only finitely many loops of length at most n which contain e .

The above definition of relatively hyperbolic is a hybrid of Farb's definition and [6, Definition 2] of Bowditch. Condition (2) of Definition 2.2 is the essence of Bowditch's definition, though he allows actions on spaces other than the coned-off Cayley graph. Farb's definition [19] is that the coned-off Cayley graph be δ -hyperbolic, but he also required that it satisfied the "Bounded Coset Penetration" (BCP) Property. It follows from [6] and [11, Theorem 6.1] that our definition is equivalent to the usual definitions (such as Farb's or Bowditch's).

Terminology 2.3 Suppose that Γ is a group which is hyperbolic relative to the collection $\{H_1, \dots, H_m\}$ of subgroups. The subgroups H_i are called *parabolic subgroups*.⁴

In this paper we are concerned with torsion-free relatively hyperbolic groups Γ where all the parabolic subgroups are free abelian.

Definition 2.4 A subgroup K of a group G is *malnormal* if for all $g \in G \setminus K$ we have $gKg^{-1} \cap K = \{1\}$.

A group G is *CSA* if any maximal abelian subgroup of G is malnormal.

⁴Alternative terminology for these subgroups is *peripheral* subgroups. Sometimes, all conjugates of the H_i are also called parabolic subgroups.

Lemma 2.5 *Suppose that Γ is a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. Then Γ is CSA.*

Proof Let A be a maximal abelian subgroup of Γ .

Since Γ is torsion-free, any conjugate of a parabolic subgroup is malnormal (see Farb [19, Example 1, page 819]). This implies that if M is a conjugate of a parabolic subgroup and A intersects M nontrivially then $A = M$, which is malnormal.

Suppose that A is a maximal abelian subgroup of Γ and that $g \in A \setminus \{1\}$. If g is not contained in a conjugate of a parabolic subgroup then a result of Osin [34, Theorem 1.14, page 10] and the comment thereafter) implies that the centraliser of $\langle g \rangle$ is virtually cyclic. Since Γ is torsion-free, this centraliser is cyclic. Therefore, in this case $A = \langle h \rangle$ for some h . Note that A is maximal cyclic in Γ . Suppose now that $\gamma \in \Gamma$ is such that $\gamma h^k \gamma^{-1} = h^j$ for some $k, j \in \mathbb{Z} \setminus \{0\}$. Then [34, Corollary 4.21, page 83] implies that $|k| = |j|$. Thus, γ^2 commutes with h^j . This implies that $\gamma^2 \in \langle h \rangle$, which in turn implies that $\gamma \in \langle h \rangle$, so A is malnormal. \square

3 Druţu and Sapir's results

In [17], Druţu and Sapir find a characterisation of relatively hyperbolic groups in terms of their asymptotic cones. In this section, we recall the definition of asymptotic cones and then briefly summarise those of Druţu and Sapir's results necessary for this paper.

3.1 Asymptotic cones

Asymptotic cones were introduced by van den Dries and Wilkie in [15] in order to recast Gromov's Polynomial Growth Theorem from [20]. See Druţu and Sapir [17] for a discussion of other results about asymptotic cones. We briefly recall the definition of asymptotic cones.

Definition 3.1 *A nonprincipal ultrafilter is a $\{0, 1\}$ -valued finitely additive measure on \mathbb{N} defined on all subsets of \mathbb{N} so that any finite set has measure 0.*

The existence of nonprincipal ultrafilters is guaranteed by Zorn's Lemma. We fix once and for all a nonprincipal ultrafilter ω .⁵ Given any bounded sequence $\{a_n\} \subset \mathbb{R}$ there is a unique number $a \in \mathbb{R}$ so that for all $\epsilon > 0$ we have $\omega(\{a_n \mid |a - a_n| < \epsilon\}) = 1$.

⁵A different choice of ultrafilter can change the resulting asymptotic cone in interesting ways, but not in a way that affects our results. Thus, we are unconcerned which ultrafilter is chosen.

We denote a by $\omega\text{-}\lim\{a_n\}$. This notion of limit exhibits most of the properties of the usual limit (see van den Dries and Wilkie [15]).

Let (X, d) be a metric space. Suppose that $\{\mu_n\}$ is a sequence of real numbers with no bounded subsequence, and that $\{x_n\}$ is a collection of points in X . Let (X_n, d_n) be the metric space which has set X and metric $(1/\mu_n)d_X$. The *asymptotic cone of X with respect to $\{x_n\}, \{\mu_n\}$ and ω* , denoted X_ω , is defined as follows. First, define the set \widetilde{X}_ω to consist of all sequences $\{y_n \mid y_n \in X_n\}$ for which $\{d_{X_n}(x_n, y_n)\}$ is a bounded sequence. Define a pseudo-metric \tilde{d} on \widetilde{X}_ω by

$$\tilde{d}(\{y_n\}, \{z_n\}) = \omega\text{-}\lim\{d_{X_n}(y_n, z_n)\}.$$

The asymptotic cone X_ω is the metric space induced by the pseudo-metric \tilde{d} on \widetilde{X}_ω :

$$X_\omega := \widetilde{X}_\omega / \sim,$$

where the equivalence relation \sim on \widetilde{X}_ω is defined by: $x \sim y$ if and only if $\tilde{d}(x, y) = 0$. The pseudo-metric \tilde{d} on \widetilde{X}_ω naturally descends to a metric on d_ω on X_ω .

3.2 Tree-graded spaces and relatively hyperbolic spaces

Definition 3.2 [17, Definition 1.10] Let Y be a complete geodesic metric space, and let \mathcal{P} be a collection of closed subsets of Y (called *pieces*). We say that the space Y is *tree-graded with respect to \mathcal{P}* if the following two conditions are satisfied:

- (T_1) Each pair of distinct pieces intersect in at most a point.
- (T_2) Every simple geodesic triangle (a simple loop composed of three geodesics) in X is contained in a single piece.

In [24] it is proved that if Y is a CAT(0) space with isolated flats and relatively thin triangles then any asymptotic cone Y_ω of Y is tree-graded with respect to its collection of maximal flats (the proof of this is essentially contained in [29]). Thus, it seemed natural to generalise the results from [24] as has been done in this paper. In [28], Hruska and Kleiner prove that if a cocompact CAT(0) space has isolated flats then it has relatively thin triangles.

One of the main results of [17] is the following.

Theorem 3.3 [17, Theorem 1.11] *A finitely generated group G is relatively hyperbolic with respect to finitely generated subgroups H_1, \dots, H_n if and only if every asymptotic cone of G (with respect to any nonprincipal ultrafilter, any sequence of scaling constants, where the basepoints are the identity of G) is tree-graded with respect to ω -limits of sequences of cosets of the subgroups H_i .*

We also need the following definition and results. In the definition below I^ω is the ultrapower of the set I with respect to the ultrafilter ω .

Definition 3.4 [17, Definition 3.19] Let (Y, dist) be a metric space and let $\mathcal{Q} = \{Q_i \mid i \in I\}$ be a collection of subsets of Y . In every asymptotic cone Y_ω , with choice of basepoints $\{x_n\}$, we consider the collection of subsets

$$\mathcal{Q}_\omega = \left\{ \lim^\omega(Q_{i_n}) \mid (i_n)^\omega \in I^\omega \text{ such that } \left\{ \frac{\text{dist}(x_n, Q_{i_n})}{d_n} \right\} \text{ is bounded} \right\}.$$

The space Y is *asymptotically tree-graded with respect to \mathcal{Q}* if every asymptotic cone Y_ω is tree-graded with respect to \mathcal{Q}_ω .

Theorem 3.5 [17, Theorem 5.1] Let Y be a metric space and let \mathcal{Q} be a collection of subsets of Y . Let $q: Y \rightarrow Y'$ be a quasi-isometry. If Y is asymptotically tree-graded with respect to \mathcal{Q} then Y' is asymptotically tree-graded with respect to $q(\mathcal{Q})$.

Theorem 3.6 [17, Theorem 4.1] Let (Y, dist) be a geodesic metric space and let $\mathcal{Q} = \{Q_i \mid i \in I\}$ be a collection of subsets of Y . The space Y is asymptotically tree-graded with respect to \mathcal{Q} if and only if the following properties are satisfied:

- (α_1) For every $\xi > 0$, the diameters of the intersections $\mathcal{N}_\xi(Q_i) \cap \mathcal{N}_\xi(Q_j)$ are uniformly bounded for all $i \neq j$.
- (α_2) For every $\theta \in [0, \frac{1}{2})$ there exists $M(\theta) > 0$ so that for every geodesic q of length l and every $Q \in \mathcal{Q}$ with $q(0), q(l) \in \mathcal{N}_\theta(Q)$ we have $q([0, l]) \cap \mathcal{N}_M(Q) \neq \emptyset$.
- (α_3) For every $k \geq 2$ there exists $\zeta > 0$, $\nu \geq 8$ and $\chi > 0$ such that every k -gon P in X with geodesic edges which is (ζ, ν, χ) -fat satisfies $P \subseteq \mathcal{N}_\chi(Q)$ for some $Q \in \mathcal{Q}$.

4 The space X

In this paragraph we define a space X , closely associated to the Cayley graph of a relatively hyperbolic group, which will be the appropriate space for our analysis of Γ -limit groups in the subsequent sections, and also in [23] and [14].

The space X is quasi-isometric to the Cayley graph. However, it has two advantages (see Theorem 4.5 and Lemma 4.6 below), which we now explain. Let Γ be a group hyperbolic relative to $\{H_1, \dots, H_m\}$ (finitely generated subgroups). Let Y be the Cayley graph of Γ with respect to some finite generating set \mathcal{A} . It is convenient to assume that \mathcal{A} contains \mathcal{B}_i , a generating set for H_i . Then the Cayley graph of H_i with

respect to \mathcal{B}_i naturally embeds in Y . This embedding is a quasi-isometric embedding. In case each H_i is free abelian of rank n_i and \mathcal{B}_i is a basis, the space X contains an isometric copy of \mathbb{R}^{n_i} (with the standard Euclidean metric) on which H_i acts properly with quotient a square torus. This is the first important property (upgrading quasi-isometric embedding to isometric embedding). If a geodesic path in Y starts and finishes near to a single coset of some H_i , then it cannot travel too far from this coset. The second important property of our space X (Lemma 4.6) is that in X such a path spends most of its time contained in the appropriate Euclidean space). Thus we improve quasi-convex to “convex and attracting”. For each of these properties, an appropriate picture to have in mind is the embedding of Euclidean space in “truncated hyperbolic space” coming from a finite-volume hyperbolic manifold. Both properties hold in this case. However, it is worth noting that away from the copies of \mathbb{R}^{n_i} , the space X is a graph (with varying edge lengths).

The space X should be considered as a hybrid of the Cayley graph of a relatively hyperbolic group and a CAT(0) space with isolated flats (although a CAT(0) space with isolated flats needn't satisfy the second property discussed above). In particular, the “flats” are Euclidean and convex. The space X has “relatively thin triangles” (Theorem 4.16) and “isolated flats” (appropriately interpreted; Lemma 4.13). Though we do not need it here, it is worth remarking that geodesics in X satisfy a fellow-travelling condition exactly analogous to hyperbolic groups (rather than merely to relatively hyperbolic groups).

Suppose that Γ is a group which is hyperbolic relative to a collection $\{H_1, \dots, H_m\}$ of (finitely generated) subgroups.

Choose a generating set \mathcal{A} for Γ which intersects each of the subgroups H_i in a generating set \mathcal{B}_i for H_i , for $1 \leq i \leq m$. Let $\mathcal{B} = \bigcup_{i=1}^m \mathcal{B}_i$. For each $i \in \{1, \dots, m\}$.

Let $d_{\mathcal{A}}$ be the word metric on Γ induced by the generating set \mathcal{A} of Γ , and let $d_{\mathcal{B}_i}$ be the word metric on H_i induced by \mathcal{B}_i .

Let Y denote the Cayley graph of Γ with respect to \mathcal{A} , where each edge is isometric to the unit interval $[0, 1]$. The group Γ acts on itself by left multiplication, which induces an isometric action on Y .

Let γH_i be a coset of some parabolic subgroup of Γ . The set \mathcal{B}_i also gives a metric on γH_i , which we denote by $d_{\mathcal{B}_i}$. Now, [17, Lemma 4.3] states that there is a constant $K \geq 0$ so that for any $x, y \in \gamma H_i$, any geodesic joining x and y in Y stays entirely in the K -neighbourhood of γH_i .

We now build a space Y^k out of Y .

Consider a coset γH_i , along with the set of edges labelled by elements of \mathcal{B}_i . The resulting subgraph $Z(\gamma, H_i)$ of Y is exactly the Cayley graph of H_i . We now form a new graph $Z(\gamma H_i)^1$, which is another copy of $Z(\gamma, H_i)$, except that each edge is isometric to the closed interval $[0, 1/4]$. Denote this new graph by $Z(\gamma, H_i)^1$, and join it to $Z(\gamma, H_i)$ by joining corresponding vertices by edges of length $1/4$. Perform this construction for each coset γH_i of a parabolic subgroup.

We define Y^j inductively, starting from Y^{j-1} . First, form $Z(\gamma, H_i)^j$ with edges of length 2^{-2j} (but otherwise isomorphic to $Z(\gamma, H_i)$), and join it to $Z(\gamma, H_i)^{j-1}$ by edges of length 2^{-2j} .

Terminology 4.1 We call the edges of length 2^{-2j} joining $Z(\gamma, H_i)^{j-1}$ to $Z(\gamma, H_i)^j$ *vertical*, and the edges which lie in some $Z(\gamma, H_i)^j$ *horizontal*.

