Beyond Limit Groups: Formal Solutions and the Profinite Topology

Jonathan Fruchter
Merton College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2023

To my family: Rachel, Yossi, Esther and Rachel

למשפחה שלי: רחל, יוסי, אסתר ורחל

Acknowledgements

Above all, I would like to express my deepest gratitude to my supervisor Martin Bridson. I thank him for his invaluable guidance throughout my mathematical pursuits and beyond. I am grateful for the countless enjoyable conversations that we have had, in which I have learned a lot of mathematics and a lot about being a mathematician.

I am grateful to Zlil Sela and Henry Wilton, for generously giving up their time to answer my questions and for sharing new ideas with me.

My thanks also go to Nikolay Nikolov, Panagiotis Papazoglou, Davide Spriano and Richard Wade who have carefully read, and helpfully commented, on my work. I am also grateful to Giles Gardam, Daniel Groves, Ashot Minasyan and Alan Reid for insightful conversations. I wish to offer special thanks to my collaborators, Simon André and Ismael Morales, with whom I have shared the pleasure of doing mathematics.

I am particularly grateful to Dario Ascari, Shaked Bader, Nadav Gropper, Gonzalo Gonzalez de Diego and Uzu Lim (my peers at the mathematical institute) for both listening to my mathematical rambles, and for sharing their own.

I owe an enormous amount to my family, Rachel, Yossi and Esther, who always support me and put me above everything else. תודה.

Last but not least, I thank Rachel for making the past years so special, and for making Oxford feel like home.

Abstract

In this thesis we explore limit groups from two different angles. One of them is model-theoretic (wherein limit groups serve as our main tool), while the other pertains to the profinite topology on limit groups (where we uncover insightful results on limit groups and residually free groups).

First, we generalize Merzlyakov's theorem about the first-order theory of free groups to acylindrically hyperbolic groups. We consequently deduce that if G is an acylindrically hyperbolic group, and $E(G)$ denotes the unique maximal finite normal subgroup of G, then G and the HNN extension $G\dot{*}_{E(G)} = \langle G, t \mid [t, g] = 1,\ \forall g \in E(G)\rangle$ (which is simply $G * \mathbb{Z}$ if $E(G)$ is trivial) have the same $\forall\exists$-theory.

The second part of this thesis focuses on limit groups over coherent right-angled Artin groups. We prove that cyclic subgroup separability is preserved under exponential completion for groups that belong to a class that includes all coherent RAAGs and toral relatively hyperbolic groups. We thus infer that the cyclic subgroups of limit groups over coherent RAAGs are closed in the profinite topology.

In the last part of the thesis, we turn to study "classical" limit groups (over free groups), as well as residually free groups. We show that the virtual second Betti number of a finitely generated, residually free group G is finite if and only if G is either free, free abelian or the fundamental group of a closed surface. Relying on these results, and employing techniques involving rank gradients of pro-p groups, we show that direct products of free and surface groups are profinitely rigid among finitely presented, residually free groups.

Statement of Originality

I declare that the work in this thesis is, to the best of my knowledge, original and my own work, except where otherwise indicated, cited, or commonly known.

Chapters I and II are preliminary chapters, that review the relevant literature. Chapter III contains material from a collaborative project with Simon André, which was published in [4]. Chapter IV contains material that was published in [46]. Chapter V contains material from the preprint [47], which is joint work with Ismael Morales.

This thesis has not been submitted for a degree at another university.

Jonathan Fruchter
Oxford, August 6th 2023

Contents

List of Figures

Introduction

Limit groups have been studied under various guises, and in multiple contexts, over the past 75 years. Some have found limit groups interesting for their defining residual properties and close relation to free groups. Others have been drawn to them since they naturally appear in the study of the first-order logic of groups. Still others have been intrigued by limit groups for their interesting structure, and the pleasing balance that they strike between being well-understood and remaining enigmatic. This dissertation celebrates limit groups from all of these perspectives.

Chapters I and II

We open this dissertation with two introductory chapters: Chapter I which covers essential basic definitions of notions that appear throughout this thesis, and Chapter II which includes a broad (yet detailed) overview of limit groups (and other related classes of groups). Although Chapter II does not contain any new results, we thought that it was important to devote the time and space to introduce the reader to the world of limit groups. This chapter is intended to be an easy read, and it contains many uncomplicated proofs which are carried out in full detail. We hope that this chapter will serve as a first introduction to limit groups for curious readers (who might not necessarily be interested in the specifics and technicalities found in later chapters).

Chapter III

Given a group G, a natural model-theoretic question is whether or not G and $G * \mathbb{Z}$ have the same first-order theory. This problem was first considered by Tarski in the case of free groups. Around 1945, he posed the following question: "*are all non-abelian free groups elementarily equivalent?*". A positive answer to this question was given by Sela in [108] (see also [71] by Kharlampovich and Myasknikov). Then, Sela generalised this result in two directions: first, he proved in [109] that every torsion-free non-elementary hyperbolic group G is elementarily equivalent to $G * \mathbb{Z}$. A few years later, he established the same result in the case where G is a non-trivial free product, different from the infinite dihedral group $D_\infty = \mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ (see [110]). In fact, in both cases he proved the following stronger result: G is elementarily embedded into $G * \mathbb{Z}$.

All these groups (namely non-elementary hyperbolic groups and non-elementary free products) have in common the property of being *acylindrically hyperbolic*, meaning that they admit a non-elementary *acylindrical action* on a hyperbolic space (for details, we refer the reader to Section I.2). The main result of Chapter III is a partial generalisation of the above-mentioned theorems of Sela to all acylindrically hyperbolic groups (see Theorem A below). This wide class of groups was introduced by Osin in [93] in order to unify several classes of negatively-curved groups considered by different authors (in particular, see [39]). This class of groups has been intensively studied in the past few years. Examples of acylindrically hyperbolic groups include, notably, all non-elementary (relatively) hyperbolic groups, all but finitely many mapping class groups of surfaces of finite type, $\mathrm{Out}(F_n)$ for $n \geq 2$, most 3-manifold groups, all non-cyclic and directly indecomposable right-angled Artin groups and many fundamental groups of graphs of groups.

Despite the intense activity around acylindrically hyperbolic groups in geometric group theory, very little is known about the first-order theory of these groups. Dahmani, Guirardel and Osin proved that acylindrically hyperbolic groups are not superstable [39, Theorem 8.1]. We remark that superstability (or more generally, stability) is a deep model theoretic notion; however, every superstable group G admits a chain of normal subgroups $1 = H_0 \lhd H_1 \lhd \cdots \lhd H_n = G$ such that each H_{i+1}/H_i is either abelian or simple [8]. This is the criterion used by Dahmani, Guirardel and Osin in their proof.

Groves and Hull adapted some of Sela's techniques to the context of acylindrically hyperbolic groups and initiated the study of solutions of systems of equations over such groups [57]. We build on their results, extend them, and derive new results throughout Chapter III.

The results presented in Chapter III are focused around the $\forall\exists$-*(first-order) theory* of acylindrically hyperbolic groups. An $\forall\exists$-*sentence* is a first-order sentence of the form $\forall \boldsymbol{x}\ \exists \boldsymbol{y}\ \psi(\boldsymbol{x}, \boldsymbol{y})$, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are two tuples of variables, and ψ is a quantifier-free formula in these variables; the set of such sentences satisfied by a group G is called the $\forall\exists$-*theory* of G. We also say that the inclusion i of a group G into an overgroup G' is an $\exists\forall\exists$-*elementary embedding* if the following condition is satisfied: for every first-order formula of the form

$$\phi(\boldsymbol{t}) : \exists \boldsymbol{x}\ \forall \boldsymbol{y}\ \exists \boldsymbol{z}\ \psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},\boldsymbol{t}),$$

where $\psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},\boldsymbol{t})$ is a quantifier-free formula, and for every tuple $\boldsymbol{g}$ of elements of G of the same arity as $\boldsymbol{t}$, if the statement $\phi(\boldsymbol{g})$ holds in G, then $\phi(i(\boldsymbol{g}))$ holds in G'. We refer the reader to Subsection I.4, where these notions are explained in more detail. Before stating our main result, recall that every acylindrically hyperbolic group G admits a unique maximal finite normal subgroup, denoted by $E(G)$ [39, Theorem 2.24]. In what follows, $G\dot{*}_{E(G)}$ denotes the HNN extension where the stable letter acts trivially, that is the group

$$G *_{E(G)} (\mathbb{Z} \times E(G)) = \langle G, t \mid [t,g] = 1,\ \forall g \in E(G)\rangle.$$

Theorem A. *Let G be an acylindrically hyperbolic group. The canonical inclusion of G into $G\dot{*}_{E(G)}$ is an $\exists\forall\exists$-elementary embedding. In particular, G and $G\dot{*}_{E(G)}$ have the same $\forall\exists$-theory.*

Remark. *This is known to be true for hyperbolic groups (see [3]).*

Remark. *Note that if the finite group $E(G)$ is trivial, the group $G\dot{*}_{E(G)}$ is simply the free product $G * \mathbb{Z}$. If $E(G)$ is non-trivial, one easily sees that $G * \mathbb{Z}$ cannot have the same $\forall\exists$-theory as G, since the existence of a non-trivial normal finite subgroup is expressible in an $\forall\exists$-sentence.*

Remark. *As an immediate consequence of Theorem A, one recovers a result of Hull and Osin [69] stating that an acylindrically hyperbolic group G with $E(G) = 1$ is mixed identity free: for every $n \in \mathbb{N}$ and every $w \in G * F_n$ there is a homomorphism $f : G * F_n \to G$ such that $f|_G = \mathrm{Id}_G$ and $f(w) \neq 1$. Note that this holds if and only if $G \models \exists(x_1,\ldots,x_n) w(x_1,\ldots,x_n) \neq 1$, which is indeed an $\forall\exists$-sentence.*

For now, it is an open question whether Theorem A above remains true if one considers the whole first-order theories of G and $G\dot{*}_{E(G)}$ instead of the $\forall\exists$ or $\exists\forall\exists$-fragments of these theories. This question can be viewed as a broad generalisation of Tarski's problem about elementary equivalence of non-abelian free groups.

The Generalised Tarski Problem. *Let G be an acylindrically hyperbolic group.*

1. *Are G and $G\dot{*}_{E(G)}$ elementarily equivalent?*

2. *Is G elementarily embedded into $G\dot{*}_{E(G)}$?*

As mentioned before, Sela proved that the answer to both of these questions is positive, under the stronger assumption that G is a torsion-free non-elementary hyperbolic group or a non-trivial and non-dihedral free product. In all other cases, including the case of hyperbolic groups with torsion, the answer is not known. Moreover, as far as we are aware, there are no examples of finitely generated groups G that are not acylindrically hyperbolic, and such that G has the same first-order theory (or even the same $\forall\exists$-theory) as $G * \mathbb{Z}$. The question of the existence of such a group is closely related to that of the preservation of acylindrical hyperbolicity under elementary equivalence among finitely generated groups; we discuss this in Section III.8. It is worth mentioning the following corollary of Theorem A (see Proposition III.8.4).

Corollary B. *Let G be an acylindrically hyperbolic group, and let H be a group that admits a non-trivial splitting over a virtually abelian group. Suppose that G and H are elementarily equivalent (or simply that they have the same $\exists\forall\exists$-theory). Then the group H is acylindrically hyperbolic.*

Remark. *As a consequence, if there exists a group G that is not acylindrically hyperbolic and such that G and $G * \mathbb{Z}$ are elementarily equivalent, then all non-trivial splittings of G (if they exist) have sufficiently complicated edge groups. For instance, if G is a generalized Baumslag-Solitar group, then G and $G * \mathbb{Z}$ are not elementarily equivalent.*

Positive theory, verbal subgroups

A first-order sentence is called *positive* if it does not involve inequalities. We say that a group G has *trivial positive theory* if every positive sentence satisfied by G is satisfied by all groups. In [84], Merzlyakov proved that non-abelian free groups have trivial positive theory. As a consequence, G has trivial positive theory if and only if it has the same positive theory as F_n, for any $n \geq 2$. Recently, in [28] and [29], Casals-Ruiz, Garreta, Kazachkov and de la Nuez González proved that many groups acting non-trivially on trees have trivial positive theory. In particular, they showed that every acylindrically hyperbolic group that acts hyperbolically and irreducibly on a tree has trivial positive theory [28, Corollary 8.2]. They also established the following quantifier elimination result [28, Theorem 6.3]: a group has trivial positive theory if and only if it has trivial positive $\forall\exists$-theory. Using this fact in conjunction with Theorem A, we prove the following result in Section III.7 (which was conjectured in [28]):

Corollary C ([28, Conjecture 9.1])**.** *Acylindrically hyperbolic groups have trivial positive theory.*

As a consequence of Corollary C, one recovers the main result of [15], due to Bestvina, Bromberg and Fujiwara (more details, and the definition of *verbal width*, appear in Section III.7).

Corollary D. *Let G be an acylindrically hyperbolic group, let $k \geq 1$ be an integer and let w be a non-trivial element of F_k. If the image of w in the abelianization $F_k/[F_k, F_k] \cong \mathbb{Z}^k$ is not primitive (i.e. it is not a proper power), then $w(G)$ has infinite width.*

It is worth noting that we do not know of any group G with non-trivial positive theory, and such that $w(G)$ has infinite width for every non-trivial w whose image in $F_k/[F_k, F_k]$ is not primitive (see [28, Section 9.8] for further discussion).

Merzlyakov's theorem

In Section III.5, we deduce Theorem A from a generalisation of Merzlyakov's theorem [84]. Assuming that a non-abelian free group F satisfies the positive first-order sentence

$$\forall \boldsymbol{x} \; \exists \boldsymbol{y} \; \Sigma(\boldsymbol{x}, \boldsymbol{y}) = 1,$$

where $\Sigma(\boldsymbol{x}, \boldsymbol{y}) = 1$ denotes a finite system of equations, Merzlyakov's theorem asserts that there exists a retraction from $\langle \boldsymbol{x}, \boldsymbol{y} \mid \Sigma(\boldsymbol{x}, \boldsymbol{y}) = 1 \rangle$ onto the free group $F(\boldsymbol{x})$ on $\boldsymbol{x}$. Upon closer inspection, this result resembles the classical implicit function theorem in the sense that it enables one to convert the relations between the tuples $\boldsymbol{x}$ and $\boldsymbol{y}$ into a function. This is why Merzlyakov's theorem is sometimes referred to as an implicit function theorem for groups. This fundamental result was one of the first steps in Sela's positive answer to Tarski's question about the elementary equivalence of non-abelian free groups.

Let us mention that previous generalisations of Merzlyakov's theorem were proved for torsion-free hyperbolic groups, for hyperbolic groups with torsion, and for π-groups (that is, pairs of the form (F, π) where $\pi : F \to Q$ is a homomorphism from a free group F to a finite group Q), respectively by Sela (see [109]), by Heil (see [67]), and by de la Nuez González (see [40]).

Given a group G and an element $g \in G$, we denote by $\mathrm{ad}(g)$ the inner automorphism $x \in G \mapsto gxg^{-1}$. Before stating our generalisation of Merzlyakov's theorem, let us introduce the following definition.

Definition. *Let G be a group, and let H be a subgroup of G. We define the subgroup $\mathrm{Aut}_G(H)$ of $\mathrm{Aut}(H)$ as follows:*

$$\mathrm{Aut}_G(H) = \{\sigma \in \mathrm{Aut}(H) \mid \exists g \in G,\ \mathrm{ad}(g)|_H = \sigma\}.$$

We prove the following version of Merzlyakov's theorem (in Section III.3, we give a more general statement allowing us to deal with finite disjunctions of finite systems of equations and inequations).

Theorem E. *Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G (called constants). Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let*

$$\Sigma(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1 \ \wedge \ \Psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) \neq 1$$

be a finite system of equations and inequations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are two tuples of variables. Let G_Σ denote the following finitely generated group, finitely presented relative to $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$:

$$\langle \boldsymbol{x},\boldsymbol{y},\boldsymbol{a} \mid R(\boldsymbol{a}) = 1,\ \Sigma(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1\rangle.$$

Let p be the arity of $\boldsymbol{x} = (x_1,\ldots,x_p)$. Suppose that G satisfies the following first-order sentence:

$$\forall \boldsymbol{x}\ \exists \boldsymbol{y}\ \Sigma(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1 \ \wedge \ \Psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) \neq 1.$$

Then, for every p-tuple $\boldsymbol{\sigma} = (\sigma_1,\ldots,\sigma_p) \in \mathrm{Aut}_G(E(G))^p$, there exists a morphism

$$\pi_{\boldsymbol{\sigma}} : G_\Sigma \to G_{\boldsymbol{\sigma}} = G *_{E(G)} \big\langle \boldsymbol{x}, E(G) \mid \mathrm{ad}(x_i)|_{E(G)} = \sigma_i,\ \forall i \in \{1,\ldots,p\}\big\rangle,$$

called a formal solution*, enjoying the following properties:*

- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \boldsymbol{x}$,
- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{a}) = \boldsymbol{a}$,
- $\Psi(\boldsymbol{x},\pi_{\boldsymbol{\sigma}}(\boldsymbol{y}),\boldsymbol{a}) \neq 1$.

Moreover, the image of $\pi_{\boldsymbol{\sigma}}$ is a subgroup of $G_{\boldsymbol{\sigma}}$ of the form

$$\langle \boldsymbol{g},\boldsymbol{a}\rangle *_{E(G)} \big\langle \boldsymbol{x}, E(G) \mid \mathrm{ad}(x_i)|_{E(G)} = \sigma_i,\ \forall i \in \{1,\ldots,p\}\big\rangle$$

for some tuple $\boldsymbol{g}$ of elements of G.

Remark. *Note that* $G_{\boldsymbol{\sigma}}$ *is isomorphic to the group* $G *_{E(G)} (F_p \times E(G))$ *obtained from* G *by adding* p *stable letters commuting with* $E(G)$*. Indeed, from the definition of* $\mathrm{Aut}_G(E(G))$*, for every* $1 \leq i \leq p$*, there exists an element* $g_i \in G$ *such that* $\mathrm{ad}(x_i)|_{E(G)} = \mathrm{ad}(g_i)|_{E(G)}$*. It follows that* $t_i = x_i g_i^{-1}$ *commutes with* $E(G)$*.*

Remark. *This theorem captures the spirit of Merzlyakov's original theorem, in the following sense: let* $\boldsymbol{g} = (g_1, \ldots, g_p)$ *be a tuple of elements of* G*, of the same arity as* $\boldsymbol{x}$*. Let* $\boldsymbol{\sigma} = (\mathrm{ad}(g_1)|_{E(G)}, \ldots, \mathrm{ad}(g_p)|_{E(G)})$*, and let* $\varphi : G_{\boldsymbol{\sigma}} \twoheadrightarrow G$ *be the retraction that maps* x_i *to* g_i *and coincides with the identity on* G*. The homomorphism* $\varphi \circ \pi_{\boldsymbol{\sigma}}$ *from* G_Σ *to* G *maps* $\boldsymbol{x}$ *to* $\boldsymbol{g}$*. Denote by* $\boldsymbol{h}$ *the image of* $\boldsymbol{y}$ *under this homomorphism. The equalities* $\Sigma(\boldsymbol{g}, \boldsymbol{h}, \boldsymbol{a}) = 1$ *hold in* G*. In other words, just as with Merzlyakov's original theorem, the theorem above gives a mechanism for associating to every tuple* $\boldsymbol{g} \in G^p$ *another tuple* $\boldsymbol{h}$ *of the same arity as* $\boldsymbol{y}$ *such that the equalities* $\Sigma(\boldsymbol{g}, \boldsymbol{h}, \boldsymbol{a}) = 1$ *hold in* G*. However, note that the image of* $\Psi(\boldsymbol{x}, \pi_{\boldsymbol{\sigma}}(\boldsymbol{y}), \boldsymbol{a})$ *by* φ *may be trivial.*

Since the statement of Theorem E is rather complicated, we hope to make it easier to digest through the following examples, which describes two formal solutions in a simple context:

Example. *Let* $G = \langle g_1, g_2, a \mid a^3 = 1,\ g_1 a g_1^{-1} = a,\ g_2 a g_2^{-1} = a^2 \rangle \simeq \mathbb{Z}/3\mathbb{Z} \rtimes F_2$*. Let* σ *be the automorphism of* $\langle a \rangle$ *that maps* a *to* a^2*, and let us consider the following first-order sentence, which is clearly satisfied by* G*:* $\forall x\ \exists y\ ([x,a] = [y,a]) \wedge (x \neq y)$*. By definition, one has:*

- $G_\Sigma = \langle x, y, a \mid a^3 = 1,\ [x,a] = [y,a] \rangle$,
- $G_{\mathrm{id}} = G *_{\langle a \rangle} \langle x, a \mid xax^{-1} = a \rangle$,
- $G_\sigma = G *_{\langle a \rangle} \langle x, a \mid xax^{-1} = a^2 \rangle$.

$$\begin{tikzcd}
& G_\Sigma \arrow[ld, "\pi_{\mathrm{id}}"'] \arrow[rd, "\pi_\sigma"] & \\
G_{\mathrm{id}} \arrow[rd] & & G_\sigma \arrow[ld] \\
& G &
\end{tikzcd}$$

The morphism π_{id} *can be defined by* $\pi_{\mathrm{id}}(x) = x$*,* $\pi_{\mathrm{id}}(a) = a$ *and* $\pi_{\mathrm{id}}(y) = g_1$*. The morphism* π_σ *can be defined by* $\pi_\sigma(x) = x$*,* $\pi_\sigma(a) = a$ *and* $\pi_\sigma(y) = g_2$*. Note that* $\pi_{\mathrm{id}}(G_\Sigma)$ *and* $\pi_\sigma(G_\Sigma)$ *are both isomorphic to* G*.*

Lastly, we would like to mention that the structure of the proof of Theorem E, which is quite different from Merzlyakov's original combinatorial proof, is inspired by Sela's geometric proof of Merzlyakov's theorem [107]. Nevertheless, both proofs rely crucially on small cancellation theory (combinatorial in one case, geometric in the other case).

Chapter IV

Casals-Ruiz, Duncan and Kazachkov defined in [30] a new class of groups $\mathcal{C}$, which was carefully designed to serve as a general framework for studying limit groups over (coherent) right-angled Artin groups (RAAGs). They succeed in showing that, as in the classical case of limit groups over free groups, the limit groups over any coherent right-angled Artin group G embed in the $\mathbb{Z}[t]$-completion of G (a group containing G in which one can interpret exponents with respect to elements in $\mathbb{Z}[t]$, see Section IV.1.2 for more details), and that this completion can be built by repeatedly extending centralisers. Casals-Ruiz, Duncan and Kazachkov [30] then pose the question: are cyclic subgroups of limit groups over coherent RAAGs closed in the profinite topology? In Chapter IV we give a positive answer to this question.

Constructing completions of groups using extensions of centralisers is not a new idea; in the 1960s, Baumslag gave a construction of the $\mathbb{Q}$-completion of a group with unique roots, which uses extensions of centralisers. Thirty odd years later, as a first step towards showing that all non-abelian free groups share the same first order theory, Sela investigated the structure of limit groups (over free groups) in [106] and showed that they admit a hierarchical structure. This structure implies that limit groups embed in the $\mathbb{Z}[t]$-completion of a free group, or in other words $F^{\mathbb{Z}[t]}$ serves as a universe for limit groups over the free group F. Similar results were obtained by Kharlampovich and Miasnikov, and in [72] they extended their argument to prove that the limit groups over any toral relatively hyperbolic group G embed in $G^{\mathbb{Z}[t]}$.

The latter work was carried out in the context of groups in the class CSA: groups whose maximal abelian subgroups are malnormal. The authors showed that extending a centraliser of a CSA group yields a CSA group, which allowed them to use induction in the process of repeatedly extending centralisers. RAAGs, even coherent ones, do not necessarily lie in the class CSA. This deficit motivated Casals-Ruiz, Duncan and Kazachkov to seek a broader setting in which similarly structured proofs would work, leading them to the class $\mathcal{C}$. They show in [30] that this class $\mathcal{C}$ contains all coherent RAAGs, as well as all toral relatively hyperbolic groups. In addition, they prove that if G is in the class $\mathcal{C}$ (and satisfies the technical condition R of Definition IV.2.5) then $G^{\mathbb{Z}[t]}$ can be built as an iterated centraliser extension over G and is fully residually G. They also show that if G is a coherent RAAG, then every limit group over G embeds in $G^{\mathbb{Z}[t]}$. They suggested that this fact, combined with the relatively simple structure of $G^{\mathbb{Z}[t]}$, ought to provide fertile ground for addressing algorithmic problems and establishing residual properties of limit groups over RAAGs. They also highlighted some specific challenges of this type, which we address in Chapter IV.

As well as building on [30], the results that we shall present here rely crucially on the simple observation that direct extensions of centralisers, which are the basis for the construction of $\mathbb{Z}[t]$-completions, have a much tamer structure than more general amalgamated products. By definition, if $C_G(u)$ denotes the centraliser of u in G, then the *direct extension* of $C_G(u)$ by $C_G(u) \times B$ is the quotient of the free product $G * B$ by relations that force the subgroups $C_G(u)$ and B to commute. This is an example of a *free product with commuting subgroups*: such a product is obtained from pairs of groups $L \leq G$ and $M \leq H$ by forming the amalgamated free product $G *_L (L \times M) *_M H$, which is abbreviated to $\langle G, H | \, [L, M] = 1 \rangle$. We shall exploit the way in which $G^{\mathbb{Z}[t]}$ is built from extensions of centralizers to prove the following theorems, repeatedly employing a criterion developed by Loginova [79] for the cyclic subgroup separability of certain free products with commuting subgroups. When G is a coherent RAAG, the special combinatorial structure of its defining graph (see Subsection IV.1.1) also plays an important part in our proof.

Theorem F. *Let G be a group in the class $\mathcal{C}$ which satisfies condition* R*, and let A be a ring. If G is cyclic subgroup separable, then the A-completion of G, G^A, is cyclic subgroup separable.*

Theorem G. *Limit groups over coherent RAAGs are cyclic subgroup separable.*

In the classic setting of limit groups over free groups, Wilton showed that all finitely generated subgroups are separable [116]. However, limit groups over coherent RAAGs are not necessarily subgroup separable: the coherent RAAG given by the presentation $L = \langle x, y, z, w | \, [x, y] = [y, z] = [z, w] = 1 \rangle$, and whose defining graph is

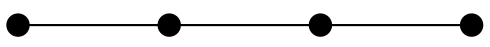

is not subgroup separable [91, Theorem 1.2].

In the last section of Chapter IV, Section IV.3, we shall discuss the word problem for free products with commuting subgroups. Given $L \leq G$ and $M \leq H$, if the word problem is solvable both in G and in H, and the membership problem is solvable for L in G and for M in H, then there is a solution to the word problem in

$$\langle G, H | \, [L, M] = 1 \rangle.$$

We shall use once again the way in which exponential completions of groups from the class $\mathcal{C}$ can be built by iterating centraliser extensions to prove the following:

Proposition H. *Let G be a group in the class $\mathcal{C}$. If G satisfies condition* R *and has a solvable word problem, then every finitely generated subgroup H of G^A has a solvable word problem.*

In particular, we recover the fact that limit groups over coherent RAAGs and toral relatively hyperbolic groups have a solvable word problem.

Chapter V

In Chapter V, the final chapter of this dissertation, we focus on *rigidity* within the class of limit groups (and the broader class of residually free groups). In particular, a recurring theme in this chapter is to utilize the homology of finite-index subgroups of a given group, in order to recognize whether it is a free or a surface group. We do so within the following classes of groups: limit groups, (finitely generated) residually free groups and word-hyperbolic fundamental groups of graphs of free groups amalgamated along cyclic subgroups. For abbreviation and readability, we refer to the latter as **HGFC-groups** (which stands for **H**yperbolic **G**raphs of **F**ree groups with **C**yclic edge groups).

Given a finitely generated group G and a field k (endowed with a trivial G-action), we denote the *i-th Betti number* of G with coefficients in k by $b_i^k(G) = \dim_k H_i(G;k)$. We focus on the *virtual i-th Betti number* of G, which allows us to study G by looking at the homology of all of its finite-index subgroups at once.

Definition. *Let G be a finitely generated group and let k be a field. The* virtual i-th Betti number *of G with coefficients in k is given by*

$$\mathrm{vb}_n^k(G) = \sup\{\dim_k H_i(H;k) \,|\, H \textit{ is a finite-index subgroup of } G\}$$

Remark. *When $k = \mathbb{Q}$, we simply write $b_i(G)$ and $\mathrm{vb}_i(G)$ to denote $b_i^{\mathbb{Q}}(G)$ and $\mathrm{vb}_i^{\mathbb{Q}}(G)$ respectively.*

One of many remarkable properties of finitely generated residually free groups, is that many group-theoretic properties can be read from their (virtual) homology. For example, the failure of a finitely generated residually free group G to satisfy certain finiteness properties, can be detected in the homology of finite-index subgroups of G [25, Theorem B]. This makes residually free groups a compelling object to study via virtual properties, and in particular via virtual homology. For more detail, we refer to Subsection II.1.6, where we give a brief overview of different properties of residually free groups which often unravel after passing to a finite index subgroup.

Virtual homology of residually free groups

Wilton showed in [120] that every one-ended HGFC-group contains a surface subgroup, and deduced that its virtual second Betti number, with coefficients in any field k, is positive. We extend this result and give a full classification of HGFC-groups by their virtual second Betti number. This classification is remarkably simple: apart from the obvious cases, the virtual second Betti number of an HGFC-group is infinite.

HGFC-groups also play a special role in the study of limit groups. Limit groups exhibit a *cyclic hierarchical structure*: they can be built up from easy-to-understand building blocks (namely free and free abelian groups), repeatedly taking cyclic amalgams and HNN extensions. In particular, HGFC-groups lie (almost) at the bottom of this hierarchical structure of a limit group. We also derive a classification for limit groups, and more generally for finitely generated residually free groups, in terms of their virtual second Betti number.

Theorem I. *Let G be a finitely generated residually free group, or a hyperbolic fundamental group of a finite graph of finitely generated free groups with infinite cyclic edge subgroups, and let k be a field. Then*

1. $\mathrm{vb}_2^k(G) = 0$ *if and only if G is free,*
2. $\mathrm{vb}_2^k(G) = 1$ *if and only if $G \cong \pi_1(\Sigma)$ where Σ is a closed, connected surface,*
3. $\mathrm{vb}_2^k(G) = \binom{d}{2}$ *if and only if $L \cong \mathbb{Z}^d$ (for $d > 2$)*
4. $\mathrm{vb}_2^k(G) = \infty$ *otherwise.*

Therefore, free, surface and free abelian groups are characterized among limit groups by their virtual second Betti number.

Remark. *By slightly modifying the proof of Theorem I, and considering maps induced on cohomology groups, the same classification holds if one replaces* $\mathrm{vb}_2^k(G)$ *by the following cohomological analogue:*

$$\sup\{\dim_k(H^2(H;k))\,|\,H \le G \textit{ of finite index}\}.$$

The idea behind the proof of Theorem I is that once a surface subgroup of infinite index appears in a limit group L, one can employ structural properties of L to replicate it and produce finite-index subgroups of L containing many surface subgroups, independent from each other in second homology. The techniques used for producing such subgroups come from the covering theory of graphs of spaces, and the methods are often inspired by previous works of Wilton [119, 120].

In fact, similar methods can be used to calculate the virtual homology of limit groups in all dimensions. In dimension 3 or higher, the proof relies on the fact that the cohomological dimension of a non-free limit group L is equal to $\max(2, k)$, where k is the maximal rank of a free abelian subgroup of L. By replicating free abelian subgroups of L within its finite-index subgroups we obtain:

Proposition J (Proposition V.2.2)**.** *Let* L *be a limit group and let* k *be a field. Then for any* $n \ge 3$, $\mathrm{vb}_n^k(L) < \infty$ *if exactly one of the following statements holds:*

1. $\mathrm{cd}(L) < n$, *and hence* $\mathrm{vb}_n^k(L) = 0$, *or*
2. L *is free abelian of rank at least* n, *and hence* $\mathrm{vb}_n^k(L) = \binom{\mathrm{rank}(L)}{n}$.

This allows us to describe the structure of closed, aspherical manifolds with a residually free fundamental group. We remark that Wilton classified which (prime, compact) 3-manifolds with incompressible toral boundary admit a residually free fundamental group; this in particular implies the case where $n = 3$ below [117, Main Theorem and Proposition 2.4]. We suspect that experts are probably also familiar with the rest of the description (see Remark V.2.5), but we record it as it easily follows from Proposition J.

Corollary K. *Let* M *be an aspherical, closed manifold of dimension* n*. Then:*

1. *if* $\pi_1(M)$ *is fully residually free and* $n \ge 3$ *then* $M \cong T^n$.
2. *if* $\pi_1(M)$ *is residually free and* $n \ge 5$ *then* M *has a finite cover that is homeomorphic to the direct product of a torus and finitely many closed surfaces.*

Virtual Betti numbers of finitely generated groups have been studied before on numerous occasions: in [22], Bridson and Kochloukova show that the virtual first Betti number of a finitely presented nilpotent-by-abelian-by-finite group is finite, and in [74], Kochloukova and Mokari show that the same holds for some abelian-by-polycyclic groups. In contrast, Agol proved that fundamental groups of closed, hyperbolic 3-manifolds are large (that is, they virtually surject onto a non-abelian free group) [2, Theorem 9.2], which implies that their virtual first Betti number is infinite. It is worth mentioning that preceding works of Agol [1], Venkataramana [114] and Cooper, Long and Reid [35] showed that under some conditions, the virtual first Betti number of arithmetic hyperbolic 3-manifold groups is infinite. In Section V.4.1 we give more examples of calculations of virtual Betti numbers of groups.

Another related work is [23], in which Bridson and Kochloukova show that the second L^2*-Betti number* of a limit group L is always 0. By Lück's approximation theorem, the second L^2-Betti number of a limit group L is equal to $\lim_{n\to\infty}\frac{\dim H_2(H_n;\mathbb{Q})}{[L:H_n]}$ where $(H_n)_{n\in\mathbb{N}}$ is any nested sequence of normal subgroups of finite-index in L with trivial intersection, i.e. $\bigcap_n H_n = \{1\}$. From the proof of Theorem I, one can easily produce a nested sequence (G_n) of finite-index normal subgroups of a given limit group L, with $\lim_{n\to\infty}[L:G_n] = \infty$, satisfying $\liminf_{n\to\infty}\frac{\dim H_2(G_n;\mathbb{Q})}{[L:G_n]} > 0$. However, for the reasons above, such a sequence (G_n) can never have trivial intersection.

The surface group conjectures

We recall a duo of famous conjectures revolving around surface groups, inspired by a question of Mel'nikov's. A group G is called a *Mel'nikov group* if it is a non-free infinite one-relator group, all of whose finite-index subgroups are one-relator groups (see [50, Definition 1.1 and following discussion]). The Surface Group Conjectures, which appear in numerous papers, claim the following:

Surface Group Conjecture A. *Let G be a residually finite Mel'nikov group. Then G is either a surface group or $G \cong \mathrm{BS}(1,n)$ (for some $n \neq 0$).*

Surface Group Conjecture B. *Let G be a Mel'nikov group, all of whose infinite-index subgroups are free. Then G is a surface group.*

It is worth noting that Fine, Kharlampovich, Myasnikov, Remeslennikov, and Rosenberger posed a third Surface Group Conjecture [44, Surface Group Conjecture C], which states that a non-free and freely indecomposable limit group, all of whose infinite-index subgroups are free, must be a surface group. The Surface Group Conjecture C was resolved by Wilton [119, Corollary 5], and it follows from [119, Theorem 3] which states that every one-ended HGFC-group which is not a surface group contains a one-ended infinite-index subgroup. Building on Wilton's result, Ciobanu, Fine and Rosenberger further confirmed that the Surface Group Conjecture B holds for a slightly larger class of cyclic amalgamations and HNN extensions of free groups [34, Theorem 3.1]. Recently, Gardam, Kielak and Logan established that both Surface Group Conjecture A and Surface Group Conjecture B are true for two-generated one-relator groups [50, Theorem 1.5]. Surface Group Conjectures A and B also appear in Baumslag's, Fine's and Rosenberger's book [11, Conjecture 2.17 and Question 2.18].

Since one-relator groups have second Betti number at most one [81], Theorem I confirms Surface Group Conjecture A in the residually free case:

Corollary L. *Let G be a residually free group all of whose finite-index subgroups are one-relator groups. Then G is either a free group or the fundamental group of a closed surface Σ with $\chi(\Sigma) \leq 0$. In particular, every residually free Mel'nikov group is a surface group.*

In fact, we show that a stronger statement holds: if G is a finitely generated residually free group all of whose finite-index subgroups H have a one-relator pro-p completion $H_{\widehat{p}}$, then G must be either free or surface. This assertion (see Theorem V.3.2) is stronger than Corollary L and will be used in the proof of Theorem N.

