ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 9 (2005) 2319–2358
Published: 21 December 2005

Limit groups for relatively hyperbolic groups II: Makanin–Razborov diagrams

Daniel Groves

Department of Mathematics, California Institute of Technology
Pasadena, CA 91125, USA

Email: groves@caltech.edu

Abstract

Let Γ be a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. We construct Makanin–Razborov diagrams for Γ. We also prove that every system of equations over Γ is equivalent to a finite subsystem, and a number of structural results about Γ–limit groups.

AMS Classification numbers Primary: 20F65

Secondary: 20F67, 20E08, 57M07

Keywords: Relatively hyperbolic groups, limit groups, $\mathbb{R}$–trees

Proposed: Benson Farb
Seconded: Walter Neumann, Martin Bridson

Received: 15 March 2005
Accepted: 3 December 2005

1 Introduction

This paper is a continuation of [19]. Throughout this paper, Γ will denote a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. For an arbitrary finitely generated group G, we wish to understand the set $\mathrm{Hom}(G,\Gamma)$ of all homomorphisms from G to Γ.

In [19] we considered a sequence of pairwise non-conjugate homomorphisms $\{h_i \colon G \to \Gamma\}$ and extracted a limiting G–action on a suitable asymptotic cone, and then extracted an $\mathbb{R}$–tree with a nontrivial G–action. This $\mathbb{R}$–tree allows much information to be obtained. In particular, in case $G = \Gamma$, we studied $\mathrm{Aut}(\Gamma)$ and also proved that Γ is Hopfian. In this paper, we continue this study, in case G is an arbitrary finitely generated group. In particular, we construct a *Makanin–Razborov diagram* for G, which gives a parametrisation of $\mathrm{Hom}(G,\Gamma)$ (see Section 6 below). We build on our work from [19], which in turn builds on our previous work of [17] and [18]. The strategy is to follow [32, Section 1], though there are extra technical difficulties to deal with.

To a system of equations Σ over Γ in finitely many variables there is naturally associated a finitely generated group G_Σ, with generators the variables in Σ, and relations the equations. The solutions to Σ in Γ are in bijection with the elements of $\mathrm{Hom}(G_\Sigma,\Gamma)$. Thus, Makanin–Razborov diagrams give a description of the set of solutions to a given system of equations over Γ. For free groups, building on the work of Makanin and Razborov, Makanin–Razborov diagrams were constructed by Kharlampovich and Miasnikov [22, 23], and also by Sela [31]. For torsion-free hyperbolic groups, Makanin–Razborov diagrams were constructed by Sela [32], and it is Sela's approach that we follow here. Alibegović [2] constructed Makanin–Razborov diagrams for limit groups.

Limit groups are hyperbolic relative to their maximal non-cyclic abelian subgroups (see [9] and [1]). Limits groups are also torsion-free. Therefore, the main result of this paper (the construction of Makanin–Razborov diagrams) generalises the main result of [2]. Alibegović has another approach to the construction of Makanin–Razborov diagrams for these relatively hyperbolic groups (see [2, Remark 3.7]).

The main results of this paper are the following five:

Theorem 5.10 *Suppose that Γ is a torsion-free relatively hyperbolic group with abelian parabolics, and that G is a finitely generated group. Then G is a Γ-limit group if and only if G is fully residually Γ.*

Proposition 5.11 *Suppose that Γ is a torsion-free relatively hyperbolic group with abelian parabolics. Then there are only countably many Γ–limit groups.*

Theorem 5.16 *Let Γ be a torsion-free relatively hyperbolic group with abelian parabolics. Then every system of equations in finitely many variables over Γ (without coefficients) is equivalent to a finite subsystem.*

In Section 3 we define an equivalence relation '$\sim$' on the set of homomorphisms $\mathrm{Hom}(G,\Gamma)$, where G is an arbitrary finitely generated group and Γ is a torsion-free relatively hyperbolic group with abelian parabolics. This equivalence relation uses conjugation, elements of $\mathrm{Mod}(G) \leq \mathrm{Aut}(G)$ and 'bending' moves (see Section 3).

Theorem 1.1 *Let Γ be a torsion-free relatively hyperbolic group with abelian parabolics, and let G be a finitely generated freely indecomposable group. There is a finite collection $\{\eta_i\colon G \to L_i\}_{i=1}^n$ of proper quotients of G such that, for any homomorphism $h\colon G \to \Gamma$ which is not equivalent to an injective homomorphism there exists $h'\colon G \to \Gamma$ with $h \sim h'$, $i \in \{1,\dots,n\}$ and $h_i\colon L_i \to \Gamma$ so that $h' = h_i \circ \eta_i$.*

The quotient groups L_i that appear in Theorem 1.1 are Γ–*limit groups*. Theorem 1.1 allows us to reduce the description of $\mathrm{Hom}(G,\Gamma)$ to a description of $\{\mathrm{Hom}(L_i,\Gamma)\}_{i=1}^n$. We then apply Theorem 1.1 to each of the L_i in turn, and so on with successive proper quotients. That this process terminates follows from the following

Theorem 5.2 *Let Γ be a torsion-free group which is hyperbolic relative to free abelian subgroups. Every decreasing sequence of Γ–limit groups:*

$$R_1 > R_2 > R_3 > \dots,$$

which are all quotients of a finitely generated group G, terminates after finitely many steps.

Theorems 1.1 and 5.2 allow us to construct *Makanin–Razborov diagrams* over Γ, which is one of the main purposes of this paper. A Makanin–Razborov diagram for G is a finite tree which encodes all of the information about $\mathrm{Hom}(G,\Gamma)$ obtained from the above process. Thus a Makanin–Razborov diagram gives a parametrisation of $\mathrm{Hom}(G,\Gamma)$. This is described in further detail in Section 6, the main result of which is the existence of Makanin–Razborov diagrams over Γ.

Sela [33, I.8] asked whether Theorem 5.16 holds and whether Makanin–Razborov diagrams exist in the context of CAT(0) groups with isolated flats. We believe that relatively hyperbolic groups with abelian parabolics are a natural setting for these questions. However, it is worth noting that for technical reasons we have restricted to abelian, rather than virtually abelian, parabolics, so we have not entirely answered Sela's questions.

An outline of this paper is as follows. In Section 2 we recall the definition of relatively hyperbolic groups, and recall the construction of limiting $\mathbb{R}$–trees from [17] and [19], as well as other useful results. In Section 3 we improve upon our version of Sela's shortening argument from [18] and [19] to deal with arbitrary sequences of homomorphisms $\{h_n\colon G \to \Gamma\}$ where G is an arbitrary finitely generated group. In Section 4 we recall Sela's construction of *shortening quotients* from [31], and adapt this construction to our setting. In Section 5 we prove Theorem 5.2, one of the main technical results of this paper. We also prove Theorems 5.11, 5.16, and a number of structural results about Γ–limit groups. Finally in Section 6 we construct Makanin–Razborov diagrams over Γ.

Acknowledgment I would like to thank Zlil Sela for providing me with the proof of [32, Proposition 1.21], which is repeated in the proof of Proposition 5.14 in this paper. I would also like to thank the referee for correcting a number of mistakes in earlier versions of this paper, in particular the use of the bending moves in shortening quotients, and for his/her careful reading(s) and numerous comments, which have substantially improved the exposition of the results in this paper.

This work was supported in part by NSF Grant DMS-0504251.

2 Preliminaries

2.1 Relatively hyperbolic groups

Relatively hyperbolic groups were first defined by Gromov in his seminal paper on hyperbolic groups [16, Subsection 8.6, p.256]. Another definition was given by Farb [15, Section 3], and further definitions given by Bowditch [7, Definitions 1 and 2, page 1]. These definitions are all equivalent (see [7, Theorem 7.10, page 44] and [8, Theorem 6.1, page 682]). Recently there has been a large amount of interest in these groups (see [1, 8, 10, 12, 13, 24, 35], among others). The definition we give here is due to Bowditch [7, Definition 2, page 1].

Definition 2.1 A group Γ with a family $\mathcal{G}$ of finitely generated subgroups is called *hyperbolic relative to* $\mathcal{G}$ if Γ acts on a δ–hyperbolic graph $\mathcal{K}$ with finite quotient and finite edge stabilisers, where the stabilisers of infinite valence vertices are the elements of $\mathcal{G}$, so that $\mathcal{K}$ has only finitely many orbits of simple loops of length n for each positive integer n.

The groups in $\mathcal{G}$ are called *parabolic subgroups of* Γ.

In this paper we will be exclusively interested in relatively hyperbolic groups which are torsion-free and have abelian parabolic subgroups.

We record the following simple lemma for later use.

Lemma 2.2 *Suppose that Γ is a torsion free relatively hyperbolic group with abelian parabolics. Then abelian subgroups of Γ are finitely generated.*

Proof Suppose that $g \in \Gamma$. If g is a hyperbolic element, then by a result of Osin (see [24, Theorem 1.14, page 10] and the comment thereafter), the centraliser of g is cyclic.

Therefore, any noncyclic abelian subgroup is contained entirely within a single parabolic subgroup. These are assumed to be abelian (and are finitely generated by the definition of 'relatively hyperbolic'). □

2.2 The limiting $\mathbb{R}$–tree

In this subsection we recall a construction from [19] (see also [17] for more details). Suppose that Γ is a torsion-free relatively hyperbolic group with abelian parabolic subgroups. In [19, Section 4], we constructed a space X on which Γ acts properly and cocompactly by isometries. For each parabolic subgroup P (of rank n, say) there is in X an isometrically embedded copy of $\mathbb{R}^n$, with the Euclidean metric, so that the action of P leaves this Euclidean space invariant and this P–action is proper and cocompact with quotient the n–torus.

Terminology 2.3 We say that two homomorphisms $h_1, h_2 \colon G \to \Gamma$ are *conjugate* if there exists $\gamma \in \Gamma$ so that $h_1 = \tau_\gamma \circ h_2$, where τ_γ is the inner automorphism of Γ induced by γ. Otherwise, h_1 and h_2 are *non-conjugate.*

Suppose now that G is a finitely generated group, and that $\{h_n \colon G \to \Gamma\}$ is a sequence of pairwise non-conjugate homomorphisms. By considering the induced actions of G on X, and passing to a limit, we extract an isometric action of G on an asymptotic cone X_ω of X. This action has no global fixed point. There is a separable G–invariant subset $\mathcal{C}_\infty \subseteq X_\omega$, and by passing to a subsequence $\{f_i\}$ of $\{h_i\}$ we may assume that the (appropriately scaled) actions of G on X converge in the G–equivariant Gromov–Hausdorff topology to the G–action on $\mathcal{C}_\infty$.

The space $\mathcal{C}_\infty$ is a *tree-graded* metric space, in the terminology of Druţu and Sapir [12, Definition 1.10, page 7]. Informally, this means that there is a collection of 'pieces' (in this case finite dimensional Euclidean spaces), and otherwise the space is 'tree-like' (see [12] for the precise definition and many properties of tree-graded metric spaces). By carefully choosing lines in the 'pieces', and projecting, an $\mathbb{R}$–tree T is extracted from $\mathcal{C}_\infty$. This tree T comes equipped with an isometric G–action with no global fixed points and the kernel of the G–action on T is the same as the kernel of the G–action on $\mathcal{C}_\infty$. For more details on this entire construction, see [19, Sections 4,5,6] and [17].

Definition 2.4 [5, Definition 1.5, page 3] Suppose that $\{h_n \colon G \to \Gamma\}$ is a sequence of homomorphisms. The *stable kernel* of $\{h_n\}$, denoted $\underrightarrow{\mathrm{Ker}}\,(h_n)$, is the set of all $g \in G$ so that $g \in \ker(h_n)$ for all but finitely many n.

A sequence of homomorphisms $\{h_i \colon G \to \Gamma\}$ is *stable* if for all $g \in G$ either (i) $g \in \ker(h_i)$ for all but finitely many i; or (ii) $g \notin \ker(h_i)$ for all but finitely many i.

The following theorem recalls some of the properties of the G–action on the $\mathbb{R}$–tree T.