The space Y^j is the union of Y^{j-1} along with $Z(\gamma, H_i)^j$, and the vertical edges joining $Z(\gamma, H_i)^{j-1}$ to $Z(\gamma, H_i)^j$. Endow Y^j with the natural path metric.

Indira Chatterji attributes a similar construction to David Epstein, who also proves an analogue of Theorem 4.5.

For an integer $k \geq 1$, let $C(\gamma, H_i)^k$ be the union of the graphs $Z(\gamma, H_i)$, $Z(\gamma, H_i)^1, \dots, Z(\gamma, H_i)^k$, along with the sets of vertical edges that join successive graphs in this sequence.

There is a natural space Y^∞ , the metric completion of $\bigcup_{s=1}^\infty Y^s$, (where we consider Y^s to be a subset of Y^{s+1}). Each coset γH_i inherits a ‘‘cone-point’’ $w_{\gamma,i}$ from this completion process. In Y^∞ , the point $w_{\gamma,i}$ lies at distance η from the coset γH_i , where $\eta = \sum_{s=1}^\infty 2^{-2s} < 1/2$. It is clear that the space Y^∞ is quasi-isometric to the coned-off Cayley graph of Γ . Let Y^∞ be Υ -hyperbolic, and suppose without loss of generality that $\Upsilon > 1$.

We consider Y^∞ to be a ‘‘coned’’ space, and each of the Y^i to be ‘‘partially coned’’ spaces.

Definition 4.2 Suppose that H_i is a parabolic subgroup of Γ and γH_i is a coset of H_i in Γ . The space $P_{\gamma,i}$ is formed from Y by performing the construction of Y^∞ to all cosets of all parabolic subgroups *except* the coset γH_i .

Lemma 4.3 Suppose that γH_i is a coset of a parabolic subgroup in Γ and that $x, y \in \gamma H_i$. Let $[x, y]$ be a geodesic between x and y in $P_{\gamma,i}$ and suppose that $[x, y]$ does not intersect γH_i except at its endpoints. Then $[x, y]$ lies entirely in the 35Υ -neighbourhood of γH_i in $P_{\gamma,i}$.

Proof Let $\widehat{[x, y]}$ be the image of $[x, y]$ in the space Y^∞ , under the inclusion $P_{\gamma, i} \subset Y^\infty$. For any R , the R -neighbourhood of $w_{\gamma, i}$ in Y^∞ naturally corresponds to the $(R - \eta)$ -neighbourhood of γH_i in $P_{\gamma, i}$, and if $R > \eta$ then outside of these balls the two spaces are locally isometric (where the local isometry is induced by the inclusion $P_{\gamma, i} \subset Y^\infty$).

Suppose that $\widehat{[x, y]}$ is not contained in the 10Υ neighbourhood of $w_{\gamma, i}$ in Y^∞ . That part of $\widehat{[x, y]}$ which lies at least 10Υ from $w_{\gamma, i}$ is a 10Υ -local-geodesic. By [7, Theorem III.H.1.13, page 405], outside the 10Υ ball around $w_{\gamma, i}$ in Y^∞ , the path $\widehat{[x, y]}$ is a $(14\Upsilon/(6\Upsilon), 2\Upsilon)$ -quasi-geodesic. However, the distance in Y^∞ which it travels outside of the 10Υ ball around $w_{\gamma, i}$ is at most 20Υ (since the path starts and finishes at distance $\eta < \frac{1}{2}$ from $w_{\gamma, i}$).

Therefore, the total distance that $\widehat{[x, y]}$ travels outside the 10Υ ball about $w_{\gamma, i}$ is at most

$$\left(\frac{7}{3}\right)20\Upsilon + 2\Upsilon < 50\Upsilon.$$

Thus $\widehat{[x, y]}$ is contained in the 35Υ ball around $w_{\gamma, i}$ in Y^∞ . As above, this implies that $[x, y]$ is contained in the 35Υ -neighbourhood of γH_i in $P_{\gamma, i}$, as required. \square

Lemma 4.4 *There exists a constant K_1 , depending only on Y , and the set $\{\gamma H_i\}$, so that for all $x, y \in \gamma H_i$*

$$d_{P_{\gamma, i}}(x, y) \leq d_{\gamma H_i}(x, y) \leq K_1 d_{P_{\gamma, i}}(x, y).$$

Proof Since $\gamma H_i \subset P_{\gamma, i}$, and since $d_{P_{\gamma, i}}$ is a path metric, the first inequality is immediate.

Let $x, y \in \gamma H_i$, and let $[x, y]$ be a geodesic between x and y in $P_{\gamma, i}$. By Lemma 4.3 above, the path $[x, y]$ lies entirely within the 35Υ -neighbourhood of γH_i .

Let c_1, c_2, \dots, c_k be points along $[x, y]$ which are such that $\eta \leq d_{P_{\gamma, i}}(c_i, c_{i+1}) \leq 1$ (this can be ensured if we choose each c_i to be either a vertex from Y or a cone-point $w_{\gamma', j}$). For each $1 \leq i \leq k$, let $b_i \in \gamma H_i$ be a point in γH_i as close as possible to c_i , and choose a geodesic $[c_i, b_i]$ (which has length at most 35Υ). Possibly $c_i = b_i$, and so the path $[c_i, b_i]$ is a constant path. Also, choose a path $[b_i, b_{i+1}] \subset \gamma H_i$ of shortest length.

Consider the paths $p_i = [b_i, b_{i+1}]$ and $q_i = [b_i, c_i, c_{i+1}, b_{i+1}]$. The path q_i has length at most $70\Upsilon + 1$. Also, unless $[c_i, c_{i+1}] \subset \gamma H_i$, the path q_i intersects γH_i only at its endpoints.

The path q_i corresponds to a path $q'_i \subset Y$, where any part of q_i which passes through a cone-point is replaced by a (shortest) path through the corresponding coset. Now,

q'_i is a relative $(70\Upsilon + 1)$ -quasi-geodesic. Also, p_i is a relative 2η -quasi-geodesic, and can be considered as a path in Y . Note that p_i penetrates γH_i while q'_i does not. Therefore, Bounded Coset Penetration implies that there is a constant $c = c(70\Upsilon + 1)$ so that p_i travels distance at most c in γH_i , which is to say that p_i has length at most c .

We have seen that each $[b_i, b_{i+1}]$ has length at most c . Therefore $d_{\gamma H_i}(x, y) \leq ck$. However, $d_{P_{\gamma,i}}(x, y) \geq \eta k$, and it suffices to take $K_1 = c/\eta$. \square

Now let $P_{\gamma,i}^{(k)}$ be the space formed from Y by adding the spaces $C(\gamma', H_j)^k$ for all cosets of parabolic subgroups except γH_i . Then for all $x, y \in \gamma H_i$ we have

$$d_{P_{\gamma,i}}(x, y) \leq d_{P_{\gamma,i}^{(k)}}(x, y),$$

and so by Lemma 4.4 we have

$$d_{\gamma H_i}(x, y) \leq K_1 d_{P_{\gamma,i}^{(k)}}(x, y).$$

Theorem 4.5 *There exists $k \geq 0$ so that each of the graphs $Z(\gamma, H_i)^k$ is isometrically embedded in Y^k .*

Proof It suffices to take $k > (\log_2 K_1)/2$, where K_1 is the constant from the conclusion of Lemma 4.4.

Suppose that there exists $u, v \in Z(\gamma, H_i)^k$ so that a geodesic $[u, v]$ between u and v does not lie entirely within $Z(\gamma, H_i)^k$. Since $Z(\gamma, H_i)^k$ is isometrically embedded in $C(\gamma, H_i)^k$, the path $[u, v]$ cannot be contained in $C(\gamma, H_i)^k$.

Suppose that x is the point furthest along $[u, v]$ so that $[u, x] \subset C(\gamma, H_i)^k$. Let y be the point furthest along $[u, v]$ so that $[x, y]$ intersects $C(\gamma, H_i)^k$ only in $Z(\gamma, H_i)$. Because of the way $C(\gamma, H_i)^k$ was built, that part of $[u, v]$ immediately before x consists entirely of edges joining the different $Z(\gamma, H_i)^i$, from $Z(\gamma, H_i)^k$ to $Z(\gamma, H_i)$. Similarly, that part of $[u, v]$ immediately after y consists of a “vertical” path from y to $Z(\gamma, H_i)$. Let x_1 be the final point in $[u, x]$ contained in $Z(\gamma, H_i)^k$ and let y_1 be the first point in $[y, v]$ contained in $Z(\gamma, H_i)^k$. Then we have $d_{\gamma H_i}(x, y) = 2^{2k} d_{Z(\gamma, H_i)^k}(x_1, y_1)$. Now let $D = d_{P_{\gamma,i}^{(k)}}(x, y)$, and note that $D \geq \eta$.

$$\begin{aligned} D + \frac{1}{2} &\leq d_{Z(\gamma, H_i)^k}(x_1, y_1) \\ &= 2^{-2k} d_{\gamma H_i}(x, y) \\ &\leq 2^{-2k} K_1 D. \end{aligned}$$

Therefore, $D + \frac{1}{2} \leq 2^{-2k} K_1 D$, which implies in particular that $2^{-2k} K_1 - 1 > 0$, contradicting our choice of k . This completes the proof. \square

We now fix k so that Theorem 4.5 holds, and consider the space Y^k .

Lemma 4.6 *There exists a function $f_1: \mathbb{N} \rightarrow \mathbb{N}$ so that if $x, y \in Y^k$ are such x and y lie in the N -neighbourhood of γH_i and $[x, y]$ does not intersect $Z(\gamma, H_i)^k$ then $d(x, y) \leq f_1(N)$.*

Proof By Theorem 4.5, and the definition of Y^k , it suffices to bound the length of a geodesic $[w, z]$ where $w, z \in \mathcal{N}_N(\gamma H_i)$ and $[w, z]$ does not intersect $C(\gamma, H_i)^k \setminus Z(\gamma, H_i)$.

For such a pair we have

$$d_{Y^k}(w, z) = d_{\mathcal{P}_{\gamma, i}^{(k)}}(w, z).$$

Denote this distance by E . Let w_1, z_1 be points in γH_i which are closest to w and z , respectively. Also, let w_2, z_2 be the points in $Z(\gamma, H_i)^k$ which are closest to w_1 and z_1 , respectively. Also, let $\eta_k = \sum_{i=1}^k 2^{-2i}$ be the distance from γH_i to $Z(\gamma, H_i)^k$. Then we have

$$\begin{aligned} E &= d_{Y^k}(w, z) \\ &\leq d_{Y^k}(w_1, z_1) + 2N \\ &= d_{Z(\gamma, H_i)^k}(w_2, z_2) + 2N + 2\eta_k \\ &= 2^{-2k} d_{\gamma H_i}(w_1, z_1) + 2N + 2\eta_k \\ &= 2^{-2k} K_1 d_{\mathcal{P}_{\gamma, i}^{(k)}}(w_1, z_1) + 2N + 2\eta_k \\ &\leq 2^{-2k} K_1 (d_{\mathcal{P}_{\gamma, i}^{(k)}}(w, z) + 2N) + 2N + 2\eta_k \\ &= 2^{-2k} K_1 (E + 2N) + 2N + 2\eta_k, \end{aligned}$$

which implies (since the choice of k from Theorem 4.5 ensures that $1 - 2^{-2k} K_1 > 0$) that

$$E \leq \frac{2^{-2k} K_1 N + 2N + 2\eta_k}{1 - 2^{-2k} K_1}.$$

This completes the proof. \square

We now assume that Γ is torsion-free and that each of the parabolic subgroups of Γ are free abelian.

Remark 4.7 Suppose that a group G is hyperbolic relative to a family \mathcal{P} of subgroups, and that some of the subgroups in \mathcal{P} are hyperbolic. Let \mathcal{P}' be the nonhyperbolic groups in \mathcal{P} . Then G is also hyperbolic relative to \mathcal{P}' .

Therefore, we assume that all parabolic subgroups of our relatively hyperbolic groups are not hyperbolic. In case parabolics are free abelian, as in this paper, this amounts to assuming that none of the parabolics are infinite cyclic (or trivial).

Suppose that the generating set for Γ intersects each parabolic subgroup in a basis (as a free abelian group). Then each of the graphs $Z(\gamma, H_i)^k$ is isomorphic to the “standard” Cayley graph of \mathbb{Z}^{n_i} (with edge of length 2^{-2k}). Fix an embedding $\phi_i: Z(\gamma, H_i)^k \hookrightarrow \mathbb{R}^{n_i}$ which is isometric on each edge, and send the vertices adjacent to the identity to (scaled) standard basis vectors in \mathbb{R}^{n_i} (and their negatives).

Using the map ϕ_i , glue a copy of \mathbb{R}^{n_i} onto each subspace $Z(\gamma, H_i)^k$ of Y^k where n_i is the rank of H_i , and \mathbb{R}^{n_i} is equipped with the standard Euclidean (ℓ_2 -) metric.

Definition 4.8 The resulting space is denoted X , and \mathcal{Q} is the collection of copies of the \mathbb{R}^{n_i} glued onto the cosets γH_i (where $i \in \{1, \dots, m\}$ and $\gamma \in \Gamma$).

The copies of \mathbb{R}^n that have been glued to Y^k to form the space X now play the role of cosets. They are isometrically embedded, and Lemma 4.6 above holds for these subspaces also, since lengths of paths are unchanged outside of $Z(\gamma, H_i)^k$, and distances can only get shorter inside $Z(\gamma, H_i)^k$.