Profinite rigidity within the class of limit groups

From a logician's point of view, limit groups are notoriously hard to distinguish from one another: they all have the same universal first-order theory. This makes the question of telling limit groups from one another a compelling one. A key ingredient in studying profinite invariants of limit groups is *cohomological goodness*, an important notion introduced by Serre in [111]. A group G is *cohomologically good* (or just *good*) if for every finite G-module M, the map $H^n(\widehat{G}; M) \to H^n(G; M)$ on cohomology groups, induced by the canonical map $G \to \widehat{G}$, is an isomorphism. Grunewald, Jaikin-Zapirain and Zalesskii showed that limit groups are good [59, Theorem 1.3]; the same holds for HGFC-groups [66,76]. Wilton used this to show that free and surface groups are profinitely rigid among limit groups in [120, Corollary D] and [121, Theorem 2], respectively. Theorem I gives an alternative proof of the latter theorem.

Corollary M. *Let $\pi_1(\Sigma)$ be the fundamental group of a closed, connected surface and let G be a limit group or a hyperbolic group that splits as a graph of free groups with cyclic edge groups. If $\widehat{G} = \widehat{\pi_1(\Sigma)}$ then $G \cong \pi_1(\Sigma)$.*

Proof. Since $\widehat{G} \cong \widehat{\pi_1(\Sigma)}$ and G and $\pi_1(\Sigma)$ are both cohomologically good,

$$\sup\{\dim_{\mathbb{F}_p}(H^2(H;\mathbb{F}_p)) | H \leq G \text{ of finite index}\} = 1.$$

Theorem I and the aforementioned Remark imply that G is a surface group. Surface groups are distinguished from one another by their abelianizations, and therefore by their finite quotients. Hence $G \cong \pi_1(\Sigma)$. □

Due to the nature of the classification of limit groups and HGFC-groups by their virtual Betti numbers, one cannot deduce further profinite rigidity results by looking at the virtual homology of limit groups (apart from the obvious cases, when G is free or free abelian). In the end of this dissertation, in Section V.4, we discuss more refined virtual homological invariants of groups which might aid in profinite recognition of more complicated limit groups.

Profinite rigidity and direct products

A celebrated result of Platonov and Tavgen [96] states that a direct product of two non-abelian free groups is not profinitely rigid: given two non-abelian free groups F and F', there exist (infinitely many non-isomorphic) finitely generated subgroups $H \leq F \times F'$ of infinite index such that the inclusion $H \hookrightarrow F \times F'$ induces an isomorphism of profinite completions. In Subsection II.1.6 we give one such example (see Lemma II.1.63). Since a finitely presented subgroup of $F \times F'$ that intersects both factors non-trivially and maps onto each factor must be of finite index [13], there are no such examples $H \leq F \times F'$ as above where H is finitely presented. This raises the question of whether direct products of free groups are profinitely rigid among finitely presented, residually free groups. Utilizing structural results about finitely generated residually free groups, as well as methods involving *rank gradients* in pro-p groups (for more details, see Subsection V.3.2), we prove the following stronger result:

Theorem N (Theorem V.3.17)**.** *Let G be a finitely presented residually free group. Let $S_1, \dots, S_n$ be free or surface groups and let $\Gamma = S_1 \times \cdots \times S_n$ be their direct product. If $\widehat{G} \cong \widehat{\Gamma}$, then $G \cong \Gamma$.*

This resolves a particular case of Bridson's conjecture [24, Conjecture 7] that direct products of free groups are profinitely rigid among finitely presented, residually finite groups.

I

Preliminaries

In this preliminary chapter we lay out foundational groundwork, encompassing essential basic notions and definitions. These include hyperbolic and acylindrically hyperbolic groups, profinite groups and the profinite topology on a group, as well as a brief introduction to first-order logic and model theory.

I.1 Conventions, notations and basic definitions

For a group G generated by a (not necessarily finite) set S, the *word length* $|g|_S$ of an element $g \in G$ is the length of the shortest word in $S \cup S^{-1}$ representing g in G. We denote the Cayley graph of G (with respect to S) by $X(G,S)$ (or just X when it is clear who are G and S) and regard X as a metric space by setting $d_S(g,h) = |g^{-1}h|_S$. The ball of radius R in $X(G,S)$ is denoted by $B_R(G,S)$ (or, again, just $B_R(G)$ when the generating set S is clear from the context). Throughout this paper, all groups acting on metric spaces act by isometries, and all metric spaces are *geodesic*, meaning that every two points x and y are connected by a path whose length coincides with $d(x,y)$.

F will usually denote a (finitely generated) free group; if S is any set, we denote the free group on S by $F(S)$. Also, if F is known to be of rank n, we denote this by a subscript F_n. We recall the definitions of *hyperbolic spaces*, *real trees* and *hyperbolic groups*:

Definition I.1.1. A geodesic metric space (X,d) is called *δ-hyperbolic* if every geodesic triangle $\Delta = (x,y,z)$ in X is *δ-slim*: every side of Δ is contained in the closed δ-neighbourhood of the union of the two other edges. The space (X, d) is called *hyperbolic* if it is δ-hyperbolic for some δ.

Definition I.1.2. A geodesic metric space (X, d) is a *real tree* (or an $\mathbb{R}$-*tree* if it is 0-hyperbolic.

Remark I.1.3. If (X, d) is 0-hyperbolic, then every geodesic triangle is in fact a tripod. In particular, every three points $x, y, z \in X$ admit a *join* point ρ that lies on each of the segments $[x, y], [x, z], [y, z]$. We also remark that every simplicial tree is a real tree, and a real tree (X, d) is a simplicial tree if and only if its set of *branch points* (that is, points whose complement has more than two connected components) is closed and discrete.

Definition I.1.4. We say that a group G is *hyperbolic* if it is finitely generated, and it admits an action $G \curvearrowright (X, d)$ by isometries on a hyperbolic space satisfying:

1. *Cocompactness*: the quotient X/G, endowed with the quotient topology, is compact;
2. *Proper discontinuity*: for every compact subset $K \subset X$, $|\{g \in G | g.K \cap K \neq \varnothing\}| < \infty$.

Remark I.1.5. If G is generated by a finite tuple S, then the action $G \curvearrowright X(G, S)$ always satisfies the cocompactness and proper discontinuity conditions above. Therefore, G is hyperbolic if and only if it admits a finite generating tuple S such that $(X(G, S), d_S)$ is hyperbolic. We remark that hyperbolicity is easily seen to be a quasi-isometry invariant, and $(X(G, S), d_S)$ is quasi-isometric to $(X(G, S'), d_{S'})$ for any finite generating tuples S and S' of G. Therefore if G is hyperbolic then any Cayley graph of G (with respect to a finite generating set) is hyperbolic.

We continue with a brief review of the types of isometries of a hyperbolic space. Let (X, d) be a δ-hyperbolic space, and recall that the *Gromov boundary* of X, ∂X, is defined as the collection of equivalence classes of quasi-isometric embeddings $\mathbb{N} \to X$ (where two embeddings are equivalent if their images lie at bounded Hausdorff distance from one another). Every isometry $f : X \to X$ belongs to exactly one of the following three types:

1. *Elliptic*, if the orbit of some $x \in X$ by f, $\{f^n(x) | n \in \mathbb{Z}\}$, is bounded; this is equivalent to the orbit of any $x \in X$ by f being bounded.

2. *Parabolic*, if it is not elliptic, and the orbit of some (equivalently any) $x \in X$ by f has a unique accumulation point on ∂X.

3. *Loxodromic* (or *hyperbolic*), if it is not elliptic, and the orbit of some (equivalently any) $x \in X$ by f has precisely two accumulation points on ∂X.

If G is a hyperbolic group, it is well-known that the classification of elements in G takes an even simpler form:

Lemma I.1.6. *Let G be a hyperbolic group generated by a finite set S, and consider its natural action on its Cayley graph $X(G,S)$. Then $g \in G$ is*

1. *elliptic, if and only if g has finite order.*

2. *loxodromic, if and only if g has infinite order.*

Remark I.1.7. We remark that g is loxodromic if and only if for some (equivalently, any) $x \in X$, the map $\mathbb{Z} \to X$ defined via $m \mapsto g^m x$ is a quasi-isometric embedding.

We also recall the definition of *graphs of groups*:

Definition I.1.8. A *graph of groups* $\mathcal{G}$ consists of the following data:

- an oriented graph $(\Gamma, \iota, \tau, \cdot^{-1})$ where $\iota : \mathrm{E}(\Gamma) \to \mathrm{V}(\Gamma)$ maps each edge e to its initial vertex $\iota(e)$, $\tau : \mathrm{E}(\Gamma) \to \mathrm{V}(\Gamma)$ maps each each e to its terminal vertex $\tau(e)$ and $\cdot^{-1} : \mathrm{E}(\Gamma) \to \mathrm{E}(\Gamma)$ is a fixed-point free involution that satisfies $\iota(e) = \tau(e^{-1})$ for every $e \in \mathrm{E}(\Gamma)$;

- for each vertex $v \in \mathrm{V}(\Gamma)$, a choice of a *vertex group* G_v;

- for each edge $e \in \mathrm{E}(\Gamma)$, a choice of an *edge group* G_e, as well as two monomorphisms $\varphi_e^{\iota} : G_e \to G_{\iota(e)}$ and $\varphi_e^{\tau} : G_e \to G_{\tau(e)}$.

If Γ in Definition I.1.8 above is finite and connected, we can associate a fundamental group to $\mathcal{G}$. The fundamental group of $\mathcal{G}$, relative to a base point $v \in \mathrm{V}(\Gamma)$, can be defined in a similar way to a standard topological fundamental group by considering equivalence classes of loops in Γ, twisted by elements of G_v at every vertex $v \in \mathrm{V}(\Gamma)$ that appears along the loop. We give below an equivalent definition, in the form of a group presentation, that will be used in later chapters.

Definition I.1.9. Let $\mathcal{G}$ be a graph of groups as in Definition I.1.8 with a finite and connected underlying graph Γ, and let T be a spanning subtree of Γ. Fix presentations $\langle X_v | R_v \rangle$ and $\langle X_e | R_e \rangle$ for every vertex group G_v and edge group G_e of $\mathcal{G}$. The *fundamental group of $\mathcal{G}$ with respect to T* is the group admitting the following presentation:

$$\pi_1(\mathcal{G}, T) = \left\langle \bigcup_{v \in \mathrm{V}(\Gamma)} X_v \cup \{t_e,\ e \in \mathrm{E}(\Gamma) \smallsetminus \mathrm{E}(T)\} \;\middle|\; \bigcup_{v \in \mathrm{V}(\Gamma)} R_v \cup R \right\rangle$$

where R is a (possibly infinite) set of relations which includes relations of two types:

- relations which identify the set of generators X_e of the edge group G_e with their images in the adjacent vertex groups under φ_e^ι and φ_e^τ whenever $e \in \mathrm{E}(T)$, and
- relations of the form $t_e^{-1} \varphi_e^\iota(x_e) t_e = \varphi_e^\tau(x_e)$ where $e \in \mathrm{E}(\Gamma) \smallsetminus \mathrm{E}(T)$ and$x_e \in X_e$.

Remark I.1.10. The resulting group $\pi_1(\mathcal{G}, T)$ does not depend on the spanning subtree T, and we will therefore omit T and denote it by $\pi_1(\mathcal{G})$. It is also clear from the definition above that the fundamental group of $\mathcal{G}$ can be obtained from the vertex and edge groups by repeatedly taking amalgamated products, followed by a sequence of HNN extensions.

I.2 Acylindrically hyperbolic groups

The aim of this Section is to familiarize the reader, without going into great depth, with acylindrically hyperbolic groups.

Let G be a group acting on a δ-hyperbolic space (X, d); as discussed in the previous section, every hyperbolic element $g \in G$ has exactly two limit points $g^{+\infty}$ and $g^{-\infty}$ on ∂X. These are represented by the quasi-isometric embeddings $n \mapsto g^n x$ and $n \mapsto g^{-n} x$ (for some $x \in X$) respectively.

Definition I.2.1. Two hyperbolic elements g and h in a group G acting on a δ-hyperbolic space (X, d) are called *independent* if $\{g^{\pm\infty}\} \cap \{h^{\pm\infty}\} = \varnothing$. We call the action of G on X *non-elementary* if there are two (or equivalently, infinitely many) independent hyperbolic elements in G.

The notion of an *acylindrical* group action on a metric space was first introduced by Bowditch in [21], and was inspired by Sela's notion of a *k-acylindrical* group action on a tree: a group action on a tree is called k-acylindrical if it contains no arcs of length greater than k which are fixed by a non-trivial element of the group (and hence, the tree contains no "cylinders"). This notion was later generalized by imposing a bound on the cardinality of a subgroup which fixes an arc of length greater than k in the tree, and coarsified in the following manner.

Definition I.2.2. A group action on a metric space $G \curvearrowright (X,d)$ is called *acylindrical* if for every $\varepsilon \geq 0$ there exist $N > 0$ and $R > 0$ such that for every $x, y \in X$ satisfying $d(x,y) \geq R$,

$$|\{g \in G \mid d(x,gx) \leq \varepsilon \text{ and } d(y,gy) \leq \varepsilon\}| \leq N.$$

Lemma I.2.5 which appears in [39] is one of the most important structural results concerning acylindrical group actions on hyperbolic spaces. In order to state it, we first define quasi-geodesic axes of loxodromic isometries:

Definition I.2.3. Let f be a loxodromic isometry of a δ-hyperbolic space (X,d). Let $x \in X$, and let ℓ be a geodesic connecting x and $f(x)$. $\ell_x = \bigcup_{n\in\mathbb{Z}} f^n(\ell)$ is called a *quasi-geodesic axis of f* (based at x).

Remark I.2.4. Note that ℓ_x contains $f^n(x)$ for every $n \in \mathbb{Z}$. Also, ℓ_x is a quasi-geodesic for every $x \in X$.

Lemma I.2.5 ([39, Lemma 6.5, Corollary 6.6])**.** *Let G be a group acting acylindrically on a hyperbolic space X and let $g \in G$ be a hyperbolic element. Then g is contained in the unique maximal and virtually cyclic subgroup $\Lambda(g)$ which consists of all $h \in H$ for which the Hausdorff distance between ℓ and $h\ell$ is finite, where ℓ is some quasi-geodesic axis of g in X. In addition, the following are equivalent for any $h \in G$:*

1. $h \in \Lambda(g)$.
2. $h^{-1}g^m h = g^k$ *for some* $0 \neq m, k \in \mathbb{Z}$.
3. $h^{-1}g^n h = g^{\pm n}$ *for some* $n \in \mathbb{N}^*$.

In addition, there exists $r \in \mathbb{N}$ such that the centralizer of g^r is given by

$$C_G(g^r) = \{h \in G \mid \exists n \in \mathbb{N},\ h^{-1}g^n h = g^n\} \subset \Lambda(g).$$

Suppose now that a group G acts acylindrically on a hyperbolic space (X,d). The action of G on X falls into exactly one of three categories [93, Theorem 1.1]:

1. the action of G is elliptic, that is every G-orbit is bounded.

2. G is virtually cyclic and contains a hyperbolic element.

3. G contains two (equivalently, infinitely many) pairwise independent hyperbolic elements.

If an action falls into category (1) or (2) above, it is termed *elementary*; groups that fall into category (3) above were studied by Osin in [93], where he termed the notion of *acylindrically hyperbolic groups*:

Definition I.2.6 ([93, Definition 1.3 and Theorem 1.2.(AH$_2$)])**.** A group G is said to be *acylindrically hyperbolic* if it admits a non-elementary and acylindrical action on a hyperbolic space.

As a matter of fact, we can always choose the hyperbolic space on which G acts to be a simplicial graph:

Theorem I.2.7 ([93, Theorem 1.2])**.** *If G is acylindrically hyperbolic then there exists a (not necessarily finite) generating set S of G such that the Cayley graph X of G with respect to S is hyperbolic and such that the natural action of G on X is non-elementary and acylindrical.*

Remark I.2.8. If the group G is not hyperbolic, then the generating set S mentioned in Theorem I.2.7 above is necessarily infinite.

I.3 Profinite groups and the profinite topology

We commence by recollecting the definitions of *inverse systems* and *inverse limits*:

Definition I.3.1. Let $(I, \leq)$ be a directed poset, that is every $i, j \in I$ admit an upper bound $k \geq i, j$. An *inverse system of groups* consists of a collection of groups G_i indexed by I together with homomorphisms $f_{ji} : G_j \to G_i$ for $j \geq i$ that satisfy the following conditions:

1. for every $i \in I$, $f_{ii} = \mathrm{Id}_{G_i}$,

2. for every $k \geq j \geq i$ in I, $f_{ji} \circ f_{kj} = f_{ki}$.

The *inverse limit* of such a system is the group

$$\varprojlim (G_i)_{i \in I} = \left\{ (g_i)_{i \in I} \in \prod_{i \in I} G_i \middle| f_{ji}(g_j) = g_i \text{ whenever } j \geq i \right\}.$$

Armed with this definition, we can define *profinite groups*:

Definition I.3.2. Let $\mathcal{C}$ be a variety of finite groups (that is, a collection of finite groups that is closed under taking subgroups, finite direct products and quotients). A *pro-$\mathcal{C}$* group is a group obtained as the inverse limit of a system of groups from $\mathcal{C}$. If $\mathcal{C}$ is the variety of all finite groups, we call such a group a *profinite group*; if $\mathcal{C}$ is the variety of all finite p-groups, we call such a group a *pro-p* group.

Notation I.3.3. We will often denote profinite groups by bold letters, e.g. $\mathbf{H}$ and $\mathbf{G}$.

To each group we can associate a *profinite completion* by considering the inverse system of finite quotients of the group:

Definition I.3.4. Let G be a group, and consider the system

$$\{G/N | N \text{ is a normal finite index finite index subgroup of } G\}$$

ordered by $G/M \leq G/N$ if N is a subgroup of M, and $f_{G/N\, G/M} : G/N \to G/M$ the corresponding quotient map. The inverse limit of this system is denoted by $\widehat{G}$ and is called the *profinite completion* of G.

Remark I.3.5. We can form the pro-$\mathcal{C}$ completion of G with respect to any variety of finite groups $\mathcal{C}$ by considering only the quotients of G that lie in $\mathcal{C}$. It is also easy to see that the map $G \to \widehat{G}$ given by $g \mapsto (gN)_{N \triangleleft_{\mathrm{f.i}} G}$ is an embedding if and only if G is residually finite. The corresponding statement also hold when considering pro-$\mathcal{C}$ groups.

Notation I.3.6. We denote the pro-$\mathcal{C}$ completion of a group G by $G_{\widehat{\mathcal{C}}}$, and the pro-$p$ completion of G by $G_{\widehat{p}}$.

Note that any profinite group $\mathbf{G}$ embeds in the product of finite groups appearing in its defining inverse system; this product can be equipped with the product topology, and hence $\mathbf{G}$ inherits the structure of a topological group. This topology on $\mathbf{G}$ makes it a compact, Hausdorff, totally disconnected topological group.

When G is residually finite (and therefore embeds in $\widehat{G}$), it inherits a topological structure from $\widehat{G}$. It is easy to see that a basis around $1 \in G$ for this topology is given by the collection of finite-index normal subgroups of G. A similar definition can be made for all groups (and not necessarily residually finite ones):

Definition I.3.7. The *profinite topology* on a group G is the topology whose neighbourhood basis around 1 consists of finite-index normal subgroups of G.

Note that in the above setting, the image of G under the inclusion map $G \to \widehat{G}$ is always dense in $\widehat{G}$.

Remark I.3.8. Finite-index normal subgroups of a group G are in fact clopens: since cosets of finite-index normal subgroups are also open, the complement of such a subgroup is open as the union of finitely many such cosets.

Corollary I.3.9. *A subgroup H of G is closed in the profinite topology if and only if it is an intersection of finite-index normal subgroups.*

We seal the discussion with a brief overview of *separability properties*:

Definition I.3.10. Let G be a group and let $H \leq G$. We say that H is *separable* in G if for every $g \in G \setminus H$ there is a finite quotient $q : G \twoheadrightarrow Q$ of G such that $q(g) \notin q(H)$.

Remark I.3.11. This is equivalent to H being closed in the profinite topology on G.

Definition I.3.12. A group G is called *subgroup separable*, or *LERF* (locally extended residually finite), if every finitely generated subgroup of G is separable in G.

We remark that subgroup separability is a strong tool for generating finite-index subgroups of a given group with additional properties, as evident from the following elementary lemma which will prove to be useful in Chapter V:

Lemma I.3.13. *Let G be a subgroup separable group and let H be a finitely generated subgroup of G. Then for every finite-index subgroup $H_0 \leq H$, there is a finite-index subgroup $G_0 \leq G$ such that $G_0 \cap H = H_0$.*

Proof. H_0 is a finitely generated subgroup of G and therefore separable in G. Let $g_1, \dots, g_n$ be coset representatives of all non-trivial cosets of H_0 in H, and for every $i \leq n$ let G_i be a finite-index subgroup of G which contains the subgroup H_0 but not the element g_i. Letting $G_0 = \bigcap_{i=1}^{n} G_i$ we have that $G_0 \cap H = H_0$. □

We also define another separability property that can be seen as a strengthening of subgroup separability:

Definition I.3.14. Let G be a group and let $H \leq G$. We say that G *virtually retracts* onto H if there is a finite-index subgroup G_0 of G containing H and a retraction $r : G_0 \to H$, that is a homomorphism r such that $r(h) = h$ for every $h \in H$. In this case we say that H is a *virtual retract* of G. If G virtually retracts onto all of its finitely generated subgroups we say that G admits *local retractions*.

Remark I.3.15. If G is a residually finite group that virtually retracts onto $H \leq G$ then H is closed in the profinite topology on G: let $r : G_0 \twoheadrightarrow H$ be a retraction where G_0 is a finite-index subgroup of G and let $\{G_i\}_{i\in\mathbb{N}}$ be a collection of finite-index subgroups of G such that $\bigcap_{i\in\mathbb{N}} G_i = \{1\}$. We have that $H = \bigcap_{i\in\mathbb{N}}((\ker r \cap G_i) \cdot G_0)$ so H is the intersection of finite-index subgroups of G. In particular, if G is a residually finite group that admits local retractions, then G is subgroup separable.

I.4 Equations over groups and first-order logic

In order to discuss first-order logic in a group theoretic context, we first need to define *equations* over groups:

Definition I.4.1. A (group theoretic) *equation* in variables $\boldsymbol{x} = (x_1, \ldots, x_n)$ is simply an element $w \in F(\boldsymbol{x})$. We will sometimes refer to the equation w as $w(\boldsymbol{x})$ or $w(\boldsymbol{x}) = 1$, depending on the context. An *inequality* (or an *inequation*) is an expression of the form $w(\boldsymbol{x}) \neq 1$ where $w \in F(\boldsymbol{x})$.

Definition I.4.2. An *equation over a group G* in variables $\boldsymbol{x}$ is an element $w \in F(\boldsymbol{x}) * G$. As before, we will often refer to w as $w(\boldsymbol{x}, \boldsymbol{a})$ where $\boldsymbol{a}$ is a tuple of elements from G. An *inequality over a group G* is an expression of the form $w(\boldsymbol{x}, \boldsymbol{a}) \neq 1$.

Remark I.4.3. We can always refer to an equation $w(\boldsymbol{x})$ as an equation over a group G. In this case we say that the equation is *without parameters.*

Notation I.4.4. Let G be a group. Given a tuple of variables $\boldsymbol{x} = (x_1, \ldots, x_n)$, $w \in F(\boldsymbol{x})$ and $\boldsymbol{g} = (g_1, \ldots, g_n) \in G^n$, we write $w(\boldsymbol{g})$ to denote the element of G obtained from w by replacing each $x_i^{\pm 1}$ with the corresponding $g_i^{\pm 1}$. For $w \in F(\boldsymbol{x}) * G$, $w(\boldsymbol{g}, \boldsymbol{a})$ is defined similarly.

Armed with this notation, we can define solutions to equations over groups:

Definition I.4.5. A *solution* to an equation $w(\boldsymbol{x}, \boldsymbol{a}) = 1$ over a group G consists of a tuple $\boldsymbol{g}$ of elements from G, of the same arity as $\boldsymbol{x}$, such that $w(\boldsymbol{g}, \boldsymbol{a}) = 1$ holds in G.

It is natural to define systems of equations: given a subset $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = \{w_i(\boldsymbol{x}, \boldsymbol{a})\}_{i\in I} \subset F(\boldsymbol{x}) * G$, we refer to the conjunction $\bigwedge_{i\in I} w_i(\boldsymbol{x}, \boldsymbol{a}) = 1$ as a *system of equations* (over G). We abbreviate and write this as $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = 1$. We say that a tuple $\boldsymbol{g}$ of elements from G is a *solution to* $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = 1$ if for every $w_i(\boldsymbol{x}, \boldsymbol{a}) \in \Sigma(\boldsymbol{x}, \boldsymbol{a})$, one has $w_i(\boldsymbol{g}, \boldsymbol{a}) = 1$

in G. Just like systems of equations, systems of inequations are conjunctions of inequations; we say that a tuple $\boldsymbol{g}$ of elements from G satisfies the system of inequations $\Phi(\boldsymbol{x}, \boldsymbol{a}) \neq 1$ in G if for every $w_i(\boldsymbol{x}, \boldsymbol{a}) \in \Phi(\boldsymbol{x}, \boldsymbol{a})$, $w_i(\boldsymbol{g}, \boldsymbol{a}) \neq 1$ holds in G.

Note that there is a one-to-one correspondence between the set of solutions to the system of equations $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = 1$ over a group G and the set of homomorphisms

$$\varphi : G_\Sigma = \langle \boldsymbol{x}, \boldsymbol{a} \mid R(\boldsymbol{a}) \cup \Sigma(\boldsymbol{x}, \boldsymbol{a}) \rangle \to G$$

(where $R(\boldsymbol{a})$ is a set of relations for which $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) \rangle$ is a presentation of the subgroup of G generated by $\boldsymbol{a}$). If $\boldsymbol{g}$ is a solution to $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = 1$, there exists a homomorphism $\varphi : G_\Sigma \to G$ mapping $\boldsymbol{x}$ to $\boldsymbol{g}$ and $\boldsymbol{a}$ to $\boldsymbol{a}$; on the other hand, given such a homomorphism φ, the tuple $\varphi(\boldsymbol{x}) \in G^p$ is a solution to $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = 1$ over G. In addition, a solution $\boldsymbol{g}$ to the system of equations $\Sigma(\boldsymbol{x}, \boldsymbol{a}) = 1$ satisfies the system of inequations $\Phi(\boldsymbol{x}, \boldsymbol{a}) \neq 1$ if and only if there exists a homomorphism $\varphi : G_\Sigma \to G$ which maps $\boldsymbol{x}$ to $\boldsymbol{g}$ and $\boldsymbol{a}$ to $\boldsymbol{a}$, and such that for every $w_i(\boldsymbol{x}, \boldsymbol{a}) \in \Phi(\boldsymbol{x}, \boldsymbol{a})$, $\varphi(w_i(\boldsymbol{x}, \boldsymbol{a})) \neq 1$. Thus one often regards the study of equations (and their solutions) over G, as the study of homomorphisms from the group G_Σ to G. Equations are also closely related to the first-order logic of groups. Equations over groups are a major part in the study of the first-order logic of groups. We continue with a short introduction to model theory, from a group theoretic perspective.

First-order logic statements take the shape of *formulas* and *sentences*; to create formulas and sentences, one has to use a *language*. The language of groups has three symbols: 1, a constant which corresponds to the trivial element in a group, $\cdot(-,-)$, a 2-ary function which corresponds to group multiplication, and $(-)^{-1}$, a 1-ary function that corresponds to taking the inverse of an element. Formulas (in any language) admit a recursive definition, but to avoid complications we will define formulas as ones that are in *conjunctive normal form*. Every formula is equivalent to one in conjunctive normal form, so we do not lose any generality by restricting our point of view.

Definition I.4.6. Let $\mathbf{z}$ be a tuple of variables. A *formula* $\varphi(\mathbf{z})$ (in the variables $\mathbf{z}$) has the form

$$\forall \boldsymbol{x}^1 \exists \boldsymbol{x}^2 \cdots \forall \boldsymbol{x}^n \bigvee_{i=1}^{k} (\Sigma_i(\mathbf{z}, \boldsymbol{x}^1, \ldots, \boldsymbol{x}^n) = 1 \wedge \Psi_i(\mathbf{z}, \boldsymbol{x}^1, \ldots, \boldsymbol{x}^n) \neq 1)$$

where each $\boldsymbol{x}^i$ is a tuple of variables, each $\Sigma_i(\mathbf{z}, \boldsymbol{x}^1, \ldots, \boldsymbol{x}^n) = 1$ is a system of equations and each $\Psi_i(\mathbf{z}, \boldsymbol{x}^1, \ldots, \boldsymbol{x}^n) \neq 1$ is a system of inequations. If the tuple $\mathbf{z}$ is empty, we call the formula a *sentence* (or a *closed formula*).

Note that given a group G, one can always choose a tuple $\boldsymbol{a}$ of elements from G, of the same arity as $\mathbf{z}$, and plug it into the expression $\varphi(\mathbf{z})$ above. If the statement one obtains is true, we say that G *satisfies* $\varphi(\boldsymbol{a})$ and denote this by $G \vDash \varphi(\boldsymbol{a})$. We will often be interested in the set of all sentences that hold in a group.

Definition I.4.7. The (first-order) *theory* of a group G is the collection of all sentences φ such that $G \vDash \varphi$. If two groups G and H have the same theory, we say that they are *elementarily equivalent* and denote this by $G \equiv H$. A homomorphism $f : H \to G$ is called *elementary* if for every formula $\varphi(\mathbf{z})$ and for any tuple $\boldsymbol{a}$ in H of the same arity as $\mathbf{z}$, $H \vDash \varphi(\boldsymbol{a}) \iff G \vDash \varphi(f(\boldsymbol{a}))$.

A special case of interest is when H is a subgroup of G, that is, when the homomorphism $f : H \to G$ is an inclusion. In this case, we say that H is an *elementary subgroup* of G (or that the inclusion map is an *elementary embedding*).

As mentioned in the Introduction, we are particularly interested in $\forall\exists$-*sentences*, that is sentences of the form $\forall \boldsymbol{x} \exists \boldsymbol{y} \varphi(\boldsymbol{x}, \boldsymbol{y})$ (where $\boldsymbol{x}$ and $\boldsymbol{y}$ are tuples of variables, and φ is a formula without any quantifiers). The subset of the theory of G that consists of all $\forall\exists$-sentences of G is called the $\forall\exists$-*theory* of G. Similar to definition I.4.7 above, one can define the notions of $\forall\exists$-*elementary equivalence* or $\forall\exists$-*elementary subgroups*. The same definition can be made for any string of alternating quantifiers; for example, in Chapter III we will talk about $\exists\forall\exists$-*elementary embeddings*: embeddings $i : H \to G$ such that for every formula $\varphi(\boldsymbol{t})$ of the form $\exists \boldsymbol{x} \forall \boldsymbol{y} \exists \mathbf{w} \psi(\boldsymbol{x}, \boldsymbol{y}, \mathbf{z}, \boldsymbol{t})$ where ψ is a quantifier-free formula, and for every tuple $\boldsymbol{a}$ of elements of H (of the same arity as $\boldsymbol{t}$),

$$H \vDash \varphi(\boldsymbol{a}) \iff G \vDash \varphi(i(\boldsymbol{a})).$$

Lastly, we introduce the language of *ultrafilters*. This will enable us to phrase certain statements related to sequences with relative ease, rather than often passing to subsequences. Loosely speaking, ultrafilters simply form "strainers" or "filters", in the sense that they give us a precise manner of saying when a subset of a set X is either "small" or "large". More precisely:

Definition I.4.8. An *ultrafilter* (on $\mathbb{N}$) is a finitely additive probability measure $\omega : 2^{\mathbb{N}} \to \{0, 1\}$. Alternatively, we can think of ω as a collection of subsets of $2^{\mathbb{N}}$ which is closed under finite intersections, closed under taking supersets and maximal in the sense that it is not a proper subset of any subset of $2^{\mathbb{N}}$ that satisfies these properties.

Definition I.4.9. An ultrafilter ω is *non-principal* if it satisfies $\omega(F) = 0$ for every finite $F \subset \mathbb{N}$.

Using the language of ultrafilters is straightforward: for a statement P, depending on an index $n \in \mathbb{N}$, we say that P holds *ω-almost-surely* if

$$\omega(\{n \in \mathbb{N} \mid P \text{ holds for } n\}) = 1.$$

We also define limits with respect to ultrafilters:

Definition I.4.10. The *ω-limit* of a sequence $(x_n)_{n\in\mathbb{N}}$ in $\mathbb{R}$ is $x \in \mathbb{R}$ if for every $\varepsilon > 0$,

$$\omega(\{n \in \mathbb{N} \mid |x - x_n| < \varepsilon\}) = 1.$$

In this case we denote $\lim_\omega(x_n) = x$. We say that $\lim_\omega(x_n) = \infty$ if $\omega(\{n \in \mathbb{N} \mid x_n > N\}) = 1$ holds for every $N \in \mathbb{N}$.

Another advantage of adopting the language of ultrafilters, is the following: every sequence of real numbers has a unique ω-limit in $\mathbb{R} \cup \{\pm\infty\}$. Note that if a sequence is bounded then its ω-limit is always a real number, but the ω-limit of an unbounded sequence can be in $\mathbb{R}$.

II

Limit groups and graphs of spaces

LIMIT groups are in the limelight of this dissertation; in this chapter we will define limit groups, discuss different approaches towards them, highlight a few historical anecdotes, and recollect some of their prominent properties. We will also define graphs of spaces and review related techniques, focusing on graphs of free groups amalgamated along cyclic subgroups.

II.1 Limit groups

Limit groups have been studied under many guises, and the first works that bring limit groups to the focus were conducted, independently, by two brothers: Gilbert Baumslag in his paper *"On generalised free products"* [10] from 1962, and more notably, Benjamin Baumslag in his paper *"Residually free groups"* [9] from 1967. In the late 1990's, limit groups gained popularity due to Sela's work on Tarski's problems revolving the first-order theory of free groups ([106] *et seq.*). Limit groups have since been studied in the context of first-order logic by numerous authors, many of which are mentioned in Chapters III and IV; they were also studied on their own, for their interesting structure and properties (for example [116] and [80]). Recently, limit groups were also used as a tool outside of the study of model theory of groups: Sela and Fujiwara [49] and Fujiwara [48] used limit groups as a tool to study exponential growth rates in groups exhibiting a hyperbolic behaviour.

We begin by defining limit groups through their *residual properties*, as in the aforementioned works of Gilbert and Benjamin Baumslag:

Definition II.1.1. Let $\mathcal{C}$ be a class of groups; a group G is called *residually-$\mathcal{C}$* if for every $1 \neq g \in G$ there is a normal subgroup $N \lhd G$ such that $G/N \in \mathcal{C}$ and $g \notin N$.

Note that another way to phrase that a group G is residually-$\mathcal{C}$ is by saying that every non-trivial element of G survives in a quotient that belongs to $\mathcal{C}$.

Definition II.1.2. Let $\mathcal{C}$ be a class of groups; a group G is called *fully residually-$\mathcal{C}$* if for every finite subset $F = \{g_1, \dots, g_n\}$ of G there is a quotient $q : G \twoheadrightarrow Q$ such that $Q \in \mathcal{C}$ and the restriction $q|_F$ is injective.

Remark II.1.3. If the class $\mathcal{C}$ is closed under taking direct products, e.g. when $\mathcal{C}$ is the class of finite groups, then the notions residually-$\mathcal{C}$ and fully residually-$\mathcal{C}$ coincide: clearly, if G is fully residually-$\mathcal{C}$ it is in particular residually-$\mathcal{C}$. On the other hand, if $F = \{g_1, \dots, g_n\}$ is a finite subset of G that does not contain the trivial element, let $q_i : G \twoheadrightarrow Q_i$ $(1 \leq i \leq n)$ be a collection of n quotients of G such that g_i survives in Q_i and $Q_i \in \mathcal{C}$. The map $q_1 \times \cdots \times q_n : G^n \twoheadrightarrow Q_1 \times \cdots \times Q_n$ is injective on F and $Q_1 \times \cdots \times Q_n \in \mathcal{C}$ which implies that G is fully residually-$\mathcal{C}$. However, in other cases, and in particular when $\mathcal{C}$ is the class of (finitely generated) free groups, being residually-$\mathcal{C}$ does not imply being fully residually-$\mathcal{C}$ as we will see in Observation II.1.6.

Finally, we define:

Definition II.1.4. A group G is called a *limit group* if it is a finitely generated, fully residually free group (that is, G is fully residually $\mathcal{C}$ where $\mathcal{C}$ is the class of [finitely generated] free groups).