Theorem 2.5 (Compare [17, Theorem 4.4] and [19, Theorem 6.4]) *Suppose that Γ is a torsion-free group that is hyperbolic relative to a collection of free abelian subgroups and that G is a finitely generated group. Let $\{h_n \colon G \to \Gamma\}$ be a sequence of pairwise non-conjugate homomorphisms. There is a subsequence $\{f_i\}$ of $\{h_i\}$ and an action of G on an $\mathbb{R}$–tree T so that if K is the kernel of the G–action on T and $L := G/K$ then:*

1. *The stabiliser in L of any non-degenerate segment in T is free abelian.*
2. *If T is isometric to a real line then L is free abelian, and for all but finitely many n the group $h_n(G)$ is free abelian.*
3. *If $g \in G$ stabilises a tripod in T then $g \in \underrightarrow{\mathrm{Ker}}\,(f_i) \subseteq K$.*
4. *Let $[y_1, y_2] \subset [y_3, y_4]$ be non-degenerate segments in T, and assume that $Stab_L([y_3, y_4])$ is nontrivial. Then*
$$Stab_L([y_1, y_2]) = Stab_L([y_3, y_4]).$$
In particular, the action of L on T is stable.

(5) *L is torsion-free.*

Thus T is isometric to a line if and only if L is abelian. If L is not abelian then $K = \underrightarrow{\mathrm{Ker}}\,(f_i)$.

We now recall the definition of *Γ–limit groups*. There are many ways of defining Γ–limit groups. We choose a geometric definition using the above construction.

Definition 2.6 (Compare [32, Definition 1.11], [19, Definition 1.2]) A *strict Γ–limit group* is a quotient G/K where G is a finitely generated group, and K is the kernel of the G–action on T, where T is the $\mathbb{R}$–tree arising from a sequence of non-conjugate homomorphisms $\{h_n \colon G \to \Gamma\}$ as described above.

A Γ–limit group is a group which is either a strict Γ–limit group or a finitely generated subgroup of Γ.

Definition 2.7 We say that a sequence of homomorphisms $\{h_n \colon G \to \Gamma\}$ *converges to a Γ–limit group L* if the subsequence in Theorem 2.5 can be taken to be $\{h_n\}$ itself and L is the resulting Γ–limit group.

Remark 2.8 There are finitely generated subgroups of torsion-free hyperbolic groups which are not finitely presented (see, for example, the construction from [25, Corollary (b), page 46]). Therefore, when Γ is a torsion-free relatively hyperbolic group with free abelian parabolic subgroups, a Γ–limit group need not be finitely presented. This presents substantial complications (many of which are already dealt with by Sela in [32, Section 1]), some of which are solved by the application of Theorem 5.6 below.

2.3 Acylindrical accessibility and JSJ decompositions

In [28], Sela studied *acylindrical* graph of groups decompositions, and proved an accessibility theorem for k–acylindrical splittings [28, Theorem 4.1]. Unlike other accessibility results such as [14, Theorem 5.1, page 456] and [3, Main theorem, page 451], Sela's result holds for finitely generated groups, rather than just for finitely presented groups.[1]

We follow [31, Section 2] to construct the abelian JSJ decomposition of a freely indecomposable strict Γ–limit group, L.

Recall the following:

[1]Since Γ–limit groups needn't be finitely presented, it is important that we can apply acylindrical accessibility.

Lemma 2.9 [19, Lemma 6.7] *Let Γ be a torsion-free group which is hyperbolic relative to abelian subgroups, and let L be a strict Γ–limit group. Abelian subgroups of L are malnormal. Each abelian subgroup of L is contained in a unique maximal abelian subgroup.*

Given Lemma 2.9, the proof of the following is identical to that of [31, Lemma 2.1].

Lemma 2.10 *Let Γ be a torsion-free group which is hyperbolic relative to abelian subgroups, and let L be a strict Γ–limit group. Let M be a maximal abelian subgroup of L, and let A be any abelian subgroup of L.*

(1) *If $L = U *_A V$ then M can be conjugated into either U or V.*
(2) *If $L = U*_A$ then either (i) M can be conjugated into U or (ii) M can be conjugated to M' and $L = U *_A M'$.*

Given Lemma 2.10, if we have an abelian splitting $L = U*_A$ where A is a subgroup of a non-elliptic maximal abelian subgroup M, we can convert it into the amalgamated free product $L = U *_A M$. Thus we will concentrate only on splittings where edge groups are abelian and every noncyclic abelian subgroup is elliptic. Just as in [31, Lemma 2.3], any such splitting can be modified by sliding operations and modifying boundary monomorphisms by conjugation to be 2–acylindrical. Therefore, we can apply acylindrical accessibility.

We now follow the analysis from [31, Section 2]. First we construct the canonical quadratic decomposition of L (see [27, Theorem 5.6, pages 98–99]. This can then be refined to construct the cyclic JSJ decomposition, and further refined to construct the abelian JSJ decomposition, which encodes all 2–acylindrical abelian splittings of L in which noncyclic abelian subgroups are elliptic. The proof of [29, Theorem 1.7] applies in this context and implies that the abelian JSJ decomposition of L is unique up to sliding moves, conjugation and modifying boundary monomorphisms by conjugation. See [32, Theorem 1.10] for a precise statement (which holds in our context without change), and [27] for definitions of terms left undefined here.

There are three types of vertex groups in the abelian JSJ decomposition of L. The first are *surface vertex groups*, which correspond to maximal quadratically hanging subgroups. The second are *abelian vertex groups* which are abelian. The third type are *rigid vertex groups*. The key feature of rigid vertex groups is that they are elliptic in any abelian splitting of L.

By Theorem 2.5 any strict Γ–limit group L admits a nontrivial stable action on any $\mathbb{R}$–tree with abelian segment stabilisers and trivial tripod stabilisers. If L

is nonabelian then this $\mathbb{R}$–tree is not a line and, by [28, Theorem 3.1], L splits over a group of the form E–by-cyclic, where E stabilises a segment. Since abelian subgroups of L are malnormal (Lemma 2.9) and segment stabilisers are abelian, any group of the form E–by-cyclic is in fact abelian. Therefore, any non-abelian strict Γ–limit group admits a nontrivial abelian splitting. This now implies the following result.

Proposition 2.11 *Suppose that L is a non-abelian, freely indecomposable strict Γ–limit group, and that L is not a surface group. The abelian JSJ decomposition of L is nontrivial.*

3 The shortening argument

In [18, Sections 3–6] and [19, Section 7] we described a version of Sela's shortening argument which worked for sequences of *surjective* homomorphisms to Γ, and described in [18, Section 3] why this notion is insufficient for all sequences of homomorphisms.

In this section we present another version of the shortening argument, which works for all sequences of homomorphisms $\{h_n \colon G \to \Gamma\}$, for *any* finitely generated group G. This version was stated but not proved in [19, Theorem 7.10], and we give the proof here.

There are two equivalent approaches to this version of the shortening argument. The first is to find a group $\widehat{G}$ which contains G and shorten using elements of $\mathrm{Mod}(\widehat{G})$, rather than just elements of $\mathrm{Mod}(G)$ (this approach was used in the proof of [18, Theorem 7.9]) . The second approach is to use the 'bending' moves of Alibegović [2]. We use the second approach here.

Definition 3.1 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 $\phi \in \mathrm{Aut}(G)$ by $\phi(a) = a$ for $a \in A$ and $\phi(b) = cbc^{-1}$ 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 $\phi \in \mathrm{Aut}(G)$ by $\phi(a) = a$ for $a \in A$ and $\phi(t) = tc$.

Definition 3.2 *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 element-wise can be naturally extended to an automorphism of the ambient group G. Such an automorphism is called a *generalised Dehn twist* of G.

Definition 3.3 Let G be a finitely generated group. We define $\mathrm{Mod}(G)$ to be the subgroup of $\mathrm{Aut}(G)$ 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 [31, Section 5] and [5, Section 1].

We will try to shorten homomorphisms by pre-composing with elements of $\mathrm{Mod}(G)$. However, as seen in [18, Section 3], this is not sufficient to get the most general result. Thus, we also define a further kind of move (very similar to Alibegović's *bending* move, [2, Section 2]).

Definition 3.4 Suppose that Γ is a torsion-free group which is hyperbolic relative to free abelian subgroups, that G is a finitely generated group, and that $h\colon G \to \Gamma$ is a homomorphism. We define two kinds of 'bending' moves as follows:

(B1) Let Λ be a graph of groups decomposition of G, and let A be an abelian vertex group in Λ. Suppose that $h(A) \leq P$, where P is a parabolic subgroup of Γ. A move of type (B1) replaces h by a homomorphism $h'\colon G \to \Gamma$ which is such that:

  (a) $h'(A) \cong h(A)$;

  (b) $h'(A) \leq P$;

  (c) h' agrees with h on all edge groups of Λ adjacent to A, and all vertex groups other than A.

(B2) Let Λ be a graph of groups decomposition of G, and let A be an abelian edge group of Λ, associated to an edge e. Suppose that $h(A) \leq P$, for some parabolic subgroup P of Γ. A move of type (B2) replaces h by a homomorphism which either

(a) conjugates the image under h of a component of $\pi_1(\Lambda \setminus e)$ by an element of P, in case e is a separating edge, or
(b) multiplies the image under h of the stable letter associated to e by an element of P, in case e is non-separating,

and otherwise agrees with h.

Definition 3.5 (Compare [5, Definition 4.2], [2, Definition 2.11]) We define the relation '$\sim$' on the set of homomorphisms $h\colon G \to \Gamma$ to be the equivalence relation generated by setting $h_1 \sim h_2$ if h_2 is obtained from h_1 by:

(1) pre-composing with an element of $\mathrm{Mod}(G)$,
(2) post-composing with an inner automorphism of Γ or
(3) a bending move of type (B1) or (B2).

Definition 3.6 Let $\mathcal{A}$ be an arbitrary finite generating set for G, and let X be the space upon which Γ acts properly, cocompactly and isometrically (defined in [19] and briefly described in Section 2), with basepoint x. For a homomorphism $h\colon G \to \Gamma$ define $\|h\|$ by

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

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

The following is the main result of this section.

Theorem 3.7 *Suppose that Γ is a torsion-free group which is hyperbolic relative to a collection of free abelian subgroups. Let G be a freely indecomposable finitely generated group and $\{h_n\colon G \to \Gamma\}$ be a sequence of non-conjugate homomorphisms which converges to a faithful action of G on the tree-graded metric space $\mathcal{C}_\infty$ as in Subsection 2.2 above. Then, for all but finitely many n, the homomorphism h_n is not short.*

Proof Suppose that the sequence $\{h_n\colon G \to \Gamma\}$ converges into a faithful G–action on $\mathcal{C}_\infty$. As described in Subsection 2.2 (see [17] and [19] for more details), a faithful G–action on an $\mathbb{R}$–tree T may be extracted from the G–action on $\mathcal{C}_\infty$. We now analyse the action of G on T more closely.

The group G is freely indecomposable and the stabiliser in G of any tripod in T is trivial, so we can apply the decomposition theorem of Sela – [28, Theorem

3.1] – and decompose T into subtrees of three types: axial, IET, and discrete (because G is freely indecomposable and tripod stabilisers are trivial, there are no Levitt-type components in T). This decomposition of T induces (via the Rips machine; see [4]) a graph of groups decomposition of G, which will allow us to shorten h_n for sufficiently large n. See [26], [28, Sections 2,3] or [18, Section 4] for more information.

Note that there are two sources for segments in T. There are segments in $\mathcal{C}_\infty$ meeting each flat in at most a point, and there are flats in $\mathcal{C}_\infty$ which are projected to lines in T. We treat these as two separate cases. However, we can make the following simplifications (let $\mathbb{P}$ be the collection of lines in T which are projections of flats in $\mathcal{C}_\infty$):

(1) 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 ([18, Proposition 4.3]).

(2) Suppose that a line $l \subset T$ is an axial subtree of T and that the line p_E is in $\mathbb{P}$. If $l \cap p_E$ contains more than a point then $l = p_E$ ([18, Proposition 4.5]).

(3) If an edge e in the discrete part of T has an intersection of positive length with $p_E \in \mathbb{P}$ then $e \subset p_E$ ([18, Lemma 4.7]).

Fix a finite generating set $\mathcal{A}$ for G. Let y be the basepoint in T, and consider the paths $[y, u.y]$ for $u \in \mathcal{A}$. If there is any IET component of T which intersects any segment $[y, u.y]$ nontrivially then we can apply [26, Theorem 5.1] and [18, Corollary 4.4] to shorten these intersections whilst leaving the remaining segments unchanged (to see that we can have G finitely generated rather than finitely presented, see [18, Remark 4.8]).