The action of Γ on X is defined in the obvious way. The stabiliser in Γ of any $Q \in \mathcal{Q}$ is a conjugate of a parabolic subgroup, which acts by translations on Q .

4.1 Properties of X

Lemma 4.9 Suppose $Q \in \mathcal{Q}$ is a copy of \mathbb{R}^{n_i} in X as above. Then Q is isometrically embedded and convex in X .

Given Lemma 4.9, it is natural to call the elements of \mathcal{Q} “flats”.

Lemma 4.10 Left multiplication of Γ on itself induces an isometric action of Γ on X . This action is proper and cocompact.

Lemma 4.11 X is asymptotically tree-graded with respect to the set \mathcal{Q} .

Proof The inclusion map $\Gamma \hookrightarrow X$ is a quasi-isometry.

We know that Γ is asymptotically tree-graded with respect to the set of cosets γH_i . By the proof of [17, Theorem 5.1, page 44], X is asymptotically tree-graded with respect to \mathcal{Q} . \square

Note also that any asymptotic cone of Γ is bi-Lipschitz homeomorphic to the analogous asymptotic cone of X (taking the same basepoints, and the same scaling factors). The utility of using X rather than just Γ is the following:

Lemma 4.12 *Suppose that X_ω is an asymptotic cone of X . Then each piece of X_ω is isometric to (\mathbb{R}^k, d) for some k , where d is the standard Euclidean metric on \mathbb{R}^k .*

In analogy with Hruska's definition of *isolated flats* for CAT(0) spaces (see Hruska [27, 2.1.2]), we note the following:

Lemma 4.13 (Isolated flats) *Let \mathcal{Q} be the collection of flats in X . Then there is a function $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for every pair of distinct flats $Q_1, Q_2 \in \mathcal{Q}$ and for every $k \geq 0$, the intersection of the k -neighbourhoods of Q_1 and Q_2 has diameter less than $\phi(k)$.*

Proof Theorem 3.6 (α_1) provides such a ϕ . □

Convention 4.14 Let $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be as in Lemma 4.13. We suppose that $\phi(k) \geq k$ for all $k \geq 0$ and that ϕ is a nondecreasing function.

We now prove a quasi-convexity result for the metric on X .

Lemma 4.15 *There exists a function $N_1: \mathbb{N} \rightarrow \mathbb{N}$ so that for any K , if $x_1, x_2, y \in X$ so that $d_X(x_1, x_2) \leq K$ and $[x_1, y]$ and $[x_2, y]$ are geodesics, then $[x_1, y]$ is contained in the $N_1(K)$ -neighbourhood of $[x_2, y]$ (and vice versa).*

Proof Choose a geodesic $[x_1, x_2]$. Then the path $[y, x_1, x_2] = [y, x_1] \cup [x_1, x_2]$ is a $(1, K)$ -quasi-geodesic.

By [17, Theorem 1.12] there are constants τ and M so that

- $[y, x_1, x_2]$ is contained in the τ -tubular neighbourhood of the M -saturation of $[y, x_2]$ (see Druţu and Sapir [17, Definition 8.9] for a definition of M -saturation);
- the points at which $[y, x_1, x_2]$ enters and leaves the τ -neighbourhood of flats in the M -saturation of $[y, x_2]$ are at bounded distance from $[a, x_1]$.

By Lemma 4.6, and the fact the flats are isometric to \mathbb{R}^n , the path $[y, x_1, x_2]$ lies in the D_2 -neighbourhood of $[y, x_2]$ for some constant D_2 . A symmetric argument on $[y, x_2, x_1]$ and $[y, x_1]$ implies that $[y, x_2, x_1]$ lies in the D_2 -neighbourhood of $[y, x_1]$. □

For our purposes, one of the most important properties of the space X is contained in the following theorem, which shows that geodesic triangles in X satisfy the *Relatively Thin Triangles Property* (see Hruska [27, Definition 3.1.1]).

Theorem 4.16 *Suppose that X is as constructed above. There exists $\delta > 0$ so that for any $a, b, c \in X$, and any $\Delta(a, b, c)$ is a geodesic triangle, either (i) $\Delta(a, b, c)$ is δ -thin in the usual sense; or else (ii) there is a unique flat $E \subset X$ so that each side of $\Delta(a, b, c)$ is contained in the δ -neighbourhood of the union of E and the other two sides.*

Proof Choose a geodesic triangle $\Delta(a, b, c)$ in X , with a choice of geodesics $[a, b]$, $[b, c]$ and $[a, c]$.

By [17, Lemma 8.16] and [17, Lemma 8.17], there is a constant α (independent of the points a, b, c) such that one of two possibilities occurs: either

- (i) there is a point $x \in X$ whose α -neighbourhood intersects all three of the geodesics $[a, b]$, $[b, c]$ and $[a, c]$ nontrivially; or
- (ii) there is a flat $E \in \mathcal{Q}$ so that the α -neighbourhood of E intersects each of the three geodesics nontrivially.

In case (i), let x_1 be a point on $[a, b]$ which is within α of x , and let x_2 be a point on $[a, c]$ which is within α of x . Then $d_X(x_1, x_2) \leq 2\alpha$.

In case (ii) let x_1 be the point on $[a, b]$ which is closest to a subject to being in the α -neighbourhood of E , and similarly for x_2 on $[a, c]$. Then [17, Corollary 8.14] implies that there is a constant D_1 so that $d_X(x_1, x_2) \leq D_1$. We assume that $D_1 \geq 2\alpha$.

So in either case, there exist points $x_1 \in [a, b]$ and $x_2 \in [a, c]$ so that $d_X(x_1, x_2) \leq D_1$. Denote by $[a, x_1]$ the subpath of $[a, b]$ from a to x_1 , and similarly for $[a, x_2] \subset [a, c]$. By Lemma 4.15, there is a constant D_2 so that $[a, x_1]$ lies in the D_2 -neighbourhood of $[a, x_2]$, and vice versa.

We use a symmetric argument on the points b and c – finding points $y_1 \in [c, a]$, $y_2 \in [c, b]$ and $z_1 \in [b, a]$, $z_2 \in [b, c]$ as with x_1 and x_2 .

Now, in case (i) above, we can take $x_1 = z_1$, $x_2 = y_1$ and $y_2 = z_2$, and we're done. In case (ii), we note that the path $[x_1, z_1] \subseteq [a, b]$ lies in the $N_1(\alpha)$ -neighbourhood of E , by Lemma 4.15, and similarly for $[x_2, y_1] \subseteq [a, c]$ and $[z_2, y_2] \subseteq [b, c]$.

Therefore, it suffices to take $\delta = \max\{D_2, N_1(\alpha)\}$. □

4.2 Projecting to flats

In this paragraph we record some results about projecting to flats which are required for the proofs in the subsequent sections.

Definition 4.17 [17, Definition 4.9] Let $x \in X$ and $A \subset X$. The *almost projection* of x onto A is the set of points $y \in A$ so that $d_X(x, y) \leq d_X(x, A) + 1$.

The following results follows immediately from [17, Corollary 8.14] and Lemma 4.11.

Lemma 4.18 *There exists a constant C_1 so that if $Q \in \mathcal{Q}$ and $x \in X$ then the almost projection of x onto Q has diameter at most C_1 .*

Lemma 4.19 *There exists a function $N_3: \mathbb{N} \rightarrow \mathbb{N}$ so that if $x_1, x_2 \in X$, $Q \in \mathcal{Q}$ and $\pi(x_1), \pi(x_2)$ are in the almost projections of x_1 and x_2 to Q , respectively then $d_X(\pi(x_1), \pi(x_2)) \leq N_3(d_X(x_1, x_2))$.*

Again, we suppose that $N_3(x) \geq x$ for all $x \geq 0$ and that N_3 is a nondecreasing function.

Recall that δ is the constant from Theorem 4.16 and that $\phi: \mathbb{N} \rightarrow \mathbb{N}$ is the function from Lemma 4.13.

Lemma 4.20 (cf [24, Lemma 2.11]) *Suppose that $\Delta = \Delta(a, b, c)$ is a geodesic triangle in X . If Δ is not $(\delta + \phi(\delta)/2)$ -thin then Δ is δ -thin relative to a unique flat $Q \in \mathcal{Q}$.*

Proof Given Lemma 4.13 and Theorem 4.16, the proof of [24, Lemma 2.11] applies directly. \square

Lemma 4.21 (cf [24, Lemma 2.21]) *Suppose that $Q \in \mathcal{Q}$, that $x, y \in Q$ and that $z \in X$. Let $[x, z]$ and $[y, z]$ be geodesics. Then there exist $u \in [x, z]$ and $v \in [y, z]$ that both lie in the 2δ -neighbourhood of Q such that*

$$d_X(u, v) \leq \phi(\delta).$$

Proof Given Theorem 4.5 and Theorem 4.16 and Lemma 4.13, the proof of [24, Lemma 2.21] applies directly. \square

Proposition 4.22 (cf [24, Proposition 2.22]) *Suppose that $Q \in \mathcal{Q}$, that $x, y \in X$ and that some geodesic $[x, y]$ does not intersect the 4δ -neighbourhood of Q . Let $\pi(x), \pi(y)$ be in the almost projections of x and y , respectively. Then*

$$d_X(\pi(x), \pi(y)) \leq N_3(\phi(3\delta) + N_1(\phi(\delta))).$$

Proof By Lemma 4.21 there exist $w_1 \in [\pi(x), y]$ and $w_2 \in [\pi(y), y]$, both in the 2δ -neighbourhood of Q such that $d_X(w_1, w_2) \leq \phi(\delta)$. By a similar argument as in the proof of Lemma 4.21 (see Groves [24]), there are $u_1 \in [\pi(x), x]$ and $u_2 \in [\pi(x), y]$ which lie outside the 2δ -neighbourhood of E so that $d_X(u_1, u_2) \leq \phi(3\delta)$. Now $[w_1, y]$ is contained in the $N_1(\phi(\delta))$ -neighbourhood of $[w_2, y]$, by Lemma 4.15. Therefore, there exists $u_3 \in [\pi(y), y]$ so that $d_X(u_2, u_3) \leq N_1(\phi(3\delta))$.

We can choose $\pi(u_1)$ and $\pi(u_3)$ in the almost projections of u_1, u_3 so that $\pi(u_1) = \pi(x)$ and $\pi(u_3) = \pi(y)$. Now, $d_X(u_1, u_3) \leq \phi(3\delta) + N_1(\phi(\delta))$, so by Lemma 4.19 we have

$$\begin{aligned} d_X(\pi(x), \pi(y)) &= d_X(\pi(u_1), \pi(u_3)) \\ &\leq N_3(d_X(u_1, u_3)) \\ &\leq N_3(\phi(3\delta) + N_1(\phi(\delta))). \end{aligned} \quad \square$$

The above result could also be easily proved for the Cayley graph of a relatively hyperbolic group, and almost projections to cosets of parabolic subgroups.

5 Asymptotic cones and compactification

In this section we start with Γ , a finitely generated group which acts properly and cocompactly by isometries on a metric space (X, d_X) , a finitely generated group G and a sequence $\{h_n: \Gamma \rightarrow G\}$ of homomorphisms. Using $\{h_n\}$ we construct a particular asymptotic cone X_ω , which is equipped with an isometric action of G with no global fixed point.

In the case of δ -hyperbolic groups and spaces, the construction we describe in this section is essentially due to Paulin [35; 37] (see also Bestvina [3] and Bridson and Swarup [8]), though was not cast there in terms of asymptotic cones. For CAT(0) spaces, this construction is performed by Kapovich and Leeb [29]. The general construction is similar. See Groves [24] for more details about this construction and van den Dries and Wilkie [15] or Druţu and Sapir [17] for many properties about asymptotic cones.

Let G be a finitely generated group and Γ a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. Let \mathcal{A} be a finite generating set for G , let X be the space constructed from a Cayley graph of Γ in Section 4, and let $x \in X$ correspond to the identity of Γ . If $h: G \rightarrow \Gamma$ is a homomorphism, define

$$\|h\| := \min_{\gamma \in \Gamma} \max_{g \in \mathcal{A}} d_X(x, (\gamma h(g)\gamma^{-1}) \cdot x),$$

and let γ_h be an element of Γ which realises this minimum.

Terminology 5.1 We say that a pair of homomorphisms $h, h': G \rightarrow \Gamma$ are *nonconjugate* if there is no inner automorphism $\tau: \Gamma \rightarrow \Gamma$ so that $h' = \tau \circ h$.

Suppose that $\{h_i: G \rightarrow \Gamma\}$ is a sequence of pairwise nonconjugate homomorphisms. Then the sequence $\{\|h_n\|\}$ does not contain a bounded subsequence. Let X_ω be the asymptotic cone, defined with respect to some nonprincipal ultrafilter ω , the sequence of basepoints $x_n = x$ and the sequence of scaling factors $\mu_n = \|h_n\|$.

The action of G on X_ω is defined by $g \cdot \{y_n\} = \{\gamma_{h_n} h_n(g) \gamma_{h_n}^{-1} \cdot y_n\}$.