Finitely generated free groups are clearly limit groups (as the identity map is a map to a free group that is injective on any finite subset). Another basic example of a limit group is the following:

Example II.1.5. $\mathbb{Z}^2$ is a limit group. Let $F = \{(a_1, b_1), \dots, (a_n, b_n)\}$ be a finite subset of $\mathbb{Z}^2$. Let $k > \sum_{i=1}^{n} |a_i|$ and define $q : \mathbb{Z}^2 \twoheadrightarrow \mathbb{Z}$ by $q((1,0)) = 1$ and $q((0,1)) = k$. Now $q((a_i, b_i)) = a_i + k \cdot b_i = q((a_j, b_j)) = a_j + k \cdot b_j$ implies $a_i = a_j$ and $b_i = b_j$ which shows that $q|_F$ is injective. A similar argument shows that every finitely generated free abelian group is a limit group.

In Subsection II.1.5 we will see that these two examples, namely finitely generated free and free abelian groups, can be seen as the basic building blocks of limit groups: every limit group can be constructed from finitely generated free and free abelian groups by repeatedly taking free products, free products with infinite cyclic amalgamation and HNN extensions.

In contrast with Remark II.1.3, we observe the following:

Observation II.1.6. Not every finitely generated, residually free group is a limit group. For example, let F be a finitely generated non-abelian free group and consider $F \times \mathbb{Z}$. Each non-trivial element in $F \times \mathbb{Z}$ projects to a non-trivial element in (at least) one of the factors, so $F \times \mathbb{Z}$ is residually free. On the other hand, let $1 \neq t \in \mathbb{Z}$ and let g and h be non-commuting elements in F; if $F \times \mathbb{Z}$ was a limit group, there would exist a quotient $q : F \times \mathbb{Z} \to Q$ such that Q is a free group and q is injective on $\{(1,1), ([a,b],1), (1,t)\}$. Since $[(a,1),(1,t)] = [(b,1),(1,t)] = 1$ and Q is free, $q((a,1)), q((b,1))$ and $q((1,t))$ are all powers of an element $g \in Q$. Therefore $q(([a,b],1)) = 1 = q((1,1))$ which contradicts the assumption that q is injective on $\{(1,1), ([a,b],1), (1,t)\}$.

Remark II.1.7. The argument in Observation II.1.6 shows that limit groups are *commutative transitive* (meaning that commutativity among non-trivial elements of a limit group is a transitive relation). In addition, the example given above is minimal in the following sense: Benjamin Baumslag proved [9, Theorems 1 and 3] that for a finitely generated residually free group G, the following three are equivalent:

1. G is a limit group,
2. G is commutative transitive,
3. G does not contain $F \times \mathbb{Z}$ as a subgroup (where F is a free group of rank 2).

Sela's study of limit groups originated in studying equations over free groups; this led him to study the space of homomorphisms from a fixed finitely generated group to a free group (recall the connection between equations in a group and homomorphisms discussed in Subsection I.1). This space of homomorphisms is not compact, but it can be compactified by adding points that correspond to homomorphisms from limit groups to a free group. We explain this idea in further detail in the next subsection.

II.1.1 Limits in the space of marked groups

As part of his study of groups of polynomial growth, Gromov came up with an idea as to how to equip a set of groups with a topology [54]. This idea led Grigorchuk to define a topological *space of marked groups* in his work on groups of intermediate growth from 1985 [53]. This topological space gives a natural setting in which to define limit groups. The contents of this Subsection are loosely based on Champetier's and Guirardel's expository work on limit groups [33], and it includes a few simple proofs that should be well-known to experts but which do not appear in the literature (to the best of our knowledge).

We sealed the previous subsection by mentioning that limit groups can be seen as "points at infinity" that compactify the space of epimorphisms from a given finitely generated group G to a free group. In this subsection, we make this assertion precise.

Given a finitely generated group G with a generating tuple $S = (g_1, \dots, g_n)$, every epimorphism f from G to a free group F yields a possible generating tuple for F, namely $f(S) = \{f(g_1), \dots, f(g_n)\}$ (note that $f(S)$ does not have to be a basis for F). This suggests the following definition:

Definition II.1.8. A *marked group* is a pair of the form (G, S) where G is a group and S is a finite tuple whose elements generate G. Given $n \in \mathbb{N}$, denote by $\mathcal{G}_n$ the set of all marked groups (G, S) with $|S| = n$.

Remark II.1.9. We emphasize that S in Definition II.1.8 above is a tuple, that is S is ordered and repetitions are allowed. In addition, if $G, (g_1, \dots, g_n)$ and $G', (g'_1, \dots, g'_n)$ are two marked groups, and the map $g_i \mapsto g'_i$ extends to an isomorphism $G \to G'$, we consider $G, (g_1, \dots, g_n)$ and $G', (g'_1, \dots, g'_n)$ as the same point in $\mathcal{G}_n$.

The next step is to define a topology on $\mathcal{G}_n$; we do so by defining a metric on $\mathcal{G}_n$. Note that a priori this metric is a pseudometric, but identifying points as in Remark II.1.9 makes it a metric. Given $(G, S), (G', S') \in \mathcal{G}_n$, set $v((G, S), (G', S'))$ to be the maximal $N \in \mathbb{N} \cup \{\infty\}$ such that $w(S) = 1$ in G if and only if $w(S') = 1$ in G' for every word $w \in F_n$ of length at most N (if (G, S) and (G', S') are isomorphic as marked groups, and they represent the same point in $\mathcal{G}_n$, set $v((G, S), (G', S')) = \infty$).

Definition II.1.10. The metric $d_n : \mathcal{G}_n \times \mathcal{G}_n \to \mathbb{R}_{\geq 0}$ is given by

$$d_n((G, S), (G', S')) = e^{-v((G,S),(G',S'))}.$$

Finally, we can give an alternative definition of limit groups:

Definition II.1.11. A group G is a *limit group* if there is $n \in \mathbb{N}$ and a generating tuple S of G of size n such that (G, S) is the limit of a sequence (G_i, S_i) in $\mathcal{G}_n$ and every G_i is a free group.

It is easy to see that Definition II.1.4 implies Definition II.1.11, that is every finitely generated, fully residually free group G is a limit of free groups in the space of marked groups: given a generating tuple S of G of size n, and given $N \in \mathbb{N}$, there exists a homomorphism from G to a free group $f_N : G \to F$ which is injective on $B_N(G, S)$.

This implies that $d_n((G,S),(F,(f_N(S)))) \leq e^{-N}$, and one can approximate (G,S) as a limit of free groups in $\mathcal{G}_n$. The converse to this claim is harder to prove, and it relies on the fact that free groups are *equationally Noetherian*; we discuss this concept in Subsection II.1.4.

We now turn to explaining how limit groups compactify the set of marked free groups in $\mathcal{G}_n$. The explanation is rather elementary, but we could not find a detailed account of it in the literature. Hence we include it here in full detail.

Note that

$$\overline{\{(G,S) \in \mathcal{G}_n | G \text{ is free}\}} = \{(G,S) \in \mathcal{G}_n | G \text{ is free}\} \cup \{(G,S) \in \mathcal{G}_n | G \text{ is a limit group}\}$$

and it is therefore enough to establish that $\mathcal{G}_n$ is a compact space. To do so, we identify $\mathcal{G}_n$ with a subset of 2^{F_n}. By Tychonoff's Theorem, 2^{F_n} is compact, so it suffices to prove the following two claims:

1. The topology that $\mathcal{G}$ inherits from 2^{F_n} coincides with the topology induced by the metric d_n.
2. $\mathcal{G}$ is a closed subset of 2^{F_n}.

Endow F_n with a basis $\{x_1,\dots,x_n\}$ and identify $\mathcal{G}$ with a subset of 2^{F_n} in the following manner: every $(G, S = (s_1,\dots,s_n)) \in \mathcal{G}_n$ can be seen as a quotient of F_n through the map $q_{(G,S)} : F_n \to G$ which sends x_i to s_i. Note that (G,S) and (G',S') are isomorphic as marked groups if and only if the maps $q_{(G,S)} : F_n \to G$ and $q_{(G',S')} : F_n \to G'$ have the same kernel, and therefore $\mathcal{G}_n$ can be identified with the set of normal subgroups of F_n (which is a subset of 2^{F_n}), which we denote by $\mathcal{N}_n$.

As mentioned earlier, $\mathcal{N}_n$ inherits the product topology from 2^{F_n}; it also comes equipped with the metric that we previously defined on $\mathcal{G}$. This metric can also be described as follows: given $K, K' \in \mathcal{N}_n$, let

$$v_{\mathcal{N}}(K,K') = \max\{N \in \mathbb{N} \cup \infty | K \cap B_N(F_n,\{x_1,\dots,x_n\}) = K' \cap B_N(F_n,\{x_1,\dots,x_n\})\}.$$

We now define

$$d_{\mathcal{N}_n}(K,K') = e^{-v_{\mathcal{N}}(K,K')}.$$

Note that indeed $v((G,S),(G',S')) = v_{\mathcal{N}}(\ker(q_{(G,S)}), \ker(q_{(G',S')}))$. We next observe:

Lemma II.1.12. *The metric induced by this topology coincides with the topology that* $\mathcal{N}_n$ *inherits as a subset of* 2^{F_n} *(endowed with the product topology):*

Proof. On one hand, the open ball of radius e^{-N} around K in $\mathcal{N}_N$ is open in the topology inherited from 2^{F_n} as it can be written as

$$\mathcal{N}_n \cap \left(\prod_{g \in K \cap B_N(F_n, \{x_1,\dots,x_n\})} \{1\} \times \prod_{g \in B_N(F_n,\{x_1,\dots,x_n\}) \smallsetminus K} \{0\} \times \prod_{g \notin B_N(F_n,\{x_1,\dots,x_n\})} \{0,1\} \right).$$

On the other hand, the topology that $\mathcal{N}_n$ inherits from 2^{F_n} is generated by sets of the following form:

$$U_g = \{K \in \mathcal{N}_n | g \in K\} \quad \text{and} \quad V_g = \{k \in \mathcal{N}_n | g \notin K\}.$$

Given $g \in F_n$, let N be such that $g \in B_N(F_n, \{x_1, \dots, x_n\})$. We have

$$U_g = \bigcup_{K \in \mathcal{N}_N,\ g \in K} B_N(K) \quad \text{and} \quad V_g = \bigcup_{K \in \mathcal{N}_N,\ g \notin K} B_N(K)$$

where the balls are taken with respect to the metric $d_{\mathcal{N}_n}$ defined above. This shows that U_g and V_g are open in the topology induced by $d_{\mathcal{N}_n}$ and the two topologies on $\mathcal{N}_n$ coincide. □

Finally, we prove:

Lemma II.1.13. *$\mathcal{N}_n$ is a closed subset of 2^{F_n}.*

Proof. Let (K_n) be a sequence in $\mathcal{N}_n$ that converges to $K \in 2^{F_n}$. We will show that K is a normal subgroup of F_n and it therefore lies in $\mathcal{N}_n$. One clearly sees that K is not empty since $1 \in K_n$ for every n. Given $g, h \in K$, a tail of the sequence (K_n) must be contained in the open set $U_{g,h} = \{A \subset 2^{F_n} | g, h \in A\}$; hence for every N large enough, $gh^{-1} \in K_N$, and $K_n \in U_{gh^{-1}} = \{A \subset 2^{F_n} | gh^{-1} \in A\}$. Since K is contained in exactly one of $U_{gh^{-1}} = \{A \subset 2^{F_n} | gh^{-1} \in A\}$ and $V_{gh^{-1}} = \{A \subset 2^{F_n} | gh^{-1} \notin A\}$ we must have that $K \in U_{gh^{-1}}$. This shows that K is a subgroup of F_n.

Similarly, given $g \in K$ and $h \in F_n$, for N large enough we have that $g \in K_N$ and hence $hgh^{-1} \in K_N$; just like before, this implies that $hgh^{-1} \in K$ and K is a normal subgroup of F_n. □

We conclude this subsection by opening up another discussion, and suggesting a definition for a *(marked) limit group over a group G*; this definition is not the common definition for an (algebraic) limit group over a group G (presented in Subsection II.1.2). However, we will see that the two definitions coincide (see Lemma II.1.20).

Recall that our discussion started by considering homomorphisms from a finitely generated group H, generated by a tuple S of size n, to a free group F; these gave rise to epimorphisms from H to free groups of rank n, or equivalently markings of F_n. Replacing F by another group G, we obtain that homomorphisms from H to G correspond to marked groups (G', S') where $|S'| = n$ and G' is a subgroup of G. This suggests the following definition:

Definition II.1.14. Let G be a group. A group H is called a *marked limit group over* G if there is an integer n such that H is the limit of a sequence of marked groups $(G_i, S_i) \in \mathcal{G}_n$, and such that $G_i \leq G$ for every $i \in \mathbb{N}$.

Remark II.1.15. With this definition, a limit group is simply a marked limit group over a free group.

Therefore, a group H is a (marked) limit group over G if and only if it admits a finite generating tuple S such that $(H, S) \in \overline{\{(G', S') \in \mathcal{G}_n | G' \leq G\}} \subset \mathcal{G}_n$ for some $n \in \mathbb{N}$. With this definition, we get that as in the case of free groups, the set of subgroups of G marked by n elements can be compactified by the collection of (marked) limit groups over G. In other words, the closure of $\{(G', S') \in \mathcal{G}_n | G' \leq G\}$ is

$$\{(G', S') \in \mathcal{G}_n | G' \leq G\} \cup \{(G', S') \in \mathcal{G}_n | G' \text{ is a limit group over } G\}.$$

II.1.2 Sequences of homomorphisms

In this subsection we discuss an *algebraic* point of view towards limit groups; in the last subsection we saw that a finitely generated, fully residually free group H gives rise to a sequence of marked free groups converging to a marking of H in some $\mathcal{G}_n$. Such a group H also gives rise to a *stable sequence of homomorphisms* (with a trivial *stable kernel*) from H to a free group F. We begin by defining the notion of a stable sequence of homomorphisms:

Definition II.1.16. Let G be a group, and let H be a finitely generated group. A sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H, G)^{\mathbb{N}}$ is called *stable* if for every $g \in H$, the sequence $(\varphi_n(g))_{n\in\mathbb{N}} \in G^{\mathbb{N}}$ is eventually always 1, or eventually never 1. The *stable kernel* of $(\varphi_n)_{n\in\mathbb{N}}$ is defined as

$$\underleftarrow{\ker}((\varphi_n)_{n\in\mathbb{N}}) = \left\{g \in H | \text{ the sequence } (\varphi_n(g))_{n\in\mathbb{N}} \text{ is eventually always } 1\right\}.$$

Remark II.1.17. Note that $\underleftarrow{\ker}((\varphi_n)_{n\in\mathbb{N}})$ is a normal subgroup of H.

Given a finitely generated, fully residually free group H, one can easily construct a stable sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H, F_2)^{\mathbb{N}}$ with a trivial stable kernel: let $H_1 \subset H_2 \subset \cdots$ be an exhaustion of H by finite subsets, that is $\bigcup_{n\in\mathbb{N}} H_n = H$. For every $n \in \mathbb{N}$, there is a homomorphism $\varphi_n : H \to F_2$ which is injective on H_n. Every $g \in H$ is contained in every H_N for N large enough, implying that the sequence $(\varphi_n(g))_{n\in\mathbb{N}}$ is eventually never 1; in other words, the sequence $(\varphi_n)_{n\in\mathbb{N}}$ is stable and $\underleftarrow{\ker}((\varphi_n)_{n\in\mathbb{N}}) = \{1\}$. This brings us to the following definition:

Definition II.1.18. Let G be a group, and let H be a finitely generated group. An *(algebraic) limit group over* G is a quotient of the form $H/\underleftarrow{\ker}((\varphi_n)_{n\in\mathbb{N}})$ for some stable sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H,G)^{\mathbb{N}}$.

Remark II.1.19. The discussion preceding Definition II.1.18 perhaps suggested that an (algebraic) limit group H over G should be defined by admitting a stable sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H,G)^{\mathbb{N}}$ with a trivial stable kernel; in the case of free groups, and more generally, whenever G is equationally Noetherian, the two definitions coincide (see Subsection II.1.4). As a matter of fact, the two definitions are equivalent if and only if G is equationally Noetherian. However, the groups we deal with in Chapter III are *not* equationally Noetherian, and we therefore stick to Definition II.1.18.

However, Definitions II.1.14 and II.1.18 do coincide; more precisely:

Lemma II.1.20. *Let G be a group. A group H is a marked limit group over G if and only if H is an (algebraic) limit group over G.*

Proof. Suppose first that H is a marked limit group over G; in other words, there is $n \in \mathbb{N}$, a generating tuple S of H of size n and a sequence of marked groups $(G_i, S_i) \in \mathcal{G}_n$ such that $G_i \leq G$ for every $i \in \mathbb{N}$ and

$$(G_i, S_i) \xrightarrow[n\to\infty]{} (H, S).$$

We will show that H is an (algebraic) limit group over G by constructing a stable sequence of homomorphisms $(\varphi_i : F_n \to G)_{i\in\mathbb{N}} \in \mathrm{Hom}(F_n,G)^{\mathbb{N}}$ for which $H = F_n/\underleftarrow{\ker}((\varphi_i)_{i\in\mathbb{N}})$.

Define $\varphi_i : F_n \to G$ by mapping the standard generating tuple of F_n to the tuple $(S_i) \in {G_i}^n \subset G^n$. This sequence is clearly stable: given $w \in F_n$ with $|g| = N$, whenever $d_n((G_i,S_i),(H,S)) \leq e^{-N}$ we have that $\varphi_i(w) = 1$ if and only if $w(S)$ is the trivial element in H; in particular, the sequence $(\varphi_i(w))_{i\in\mathbb{N}}$ is either eventually trivial, or eventually never trivial (depending on whether $w(S)$ is trivial in H), and $(\varphi_i : F_n \to G)_{i\in\mathbb{N}}$ is stable. This also implies that the kernel of the epimorphism $\varphi : F_n \to H$ which maps the standard generating tuple of F_n to S is exactly $\underleftarrow{\ker}((\varphi_i)_{i\in\mathbb{N}})$; hence $H = F_n/\underleftarrow{\ker}((\varphi_i)_{i\in\mathbb{N}})$.

For the converse, suppose that there is a group H' generated by a tuple $(g_1, \dots, g_n)$ and a stable sequence of homomorphisms $(\varphi_i : H' \to G)_{i \in \mathbb{N}} \in \mathrm{Hom}(H', G)^{\mathbb{N}}$ such that $H = H' / \underleftarrow{\ker}((\varphi_i)_{i \in \mathbb{N}})$. We will show that the marking

$$\left(H = H' / \underleftarrow{\ker}((\varphi_i)_{i \in \mathbb{N}}), S = \left(g_1 \cdot \underleftarrow{\ker}((\varphi_i)_{i \in \mathbb{N}}), \dots, g_n \cdot \underleftarrow{\ker}((\varphi_i)_{i \in \mathbb{N}}) \right) \right)$$

of H is the limit of the sequence $(G_i = \varphi_i(H'), S_i = (\varphi_i(g_1), \dots, \varphi_i(g_n)))_{i \in \mathbb{N}}$ in $\mathcal{G}_n$.

It is enough to show that for every $N \in \mathbb{N}$, there is $j_N \in \mathbb{N}$ such that for every $w \in F_n$ of length at most N, $w(S) = 1$ in H if and only if $w(S_i) = 1$ in G_i for $i \geq j_N$ (as this implies that $d_n((H, S), (G_i, S_i)) \leq e^{-N}$). For every $w \in F_n$ there is $j_w \in \mathbb{N}$ such that $w(S_i) = 1$ in G_i if and only if $w(g_1, \dots, g_n) \in \underleftarrow{\ker}((\varphi_i)_{i \in \mathbb{N}})$, that is $w(S) = 1$ in H. Since $|B_N(F_n)| < \infty$, setting $j_N = \max_{w \in B_N(F_n)} j_w$ gives the desired result. □

Stable sequences of homomorphisms are by no means special: a standard diagonalization argument shows that every sequence of homomorphisms $(\varphi_n)_{n \in \mathbb{N}} \in \mathrm{Hom}(H, G)^{\mathbb{N}}$ has a stable subsequence (as long as H is countable, which is always the case since we assume that H is finitely generated). The process of extracting a stable subsequence via diagonalization involves choice, and different stable subsequences can admit different stable kernels. Using ultrafilters (see Subsection I.4) we avoid this ambiguity, while still working with general (and not necessarily stable) sequences of homomorphisms. Adopting the language of ultrafilters also saves us from often passing to subsequences, and allows us to phrase statements with relative ease. We fix a non-principal ultrafilter ω and define:

Definition II.1.21. Let G be a group, let H be a finitely generated group and let $(\varphi_n)_{n \in \mathbb{N}} \in \mathrm{Hom}(H, G)^{\mathbb{N}}$ be a sequence of homomorphisms. The *stable kernel of* $(\varphi_n)_{n \in \mathbb{N}}$ *(with respect to* ω*)*, $\underleftarrow{\ker}_\omega((\varphi_n)_{n \in \mathbb{N}})$, is the collection of all $g \in H$ for which $\varphi_n(g) = 1$ ω-almost-surely, that is

$$\underleftarrow{\ker}_\omega((\varphi_n)_{n \in \mathbb{N}}) = \{ g \in H \mid g \in \ker(\varphi_n) \ \omega\text{-almost-surely} \}.$$

Subsequently, we give a definition for limit groups over a group G as follows; this definition is simply a rephrasing of Definition II.1.18 in the language of ultrafilters:

Definition II.1.22. Keeping the notation from Definition II.1.21, an *(algebraic) limit group over* G is a group of the form $L = H / \underleftarrow{\ker}_\omega((\varphi_n)_{n \in \mathbb{N}})$.

We use the following terminology and notation when discussing (algebraic) limit groups over a group G:

Notation II.1.23. The sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H,G)^{\mathbb{N}}$ is called the *defining sequence* of homomorphisms for L. The corresponding quotient map $H \twoheadrightarrow L$ obtained by quotienting out the stable kernel $\underleftarrow{\ker}_\omega((\varphi_n)_{n\in\mathbb{N}})$ is denoted by φ_∞ and we refer to it as the *limit map* associated with the sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H,G)^{\mathbb{N}}$.

In the following subsection we explore how such a sequence of homomorphisms can be used to construct an action of a limit group on an $\mathbb{R}$-tree, given that G admits a suitable action on a hyperbolic space. In subsection II.1.4 we will discuss how the algebraic definition for a limit group over G given in this subsection relates to equational Noetherianity and being fully residually-G.

II.1.3 Limiting $\mathbb{R}$-trees

The main advantage of embracing the approach in Subsection II.1.2 is that it enables one, under some conditions, to extract an action of a limit group on a real tree. In what follows, we assume that all metric spaces are geodesic (recall that a metric space (X,d) is geodesic if any two points $x,y\in X$ are connected by a geodesic, i.e. a path whose length is $d(x,y)$). A *pointed metric space* is a triplet (X,d,o) where (X,d) is a metric space and $o \in X$.

Definition II.1.24. Let $(X_n,d_n,o_n)_{n\in\mathbb{N}}$ be a sequence of pointed metric spaces. The *limit of* $(X_n,d_n,o_n)_{n\in\mathbb{N}}$ *with respect to an ultrafilter* ω is a triplet $(X_\omega,d_\omega,o_\omega)$ where

$$X_\omega = \Big(\prod_{n\in\mathbb{N}} X_n\Big)/\omega = \frac{\{(x_n)_{n\in\mathbb{N}} \mid \lim_\omega d_n(x_n,o_n) < \infty\}}{(x_n)_{n\in\mathbb{N}}\ (y_n)_{n\in\mathbb{N}} \iff \lim_\omega d_n(x_n,y_n) = 0},$$

$d_\omega : X_\omega \times X_\omega \to \mathbb{R}_{\geq} 0$ is given by

$$d_\omega([(x_n)_{n\in\mathbb{N}}],[(y_n)_{n\in\mathbb{N}}]) = \lim_\omega d_n(x_n,y_n)$$

and o_ω is the equivalence class of the sequence $(o_n)_{n\in\mathbb{N}}$.

Remark II.1.25. We will often abuse notation and refer to the equivalence class of a sequence $(x_n)_{n\in\mathbb{N}}$ in X_ω as $(x_n)_{n\in\mathbb{N}}$ instead of $[(x_n)_{n\in\mathbb{N}}]$. If a sequence $(x_n)_{n\in\mathbb{N}}$ lies in X_ω we call it a *visible sequence.*

It is straightforward to verify that $(X_\omega,d_\omega,o_\omega)$ is a pointed metric space (that is, d_ω is well-defined and satisfies the required conditions). We continue by showcasing a simple instance in which $(X_\omega,d_\omega,o_\omega)$ inherits some properties from the spaces in the sequence $(X_n,d_n,o_n)_{n\in\mathbb{N}}$; this will prove to be of great importance in Chapter III.

Lemma II.1.26. *Let* $(X_n, d_n, o_n)_{n\in\mathbb{N}}$ *be a sequence of pointed metric spaces such that each* (X_n, d_n) *is* δ_n*-hyperbolic. If* $\delta_\omega = \lim_\omega \delta_i < \infty$ *then* (X_ω, d_ω) *is a* δ_ω*-hyperbolic space. In particular, if* $\delta_\omega = 0$ *then* (X_ω, d_ω) *is a real tree.*

Proof. It suffices to verify that X_ω is a geodesic space that satisfies the Gromov four-point condition with respect to δ_∞, that is

$$((x_n)_{n\in\mathbb{N}}, (z_n)_{n\in\mathbb{N}})_{(p_n)_{n\in\mathbb{N}}} \geq \{((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})_{(p_n)_{n\in\mathbb{N}}}, ((y_n)_{n\in\mathbb{N}}, (z_n)_{n\in\mathbb{N}})_{(p_n)_{n\in\mathbb{N}}}\} - \delta_\omega.$$

It is easy to see that (X_ω, d_ω) is a geodesic space: given $(x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}} \in X_\omega$, let $t_n : [0, d_n(x_n, y_n)] \to X_n$ be a geodesic from x_n to y_n in (X_n, d_n). Define $t : [0, d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})] \to X_\omega$ by

$$s \mapsto t_n\left(\frac{d_n(x_n, y_n)}{d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})} \cdot s\right).$$

It is straightforward to verify that t is the suitable geodesic in (X_ω, d_ω). First, note that for every $s \in [0, d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})]$ with $t(s) = (s_n)_{n\in\mathbb{N}}$ we have that s_n is on the geodesic between x_n and y_n so

$$\lim_\omega d_n(o_n, s_n) \leq \lim_\omega (d_n(o_n, x_n) + d_n(x_n, s_n)) \leq \lim_\omega (d_n(o_n, x_n) + d_n(x_n, y_n)) < \infty$$

and t is well-defined. Furthermore, given $s_1 < s_2 \in [0, d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})]$ we have that

$$\begin{aligned}
d_\omega(t(s_1), t(s_2)) &= \lim_\omega d_n\left(\left(\frac{d_n(x_n, y_n)}{d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})} \cdot s_1\right), \left(\frac{d_n(x_n, y_n)}{d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})} \cdot s_2\right)\right) \\
&= \lim_\omega \left(\frac{d_n(x_n, y_n)}{d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})} \cdot (s_2 - s_1)\right) \\
&= \frac{\lim_\omega (d_n(x_n, y_n))}{d_\omega((x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}})} \cdot (s_2 - s_1) = s_2 - s_1
\end{aligned}$$

and t minimizes distances.

Finally, given $x_n, y_n, z_n, p_n \in X_n$ for every $n \in \mathbb{N}$ we have that

$$(x_n, z_n)_{p_n} \geq \{(x_n, y_n)_{p_n}, (y_n, z_n)_{p_n}\} - \delta_n$$

for every $n \in \mathbb{N}$ since (X_n, d_n) is δ_n-hyperbolic. Letting $n \to \infty$ we get that (X_ω, d_ω) satisfies Gromov's four-point condition for $\delta_\omega = \lim_\omega \delta_n$. □

Remark II.1.27. The proof above also shows that if geodesics in (X_ω, d_ω) are unique (e.g. when (X_ω, d_ω) is an $\mathbb{R}$-tree), then every geodesic in (X_ω, d_ω) can be approximated by geodesics in the spaces (X_n, d_n).

Suppose now that H is a finitely generated group, and let $S = (s_1, \ldots, s_n)$ be a finite generating tuple of H. Let (X_n, d_n, o_n) be a sequence of pointed metric spaces, and suppose in addition that H acts on each (X_n, d_n, o_n). Under a mild assumption, we get that H acts on X_ω:

Lemma II.1.28. *If $\lim_\omega d_n(o_n, s.o_n) < \infty$ for every $s \in S$, then H acts on X_ω by $h.(x_n)_{n\in\mathbb{N}} = (h.x_n)_{n\in\mathbb{N}}$.*

Proof. We just need to verify that for every $(x_n)_{n\in\mathbb{N}}$ and $h \in H$ we have that $(h.x_n)_{n\in\mathbb{N}} \in X_\omega$, that is $\lim_\omega d_n(o_n, h.x_n) < \infty$. The other properties that define a group action easily follow.

Write $h = s_1 \cdots s_k$ where $s_i \in S$ for $1 \le i \le k$ and note that

$$\begin{aligned} d_n(o_n, h.o_n) &= d_n(o_n, (s_1\cdots s_k).o_n) \\ &\le d_n(o_n, s_1.o_n) + d_n(s_1.o_n, (s_1 \cdot s_2).o_n) + \cdots + d_n((s_1\cdots s_{k-1}).o_n, (s_1\cdots s_k).o_n) \\ &= \sum_{i=1}^{k} d_n(o_n, s_i.o_n). \end{aligned}$$

and therefore

$$\begin{aligned} \lim_\omega d_n(o_n, h.x_n) &\le \lim_\omega d_n(o_n, h.o_n) + \lim_\omega d_n(h.o_n, h.x_n) \\ &= d_n(o_n, h.o_n) + \lim_\omega d_n(o_n, x_n) \\ &\le \lim_\omega d_n(o_n, x_n) + \sum_{i=1}^{k} \lim_\omega d_n(o_n, s_i.o_n) < \infty \end{aligned}$$

as desired. □

Corollary II.1.29. *Let G be a group acting on a sequence of pointed metric spaces $(X_n, d_n, o_n)_{n\in\mathbb{N}}$, let H be a group generated by a finite tuple S and let $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H,G)^{\mathbb{N}}$ be a sequence of homomorphisms. For each n, we get an action $H \curvearrowright X_n$ by setting $h.x = \varphi_n(h).x$ for $h \in H$ and $x \in X_n$. Therefore, by Lemma II.1.28, if $\lim_\omega d_n(o_n, s.o_n) < \infty$ for every $s \in S$ then H acts on X_ω.*

The remainder of this subsection will focus on how to extract a (non-trivial) action of a limit group on a limiting $\mathbb{R}$-tree (that is, an action without a global fixed point). We therefore fix a group G admitting a non-trivial action on a δ-hyperbolic space (X, d) and a limit group $L = H/\underleftarrow{\ker}_\omega((\varphi_n)_{n\in\mathbb{N}})$. We also assume that H is generated by a finite tuple S. The limiting real tree that we obtain will be a limit of a sequence of pointed metric spaces, where each element in the sequence is X equipped with a rescaling of the metric d. We therefore define:

Definition II.1.30. Let o be a choice of a basepoint of (X, d), and let $n \in \mathbb{N}$. The *scaling factor of* φ_n *at* o is given by

$$\|\varphi_n\|_o = \max_{s\in S} d(o, \varphi_n(s).o).$$

Combining Lemmas II.1.26 and II.1.28 we obtain:

Corollary II.1.31. *If there is a sequence* $(o_n)_{n\in\mathbb{N}}$ *of points in* X *such that* $\lim_\omega \|\varphi_n\|_{x_n} = \infty$ *then the sequence* $(X, d/\|\varphi_n\|_{o_n}, o_n)$ *converges to a real tree* $(X_\omega, d_\omega, o_\omega)$ *on which* H *acts by* $h.(x_n)_{n\in\mathbb{N}} = (\varphi_n(h).o_n)_{n\in\mathbb{N}}$. *Furthermore,* $H \curvearrowright X_\omega$ *gives rise to an action* $L \curvearrowright X_\omega$.

Proof. Since (X, d) is δ-hyperbolic, $(X, d/\|\varphi_n\|_{o_n})$ is $\delta/\|\varphi_n\|_{o_n}$-hyperbolic; by Lemma II.1.26, since $\lim_\omega \|\varphi_n\|_{x_n} = \infty$, (X_ω, d_ω) is a real tree.

It is enough to show that the boundedness condition of Lemma II.1.28 holds. Indeed, given $s \in S$ we have that

$$\|\varphi_n\|_{o_n}(o_n, \varphi_n(s).o_n) \le \frac{\max_{s\in S} d(o, \varphi_n(s).o)}{\|\varphi_n\|_o} = 1$$

as desired. Lastly, since $\underleftarrow{\ker}_\omega((\varphi_n)_{n\in\mathbb{N}}$ acts trivially on X_ω, the action $H \curvearrowright X_\omega$ induces an action $L \curvearrowright X_\omega$. □

Corollary II.1.31 is still not enough: there is nothing preventing the action $H \curvearrowright X_\omega$ from having a global fixed point. Bestvina introduced a method for overcoming this problem in [17] (where the spaces considered are hyperbolic n-spaces), which was later generalised by Paulin in [94] (to accommodate any hyperbolic space). This method is often referred to as the "Bestvina-Paulin trick", and it revolves around carefully choosing the sequence basepoints $(o_n)_{n\in\mathbb{N}}$. We amass the relevant definitions for explaining this method, and begin with an absolute version of the scaling factor defined above:

Definition II.1.32. Keeping the notation above, the *scaling factor* of a homomorphism $\varphi_n : H \to G$ is

$$\|\varphi_n\| = \inf_{x\in X} \max_{s\in S} d(x, \varphi_n(s).x).$$

Remark II.1.33. Note that $\|\varphi_n\|$ does not depend on a choice of a basepoint o of X.

The scaling factor of a homomorphism will play a crucial role in Chapter III, where we prove a version of the *shortening argument* which essentially says that if the homomorphisms in the sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H, G)^{\mathbb{N}}$ satisfy a certain minimality condition (that's related to their scaling factor), the stable kernel of the sequence $\underleftarrow{\ker}_\omega((\varphi_n)_{n\in\mathbb{N}})$ cannot be trivial.

Definition II.1.34. The limit group L is called *divergent* if $\underleftarrow{\ker}_\omega(\|\varphi_n\|) = \infty$.

We conclude with the following Theorem:

Theorem II.1.35 (Bestvina-Paulin trick)**.** *If L is a divergent limit group then there exists a choice of basepoints $(o_n)_{n\in\mathbb{N}}$ in X such that the sequence $(X, d/\|\varphi_n\|, o_n)$ converges to a real tree $(X_\omega, d_\omega, o_\omega)$ on which L acts non-trivially.*

Proof. For every $n \in \mathbb{N}$, choose $o_n \in X$ that satisfies $d(o_n, \varphi_n(s).o_n) \leq \|\varphi_n\| + {}^1\!/\!_n$ for every $s \in S$ (this is possible since $\|\varphi_n\| = \inf_{x\in X} \max_{s\in S} d(x, \varphi_n(s).x)$). As in Corollary II.1.31, we obtain an action $H \curvearrowright X_\omega$ which gives rise to an action $L \curvearrowright X_\omega$; it is enough to show that the action of H on X_ω is non-trivial, that is, no $(x_n)_{n\in\mathbb{N}} \in X_\omega$ is fixed by all of H.

Since S is finite, there exists $s \in S$ such that $d(o_n, \varphi_n(s).o_n) \leq \|\varphi_n\| + {}^1\!/\!_n$ holds ω-almost surely. Therefore, for every $(x_n)_{n\in\mathbb{N}} \in X_\omega$ we have that

$$d/(\|\varphi_n\|)(x, \varphi_n(s).x) \geq \frac{d(o_n, \varphi_n(s).o_n) - 1/n}{\|\varphi_n\|} \geq 1 - \frac{1}{n}$$

ω-almost surely, so $d_\omega((x_n)_{n\in\ltimes}, s.(x_n)_{n\in\ltimes}) \geq 1$ and H does not fix $(x_n)_{n\in\mathbb{N}}$. ☐

Remark II.1.36. Note that H and L are both finitely generated, and therefore countable. If X_ω is not a line, the valence of every vertex of X_ω is uncountable and therefore the action $H \curvearrowright X_\omega$ is not minimal. If needed, one can always reduce to a subtree on which the action is minimal.

II.1.4 Equational Noetherianity and residual properties

We begin by defining *equational Noetherianity*, a notion that was introduced by Baumslag, Myasnikov and Remeslennikov in [12]:

Definition II.1.37. A group G is called *equationally Noetherian* if the following holds: every system of equations $\Sigma \subset G * F(\boldsymbol{x})$ with finitely many variables $\boldsymbol{x} = (x_1, \ldots, x_n)$ (and parameters from G) admits a finite subsystem $\Sigma_0 \subset \Sigma$, such that the sets of solutions to Σ and Σ_0,

$$V_G(\Sigma) = \{\boldsymbol{g} \in G^n | \forall \sigma \in \Sigma,\ \sigma(\boldsymbol{g}) = 1\} \quad \textit{and} \quad V_G(\Sigma_0) = \{\boldsymbol{g} \in G^n | \forall \sigma_0 \in \Sigma,\ \sigma_0(\boldsymbol{g}) = 1\}$$

respectively, coincide. In this case, we say that the two systems Σ and Σ_0 are *equivalent*.