Suppose that some segment $[y, u.y]$ has an intersection of positive length with some axial component $l \subset T$ so that l is not contained in any $p_E \in \mathbb{P}$. Then [18, Theorem 5.1] can be used to shorten those segments $[y, u_i.y]$ intersecting the orbit of p_E nontrivially, and leaving other segments unchanged.

Suppose that $[y, u.y]$ intersects some line p_E nontrivially, and that p_E is an axial component of T. The only place where the proof of [18, Theorem 5.2] breaks down is that the images $h_i(G)$ may not intersect parabolics in its image in denser and denser subsets (when measured with the scaled metric). However, this is exactly what the bending move (B1) is designed to deal with.

We have the following analogue of [18, Proposition 5.4]: Let E denote the flat in $\mathcal{C}_\infty$ which projects to p_E. The subgroup $\mathrm{Stab}_G(E)$ is an abelian subgroup of G. There is a sequence of flats $E_i \subset X_i$ so that $E_i \to E$ in the Gromov topology.

The subgroups $h_i(\mathrm{Stab}_G(E))$ are abelian, and fix the flat E_i, for sufficiently large i. Thus $h_i(\mathrm{Stab}_g(E))$ is contained in a unique maximal abelian subgroup A_{E_i} of Γ. If we fix a finite subset W of $\mathrm{Stab}_G(E)$ and $\epsilon > 0$, then for sufficiently large i, there is an automorphism $\sigma_i\colon A_{E_i} \to A_{E_i}$ so that

(1) For every $w \in W$, and every $r_i \in E_i$,
$$d_{X_i}(r_i, h_i(\sigma_i(w)).r_i) < \epsilon.$$

(2) For any $k \in \mathrm{Stab}_G(E)$ which acts trivially on E we have $\sigma_i(h_i(k)) = h_i(k)$.

The proof of the existence of such a σ_i is the same as the proof of [18, Proposition 5.4]. Such a σ_i induces a move of Type (B1) in a straightforward manner, since the adjacent edge groups to the vertex group $\mathrm{Stab}_G(E)$ contain elements which act trivially on E, therefore we replace h_i by the homomorphism which agrees with h_i on all edge groups and on all vertex groups which are not $\mathrm{Stab}_G(E)$, and replaces $h_i|_{\mathrm{Stab}_G(E)}$ by $\sigma_i \circ h_i|_{\mathrm{Stab}_G(E)}$.

We now construct shortening elements for all but finitely many of the intervals $[y_n, h_n(u).y_n]$ by following the proof of [26, Theorem 5.1] (see [18, Section 5] for more details).

Finally, we are left with the case where $[y, u.y]$ is contained entirely in the discrete part of T. We follow the proof of [18, Theorem 6.1], which in turn followed [26, Section 6]. This argument naturally splits into a number of cases.

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

Case 1a e is not contained entirely in a line $p_E \in \mathbb{P}$ and $\bar{e} \in T/G$ is a splitting edge.

This case follows directly as in [18, Section 6].

Case 1b e is not completely contained in a line $p_E \in \mathbb{P}$ and $\bar{e}$ is not a splitting edge.

This case also follows directly as in [18, Section 6].

Case 1c e is contained in a line $p_E \in \mathbb{P}$ and $\bar{e}$ is a splitting edge.

In this case, we have a graph of groups decomposition $G = H_1 *_{A_E} H_2$, where $A_E = \mathrm{Stab}_G(E)$ (and E is the flat in $\mathcal{C}_\infty$ which projects to p_E). The Dehn twist which is found in [18, Section 6] is naturally replaced by a bending move of type (B2).

Case 1d e is contained in a line $p_E \in \mathbb{P}$ and $\bar{e}$ is not a splitting edge.

Once again the Dehn twist is replaced by a bending move of Type (B2).

Case 2 y is a vertex of T.

Once again here there are four cases, depending on whether on edge adjacent to y is or is not a splitting edge and is or is not contained in a line $p_E \in \mathbb{P}$. In case an edge e is not contained in a line $p_E \in \mathbb{P}$, we proceed exactly as in [18], following [26] directly. In case $e \subset p_E$, we replace the shortening Dehn twists by bending moves of type (B2) as in Case 1 above.

Therefore, in any case, we can find moves which shorten all but finitely many of the h_i, as required. □

4 Shortening quotients

In this section we recall the concept of *shortening quotients* from [31, Section 5] and [32, Definition 1.14], and generalise this notion to the relatively hyperbolic setting. The use of bending moves means that the output of this section is slightly different to that in the hyperbolic case (though it will be good enough for our purposes).

Let G be a finitely generated group, Γ a torsion-free relatively hyperbolic group with abelian parabolics and $\{h_n\colon G \to \Gamma\}$ a stable sequence of homomorphisms, with associated Γ–limit group L.[2] Suppose that L is nonabelian and freely indecomposable. Let $\pi\colon G \to L$ be the canonical quotient map. The shortening procedure (given below) constructs a sequence of homomorphisms $\{\nu_i\colon F \to \Gamma\}$, where F is a finitely generated free group. The sequence $\{\nu_i\}$ has a subsequence converging to a Γ–limit group Q, which comes equipped with a canonical epimorphism $\eta : L \to Q$. See Lemma 4.1 and Proposition 4.2 for more information.

We follow the construction from [31, Section 3] and [30, Section 3] (see [18, Section 7] for more details in the current context). See also [34] for more information on the Bass–Serre theory, in particular [34, I.5] for the standard presentation of the fundamental group of a graph of groups.

Given the situation described in the previous paragraph, we now describe the construction of $\{\nu_i\}$, Q and η. Let Λ_L be an abelian JSJ decomposition for L, with vertex groups $V^1, \ldots, V^m$ and edge groups $E^1, \ldots, E^s$. Let $t^1, \ldots, t^b$ be the Bass–Serre generators for L with respect to some (fixed) maximal subtree

[2] *Stable* sequences of homomorphisms were defined in Definition 2.4.

of Λ_L. Note that L is finitely generated, so each of the vertex groups V_i is generated by a finite set together with the adjacent edge groups.

As in [31, Theorem 3.2, comment at the bottom of page 45] we have not yet proved that the edge groups are finitely generated (though this will eventually turn out to be the case; see Corollary 5.12 below).

We can 'approximate' the finitely generated group L by finitely presented groups U_n, each equipped with a graph of groups decomposition Λ_n which is a 'lift' of Λ_L. We describe this approximation now, following [31, Theorem 3.2] and [18, Theorem 7.9].

Let $g_1, \dots, g_k$ generate G and let $\pi\colon G \to L$ be the canonical quotient map associated to the convergent subsequence of $\{h_n\colon G \to \Gamma\}$. Note that L is clearly generated by $\{\pi(g_1), \dots, \pi(g_k)\}$. Now choose elements

$$v_1^1, \dots, v_{l_1}^1, \dots, v_1^m, \dots, v_{l_m}^m, t_1, \dots, t_b \in G,$$

so that (i) $\pi(v_k^i) \in V^i$; (ii) $\pi(t_k) = t_k$; and (iii) for each generator g_j of G there is a word:

$$g_j = w_j(v_1^1, \dots, v_{l_m}^m, t_1, \dots, t_b).$$

Thus G is generated by $\{v_1^1, \dots, v_{l_m}^m, t_1, \dots, t_b\}$. We may also assume, by adding finitely many more elements to our list if necessary, that each V^i is generated by $\{\pi(v_1^i), \dots, \pi(v_{l_i}^i)\}$, together with the edge groups E^j which are adjacent to V^i. In case the edge groups adjacent to V^i are all finitely generated, we assume that V^i is in fact generated by $\{\pi(v_1^i), \dots, \pi(v_{l_i}^i)\}$. For each j, let $e_1^j, e_2^j, \dots \in G$ be a set of elements for which $\pi(e_p^j) \in E^j$ so that E^j is generated by $\{\pi(e_1^j), \pi(e_2^j), \dots\}$.

We now define the groups U_n. First, define the group H_n to be the group with the generating set:

$$\{x_1^1, \dots, x_{l_1}^1, \dots, x_1^m, \dots, x_{l_m}^m, y_1, \dots, y_b, z_1^1, \dots, z_1^s, z_n^1, \dots, z_n^s\},$$

together with the relations $[z_{p_1}^j, z_{p_2}^j] = 1$ for $j = 1, \dots, s$, and $1 \le p_1, p_2 \le n$. Clearly there exists a natural epimorphism $\sigma_n\colon H_n \to L$ defined by $\sigma_n(x_p^i) = \pi(v_p^i), \sigma_n(y_r) = t^r$ and $\sigma_n(z_d^j) = \pi(e_d^j)$. We define U_n to be the quotient of H_n defined by adding as relations each word w in the given generators of H_n for which

(1) $\sigma_n(w) = 1$,

(2) the length of w in the given generating set for H_n is at most n and

(3) for some fixed index $i \in \{1, \dots, m\}$, the word w is a word in

(a) the generators $x^i_1, \dots, x^i_{l_i}$,

(b) the elements $z^j_1, \dots, z^j_n$ for any of the indices $j \in \{1, \dots, s\}$ for which $E^j < V^i$ and

(c) the words $y_r z^j_1 y_r^{-1}, \dots, y_r z^j_n y_r^{-1}$ for any pair of indices (j, r) for which $t^r E^j (t^r)^{-1} < V^i$.

There exists a natural map $\kappa_n \colon U_n \to U_{n+1}$. Since L is the quotient of G by the stable kernel of $\{h_i \colon G \to \Gamma\}$, for each n there exists $k_n > k_{n-1}$ so that the homomorphism $h_{k_n} \colon G \to \Gamma$ induces a homomorphism $\lambda_n \colon U_n \to \Gamma$ defined by $\lambda_n(x^i_p) = h_{k_n}(v^i_p)$, $\lambda_n(y_r) = h_{k_n}(t_r)$ and $\lambda_n(z^j_d) = h_{k_n}(e^j_d)$. Also, the homomorphism $\sigma_n \colon H_n \to L$ factors through U_n, and we denote the associated homomorphism from U_n to L by σ_n.

Since the second set of defining relations of U_n consists of words whose letters are mapped by σ_n into the same vertex group of Λ_L, each of the groups U_n admits an abelian splitting Λ_n which projects by σ_n into the abelian decomposition Λ_L of L. That is to say each of the vertex groups V^i_n in Λ_n satisfies $\sigma_n(V^i_n) \le V^i$, each of the edge groups $E^j_n = \langle z^j_1, \dots, z^j_n \rangle$ satisfies $\sigma_n(E^j_n) \le E^j$ and each of the Bass–Serre generators in Λ_n satisfies $\sigma_n(y_r) = t^r$.

By the *surface* vertex groups of Λ_n we mean the vertex groups in Λ_n which project to surface vertex groups of Λ_L. Since such groups are finitely presented, for all but finitely many n each surface vertex group in Λ_n is isomorphic to the corresponding surface vertex group of Λ_L, via the natural projection from U_n to L. Denote by $\mathrm{Mod}(U_n)$ the subgroup of $\mathrm{Aut}(U_n)$ generated by

(1) inner automorphisms,

(2) Dehn twists in edge groups of Λ_n,

(3) Dehn twists in cyclic groups which arise as the edge group in a splitting of U_n obtained by cutting the surface associated to a surface vertex group of Λ_n along a weakly essential s.c.c.[3] and

(4) generalised Dehn twists induced by the decomposition Λ_n.

Let W_n be the subgroup of U_n generated by the x^i_p's and the y_r's. Clearly the homomorphism $\kappa_n \colon U_n \to U_{n+1}$ restricts to an epimorphism $\kappa_n \colon W_n \to W_{n+1}$, and λ_n restricts to a homomorphism from W_n to Γ.

[3]Following [27], a *weakly essential s.c.c* on a surface is a simple closed curve which is not null-homotopic, not boundary parallel and not the core of a Möbius band.