Lemma 5.2 *The action of G on X_ω has no global fixed point.*

Proof See Lemma 3.9 of [24], where the proof does not use the CAT(0) property. \square

5.1 The subspace \mathcal{C}_∞ and the action of G

We now change our focus somewhat, and consider the Gromov–Hausdorff topology rather than asymptotic cones. The reason we make this change is that it allows us to replace various statements about “ ω -almost all i ” by the stronger “for all but finitely many i ”. It is possible to make the entire construction in the context of Gromov–Hausdorff convergence, avoiding asymptotic cones and ultrafilters entirely. However, this approach requires considering “approximate convex hulls” as in [8], and is more cumbersome. Thus, we preferred to start with asymptotic cones, and only now switch to Gromov–Hausdorff convergence.

Definition 5.3 Define the subspace $\mathcal{C}_\infty \subset X_\omega$ to be the union of (i) the geodesic segments $[x_\omega, g \cdot x_\omega]$ for all $g \in G$; and (ii) the flats $Q_\omega \subseteq X_\omega$ which contain a simple geodesic triangle contained in $\Delta(g_1 \cdot x_\omega, g_2 \cdot x_\omega, g_3 \cdot x_\omega)$ for some g_1, g_2, g_3 .

Lemma 5.4 *The space \mathcal{C}_∞ is (i) separable; (ii) G -invariant; (iii) convex in X_ω ; and (iv) tree-graded with pieces isometric to (\mathbb{R}^n, ℓ_2) , for some n (where n may vary according to the piece).*

Suppose that $\{(Y_n, \lambda_n)\}_{n=1}^\infty$ and (Y, λ) are pairs consisting of metric spaces, together with actions $\lambda_n: G \rightarrow \text{Isom}(Y_n)$, $\lambda: G \rightarrow \text{Isom}(Y)$. Recall (cf [4, Section 3.4, page 16]) that $(Y_n, \lambda_n) \rightarrow (Y, \lambda)$ in the G -equivariant Gromov topology if and only if: for any finite subset K of Y , any $\epsilon > 0$ and any finite subset P of G , for sufficiently large n , there are subsets K_n of Y_n and bijections $\rho_n: K_n \rightarrow K$ such that for all $s_n, t_n \in K_n$ and all $g_1, g_2 \in P$ we have

$$|d_Y(\lambda(g_1) \cdot \rho_n(s_n), \lambda(g_2) \cdot \rho_n(t_n)) - d_{Y_n}(\lambda_n(g_1) \cdot s_n, \lambda_n(g_2) \cdot t_n)| < \epsilon.$$

To a homomorphism $h: G \rightarrow \Gamma$, we naturally associate a pair (X_h, λ_h) as follows: let $X_h = X$, endowed with the metric $(1/\mu_h)d_X$; and let $\lambda_h = \iota \circ h$, where $\iota: \Gamma \rightarrow \text{Isom}(X)$ is the fixed homomorphism.

Let $\lambda_\infty: G \rightarrow \text{Isom}(\mathcal{C}_\infty)$ denote the action of G on \mathcal{C}_∞ .

Proposition 5.5 [24, Lemma 3.15] *If there is a separable G -invariant subspace \mathcal{C} of X_ω which contains the basepoint x_ω of X_ω then there is a subsequence $\{f_i\}$ of $\{h_i\}$ so that $(X_{f_i}, \lambda_{f_i}) \rightarrow (\mathcal{C}_\infty, \lambda_\infty)$ in the G -equivariant Gromov topology.*

For the remainder of the section and the next, we assume that we have passed to the convergent subsequence $\{f_i\}$ of $\{h_i\}$. In this vein, we denote X_{f_i} by X_i , and λ_{f_i} by λ_i .

Lemma 5.6 [24, Corollary 3.17] *Let \mathcal{F}_∞ be the set of flats in \mathcal{C}_∞ . For each $E \in \mathcal{F}_\infty$ there is a sequence $\{E_i \subset X_i\}$ so that $E_i \rightarrow E$ in the G -equivariant Gromov topology.*

Observation 5.7 The action of G on \mathcal{C}_∞ has no global fixed point.

Definition 5.8 Suppose that G and Γ are finitely generated groups and $\{h_i: G \rightarrow \Gamma\}$ is a sequence of pairwise nonconjugate homomorphisms, leading to an isometric action of G on \mathcal{C}_∞ , where \mathcal{C}_∞ is constructed from X_ω , the asymptotic cone of Γ , as above. Let K_∞ be the kernel of the action of G on \mathcal{C}_∞ :

$$K_\infty = \{g \in G \mid \forall y \in \mathcal{C}_\infty, g \cdot y = y\}.$$

The *strict Γ -limit group* is $L_\infty = G/K_\infty$.

A *Γ -limit group* is a group which is either a strict Γ -limit group as above or else a finitely generated subgroup of Γ .

The following result is clear from the definition of the Gromov topology.

Lemma 5.9 *Suppose that the sequence of homomorphisms $\{f_i: G \rightarrow \Gamma\}$ gives rise to a sequence of actions converging to an action of G on \mathcal{C}_∞ , and that K_∞ is the kernel of the action of G on \mathcal{C}_∞ . Then $\varinjlim \text{Ker}(f_i) \subset K_\infty$.*

The following results give information about the flats in \mathcal{C}_∞ , and their stabilisers in G .

Proposition 5.10 (cf [24, Lemma 3.18]) *Suppose $g \in G$ leaves a flat $E \subseteq \mathcal{C}_\infty$ (set-wise) invariant, and that $\{E_j\}$ converges to E . Then for all but finitely many i we have $f_i(g) \cdot E_i = E_i$.*

Proof The proof of [24, Lemma 3.18] applies directly. \square

Proposition 5.11 (cf [24, Lemma 3.19]) *Suppose $g \in \text{Stab}_G(E)$ for some flat $E \subseteq \mathcal{C}_\infty$. Then g acts (possibly trivially) by translation on E .*

Proof The proof of [24, Lemma 3.19] applies directly, once we notice that an element of $\gamma \in \Gamma$ which leaves a flat in X invariant lies in a conjugate of a parabolic subgroup and acts by Euclidean translations on the flat. \square

6 The \mathbb{R} -tree T

6.1 Constructing the \mathbb{R} -tree

We now follow [24] to construct from \mathcal{C}_∞ an \mathbb{R} -tree T equipped with an isometric G -action with no global fixed point. Given the construction of \mathcal{C}_∞ in the previous section, the construction of T is exactly the same as in [24]. We repeat the definition of T here.

Let \mathcal{F}_∞ be the collection of all pieces in \mathcal{C}_∞ . By Definition 3.2, for any $g \in G$ exactly one of the following holds: (i) $g \cdot E = E$; (ii) $|g \cdot E \cap E| = 1$; or (iii) $g \cdot E \cap E = \emptyset$. By Proposition 5.10 and Proposition 5.11, $\text{Stab}(E)$ is a countable abelian group, acting by translations on E (possibly not faithfully).

Let \mathcal{D}_E be the set of directions of the translations of E by elements of $\text{Stab}(E)$.

For each element $g \in G \setminus \text{Stab}(E)$, let $l_g(E)$ be the (unique) point where any geodesic from a point in E to a point in $g \cdot E$ leaves E , and let \mathcal{L}_E be the set of all $l_g(E) \subset E$. Note that if $g \cdot E \cap E$ is nonempty (and $g \notin \text{Stab}(E)$) then $g \cdot E \cap E = \{l_g(E)\}$.

Since G is finitely generated, and hence countable, both sets \mathcal{D}_E and \mathcal{L}_E are countable. Given a (straight) line $p \subset E$, let χ_E^p be the projection from E to p . Since \mathcal{L}_E is countable, there is a line $p_E \subset E$ such that

- (1) the direction of p_E is not orthogonal to a direction in \mathcal{D}_E ;
- (2) if x and y are distinct points in \mathcal{L}_E , then $\chi_E^{p_E}(x) \neq \chi_E^{p_E}(y)$.

Project E onto p_E using $\chi_E^{p_E}$. The action of $\text{Stab}(E)$ on p_E is defined in the obvious way (using projection) – this is an action since the action of $\text{Stab}(E)$ on E is by translations. Connect $\mathcal{C}_\infty \setminus E$ to p_E in the obvious way – this uses the following:

Observation 6.1 Suppose S is a component of $\mathcal{C}_\infty \setminus E$. Then there is a (unique) point $x_S \in E$ so that S is a component of $\mathcal{C}_\infty \setminus \{x_S\}$.

Glue such a component S to p_E at the point $\chi_E^{p_E}(x_S)$.

Perform this projecting and gluing construction in an equivariant way for all flats $E \subseteq \mathcal{C}_\infty$ – so that for all $E \subseteq \mathcal{C}_\infty$ and all $g \in G$ the direction of the lines $p_{g.E}$ and $g.p_E$ is the same (this is possible since the action of $\text{Stab}(E)$ on E is by translations, so doesn't change directions).

Having done this for all flats $E \subseteq \mathcal{C}_\infty$, we arrive at a space T which we endow with the (obvious) path metric.

An isometric action of G on T is naturally induced by the action of G on X_ω .

The space T has a distinguished set of geodesic lines, namely those of the form p_E , for $E \in \mathcal{F}_\infty$. Denote the set of such geodesic lines by \mathbb{P} .

The following lemma is [24, Lemma 4.2], and the proof there holds in the current situation.

Lemma 6.2 *The space T is an \mathbb{R} -tree and has an isometric G -action with no global fixed point.*

Remark 6.3 Since $K_\infty \leq G$ acts trivially on \mathcal{C}_∞ , it also acts trivially on T , and the action of G on T induces an isometric action of L_∞ on T .

6.2 The actions of G and L_∞ on T

Let G be a finitely generated group, and Γ a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. Suppose that $\{h_i: G \rightarrow \Gamma\}$ is a sequence of pairwise nonconjugate homomorphisms. Let X_ω , \mathcal{C}_∞ and T be as in Section 5 and Section 5.1 and Section 6.1, respectively. Let $\{f_i: G \rightarrow \Gamma\}$ be the subsequence of $\{h_i\}$ as in the conclusion of Proposition 5.5. Let K_∞ be the kernel of the action of G on \mathcal{C}_∞ and let $L_\infty = G/K_\infty$ be the associated strict Γ -limit group.

Definition 6.4 [26, page 2] An action of a finitely generated group G on an \mathbb{R} -tree T is called *superstable* if whenever there are nontrivial arcs $I \not\subset J$ so that the stabilizers of I and J in G are different, the stabilizer of J is trivial.

Theorem 6.5 (cf [40, Lemma 1.3; 24, Theorem 4.4]) *In the above situation, the following properties hold.*

- (1) *Suppose that $[A, B]$ is a nondegenerate segment in T . Then $\text{Stab}_{L_\infty}[A, B]$ is an abelian subgroup of L_∞ .*
- (2) *If T is isometric to a real line then for all but finitely many n the group $f_n(G)$ is free abelian. Furthermore, in this case L_∞ is free abelian.*

- (3) If $g \in G$ fixes a tripod in T pointwise then $g \in \underline{\text{Ker}}(f_i)$.
- (4) Let $[y_1, y_2] \subset [y_3, y_4]$ be a pair of nondegenerate segments of T and assume that $\text{Stab}_{L_\infty}[y_3, y_4]$ is nontrivial. Then

$$\text{Stab}_{L_\infty}[y_1, y_2] = \text{Stab}_{L_\infty}[y_3, y_4].$$

In particular, the action of L_∞ on the \mathbb{R} -tree T is superstable.

- (5) Let $g \in G \setminus K_\infty$. Then for all but finitely many n we have $g \notin \ker(f_n)$.
- (6) L_∞ is torsion-free.
- (7) If T is not isometric to a real line then $\{f_i\}$ is a stable sequence of homomorphisms.

Proof The proof of [24, Theorem 4.4] relies on a number of different results. In each case, we have an exact analogue in the setting of Theorem 6.5 here.

The results we need are: Proposition 5.10, Lemma 5.9, Lemma 5.6, Lemma 4.20, Lemma 4.21, Proposition 5.5, Proposition 4.22 and the fact that stabilisers in Γ of flats in X are malnormal (see Farb [19, Example 1, page 819]).

Given these results, the proof of [24, Theorem 4.4] applies directly.⁶ The only change is that some of the constants have changed, so some of the counting has to be changed. This is straightforward. \square

We need the following lemma later when we describe the shortening argument.

Lemma 6.6 [24, Lemma 4.5] Suppose $\alpha, \beta \in X$ and $g \in G$ are such that there is a segment of length at least

$$6\phi(4\delta) + 4 \max\{d_X(g\alpha, \alpha), d_X(g\beta, \beta)\}$$

in a geodesic $[\alpha, \beta]$ which is within δ of a flat $E \in \mathcal{Q}$. Then $g \in \text{Fix}(E)$.

Proof Given Theorem 4.16 and Lemma 4.13, the proof of [24, Lemma 4.5] applies without change. \square

The following are two immediate applications of the above construction of the \mathbb{R} -tree T , and of Theorem 6.5. See Groves [24] for proofs which apply without change in the current setting.

⁶Note in particular that the analogue of [24, Lemma 4.5] holds in this setting. We state this separately below.

Theorem 6.7 (cf [24, Theorem 5.1]) *Suppose that Γ is a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. A group L is a Γ -limit group in the sense of Definition 1.2 if and only if it is a Γ -limit group in the sense of Definition 5.8.*

Theorem 6.8 (cf [24, Theorem 5.9]) *Suppose that Γ is a torsion-free group which is hyperbolic relative to a collection of finitely generated free abelian subgroups, and suppose that $\text{Out}(\Gamma)$ is infinite. Then Γ admits a nontrivial splitting over a finitely generated free abelian group.*

The results of [9] now imply the following result. (This is straightforward to prove using Theorem 6.5; see Sela [40].)