The proof of the following lemma highlights why "equational Noetherianity" is a suitable name for the property described in Definition II.1.37:

Lemma II.1.38. *Finitely generated free groups are equationally Noetherian.*

Proof. Equational Noetherianity is inherited by subgroups; we will therefore prove the lemma for a free group of rank 2, $F_2 = \langle x_1, x_2 \rangle$.

F_2 embeds in $\mathrm{SL}_2(\mathbb{Z})$ by mapping

$$x_1 \longmapsto \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad x_2 \longmapsto \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}.$$

We denote this embedding by $\varphi : F_2 \to \mathrm{SL}_2(\mathbb{Z})$. The fact that the subgroup of $\mathrm{SL}_2(\mathbb{Z})$ generated by the two matrices above is a free group of rank 2 can be easily verified by considering the action of $\mathrm{SL}_2(\mathbb{Z})$ on the real plane $\mathbb{R}^2$ and using the Ping Pong lemma with the sets $\{(x,y) \in \mathbb{R}^2 | |x| > |y|\}$ and $\{(x,y) \in \mathbb{R}^2 | |x| < |y|\}$.

Let $\Sigma \subset F_2 * F_n$ be a system of equations in F_2 with variables $\boldsymbol{y} = (y_1, \dots, y_n)$. Let $\Sigma' \subset \mathrm{SL}_2(\mathbb{Z}) * F_n$ be the corresponding system of equations in $\mathrm{SL}_2(\mathbb{Z})$, that is for every $\sigma \in \Sigma$ define $\sigma' \in \Sigma'$ to be the same equation where the parameters from F_2 are replaced with their image under φ in $\mathrm{SL}_2(\mathbb{Z})$. Note that $\sigma(\boldsymbol{g}) = 1$ in F_2 (for $\boldsymbol{g} \in F_2^n$) if and only if

$$\sigma'(\varphi(\boldsymbol{g})) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Regarding the n variables in each σ' as matrices with four entires, we get that σ' gives rise to four polynomial equations (with coefficients in $\mathbb{Z}$) in $4n$ variables $\boldsymbol{z} = (z_1^1, z_2^1, z_3^1, z_4^1, \dots, z_1^n, z_2^n, z_3^n, z_4^n)$, obtained by comparing the different entries with those of the identity matrix. Therefore, the system of equations Σ gives rise to a system of polynomial equations Ψ over $\mathbb{Z}$.

Consider the ideal (Ψ) in $\mathbb{Z}[\boldsymbol{z}]$; by Hilbert's basis theorem, $\mathbb{Z}[\boldsymbol{z}]$ is Noetherian, and therefore the ideal (Ψ) is finitely generated. Write $(\Psi) = (\psi_1, \dots, \psi_k)$. Each ψ_i was obtained from some $\sigma_i' \in \Sigma'$. It follows that $\{\sigma_1', \dots, \sigma_k'\} \subset \Sigma'$ is equivalent to Σ', and therefore $\{\sigma_1, \dots, \sigma_k\} \subset \Sigma$ is equivalent to Σ. □

Remark II.1.39. The same argument shows that if a group G is linear over a field (or more generally, a commutative Noetherian unity ring) then G is equationally Noetherian (cf. [12, Theorem B1]). In particular, right-angled Artin groups are equationally Noetherian, a fact that will come into play in Chapter IV.

We point out that other classes of groups that are closely related to free groups have also been shown to be equationally Noetherian; these include hyperbolic groups (proven by Sela [109] in the torsion-free case, and generalized by Reinfeldt and Weidmann [97] to all hyperbolic groups), torsion-free toral relatively hyperbolic groups (proven by Groves [55]) and more generally any group that is hyperbolic relative to equationally Noetherian groups (proven by Groves and Hull [57, Theorem D]).

We explained in Subsection II.1.1 why a finitely generated group H, which is fully residually-G, is in fact a limit group over G (we did this in the case where G is a free group, but the same argument works for any group G); we now turn to prove the converse for an equationally Noetherian group G.

Lemma II.1.40. *Let G be an equationally Noetherian group and let H be a limit group over G. Then H is fully residually-G.*

Proof. Let S be a finite generating tuple of H such that (H,S) is the limit of a sequence $(G_i,S_i)_{i\in\mathbb{N}}$ in $\mathcal{G}_n$ and $G_i\leq G$ for every $i\in\mathbb{N}$. Let $\Sigma\subset F_n$ be the kernel of the homomorphism $F_n\to H$ given by mapping the standard generating tuple of F_n to S. Since G is equationally Noetherian, there is a finite subset $\Sigma_0\subset\Sigma$ such that $V_G(\Sigma)=V_G(\Sigma_0)$. In particular, for every $\boldsymbol{g}\in G^n$, if $\sigma_0(\boldsymbol{g})=1$ for every $\sigma_0\in\Sigma_0$, then $\sigma(\boldsymbol{g})=1$ for every $\sigma\in\Sigma$.

We next show that for sufficiently large i, the map which sends S to S_i extends to a homomorphism $f_i:H\to G_i$. It is enough to show that for every $\sigma\in\Sigma$, $\sigma(S_i)$ is trivial in G_i. Given $\sigma\in\Sigma$, since $\sigma(S)=1$ in H we have that $\sigma(S_i)=1$ in G_i for sufficiently large i; since Σ_0 is finite, we have that, for sufficiently large i, $\sigma_0(S_i)=1$ in G_i for every $\sigma_0\in\Sigma_0$. As discussed in the previous paragraph, it follows that, for sufficiently large i, $\sigma(S_i)=1$ in G_i for every $\sigma\in\Sigma$ and the map which sends S to S_i extends to a homomorphism $f_i:H\to G_i$.

Finally, given a finite subset $F\subset H$, set $N=\max_{g\in F}|g|$. For i large enough we get that $d_n((H,S),(G_i,S_i))<e^{-N}$, which implies that the homomorphism $f_i:H\to G_i$ is injective on F. ☐

The proof of Lemma II.1.40 also implies the following: if $L=H/\underleftarrow{\ker}_\omega((\varphi_n)_{n\in\mathbb{N}})$ is a limit group over an equationally Noetherian group G with a defining sequence of homomorphisms $(\varphi_n)_{n\in\mathbb{N}}\in\mathrm{Hom}(H,G)^{\mathbb{N}}$, then φ_n factors via the limit map $\varphi_\infty: H\twoheadrightarrow L$ ω-almost surely. It follows that in this case, there is a stable sequence of homomorphisms $(\phi_n)_{n\in\mathbb{N}}\in\mathrm{Hom}(L,G)^{\mathbb{N}}$ whose stable kernel is trivial. As a matter of fact, this condition is equivalent to equational Noetherianity:

Lemma II.1.41. *Let G be a countable group. The following are equivalent:*

1. *G is equationally Noetherian,*
2. *For every finitely generated group H and for every sequence of homomorphisms $(\varphi_n)_{n\in\mathbb{N}} \in \operatorname{Hom}(H,G)^{\mathbb{N}}$, φ_n factors via the limit map φ_∞ ω-almost surely,*
3. *For every finitely generated group H and for every sequence of homomorphisms $(\varphi_n)_{n\in\mathbb{N}} \in \operatorname{Hom}(H,G)^{\mathbb{N}}$, there exists $n \in \mathbb{N}$ such that φ_n factors via the limit map φ_∞.*

Remark II.1.42. The assumption that G is countable is only required for the implication *3.* $\Longrightarrow$ *1.*

Proof. The fact that *1.* implies *2.* follows from Lemma II.1.40; it is obvious that *2.* implies *3.* We will show that *3.* implies *1* by proving the contrapositive statement.

Suppose now that *3.* holds, and let $\Sigma \subset G * F_n$ be a system of equations with parameters from G in variables $\boldsymbol{x} = (x_1, \dots, x_n)$. Since G is countable, $G * F_n$ is countable and therefore Σ admits an exhaustion by finite subsets $\Sigma_1 \subset \Sigma_2 \subset \cdots$. We will show that Σ is equivalent to Σ_i for some i. Suppose, to obtain a contradiction, that Σ is not equivalent to any of the Σ_i's. In particular, for each i there exists $\boldsymbol{g}_i = (g_1^i, \dots, g_n^i) \in V_G(\Sigma_i) \setminus V_G(\Sigma)$. Therefore, the map which sends $\boldsymbol{x}$ to $\boldsymbol{g}_i$ extends to a homomorphism $\varphi_i : G * F_n \to G$ that satisfies:

1. φ_i is the identity on G,
2. $\varphi_i(\sigma) = 1$ for every $\sigma \in \Sigma_i$, that is $\Sigma_i \subset \ker(\varphi_i)$ (since $\varphi_i(\boldsymbol{x}) = \boldsymbol{g}_i \in V_G(\Sigma_i)$),
3. $\Sigma \not\subset \ker(\varphi_i)$ (since $\varphi_i(\boldsymbol{x}) = \boldsymbol{g}_i \notin V_G(\Sigma)$, there exists $\sigma \in \Sigma$ such that $\varphi_i(\sigma) = \sigma(\boldsymbol{g}_i)$ is non-trivial).

Since $\Sigma_i \subset \ker(\varphi_j)$ for every $j \geq i$, we also have that $\Sigma = \bigcup_{i\in\mathbb{N}} \Sigma_i \subset \ker(\varphi_\infty)$.

Consider now the sequence $(\varphi_i|_{F_n}) \in \operatorname{Hom}(F_n, G)^{\mathbb{N}}$, and by our assumption $\varphi_i|_{F_n}$ factors through $\varphi_\infty|_{F_n}$ for some i. It follows that φ_i factors via φ_∞ since $\varphi_i|_G = \varphi_\infty|_G = \mathrm{Id}_G$. Therefore $\Sigma \subset \ker(\varphi_\infty) \subset \ker(\varphi_i)$, which contradicts 3. above. ☐

Remark II.1.43. In [57, Theorem B], Groves and Hull prove that if G is acylindrically hyperbolic, then the failure of G to be equationally Noetherian can be detected by a *non-divergent* sequence of homomorphisms from a finitely generated group to G that fails to factor through the corresponding limit map.

In light of Lemma II.1.41, we see how the lack of equational Noetherianity can pose a problem when working with limiting real trees; as mentioned earlier, parts of Chapter III are devoted to overcoming these obstacles.

We have established that for a countable group G, being equationally Noetherian is equivalent to every limit group over G being fully residually-G. We conclude the discussion with a few historical anecdotes, yielding some examples of limit groups (other than free and free abelian groups that were discussed in Example II.1.5) and preparing the grounds for Subsection II.1.5 where we discuss the structure of limit groups. Gilbert Baumslag proved the following in 1962 (the following is a rephrasing of Baumslag's original statement):

Proposition II.1.44 (cf. [10, Proposition 1])**.** *Let F be a free group, and let*

$$g_1, \dots, g_n, u \in F.$$

If $[u, g_i] \neq 1$ for every $1 \leq i \leq n$, then there exists $N \in \mathbb{N}$ such that for every $k_1, \dots, k_n$ with $|k_i| \geq N$,

$$g_1 \cdot u^{k_1} \cdot g_2 \cdot u^{k_2} \cdots g_n \cdot u^{k_n} \neq 1.$$

Remark II.1.45. Proposition II.1.44 above can be proved using a standard Ping Pong argument.

We continue with a short list of examples, based on Benjamin and Gilbert Baumslag's works [9, 10].

Example II.1.46. The fundamental group of an orientable surface Σ of genus 2, $\pi_1(\Sigma) = \langle a, b, c, d | \ [a, b] = [c, d] \rangle$ is fully residually free, and hence a limit group. To prove this, define $f : \pi_1(\Sigma) \to F_2 = \langle x_1, x_2 \rangle$ by mapping a and c to x_1 and b and d to x_2. Let $\varphi : \pi_1(\Sigma) \to \pi_1(\Sigma)$ be the automorphism of $\pi_1(\Sigma)$ given by a Dehn twist by $[a, b]$, that is the unique automorphism of $\pi_1(\Sigma)$ extending the map

$$\varphi(g) = \begin{cases} g & g \in \langle a, b \rangle \\ [a, b] g [a, b]^{-1} & g \in \langle c, d \rangle \end{cases}$$

Now, given $g \in \pi_1(\Sigma)$, write $g = g_1 \cdot h_1 \cdots g_k \cdot h_k$ where for every $1 \leq i \leq k$ $g_i \in \langle a, b \rangle$ and $h_i \in \langle c, d \rangle$; we may assume further that no g_i and h_i are trivial or lie in $\langle [a, b] \rangle$ (except for, perhaps, g_1 and h_k). By Proposition II.1.44, for $f(g_1), f(h_1), \dots, f(g_k), f(h_k), f([a, b])$ there exists $N = N(g) \in \mathbb{N}$ such that (note that the only elements commuting with $[a, b]$ in $\pi_1(\Sigma)$ are powers of $[a, b]$)

$$\begin{aligned} f \circ \varphi^n(g) = & f(g_1) \big(\cdot f([a, b])^n \cdot f(h_1) \cdot f([a, b])^{-n} \big) \cdots \\ & \cdots f(g_k) \cdot \big(f([a, b])^n \cdot f(h_k) \cdot f([a, b])^{-n} \big) \neq 1 \end{aligned}$$

as long as $n \geq N(g)$. Finally, given any finite subset $E \subset \pi_1(\Sigma)$, we have that $f \circ \varphi^n : \pi_1(\Sigma) \to F_2$ must be injective on E for

$$n \geq \max_{g \in (E \cup e^{-1}) \cdot (E \cup e^{-1})} N(g).$$

Remark II.1.47*.* The *double* of a free group $F = \langle x_1, \dots, x_n \rangle$ along an element $w \in F$ is given by

$$F *_w F = \langle x_1^1, \dots, x_n^1, x_1^2, \dots, x_n^2 \mid w(x_1^1, \dots, x_n^1) = w(x_1^2, \dots, x_n^2) \rangle.$$

The exact same proof as in the example above shows that whenever w is not a proper power in F, $F *_w F$ is a limit group.

In fact, the exact same proof also shows that *generalised doubles over limit groups*, which were introduced by Champetier and Guirardel (see [33, Definition 4.4]) are limit groups themselves. We record here the definition of a generalised double: a generalised double over a group G is a group H that admits a splitting $A *_C B$ (or $A * C$) satisfying:

1. A and B are finitely generated,
2. C is a non-trivial abelian group that is maximal abelian in both A and B (or just in A if $H = A*_C$),
3. there is a surjective homomorphism $f : H \to G$ such that $f|_A$ and $f|_B$ (or just $f|_A$ if $H = A*_C$) is injective.

We therefore extablished that iterated generalised doubles over a free group are limit groups, yielding a large family of examples of limit groups; as a matter of fact, following Sela's work on the structure of limit groups, Champetier and Guirardel proved the (much stronger) converse of the aforementioned result:

Theorem II.1.48 ([33, Theorem 4.6])**.** *Every limit group is an iterated generalised double over a free group.*

We continue by mentioning another family of examples of limit groups discussed in [10] and [9].

Definition II.1.49. Let G be a group and let $u \in G$. Let C be another copy of the centraliser $C_G(u)$ of u, and let $H = C \times \mathbb{Z}$. The *free extension of the centraliser* $C_G(u)$ is the group

$$G(u) = G *_{C_G(u) = C} H.$$

Example II.1.50. Every free extension of a centraliser in a non-abelian free group is a limit group. Indeed, let F be a free group and let $C_F(u)$ be the centraliser of an element $u \in F$; as in Definition II.1.49 write $H = C \times \mathbb{Z} = C \times \langle t \rangle$. Define $\varphi_n : F(u) \to F$ to be the identity on F and map t to u^n. Using Proposition II.1.44 as in Example II.1.46 we obtain that for every finite $E \subset F(u)$, φ_n is injective on E for n large enough.

As before, we get that a free centraliser extension of a limit group is a limit group; therefore, every iterated free extension of centralisers over a free group F is a limit group. As in the case of generalised doubles, the following stronger statement holds:

Theorem II.1.51 ([70]). *Every limit group is a subgroup of an iterated free extension of centralisers over a free group.*

We will further discuss extensions of centralisers in Chapter IV, where we exploit the fact that they can be seen as *free products with commuting subgroups*.

We have just discussed two *hierarchical structures* that can be associated to a limit group. We now turn to present additional hierarchical structures, and recollect some finer properties of limit groups.

II.1.5 Structure and properties

In what follows we assume that limit groups are always defined over finitely generated free groups. We begin with a list of classical and elementary properties of limit groups. Most of these appear in some of the texts cited in the previous Subsections, and in particular in [33]; we remark that the proofs are elementary, and are included here for completeness.

Lemma II.1.52. *Let L be a limit group. The following holds:*

1. *L is torsion-free,*

2. *Every 2-generated subgroup of L is either free abelian or free; in particular, if $g, h \in L$ then $\langle g, h \rangle \in \{\{1\}, \mathbb{Z}, \mathbb{Z}^2, F_2\}$,*

3. *L is commutative transitive and* conjugacy separated abelian (CSA)*, that is every maximal abelian subgroup of L is malnormal,*

4. *if L is non-abelian then it has the same $\forall$-theory as a non-abelian free group.*

5. *if a finitely generated group H has the same $\forall$-theory as a non-abelian free group then it is a non-abelian limit group.*

Proof. We prove properties *1.-5.*:

1. Every $1 \neq g \in L$ survives in a free quotient, and is therefore not a torsion element.

2. Let $g, h \in L$; if $\langle g, h\rangle$ is abelian then by *1.* it must be either trivial, $\mathbb{Z}$ or $\mathbb{Z}^2$. Otherwise, there is a map f from L to a free group F which is injective on $\{g, h, [g, h]\}$; the subgroup of F generated by $\{f(g), f(h)\}$ must therefore be a free group of rank 2. This implies that $\langle g, h\rangle \cong F_2$.

3. We already saw that L is commutative transitive in Remark II.1.7. We will prove the stronger assertion that L is CSA. Let A be a maximal abelian subgroup of L and let $g \notin A$. We need to show that if there exists $1 \neq h \in gAg^{-1} \cap A$ then it must be that $g \in A$. Since A is maximal abelian, it is enough to show that g commutes with every $a \in A$, and since L is commutative transitive, it is enough to show that g commutes with a single element of A. We will show that g commutes with h.

   As before, there is a map $f : L \to F$ which is injective on $\{1, g, h, [g, h]\}$. Since $g^{-1}hg \in A$, $g^{-1}hg$ and h commute and therefore $[f(g^{-1}hg), f(h)] = 1$. Hence $f(g^{-1}hg)$ and $f(h)$ lie in a cyclic subgroup C of F, and since the centraliser of C is itself cyclic, up to replacing C with its centraliser we can assume that this cyclic subgroup contains $f(g)$. Hence $f([g, h]) = 1$ and we must have that $[g, h] = 1$.

4. A priori it is not clear that all free groups have the same $\forall$-theory; this easily follows from the following simple observation: if $G \leq H$, then the $\forall$-theory of H is contained in that of G. Now, for every free group F_n we have that $F_2 \leq F_n \leq F_2$ so F_2 and F_n have the same $\forall$-theory. By *2.*, we have that $F_2 \leq L$. Therefore, it is enough to show that the $\forall$-theory of F_2 is contained in the $\forall$-theory of L.

   Let Φ be a $\forall$-sentence and assume that $F_2 \vDash \Phi$; we can further assume that Φ has the following form

   $$\Phi = \forall \boldsymbol{x} \bigvee_{i=1}^{k} \Sigma_i(\boldsymbol{x}) = 1 \wedge \Psi_i(\boldsymbol{x}) \neq 1,$$

   where $\boldsymbol{x} = (x_1, \dots, x_n)$ is a tuple of elements, and Σ_i and Ψ_i for $1 \leq i \leq k$ are, respectively, systems of equations and inequations. We remind the reader that by writing $\Sigma_i(\boldsymbol{x}) = 1$ we mean $\bigwedge_{\sigma\in\Sigma_i}\sigma(\boldsymbol{x}) = 1$ (and the same when we write $\Psi_i(\boldsymbol{x}) \neq 1$).

Let $\boldsymbol{g} \in L^n$ and let $E = \{1\} \cup \{\sigma(\boldsymbol{g}) | \exists 1 \leq i \leq k,\ \sigma \in \Sigma_i\} \cup \{\psi(\boldsymbol{g}) | \exists 1 \leq i \leq k,\ \psi \in \Psi_i\}$, that is, E is the set of all words appearing in Φ evaluated at $\boldsymbol{g}$ (and the identity). There is a homomorphism $f_E : L \to F_2$ that is injective on E. Therefore

$$\chi(\boldsymbol{g}) = 1 \text{ in } L \iff \chi(f_E(\boldsymbol{g})) = 1 \text{ in } F_2$$

for every $\chi \in \bigcup_{i=1}^k \Sigma_i \cup \Psi_i$. Taking boolean combinations we have that

$$L \vDash \bigvee_{i=1}^k \Sigma_i(\boldsymbol{g}) = 1 \wedge \Psi_i(\boldsymbol{g}) \neq 1 \iff F_2 \vDash \bigvee_{i=1}^k \Sigma_i(f_E(\boldsymbol{g})) = 1 \wedge \Psi_i(f_E(\boldsymbol{g})) \neq 1$$

and since $F_2 \vDash \Phi$ it follows that $L \vDash \bigvee_{i=1}^k \Sigma_i(\boldsymbol{g}) = 1 \wedge \Psi_i(\boldsymbol{g}) \neq 1$. Repeating this argument for every $\boldsymbol{g} \in L^n$ we get that $L \vDash \Phi$.

5. For this assertion we adopt the approach of Champetier and Guirardel [33, Theorem 5.1]. Since every $\forall$-sentence is the negation of an $\exists$-sentence and vice-versa, it is enough to show that if H has the same $\exists$-theory as F_2 then H is a limit group. Suppose that H is generated by a tuple S of arity n and work in $\mathcal{G}_n$. It is enough to find, for every $R \in \mathbb{N}$, a free group F and a generating tuple S' of F of arity n such that the balls of radius R in $X(H,S)$ and $X(F,S')$ coincide. The idea behind the proof is simple: the ball of radius R in $X(H,S)$ can be encoded by a collection of equations and inequalities in H.

   More precisely, given $R \in \mathbb{N}$, let Σ_R be the collection of all words in the variables $\boldsymbol{x} = (x_1, \dots, x_n)$. Define

   $$\Sigma^+ = \{\sigma_1 \cdot \sigma_2^{-1} | \sigma_1, \sigma_2 \in \Sigma_R,\ \sigma_1(S) = \sigma_2(S) \text{ in } H\}$$

   and similarly

   $$\Sigma^- = \{\sigma_1 \cdot \sigma_2^{-1} | \sigma_1, \sigma_2 \in \Sigma_R,\ \sigma_1(S) \neq \sigma_2(S) \text{ in } H\}.$$

   The collections of equations and inequalities Σ^+ and Σ^- encode the ball of radius R in H. We have that $H \vDash \bigwedge_{\sigma\in\Sigma^+}\sigma(S) = 1 \wedge \bigwedge_{\sigma\in\Sigma^-}\sigma(S) \neq 1$ and therefore the following sentence

   $$\Phi_R = \exists \boldsymbol{x} \bigwedge_{\sigma\in\Sigma^+}\sigma(\boldsymbol{x}) = 1 \wedge \bigwedge_{\sigma\in\Sigma^-}\sigma(\boldsymbol{x}) \neq 1$$

   lies in the $\exists$-theory of H. Therefore, $F_2 \vDash \Phi_R$ and there is a tuple S' in F_2 that witnesses this fact. Therefore the balls of radius R of $X(H,S)$ and $X(F = \langle S' \rangle, S')$ are isomorphic which completes the proof.

□

Remark II.1.53. The proofs of *4.* and *5.* above also work for limit groups over an equationally Noetherian group G.

Another important property of limit groups proven by Sela is the following:

Theorem II.1.54 ([105, Theorem 3.1])**.** *Every non-abelian limit group admits a non-trivial cyclic splitting.*

The proof of Theorem II.1.54 is highly involved and is beyond the scope of this preliminary subsection; however, reading Chapter III will give the reader a clear idea of how to prove it.

This property of limit groups was used by Sela to show that limit groups admit a *cyclic hierarchical structure* in [106] (or *analysis lattice* in the language of [106]). Before defining hierarchies more precisely, we briefly sketch Sela's argument: given a limit group L, take its Grushko decomposition $L_1 * \cdots * L_k * F$; if some L_i is not a free abelian or a surface group, take a cyclic JSJ decomposition of L_i (see Section II.3 for more details; also, note that such a splitting is not trivial by Theorem II.1.54) and repeat the process for the vertex groups of this decomposition. Note that accessibility arguments bound the complexity of such a JSJ splitting, so that it has finitely many vertices and edges. One now repeats the process, which stops when the terminal groups are all either free, free abelian or surface groups. Sela showed that this process eventually terminates by proving that a certain measure of complexity decreases in this iterative process (roughly speaking, the minimal first Betti number of a limit group in which a group embeds decreases as one travels down a branch of the analysis lattice) [106, Proposition 4.3]. An important Corollary of this argument is that limit groups can be built from free abelian groups by repeatedly taking free products, amalgamations and HNN extensions; this implies:

Theorem II.1.55. *Limit groups are finitely presented.*

Corollary II.1.56. *Since every finitely generated subgroup of a limit group is itself a limit group, limit groups are coherent.*

We now turn to define cyclic group hierarchies:

Definition II.1.57. A *cyclic hierarchy* of a group G is a set $\mathcal{H}(G)$ of subgroups of G obtained by iterating the following procedure, starting with G: for $H \in \mathcal{H}(G)$, if H admits a splitting $\mathcal{G}(H)$ with (possibly trivial) cyclic edge groups, add each of the vertex groups of this splitting to $\mathcal{H}(G)$.

As previously mentioned, a priori the process described in Definition II.1.57 above does not necessarily terminate. If this process comes to a halt after finitely many steps (in which case $\mathcal{H}(G)$ is finite), as in the case of limit groups, G is said to have a *finite hierarchy*. The groups appearing at the bottom of the hierarchy, that is groups which do not split over a cyclic subgroup, are called *absolutely rigid*. Therefore, in modern terms, the discussion above can be summarized as follows:

Theorem II.1.58. *Limit groups have a finite hierarchy. Moreover, absolutely rigid groups appearing in the hierarchy of a limit group are free abelian.*

Note that every abelian subgroup of a limit group, which is not isomorphic to $\mathbb{Z}$ or $\mathbb{Z}^2$, must be elliptic in its cyclic JSJ decomposition. Therefore, such subgroups of a limit group must appear as terminal vertex groups in its hierarchy. By [106, Proposition 4.3], the first Betti number of a group appearing at the bottom of the hierarchy of a limit group L is bounded by $b_1(L) < \infty$; we deduce:

Corollary II.1.59. *Abelian subgroups of limit groups are finitely generated and free abelian.*

We further deduce:

Corollary II.1.60. *Let L be a limit group. Then,*

1. *If all of the abelian subgroups of L are cyclic (which by Theorem II.1.59 is equivalent to L not containing $\mathbb{Z}^2$ as a subgroup) then L is hyperbolic. This follows from Corollary II.1.59 above, the Bestvina-Feighn Combination Theorem [16], commutative transitivity (see Lemma II.1.52) and the fact that L can be built from infinite cyclic groups by taking free products, amalgamations and HNN extensions.*
2. *If L is a one-ended limit group and $H \in \mathcal{H}(L)$ is a one-ended group with no one-ended groups below it in the hierarchy, then one of the following holds:*
   - *(a) H is absolutely rigid, and therefore free abelian, or*
   - *(b) H is a hyperbolic group that splits as a graph of free groups amalgamated along infinite cyclic subgroups.*

Another important type of structure that can be associated to limit groups is that of *ω-residually free towers*. ω-residually free towers are not used explicitly in this dissertation, but they form an indispensable tool in the study of limit groups so we chose to mention them here.

Definition II.1.61 ([106, Definition 6.1]). An *ω-residually free tower* X of height $h \in \mathbb{N} \cup \{0\}$ is a space constructed with the following iterative process:

1. if $h = 0$, then X is a wedge sum of finitely many finite graphs, finite-dimensional tori and closed hyperbolic surfaces of Euler characteristic less than -1.
2. if $h > 0$ then X is obtained from a tower Y of heigh $h - 1$ by attaching a *block*; blocks are in one of the following two forms:
   (a) *Quadratic block*: a compact, hyperbolic surface with boundary Σ, with connected components being either punctured tori or admitting Euler characteristic less than -1, is attached to Y by identifying the boundary of Σ with curves on Y. The resulting space is X. We also require that there is a retraction $r : X \to Y$ such that $r_*(\partial(\Sigma))$ is non-abelian.
   (b) *Abelian block*: an n-torus T is attached to Y in the following way: fix a coordinate circle t in T and a curve γ in Y such that γ generates a maximal abelian subgroup in $\pi_1(Y)$. Let T^2 be a 2-torus and identify its coordinate circles $S^1 \times \{0\}$ and $S^1 \times \{1\}$ with t and γ respectively.

It is not hard to see that if X is an ω-residually free tower then $\pi_1(X)$ is a limit group; the proof is incredibly similar to Examples II.1.46 and II.1.50. It is much harder to see that every limit group is a subgroup of an ω-residually free tower [107]. Finally, it is extremely difficult to see that a finitely generated group G is elementarily equivalent to F_2 if and only if it is the fundamental group of an ω-residually free tower X constructed without using any tori [108].

In addition to having pleasing structural properties, limit groups also have good separability properties. This is perhaps not surprising, since limit groups are closely related to free groups, and free groups admit good separability properties. We recall Marshall Hall's famous theorem which states that given a finitely generated free group F and a finitely generated subgroup H, there is a finite-index subgroup $F' \leq F$ in which H is a free factor. In particular, F' retracts onto H so F admits local retractions and is subgroup separable (see Section I.3). Wilton proved that a similar phenomenon occurs in limit groups:

Theorem II.1.62 ([116, Theorems A and B]). *Limit groups are subgroup separable and admit local retractions.*

These results will be extensively used in Chapter V.

II.1.6 Residually free groups

Unlike fully residually free groups, residually free groups can have a wild structure. Many of the pathologies of residually free groups already appear within (the seemingly simple) $F_2 \times F_2$; these pathologies include, for example, incoherence and not being subgroup separable. The following lemma illustrates some of these ideas:

Lemma II.1.63. *$F_2 \times F_2$ has a finitely generated subgroup H of infinite index, which is dense in the profinite topology. In particular, it is not subgroup separable.*

Proof. The construction of H presented below relies on the fact that Thompson's group V, which is one of the most famous examples of infinite simple groups, admits a finite presentation with two generators u and v and 7 relations [19, Theorem 1.3]. We remark that it is enough to use a group that admits a finite presentation on 2 generators that is not simple, but has no finite-index subgroups.

Fix an epimorphism $f : F_2 \twoheadrightarrow V$ and let H be the corresponding fibre product, that is

$$H = \{(g, h) \in F_2 \times F_2 | f(g) = f(h)\}.$$

Since V is finitely presented, H is finitely generated [13, Lemma 2], and H is easily seen to be of infinite index since V is infinite. Finally, to show that H is profinitely dense in $F_2 \times F_2$, it is enough to show that it is not contained in a proper finite index subgroup G of $F_2 \times F_2$. Indeed, if it were contained in such G, then $G/(\ker f \times \ker f)$ would be a proper finite index subgroup of $V \times V$, but these do not exist since V is simple. ☐

Residually free groups are closely tied to direct products of limit groups, and arguments similar to the one above are common in the study of residually free groups. Our segue into the exploration of the relation between residually free groups and direct products of limit groups will be facilitated by a surprising property shared by all finitely generated groups that are *not* limit groups. The reason this proof works only for groups that are not limit groups, is that a limit quotient of a group that is not a limit group must be a proper quotient.

Theorem II.1.64 ([106, Theorem 7.2])**.** *Let G be a finitely generated group that is not a limit group. Then there are finitely many limit groups $L_1, \dots, L_n$, and quotient maps $q_i : G \twoheadrightarrow L_i$, such that every epimorphism f from G to a limit group L factors through one of the maps $q_1, \dots, q_n$.*

Definition II.1.65. The collection of quotients $\{q_i : G \twoheadrightarrow L_i\}_{1 \le i \le n}$ mentioned in Theorem II.1.64 is called a *factor set* of G.

The proof of Theorem II.1.64 below follows the proofs of [106, Lemmas 5.4 and 5.5, Proposition 5.6 and Theorem 5.7], where the theorem is proved under the assumption that G *is* a limit group. We remark that the proof in the case where G is a limit group is a lot more intricate.

Proof of Theorem II.1.64. The factor set $\{q_i : G \twoheadrightarrow L_i\}_{1 \le i \le n}$ that we construct consists of maximal elements in a poset Q_G of limit quotients of G. We only consider such quotients up to equivalence, where two quotients $q_1 : G \twoheadrightarrow L_1$ and $q_2 : G \twoheadrightarrow L_2$ are equivalent if there is an isomorphism $f : Q_1 \to Q_2$ such that $f \circ q_1 : L \twoheadrightarrow L_2$ and $q_2 : L \twoheadrightarrow L_2$ are the same. We say that

$$(q_1 : G \twoheadrightarrow L_1) \le (q_2 : G \twoheadrightarrow L_2)$$

if and only if the quotient map q_1 factors via q_2. A straightforward application of Zorn's lemma shows that Q_G must contain a maximal element. Indeed, a diagonalization argument shows that every chain $q_1 < q_2 < \cdots$ in Q_G admits an upper bound. For each $q_i : G \twoheadrightarrow L_i$ there is an associated sequence of homomorphisms $(f^i_n : G \to F)_{n \in \mathbb{N}}$; we extract a subsequence $(h^i_n : G \to F)_{n \in \mathbb{N}}$ such that h^i_n is injective on $\{g \in G | g \notin \ker q_i \text{ and } |g|_S \le n\}$. The diagonal sequence $(h^i_i : G \to F)$ gives rise to a limit group L and a quotient $q : G \twoheadrightarrow L$, and since L is finitely presented, q_i factors via q for large enough i. Therefore every q_i factors via q and $q : G \twoheadrightarrow L$ is an upper bound for the chain. Since G is not a limit group, q is a proper quotient and it is an element of Q_G.

Therefore, every quotient $q : G \twoheadrightarrow L$ factors through a maximal element in Q_G, and it is enough to show that there are only finitely many such maximal elements. An identical diagonalization argument gives this result: if $\{q_i : L \twoheadrightarrow L_i\}_{i \in \mathbb{N}}$ is a collection of maximal elements in Q_G, and $q_\infty : L \twoheadrightarrow L_\infty$ is obtained from this collection by the same diagonalization process, then all but finitely many q_i factor through q_∞; since q_i is maximal, it follows that $q_i = q_\infty$ and $L_i = L_\infty$.

□

Theorem II.1.64 is a rare instance in which being fully residually free presents an obstruction to a proof, rather than being a simplifying assumption. However, a similar property does hold for (non-free) limit groups. This property is commonly known as having "finite width", and this name is derived from applying this corollary to the study of homomorphisms from a fixed finitely generated group G to a free group. All such homomorphisms can be encoded within a tree-like diagram called the *Makanin Razborov diagram*, and the fact that limit groups have finite width implies that each branching point of the Makanin Razborov diagram is of finite valence. As mentioned earlier, the difficulty in establishing a similar result for limit groups lies in the fact that the identity map $\mathrm{Id} : L \twoheadrightarrow L$ is the unique maximal element in the poset of limit quotients of L. One overcomes this obstacle by looking at *shortening quotients*, which are quotients obtained from sequences of *short* homomorphisms (see Section III.1 in Chapter III).

Theorem II.1.66 ([106, Theorem 5.7])**.** *To any non-free limit group L one can associate a finite collection of proper quotients $\{q_i : L \twoheadrightarrow L_i\}_{1\le i\le n}$ such that each L_i is a limit group, and each homomorphism $f : L \to F$ factors through (at least) one of the maps q_i after precomposing L with an automorphism.*

Remark II.1.67. The automorphisms of L appearing in Theorem II.1.66 are *modular automorphisms*; these are defined in Section III.1 in Chapter III.

