

ANNALES MATHÉMATIQUES



BLAISE PASCAL

RICHARD WEIDMANN & CORNELIUS REINFELDT

Makanin–Razborov diagrams for hyperbolic groups

Volume 26, n° 2 (2019), p. 119-208.

http://ambp.centre-mersenne.org/item?id=AMBP_2019__26_2_119_0



Cet article est mis à disposition selon les termes de la licence
CREATIVE COMMONS ATTRIBUTION 4.0.
<https://creativecommons.org/licenses/4.0/>

L'accès aux articles de la revue « Annales mathématiques Blaise Pascal »
(<http://ambp.centre-mersenne.org/>), implique l'accord avec les conditions
générales d'utilisation (<http://ambp.centre-mersenne.org/legal/>).

*Publication éditée par le laboratoire de mathématiques Blaise Pascal
de l'université Clermont Auvergne, UMR 6620 du CNRS
Clermont-Ferrand — France*



Publication membre du
Centre Mersenne pour l'édition scientifique ouverte
<http://www.centre-mersenne.org/>

Makanin–Razborov diagrams for hyperbolic groups

RICHARD WEIDMANN
CORNELIUS REINFELDT

Abstract

We give a detailed account of Zlil Sela’s construction of Makanin–Razborov diagrams describing $\text{Hom}(G, \Gamma)$ where G is a finitely generated group and Γ is a hyperbolic group. We also deal with the case where Γ has torsion.

Les diagrammes Makanin–Razborov pour les groupes hyperboliques

Résumé

Nous proposons une présentation détaillée de la construction des diagrammes de Makanin–Razborov de Zlil Sela qui décrivent $\text{Hom}(G, \Gamma)$ pour un groupe G de type fini et un groupe hyperbolique Γ . De plus, nous traitons le cas où Γ est un groupe ayant de la torsion.

1. Introduction

The question whether one can decide if a system of equations (with constants) in a free group has a solution was answered affirmatively by Makanin [31] who described an algorithm that produces such a solution if it exists and says *No* otherwise. In his groundbreaking work he introduced a rewriting process for systems of equations in the free semigroup. This process was later refined by Razborov to give a complete description of the set of solutions for a system of equations in a free group [32, 33]. This description is now referred to as Makanin–Razborov diagrams.

Rips recognized that the Makanin process can be adapted to study group actions on real trees which gave rise to what is now called the Rips machine, a structure theorem for finitely presented groups acting on real trees similar to Bass–Serre theory for groups acting on simplicial trees, see [5, 17, 21]. This has been generalized to finitely generated groups by Sela [38] and further refined by Guirardel [22]. Recently Dahmani and Guirardel [11] have in turn used the geometric ideas underlying the Rips theory to provide an alternative version of Makanin’s algorithm.

Razborov’s original description of the solution set has been refined independently by Kharlampovich and Myasnikov [26, 27] and Sela [40]; this description has been an important tool in their solutions to the Tarski problems regarding the elementary theory of free groups. Kharlampovich and Myasnikov modified Razborov’s methods to obtain

Keywords: Makanin–Razborov Diagrams.

their own version of the rewriting process while Sela used the Rips machine extensively, bypassing much of the combinatorics of the process.

While it seems that the full potential of the ideas underlying the Makanin process has not yet been realized there are a number of generalizations of the above results. Sela [41] has shown the existence of Makanin–Razborov diagrams for torsion-free hyperbolic groups, which was then generalized to torsion-free relatively hyperbolic groups with finitely generated Abelian parabolic subgroups by Groves [19]. Makanin–Razborov diagrams for free products have been constructed independently by Jaligot and Sela following Sela’s geometric approach and by Casals-Ruiz and Kazachkov [10] following the combinatorial approach of Kharlampovich and Myasnikov. Casals-Ruiz and Kazachkov also constructed Makanin–Razborov diagrams for graph groups [10].

It is the purpose of this article to give a detailed account of Sela’s construction of Makanin–Razborov diagrams for hyperbolic groups while also removing the torsion-freeness assumption. It should be noted that Dahmani and Guirardel’s version of the Makanin algorithm also applies to hyperbolic groups with torsion.

We are only dealing with equations without constants, no new ideas are needed to deal with them. Thus we are interested in finding all tuples $(x_1, \dots, x_n) \in \Gamma^n$ that satisfy equations

$$w_i(x_1, \dots, x_n) = 1$$

for $i \in I$ where w_i is some word in the $x_j^{\pm 1}$ and Γ is a hyperbolic group. It is clear that these solutions are in 1-to-1 correspondence to homomorphisms from

$$G = \langle x_1, \dots, x_n \mid w_i(x_1, \dots, x_n), i \in I \rangle$$

to Γ . Thus parametrizing the set of solutions to the above system of equations is equivalent to parametrizing $\text{Hom}(G, \Gamma)$.

Theorem 1.1. *Let Γ be a hyperbolic group and G be a finitely generated group. Then there exists a finite directed rooted tree T with root v_0 satisfying*

- (1) *The vertex v_0 is labeled by G .*
- (2) *Any vertex $v \in VT$, $v \neq v_0$, is labeled by a group G_v that is fully residually Γ .*
- (3) *Any edge $e \in ET$ is labeled by an epimorphism $\pi_e : G_{\alpha(e)} \rightarrow G_{\omega(e)}$*

such that for any homomorphism $\phi : G \rightarrow \Gamma$ there exists a directed path e_1, \dots, e_k from v_0 to some vertex $\omega(e_k)$ such that

$$\phi = \psi \circ \pi_{e_k} \circ \alpha_{k-1} \circ \dots \circ \alpha_1 \circ \pi_{e_1}$$

where $\alpha_i \in \text{Aut } G_{\omega(e_i)}$ for $1 \leq i \leq k$ and $\psi : G_{\omega(e_k)} \rightarrow \Gamma$ is locally injective.

Here a homomorphism is called locally injective if it is injective when restricted to 1-ended and finite subgroups, thus it is injective on the vertex groups of any Dunwoody/Linnell decomposition.

The proof broadly follows Sela's proof in the torsion-free case but is also partly inspired by Bestvina and Feighn's exposition of Sela's construction in the case of free groups [6] and by Groves's adaption of Sela's work to the relatively hyperbolic case. We further rely on Guirardel's version of the Rips machine.

When following Sela's strategy to prove Theorem 1.1 it is almost unavoidable to also prove the following results that are all proven in Section 7 and follow from the same stream of ideas. All results are due to Sela in the torsion-free case. Theorem 1 reconciles two related notions of Γ -limit groups in the case of hyperbolic Γ , see Section 2.1 for definitions.

Corollary 7.6. *Let Γ be a hyperbolic group. Then finitely generated Γ -limit groups are fully residually Γ .*

Recall that a group G is called *Hopfian* if every epimorphism $G \rightarrow G$ is an isomorphism.

Corollary 7.9. *Hyperbolic groups are Hopfian.*

A group G is called *equationally Noetherian* if for every system of equations in G there exists a finite subsystem that has the same set of solutions.

Corollary 7.13. *Hyperbolic groups are equationally Noetherian.*

In Section 2 we introduce Γ -limit groups, these are the groups that occur as vertex groups in the Makanin–Razborov diagram. Moreover, we illustrate how Γ -limit groups admit actions on real trees if Γ is a hyperbolic group and study the stability properties of these actions. In Section 3 and Section 4 we discuss the Rips machine and the JSJ-decomposition of Γ -limit groups before we give a detailed discussion of the shortening argument in Section 5. The shortening argument in particular implies that the Makanin–Razborov diagrams are locally finite. In Section 6 we then construct Makanin–Razborov diagrams for hyperbolic groups under the additional assumption that Γ is equationally Noetherian. In Section 7 we then discuss Sela's shortening quotients and prove that all hyperbolic groups are in fact equationally Noetherian, i.e. that the construction in Section 6 applies to all hyperbolic groups. The authors would like to thank Abderezak Ould Houcine for pointing out the content of Lemma 2.22(3) and Lemma 6.1.

A first version of this paper circulated in 2010. We do not discuss the numerous more recent developments in the field.

2. Γ -limit groups and their actions on real trees

In this chapter we introduce the concept of a Γ -limit group. While we will later on solely be interested in the case where Γ is hyperbolic, the notion of a Γ -limit group can be formulated for any group Γ . In the hyperbolic case, a Γ -limit group naturally admits actions on metric \mathbb{R} -trees, which arise as limits of group actions on the Cayley graph of Γ (w.r.t. a fixed finite generating set). These Γ -limit groups will occur in the Makanin–Razborov diagrams. In Section 2.2 we show in detail how these limit actions arise, and in Section 2.3 we prove important stability properties of these actions. We will then study the structure of virtually Abelian subgroups of Γ -limit groups. Most of the material is standard except that we have to deal with torsion.

2.1. Γ -limit groups

Throughout this section Γ is an arbitrary group. We will discuss sequences of homomorphisms from a given (f.g.) group G to Γ . We denote by $\text{Hom}(G, \Gamma)$ the set of all such homomorphisms. Moreover, if $(\varphi_i)_{i \in \mathbb{N}}$ is a sequence of homomorphisms from G to Γ , we simply write $(\varphi_i) \subset \text{Hom}(G, \Gamma)$.

With these notations we give a definition of a Γ -limit group which is analogous to the definition of Bestvina and Feighn [6] in the case where Γ is free.

Definition 2.1. Let G be a group and $(\varphi_i) \subset \text{Hom}(G, \Gamma)$. The sequence (φ_i) is *stable* if for any $g \in G$ either $\varphi_i(g) = 1$ for almost all i or $\varphi_i(g) \neq 1$ for almost all i . If (φ_i) is stable then the *stable kernel* of the sequence, denoted by $\underline{\ker}(\varphi_i)$, is defined as

$$\underline{\ker}(\varphi_i) := \{g \in G \mid \varphi_i(g) = 1 \text{ for almost all } i\}.$$

We then call the quotient $G/\underline{\ker}(\varphi_i)$ the Γ -limit group associated to (φ_i) and the projection $\varphi : G \rightarrow G/\underline{\ker}(\varphi_i)$ the Γ -limit map associated to (φ_i) .

We call a quotient map $\varphi : G \rightarrow G/N$ a Γ -limit map if it is the Γ -limit map associated to some stable sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$. We will denote the Γ -limit group G/N by L_φ .

We further call the sequence (φ_i) *stably injective* if (φ_i) is stable and $\underline{\ker}(\varphi_i) = 1$.

Remark 2.2. Note that any subgroup H of a Γ -limit group $L = L_\varphi$ is a Γ -limit group. This can be seen by considering the stable sequence $(\varphi_i|_{\varphi^{-1}(H)}) \subset \text{Hom}(\varphi^{-1}(H), \Gamma)$ with associated Γ -limit group H .

Clearly Γ and every subgroup of Γ are Γ -limit groups.

It turns out that Γ -limit groups are closely related to fully residually Γ groups: A group G is called *residually Γ* if for every non-trivial element $g \in G$ there exists a

homomorphism $\varphi : G \rightarrow \Gamma$ s.th. $\varphi(g) \neq 1$. Further, G is called *fully residually* Γ if for any finite set $S \subset G$ there exists a homomorphism $\varphi : G \rightarrow \Gamma$ such that $\varphi|_S$ is injective. Note that any subgroup of Γ is fully residually Γ . We can immediately verify that in the case of countable groups, every fully residually Γ group is a Γ -limit group:

Lemma 2.3. *If G is countable and fully residually Γ , then G is a Γ -limit group.*

Proof. Choose a surjective map $f : \mathbb{N} \rightarrow G$. For each $i \in \mathbb{N}$ let $M_i := \{f(j) \mid j \leq i\}$ and choose $\varphi_i : G \rightarrow \Gamma$ such that $\varphi_i|_{M_i}$ is injective. Clearly the sequence (φ_i) is stably injective. Thus $G = G/\varinjlim(\varphi_i)$ is a Γ -limit group. \square

It turns out that in many situations, in particular in the case where Γ is hyperbolic, a partial converse is true as well, namely f.g. Γ -limit groups are fully residually Γ , see Corollary 7.6 below. This is more generally true whenever Γ is *equationally Noetherian*, which we will prove in Section 6.1.

From now on we will almost exclusively consider finitely generated groups.

2.2. Limit actions on real trees

Throughout this section G is a finitely generated group endowed with a finite generating set S_G . Call a pseudo-metric d on G (left) G -invariant if

$$d(g, h) = d(kg, kh) \quad \text{for all } g, h, k \in G.$$

We consider $\mathcal{A}(G)$, the space of all G -invariant pseudo-metrics on G , with the compact-open topology. Thus a sequence (d_i) of G -invariant pseudo-metrics on G converges in $\mathcal{A}(G)$ iff the sequence $(d_i(1, g))$ converges in \mathbb{R} for all $g \in G$.

Lemma 2.4. *Let $(d_i) \subset \mathcal{A}(G)$ be a sequence of G -invariant pseudo-metrics. If there exists $\lambda \in \mathbb{R}$ such that for each $i \in \mathbb{N}$ and $s \in S_G$, $d_i(1, s) \leq \lambda$, then (d_i) has a subsequence which converges in $\mathcal{A}(G)$.*

Proof. For $k \in \mathbb{N}$, let $B_k := \{g \in G \mid d_{S_G}(1, g) \leq k\}$ (where d_{S_G} denotes the word metric on G w.r.t. S_G). If λ is as above, it follows that for all $k \in \mathbb{N}$ and $g \in B_k$,

$$d_i(1, g) \leq k\lambda.$$

As B_k is finite, the compactness of the cube $[0, k\lambda]^{|B_k|}$ then implies that there is a subsequence $(d_{i_j,k})_{j \in \mathbb{N}} \subset (d_i)$ such that for all $g \in B_k$, the sequence

$$\left(d_{i_j,k}(1, g) \right)_{j \in \mathbb{N}} \subset \mathbb{R}$$

converges. Moreover, each sequence $(d_{i_j,k})_{j \in \mathbb{N}}$ may be chosen as a subsequence of $(d_{i_j,k-1})_{j \in \mathbb{N}}$. It is now obvious that the diagonal sequence $(d_{i_{j,k}})_{k \in \mathbb{N}}$ converges in $\mathcal{A}(G)$. \square

Remark 2.5. Note that the condition that $d_i(1, s) \leq \lambda$ for all $s \in S$ and $i \in \mathbb{N}$ is essential, i.e. in general, sequences of pseudo-metrics do not have convergent subsequences. In the literature this is often bypassed by projectivizing the space $\mathcal{A}(G)$ and thereby compactifying it. We will not adopt this point of view.

In the following we study sequences of pseudo-metrics on a group G that are induced by actions on based hyperbolic G -spaces.

Throughout this article we use the following definition of δ -hyperbolicity: A pseudo-metric space (X, d) (resp. a pseudo-metric d) is called δ -hyperbolic for $\delta \geq 0$ if for any $x, y, z, t \in X$,

$$(x|y)_t \geq \min((x|z)_t, (y|z)_t) - \frac{\delta}{3}, \quad (2.1)$$

where $(x|y)_t$ is the Gromov product of x and y with respect to t given by

$$(x|y)_t := \frac{1}{2} (d(x, t) + d(y, t) - d(x, y)). \quad (2.2)$$

Note that the constant being $\frac{\delta}{3}$ rather than δ is slightly non-standard. For a geodesic metric space this choice of constant implies the δ -thinness of geodesic triangles, i.e. that for any geodesic triangle $[x, y] \cup [y, z] \cup [z, x]$ we have

$$[x, y] \subset N_\delta([y, z] \cup [z, x]),$$

see [1]. The definition via the Gromov product has the advantage that it also applies to pseudo-metric spaces that are not geodesic.

In the following we study group actions. An *action* of a group G on a metric space X is a homomorphism

$$\rho : G \rightarrow \text{Isom}(X)$$

from G to the isometry group of X . A (based) G -space is a tuple (X, x_0, ρ) consisting of a metric space X , a base point $x_0 \in X$ and an action ρ of G on X . If $g \in G$ and $x \in X$, for convenience we will denote the element $\rho(g)(x) \in X$ simply by gx if the action ρ is understood. Moreover, when we want to explicitly state the action, we use the notation ${}_\rho gx := \rho(g)(x)$ to improve readability.

We will study pseudo-metrics that are induced by actions on metric spaces. Let $X = (X, x, \rho)$ be a based G -space. Then the G -action ρ on X induces a pseudo-metric

$$d_\rho^x : G \times G \rightarrow \mathbb{R}_{\geq 0}$$

on G , given by

$$d_\rho^x(g, h) = d_X({}_\rho gx, {}_\rho hx).$$

Note that this pseudo-metric is G -invariant, and that d_ρ^x is δ -hyperbolic if d_X is δ -hyperbolic. We will denote d_ρ^x simply by d_ρ if the basepoint is clear from the context.

Lemma 2.6. For $i \in \mathbb{N}$, let $\delta_i \geq 0$ and $X_i = (X_i, x_i, \rho_i)$ a based δ_i -hyperbolic G -space. Assume that the sequence $(d_{\rho_i}^{x_i})$ converges in $\mathcal{A}(G)$ to a limit sequence d_∞ and

$$\lim_{i \rightarrow \infty} \delta_i = 0.$$

Then there exists a based real G -tree (T, x, ρ) such that T is spanned by ${}_\rho Gx$, and $d_\rho^x = d_\infty$.




FIGURE 2.1. A quadrilateral degenerating to a tree

Proof. As $\lim_{i \rightarrow \infty} \delta_i = 0$, it follows that d_∞ is 0-hyperbolic. We obtain a G -action on the 0-hyperbolic metric space $(\widehat{G}, \widehat{d}_\infty)$ obtained from the pseudo-metric space (G, d_∞) by metric identification, i.e. by identifying points of distance 0. Then there exists a real G -tree T satisfying

- T admits a G -equivariant isometric embedding $\eta : \widehat{G} \rightarrow T$,
- $\eta(\widehat{G})$ spans T , i.e. no proper subtree of T contains $\eta(\widehat{G})$.

We sketch the way T is constructed, for details we refer to Lemma 2.13 of [3]: Start with the metric space $(\widehat{G}, \widehat{d}_\infty)$ and for any $x, y \in \widehat{G}$, add a segment of length $\widehat{d}_\infty(x, y)$ between x and y . Finally, identify the initial segments of $[x, y]$ and $[x, z]$ of length $(y|z)_x$ for any $x, y, z \in \widehat{G}$. By construction, the induced pseudo-metric d_ρ^x of the G -action on T with basepoint $x := \eta(1)$ is equal to d_∞ . \square

Later, we will need to scale the pseudo-metrics induced by sequences of actions of δ -hyperbolic G -spaces to ensure the existence of a converging subsequence. The scaling factor is the norm of the action as defined below. Note that the definition depends on the choice of the generating set S_G .

Definition 2.7. Let $X = (X, x, \rho)$ a based G -space. The norm of the action ρ with respect to the base point x (and the fixed generating set S_G), denoted by $|\rho|_x$, is defined as

$$|\rho|_x := \sum_{s \in S_G} d_X(x, \rho s x).$$

We may simply write $|\rho|$ if the basepoint x is clear from the context. Using this definition, we obtain the following corollary as a consequence of Lemma 2.6.

Corollary 2.8. Let $\delta \geq 0$ and for $i \in \mathbb{N}$, let $X_i = (X_i, x_i, \rho_i)$ be a based δ -hyperbolic G -space such that

$$\lim_{i \rightarrow \infty} |\rho_i|_{x_i} = \infty.$$

Then there exists a based real G -tree (T, x, ρ) such that the induced sequence $(\frac{1}{|\rho_i|} d_{\rho_i})$ of scaled pseudo-metrics on G has a subsequence converging in $\mathcal{A}(G)$ to d_ρ , and T is spanned by ${}_\rho Gx$.

Proof. For each $s \in S_G$ and $i \in \mathbb{N}$, we have $\frac{1}{|\rho_i|} d_{\rho_i}(x_i, \rho_i s x_i) \leq \frac{1}{|\rho_i|} \cdot |\rho_i| = 1$. Hence by Lemma 2.4, the sequence $(\frac{1}{|\rho_i|} d_{\rho_i})$ has a convergent subsequence. Moreover, the pseudo-metric $\frac{1}{|\rho_i|} d_{\rho_i}$ is $\frac{\delta}{|\rho_i|}$ -hyperbolic and

$$\lim_{i \rightarrow \infty} \frac{\delta}{|\rho_i|} = 0.$$

So the claim follows from Lemma 2.6. □

Note that the action of G on the limit tree T may not be minimal. However, Theorem 2.11 below shows that the minimality of the action can be guaranteed if the basepoints are chosen appropriately. Before we show this, we introduce the useful concept of approximating sequences.

Definition 2.9. Let $(X_i) = ((X_i, x_i, \rho_i))$ be a sequence of metric G -spaces. Assume that the sequence (d_{ρ_i}) converges to a pseudo-metric d_ρ induced by the based G -space $X = (X, x, \rho)$. For a point $t \in X$, an approximating sequence of t in (X_i) is a sequence (t_i) with $t_i \in X_i$ for each i such that

$$\lim_{i \rightarrow \infty} ({}_{\rho_i} g x_i \mid_{\rho_i} h x_i)_{t_i} = ({}_\rho g x \mid_\rho h x)_t \tag{2.3}$$

for any $g, h \in G$.

It is easy to see that every point ${}_\rho g x$ in the orbit of the basepoint x is approximated by the sequence $({}_{\rho_i} g x_i)$. In particular, the sequence (x_i) of basepoints approximates x . However, in general the limit space may contain points which do not admit an approximating sequence. But the following lemma implies that this cannot occur under the hypothesis of Lemma 2.6.

Lemma 2.10. *Let $(X_i) = (X_i, x_i, \rho_i)$ be a sequence of geodesic G -spaces, where each X_i is δ_i -hyperbolic and*

$$\lim_{i \rightarrow \infty} \delta_i = 0.$$

Assume that (d_{ρ_i}) converges to d_ρ where $T = (T, x, \rho)$ is a G -tree spanned by ${}_\rho Gx$. Then the following hold.

- (1) *Every $t \in T$ has an approximating sequence.*
- (2) *If (t_i) and (\bar{t}_i) are approximating sequences for some $t \in T$, then*

$$\lim_{i \rightarrow \infty} d_{X_i}(t_i, \bar{t}_i) = 0.$$

- (3) *If (t_i) is an approximating sequence for t then $({}_{\rho_i} g t_i)$ is an approximating sequence for ${}_\rho g t$.*
- (4) *If (t_i) and (y_i) are approximating sequences for t and y then*

$$\lim_{i \rightarrow \infty} d_{X_i}(t_i, y_i) = d_T(t, y).$$

Proof. Let $t \in T$. As T is spanned by ${}_\rho Gx$ there exist $g_1, g_2 \in G$ such that $t \in [{}_\rho g_1 x, {}_\rho g_2 x]$. Fix such g_1, g_2 . For each i choose $t_i \in [{}_{\rho_i} g_1 x_i, {}_{\rho_i} g_2 x_i]$ s.th.

$$\frac{d_{X_i}(t_i, {}_{\rho_i} g_1 x_i)}{d_{X_i}({}_{\rho_i} g_1 x_i, {}_{\rho_i} g_2 x_i)} = \frac{d_T(t, {}_\rho g_1 x)}{d_T({}_\rho g_1 x, {}_\rho g_2 x)}.$$

This choice clearly implies that (2.3) holds for $g \in \{g_1, g_2\}$. Now consider an arbitrary $g \in G$. To prove (1) we need to show that $\lim_{i \rightarrow \infty} d_{X_i}(t_i, {}_{\rho_i} g x_i) = d_T(t, {}_\rho g x)$.

Note that $({}_\rho g_1 x | {}_\rho g_2 x)_t = 0$ (cf. (2.2)). Since T is 0-hyperbolic, this implies w.l.o.g. (possibly after exchanging g_1 and g_2) that $({}_\rho g_2 x | {}_\rho g x)_t = 0$ as in Figure 2.2. This implies that $({}_\rho g_2 x | {}_\rho g x)_t = 0$.




FIGURE 2.2. An approximating sequence (t_i) of t .

Now choose $t'_i \in [\rho_i g_2 x_i, \rho_i g x_i]$ such that

$$d_{X_i}(\rho_i g_2 x_i, t'_i) = d_{X_i}(\rho_i g_2 x_i, t_i).$$

It is easily verified that $\lim_{i \rightarrow \infty} d_{X_i}(t_i, t'_i) = 0$. This implies that

$$\begin{aligned} \lim_{i \rightarrow \infty} d_{X_i}(t_i, \rho_i g x_i) &= \lim_{i \rightarrow \infty} d_{X_i}(t'_i, \rho_i g x_i) \\ &= \lim_{i \rightarrow \infty} (d_{X_i}(\rho_i g_2 x_i, \rho_i g x_i) - d_{X_i}(t'_i, \rho_i g_2 x_i)) \\ &= d_T(\rho g_2 x, \rho g x) - d_T(t, \rho g_2) \\ &= d_T(t, \rho g x). \end{aligned}$$

Thus (t_i) is an approximating sequence of t and (1) is established.

To prove (2) note first that it suffices to deal with the case where (t_i) is constructed as in the proof of (1), in particular $t_i \in [\rho_i g_1 x_i, \rho_i g_2 x_i]$ for all i and some fixed $g_1, g_2 \in G$. As (\bar{t}_i) is an approximating sequence for t it follows that

$$\lim_{i \rightarrow \infty} d_{X_i}(\rho_i g_k x_i, \bar{t}_i) = \lim_{i \rightarrow \infty} d_{X_i}(\rho_i g_k x_i, t_i)$$

for $k = 1, 2$. As X_i is δ_i -hyperbolic with $\lim_{i \rightarrow \infty} \delta_i = 0$ it follows easily that $\lim_{i \rightarrow \infty} d_{X_i}(\bar{t}_i, t_i) = 0$.

(3) is trivial and (4) follows from (2) and the fact that we can construct approximating sequences for t and y as in the proof of (1) by choosing g_1, g_2 such that both t and y lie on $[\rho g_1 x, \rho g_2 x]$. \square

We conclude the section with the following theorem, it guarantees that a limit action is minimal provided that base points are centrally located, meaning that the norm is smallest for the chosen base points.

Theorem 2.11. *Let (X_i, x_i, ρ_i) for $i \in \mathbb{N}$ and (T, x, ρ) be as in Lemma 2.6. If for any i and any $y \in X_i$,*

$$|\rho_i|_y \geq |\rho_i|_{x_i}, \tag{2.4}$$

then the limit action of G on T is minimal.

Proof. The proof is by contradiction.

Assume that $T' \subset T$ is a proper G -invariant subtree. Recall that T is spanned by the orbit of the base point x . This implies that $x \notin T'$ as otherwise the orbit ρGx and therefore T would be contained in T' . Let p_x be the nearest point projection of x to T' .

It follows that for any $g \in G$ we have $\rho g p_x = p_x$ or $[x, \rho g x] = [x, p_x] \cup [p_x, \rho g p_x] \cup [\rho g p_x, \rho g x]$. It follows that

$$d_T(x, \rho g x) \geq d_T(p_x, \rho g p_x)$$

for all $g \in G$. Moreover for some $s \in S_G$ we have ${}_{\rho}sx \neq x$ as otherwise ${}_{\rho}Gx = x$ and therefore $T = \{x\}$ as T is spanned by ${}_{\rho}Gx$. This implies in particular that for some $s \in S_G$ we have $d_T({}_{\rho}sx, x) > d_T({}_{\rho}sp_x, p_x)$, it follows that

$$|\rho|_{p_x} = \sum_{s \in S_G} d_T(p_x, {}_{\rho}sp_x) < \sum_{s \in S_G} d_T(x, {}_{\rho}sx) = |\rho|_x.$$

Let (p_x^i) be an approximating sequence for p_x . Then,

$$\lim_{i \rightarrow \infty} |\rho_i|_{p_x^i} = |\rho|_{p_x} < |\rho|_x = \lim_{i \rightarrow \infty} |\rho_i|_{x_i}.$$

Therefore, for large i , we get $|\rho_i|_{p_x^i} < |\rho_i|_{x_i}$, contradicting the assumption on the x_i . \square

We conclude this section by explaining how the previous statements imply that most Γ -limit groups act on real trees if Γ is a hyperbolic group.

Let Γ be a hyperbolic group with finite generating set S_Γ and $X := \text{Cay}(\Gamma, S_\Gamma)$. As before G is a group with finite generating set S_G . Let δ be the hyperbolicity constant of X . Recall that Γ is a subgroup of $\text{Isom}(X)$ if we consider the natural left action of Γ on X . Thus any $\varphi \in \text{Hom}(G, \Gamma)$ defines an action of G on X , namely

$$\varphi : G \rightarrow \Gamma \subset \text{Isom}(X).$$

It follows that $(X, 1, \varphi)$ is a based G -space for any $\varphi \in \text{Hom}(G, \Gamma)$. This action induces the δ -hyperbolic pseudo-metric $d_\varphi := d_\varphi^1$ on G given by

$$d_\varphi(g, h) = d_X(\varphi g \cdot 1, \varphi h \cdot 1) = d_X(\varphi(g), \varphi(h)) = |\varphi(g^{-1}h)|_{S_\Gamma}.$$

We can now define the norm of φ , cf. Definition 2.7:

Definition 2.12. We call $|\varphi| := |\varphi|_1 = \sum_{s \in S_G} d_X(1, \varphi(s))$ the norm of φ .

For any $g \in \Gamma$ let $c_g : \Gamma \rightarrow \Gamma$ be the inner automorphism given by $c_g(h) := g^{-1}hg$ for all $h \in \Gamma$. It is now clear that $d_{c_g \circ \varphi} = d_\varphi^g$ as

$$d_{c_g \circ \varphi}(h, k) = d_X(g^{-1}\varphi(h)g, g^{-1}\varphi(k)g) = d_X(\varphi(h)g, \varphi(k)g) = d_\varphi^g(h, k)$$

for all $h, k \in G$. It follows in particular that $|c_g \circ \varphi| = |\varphi|_g$.

We call a homomorphism $\varphi \in \text{Hom}(G, \Gamma)$ *conjugacy-short* if $|\varphi| \leq |\varphi|_g$ for all $g \in \Gamma$. We now consider a sequence of homomorphisms $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ such that the following hold:

- (1) φ_i is conjugacy-short for all $i \in \mathbb{N}$.
- (2) (φ_i) does not contain a subsequence (φ_{j_i}) such that $\ker \varphi_{j_i} = \ker \varphi_{j_{i'}}$ for all $i, i' \in \mathbb{N}$.

The second condition implies that we may assume that the φ_i are pairwise distinct after passing to a subsequence. It follows in particular that $\lim_{i \rightarrow \infty} |\varphi_i| = \infty$ as for any K there are only finitely many homomorphisms of norm at most K .

Let now $X_i = (X_i, d_{X_i})$ be the metric space obtained from X by scaling with the factor $\frac{1}{|\varphi_i|}$, thus $X_i = X$ (the underlying sets) and $d_{X_i} = \frac{1}{|\varphi_i|} d_X$. Clearly G also acts on X_i by isometries where the action on the underlying sets $X = X_i$ coincide, hence we obtain a based G -space $(X_i, 1, \rho_i)$ where $\rho_i : G \rightarrow \text{Isom}(X_i)$ is given by

$$\rho_i g x = \varphi_i(g)x \text{ for all } x \in X_i = X.$$

It is immediate that $d_{\rho_i} = \frac{1}{|\varphi_i|} d_{\varphi_i}$. Moreover d_{ρ_i} is δ_i -hyperbolic with $\delta_i := \frac{\delta}{|\varphi_i|}$ and $\lim_{i \rightarrow \infty} \delta_i = 0$.

The following theorem is a consequence of Theorem 2.6 and the above discussion. Here we call a G -space *non-trivial* if it does not have a global fixed point.

Theorem 2.13. *Let $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ be a sequence of conjugacy-short homomorphisms. Then one of the following holds.*

- (1) (φ_i) contains a constant subsequence.
- (2) A subsequence of $(\frac{1}{|\varphi_i|} d_{\varphi_i})$ converges to d_ρ^X for some non-trivial, minimal real based G -tree (T, x, ρ) .

Proof. It follows from Lemma 2.4 that $(d_{\rho_i}) = (\frac{1}{|\varphi_i|} d_{\varphi_i})$ has a convergent subsequence as

$$\left(\frac{1}{|\varphi_i|} d_{\varphi_i} \right) (1, \varphi_i(s')) = \frac{d_{\varphi_i}(1, \varphi_i(s'))}{\sum_{s \in S_G} d_{\varphi_i}(1, \varphi_i(s))} \leq 1$$

for all $s' \in S_G$. Assume that (φ_i) does not contain a constant subsequence. Then d_{ρ_i} is δ_i -hyperbolic with $\lim \delta_i = 0$, and it follows from Lemma 2.6 that the limit metric coincides with d_ρ for some real based G -tree (T, x, ρ) . As $|\varphi_i| \leq |\varphi_i|_g$ for all $i \in \mathbb{N}$ and $g \in \Gamma$ it further follows from Theorem 2.11 that this action is minimal. Thus we need to argue that T has more than one point. This however follows from the fact that

$$\sum_{s \in S_G} d_T(x, \rho s x) = \lim_{i \rightarrow \infty} \sum_{s \in S_G} \frac{1}{|\varphi_i|} d_{\varphi_i}(1, s) = \lim_{i \rightarrow \infty} \frac{1}{|\varphi_i|} \sum_{s \in S_G} d_{\varphi_i}(1, s) = 1. \quad \square$$

Remark 2.14. Let (φ_i) and (T, x, ρ) be as in the conclusion of Theorem 2.13(2). By Lemma 2.10 any $t \in T$ admits an approximating sequence, i.e. a sequence (t_i) with $t_i \in X_i = X$ such that for all $g \in G$ we have

$$\lim_{i \rightarrow \infty} d_{X_i}(\varphi_i(g), t_i) = \lim_{i \rightarrow \infty} \frac{1}{|\varphi_i|} d_X(\varphi_i(g), t_i) = d_T(\rho g x, t).$$

As $\lim_{i \rightarrow \infty} |\varphi_i| = \infty$ this clearly implies that for each $i \in \mathbb{N}$, t_i may be chosen to be a vertex in X .