Lemma 6.9 (cf [24, Corollary 5.7]) *Suppose that Γ is a torsion-free group hyperbolic relative to a collection of finitely generated free abelian subgroups, and suppose that L is a Γ -limit group. Then*

- (1) *any finitely generated subgroup of L is a Γ -limit group;*
- (2) *L is torsion-free;*
- (3) *L is commutative-transitive and CSA;*
- (4) *every solvable subgroup of L is abelian.*

7 The shortening argument

Definition 7.1 (Dehn twists) Let G be a finitely generated group. A *Dehn twist* is an automorphism of one of the following two types:

- (1) Suppose that $G = A *_C B$ and that c is contained in the centre of C . Then define $\phi \in \text{Aut}(G)$ by $\phi(a) = a$ for $a \in A$ and $\phi(b) = bcb^{-1}$ for $b \in B$.
- (2) Suppose that $G = A *_C$, that c is in the centre of C , and that t is the stable letter of this HNN extension. Then define $\phi \in \text{Aut}(G)$ by $\phi(a) = a$ for $a \in A$ and $\phi(t) = tc$.

Definition 7.2 (Generalised Dehn twists) Suppose G has a graph of groups decomposition with abelian edge groups, and A is an abelian vertex group in this decomposition. Let $A_1 \leq A$ be the subgroup generated by all edge groups connecting A to other vertex groups in the decomposition. Any automorphism of A that fixes A_1 elementwise can be naturally extended to an automorphism of the ambient group G . Such an automorphism is called a *generalised Dehn twist* of G .

Definition 7.3 Let G be a finitely generated group. We define $\text{Mod}(G)$ to be the subgroup of $\text{Aut}(G)$ generated by

- (1) inner automorphisms;
- (2) Dehn twists arising from splittings of G with abelian edge groups;
- (3) generalised Dehn twists arising from graph of groups decompositions of G with abelian edge groups.

Similar definitions are made in [40, Section 5] and [4, Section 1].

Suppose that Γ is a torsion-free group which is hyperbolic relative to abelian subgroups, acting by isometries on the space X constructed in Section 4, with basepoint $x \in X$. Suppose also that G is a finitely generated group, with finite generating set \mathcal{A} . Let $h: G \rightarrow \Gamma$ be a homomorphism. Recall that in Section 5 we defined the *length* of h by

$$\|h\| := \max_{g \in \mathcal{A}} \{d_X(x, h(g).x)\}.$$

Definition 7.4 (cf [4, Definition 4.2]) We define an equivalence relation on the set of homomorphisms $h: G \rightarrow \Gamma$ by setting $h_1 \sim h_2$ if there is $\alpha \in \text{Mod}(G)$ and $\gamma \in \Gamma$ so that $h_1 = \tau_\gamma \circ h_2 \circ \alpha$, where τ_γ is the inner automorphism of Γ induced by γ .

A homomorphism $h: G \rightarrow \Gamma$ is *short* if for any h' such that $h \sim h'$ we have $\|h\| \leq \|h'\|$.

The following is one of the main technical results of this paper.

Theorem 7.5 (Shortening argument) *Suppose that Γ is a nonabelian, freely indecomposable, torsion-free group which is hyperbolic relative to abelian subgroups, and suppose that the sequence of automorphisms $\{h_n: \Gamma \rightarrow \Gamma\}$ converges to a faithful action $\eta: \Gamma \rightarrow \text{Isom}(\mathcal{C}_\infty)$ as above. Then for all but finitely many n the homomorphism h_n is not short.*

In order to “shorten” arbitrary homomorphisms, rather than just automorphisms, we need to introduce two new “bending” moves. This is undertaken in [23] (using ideas inspired by Alibegović [2]).

We now show how Theorem 7.5 implies Theorem 0.1.

Proof of Theorem 0.1 (assuming Theorem 7.5) If Γ is abelian then the theorem is clear, since in this case $\text{Mod}(\Gamma) = \text{Aut}(\Gamma)$. Thus we assume that Γ is nonabelian.

For each coset $C_i = \rho_i \text{Mod}(\Gamma)$ of $\text{Mod}(\Gamma)$ in $\text{Aut}(\Gamma)$ choose a representative $\hat{\rho}_i$ which is shortest amongst all representatives of C_i . That is to say, each of the automorphisms $\hat{\rho}_i$ is short.

However, by Theorem 7.5 we cannot have an infinite sequence $\{\widehat{\rho}_n: \Gamma \rightarrow \Gamma\}$ of nonequivalent short automorphisms, since then some subsequence will converge to a faithful action of Γ on a space \mathcal{C}_∞ . Hence $\text{Mod}(\Gamma)$ has finite index in $\text{Aut}(\Gamma)$ as required. \square

The remainder of this paper is devoted to proving Theorem 7.5. Before launching into the proof of Theorem 7.5, we need to recall some of the theory of groups acting on \mathbb{R} -trees.

8 Isometric actions on \mathbb{R} -trees

In this section we recall a result of Sela from [39]. Given a finitely generated group G and an \mathbb{R} -tree T with an isometric G -action, Theorem 8.1 below gives a decomposition of T which induces a graph of groups decomposition of G . In the case that G is finitely presented, this result follows immediately from Rips Theory; see Bestvina and Feighn [5].

There are two sets of terminology in English for the components of the above-mentioned decomposition⁷. Since we are quoting Sela's result, we use his (Rips') terminology. However, we assume that all axial components are isometric to a real line. Using Rips and Sela's definition of axial [38, Section 10], one other case could arise in the arguments that follow (where our group splits as $A *_{[a,b]} \langle a, b \rangle$). Just as noted in [38, Section 4, page 346], we can treat this case as an IET component. Thus, without further mention, we consider all axial components to be isometric to a real line.

The following theorem of Sela is used to decompose our limiting \mathbb{R} -trees. We remark that Sela claimed the following result for stable actions (a weaker condition than superstability). Guirardel [26] provided a counterexample to this, but pointed out that Sela's proof works under the assumption of superstability.

Theorem 8.1 ([39, Theorem 3.1]; see also [40, Theorem 1.5]) *Let G be a freely indecomposable finitely generated group which admits a superstable isometric action on a real tree Y . Assume that the stabiliser in G of each tripod in T is trivial.*

- (1) *There exist canonical orbits of subtrees of T , denoted T_1, \dots, T_k , with the following properties:*
 - (i) *For each $g \in G$ and each $i, j \in \{1, \dots, k\}$ with $i \neq j$, the subtree $g.T_i$ intersects T_j in at most a single point.*

⁷There is also the work in French by Gaboriau, Levitt and Paulin, with its attendant French terminology; see Paulin [36], for example.

- (ii) For each $g \in G$ and each $i \in \{1, \dots, k\}$, the subtree $g.T_i$ is either equal to T_i or intersects T_i in at most a single point.
 - (iii) The action of $\text{Stab}_G(T_i)$ on T_i is of axial or IET type.
- (2) G is the fundamental group of a graph of groups with
- (i) vertices corresponding to orbits of branching points with nontrivial stabiliser in T ;
 - (ii) vertices corresponding to the orbits of the canonical subtrees T_1, \dots, T_k which are of axial or IET type. The groups associated with these vertices are conjugates of the stabilisers of these subtrees. To a stabiliser of an IET component there exists an associated 2-orbifold, \mathcal{O} . Any element of $\pi_1(\mathcal{O})$ which corresponds to a boundary component or branching point in \mathcal{O} stabilises a point in T . For each stabiliser of an IET subtree we add edges that connect the vertex stabilised by it and the vertices stabilised by its boundary components and branching points;
 - (iii) edges corresponding to orbits of edges between branching points with non-trivial stabiliser in the discrete part of T (see Terminology 8.2 below) with edge groups which are conjugates of the stabilisers of these edges;
 - (iv) edges corresponding to orbits of points of intersection between the orbits of T_1, \dots, T_k .

Terminology 8.2 Let G and T be as in Theorem 8.1 above. The *discrete part* of T is the union of the metric closures of the connected components of $T \setminus (\bigcup_{i=1}^k GT_i)$.

Remark 8.3 In the theory of stable isometric actions on \mathbb{R} -trees, there is one further type of component arising in the decomposition of T . This is called “thin” in [5] and was discovered and investigated by Levitt [32]. However, in case G is freely indecomposable and the stabiliser of any nondegenerate tripod is trivial (both of these conditions hold in all cases considered in this paper), the limiting \mathbb{R} -tree has no thin components.

9 The shortening argument – Outline

In this section we outline the proof of Theorem 7.5 (the complete proof is contained in this and the subsequent two sections):

Theorem 7.5 (Shortening argument) *Suppose that Γ is a nonabelian, freely indecomposable, torsion-free relatively hyperbolic group with abelian parabolics, and suppose that the sequence of automorphisms $\{h_n: \Gamma \rightarrow \Gamma\}$ converges to an action $\eta: \Gamma \rightarrow \text{Isom}(\mathcal{C}_\infty)$ as above. Then for all but finitely many n the homomorphism h_n is not short.*

Remark 9.1 Although we call the above theorem the “shortening argument”, at least for hyperbolic groups the shortening argument is really a collection of ideas applicable in myriad situations. The above theorem is enough to prove Theorem 0.1. In order to build Makanin–Razborov diagrams in [23], we need to be able to shorten homomorphisms which are not surjective. The difficulty is that the images of the homomorphisms may not be “denser and denser” in the rescaled metric. In the general case, precomposition by elements of $\text{Mod}(\Gamma)$ is not sufficient and extra “bending” moves are required to shorten (see also Alibegović [2]). This is the reason why Theorem 10.2 and Theorem 11.1 below are stated for automorphisms.

Let $\{h_n: \Gamma \rightarrow \Gamma\}$ be a sequence of pairwise nonconjugate automorphisms. Since Γ is nonabelian, the action of Γ on the limiting space \mathcal{C}_∞ is faithful, and the action of Γ on the associated \mathbb{R} -tree T is also faithful. We prove that for all but finitely many n , the homomorphism h_n is not short.

Since the action of Γ on T is faithful, Γ is a strict Γ -limit group, and by Theorem 6.5 (3) the stabiliser in Γ of any tripod in T is trivial.

The approach to proving Theorem 7.5 is as follows: we consider the finite generating set \mathcal{A}_1 of Γ , and the basepoint y of T . We consider the paths $[y, u, y]$ where $u \in \mathcal{A}_1$. These paths can travel through various types of subtrees of T ; the IET subtrees, the axial subtrees, and the discrete part of T .⁸ Depending on the types of subtrees which have positive length intersection with $[y, u, y]$, we need various types of arguments which allow us to shorten the homomorphisms which “approximate” the action of Γ on \mathcal{C}_∞ .

Mostly, we follow the shortening argument as developed in [38]. There are two main obstacles to implementing this strategy in the context of torsion-free relatively hyperbolic groups with abelian parabolics. Note that the automorphisms $h_n: \Gamma \rightarrow \Gamma$ actually approximate the action of Γ on \mathcal{C}_∞ , from which the action of Γ on T is extracted. The two main impediments are: (i) those lines $p_E \in \mathbb{P}$ which correspond to flats $E \in \mathcal{C}_\infty$; and (ii) that triangles in the approximating spaces are only relatively thin, not actually thin.

9.1 IET components

The following theorem of Rips and Sela deals with IET components.

Theorem 9.2 [38, Theorem 5.1, pages 346–347] *Let G be a finitely presented freely indecomposable group and assume that $G \times T \rightarrow T$ is a small stable action of G on*

⁸By Remark 8.3, there are no thin components in the decomposition of T .

a real tree T with trivial stabilisers of tripods. Let U be a finite subset of G and let $y \in T$. Then there exists $\phi_I \in \text{Mod}(G)$ such that for any $u \in U$, if $[y, u(y)]$ has an intersection of positive length with some IET-component of T then

$$d_Y(y, \phi_I(u) \cdot y) < d_T(y, u(y)),$$

and otherwise $\phi_I(u) = u$.

It is worth noting that in [38] a more restrictive class of automorphisms is used to shorten the homomorphisms. Since it is a more restrictive class, their results still hold using our definition of $\text{Mod}(G)$.

Proposition 9.3 *Suppose that Y is an IET subtree of T and that $p_E \in \mathbb{P}$ is a line in T . Then the intersection $Y \cap p_E$ contains at most a point.*

Proof Since Y is an IET subtree, if σ is a nondegenerate arc in Y and $\epsilon > 0$ then there exists $\gamma \in \text{Stab}(Y)$ so that $\gamma \cdot \sigma \cap \sigma$ has positive length and there is some $x \in \sigma$ such that $d_T(x, \gamma \cdot x) < \epsilon$.