Corollary II.1.68. *Every finitely generated residually free group embeds into a direct product of finitely many limit groups.*

Proof. Let G be a finitely generated residually free group. If G is a limit group, the assertion clearly holds. Otherwise, let $\{q_i : G \twoheadrightarrow L_i\}_{1\le i\le n}$ be the factor set associated to G and consider the map $q = q_1 \times \cdots q_n : G \to L_1 \times \cdots L_n$. It is enough to show that $\ker q = \bigcap_{i=1}^{n} \ker q_i$ is trivial. Indeed, if $1 \ne g \in G$, since G is residually free there is a free quotient $q' : G \twoheadrightarrow F$ of G in which g survives. By Theorem II.1.64 q' factors through q_i for some $1 \le i \le n$, and therefore $q_i(g) \ne 1$. It follows that $g \notin \ker q$. ☐

Direct products of free groups, or more generally direct products of limit groups, are fairly tractable; however, their subgroups can be very complicated. Baumslag and Roseblade studied subgroups of direct products of free groups, and proved:

Theorem II.1.69 ([13, Theorem 1])**.** *There are uncountably many non-isomorphic subgroups of $F_2 \times F_2$.*

It follows that unlike limit groups, $F_2 \times F_2$ has (uncountably many) subgroups that are not finitely presented (in fact, it has uncountably many subgroups that are not even recursively presented). However, under additional *finiteness* assumptions, the situation becomes completely different:

Theorem II.1.70 ([13, Theorem 2]). *If $H \leq F_2 \times F_2$ is finitely presented, then it is virtually a direct product of two free groups of finite rank.*

Bridson, Howie, Miller and Short took these ideas further in [25], and described how groups can be embedded into a direct product of limit groups; these embeddings are dictated by the *finiteness properties* of the embedded subgroups. Before we state their results, we briefly remind the reader of some *finiteness properties of groups*.

Definition II.1.71. A group G is said to be of *type*

- $\mathtt{F}_n$ $(n \geq 1)$, if there is an aspherical CW-complex X with a finite n'th skeleton such that $G = \pi_1(X)$.
- $\mathtt{F}_\infty$, if it is of type F_n for every $n \geq 1$.
- $\mathtt{F}$, if it is the fundamental group of a finite, aspherical CW-complex.
- $\mathtt{FP}_n(R)$ $(n \geq 1$ and R is a ring), if there is an exact sequence of $R[G]$-modules $\cdots \to M_1 \to M_0 \to R \to 0$ in which $M_0, \ldots, M_n$ are finitely generated, projective $R[G]$-modules.
- $\mathtt{FP}_\infty(R)$, if it is of type $\mathtt{FP}_n(R)$ for every $n \geq 1$.
- $\mathtt{FP}(R)$, if there is an exact sequence $0 \to M_n \to \cdots M_0 \to R \to 0$ in which every M_i is a finitely generated, projective $R[G]$-module.

It is worth noting that among residually free groups, being of type $\mathtt{FP}_2(\mathbb{Q})$ is equivalent to being finitely presented [26, Theorem D]; this strengthens the analogy between Theorem II.1.70 and the following:

Theorem II.1.72 ([25, Theorem A]). *Let $L_1, \cdots, L_n$ be limit groups and let $S \leq L_1 \times \cdots \times L_n$ be a subgroup of type $\mathtt{FP}_n(\mathbb{Q})$. Then S is virtually a direct product of n (or fewer) limit groups. In particular, every residually free group of type $\mathtt{FP}_\infty(\mathbb{Q})$ is virtually a direct product of finitely many limit groups.*

Another result of Bridson, Howie, Miller and Short allows one to regard finitely presented residually free groups as a subgroup of a direct product of limit groups that contains a term of the lower central series. This result will be used in Chapter V.

Theorem II.1.73 (Propositions 3.1 and 6.4, [25])**.** *Let G be a finitely presented subgroup of a direct product of limit groups $L_1 \times \cdots \times L_n$, that projects onto each L_i and such that $G \cap L_i \neq 1$ for all $1 \leq i \leq n$. Then there exists a finite-index subgroup $E \leq L_1 \times \cdots \times L_n$ and a positive integer N such that the N-th term in the lower central series of E, $\gamma_N E$, is contained in $E \cap G$.*

Residually free groups also admit nice separability properties, which are again linked to finiteness properties. The following theorem will also play an important role in Chapter V:

Theorem II.1.74 ([27, Theorems A and B])**.** *Let G be a finitely generated residually free group, and let $H \leq G$. If H is finitely presented, then H is separable in G; if furthermore H is of type $\mathtt{FP}_\infty(\mathbb{Q})$ then G virtually retracts onto H.*

II.2 Graphs of spaces

As evident from the last section, splittings (and in particular cyclic splittings) are in abundance when working with limit groups. In this section, we introduce the notion of a *graph of spaces* which will allow us to incorporate additional geometric techniques when working with group splittings. The standard reference for graphs of spaces is Scott's and Wall's "Topological methods in group theory" [102, Chapter 5]; we adopt some of their terminology, along with more modern terminology following Wilton's works ([116], [119]).

Definition II.2.1. A *graph of spaces* X consists of the following data:

- a graph Ξ,
- for each vertex $v \in \mathrm{V}(\Xi)$ a connected CW-complex X_v,
- for each edge $e \in \mathrm{E}(\Xi)$ a connected CW-complex X_e and two π_1-injective maps $\partial_e^{\pm} : X_e \to X_{v^{\pm}}$ (called *attaching maps*) where v^+ and v^- are the endpoints of the edge e.

We will usually assume all graphs of spaces to have a finite underlying graph. Note that each graph of spaces X has a topological space naturally associated to it:

Definition II.2.2. The *geometric realization* of X is defined as

$$\left(\bigsqcup_{v \in \mathrm{V}(\Xi)} X_v \sqcup \bigsqcup_{e \in \mathrm{E}(\Xi)} X_e \times [-1,1] \right) \Big/ \sim$$

where the equivalence relation $\sim$ identifies $(x, \pm 1) \in X_e \times [-1,1]$ with $\partial_e^{\pm}(x) \in X_{v^{\pm}}$.

Remark II.2.3. We will often abuse notation and use X to refer to the geometric realization of X.

With this in mind, we define:

Definition II.2.4. The *fundamental group* $\pi_1(X)$ of a graph of spaces X is simply the fundamental group of the geometric realization of X.

By the Seifert–Van Kampen theorem, the fundamental group of X admits a graph of groups decomposition: $\pi_1(X)$ is the fundamental group of the graph of groups with underlying graph Ξ and whose vertex and edge groups are $\{\pi_1(X_v) | v \in \mathrm{V}(\Xi)\}$ and $\{\pi_1(X_e) | e \in \mathrm{e}(\Xi)\}$ respectively. The edge maps of this graph of groups are given, up to conjugation, by $(\partial_e^{\pm})_* : \pi_1(X_e) \to \pi_1(X_{v^{\pm}})$. We will denote this graph of groups decomposition by $\mathcal{G}(X)$ and refer to the vertex and edge groups of $\mathcal{G}(X)$ as G_v and G_e.

II.2.1 Covering spaces and precovers

The (finite-index) subgroups of the fundamental group G of a graph of spaces X correspond to (finite-sheeted) coverings of X. In Chapter V we will construct subgroups of fundamental groups of graphs of spaces geometrically, heavily relying on the contents of this subsection. Most of the theory covered in this subsection appears in works of Wise [123] and Wilton [119], and some key ideas date back to Gitik's work [51]. We begin by explaining how a covering space of the geometric realization of a graph of spaces admits a graph of spaces decomposition.

Let X be a graph of spaces with an underlying graph Ξ, vertex spaces $\{X_v\}$, edge spaces $\{X_e\}$ and attaching maps $\partial_e^{\pm} : X_e \to X_{v^{\pm}}$. Suppose now that $\widehat{X}$ is a covering of (the geometric realization of) X. $\widehat{X}$ inherits a graph of spaces structure from X, with the following vertex and edge spaces:

- The vertex spaces of $\widehat{X}$ are the connected components of the preimages of the vertex spaces of X under the covering map; each of these forms a covering space of a vertex space of X.

- The edge cylinders of $\widehat{X}$ are the connected components of the preimages of the edge cylinders of X under the covering map; again, each edge space of $\widehat{X}$ is a covering space of an edge space of X.

From this, it is evident that the underlying graph $\widehat{\Xi}$ of $\widehat{X}$ can be obtained from $\widehat{X}$ by collapsing each vertex space to a point and each edge cylinder to an arc. Defining the attaching maps of $\widehat{X}$ is a bit more subtle. The following definition is due to Wise [123].

Definition II.2.5. An *elevation* of an attaching map $\partial_e^\pm : X_e \to X_{v^\pm}$ is defined as follows: let $\widehat{X}_{\widehat{v}}$ be a vertex space of $\widehat{X}$ which lies above X_v. Let $\{\widehat{X}_{\widehat{e}_1}, \dots, \widehat{X}_{\widehat{e}_k}\}$ be the edge spaces of $\widehat{X}$ such that each $\widehat{e}_i \in \mathrm{E}(\widehat{\Xi})$ is in the preimage of $e \in \mathrm{E}(\Xi)$ and incident to $\widehat{v}$. This data fits into the following commuting diagram:

$$\begin{tikzcd}
\bigsqcup_{i=1}^k \widehat{X}_{\widehat{e}_i} \arrow[r, dashed, "\sqcup_{i=1}^k \widehat{\partial}^\pm_{\widehat{e}_i}"] \arrow[d] & \widehat{X}^\pm_{\widehat{v}} \arrow[d] \\
X_e \arrow[r, "\partial_e^\pm"] & X_v{}^\pm
\end{tikzcd}$$

where the maps $\bigsqcup_{i=1}^k \widehat{\partial}^\pm_{\widehat{e}_i}$ form the *pullback* of the attaching maps $\partial_e^\pm$ and the covering map. The pullback can be thought of as satisfying the following property: fix a bsaepoint x of X_e, and for every $1 \le i \le k$ fix a basepoint $\widehat{x}_i$ in the preimage of x under the covering map $\widehat{X}_{\widehat{e}_i} \twoheadrightarrow X_e$. Every $\widehat{\partial}^\pm_{\widehat{e}_i} : \widehat{X}_{\widehat{e}_i} \to \widehat{X}^\pm_{\widehat{v}}$ is a lift of the map

$$(\widehat{X}_{\widehat{e}_i}, \widehat{x}_i) \twoheadrightarrow (X_e, x) \xrightarrow{\partial_e{}^\pm} (X_v{}^\pm, \partial_e{}^\pm(x))$$

to $\widehat{X}^\pm_{\widehat{v}}$ which is minimal in the following sense: for every intermediate cover

$$(\widehat{X}_{\widehat{e}_i}, \widehat{x}_i) \twoheadrightarrow (X'_e, x') \twoheadrightarrow (X_e, x)$$

the map $(X'_e, x'_e) \twoheadrightarrow (X_e, x) \xrightarrow{\partial_e^\pm} (X_v^\pm, \partial_e^\pm(x))$ does not lift to a map

$$(X'_e, x') \to (\widehat{X}_{\widehat{v}}, \widehat{\partial}^\pm_{\widehat{e}_i}(\widehat{x}_i)).$$

This is better illustrated in Example II.2.8. Finally, an elevation is a restriction of this map to one of the $\widehat{X}_{\widehat{e}_i}$, and each elevation is an attaching map of $\widehat{X}$.

Our next goal is to explain how can one construct covering spaces of graphs of spaces from a collection of covers of the different vertex spaces; such a construction is reminiscent of assembling a jigsaw puzzle, and the following definition is essential for explaining when two pieces can "interlock".

Definition II.2.6. The *degree* of an elevation $\partial^\pm_{\widehat{e}_i}$ is the conjugacy class of $\pi_1(\widehat{X}_{\widehat{e}_i})$ in $\pi_1(X_e)$.

Remark II.2.7. If $\pi_1(X_e)$ is abelian, conjugacy classes of subgroups contain a single element, and the degree of an elevation can be simply thought of as $\pi_1(\widehat{X}_{\widehat{e}_i})$ itself. When X_e is a circle, the degree of $\partial^{\pm}_{\widehat{e}_i}$ can be thought of as the degree of the covering map $\widehat{X}_{\widehat{e}_i} \to X_e$.

Note that covers of different vertex spaces that are attached by an edge cylinder in X can be attached by some edge cylinder in a cover $\widehat{X}$ of X if and only if the corresponding elevations to $\widehat{X}$ have the same degree. The following example illustrates how to assemble a covering space of a graph of spaces from covering spaces of the vertices.

Example II.2.8. Consider the group $G = \langle a, b, a', b' | [a,b]^2 = [a', b']^2 \rangle$. G admits a graph of spaces decomposition X, where the two vertex spaces are surfaces with boundary and the single edge space is a circle. This is better illustrated with the following drawing:

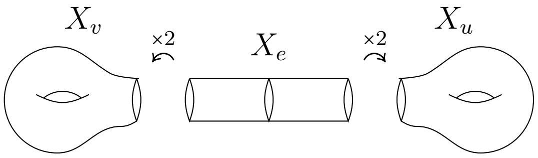

Figure II.1: A graph of spaces decomposition X of G. The "×2" inscriptions signify that the attaching maps $X_e \to X_v$ and $X_e \to X_u$ wrap the boundary components of X_v and X_u twice (and hence both of are degree two).

For illustrative purposes, we construct a cover $\widehat{X}$ of X from the ground up, that is we construct covers of X_v, X_u and X_e that fit into commuting diagrams as in Definition II.2.5 above. Each of X_v and X_u admit a 4-fold cover that is a surface with two boundary components; we denote these by $\widehat{X}_v$ and $\widehat{X}_u$. Each attaching map admits four elevations with target in $\widehat{X}_v$ or $\widehat{X}_u$, and each of these elevations is of degree one. This information fits in the pullback diagram in Figure II.2.

Note that $\bigsqcup_{i=1}^4 \widehat{\partial}^+_{\widehat{e}_i}$ indeed forms a pullback: any space that maps to both X_e and $\widehat{X}_{\widehat{v}}$ and fits in a suitable commuting diagram must factor through one of the four $\widehat{X}_{\widehat{e}_i}$.

It is left for us to choose a pairing between the two sets of elevations $\{\widehat{\partial}^+_{\widehat{e}_i}\}_{1\leq i \leq 4}$ and $\{\widehat{\partial}^-_{\widehat{e}_i}\}_{1\leq i \leq 4}$; any pairing will yield a 4-fold covering space of X. In Figure II.3 below we illustrate the geometric realization of one such possible 4-fold cover $\widehat{X}$.

We are now ready to define *precovers*, which are spaces that admit a construction similar to that in Examples II.2.8 but are not covering spaces. These were first defined

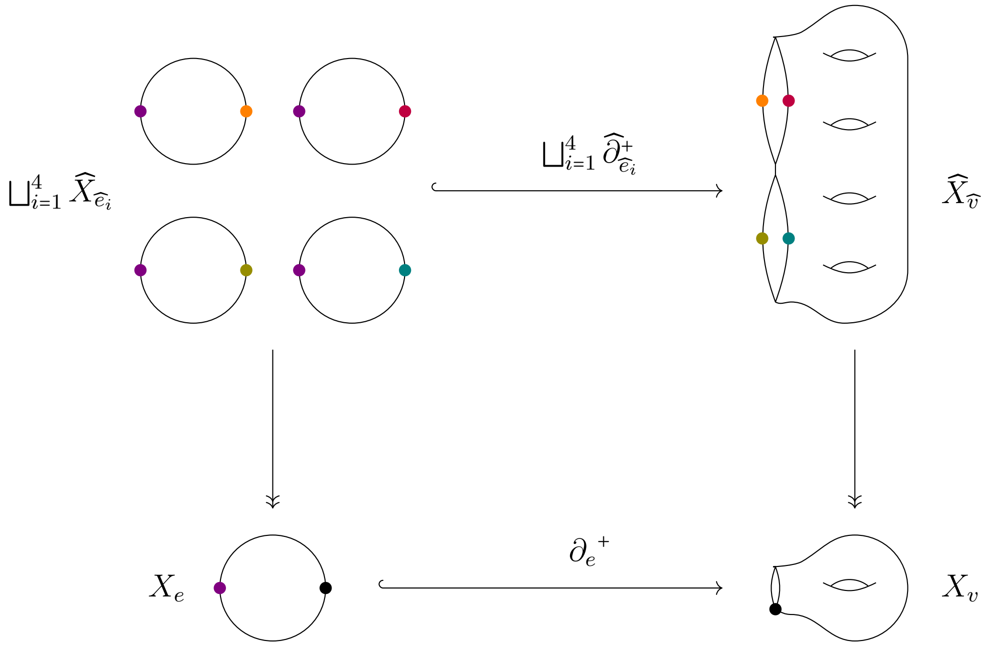

Figure II.2: A cover $\widehat{X}_{\widehat{v}}$ of X_v, and the corresponding pullback diagram. Each of the enlarged dots maps to another enlarged dot (of the same colour, when possible).

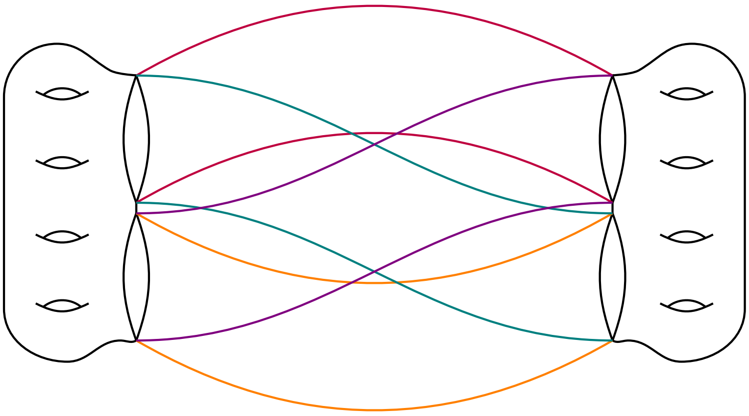

Figure II.3: The geometric realization of a 4-fold cover $\widehat{X}$ of X. For interpretability, different edge cylinders appear in different colours.

by Gitik in [51], and the definition that we use first appeared in Wilton's work [116]. Informally, a precover X' of X is a graph of spaces which partially covers X: the vertex and edge spaces of X' cover those of X and the degrees of the two ends of each attaching map coincide. A good picture to keep in mind when thinking of a precover, is that of the *core graph* corresponding to a finitely generated subgroup of a free group; Lemma II.2.12 should make the analogy clearer. We begin with a constructive definition of precovers:

Definition II.2.9. Let X be a graph of spaces. Construct another graph of spaces X' in the following manner: let $\{X'_{v'} \to X_v\}$ be a collection of covering maps of vertex spaces of X and for each $e \in \mathrm{E}(\Xi)$ let $\{\partial^+_{e'_i}\}_{i\in I_e}$ and $\{\partial^-_{e'_i}\}_{i\in I_e}$ (where I_e is some indexing set) be subsets of the elevations of ∂^+_e and ∂^-_e to $\bigsqcup_{v'} X_{v'}$. Assume further that the degrees of $\partial^+_{e'_i}$ and $\partial^-_{e'_i}$ coincide for every $e \in \mathrm{E}(\Xi)$ and every $i \in I_e$. Now let X' be the graph of spaces obtained by gluing the collection $\{X'_{v'}\}$ along the attaching maps $\{\partial^\pm_{e'_i}\}$. A graph of spaces obtained in this manner is called a *precover* of X. Note that X' comes equipped with a locally injective map $X' \to X$. Elevations of attaching maps of X to X' which are not attaching maps of X' are called *hanging elevations.*

One can also use the following equivalent definition:

Definition II.2.10. Let X and X' be graphs of spaces. We say that X' is a *precover* of X if there is a locally injective map $p' : X' \to X$ that satisfies the following properties:

1. if $X'_{v'}$ is a vertex space of X', then $f|_{X'_{v'}}$ is a covering map that maps $X'_{v'}$ onto some vertex space X_v of X;
2. similarly, if $X'_{e'}$ is an edge space of X', then $f|_{X'_{e'}}$ is a covering map that maps $X'_{e'}$ onto some edge space X_e of X;
3. for any attaching map ${\partial'_{e'}}^{\pm} : X'_{e'} \to X'_{v'^\pm}$, the following diagram commutes:

$$\begin{tikzcd}
X'_{e'} \arrow[r, "{\partial'_{e'}}^{\pm}"] \arrow[d, "f|_{X'_{e'}}"'] & X'_{v'^\pm} \arrow[d, "f|_{X'_{v'^\pm}}"] \\
X_e \arrow[r, "\partial^\pm_e"] & X_{v^\pm}
\end{tikzcd}$$

Example II.2.11. We keep the setting of Example II.2.8, and construct a precover X' of X that is a closed and orientable surface of genus 9. Let $\widehat{X}_{\widehat{v}}$ and $\widehat{X}_{\widehat{u}}$ be the covers of X_v and X_u given in Example II.2.8, and note that all of the elevations of ${\partial_e}^{\pm}$ to $\widehat{X}_{\widehat{v}}$ and $\widehat{X}_{\widehat{u}}$ are of degree 2; we choose two elevations $\widehat{\partial}^+_{\widehat{e}_{1,2}}$ to $\widehat{X}_{\widehat{v}}$ and two elevations $\widehat{\partial}^-_{\widehat{e}_{1,2}}$ to $\widehat{X}_{\widehat{u}}$ and glue them together as in Figure II.4 below.

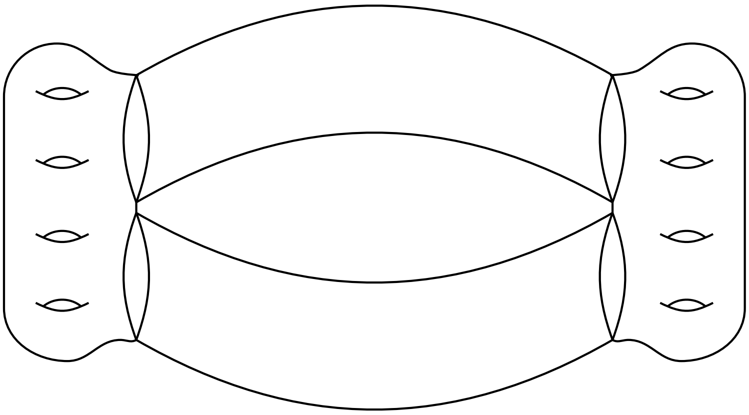

Figure II.4: The precover X' is a closed and orientable surface of genus 9.

The key observation regarding precovers is that precovers of X give rise to subgroups of X: each precover X' can be extended to a covering space by attaching pieces that do not contribute additional information to the fundamental group, resulting in a covering space with fundamental group $\pi_1(X')$. For the sake of completeness, we include a proof of this assertion below:

Lemma II.2.12 ([116, Proposition 2.19] and [119, Lemmas 15 and 16])**.** *If X' is a precover of X then the natural map $X' \to X$ can be extended to a covering map $\widehat{X} \to X$ such that the induced map $\pi_1(X') \to \pi_1(\widehat{X})$ is an isomorphism. Therefore the induced map $\pi_1(X') \to \pi_1(X)$ is injective.*

Proof. Note that a precover without any hanging elevations is simply a cover; we will therefore "cap off" each hanging elevation of X' and obtain the suitable cover $\widehat{X}$ of X. Let $\partial_{e'} : X'_{e'} \to X'_{v'}$ be a hanging elevation of an attaching map $\partial_e : X_e \to X_v$ of X to X'; we will attach another precover of X to X' via $\partial_{e'}$. Repeating this for all hanging elevations to X' will result in the desired cover (if there are infinitely many hanging elevations to X', one can simply take the corresponding ascending union).

Let $\widehat{X}_{e'}$ be the covering space of X with fundamental group $\pi_1(X'_{e'})$, and recall that it inherits a graph of spaces decomposition from X. This decomposition contains an edge space with fundamental group $\pi_1(X'_{e'})$. Detach the corresponding edge cylinder by unidentifying its ends with their target in the vertex groups of $\widehat{X}_{e'}$, and note that each connected component Y of the resulting space satisfies:

1. $\pi_1(Y) = \pi_1(X'_{e'})$, and

2. there is a single hanging elevation ∂_Y of either ${\partial_e}^+$ or ${\partial_e}^-$ to Y (where ${\partial_e}^\pm$ are the two attaching maps corresponding to the edge $e \in \mathrm{E}(\Xi)$).

Assume without loss of generality that $\partial_{e'}$ is an elevation of ${\partial_e}^+$, and choose such Y with ∂_Y an elevation of ${\partial_e}^-$ to Y. Glue Y to X' (note that the degree of ∂_Y coincides with that of $\partial_{e'}$). This yields a precover X'' of X that extends X', and such that the inclusion map $X' \to X''$ induces a monomorphism of fundamental groups, as desired. $\square$

The next subsection will focus on a certain kind of graphs of spaces: graphs of graphs amalgamated along circles.

II.2.2 Graphs of free groups with cyclic edges

Graphs of free groups amalgamated along cyclic subgroups play a special role in the study of limit groups: they are the groups which lie at the second level of the hierarchy of a hyperbolic limit group, above free groups (for more details, see Subsection II.1.5). These groups also admit a natural graph of spaces decomposition, where the vertex spaces are graphs and the edge spaces are circles.

Remark II.2.13. Since the spine of a (connected) compact surface with boundary is a (connected) finite graph, we will sometimes (when applicable) refer to vertex spaces of such a graph of spaces decomposition as surfaces (with boundary), attached to other vertices along boundary components. This is better illustrated in Example II.2.8.

Note that by Theorem II.1.62, every finitely generated subgroup of a limit group is a local retract, and therefore quasiconvex (since a retract of a group can always be quasi-isometrically embedded into the ambient group). Since quasiconvex subgroups of hypebrolic groups are hyperbolic, the (fundamental groups of the) graphs of free groups mentioned in the previous paragraph are hyperbolic themselves. We therefore focus on *hyperbolic* graphs of free groups amalgamated along cyclic subgroups. Note that since hyperbolic groups are finitely presented, these graphs of free groups must have finitely generated vertex groups and a finite underlying graph.

Notation II.2.14. For abbreviation, we will refer to **H**yperbolic fundamental groups of **G**raphs of **F**ree groups with infinite **C**yclic edge groups as *HGFC-groups.*

We continue by mentioning a few separability properties of HGFC-groups; these make HGFC-groups a convenient object to be working with, especially in the context of constructing finite-sheeted covers and precovers. By work of Wise, HGFC-groups admit nice separability properties:

Theorem II.2.15 ([123]). *Let G be the fundamental group of a finite graph of finitely generated free groups amalgamated along infinite cyclic subgroups. Then G is subgroup separable unless there exists $g \in G$ such that g^n is conjugate to g^m for $n \neq \pm m$. In particular, if G is an HGFC-group, it is subgroup separable.*

Hsu and Wise went on and showed in [68] that such groups are in fact virtually compact special (the definition of a special group is highly technical and, apart for its consequences in this context, does not play a role in this dissertation; we therefore refer the reader to Haglund's and Wise's work [66] for the precise definition). The benefits that follow from being virtually compact special include:

Theorem II.2.16 (follows from [66, Corollary 6.7 and Theorem 7.3]). *Hyperbolic compact special groups virtually retract onto their quasiconvex subgroups.*

In addition, HGFC-groups enjoy the following property:

Theorem II.2.17 (follows from [18, Theorem D]). *Every HGFC-group is locally quasiconvex.*

These together imply:

Corollary II.2.18. *Let G be an HGFC-group. Then G has a finite-index subgroup which admits local retractions.*

Constructing precovers of (graphs of spaces associated to) HGFC-groups is one of the main techniques used in Chapter V; in particular, we will be interested in constructing precovers that are surfaces (or branched surfaces, in which the neighbourhood of a point on the boundary of a vertex space is not isomorphic to a plane). Recall that in Example II.2.8, we started with a graph of spaces that resembled a surface, but its attaching maps were not homeomorphisms; that is, each attaching map wrapped a boundary component of the surface multiple times. We "unwrapped" these attaching maps in a finite cover, and as a result constructed a precover that was homeomorphic to a closed surface in Example II.2.12. This is a common technique used when working with graphs of free groups. In the next subsection we will prove a standard lemma (Lemma II.2.28) that generalizes the constructions in Examples II.2.8 and II.2.12.

II.2.3 Peripheral structures and surface subgroups

In this subsection we discuss local-global interplay in graphs of free groups with cyclic edge groups. In other words, we explore how the *global* properties in the group in question are related to the *local* properties one sees when looking at a vertex and its adjacent edges. Recall that a collection of non-trivial words $\underline{w}$ of a free group F is called a *multiword*.

Definition II.2.19. A *peripheral structure* on a free group F is a set of pairwise non-conjugate maximal cyclic subgroups of F.

Remark II.2.20. Note that every multiword $\underline{w} = (w_1, \dots, w_n)$ gives rise to a peripheral structure on F: for each w_i pick a maximal cyclic subgroup C_i containing w_i, and if some C_i and C_j are conjugate simply remove one of them from the list. Sometimes we will think of $[\underline{w}]$ as a collection of embedded circles in a graph Γ with $\pi_1(\Gamma) = F$.

Notation II.2.21. We denote the peripheral structure induced by a multiword $\underline{w}$ in F by $[\underline{w}]$. We will refer to a free group accompanied by a peripheral structure as a *pair* and denote this data by $(F, [\underline{w}])$.

We remark that peripheral structures can be defined on any group (as a finite collection of conjugacy classes of subgroups).

An important observation here is the following: if G admits a splitting as a graph of free groups with cyclic edge groups, then every vertex group G_v of G comes equipped with a peripheral structure induced by the edge groups corresponding to the edges adjacent to v.

Notation II.2.22. In the above setting, we denote the peripheral structure on G_v by $[\underline{w}_v]$. We refer to the pair $(G_v, [\underline{w}_v])$ as the *induced pair at* v.

Since we will be interested in constructing finite-sheeted covers and precovers of an HGFC-group G, we will also be interested in how peripheral structures behave in such circumstances. Following the discussion in Subsection II.2.1 (and Wilton's work [119]), we define:

Definition II.2.23. Let $[\underline{w}]$ be a peripheral structure on a free group F, and let F' be a finite-index subgorup of F. The *induced pair for* F' is the pair $(F', [\underline{w}'])$ obtained from a pullback as in Definition II.2.5. Alternatively, one can think of $[\underline{w}']$ as the collection of conjugacy classes of maximal cyclic subgroups of F', that are conjugate in F into one of the conjugacy classes in $[\underline{w}']$.

We continue by describing three types of pairs of particular interest:

Definition II.2.24. A pair $(F, [\underline{w}])$ is said to be

- *one-ended* if F does not split freely relative to the elements of $[\underline{w}]$, that is, the elements of $[\underline{w}]$ are hyperbolic in every free splitting of F,
- *rigid* if the elements of $[\underline{w}]$ are hyperbolic in every cyclic splitting of F, and
- of *surface type* if there is a closed surface with boundary Σ and an isomorphism $F \cong \pi_1\Sigma$ which identifies $[\underline{w}]$ with the conjugacy classes of the cyclic subgroups of $\pi_1\Sigma$ that correspond to $\partial\Sigma$.

Remark II.2.25. If $(F, [\underline{w}])$ is of surface type, and the corresponding surface is Σ, we will sometimes refer to $(F, [\underline{w}])$ as $(\Sigma, \partial\Sigma)$. If Σ is a thrice-punctured sphere, then the pair $(F, [\underline{w}])$ is also rigid; we will refer to pairs of this kind as pairs of surface type and not as rigid pairs.

All of these types play a role when studying HGFC-groups through their induced pairs at the different vertices; they will also appear in Subsection II.3 where we discuss JSJ decompositions. A classical lemma that fits well with this theme is Shenitzer's lemma, which states that an amalgamated product of two free groups along an infinite cyclic subgroup is free if and only if the edge group maps onto a primitive element in one of the vertex groups. This was generalized by Wilton as follows:

Theorem II.2.26 (Relative Shenitzer's Lemma: [119, Theorem 18] , see also [41, Corollary 7.3] and [113, Main Theorem])**.** *Let G be a finitely generated group which is the fundamental group of a graph of groups with infinite cyclic edge groups. Then G is one-ended if and only if each of its induced pair at the different vertices is one-ended.*

Before proving a couple of related lemmas, we would like to introduce a "normalization procedure" which, given a hyperbolic group G and a splitting of G as a graph of free groups amalgamated along cyclic edges, outputs a different graph of groups decomposition of G that will prove to be very convenient to work with. The idea behind this procedure is not new, and it has been used before specifically in the context of making JSJ decompositions canonical (we refer the reader to Section II.3). We remark that a detailed algebraic description of Procedure II.2.27 below, in the case where G is a free group, appears in [32, Section 2.4]; note that in our context G is not necessarily a free group, but it does not exhibit any of the possible obstructions mentioned in Cashen's work. A more general approach towards making group splittings "canonical" can be found in Guirardel's and Levitt's work [64]. The graph of groups decomposition described in Procedure II.2.27 below coincides with the decomposition one obtains from [64, Example 3.1 and Subsection 4.1].

Procedure II.2.27. Let G be a hyperbolic group that splits as a finite graph of finitely generated free groups amalgamated along cyclic subgroups, and let X be a graph of spaces decomposition of G in which all of the vertex spaces are graphs (or surfaces with boundary whenever the induced pair at a vertex is of surface type). We will replace X with a graph of spaces that will satisfy the following:

1. the underlying graph is bipartite, with each edge connecting a non-cyclic vertex (that is, a vertex with a non-cyclic vertex group) to a *cyclic vertex* (that is, a vertex with a cyclic vertex group $\mathbb{Z}$),
2. each cyclic vertex group is a maximal cyclic subgroup of G, and the degree (in the underlying graph Ξ of X) of each cyclic vertex is at least 2,
3. for each non-cyclic vertex v consider the induced pair $(G_v, [\underline{w}_v])$; then for every $C \in [\underline{w}_v]$ there is exactly one edge incident to v whose edge group is mapped onto C. In particular, for every edge e, the edge map that maps the corresponding edge group G_e into a non-cyclic vertex G_v, maps G_e onto a maximal cyclic subgroup of the vertex group.

To do so, we begin with a folding procedure: subdivide each edge space of X, adding a cyclic vertex in the middle of the edge cylinder. Let e be an edge of the resulting graph, connecting a non-cyclic vertex v to a cyclic vertex c. If the image of G_e in G_v is not a maximal cyclic subgroup of G_v, we subdivide e edge yet again by adding another cyclic vertex group c'; the two edges adjacent to c' identify $G_{c'}$ with a maximal cyclic subgroup of G_v, and G_c with a proper subgroup of $G_{c'}$. This takes care of the second part of item 3. above.

We next carry out a folding procedure iteratively, going over all non-cyclic vertices of X one at a time: if for some X_v there are multiple edge cylinders with the same image in X_v, we fold these edges and identify them to obtain a single edge; we are also identifying the cyclic vertices lying at the other end of these edges in the process. Note that in order to do so, we might need to "twist" the edge cylinders attached to these cyclic vertices, so that the attaching maps that attach them to X_v line up. An important thing to notice here, is that this folding is a homotopy equivalence. Since all of the attaching maps are isomorphisms, the only obstruction one might encounter is the following: there are two edges that connect the same loop in X_v to a single cyclic vertex X_c, and these are identified when we fold. Such a fold is clearly not a π_1-isomorphism. However, if a boundary component of some X_v is connected by two edges to a cyclic vertex, it follows that G must contain $\mathbb{Z}^2$ or $\mathrm{BS}(1,-1)$ as a subgroup, which is impossible since G is hyperbolic. We repeat this for every non-cyclic vertex $v \in \mathrm{V}(\Xi)$. The resulting graph now satisfies conditions 2. and 3. above.

To finish, we still need to make the graph bipartite (as in 1. above): in the beginning of the process, we (perhaps) divided some of the edges multiple times, adding up to 3 cyclic vertices between two vertices of the original graph of spaces decomposition X. However, note that if two cyclic vertices c and c' are connected by an edge, and both of the edge maps into G_c and $G_{c'}$ are not isomorphisms, then G contains a Baumslag-Solitar subgroup contradicting the fact that it is hyperbolic. Therefore, if any edge connects two cyclic vertices, it can be collapsed; collapsing all such edges yields the desired decomposition.

We first use Procedure II.2.27 to generalize the "unwrapping technique" from Examples II.2.8 and II.2.12.

Lemma II.2.28. *Let G be a group that splits as a finite graph of finitely generated free groups amalgamated along infinite cyclic subgroups. If G is subgroup separable then it admits a finite-index subgroup G' that admits a graph of spaces decomposition X' with the three properties guaranteed in Procedure II.2.27, and in which the attaching maps at each cyclic vertex are homeomorphisms.*

Proof. We associate to G a graph of spaces X. Replacing X with the graph of spaces decomposition of G obtained by applying Procedure II.2.27 gives a graph of spaces that satisfies the first 3 properties described above.

We obtain the desired finite-index subgroup of G by using subgroup separability. Note that if some edge inclusion into a cyclic vertex group is not an isomorphism, or equivalently if an attaching map $\partial_e : X_e \to X_c$ into some cyclic vertex X_c of X is not a homeomorphism, then the image of the generator of $\pi_1(X_e)$ under $\partial_{e*} : \pi_1(X_e) \to \pi_1(X_c)$ is a proper power in $\pi_1(X_c)$.