Thus we have the following commutative diagram, where $\iota_n\colon W_n \to U_n$ is inclusion.

$$\begin{tikzcd}
W_1 \arrow[r, "\kappa_1"] \arrow[d, "\iota_1"'] & W_2 \cdots \arrow[r, "\kappa_{n-2}"] \arrow[d, "\iota_2"'] & W_{n-1} \arrow[r, "\kappa_{n-1}"] \arrow[d, "\iota_{n-1}"'] & W_n \arrow[d, "\iota_n"'] \\
U_1 \arrow[r, "\kappa_1"] \arrow[d, "\lambda_1"'] & U_2 \cdots \arrow[r, "\kappa_{n-2}"] \arrow[d, "\lambda_2"'] & U_{n-1} \arrow[r, "\kappa_{n-1}"] \arrow[d, "\lambda_{n-1}"'] & U_n \arrow[d, "\lambda_n"'] \\
\Gamma & \Gamma & \Gamma & \Gamma
\end{tikzcd}$$

The Γ–limit group L is the direct limit of the sequence $\{(W_i, \kappa_i)\}$. Denote the free group W_1 by F, and let $\phi\colon F \to L$ be the canonical quotient map.

Define an equivalence relation on $\mathrm{Hom}(W_n, \Gamma)$ analogously to Definition 3.5, using (i) automorphisms $\phi \in \mathrm{Mod}(U_n)$; (ii) bending moves (where we restrict attention to bending moves defined using the decomposition Λ_n of U_n); and (iii) conjugations in Γ. Also, define the norm $\|h\|$ of $h \in \mathrm{Hom}(W_n, \Gamma)$, using the generating set $\{x_1^1, \ldots, x_{l_m}^m, y_1, \ldots, y_r\}$. For each $n \geq 1$, choose a homomorphism $\widehat{\lambda_n}\colon W_n \to \Gamma$ so that (i) $\widehat{\lambda_n} \sim \lambda_n|_{W_n}$; and (ii) $\|\widehat{\lambda_n}\|$ is minimal amongst all homomorphisms equivalent to $\lambda_n|_{W_n}$.

Let ξ_n be the given epimorphism from F to W_n (obtained by identifying their generating sets in the obvious way), and let $\nu_n = \widehat{\lambda_n} \circ \xi_n\colon F \to \Gamma$. Passing to a subsequence of $\{\nu_n\}$, we obtain an associated Γ–limit group Q, with canonical quotient map $\nu_\infty\colon F \to Q$.

Lemma 4.1 *The natural map between the generating set of L and the generating set of Q extends to an epimorphism $\eta : L \to Q$. Furthermore, η is injective on each rigid vertex group of Λ_L.*

Proof The group L is the direct limit of the sequence $\{(W_i, \kappa_i)\}$, so the kernel of the map from F to L is $\underrightarrow{\mathrm{Ker}}\,(\xi_n)$.

The Γ–limit group Q is obtained from the sequence $\{\nu_n = \widehat{\lambda_n} \circ \xi_n\}$. It is clear that $\underrightarrow{\mathrm{Ker}}\,(\nu_n) \subseteq \underrightarrow{\mathrm{Ker}}\,(\xi_n)$. This proves that there is an epimorphism $\eta\colon L \to Q$ as required by the statement of the lemma.

Finally, we prove that the rigid vertex groups of Λ_L are mapped injectively into Q by η. First note that rigid vertex groups are non-abelian, by definition. Let V^i be a rigid vertex group of L, and let x, y be distinct elements of V^i. Suppose that $u, v \in F$ are such that $\phi(u) = x$, $\phi(v) = y$. By the properties of the above construction, there exists $j_0 \geq 1$ so that for all $j \geq j_0$ the following hold:

(i) $\xi_j(u), \xi_j(v) \in V_j^i$; (ii) $\xi_j(u) \neq \xi_j(v)$; and (iii) $\lambda_j(\xi_j(u)) \neq \lambda_j(\xi_j(v))$.

Now, the effect of all Dehn twists and generalised Dehn twists arising from Λ_j, and also of the bending moves arising from Λ_j, is conjugation on V_j^i (by the same element for all of V_j^i). Hence $\widehat{\lambda_j}(\xi_j(u)) \neq \widehat{\lambda_j}(\xi_j(v))$, which is to say that $\nu_j(u) \neq \nu_j(v)$. Since this is true for all $j \geq j_0$, we see that $\nu_\infty(u) \neq \nu_\infty(v)$. By construction, $\eta(x) = \nu_\infty(u)$ and $\eta(y) = \nu_\infty(v)$. Therefore, η is injective on V^i, as required. □

The group Q is called the *shortening quotient* of L associated to $\{h_n \colon G \to \Gamma\}$.

Although we speak of *the* shortening quotient, Q and η depend on the choices of shortest homomorphism in the equivalence class of λ_n, and also on the convergent subsequence of $\{\nu_n\}$ chosen.

We now record some important properties of shortening quotients for use later in this paper.

Proposition 4.2 *Suppose that G is a finitely generated group, that Γ is a torsion-free relatively hyperbolic group with abelian parabolics, and that $\{h_n \colon G \to \Gamma\}$ is a sequence of homomorphisms that converges into a faithful action of a non-abelian, freely indecomposable strict Γ–limit group L on an $\mathbb{R}$–tree.*

Suppose that $\{\nu_{n_t} \colon F \to \Gamma\}$ is a (convergent) sequence of homomorphisms so that $\nu_{n_t} = \widehat{\lambda_{n_t}} \circ \xi_{n_t}$, where $\widehat{\lambda_{n_t}} \colon W_{n_t} \to \Gamma$ is a short homomorphism, as in the construction of shortening quotients above. Let Q be the resulting shortening quotient, with canonical epimorphism $\eta \colon L \to Q$ as in Lemma 4.1 above.

1. *If the sequence $\{\nu_{n_t}\}$ is not contained in finitely many conjugacy classes then η is not injective, so Q is a proper quotient of L.*
2. *If the sequence $\{\nu_{n_t}\}$ is contained in finitely many conjugacy classes, then there is some subsequence $\{h_r\}$ of $\{h_{n_t}\}$ so that each h_r factors through the canonical epimorphism $\pi \colon G \to L$.*

Proof First we prove (1). Suppose that Q is not a proper quotient of L, and that $\{\nu_{n_t}\}$ is not contained in finitely many conjugacy classes. Then the limiting action of Q on the $\mathbb{R}$–tree T, induces a faithful action of L on T. This, in turn, induces an abelian splitting Θ of L. We use the splitting Θ to shorten all but finitely many of the homomorphisms $\widehat{\lambda_{n_t}}$, which contradicts the fact that they were chosen to be short. We argue as in Section 3.

Consider a generator u of L, the basepoint y of T and the segment $[y, u.y]$. If $[y, u.y]$ intersects an IET component of T in more than a point, then [26, Theorem 5.1] finds a mapping class of the associated surface in the splitting of L which shortens the limiting action on T. By [18, Corollary 4.4], the length of $[y, u.y]$ in T is the same as its the pre-image in $\mathcal{C}_\infty$, so we use the same automorphism to shorten the limiting action on $\mathcal{C}_\infty$. The surface group found above is conjugate into a surface vertex group of the JSJ decomposition of L, and so the shortening automorphism is a product of Dehn twists in s.c.c. on this surface. Since a surface vertex is finitely presented, all but finitely many of the U_n have an isomorphic copy of this surface vertex, and we can shorten the approximating homomorphisms. Thus we may suppose that $[y, u.y]$ does not intersect any IET components in more than a point.

If $[y, u.y]$ intersects an axial subtree in a segment of positive length then we apply the shortening moves in Section 3 to shorten the limiting action on $\mathcal{C}_\infty$. Once again, these shortening moves can easily be lifted to all but finitely many of the U_n (by using the fact that the splitting Λ_L of L essentially encodes all abelian splittings of L). Thus the appropriate shortening moves will shorten all but finitely many of the approximating homomorphisms from W_n to Γ. It is worth noting that even though we do not know that abelian subgroups of L are finitely generated, the generalised Dehn twists that we need can be chosen to be supported on a finitely generated subgroup, which is finitely presented and will lift to all but finitely many of the U_n.

Finally, there is the case that $[y, u.y]$ is entirely contained in the discrete part of T. In this case, the limiting action is not shortened, but the appropriate Dehn twist, generalised Dehn twist or bending move can be lifted to all but finitely many of the U_n, and used to shorten the approximating homomorphisms, just as in Section 3. The bending moves in any of the cases in the proof are naturally interpreted using the homomorphisms from W_n to Γ (and the splitting Λ_n of U_n). Thus (1) is proved by following the proof of Theorem 3.7.

We now prove (2). Suppose that the elements of the sequence $\{\nu_{n_t}\}$ belong to finitely many conjugacy classes. Since each of the $\widehat{\lambda}_{n_t}$ is short, the sequence $\|\nu_{n_t}\|$ is bounded, and by passing to a subsequence we may assume that $\{\nu_{n_t}\}$ is constant, and equal to ν, say. For ease of notation, we refer to the constant subsequence of $\{\nu_{n_t}\}$ by $\{\nu_i\}$, and the associated subsequence of $\{h_{n_t}\}$ by $\{h_i\}$.

We have $Q \cong \nu(F) \le \Gamma$. By Lemma 4.1, each of the rigid vertex groups of Λ_L embed in Q, and hence in Γ. Let V_j be a rigid vertex group in Λ_L, and denote by V_j' the isomorphic image of V_j in Γ given by the above embedding. By Lemma 2.2, the edge groups in Λ_L adjacent to rigid vertex groups are finitely generated. The edge groups adjacent to surface groups are finitely generated, and Λ_L does not have a pair of adjacent abelian vertex groups. Therefore, all edge groups in Λ_L are finitely generated, and therefore all vertex groups are also finitely generated.

Since all of the edge groups of Λ_L are finitely generated, we assumed during the construction of the U_i and W_i that the vertex group V_j of Λ_L was generated by $\{\pi(v_1^j), \ldots, \pi(v_{l_j}^j)\}$, for each $1 \leq j \leq m$. Now, for each $1 \leq j \leq m$ and each $1 \leq k \leq l_j$, we have $h_i(v_k^j) = \lambda_i(x_k^j)$. Therefore, if $Y_j = \langle v_1^j, \ldots, v_{l_j}^j \rangle \leq G$ and $Z_j^i = \langle x_1^j, \ldots, x_{l_j}^j \rangle \leq W_i$ then

$$h_i(Y_j) = (\lambda_i \circ \iota_i)(Z_j^i).$$

By construction, $(\lambda_i \circ \iota_i) \sim \widehat{\lambda}_i$.

The fact that $Q \cong \nu(F) \leq \Gamma$ implies that $\widehat{\lambda}_i(Z_j^i)$ is isomorphic to V_j'. Now, $(\lambda_i \circ \iota_i)$ and $\widehat{\lambda}_i$ are related by a sequence of shortening moves, each of which acts by conjugation on Z_j^i. Thus $(\lambda_i \circ \iota_i)(Z_j^i)$ is conjugate to V_j' in Γ. The same is therefore true of $h_i(Y_j)$.

Since the edge groups of Λ are finitely generated (and abelian), and since the non-rigid vertex groups are finitely presented, a presentation for L can be found which consists of presentations for the rigid vertex groups, and finitely many other generators and relations. The sequence $\{h_n \colon G \to \Gamma\}$ converges to $\pi \colon G \to L$, so each of these finitely many extra relations will be sent to 1 by h_l for sufficiently large l. However, we have already seen that the relations corresponding to the rigid vertex groups are mapped to 1 by all the h_i. This shows that all but finitely many of the h_i factor through $\pi \colon G \to L$, as required. This proves (2). □

5 Γ–limit groups

In this section we follow [32, Section 1] in order to understand Γ–limit groups, and to understand $\mathrm{Hom}(G, \Gamma)$, where G is an arbitrary finitely generated group. The main technical results of this section are Theorem 5.2 and Theorem 5.6.

These technical results are then applied to yield various applications: Theorem 5.10, Proposition 5.11, Theorem 5.16, and in the next section the construction of Makanin–Razborov diagrams.

Remark 5.1 A previous version of this paper claimed that any strict Γ–limit group which is freely indecomposable and nonabelian admits a principal cyclic splitting. However, this is false.

To see this, consider the following example. Let Γ be the fundamental group of a finite volume hyperbolic n–manifold M with torus cusps, for $n \geq 3$, so that Γ admits no splitting over $\mathbb{Z}$. Let M' be the double of M along one of its cusps, and $\Gamma' = \pi_1(M')$. By repeatedly Dehn twisting along a simple closed curve in the attaching torus, it is not difficult to see that Γ' is a strict Γ–limit group. However, the JSJ decomposition of Γ' is given by two copies of Γ attached by a (noncyclic) abelian edge group. It is clear that Γ' admits no principal cyclic splitting.