Suppose that $Y \cap p_E$ contains more than a point. By the above remark there exists $\gamma \in \text{Stab}(Y)$ for which $\gamma \cdot p_E \cap p_E$ contains more than a point. Hence $\gamma \cdot p_E = p_E$, and $p_E \subset Y$. This, combined with the above fact about IET components, implies that the action of $\text{Stab}(p_E)$ on p_E is indiscrete. However, this implies that it contains a noncyclic free abelian group, which cannot be a subgroup of $\text{Stab}(Y)$ when Y is an IET subtree. This contradiction proves the proposition. \square

Corollary 9.4 *Let T be an \mathbb{R} -tree arising from some \mathcal{C}_∞ as above. Suppose Y is an IET subtree of T and $\sigma \subset Y$ is a nondegenerate segment. Then there is a segment $\hat{\sigma} \subset \mathcal{C}_\infty$, of the same length as σ , which corresponds to σ under the projection from \mathcal{C}_∞ to T .*

9.2 Non-IET subtrees, technical results and the proof of Theorem 7.5

An entirely analogous argument to that of Proposition 9.3 proves:

Proposition 9.5 *Suppose that a line $l \subset T$ is an axial subtree and the line $p_E \subset T$ is associated to a flat $E \subset \mathcal{C}_\infty$. If $l \cap p_E$ contains more than a point then $l = p_E$.*

Corollary 9.6 *Let T be an \mathbb{R} -tree arising from some \mathcal{C}_∞ as above. Suppose l is an axial component of T so that $l \not\subset \mathbb{P}$ and $\sigma \subset l$ is a nondegenerate segment. Then there is a segment $\hat{\sigma} \subset \mathcal{C}_\infty$, of the same length as σ , which corresponds to σ under the projection from \mathcal{C}_∞ to T .*

Lemma 9.7 *If an edge e in the discrete part of T has an intersection of positive length with some line p_E then $e \subset p_E$.*

Proof Suppose that e contains a nontrivial segment from p_E but that $e \not\subset p_E$. Let C be the edge stabiliser of e . Since Γ is freely indecomposable, C is nontrivial, and since Γ is torsion-free, C is infinite. Let $\gamma \in C$. Then γ leaves more than one point of p_E invariant, so leaves all of p_E invariant. Thus γ leaves $E \subset C_\infty$ invariant. Also, since $e \not\subset p_E$, γ leaves some point $v \in C_\infty \setminus E$ invariant.

By Proposition 5.10, if $\{E_i\}$ is a sequence of flats ($E_i \subset X_i$) which converges to E , then for all but finitely many n the element $h_n(\gamma)$ leaves E_n invariant. By choosing an n large enough, $h_n(\gamma).E_n = E_n$, and furthermore if $\{v_i\}$ represents v , then $h_n(\gamma)$ moves v_n a distance which is much smaller than the distance from v_n to E_n . In particular, we can ensure that the geodesic $[v_n, h_n(\gamma).v_n]$ does not intersect the 4δ -neighbourhood of E_n . Then by Proposition 4.22, if $\pi: X_n \rightarrow E_n$ is the projection map then $d_X(\pi(v_n), \pi(h_n(\gamma).v_n)) \leq N_3(\phi(3\delta) + N_1(\phi(\delta)))$. However, since $h_n(\gamma).E_n = E_n$, we know that $\pi(h_n(\gamma).v_n) = h_n(\gamma).\pi(v_n)$. Thus $h_n(\gamma)$ moves $\pi(v_n)$ a distance at most $N_3(\phi(3\delta) + N_1(\phi(\delta)))$.

Repeating this argument with a large enough subset of C (a subset larger than the maximal size of an intersection of any orbit $\Gamma.u$ with a ball of radius $N_3(\phi(3\delta) + N_1(\phi(\delta)))$), we obtain a (finite) bound on the size of C . However, C is infinite, as noted above. This contradiction finishes the proof. \square

The following Theorems 10.1, 10.2 and 11.1 are the technical results needed to prove Theorem 7.5.

Theorem 10.1 *Let G be a finitely generated freely indecomposable group and assume that $G \times T \rightarrow T$ is a small stable action of G on an \mathbb{R} -tree T with trivial stabilisers of tripods. Let U be a finite subset of G and let $y \in T$. Then there exists $\phi_A \in \text{Mod}(G)$ so that for any $u \in U$, if $[y, u.y]$ has an intersection of positive length with some axial component of T then*

$$d_T(y, \phi_A(u).y) < d_T(y, u.y),$$

and otherwise $\phi_A(u) = u$.

As far as I am aware, Theorem 10.1 has not appeared in print. However, its statement and proof are very similar to those of Theorem 9.2, and it is certainly known at least to Sela [40, Section 5] and to Bestvina and Feighn [4, Exercise 11].

Remark 9.8 Theorem 9.2 is stated for finitely presented groups. The only time in the proof when it is required that G be finitely presented rather than just finitely generated is when a result of Morgan from [33] is quoted.

Specifically, Rips and Sela show that when G is freely indecomposable and finitely presented and acts on an \mathbb{R} -tree T with trivial tripod stabilisers then, for each $g \in G$ and any $y \in T$, the path $[y, \gamma \cdot y]$ cuts only finitely many components of axial or IET type in Y (see Rips and Sela [38, pages 350–351]).

However, this is also true when G is only assumed to be finitely generated, rather than finitely presented (but all other assumptions apply). This follows from the arguments in [39, Section 3]. The action of G on T can be approximated by actions of finitely presented groups G_i on \mathbb{R} -trees Y_i . For large enough k , the IET and axial components of Y_k map isometrically onto the IET and axial components of T (see Sela [39, Section 3] for details).

Therefore, Theorem 9.2 stills holds when G is assumed to be finitely generated, but not necessarily finitely presented. Similarly, Theorem 10.1 above, whose proof mimics the proof of Theorem 9.2, holds for finitely generated groups G . However, in this paper we can assume G is finitely presented.

We now state the further technical results which are required for the proof of Theorem 7.5. These technical results are proved in the subsequent two sections.

Theorem 10.2 *Let Γ be a freely indecomposable, torsion-free, nonabelian relatively hyperbolic group with abelian parabolics. Suppose that $h_n: \Gamma \rightarrow \Gamma$ is a sequence of automorphisms converging to a faithful action of Γ on a limiting space \mathcal{C}_∞ and let T be the \mathbb{R} -tree associated to \mathcal{C}_∞ . Let U be a finite subset of Γ . Let $y \in T$, let $\hat{y} \in \mathcal{C}_\infty$ project to $y \in T$ and let $\{\hat{y}_m\}$ be a sequence representing \hat{y} . Let p_E be an axial component of T . There exists m_0 so that: for all $m \geq m_0$ there is $\phi_{p_E, m} \in \text{Mod}(\Gamma)$ so that for any $u \in U$, if $[y, u \cdot y]$ has an intersection of positive length with a line in the Γ -orbit of p_E then*

$$d_{X_m}(\hat{y}_m, (h_m \circ \phi_{p_E, m})(u)) \cdot \hat{y}_m < d_{X_m}(\hat{y}_m, h_m(u) \cdot \hat{y}_m),$$

and otherwise $\phi_{p_E, m}(u) = u$.

Theorem 11.1 *Let Γ be a freely indecomposable torsion-free relatively hyperbolic group with abelian parabolics. Suppose that $h_n: \Gamma \rightarrow \Gamma$ is a sequence of automorphisms converging to a faithful action Γ on a limiting space \mathcal{C}_∞ with associated \mathbb{R} -tree T . Suppose further that U is a finite subset of Γ . Let $y \in T$, let $\hat{y} \in \mathcal{C}_\infty$ project to $y \in T$ and let $\{\hat{y}_m\}$ be a sequence representing \hat{y} . There exists m_0 so that: for all $m \geq m_0$ there is $\phi_{D, m} \in \text{Mod}(\Gamma)$ so that for any $u \in U$ which does not fix y and with $[y, u \cdot y]$ supported only in the discrete parts of T we have*

$$d_{X_m}(\hat{y}_m, (h_m \circ \phi_{D, m})(u) \cdot \hat{y}_m) < d_{X_m}(\hat{y}_m, u \cdot \hat{y}_m).$$

Armed with Proposition 9.3, and assuming Theorems 10.1, 10.2 and 11.1, we now prove Theorem 7.5.

Proof of Theorem 7.5 We have already noted that the action of Γ on T is faithful, that T is not isometric to a real line, and that the stabiliser in Γ of any tripod in T is trivial. Also, $\bigcap_{n \geq 1} \text{Ker}(h_n) = \{1\}$.

We suppose (by passing to a subtree if necessary) that the tree T is minimal. As noted in Remark 8.3 above, T contains no thin components.

Let $U = \mathcal{A}_1$ be the fixed generating set of Γ used to define $\|f\|$ for a homomorphism $f: \Gamma \rightarrow \Gamma$, let y be the image in T of the basepoint $x_\omega \in \mathcal{C}_\infty$ and let $\hat{y}_m = x$ for each m .

Let ϕ_I be the automorphism of Γ given by Theorem 9.2 and ϕ_A the automorphism from Theorem 10.1.

Suppose that $u \in U$ is such that $[y, u \cdot y]$ has an intersection of positive length with an IET component of T . Then Theorem 9.2 and Corollary 9.4 guarantee that for all but finitely many n we have $\|h_n \circ \phi_I\| < \|h_n\|$ so h_n is not short. Similarly, if $[y, u \cdot y]$ has an intersection of positive length with an axial component which is not contained in any $p_E \in \mathcal{P}$ then Theorem 10.1 and Corollary 9.6 imply that for all but finitely many n we have $\|h_n \circ \phi_A\| < \|h_n\|$ so also in this case h_n is not short.

Suppose then that $[y, u \cdot y]$ has an intersection of positive length with a line in the Γ -orbit of some p_E , and suppose that p_E is an axial component of T . Then by Theorem 10.2 for all but finitely many n there exists an automorphism $\phi_{p_E, n} \in \text{Mod}(\Gamma)$ so that $\|h_n \circ \phi_{p_E, n}\| < \|h_n\|$, so h_n is not short.

Finally, suppose that all of the segments $[y, u \cdot y]$ are entirely contained in the discrete part of T . Then by Theorem 11.1 for all but finitely many n there exists $\phi_{D, n} \in \text{Mod}(\Gamma)$ so that $\|h_n \circ \phi_{D, n}\| < \|h_n\|$, and once again h_n is not short.

This completes the proof of the theorem. \square

10 Axial components

The purpose of this section is to prove the following two theorems.

Theorem 10.1 *Let G be a finitely generated freely indecomposable group and assume that $G \times T \rightarrow T$ is a small stable action of G on an \mathbb{R} -tree T with trivial stabilisers of tripods. Let U be a finite subset of G and let $y \in T$. Then there exists $\phi_A \in \text{Mod}(G)$*

so that for any $u \in U$, if $[y, u \cdot y]$ has an intersection of positive length with some axial component of T then

$$d_T(y, \phi_A(u) \cdot y) < d_T(y, u \cdot y),$$

and otherwise $\phi_A(u) = u$.

Theorem 10.2 *Let Γ be a freely indecomposable, torsion-free, nonabelian relatively hyperbolic group with abelian parabolics. Suppose that $h_n: \Gamma \rightarrow \Gamma$ is a sequence of automorphisms converging to a faithful action of Γ on a limiting space \mathcal{C}_∞ and let T be the \mathbb{R} -tree associated to \mathcal{C}_∞ . Let U be a finite subset of Γ . Let $y \in T$, let $\hat{y} \in \mathcal{C}_\infty$ project to $y \in T$ and let $\{\hat{y}_m\}$ be a sequence representing \hat{y} . Let p_E be an axial component of T . There exists m_0 so that: for all $m \geq m_0$ there is $\phi_{p_E, m} \in \text{Mod}(\Gamma)$ so that for any $u \in U$, if $[y, u \cdot y]$ has an intersection of positive length with a line in the Γ -orbit of p_E then*

$$d_{X_m}(\hat{y}_m, (h_m \circ \phi_{p_E, m})(u) \cdot \hat{y}_m) < d_{X_m}(\hat{y}_m, h_m(u) \cdot \hat{y}_m),$$

and otherwise $\phi_{p_E, m}(u) = u$.

To prove Theorem 10.1 we follow the proof of [38, Theorem 5.1] (which is Theorem 9.2 in this paper). First, we need the following result, the (elementary) proof of which we include because of its similarity to Proposition 10.4 below.

Proposition 10.3 *Suppose that $\rho: P \times \mathbb{R} \rightarrow \mathbb{R}$ is an orientation-preserving, indiscrete isometric action of $P \cong \mathbb{Z}^n$ on the real line \mathbb{R} . For any finite subset W of P and any $\epsilon > 0$ there exists an automorphism $\sigma: P \rightarrow P$ such that*

(1) *for every $w \in W$ and every $r \in \mathbb{R}$*

$$d_{\mathbb{R}}(r, \sigma(w) \cdot r) < \epsilon;$$

(2) *for any $k \in \ker(\rho)$ we have $\sigma(k) = k$.*

Proof There is a direct product decomposition $P = A \oplus B$ where A is the kernel of the action of P on \mathbb{R} , and B is a finitely generated free abelian group which has a free, indiscrete and orientation preserving action on \mathbb{R} . The automorphism σ we define fixes A elementwise, so we can assume that all elements of W lie in B (since elements of A fix \mathbb{R} pointwise). Thus, we need only prove the lemma in case the action is faithful.

Since the action of B on \mathbb{R} is indiscrete and free, the translation lengths of elements of a basis of B are independent over \mathbb{Z} . In particular, there is a longest translation length amongst the translation lengths of a basis of B . Suppose that $b_1 \in B$ is the

element of the basis with largest translation length, and that b_2 has the second largest. Denote these translation lengths by $|b_1|$ and $|b_2|$, respectively. Since $|b_1|$ and $|b_2|$ are independent over \mathbb{Z} , there is $n \in \mathbb{Z}$ so that $0 < |b_1 + nb_2| < |b_2|$. Replace b_1 by $b_1 + nb_2$. This is an automorphism of P , fixing A elementwise.

Proceeding in this manner, we can make each of the elements of a basis as small as we wish, and so given W and $\epsilon > 0$, we can make each of the elements of W (considered as a word in the basis of B) have translation length less than ϵ , as required. \square

Proof of Theorem 10.1 By Remark 9.8, each of the segments $[y, u \cdot y]$ for $u \in U$ cuts only finitely many components of T of axial or IET type. Let ϵ be the minimum length of a (nondegenerate) interval of intersection between $[y, u \cdot y]$ and an axial component of T , for all $u \in U$.