Let $D = \{\partial_{e_i} : X_{e_i} \to X_{c_i}\}_{1 \le i \le n}$ be the (finite) list of all attaching maps into cyclic vertices of X that are not homeomorphisms. For each i, denote by g_i the generator of $\pi_1(X_{e_i})$. Note that ${\partial_{e_i}}_*(g_i) = g_i^{k_i}$ generates a finite-index subgroup of $\pi_1(X_{c_i})$, and by Lemma I.3.13 there is a finite-index subgroup G_i of G in which $g_i^{k_i}$ is not a power. It follows that in the corresponding finite-sheeted cover X_i of X, there is an elevation of ∂_{e_i} that is a homeomorphism (and the fundamental group of its image in the corresponding cyclic vertex is exactly $\langle g_i^{k_i} \rangle$). In fact, in every finite cover of X_i the elevation that corresponds to a power of g_i must be an isomorphism.

Consider now the subgroup $H = \bigcap_{i=1}^{n} G_i$; it is a finite-index subgroup of G, and it admits a corresponding graph of spaces decomposition Y in which every elevation that corresponds to a power of some g_i must be an isomorphism. However, there can be other elevations to Y (corresponding to conjugates of powers of g_i in G, that are not conjugate to a power of g_i in H) that are not isomorphisms.

To resolve this, we simply pass to a normal cover. In such a cover, if one elevation of some ∂_{e_i} is an isomorphism, then every elevation of ∂_{e_i} must be an isomorphism. The normal cover that we take, is the one corresponding to the normal core of H, that is we take

$$G' = \bigcap_{g \in G} gHg^{-1}.$$

Note that G' is a finite index subgroup as the kernel of a map from G to a finite symmetric group. More precisely, G' is the kernel of the map $G \to \mathrm{Sym}(G/H)$ (that corresponds to the action of G on the left cosets of H by left multiplication). $\square$

Another simple observation is the following (cf. [120, Lemma 5.10]):

Lemma II.2.29. *Let G be an HGFC-group. If for every vertex v of G the induced pair $(G_v, [\underline{w}_v])$ is of surface type, then G has a surface subgroup (that is, a subgroup that is isomorphic to the fundamental group of a closed, connected surface).*

Proof. Note that since each induced pair is one-ended, G must be one-ended itself by Lemma II.2.26. By Lemma II.2.28, up to replacing G with a finite-index subgroup, we may assume that all of the attaching maps of G (or the corresponding graph of spaces X) at cyclic are isomorphisms (or homeomorphisms); we may also assume that each attaching map with target in a non-cyclic vertex space X_v identifies an edge space with a boundary compoent of X_v. Note that all of the "singularities" (that is, points that do not admit a neighborhood homeomorphic to the Euclidean plane) in this graph of spaces are concentrated around cyclic vertices.

Therefore, if G is not a surface group itself, it follows that there is a cyclic vertex space X_c with degree greater than 2 in the underlying graph Ξ of X. We will use this to construct a precover X' of X that satisfies the properties described in Procedure II.2.27, and in which every cyclic vertex space X_c has degree 2 in the underlying graph Ξ of X. The geometric realization of X' will therefore be homeomorphic to a closed, orientable hyperbolic surface, yielding a surface subgroup of G.

We amass a collection of vertex spaces of X: we choose 2 copies of each surface vertex space of X, and $\deg(c)$ copies of each cyclic vertex space X_c of X. Let $\partial_e : X_e \to X_v$ be an attaching map of X. If X_v is a surface vertex, there are two hanging elevations in our collection of vertex spaces; these will be paired with hanging elevations to cyclic vertex spaces. If X_v is a cyclic vertex, there are $\deg(v)$ such hanging elevations. We can pair the hanging elevations, making sure that for every cyclic vertex X_v of X, we pair exactly two elevations with targets in X_v. The resulting precover X' is a closed surface, and each of its connected components yields a surface subgroup of G. □

One would probably expect that a graph of surfaces with boundary, attached along their boundary components, must contain a surface subgroup. Surprisingly, Wilton showed that the same holds for every HGFC-group. The proof is, unfortunately, reaching far beyond the scope of this dissertation. We have mentioned local-global interplay when studying graphs of free groups with cyclic edges, and Wilton's proof includes an additional layer of local-global interplay in which pairs now play the global role (and *Whitehead graphs*, or links, at the different vertex spaces, form the local pieces). It culminates in the use of linear programming techniques which allow one to assemble multiple Whitehead graphs into a pair of surface type, and use these pairs to construct a surface as in Lemma II.2.29. We remark that the formulation of Wilton's result that we state below is relatively complicated since we wish to avoid using terminology that was not defined in this section.

Theorem II.2.30 ([120, Theorem 5.11]). *If $(F = \pi_1(\Gamma), [\underline{w}])$ is a one-ended pair, then there exists a pair $(\Sigma, \partial\Sigma)$ of surface type, that satisfies the following:*

1. *realizing $(\Sigma, \partial\Sigma)$ as a graph accompanied by embedded loops, there is a combinatorial map $f : \Sigma \to \Gamma$ whose restriction to each component of $\partial\Sigma$ is a covering map that covers some w_i,*
2. *f induces a monomorphism on fundamental groups,*
3. *for each w_i and w_j, the number of preimages of a point on w_i under f coincides with the number of preimages of a point on w_j under f,*
4. *the map f factors through an immersion*
$$(\Sigma, \delta\Sigma) \xrightarrow{f'} (F', [\underline{w}']) \overset{i}{\looparrowright} (F, [\underline{w}])$$
(where both maps above satisfy the properties of 1.) such that f' identifies $\partial\Sigma$ with $[\underline{w}']$.

Corollary II.2.31 ([120, Theorem A])**.** *Every one-ended HGFC-group G has a surface subgroup. This can be seen refining the corresponding graph of spaces X as in Lemma II.2.29, replacing each non-cyclic vertex space by a pair of surface type following Theorem II.2.30, and proceeding as in the proof of Lemma II.2.29.*

II.3 JSJ decompositions

We conclude our background review with a brief overview of *JSJ decompositions*. Roughly speaking, a JSJ decomposition of a group is a universal splitting, that serves as a "dictionary" in which one can look up the different splittings of the group. The different properties of JSJ decompositions often depend on the type of splittings that one considers (that is, what are the edge groups that one allows in such splittings). We will mention relevant results when the edge stabilizers are either finite (of bounded cardinality) or two-ended. These results are used in Chapters III and V respectively. We begin with the definition of a JSJ decomposition:

Definition II.3.1. Let G be a group, and let $\mathcal{E}$ be a family of subgroups of G that is closed under conjugation and taking subgroups. A *JSJ decomposition of G over $\mathcal{E}$* is a tree T on which G acts, and which satisfies the following:

1. every edge stabilizer of the action $G \curvearrowright T$ is in $\mathcal{E}$,
2. every edge stabilizer of the action $G \curvearrowright T$ is elliptic with respect to any action of G on a tree T' with edge stabilizers in $\mathcal{E}$,
3. if an action of G on a tree S also satisfies condition 2. above, then there is a G-equivariant map $T \to S$.

Remark II.3.2. We also refer to the graph of groups corresponding to the action $G \curvearrowright T$ in Definition II.3.1 above as a JSJ decomposition of G (over $\mathcal{E}$).

Example II.3.3. If $\mathcal{E}$ consists only of the trivial group, the Grushko decomposition of a finitely generated group G is a JSJ decomposition of G over $\mathcal{E}$.

We remark that a JSJ decomposition of a group G does not always exist. On the other hand, if a JSJ decomposition of G exists, it is not necessarily unique. In some cases, and in particular when G is finitely presented, all of its JSJ decompositions (over the same class of elliptic subgroups $\mathcal{E}$) are related, and can be seen as points in a space called the *JSJ deformation space*. We refer the curious reader to [65, Section 2.3].

A general strategy for constructing a JSJ decomposition of a group G over a family of elliptic subgroups $\mathcal{E}$ is the following: take a splitting $\mathcal{G}$ of G over groups in $\mathcal{E}$; look at the vertex groups of this splitting, and take splittings of these vertex over $\mathcal{E}$ in which the edge groups of $\mathcal{G}$ are elliptic. Repeat this process. It is not clear that this process comes to an end, and a common way to show that it must terminate is using *accessibility* arguments.

Definition II.3.4. A group G is *accessible* with respect to a family of subgroups $\mathcal{E}$ if there is a bound on the number of edges that can appear in a graph of groups decomposition of G with edges in $\mathcal{E}$ (assuming that the decomposition is reduced).

One of the first accessibility results was proved by Linnell [77, Theorem 1], and it states the following: let $C \in \mathbb{N}$, then every finite group G is accessible with respect to the family $\mathcal{E}_C$ of finite subgroups of G of order at most C. This implies:

Corollary II.3.5. *Let G be a finitely generated group and let $C \in \mathbb{N}$. Then there exists a JSJ decomposition of G over the family $\mathcal{E}_C$ of finite subgroups of G of order at most C.*

This result will be used in Chapter III. We conclude by mentioning two additional versions of JSJ decompositions in which the family of elliptic subgroups consists of infinite cyclic groups. Unlike other JSJ decompositions mentioned above, these JSJ decompositions are *canonical* (which means that their construction is invariant under taking an automorphism of the group). They are both obtained by studying cut points in the boundary (or a boundary-like space) of the group in question. The first theorem we mention was proved by Bowditch, and it deals with splittings of one-ended hyperbolic groups relative to two-ended subgroups. We state a simpler version, for torsion-free groups, in which the elliptic subgroups are all isomorphic to $\mathbb{Z}$.

Theorem II.3.6 ([20, Theorem 0.1])**.** *Let G be a one-ended, torsion-free hyperbolic group and let $\mathcal{E}$ be the family of all infinite cyclic subgroups of G. Then G admits a canonical JSJ-decomposition $\mathcal{G}$ over $\mathcal{E}$ satisfying the following properties:*

1. *each vertex group of $\mathcal{G}$ is of one of the following types:*
    - *cyclic, that is $G_v \cong \mathbb{Z}$,*
    - *surface type, for which the induced pair $(G_v, [\underline{w}_v])$ is of surface type,*

- *rigid, that is the vertex group does not admit a cyclic splitting relative to the adjacent edge groups.*

2. *the graph is bipartite, with edges adjoining cyclic vertices to non-cyclic vertices.*

Theorem II.3.7 ([32, Theorem 4.25])**.** *Let F be a free group, let $\underline{w}$ be a multiword in F and let $\mathcal{E}$ be the family of all infinite cyclic subgroups of F. Then there is a canonical* relative *JSJ decomposition $\mathcal{F}$ of F over $\mathcal{E}$ relative to $\underline{w}$ (that is, a maximal universal splitting with respect to all splittings of F in which the elements of $\underline{w}$ are elliptic) satisfying the following properties:*

1. *each vertex group of $\mathcal{F}$ is of one of the following types:*
    - *cyclic, that is $F_v \cong \mathbb{Z}$,*
    - *surface type, for which the induced pair $(F_v, [\underline{w}_v])$ is of surface type,*
    - *rigid, for which the induced pair $(F_v, [\underline{w}_v])$ is rigid.*
2. *the graph is bipartite, and each edge adjoins a cyclic vertex to non-cyclic vertex.*
3. *if F_v is a non-cyclic vertex group, then the adjacent edge groups map onto maximal cyclic subgroups of F_v that are non-conjugate in F_v.*

As mentioned earlier, note that the structure of these JSJ decompositions is similar to the graph of groups decomposition obtained in Procedure II.2.27. This normalization process is what makes these splittings canonical.

Remark II.3.8. With both of these theorems, it follows that if G splits over $\langle g \rangle$ (or F splits over $\langle g \rangle$ relative to $\underline{w}$) then g is conjugate into a cyclic or a surface type vertex of the JSJ decomposition. In addition, if G is a one-ended, torsion-free HGFC-group and $v \in \mathrm{V}(\mathcal{G})$ is a rigid vertex in its canonical JSJ decomposition obtained from Theorem II.3.6, then G_v is a free group and the induced pair $(G_v, [\underline{w}_v])$ is rigid.

III

Formal solutions and the first-order theory of acylindrically hyperbolic groups

MERZLYAKOV'S theorem about the first-order theory of non-abelian free groups is often considered as the first milestone in the study of the logic of free groups. We generalise Merzlyakov's theorem to all acylindrically hyperbolic groups, a vast class of groups that is yet to be studied from a model-theoretic point of view. As a corollary, we deduce that if G is an acylindrically hyperbolic group and $E(G)$ denotes the unique maximal finite normal subgroup of G, then G and the HNN extension $G\dot{\ast}_{E(G)}$, which is simply the free product $G * \mathbb{Z}$ when $E(G)$ is trivial, have the same $\forall\exists$-theory. As a consequence, we confirm the following conjecture, formulated by Casals-Ruiz, Garreta and de la Nuez González: acylindrically hyperbolic groups have trivial positive theory. In particular, one recovers a result proved by Bestvina, Bromberg and Fujiwara, stating that, with only the obvious exceptions, verbal subgroups of acylindrically hyperbolic groups have infinite width.

We begin with a quick reminder of some preliminary notions appearing in Chapters I and II.

Definition III.0.1. A group G is said to be *acylindrically hyperbolic* if it admits a non-elementary and acylindrical action on a hyperbolic space X: for every $\varepsilon \geq 0$ there exist $N > 0$ and $R > 0$ such that for every $x, y \in X$ satisfying $d(x,y) \geq R$,

$$|\{g \in G \mid d(x,gx) \leq \varepsilon \text{ and } d(y,gy) \leq \varepsilon\}| \leq N.$$

Remark III.0.2. As mentioned in Section I.2, we can, and therefore will, always assume that X is a graph.

Recall that a limit group L over a group G is obtained from a sequence of homomorphisms $\mathrm{Hom}(H,G)^{\mathbb{N}}$ (where H is a finitely generated group) by taking the quotient of H by the stable kernel associated to the sequence,

$$\underleftarrow{\ker}_{\omega}((\varphi_n)_{n\in\mathbb{N}}) = \{g \in H \mid g \in \ker(\varphi_n)\ \omega\text{-almost surely}\}.$$

Fixing a finite generating set S of H, we associate a *scaling factor* to every morphism in the sequence $(\varphi_n)_{n\in\mathbb{N}}$, defined by

$$\|\varphi_n\| = \min_{y\in X} \max_{s\in S} d(y, \varphi_n(s)y).$$

Note that X is a simplicial graph, so the scaling factor above is well-defined.

If the limit group is *divergent*, that is if $\lim_{\omega}(\|\varphi_n\|) = \infty$, then L comes armed with a limiting action on a real tree T. We refer to Subsection II.1.3 for further details about the construction of T. The vast majority of this chapter revolves around studying such actions of limit groups on their limiting $\mathbb{R}$-trees.

III.1 Decompositions of limit group and approximations

We focus our attention on divergent limit groups over an acylindrically hyperbolic group G. A group acting on a simplicial tree admits a graph of groups decomposition, but the same is not always true for a group acting on a real tree. In the first half of this section, we gather a few results that imply that divergent limit groups over G do split (for a related discussion, see Subsection II.1.5). We also give a description of this action, that is commonly known as the output of the *Rips Machine*.

In the second half of this section, we introduce *approximations* of limit groups over G. Since G is not necessarily equationally Noetherian, the defining sequence $\mathrm{Hom}(H,G)^{\mathbb{N}}$ of a limit group L over G does not necessarily factor through L (ω-almost surely); see Subsection II.1.4 for additional details. We therefore seek to approximate L by finitely presented groups (through which $\mathrm{Hom}(H,G)^{\mathbb{N}}$ will factor ω-almost surely), that admit splittings which mimic the different splittings of L. This is an important ingredient in the next section, Section III.1.4, where we prove a version of Sela's *shortening argument* for acylindrically hyperbolic groups.

Before getting down to the matter at hand, we give a brief overview of *relative group presentations* as these will appear frequently throughout this chapter.

III.1.1 Relative group presentations

A *relative presentation* of a group G with respect to a subgroup $H \subset G$ and a subset $X \subset G$ is a presentation of the form

$$G = \langle H, X \mid \mathcal{R} \rangle,$$

which is obtained from H by adding the set of generators X and the set of relations $\mathcal{R}$. Thus $G = (H * F(X))/\langle\langle \mathcal{R} \rangle\rangle$, where $F(X)$ is the free group on X and $\langle\langle \mathcal{R} \rangle\rangle$ is the normal closure of $\mathcal{R}$ in $H * F(X)$. In the case where the set $\mathcal{R}$ is finite, one says that G is *finitely presented* relative to H.

III.1.2 Graphs of actions and the Rips machine

Under certain conditions, a group acting on a real tree splits as a *graph of actions*. This splitting endows the group with an action on a simplicial tree which is generally easier to understand than an action on a real tree. Groves and Hull proved in [57] that divergent limit groups over acylindrically hyperbolic groups (along with their canonical actions on real trees) satisfy the desired conditions which are required to invoke Guirardel's version of the Rips machine (see [62]). We present the relevant definitions and results from both works. For a more penetrating discussion about the Rips machine, we refer the curious reader to [62].

Definition III.1.1. [62, Definition 1.2] A *graph of actions* $\mathcal{G}$ consists of:

1. an underlying graph of groups $\mathbb{A} = \big(A, (A_v)_{v \in \mathrm{V}(A)}, (A_e)_{e \in \mathrm{E}(A)}, (i_e)_{e \in \mathrm{E}(A)}\big)$,
2. a collection of real trees $(T_v, d_v)_{v \in \mathrm{V}(A)}$ such that A_v acts on T_v,
3. a collection of points $(p_e \in T_{t(e)})_{e \in \mathrm{E}(A)}$ such that every p_e is fixed by $i_e(A_e)$, called *attaching points*,
4. a function $\ell : \mathrm{E}(A) \to \mathbb{R}_{\geq 0}$ assigning *lengths* to the edges of A, and such that $\ell(e) = \ell(\bar{e})$ for every $e \in \mathrm{E}(A)$.

We usually present the information above as a tuple and write

$$\mathcal{G} = \mathcal{G}(\mathbb{A}) = \big(\mathbb{A}, (T_v)_{v \in \mathrm{V}(A)}, (p_e)_{e \in \mathrm{E}(A)}, \ell\big).$$

A graph of actions $\mathcal{G}(\mathbb{A})$ enables one to canonically construct a real tree $T_{\mathcal{G}}$ on which $G = \pi_1(\mathbb{A})$ acts: replace each vertex $\tilde{v}$ of the Bass-Serre tree T_A corresponding to $\mathbb{A}$ by a copy of T_v (where v is the image of $\tilde{v}$ under the quotient map $q : T_A \to$

$G\backslash T_A = A$), and replace any edge $\tilde{e}$ of T_A by a segment of length $\ell(e)$ (where $e = q(\tilde{e})$). We also ask that if $t(\tilde{e}) = \tilde{v}$ in T_A then $t(\tilde{e}) = \tilde{p}_e$ in $T_{\mathcal{G}}$, that is we attach the tree T_v via the attaching point p_e. The action of $\pi_1(\mathbb{A})$ on T_A extends naturally to an action $\pi_1(\mathbb{A}) \curvearrowright T_{\mathcal{G}}$. We next define notions of stability concerning group actions on real trees which will allow us to describe the output of the Rips machine.

Definition III.1.2. Suppose that L is a group acting on a real tree T.

1. A subtree $T' \subset T$ is called *stable* if for every non-degenerate subtree $T'' \subset T'$, the pointwise estabilizers of T' and T'' coincide, that is $\mathrm{Stab}_L(T') = \mathrm{Stab}_L(T'')$. Otherwise, T' is called *unstable*. An action on a real tree is *stable* if any non-degenerate arc contains a non-degenerate stable subarc.

2. The action $L \curvearrowright T$ is said to satisfy the *ascending chain condition* if for any sequence of nested arcs $I_1 \supset I_2 \supset \cdots$ in T whose lengths tend to 0, the corresponding sequence of stabilizers $\mathrm{Stab}_L(I_1) \subset \mathrm{Stab}_L(I_2) \subset \cdots$ eventually stabilizes.

We are now ready to state a relative version of the Rips machine; we remark that we use the term *relative*, because the decomposition is taken relative to a subgroup U. Further information about relative group presentations appears in Subsection III.1.1.

Theorem III.1.3. *[62, Theorem 5.1] Let L be a group acting minimally and non-trivially on a real tree T by isometries. Let U be a subgroup of L such that L is finitely generated over U and such that U fixes a point in a real tree T on which L acts. Assume in addition that the action of L on T satisfies the ascending chain condition, and that for any unstable arc $I \subset T$,*

1. $\mathrm{Stab}_L(I)$ *is finitely generated, and*

2. $\mathrm{Stab}_L(I)$ *is not a proper subgroup of any conjugate of itself.*

Then one of the following holds.

1. L splits over the stabilizer of an unstable arc and U is contained in one of the factors.

2. L splits over the stabilizer N of an infinite tripod and U is contained in one of the factors, and the normalizer of N contains a non-abelian free group generated by two hyperbolic elements whose axes do not intersect.

3. The action $L \curvearrowright T$ decomposes as a graph of actions $\mathbb{R}_L$ where each vertex action is either *(a)* simplicial*: a simplicial action on a simplicial tree,*

(b) of Seifert-type*: the action of L_v has kernel N_v and the faithful action of L_v/N_v is dual to an arational measured foliation on a compact 2-orbifold with boundary,*

(c) or axial*: T_v is a line, and the image of L_v in* Isom(T_v) *is a finitely generated group acting with dense orbits on T_v.*

We are (obviously) interested in decompositions of divergent limit groups over acylindrically hyperbolic groups; the following lemma that appears in [57], also known as the *stability lemma*, implies that such limit groups do indeed satisfy the stability conditions required for applying Theorem III.1.3. Recall that L is a divergent limit group with defining sequence of homomorphisms $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H,G)^{\mathbb{N}}$. In the following lemma, δ denotes the hyperbolicity constant of a hyperbolic space on which G acts acylindrically and non-elementarily, and N and R denote the acylindricity constants appearing in Definition I.2.2.

Lemma III.1.4 ([57, Lemma 4.7])**.** *There is a constant C depending only on δ, N and R such that the action of L on the corresponding real tree T satisfies the following conditions.*

1. *If $A \subset L$ stabilizes a non-trivial arc of T, or if A preserves a line in T and fixes its ends, then A is an extension of an abelian group by a finite group of order $\leq C$.*

2. *The stabilizer of a tripod in T is of order $\leq C$.*

3. *The stabilizer of an unstable arc $I \subset T$ is of order $\leq C$.*

4. *If $K \subset L$ is* locally stably elliptic*, that is for every finitely generated subgroup $K' \subset K$, the action of $\varphi_n(\tilde{K}')$ (where $\tilde{K}'$ is a lift of K' to H) on X_n is elliptic ω-almost surely, then the order of K is $\leq C$.*

Corollary III.1.5. *The fact that stabilizers of unstable arcs are finite implies that the action of L on T satisfies the ascending chain condition and the rest of the conditions required for Theorem III.1.3. Hence if L does not split non-trivially over a finite subgroup of order $\leq C$, it must split as a graph of actions as in Theorem III.1.3.*

III.1.3 Approximations of limit groups

We keep the notation from the previous subsection, that is G is a group that admits an acylindrical and non-elementary action on a δ-hyperbolic simplicial graph (X, d) and L is a limit group with defining sequence $(\varphi_n)_{n\in\mathbb{N}} \in \mathrm{Hom}(H, G)^{\mathbb{N}}$, where H is a finitely generated group; we also assume that the sequence $(\varphi_n)_{n\in\mathbb{N}}$ is divergent.

Standing Assumption III.1.6. *In what follows, we assume that H is finitely presented over an infinite finitely generated (but not necessarily finitely presented) subgroup $U \subset H$. We denote by S a finite generating set of H. In addition, we suppose that U acts elliptically on the limiting tree T, and that the restriction of the limit map φ_∞ to U is injective, which allows us to identify U with its image under φ_∞.*

In this subsection, we aim to prepare the grounds for proving a version of the shortening argument for acylindrically hyperbolic groups (slightly different from the version proved in [57] by Groves and Hull). First, let us recall that the group G is not equationally Noetherian in general. Therefore, the sequence $(\varphi_n : H \to G)_{n\in\mathbb{N}}$ defining the G-limit group L does not factor through the quotient map φ_∞ *a priori*. This is a major obstacle to generalising the standard proof of the shortening argument, since we cannot rely on the splitting of L as a graph of actions outputted by the Rips machine in order to shorten the morphisms of the sequence. Before we explain our approach for overcoming this difficulty (which is, in some sense, similar to the approach taken in [57, Lemma 6.3], coming from [106, Theorem 3.2]), we begin by defining *approximations* of limit groups.

Definition III.1.7. Given a finite set of relations $\mathcal{R} \subset \ker(\varphi_\infty)$, we define the *$\mathcal{R}$-approximation* A of L as $A = H/\langle\langle \mathcal{R} \rangle\rangle$. In general, we call a group A obtained in this manner an *approximation* of L.

Remark III.1.8. Since the set $\mathcal{R}$ is a subset of $\ker(\varphi_\infty)$, the action of H on the limiting tree T gives rise to an action of A on T. Both quotient maps $H \twoheadrightarrow A$ and $A \twoheadrightarrow L$ are equivariant with respect to the corresponding actions on T.

Our motivations for introducing approximations of limit groups are the following.

1. Since H is finitely presented over U (see Standing Assumption III.1.6), every $\mathcal{R}$-approximation A of L is finitely presented over U. Therefore, the homomorphisms in the sequence $(\varphi_n : H \to G)_{n\in\mathbb{N}}$ factor ω-almost surely through the quotient map $q : H \twoheadrightarrow A$ (since each of the defining relations in the finite presentation of A over U is mapped to 1 under φ_n ω-almost surely). We will denote

the maps arising from this factorization by θ_n, that is $\varphi_n = \theta_n \circ q$, and refer to the corresponding limit map by either $\theta_\infty : A \twoheadrightarrow L$ or simply $\pi : A \twoheadrightarrow L$. This factorization will be crucial in our proof of the general version of the shortening argument.

2. Suppose that L admits a nice splitting as a graph of groups (below we give a list of such splittings). We will see that, provided that the finite set of relations $\mathcal{R} \subset \ker(\varphi_\infty)$ is carefully chosen, the approximation A admits a splitting that mimics the splitting of L, in a precise sense.

In the proof of the shortening argument, as well as in the proof of Merzlyakov's theorem, we will consider different splittings of L:

- if L splits non-trivially relative to U over a finite subgroup of order $\leq C$ (the constant appearing in Lemma III.1.4), a reduced JSJ splitting of L relative to U over finite subgroups of order $\leq C$ (see Section II.3 and Definition II.3.1), denoted by $\mathbb{J}_L$ (see Proposition III.1.11 and Corollary III.1.13);
- if L does not split non-trivially relative to U over a finite subgroup of order $\leq C$, a splitting of L as a graph of actions outputted by the Rips machine as in Theorem III.1.3, denoted by $\mathbb{R}_L$ (see Proposition III.1.11 and Corollary III.1.14). This is our main motivation for approximating limit groups;
- more generally, a splitting $\mathbb{RJ}_L$ of L obtained from $\mathbb{J}_L$ by replacing the unique vertex u fixed by U with $\mathbb{R}_{L_u}$ (see Proposition III.1.11 and Corollary III.1.15).

Before we construct approximations of L equipped with splittings that mimic one of the aforementioned splittings, we define with more details the sense in which an approximation of L mimics a certain splitting.

Definition III.1.9. Let A be an approximation of L as in Definition III.1.7. Let $\pi : A \twoheadrightarrow L$ be the natural epimorphism (obtained by quotienting out by the image of $\ker(\varphi_\infty)$ in A). Suppose that L splits as a finite graph of groups $\mathbb{S}_L$. We say that A is an $\mathbb{S}_L$-*approximation* of L if the following four conditions hold:

1. A splits as a graph of groups $\mathbb{S}_A$ with the same underlying graph as $\mathbb{S}_L$, and in which all the edge groups are finitely presented and all the vertex groups are finitely presented (relative to U);
2. π induces an isomorphism of graphs, denoted by f, between the underlying graph of $\mathbb{S}_A$ and the underlying graph of $\mathbb{S}_L$;

3. if $i : A_e \hookrightarrow A_v$ denotes the inclusion of an edge group of $\mathbb{S}_A$ into an adjacent vertex group, and $j : L_{f(e)} \hookrightarrow L_{f(v)}$ denotes the corresponding inclusion in the graph of groups $\mathbb{S}_L$, then the following diagram commutes:

$$\begin{tikzcd}
A_e \arrow[r, "i"] \arrow[d, "\pi"] & A_v \arrow[d, "\pi"] \\
L_{f(e)} \arrow[r, "j"] & L_{f(v)}
\end{tikzcd}$$

4. π maps every edge group A_e of $\mathbb{S}_A$ into the corresponding edge group $L_{f(e)}$ of $\mathbb{S}_L$.

For readability, we omit the isomorphism f and denote the vertex $f(v)$ and the edge $f(e)$ by v and e respectively.

Remark III.1.10. The second and third conditions in the definition above can be phrased, equivalently, as follows: there exists a π-equivariant isomorphism of graphs between the Bass-Serre trees of the splittings $\mathbb{S}_A$ and $\mathbb{S}_L$.

A similar method of approximating limit groups appears in [106, Theorem 3.2] (see also [56], [97, Lemma 6.1] and [57, Lemma 6.3]). Note that in these papers, in the process of approximating a limit group, one constructs countably many approximations; each of them approximates the limit group L to a greater extent than its predecessors. In addition, one does not obtain a concrete approximation of L catered to satisfy a desired property. Below we give an alternative construction, which approximates (only) a specific properties of the limit group L and use it to approximate the properties required for the proofs of the shortening argument and Merzlyakov's theorem.

Proposition III.1.11. *Suppose that L splits as a graph of groups $\mathbb{S}_L$ in which all the edge groups are virtually abelian. Then there exists an $\mathbb{S}_L$-approximation A of L, which in addition satisfies the following two properties.*

1. *If L_e is a finitely generated edge group of $\mathbb{S}_L$, then the quotient map $\pi : A \twoheadrightarrow L$ is an isomorphism.*

2. *Let L_v be a vertex group of $\mathbb{S}_L$; if all the edge groups of $\mathbb{S}_L$ adjacent to L_v are finitely generated, then the quotient map $\pi : A \twoheadrightarrow L$ maps A_v onto L_v. Moreover, if L_v is finitely presented (relative to U), then the map $\pi|_{A_v} : A_v \to L_v$ is an isomorphism.*

In addition, for any finite set of relations $\mathcal{F} \subset \ker(\varphi_\infty)$, we can choose the approximation A such that the image of $\mathcal{F}$ in A is trivial.

Proof. We choose to construct the vertex and edge groups of $\mathbb{S}_A$ before constructing the group A itself. By doing so, we hope to give the reader a better understanding of the structure of the approximation A. Since the proof of this proposition is quite intricate, we divide it into four steps.

Step 1. We begin by fixing explicit presentations of the edge and vertex groups of the graph of groups $\mathbb{S}_L$ that will be used throughout the proof. Since L is finitely generated, every vertex group L_v of $\mathbb{S}_L$ is finitely generated relative to its adjacent edge groups. Fix a presentation $\langle S \mid P_U \cup P\rangle$ of H, where P_U consists of relations involving only elements from U and P is finite. We also fix a finite generating set X_U of U.

For each edge $e \in \mathrm{E}(\mathbb{S}_L)$, fix a presentation $\langle X_e \mid R_e\rangle$ of the edge group L_e. If L_e is finitely generated (and hence finitely presented, as a virtually abelian group), we choose this presentation to be finite. Note that otherwise, X_e and R_e are both infinite. Then, for each vertex $v \in \mathrm{V}(\mathbb{S}_L)$, fix a presentation of the vertex group L_v of the form

$$\langle X_v = Y_v \cup X^v_{e_1} \cup \cdots \cup X^v_{e_{n_v}} \mid R_v = Q_v \cup R^v_{e_1} \cup \cdots \cup R^v_{e_{n_v}}\rangle,$$

where

- $e_1, \ldots, e_{n_v}$ are the edges adjacent to v in the underlying graph of $\mathbb{S}_L$;
- each $\langle X^v_{e_i} \mid R^v_{e_i}\rangle$ is a copy of the corresponding edge group within L_v.
- Recall that L is finitely generated; therefore, L_v is finitely generated relative to its adjacent edge groups. Hence, the set Y_v can be chosen to be finite (and if L_v contains U we choose Y_v to be the union of X_U and a finite set).
- Q_v is a (possibly infinite) set of relations.

In addition, we fix a presentation of L as the fundamental group of $\mathbb{S}_L$, that is

$$L = \left\langle \bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} X_v \cup \{t_e,\ e \in \mathrm{E}\} \;\middle|\; \bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} R_v \cup R \right\rangle$$

where

- $\mathrm{E} = \mathrm{E}(\mathbb{S}_L) \backslash T_{\mathbb{S}_L}$ for a spanning subtree $T_{\mathbb{S}_L}$ of the underlying graph of $\mathbb{S}_L$, and

- R is a (possibly infinite) set of relations which includes relations of two types:
  - relations which identify the set of generators X_e of the edge group L_e with their images in the adjacent vertex groups whenever $e \in T_{\mathbb{S}_L}$.
  - relations of the form $t_e^{-1} i_e(x_e) t_e = i_{\overline{e}}(x_e)$ where $e \in \mathrm{E}$, $x_e \in X_e$ and the maps i_e and $i_{\overline{e}}$ are the inclusion maps of the edge group L_e into the adjacent vertex groups.

Step 2. Having fixed the relevant presentations, we seek to pick a finite set $X_L \subset \bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} X_v \cup \{t_e,\ e \in \mathrm{E}\}$ that generates L (relative to U). The elements in X_L will be used to define the vertex and edge groups of $\mathbb{S}_A$. We choose X_L to be extensive enough so that each of the relations in the finite set $\varphi_\infty(P \cup \mathcal{F})$ can be written as a product of conjugates of relations from the presentation of L above as the fundamental group of $\mathbb{S}_L$, involving only elements from X_L.

For every $s \in S$, write $\varphi_\infty(s)$ as a product of generators appearing in the presentation of L above. Let X_S be the finite subset of the generating set $\bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} X_v \cup \{t_e,\ e \in \mathrm{E}\}$ of L composed of the generators appearing in these products.

Similarly, each relation r in the finite set $\varphi_\infty(P \cup \mathcal{F})$ can be written as a product of conjugates of relations appearing in the presentation of L above. Let R_L be the finite set of relations that participate in such products, and let X_R be the finite subset of $\bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} X_v \cup \{t_e,\ e \in \mathrm{E}\}$ which consists of all the generators of L participating in the products of conjugates of relations from R_L described above.

Lastly, let $X_L = X_S \cup X_R$.

Step 3. We finally construct the edge and vertex groups of the splitting $\mathbb{S}_A$. For every $e \in \mathrm{E}(\mathbb{S}_L)$ we define A_e as follows: if L_e is finitely generated (and hence, finitely presented), let $A_e = L_e$. We fix an alternative notation for the finite presentation of A_e: $\langle X'_e \mid R'_e \rangle$. If L_e is not finitely presented, let A_e be the subgroup of L_e generated by $L_e \cap X_L$; note that A_e is finitely presented (as a finitely generated virtually abelian group) and fix a finite presentation $\langle X'_e \mid R'_e \rangle$ of A_e. Up to modifying the original presentation of L_e, we may assume that X'_e is a subset of X_e.

For every $v \in \mathrm{V}(\mathbb{S}_L)$ we define A_v as follows: if L_v is finitely presented over U, we let $A_v = L_v$ (and fix an alternative notation for the presentation of A_v: $\langle X'_v \mid R'_v \rangle$); otherwise, we set A_v to be the group admitting the following presentation:

$$\left\langle X'_v = Y_v \cup (X'_{e_1})^v \cup \cdots \cup (X'_{e_{n_v}})^v \mid \ R'_v = Q'_v \cup \cdots \cup (R'_{e_{n_v}})^v \right\rangle,$$

where

- $e_1, \dots, e_{n_v}$ are the edges adjacent to v in the underlying graph of $\mathbb{S}_L$;
- Y_v is as in the presentation of L_v;
- each $\langle (X'_{e_i})^v \mid (R'_{e_i})^v \rangle$ is a copy of A_{e_i} within A_v;
- if L_v does not contain U, one has $Q'_v = Q_v \cap R_L$ if Q_v is infinite, and $Q'_v = Q_v$ otherwise. If L_v contains U, we pick Q'_v in the same manner, but include in Q'_v the (possibly infinite) subset of Q_v which consists of relations involving only elements from U.

Recall that R is the set of relations from the presentation of L above. Let R' be the finite set of relations which involve only the generators X'_e of A_e for each $e \in \mathrm{E}(\mathbb{S}_L)$. Let $\mathcal{R} = \bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} R'_v \cup R'$, where all of the relations in the union are written with the letters of the generating set S of H. Note that $\mathcal{R}$ is finite. Now, let A be the $\mathcal{R}$-approximation of L, that is define $A = H / \langle\langle \mathcal{R} \rangle\rangle$.