It is possible to adapt the proof of [31, Theorem 3.2] to this context to prove that the above example is one of the few ways in which a strict Γ–limit group can fail to have a principal cyclic splitting. However, rather than make this long diversion, we prefer to work entirely with the abelian JSJ decomposition of a strict Γ–limit group L, which we already know to be nontrivial, so long as L is freely indecomposable, non-abelian and not a surface group, by Proposition 2.11.

Let G be a fixed finitely generated group. Define an order on the set of Γ–limit groups that are quotients of G as follows: suppose R_1 and R_2 are both Γ–limit groups that are quotients of G, and that $\eta_i \colon G \to R_i$ are the (fixed) canonical quotient maps. We say $R_1 > R_2$ if there exists an epimorphism with non-trivial kernel $\tau \colon R_1 \to R_2$ so that $\eta_2 = \tau \circ \eta_1$. We say that R_1 and R_2 are *equivalent* if there is an isomorphism $\tau \colon R_1 \to R_2$ so that $\eta_2 = \tau \circ \eta_1$.

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

Theorem 5.2 (Compare [32, Theorem 1.12]) *Let Γ be a torsion-free group which is hyperbolic relative to free abelian subgroups. Every decreasing sequence of Γ–limit groups:*

$$R_1 > R_2 > R_3 > \ldots,$$

which are all quotients of a finitely generated group G, terminates after finitely many steps.

For limit groups, the analogous result has a short proof using the fact that free groups are linear (see [5, Corollary 1.9, page 4]). However (as observed by M. Kapovich; see, for example, [6, Subsection 1.4]), not all hyperbolic groups are linear and therefore not all relatively hyperbolic groups are linear.

It is worth noting that we do not insist that the Γ–limit groups R_n in the statement of Theorem 5.2 be strict Γ–limit groups. This will be important later, because to study a single homomorphism, we consider a constant sequence, which leads to a Γ–limit group which need not be strict.

Before we prove Theorem 5.2, we prove the following lemma (implicit in [32, page 7]):

Lemma 5.3 *Let Ξ be a finitely generated group, let L be a Ξ–limit group and suppose that L is d–generated. Then L can be obtained as a limit of homomorphisms $\{f_n\colon F_d \to \Xi\}$, where F_d is the free group of rank d.*

Proof The group L is obtained as the limit of a (convergent) sequence of homomorphisms $\{h_n\colon G \to \Xi\}$, where G is a finitely generated group. Let $\eta\colon G \to L$ be the canonical quotient map. Let $\{r_1, \ldots, r_d\}$ be a generating set for R and let $\{y_1, \ldots, y_d\} \subset G$ be such that $\eta(y_i) = r_i$ for $1 \le i \le d$. Let $G_d \le G$ be the subgroup generated by $\{y_1, \ldots, y_d\}$, and let $\pi\colon F_d \to G_d$ be the natural quotient map. Define $h_n = f_n \circ \pi$. It is easy to see that the sequence $\{h_n\colon F_d \to \Xi\}$ converges to L. □

Proof of Theorem 5.2 We follow the outline of the proof of [32, Theorem 1.12]. In order to obtain a contradiction, we suppose that there exists a finitely generated group G and an infinite descending sequence of Γ–limit groups:

$$R_1 > R_2 > R_2 > \ldots,$$

all quotients of G.

Without loss of generality, we may assume $G = F_d$, the free group of rank d. Let $\{f_1, \ldots, f_d\}$ be a basis for F_d, and let C be the Cayley graph of F_d with respect to this generating set. We construct a particular decreasing sequence of Γ–limit groups as follows. Let T_1 be a Γ–limit group with the following properties:

(1) T_1 is a proper quotient of F_d.
(2) T_1 can be extended to an infinite decreasing sequence of Γ–limit groups: $T_1 > L_2 > L_3 > \ldots$.
(3) The map $\eta_1\colon F_d \to T_1$ maps to the identity the maximal number of elements in the ball of radius 1 about the identity in C among all possible epimorphisms from F_d to a Γ–limit group L that satisfies the first two conditions.

Continue to define the sequence inductively. Suppose that the finite sequence $T_1 > T_2 > \ldots > T_{n-1}$ has been constructed, and choose T_n to satisfy:

(1) T_n is a proper quotient of T_{n-1}.

(2) The finite decreasing sequence of Γ–limit groups $T_1 > T_2 > \ldots > T_n$ can be extended to an infinite decreasing sequence.

(3) The map $\eta_n \colon F_d \to T_n$ maps to the identity the maximal number of elements in the ball of radius n about the identity in C among all possible maps from F_d to a Γ–limit group L_n satisfying the first two conditions.

Since each of the Γ–limit groups T_n is a quotient of F_d, each T_n is d–generated. Let $\{r_{1,n}, \ldots, r_{d,n}\}$ be a generating set for T_n. By Lemma 5.3, T_n can be obtained as a limit of a sequence of homomorphisms $\{v_i^n \colon F_d \to \Gamma\}$, with the quotient map $\eta_n \colon F_d \to T_n$ sending f_i to $r_{i,n}$.

For each n, choose a homomorphism $v_{i_n}^n : F_d \to \Gamma$ for which:

(1) Every element in the ball of radius n about the identity in C that is mapped to the identity by $\eta_n \colon F_d \to T_n$ is mapped to the identity by $v_{i_n}^n$. Every such element that is mapped to a nontrivial element by η_n is mapped to a nontrivial element by $v_{i_n}^n$.

(2) There exists an element $f \in F_d$ that is mapped to the identity by $\eta_{n+1} \colon F_d \to T_{n+1}$ for which $v_{i_n}^n(f) \neq 1$.

Denote the homomorphism $v_{i_n}^n$ by h_n. By construction, the set of homomorphisms $\{h_n \colon F_d \to \Gamma\}$ does not belong to a finite set of conjugacy classes. Therefore, from the sequence $\{h_n\}$ we can extract a subsequence that converges to a (strict) Γ–limit group, denoted T_∞. By construction, the Γ–limit group T_∞ is the direct limit of the sequence of (proper) epimorphisms:

$$F_d \to T_1 \to T_2 \to \cdots$$

Let $\eta_\infty \colon F_d \to T_\infty$ be the canonical quotient map.

The Γ–limit group T_∞ is a proper quotient of each of the Γ–limit groups T_n. For each index n, the group T_n was chosen to maximise the number of elements in the ball of radius n about the identity in C that are mapped to the identity in Γ among all Γ–limit groups that are proper quotients of T_{n-1} and that admit an infinite descending chain of Γ–limit groups. Therefore, it is not difficult to see that T_∞ does not admit an infinite descending chain of Γ–limit groups.

The proof of the following proposition is similar to that of [32, Proposition 1.16].

Proposition 5.4 *There is a subsequence $\{h_{n_t}\}$ of $\{h_n\colon F_d \to \Gamma\}$ so that each h_{n_t} factors through $\eta_\infty\colon F_d \to T_\infty$.*

Proof Let $T_\infty = H_1 * \cdots * H_p * F$ be the Grushko decomposition of T_∞. Consider H_1, and a finitely generated subgroup F^1 of F_d so that $\eta(F^1) = H_1$. Use the sequence $\{h_n|_{F^1}\colon F^1 \to \Gamma\}$ to construct a shortening quotient of H_1, proper if possible. Repeat for each of the H_i in turn, starting with the subsequence of $\{h_n\}$ used to construct the shortening quotient of H_{i-1} as a starting point for constructing a shortening quotient of H_i. In this way, we eventually get a quotient L_1 of T_∞ (map F to itself in this quotient).

For each of the H_i which admitted a proper shortening quotient, take the Grushko decomposition of the corresponding free factor in L_1, and for each of the non-free factors H_j' in this decomposition, attempt to find a proper shortening quotient of H_j', passing to finer and finer subsequences each time.

Continuing in this manner, we obtain a sequence of quotients:

$$T_\infty \to L_1 \to \cdots \to L_s$$

which are all proper quotients, except possibly in the case $s = 1$ and $T_\infty \to L_1$ is an isomorphism. This process will terminate after finitely many steps, since T_∞ does not admit an infinite descending chain of proper Γ–limit quotients.

We now have a subsequence $\{f_i\}$ of $\{h_n\}$, the final subsequence used in constructing a shortening quotient ($\{f_i\}$ is contained in all of the other subsequences used in this process). By construction, L_s is a free product of finitely generated free groups, and Γ–limit groups which do not admit a proper shortening quotient starting from the sequence $\{f_i\}$.

Let $L_s = K_1 * \ldots * K_q * F''$ be the given free product decomposition of L_s, where the K_i are not free, and let $\sigma_i\colon F_{k_i} \to K_i$ be the canonical quotient map. By Proposition 4.2, there is a subsequence $\{f_i'\}$ of $\{f_i\}$ so that all $f_i'|_{F_{k_1}}\colon F_{k_1} \to \Gamma$ factor through σ_1. Similarly, there is a further subsequence of $\{f_i'\}$ which, when suitably restricted, factor through σ_2. By passing to further subsequences, we can see that there is a subsequence of the f_i which all factor through $\eta_s\colon F_d \to L_s$.

We have already noted that a subsequence of the homomorphisms $\{h_n\}$ can be assumed to factor through L_s. Thus, by induction on j, suppose that there is a further subsequence of $\{h_n\}$ which factors through the canonical quotient map $F_d \to L_{j+1}$. The rigid vertex groups of the non-free factors of the Grushko decomposition of L_j inject into L_{j+1}, by the shortening argument. Since abelian subgroups of Γ are finitely generated (Lemma 2.2), we may assume by induction that the abelian subgroups of L_{j+1} are finitely generated. Thus the abelian vertex groups of the JSJ decomposition of a Grushko factor of L_j, as well as the edge groups, are finitely generated. This implies that the abelian subgroups of L_j are all finitely generated. The canonical quotient map $L_j \to L_{j+1}$ embeds the rigid vertex groups in the JSJ decomposition of L_j into L_{j+1}. Thus we may assume that a subsequence of $\{h_n\}$ factors through the canonical quotient map from a finitely generated subgroup of F_d to the rigid vertex groups of L_j. The other vertex groups in the JSJ decomposition of L_j are finitely presented, so there is some further subsequence of $\{h_n\}$ which factors through the canonical quotient to these vertex groups (the proof of these claims are almost identical to the proof of Proposition 4.2.(2)) . It now follows that some subsequence of $\{h_n\}$ factors through the canonical quotient $F_d \to L_j$. Thus, by induction, a subsequence of $\{h_n\}$ factors through $\eta_\infty\colon F_d \to T_\infty$, as required. □

The homomorphisms $\{h_n\colon F_d \to \Gamma\}$ were chosen so that for every index n there exists some element $f \in F_d$ for which $\eta_{n+1}(f) = 1$ and $h_n(f) \neq 1$. Now, $\eta_\infty\colon F_d \to T_\infty$ is the direct limit of the sequence

$$F_d \to T_1 \to T_2 \to \cdots$$

where the induced map from F_d to T_i is η_i. Thus, for all $n \geq 1$ and all $f \in F_d$, if $\eta_{n+1}(f) = 1$ then $\eta_\infty(f) = 1$. By Proposition 5.4 it is possible to extract a subsequence $\{h_{n_t}\colon F_d \to \Gamma\}$ that factors through the map $\eta_\infty\colon F_d \to T_\infty$, which is to say that there is a homomorphism $\pi_t\colon T_\infty \to \Gamma$ so that $h_{n_t} = \pi_t \circ \eta_\infty$. Hence, for every index t, and every element $f \in F_d$, if $\eta_{n_t+1}(f) = 1$ then $\eta_\infty(f) = 1$, which implies that $h_{n_t}(f) = 1$, in contradiction to the way that the homomorphisms h_n were chosen. This finally ends the proof of Theorem 5.2. □

Corollary 5.5 *Let Γ be a torsion-free relatively hyperbolic group with abelian parabolics, and let L be a Γ–limit group. Then L is Hopfian.*

Corollary 5.5 generalises, and gives an independent (though similar) proof of, [19, Theorem A], and implies that the relation defined on Γ–limit groups which are quotients of a fixed group G is a partial order. The proof of the following theorem is similar to that of [32, Theorem 1.17], and closely follows that of Proposition 5.4 above.