The action of G on T induces a graph of groups decomposition Λ of G as in Theorem 8.1. Let T_i be an axial component of T . There is a vertex group of Λ corresponding to the G -orbit of T_i , with vertex group a conjugate of $\text{Stab}(T_i)$. By Theorem 6.5 and Lemma 6.9, $\text{Stab}(T_i)$ is a free abelian group. The vertex groups adjacent to $\text{Stab}(T_i)$ (those that are separated by a single edge) stabilise a point in the orbit of a branching point in T_i with nontrivial stabiliser. Recall that G is freely indecomposable, so all edge groups are nontrivial.

Let q_1 be the point on T_i closest to y (if $y \in T_i$ then $q_1 = y$). Choose points $q_2, \dots, q_m \in T_i$ in the orbits of the branching points corresponding to the adjacent vertex groups such that $d_T(q_i, q_j) < \epsilon/20$. We can do this since the action of $\text{Stab}(T_i)$ on T_i has all orbits dense, since T_i is an axial component.

The proof of Theorem 10.1 now proceeds exactly as the proof of [38, Theorem 5.1] (start with Case 1 on page 351). \square

Proof of Theorem 10.2 Since $p_E \in \mathbb{P}$ is an axial component of T , there is a vertex group corresponding to the conjugacy class of $\text{Stab}(p_E)$ in the graph of groups decomposition which the (faithful) action of Γ on T induces (see Theorem 8.1).

Now, the stabiliser in Γ of p_E is exactly the stabiliser in Γ of E , when Γ acts (also faithfully) on \mathcal{C}_∞ . By [24, Corollary 3.17], there is a sequence of flats E_i in the approximating spaces X_i so that $E_i \rightarrow E$ in the Γ -equivariant Gromov topology. By Proposition 5.10, if $\gamma \in \text{Stab}_\Gamma(E)$ then for all but finitely i we have $h_i(\gamma) \in \text{Stab}(E_i)$. For such an i , the element $h_i(\gamma)$ is contained in a unique noncyclic maximal abelian subgroup A_i of Γ . However, h_i is an automorphism, so γ is contained in a unique noncyclic maximal abelian subgroup A_γ of Γ , and $A_i = h_i(A_\gamma)$.

If γ' is another element of $\text{Stab}_\Gamma(E)$, then $[\gamma, \gamma'] = 1$, and it is not difficult to see that $A_\gamma = A_{\gamma'}$. Also, if $\gamma_0 \in A_\gamma$ then $\gamma_0 \in \text{Stab}_\Gamma(E)$. Hence $A_\gamma = \text{Stab}_\Gamma(E)$. We denote the subgroup $\text{Stab}_\Gamma(E)$ by A_E .

We now prove Theorem 10.2 by finding an analogue of Proposition 10.3 in the flats E_i and then once again following the proof from [38].

Proposition 10.4 *Let W be a finite subset of A_E . For any $\epsilon > 0$ there exists i_0 so that for all $i \geq i_0$, there is an automorphism $\sigma_i: A_E \rightarrow A_E$ so that*

- (1) *for every $w \in W$, and every $r_i \in E_i$,*

$$d_{X_i}(r_i, h_i(\sigma_i(w)) \cdot r_i) < \epsilon;$$

- (2) *for any $k \in A_E$ which acts trivially on E we have $\sigma_i(k) = k$.*

Proof of Proposition 10.4 The group A_E admits a decomposition $A_E = A_0 \oplus A_1$, where A_0 acts trivially on E , and A_1 acts freely on E . Choose a basis \mathcal{B} of A_E consisting of a basis for A_0 and a basis for A_1 . Let k_W be the maximum word length of any element of W with respect to the chosen basis.

Since the $h_i: \Gamma \rightarrow \Gamma$ are automorphisms, for sufficiently large i and any $a \in E_i$, the set $h_i(A_E) \cdot a \subset E_i$ forms an $\epsilon/(20k_W)$ -net in E_i (where distance is measured in the metric $1/\|h_i\|$ on X_i). Choose a (possibly larger) i so that also the action of $h_i(\mathcal{B})$ on E_i approximates the action of \mathcal{B} on E to within $\epsilon/(20k_W)$ (note that since the action of A_E on E and the action of $h_i(A_E)$ on E_i are both by translations, and translations of Euclidean space move every point the same distance, there are arbitrarily good approximations for the action of any finite subset of A_E on the whole of E).

The remainder of the proof proceeds just as the proof of Proposition 10.3 above, although in the step where we replace b_1 by $b_1 + nb_2$, we cannot insist that b_2 acts nontrivially on E . However, we of course can insist that b_1 acts nontrivially on E , since otherwise it moves all points of E a distance at most $\epsilon/20$. Therefore, such an automorphism is nonetheless a generalised Dehn twist. \square

Given Proposition 10.4, the proof of Theorem 10.2 once again follows the proof of [38, Theorem 5.1, pages 350–353], although in this case we have to choose approximations to the action of Γ on \mathcal{C}_∞ (the important point here is that the sets $h_i(A_E) \cdot a$, for any $a \in E_i$, get denser and denser in E_i , when considered in the scaled metric $(1/\|h_i\|)d_X$ of X_i). These small changes are straightforward, but do lead to the different shortening automorphisms $\phi_{p_E, m}$ in the statement of Theorem 10.2. \square

11 The discrete case

In this section we shorten the approximations to paths of the form $[\hat{y}, u.\hat{y}]$, where $\hat{y} \in \mathcal{C}_\infty$ projects to $y \in T$ and $[y, u.y]$ is entirely supported in the discrete part of T . The lengths of the limiting paths $[\hat{y}, u.\hat{y}]$ and $[y, u.y]$ are unchanged.

The purpose of this section is to prove the following:

Theorem 11.1 *Let Γ be a freely indecomposable torsion-free relatively hyperbolic group with abelian parabolics. Suppose that $h_n: \Gamma \rightarrow \Gamma$ is a sequence of automorphisms converging to a faithful action Γ on a limiting space \mathcal{C}_∞ with associated \mathbb{R} -tree T . Suppose further that U is a finite subset of Γ . Let $y \in T$, let $\hat{y} \in \mathcal{C}_\infty$ project to $y \in T$ and let $\{\hat{y}_m\}$ be a sequence representing \hat{y} . There exists m_0 so that: for all $m \geq m_0$ there is $\phi_{D,m} \in \text{Mod}(\Gamma)$ so that for any $u \in U$ which does not fix y and with $[y, u.y]$ supported only in the discrete parts of T we have*

$$d_{X_m}(\hat{y}_m, (h_m \circ \phi_{D,m})(u).\hat{y}_m) < d_{X_m}(\hat{y}_m, u.\hat{y}_m).$$

The proof of Theorem 11.1 follows [38, Section 6].

By Lemma 9.7, if e is a discrete edge in T then either $e \in p_E$ for some flat $E \subset \mathcal{C}_\infty$, or \mathcal{C}_∞ contains a well-defined, canonical, isometric image \hat{e} of e , so that \hat{e} projects to e .

We have a sequence of automorphisms $\{h_n: \Gamma \rightarrow \Gamma\}$, converging to a faithful action of Γ on a limiting space \mathcal{C}_∞ , with associated \mathbb{R} -tree T .

There are a number of different cases to consider:

Case 1 y is contained in the interior of an edge e .

Case 1a e is not completely contained in a line of the form p_E and $\bar{e} \in T/\Gamma$ is a splitting edge.

Note that because e is not contained in any p_E , there is a single point $\hat{y} \in \mathcal{C}_\infty$ which corresponds to $y \in T$.

This case is very similar to the Case 1a on pages 355–356 of [38]. In this case we have a decomposition $\Gamma = A *_C B$ where C is a finitely generated free abelian group properly contained in both A and B .

Given $u \in U$ we can write

$$u = a_u^1 b_u^1 \cdots a_u^{n_u} b_u^{n_u},$$

where $a_u^i \in A$ and $b_u^i \in B$ (possibly a_u^1 and/or $b_u^{n_u}$ are the identity). Let $\{z_1, \dots, z_n\}$ be a generating set for Z .

Let ϵ be the minimum of

- (1) the length of the shortest edge in the discrete part of T ;
- (2) the distance between y and the vertices of e .

Recall that triangles in X are relatively δ -thin, and the function ϕ comes from the definition of isolated flats. Let C_0 be the maximum size of an intersection of an orbit $\Gamma.z$ with a ball of radius $10\delta + 2\phi(3\delta)$ in X (where distance is measured in d_X).

Now take F to be the finite subset of G containing 1 and

$$za_u^i z^{-1}, \quad zb_u^i z^{-1},$$

where $z \in C$ has word length at most $10C_0$.

For large enough m we have, for all $\gamma_1, \gamma_2 \in F$,

$$(1) \quad |d_{X_m}(h_m(\gamma_1) \cdot \hat{y}_m, h_m(\gamma_2) \cdot \hat{y}_m) - d_{C_\infty}(\gamma_1 \cdot \hat{y}, \gamma_2 \cdot \hat{y})| < \epsilon_1,$$

where $\epsilon_1 = \epsilon / (8000C_0)$.

Let $w_m \in [\hat{y}_m, h_m(a_u^1) \cdot \hat{y}_m]$ and $w'_m \in [\hat{y}_m, h_m(b_u^1) \cdot \hat{y}_m]$ satisfy

$$(2) \quad \frac{\epsilon}{2} - \frac{\epsilon}{1000} \leq d_{X_m}(w_m, \hat{y}_m) = d_{X_m}(w'_m, \hat{y}_m) \leq \frac{\epsilon}{2} + \frac{\epsilon}{1000}.$$

Lemma 11.2 For some $z \in C$ of word length at most $10C_0$ we have, for all but finitely many m ,

$$d_{X_m}(\hat{y}_m, h_m(z) \cdot w_m) < d_{X_m}(\hat{y}_m, w_m) - 8\delta_m,$$

and
$$d_{X_m}(\hat{y}_m, h_m(z) \cdot w'_m) < d_{X_m}(\hat{y}_m, w'_m) + 8\delta_m.$$

Proof Let W be the set of all elements $z \in C$ of word length at most $10C_0$ in the generators $\{z_1, \dots, z_n\}$ and their inverses.

First suppose that for all but finitely many i we have $h_i(W) \subseteq \text{Stab}_\Gamma(E_i)$. Then since the edge containing y is not completely contained in a single p_E , we can assume that each element of W fixes a point outside of E . Now, using Proposition 4.22, there is a point in E_i which is moved at most $N_3(\phi(3\delta) + N_1(\phi(\delta)))$ by each element of $h_i(W)$. This gives a bound on the size of $h_i(W)$ which does not depend on i (so long as i is large enough). However, this contradicts the choice of $W \subseteq \Gamma$. Therefore it is not the case that $h_i(W) \subseteq \text{Stab}_\Gamma(E_i)$ for all but finitely many i .

By the argument in the paragraphs after the proof of [24, Lemma 4.5], for all but finitely many k , the elements $h_i(z)$ act approximately like translations. Since W is closed under inverses, and we have chosen W large enough that some element “translates” by at least $10\delta_m$, we can choose some $z \in W$ which satisfies the conclusion of the lemma. \square

In order to finish Case 1a, we follow the proofs of Proposition 6.3 and Theorem 6.4 from [38]. The only additional thing needed in this case is to force w_m to lie close to each $[\hat{y}_m, h_m(a_u^i) \cdot \hat{y}_m]$. We do this by applying Lemma 6.6 and the arguments in the paragraphs in [24] which follow the proof of [24, Lemma 4.5]. It is for this reason that we left some flexibility as to the choice of w_m and w'_m in (2) above.

In this manner, we can ensure that w_m lies within 2δ of each geodesic segment $[\hat{y}_m, h_m(a_u^i) \cdot \hat{y}_m]$, and similarly for w'_m . We can now follow the proof [38, Proposition 6.3]. The proof of [38, Theorem 6.4] is not included in [38] (or in [39] as claimed in [38]). However, it is straightforward, so we omit it here also.

The automorphism we use to shorten in this case is

$$\begin{aligned} \forall a \in A \quad \phi(a) &= zaz^{-1}, \\ \forall b \in B \quad \phi(b) &= z^{-1}bz, \end{aligned}$$

where a is as in Lemma 11.2. This completes the proof in Case 1a. It is worth noting here that we are shortening the actions on X_i which approximate the action on \mathcal{C}_∞ . However, this does not affect the analogy between the proofs here and those in [38].

Case 1b e is not completely contained in a single p_E and $\bar{e} \in T/\Gamma$ is not a splitting edge.

In this case we have a decomposition $\Gamma = A * C$, where C is a finitely generated free abelian group.

In the same way as we adapted the proof of Case 1a from [38] above, we may adapt the proof of Case 1b from [38]. The key point is that we allow a small amount of flexibility in the choice of w_m and w'_m . Doing this, we may ensure that even though the approximating triangles we consider are only relatively thin, rather than actually thin, all of the features we need to apply the proof from [38] still hold, because we can make sure that we are not near the “fat” part of any triangle. Proceeding with this idea in mind, the proof from [38] can be adapted without difficulty.