Step 4. We now show that A satisfies the desired properties. First, note that A admits the presentation $\langle S \mid P_U \cup P \cup \mathcal{R} \rangle$ in the generators of H. By expressing this presentation in terms of the generators of X_L, we obtain the following presentation of A:

$$A = \left\langle \bigcup_{v \in \mathrm{V}(\mathbb{S}_L)} X'_v \cup \{t_e,\ e \in \mathrm{E}\} \;\middle|\; R_L \cup \mathcal{R} \right\rangle.$$

But since R_L is contained in $\mathcal{R}$ by definition of the R'_v, one can omit R_L in the previous presentation of A. Hence, A is simply the fundamental group of the graph of groups $\mathbb{S}_A$ obtained from $\mathbb{S}_L$ by replacing each vertex group L_v with the group A_v, and each edge group L_e with the group A_e. In addition, all of the relations in $\mathcal{F}$ hold in A. This shows that condition (1) of Definition III.1.9 holds.

Next, note that the map π', defined by mapping each generator in the presentation of A above to the corresponding generator in the presentation of L as the fundamental group of $\mathbb{S}_L$, coincides with the natural epimorphism $\pi : A \twoheadrightarrow L$ obtained by quotienting out the image of $\ker(\varphi_\infty)$ in A. Indeed, denote by q the quotient map $H \twoheadrightarrow A$ and observe that for every $s \in S$ one has $\pi' \circ q(s) = \varphi_\infty(s)$; this implies that $\pi' = \pi$. Last, properties (2), (3) and (4) appearing in Definition III.1.9 are clearly satisfied.

To finish, let us check that properties *1.* and *2.* of Proposition III.1.11 hold: for property *1.*, recall that whenever an edge group L_e of $\mathbb{S}_L$ is finitely generated, we defined A_e to be L_e. For property *2.*, recall that if all the edge groups adjacent to a vertex group L_v of $\mathbb{S}_L$ are finitely generated, then the generators X'_v of A_v correspond to the generators X_v of L_v. This implies that the restriction of the map $\pi : A \twoheadrightarrow L$ to A_v is a surjection. If in addition L_v is finitely presented (over U), then A_v and L_v admit the same presentation and π maps A_v isomorphically to L_v. □

Remark III.1.12. Suppose that L admits a splitting $\mathbb{S}_L$, and let $\{h_1, \dots, h_k\}$ be a finite set of elements of L. Write each element h_i as a product $s_{i,1} \cdots s_{i,m_i}$ of generators appearing in the presentation of L as the fundamental group of $\mathbb{S}_L$. By choosing the finite set of relations $\mathcal{F}$ in Proposition III.1.11 above wisely, we can make sure that L has an $\mathbb{S}_L$-approximation A such that each h_i has a preimage a_i in A that admits the same decomposition as h_i as a product of generators. More precisely, let $\tilde{h}_i, \tilde{s}_{i,1}, \dots, \tilde{s}_{i,m_i}$ be lifts of $h_i, s_{i,1}, \dots, s_{i,m_i}$ to H, for $1 \leq i \leq k$. Let

$$\mathcal{F} = \{\tilde{h}_i^{-1} \tilde{s}_{i,1} \cdots \tilde{s}_{i,m_i},\ 1 \leq i \leq k\} \subset \ker(\varphi_\infty),$$

and let A be an $\mathbb{S}_L$-approximation of L in which the relations in $\mathcal{F}$ hold. Then by Proposition III.1.11 above, all of the generators $s_{i,j}$ appear in the presentation of A as the fundamental group of $\mathbb{S}_A$, and the image of h_i in A can be simply written as $s_{i,1} \cdots s_{i,m_i}$. We will use this method in our proof of the general version of the shortening argument.

We now deduce from Proposition III.1.11 a series of three corollaries.

Corollary III.1.13. *Suppose that L admits a splitting $\mathbb{S}_L$ in which all the edge groups are finite (for instance, $\mathbb{S}_L$ can be a reduced JSJ splitting of L relative to U over finite subgroups of order $\leq C$, denoted by $\mathbb{J}_L$). In this case, all the vertex groups of $\mathbb{S}_L$ are finitely generated. Then there exists an $\mathbb{S}_L$-approximation A of L, whose splitting is denoted by $\mathbb{S}_A$, such that $\mathbb{S}_A$ and $\mathbb{S}_L$ share the same edge groups, and the vertex groups of $\mathbb{S}_A$ surject onto those of $\mathbb{S}_L$.*

Proof. Since edge groups of $\mathbb{S}_L$ are finite, they are virtually abelian and finitely generated. Thus the existence of A is an immediate consequence of Proposition III.1.11. □

As mentioned earlier, the main motivation for defining approximations of limit groups is to approximate in an accurate manner splittings of L as a graph of actions outputted by the Rips machine III.1.3. The following lemma proves that the approximations given by Proposition III.1.11 capture many of the properties of such splittings.

Corollary III.1.14. *Suppose that L does not split non-trivially over a finite subgroup of order $\leq C$, and let $\mathbb{R}_L$ be a splitting of L as a graph of actions outputted by the Rips machine. Let A be an $\mathbb{R}_L$-approximation of L given by Proposition III.1.11 and let $\mathbb{R}_A$ denote its splitting. Then the following hold:*

1. *every edge group of $\mathbb{R}_A$ is finitely presented and virtually abelian, and the quotient map $\pi : A \twoheadrightarrow L$ maps every edge group of $\mathbb{R}_A$ into the corresponding edge group of $\mathbb{R}_L$;*
2. *if L_v is a simplicial vertex group of $\mathbb{R}_L$ and A_v is the corresponding vertex group of $\mathbb{R}_A$, then π maps A_v onto a finitely generated (relative to U) subgroup of L_v;*
3. *if L_v is a Seifert-type vertex group of $\mathbb{R}_L$ and A_v is the corresponding vertex group of $\mathbb{R}_A$, then π maps A_v isomorphically to L_v;*
4. *if L_v is an axial vertex group of $\mathbb{R}_L$ and A_v is the corresponding group of $\mathbb{R}_A$, then A_v is finitely presented (relative to U) and virtually abelian, and π maps A_v into L_v.*

Proof. Recall that the edge groups of $\mathbb{R}_L$ are virtually abelian, and that each edge group A_e of $\mathbb{R}_A$ is finitely presented (condition (1) in Definition III.1.9). We also have that π maps A_e into the corresponding edge group L_e (condition (4) in Definition III.1.9). Hence, A_e is virtually abelian and the first assertion above holds.

Next, note that the second assertion is an immediate consequence of the construction of A in Proposition III.1.11.

For *3.*, let L_v be a Seifert-type vertex group of $\mathbb{R}_L$. Note that L_v is finitely presented (relative to U). We will prove that π maps the corresponding vertex group A_v of $\mathbb{R}_A$ isomorphically to L_v. By the second assertion of Proposition III.1.11, it is enough to show that L_v does not contain infinitely generated abelian subgroups, and thus that all the edge groups adjacent to L_v are finitely generated. This follows from the following easy observation: since L_v is a Seifert-type vertex group, it is hyperbolic, and therefore all of its abelian subgroups are virtually cyclic.

Lastly, for *4.*, note that any subgroup of L_v which is finitely generated (relative to U) is finitely presented (relative to U) since L_v is virtually abelian. This implies that, as part of the construction of A_v in the proof of Proposition III.1.11, A_v is in fact a subgroup of L_v which is finitely presented (relative to U) and π maps A_v injectively to L_v. □

The last corollary is a combination of Corollaries III.1.13 and III.1.14, and it can be proved in a similar way.

Corollary III.1.15. *Let $\mathbb{J}_L$ be a reduced JSJ splitting of L over finite groups of order $\le C$ relative to U, and let $\mathbb{RJ}_L$ be a splitting of L obtained from $\mathbb{J}_L$ by replacing the unique vertex u fixed by U with $\mathbb{R}_{L_u}$. Note that L_u can be viewed as a limit group whose defining sequence of homomorphisms is $(\varphi_n|_{H_u})_{n\in\mathbb{N}}$, where H_u denotes a finitely generated subgroup of H that contains U and such that $\varphi_\infty(H_u) = L_u$. Then any $\mathbb{RJ}_L$-approximation A of L outputted by Proposition III.1.11 admits a splitting $\mathbb{RJ}_A$ satisfying the following conditions.*

1. *Denote by $\mathbb{R}_{A_u}$ the subgraph of $\mathbb{RJ}_A$ which corresponds to the subgraph $\mathbb{R}_{L_u}$ of $\mathbb{RJ}_L$. Denote by A_u the fundamental group of $\mathbb{R}_{A_u}$ and note that A_u is a lift of L_u to A. Then A_u is an $\mathbb{R}_{L_u}$-approximation of L_u. Furthermore, the splitting $\mathbb{R}_{A_u}$ of A_u enjoys the properties described in Corollary III.1.14.*
2. *Let $\mathbb{J}_A$ be the splitting of A obtained by collapsing the subgraph $\mathbb{R}_{A_u}$ of $\mathbb{RJ}_A$ to a point. Then A, equipped with the splitting $\mathbb{J}_A$, is a $\mathbb{J}_L$-approximation of L.*

Remark III.1.16. Note that the splitting $\mathbb{RJ}_L$ is not unique in general since the finite edge groups adjacent to the vertex u in $\mathbb{J}_L$ may fix several vertex groups in the splitting $\mathbb{R}_{L_u}$ of L_u.

III.1.4 The shortening argument

The *shortening argument* encompasses a wide array of results, all of which share a similar nature: shortening homomorphisms. The classical result asserts that given a sequence of homomorphisms from a finitely generated group to another group from a certain class, either one can shorten the homomorphisms (in some sense), or the stable kernel of the sequence is non-trivial. For the class of acylindrically hyperbolic groups, a version of the shortening argument is proven in [57, Theorem 5.29]. In this section, we provide two additional versions of this result which will help us deal with inequalities in the proof of Merzlyakov's theorem. We first define the notion of a *short* homomorphism.

Definition III.1.17. Recall that the *scaling factor* (or *length*) of a homomorphism $\varphi : H \to G$ is defined by

$$\|\varphi\| = \inf_{y\in X} \max_{s\in S} d(y, \varphi(s)y),$$

where S is a finite generating set of H. We call φ *short relative to U* if for every homomorphism $\phi : H \to G$ whose restriction to U coincides with $\varphi|_U$ up to conjugation (that is, there exists $g \in G$ such that $\phi(h) = g\varphi(h)g^{-1}$ for every $h \in U$), one has:

$$\|\varphi\| \leq \|\phi\|.$$

We would like to point out once again the main difficulty which stands in our way: unlike hyperbolic groups, acylindrically hyperbolic groups are not equationally Noetherian in general. Therefore, the sequence of homomorphisms $(\varphi_n)_{n\in\mathbb{N}}$ does not necessarily factor via the limit group L (ω-almost surely), and one cannot use automorphisms of L in order to shorten the homomorphisms in the sequence $(\varphi_n)_{n\in\mathbb{N}}$. To combat this, we use approximations of L which were defined in the previous section (see Definitions III.1.7 and III.1.9). Since the homomorphisms in the sequence $(\varphi_n)_{n\in\mathbb{N}}$ do factor via an approximation A of L ω-almost surely, we can use automorphisms of A in order to shorten the sequence $(\varphi_n)_{n\in\mathbb{N}}$. The automorphisms which we use are lifts of a certain type of *modular* automorphisms of L (see Definitions III.1.21 and III.1.23) to A.

Before stating and proving our two versions of the shortening argument, we begin by collecting a few definitions and results.

Definition III.1.18. [97, Definition 3.13] Let G be a group which splits as a graph of groups $\mathbb{S}$ and let G_v be one of its vertex groups. Suppose that $\alpha_v \in \mathrm{Aut}(G_v)$ satisfies the following property: for every edge group G_{e_i} adjacent to G_v there exists an element $c_{e_i} \in G_v$ such that α_v restricts to conjugation by c_{e_i} on G_{e_i}. Recall that each element of G can be realized as a loop in the graph of groups $\mathbb{S}$. The homomorphism $\alpha : G \to G$ defined by

$$[a_0, e_1, a_1, \ldots, e_k, a_k] \mapsto [b_0, e_1, b_1, \ldots, e_k, b_k]$$

where

$$b_i = \begin{cases} a_i & a_i \notin A_v \\ c_{e_i^{-1}}^{-1}\alpha_v(a_i)c_{e_{i+1}} & a_i \in A_v \end{cases}.$$

is called a *natural extension* of α_v.

Remark III.1.19. Note that the elements c_{e_i} are not unique in general, and hence the morphism α above is not uniquely defined by α_v.

The following short lemma shows that such a natural extension α is an automorphism of G.

Lemma III.1.20. *Let G be a group that splits as a graph of groups $\mathbb{S}$; let G_v be one of its vertex groups. Let $\alpha_v \in \mathrm{Aut}(G_v)$ satisfy the properties appearing in Definition III.1.18 above and let $\alpha : G \to G$ be a natural extension of α_v. Then α is a well-defined automorphism of G whose restriction to G_v is α_v, and whose restriction to every edge group of $\mathbb{S}$ is a conjugation (by some element, depending on the edge).*

Proof. Since G can be realized as a sequence of amalgamated products followed by a sequence of HNN extensions, it is enough to prove the lemma in the case where $\mathbb{S}$ has only one edge.

First case. Suppose that $G = A *_C B$, and assume that α_v is an automorphism of A such that $\alpha_v|_C = \mathrm{ad}(a)$ for some $a \in A$. Define α as in Definition III.1.18, that is: $\alpha|_A = \alpha_v$ and $\alpha|_B = \mathrm{ad}(a)$. This endomorphism is well-defined, and it is clearly surjective since its image contains $\alpha_v(A) = A$ and aBa^{-1}, which generate G. Let us prove that α is injective. Consider a non-trivial element $g = a_1b_1a_2b_2\cdots a_nb_n \in G$ written in normal form. The elements a_i and b_i do not belong to C, except maybe a_1 or b_n. One can write $\alpha(g) = a'_1b_1a'_2b_2\cdots a'_nb_na'_{n+1}$ with $a'_1 = \alpha_v(a_1)a$, $a'_i = a^{-1}\alpha_v(a_i)a$ for $1 < i \leq n$ and $a'_{n+1} = a^{-1}$. Observe that a'_i does not belong to C for $1 < i \leq n$, otherwise $a'_i = a^{-1}\alpha_v(a_i)a = c \in C$, thus $\alpha_v(a_i) = aca^{-1} = \alpha_v(c)$. It follows that $a_i = c$; this is a contradiction. Hence, the previous decomposition of $\alpha(g)$ is in normal form, which proves that $\alpha(g)$ is not trivial.

Second case. Suppose that $G = \langle A, t \mid tct^{-1} = \sigma(c),\ \forall c \in C_1\rangle$, where σ denotes an isomorphism between two subgroups C_1 and $C_2 = \sigma(C_1)$ of A. Suppose that α_v is an automorphism of A such that $\alpha_{v|C_i} = \mathrm{ad}(a_i)$ for some $a_i \in A$, for $1 \leq i \leq 2$. Define α as follows: $\alpha|_A = \alpha_v$ and $\alpha(t) = a_2ta_1^{-1}$. As in the first case, one easily sees that α is well-defined and surjective. The injectivity follows from Britton's lemma, by a similar argument to the one appearing above. □

We next define the modular group of a limit group (see [57, Definition 5.22]).

Definition III.1.21. Suppose that L admits a splitting as a graph of actions $\mathbb{R}_L$ outputted by the Rips machine. The *modular group* $\mathrm{Mod}_{\mathbb{R}_L}(L)$ associated with the splitting $\mathbb{R}_L$ is the subgroup of $\mathrm{Aut}(L)$ generated by the following automorphisms.

1. Inner automorphisms.

2. Dehn twists over the virtually abelian edge groups of $\mathbb{R}_L$: if L_e is an edge group of $\mathbb{R}_L$ and $c \in Z(L_e)$ then the *Dehn twist by c* is the automorphism of L given

by

$$\begin{cases} \tau_c(a) = a,\ \tau_c(b) = cbc^{-1} & \text{if } L = A_1 *_{A_e} A_2,\ a \in A_1 \text{ and } b \in A_2 \\ \tau_c(a) = a,\ \tau_c(t) = tc & \text{if } L = A *_{L_e} \text{ with stable letter } t \text{ and } a \in A. \end{cases}$$

3. Natural extensions of automorphisms of Seifert-type vertex groups that satisfy the following condition: recall that if L_v is a Seifert-type vertex group of $\mathbb{R}_L$ then it fits into a short exact sequence

$$1 \longrightarrow N \longrightarrow L_v \xrightarrow{q} \pi_1(\mathcal{O})$$

where N is finite and $\mathcal{O}$ is a 2-orbifold. If α_v is an extendable automorphism of L_v, then α_v restricts to conjugation on the preimage of (the orbifold fundamental group of) each boundary component of $\mathcal{O}$ under q. Note further than N is a characteristic subgroup of L_v, and therefore there exists a natural map

$$\mathrm{Aut}(L_v) \longrightarrow \mathrm{Aut}(L_v/N) = \mathrm{Aut}(\pi_1(\mathcal{O})).$$

We require that α_v is sent via this map to a mapping class of $\mathcal{O}$ (which corresponds, up to isotopy, to a homeomorphism of $\mathcal{O}$ that fixes the boundary and the conical points).

4. Natural extensions of automorphisms of axial vertex groups, which satisfy the following condition. Denote by L_v an axial vertex group of $\mathbb{R}_L$, then by [57, Lemma 5.1] every subgroup $B \leq L_v$ is virtually abelian, and has a unique maximal subgroup B^+ of index at most 2 which is finite-by-abelian. Denote by $E(L_v)$ the subgroup of L_v generated by its adjacent edge groups. We allow natural extensions of automorphisms α_v of L_v for which:

    (a) α_v fixes the subgroup P_v^+ of L_v which consists of all $g \in L_v$ such that $g \in \ker(\phi)$ for every homomorphism $\phi : L_v \to \mathbb{Z}$ satisfying $E(L_v) \cap A_v^+ \subset \ker(\phi)$, and

    (b) α_v restricts to conjugation on every subgroup $B \leq L_v$ for which $B^+ = P_v^+$.

Remark III.1.22. Note that every modular automorphism of one of the types (1)-(4) above restricts to conjugation on every finite subgroup of L, and hence modular automorphisms always restrict to conjugation on finite subgroups of L.

Our next goal is to show that given a modular automorphism α of L, under some restrictions, one can find an approximation A of L and a lift $\beta \in \operatorname{Aut}(A)$ of α. As Lemma III.1.24 will show, every modular automorphism of types (1)-(3) above admits a lift to an approximation of L. This follows from the extent to which one can approximate the edge groups and the Seifert-type vertex groups of the splitting of L as a graph of actions, as evident in Corollary III.1.14. However, dealing with axial vertex groups is slightly more complicated. We therefore discuss further the structure of axial vertex groups of L and describe a few properties of modular automorphisms of type (4) used in the proof of the shortening argument. We follow [97, Subsection 4.2.1] and refer the reader to [97] for further details.

Suppose that L admits a splitting as a graph of actions $\mathbb{R}_L$ outputted by the Rips machine and that L_v is an axial vertex group of L. Denote by $E \leq L_v$ the torsion subgroup of L_v and let π_E be the quotient map $L_v \twoheadrightarrow L_v/E$. Recall that L_v has a subgroup L_v^+ of index at most 2 which is finite-by-abelian, and let $(L_v/E)^+$ be the image of L_v^+ in L_v/E. The group $(L_v/E)^+$ admits a decomposition $(L_v/E)^+ = A \oplus B$ where A is a finitely generated free abelian group and B is the torsion-free (and abelian) kernel of the action of $(L_v/E)^+$ on the line $T_v \subset T$. Let $\tilde{A} = \pi_E^{-1}(A)$ and $\tilde{B} = \pi_E^{-1}(B)$. As in [97, Subsection 4.2.1], there exists an element $s \in L_v$ such that $L_v = \langle A, B, s \rangle$ and every $g \in L_v$ can be written as a product of the form

$$g = abs^{\eta}$$

where $a \in \tilde{A}$, $b \in \tilde{B}$ and $\eta \in \{0, 1\}$.

We now define the subgroup $\operatorname{Aut}^*(L_v)$ of $\operatorname{Aut}(L_v)$ to be the subgroup which consists of all the automorphisms $\alpha_v \in \operatorname{Aut}(L_v)$ which satisfy the following three properties:

1. α_v preserves $\tilde{A}$;
2. α_v restricts to the identity on $\langle \tilde{B}, s \rangle$;
3. consider the action of L_v on the line $T_v \subset T$. Then for every $x \in T_v$, α_v restricts to conjugation on the stabilizer $(L_v)_x$ of x.

This leads us to define the following subgroup of $\operatorname{Mod}_{\mathbb{R}_L}(L)$.

Definition III.1.23. The group $\operatorname{Mod}^*_{\mathbb{R}_L}(L)$ is the subgroup of $\operatorname{Mod}_{\mathbb{R}_L}(L)$ generated by:

1. modular automorphisms of types (1)-(3) (see Definition III.1.21);

2. modular automorphisms α of type (4) which satisfy the following: if α is a natural extension of $\alpha_v \in \operatorname{Aut}(L_v)$ for an axial vertex group L_v of L, then $\alpha_v \in \operatorname{Aut}^*(L_v)$.

The motivation behind Definition III.1.23 comes from the fact that by [97, Subsection 4.2.1] it is enough to use modular automorphisms which lie in $\operatorname{Mod}^*_{\mathbb{R}_L}(L)$ in the proof of the shortening argument. We finish this discussion with the following easy observation.

Observation: Suppose that L_v is an axial vertex group of L; since the torsion subgroup E of L_v is finite (by Lemma III.1.4) and since the subgroup A of $(L_v/E)^+$ is finitely generated, there are finitely generated subgroups of L_v which contain $\tilde{A}$. In addition, let $\alpha_v \in \operatorname{Aut}^*(L_v)$ and suppose that L'_v is any subgroup of L_v which contains $\tilde{A}$, then the restriction of α_v to L'_v is an automorphism.

Proposition III.1.24. *Suppose that L does not split non-trivially over a subgroup of order $\leq C$ and let $\mathbb{R}_L$ be the graph of actions decomposition of L outputted by the Rips machine. Let $\alpha \in \operatorname{Mod}^*_{\mathbb{R}_L}(L)$. Then there exists an $\mathbb{R}_L$-approximation A of L and $\beta \in \operatorname{Aut}(A)$ such that the solid part of following diagram commutes ω-almost surely.*

$$\begin{tikzcd}
H \arrow[rrr, "\varphi_n"] \arrow[dr, two heads, "q"] \arrow[ddr, two heads, "\varphi_\infty"'] & & & G \\
& A \arrow[urr, "\theta_n"] \arrow[r, "\beta"'] \arrow[d, two heads, "\theta_\infty"] & A \arrow[ur, dotted, "\theta_n"'] \arrow[d, two heads, "\theta_\infty"] & \\
& L \arrow[r, "\alpha"] & L &
\end{tikzcd}$$

Moreover, for every $h \in H$ such that $\varphi_n(h) \neq 1$ ω-almost surely, $\theta_n \circ \beta \circ q(h) \neq 1$ ω-almost surely.

Proof. Recall that the maps φ_n factor via A ω-almost surely and that the maps arising from this factorization are denoted by θ_n, that is $\varphi_n = \theta_n \circ q$ ω-almost surely. The limit map associated to the sequence $(\theta_n)_{n\in\mathbb{N}}$ is denoted by θ_∞. The fact that $\varphi_n = \theta_n \circ q$ ω-almost surely implies that the map $\theta_\infty : A \twoheadrightarrow L$ can also be obtained by quotienting out the image of $\ker(\varphi_\infty)$ in A.

Write $\alpha = \alpha_k \circ \cdots \circ \alpha_1$ where for every $1 \leq i \leq k$, $\alpha_i \in \operatorname{Mod}_{\mathbb{R}_L}(L)$ is a modular automorphism of L of one of the types (1)-(4) appearing in Definition III.1.21. Furthermore, if α_i is a modular automorphism of type (4), we assume that it satisfies the condition appearing in Definition III.1.23. It is enough to find an $\mathbb{R}_L$-approximation A of L and an automorphism $\beta \in \operatorname{Aut}(A)$ for which

$$\theta_\infty \circ \beta = \alpha \circ \theta_\infty.$$

In fact, it is enough to show that there is an $\mathbb{R}_L$-approximation A of L and automorphisms $\beta_1, \dots, \beta_k$ of L such that the following holds for every $1 \le j \le k$:

$$\theta_\infty \circ \beta_j = \alpha_j \circ \theta_\infty.$$

We begin with the construction of the approximation A of L. We define a finite subset C of L as follows:

1. Whenever α_i is a Dehn twist for $1 \le i \le k$, we add to C an element c which lies in an edge group of $\mathbb{R}_L$ and such that α_i is a Dehn twist by c.

2. We keep the notations from the discussion appearing before Definition III.1.23. Suppose now that α_i is a modular automorphism of type (4); denote by $\alpha_v \in \mathrm{Aut}^*(L_v)$ an automorphism of an axial vertex group L_v of L such that α_i is a natural extension of α_v. Let $L_{e_1}, \dots, L_{e_k}$ be the edge groups adjacent to L_v and recall that α_v restricts to conjugation by elements $c_1, \dots, c_k \in L_v$ on $L_{e_1}, \dots, L_{e_k}$ respectively. Let L'_v be a finitely generated subgroup of L_v (over U) which contains $\tilde{A}$ and $c_1, \dots, c_k$; such a group exists by the observation following Definition III.1.23. Denote by S'_v a finite set of generators of L'_v. We add the elements in S'_v to C.

Let A be an $\mathbb{R}_L$-approximation of L which satisfies the following condition: for every $c \in C$ belonging to an edge group L_e of $\mathbb{R}_L$, there is an element $c' \in A_e$ such that $\theta_\infty(c') = c$. Such an approximation A of L exists by Remark III.1.12. Recall that by Lemma III.1.9, A admits a splitting $\mathbb{R}_A$ which satisfies the following: θ_∞ maps every vertex group or edge group of A_v to the corresponding vertex or edge group of L_v. In addition, θ_∞ maps every stable letter in the presentation of A as the fundamental group of $\mathbb{R}_A$ to the corresponding stable letter in the presentation of L as the fundamental group of $\mathbb{R}_L$.

We next construct the automorphisms $\beta_1, \dots, \beta_k$ of A. Let $1 \le j \le k$. We divide the construction of β_j into cases, depending on the type of the modular automorphism $\alpha_j \in \mathrm{Mod}_{\mathbb{R}_L}(L)$.

1. α_j is a modular automorphism of L of type (1), that is conjugation by some element $g \in L$. Let $g' \in A$ be such that $\theta_\infty(g') = g$ and set β_j to be conjugation by g'. It is clear that the desired equality holds.

2. α_j is a modular automorphism of L of type (2), that is a Dehn twist by some element $c \in C$ which lies in an edge group L_e of $\mathbb{R}_L$. By the manner in which the

approximation A was chosen, there is an element $c' \in A_e$ such that $\theta_\infty(c') = c$. Let β_j be a Dehn twist by c'. Assume that by collapsing every edge except for e, A splits as an amalgamated product over A_e, that is $A = A_1 *_{A_e} A_2$; the case where A splits as an HNN extension is similar. It follows that L splits as an amalgamated product over L_e; write $L = L_1 *_{L_e} L_2$. In addition, by the properties of the approximation A, one has $\theta_\infty(A_i) \subset L_i$ for $i \in \{1,2\}$.

Let $g \in A$ and write g as an alternating product of elements from A_1 and A_2, that is $g = a_1 b_1 a_2 \cdots a_m b_m$ with $a_i \in A_1$ and $b_i \in A_2$ for $1 \le i \le m$. We have that

$$\begin{aligned}
\theta_\infty \circ \beta_j(g) &= \theta_\infty \circ \beta_j(a_1 b_1 a_2 \cdots a_m b_m) \\
&= \theta_\infty(a_1(c' b_1 (c')^{-1}) a_2 \cdots a_m (c' b_m (c')^{-1})) \\
&= \theta_\infty(a_1)(c\theta_\infty(b_1)c^{-1})\theta_\infty(a_2)\cdots\theta_\infty(a_m)(c\theta_\infty(b_m)c^{-1}) \\
&= \alpha_j(\theta_\infty(a_1)\theta_\infty(b_1)\theta_\infty(a_2)\cdots\theta_\infty(a_m)\theta_\infty(b_m)) \\
&= \alpha_j \circ \theta_\infty(g).
\end{aligned}$$

3. α_j is a modular automorphism of type (3), that is a natural extension of an automorphism α_v of a vertex group L_v of $\mathbb{R}_L$ of Seifert-type, as described in Definition III.1.21. Let $e_1^v, \cdots, e_\ell^v \in \mathrm{E}(\mathbb{R}_L)$ be an enumeration of the edges of $\mathbb{R}_L$ which are adjacent to v and recall that α_j restricts to conjugation by some $c_{e_i^v} \in L_v$ on $L_{e_i^v}$ for every $1 \le i \le \ell$. In addition, by Corollary III.1.14, θ_∞ maps A_v isomorphically to L_v. This implies that α_v is an isomorphism of A_v, and that there are elements $c'_{e_1^v}, \ldots, c'_{e_\ell^v} \in A_v$ such that $\theta_\infty(c'_{e_i^v}) = c_{e_i^v}$ and α_v restricts to conjugation by $c'_{e_i^v}$ on $A_{e_i^v}$ for every $1 \le i \le \ell$. Let β_j be the natural extension of α_v to A, with respect to the elements $c'_{e_1^v}, \ldots, c'_{e_\ell^v}$. Now let $g \in A$ and write g as a loop in the graph of groups $\mathbb{R}_A$, that is $g = [a_0, e_1, a_1, \ldots, e_k, a_k]$. Then $\beta_j(g) = [b_0, e_1, b_1, \ldots, e_k, b_k]$ where

$$b_i = \begin{cases} a_i & a_i \notin A_v \\ {c'_{e_i}}^{-1} \alpha_v(a_i) c'_{e_{i+1}} & a_i \in A_v \end{cases}.$$

To finish, note that $\theta_\infty(g)$ can be written as a loop $[\theta_\infty(a_0), e_1, \theta_\infty(a_1), \ldots, e_k, \theta_\infty(a_k)]$ in $\mathbb{R}_L$, which implies that $\alpha_j \circ \theta_\infty(g) = [c_0, e_1, c_1, \ldots, e_k, c_k]$ where

$$c_i = \begin{cases} \theta_\infty(a_i) & \theta_\infty(a_i) \notin A_v \\ c^{-1}_{e_i^{-1}} \alpha_v(\theta_\infty(a_i)) c_{e_{i+1}} & a_i \in A_v \end{cases}.$$

One easily sees that $\theta_\infty(b_i) = c_i$ for $1 \le i \le k$, which implies that

$$\theta_\infty \circ \beta_j \circ \cdots \circ \beta_1 = \alpha_j \circ \cdots \circ \alpha_1 \circ \theta_\infty.$$

4. α_j is a modular automorphism of L of type (4), and which satisfies the condition appearing in Definition III.1.23. In particular, α_j is a natural extension of an automorphism $\alpha_v \in \mathrm{Aut}^*(L_v)$ of an axial vertex group L_v of L. Note that θ_∞ maps A_v into L_v; we identify A_v with its image in L_v. By the manner in which the set C was defined, and by the observation following Definition III.1.23, we have that the restriction of α_v to A_v is an automorphism. One can continue as in (3) above, by taking a natural extension of $\alpha_v|_{A_v}$ to A.

Finally, let $\beta = \beta_k \circ \cdots \circ \beta_1$. The construction of the automorphisms $\beta_1, \ldots, \beta_k$ implies that the diagram appearing in the statement of this proposition does commute ω-almost surely. Lastly, let $h \in H$ be such that $\varphi_n(h) \neq 1$ ω-almost surely. It follows that $\theta_\infty \circ q(h) \neq 1$. Therefore $\alpha \circ \theta_\infty \circ q(h) \neq 1$ which implies that $\theta_\infty \circ \beta \circ q(h) \neq 1$. Hence $\theta_n \circ \beta \circ q(h) \neq 1$ ω-almost surely. □

Remark III.1.25. Note that since the action of U on T is elliptic, the modular automorphism α of L restricts to conjugation on U. The proof of Proposition III.1.24 above implies that β also restricts to conjugation on U.

Theorem III.1.26 (The shortening argument)**.** *Suppose that L does not split non-trivially over a finite subgroup of order $\leq C$, then ω-almost surely the homomorphisms φ_n are not short relative to U. More explicitly, denote by $\mathbb{R}_L$ the splitting of L as a graph of actions outputted by the Rips machine. Then there is an $\mathbb{R}_L$-approximation A of L admitting a splitting $\mathbb{R}_A$, and an automorphism $\beta \in \mathrm{Aut}(A)$, for which the following holds: denote by q the quotient map $H \twoheadrightarrow A$ and let $(\theta_n : A \to G)_{n\in\mathbb{N}}$ be such that $\varphi_n = \theta_n \circ q$ ω-almost surely. Then the sequence $(\phi_n = \theta_n \circ \beta \circ q : H \to G)_{n\in\mathbb{N}}$ satisfies the following:*

1. *$\phi_n|_U$ coincides with $\varphi_n|_U$ up to conjugation ω-almost surely;*

2. *$||\phi_n|| < ||\varphi_n||$ ω-almost surely;*

3. *for every $h \in H$ such that $\varphi_n(h) \neq 1$ ω-almost surely, $\phi_n(h) \neq 1$ ω-almost surely.*

Remark III.1.27. Note that conditions *1.* and *2.* above imply that φ_n is not short relative to U ω-almost surely. Furthermore, condition *2.* can be equivalently phrased as follows: $||\theta_n \circ \beta|| < ||\theta_n||$ ω-almost surely, where the lengths are taken with respect to the set $q(S)$. Condition *3.* is equivalent to each of the following two conditions:

(*3.'*) $\underleftarrow{\ker}_\omega((\phi_n)_{n\in\mathbb{N}}) \subset \underleftarrow{\ker}_\omega((\varphi_n)_{n\in\mathbb{N}})$;

(*3."*) $\beta^{-1}(\underleftarrow{\ker}_\omega((\theta_n)_{n\in\mathbb{N}})) \subset \underleftarrow{\ker}_\omega((\theta_n)_{n\in\mathbb{N}})$.

Remark III.1.28. The version of the shortening argument appearing in [57, Theorem 5.29] satisfies conditions *1.* and *2.* above, but does not necessarily satisfy condition *3.* To obtain condition *3.* we approximate a single modular automorphism of L, rather than a sequence of modular automorphisms, and the result follows from Lemma III.1.24.

The second version of the shortening argument that will be proved is a strengthened version of Theorem III.1.26 which accommodates the use of JSJ decompositions of limit groups over finite groups of order less than C. Note that in this version of the shortening argument, we assume that a single vertex group of a JSJ decomposition of L over groups of order $\leq C$ admits a splitting outputted by the Rips machine, and we shorten the homomorphisms with respect to the generators of this vertex group.

Theorem III.1.29. *Let $\mathbb{J}_L$ be a reduced JSJ splitting of L over finite groups of order $\leq C$ relative to U and let u be the vertex fixed by U. Let $\mathbb{R}_{L_u}$ be the splitting of L_u as a graph of actions outputted by the Rips machine, and let $\mathbb{RJ}_L$ be the splitting of L obtained from $\mathbb{J}_L$ by replacing u with $\mathbb{R}_{L_u}$. Let H_u be a finitely generated subgroup of H containing U and such that $\varphi_\infty(H_u) = L_u$. Let S_u be a finite generating set of H_u. Then the following hold:*

1. *there exist an $\mathbb{RJ}_L$-approximation A of L admitting a splitting $\mathbb{RJ}_A$ as in Corollary III.1.15 and a sequence $(\theta_n : A \to G)_{n\in\mathbb{N}}$ that satisfies $\varphi_n = \theta_n \circ q$ ω-almost surely (where q is the quotient map $H \twoheadrightarrow A$);*

2. *denote by $\mathbb{R}_{A_u}$ the subgraph of $\mathbb{RJ}_A$ corresponding to the subgraph $\mathbb{R}_{L_u}$ of $\mathbb{RJ}_L$. Denote by A_u the fundamental group of $\mathbb{R}_{A_u}$. There exists an automorphism β_u of A_u that admits a natural extension $\beta \in \mathrm{Aut}(A)$, and such that the sequence $(\phi_n = \theta_n \circ \beta \circ q : H \to G)_{n\in\mathbb{N}}$ satisfies the following properties:*
    - *(a) $\phi_n|_U$ coincides with $\varphi_n|_U$ up to conjugation ω-almost surely;*
    - *(b) $\|\phi_n|_{H_u}\| < \|\varphi_n|_{H_u}\|$ ω-almost surely (where the lengths are taken with respect to the S_u);*
    - *(c) for every $h \in H$ such that $\varphi_n(h) \neq 1$ ω-almost surely, $\phi_n(h) \neq 1$ ω-almost surely.*

Remark III.1.30. As in Theorem III.1.26, the condition *2.(b)* above is equivalent to

$$\|(\theta_n \circ \beta)|_{q(H_u)}\| < \|\theta_n|_{q(H_u)}\|.$$

We need the following lemma in order to prove Theorems III.1.26 and III.1.29.