We conclude this section with an important example that shows how distinct sequences of homomorphisms in $\text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ give rise to distinct actions on real trees. We consider $G = \mathbb{Z}^2$ with the generating set $S_G = \{(1, 0), (0, 1)\}$ and $\Gamma = \mathbb{Z}$ with the generating set $S_\Gamma = \{1\}$.

- (1) Consider $(\varphi_i) \subset \text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ where $\varphi_i((1, 0)) = 1$ and $\varphi_i((0, 1)) = i$. Then $|\varphi_i| = i + 1$ and therefore

$$\lim_{i \rightarrow \infty} \frac{1}{|\varphi_i|} d_{\varphi_i}(0, (x, y)) = \lim_{i \rightarrow \infty} \frac{x + iy}{i + 1} = y.$$

Thus $(\frac{1}{|\varphi_i|} d_{\varphi_i})$ converges to d_ρ^0 for the real \mathbb{Z}^2 -tree $(\mathbb{R}, 0, \rho)$ with $\rho(1, 0)r = r$ and $\rho(0, 1)r = r + 1$ for all $r \in \mathbb{R}$. It follows in particular that the subgroup $\langle (1, 0) \rangle$ acts trivially on \mathbb{R} .

- (2) Let $\lambda_1, \lambda_2 \in \mathbb{R}$ with $\lambda_1 + \lambda_2 = 1$ be linearly independent over \mathbb{Q} and choose sequences (a_i) and (b_i) of integers such that

$$\lim_{i \rightarrow \infty} \frac{a_i}{a_i + b_i} = \lambda_1 \text{ (and therefore } \lim_{i \rightarrow \infty} \frac{b_i}{a_i + b_i} = \lambda_2).$$

Consider $(\varphi'_i) \subset \text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ where $\varphi'_i((1, 0)) = a_i$ and $\varphi'_i((0, 1)) = b_i$. Clearly $|\varphi'_i| = a_i + b_i$ and therefore

$$\lim_{i \rightarrow \infty} \frac{1}{|\varphi'_i|} d_{\varphi'_i}(0, (x, y)) = \lim_{i \rightarrow \infty} \frac{xa_i + yb_i}{a_i + b_i} = x\lambda_1 + y\lambda_2.$$

Thus $(\frac{1}{|\varphi'_i|} d_{\varphi'_i})$ converges to $d_{\rho'}^0$ for the real \mathbb{Z}^2 -tree $(\mathbb{R}, 0, \rho')$ with $\rho'((1, 0))r = \lambda_1 + r$ and $\rho'((0, 1))r = \lambda_2 + r$ for all $r \in \mathbb{R}$. It follows in particular that any non-trivial element of \mathbb{Z}^2 act non-trivially on \mathbb{R} .

2.3. Stability of limit actions of limit groups

Throughout this section let G be a f.g. group and Γ be an infinite hyperbolic group, both equipped with word metrics relative to fixed finite generating sets S_G and S_Γ respectively.

Call a sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ *convergent* if the sequence $(\frac{d_{\varphi_i}}{|\varphi_i|})$ converges in $\mathcal{A}(G)$. Now let (φ_i) be a convergent sequence and assume that (φ_i) does not contain a constant subsequence. Then we obtain a G -tree (T, x, ρ) as in Corollary 2.8. We will simply say that the sequence (φ_i) converges to the limit G -tree T .

Moreover, if (φ_i) is stable with Γ -limit map φ , then the action ρ factors through φ , i.e. we obtain an action $\rho_\infty : L_\varphi \rightarrow \text{Isom}(T)$ s.th. $\rho = \rho_\infty \circ \varphi$. We call the space (T, x, ρ_∞)

the *limit L_φ -tree* of the sequence (φ_i) , resp. (T, x, ρ) the limit G -tree. (We may simply speak of the limit tree, or limit action, of (φ_i) if it is clear from the context whether we refer to the G -action or the L_φ -action.)

Now consider an arbitrary stable sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ with associated Γ -limit map φ . We may assume that φ_i is conjugacy-short for all $i \in \mathbb{N}$ as replacing φ_i with $c_g \circ \varphi_i$ for some $g \in \Gamma$ does not change the kernel and therefore preserves the stability of (φ_i) .

If (φ_i) contains a constant subsequence, i.e. contains a homomorphism φ' infinitely many times, then the stability of (φ_i) implies that $\ker \varphi' = \underline{\ker}(\varphi_i)$. It then follows that

$$L_\varphi = \text{Im}(\varphi) \cong \text{Im}(\varphi') \leq \Gamma$$

is isomorphic to a subgroup of Γ .

If such a sequence does not exist, then Theorem 2.13 implies that (φ_i) has a subsequence (φ_{j_i}) such that the sequence of pseudo-metrics $(\frac{1}{|\varphi_{j_i}|} d_{\varphi_{j_i}})$ converges to a pseudo-metric d_ρ^x for some minimal non-trivial G -tree (T, x, ρ) . Thus, (φ_{j_i}) is a converging subsequence. We will call a sequence with the above properties *strict*:

Definition 2.15. A stable sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ is called *strict* if the following are satisfied.

- (1) φ_i is conjugacy-short for all $i \in \mathbb{N}$.
- (2) (φ_i) does not have a constant subsequence.

Thus the above discussion implies that any strict sequence has a subsequence that converges to a limit G -tree T .

Remark. Note that the limit action does depend on the choice of the subsequence. Indeed, consider the sequences $(\varphi_i) \subset \text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ and $(\varphi'_i) \subset \text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ discussed at the end of the previous section. Clearly both sequences are stable and

$$\underline{\ker}(\varphi_i) = \underline{\ker}(\varphi'_i) = 1.$$

Let now $(\bar{\varphi}_i) \subset \text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ be the sequence with $\bar{\varphi}_{2i} = \varphi_i$ and $\bar{\varphi}_{2i+1} = \varphi'_i$ for all $i \in \mathbb{N}$. Again $(\bar{\varphi}_i)$ is stable with $\underline{\ker}(\bar{\varphi}_i) = 1$. While $(\bar{\varphi}_i)$ is not convergent the two subsequences (φ_i) and (φ'_i) are. However they converge to distinct actions of \mathbb{Z}^2 on \mathbb{R} .

Definition 2.16 ([5]). Let T be a G -tree. A non-degenerate subtree $S \subset T$ is called *stable* if for every non-degenerate subtree $S' \subset S$, $\text{stab}_G(S') = \text{stab}_G(S)$. Otherwise S is called *unstable*. The tree T is stable if every non-degenerate subtree of T contains a stable subtree.

The following theorem is the main result of this section. Recall that a *tripod* is a tree spanned by three points and a tripod is called *non-degenerate* if it is not spanned by two points.

If P is a class of groups then we say that a group G is finite-by- P if G contains a finite normal subgroup N such that G/N is in P .

Theorem 2.17. *Let $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ be a convergent strict sequence with induced Γ -limit map φ , and $L = L_\varphi$. Let T be the limit L -tree. Then the following hold for the action of L on T .*

- (1) *The stabilizer of any non-degenerate tripod is finite.*
- (2) *The stabilizer of any non-degenerate arc is finite-by-Abelian.*
- (3) *Every subgroup of L which leaves a line in T invariant and fixes its ends is finite-by-Abelian.*
- (4) *The stabilizer of any unstable arc is finite.*

Before we proceed with the proof of Theorem 2.17 we recall some useful facts about torsion subgroups of hyperbolic groups.

Proposition 2.18. *Let Γ be a hyperbolic group. Then the following hold.*

- (1) *There exists a constant $N = N(\Gamma)$ such that every torsion subgroup of Γ has at most N elements.*
- (2) *There exists a constant $L = L(\Gamma)$ such that for every subgroup $H \leq \Gamma$, one of the following holds.*
 - (a) *H is a finite group (of order at most $N(\Gamma)$).*
 - (b) *For any generating set S of H there exists a hyperbolic element $\gamma \in H$ such that $|\gamma|_S \leq L$, where $|\cdot|_S$ denotes the word length on H relative to S .*

Proof. Note first that torsion subgroups of hyperbolic groups are finite, see e.g. [18, Chapitre 8, Corollaire 36]. Thus (1) follows from the fact that for any hyperbolic group Γ there exists $N(\Gamma)$ such that any finite group is of order at most $N(\Gamma)$, see [7, 8].

(2) is essentially due to M. Koubi [28]. Proposition 3.2 of [28] implies that there exists a finite set $\bar{S} \subset \Gamma$ such that any set $S \subset \Gamma$ is either conjugate to a subset of \bar{S} or that there exists a word w in $S \cup S^{-1}$ of length at most 2 such that w represents a hyperbolic element of Γ . Now for each subset S of \bar{S} let $L(S) = 0$ if $\langle S \rangle$ is finite and let $L(S)$ be the length of a shortest word in $S \cup S^{-1}$ that represents a hyperbolic element otherwise.

Because of (1) such an element always exists. The conclusion now follows by putting $L(\Gamma) := \max(2, \max_{S \subset \bar{S}}(L(S)))$. \square

It turns out that the existence of a uniform bound on the order of finite subgroups of a hyperbolic group Γ ensures the same for Γ -limit groups.

Lemma 2.19. *Let Γ be hyperbolic and L be a Γ -limit group. Then the following hold.*

- (1) *Every torsion subgroup of L has at most $N(\Gamma)$ elements.*
- (2) *A subgroup $A \leq L$ is finite-by-Abelian iff all f.g. subgroups of A are finite-by-Abelian.*

Proof. Choose a stable sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ with induced Γ -limit map $\varphi : G \rightarrow L = L_\varphi$.

We prove (1) by contradiction. Thus assume there exists a torsion subgroup $E \leq L$ that contains $N(\Gamma) + 1$ pairwise distinct elements $g_0, \dots, g_{N(\Gamma)}$. For each $k = 0, \dots, N(\Gamma)$, pick $\tilde{g}_k \in G$ s.th. $\varphi(\tilde{g}_k) = g_k$.

The stability of (φ_i) implies that $\varphi_i(\tilde{g}_m) \neq \varphi_i(\tilde{g}_n)$ for large i and $0 \leq n \neq m \leq N(\Gamma)$. Thus Proposition 2.18 (1) implies that $\langle \varphi_i(\tilde{g}_0), \dots, \varphi_i(\tilde{g}_{N(\Gamma)}) \rangle$ is infinite for large i .

Proposition 2.18 (2) then implies that for large i there exists a word w_i in $\tilde{g}_0, \dots, \tilde{g}_{N(\Gamma)}$ of length at most $L(\Gamma)$ such that $\varphi_i(w_i)$ is of infinite order. Now there are only finitely many such words. Thus there exists a word w such that $w = w_i$ for infinitely many i . As E is assumed to be a torsion group it follows that $\varphi(w)^k = 1$ for some k , i.e. that $w^k \in \ker \varphi_i$ and therefore $w^k \in \ker \varphi_i$ for almost all i , a contradiction. Thus (1) is proven.

We now show (2). Clearly, if A is finite-by-Abelian, so are all f.g. subgroups. Thus we need to show that if the commutator subgroup of A is infinite, i.e. contains $N(\Gamma) + 1$ distinct elements $g_0, \dots, g_{N(\Gamma)}$, then the same is true for some f.g. subgroup of A . This however is obvious as any element of the commutator subgroup of some group is the product of finitely many commutators and therefore lies in the commutator subgroup of a finitely generated subgroup. \square

Let now Γ be a hyperbolic group and $X := \text{Cay}(\Gamma, S_\Gamma)$ its Caleygraph. For any hyperbolic element $\gamma \in \Gamma$ we denote its fixed points in ∂X by p_+^γ and p_-^γ and we denote by A_γ the union of all geodesics connecting p_+^γ and p_-^γ . We call A_γ the *axis* of γ . A_γ is a γ -invariant subset of X that is contained in the 2δ -neighborhood of any geodesic connecting p_+^γ and p_-^γ .

Lemma 2.20. *Let Γ be a hyperbolic group and $X := \text{Cay}(\Gamma, S_\Gamma)$. Then there exists a constant K such that for any hyperbolic element $\gamma \in \Gamma$ and $x \in X$ we have*

$$d_X(x, \gamma x) \geq 2d_X(x, A_\gamma) - K.$$

Proof. If the translation length of γ is greater than 10δ then the definition of hyperbolicity easily implies that this inequality holds for $K = 0$ as the path $[x, p_x] \cup [p_x, \gamma p_x] \cup [\gamma p_x, \gamma x]$ is a $(1, 4\delta)$ -quasigeodesic joining x and γx . For any fixed hyperbolic element γ the existence of such a constant K is also easily established. As there are only finitely many conjugacy classes of elements with translation length at most 10δ this guarantees the existence of a uniform K as above. \square

The following proposition is the main step in the proof of Theorem 2.17 (2) and (3).

Proposition 2.21. *Let Γ be hyperbolic with finite generating set S_Γ and $X := \text{Cay}(\Gamma, S_\Gamma)$. Then there exist constants $C, \varepsilon \in \mathbb{R}_{>0}$ such that for any set $Y \subset \Gamma$ the following holds.*

If there exist $x_1, x_2 \in X$ such that

- (1) $d(x_1, x_2) > C$,
- (2) $d(x_i, hx_i) < \varepsilon \cdot d(x_1, x_2)$ for all $h \in Y$ and $i \in \{1, 2\}$,

then $\langle Y \rangle$ is either finite or finite-by- \mathbb{Z} .

Proof. Let K be the constant from Lemma 2.20. We may assume that $K \geq \delta$ where δ is the hyperbolicity constant of X .

Let M be the number of elements of Γ of word length at most 20δ . We will show that the claim of the proposition holds for $C := 1000 \cdot K$ and $\varepsilon := \frac{1}{1000(M \cdot L(\Gamma) + 1)}$.

Let $Y \subset \Gamma$ and assume that $x_1, x_2 \in X$ satisfy the conditions of the proposition with C, ε as above. If $\langle Y \rangle$ is finite there is nothing to show. Thus we may assume that $\langle Y \rangle$ is infinite. By Proposition 2.18(2) there exists a word w_0 of length at most $L(\Gamma)$ in $Y \cup Y^{-1}$ that represents a hyperbolic element of Γ . As before we denote by $p_+^{w_0}$ and $p_-^{w_0}$ the fixed points of w_0 in ∂X . We will show that any $y \in Y$ also fixes $p_+^{w_0}$ and $p_-^{w_0}$.

The triangle inequality and the assumption of the proposition imply that for any $w \in \langle Y \rangle$ we have

$$d(x_i, wx_i) \leq |w|_Y \cdot \varepsilon \cdot d(x_1, x_2) = |w|_Y \cdot \frac{d(x_1, x_2)}{1000(M \cdot L(\Gamma) + 1)}$$

for $i = 1, 2$. It follows that for $j \in \{0, \dots, M\}$, $y \in Y \cup \{1\}$ and $i \in \{1, 2\}$ we have

$$d(x_i, (w_0^j y)_x) \leq (M \cdot L(\Gamma) + 1) \cdot \frac{d(x_1, x_2)}{1000(M \cdot L(\Gamma) + 1)} = \frac{d(x_1, x_2)}{1000}.$$

The case $j = 1$ and $y = 1$ gives

$$d(x_i, w_0 x_i) \leq \frac{d(x_1, x_2)}{1000}$$

for $i = 1, 2$. It follows from Lemma 2.20 that

$$\begin{aligned} d(x_i, A_{w_0}) &\leq \frac{1}{2}(d(x_i, w_0 x_i) + K) = \frac{1}{2} \left(d(x_i, w_0 x_i) + \frac{C}{1000} \right) \\ &\leq \frac{1}{2} \left(\frac{d(x_1, x_2)}{1000} + \frac{d(x_1, x_2)}{1000} \right) = \frac{d(x_1, x_2)}{1000} \end{aligned}$$

for $i = 1, 2$.




FIGURE 2.3. The midpoint z is moved by at most 20δ by $w_0 \cdot (w_0^j y) w_0^{-1} (w_0^j y)^{-1}$

It follows that for $i = 1, 2, j \in \{0, \dots, M\}$ and $y \in Y \cup \{1\}$ we have

$$\begin{aligned} d(x_i, w_0^j y A_{w_0}) &= d((w_0^j y)^{-1} x_i, A_{w_0}) \leq d((w_0^j y)^{-1} x_i, x_i) + d(x_i, A_{w_0}) \\ &= d(x_i, w_0^j y x_i) + d(x_i, A_{w_0}) \leq \frac{d(x_1, x_2)}{1000} + \frac{d(x_1, x_2)}{1000} = \frac{d(x_1, x_2)}{500}, \end{aligned}$$

thus x_i is contained in the $\frac{d(x_1, x_2)}{500}$ -neighborhoods of both A_{w_0} and $w_0^j y A_{w_0} = A_{(w_0^j y) w_0 (w_0^j y)^{-1}}$.

Let now $z \in VX = \Gamma$ be the midpoint of a geodesic segment $[x_1, x_2]$ between x_1 and x_2 . The 2δ -thinness of geodesic triangles imply that the 2δ -neighborhoods of both A_{w_0} and $w_0^j y A_{w_0}$ contain the geodesic segment $[y_1, y_2] \subset [x_1, x_2]$ that is centered at z and of length $\frac{99}{100} d(x_1, x_2)$. It now follows easily that

$$d(z, w_0 \cdot (w_0^j y) w_0^{-1} (w_0^j y)^{-1} \cdot z) \leq 20\delta.$$

As there are M points in the 20δ -neighborhood of z it follows that there are $j_1 < j_2 \in \{0, \dots, M\}$ such that

$$w_0 (w_0^{j_1} y) w_0^{-1} (w_0^{j_1} y)^{-1} \cdot z = w_0 (w_0^{j_2} y) w_0^{-1} (w_0^{j_2} y)^{-1} \cdot z$$

and therefore

$$w_0 (w_0^{j_1} y) w_0^{-1} (w_0^{j_1} y)^{-1} = w_0 (w_0^{j_2} y) w_0^{-1} (w_0^{j_2} y)^{-1}.$$

After conjugation with $w_0^{j_1}$ we get

$$[w_0, y] = w_0 y w_0^{-1} y^{-1} = w_0 w_0^{j_0} y w_0^{-1} y^{-1} w_0^{-j_0} = w_0^{j_0} [w_0, y] w_0^{j_0}$$

for $j_0 = j_2 - j_1 \in \{1, \dots, M\}$. As $[w_0, y]$ commutes with $w_0^{j_0}$ it follows that $[w_0, y]$ preserves the fixed points of $w_0^{j_0}$, i.e. preserves $p_+^{w_0}$ and $p_-^{w_0}$. As $[w_0, y] = w_0 \cdot y w_0 y^{-1}$ and as w_0 preserves $p_+^{w_0}$ and $p_-^{w_0}$ it follows that $y w_0 y^{-1}$ also preserves $p_+^{w_0}$ and $p_-^{w_0}$, thus y preserves $\{p_+^{w_0}, p_-^{w_0}\}$ as a set. It now easily follows from assumption (2) that y in fact preserves $p_+^{w_0}$ and $p_-^{w_0}$ pointwise. This implies that $\langle Y \rangle$ preserves $p_+^{w_0}$ and $p_-^{w_0}$. Thus $\langle Y \rangle$ is a 2-ended group that acts end-preservingly on itself, thus $\langle Y \rangle$ is finite-by- \mathbb{Z} by Lemma 4.1 of [43]. \square

We can now prove Theorem 2.17.

Proof of Theorem 2.17. We first prove (1). Let D be a non-degenerate tripod spanned by $x, y, z \in T$. By Lemma 2.19 it suffices to show that $H := \text{stab}_L(D)$ is a torsion group. Let $h \in H$ and pick $\tilde{h} \in G$ such that $\varphi(\tilde{h}) = h$. We need to show that $\varphi_i(\tilde{h})$ is of finite order for large i as this implies that $\tilde{h}^{N(\Gamma)!} \in \ker(\varphi_i)$ which implies that h is a torsion element.

We follow the argument from the proof of Lemma 4.1 in [35]. We argue by contradiction, thus after passing to a subsequence we may assume that $\varphi_i(\tilde{h})$ is hyperbolic for all $i \in \mathbb{N}$.

Let $(x_i), (y_i)$ and (z_i) be approximating sequences of x, y and z . By Remark 2.14 we may assume that the x_i, y_i and z_i are vertices of X , i.e. elements of Γ . By Lemma 2.20 there exists a constant K such that

$$d_X(x_i, A_{\varphi_i(\tilde{h})}) \leq \frac{1}{2}(d_X(x_i, \varphi_i(\tilde{h})x_i) + K)$$

for all $i \in \mathbb{N}$ and therefore

$$\begin{aligned} \lim_{i \rightarrow \infty} \frac{1}{|\varphi_i|} d_X(x_i, A_{\varphi_i(\tilde{h})}) &\leq \lim_{i \rightarrow \infty} \frac{1}{2|\varphi_i|} (d(x_i, \varphi_i(\tilde{h})x_i) + K) \\ &= \frac{1}{2} d_T(x, \varphi(\tilde{h})x) = \frac{1}{2} d_T(x, hx) = 0. \end{aligned}$$

It follows that $d_T(x, A_h) = \lim_{i \rightarrow \infty} \frac{1}{|\varphi_i|} d_X(x_i, A_{\varphi_i(\tilde{h})}) = 0$, i.e. x lies on A_h . The same argument shows that y and z lie on A_h contradicting the assumption that x, y and z span a non-degenerate tripod as A_h is a line.

To prove (2), assume that $H \leq G$ stabilizes a non-degenerate arc $[x, y] \subset T$. Let (x_i) and (y_i) be approximating sequences of x , respectively y . Clearly $\lim_{i \rightarrow \infty} d_X(x_i, y_i) = \infty$ and

$$\lim_{i \rightarrow \infty} \frac{d_X(x_i, \varphi_i(h)x_i)}{d_X(x_i, y_i)} = \lim_{i \rightarrow \infty} \frac{d_X(y_i, \varphi_i(h)y_i)}{d_X(x_i, y_i)} = 0$$




FIGURE 2.4. Tripod stabilizers are finite

for any $h \in H$. Let now U be a f.g. subgroup of H and let \tilde{S} be a finite generating set of U . The above discussion implies that the hypothesis of Proposition 2.21 is satisfied for $S := \varphi_i(\tilde{S})$ and i sufficiently large. Thus $\varphi_i(U) = \langle S \rangle$ is finite-by-Abelian, i.e. $\varphi_i([U, U]) = [\varphi_i(U), \varphi_i(U)]$ is of order at most $N(\Gamma)$ for i sufficiently large.

The stability of the sequence (φ_i) now implies that $\varphi([U, U])$ is a torsion group and therefore finite by Lemma 2.19(1), in particular $\varphi(U) \leq L$ is finite-by-Abelian. Thus $\varphi(H)$ is finite-by-Abelian by Lemma 2.19(2).

The proof of (3) is similar to that of (2). Assume that $H \leq G$ acts orientation-preservingly on a line $Y \subset T$ with ends x and y . Choose sequences $(x^k)_{k \in \mathbb{N}}$ and $(y^k)_{k \in \mathbb{N}}$ of points on Y that converge to x , respectively y .

Clearly $\lim_{k \rightarrow \infty} d_T(x^k, y^k) = \infty$ and therefore

$$\lim_{k \rightarrow \infty} \frac{d_T(x^k, \varphi(h)x^k)}{d_T(x^k, y^k)} = 0 = \lim_{k \rightarrow \infty} \frac{d_T(y^k, \varphi(h)y^k)}{d_T(x^k, y^k)}$$

for all $h \in H$ as $d_T(x^k, \varphi(h)x^k) = d_T(y^k, \varphi(h)y^k)$ is just the translation length of $\varphi(h)$ and therefore independent of k .

For each k choose approximating sequences $(x_i^k)_{i \in \mathbb{N}}$ of x^k and $(y_i^k)_{i \in \mathbb{N}}$ of y^k . Now fix $h \in H$ and $k \in \mathbb{N}$. It follows from the definition of approximating sequences that

$$\lim_{i \rightarrow \infty} \frac{d(x_i^k, \varphi_i(h)x_i^k)}{d(x_i^k, y_i^k)} = \frac{d_T(x^k, \varphi(h)x^k)}{d_T(x^k, y^k)}$$

and

$$\lim_{i \rightarrow \infty} \frac{d(y_i^k, \varphi_i(h)y_i^k)}{d(x_i^k, y_i^k)} = \frac{d_T(y^k, \varphi(h)y^k)}{d_T(x^k, y^k)}.$$

As the right hand sides of these equations tend to 0 as k tends to ∞ it follows that for some subsequence (φ_{m_i}) of (φ_i) we get

$$\lim_{i \rightarrow \infty} \frac{d(x_{m_i}^i, \varphi_{m_i}(h)x_{m_i}^i)}{d(x_{m_i}^i, y_{m_i}^i)} = 0 = \lim_{i \rightarrow \infty} \frac{d(y_{m_i}^i, \varphi_{m_i}(h)y_{m_i}^i)}{d(x_{m_i}^i, y_{m_i}^i)}.$$

As H is countable a diagonal argument shows that we can assume that this holds for all $h \in H$ after passing to a subsequence. Thus we can argue as in the proof of (2).

To prove (4), let $[y_1, y_2] \subsetneq [y_3, y_4]$ and

$$\gamma \in \text{stab}_L([y_1, y_2]) \setminus \text{stab}_L([y_3, y_4]).$$

As γ does not fix both y_3 and y_4 , we may assume $\gamma(y_3) \neq y_3$.



FIGURE 2.5. Unstable arcs have finite stabilizers

Note that for each $\bar{\gamma} \in \text{stab}_L([y_3, y_4])$ we have

$$\bar{\gamma}(\gamma(y_3)) = [\bar{\gamma}, \gamma](\gamma(\bar{\gamma}(y_3))) = [\bar{\gamma}, \gamma](\gamma(y_3)).$$

As the commutator subgroup of $\text{stab}_L([y_1, y_2])$ is finite by (2) it follows that $\{[\bar{\gamma}, \gamma] \bar{\gamma} \in \text{stab}_L([y_3, y_4])\}$ and therefore the $\text{stab}_L([y_3, y_4])$ -orbit of $\gamma(y_3)$ is finite. It follows that a finite index subgroup U of $\text{stab}_L([y_3, y_4])$ fixes $\gamma(y_3)$ and therefore also the tripod spanned by y_3, y_2 and $\gamma(y_3)$. By (1) the subgroup U is finite. Thus $\text{stab}_L([y_3, y_4])$ is finite. \square

2.4. Virtually Abelian subgroups of Γ -limit groups

Call a group *almost Abelian* if it contains a finite-by-Abelian subgroup of finite index. It is trivial that any virtually Abelian group is almost Abelian and the converse holds for finitely generated groups, see Theorem 8.40 of [9].

While we will later see that almost Abelian subgroups of finitely generated Γ -limit groups (with Γ hyperbolic) are finitely generated and therefore virtually Abelian, this is can also be verified directly exploiting the ideas from the previous section.

Note that subgroups of almost Abelian groups are almost Abelian and that almost Abelian subgroups of hyperbolic groups are 2-ended. Throughout this section Γ is a hyperbolic group. We study not necessarily finitely generated almost Abelian Γ -limit groups.

Lemma 2.22. *Let A be an infinite Γ -limit group. Then the following hold.*

- (1) *If A is almost Abelian, it is either finite-by-Abelian or contains a unique finite-by-Abelian subgroup A^+ of index 2.*
- (2) *If A is finite-by-Abelian with center $Z(A)$ then $|A : Z(A)| < \infty$.*
- (3) *A is virtually Abelian iff A is almost Abelian.*
- (4) *A is virtually Abelian iff all f.g. subgroups of A are virtually Abelian.*

Proof. Assume that $A = \langle a_0, a_1, \dots \rangle$. As A is infinite it contains an element of infinite order by Proposition 2.18, thus we may assume that a_0 is of infinite order. Choose a stable sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ with induced Γ -limit map $\varphi : G \rightarrow A = L_\varphi$. For each i pick a lift $\tilde{a}_i \in G$ such that $\varphi(\tilde{a}_i) = a_i$. For $k \in \mathbb{N}$ put $A_k = \langle a_0, \dots, a_k \rangle$ and $\tilde{A}_k := \langle \tilde{a}_0, \dots, \tilde{a}_k \rangle$, clearly $A = \bigcup_{k \in \mathbb{N}} A_k$. Note that all A_k and \tilde{A}_k are infinite as a_0 and therefore also \tilde{a}_0 is of infinite order.

(1). Let A be almost Abelian. Note first that all A_k are finitely generated and almost Abelian and therefore finitely presented. Choose relators $r_1, \dots, r_{m_k} \in F(a_0, \dots, a_k)$ such that

$$A_k = \langle a_0, \dots, a_k \mid r_1, \dots, r_{m_k} \rangle.$$

Let \tilde{r}_l be the word obtained from r_l by replacing occurrences of $a_j^{\pm 1}$ by $\tilde{a}_j^{\pm 1}$.

As $\varphi(\tilde{r}_l) = 1$ for all l it follows that $\tilde{r}_l \in \ker(\varphi_i)$ and therefore $\varphi_i(\tilde{r}_l) = 1$ for all l and large i . This implies that there exists m_k such that $A_{k,i} := \varphi_i(\tilde{A}_k)$ is a quotient of A_k for $i \geq m_k$ and therefore almost Abelian. We may further assume that for $i \geq m_k$ $\varphi_i(\tilde{a}_0)$ is of infinite order. For $i \geq m_k$ the group $A_{k,i}$ is therefore an infinite almost Abelian subgroup of some hyperbolic group and hence 2-ended.

For $i \geq m_k$ let $A_{k,i}^+$ be the subgroup of $A_{k,i}$ consisting of all elements that preserve the ends of $A_{k,i}$. Clearly, $|A_{k,i} : A_{k,i}^+| \leq 2$. Moreover put $\tilde{A}_{k,i}^+ := \varphi_i^{-1}(A_{k,i}^+) \cap \tilde{A}_k$, again it follows that $|\tilde{A}_k : \tilde{A}_{k,i}^+| \leq 2$.

As \tilde{A}_k is finitely generated it contains only finitely many subgroups of index 2. A diagonal argument therefore shows that after passing to a subsequence we can assume that for each k there exists a subgroup $\tilde{A}_k^+ \leq \tilde{A}_k$ and $n_k \geq m_k \in \mathbb{N}$ such that $\tilde{A}_k^+ = \tilde{A}_{k,i}^+$ for all $i \geq n_k$. As the images $\varphi_i(\tilde{A}_k^+)$ act orientation preservingly on an axis of Γ for $i \geq n_k$ it follows from Proposition 2.21 that $A_{k,i}^+ := \varphi_i(\tilde{A}_k^+)$ is finite-by- \mathbb{Z} for $i \geq n_k$. Let $A_k^+ = \varphi(\tilde{A}_k^+)$. It follows as in the proof of Theorem 2.17(2) that $[A_k^+, A_k^+]$ is finite, i.e. that also A_k^+ is finite-by-Abelian.

Clearly, A_k^+ is of index at most 2 in A_k . It is further easily verified that $A_k^+ \leq A_{k+1}^+$ as for large i we have $A_{k,i}^+ \leq A_{k+1,i}^+$ and therefore $\widetilde{A}_{k,i}^+ \leq \widetilde{A}_{k+1,i}^+$. It follows that $A^+ := \bigcup_{k \in \mathbb{N}} A_k^+$ is finite-by-Abelian by Lemma 2.19 and a subgroup of $A = \bigcup_{k \in \mathbb{N}} A_k$ of index at most 2.

It remains to show the uniqueness of A^+ if $|A : A^+| = 2$. Let $U \neq A^+$ be a finite-by-Abelian index 2 subgroup of A . Pick $k \in \mathbb{N}$ and put $\widetilde{U}_k := \widetilde{A}_k \cap \varphi^{-1}(U \cap A_k)$. Then \widetilde{U}_k is of index 2 in \widetilde{A}_k and distinct from \widetilde{A}_k^+ if k is sufficiently large. Therefore \widetilde{U}_k contains an element $g \in \widetilde{A}_k \setminus \widetilde{A}_k^+$. Thus for large i , $\varphi_i(g)$ swaps the ends of $A_{k,i}$. Thus $\varphi_i(\widetilde{U}_k)$ contains an infinite dihedral group and cannot be finite-by-Abelian. As this holds for all (large enough) i , it follows easily that U is not finite-by-Abelian, which is a contradiction. This proves (1).

(2). If A is finite-by-Abelian and all choices are as in the proof of (1), then $A_{k,i} = A_{k,i}^+$ is finite-by- \mathbb{Z} for each k and $i \geq n_k$. As the order of finite subgroups is bounded by $N(\Gamma)$ there are only finitely many isomorphism classes of subgroups that are met by the $A_{k,i}$. In particular there exists a constant M such that for all k and i the center $Z_{k,i}$ of $A_{k,i}$ is a subgroup of index at most M . Define $\widetilde{Z}_{k,i} := \widetilde{A}_k \cap \varphi_i^{-1}(Z_{k,i})$, for large i we have $|\widetilde{A}_k : \widetilde{Z}_{k,i}| \leq M$, as there are only finitely many subgroups of given (finite) index. Thus we can assume, after passing to a subsequence, that $\widetilde{Z}_{k,i} = \widetilde{Z}_{k,j}$ for all i, j , we denote this group by \widetilde{Z}_k . It follows that $\varphi(\widetilde{Z}_k)$ is central in A_k and of index at most M , in particular $|A_k : Z(A_k)| \leq M$ for all k (here $Z(A_k)$ is the center of A_k). After passing to a subsequence we can assume that $|A_k : Z(A_k)| = M_0$ for all k and some $M_0 \leq M$. Clearly $Z(A_{k+1}) \cap A_k \leq Z(A_k)$. Now as $|A_{k+1} : Z(A_{k+1})| = M_0$ it follows that

$$|A_{k+1} \cap A_k : Z(A_{k+1}) \cap A_k| = |A_k : Z(A_{k+1}) \cap A_k| \leq M_0$$

and therefore $Z(A_k) = Z(A_{k+1}) \cap A_k$, in particular $Z(A_k) \leq Z(A_{k+1})$. Thus $Z := \bigcup_{k \in \mathbb{N}} Z(A_k)$ is a central subgroup (in fact the center) of $A = \bigcup_{k \in \mathbb{N}} A_k$. Clearly $|A : Z| = M_0$.