We now deal with the two cases where y is contained in the interior of the edge e and $e \subset p_E$ for some $p_E \in \mathbb{P}$. Using Lemma 4.21 and Proposition 4.22, the following result is not difficult to prove:

Proposition 11.3 *Suppose that X is the space constructed in Section 4. There exists a constant N_4 , depending only on X so that if $E_1, E_2 \in \mathcal{Q}$ are maximal flats in X then there is a set J_{E_1, E_2} so that*

- (1) $\text{Diam}(J_{E_1, E_2}) \leq N_4$;
- (2) *if $x \in E_1$ and $u \in E_2$ then any geodesic between x and y intersects J_{E_1, E_2} nontrivially.*

Recall that \mathcal{Q} is the family of maximal flats from the definition of X , that triangles in X are relatively δ -thin (Theorem 4.16), and that ϕ is the function from Lemma 4.13. Following Convention 4.14, we assume without loss of generality that for all $k \geq 0$ we have $\phi(k) \geq k$ and also that ϕ is a nondecreasing function.

Choose compact fundamental domains for the action of $\text{Stab}_\Gamma(E)$ on E , for each conjugacy class of maximal flat in X , and let K_F be the maximal diameter of these fundamental domains. Also, let K_X be the diameter of a compact set D for which $\Gamma \cdot D = X$. For the remainder of Case 1, we replace the constant δ by

$$\max \{ \delta, 1000K_F, 1000K_X, 1000(7\delta + 14\phi(4\delta)) \}.$$

The stabiliser of the edge e is a subgroup of $\text{Stab}_\Gamma(E)$. Since p_E is not an axial component, the action of $\text{Stab}(E)$ on E is either trivial or factors through a infinite cyclic group. If $\bar{e} \in T/\Gamma$ is a splitting edge, then necessarily the action of $\text{Stab}(E)$ on E is trivial.

Case 1c e is completely contained in some p_E , and $\bar{e} \in T/\Gamma$ is a splitting edge.

Let $A_E = \text{Stab}_\Gamma(E)$. Then, $A_E = A_0 \oplus A_1$, where A_0 acts trivially on E and A_1 acts freely on E . Since p_E is a splitting edge, $A_1 = \{1\}$.

We have a decomposition $\Gamma = H_1 *_{A_E} H_2$.

The subgroup H_1 fixes a point in p_E , but does not fix all of p_E . Thus, H_1 fixes a point $v_1 \in E$. Similarly, H_2 fixes a point $v_2 \in E$, but does not fix all of E . We choose points $\hat{y}_m \subset E_m$ so that (i) $\{\hat{y}_m\}$ represents $\hat{y} \in \mathcal{C}_\infty$ which projects to $y \in T$; (ii) each \hat{y}_m lies in the orbit $\Gamma \cdot x$; and (iii) subject to the first two conditions, \hat{y}_m lies as close as possible to the line $[v_1^m, v_2^m]$, where $\{v_i^m\} \rightarrow v_i$, $i = 1, 2$.

We proceed as in Case 1a. However, this time we cannot find a single automorphism to shorten the $\|h_i\|$, but we use the fact that the sets $h_i(A_E) \cdot a \subset E_i$ are denser and denser (when distance is measured in the metrics $(1/\|h_i\|)d_X$) to find, for all but finitely many i , a Dehn twist $\phi_{e,i}$ which shortens the action on X_i . This proceeds in a similar way to Case 1a above, using the ideas in Proposition 10.4 and the proof of Theorem 10.2 above.

Case 1d e is completely contained in some p_E and $\bar{e} \in T/\Gamma$ is not a splitting edge.

There are two cases here. As in Case 1b, we have a decomposition $\Gamma = A *_C$, where C is a finitely generated free abelian group. Let t be the stable letter of this HNN extension, and suppose that $C \leq \text{Stab}(E)$, a maximal flat in \mathcal{C}_∞ . The two cases are where $f \in \text{Stab}(E)$, and when $f \notin \text{Stab}(E)$.

Each of these cases follow the proof of Case 1b above (and therefore Case 1b from [38]) in the same way as Case 1c followed the proof of Case 1a.

Case 2 y is a vertex of T .

In this case, we do not shorten the approximations to a particular edge, but each of the edges adjacent to y . As before, we largely follow [38, Section 6].

There are four cases again, when the edge is splitting, and nonsplitting, coupled with the cases where the edge is contained in some p_E and when it is not.

These follow the proofs from [38] just as in Case 1 above. Note that the shortening automorphisms fix elementwise $\text{Stab}_\Gamma(\hat{y})$.

Proof of Theorem 11.1 If y is contained in the interior of an edge, then apply Case 1 above to find a sequence of automorphisms which shorten the h_n .

If y is a vertex in T , then we shorten the h_n on each of the adjacent edges separately using Case 2 and [38, Section 6]. \square

This finally completes the proof of Theorem 7.5.

References

- [1] **E Alibegović**, *A combination theorem for relatively hyperbolic groups*, Bull. London Math. Soc. 37 (2005) 459–466 MR2131400
- [2] **E Alibegović**, *Makanin–Razborov diagrams for limit groups*, Geom. Topol. 11 (2007) 643–666 MR2302499
- [3] **M Bestvina**, *Degenerations of the hyperbolic space*, Duke Math. J. 56 (1988) 143–161 MR932860
- [4] **M Bestvina**, **M Feighn**, *Notes on Sela’s work: Limit groups and Makanin–Razborov diagrams*, Preprint, Available online at <http://www.math.utah.edu/~bestvina/eprints/notes1.pdf>
- [5] **M Bestvina**, **M Feighn**, *Stable actions of groups on real trees*, Invent. Math. 121 (1995) 287–321 MR1346208

- [6] **B Bowditch**, *Relatively hyperbolic groups*, Preprint Available at <http://www.warwick.ac.uk/~masgak/papers/bhb-relhyp.pdf>
- [7] **MR Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, *Grund. der Math. Wissenschaften* 319, Springer, Berlin (1999) MR1744486
- [8] **MR Bridson, G A Swarup**, *On Hausdorff–Gromov convergence and a theorem of Paulin*, *Enseign. Math. (2)* 40 (1994) 267–289 MR1309129
- [9] **C Champetier, V Guirardel**, *Limit groups as limits of free groups*, *Israel J. Math.* 146 (2005) 1–75 MR2151593
- [10] **I Chatterji, K Ruane**, *Some geometric groups with rapid decay*, *Geom. Funct. Anal.* 15 (2005) 311–339 MR2153902
- [11] **F Dahmani**, *Classifying spaces and boundaries for relatively hyperbolic groups*, *Proc. London Math. Soc. (3)* 86 (2003) 666–684 MR1974394
- [12] **F Dahmani**, *Combination of convergence groups*, *Geom. Topol.* 7 (2003) 933–963 MR2026551
- [13] **F Dahmani**, *Accidental parabolics and relatively hyperbolic groups*, *Israel J. Math.* 153 (2006) 93–127 MR2254639
- [14] **F Dahmani, D Groves**, *The isomorphism problem for toral relatively hyperbolic groups*, *Publ. Math. Inst. Hautes Études Sci.* (2008) 211–290 MR2434694
- [15] **L van den Dries, A J Wilkie**, *Gromov’s theorem on groups of polynomial growth and elementary logic*, *J. Algebra* 89 (1984) 349–374 MR751150
- [16] **C Druţu, M Sapir**, *Relatively hyperbolic groups with rapid decay property*, *Int. Math. Res. Not.* (2005) 1181–1194 MR2147058
- [17] **C Druţu, M Sapir**, *Tree-graded spaces and asymptotic cones of groups*, *Topology* 44 (2005) 959–1058 MR2153979 With an appendix by D Osin and Sapir
- [18] **C Druţu, M V Sapir**, *Groups acting on tree-graded spaces and splittings of relatively hyperbolic groups*, *Adv. Math.* 217 (2008) 1313–1367 MR2383901
- [19] **B Farb**, *Relatively hyperbolic groups*, *Geom. Funct. Anal.* 8 (1998) 810–840 MR1650094
- [20] **M Gromov**, *Groups of polynomial growth and expanding maps*, *Inst. Hautes Études Sci. Publ. Math.* (1981) 53–73 MR623534
- [21] **M Gromov**, *Hyperbolic groups*, from: “Essays in group theory”, (S M Gersten, editor), *Math. Sci. Res. Inst. Publ.* 8, Springer, New York (1987) 75–263 MR919829
- [22] **D Groves**, *Limits of (certain) CAT(0) spaces II: The Hopf property and the shortening argument* [arXiv:math.GR/0408080](https://arxiv.org/abs/math.GR/0408080)
- [23] **D Groves**, *Limit groups for relatively hyperbolic groups. II. Makanin–Razborov diagrams*, *Geom. Topol.* 9 (2005) 2319–2358 MR2209374

- [24] **D Groves**, *Limits of (certain) CAT(0) groups. I. Compactification*, *Algebr. Geom. Topol.* 5 (2005) 1325–1364 MR2171812
- [25] **D Groves, J F Manning**, *Dehn filling in relatively hyperbolic groups*, *Israel J. Math.* 168 (2008) 317–429 MR2448064
- [26] **V Guirardel**, *Actions of finitely generated groups on \mathbb{R} -trees*, *Ann. Inst. Fourier (Grenoble)* 58 (2008) 159–211 MR2401220
- [27] **G C Hruska**, *Nonpositively curved spaces with isolated flats*, PhD thesis, Cornell University (2002)
- [28] **G C Hruska, B Kleiner**, *Hadamard spaces with isolated flats*, *Geom. Topol.* 9 (2005) 1501–1538 MR2175151 With an appendix by the authors and M Hindawi
- [29] **M Kapovich, B Leeb**, *On asymptotic cones and quasi-isometry classes of fundamental groups of 3-manifolds*, *Geom. Funct. Anal.* 5 (1995) 582–603 MR1339818
- [30] **O Kharlampovich, A Myasnikov**, *Irreducible affine varieties over a free group. I, II*, *J. Algebra* 200 (1998) 472–516, 517–570 MR1610660
- [31] **O Kharlampovich, A Myasnikov**, *Elementary theory of free non-abelian groups*, *J. Algebra* 302 (2006) 451–552 MR2293770
- [32] **G Levitt**, *La dynamique des pseudogroupes de rotations*, *Invent. Math.* 113 (1993) 633–670 MR1231840
- [33] **J W Morgan**, *Ergodic theory and free actions of groups on \mathbb{R} -trees*, *Invent. Math.* 94 (1988) 605–622 MR969245
- [34] **D V Osin**, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, *Mem. Amer. Math. Soc.* 179 (2006) vi+100 MR2182268
- [35] **F Paulin**, *Outer automorphisms of hyperbolic groups and small actions on \mathbb{R} -trees*, from: “Arboreal group theory”, (R Alperin, editor), *MSRI Publications* 19 (1991) 331–341
- [36] **F Paulin**, *Actions de groupes sur les arbres*, from: “Séminaire Bourbaki, Vol. 1995/96, Exp. 808”, *Astérisque* 241 (1997) 97–137 MR1472536
- [37] **F Paulin**, *Sur la théorie élémentaire des groupes libres (d’après Sela)*, *Astérisque* 294 (2004) ix, 363–402 MR2111650
- [38] **E Rips, Z Sela**, *Structure and rigidity in hyperbolic groups. I*, *Geom. Funct. Anal.* 4 (1994) 337–371 MR1274119
- [39] **Z Sela**, *Acyindrical accessibility for groups*, *Invent. Math.* 129 (1997) 527–565 MR1465334
- [40] **Z Sela**, *Diophantine geometry over groups. I. Makanin–Razborov diagrams*, *Publ. Math. Inst. Hautes Études Sci.* (2001) 31–105 MR1863735

- [41] **Z Sela**, *Diophantine geometry over groups and the elementary theory of free and hyperbolic groups*, from: “Proceedings of the International Congress of Mathematicians, Vol. II (Beijing, 2002)”, Higher Ed. Press, Beijing (2002) 87–92 MR1957023
- [42] **Z Sela**, *Diophantine geometry over groups. II. Completions, closures and formal solutions*, Israel J. Math. 134 (2003) 173–254 MR1972179
- [43] **Z Sela**, *Diophantine geometry over groups. IV. An iterative procedure for validation of a sentence*, Israel J. Math. 143 (2004) 1–130 MR2106978
- [44] **Z Sela**, *Diophantine geometry over groups. III. Rigid and solid solutions*, Israel J. Math. 147 (2005) 1–73 MR2166355
- [45] **Z Sela**, *Diophantine geometry over groups. V₁. Quantifier elimination. I*, Israel J. Math. 150 (2005) 1–197 MR2249582
- [46] **Z Sela**, *Diophantine geometry over groups. V₂. Quantifier elimination. II*, Geom. Funct. Anal. 16 (2006) 537–706 MR2238944
- [47] **Z Sela**, *Diophantine geometry over groups. VI. The elementary theory of a free group*, Geom. Funct. Anal. 16 (2006) 707–730 MR2238945
- [48] **Z Sela**, *Diophantine geometry over groups VIII: The elementary theory of a hyperbolic group*, Preprint Available at <http://www.ma.huji.ac.il/~zlil/stab.pdf>
- [49] **A Szczepański**, *Relatively hyperbolic groups*, Michigan Math. J. 45 (1998) 611–618 MR1653287
- [50] **A Yaman**, *A topological characterisation of relatively hyperbolic groups*, J. Reine Angew. Math. 566 (2004) 41–89 MR2039323

Department of Mathematics, University of Illinois at Chicago
851 S Morgan St, Chicago, IL 60607-7045, USA

groves@math.uic.edu

<http://www.math.uic.edu/~groves>

Received: 20 March 2008 Revised: 17 December 2008