Theorem 5.6 *Suppose that G is a finitely generated group, and Γ a torsion-free relatively hyperbolic group with abelian parabolics, and let $\{h_n\colon G \to$*

$\Gamma\}$ be a sequence of homomorphisms converging to a Γ–limit group L, with canonical quotient $\pi\colon G \to L$. There exists a subsequence $\{h_{n_t}\}$ of $\{h_n\}$ so that each $h_{n_t}\colon G \to \Gamma$ factors through $\pi\colon G \to L$.

Proof The sequence $\{h_n\}$ converges to a faithful action of L on an $\mathbb{R}$–tree T. As in the proof of Theorem 5.4, take the Grushko decomposition of L, and attempt to form proper shortening quotients of the non-free factors. For each such proper quotient, take the Grushko decomposition and attempt to form further proper shortening quotients from the non-free factors, starting with a finer subsequence of $\{h_n\}$ at each stage. We eventually obtain a sequence (finite by Theorem 5.2)

$$G \to L \to L_2 \to \cdots \to L_s.$$

The rest of the proof is the same as that of Proposition 5.4. □

Theorem 5.6 is the tool which we will use in the study of Γ–limit groups as a replacement for the finite presentability of limit groups.

Corollary 5.7 *Let Γ be a torsion-free group which is hyperbolic relative to abelian subgroups, and suppose that L is a Γ–limit group. Then there is a sequence of homomorphisms $\{\rho_n\colon L \to \Gamma\}$ which converge to L.*

Proof In case L is not a strict Γ–limit group, it is isomorphic to a finitely generated subgroup of Γ, so we can take a constant sequence of embeddings. In case L is a strict Γ–limit group, the result follows from Theorem 5.6. □

Corollary 5.7 simplifies the definition of shortening quotients considerably. In particular, the groups U_n and W_n are no longer required. We briefly describe the simplified construction of shortening quotients.

Start with a non-abelian and freely indecomposable strict Γ–limit group L, and a sequence of homomorphisms $\{h_n\colon L \to \Gamma\}$ which converges to L. For each n, let $\widehat{h_n}$ be a short homomorphism which is equivalent to h_n (where now 'short' is as defined in Definition 3.5). Passing to a convergent subsequence of $\{\widehat{h_n}\colon L \to \Gamma\}$, we find a Γ–limit group Q which is a quotient of L.

In this context, Proposition 4.2.(1) follows immediately from Theorem 3.7. The homomorphisms h_n all have L as the domain, so Proposition 4.2.(2) is tautological in this context.

We record the simpler version of Proposition 4.2 here:

Proposition 5.8 *Suppose that Γ is a torsion-free relatively hyperbolic group with abelian parabolics, that L is a non-abelian and freely-indecomposable strict Γ–limit group, and that $\{h_n\colon L \to \Gamma\}$ is a sequence of homomorphisms converging to L. Suppose further that $\widehat{h_{n_t}}\colon L \to \Gamma$ is a convergent sequence so that $\widehat{h_{n_t}}$ is short and equivalent to h_{n_t}, for each t. Let Q be the resulting shortening quotient, with canonical epimorphism $\eta\colon L \to Q$.*

If the sequence $\widehat{h_{n_t}}$ is not contained in finitely many conjugacy classes then Q is a proper quotient of L.

Definition 5.9 Let Γ and G be finitely generated groups. We say that G is *fully residually* Γ if for every finite $\mathcal{F} \subset G$ there is a homomorphism $h\colon G \to \Gamma$ which is injective on $\mathcal{F}$.

Theorem 5.10 *Let Γ be a torsion-free relatively hyperbolic group with abelian parabolics. A finitely generated group G is a Γ–limit group if and only if it is fully residually Γ.*

Proof If G is a fully residually Γ then it is certainly a Γ–limit group. Conversely, suppose that G is a Γ–limit group. By Corollary 5.7, there is a sequence of homomorphisms $\{h_n\colon G \to \Gamma\}$ which converges to G. Using this sequence, it is clear that G is fully residually Γ. □

Proposition 5.11 *Let Γ be a torsion-free relatively hyperbolic group with abelian parabolics. Then there are only countably many Γ–limit groups.*

Proof If Γ is abelian, then all Γ–limit groups are finitely generated abelian groups, of which there are only countably many.

Suppose then that Γ is nonabelian. Let L be a Γ–limit group, and consider a sequence of shortening quotients of L:

$$L \to L_2 \to \ldots \to L_s$$

constructed as in the proof of Theorem 5.6, starting with a sequence $\{h_n\colon L \to \Gamma\}$ which converges to L. Note that none of the non-free factors of L_s admits a proper shortening quotient, and if $s > 1$ then each map in the sequence is a proper quotient. We will prove by induction that there are only countably many Γ–limit groups.

For the base case, note that L_s is a free product of a finitely generated free group and a finite collection of finitely generated subgroups of Γ. There are only countably many such groups.

Assume by induction there are only countably many groups which can be L_{j+1} in the sequence. Now, L_j is a free product of a finitely generated free group and freely indecomposable Γ–limit groups. Some of the free factors of L_j embed into L_{j+1} (if they do not admit a proper shortening quotient). If H_k is one of the other non-free Grushko factors of L_j, then H_k admits an abelian JSJ decomposition Λ_{H_k}. By Lemma 4.1, the rigid vertex groups of Λ_{H_k} embed into L_{j+1}. By induction and Lemma 2.2 we may assume that abelian subgroups of L_{j+1} are finitely generated. By the argument in the proof of Proposition 4.2, we see that the edge groups and vertex groups of Λ_{H_k} are finitely generated. Therefore each free factor of L_j is either (i) a finitely generated free group; (ii) a finitely generated subgroup of L_{j+1}; or (iii) can be formed by taking finitely many HNN extensions and amalgamated free products of finitely generated subgroups of L_{j+1} and finitely generated free groups over finitely generated abelian subgroups. There are only countably many such constructions, and so there are only countably many choices for L_j. This also implies that abelian subgroups of L_j are finitely generated (which we assumed by induction for L_{j+1}).

Therefore, by induction, there are only countably many choices for L. This completes the proof of the proposition. □

The following result was proved in the course of proving Proposition 5.11.

Corollary 5.12 *Any abelian subgroup of a Γ–limit group is a finitely generated free abelian group.*

Proposition 5.13 (Compare Proposition 1.20, [32]) *Let G be a finitely generated group and Γ a torsion-free relatively hyperbolic group with abelian parabolics. Let $R_1, R_2 \dots$ be a sequence of Γ–limit groups that are all quotients of G so that*

$$R_1 < R_2 < \cdots .$$

Then there exists a Γ–limit group R, a quotient of G, so that $R > R_m$ for all m.

Proof For each m, choose a homomorphism $h_m \colon G \to \Gamma$ that factors through the quotient map $\eta_m \colon G \to R_m$ as $h_m = h'_m \circ \eta_m$ and so that h'_m is injective on the ball of radius m in R_m. This is possible by Theorem 5.6.

A subsequence of $\{h_m\}$ converges to a Γ–limit group R, which is a quotient of G. By Theorem 5.6, we may assume that each element of this subsequence factors through the canonical quotient map $\eta \colon G \to R$.

We prove that $R > R_m$ for each m. We have quotient maps $\eta\colon G \to R$, and $\eta_i\colon G \to R_i$. Since $R_i < R_{i+1}$ there exists $\tau_i\colon R_{i+1} \to R_i$ so that $\eta_i = \tau_i \circ \eta_{i+1}$. In particular, $\ker(\eta_{i+1}) \subseteq \ker(\eta_i)$.

Let $\mathcal{A}$ be the fixed finite generating set for G. We attempt to define a homomorphism $\kappa_i\colon R \to R_i$ as follows: for $a \in \mathcal{A}$, define $\kappa_i(\eta(a)) = \eta_i(a)$. This is well-defined if and only if $\ker(\eta) \subseteq \ker(\eta_i)$. Therefore, suppose that $g \in \ker(\eta)$. Since each h_j factors through η, we have $h_j(g) = 1$ for all j. Suppose that g lies in the ball of radius n about the identity in the Cayley graph of G. Then for all j, the element $\eta_j(g)$ lies in the ball of radius n about the identity in R_j. Since $h_j(g) = 1$, for all j, and by the defining property of the h_j, if $k \geq n$ then $\eta_k(g) = 1$. Thus since for all j we have $\ker(\eta_{j+1}) \subseteq \ker(\eta_j)$ we have $\eta_i(g) = 1$, as required. We have constructed a homomorphism $\kappa_i\colon R \to R_i$ so that $\eta_i = \kappa_i \circ \eta$, which is to say that $R > R_i$. This finishes the proof. □

Propositions 5.11 and 5.13 imply that there are maximal elements for the set of Γ–limit groups which are quotients of a fixed finitely generated group G, under the order described before Theorem 5.2.

Recall that we say that two Γ–limit groups which are quotients of G, $\eta_1\colon G \to R_1$ and $\eta_2\colon G \to R_2$ are *equivalent* if there is an isomorphism $\tau\colon R_1 \to R_2$ so that $\eta_2 = \eta_1 \circ \tau$.

Proposition 5.14 (Compare [32, Proposition 1.21]) *Let G be a finitely generated group and Γ a torsion-free group hyperbolic relative to free abelian subgroups. Then there are only finitely many equivalence classes of maximal elements in the set of Γ–limit groups that are quotients of G.*

Proof The following proof was explained to me by Zlil Sela in the context of torsion-free hyperbolic groups. The same proof works in the current context.

Suppose on the contrary that there are infinitely many non-equivalent maximal Γ–limit groups $R_1, R_2, \ldots$, each a quotient of G. Let $\eta_i\colon G \to R_i$ be the canonical quotient map. Fixing a finite generating set $\mathcal{A}$ for G, we fix a finite generating set for each of the R_i, and hence obtain maps $\nu_i\colon F_d \to R_i$, where $d = |\mathcal{A}|$. There is a fixed quotient map $\pi\colon F_d \to G$ so that for each i we have $\nu_i = \eta_i \circ \pi$.

For each i, consider the set of words of length 1 in F_d that are mapped to the identity by ν_i. This set is finite for each i, and there is a bound on its size, so there is a subsequence of the R_i so that this set is the same for all i. Starting with this subsequence, consider those words of length 2 in F_d which are mapped to the identity by ν_i, and again there is a subsequence for which this (bounded) collection is the same for all i. Continue with this process for all lengths of words in F_d, passing to finer and finer subsequences, and consider the diagonal subsequence. We continue to denote this subsequence by $R_1, R_2, \ldots$.

Now, for each i, choose a homomorphism $h_i\colon F_d \to \Gamma$ so that for words w of length at most i in F_d, we have $h_i(w) = 1$ if and only if $\nu_i(w) = 1$, and so that h_i factors through the quotient map $\pi\colon F_d \to G$. This is possible because each R_i is a Γ–limit group which is a quotient of G.

A subsequence of $\{h_i\colon F_d \to \Gamma\}$ converges into a Γ–limit group M, which is a quotient of G since all h_i factor through π. Let $\psi\colon F_d \to M$ be the canonical quotient, and $\phi\colon G \to M$ the quotient for which $\psi = \phi \circ \pi$. Note that a word w of length at most i in F_d maps to the identity under ψ if and only if $\nu_i(w) = 1$.

Now, $R_1, R_2, \ldots$ are non-equivalent maximal Γ–limit quotients, so (possibly discarding one R_i which is equivalent to M) are all non-equivalent to M. Therefore, for each i there does not exist a homomorphism $\mu\colon M \to R_i$ so that $\nu_i = \mu \circ \psi$. That is to say that for each i there exists $u_i \in F_d$ so that $\psi(u_i) = 1$ but $\nu_i(u_i) \neq 1$.

Let $\{\tau_i\colon F_d \to \Gamma\}$ be a sequence of homomorphisms that all factor through $\pi\colon F_d \to G$ so that

- a word $w \in F_d$ of length at most i satisfies $\tau_i(w) = 1$ if and only if $\nu_i(w) = 1$ and
- $\tau_i(u_i) \neq 1$.