(3). If A is almost Abelian then A is virtually Abelian by (1) and (2). Moreover if A is virtually Abelian then the proof of (1) goes through unchanged as images of virtually Abelian groups are virtually Abelian, it follows that A is almost Abelian.

(4). Clearly, if A is virtually Abelian, so is every f.g. subgroup. Conversely, assume that all finitely generated subgroups of A are virtually Abelian. This implies in particular that A_k is virtually Abelian for all k . If infinitely many A_k are finite-by-Abelian, then each f.g. subgroup is finite-by-Abelian as a subgroup of some A_k , and the claim follows from Lemma 2.19. So assume that (for large enough k) A_k is not finite-by-Abelian. By (1) A_k contains a unique finite-by-Abelian subgroup U_k of index 2. The uniqueness of U_k implies that $U_k \leq U_{k+1}$, as $U_{k+1} \cap A_k$ is a finite-by-Abelian subgroup of A_k of index

2 and therefore equal to U_k . It follows that $U = \bigcup k \in \mathbb{N} U_k$ is of index 2 in A , and finite-by-Abelian by Lemma 2.19. The assertion follows. \square

We are now able to establish the following properties of virtually Abelian subgroups which will be important later on.

Lemma 2.23. *Let L be a Γ -limit group. Then the following hold.*

- (1) *If L is finite-by-Abelian then the subgroup $E \leq L$ that is generated by the torsion elements of L is finite and therefore of order at most $N(\Gamma)$.*
- (2) *If $B \leq L$ and $A \leq L$ are virtually Abelian and $|A \cap B| = \infty$ then $\langle A, B \rangle$ is virtually Abelian.*

Proof. (1). Let $\{g_0, g_1, \dots\} \subset L$ be the set of torsion elements of L . Again, let $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ be a stable sequence with induced limit map φ s.th. $L = L_\varphi$. For each $k \in \mathbb{N}$, pick $\tilde{g}_k \in G$ satisfying $\varphi(\tilde{g}_k) = g_k$. Note that for each k , $\varphi_i(\tilde{g}_k)$ is of finite order for large enough i . Put $\tilde{E}_k := \langle \{\tilde{g}_0, \dots, \tilde{g}_k\} \rangle$. We show that each $E_k := \varphi(\tilde{E}_k)$ is finite, hence of order at most $N(\Gamma)$. As $E_k \leq E_{k+1}$ for each k , this clearly implies that $E = \bigcup_{k \in \mathbb{N}} E_k$ is finite.

Fix $k \in \mathbb{N}$. By Proposition 2.21, $\varphi_i(\tilde{E}_k)$ is finite or finite-by- \mathbb{Z} for sufficiently large i . Assume that $\varphi_i(\tilde{E}_k)$ is finite-by- \mathbb{Z} for large i . Hence it acts invariantly on an axis in Γ and this action is orientation-preserving. If the image is infinite, it is therefore isomorphic to an HNN-extension (with surjective boundary monomorphisms), i.e.

$$\varphi_i(E_k) \cong F_i *_{F_i} .$$

But this HNN-extension is not generated by torsion elements, which is a contradiction. Thus $\varphi_i(\tilde{E}_k)$ is finite for large i , hence of order at most $N(\Gamma)$. This implies that $\ker \varphi_i \cap \tilde{E}_k$ is of index at most $N(\Gamma)$ in \tilde{E}_k . As \tilde{E}_k is f.g., there are only finitely many such kernels. The stability of (φ_i) then implies that $\ker \varphi_i \cap \tilde{E}_k$ eventually stabilizes. It follows that $E_k = \varphi(\tilde{E}_k) = \varphi_i(\tilde{E}_k)$ (for large i) is finite, hence (1) is proven.

(2). Let $A = \langle a_0, a_1, \dots \rangle$ and $B = \langle b_0, b_1, \dots \rangle \leq L$ be virtually Abelian such that $A \cap B$ is infinite. As L does not contain infinite torsion subgroups (cf. Lemma 2.19) it follows that $A \cap B$ contains an element of infinite order, so we assume w.l.o.g. that $a_0 = b_0$ is of infinite order. Now for each k we choose $\tilde{a}_k, \tilde{b}_k \in G$ s.th. $\varphi(\tilde{a}_k) = a_k$ and $\varphi(\tilde{b}_k) = b_k$. Define $\tilde{A}_k := \langle \tilde{a}_0, \dots, \tilde{a}_k \rangle$ and $\tilde{B}_k := \langle \tilde{b}_0, \dots, \tilde{b}_k \rangle$. The same argument as in the proof of Lemma 2.22 shows for each k and sufficiently large i both $\varphi_i(\tilde{A}_k)$ and $\varphi_i(\tilde{B}_k)$ are virtually Abelian, hence 2-ended. Now for large i the element $\varphi_i(\tilde{a}_0) = \varphi_i(\tilde{b}_0)$ is of infinite order, which implies that $\varphi_i(\langle \tilde{A}_k, \tilde{B}_k \rangle)$ lies in the unique maximal 2-ended subgroup of Γ

containing $\varphi_i(\tilde{a}_0)$. Hence $\varphi_i(\langle \tilde{A}_k, \tilde{B}_k \rangle)$ is 2-ended for large enough i . It follows easily that $\langle A_k, B_k \rangle$ is virtually Abelian. Now $\langle A, B \rangle = \bigcup_{k \in \mathbb{N}} \langle A_k, B_k \rangle$, so the result follows from Lemma 2.22(4). \square

We get the following immediate consequences.

Corollary 2.24. *Let L be a Γ -limit group and $a \in L$ be an element of infinite order. Then*

$$A := \langle \{a' \in L \mid \langle a, a' \rangle \text{ is virtually Abelian} \} \rangle$$

is the unique maximal virtually Abelian subgroup of L containing a .

Proof. Let $\{a_0, a_1, \dots\}$ be the set of those elements that satisfy that $\langle a, a_i \rangle$ is virtually Abelian. Applying Lemma 2.23(2) repeatedly implies that for each k , $A_k := \langle a, a_0, \dots, a_k \rangle$ is virtually Abelian. Thus A is virtually Abelian by Lemma 2.22(4). The uniqueness and maximality of A are trivial. \square

Corollary 2.25. *Let A be a maximal virtually Abelian subgroup of a Γ -limit group L and $g \in L$. If $gAg^{-1} \cap A$ is infinite then $g \in A$.*

Proof. Suppose that $gAg^{-1} \cap A$ is infinite. It follows from Lemma 2.23 that $\langle A, gAg^{-1} \rangle$ is virtually Abelian and therefore equal to A as A is maximal. Choose an element $a \in A$ of infinite order. Then $\langle a, gag^{-1} \rangle \leq A$ is virtually Abelian. Pick lifts \tilde{g}, \tilde{a} of g, a in L . Then $\varphi_i(\langle \tilde{a}, \tilde{g}\tilde{a}\tilde{g}^{-1} \rangle)$ is virtually Abelian and therefore 2-ended for large i . Thus $\varphi_i(\tilde{g})$ preserves or exchanges the ends of $\langle \varphi_i(\tilde{a}) \rangle$. It follows that $\langle \varphi_i(\tilde{g}), \varphi_i(\tilde{a}) \rangle$ is 2-ended for large i , thus $\langle a, g \rangle$ is virtually Abelian. The statement follows now from Corollary 2.24. \square

3. The structure of groups acting on real trees

Bass–Serre theory clarifies the algebraic structure of groups acting on simplicial trees. The structure of groups acting on real trees is more complicated but still fairly well understood provided that the action satisfies certain properties. This theory is mainly based on ideas of Rips who in turn applied ideas from the Makanin–Razborov rewriting process. Rips (unpublished) described the structure of finitely presented groups acting freely on real trees, see [17] for an account of his ideas. This was then generalized to stable actions by Bestvina and Feighn [5]. Sela [38] then proved a version for finitely generated groups under stronger stability assumptions; the version we present is a generalization of Sela’s result due to Guirardel [22].

We first fix notations for graphs of groups and recall the notion of a graph of actions. We then briefly study 2-orbifolds and describe certain actions of 2-orbifold groups on real trees. We then formulate the structure theorem of [22] in those terms.

3.1. Graphs of groups and the Bass–Serre tree

In this section we fix the notations for basic Bass–Serre theory as we will need precise language later on. For details see Serre’s book [42] or [24] for more similar notation.

A graph A is understood to consist of a vertex set VA , a set of oriented edges EA , a fixed point free involution $^{-1} : EA \rightarrow EA$ and a map $\alpha : EA \rightarrow VA$ which assigns to each edge e its initial vertex $\alpha(e)$. Moreover, we will denote $\alpha(e^{-1})$ alternatively by $\omega(e)$ and call $\omega(e)$ the terminal vertex of e .

A graph of groups \mathbb{A} then consists of an underlying graph A and the following data.

- (1) For each $v \in VA$, a vertex group A_v .
- (2) For each $e \in EA$, an edge group $A_e = A_{e^{-1}}$.
- (3) For each $e \in EA$, an embedding $\alpha_e : A_e \rightarrow A_{\alpha(e)}$.

Again, the embedding $\alpha_{e^{-1}}$ will alternatively be denoted by ω_e . The maps α_e and ω_e are called the *boundary monomorphisms* of the edge e .

An \mathbb{A} -path from $v \in VA$ to $w \in VA$ of length k is a sequence

$$a_0, e_1, a_1, \dots, e_k, a_k$$

where e_1, \dots, e_k is an edge path in A from v to w , $a_0 \in A_v$ and $a_i \in A_{\omega(e_i)}$ for $i = 1, \dots, k$. For two \mathbb{A} -paths $p = a_0, e_1, \dots, e_k, a_k$ and $q = b_0, f_1, \dots, f_l, b_l$ satisfying that $\omega(e_k) = \alpha(f_1)$, we define a product pq by

$$pq := a_0, e_1, \dots, e_k, a_k b_0, f_1, \dots, f_l, b_l.$$

An equivalence relation \sim on the set of \mathbb{A} -paths is defined as the relation generated by the elementary equivalences $a, e, b \sim a\alpha_e(c), e, \omega_e(c^{-1})b$ and $a, e, \omega_e(c), e^{-1}, b \sim a\alpha_e(c)b$. We denote the \sim -equivalence class of an \mathbb{A} -path p by $[p]$. We call an \mathbb{A} -path *reduced* if it cannot be shortened by an elementary equivalence.

Given a base vertex $v_0 \in VA$, the fundamental group of \mathbb{A} with respect to v_0 , $\pi_1(\mathbb{A}, v_0)$, is the set of equivalence classes of \mathbb{A} -paths from v_0 to v_0 , with the multiplication given by $[p][q] := [pq]$.

Recall that the edge and vertex groups of \mathbb{A} embed into $\pi_1(\mathbb{A})$. These embeddings are unique up to conjugacy. If $G = \pi_1(\mathbb{A}, v_0)$ or $G \cong \pi_1(\mathbb{A}, v_0)$ then we call \mathbb{A} a splitting of G . We moreover say that G splits over a subgroup H if $G = \pi_1(\mathbb{A}, v_0)$ such that H is an edge group of \mathbb{A} .

If $p = a_0, e_1, \dots, e_k, a_k$ then we say that $p \approx q$ if $q \sim a_0, e_1, \dots, e_k, a_k a$ for some $a \in A_v$, this defines an equivalence relation on the set of \mathbb{A} -path starting at v_0 . We denote the \approx -equivalence class of p by $[[p]]$. The \approx -equivalence classes are precisely the vertices

of the Bass–Serre tree $(\widetilde{\mathbb{A}}, v_0)$. We will usually simply write $\widetilde{\mathbb{A}}$ rather than $(\widetilde{\mathbb{A}}, v_0)$. For any vertex \tilde{v} of $\widetilde{\mathbb{A}}$, we will denote the projection of \tilde{v} to VA by $\downarrow\tilde{v}$. Thus $\downarrow\tilde{v} = v$ if $\tilde{v} = \llbracket p \rrbracket$ where p is an \mathbb{A} -path from v_0 to v .

If we choose for each $e \in EA$ a set C_e of left coset representatives of $\alpha_e(A_e)$ in $A_{\alpha(e)}$, then each \mathbb{A} -path q is equivalent to a unique reduced \mathbb{A} -path $q' = a_0, e_1, \dots, e_k, a_k$ such that $a_{i-1} \in C_{e_i}$ for $1 \leq i \leq k$. We say that q' is in *normal form* (relative to the set $\{C_e \mid e \in EA\}$, which we usually don't mention explicitly).

Any vertex $\tilde{v} \in V\widetilde{\mathbb{A}}$ is represented by a unique reduced \mathbb{A} -path $p_{\tilde{v}} = a_0, e_1, a_1, \dots, a_{k-1}, e_k, 1$ which is in normal form. We call $p_{\tilde{v}}$ the *representing path* of \tilde{v} . Note that any normal form \mathbb{A} -path p representing \tilde{v} is of the form $p = p_{\tilde{v}}a$ for a unique $a \in A_{\downarrow\tilde{v}}$.

The edge set $E\widetilde{\mathbb{A}}$ is then the set of pairs $(\tilde{v}_1, \tilde{v}_2)$ of vertices satisfying

$$p_{\tilde{v}_1}a_1, e, 1 \sim p_{\tilde{v}_2}a_2 \quad (3.1)$$

for some $e \in EA$, $a_1 \in A_{\downarrow\tilde{v}_1}$ and $a_2 \in A_{\downarrow\tilde{v}_2}$.

Note that if $(\tilde{v}_1, \tilde{v}_2) \in E\widetilde{\mathbb{A}}$, then also $(\tilde{v}_2, \tilde{v}_1) \in E\widetilde{\mathbb{A}}$ as (3.1) is equivalent to

$$p_{\tilde{v}_1}a_1 \sim p_{\tilde{v}_1}a_1, e, 1, e^{-1}, 1 \sim p_{\tilde{v}_2}a_2, e^{-1}, 1.$$

Thus the map

$$^{-1} : E\widetilde{\mathbb{A}} \rightarrow E\widetilde{\mathbb{A}}, \quad (\tilde{v}_1, \tilde{v}_2) \mapsto (\tilde{v}_2, \tilde{v}_1)$$

is an involution on $E\widetilde{\mathbb{A}}$, which is fixed point free as $\tilde{v}_1 \neq \tilde{v}_2$ if $(\tilde{v}_1, \tilde{v}_2) \in E\widetilde{\mathbb{A}}$. For any $\tilde{e} = (\tilde{v}_1, \tilde{v}_2) \in E\widetilde{\mathbb{A}}$ we put $\alpha(\tilde{e}) = \omega(\tilde{e}^{-1}) = \tilde{v}_1$. Moreover, for $\tilde{e} = (\tilde{v}_1, \tilde{v}_2)$ as above, we denote the edge $e \in EA$ (cf. (3.1)), by $\downarrow\tilde{e}$.

With the above notations, we obtain a natural action of $\pi_1(\mathbb{A}, v_0)$ on $\widetilde{\mathbb{A}}$ in the following way. For $[q] \in \pi_1(\mathbb{A}, v_0)$ put

$$\begin{aligned} \llbracket p \rrbracket &:= \llbracket qp \rrbracket && \text{for } \llbracket p \rrbracket \in V\widetilde{\mathbb{A}} \\ [q](\llbracket p_1 \rrbracket, \llbracket p_2 \rrbracket) &:= (\llbracket qp_1 \rrbracket, \llbracket qp_2 \rrbracket) && \text{for } (\llbracket p_1 \rrbracket, \llbracket p_2 \rrbracket) \in E\widetilde{\mathbb{A}}. \end{aligned}$$

With this G -action on $\widetilde{\mathbb{A}}$, for every $\tilde{v} \in V\widetilde{\mathbb{A}}$ the map

$$\theta_{\tilde{v}} : A_{\downarrow\tilde{v}} \rightarrow \text{stab}_{\widetilde{\mathbb{A}}} \tilde{v}, \quad h \mapsto [p_{\tilde{v}} h p_{\tilde{v}}^{-1}]$$

is an isomorphism between the vertex group $A_{\downarrow\tilde{v}}$ and the stabilizer of the vertex \tilde{v} in $\widetilde{\mathbb{A}}$. Likewise, for an edge $\tilde{e} = (\tilde{v}_1, \tilde{v}_2) \in E\widetilde{\mathbb{A}}$, the map

$$\theta_{\tilde{e}} := \theta_{\tilde{v}_1} \circ c_{a_1} \circ \alpha_e : A_{\downarrow\tilde{e}} \rightarrow \text{stab}_{\widetilde{\mathbb{A}}} \tilde{e}, \quad (3.2)$$

where a_1 is as in (3.1) and c_{a_1} denotes conjugation by a_1 , is an isomorphism between the edge group $A_{\downarrow\tilde{e}}$ and the stabilizer of the edge \tilde{e} . It follows easily from (3.1) that $\theta_{\tilde{e}} = \theta_{\tilde{e}^{-1}}$.

Sometimes we will need to refine a given splitting of a group, i.e. increase the complexity of a graph of groups decomposition by splitting some vertex group in a way that is compatible with the existing splitting. The following is obvious.

Definition & Lemma 3.1. *Let \mathbb{A} be a graph of groups, $v \in VA$, and \mathbb{A}^v a graph of groups such that $A_v = \pi_1(\mathbb{A}^v, v_0^v)$ for some $v_0^v \in A^v$. Suppose that for each edge $e \in EA$ with $\alpha(e) = v$, $\alpha_e(A_e)$ is conjugate into a vertex group $A_{w_e}^v$ for some vertex $w_e \in VA^v$.*

Then the graph of groups \mathbb{A}' defined below is called the refinement of \mathbb{A} by \mathbb{A}^v . The underlying graph A' has vertex set $VA' = (VA \setminus \{v\}) \cup VA^v$ and edge set $EA' = EA \cup EA^v$. Moreover for each edge $e \in EA'$ the attaching map α' and boundary monomorphism α'_e are as follows.

- (1) *If $e \in EA$ and $\alpha(e) \neq v$ then $\alpha'(e) = \alpha(e)$ and $\alpha'_e = \alpha_e$.*
- (2) *If $e \in EA^v$ then $\alpha'(e) = \alpha^v(e)$ and $\alpha'_e = \alpha_e^v$.*
- (3) *If $e \in EA$ and $\alpha(e) = v$ then $\alpha'(e) = w_e$ and $\alpha'_e : A_e \rightarrow A_{w_e}^v$ is such that $i_{w_e} \circ \alpha'_e$ is in A_v conjugate to α_e where i_{w_e} is the (up to conjugacy) unique inclusion of A_{w_e} in A_v .*

If \mathbb{A}' is a refinement of \mathbb{A} then $\pi_1(\mathbb{A}') \cong \pi_1(\mathbb{A})$. The operation inverse to a refinement is called a collapse.

3.2. Graphs of actions

In this section we recall the notion of a graph of actions, see [29]. This is a way of decomposing an action of a group on a real tree into pieces. In the structure theorem these pieces will be of very simple types.

Definition 3.2. A graph of actions is a tuple

$$\mathcal{G} = \mathcal{G}(\mathbb{A}) = (\mathbb{A}, (T_v)_{v \in VA}, (p_e^\alpha)_{e \in EA}, l)$$

where

- \mathbb{A} is a graph of groups,
- for each $v \in VA$, $T_v = (T_v, d_v)$ is a real A_v -tree,
- for each $e \in EA$, $p_e^\alpha \in T_{\alpha(e)}$ is a point fixed by $\alpha_e(A_e)$,
- $l : EA \rightarrow \mathbb{R}_{\geq 0}$ is a function satisfying $l(e) = l(e^{-1})$ for all $e \in EA$.

We call $l(e)$ the *length of e* . If $l = 0$ then we omit l , i.e. we write $\mathcal{G} = \mathcal{G}(\mathbb{A}) = (\mathbb{A}, (T_v)_{v \in VA}, (p_e^\alpha)_{e \in EA})$.

The points p_e^α are called *attaching points*. In the following we will denote $p_{e^{-1}}^\alpha$ alternatively by p_e^ω . Note that $p_e^\omega \in T_{\omega(e)}$ and that p_e^ω is fixed by $\omega_e(A_e)$.

Associated to any graph of actions \mathcal{G} is a real tree $T_{\mathcal{G}}$ obtained by replacing the vertices of the Bass–Serre tree (\mathbb{A}, v_0) by copies of the trees T_v and any lift $\tilde{e} \in E\tilde{\mathbb{A}}$ of $e \in EA$ by a segment of length $l(e)$. $T_{\mathcal{G}}$ comes with a natural $\pi_1(\mathbb{A}, v_0)$ -action. In the remainder of this section we will give a detailed description of the construction of $T_{\mathcal{G}}$ and its natural metric.

Let \mathcal{G} be a graph of actions as above. Choose a base vertex $v_0 \in VA$ and sets of left coset representatives C_e of $\alpha_e(A_e)$ in $A_{\alpha(e)}$ for all $e \in EA$.

For any $\tilde{v} \in V\tilde{\mathbb{A}}$, define $T_{\tilde{v}} := T_{\downarrow\tilde{v}} \times \{\tilde{v}\}$ to be a copy of $T_{\downarrow\tilde{v}}$ with the induced metric $d_{\tilde{v}}$ given by $d_{\tilde{v}}((x_1, \tilde{v}), (x_2, \tilde{v})) := d_{\downarrow\tilde{v}}(x_1, x_2)$. We further put

$$T_{\mathcal{G}}^V := \bigcup_{\tilde{v} \in V\tilde{\mathbb{A}}} T_{\tilde{v}}.$$

For any edge $\tilde{e} = (\tilde{v}_1, \tilde{v}_2) \in E\tilde{\mathbb{A}}$ and a_1, a_2 and e as in (3.1), we then define

$$p_{\tilde{e}}^\alpha := (a_1 p_e^\alpha, \tilde{v}_1) \in T_{\tilde{v}_1}. \quad (3.3)$$

Again, for an edge \tilde{e} we denote $p_{\tilde{e}^{-1}}^\alpha$ alternatively by $p_{\tilde{e}}^\omega$. With notations as in (3.1) we get $p_{\tilde{e}}^\omega = p_{\tilde{e}^{-1}}^\alpha = (a_2 p_{e^{-1}}^\alpha, \tilde{v}_2)$. We call $p_{\tilde{e}}^\alpha$ and $p_{\tilde{e}}^\omega$ the *attaching points* of the edge \tilde{e} .

For any $(x, \tilde{v}) \in T_{\mathcal{G}}^V$ and $g = [q] \in \pi_1(\mathbb{A}, v_0)$, put

$$g(x, \tilde{v}) := (ax, g\tilde{v})$$

if $qp_{\tilde{v}} \sim p_{g\tilde{v}}a$. It follows from the above definitions that this defines an action of $\pi_1(\mathbb{A}, v_0)$ on $T_{\mathcal{G}}^V$ with the following properties:

$$(1) \quad d_{\tilde{v}}(x, y) = d_{g\tilde{v}}(gx, gy) \text{ for all } x, y \in T_{\tilde{v}} \text{ and } g \in \pi_1(\mathbb{A}, v_0).$$

$$(2) \quad gp_{\tilde{e}}^\alpha = p_{g\tilde{e}}^\alpha \text{ for all } \tilde{e} \in E\tilde{\mathbb{A}} \text{ and } g \in \pi_1(\mathbb{A}, v_0).$$

For any $\tilde{e} \in E\tilde{\mathbb{A}}$ define $T_{\tilde{e}} := [0, l(\downarrow\tilde{e})] \times \{\tilde{e}\}$ to be a copy of the real interval $[0, l(\downarrow\tilde{e})]$. Let $d_{\tilde{e}}$ be the standard metric on $T_{\tilde{e}}$, i.e. $d_{\tilde{e}}((x, \tilde{e}), (y, \tilde{e})) = |x - y|$ for all $(x, \tilde{e}), (y, \tilde{e}) \in T_{\tilde{e}}$. We define

$$T_{\mathcal{G}}^E := \bigcup_{\tilde{e} \in E\tilde{\mathbb{A}}} T_{\tilde{e}}.$$

Note that $T_{\tilde{e}}$ consists of a single point if $l(\downarrow\tilde{e}) = 0$. Now for any $(x, \tilde{e}) \in T_{\mathcal{G}}^E$ and $g \in \pi_1(\mathbb{A}, v_0)$, put

$$g(x, \tilde{e}) := (x, g\tilde{e}).$$

This clearly defines an action of $\pi_1(\mathbb{A}, v_0)$ on $T_{\mathcal{G}}^E$, which satisfies $d_{\tilde{e}}(x, y) = d_{g\tilde{e}}(gx, gy)$ for all $x, y \in T_{\tilde{e}}$ and $g \in \pi_1(\mathbb{A}, v_0)$.

We can now define the tree $T_{\mathcal{G}}$. We put

$$T_{\mathcal{G}} := (T_{\mathcal{G}}^V \cup T_{\mathcal{G}}^E) / \sim$$

where \sim is the equivalence relation generated by

- (1) $p_{\tilde{e}}^{\alpha} \sim (0, \tilde{e})$ for any $\tilde{e} \in E\tilde{\mathbb{A}}$,
- (2) $(k, \tilde{e}) \sim (l(\downarrow\tilde{e}) - k, \tilde{e}^{-1})$ for each $\tilde{e} \in E\tilde{\mathbb{A}}$ and $k \in [0, l(\downarrow\tilde{e})]$.

The definition of $T_{\mathcal{G}}$ ensure that the images of $T_{\alpha(\tilde{e})}$ and $T_{\omega(\tilde{e})}$ in $T_{\mathcal{G}}$ are joint by a segment of length $l(\downarrow\tilde{e})$ for any $\tilde{e} \in E\tilde{\mathbb{A}}$, see Figure 3.1. It also takes care of the fact that in $\tilde{\mathbb{A}}$ each geometric edge occurs with both orientations.




FIGURE 3.1. $T_{\alpha(\tilde{e})}$ and $T_{\omega(\tilde{e})}$ are joint by a segment of length $l(\downarrow\tilde{e})$.

It is clear from the above observations that this equivalence relation is preserved by the $\pi_1(\mathbb{A}, v_0)$ -action on $T_{\mathcal{G}}^V \cup T_{\mathcal{G}}^E$, thus it induces a $\pi_1(\mathbb{A}, v_0)$ -action on $T_{\mathcal{G}}$.

There exists a unique $\pi_1(\mathbb{A}, v_0)$ -invariant path metric $d_{\mathcal{G}}$ on $T_{\mathcal{G}}$ such that $d_{\mathcal{G}}(y_1, y_2) = d_{\tilde{v}}(y_1, y_2)$ if $y_1, y_2 \in T_{\tilde{v}}$, $d_{\mathcal{G}}(y_1, y_2) = d_{\tilde{e}}(y_1, y_2)$ if $y_1, y_2 \in T_{\tilde{e}}$. If $y_1, y_2 \in T_{\mathcal{G}}$ do not lie in the same edge or vertex space then $d_{\mathcal{G}}(y_1, y_2)$ is computed as follows:

- (1) If $y_1 = (x_1, \tilde{v}_1), y_2 = (x_2, \tilde{v}_2) \in T_{\mathcal{G}}^V$ and $\tilde{e}_1, \dots, \tilde{e}_k$ is a reduced path in $\tilde{\mathbb{A}}$ from \tilde{v}_1 to \tilde{v}_2 then

$$d_{\mathcal{G}}(y_1, y_2) = d_{\tilde{v}_1}(y_1, p_{\tilde{e}_1}^{\alpha}) + \sum_{i=1}^k l(\downarrow\tilde{e}_i) + \sum_{i=1}^{k-1} d_{\omega(\tilde{e}_i)}(p_{\tilde{e}_i}^{\omega}, p_{\tilde{e}_{i+1}}^{\alpha}) + d_{\tilde{v}_2}(p_{\tilde{e}_k}^{\omega}, y_2)$$

- (2) If $y_1 = (x_1, \tilde{e}_1) \in T_{\mathcal{G}}^E$, $y_2 = (x_2, \tilde{v}_2) \in T_{\mathcal{G}}^V$ and $\tilde{e}_1, \dots, \tilde{e}_k$ is a reduced path in $\tilde{\mathbb{A}}$ with $\omega(\tilde{e}_k) = \tilde{v}_2$ then

$$d_{\mathcal{G}}(y_1, y_2) = (l(\downarrow \tilde{e}_1) - x_1) + \sum_{i=2}^k l(\downarrow \tilde{e}_i) + \sum_{i=1}^{k-1} d_{\omega(\tilde{e}_i)}(p_{\tilde{e}_i}^{\omega}, p_{\tilde{e}_{i+1}}^{\alpha}) + d_{\tilde{v}_2}(p_{\tilde{e}_k}^{\omega}, y_2)$$

- (3) If $y_1 = (x_1, \tilde{e}_1), y_2 = (x_2, \tilde{e}_k) \in T_{\mathcal{G}}^E$ and $\tilde{e}_1, \dots, \tilde{e}_k$ is a reduced path in $\tilde{\mathbb{A}}$ then

$$d_{\mathcal{G}}(y_1, y_2) = (l(\tilde{e}_1) - x_1) + \sum_{i=2}^{k-1} l(\downarrow \tilde{e}_i) + \sum_{i=1}^{k-1} d_{\omega(\tilde{e}_i)}(p_{\tilde{e}_i}^{\omega}, p_{\tilde{e}_{i+1}}^{\alpha}) + x_2$$

Recall that the restriction of $d_{\mathcal{G}}$ to any vertex tree $T_{\tilde{v}}$, resp. edge segment $T_{\tilde{e}}$, equals $d_{\tilde{v}}$, resp. $d_{\tilde{e}}$. It therefore follows from (3.3) that the distance of two points in $T_{\mathcal{G}}$ can be computed entirely in terms of the metrics d_v of the vertex trees of \mathcal{G} and its length function l . The case we are mostly interested in is case 1 above, i.e. the case where y_1 and y_2 are contained in vertex trees $T_{\tilde{v}_1}$ and $T_{\tilde{v}_2}$. If $p = a_0, e_1, a_1, \dots, a_k$ is a reduced \mathbb{A} -path equivalent to $p_{\tilde{v}_1}^{-1} p_{\tilde{v}_2}$, then $d_{\mathcal{G}}(y_1, y_2)$ can be computed as

$$d_{\mathcal{G}}(y_1, y_2) = d_{\alpha(e_1)}(x_1, a_0 p_{e_1}^{\alpha}) + \sum_{i=1}^k l(e_i) + \sum_{i=1}^{k-1} d_{\omega(e_i)}(p_{e_i}^{\omega}, a_i p_{e_{i+1}}^{\alpha}) + d_{\omega(e_k)}(p_{e_k}^{\omega}, a_k x_2). \quad (3.4)$$

We say that a G -tree T splits as a graph of actions $\mathcal{G}(\mathbb{A})$ if $G \cong \pi_1(\mathbb{A})$ and there is a G -equivariant isometry from T to $T_{\mathcal{G}}$.

Remark 3.3. Let \mathcal{G} be a graph of actions, $e \in EA$ and $g \in A_{\alpha(e)}$. Assume that \mathcal{G}' is the graph of actions obtained from \mathcal{G} by replacing the attaching point $p_e^{\alpha} \in T_{\alpha(e)}$ by $g p_e^{\alpha}$ and the embedding $\alpha_e : A_e \rightarrow A_{\alpha(e)}$ by $i_g \circ \alpha_e$. Then $T_{\mathcal{G}}$ also splits as the graph of actions \mathcal{G}' .

It follows from the remark that in a graph of actions splitting of a tree T we are free to alter the attaching points within their orbits of the vertex actions. In particular, if a vertex group A_v acts with dense orbits on T_v , the attaching points in A_v can be chosen to be arbitrarily close to each other. This will turn out useful in Section 5.

3.3. 2-orbifolds

We will only be interested in 2-orbifolds that are non-spherical. Thus for us a 2-orbifold \mathcal{O} is a quotient of \mathbb{R}^2 by a group Γ of homeomorphisms that acts properly discontinuously. In particular all surfaces different from the 2-sphere and the projective plane occur this way. Our discussion of 2-orbifolds is very brief and informal; for a more detailed discussion

the reader is referred to the article of P. Scott [37] or the book of Zieschang, Vogt and Coldewey [46]. Note that we can assume that \mathbb{R}^2 is either endowed with the standard Euclidean metric or with a metric of constant curvature -1 making it isometric to the hyperbolic plane, and that Γ acts by isometries.

A 2-orbifold $\mathcal{O} = \mathbb{R}^2/\Gamma$ is topologically a surface but points whose lifts in \mathbb{R}^2 have non-trivial stabilizers are endowed with a marking that reflects the structure of the stabilizer. The stabilizers are in $\text{Homeo}(\mathbb{R}^2)$ conjugate to a finite group of Euclidean isometries and are therefore of one of the following types:

- (1) \mathbb{Z}_p and generated by a rotation of order p .
- (2) A dihedral group D_{2p} and generated by two reflections whose product is a rotation of order p with $p \in \mathbb{N}$.
- (3) A group of order two being generated by a reflection.

In the quotient, points of the first type are depicted as points labeled by p and are referred to as cone points, points of the last type are depicted as points of fat lines and are referred to as points of the reflection boundary and points lifting to points with dihedral stabilisers occur as points incident to two reflection lines and are labeled by p if the dihedral group was of type D_{2p} . Thus 2-orbifolds can locally be depicted as in Figure 3.2. We further say that a 2-orbifold is of *cone type* if it has no reflection boundary.




FIGURE 3.2. The local structure of a 2-orbifold

There is a well-known definition for the fundamental group of a 2-orbifold which generalizes the case of a surface, see [37]. In the case of an orbifold \mathcal{O} of cone type with cone points of orders p_1, \dots, p_k this is easily done as follows. First let $S_{\mathcal{O}}$ be the surface (with boundary) obtained by replacing each cone point with a boundary component. Then define $\pi_1(\mathcal{O})$ to be the quotient of $\pi_1(S_{\mathcal{O}})$ obtained by adding the relations $\gamma_i^{p_i}$ where γ_i is the homotopy class represented by the boundary component corresponding to the i th cone point. This fundamental group recovers the group Γ just as the fundamental group of

a manifold recovers the group of Decktransformations on the universal covering, in fact a completely analogous theory of coverings can be developed for 2-orbifolds, see [37].

We do not consider points on the reflection boundary to lie in the boundary of a 2-orbifold. Thus 2-orbifolds as we defined them are open and have no boundary. If however Γ is finitely generated then we can, as in the surface case, find a compact suborbifold of $\mathcal{O} = \mathbb{R}^2/\Gamma$ with boundary whose complement consist of ends of type $(0, 1) \times S^1$ or quadrilaterals $(0, 1) \times [0, 1]$ where $(0, 1) \times \{0\}$ and $(0, 1) \times \{1\}$ belong to the reflection boundary. This compact suborbifold is a natural compactification of the original orbifold and we will ignore the difference between the orbifold and this compact subspace. Thus orbifolds have boundary components which are either loops or segments joining points on the reflection boundary. Figure 3.3 is an example of a 2-orbifold with two boundary components, one of each type.




FIGURE 3.3. A 2-orbifold with a dihedral and a cyclic boundary component

The subgroup of Γ corresponding to the boundary components are either infinite cyclic groups or infinite dihedral groups if the boundary component is a circle or line, respectively. In Figure 3.3 they are represented by dotted lines, in particular the points on the reflection boundary are not part of the boundary. We call the subgroups of an orbifold group corresponding to a boundary component *parabolic* or *peripheral*.

Given a simple closed curve on a 2-orbifold that does not bound a disk with at most one cone point, cutting the orbifold along this line corresponds to decomposing the orbifold group into an HNN-extension or amalgamated product over the infinite cyclic group. We call a simple closed curve of this type *essential*.




FIGURE 3.4. Curves representing HNN-extensions over \mathbb{Z} and D_∞

Similarly cutting along a segment joining two points on the reflection boundary that does not bound a disk without cone points in the interior and at most one exceptional point on the reflection boundary corresponds to a splitting along an infinite dihedral group. Again we call such segments *essential*.




FIGURE 3.5. Curves representing amalgamated products over \mathbb{Z} and D_∞

Note that the examples of essential curves and segments depicted in Figure 3.4 and Figure 3.5 are orbifolds without boundary, however a similar statement holds for orbifolds with boundary: Given an essential simple closed curve or an essential segment joining two points on the reflection boundary that does not meet the boundary of the orbifold, this curve induces a splitting of the orbifold group along an infinite cyclic or infinite dihedral group such that all peripheral subgroups are elliptic in this splitting.

In the remainder of this section we will describe an important class of actions of fundamental groups of cone type orbifold groups on real trees. We start with actions of fundamental groups of surfaces with boundary. While these groups are algebraically

just free groups the action is constructed using the surface geometry. The orbifold group actions are then constructed as quotients of those actions.

Surfaces with boundary can be described by unions of bands that have a foliation with a transverse measure. We briefly describe their structure, see [5] for a detailed discussion. Suppose that the following data is given.

- (1) Γ is the union of disjoint closed intervals $I_1, \dots, I_k \subset \mathbb{R}$.
- (2) $J_1, \dots, J_m \subset \mathbb{R}$ are disjoint closed intervals. For each $i \in \{1, \dots, m\}$ the length of J_i is r_i and $f_i^+, f_i^- : J_i \rightarrow \Gamma$ are isometric embeddings such that the following hold:
 - (a) All but finitely many points of Γ lie in the image of precisely two of the maps f_i^+, f_i^- .
 - (b) No point of Γ lies in the image of four of the maps f_i^+, f_i^- .

From this data we construct a surface with boundary by gluing bands of width r_i to Γ using the maps f_i^+ and f_i^- as follows. For each i define the band B_i as $B_i := J_i \times [0, 1]$ and define $f_i : J_i \times \{0, 1\} \rightarrow \Gamma$ as $f_i(t, 0) = f_i^+(t)$ and $f_i(t, 1) = f_i^-(t)$ for all $t \in J_i$ and $1 \leq i \leq m$.

The associated band complex is then defined as

$$\Sigma := (\Gamma \sqcup B_1 \sqcup \dots \sqcup B_m) / \sim$$

where \sim is the equivalence relation generated by $f_i(x) \sim x$ for all $x \in J_i \times \{0, 1\}$ and $1 \leq i \leq m$. Clearly Σ is a surface with boundary, see Figure 3.6 for a once-punctured torus.




FIGURE 3.6. A union of (two) bands homeomorphic to a punctured torus

Now each band B_i has a natural foliation where the leaves are the sets $\{t\} \times [0, 1]$ with $t \in J_i$. These foliations of the B_i extend to a foliation of Σ where leaves of the B_i are joint along common endpoints in Γ to leaves of Σ . Thus the leaves of the foliation of Σ are graphs with vertices of valency at most 3. Moreover there are at most $2m$ vertices of valency 3 and each boundary component is a subset of some leaf. In particular all but finitely many leaves are homeomorphic to \mathbb{R} or S^1 . We will be mostly interested in cases where the foliation of Σ has the property that every leaf is dense and either homeomorphic to \mathbb{R} or consists of a single boundary component of Σ with infinite rays attached to them. Such foliations are called *arational*. If we assume that the foliation depicted in Figure 3.6 is arational then there is precisely one leaf that is not homeomorphic to \mathbb{R} , it consists of the single boundary component with two rays attached to it.

Now the foliation of Σ induces a foliation of the universal covering $\tilde{\Sigma}$ and the action of $\pi_1(\Sigma)$ on $\tilde{\Sigma}$ by deck transformations respects this foliation, i.e. gives an action on the leaf space T . It is clear that any element corresponding to a boundary component of Σ acts with a fixed point on T as it preserves some lift of the leaf containing the corresponding boundary component. T comes with a pseudo-metric where the distance between two leaves x and y is the length of the shortest path starting in x and ending in y . The length of the curve is here measured with respect to the natural transverse measure of the foliation, in particular for any curve contained in the lift of some band the transverse measure is simply the length of the projection to the base; see [5] for details. If the foliation is arational then the pseudo-metric is a metric turning T into a real tree; for the remainder of this section we assume this to be the case.

Let $\gamma_1, \dots, \gamma_k$ be homotopy classes representing the boundary components. By construction the γ_i act with a fixed point on T . Recall that we obtain a 2-orbifold group $H = \pi_1(\mathcal{O})$ from $\pi_1(\Sigma)$ by factoring out relations of the type $\gamma_i^{r_i}$ for $1 \leq i \leq k$ with $r_i \in \mathbb{N}_0$. Denote the kernel of the projection $\pi_1(\Sigma) \rightarrow H$ by N . Now H clearly acts on $T' := T/N$ and T' is again a real tree as N is generated by elliptic elements. We say that this action of H is dual to an arational foliation on the cone type orbifold \mathcal{O} .

3.4. The structure theorem

In this section we state the structure theorem for finitely generated groups acting on \mathbb{R} -trees as it appears in [22]. This theorem (and its relatives) are usually simply referred to as the Rips machine.

We recall from [22] that a G -tree T satisfies the *ascending chain condition* if for any sequence of arcs $I_1 \supset I_2 \supset \dots$ in T whose lengths converge to 0, the sequence of the stabilizers of the segments is eventually constant.

The following theorem is a generalization of Sela’s version of the Rips machine for finitely generated groups [38]. Other than Sela’s original version it allows us to deal with torsion.

Theorem 3.4 (Main Theorem of [22]). *Consider a non-trivial action of a f.g. group G on an \mathbb{R} -tree T by isometries. Assume that*

- *T satisfies the ascending chain condition,*
- *for any unstable arc $J \subset T$,*
 - *$\text{stab}(J)$ is finitely generated*
 - *$\text{stab}(J)$ is not a proper subgroup of any conjugate of itself, i.e. $\forall g \in G, (\text{stab}(J))^g \subset \text{stab}(J) \Rightarrow (\text{stab}(J))^g = \text{stab}(J)$.*

Then either G splits over the stabilizer of an unstable arc or over the stabilizer of an infinite tripod, or T splits as a graph of actions

$$\mathcal{G} = (\mathbb{A}, (T_v)_{v \in VA}, (p_e^\alpha)_{e \in EA})$$

where each vertex action of A_v on the vertex tree T_v is either

- *simplicial: a simplicial action on a simplicial tree,*
- *of orbifold (or Seifert) type: the action of A_v has kernel N_v and the faithful action of A_v/N_v is dual to an arational measured foliation on a closed 2-orbifold with boundary, or*
- *axial: T_v is a line and the image of A_v in $\text{Isom}(T_v)$ is a finitely generated group acting with dense orbits on T_v .*

Note that if the G -tree T admits a splitting as a graph of actions \mathcal{G} as in Theorem 3.4 and if \mathbb{A} contains a nondegenerate simplicial vertex tree, we get a refined splitting of T as a graph of actions

$$\mathcal{G}' = (\mathbb{A}', (T_v)_{v \in VA'}, (p_e^\alpha)_{e \in EA'}, l)$$

such that any vertex tree is either of axial type, or of orbifold type or is degenerate, i.e. consists of a single point. This is easily achieved by decomposing each simplicial vertex tree using Bass–Serre theory, possibly after subdividing some edges to ensure that the original attaching points are vertices. Note that if an edge e of the refined graph of actions has non-zero length then both $T_{\alpha(e)}$ and $T_{\omega(e)}$ are degenerate.

4. The virtually Abelian JSJ-decomposition of Γ -limit groups

Throughout this chapter Γ is a hyperbolic group. We first study basic properties of virtually Abelian splittings of Γ -limit groups and then discuss virtually Abelian JSJ-decompositions of finitely generated Γ -limit groups, splittings that reveal all virtually Abelian splittings simultaneously. We closely follow Sela's construction of the JSJ-decomposition [38] but similar to the discussion of Bestvina and Feighn [6] do not require the JSJ to be unfolded. The fact that the JSJ can be chosen to be unfolded can be established using the acylindricity of the splittings. This has been observed by Rips and Sela [36] in the case of splittings over cyclic groups. That this can also be done in the current setting has been established independently by Guiradel and Levitt [23] and the first author in his PhD thesis [34].

4.1. Modifying splittings

Recall that a graph of groups is *minimal* if the corresponding Bass–Serre tree contains no proper invariant subtree. A finite graph of groups is minimal iff there is no surjective boundary monomorphism into a vertex group of a valence 1 vertex. A graph of groups is called *reduced* if no boundary monomorphism is surjective. Thus any reduced finite graph of groups is minimal.

The JSJ-decompositions we construct reveal virtually Abelian splittings only up to certain modifications. We start by introducing these modifications, they all leave the fundamental group unchanged.

Definition 4.1. Let \mathbb{A} be a graph of groups. A *splitting move* on \mathbb{A} is one of the following modifications of \mathbb{A} .

- (1) *Boundary slide:* Let $e \in EA$. A *boundary slide* (of the boundary monomorphism α_e) is the replacement of α_e by $c_g \circ \alpha_e$ for an element $g \in A_{\alpha(e)}$.
- (2) *Edge slide:* Let $e_1 \neq e_2 \in EA$ such that $\omega(e_1) = \alpha(e_2)$. Suppose that $\omega_{e_1}(A_{e_1})$ is in $A_{\omega(e_1)}$ conjugate to a subgroup of $\alpha_{e_2}(A_{e_2})$.

Then we first perform a boundary slide such that $\omega_{e_1}(A_{e_1}) \leq \alpha_{e_2}(A_{e_2})$ and then replace e_1 with an edge e'_1 with edge group $A_{e'_1} = A_{e_1}$ such that

- (a) $\alpha(e'_1) = \alpha(e_1)$ and $\alpha_{e'_1} = \alpha_{e_1}$.
- (b) $\omega(e'_1) = \omega(e_2)$.
- (c) $\omega_{e'_1} = \omega_{e_2} \circ \alpha_{e_2}^{-1} \circ \omega_{e_1}$.

The combination of the initial boundary slide and the subsequent modification is called an *edge slide* of e_1 over e_2 .

- (3) *Folding/Unfolding*: Let $e \in EA$ and $\alpha_e(A_e) < C < A_{\alpha(e)}$. A *folding* along e is the replacement of $A_{\omega(e)}$ by $C *_{A_e} A_{\omega(e)}$, A_e by C and the corresponding replacement of the boundary monomorphisms α_e and ω_e . The inverse of a folding along e is an *unfolding* along e .

Note that we usually only consider graphs of groups up to boundary slides.

An edge slide can be alternatively defined on the Bass–Serre tree: Given two inequivalent edges f_1 and f_2 such that $\omega(f_1) = \alpha(f_2)$ and $\text{stab } f_1 \leq \text{stab } f_2$ we slide f_1 over f_2 in an equivariant way. One can also think of it as first (equivariantly) subdividing f_1 into f'_1 and f''_1 and then (equivariantly) identifying f''_1 with f_2 . Note that the action on the vertex set of the tree is unchanged under this operation.




FIGURE 4.1. An edge slide as seen in the tree

The following lemma is trivial as the existence of a surjective boundary monomorphism is preserved by boundary slides and edge slides.

Lemma 4.2. *Let \mathbb{A} be a graph of groups. Assume that \mathbb{A}' is obtained from \mathbb{A} by boundary slides and edge slides. Then \mathbb{A} is reduced and minimal iff \mathbb{A}' is reduced and minimal.*

4.2. Virtually Abelian splittings of Γ -limit groups

In this section we study splittings of Γ -limit groups as fundamental groups of graphs of groups with virtually Abelian edge groups. We call such splittings *virtually Abelian splittings*.

In the following we call a virtually Abelian group *large* if it contains a one-ended subgroup. Note that for finitely generated virtually Abelian groups, being large is equivalent to being one-ended. A crucial observation in this section will be that any finite virtually Abelian splitting of a Γ -limit group can be modified by boundary slides and some further simple modifications such that all large virtually Abelian subgroups are elliptic.

Lemma 4.3. *Let L be a Γ -limit group with virtually Abelian graph of groups decomposition \mathbb{A} . Let $M \leq L$ be a maximal large one-ended virtually Abelian subgroup which is not elliptic in \mathbb{A} .*

Then M acts with an invariant line $T \subset \tilde{\mathbb{A}}$ and the following hold: If e_1 and e_2 are edges in T and $g \in L$ s.t. $ge_1 = e_2$, then $g \in M$. In particular $\text{stab}_M(e) = \text{stab}_L(e)$ for any edge e of T and $g \in M$ iff $gT = T$.

Proof. As M does not contain a non-Abelian free group it either acts with a fixed point or with an invariant line or parabolically, i.e. fixes a unique end of T . By assumption M does not act with a fixed point, thus we need to show that M does not act parabolically.

Assume to the contrary that M acts parabolically on $\tilde{\mathbb{A}}$. Thus M preserves a unique end and it follows that either M is a strictly ascending HNN extension or that every element of M is elliptic. Being virtually Abelian, M can not be a strictly ascending HNN-extension, hence every $g \in M$ is elliptic. Let e_1, e_2, \dots be a ray in $\tilde{\mathbb{A}}$ representing the fixed end. We get an infinite ascending sequence of stabilizers

$$\text{stab}_M(e_1) \leq \text{stab}_M(e_2) \leq \dots,$$

such that $M = \bigcup \text{stab}_M(e_i)$. As M is infinite and there is a uniform bound for the order of finite Γ -limit groups it follows that there exists i_0 such that $\text{stab}_M(e_i)$ is infinite for $i \geq i_0$. This implies in particular that $\text{stab}_M(e_i) = \text{stab}_L(e_i)$ for $i \geq i_0$ as M is the unique maximal virtually Abelian subgroup of L containing $\text{stab}_M(e_i)$. As EA is finite there exists $i > j \geq i_0$ and $g \in L$ such that $e_i = ge_j$. Now

$$\begin{aligned} \text{stab}_L(e_j) &= \text{stab}_L(e_i) \cap \text{stab}_L(e_j) \\ &= \text{stab}_L(ge_j) \cap \text{stab}_L(e_j) \\ &= g \text{stab}_L(e_j)g^{-1} \cap \text{stab}_L(e_j), \end{aligned}$$

thus by Corollary 2.25 g is contained in the maximal virtually Abelian subgroup containing $\text{stab}_L(e_j)$, i.e. $g \in M$. But g acts without fixed point, which contradicts the assumption on the action of M .

Thus M preserves a line $T \subset \tilde{\mathbb{A}}$. As M is large it follows that $\text{stab}_M(e)$ is infinite for all $e \in T$. M is the unique maximal virtually Abelian subgroup of L containing $\text{stab}_L(e)$ for any edge e of T . It follows that $\text{stab}_L(e_1) = \text{stab}_M(e_1) = \text{stab}_M(e_2) = \text{stab}_L(e_2)$ for any two edges e_1, e_2 of T . Thus the above arguments show that $g \in M$ if $ge_1 = e_2$. \square

Proposition 4.4. *Let L be a one-ended Γ -limit group with finite virtually Abelian graph of groups decomposition \mathbb{A} . Then after finitely many edge slides we can assume that for any large maximal virtually Abelian subgroup M of L one of the following holds.*

- (1) M is elliptic.
- (2) M is the unique maximal virtually Abelian subgroup containing some edge group A_e and M is of type $A_1 *_{A_e} A_2$ where $A_1 \leq A_{\alpha(e)}$, $A_2 \leq A_{\omega(e)}$ and $|A_1 : \alpha_e(A_e)| = |A_2 : \omega_e(A_e)| = 2$.

- (3) M is the unique maximal virtually Abelian subgroup containing the edge group A_e of some loop edge e . Furthermore $M = A_e *_{A_e}$ where the stable letter is the element corresponding to the loop edge e , in particular $\alpha_e(A_e) = \omega_e(A_e) \leq A_{\alpha(e)}$.

Proof. Assume that M is not elliptic. By Lemma 4.3, there is an M -invariant line $T \subset \tilde{\mathbb{A}}$. Let f_1, f_2, \dots, f_k be an edge path in T that is a fundamental domain for the action of M on T .

If $k = 1$, i.e. if the edge path consists of a single edge, then there is nothing to show as we are either in situation (2) or in situation (3). Thus we can assume that $k \geq 2$.

It follows from Lemma 4.3 that $\text{stab}_L(f_1) = \text{stab}_L(f_2)$. Thus we can L -equivariantly slide f_1 over f_2 . The new fundamental domain for M has only $k - 1$ edges. After finitely many slides the fundamental domain for M consists of a single edge, which implies that the claim of the proposition holds for M .

The conclusion now follows from the observation that an edge is slid over another edge only if their edge groups are contained in the same maximal virtually Abelian subgroup. Thus the above process for one maximal virtually Abelian subgroup does not affect the validity of the conclusion of the proposition for the other. Thus we can use finitely many edge slides to obtain the desired conclusion for all maximal large virtually Abelian subgroups simultaneously. This process is finite as there are only finitely many edge groups and therefore only finitely many conjugacy classes of large maximal virtually Abelian subgroups that act non-elliptically. \square

We now explain how splittings that satisfy the conclusion of Proposition 4.4 can be modified such that afterwards all large virtually Abelian subgroups are elliptic. We call a virtually Abelian splitting with this property *compatible*.

Let \mathbb{A} be a finite virtually Abelian splitting of L satisfying the conclusion of Proposition 4.4. Let M be a large maximal virtually Abelian subgroup of L that is not elliptic.

Assume first that M satisfies (2) of Proposition 4.4. Then we subdivide the edge e into edge e_1 and e_2 such that $\alpha(e_1) = \alpha(e)$, $\omega(e_2) = \omega(e)$ and $\omega(e_1) = \alpha(e_2) = v'$ is a new vertex. Moreover $A_{e_1} = A_1$, $A_{e_2} = A_2$ and $A_{v'} = A_1 *_{A_e} A_2$ and the boundary monomorphisms are the natural ones, see Figure 4.2 for an illustration of both the case where e is a non-loop edge, and where e is a loop edge.

If M satisfies (3) of Proposition 4.4 then we remove the edge e and add a new edge e' with $A_{e'} = A_e$ such that $\alpha(e') = \alpha(e)$, and $\omega(e') = w$ is a new vertex with vertex group $A_w = A_e *_{A_e}$. The boundary monomorphisms are the natural embeddings of A_e , see Figure 4.3.

We now argue that for a finitely generated Γ -limit group L there exists an upper bound for the complexity of a minimal (and reduced) virtually Abelian splitting of L , which only depends on the rank of L and $N(\Gamma)$. For a given graph of groups \mathbb{A} , we define its




FIGURE 4.2. A new vertex with maximal virtually Abelian vertex group




FIGURE 4.3. A new vertex with maximal virtually Abelian vertex group

complexity $C(\mathbb{A})$ by

$$C(\mathbb{A}) := e(A) + \beta_1(A),$$

where $e(A) := |EA|/2$ denotes the number of edge pairs of the graph A underlying \mathbb{A} , and $\beta_1(A)$ is the first Betti number of A . While the Betti number is bounded from above by the rank of L , a bound for $e(A)$ can be obtained from Theorem 4.5 and Lemma 4.6 below. Recall that a graph of groups is called (k, C) -acylindrical if the stabilizer of any segment $[v, w]$ in the Bass–Serre tree with $d(v, w) > k$ is of order at most C .

The following theorem from [45] provides a bound for the complexity of (k, C) -acylindrical splittings. It is a generalization of Sela’s acylindrical accessibility theorem [38] which deals with the case $C = 1$, see also [44].

Theorem 4.5. *Let \mathbb{A} be a reduced and minimal (k, C) -acylindrical graph of groups with $k \geq 1$. Then*

$$e(A) \leq (2k + 1) \cdot C \cdot (\text{rank } \pi_1(\mathbb{A}) - 1).$$

Thus Theorem 4.5 together with Lemma 4.6 below provide a bound for the complexity of reduced virtually Abelian compatible splittings of Γ -limit groups.

For an infinite virtually Abelian subgroup $H \leq L$ we denote by $\text{MA}(H)$ the conjugacy class of K where K is the unique maximal virtually Abelian subgroup containing H . Note that $\text{MA}(H) = \text{MA}(H')$ if H and H' are conjugate. Thus it makes sense to speak of $\text{MA}(A_e)$.

Lemma 4.6. *Let \mathbb{A} be a compatible virtually Abelian splitting of L and assume that all edge groups are infinite. Then \mathbb{A} can be modified by a finite sequence of edge slides to be $(2, N(\Gamma))$ -acylindrical.*

Proof. Note first that for any $e \in EA$ the representatives of $\text{MA}(A_e)$ are elliptic. If the representatives are large this holds as \mathbb{A} is assumed to be compatible; otherwise they contain an elliptic subgroup of finite index (a conjugate of A_e) which also implies that they are elliptic.

For any conjugacy class of maximal virtually Abelian elliptic subgroups $[M]$ we choose a vertex $v_{[M]}$ such that M is conjugate into $A_{v_{[M]}}$. By a finite sequence of edge slides any edge e can be slid such that it is incident to $v_{\text{MA}(A_e)}$. The $(2, N(\Gamma))$ -acylindricity of the obtained splitting is easily verified: Otherwise there exists a segment Y of length 3 in the Bass–Serre tree fixed by some infinite virtually Abelian group H . The construction implies that any edge of Y must be incident to a vertex that is a lift of $v_{[M]}$. As Y is of length 3 there must be two distinct such vertices, say v_1 and v_2 . Choose $g \in L$ such that $gv_1 = v_2$. It follows that $H \subset \text{stab}(v_1) \cap \text{stab}(v_2) = \text{stab}(v_1) \cap g \text{stab}(v_1)g^{-1}$. As H is infinite and $\text{stab}(v_1)$ is maximal virtually Abelian it follows from Corollary 2.25 that $g \in \text{stab}(v_1)$, a contradiction. Thus the obtained graph of groups is $(2, N(\Gamma))$ -acylindrical. \square

The construction of the $(2, N(\Gamma))$ -acylindrical graph of groups in the proof of Lemma 4.6 depends on the choice of the vertices $v_{[M]}$ and the output is therefore not canonical. We will obtain a canonical splitting (up to boundary slides) by performing the following *normalization process* for a given compatible virtually Abelian splitting \mathbb{A} of L . As before we assume that all edge groups are infinite.

This process is only a slight modification of the proof of Lemma 4.6. Choose M_1, \dots, M_k such that $[M_1], \dots, [M_k]$ is the collection of those conjugacy classes of maximal virtually Abelian subgroups which appear as $\text{MA}(A_e)$ for some $e \in EA$. This collection is clearly finite as EA is finite.

For each $i = 1, \dots, k$ choose a vertex v_i such that M_i is conjugate into A_{v_i} and introduce a new vertex $v_{[M_i]}$ and an edge e_i with $\alpha(e_i) = v_i$ and $\omega(e_i) = v_{[M_i]}$ such that $A_{v_{[M_i]}} = A_{e_i} = M_i$, $\omega_{e_i} = \text{id}_{M_i}$ and that $\text{Im}(\alpha_{e_i}) \leq L$ is conjugate to M_i . We then slide as in the proof of Lemma 4.6, i.e. slide every edge e such that it is incident to $v_{\text{MA}(A_e)}$.

Finally, we minimize the obtained graph of groups by removing unnecessary valence 1 vertices and corresponding edges.

It is clear that, up to boundary slides, the obtained graph of groups does not depend on the choice of the v_i . We say that \mathbb{A} is in *normal form* if it is the output of the normalization process for some graph of groups \mathbb{A}' (or equivalently, if the normalization process of \mathbb{A} reproduces \mathbb{A}). We further call the vertices v_1, \dots, v_k in the above construction the *characteristic vertices* of the normal form. By construction any edge of a graph of groups in normal form is incident to exactly one characteristic vertex, thus the underlying graph is bipartite. A graph of groups in normal form is $(2, N(\Gamma))$ -acylindrical by the same argument as in the proof of Lemma 4.6 and minimal by construction. Note that a graph of groups in normal form may be non-reduced as some edge groups may be maximal virtually abelian subgroups of L . It is however easily verified that any graph of groups in normal form can be obtained by applying the normalization process to some reduced, minimal and $(2, N(\Gamma))$ -acylindrical graph of groups \mathbb{A}' . As by construction $C(\mathbb{A}) \leq 2C(\mathbb{A}')$ it follows, using Theorem 4.5, that there is a global upper bound for the complexity of all normal form splittings of a given finitely generated one-ended Γ -limit group L .

4.3. Morphisms of graphs of groups

For a based simplicial G -tree $T = (T, \tilde{v}_0)$ and a based simplicial H -tree $Y = (Y, \tilde{u}_0)$ a morphism from T to Y is a pair (φ, f) where $\varphi : G \rightarrow H$ is a homomorphism and $f : T \rightarrow Y$ is a simplicial map such that $f(\tilde{v}_0) = \tilde{u}_0$ and that

$$f(gx) = \varphi(g)f(x)$$

for all $x \in T$ and $g \in G$. Any such morphism can be encoded on the level of the associated graphs of groups. We will discuss such morphisms for graphs of groups and make some basic observations.

A *morphism* from a graph of groups \mathbb{A} to a graph of groups \mathbb{B} is a tuple

$$\mathfrak{f} = (f, \{\psi_v \mid v \in VA\}, \{\psi_e \mid e \in EA\}, \{o_e \mid e \in EA\})$$

where

- (1) $f : A \rightarrow B$ is a graph morphism.
- (2) ψ_v is a homomorphism from A_v to $B_{f(v)}$ for all $v \in VA$.
- (3) ψ_e is a homomorphism from A_e to $B_{f(e)}$ for all $e \in EA$.
- (4) $o_e \in B_{f(\alpha(e))}$ for all $e \in EA$.

$$(5) \psi_{\alpha(e)} \circ \alpha_e = i_{o_e} \circ \alpha_{f(e)} \circ \psi_e \text{ for all } e \in EA.$$

We will denote $\sigma_{e^{-1}}$ alternatively by t_e . A morphism from \mathbb{A} to \mathbb{B} induces a homomorphism

$$\tilde{f}_* : \pi_1(\mathbb{A}, v_0) \rightarrow \pi_1(\mathbb{B}, f(v_0))$$

given by

$$[a_0, e_1, a_1, \dots, a_{k-1}, e_k, a_k] \mapsto [b_0, f(e_1), b_1, \dots, b_{k-1}, f(e_k), b_k]$$

where $b_i = t_{e_i} \psi_{\omega(e_i)}(a_i) \sigma_{e_{i+1}}$ for $i = 0, \dots, k$ (with $t_{e_0} = t_{e_{k+1}} = 1$).

We will write ψ_v^\dagger or ψ_e^\dagger instead of ψ_v or ψ_e if we want to make explicit that the maps come from the morphism \tilde{f} . We will further say that a morphism \tilde{f} is surjective if \tilde{f}_* is surjective.

The morphism \tilde{f} determines a morphism $\tilde{f} : \widetilde{(\mathbb{A}, v_0)} \rightarrow \widetilde{(\mathbb{B}, u_0)}$ (that maps the base point \tilde{v}_0 to \tilde{u}_0). The pair (\tilde{f}_*, \tilde{f}) is a morphism from the $\pi_1(\mathbb{A}, v_0)$ -tree $\widetilde{(\mathbb{A}, v_0)}$ to the $\pi_1(\mathbb{B}, u_0)$ -tree $\widetilde{(\mathbb{B}, u_0)}$, see [24] for details. Moreover any morphism from a G -tree T to an H -tree Y occurs this way.

The following simple proposition will be useful in subsequent sections. In the statement of the proposition we identify the fundamental group of \mathbb{A} and the fundamental group of $\bar{\mathbb{A}}$ obtained from \mathbb{A} by edge collapses and subdivision in the natural way.

Proposition 4.7. *Let \mathbb{A} and \mathbb{B} be finite graphs of groups and*

$$\eta : \pi_1(\mathbb{A}, v_0) \rightarrow \pi_1(\mathbb{B}, u_0)$$

be an isomorphism. Suppose further that there exists a map $h : VA \rightarrow VB$ such that the following hold:

$$(1) h(v_0) = u_0$$

$$(2) \eta(A_v) \text{ is conjugate to a subgroup of } B_{h(v)} \text{ for all } v \in VA.$$

Then there exists a graph of groups $\bar{\mathbb{A}}$ obtained from \mathbb{A} by collapses of edges followed by subdivisions of edges and a morphism $\tilde{f} : \bar{\mathbb{A}} \rightarrow \mathbb{B}$ such that $f(v) = h(v)$ for all $v \in VA$ and $\tilde{f}_ = \eta$.*

Proof. Let $T_A = \widetilde{(\mathbb{A}, v_0)}$ and $T_B = \widetilde{(\mathbb{B}, u_0)}$ be the Bass–Serre trees with base points \tilde{v}_0 and \tilde{u}_0 . It suffices to show that, after (equivariant) collapses and subdivisions of edges of T_A , there exists a morphism (η, f) from the $\pi_1(\mathbb{A}, v_0)$ -tree T_A to the $\pi_1(\mathbb{B}, u_0)$ -tree T_B such that $f(\tilde{v}_0) = \tilde{u}_0$ and that $p_B(f(v)) = h(p_A(v))$ for all $v \in VT_A$ where $p_A : T_A \rightarrow A$ and $p_B : T_B \rightarrow B$ are the canonical quotient maps.

Pick a maximal tree Y_A in A and a lift \tilde{Y}_A to T_A such that the lift of v_0 is \tilde{v}_0 . By assumption we can choose for each vertex $v \in \tilde{Y}_A$ a vertex $w_v \in T_B$ such that $\eta(\text{stab } v) \leq \text{stab } w_v$ and




FIGURE 4.4. The restriction of f to \tilde{Y}_A

that $h(p_A(v)) = p_B(w_v)$; we may further assume that $w_{\tilde{v}_0} = \tilde{u}_0$. We now define a map $f : VT_A \rightarrow VT_B$ by $gv \mapsto \eta(g)w_v$ for all $v \in \tilde{Y}_A$ and $g \in \pi_1(\mathbb{A}, v_0)$. The map is easily extended to the edges of T_A by mapping an edge $e = (v_1, v_2)$ to the reduced edge path p from $f(v_1)$ to $f(v_2)$. Here we apply $k - 1$ subdivisions to the edge e if the length k of p is greater than 1, or collapse e in case $k = 0$.

To make the map simplicial we need to collapse e if $f(v_1) = f(v_2)$ and subdivide it $d_{T_B}(f(v_1), f(v_2)) - 1$ times otherwise. \square

Remark 4.8. It follows from the above proof that the number of subdivisions applied to an edge $e = (v_1, v_2) \in EA$ is bounded by $\text{diam}_{\mathbb{B}}(\text{Fix } \eta(A_e)) - 1$ as $\eta(A_e)$ fixes the segment $[f(v_1), f(v_2)]$.

4.4. The virtually Abelian JSJ-decomposition of a Γ -limit group

In this section we establish the existence of virtually Abelian JSJ-decompositions of finitely generated Γ -limit groups. A virtually Abelian JSJ-decomposition of a group G is a splitting of G in which all compatible virtually Abelian splittings of G are apparent.

In the following we say that a vertex group A_v of a graph of groups \mathbb{A} is a QH-vertex group (quadratically hanging vertex groups) if the following hold.

- (1) A_v is finite-by-orbifold, i.e. there exists an orbifold \mathcal{O} with fundamental group $O = \pi_1(\mathcal{O})$, some finite group E and a short exact sequence

$$1 \rightarrow E \rightarrow A_v \xrightarrow{\pi} O \rightarrow 1.$$

- (2) For any edge $e \in EA$ s.t. $\alpha(e) = v$ there exists a peripheral subgroup O_e of O such that $\alpha_e(A_e)$ is in A_v conjugate to a finite index subgroup of $\pi^{-1}(O_e)$.

We will also say that a subgroup of G is a QH-subgroup if it is conjugate to a QH-vertex group of some splitting \mathbb{A} of G .

It is a trivial but important observation that any essential simple closed curve or essential segment joining two points on the reflection boundary of \mathcal{O} induces a splitting of $G = \pi_1(\mathbb{A})$ over a 2-ended group. We call such a splitting *geometric* with respect to the QH-subgroup A_v .

We can depict the splitting \mathbb{A} by drawing the orbifold \mathcal{O} for the vertex group A_v and depicting all non-QH vertex groups as balls joint by edges.