By Theorem 5.4 there is a subsequence $\{\tau_{n_i}\}$ of $\{\tau_i\}$ which converges into a Γ–limit group (which must be M) so that each τ_{n_i} factors through $\psi\colon F_d \to M$. Therefore, there is $r_i\colon M \to \Gamma$ so that $\tau_{n_i} = r_i \circ \psi$. However, we have that $\psi(u_{n_i}) = 1$, but $1 \neq \tau_{n_i}(u_{n_i}) = r_i(\psi(u_{n_i})) = r_i(1) = 1$, a contradiction. This contradicts the existence of $R_1, R_2, \ldots$, and finishes the proof. □

One of the motivations for the work done in this paper and its predecessors [17, 18, 19] were the following questions asked by Sela [33, I.8]: Let G be a CAT(0) group with isolated flats (see [21]).

I.8(i) Is G Hopf?

I.8(ii) Is it true that every system of equations over G is equivalent to a finite system?

I.8(iii) Is it possible to associate a Makanin–Razborov diagram to a system of equations over G?

I.8(iv) Is it possible to decide if a system of equations over G has a solution?

We believe that relatively hyperbolic groups with abelian parabolics form a more natural context for this question than CAT(0) groups with isolated flats. Note that 'virtually abelian' rather than 'abelian' would be a natural class (containing CAT(0) with isolated flats), but our methods only work for abelian parabolics. For torsion-free relatively hyperbolic groups with abelian parabolics, I.8(i) was answered by the author in [19] (and also follows from Corollary 5.5 above), I.8(ii) is the content of the next theorem, and I.8(iii) is answered in Section 6 below. We also remark that Dahmani [11, Theorem 0.2] has answered I.8(iv) for a class of groups which includes torsion-free relatively hyperbolic groups with virtually abelian parabolics.

Definition 5.15 Let G be a finitely generated group. Two systems of equations Φ and Φ' in finitely many variables over G are *equivalent* if the sets of solutions of Φ and Φ' are the same in G^n (where the number of variables is n).

Theorem 5.16 (Compare [32, Theorem 1.22]) *Suppose that Γ is a torsion-free relatively hyperbolic group with abelian parabolics. Then every system of equations in finitely many variables over Γ (without coefficients) is equivalent to a finite subsystem.*

Proof We follow the proof of [32, Theorem 1.22]. Let F be the free group with basis the variables of Σ. We implicitly use the fact that a solution to Σ over Γ corresponds to a homomorphism from F to Γ which sends each of the equations in Σ (considered as words in F) to $1 \in \Gamma$.

Let Σ be a system of equations in finitely many variables over Γ. We iteratively construct a directed locally finite tree as follows. Start with the first equation σ_1 in Σ, and associate with it a one relator group $G_1 = F/\langle \sigma_1 \rangle^F$. By Proposition 5.14, to G_1 is associated a finite number of maximal Γ–limit quotients. Denote these quotients of G_1 by $R_1, \ldots, R_m$. Place G_1 at the root node of a tree, and a directed edge from G_1 to each R_i. Note that if G_1 is a Γ–limit group then it is a maximal Γ–limit quotient of itself, and we do not need any new vertices at this stage.

Now let σ_2 be the second equation in Σ, and consider each R_i in turn. If σ_2 represents the trivial element of R_i, leave it unchanged. If σ_2 is nontrivial in R_i, define $\widehat{R_i} = R_i/\langle \sigma_2 \rangle^{R_i}$. With $\widehat{R_i}$, we associate its finite collection of maximal Γ–limit quotients, and extend the locally finite tree by adding new vertices for these quotients of R_i, and directed edges joining R_i to each of its quotients.

Continue this procedure iteratively. By Theorem 5.2, each branch of this locally finite tree is finite, and therefore by Konig's Lemma the entire tree is finite. This implies that the construction of this tree terminates after finitely many steps.

Let Σ' be the (finite) subset of Σ consisting of those equations σ_i considered before the above procedure terminates. We claim that the system Σ is equivalent to the subsystem Σ'. Certainly any solution to Σ is a solution to Σ'. Suppose that there is a solution to Σ' which is not a solution to Σ. Then there is a homomorphism $h\colon F \to \Gamma$ which sends each element of Σ' to 1 but so that $h(\sigma_j) \neq 1$ for some $\sigma_j \in \Sigma \smallsetminus \Sigma'$.

Let F be the free group with basis the variables of Σ. Any homomorphism from F to Γ which sends σ_1 to 1 factors through G_1, and then in turn factors through one of the maximal Γ–limit quotients $R_1, \ldots, R_m$ of G_1. Similarly, any homomorphism from F to Γ which sends σ_1 and σ_2 to 1 either factors through one of the R_i (in case σ_2 is trivial in R_i) or else through one of the maximal Γ–limit quotients of $\widehat{R_i} = R_i / \langle \sigma_2 \rangle^{R_i}$.

Arguing in this manner, we see that any homomorphism from F to Γ which sends all of Σ' to 1 factors through G_1 and then through some branch in the tree. Thus we find some terminal vertex group V of this tree so that $\sigma_j \neq 1$ in V (since $h(\sigma_j) \neq 1$). This implies that the procedure does not terminate when we consider V. This contradiction implies that all solutions to Σ' are solutions to Σ, as required. □

Guba [20] proved the analogous theorem for free groups, whilst Sela [32, Theorem 1.22] proved it for torsion-free hyperbolic groups.

6 Makanin–Razborov diagrams

In this final section, we describe the construction of *Makanin–Razborov* diagrams for Γ, which give a description of the set $\mathrm{Hom}(G, \Gamma)$, where G is an arbitrary finitely generated group. This is analogous to the constructions in [31, Section 5] and [32, Section 1].

Let R be a freely indecomposable Γ–limit group, and let $r_1, \ldots, r_m \in R$ be a fixed generating set for R. We assume that we always use the generating set

$\{r_1, \ldots, r_m\}$ to define the length of homomorphisms, and hence to find short homomorphisms.

Following [31, page 63] and [32, page 13] we say that two proper shortening quotients S_1, S_2 of R are *equivalent* if there is an isomorphism $\tau\colon S_1 \to S_2$ so that the canonical quotient maps $\eta_i\colon R \to S_i$, for $i = 1, 2$ satisfy $\eta_2 = \tau \circ \eta_1$. This defines an equivalence relation on the set of shortening quotients of R, paired with the canonical quotient maps: $\{(S_i, \eta_i\colon R \to S_i)\}$.

Let $SQ(R, r_1, \ldots, r_m)$ be the set of (proper) shortening quotients of R. On the set $SQ(R, r_1, \ldots, r_m)$ we define a partial order as follows: given two proper shortening quotients S_1, S_2 of R, along with canonical quotients $\eta_i\colon R \to S_i$, for $i = 1, 2$, we say that $S_1 > S_2$ if there exists a proper epimorphism $\nu\colon S_1 \to S_2$ so that $\eta_2 = \nu \circ \eta_1$.

Lemma 6.1 (Compare [32, Lemma 1.23]) *Let L be a freely-indecomposable Γ–limit group. Let $S_1 < S_2 < S_3 < \cdots$ (where $S_j \in SQ(L, r_1, \ldots, r_m)$) be a properly increasing sequence of (proper) shortening quotients of L. Then there exists a proper shortening quotient $R \in SQ(L, r_1, \ldots, r_m)$ so that for each j we have $R > S_j$.*

Proof Restricting to short homomorphisms throughout, the proof is identical to that of Proposition 5.13 above. We remark that the sequence of short homomorphisms $\{h_m\}$ obtained to find the shortening quotient R as in the proof of Proposition 5.13 is not contained in finitely many conjugacy classes. This is because the sequence of kernels $\ker(\eta_i)$ (where $\eta_i\colon L \to S_i$ is the canonical epimorphism) is strictly decreasing.

Thus, R is indeed a proper shortening quotient of L, by Proposition 4.2. □

Lemma 6.2 (Compare [32, Lemma 1.24]) *Let L be a freely-indecomposable Γ–limit group. The set, $SQ(L, r_1, \ldots, r_m)$, of (proper) shortening quotients of L contains only finitely many equivalence classes of maximal elements with respect to the partial order.*

Proof Once again, restricting throughout to short homomorphisms, the proof is almost identical to that of Proposition 5.14 above.

Once again, we need to ensure that the shortening quotient M obtained is a proper quotient of L. This will follow from Proposition 4.2 if we can ensure that the (short) homomorphisms $\{h_i\colon F_d \to \Gamma\}$ obtained do not belong to finitely many conjugacy classes. In order to ensure this, note that for each $i \neq j$, there exists $w_{ij} \in F_d$ so that $\nu_i(w_{ij}) = 1$ but $\nu_j(w_{ij}) \neq 1$, since R_i and R_j are maximal and inequivalent.

Now choose a short homomorphism $h_i\colon F_d \to \Gamma$ so that (i) for words w of length at most i in F_d we have $h_i(w) = 1$ if and only if $\nu_i(w) = 1$; (ii) h_i factors through $\pi\colon F_d \to G$; (iii) $h_i(w_{ij}) = 1$ for $j = 1, \ldots, i-1$; and (iv) $h_i(w_{ki}) \neq 1$ for $k = i+1, \ldots, 2i$.

Conditions (iii) and (iv) ensure that for each i the set $\{h_i, \ldots, h_{2i}\}$ have distinct kernels and so belong to different conjugacy classes. Therefore, the h_i do not belong to finitely many conjugacy classes.

The rest of the proof is identical to that of Proposition 5.14 above. □

We can now use shortening quotients to 'encode and simplify' all homomorphisms from a freely-indecomposable Γ–limit group into Γ.

Proposition 6.3 (Compare [32, Proposition 1.25]) *Suppose that R is a freely-indecomposable Γ–limit group. Let $r_1, \ldots, r_m \in R$ be a generating set for R, and let $M_1, \ldots, M_k$ be a set of representatives of the (finite) set of equivalence classes of maximal (proper) shortening quotients in $SQ(R, r_1, \ldots, r_m)$, equipped with the canonical quotient maps $\eta_i\colon R \to M_i$, for $i = 1, \ldots, k$.*

Let $h\colon R \to \Gamma$ be a homomorphism which is not equivalent to an embedding. Then there exist a (not necessarily unique) index $1 \leq i \leq k$, a homomorphism $\widehat{h}\colon R \to \Gamma$ so that $\widehat{h} \sim h$, and a homomorphism $h_{M_i}\colon M_i \to \Gamma$ so that $\widehat{h} = h_{M_i} \circ \eta_i$.

Proof Choose $\widehat{h} \sim h$ so that $\widehat{h}$ is short. By hypothesis, $\widehat{h}$ is not injective. The constant sequence $\widehat{h}, \widehat{h}, \ldots$ converges into a proper shortening quotient S of R. Now, $S \cong \widehat{h}(R)$, and the canonical quotient map is just $\widehat{h}$. By Lemma 6.2, there exists some M_i so that $M_i > S$ or M_i is equivalent to S. The result now follows. □

We use Propositions 5.14 and 6.3 to prove Theorem 1.1.

Proof of Theorem 1.1 Let Γ be a torsion-free relatively hyperbolic group with abelian parabolics, G a freely indecomposable finitely generated group and $h\colon G \to \Gamma$ a homomorphism.

The image $h(G)$ is a Γ–limit group. By Proposition 5.14, there are finitely many equivalence classes of maximal Γ–limit groups which are quotients of G, and h factors through one of these maximal Γ–limit groups. If G is not a Γ–limit group, then these maximal Γ–limit quotients of G are proper quotients, and the theorem is proved in this case.

Thus we may assume that G is a Γ–limit group, and the result now follows from Proposition 6.3. □

Remark 6.4 It is reasonable in the hypothesis to restrict to freely indecomposable groups G, because if $G = A * B$ then $\mathrm{Hom}(G,\Gamma) = \mathrm{Hom}(A,\Gamma) \times \mathrm{Hom}(B,\Gamma)$. Thus a homomorphism $h\colon G \to \Gamma$ induces a pair of homomorphisms, $h_A\colon A \to \Gamma$ and $h_B\colon B \to \Gamma$, and vice versa.

Hence, it is easy to understand the homomorphisms from an arbitrary finitely generated group G to Γ in terms of the sets of homomorphisms of the free factors in the Grushko decomposition of G to Γ.

Finally, we now construct *Makanin–Razborov diagrams* over Γ. Let G be an arbitrary finitely generated group. A Makanin–Razborov diagram is a finite (directed) tree associated to a finitely generated group G which encodes the set $\mathrm{Hom}(G,\Gamma)$.