FIGURE 4.5. A QH-subgroup with a simple closed curve representing a splitting of $\pi_1(\mathbb{A})$ as an HNN-extension

Recall that a 1-edge splitting \mathbb{A}_1 , i.e. a splitting as an amalgamated product or an HNN-extension, of a group G is called elliptic with respect to another splitting \mathbb{A}_2 if the edge group of \mathbb{A}_1 is conjugate to a subgroup of a vertex group of \mathbb{A}_2 . Otherwise \mathbb{A}_1 is called hyperbolic with respect to \mathbb{A}_2 . It is an important observation [36] that if G is one-ended and \mathbb{A}_1 and \mathbb{A}_2 are 1-edge splittings of G over 2-ended groups then the two splittings are either both elliptic or both hyperbolic with respect to each other. In the first case we say that \mathbb{A}_1 and \mathbb{A}_2 are elliptic-elliptic and in the latter that they are hyperbolic-hyperbolic.

The following theorem is the key step in the proof of the JSJ-decomposition, it implies in particular the existence of a splitting of a Γ -limit group that encodes all hyperbolic-hyperbolic splittings over 2-ended subgroups.

Theorem 4.9. *Let G be a f.g. one-ended group. Assume that there exists N such that any finite subgroup of G has order at most N .*

Then there exists a reduced, minimal $(2, N)$ -acylindrical graph of groups decomposition \mathbb{A} of G such that the following hold:

- (1) *Any 1-edge splitting of G over a 2-ended group that is hyperbolic-hyperbolic with respect to another 1-edge splitting over a 2-ended group is geometric with respect to some QH-vertex group of \mathbb{A} .*

- (2) *If \mathbb{B} is a virtually Abelian splitting of G such that no edge group of \mathbb{B} is 2-ended and hyperbolic-hyperbolic with respect to another splitting over a 2-ended group then the QH-vertex groups of \mathbb{A} are elliptic with respect to \mathbb{B} .*
- (3) *Any QH-subgroup of G is conjugate to a subgroup of a QH-vertex group A_v , moreover this subgroup corresponds to a suborbifold of the orbifold corresponding to v . In particular every maximal QH-subgroup is conjugate to a vertex group of \mathbb{A} .*

Proof. For the proof we can essentially refer to the construction of the JSJ-decomposition of Dunwoody and Sageev [15] and Fujiwara and Papasoglu [16], but we need a few adaptations.

First, note that any given splitting can be modified to be $(2, N)$ -acylindrical while preserving the conjugacy classes of the QH-vertex groups. Indeed, we first refine the splitting by replacing each QH-vertex group A_v by a tree of groups consisting of a vertex x_v with vertex group A_v and for each peripheral subgroup P an edge e_P with edge group P and a vertex x_P with vertex group P such that $\alpha(e_P) = x_v$ and $\omega(e_P) = x_P$ with the boundary monomorphisms being the obvious ones. We may refine in such a way that all previous edges are incident to one of the vertices of type x_P . We then collapse all edges not incident to a vertex of type x_v . The $(2, N)$ -acylindricity of the obtained splitting follows from the fact that for any QH-vertex group A_v and peripheral subgroups P_1 and P_2 corresponding to distinct boundary components the following hold.

- (1) $gP_1g^{-1} \cap P_1$ is finite for all $g \in A_v \setminus P_1$.
- (2) $gP_1g^{-1} \cap P_2$ is finite for all $g \in A_v$.

This observation implies in particular that the number of QH-vertex groups of any splitting of G is bounded in terms of N and the rank of G because of Theorem 4.5. This also bounds the complexity (genus and number of exceptional points) of the orbifolds corresponding to the QH-vertex groups as orbifolds of large complexity can be cut along essential simple closed curves or segments to produce a splitting with a higher number of QH-vertex groups that is again $(2, N)$ -acylindrical.

Now, Dunwoody and Sageev [15] and Fujiwara and Papasoglu [16] construct splittings such that arbitrary finite collections of hyperbolic-hyperbolic splittings over 2-ended groups can be seen as geometric splittings in QH-vertex groups. In their proofs they then assume that G is finitely presented so they can apply Bestvina–Feighn accessibility [4] to guarantee termination of the construction. In our case we can exploit the fact that the obtained splittings are $(2, N)$ -acylindrical after the modifications discussed above. The

same argument was applied in the construction of the quadratic decomposition in [36]. (3) also follows as in [15, 16, 36].

To see (2) suppose that a QH-vertex group A_v of \mathbb{A} is non-elliptic in \mathbb{B} . Then there exists a geometric 1-edge splitting \mathbb{D} corresponding to some essential simple closed curve or essential segment on the orbifold such that the single edge group D_e of \mathbb{D} is non-elliptic in \mathbb{B} , see Corollary 4.12 of [16]. It follows that \mathbb{D} is hyperbolic with respect to the 1-edge splitting corresponding to some edge $f \in EB$, i.e. the splitting obtained by collapsing all edges of B but f . Note that B_f is elliptic in \mathbb{D} : If B_f is 2-ended this holds by assumption on \mathbb{B} and if B_f is large then B_f must act on $\widetilde{\mathbb{D}}$ with infinite pointwise stabilizer contradicting the $(2, N)$ -acylindricity of \mathbb{D} . It follows, see Remark 2.3 of [16] that G splits over a subgroup of D_e of infinite index, i.e. that G splits over a finite group. This contradicts the one-endedness of G . \square

Before we proceed with the virtually Abelian JSJ-decomposition, we introduce some notation concerning virtually Abelian vertex groups of splittings. Recall that if M is an infinite virtually Abelian subgroup of a Γ -limit group then M^+ denotes the unique maximal finite-by-Abelian subgroup of M which is of index at most 2.

Definition 4.10. Let \mathbb{A} be a graph of groups decomposition of a finitely generated 1-ended Γ -limit group L . Let $v \in VA$ such that A_v is virtually Abelian.

Let Δ be the set of homomorphisms $\eta : A_v^+ \rightarrow \mathbb{Z}$ such that $\eta(\alpha_e(A_e^+)) = 0$ for all $e \in EA$ with $\alpha(e) = v$. We then define

$$P_v^+ := \{g \in A_v^+ \mid \eta(g) = 0 \text{ for all } \eta \in \Delta\}.$$

A simple homology argument shows that A_v^+/P_v^+ is a finitely generated free Abelian group whose rank is bounded from above in terms of rank L and the valence of v in A . Together with Theorem 4.9 this implies in particular that JSJ-decompositions in the sense of the following definition exist.

Definition 4.11. Let L be a finitely generated one-ended Γ -limit group and \mathbb{A} be a virtually Abelian compatible splitting of L . Then \mathbb{A} is called a virtually Abelian JSJ-decomposition of L if the following hold.

- (1) Every splitting over a 2-ended group that is hyperbolic-hyperbolic with respect to another splitting over a 2-ended group is geometric with respect to a QH-subgroup of \mathbb{A} .
- (2) Any edge group of \mathbb{A} that can be unfolded to be finite-by-Abelian is finite-by-Abelian.

- (3) For any virtually Abelian vertex group A_v , the rank of A_v^+/P_v^+ cannot be increased by unfoldings.
- (4) \mathbb{A} is in normal form and of maximal complexity among all virtually Abelian compatible splittings of L that satisfy (1)–(3).

In the following we refer to vertex groups of a virtually Abelian JSJ-decomposition which are neither QH nor virtually Abelian as *rigid*. Note that any virtually Abelian JSJ-decomposition can be obtained from a splitting as in Theorem 4.9 by refinements of non-QH-subgroups, unfoldings and the normalization process. This is true as the maximal QH-subgroups must be elliptic by Definition 4.11(2).

The following theorem describes the basic properties of virtually Abelian JSJ-decompositions of Γ -limit groups. In the following we say that a graph of groups \mathbb{B} is *visible* in \mathbb{A} if \mathbb{B} is obtained from \mathbb{A} by a sequence of collapses.

Theorem 4.12. *Let L be a finitely generated one-ended Γ -limit group and let \mathbb{A} be a virtually Abelian JSJ-decomposition of L . Then the following hold.*

- (1) *Let \mathbb{B} be a virtually Abelian compatible splitting of L such that all maximal QH-subgroups are elliptic. Assume further that \mathbb{B} is either in normal form or a 1-edge splitting. Then \mathbb{B} is visible in a graph of groups obtained from \mathbb{A} by unfoldings followed by foldings and edge slides.*
- (2) *Any other JSJ-decomposition \mathbb{B} of L can be obtained from \mathbb{A} by a sequence of unfoldings and foldings.*

Remark. While Theorem 4.12(2) implies that any finite collection of JSJ-decompositions has a common unfolding we do not claim that the JSJ can be chosen to be unfolded, i.e. such that no further unfolding is possible.

Proof. We first prove (1). Let $T_{\mathbb{A}}$ and $T_{\mathbb{B}}$ be the respective Bass–Serre-trees of \mathbb{A} and \mathbb{B} . For each $v \in VA$, by restricting the G -action on $T_{\mathbb{B}}$ to A_v , we obtain a (possibly trivial) splitting \mathbb{A}^v of A_v corresponding to a minimal A_v -invariant subtree. Note that by assumption these splittings are trivial for QH-vertex groups.

Denote by \mathbb{A}' the graph of groups obtained by refining \mathbb{A} in each vertex v by \mathbb{A}^v , and normalizing this refined graph of groups. By construction, neither the Betti number nor the number of edges of \mathbb{A} decrease by the refinement. As the complexity cannot increase by the maximality assumption, it follows that both the Betti number and the number of edges remain unchanged, in particular $C(\mathbb{A}') = C(\mathbb{A})$. We show that \mathbb{A}' can be obtained from \mathbb{A} by unfoldings and that \mathbb{B} is visible in \mathbb{A}' after foldings and edge slides.

By construction, all vertex groups of \mathbb{A}' are elliptic in \mathbb{A} . Therefore, by Lemma 4.7, there is a graph of groups $\bar{\mathbb{A}}$ obtained from \mathbb{A}' by collapses and subdivisions and a morphism $\bar{f} : \bar{\mathbb{A}} \rightarrow \mathbb{A}$ with $\bar{f}_* = \text{id}_G$, which maps the characteristic vertices of \mathbb{A}' to the characteristic vertices of \mathbb{A} . The $(2, N(\Gamma))$ -acylindricity of \mathbb{A} implies that the stabilizer of a segment of length 2 in the Bass–Serre tree $\tilde{\bar{\mathbb{A}}}$ can only have infinite stabilizer if its midpoint corresponds to a characteristic vertex. Thus the proof of Proposition 4.7 implies that no subdivision is necessary as one endpoint of each edge path that occurs as an image of an edge must correspond to a characteristic vertex and as characteristic and non-characteristic vertices alternate.

Moreover, as \mathbb{A} is minimal, $T_{\mathbb{A}}$ does not contain a proper G -invariant subtree, so f is surjective. This implies that no edges are collapsed as $\#EA = \#EA'$. It follows that $\bar{\mathbb{A}} = \mathbb{A}'$, and we obtain a morphism, again called \bar{f} , from \mathbb{A}' to \mathbb{A} whose underlying graph morphism f is a graph isomorphism.

For each $e \in EA'$, we have $A'_e \leq A_{f(e)}$ (we identify \mathbb{A}' with $\bar{f}(\mathbb{A}')$, in particular we write A'_e instead of $\psi'_e(A'_e)$). Assume that for some e , A'_e is a proper subgroup of $A_{f(e)}$, and assume w.l.o.g. that $\alpha(e)$ is the characteristic vertex of $\text{MA}(A'_e)$. Then we can perform a fold along e , replacing A'_e by $A_{f(e)}$ and $A'_{\omega(e)}$ by $A_{f(e)} *_{A'_e} A'_{\omega(e)}$. After applying finitely many such folds we get that $A'_e \cong A_{f(e)}$ for all $e \in EA'$. As $\bar{f}_* : \pi_1(\bar{\mathbb{A}}) \rightarrow \pi_1(\mathbb{A})$ is an isomorphism, it follows that \mathbb{A}' , after the foldings, is isomorphic to \mathbb{A} . Conversely, \mathbb{A}' can be obtained from \mathbb{A} by unfoldings.

We now show that \mathbb{B} is visible in a splitting obtained from \mathbb{A}' by foldings and edge slides. By construction all vertex groups of \mathbb{A}' are elliptic in \mathbb{B} . Again there is a graph of groups $\bar{\mathbb{A}}$, which is obtained from \mathbb{A}' by collapses of edges, and a morphism $\bar{f} : \bar{\mathbb{A}} \rightarrow \mathbb{B}$. The argument that no subdivision is necessary is the same as above.

After foldings that replace the edge groups \bar{A}_e with $B_{f(e)}$ we can assume that \bar{f} is bijective on edge groups. Assume now that $f(e_1) = f(e_2)$ for some $e_1, e_2 \in E\bar{\mathbb{A}}$, in particular $\bar{A}_{e_1} = \bar{A}_{e_2} = B_{f(e_1)}$. Possibly after changing the orientation of e_1 and e_2 we can assume that $\alpha(e_1) = \alpha(e_2) = v_{\text{MA}(A_{e_1})}$.

We can now alter $\bar{\mathbb{A}}$ by identifying e_1 and e_2 by a Stallings fold of type IA or IIIA, see [4], clearly \bar{f} factors through this fold. Note that this fold can also be thought of as first sliding e_1 over e_2 and then collapsing e_1 . After finitely many such operations we obtain a graph of groups $\hat{\bar{\mathbb{A}}}$ and a morphism $\hat{\bar{f}} : \hat{\bar{\mathbb{A}}} \rightarrow \mathbb{B}$ such that $\hat{\bar{f}}$ is a graph isomorphism, that is bijective on edge groups and induces an isomorphism on the level of the fundamental group. Thus $\hat{\bar{\mathbb{A}}}$ is isomorphic to \mathbb{B} . As $\hat{\bar{\mathbb{A}}}$ has been obtained from \mathbb{A}' by edge collapses, foldings and edge slides it follows that \mathbb{B} is visible in a splitting obtained from \mathbb{A}' by foldings and edge slides as the collapses can be performed last. This concludes the proof of (1).

If \mathbb{B} is itself a JSJ-decomposition of L then the same argument that showed that \mathbb{A} can be obtained from \mathbb{A}' by foldings shows that \mathbb{B} can be obtained from \mathbb{A}' by foldings. This proves (2). \square

We conclude this section by defining the modular group of a one-ended Γ -limit group with respect to a given virtually Abelian splitting, using the following definition of natural extensions of vertex automorphisms and the subsequent definition of Dehn twists.

Definition 4.13. Let G be a group with a splitting $G = \pi_1(\mathbb{A}, v_0)$, and $v \in VA$. Assume that $\sigma_v \in \text{Aut}(A_v)$ is an automorphism such that for each $e \in EA$ with $\alpha(e) = v$ there exists an element $\gamma_e \in A_v$ such that $\sigma_v(a) = \gamma_e a \gamma_e^{-1}$ for all $a \in \alpha_e(A_e)$. Then the map

$$\phi : G \rightarrow G, [a_0, e_1, a_1, \dots, e_n, a_n] \mapsto [\bar{a}_0, e_1, \bar{a}_1, \dots, e_n, \bar{a}_n]$$

where

$$\bar{a}_k = \begin{cases} a_k & a_k \notin A_v \\ \gamma_{e_k}^{-1} \sigma_v(a_k) \gamma_{e_{k+1}} & a_k \in A_v \end{cases}$$

($\gamma_{e_0}^{-1} = \gamma_{e_{n+1}} = 1$) is a well-defined automorphism of G . We call it a natural extension of σ_v (with respect to the base vertex v_0), and say that σ_v is naturally extendable.

Note that a natural extension of a vertex automorphism σ_v is not unique as the γ_e are not uniquely determined by σ_v .

Definition 4.14. Let \mathbb{A} be a graph of groups, $e \in EA$ and $g \in A_{\alpha(e)}$ such that $g a g^{-1} = a$ for all $a \in \omega_e(A_e)$. Then we call the automorphism \mathfrak{f}_* of $\pi_1(\mathbb{A})$ induced by the morphism

$$\mathfrak{f} := (\text{id}_A, \{\psi_v = \text{id}_{A_v} \mid v \in VA\}, \{\psi_e = \text{id}_{A_e} \mid e \in EA\}, \{o_e \mid e \in EA\})$$

with $o_{e^{-1}}^{-1} = t_e = g$ and $o_f = 1$ for $f \neq e^{-1}$ the Dehn twist along e by g .

Note that in the case of 1-edge splittings this just recovers the usual Definition of a Dehn twist. In the case of an amalgamated product, i.e. a splitting with a single non-loop edge e , the automorphism $\mathfrak{f}_* : \pi_1(\mathbb{A}, \alpha(e)) \rightarrow \pi_1(\mathbb{A}, \alpha(e))$ is given by

$$\begin{aligned} \mathfrak{f}_*([a_0, e, a_1, e^{-1}, a_2, \dots, a_{2k-2}, e, a_{2k-1}, e^{-1}, a_{2k}]) \\ = [a_0, e, g a_1 g^{-1}, e^{-1}, a_2, \dots, a_{2k-2}, e, g a_{2k-1} g^{-1}, e^{-1}, a_{2k}]. \end{aligned}$$

Definition 4.15. Let \mathbb{A} be a virtually Abelian splitting of a one-ended group L . Then $\text{Mod}_{\mathbb{A}}(L) \leq \text{Aut}(L)$ is the group generated by the following automorphisms.

- (1) Inner automorphisms of L .

- (2) If $e \in EA$ with A_e finite-by-Abelian and \mathbb{A}^e is the 1-edge splitting obtained from \mathbb{A} by collapsing all edges but e : A Dehn twist along e by an elliptic (with respect to \mathbb{A}) element $g \in Z(M^+)$ where M is the maximal virtually Abelian subgroup containing A_e .
- (3) Natural extensions of automorphisms of QH-subgroups.
- (4) Natural extensions of automorphisms of maximal virtually Abelian vertex groups A_v which restrict to the identity on P_v^+ and to conjugation on each virtually Abelian subgroup $U \leq A_v$ with $U^+ = P_v^+$.

Remark. Note that the extendable automorphisms of the QH-subgroups all correspond to automorphisms of the corresponding orbifold. We also note that the above Dehn twists along e by g can be realized in the following way: If g is contained in A_v , we may slide the edge e (by a finite sequence of edge slides) to obtain a graph of groups \mathbb{A}' where $v = \alpha'(e)$. This induces an isomorphism φ from $\pi_1(\mathbb{A})$ to $\pi_1(\mathbb{A}')$ for the obtained graph of groups \mathbb{A}' . Now, consider \bar{f} as in Definition 4.14, inducing the Dehn twist along e by g in A' . Then $\varphi^{-1} \circ \bar{f}_* \circ \varphi$ is the desired map.

It turns out that if \mathbb{A} is a virtually Abelian JSJ-decomposition of L then $\text{Mod}_{\mathbb{A}}(L)$ contains all other modular groups.

Lemma 4.16. *Let \mathbb{A} be a virtually Abelian splitting of a one-ended Γ -limit group L and assume that \mathbb{A}' is obtained from \mathbb{A} by edge slides and boundary slides. Then $\text{Mod}_{\mathbb{A}}(L) = \text{Mod}_{\mathbb{A}'}(L)$.*

Proof. It is obvious that boundary slides preserve the modular group. So assume that \mathbb{A}' is obtained from \mathbb{A} by an edge slide of e_1 over e_2 , i.e. e_1 is replaced by e'_1 , using the notations of Definition 4.1. It is easy to see that all natural extensions of vertex automorphisms, as well as any Dehn twist along an edge distinct from e_2 are unaffected by the edge slide.

Now assume that α is a Dehn twist along e_2 by $g \in L$. Then, in \mathbb{A}' , α appears as the product of the Dehn twists by g along e_2 and by $\omega_{e'_1}^{-1} \circ \omega_{e_2}(g)$ along e'_1 . \square

Proposition 4.17. *Let L be a finitely generated one-ended Γ -limit group, \mathbb{A} be a virtually Abelian JSJ-decomposition of L and \mathbb{B} be a virtually Abelian splitting of L . Then*

$$\text{Mod}_{\mathbb{B}}(L) \leq \text{Mod}_{\mathbb{A}}(L).$$

Proof. We first deal with the case where \mathbb{B} is compatible. Let $\phi \in \text{Mod}_{\mathbb{B}}(L)$. While there is nothing to prove if ϕ is an inner automorphism of L , we need to show that the Dehn twists and the natural extensions of vertex automorphisms arising in $\text{Mod}_{\mathbb{B}}(L)$ are contained in $\text{Mod}_{\mathbb{A}}(L)$.

First, assume that ϕ is a Dehn twist along an edge $e \in EB$. By Theorem 4.12, the induced 1-edge splitting of L with edge e and edge group B_e is visible in \mathbb{A} after unfoldings, foldings and edge slides. As unfolding and folding the edge group does not alter the maximal virtually Abelian subgroup M that contains A_e , M^+ is also unchanged. Thus the presence of the Dehn twist is preserved by foldings and unfoldings unless an edge is folded such that the corresponding edge group ceases to be finite-by-Abelian. But since the edge groups of the JSJ-decomposition are finite-by-Abelian if possible they do exist in the JSJ if they exist in any splitting. Moreover, edge slides preserve the Dehn twist by Lemma 4.16.

Now assume that ϕ is a natural extension of an automorphism $\sigma \in \text{Aut}(B_v)$ for some $v \in VB$. If v is a QH-subgroup, there is nothing to show as these automorphisms lift to automorphisms of the QH-vertex group of \mathbb{A} containing the QH-vertex group of \mathbb{B} . Thus we can assume that B_v is virtually Abelian. We may assume that B_v is maximal virtually Abelian, as otherwise the maximal virtually Abelian subgroup B_v would be conjugate into another vertex group $B_{v'}$ as \mathbb{B} is assumed to be compatible and we may assume that ϕ is a natural extension of an automorphism of $B_{v'}$. It follows that there is a vertex $u \in VA$ such that $A_u = B_v$. We show that ϕ arises as a natural extension of σ in \mathbb{A} . It clearly suffices to show that $P_u^+ \subset P_v^+$.

It follows from Theorem 4.12 that after unfolding \mathbb{A} we get a graph of groups \mathbb{A}' such that there exists a morphism from \mathbb{A}' to \mathbb{B} . Denote the image of the vertex u in \mathbb{A}' by u' . The existence of this morphism implies that $P_{u'}^+ \subset P_v^+$, thus it suffices to show that $P_u^+ = P_{u'}^+$. Now both A_u^+/P_u^+ and $A_{u'}^+/P_{u'}^+$ are finitely generated free Abelian groups and as \mathbb{A} is a JSJ-decomposition it follows that $\text{rank } A_u^+/P_u^+ \geq \text{rank } A_{u'}^+/P_{u'}^+$. As the quotient map $\theta : A_u^+ \rightarrow A_{u'}^+/P_{u'}^+$ factors through A_u^+/P_u^+ , this implies that θ is an isomorphism as f.g. free Abelian groups are hopfian. Thus $P_u^+ \subset P_v^+$.

If \mathbb{B} is not compatible, we can use edge slides to assure that \mathbb{B} satisfies the conclusion of Proposition 4.4, by Lemma 4.16 the slides do not change $\text{Mod}_{\mathbb{B}}(L)$. We will further modify \mathbb{B} by performing the modifications discussed following the proof of Proposition 4.4 to produce a compatible splitting. This modification possibly increases but does not decrease $\text{Mod}_{\mathbb{B}}(L)$. The Dehn twists along the edges that are being collapsed now occur as natural extensions of automorphisms of virtually Abelian vertex groups. Thus the claim follows from the case of \mathbb{B} being compatible. \square

Corollary 4.18. *Let L be a finitely generated one-ended Γ -limit group and \mathbb{A}, \mathbb{A}' virtually Abelian JSJ-decompositions of L . Then $\text{Mod}_{\mathbb{A}}(L) = \text{Mod}_{\mathbb{A}'}(L)$. In particular we can define*

$$\text{Mod}(L) := \text{Mod}_{\mathbb{A}}(L)$$

where \mathbb{A} is an arbitrary virtually Abelian JSJ-decomposition of L .

5. Γ -factor sets of one-ended groups

The term of *factor sets* was coined in [6] in the context of free groups. In this chapter we define Γ -factor sets for an arbitrary group Γ , which are a natural generalization of factor sets. The goal of this chapter is the proof of Theorem 5.2, which states that if Γ is hyperbolic, any finitely generated one-ended group G admits a Γ -factor set. This is the main step in the construction of Makanin–Razborov diagrams.

5.1. Γ -Factor sets

Definition 5.1. Let G and Γ be groups and $H \leq \text{Aut}(G)$. A Γ -factor set of G relative to H is a finite set of proper quotient maps $\{q_i : G \rightarrow \Gamma_i\}$ such that for each non-injective homomorphism $q : G \rightarrow \Gamma$, there exists some $\alpha \in H$ such that $q \circ \alpha$ factors through some q_i .

The following is the main theorem of Section 5.

Theorem 5.2. *Let G be a f.g. one-ended group and Γ be a hyperbolic group. Then the following hold.*

- (1) *If G is not fully residually Γ then there exists a Γ -factor set of G relative to $\{\text{id}\}$.*
- (2) *If G is fully residually Γ then there exists a Γ -factor set of G relative to $\text{Mod}(G)$.*

The proof of the first part of Theorem 5.2 is trivial. Indeed, if G is not fully residually Γ , then there is a finite set $S \subset G \setminus \{1\}$ such that for any $\varphi \in \text{Hom}(G, \Gamma)$, $S \cap \ker \varphi \neq \emptyset$. Thus the set of quotient maps

$$\{q_s : G \rightarrow G/\langle\langle s \rangle\rangle \mid s \in S\}$$

is a Γ -factor set relative to $\{\text{id}\}$.

If G is fully residually Γ the argument is significantly more involved and it turns out to be crucial to allow for precomposition with modular automorphisms as it allows us to only consider short homomorphisms in the sense of Definition 5.3 below. Recall that by Lemma 2.3 G is a Γ -limit group, thus $\text{Mod}(G)$ is defined (cf. Corollary 4.18). For the remainder of Section 5, we fix a f.g. one-ended group G and a hyperbolic group Γ with fixed finite generating sets S_G and S_Γ respectively.

Definition 5.3. A homomorphism $\varphi : G \rightarrow \Gamma$ is called short relative to $H \leq \text{Aut}(G)$, if for every $\alpha \in H$ and $g \in \Gamma$,

$$|\varphi| \leq |c_g \circ \varphi \circ \alpha|$$

(where c_g denotes conjugation by g and $|\cdot|$ is as in Definition 2.12).

Remark 5.4. It follows that a homomorphism that is short relative to any subgroup of $\text{Aut}(G)$ is conjugacy short and therefore satisfies (2.4) of Theorem 2.11. Thus every convergent sequence of pairwise distinct short homomorphisms from G to Γ yields a non-trivial limit action of G on a real tree.

For the remainder of this chapter we will not always explicitly mention $\text{Mod}(G)$, i.e. *short* will always mean *short relative to $\text{Mod}(G)$* and Γ -*factor set* will mean Γ -*factor set relative to $\text{Mod}(G)$* . It will always be obvious that the constructed automorphisms are indeed modular automorphisms.

The proof of the second claim of Theorem 5.2 is by contradiction, i.e. we assume that G does not admit a Γ -factor set. This assumption implies the following.

Lemma 5.5. *Suppose that G is fully residually Γ and that G does not admit a Γ -factor set. Then there exists a stably injective convergent sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ of pairwise distinct non-injective short homomorphisms.*

Proof. For $i \in \mathbb{N}$, let $B_i \subset G$ be the ball of radius i with center 1 in G (with respect to the word metric). For each i , there is a non-injective $\varphi_i \in \text{Hom}(G, \Gamma)$ such that $B_i \cap \ker \varphi_i = \{1\}$, as otherwise the set of quotient maps

$$\{q_g : G \rightarrow G/\langle\langle g \rangle\rangle \mid g \in B_i \setminus \{1\}\}$$

would be a Γ -factor set of G . Moreover, since the definition of the factor set allows precomposition by a modular automorphism of G , each φ_i can be chosen to be short.

Clearly, the obtained sequence (φ_i) of short homomorphisms is stably injective. Since each φ_i is non-injective, it occurs only finitely many times in the sequence. Thus (φ_i) has a convergent subsequence of pairwise distinct non-injective homomorphisms, see Lemma 2.6. \square

In the remainder of this chapter we prove the following proposition, which yields a contradiction to the conclusion of Lemma 5.5 and therefore implies Theorem 5.2.

Proposition 5.6. *Let (φ_i) be a stably injective convergent sequence of pairwise distinct homomorphisms from G to Γ . Then the φ_i are eventually not short.*

Remark 5.7. If $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ is a stably injective sequence, we can associate to (φ_i) a sequence $(\widehat{\varphi}_i)$ where $\widehat{\varphi}_i = c_{g_i} \circ \varphi_i \circ \alpha_i$, for some $\alpha_i \in \text{Mod}(G)$ and $g_i \in \Gamma$, and $\widehat{\varphi}_i$ is short for each $i \in \mathbb{N}$. After passing to a subsequence we can again assume that $(\widehat{\varphi}_i)$ is stable. Proposition 5.6 implies that $Q := G/\overrightarrow{\ker}(\widehat{\varphi}_i)$ is a proper quotient of G . This is an instance of a *shortening quotient* discussed in Section 7.1.

Section 5.2 is entirely devoted to the proof of Proposition 5.6.

5.2. The shortening argument

Let (φ_i) be as in Proposition 5.6. By Theorem 2.6, (φ_i) converges to a non-trivial action of G on an \mathbb{R} -tree $T = (T, x_0)$.

Since $\varinjlim \ker(\varphi_i) = 1$ it follows that the action of $G = G/\varinjlim \ker(\varphi_i)$ on T satisfies the stability assertions of Theorem 2.17. This implies that T satisfies the assumptions of Theorem 3.4. As G is assumed to be one-ended and the stabilizers of unstable arcs are finite, G does not split over the stabilizer of an unstable arc. It follows that T splits as a graph of actions. We will use this decomposition of the action of G on T to show that for large i , the homomorphisms φ_i are not short, more precisely we will construct modular automorphisms $\alpha_i \in \text{Mod}(G)$ such that $|\varphi_i \circ \alpha_i| < |\varphi_i|$ for large i .

Let $\mathcal{G} = \mathcal{G}(\mathbb{A})$ be the graph of actions decomposition of T given by Theorem 3.4. We identify G with $\pi_1(\mathbb{A}, v_0)$ and assume that the basepoint x_0 is given by $x_0 = (\bar{x}_0, \tilde{v}_0)$, where $\tilde{v}_0 = \llbracket 1 \rrbracket$ is the base vertex of $\tilde{\mathbb{A}} = \overline{(\mathbb{A}, v_0)}$ and $\bar{x}_0 \in T_{v_0}$. Moreover, we denote the metric $d_{\mathcal{G}}$ of \mathcal{G} simply by d . By Remark 3.3 we can assume that the following hold for any $e, f \in EA$.

- (1) If $\alpha(e) = v_0$ and \bar{x}_0 is A_{v_0} -equivalent to p_e^α then $\bar{x}_0 = p_e^\alpha$.
- (2) If $\alpha(e) = \alpha(f)$ and p_e^α is $A_{\alpha(e)}$ -equivalent to p_f^α then $p_e^\alpha = p_f^\alpha$.

In the construction of the shortening automorphisms we need to deal with each of the different types of vertex trees of \mathcal{G} , we will do this in the following three sections before plugging it all together to conclude. The shortening argument first appeared in [35] but the underlying ideas are also implicit in the work of Razborov. In order to deal with torsion a number of additional issues need to be addressed.

5.2.1. Axial components

The purpose of this section is to prove the following.

Theorem 5.8. *Let $v_1 \in VA$ be an axial vertex. For any finite subset $S \subset G$ there exists some $\phi \in \text{Mod}(G)$ such that for any $g \in S$ the following hold.*

- If $[x_0, gx_0]$ has a nondegenerate intersection with a vertex space $T_{\tilde{v}}$ and $\downarrow \tilde{v} = v_1$, then $d(x_0, \phi(g)x_0) < d(x_0, gx_0)$,
- otherwise, $d(x_0, \phi(g)x_0) = d(x_0, gx_0)$.

The main step is to construct an automorphism of the axial vertex group that shortens the action on its vertex tree and that can be extended to an automorphism of G . We start by studying the algebraic structure of axial groups.

Let $G_A = A_{v_1}$ be a vertex group with an axial action on a tree $T_A = T_{v_1}$. We assume that the group G_A does not preserve the ends of T_A , i.e. that G_A contains elements that act by reflections. We denote the index 2 subgroup of G_A which preserves the ends by G_A^+ . By Theorem 2.17, G_A^+ is finite-by-Abelian. The case where G_A preserves the ends follows by considering only G_A^+ .

Let $E := \langle \{g \in G_A^+ \mid |g| < \infty\} \rangle$. E is normal in G_A of order at most $N(\Gamma)$ by Lemma 2.23. Put $H := G_A/E$ and denote by π the corresponding quotient map

$$\pi : G_A \rightarrow H = G_A/E.$$

As E is contained in the kernel of the action of G_A on T_A , the action factors through π and induces an action of H on T_A . Let $H^+ = \pi(G_A^+)$. As G_A^+ is finite-by-Abelian and any finite subgroup of G_A^+ lies in E , H^+ is Abelian. Moreover, H^+ is torsion-free as for any $g \in G_A^+$ s.t. $g^k \in E$ for some $k \geq 1$ we have $g^{k \cdot |E|} = 1$, and so $g \in E$ by construction.

Since the image of G_A^+ (and therefore also the image of H) in $\text{Isom}(T_A)$ is f.g. by Theorem 3.4, there is a decomposition

$$H^+ = A \oplus B$$

where A is f.g. free Abelian and B is the torsion-free Abelian kernel of the action of H^+ on T_A . Put $\tilde{A} = \pi^{-1}(A)$ and $\tilde{B} = \pi^{-1}(B)$.

Lemma 5.9. *H is the semi-direct product $\mathbb{Z}_2 \ltimes H^+$. The action of $\mathbb{Z}_2 = \langle s \mid s^2 \rangle$ on H^+ is given by $shs^{-1} = h^{-1}$ for all $h \in H^+$.*

Proof. Let s be an arbitrary element of $H \setminus H^+$ and let \tilde{s} be a lift of s to G_A , clearly $\tilde{s} \in G_A \setminus G_A^+$. It follows as in the proof of Lemma 2.23 that for large i the element $\varphi_i(\tilde{s})$ is of finite order as it either lies in a finite group or in a 2-ended group exchanging the ends. Thus \tilde{s} is of finite order, i.e. $\tilde{s}^2 \in E$. It follows that $s^2 = \pi(\tilde{s}^2) = 1$. This proves that $\mathbb{Z}_2 \ltimes H^+ = H$.

We show that the action is as desired. Let $h \in H^+$. Choose \tilde{h} such that $\pi(\tilde{h}) = h$. For large i the group $\varphi_i(\langle \tilde{h}, \tilde{s} \rangle)$ is 2-ended with $\varphi_i(\tilde{s})$ exchanging ends and \tilde{h} preserving ends. Here it is easily verified that $\varphi_i(\tilde{h}\tilde{s}\tilde{h}\tilde{s}^{-1})$ is of finite order for large i ; thus $\tilde{h}\tilde{s}\tilde{h}\tilde{s}^{-1} \in E$, i.e. $hshs^{-1} = 1$. It follows that $shs^{-1} = h^{-1}$ as desired. \square

As \tilde{B} is the kernel of the action of G_A on T_A it follows that \tilde{A} normalises \tilde{B} , i.e. \tilde{A} acts on \tilde{B} by conjugation. We can immediately see that the kernel of this action is of finite index in \tilde{A} :

Lemma 5.10. *Let $K_A^{\tilde{B}}$ be the kernel of the action of \tilde{A} on \tilde{B} . Then $|\tilde{A} : K_A^{\tilde{B}}| < \infty$.*

Proof. As G_A^+ is finite-by-Abelian, by Lemma 2.22 (2), $Z(G_A^+)$ is of finite index in G_A^+ . Therefore, $Z(G_A^+) \cap \tilde{A}$ is of finite index in \tilde{A} and clearly this group is contained in $K_A^{\tilde{B}}$. \square

Let now \tilde{s} be a lift of s to G_A and $\text{Aut}_{\tilde{s}}(G_A) \leq \text{Aut}(G_A)$ be the subgroup consisting of those automorphisms that restrict to the identity on $\langle \tilde{B}, \tilde{s} \rangle$ and preserve \tilde{A} . Let $\text{Aut}_{\tilde{s}}^*(G_A) \leq \text{Aut}_{\tilde{s}}(G_A)$ be the subgroup consisting of those automorphisms α that conjugate all point stabilizers pointwise. Thus for any $\alpha \in \text{Aut}_{\tilde{s}}^*(G_A)$ and $x \in T_A$ there exists some g_x such that $\alpha(g) = g_x g g_x^{-1}$ for all $g \in \text{stab } x$.

Lemma 5.11. $\text{Aut}_{\tilde{s}}^*(G_A)$ is a finite index subgroup of $\text{Aut}_{\tilde{s}}(G_A)$.

Proof. Note first that for any $x \in T_A$ either $\text{stab } x = \tilde{B}$ or $\text{stab } x = \langle \tilde{B}, w\tilde{s} \rangle$ for some $w \in \tilde{A}$. As all elements of $\text{Aut}_{\tilde{s}}(G_A)$ act trivially on \tilde{B} by definition we can ignore the first case.

Note further that there are only finitely many conjugacy classes of stabilizers of type $\langle \tilde{B}, w\tilde{s} \rangle$; this follows from the fact that if A is free Abelian of rank n then $\langle A, s \rangle$ has precisely 2^n conjugacy classes of reflections. As $\text{Aut}_{\tilde{s}}(G_A)$ acts on conjugacy classes of stabilizers this implies that there exists a finite index subgroup $\text{Aut}'_{\tilde{s}}(G_A)$ of $\text{Aut}_{\tilde{s}}(G_A)$ that preserves conjugacy classes of stabilizers.

For any point stabilizer $C = \langle \tilde{B}, w\tilde{s} \rangle$ let $\text{Aut}_{\tilde{s}}^C(G_A) \leq \text{Aut}'_{\tilde{s}}(G_A)$ be the subgroup consisting of those automorphisms α that conjugate C , i.e. for which $\alpha(c) = gcg^{-1}$ for all $c \in C$ and some fixed g . As $\alpha(hch^{-1}) = \alpha(h)\alpha(c)\alpha(h)^{-1}$ it follows that $\text{Aut}_{\tilde{s}}^C(G_A) = \text{Aut}_{\tilde{s}}^{hCh^{-1}}(G_A)$ for every conjugate hCh^{-1} of C .

To prove the claim of the lemma it suffices to show that for any such C the group $\text{Aut}_{\tilde{s}}^C(G_A)$ is of finite index in $\text{Aut}'_{\tilde{s}}(G_A)$. Indeed as there are only finitely many conjugacy classes and the intersection of finitely many subgroups of finite index is of finite index, this proves the claim.

Let now $C = \langle \tilde{B}, w\tilde{s} \rangle$. Suppose that there exists a sequence $(\alpha_i)_{i \in \mathbb{N}} \subset \text{Aut}'_{\tilde{s}}(G_A)$ such that $\alpha_i \text{Aut}_{\tilde{s}}^C(G_A) \neq \alpha_j \text{Aut}_{\tilde{s}}^C(G_A)$ for $i \neq j$. For each i choose $f_i \in \tilde{A}$ and $e_i \in E$ such that

$$\alpha_i(w\tilde{s}) = f_i w \tilde{s} e_i f_i^{-1}.$$

Such elements f_i and e_i exist as by assumption $\alpha_i(w) \in \tilde{A}$ and C is conjugate to $\alpha_i(C)$. After passing to a subsequence we can assume that $e_i = e$ for all $i \in \mathbb{N}$ and some fixed $e \in E$. Moreover, we may assume that $f_i b f_i^{-1} = f_j b f_j^{-1}$ for all $i, j \in \mathbb{N}$ and $b \in \tilde{B}$, this follows as the kernel of the action of \tilde{A} on \tilde{B} by conjugation is of finite index in \tilde{A} .

It follows that for all i, j we have

$$\alpha_j(w\tilde{s}) = (f_j f_i^{-1}) \alpha_i(w\tilde{s}) (f_j f_i^{-1})^{-1}$$

which implies that the restriction of $\alpha_j \circ \alpha_i^{-1}$ to $\alpha_i(C)$ is conjugation by $f_j f_i^{-1}$. As $\alpha_i(C)$ is conjugate to C this implies that $\alpha_j \circ \alpha_i^{-1} \in \text{Aut}_{\tilde{s}}^C(G_A)$, a contradiction. \square

Any $\alpha \in \text{Aut}_{\tilde{s}}(G_A)$ restricts to an automorphism of \tilde{A} and therefore induces an automorphism of $A = \tilde{A}/E$. Denote the subgroup of $\text{Aut}(A)$ induced in this fashion by $K_{\tilde{s}}$. Moreover let $K_{\tilde{s}}^*$ be the subgroup of $\text{Aut}(A)$ induced by $\text{Aut}_{\tilde{s}}^*(G_A)$

Lemma 5.12. *Let \tilde{s} be as above. Then $K_{\tilde{s}}$ is of finite index in $\text{Aut}(A)$.*

Proof. Suppose that A is free Abelian of rank n and let a_1, \dots, a_n be a basis of A . The proof is by contradiction thus we assume that $K_{\tilde{s}}$ is of infinite index in $\text{Aut}(A)$. Choose a sequence (α_i) of elements of $\text{Aut}(A)$ that represent pairwise distinct cosets of $K_{\tilde{s}}$, i.e. that $\alpha_i \circ \alpha_j^{-1} \notin K_{\tilde{s}}$ for all $i \neq j$. For each $i \in \mathbb{N}$ let $P_i = (x_1^i, \dots, x_n^i)$ where x_k^i is a lift of $\alpha_i(a_k) \in A$ to \tilde{A} for $1 \leq k \leq n$.

After passing to a subsequence we can assume that for all $i, j \in \mathbb{N}$ and $1 \leq k, l \leq n$ the following hold:

- (1) $[x_k^i, x_l^i] = [x_k^j, x_l^j]$.
- (2) The actions of x_k^i and x_k^j on \tilde{B} coincide.
- (3) $\tilde{s}x_k^i\tilde{s}^{-1}x_k^i = \tilde{s}x_k^j\tilde{s}^{-1}x_k^j$.

This however implies that for i, j the map $x_k^j \mapsto x_k^i$ for $1 \leq k \leq n$ extends to an automorphism $\alpha \in \text{Aut}(G_A, \langle \tilde{B}, \tilde{s} \rangle)$. Now this automorphism induces $\alpha_i \circ \alpha_j^{-1}$ on A contradicting our assumption that $\alpha_i \circ \alpha_j^{-1} \notin K_{\tilde{s}}$. \square

As an immediate consequence of Lemma 5.11 and Lemma 5.12 we get the following.

Corollary 5.13. *Let \tilde{s} be as above. Then $K_{\tilde{s}}^*$ is of finite index in $\text{Aut}(A)$.*

The following proposition is the main technical result of this section.

Proposition 5.14. *Let G_A be as above and $x, x_1, \dots, x_k \in T_A$. For each finite $S \subset G_A$ and $\epsilon > 0$, there exist elements $\gamma_1, \dots, \gamma_k \in G_A$ and an automorphism σ of G_A such that the following hold.*

- (1) For each $g \in S$,

$$d(x, \sigma(g)x) < \epsilon. \tag{5.1}$$
- (2) $\sigma(g) = \gamma_i g \gamma_i^{-1}$ for $1 \leq i \leq k$ and $g \in \text{stab } x_i$.
- (3) $d(x, \gamma_i x_i) < \epsilon$ ($i = 1, \dots, k$).

We can moreover assume that $\gamma_i = \gamma_j$ if $x_i = x_j$ and that $\gamma_i x_i = x_i$ if $x_i = x$.

Proof. Possibly after choosing a different reflection s we can assume that a lift \tilde{s} of s fixes a point $p_{\tilde{s}}$ such that $d(x, p_{\tilde{s}}) \leq \epsilon/4$.

Let a_1, \dots, a_n be a basis of the free Abelian group A ; recall that A acts on T_A by translations with dense orbits. The Euclidean algorithm guarantees the existence of a sequence $(\alpha_i) \subset \text{Aut}(A)$ such that the translation length of $\alpha_i(a_k)$ ($1 \leq k \leq n$) converges to 0 for $i \rightarrow \infty$. This implies that the translation lengths of $\alpha_i(a)$ converge to 0 for any $a \in A$. As $|\text{Aut}(A) : K_s^*| < \infty$ we can choose $(\alpha_i) \subset K_s^*$. For any $i \in \mathbb{N}$ let $\tilde{\alpha}_i$ be a lift of α_i to $\text{Aut}_{\tilde{s}}^*(G_A)$.

Now any element $g \in S$ can be written as $\tilde{a}_g \tilde{b}_g \tilde{s}^{\eta_g}$ where $\tilde{a}_g \in \tilde{A}$, $\tilde{b}_g \in \tilde{B}$ and $\eta_g \in \{0, 1\}$. As $\tilde{\alpha}_i \in \text{Aut}_{\tilde{s}}^*(G_A)$ it follows that $\tilde{\alpha}_i(g) = \tilde{\alpha}_i(\tilde{a}_g) \tilde{b}_g \tilde{s}^{\eta_g}$ for all $i \in \mathbb{N}$ and $g \in S$. Moreover as $\tilde{\alpha}_i(\tilde{a}_g)$ is a lift of $\alpha_i(\pi(\tilde{a}_g))$ it follows that the translation length of $\tilde{\alpha}_i(\tilde{a}_g)$ converges to 0. Thus we get

$$\lim_{i \rightarrow \infty} d(x, \tilde{\alpha}_i(g)x) = \lim_{i \rightarrow \infty} d(x, \tilde{\alpha}_i(\tilde{a}_g) \tilde{b}_g \tilde{s}^{\eta_g} x)$$

and

$$\begin{aligned} \lim_{i \rightarrow \infty} d(x, \tilde{s}^{\eta_g} x) + \lim_{i \rightarrow \infty} d(\tilde{s}^{\eta_g} x, \tilde{b}_g \tilde{s}^{\eta_g} x) + \lim_{i \rightarrow \infty} d(\tilde{b}_g \tilde{s}^{\eta_g} x, \tilde{\alpha}_i(\tilde{a}_g) \tilde{b}_g \tilde{s}^{\eta_g} x) \\ \leq \epsilon/2 + 0 + 0 = \epsilon/2 \end{aligned}$$

for all $g \in S$. This implies that for sufficiently large i assertions (1) and (2) are satisfied for $\sigma = \tilde{\alpha}_i$.

If $\text{stab } x_i = \tilde{B}$ then γ_i can be replaced by $\gamma_i h$ with $h \in K_A^{\tilde{B}}$ while preserving (2). As $K_A^{\tilde{B}}$ acts on T_A with dense orbits this ensures the existence of some γ_i such that both (2) and (3) are satisfied.

If $\text{stab } x_i$ is of type $\langle \tilde{B}, \tilde{a}\tilde{s} \rangle$ for some $\tilde{a} \in \tilde{A}$ then the fixed point of $\tilde{\alpha}_i(\text{stab } x_i) = \langle \tilde{B}, \tilde{\alpha}_i(\tilde{a})\tilde{s} \rangle$ converges to $p_{\tilde{s}}$ as the translation length of $\tilde{\alpha}_i(\tilde{a})$ converges to 0. As this fixed point equals $\gamma_i x_i$ and as $d(x, p_{\tilde{s}}) \leq \epsilon/4$ assertion (3) follows for large i .

Now assume that $x = x_i$ for some i . Then either $\text{stab } x = \tilde{B}$ or we can choose \tilde{s} such that $\text{stab } x = \langle \tilde{B}, \tilde{s} \rangle$. As in both cases the $\tilde{\alpha}_i$ restrict to the identity on $\text{stab } x = \text{stab } x_i$ we can choose $\gamma_i = 1$. Moreover, it is trivial that we can choose $\gamma_i = \gamma_j$ whenever $x_i = x_j$. The claim follows. \square

Proof of Theorem 5.8. For any $g \in S$ and \tilde{v} with $\downarrow \tilde{v} = v_1$ the intersection $T_{\tilde{v}} \cap [x_0, gx_0]$ is either empty or a (possibly degenerate) segment. If all such intersections are degenerate for all $g \in S$, there is nothing to show as the theorem holds for $\phi = \text{id}_G$. Thus we can assume that at least one such intersection is non-degenerate. Let $r > 0$ be the length of the shortest non-degenerate segment that occurs this way.

Recall from (3.4) that if q is a normal form \mathbb{A} -path

$$q = a_0, e_1, a_1, e_2, \dots, e_k, a_k \tag{5.2}$$

and $g = [q]$ then

$$d(x_0, gx_0) = d_{v_0}(\bar{x}_0, a_0 p_{e_1}^\alpha) + \sum_{j=1}^{k-1} d_{\omega(e_j)}(p_{e_j}^\omega, a_j p_{e_{j+1}}^\alpha) + d_{v_0}(p_{e_k}^\omega, a_k \bar{x}_0). \tag{5.3}$$




FIGURE 5.1. The path for an element $g = [a_0, e_1, \dots, e_3, a_3]$

Choose a point $p \in T_{v_1}$, if $v_0 = v_1$ then we choose $p := \bar{x}_0$. Let S_{v_1} be the set of elements of A_{v_1} that occur in the normal forms of the elements in S . By Proposition 5.14, there is an automorphism σ of A_{v_1} and for each $e \in EA$ with $\alpha(e) = v_1$ an element $\gamma_e \in A_{v_1}$ such that the following hold.

- $d_{v_1}(p, \sigma(a)p) < \frac{r}{6}$ for all $a \in S_{v_1}$.
- For $e \in EA$ with $\alpha(e) = v_1$, the restriction of σ to $\alpha_e(A_e) \leq \text{stab } p_e^\alpha$ is conjugation by γ_e , and

$$d_{v_1}(p, \gamma_e p_e^\alpha) < \frac{r}{6}. \tag{5.4}$$

- $\gamma_e = \gamma_f$ if $p_e^\alpha = p_f^\alpha$ and $\gamma_e p_e^\alpha = p_e^\alpha$ if $p_e^\alpha = p$.

Fix such an automorphism σ and let $\phi \in \text{Aut}(G)$ be a natural extension of σ (cf. Definition 4.13). Thus if $g = [a_0, e_1, \dots, e_k, a_k]$ as before we get $\phi(g) = [\bar{a}_0, e_1, \bar{a}_1, e_2, \dots, e_k, \bar{a}_k]$ where $\bar{a}_i = \gamma_{e_i}^{-1} \sigma(a_i) \gamma_{e_{i+1}}$ if $a_i \in A_{v_1}$ (and $\gamma_{e_0}^{-1} = \gamma_{e_{k+1}} = 1$) and $\bar{a}_i = a_i$ otherwise. In particular we have

$$d(x_0, \phi(g)x_0) = d_{v_0}(\bar{x}_0, \bar{a}_0 p_{e_1}^\alpha) + \sum_{j=1}^{k-1} d_{\omega(e_j)}(p_{e_j}^\omega, \bar{a}_j p_{e_{j+1}}^\alpha) + d_{v_0}(p_{e_k}^\omega, \bar{a}_k \bar{x}_0). \tag{5.5}$$

In the following we compare the summands occurring in (5.3) to those occurring in (5.5). If $a_i = \bar{a}_i$ the corresponding summands clearly coincide. Thus we can assume that $a_i \neq \bar{a}_i \in A_{v_1}$. We distinguish two cases.

Case 1. If $i \in \{1, \dots, k-1\}$ then we get

$$\begin{aligned} d_{\omega(e_i)}(p_{e_i}^\omega, \bar{a}_i p_{e_{i+1}}^\alpha) &= d_{\omega(e_i)}(p_{e_i}^\omega, \gamma_{e_i}^{-1} \sigma(a_i) \gamma_{e_{i+1}} p_{e_{i+1}}^\alpha) = d_{\omega(e_i)}(\gamma_{e_i}^{-1} p_{e_i}^\omega, \sigma(a_i) \gamma_{e_{i+1}} p_{e_{i+1}}^\alpha) \\ &\leq d_{\omega(e_i)}(\gamma_{e_i}^{-1} p_{e_i}^\omega, p) + d_{\omega(e_i)}(p, \sigma(a_i) p) + d_{\omega(e_i)}(\sigma(a_i) p, \sigma(a_i) \gamma_{e_{i+1}} p_{e_{i+1}}^\alpha) \\ &\leq \frac{r}{6} + \frac{r}{6} + d_{\omega(e_i)}(p, \gamma_{e_{i+1}} p_{e_{i+1}}^\alpha) \leq \frac{2r}{6} + \frac{r}{6} = \frac{r}{2}. \end{aligned}$$

If moreover $d_{\omega(e_i)}(p_{e_i}^\omega, a_i p_{e_{i+1}}^\alpha) = 0$ then $p_{e_i}^\omega$ and $p_{e_{i+1}}^\alpha$ are A_{v_1} -equivalent and therefore $p_{e_i}^\omega = p_{e_{i+1}}^\alpha$ by assumption, thus $a_i \in \text{stab } p_{e_i}^\omega$ and $\gamma_{e_i}^{-1} = \gamma_{e_{i+1}}$, by assumption this implies that $\sigma(a_i) = \gamma_{e_{i+1}} a_i \gamma_{e_{i+1}}^{-1}$. The above computation therefore implies that

$$d_{\omega(e_i)}(p_{e_i}^\omega, \bar{a}_i p_{e_{i+1}}^\alpha) = d_{\omega(e_i)}(p_{e_i}^\omega, p_{e_{i+1}}^\alpha) = 0.$$

Case 2. If $i = 0$ (the case $i = k$ is analogous) then we get

$$\begin{aligned} d_{v_0}(\bar{x}_0, \bar{a}_0 p_{e_1}^\alpha) = d_{v_0}(p, a_0 \gamma_{e_1} p_{e_1}^\alpha) &\leq d_{v_0}(p, a_0 p) + d_{v_0}(a_0 p, a_0 \gamma_{e_1} p_{e_1}^\alpha) \\ &\leq \frac{r}{6} + d_{v_1}(p, \gamma_{e_1} p_{e_1}^\alpha) \leq \frac{r}{6} + \frac{r}{6} = \frac{r}{3}. \end{aligned}$$

If moreover $d_{v_1}(p, a_0 p_{e_1}^\alpha) = 0$ then p and $p_{e_1}^\alpha$ are A_{v_1} equivalent and therefore $p = p_{e_1}^\alpha$ by assumption. Thus $a_0 \in \text{stab } p$ and therefore $\gamma_{e_1} \in \text{stab } p$. Thus

$$d_{v_1}(p, \bar{a}_0 p_{e_1}^\alpha) = d_{v_1}(p, a_0 \gamma_{e_1} p) = d_{v_1}(p, p) = 0.$$




FIGURE 5.2. The path for $\phi(g)$ if precisely \tilde{v}_2 is of type v_1

Comparing the summands in (5.3) to those occurring in (5.5) shows that each summand is preserved unless it corresponds to a non-degenerate intersection with some vertex tree $T_{\tilde{v}}$ with $\downarrow\tilde{v} = v_1$ (of length at least r) in which case it is replaced by at most $\frac{r}{2}$. This proves the claim. \square

5.2.2. Orbifold components

Analogous to Theorem 5.8, we prove the existence of a shortening automorphism provided there exists a vertex of orbifold type.

Theorem 5.15. *Let $v_1 \in VA$ be an orbifold type vertex. For any finite subset $S \subset G$ there exists some $\phi \in \text{Mod}(G)$ such that the following hold for any $g \in S$.*

- *If $[x_0, gx_0]$ has a nondegenerate intersection with a vertex space $T_{\tilde{v}}$ and $\downarrow\tilde{v} = v_1$, then $d(x_0, \phi(g)x_0) < d(x_0, gx_0)$,*
- *otherwise, $d(x_0, \phi(g)x_0) = d(x_0, gx_0)$.*

The proof of Theorem 5.15 follows from the following proposition in exactly the same way as Theorem 5.8 follows from Proposition 5.14.

Proposition 5.16. *Let $v \in VA$ be an orbifold type vertex and $x, x_1, \dots, x_k \in T_v$. For each finite $S \subset A_v$ and $\epsilon > 0$, there exist elements $\gamma_1, \dots, \gamma_k \in A_v$ and an automorphism $\sigma \in \text{Aut}(A_v)$ such that the following hold.*

- (1) *for each $g \in S$,*

$$d(x, \sigma(g)x) < \epsilon. \tag{5.6}$$
- (2) *$\sigma(g) = \gamma_i g \gamma_i^{-1}$ for $1 \leq i \leq k$ and $g \in \text{stab } x_i$.*
- (3) *$d(x, \gamma_i x_i) < \epsilon$ ($i = 1, \dots, k$).*

We can moreover assume that $\gamma_i = \gamma_j$ if $x_i = x_j$ and that $\gamma_i x_i = x_i$ if $x_i = x$.

The remainder of this section is dedicated to the proof of Proposition 5.16. The argument in this case is essentially due to Rips and Sela [35] who give a proof of Proposition 5.16 in the case where the action of A_v on T_v has trivial kernel. Thus we only need to address the case where this kernel is non-trivial.

If \mathcal{H} is a family of subgroups of G then we will denote by $\text{Aut}_{\mathcal{H}}(G)$ the subgroup of $\text{Aut}(G)$ consisting of those automorphisms that act on each $H \in \mathcal{H}$ by conjugation with an element of G .

Lemma 5.17. *Let G be a f.p. group, $\mathcal{H} := \{H_1, \dots, H_k\}$ a finite collection of cyclic, malnormal subgroups of G . Suppose that \tilde{G} is an extension of some finite group E by G , i.e. that there is the short exact sequence*

$$1 \longrightarrow E \longrightarrow \tilde{G} \xrightarrow{\pi} G \longrightarrow 1.$$

Put $\tilde{\mathcal{H}} := \{\tilde{H}_i := \pi^{-1}(H_i) \mid 1 \leq i \leq k\}$. Let S be the group of those automorphisms $\sigma \in \text{Aut}_{\mathcal{H}}(G)$ that lift to $\text{Aut}_{\tilde{\mathcal{H}}}(\tilde{G})$, i.e. for which there exists $\tilde{\sigma} \in \text{Aut}_{\tilde{\mathcal{H}}}(\tilde{G})$ such that $\pi \circ \tilde{\sigma} = \sigma \circ \pi$.

Then S has finite index in $\text{Aut}_{\mathcal{H}}(G)$.

Proof. The proof is in two steps. We first prove that the subgroup S_1 of $\text{Aut}_{\mathcal{H}}(G)$ consisting of automorphisms that lift to automorphisms of \tilde{G} is of finite index in $\text{Aut}_{\mathcal{H}}(G)$. We then show that S is a finite index subgroup of S_1 .

Assume that $|\text{Aut}_{\mathcal{H}}(G) : S_1| = \infty$. Then there exists a sequence $(\alpha_i) \subset \text{Aut}(G)$ of automorphisms such that $\alpha_i S_1 \neq \alpha_j S_1$ for $i \neq j$. Let $\langle S_E \mid R_E \rangle$ be a presentation of E and

$$\langle s_1, \dots, s_m \mid r_1, \dots, r_g \rangle$$

be a finite presentation of G . Every automorphism α_i of G gives rise to a (not unique) presentation of \tilde{G} as

$$\mathcal{P}_i(\tilde{G}) = \langle S_E, s_1, \dots, s_m \mid R_E, r_1 e^1, \dots, r_g e^g, \{s_i e s_i^{-1} = f_{i,e}\} \rangle,$$

where the e^i and $f_{i,e}$ lie in E and the generators s_i correspond to chosen lifts of the images of the generators of G under α_i . Since there are only finitely many such presentations, there are $i, j \in \mathbb{N}$ s.t. $i \neq j$ and $\mathcal{P}_i(\tilde{G}) = \mathcal{P}_j(\tilde{G})$. It follows that $\alpha_i^{-1} \alpha_j \in S_1$ and therefore $\alpha_i S_1 = \alpha_j S_1$, contradicting the above assumption. Thus $|\text{Aut}_{\mathcal{H}}(G) : S_1| < \infty$.

Let now S_2 be the subgroup of $\text{Aut}(\tilde{G})$ consisting of all lifts of automorphisms of S_1 . It clearly suffices to show that $S_2 \cap \text{Aut}_{\tilde{\mathcal{H}}}(\tilde{G})$ is of finite index in S_2 .

For any $\tilde{\alpha} \in S_2$ and $i = 1, \dots, k$ we have $\alpha(\tilde{H}_i) = c_i \tilde{H}_i c_i^{-1}$ for some $c_i \in \tilde{G}$ (but the conjugation is not pointwise in general). Since H_i is malnormal, c_i is unique up to elements of \tilde{H}_i . It follows that α induces a well-defined outer automorphism $\sigma_i(\alpha)$ of \tilde{H}_i represented by the automorphism

$$g \mapsto c_i^{-1} \alpha(g) c_i \quad \text{for all } g \in \tilde{H}_i.$$

Now $S_2 \cap \text{Aut}_{\tilde{\mathcal{H}}}(\tilde{G})$ is the kernel of the homomorphism

$$S_2 \rightarrow \prod_{i=1}^k \text{Out}(\tilde{H}_i), \alpha \mapsto (\sigma_1(\alpha), \dots, \sigma_k(\alpha)).$$

This clearly proves the assertion as $\text{Out}(H_i)$ is finite for all i . □

We can now proceed with the proof of Proposition 5.16.

Proof of Proposition 5.16. Denote by E the kernel of the action of A_v on T_v . Since T_v is not a line, E stabilizes a tripod, and therefore is finite by Theorem 2.17. Denote by $\pi : A_v \rightarrow P := A_v/E$ the quotient map. Then P acts faithfully on T_v , and by Theorem 3.4, P is the fundamental group of a compact 2-orbifold Σ with boundary. In particular, P is finitely presented.

Now let x be an arbitrary point of T_v . Applying Proposition 5.2 of [35], we get an infinite sequence (α_i) of automorphisms of P satisfying the following.

- For each $g \in S$, $\lim_{i \rightarrow \infty} d(x, \alpha_i(\pi(g))x) = 0$.
- For each (infinite cyclic) peripheral subgroup Z of P the restriction $\alpha_i|_Z$ is conjugation with an element $c_i^{(e)} \in P$.
- For any peripheral subgroup Z of P the distance between x and the (unique) fixed point of $\alpha_i(Z)$ tends to 0.

Let $F \subset EA$ be the set of edges whose initial vertex is v . For each $e \in F$, the image Z_e of $\alpha_e(A_e)$ under π in $P = \pi(A_v) = \pi_1(\Sigma)$ correspond to a loop in Σ that is homotopic to a boundary component. This implies that $\pi(\alpha_e(A_e))$ is malnormal in P .

Define $S \leq \text{Aut}_{\mathcal{H}}(P)$ as in Lemma 5.17 with respect to the collection of subgroups $\mathcal{H} := \{Z_e \mid e \in F\}$ of P and $\tilde{\mathcal{H}} := \{\tilde{Z}_e := \pi^{-1}(Z_e) \mid e \in F\}$.

Then by Lemma 5.17, $|\text{Aut}_{\mathcal{H}}(P) : S| < \infty$. It follows that there is a subsequence $(\alpha_{i_j}) \subset (\alpha_i)$ such that all α_{i_j} are in the same left coset C of S . Fix a representative $\gamma \in C$. Then the sequence (α'_i) given by

$$\alpha'_j := \gamma^{-1} \alpha_{i_j}$$

is in S and $\lim_{i \rightarrow \infty} |\alpha'_i| = 0$. Choosing i large enough and extending α'_i to A_v , gives the desired automorphism. \square

5.2.3. Simplicial components

Assume now that \mathcal{G} has a nondegenerate simplicial vertex tree. Thus \mathcal{G} can be refined in a simplicial type vertex yielding a (refined) graph of actions with non-zero length function l such that all vertices that are adjacent to edges of non-zero length have degenerate vertex trees. We denote this graph of actions again by \mathcal{G} . We can still assume that the base point x_0 is contained in a vertex tree $\tilde{v}_0 = \llbracket 1 \rrbracket$, i.e. that $x_0 = [\bar{x}_0, \tilde{v}_0]$. Indeed, if x_0 is contained in the interior of an edge segment $T_{\tilde{e}}$, we can split the corresponding edge $\downarrow \tilde{e} \in EA$ by introducing a valence 2 vertex with vertex group $A_{\downarrow \tilde{e}}$ and degenerate vertex tree such that x_0 is precisely a lift of this vertex tree.

We will construct a Dehn twist automorphism on an edge with non-zero length such that powers of this Dehn twist shorten the action of G on X induced by φ_i (for large enough i). Other than in the axial and orbifold cases these automorphisms do not shorten the action on the limit tree.

The following proposition is the key observation needed for the construction of the shortening automorphisms. In the following $d_i := \frac{d_X}{|\varphi_i|}$ is the scaled metric on the Cayley graph X . Thus

$$\lim_{i \rightarrow \infty} d_i(\varphi_i(g), \varphi_i(h)) = d_\varphi(g, h)$$

for all $g, h \in G$. Recall that (X, d_i) is δ_i -hyperbolic with $\lim \delta_i = 0$. Also recall from (3.2) that for any $c \in A_e$ and lift \tilde{e} of e the element $\theta_{\tilde{e}}(c)$ is a natural lift of c to the stabilizer of \tilde{e} .

Proposition 5.18. *Let $e \in EA$ be an edge with positive length and $c \in Z(A_e)$ of infinite order. There exists a sequence $(m_i) \subset \mathbb{Z}$ such that for any lift \tilde{e} of e^ε , $\varepsilon \in \{-1, 1\}$, the following holds.*

If $(y_i), (z_i) \subset X$ are approximating sequences of $T_{\alpha(\tilde{e})}$ and $T_{\omega(\tilde{e})}$ (which are single points) respectively, then

$$\lim_{i \rightarrow \infty} d_i(y_i, \varphi_i(\theta_{\tilde{e}}(c)^{\varepsilon \cdot m_i})z_i) = 0. \quad (5.7)$$

Proof. We assume that $\varepsilon = 1$, the case where $\varepsilon = -1$ is an immediate consequence due to the equivariance of the action. Fix some lift \tilde{e} of e and put $c_{\tilde{e}} := \theta_{\tilde{e}}(c)$. For large i the element $\varphi_i(c_{\tilde{e}})$ is hyperbolic and we define A_i to be the axis of $\varphi_i(c_{\tilde{e}})$ in X , i.e. the union of all geodesics joining the ends fixed by $\varphi_i(c_{\tilde{e}})$. A_i is easily seen to be in the $4\delta_i$ -neighbourhood of any of these geodesics with respect to the metric d_i .

For each i let y'_i and z'_i be points on A_i closest to y_i and z_i respectively. It is clear that $\lim_{i \rightarrow \infty} d_i(y_i, y'_i) = 0$ and $\lim_{i \rightarrow \infty} d_i(z_i, z'_i) = 0$ as $c_{\tilde{e}}$ fixes $T_{\tilde{v}_1}$ and $T_{\tilde{v}_2}$ in the limit action.

Moreover, there are integers m_i such that

$$d(y'_i, \varphi_i(c_{\tilde{e}})^{m_i} z'_i) \leq l(\varphi_i(c_{\tilde{e}})) + 8\delta$$

where $l(\varphi_i(c_{\tilde{e}}))$ denotes the translation length of $\varphi_i(c_{\tilde{e}})$ (note that we used the non-scaled metric d here).

As $d(y'_i, \varphi_i(c_{\tilde{e}})^{m_i} z'_i)$ is globally bounded from above, it follows that

$$\lim_{i \rightarrow \infty} d_i(y'_i, \varphi_i(c_{\tilde{e}})^{m_i} z'_i) = 0$$

and therefore also $\lim_{i \rightarrow \infty} d_i(y_i, \varphi_i(c_{\tilde{e}})^{m_i} z_i) = 0$.