There is a root vertex, labelled by G. All other vertices are labelled by Γ–limit groups. We start with G.

Suppose first that G is not a Γ–limit group, then by Proposition 5.14 there are finitely many equivalence classes of maximal Γ–limit quotients of G. Let $R_1, \ldots, R_s$ be a collection of representatives of these equivalence classes. Add a new vertex for each R_i, and a directed edge joining G to R_i, labelled by the canonical quotient map from G to R_i.

These are the only edges emanating from G (unless G is a Γ–limit group). All other edges will be between a pair of Γ–limit groups.

If G *is* a Γ–limit group, it is analysed in exactly the same way as all of the other vertices. We proceed with this analysis now.

There are two kinds of edges. Suppose that the Γ–limit group R is the label of a vertex, and the Grushko decomposition of R is $R = H_1 * \ldots * H_k * F$, where each H_i is freely indecomposable, noncyclic and finitely generated, and F is a finitely generated free group. Then we add new vertices, one for each H_i and one for F, and directed edges from R to each H_i and to F. We do not label this first kind of edge.

We now describe the second kind of edge. Suppose that R is a noncyclic freely indecomposable Γ–limit group labelling some vertex. By Lemma 6.2, R admits only finitely many equivalence classes of maximal proper shortening quotients. Add a new vertex for each equivalence class, and a directed edge from R to each of these quotients, labelled by the canonical quotient map.

The vertices labelled by finitely generated free groups are terminal vertices of the tree.

We thus build the Makanin–Razborov diagram, successively taking Grushko decompositions and maximal proper shortening quotients, if possible.

This tree obviously has finite width, and by Theorem 5.2 each branch has finite length. Therefore, by Konig's Lemma, the tree is finite.

We now describe how the Makanin–Razborov diagram for G encodes the set of all homomorphisms from G to Γ. Denote the Makanin–Razborov diagram for G by MR_G.

Any homomorphism from G to Γ factors through one of the R_i. Similarly, if L is the label of a vertex, and L is noncyclic and freely indecomposable, with $M_1, \ldots, M_k$ the maximal proper shortening quotients of L, then any homomorphism from L to Γ which is not equivalent to an embedding factors through one of the M_i.

For any homomorphism $h\colon G \to \Gamma$ we can find a sub-tree Δ_h of MR_G. Each edge of Δ_h is labelled as in MR_G (or not labelled, as the case may be). Each vertex v has two labels: the group H_v which labels it in MR_G, and also a homomorphism $h_v\colon H_v \to \Gamma$.

The diagram Δ_h encodes a factorisation of h, and is constructed as follows: The root vertex of Δ_h is the root vertex of MR_G, labelled by G and h.

If G is not a Γ–limit group, then there is some maximal Γ–limit quotient R_i of G (a proper quotient), equipped with the canonical quotient $\eta\colon G \to R_i$, so that $h = h' \circ \eta$ for some homomorphism $h'\colon R_i \to \Gamma$. Add to Δ_h the edge labelled by η, and vertex labelled R_i (which we label by R_i and h').

From R_i (respectively G, in case G is a Γ–limit group), we add the vertices corresponding to the Grushko decomposition. These vertices are labelled by the Grushko factors, and the homomorphism induced on these free factors by h' (respectively h), as in Remark 6.4.

From a free Grushko factor F_n, of rank n, the set $\mathrm{Hom}(F_n, \Gamma)$ is naturally parametrised by Γ^n. The vertex labelled by F_n is a terminal vertex of MR_G, and of Δ_h also (whatever the other label may be).

Suppose that v is a vertex of Δ_h, labelled by a noncyclic freely indecomposable Γ–limit group L, and an associated homomorphism $h_L\colon L \to \Gamma$. Then either

(i) h_L is equivalent to an embedding $\widehat{h_L}\colon L \to \Gamma$, or

(ii) h_L is equivalent to some homomorphism $\widehat{h_L}$ which factors through one of the maximal proper shortening quotients of L.

In the first case, the associated vertex is a terminal vertex of Δ_h (even though it may not be a terminal vertex of MR_G; see Remark 6.5 below). In the second case, let M be the maximal shortening quotient, $\rho_M\colon L \to M$ the canonical quotient, and $\widehat{h_L} = h_M \circ \rho_M$. To Δ_h we add the edge labelled by ρ_M, and the vertex from MR_G labelled by M, which we label by M and h_M.

We then take the Grushko decomposition of M, and proceed by repeating the above analysis. In this way, we obtain a tree Δ_h which encodes a factorisation of h. At each vertex v, with labels $(G_v, h_v\colon G_v \to \Gamma)$, one of the following is performed:

(1) Take the Grushko decomoposition of G_v, and consider the homomorphisms induced by h_v on the free factors in turn.
(2) Replace h_v by an equivalent homomorphism $\widehat{h_v}\colon G_v \to \Gamma$ which is an injection, and stop the analysis of the branch at this vertex.
(3) Replace h_v by an equivalent homomorphism $\widehat{h_v}\colon G_v \to \Gamma$, where $\widehat{h_v} = h_M \circ \eta$, for some maximal proper shortening quotient $\eta\colon G_v \to M$ of G_v, and proceed by analysing (M, h_M).
(4) Stop the analysis of a branch at a free group, since $\mathrm{Hom}(F, \Gamma)$ is easily understood.

Note that in general the diagram Δ_h need not be unique.

Remark 6.5 It is worth remarking that the terminal vertices of Δ_h need not be terminal vertices of MR_G. This is because it is possible that some freely indecomposable finitely generated subgroup H of Γ is also a strict Γ–limit group which admits a proper shortening quotient. Thus, some homomorphisms from H to Γ are not equivalent to an injection, but some are.

This phenomena arises in the case of Makanin–Razborov diagrams over torsion-free hyperbolic groups (as in [32]), but not in Makanin–Razborov diagrams over free groups (because finitely generated subgroups of free groups are free, and free groups are always terminal vertices).

Remark 6.5 captures one of the key differences between the Makanin–Razborov diagrams for free groups (as constructed in [31, Section 5]) and those constructed in this paper for torsion-free relatively hyperbolic groups with abelian parabolics.

There is one other difference which is worthy of remark. In the case of free groups (and of torsion-free hyperbolic groups), every non-injective homomorphism from a Γ–limit group L to the target group Γ is equivalent to one which factors through a maximal shortening quotient of L. This equivalent homomorphism can be realised by pre-composing with an element of $\mathrm{Mod}(L)$, and post-composing with an inner automorphism of Γ. Since this conjugation does not change the kernel, we see that for every $h\colon L \to \Gamma$ there is a maximal shortening quotient M of L, and $\phi \in \mathrm{Mod}(L)$ so that $h \circ \phi$ factors through M. Applying this analysis repeatedly leads to a factorisation of any homomorphism as a composition of modular automorphisms and canonical quotient maps, followed finally by a product of embeddings and maps from finitely generated free groups to the target group Γ.

If Γ is a torsion-free relatively hyperbolic group with abelian parabolics, the situation is slightly more complicated. It is still the case that if L is a Γ–limit group then for any $h\colon L \to \Gamma$ which is not equivalent to an embedding, there is a maximal shortening quotient M of L and $\hat{h} \sim h$ so that $\hat{h}$ factors through M. However, the equivalence between h and $\hat{h}$ now involves bending moves, as well as pre-composition with an element of $\mathrm{Mod}(L)$ and conjugation in Γ. Since the bending moves can change the kernel of a map,[4] it is not necessarily the case that there is $\phi \in \mathrm{Mod}(L)$ so that $h \circ \phi$ factors through some maximal shortening quotient. Therefore, we do not get a factorisation analogous to that for free and torsion-free hyperbolic groups.

The construction of Makanin–Razborov diagrams for torsion-free relatively hyperbolic groups with abelian parabolics (essentially) answers a question asked by Sela [33, Problem I.8(iii)]. See the discussion above Theorem 5.16.

References

[1] **E Alibegović**, *A combination theorem for relatively hyperbolic groups*, Bull. London Math. Soc. 37 (2005) 459–466

[2] **E Alibegović**, *Makanin–Razborov diagrams for limit groups*, preprint

[3] **M Bestvina**, **M Feighn**, *Bounding the complexity of simplicial group actions on trees*, Invent. Math. 103 (1991) 449–469

[4] **M Bestvina**, **M Feighn**, *Stable actions of groups on real trees*, Invent. Math. 121 (1995) 287–321

[4] I thank the referee for pointing out this important fact.

[5] **M Bestvina**, **M Feighn**, *Notes on Sela's work: Limit groups and Makanin–Razborov diagrams*, preprint

[6] **M Bestvina**, *Questions in geometric group theory*, available at `http://www.math.utah.edu/~bestvina`

[7] **B Bowditch**, *Relatively hyperbolic groups*, preprint

[8] **F Dahmani**, *Classifying spaces and boundaries for relatively hyperbolic groups*, Proc. London Math. Soc. (3) 86 (2003) 666–684

[9] **F Dahmani**, *Combination of convergence groups*, Geom. Topol. 7 (2003) 933–963

[10] **F Dahmani**, *Accidental parabolics and relatively hyperbolic groups*, Israel J. Math. to appear

[11] **F Dahmani**, *On equations in relatively hyperbolic groups*, preprint

[12] **C Druţu**, **M Sapir**, *Tree-graded spaces and asymptotic cones of groups*, Topology 44 (2005) 959–1058

[13] **C Druţu**, **M Sapir**, *Relatively hyperbolic groups with rapid decay property*, Int. Math. Res. Not. (2005) 1181–1194

[14] **M J Dunwoody**, *The accessibility of finitely presented groups*, Invent. Math. 81 (1985) 449–457

[15] **B Farb**, *Relatively hyperbolic groups*, Geom. Funct. Anal. 8 (1998) 810–840

[16] **M Gromov**, *Hyperbolic groups*, from: "Essays in group theory", Math. Sci. Res. Inst. Publ. 8, Springer, New York (1987) 75–263

[17] **D Groves**, *Limits of (certain) CAT(0) groups. I. Compactification*, Algebr. Geom. Topol. 5 (2005) 1325–1364

[18] **D Groves**, *Limits of (certain) $CAT(0)$ groups, II: The Hopf property and the shortening argument*, `arXiv:math.GR/0408080`

[19] **D Groves**, *Limit groups for relatively hyperbolic groups, I: The basic tools*, `arXiv:math.GR/0412492`

[20] **V S Guba**, *Equivalence of infinite systems of equations in free groups and semigroups to finite subsystems*, Mat. Zametki 40 (1986) 321–324, 428

[21] **G C Hruska**, *Nonpositively curved spaces with isolated flats*, PhD thesis, Cornell University (2002)

[22] **O Kharlampovich**, **A Myasnikov**, *Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz*, J. Alg. 200 (1998) 472–516

[23] **O Kharlampovich**, **A Myasnikov**, *Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups*, J. Alg. 200 (1998) 517–570

[24] **D Osin**, *Relatively hyperbolic groups: Intrinsic geometry, algebraic properties, and algorithmic properties*, preprint

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

[26] **E Rips**, **Z Sela**, *Structure and rigidity in hyperbolic groups. I*, Geom. Funct. Anal. 4 (1994) 337–371

[27] **E Rips**, **Z Sela**, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. (2) 146 (1997) 53–109

[28] **Z Sela**, *Acylindrical accessibility for groups*, Invent. Math. 129 (1997) 527–565

[29] **Z Sela**, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank* 1 *Lie groups. II*, Geom. Funct. Anal. 7 (1997) 561–593

[30] **Z Sela**, *Endomorphisms of hyperbolic groups. I. The Hopf property*, Topology 38 (1999) 301–321

[31] **Z Sela**, *Diophantine geometry over groups. I. Makanin–Razborov diagrams*, Publ. Math. Inst. Hautes Études Sci. (2001) 31–105

[32] **Z Sela**, *Diophantine geometry over groups VIII: The elementary theory of a hyperbolic group*, preprint

[33] **Z Sela**, *Diophantine geometry over groups: a list of research problems*, `http://www.ma.huji.ac.il/~zlil/problems.dvi`

[34] **J-P Serre**, *Trees*, Springer–Verlag, Berlin (1980)

[35] **A Yaman**, *A topological characterisation of relatively hyperbolic groups*, J. Reine Angew. Math. 566 (2004) 41–89