To conclude, it suffices to show that the choice of the m_i does not depend on the choice of the lift of e . Indeed this follows immediately from the fact that if $\tilde{e}' = h\tilde{e}$ is another



FIGURE 5.3.

lift of e then $(\varphi_i(h)y_i), (\varphi_i(h)z_i) \subset X$ are approximating sequences of $T_{\alpha(\bar{e})} = hT_{\alpha(\bar{e})}$ and $T_{\omega(\bar{e})} = hT_{\omega(\bar{e})}$ respectively, and $c_{\bar{e}} = hc_{\bar{e}}h^{-1}$. \square

From now on let $e \in EA$ be a fixed edge with positive length. Let $c \in Z(A_e)$ be of infinite order; the existence of such a c follows immediately from Lemma 2.22(2) as edge groups of edges with positive length are finite-by-Abelian. Let $(m_i)_{i \in \mathbb{N}}$ as in Proposition 5.18. Define $\sigma : G \rightarrow G$ to be the Dehn twist automorphism along e by $\omega_e(c)$; it is easily verified that σ is given by

$$[a_0, e_1, a_1, \dots, e_n, a_n] \mapsto [\bar{a}_0, e_1, \bar{a}_1, \dots, e_n, \bar{a}_n] \tag{5.8}$$

where

$$\bar{a}_k = \begin{cases} \omega_{e_k}(c^{-1})a_k & e_k = e, \\ \omega_{e_k}(c)a_k & e_k = e^{-1} \\ a_k & e_k \neq e^{\pm 1}. \end{cases}$$

Proposition 5.19. *Let $g = [q] \in G$. If q is reduced and contains an edge $e^{\pm 1}$, then*

$$d_i(1, \varphi_i \circ \sigma^{m_i}(g)) < d_i(1, \varphi_i(g))$$

for sufficiently large i . Otherwise, $d_i(1, \varphi_i \circ \sigma^{m_i}(g)) = d_i(1, \varphi_i(g))$.

As the tree $T_{\mathcal{G}}$ is minimal it follows that for any generating set S of G the normal form of at least one element of S contains an edge $e^{\pm 1}$. Thus we obtain the following immediate corollary.

Corollary 5.20. $|\varphi_i \circ \sigma^{m_i}| < |\varphi_i|$ for sufficiently large i .

Proof of Proposition 5.19. Let $q = a_0, e_1, a_1, e_2, \dots, a_n$ be a normal form \mathbb{A} -path s.th. $g = [q]$, and $\tilde{e}_1, \dots, \tilde{e}_n$ be the reduced edge path in $\tilde{\mathbb{A}}$ from \tilde{v}_0 to $g\tilde{v}_0$. Then

$$d_{\mathcal{G}}(x_0, gx_0) = d_{\tilde{v}_0}(\bar{x}_0, p_{\tilde{e}_1}^\alpha) + \sum_{k=1}^{n-1} d_{\omega(\tilde{e}_k)}(p_{\tilde{e}_k}^\omega, p_{\tilde{e}_{k+1}}^\alpha) + \sum_{k=1}^n d_{\tilde{e}_k}(p_{\tilde{e}_k}^\alpha, p_{\tilde{e}_k}^\omega) + d_{g\tilde{v}_0}(p_{\tilde{e}_n}^\omega, g\bar{x}_0).$$




FIGURE 5.4. The segment $[x_0, gx_0] \subset T$ with $g = [a_0, e_1, \dots, e_3, a_3]$

Now for each k , let $(p_{k,i}^\alpha)_{i \in \mathbb{N}}$ and $(p_{k,i}^\omega)_{i \in \mathbb{N}}$ be approximating sequences of $p_{e_k}^\alpha$ and $p_{e_k}^\omega$ respectively. Recall that (1) and $(\varphi_i(g))$ are approximating sequences of x_0 and gx_0 respectively. Thus by Lemma 2.10, for any $\epsilon > 0$ and large enough i , we have

$$d_X(1, \varphi_i(g)) \geq d_i(1, p_{1,i}^\alpha) + \sum_{k=1}^{n-1} d_i(p_{k,i}^\omega, p_{k+1,i}^\alpha) + \sum_{k=1}^n d_i(p_{k,i}^\alpha, p_{k,i}^\omega) + d_i(p_{n,i}^\omega, \varphi_i(g)) - \epsilon. \quad (5.9)$$




FIGURE 5.5. The segment $[1, \varphi_i(g)] \subset X$ with $g = [a_0, e_1, \dots, e_3, a_3]$

For each $i \in \mathbb{N}$ we put $\bar{a}_0^i := a_0$ and for $k \in \{1, \dots, n\}$ put

$$\bar{a}_k^i = \begin{cases} \omega_{e_k}(c^{-m_i})a_k & \text{if } e_k = e \\ \omega_{e_k}(c^{m_i})a_k & \text{if } e_k = e^{-1} \\ a_k & \text{if } e_k \neq e^{\pm 1}. \end{cases}$$

Moreover for $0 \leq k \leq n$ and $i \in \mathbb{N}$ we define

$$q_k := a_0, e_1, a_1, \dots, a_k, \\ \bar{q}_k^i := \bar{a}_0^i, e_1, \bar{a}_1^i, \dots, \bar{a}_k^i.$$

This implies that

$$[\bar{q}_n^i] = [\bar{a}_0^i, e_1, \bar{a}_1^i, \dots, e_n, \bar{a}_n^i] = \sigma^{m_i}(g).$$

Further, for each k and i , put

$$\bar{p}_{k,i}^\alpha := \varphi_i([\bar{q}_{k-1}^i q_{k-1}^{-1}]) p_{k,i}^\alpha$$

and

$$\bar{p}_{k,i}^\omega := \varphi_i([\bar{q}_k^i q_k^{-1}]) p_{k,i}^\omega.$$

Note that

$$\varphi_i([\bar{q}_n^i q_n^{-1}]) \cdot \varphi_i(g) = \varphi_i(\sigma^{m_i}(g) \cdot g^{-1}) \cdot \varphi_i(g) = \varphi_i \circ \sigma_i^{m_i}(g).$$

Using the triangle inequality, this implies that for large i we get

$$d_i(1, \varphi_i \circ \sigma^{m_i}(g)) \leq d_i(1, \bar{p}_{1,i}^\alpha) + \sum_{k=1}^{n-1} d_i(\bar{p}_{k,i}^\omega, \bar{p}_{k+1,i}^\alpha) + \sum_{k=1}^n d_i(\bar{p}_{k,i}^\alpha, \bar{p}_{k,i}^\omega) + d_i(\bar{p}_{n,i}^\omega, \varphi_i \circ \sigma^{m_i}(g)) + \epsilon.$$




FIGURE 5.6. The segment $[1, \varphi_i \circ \sigma^{m_i}(g)] \subset X$ with $g = [a_0, e_1, \dots, e_e, a_3]$

The G -equivariance of the metric d_i immediately implies

- (1) $d_i(1, p_{1,i}^\alpha) = d_i(1, \bar{p}_{1,i}^\alpha)$,
- (2) $d_i(p_{n,i}^\omega, \varphi_i(g)) = d_i(\bar{p}_{n,i}^\omega, \varphi_i \circ \sigma^{m_i}(g))$,
- (3) $d_i(\bar{p}_{k,i}^\omega, \bar{p}_{k+1,i}^\alpha) = d_i(p_{k,i}^\omega, p_{k+1,i}^\alpha)$ for any k ,
- (4) $d_i(\bar{p}_{k,i}^\alpha, \bar{p}_{k,i}^\omega) = d_i(p_{k,i}^\alpha, p_{k,i}^\omega)$ whenever $\downarrow \tilde{e}_k \neq e^{\pm 1}$.

Assume that for some k , $\downarrow \tilde{e}_k = e^\varepsilon$ with $\varepsilon \in \{-1, 1\}$. Then $\theta_{\tilde{e}_k}(c)$ can be written as

$$\theta_{\tilde{e}_k}(c) = a_0, e_1, \dots, a_{k-1}, e_k, \omega_{e_k}(c), e_k^{-1}, a_{k-1}^{-1}, \dots, e_1^{-1}, a_0^{-1},$$

and it is easy to verify that $\bar{q}_k^i q_k^{-1} \sim \bar{q}_{k-1}^i q_{k-1}^{-1} \theta_{\bar{e}_k}(c)^{\varepsilon \cdot m_i}$. Therefore by Proposition 5.18

$$\begin{aligned} \lim_{i \rightarrow \infty} d_i(\bar{p}_{k,i}^\alpha, \bar{p}_{k,i}^\omega) &= \lim_{i \rightarrow \infty} d_i(\varphi_i([\bar{q}_{k-1}^i q_{k-1}^{-1}]) p_{k,i}^\alpha, \varphi_i([\bar{q}_k^i q_k^{-1}]) p_{k,i}^\omega) \\ &= \lim_{i \rightarrow \infty} d_i(\varphi_i([\bar{q}_{k-1}^i q_{k-1}^{-1}]) p_{k,i}^\alpha, \varphi_i(\bar{q}_{k-1}^i q_{k-1}^{-1} \theta_{\bar{e}_k}(c)^{\varepsilon \cdot m_i}) p_{k,i}^\omega) \\ &= \lim_{i \rightarrow \infty} d_i(p_{k,i}^\alpha, \varphi_i(\theta_{\bar{e}_k}(c)^{\varepsilon \cdot m_i}) p_{k,i}^\omega) \\ &= 0 \end{aligned}$$

Comparing this to (5.9) we obtain

$$\lim_{i \rightarrow \infty} (d_i(1, \varphi_i \circ \sigma^{m_i}(g))) = d(x_0, g x_0) - s \cdot l(e). \quad \square$$

5.2.4. The shortening automorphism

In view of the previous sections, we are now able to conclude the proof of Proposition 5.6. Let $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ a converging stable sequence of pairwise distinct homomorphisms with associated Γ -limit map φ . Assume that all φ_i are short (with respect to fixed finite generating sets of G and Γ). Then, by Theorems 2.6 and 2.11 we obtain a non-trivial limit G -tree T , which splits as a graph of actions \mathcal{G} by Theorem 3.4.

If \mathcal{G} contains an axial vertex space or an orbifold type vertex space it follows from Theorem 5.8, respectively Theorem 5.15, that we can shorten the action on T by precomposing with an automorphism $\alpha \in \text{Aut}(G)$. As this action is being approximated by the action on X via the φ_i it follows that for large i these actions can be shortened likewise by precomposing with α . This proves that for large i the φ_i are not short, which is the claim of the theorem.

In the remaining case \mathcal{G} contains a simplicial vertex space and the Proposition follows immediately from Corollary 5.20.

6. Makanin–Razborov diagrams

In this chapter we use the existence of Γ -factor sets proven in the previous chapter to give a complete description of the set of all homomorphisms from a finitely generated group G to an equationally Noetherian hyperbolic group Γ . In Section 7.2 we will see that this assumption was vacuous as all hyperbolic groups are equationally Noetherian.

6.1. Equationally Noetherian groups

Let Γ be a group and $F(x_1, \dots, x_n)$ be a free group of rank n . We then define

$$\Gamma[x_1, \dots, x_n] := \Gamma * F(x_1, \dots, x_n).$$

For any $\eta \in \Gamma[x_1, \dots, x_n]$ and $(\gamma_1, \dots, \gamma_n) \in \Gamma^n$ we define $\eta(\gamma_1, \dots, \gamma_n)$ to be the element of Γ obtained from η by substituting any occurrence of x_i by γ_i . We say that $(\gamma_1, \dots, \gamma_n)$ satisfies the equation η if $\eta(\gamma_1, \dots, \gamma_n) = 1$.

For any set $S \subset \Gamma[x_1, \dots, x_n]$ the *variety* of S is defined as

$$\text{rad}(S) := \{(\gamma_1, \dots, \gamma_n) \in \Gamma^n \mid \eta(\gamma_1, \dots, \gamma_n) = 1 \text{ for all } \eta \in S\},$$

thus $\text{rad}(S)$ is the set of all n -tuples of elements of Γ that satisfy all equations of S simultaneously.

Γ is called *equationally Noetherian* if for every $n \in \mathbb{N}$ and any subset $S \subset \Gamma[x_1, \dots, x_n]$ there exists a finite subset $S_0 \subset S$ such that $\text{rad}(S) = \text{rad}(S_0)$.

It was shown by Guba [20] that f.g. free groups are equationally Noetherian. His proof exploited the fact that free groups are linear which allows him to appeal to some classical algebraic geometry. This fact was used in [2] to show that large classes of linear groups are equationally Noetherian. Note however that hyperbolic groups are not necessarily linear [25]. Thus we can in general not appeal to linearity to establish that hyperbolic groups are equationally Noetherian.

The following simple fact was pointed out to the authors by Abderezak Ould Houcine, it says that to check whether a finitely generated group Γ is equationally Noetherian it suffices to check systems of equations without constants:

Lemma 6.1. *Let Γ be a finitely generated group. Suppose that for any $n \in \mathbb{N}$ and any subset $S \subset F(x_1, \dots, x_n) \subset \Gamma[x_1, \dots, x_n]$ there exists a finite subset $S_0 \subset S$ such that $\text{rad}(S) = \text{rad}(S_0)$. Then Γ is equationally Noetherian.*

Proof. Let $\Gamma = \langle g_1, \dots, g_k \rangle$ and $S \subset \Gamma[x_1, \dots, x_n]$. Consider the epimorphism

$$\beta : F(y_1, \dots, y_k, x_1, \dots, x_n) \rightarrow \Gamma[x_1, \dots, x_n]$$

given by $\beta(y_i) = g_i$ for $1 \leq i \leq k$ and $\beta(x_i) = x_i$ for $1 \leq i \leq n$. Put $\tilde{S} = \beta^{-1}(S)$. By hypothesis there exists a finite set $\tilde{S}_0 \subset \tilde{S}$ such that $\text{rad}(\tilde{S}) = \text{rad}(\tilde{S}_0)$. Clearly $S_0 := \beta(\tilde{S}_0)$ is a finite subset of S , as further $\text{rad}(S) = \text{rad}(S_0)$ this implies the claim. \square

Lemma 6.2. *If Γ is equationally Noetherian then for any sequence*

$$G_1 \rightarrow G_2 \rightarrow G_3 \rightarrow \dots$$

of epimorphisms of finitely generated groups the associated embeddings

$$\text{Hom}(G_1, \Gamma) \leftarrow \text{Hom}(G_2, \Gamma) \leftarrow \text{Hom}(G_3, \Gamma) \leftarrow \dots$$

eventually become bijections.

Proof. Given a finitely generated group $G = \langle x_1, \dots, x_n \mid R \rangle$ and a group Γ there is a one-to-one correspondence between $\text{Hom}(G, \Gamma)$ and $\text{rad}(R)$. Indeed if $\phi \in \text{Hom}(G, \Gamma)$ then

$(\phi(x_1), \dots, \phi(x_n)) \in \text{rad}(R)$ and for each tuple $(\gamma_1, \dots, \gamma_n) \in \text{rad}(R)$ the map $x_i \mapsto \gamma_i$ for $i = 1, \dots, n$ extends to a homomorphism $G \rightarrow \Gamma$.

Choose presentations $\langle x_1, \dots, x_n \mid R_i \rangle$ of G_i such that $R_i \subset R_{i+1}$. Put

$$R_\infty = \bigcup_{i \in \mathbb{N}} R_i$$

and $G_\infty := \langle x_1, \dots, x_n \mid R_\infty \rangle$, i.e. G_∞ is the direct limit of the G_i . As Γ is equationally Noetherian it follows that $\text{rad}(R_\infty) = \text{rad}(R_i)$ for some i . The claim follows. \square

Corollary 6.3. *If Γ is equationally Noetherian then any sequence*

$$G_1 \rightarrow G_2 \rightarrow G_3 \rightarrow \dots$$

of epimorphisms of finitely generated groups that are residually Γ eventually stabilizes.

Proof. Because of Lemma 6.2 it suffices to show that if two groups G and G' are residually Γ and $\pi : G \rightarrow G'$ is a non-injective epimorphism then $\pi_* : \text{Hom}(G', \Gamma) \rightarrow \text{Hom}(G, \Gamma)$ is not surjective. This is obvious as for $k \in \ker \pi \setminus \{1\}$ there is a homomorphism $\phi : G \rightarrow \Gamma$ such that $\phi(k) \neq 1$ as G is residually Γ . Clearly ϕ does not lie in the image of π_* . \square

Corollary 6.4. *Suppose that Γ is equationally Noetherian. Then any Γ -limit group is fully residually Γ .*

Proof. Let $L = F_k / \varinjlim \langle \varphi_i \rangle$ be a Γ -limit group. Choose a sequence

$$G_0 = F_k \rightarrow G_1 \rightarrow G_2 \rightarrow \dots$$

of finitely presented groups such that L is their direct limit. By Corollary 6.3 there exists some G_{i_0} such that any homomorphism $\varphi : G_{i_0} \rightarrow \Gamma$ factors through L . After passing to a subsequence we can further assume that any φ_i factors through G_{i_0} as G_{i_0} is finitely presented. Thus the sequence factors in fact through L , i.e. there exists a stable sequence $(\eta_i) \subset \text{Hom}(L, \Gamma)$ such that $\varinjlim \langle \eta_i \rangle = 1$, i.e. that $L = L / \varinjlim \langle \eta_i \rangle$. This clearly implies that for any finite set $M \subset L$, $\eta_i|_M$ is injective for sufficiently large i . \square

We will need the following simple lemma, its proof is identical to that in the case of a free group, see [6].

Lemma 6.5. *Let Γ be an equationally Noetherian group and G be a finitely generated group. Then there exist finitely many groups L_1, \dots, L_k and epimorphisms $q_i : G \rightarrow L_i$ such that the following hold.*

- (1) L_i is fully residually Γ for $i = 1, \dots, k$.
- (2) For any homomorphism $\phi : G \rightarrow \Gamma$ there exists $i \in \{1, \dots, k\}$ such that ϕ factors through q_i .

Proof. Let \widehat{G} be the universal residually Γ quotient of G , i.e. $\widehat{G} = G/N$ where N is the intersection of all kernels of homomorphisms from G to Γ . Clearly any homomorphism $\phi : G \rightarrow \Gamma$ factors through the canonical projection $\pi : G \rightarrow \widehat{G} = G/N$. It therefore suffices to show that there exists a finite collection of groups L_1, \dots, L_k that are fully residually Γ and epimorphisms $\widehat{q}_i : \widehat{G} \rightarrow L_i$ for $1 \leq i \leq k$ such that any homomorphism $\widehat{\phi} : \widehat{G} \rightarrow \Gamma$ factors through some \widehat{q}_i .

If \widehat{G} is fully residually Γ then the claim is trivial. Thus we can assume that \widehat{G} is not fully residually Γ . It follows that there exists a finite set $M = \{g_1, \dots, g_k\} \subset \widehat{G} \setminus \{1\}$ such that $\ker \phi \cap M \neq \emptyset$ for any homomorphism $\widehat{\phi} : \widehat{G} \rightarrow \Gamma$. It follows that any $\widehat{\phi} : \widehat{G} \rightarrow \Gamma$ factors through one of the epimorphisms $q_i : \widehat{G} \rightarrow L_i$ where L_i is the universal residually Γ quotient of $\widehat{G}/\langle\langle g_i \rangle\rangle$ for $1 \leq i \leq k$ and q_i is the canonical quotient map.

If L_i is not fully residually Γ we repeat this construction for L_i , after finitely many iterations this must terminate by Lemma 6.2. Thus we get a finite directed tree of epimorphisms such that any homomorphism factors through one branch, the assertion follows by choosing as epimorphisms the compositions of epimorphisms along maximal (directed) branches of this tree. \square

6.2. Dunwoody decompositions

Recall that a group G is called accessible if there exists a reduced graph of groups \mathbb{A} such that the following hold.

- (1) $\pi_1(\mathbb{A}) \cong G$.
- (2) Any edge group of \mathbb{A} is finite.
- (3) Any vertex group of \mathbb{A} is one-ended or finite.

We call any such \mathbb{A} a Dunwoody decomposition of G . Note that the graph of groups \mathbb{A} is far from being unique for a given accessible group G . However the maximal vertex groups are unique up to conjugacy; indeed they are precisely the maximal one-ended subgroups of G .

It is the Dunwoody accessibility theorem [12] that states that all finitely presented groups are accessible. It turns out that f.g. groups are in general not accessible but the particular case that we will need is covered by the following theorem of P. Linnell [30] of which Theorem 4.5 is a generalization.

Theorem 6.6. *Let G be a f.g. group. Suppose that there exists some constant C such that any finite subgroup H of G is of order at most C . Then G is accessible.*

Let now Γ be a hyperbolic group and L a Γ -limit group. As the order of finite subgroups of L is bounded by $N(\Gamma)$ it follows that Theorem 6.6 applies to L , i.e. L admits a Dunwoody decomposition \mathbb{D} . As modular automorphisms of the vertex groups of \mathbb{D} restrict to the identity on all finite subgroups it follows that they extend to automorphisms of L .

In the following we call the subgroup of $\text{Aut } L$ consisting of those automorphisms that restrict (up to conjugation) to modular automorphism of the vertex groups of \mathbb{D} the modular group of L and denote it by $\text{Mod}(L)$. In the case of a one-ended group this recovers our original definition.

Let now G be an accessible group and Γ be a group. Then we call a homomorphism $\psi : G \rightarrow \Gamma$ *locally injective* if ψ is injective when restricted to the vertex groups of some (and therefore all) Dunwoody decomposition of G . Note that this is equivalent to saying that ψ is injective when restricted to 1-ended and finite subgroups of G .

6.3. MR-diagrams for equationally Noetherian hyperbolic groups

In this section we give a proof of the main theorem of this article, i.e. the description of $\text{Hom}(G, \Gamma)$ for some finitely generated group G and some hyperbolic group Γ under the additional assumption that Γ is equationally Noetherian. It will then be the purpose of Section 7 to establish that all hyperbolic groups have this property.

Theorem 6.7. *Let Γ be an equationally Noetherian hyperbolic group and G be a finitely generated group. Then there exists a finite directed rooted tree T with root v_0 satisfying*

- (1) *The vertex v_0 is labeled by G .*
- (2) *Any vertex $v \in VT$, $v \neq v_0$, is labeled by a Γ -limit group G_v .*
- (3) *Any edge $e \in ET$ is labeled by an epimorphism $\pi_e : G_{\alpha(e)} \rightarrow G_{\omega(e)}$*

such that for any homomorphism $\phi : G \rightarrow \Gamma$ there exists a directed path e_1, \dots, e_k from v_0 to some vertex $\omega(e_k)$ such that

$$\phi = \psi \circ \pi_{e_k} \circ \alpha_{k-1} \circ \dots \circ \alpha_1 \circ \pi_{e_1}$$

where $\alpha_i \in \text{Mod } G_{\omega(e_i)}$ for $1 \leq i \leq k$ and ψ is locally injective.

Remark 6.8. In case G is fully residually Γ , the factorization of homomorphisms from G to Γ as in Theorem 6.7 requires modular automorphisms of G before the first proper quotient map. Thus in this case the diagram has precisely one edge e satisfying $\alpha(e) = v_0$, and $\pi_e : G = G_{\alpha(e)} \rightarrow G_{\omega(e)}$ is an isomorphism.

Proof of Theorem 6.7. In view of Lemma 6.5 and Corollary 6.3 it clearly suffices to show that any group G that is fully residually Γ admits a finite set $\{q_i : G \rightarrow \Gamma_i\}$ of proper quotient maps such that any homomorphism $\varphi : G \rightarrow \Gamma$ which is not locally injective factors through some q_i after precomposition with an element of $\text{Mod } G$.

Choose a Dunwoody decomposition \mathbb{D} of G . For each vertex group D_v there is a factor set

$$S_v = \{q_i^v : D_v \rightarrow D_v/N_v^i\}.$$

This follows from Theorem 5.2 if D_v is one-ended and is trivial if D_v is finite. For each v and i let \bar{N}_v^i be the normal closure of N_v^i in G . We now define the factor set for G to be

$$\{Q_i^v : G \rightarrow G/\bar{N}_v^i \mid v \in VD, q_i^v \in S_v\}.$$

To see that this is a factor set let $\varphi : G \rightarrow \Gamma$ be a non-locally injective homomorphism. Choose $v \in VD$ such that $\varphi|_{D_v}$ is non-injective. Thus there exists $\alpha \in \text{Mod}(D_v)$ such that $\varphi|_{D_v} \circ \alpha : D_v \rightarrow \Gamma$ factors through some q_i^v . As α extends to an automorphism $\alpha' \in \text{Mod}(G)$ it follows that $\varphi \circ \alpha'$ factors through Q_i^v . \square

7. Shortening quotients and applications

In the previous section we have constructed Makanin–Razborov diagrams for equationally Noetherian hyperbolic groups. It is the purpose of this last chapter to establish that all hyperbolic groups are equationally Noetherian, i.e. that the construction of the Makanin–Razborov diagrams applies to all hyperbolic groups.

7.1. Shortening quotients

In Section 5, see Remark 5.7, we have seen that if $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ is a stable sequence such that $\overrightarrow{\ker}(\varphi_i) = 1$, i.e. that $L = G/\overrightarrow{\ker}(\varphi_i) = G$, then we can construct a proper quotient $G/\overrightarrow{\ker}(\widehat{\varphi}_i)$ of $G = L$ where the $\widehat{\varphi}_i$ are the shortened φ_i . This quotient is clearly again a Γ -limit group and is called a shortening quotient.

This construction only works if G is fully residually Γ . It is the main purpose of this section to construct shortening quotients for arbitrary Γ -limit groups. In the end, see Corollary 7.6, it will turn out that all Γ -limit groups are fully residually Γ . We will first treat one-ended Γ -limit groups and then deal with the general case.

Let $L = G/\overrightarrow{\ker}(\varphi_i)$ be a one-ended Γ -limit group and \mathbb{A} be an almost Abelian JSJ-decomposition of L . Lemma 7.1 below guarantees that we can approximate L by a sequence of groups (W_i) that are endowed with splittings that approximate \mathbb{A} .

Lemma 7.1. *Let G be a finitely presented group and $L = G/\overrightarrow{\ker}(\varphi_i)$ be a one-ended Γ -limit group with associated Γ -limit map $\varphi : G \rightarrow L$. Let \mathbb{A} be an almost Abelian JSJ-decomposition of L , in particular $L = \pi_1(\mathbb{A}, v_0)$.*

Then there exists a sequence of graphs of groups (\mathbb{A}^i) with underlying graph A and finitely presented fundamental groups $W_i = \pi_1(\mathbb{A}^i, v_0)$, surjective morphisms $f^i : \mathbb{A}^i \rightarrow \mathbb{A}^{i+1}$ and $h^i : \mathbb{A}^i \rightarrow \mathbb{A}$ and an epimorphism $\gamma : G \rightarrow W_0$ such that the following hold.

- (1) $\varphi = h_*^0 \circ \gamma : G \rightarrow L$.
- (2) $h^i = h^{i+1} \circ f^i$ for all i .
- (3) L is the direct limit of the sequence W_i , i.e.

$$\overrightarrow{\ker}(\varphi_i) = \bigcup_{k=1}^{\infty} \ker(f_*^k \circ f_*^{k-1} \circ \dots \circ f_*^1 \circ f_*^0 \circ \gamma).$$

- (4) *If A_v is an orbifold type vertex group then $\psi_v^{h^i} : A_v^i \rightarrow A_v$ is an isomorphism for all i .*
- (5) *If A_v is of axial type then $\psi_v^{h^i} : A_v^i \rightarrow A_v$ is injective for all i .*
- (6) *The maps $\psi_e^{h^i} : A_e^i \rightarrow A_e$ are injective for all i and $e \in EA$.*
- (7) *For any $v \in VA$ we have $\bigcup \psi_v^{h^i}(A_v^i) = A_v$.*
- (8) *For any $e \in EA$ we have $\bigcup \psi_e^{h^i}(A_e^i) = A_e$.*

Proof. This is a simple application of foldings as discussed in [4] and Dunwoody's vertex morphisms [13]. Let $T = \widetilde{\mathbb{A}}$ be the Bass–Serre tree corresponding to \mathbb{A} , thus T is an L -tree. The Dunwoody Resolution Lemma guarantees that there is a G -tree Y with finitely generated edge and vertex stabilizers and a surjective morphism (id, p) from Y to the L -tree T , see [12, 14] or [4].

After applying finitely many folds to the G -tree Y we obtain a G -tree Y' such that the induced map on the graphs of groups $Y'/G \rightarrow T/L = \mathbb{A}$ is bijective on the level of graphs and surjective for the edge and vertex groups of \mathbb{A} that are finitely generated, see [4].

We now apply vertex morphisms to quotient out the kernels of the homomorphisms of edge groups and of the vertex groups whose targets are QH-subgroups or almost Abelian groups. This clearly adds only finitely many relations. Denote the resulting graph of groups by \mathbb{A}_0 , the morphism from \mathbb{A}_0 to \mathbb{A} clearly satisfies (4)–(6).

We can now continue to apply folds of type IIA and IIB, see [4], and get a sequence of graphs of groups satisfying (7) and (8). At each step we further add all relators to

edge groups and vertex group mapped to almost Abelian vertex groups which makes this sequence preserve properties (4)–(6). Finally we add at each step the shortest relator to the vertex groups that are mapped to rigid vertex groups that does not yet hold. This then implies that all relations of L will hold eventually, i.e. that (3) holds. Parts (1) and (2) hold by construction. \square

Let now G be finitely presented with fixed finite generating set S , $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ a stable sequence such that $L = G/\underline{\ker}(\varphi_i)$ is a one-ended Γ -limit group. Let \mathbb{A} be an almost Abelian JSJ-decomposition of L . Choose sequences (\mathbb{A}_i) , (h^i) and (f^i) , (W_i) and γ as in Lemma 7.1.

Let

$$\xi_k := f_*^{k-1} \circ \dots \circ f_*^1 \circ f_*^0 \circ \gamma : G \rightarrow W_k$$

be the epimorphism induced by the f^i and γ . Then $S_i := \xi_i(S)$ is a generating set of W_i . After replacing (φ_i) by a subsequence we can assume that φ_i factors through ξ_i for all i , i.e. that

$$\varphi_i = \lambda_i \circ \xi_i$$

for some $\lambda_i : W_i \rightarrow \Gamma$. This is clearly possible as the finitely many defining relations of W_i lie in $\underline{\ker}(\varphi_i)$ and therefore in the kernel of φ_j for sufficiently large j .

Let now $\widehat{\lambda}_i : W_i \rightarrow \Gamma$ be the homomorphism obtained from shortening λ_i by precomposition with elements of $\text{Mod}_{\mathbb{A}_i}(W_i)$ and postcomposition with an inner automorphism. Here shortness is measured with respect to the generating set S_i . We then put $\eta_i := \widehat{\lambda}_i \circ \xi_i$. After passing to a subsequence we can assume that (η_i) is stable. We then put

$$Q := G/\underline{\ker}(\eta_i)$$

and call Q a shortening quotient of L . It is clear from the construction that Q is a quotient of L . Indeed if $g \in \underline{\ker}(\varphi_i)$ then $g \in \ker \xi_i$ for large i as we assume that (3) of Lemma 7.1 is satisfied. Thus $g \in \ker \eta_i = \ker \widehat{\lambda}_i \circ \xi_i$ for large i and therefore $g \in \underline{\ker}(\eta_i)$. We denote the projection from L to Q by π , thus we have $\eta = \pi \circ \varphi$ if η and φ are the Γ -limit maps associated to the sequences (η_i) and (φ_i) .

Proposition 7.2. *Let $L = G/\underline{\ker}(\varphi_i)$ and $Q = G/\underline{\ker}(\eta_i)$ be as above and \mathbb{A} be an almost Abelian JSJ-decomposition of L . Then the following hold.*

- (1) *The epimorphism $\pi : L \rightarrow Q$ is injective on rigid vertex groups of \mathbb{A} .*
- (2) *If (η_i) is not contained in finitely many conjugacy classes then Q is a proper quotient of L .*

(3) *If all almost Abelian subgroups of Q are finitely generated then the following hold:*

- (a) *If a subsequence of (η_i) factors through $\eta : G \rightarrow Q$ then a subsequence of (φ_i) factors through $\varphi : G \rightarrow L$.*
- (b) *All almost Abelian subgroups of L are finitely generated.*

Proof. (1). Let $g \in L$ be an element that is conjugate to an element h of a rigid vertex group A_v of \mathbb{A} such that $\pi(g) = 1$. We need to show that $g = 1$.

As we assume that (7) of Lemma 7.1 is satisfied it follows that for some i_0 there exists $g_{i_0} \in W_{i_0}$ that is conjugate to some $k_{i_0} \in A_v^{i_0}$ such that $h_*^{i_0}(g_{i_0}) = g$. Choose $\tilde{g} \in G$ such that $\xi_{i_0}(\tilde{g}) = g_{i_0}$ and put $g_i = \xi_i(\tilde{g})$ for all i . Note that for $i \geq i_0$ the element g_i is conjugate to some element $k_i \in A_v^i$.

Now $\pi(g) = 1$ means that $\eta_i(\tilde{g}) = \widehat{\lambda}_i \circ \xi_i(\tilde{g}) = \widehat{\lambda}_i(g_i) = 1$ for large i , in particular any element conjugate to g_i lies in the kernel of $\widehat{\lambda}_i$. Recall that $\widehat{\lambda}_i = \lambda_i \circ \alpha_i$ for some modular automorphism α_i of W_i . As g_i is conjugate to the element $k_i \in A_v^i$ and as modular automorphisms act on rigid groups by conjugation it follows that $\alpha_i(g_i)$ is also conjugate to k_i and therefore conjugate to g_i . Thus $\widehat{\lambda}_i \circ \alpha_i(g_i) = \lambda_i(g_i) = 1$. This implies that $\varphi_i(\tilde{g}) = \lambda_i \circ \xi_i(\tilde{g}) = \lambda_i(g_i) = 1$ for large i , it follows that $\tilde{g} \in \ker(\varphi_i)$. Thus $g = 1$.

(2). Assume to the contrary that (η_i) contains infinitely many conjugacy classes and that $L = Q$, i.e. that $L = Q = G/\ker(\eta_i)$. After passing to a subsequence we can assume that (η_i) converges to an action (T, x, ρ) of L on an \mathbb{R} -tree T satisfying the assumptions of Theorem 3.4. Let now \mathcal{G} be the graph of actions decomposition corresponding to this action. As in Section 5 we distinguish 3 different cases.

If the graph of actions has an orbifold type vertex then there is an automorphism of its vertex group that extends to a modular automorphism α of L which shortens the action of L on T , i.e. for which $|\rho \circ \alpha|_x \leq |\rho|_x$. Recall that $\eta = h_*^k \circ \xi_k$ for all k . Now this orbifold type vertex group corresponds to a suborbifold of one of the orbifold type vertices of the JSJ-decomposition of L . Thus α can be lifted to any W_i as the morphisms h^i are isomorphisms when restricted to QH-subgroups. Thus there exists $\alpha_i \in \text{Mod } W_i$ such that $\alpha \circ h_*^i = h_*^i \circ \alpha_i$. It follows that

$$|h_*^i \circ \alpha_i \circ \xi_i|_x < |\rho \circ \alpha|_x < |\rho|_x$$

and therefore $|\widehat{\lambda}_i \circ \alpha_i \circ \xi_i| < |\widehat{\lambda}_i \circ \xi_i| = |\eta_i|$, contradicting the shortness of the $\widehat{\lambda}_i$.

If the action has an axial type vertex then we can choose α as in the case of an orbifold type vertex but the lifting is slightly more subtle. Note first that the vertex group corresponding to this axial vertex space is also a vertex group of the JSJ-decomposition of L . This is true as the group must be elliptic in the JSJ as it is an almost Abelian subgroup

that is not 2-ended. We can assume that it is a vertex group as we could otherwise refine the JSJ contradicting its maximality.

Now the morphisms $\psi_v^{h_i} : A_v^i \rightarrow A_v$ are not necessarily surjective on almost Abelian vertex groups. However for large i the group $\psi_v^{h_i}(A_v^i) \leq A_v$ contains all generators, and therefore all elements, of A_v that act non-trivially on the axial tree. This means we can define an automorphism of A_v^i that extends to an automorphism α_i of W_i such that $\widehat{\lambda}_i \circ \alpha_i$ is shorter than $\widehat{\lambda}_i$ by analyzing the action of A_v^i on the axial tree via $\psi_v^{h_i}$. Note that this is easier than in Section 5 as the group A_v^i is finitely generated.

In the simplicial case the edge group along which the Dehn twist is performed either corresponds to an edge group of the JSJ or to a simple closed curve in a QH-subgroup of the JSJ. In both cases we can simply lift the Dehn twists to W_i and thereby shorten the homomorphism η_i . In the case where the edge group corresponds to a simple closed curve of a QH-subgroup this is obvious, in the other case it follows as an element that is central in the edge group of the graph of actions is also central in the corresponding edge group of \mathbb{A}^i .

(3). Note first that any edge group of \mathbb{A} , the JSJ-decomposition of L , is contained in either a rigid or an orbifold type vertex group. It follows that all edge groups of \mathbb{A} are finitely generated. Indeed if the edge group is contained in a rigid vertex group then it follows from (1) that the edge group embeds into Q and is therefore finitely generated. Otherwise the edge group is virtually cyclic and the modular automorphisms act on the group by conjugation, it therefore follows as in the proof of (1) that it is embedded into Q and is therefore finitely generated.

As $L = \pi_1(\mathbb{A})$ is finitely generated and all edge groups of \mathbb{A} are finitely generated it follows that also all vertex groups of \mathbb{A} are finitely generated. Thus there exist i_0 such that for $i \geq i_0$ the morphism h^i is bijective on all edge groups and non-rigid vertex groups. On the rigid vertex groups h^i is surjective, i.e. the morphism consists just of vertex morphisms in the sense of Dunwoody [13]. As almost all ξ_i factor through W_{i_0} we can pass to a subsequence and assume that $i_0 = 0$.

Suppose now that η_i factors through Q . For any $v \in VA$ denote the kernel of the map $\psi_v^{h_0} : A_v^0 \rightarrow A_v$ by K_v . As η_i factors through Q and therefore through L it follows that $K_v \subset \ker \eta_i$ for any vertex group A_v . As the η_i and the φ_i only differ by precomposition with an automorphism that acts by conjugation on rigid vertex groups this implies that $K_v \subset \ker \varphi_i$ for all rigid vertex groups A_v . Thus φ_i factors through L as all other relations of L already hold in W_0 . The second assertion follows immediately from the proof. \square

The above construction only works for one-ended Γ -limit groups as we need the existence of an almost Abelian JSJ-decomposition. In the remainder of this section we

will show that the concept of a shortening quotient generalizes naturally to all Γ -limit groups.

Let now G be a finitely presented group and $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ be a stable sequence. Put $L := G/\underline{\ker}(\varphi_i)$ and denote the associated Γ -limit map by φ . Let \mathbb{D} be a Dunwoody decomposition of L , i.e. $L = \pi_1(\mathbb{D}, v_0)$, all edge groups of \mathbb{D} are finite and no vertex group splits over finite groups. Thus every vertex group is either finite or one-ended.

As in the proof of Lemma 7.1 we see that there is a graph of groups \mathbb{D}' whose underlying graph D is the same graph that is underlying \mathbb{D} , a morphism $f : \mathbb{D}' \rightarrow \mathbb{D}$ and an epimorphism $\gamma : G \rightarrow \pi_1(\mathbb{D}', v_0)$ such that the following hold.

- (1) $\varphi = f_* \circ \gamma$.
- (2) $\pi_1(\mathbb{D}', v_0)$ is finitely presented.
- (3) The morphism f is bijective on edge groups, thus f only consists of a collection of vertex morphisms on some vertices.

Now as $\pi_1(\mathbb{D}', v_0)$ is finitely presented almost all φ_i factor through γ . Thus after omitting finitely many elements from (φ_i) we can assume that for all i there exists $\bar{\varphi}_i : \pi_1(\mathbb{D}', v_0) \rightarrow L$ such that $\varphi_i = \bar{\varphi}_i \circ \gamma$.

For each vertex $v \in VD$ we get a stable sequence $(\bar{\varphi}_i^v)$ where $\bar{\varphi}_i^v : D'_v \rightarrow \Gamma$ is the restriction of $\bar{\varphi}_i$ to D'_v . Note that this restriction is only unique up to inner automorphisms of Γ unless we choose a preferred conjugate of D'_v in $\pi_1(\mathbb{D}', v_0)$. Independently of these conjugacy factors the obtained sequence is stable for all $v \in VD$ and we have $D_v = D'_v/\underline{\ker}(\bar{\varphi}_i^v)$.

Now for every one-ended D_v we can apply the construction of the shortening quotient to the sequence $(\bar{\varphi}_i^v)$ and obtain (after passing to a subsequence) a new stable sequence $(\bar{\eta}_i^v) \subset \text{Hom}(D'_v, \Gamma)$ such that $\underline{\ker}(\bar{\varphi}_i^v) \leq \underline{\ker}(\bar{\eta}_i^v)$ and that all conclusions of Proposition 7.2 hold for the quotient map

$$\pi_v : D_v = D'_v/\underline{\ker}(\bar{\varphi}_i^v) \rightarrow Q_v := D'_v/\underline{\ker}(\bar{\eta}_i^v).$$

If D_v is finite we put $\bar{\eta}_i^v = \bar{\varphi}_i^v$ for all i .

Now as the shortening automorphisms act on finite subgroups by conjugation it follows that for each i there exists a (not unique) homomorphism $\bar{\eta}_i : \pi_1(\mathbb{D}', v_0) \rightarrow \Gamma$ such that the restriction of $\bar{\eta}_i$ to D'_v is conjugate to $\bar{\eta}_i^v$ for all $v \in VD$.

We put $\eta_i = \bar{\eta}_i \circ \gamma : G \rightarrow \Gamma$. After passing to a subsequence we can assume that (η_i) is stable and we put $Q := G/\underline{\ker}(\eta_i)$. As $\underline{\ker}(\varphi_i) \leq \underline{\ker}(\eta_i)$ by construction we have a natural epimorphism $\pi : L \rightarrow Q$. As in the one-ended case it follows that $\eta = \pi \circ \varphi$ if η and φ are the Γ -limit maps associated to the sequences (η_i) and (φ_i) , respectively.

It is clear that the epimorphism π maps the vertex groups D_v of \mathbb{D} to subgroups of Q that are isomorphic to their shortening quotients Q_v , but we do not claim that the Q_v are vertex groups of the Dunwoody decomposition of Q .

Theorem 7.3. *Let G be a finitely presented group and $L = G/\overrightarrow{\ker}(\varphi_i)$ a Γ -limit group. Let (η_i) be as above and $Q = G/\overrightarrow{\ker}(\eta_i)$. Let $\pi : L \rightarrow Q$ be the natural quotient map. Then one of the following holds.*

- (1) $\ker \pi \neq 1$.
- (2) *A subsequence of (η_i) factors through η and all almost Abelian subgroups of Q are finitely generated.*

If moreover all almost Abelian subgroups of Q are finitely generated then the following hold.

- (a) *If a subsequence of (η_i) factors through $\eta : G \rightarrow Q$ then a subsequence of (φ_i) factors through $\varphi : G \rightarrow L$.*
- (b) *Almost Abelian subgroups of L are finitely generated.*

Proof. We first prove (1). Assume that $\ker \pi = 1$, i.e. that $L = G/\overrightarrow{\ker}(\varphi_i) = G/\overrightarrow{\ker}(\eta_i) = Q$. Thus for each $v \in VD$, the epimorphism $\pi_v : D_v \rightarrow Q_v$ is an isomorphism, hence by Proposition 7.2, $(\bar{\eta}_i^v)$ contains only finitely many conjugacy classes. After passing to a subsequence we can assume that for each v , all $\bar{\eta}_i^v$ are conjugate, i.e. that $\overrightarrow{\ker}(\bar{\eta}_i^v) = \ker \bar{\eta}_i^v$ for all i and all $v \in VD$, in particular $D_v \cong Q_v = G/\overrightarrow{\ker}(\bar{\eta}_i^v) \cong \bar{\eta}_i^v(G) \leq \Gamma$. As almost Abelian subgroups of hyperbolic groups are 2-ended it follows that all almost Abelian subgroups of vertex groups of \mathbb{D} and therefore of $\pi_1(\mathbb{D}, v_0) = L = Q$ are finitely generated.

We will now show that a subsequence of (η_i) factors through $\eta = \pi \circ \varphi$ where η is the limit map associated to $(\eta_i) = (\bar{\eta}_i \circ \gamma)$. As all η_i factor through γ it follows that the associated Γ -limit map η also factors through γ . Choose $\bar{\eta}$ such that $\eta = \bar{\eta} \circ \gamma$. As $\ker \pi = 1$ it follows that $\ker \eta = \ker \pi \circ \varphi = \ker \varphi$. Thus the kernel of $\bar{\eta}$ is normally generated by the stable kernels $\overrightarrow{\ker}(\bar{\eta}_i^v)$. By the above remark there is a subsequence of $(\bar{\eta}_i)$ for which $\overrightarrow{\ker}(\bar{\eta}_i^v) = \ker \bar{\eta}_i^v$ for all $v \in VD$ and i , it follows that this subsequence factors through η .

Now suppose that all almost Abelian subgroups of Q are finitely generated. The shortening quotient Q_v of D_v embeds into Q for all v , thus all almost Abelian subgroups of Q_v are finitely generated. It thus follows from Proposition 7.2 that all almost Abelian subgroups of D_v and therefore also $L = \pi_1(D_v)$ are finitely generated, this proves (b). The proof of (a) is similar to the proof of the first part; it suffices to show for a subsequence

of (φ_i) we have $\ker(\overrightarrow{\varphi_i^v}) \subset \ker \overrightarrow{\varphi_i^v}$. By assumption a subsequence of (η_{j_i}) of (η_i) factors through η which implies that the sequences $(\overrightarrow{\eta_{j_i}^v})$ factor through $D'_v \rightarrow Q_v = D'_v / \ker(\overrightarrow{\eta_{j_i}^v})$ for all $v \in DV$. As almost Abelian subgroups of Q_v are finitely generated it follows from Proposition 7.2 that a subsequence of $(\overrightarrow{\varphi_i^v})$ factors through $D'_v \rightarrow D_v = D'_v / \ker(\overrightarrow{\varphi_i^v})$ for all $v \in VD$, thus a subsequence of $(\overrightarrow{\varphi_i})$ factors $\pi_1(\mathbb{D}', v_0) \rightarrow L = \pi_1(\mathbb{D}, v_0)$, i.e. a subsequence of $(\varphi_i) = (\overrightarrow{\varphi_i} \circ \gamma)$ factors through $\varphi : G \rightarrow L$. \square

If $L = G / \ker(\overrightarrow{\varphi_i})$ and $Q = G / \ker(\overrightarrow{\eta_i})$ are as in Theorem 7.3 then we call Q a *shortening quotient* of L and we say that (η_i) is obtained from (φ_i) by shortening or by the shortening procedure. It will be important in the next section that η factors through φ if η and φ are the Γ -limit maps corresponding to (η_i) and (φ_i) .

7.2. Hyperbolic groups are equationally Noetherian

In this chapter we show that hyperbolic groups are equationally Noetherian. We fix a hyperbolic group Γ . Crucial to the argument is a partial order on Γ -limit maps defined as follows.

Definition 7.4. Let G be f.g. and $\varphi : G \rightarrow L_\varphi, \eta : G \rightarrow L_\eta$ be Γ -limit maps. We say that $\eta \leq \varphi$ if $\eta = \pi \circ \varphi$ for some epimorphism $\pi : L_\varphi \rightarrow L_\eta$. We further say $\eta < \varphi$ if $\eta \leq \varphi$ and $\varphi \not\leq \eta$.

Fix a f.g. group G . By definition, every homomorphism $\eta : G \rightarrow \Gamma$ is a Γ -limit map, arising from the constant sequence (η) . Note further that the relation \leq on the set of all Γ -limit maps from G is transitive. We will show that there are only finitely many maximal Γ -limit maps with respect to \leq . The main technical step is the proof of the following theorem.

Theorem 7.5. *Let $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ be a stable sequence and $\varphi : G \rightarrow G / \ker(\overrightarrow{\varphi_i})$ the corresponding Γ -limit map. Then a subsequence of (φ_i) factors through φ .*

As an immediate consequence of Theorem 7.5 we get the following.

Corollary 7.6. *Finitely generated Γ -limit groups are fully residually Γ .*

Proof. Let G be f.g., $(\eta_i) \subset \text{Hom}(G, \Gamma)$ a stable sequence with stable limit map $\eta : G \rightarrow L := G / \ker(\overrightarrow{\eta_i})$. Let $E = \{g_1, \dots, g_k\} \subset L$, we need to show that there is a homomorphism from L to Γ that maps E injectively.

Choose $\tilde{E} = \{\tilde{g}_1, \dots, \tilde{g}_k\} \subset G$ such that $\eta(\tilde{g}_j) = g_j$ for $j = 1, \dots, k$. As $\eta|_{\tilde{E}}$ is injective, there exists an $i_0 \in \mathbb{N}$ such that for $i \geq i_0$, $\eta_i|_{\tilde{E}}$ is injective. Moreover, by Theorem 7.5, there is an $i \geq i_0$ such that $\eta_i = \tilde{\eta}_i \circ \eta$ for some $\tilde{\eta}_i \in \text{Hom}(L, \Gamma)$. Clearly, $\tilde{\eta}_i|_E$ is injective. \square

Theorem 7.5 is an immediate consequence of Lemma 7.7 and Lemma 7.8.

Lemma 7.7. *Let $(\varphi_i) \subset \text{Hom}(F_k, \Gamma)$ be a stable sequence and φ its associated Γ -limit map. Then one of the following holds.*

(1) *There exists an infinite descending sequence of Γ -limit maps*

$$\varphi > \eta^1 > \eta^2 > \eta^3 > \dots$$

(2) *For infinitely many i , φ_i factors through φ .*

Proof. Assume that (1) does not hold. Let (η_i^1) be a stable sequence obtained from (φ_i) by shortening (and passing to a subsequence) and let η^1 be the corresponding Γ -limit map. If $\eta_1 < \varphi$ then we choose a sequence (η_i^2) with Γ -limit map η^2 by shortening (η_i^1) and so on. By assumption this process terminates, i.e. for some s we have $\ker \eta^s = \ker \eta^{s+1}$.

By Theorem 7.3 (2) a subsequence of $(\bar{\eta}_i^{s+1})$ factors through η^{s+1} and all almost Abelian subgroups of $L_{\eta^{s+1}}$ are finitely generated. Applying Theorem 7.3 (a) and (b) $s + 1$ times implies that a subsequence of (φ_i) factors through φ , i.e. that (2) occurs. \square

Lemma 7.8. *There exists no infinite descending sequence of Γ -limit maps.*

Proof. Assume that an infinite descending sequence of Γ -limit maps exists. For each $k \in \mathbb{N}$, choose a stable sequence $(\eta_i^k) \subset \text{Hom}(G, \Gamma)$ with associated Γ -limit map $\eta^k : G \rightarrow G/\overrightarrow{\ker(\eta_i^k)}$ such that

(1) $\eta^1 > \eta^2 > \dots$ is an infinite descending sequence of Γ -limit maps,

(2) for each $n > 1$, if $\bar{\eta}^n$ is a Γ -limit map such that $\bar{\eta}^n < \eta^{n-1}$ and there is an infinite descending sequence $\eta^{n-1} > \bar{\eta}^n > \dots$ of Γ -limit maps, then

$$|\ker \bar{\eta}^n \cap B_n| \leq |\ker \eta^n \cap B_n|,$$

where B_n is the Ball of radius n in G around the identity with respect to some fixed finite generating set.

It is clear that such a sequence exists, as the η^i can be chosen inductively to satisfy property 2. For each n choose an index i_n such that

(3) $\ker \eta_{i_n}^n \cap B_n = \ker \eta^n \cap B_n$,

(4) $\ker \eta^{n+1} \not\subseteq \ker \eta_{i_n}^n$.

As $\eta^{n+1} < \eta^n$, it is clear that these conditions are satisfied if i_n is chosen sufficiently large.

By construction, the diagonal sequence $(\eta_{i_n}^n)_{n \in \mathbb{N}} \subset \text{Hom}(G, \Gamma)$ is stable. Denote its Γ -limit map by η^∞ , clearly $\eta^\infty < \eta^n$ for all n .

It suffices to show that η^∞ does not allow an infinite descending sequence of Γ -limit maps

$$\eta^\infty > \varphi^1 > \varphi^2 > \dots$$

as it follows then from Lemma 7.7 that infinitely many $\eta_{i_n}^n$ factor through η^∞ , which clearly contradicts condition 4 of the construction, as $\ker \eta^{n+1} \not\subseteq \ker \eta_{i_n}^n$ implies that $\eta_{i_n}^n$ does not factor through η^{n+1} and therefore not through η^∞ .

So assume that an infinite descending sequence $\eta^\infty > \varphi^1 > \varphi^2 > \dots$ exists. Choose an element $g \in G$ with $\eta^\infty(g) \neq 1$, but $\varphi^1(g) = 1$. Assume that $|g| = n$. Then

$$\eta^1 > \eta^2 > \dots > \eta^{n-1} > \varphi^1 > \varphi^2 \dots$$

is an infinite descending sequence of Γ -limit maps and

$$|\ker \varphi^1 \cap B_n| > |\ker \eta^n \cap B_n|,$$

in contradiction to condition (2). □

In [39] proved that torsion-free hyperbolic groups are Hopfian, essentially the same argument works in our situation; see also [19] for the case of torsion-free toral relatively hyperbolic groups.

Corollary 7.9. *Hyperbolic groups are Hopfian.*

Proof. Let Γ be a hyperbolic group. We need to show that any epimorphism $\eta : \Gamma \rightarrow \Gamma$ is an isomorphism, i.e. has trivial kernel.

Note that $\eta^n : \Gamma \rightarrow \Gamma$ (the n th power of η) is also an epimorphism. If η has non-trivial kernel then $\ker \eta^{n+1} < \ker \eta^n$ for all n . Thus we have an infinite sequence

$$id > \eta > \eta^2 > \eta^3 > \dots$$

of Γ limit maps (recall that all homomorphisms to Γ are Γ -limit maps coming from the constant sequence), a contradiction to Lemma 7.8. □

We can now establish the existence of maximal Γ -limit quotients.

Theorem 7.10. *Let*

$$\eta^1 < \eta^2 < \eta^3 < \dots$$

be an infinite ascending sequence of Γ -limit maps. There exists a Γ -limit map η such that for every $n \in \mathbb{N}$, $\eta^n < \eta$.

Proof. Assume that for each n , η^n is the Γ -limit map of $(\eta_i^n)_{i \in \mathbb{N}}$. Choose an index i_n such that $\ker \eta_{i_n}^n \cap B_n = \ker \eta^n \cap B_n$.

By construction, the sequence $(\eta_{i_n}^n)_{n \in \mathbb{N}}$ is stable. Denote its associated Γ -limit map by η . We claim that for every n , $\eta^n < \eta$. Assume that for some n_0 , $\eta^{n_0} \not< \eta$. Then there is an element $g \in \ker \eta$ such that $g \notin \ker \eta^{n_0}$. It follows that $g \notin \ker \eta^n$ for all $n \geq n_0$, and so for each $n \geq \max\{n_0, |g|\}$, $g \notin \ker \eta_{i_n}^n$. This implies that $g \notin \ker \eta$, a contradiction. \square

Theorem 7.11. *Let G be f.g. There are only finitely many maximal Γ -limit maps from G .*

Proof. The proof is by contradiction. If there are infinitely many maximal Γ -limit maps then it is easily verified that there is a sequence (η^i) of pairwise distinct maximal Γ -limit maps such that for each $j, k \geq i$ we have

$$\ker \eta^j \cap B_i = \ker \eta^k \cap B_i,$$

where B_i is the Ball of radius i in G around the identity with respect to some fixed finite generating set.

For each i choose $\eta_i : G \rightarrow \Gamma$ such that $\ker \eta^i \cap B_i = \ker \eta_i \cap B_i$. The sequence (η^i) is clearly stable. Let $\eta : G \rightarrow G/\underline{\ker}(\eta_i)$ be the corresponding Γ -limit map.

After possibly removing a single η^i from the sequence we can assume that $\eta \not\geq \eta^i$ for all i as we would otherwise get a contradiction to the maximality of the η^i . Thus for each i there exists $g_i \in G$ such that $\eta^i(g_i) \neq 1$ and $\eta(g_i) = 1$.

It follows that there exists a stable sequence $(\varphi_i) \subset \text{Hom}(G, \Gamma)$ such that for each i we have $\ker \varphi_i \cap B_i = \ker \eta^i \cap B_i$ and $\varphi_i(g_i) \neq 1$, in particular no φ_i factors through η as $\eta(g_i) = 1$.

Let $\varphi : G \rightarrow G/\underline{\ker}(\varphi_i)$ be the associated Γ -limit map, we clearly get $\varphi = \eta$. By Theorem 7.5 a subsequence of (φ_i) factors through $\varphi = \eta$, a contradiction. \square

Lemma 7.12. *Let Γ be a hyperbolic group and $\varphi : F(x_1, \dots, x_n) \rightarrow H$ be an epimorphism.*

Assume that $S \subset F_n = F(x_1, \dots, x_n) \subset \Gamma[x_1, \dots, x_n]$ is such that for every finite $S_0 \subset S$,

$$\text{rad}(\ker \varphi \cup S) \subsetneq \text{rad}(\ker \varphi \cup S_0).$$

Then there is a Γ -limit map $\eta : F_n \rightarrow F_n/\underline{\ker}(\eta_i)$ such that $\ker \varphi \leq \ker \eta$ and that

$$\text{rad}(\ker \eta \cup S) \subsetneq \text{rad}(\ker \eta \cup S_0)$$

for every finite $S_0 \subset S'$.

Proof. By Theorem 7.11, there are only finitely many maximal Γ -limit maps $\varphi_1, \dots, \varphi_k : H \rightarrow H_i$, put $\eta_i = \varphi_i \circ \varphi$ for $1 \leq i \leq k$. Note that

$$\bigcup \text{rad}(\ker \eta_i) = \text{rad}(\ker \varphi)$$

as any homomorphism from F_n to Γ that factors through φ must factor through some maximal Γ -limit map and therefore through some η_i .

Assume that for each i there is a finite set $S_0^i \subset S$ such that $\text{rad}(\ker \eta_i \cup S_0^i) = \text{rad}(\ker \eta_i \cup S)$. Putting $S_0 := \bigcup S_0^i$, we get

$$\begin{aligned} \text{rad}(\ker \varphi \cup S_0) &= \bigcup \text{rad}(\ker \eta_i \cup S_0) \\ &= \bigcup \text{rad}(\ker \eta_i \cup S) \\ &= \text{rad}(S) \cap \bigcup \text{rad}(\ker \eta_i) \\ &= \text{rad}(S) \cap \text{rad}(\ker \varphi) \\ &= \text{rad}(S \cup \ker \varphi), \end{aligned}$$

which is a contradiction. Thus for some i_0 such a set $S_0^{i_0}$ does not exist and the conclusion follows by putting $\eta = \eta_{i_0}$. \square

Corollary 7.13. *Hyperbolic groups are equationally Noetherian.*

Proof. Let Γ be a hyperbolic group. Because of Lemma 6.1 it suffices to check that for any set $S \subset F_n = F(x_1, \dots, x_n) \subset \Gamma[x_1, \dots, x_n]$ there exists a finite subset $S_0 \subset S$ such that $\text{rad}(S) = \text{rad}(S_0)$.

Assume that $n \in \mathbb{N}$ and $S = \{w_1, w_2, \dots\} \subset F_n$ such that $\text{rad}(S) \subsetneq \text{rad}(S_0)$ for every finite $S_0 \subset S$. We show that this implies the existence of an infinite descending sequence of Γ -limit maps, contradicting Lemma 7.8. Let $\varphi_1 : F_n \rightarrow F_n / \langle\langle w_1 \rangle\rangle$ and η_1 a Γ -limit map with $\ker \varphi_1 \leq \ker \eta_1$ as in Lemma 7.12. Then inductively for each i , pick $w_{j_i} \in S \setminus \ker \eta_{i-1}$ and put

$$\varphi_i : F_n / \langle\langle \ker \eta_i \cup w_{j_i} \rangle\rangle$$

and apply Lemma 7.12 to obtain a Γ -limit map η_i with $\ker \varphi_i \leq \ker \eta_i$. Then all η_i are Γ -limit maps and

$$\eta_1 > \eta_2 > \eta_3 > \dots \quad \square$$

References

- [1] Juan Miguel Alonso, Tom Brady, Daryl Cooper, Vincent Ferlini, Martin Lustig, Michael Mihalic, Michael Shapiro, and Hamish Short. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint*, pages 3–63. World Scientific, 1991.
- [2] Gilbert Baumslag, Alexei Myasnikov, and Vladimir Remeslennikov. Algebraic geometry over groups I: Algebraic sets and ideal theory. *J. Algebra*, 219(1):16–79, 1999.

- [3] Mladen Bestvina. \mathbb{R} -trees in topology, geometry, and group theory. In *Handbook of geometric topology*, pages 55–91. Elsevier, 2002.
- [4] Mladen Bestvina and Mark Feighn. Bounding the complexity of simplicial group actions on trees. *Invent. Math.*, 103(3):449–469, 1991.
- [5] Mladen Bestvina and Mark Feighn. Stable actions of groups on real trees. *Invent. Math.*, 121(2):287–321, 1995.
- [6] Mladen Bestvina and Mark Feighn. Notes on Sela’s work: Limit groups and Makanin–Razborov diagrams. In *Geometric and cohomological methods in group theory*, volume 358 of *London Mathematical Society Lecture Note Series*, pages 1–29. Cambridge University Press, 2009.
- [7] M. Bogopolski and Victor N. Gerasimov. Finite subgroups of hyperbolic groups. *Algebra Logic*, 34(6):343–345, 1996.
- [8] Noel Brady. A note on finite subgroups of hyperbolic groups. *Int. J. Algebra Comput.*, 10(4):399–405, 2000.
- [9] Martin Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*. Springer, 1999.
- [10] Montserrat Casals-Ruiz and Ilya V. Kazachkov. On systems of equations over free products of groups. *J. Algebra*, 333(1):368–426, 2011.
- [11] François Dahmani and Vincent Guirardel. Foliations for solving equations in groups: free, virtually free and hyperbolic groups. *J. Topol.*, 3(2):343–404, 2010.
- [12] Martin Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81:449–457, 1985.
- [13] Martin Dunwoody. Groups acting on protrees. *J. Lond. Math. Soc.*, 56:125–136, 1997.
- [14] Martin Dunwoody and Roger A. Fenn. On the finiteness of higher knot sums. *Topology*, 26:337–343, 1987.
- [15] Martin Dunwoody and Michah Sageev. Jsj-splittings for finitely presented groups over slender groups. *Invent. Math.*, 135(1):25–44, 1999.
- [16] Koji Fujiwara and Panos Papasoglu. Jsj-decomposition of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.
- [17] Damien Gaboriau, Gilbert Levitt, and Frédéric Paulin. Pseudogroups of isometries of \mathbb{R} and Rips’ theorem on free actions on \mathbb{R} -trees. *Isr. J. Math.*, 87(1-3):403–428, 1994.

- [18] Etienne Ghys and Pierre de la Harpe, editors. *Sur les groupes hyperboliques d'après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser, 1990.
- [19] Daniel Groves. Limit groups for relatively hyperbolic groups. II. Makanin–Razborov diagrams. *Geom. Topol.*, 9:2319–2358, 2005.
- [20] Victor S. Guba. Equivalence of infinite systems of equations in free groups and semigroups to finite subsystems. *Mat. Zametki*, 48(3):321–324, 1986.
- [21] Vincent Guirardel. Approximations of stable actions on \mathbb{R} -trees. *Comment. Math. Helv.*, 73(1):89–121, 1998.
- [22] Vincent Guirardel. Actions of finitely generated groups on \mathbb{R} -trees. *Ann. Inst. Fourier*, 58(1):159–211, 2008.
- [23] Vincent Guirardel and Gilbert Levitt. *JSJ decompositions of groups*, volume 395 of *Astérisque*. Société Mathématique de France, 2017.
- [24] Illya Kapovich, Alexei Myasnikov, and Richard Weidmann. Foldings, graphs of groups and the membership problem. *Int. J. Algebra Comput.*, 15(1):95–128, 2005.
- [25] Michael Kapovich. Representations of polygons of finite groups. *Geom. Topol.*, 9:1915–1951, 2005.
- [26] Olga Kharlampovich and Alexei Myasnikov. Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz. *J. Algebra*, 200(2):472–516, 1998.
- [27] Olga Kharlampovich and Alexei Myasnikov. Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups. *J. Algebra*, 200(2):517–570, 1998.
- [28] Malik Koubi. Croissance uniforme dans les groupes hyperboliques. *Ann. Inst. Fourier*, 48(5):1441–1453, 1998.
- [29] Gilbert Levitt. Graphs of actions on \mathbb{R} -trees. *Comment. Math. Helv.*, 69(1):28–38, 1994.
- [30] Peter A. Linnell. On accessibility of groups. *J. Pure Appl. Algebra*, 30:39–46, 1983.
- [31] Gennadiĭ S. Makanin. Equations in a free group. *Izv. Akad. Nauk SSSR, Ser. Mat.*, 46:1188–1273, 1982.
- [32] Aleksandr A. Razborov. Systems of equations in a free group. *Izv. Akad. Nauk SSSR, Ser. Mat.*, 48(4):779–832, 1984.
- [33] Aleksandr A. Razborov. *On Systems of Equations in a Free Group*. PhD thesis, 1987.

- [34] Cornelius Reinfeldt. *Limit groups and Makanin–Razborov diagrams for hyperbolic groups*. PhD thesis, 2010.
- [35] Eliyahu Rips and Zlil Sela. Structure and rigidity in hyperbolic groups I. *Geom. Funct. Anal.*, 4(3):337–371, 1994.
- [36] Eliyahu Rips and Zlil Sela. Cyclic splittings of finitely presented groups and the canonical JSJ-decomposition. *Ann. Math.*, 146(1):53–109, 1997.
- [37] Peter Scott. The geometries of 3-manifolds. *Bull. Lond. Math. Soc.*, 15(5):401–487, 1983.
- [38] Zlil Sela. Acylindrical accessibility. *Invent. Math.*, 129(3):527–565, 1997.
- [39] Zlil Sela. Endomorphisms of hyperbolic groups. I. The Hopf property. *Topology*, 38(2):301–321, 1999.
- [40] Zlil Sela. Diophantine geometry over groups. I: Makanin–Razborov diagrams. *Publ. Math., Inst. Hautes Étud. Sci.*, 93:31–105, 2001.
- [41] Zlil Sela. Diophantine geometry over groups. VII: The elementary theory of a hyperbolic group. *Proc. Lond. Math. Soc.*, 99(1):217–273, 2009.
- [42] Jean-Pierre Serre. *Trees*, volume 146. Springer, 1980.
- [43] Charles T. C. Wall. Poincaré Complexes: I. *Ann. Math.*, 86:213–245, 1967.
- [44] Richard Weidmann. The Nielsen method for groups acting on trees. *Proc. Lond. Math. Soc.*, 85(1):93–118, 2002.
- [45] Richard Weidmann. On accessibility of finitely generated groups. *Q. J. Math.*, 63(1):211–225, 2012.
- [46] Heiner Zieschang, Elmar Vogt, and Hans-Dieter Coldewey. *Surfaces and planar discontinuous groups*, volume 835 of *Lecture Notes in Mathematics*. Springer, 1980.

RICHARD WEIDMANN
Mathematisches Seminar
Christian-Albrechts-Universität zu Kiel
Ludewig-Meyn Str. 4
24098 Kiel
GERMANY
weidmann@math.uni-kiel.de

CORNELIUS REINFELDT
gestigon GmbH
Maria-Goeppert-Str. 17
23562 Lübeck
GERMANY
mailmir@conninho.de