Lemma III.1.31. *Suppose that L does not split non-trivially over a finite subgroup of order at most C, and that $e \subset T$ (where T is the limiting tree on which L acts) is an edge of a simplicial subtree of T; denote its stabilizer by L_e. Then there is an element $\tilde{c}_e \in H$ whose image c_e in L is contained in $Z(L_e)$ and such that $\varphi_n(\tilde{c}_e)$ is hyperbolic ω-almost surely.*

To prove this lemma, we use the following standard result which states that if an isometry of a hyperbolic space does not move two distant points by a big amount, then it acts as a quasi-translation on the geodesic connecting them.

Lemma III.1.32. *[95, Lemma 3.5] Let (X, d) be a δ-hyperbolic space and let $g : X \to X$ be an isometry. Suppose that for $x, y \in X$*

$$d(x, g(x)) + d(y, g(y)) < 2d(x, y) - 4\delta.$$

Then for some $\lambda \in \mathbb{R}$ such that $|\lambda| \le \max\{d(x, g(x)), d(y, g(y))\}$ the isometry g acts by $(\lambda, 2\delta)$-quasi-translation on a subgeodesic of $[x, y]$: for every $p \in [x, y]$ at distance greater than $\max\{d(x, g(x)), d(y, g(y))\}$ from both x and y (if such a point p exists), one has

$$d(g(p), p_\lambda) < 2\delta \quad \textit{and} \quad d(g^{-1}(p), p_{-\lambda}) < 2\delta$$

where p_λ and $p_{-\lambda}$ are the points on $[x, y]$ which lie at distance λ from p.

Proof of Lemma III.1.31. Since L does not split non-trivially over a finite subgroup of order $\le C$, Lemma III.1.4 implies that there exists $c \in L_e$ of infinite order and such that $\varphi_n(\tilde{c})$ is hyperbolic ω-almost surely. Fix $\varepsilon = 8\delta$ and let N and R be the corresponding acylindricity constants. We claim that $c_e = c^{N!}$ lies in $Z(L_e)$. Let $g \in L_e$, write $e = [x, y]$ and let $(x_n)_{n \in \mathbb{N}}$ and $(y_n)_{n \in \mathbb{N}}$ be approximating sequences for x and y respectively; let $\tilde{g}$ be a lift of g to H. Since g fixes both x and y in the limiting tree, $\varphi_n(\tilde{g})$ must displace both x_n and y_n by a distance which is significantly smaller than $d(x_n, y_n)$ ω-almost surely. More precisely, we have that

$$\begin{aligned}
&d(x_n, y_n) > 100 \cdot R \\
&d(x_n, y_n) > 100 \cdot \max\{d(x_n, \varphi_n(\tilde{g})x_n), d(y_n, \varphi_n(\tilde{g})y_n)\} \\
&d(x_n, y_n) > 100 \cdot \max\left\{d(x_n, \varphi_n(\tilde{c})^j x_n), d(y_n, \varphi_n(\tilde{c})^j y_n)\right\}
\end{aligned}$$

holds for all $j \in \{1, \dots, N+1\}$ ω-almost surely. Choose two points $p_n, q_n \in e$ satisfying $d(x_n, p_n) < d(x_n, q_n)$, $d(p_n, q_n) > R$ and

$$\min\{d(x_n, p_n), d(q_n, y_n)\} > 10 \cdot \left(\max_{h \in \{g, c, c^2, \dots, c^{N+1}\}} \left\{d\left(x_n, \varphi_n(\tilde{h})x_n\right), d\left(y_n, \varphi_n(\tilde{h})y_n\right)\right\} \right)$$

ω-almost surely. It follows from Lemma III.1.32 that each of $g, c, c^2, \dots, c^{N+1}$ acts on a subsegment of $[x, y]$ which contains p_n and q_n by 2δ-quasi-translation, and therefore both $d(p_n, [\varphi_n(\tilde{g}), \varphi_n(\tilde{c})^j] p_n) \le 8\delta$ and $d(q_n, [\varphi_n(\tilde{g}), \varphi_n(\tilde{c})^j] q_n) \le 8\delta$ hold for every $j \in \{1, \dots, N+1\}$ ω-almost surely. The acylindricity condition implies that not all of the $N+1$ commutators can be distinct and there are $i, j \in \{1, \dots, N+1\}$ for which

$$\begin{aligned}\left[\varphi_n(\tilde{g}), \varphi_n(\tilde{c})^j\right] &= \varphi_n(\tilde{g})\varphi_n(\tilde{c})^j\varphi_n(\tilde{g})^{-1}\varphi_n(\tilde{c})^{-j} \\ &= \varphi_n(\tilde{g})\varphi_n(\tilde{c})^i\varphi_n(\tilde{g})^{-1}\varphi_n(\tilde{c})^{-i} \\ &= \left[\varphi_n(\tilde{g}), \varphi_n(\tilde{c})^i\right].\end{aligned}$$

In particular, $\varphi_n(\tilde{g})$ commutes with $\varphi_n(\tilde{c})^{j-i}$ ω-almost surely. Since $|j-i| \le N$, the element $c_e = c^{N!}$ is a power of c^{j-i} and $\varphi_n(\tilde{g})$ commutes with $\varphi_n(\tilde{c}_e)$ ω-almost surely; hence g commutes with c_e. □

Proof of Theorem III.1.26. The proof of this theorem is divided in two:

1. using results from [97] and [100] and explaining how they adapt to our setting, we first find a suitable modular automorphism $\alpha \in \mathrm{Mod}^*_{\mathbb{R}_L}(L)$,
2. then, by means of Proposition III.1.24, we find an $\mathbb{R}_L$-approximation A of L and an automorphism β of A which satisfy the desired properties.

The idea behind the proof amounts to finding a finite sequence of modular automorphisms of L, each of which shortens the actions of the generators S of H over U with respect to the different vertex actions in the graph of actions decomposition $\mathbb{R}_L$ of L. The construction of the modular automorphism $\alpha \in \mathrm{Mod}^*_{\mathbb{R}_L}(L)$ relies on the proofs appearing in [97] for the axial and Seifert-type cases, and on the proof appearing in [100] for the simplicial case. We begin with vertex actions $L_v \curvearrowright T_v$ which admit dense orbits, namely axial and Seifert-type vertex actions. Recall that if L_v is of Seifert-type, then the index of the 2-orbifold subgroup of L_v is at most C. Therefore, by [97, Subsections 4.2.1 and 4.2.2], for every such vertex action and every finite subset $F \subset H$, there exists a modular automorphism $\alpha_v^F \in \mathrm{Mod}^*_{\mathbb{R}_L}(L)$ of type (3) or (4) which satisfies the following: denote by $o = (o_n)_{n\in\mathbb{N}}$ the base point of T, then for every $f \in F$,

$$d_T(o, \alpha_v^F(\varphi_\infty(f))o) < d_T(o, \varphi_\infty(f)o)$$

whenever $[o, \varphi_\infty(f)o]$ has a non-degenerate intersection with a translate of T_v in T, and $d_T(o, \alpha_v^F(\varphi_\infty(f))o) = d_T(o, \varphi_\infty(f)o)$ otherwise.

This allows us to shorten the actions (on the real tree T) of all the generators s for which $[o, \varphi_\infty(s)o]$ intersects (a translate of) an axial or a Seifert-type component of T non-degenerately: let $v_1, \cdots, v_m$ be an enumeration of the axial and Seifert-type vertices of $\mathbb{R}_L$; we define a sequence $(\alpha_1, \dots, \alpha_m) \in \mathrm{Mod}^*_{\mathbb{R}_L}(L)^m$ iteratively. Let $\alpha_1 = \alpha_{v_1}^{\varphi_\infty(S)}$, and after $\alpha_1, \dots, \alpha_i$ were defined let

$$\alpha_{i+1} = \alpha_{v_{i+1}}^{\alpha_i \circ \cdots \alpha_1(S)}.$$

Now note that for every $s \in S$ such that $[o, \varphi_\infty(f)o]$ has a non-degenerate intersection with (a translate of) an axial or a Seifert-type component of T, we have that

$$d_T(o, \alpha_m \circ \cdots \circ \alpha_1 \circ \varphi_\infty(s)o) < d_T(o, \varphi_\infty(s)o).$$

Bring to mind that the modular automorphism $\alpha_m \circ \cdots \circ \alpha_1$ of L does not necessarily shorten the actions of S on T: it could be that for every $s \in S$, $[o, \varphi_\infty(s)o]$ is contained entirely in the discrete part of T, that is, it does not intersect non-degenerately (translates of) axial and Seifert-type components of T. We therefore adapt [100, Theorem 6.1] to our settings. This theorem states that for every finite set $F \subset H$ there is a modular automorphism $\alpha^F_{\text{sim}} \in \mathrm{Mod}^*_{\mathbb{R}_L}(L)$ of L, which can be written as a composition of Dehn twists about elements which lie in the edge groups of $\mathbb{R}_L$, and which satisfies the following: for every $f \in F$ which does not fix o, and such that $[o, \varphi_\infty(f)o]$ lies entirely in the discrete part of T, let $f_{\alpha^F_{\text{sim}}}$ be a lift of $\alpha^F_{\text{sim}}(\varphi_\infty(f))$ to H; then

$$d\left(o_n, \varphi_n\left(f_{\alpha^F_{\text{sim}}}\right) o_n\right) < d(o_n, \varphi_n(f)o_n),$$

ω-almost surely. In addition, for every $f \in F$,

$$d_T(o, \alpha^F_{\text{sim}}(\varphi_\infty(f))o) = d_T(o, \varphi_\infty(f)o).$$

Note that in this case the modular automorphism α^F_{sim} does not shorten the actions of the elements in F on T, but rather shortens the actions of the elements in F directly on the spaces X_n.

In the proof of Theorem [100, Theorem 6.1], one finds Dehn twists over the edge groups of $\mathbb{R}_L$, where each Dehn twist does not affect the displacement of o by the elements of F in T, but does affect the displacement of o_n by $\varphi_n(F)$ in X_n. The construction of a Dehn twist over the edge group corresponding to an edge e of $\mathbb{R}_L$ is divided into three cases (below, $\tilde{e}$ denotes an edge of a simplicial subtree of T which corresponds to e as in the paragraph following Definition III.1.1):

1. o lies in the interior of $\tilde{e}$ and L splits as an amalgamated product over the stabilizer of $\tilde{e}$;
2. o lies in the interior of $\tilde{e}$ and L splits as an HNN extension over the stabilizer of $\tilde{e}$;
3. o does not lie in the interior of $\tilde{e}$, and is one of its vertices.

In the last case, one considers all such edges $\tilde{e}$ in (simplicial subtrees of) T which are adjacent to o and shortens the action of the generators with respect to all of these edges simultaneously. The proof appearing in [100] can be transitioned almost seamlessly to our setting. In [100], the stabilizers of edges in the limiting tree are cyclic, whereas in our case they are virtually abelian. Therefore, in our case we cannot take Dehn twists by any element in an edge group L_e of $\mathbb{R}_L$, and instead take Dehn twists by a power of an element $c_e \in Z(L_e)$, of infinite order, which exists by Lemma III.1.31. The only other parts which do not carry over to our settings are Lemmas 6.2, 6.5, 6.8 and 6.11 in [100] which assert that there are elements of infinite order in the edge groups of $\mathbb{R}_L$ which satisfy the following: let $\tilde{c} \in H$ be a lift of such an element, then $\varphi_n(\tilde{c})$ displaces certain points in X_n by a distance bounded from below by 10δ or 20δ ω-almost surely. These are the elements by which one takes the Dehn twists. We can easily overcome this: by [21, Lemma 2.2] (recall that X is a simplicial graph), there is $\eta > 0$ which depends only on the hyperbolicity constant of X and the acylindricity constants of the action $G \curvearrowright X$ such that

$$\eta < \ell(g) = \lim_{n\to\infty} \frac{1}{n} d(g^n o, o)$$

(where o is the basepoint of X) for every hyperbolic $g \in G$. In addition, denote $||g|| = \inf_{x\in X}\{d(x, gx)\}$ and by [36, Lemma 10.6.4] we have that $\ell(g) \le ||g||$ and clearly $\ell(g^n) = n\ell(g)$. Therefore, given $D > 0$ there exists N such that

$$D < N\eta < N\ell(g) = \ell(g^N) \le ||g^N|| = \inf_{x\in X}\{d(x, g^N x)\}$$

for every hyperbolic $g \in G$. We also remark that one might have to enlarge the constant C_0 appearing in [100] to accommodate the choice of N above. This implies that the proof of [100, Theorem 6.1] can be carried out in our setting.

Now let

$$\alpha = \alpha_{\text{sim}}^{\alpha_m \circ \cdots \circ \alpha_1 \circ \varphi_\infty(S)} \circ \alpha_m \circ \cdots \circ \alpha_1 \in \text{Mod}^*_{\mathbb{R}_L}(L).$$

For every $s \in S$, denote by $s_\alpha \in H$ a lift of $\alpha(\varphi_\infty(s))$ to H. Since $d_T(o, \alpha(\varphi_\infty(s))o) < d_T(o, \varphi_\infty(s)o)$ whenever $[o, \varphi_\infty(s)o]$ intersects a translate of an axial or a Seifert type component of T non-degenerately, the following holds ω-almost surely:

$$d(o_n, \varphi_n(s_\alpha)o_n) < d(o_n, \varphi_n(s)o_n).$$

By Proposition III.1.24 there exists an $\mathbb{R}_L$-approximation A of L and $\beta \in \mathrm{Aut}(A)$ such that the following diagram commutes ω-almost surely.

$$\begin{tikzcd}
H \arrow[rrr, "\varphi_n"] \arrow[dr, two heads, "q"] \arrow[ddr, two heads, "\varphi_\infty"'] & & & G \\
& A \arrow[r, "\beta"'] \arrow[urr, "\theta_n"] \arrow[d, two heads, "\theta_\infty"] & A \arrow[d, two heads, "\theta_\infty"] & \\
& L \arrow[r, "\alpha"] & L &
\end{tikzcd}$$

We claim that the approximation A and its automorphism β satisfy the properties described in the theorem.

By Remark III.1.25, since the action of U on T is elliptic, β restricts to conjugation on U. Hence condition *1.* holds. For *2.*, note that for every $s \in S$, $\alpha(\varphi_\infty(s))$ and $\beta(q(s))$ share the same lift in H; setting $\varphi_n = \theta_n \circ \beta \circ q$, it follows that

$$d(o_n, \theta_n \circ \beta \circ q(s)o_n) = d(o_n, \phi_n(s)o_n) < d(o_n, \varphi_n(s)o_n),$$

and in particular, since o_n realizes the infimum $\inf_{x \in X} \max_{s \in S} d(x, \varphi_n(s)x)$,

$$\begin{aligned}
\|\phi_n\| &= \inf_{x \in X} \max_{s \in S} d(x, \phi_n(s)x) \\
&\leq \max_{s \in S} d(o_n, \phi_n(s)o_n) \\
&< \max_{s \in S} d(o_n, \varphi_n(s)o_n) \\
&= \|\varphi_n\|
\end{aligned}$$

ω-almost surely. Lastly, Property *3.* follows directly from Proposition III.1.24. ☐

The proof of Theorem III.1.29 is very similar to the proof of Theorem III.1.26.

Proof of Theorem III.1.29. The proof is identical to that of Theorem III.1.26, with one change (applied to both Theorem III.1.26 and Proposition III.1.24): one has to take natural extensions of automorphisms of axial and Seifert-type vertex group with respect to the entire graph of groups decomposition $\mathbb{RJ}_L$ of L (and not with respect to a graph of actions outputted by the Rips machine) instead of modular automorphisms of L with respect to $\mathbb{R}_L$. ☐

III.2 Test sequences

The goal of this section is to define *test sequences* and prove important preliminary results about them; these will play a crucial role in the proof of Merzlyakov's theorem (Theorem E). This section culminates in the use of small cancellation arguments that allow us to show that a certain subtree of the limiting tree T_L associated to a limit group L must be simplicial.

III.2.1 Transverse coverings

We will use the following definitions from [61]:

Definition III.2.1 ([61, Definition 4.6]). Let T be a real tree endowed with an action of a group G, and let $(Y_j)_{i\in J}$ be a G-invariant family of non-degenerate closed subtrees of T. We say that $(Y_j)_{j\in J}$ is a *transverse covering* of T if the following two conditions hold:

- *Transverse intersection*: if $Y_i \cap Y_j$ contains more than one point, then $Y_i = Y_j$.
- *Finiteness condition*: every arc of T is covered by finitely many Y_j.

Definition III.2.2 ([61, Definition 4.8]). Let T be a real tree, and let $(Y_j)_{j\in J}$ be a transverse covering of T. The *skeleton* of $(Y_j)_{j\in J}$ is the bipartite simplicial tree S defined as follows:

1. $V(S) = V_0(S) \sqcup V_1(S)$ where $V_1(S) = \{Y_j \mid j \in J\}$ and $V_0(S)$ is the set of points $x \in T$ that belong to at least two distinct subtrees Y_i and Y_j,
2. there is an edge $\varepsilon = (Y_j, x)$ between $Y_j \in V_1(S)$ and $x \in V_0(S)$ if and only if x, viewed as a point of T, belongs to Y_j, viewed as a subtree of T.

Remark III.2.3. The stabilizer of a vertex of S is the stabilizer G_{Y_i} or G_x of the corresponding subtree or point of T. The stabilizer of an edge $\varepsilon = (Y_j, x)$ is $G_{Y_i} \cap G_x$. Moreover, the action of G on S is minimal, provided that the action of G on T is minimal (see [61, Lemma 4.9]).

III.2.2 Bounding the number of branch points

In this short subsection, we state a theorem of Guirardel and Levitt that will allow us to both

- bound the number of L-orbits of branch points in the real tree associated to the limit group L, and
- bound the number of L-orbits of directions at each branch point in the real tree associated to L.

Let T be a tree, and let G be a group acting on T. Recall that an arc $I \subset T$ is called stable if for every non-degenerate subarc $J \subset I$, the stabilizer of J is equal to the stabilizer of I (see Definition III.1.2). Otherwise, I is called unstable.

Definition III.2.4. An action on a real tree is *K-superstable* if every arc whose pointwise stabilizer has order greater than K is stable.

Let T be a real tree, and let x be a point of T. A *direction* at x is a connected component of $T \smallsetminus \{x\}$. We say that x is a *branch point* if there are at least three directions at x. The following result appears in work in preparation [63] by Guirardel and Levitt (improving [60]).

Theorem III.2.5 ([63, Theorem 4.27])**.** *Let L be a group acting on a real tree T_L. Suppose that L is finitely generated relative to a countable subgroup Γ that is elliptic in T_L. Suppose that the following two conditions are satisfied:*

1. *the action is K-superstable for some constant K;*
2. *arc stabilizers are finitely generated.*

Then every point stabilizer is finitely generated relative to Γ. Furthermore, the number of orbits of branch points in T_L is finite, and the number of orbits of directions at branch points in T_L is finite.

III.2.3 Geometric small cancellation

Small cancellation theory is a fruitful notion that is often used to provide interesting (counter) examples or groups that exhibit a strange behaviour. The original proof of Merzlyakov's theorem relies implicitly on combinatorial small cancellation, and Sela's geometric proof of Merzlyakov's theorem utilizes metric small cancellation conditions. We will use a geometric analogue of small cancellation in our proof.

Let (X, d) be a δ-hyperbolic simplicial graph, let G be a group acting on (X, d) by isometries, and let g be an element of G. We define the *translation length* of g by $||g|| = \inf_{x \in X} d(x, gx)$. If g is loxodromic, the *quasi-axis* of g, denoted by $A(g)$, is the union of all geodesics joining $g^{-\infty} \in \partial X$ and $g^{+\infty} \in \partial X$ (note that this definition is slightly different to the one given in Subsection I.2). By [37, Lemma 2.26], the quasi-axis $A(g)$ is 11δ-quasi-convex.

Definition III.2.6. Let g and g' be loxodromic elements of G. The *fellow traveling constant* of g and g', $\Delta(g, g')$, is defined as:

$$\Delta(g, g') = \operatorname{diam}\left(A(g)^{+100\delta} \cap A(g')^{+100\delta}\right) \in \mathbb{N} \cup \{\infty\},$$

where $A(g)^{+100\delta}$ is the 100δ-neighbourhood of $A(g)$ in (X, d), and $A(g')^{+100\delta}$ is defined similarly.

Recall that if G acts acylindrically on (X, d), then every loxodromic element $g \in G$ is contained in a unique maximal infinite virtually cyclic subgroup $\Lambda(g)$ of G (see Lemma I.2.5). Moreover, there exists a constant $N(g) \geq 0$ such that every element $h \in G$ satisfying $\Delta(g, hgh^{-1}) \geq N(g)$ belongs to $\Lambda(g)$ (see for example [38]). Lastly, if g and h are loxodromic elements of G, then either $\Lambda(g) = \Lambda(h)$ or $\Lambda(g) \cap \Lambda(h)$ is finite.

Definition III.2.7. Let $\varepsilon > 0$. We say that a hyperbolic element $g \in G$ satisfies the *ε-small cancellation condition* if the following holds: for every $h \in G$, if

$$\Delta(g, hgh^{-1}) > \varepsilon \|g\|,$$

then h and g commute (and therefore h belongs to $\Lambda(g)$).

Remark III.2.8. The definition above implies that if g satisfies the ε-small cancellation condition then g is central in $\Lambda(g)$.

Definition III.2.9. Let $\varepsilon > 0$. We say that a tuple $(g_1, \ldots, g_p) \in G^p$ of loxodromic elements of G satisfies the *ε-small cancellation condition* if the following condition holds: for every $h \in G$, and for every $i, j \in \{1, \ldots, p\}$, if

$$\Delta(g_i, hg_jh^{-1}) > \varepsilon \min(\|g_i\|, \|g_j\|),$$

then $i = j$, and the elements h and g_i commute. In particular, h belongs to $\Lambda(g_i)$.

Remark III.2.10. As before, in this case g_i is central in $\Lambda(g_i)$ and for every $i, j \in \{1, \ldots, p\}$ we have that $\Lambda(g_i) \neq \Lambda(g_j)$, i.e. $\Lambda(g_i) \cap \Lambda(g_j) = E(G)$.

III.2.4 Test sequences and preliminary lemmas

Let G be an acylindrically hyperbolic group. By Theorem I.2.7, there exists a generating set S of G such that the Cayley graph $X = X(G,S)$ of G with respect to S is δ-hyperbolic for some $\delta \geq 0$. Furthermore, the natural action of G on X is acylindrical and non-elementary. We denote by d the word metric on X (with respect to S).

Definition III.2.11. Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G. Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let $\Sigma(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1$ be a finite system of equations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are tuples of variables. Denote $G_\Sigma = \langle \boldsymbol{x},\boldsymbol{y},\boldsymbol{a} \mid R(\boldsymbol{a}) = 1, \Sigma(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1\rangle$. Let $p = |\boldsymbol{x}|$ be the arity of $\boldsymbol{x}$, and let x_i denote the i-th component of $\boldsymbol{x}$. Let $(\sigma_1,\ldots,\sigma_p)$ be a p-tuple of elements of $\mathrm{Aut}_G(E(G))$ (recall that $\mathrm{Aut}_G(E(G))$ was defined in the Introduction, and it consists of all automorphisms of G that restrict to conjugation on $E(G)$). A sequence of homomorphisms $(\varphi_n : G_\Sigma \to G)_{n\in\mathbb{N}}$ is called a $(\sigma_1,\ldots,\sigma_p)$-*test sequence* if the following four conditions hold:

1. for every n, the morphism φ_n restricts to conjugation on $\boldsymbol{a}$,
2. for every $i \in \{1,\ldots,p\}$, the translation length $||\varphi_n(x_i)||$ of $\varphi_n(x_i)$ tends to infinity as n tends to infinity,
3. for every $i,j \in \{1,\ldots,p\}$, there exists a real number $r_{i,j} \in (0,\infty)$ such that the ratio
$$\frac{||\varphi_n(x_i)||}{||\varphi_n(x_j)||}$$
tends to $r_{i,j}$ as n tends to infinity,
4. there exists a sequence of positive real numbers $(\varepsilon_n)_{n\in\mathbb{N}}$ converging to 0 such that, for every n, the tuple $\varphi_n(\boldsymbol{x})$ satisfies the ε_n-small cancellation condition (see Definition III.2.9). In addition, the following equality holds for every $1 \leq i \leq p$:
$$\Lambda(\varphi_n(x_i)) = \langle \varphi_n(x_i), E(G) \mid \mathrm{ad}(\varphi_n(x_i))|_{E(G)} = \sigma_i\rangle.$$
In particular, the image of $\varphi_n(x_i)$ in $\Lambda(\varphi_n(x_i))/E(G)$ has no roots.

In the particular case where $(\sigma_1,\ldots,\sigma_p) = (\mathrm{id}_{E(G)},\ldots,\mathrm{id}_{E(G)})$, one simply says that $(\varphi_n)_{n\in\mathbb{N}}$ is a *test sequence*.

Remark III.2.12. Note that a test sequence is always divergent: by item 2. above, the translation length $||\varphi_n(x_i)||$ goes to infinity as n goes to infinity. $||\varphi_n(x_i)||$ is also at most the scaling factor $||\varphi_n||$ (see Definition II.1.32) since x_i belongs to the generating set $\{\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}\}$ of G_Σ.

Remark III.2.13. Let $(\varphi_n : G_\Sigma \to G)_{n\in\mathbb{N}}$ be a $(\sigma_1, \ldots, \sigma_p)$-test sequence. Let U be the subgroup of G_Σ generated by $\boldsymbol{x}$ and $\boldsymbol{a}$. Since the translation length $||\cdot||$ is constant on conjugacy classes, one easily sees that any sequence $(\theta_n : G_\Sigma \to G)_{n\in\mathbb{N}}$ such that θ_n coincides on U with φ_n up to conjugation is also a $(\sigma_1', \ldots, \sigma_p')$-test sequence for some $(\sigma_1', \ldots, \sigma_p') \in \mathrm{Aut}_G(E(G))^p$.

Remark III.2.14. Note that any subsequence of a $(\sigma_1, \ldots, \sigma_p)$-test sequence is also a $(\sigma_1, \ldots, \sigma_p)$-test sequence.

The following easy lemma will be useful in the sequel.

Lemma III.2.15. *Let $(\varphi_n)_{n\in\mathbb{N}}$ be a $(\sigma_1, \ldots, \sigma_p)$-test sequence. For every infinite subset $A \subset \mathbb{N}$ and every integer $1 \leq i \leq p$ we have*

$$\bigcap_{n\in A} \Lambda(\varphi_n(x_i)) = E(G).$$

Proof. Suppose that $g \in G$ belongs to $\Lambda(\varphi_n(x_i))$ for every $n \in A$. Therefore, for every $n \in A$ there exists an integer k_n and an element $g_n \in E(G)$ such that $g = \varphi_n(x_i)^{k_n} g_n$. Now, observe that k_n must be equal to 0 for every n large enough, otherwise (up to extracting a subsequence) $||\varphi_n(x_i)^{k_n}||$ goes to infinity, and so does the constant $||g||$, which is a contradiction. It follows that g belongs to $E(G)$. ☐

Keeping the same notations as above, let $(\varphi_n : G_\Sigma \to G)_{n\in\mathbb{N}}$ be a $(\sigma_1, \ldots, \sigma_p)$-test sequence. Let L be the quotient of G_Σ by the stable kernel of the sequence $(\varphi_n)_{n\in\mathbb{N}}$ and let $\varphi_\infty : G_\Sigma \twoheadrightarrow L$ be the corresponding limiting map. Let S be a finite generating set of L containing the images of $\boldsymbol{x}$ and $\boldsymbol{a}$ in L (still denoted by $\boldsymbol{x}$ and $\boldsymbol{a}$), and let (X, d) be a hyperbolic Cayley graph of G on which G acts acylindrically and non-elementarily. By Remark III.2.12, the sequence $(\varphi_n)_{n\in\mathbb{N}}$ is divergent, and hence the construction of the real tree described in Subsection II.1.3 is applicable. From now on, we denote the real tree corresponding to L by T_L (and denote its basepoint by o).

Lemma III.2.16. *Define $K = \{g \in G_\Sigma \mid \varphi_n(g) \in E(G)$ ω-almost surely$\}$ and $F = \varphi_\infty(K)$. Then F is finite. Moreover, it is the unique maximal finite subgroup of L normalized by $\varphi_\infty(x_i)$, for every element x_i of $\boldsymbol{x} = (x_1, \ldots, x_p)$.*

Proof. Note that F is clearly finite. Indeed, for every finite subset $K' \subset K$ one has that $|\varphi_\infty(K')| \leq |E(G)|$ because $\varphi_n(K')$ is contained in $E(G)$ ω-almost surely, and hence $|\varphi_\infty(K)| \leq |E(G)|$. We will prove that F is the unique maximal finite subgroup of L normalized by $\varphi_\infty(x_i)$ for every $1 \leq i \leq p$. To do so, we will first prove the existence of a unique maximal finite subgroup of L normalized by $\varphi_\infty(x_i)$, which we denote by F_i. Then, we will prove that $F_i = F$.

Let $\{F_{i,j}\}_{j\in J}$ be the collection of all the finite subgroups of L that are normalized by $\varphi_\infty(x_i)$, let $A_i = \bigcup_{j\in J} F_{i,j}$ and let F_i be the subgroup of L generated by A_i. Let $K_i = \varphi_\infty^{-1}(F_i)$. We claim that K_i is contained in K, that is, that $\varphi_n(k)$ belongs to $E(G)$ ω-almost surely for every $k \in K_i$.

Suppose that $k \in K_i$ is a preimage of an element of A_i. From the definition of A_i, $\varphi_\infty(k)$ is contained in a finite subgroup of L normalized by $\varphi_\infty(x_i)$. Therefore, there exists an integer $m \geq 1$ such that $\varphi_\infty([k, x_i^m]) = 1$. It follows that $\varphi_n([k, x_i^m])$ is trivial ω-almost surely. By Lemma I.2.5, this implies that $\varphi_n(k)$ belongs to $\Lambda(\varphi_n(x_i))$ ω-almost surely. Recall that by the definition of a test sequence, $\Lambda(\varphi_n(x_i))$ is generated by $\varphi_n(x_i)$ and $E(G)$. In particular, $\Lambda(\varphi_n(x_i))$ is $E(G)$-by-$\mathbb{Z}$. This shows that $\varphi_n(k)$ is contained in $E(G)$, since $\varphi_n(k)$ is of finite order by definition of A_i. Note that since F_i is generated by A_i, this fact remains true if k is any element of K_i (and not necessarily a preimage of an element of A_i). It follows that K_i is contained in K. Since F is finite, F_i is finite as well. Moreover, as evident from our construction, F_i is the unique maximal finite subgroup of L normalized by $\varphi_\infty(x_i)$.

We next prove that K_i and K coincide (and in particular K_i does not depend on i). Recall that K_i is a subset of K. Let $k \in K$; we will prove that k belongs to K_i. Let K_i' be the subgroup of G_Σ defined by

$$K_i' = \langle \{x_i^\ell k x_i^{-\ell},\ \ell \in \mathbb{N}\}\rangle.$$

From the definition of K, $\varphi_n(k)$ lies in $E(G)$ ω-almost surely, and hence $\varphi_n(x_i^\ell k x_i^{-\ell})$ belongs to $E(G)$ ω-almost surely (since $E(G)$ is normal in G). It follows that $\varphi_n(K_i')$ is a subgroup of $E(G)$ ω-almost surely. In addition, this subgroup is normalized by $\varphi_n(x_i)$ by construction. As a consequence, $\varphi_\infty(K_i')$ is contained in F_i. In particular, $\varphi_\infty(k)$ belongs to F_i, which implies that k belongs to K_i. Thus $K_i = K$ for every $1 \leq i \leq p$ and the finite groups $F_1, \ldots, F_p$ are all equal to F. ☐

In what follows, we abuse notation and use x_i to denote not just the element of G_Σ, but also its image in L under φ_∞. Keeping the previous notation and assumptions, let $Y \subset L$ be the stabilizer of the base point o of T_L. Note that Y contains each element of the tuple $\boldsymbol{a}$. Indeed, each morphism φ_n restricts to a conjugation on $\boldsymbol{a}$.

Lemma III.2.17. *Suppose that* $\Gamma = \langle \boldsymbol{x} \cup Y \rangle \leq L$ *does not fix a point in* T_L*. Then the minimal subtree* $T_\Gamma \subset T_L$ *of* Γ *is simplicial. Moreover, the graph of groups* T_Γ/Γ *admits the following description:*

- *its underlying graph is a rose,*
- *its unique vertex group is* Y*,*
- *each of its edge group coincides with* F *(which is the finite subgroup of* L *defined in the previous lemma), and*
- *the stable letters are the components* $x_1, \ldots, x_p$ *of* $\boldsymbol{x}$*.*

In particular, it follows that Γ *admits a splitting of the form*

$$\Gamma = Y *_F \langle \boldsymbol{x}, F \mid \mathrm{ad}(x_i)|_F = \alpha_i, \forall i \in \{1, \ldots, p\} \rangle,$$

where α_i *denotes the automorphism of* F *induced by the action of the stable letter* x_i *on* F*.*

Proof. Suppose that Γ does not fix a point of T_L and let $T_\Gamma \subset T_L$ be the minimal subtree of Γ. We first prove that each x_i acts hyperbolically on T_Γ. Note that there exists some $1 \leq i \leq p$ such that $\frac{||\varphi_n(x_i)||}{||\varphi_n||}$ does not approach 0 as n goes to infinity. Otherwise, Γ (which is generated by Y and $\boldsymbol{x}$) would be elliptic in T_L. Moreover, by the third condition of Definition III.2.11, $||\varphi_n(x_i)||/||\varphi_n(x_j)||$ tends to a real number $r_{i,j} \in (0, \infty)$ for every $1 \leq i, j \leq p$. Consequently, every x_i is hyperbolic. Denote by ℓ_i the limit of the sequence

$$\left(\frac{||\varphi_n(x_i)||}{||\varphi_n||} \right)_{n \in \mathbb{N}}$$

and note that one has $0 < \ell_i < \infty$.

Next, consider the subset T of T_Γ defined as the union of the axes of all of the hyperbolic elements of Γ. We prove that T is connected, i.e. that T is a subtree of T_Γ. Let p_1 and p_2 be two points of T. By the definition of T, there exist two hyperbolic elements γ_1 and γ_2 of Γ such that p_1 belongs to the axis of γ_1 and p_2 belongs to the axis of γ_2. If these two axes are not disjoint, we are done. Otherwise, if the axes are disjoint, then it is well-known that $\gamma_1\gamma_2$ is hyperbolic and that its axis $A(\gamma_1\gamma_2)$ intersects both $A(\gamma_1)$ and $A(\gamma_2)$. Hence, T is a subtree of T_Γ. In addition, T is Γ-invariant since $\gamma_1 A(\gamma_2) = A(\gamma_1\gamma_2\gamma_1^{-1})$ for every $\gamma_1 \in \Gamma$ and every hyperbolic element $\gamma_2 \in \Gamma$. The minimality of T_Γ also implies that $T = T_\Gamma$. We will prove that T_Γ is a simplicial tree. To do so, we prove a duo of auxiliary claims.

Claim 1: *for every $1 \leq i, j \leq p$ and for every $s \in Y$, if the intersection of the axes of x_j and $s x_i s^{-1}$ contains two distinct points v and w, then $i = j$ and s belongs to F.*

Proof of Claim 1. By assumption, the arc $[v, w]$ is contained in the intersection of the axes of $s x_i s^{-1}$ and x_j. Let η be the length of $[v, w]$ in the limiting tree T_Γ. Let $\overline{s}$ be a preimage of s in G_Σ. The fellow traveling constant $\Delta(\varphi_n(\overline{s} x_i \overline{s}^{-1}), \varphi_n(x_j))$ is close to $\eta \|\varphi_n\|$ ω-almost surely. As a consequence, ω-almost surely, the following inequality holds:

$$\Delta(\varphi_n(\overline{s} x_i \overline{s}^{-1}), \varphi_n(x_j)) \geq \frac{\eta \|\varphi_n\|}{2}.$$

Moreover, the translation length $\|\varphi_n(x_i)\|$ is equivalent to $\ell_i \|\varphi_n\|$. Thus, ω-almost surely, one has that

$$\Delta(\varphi_n(\overline{s} x_i \overline{s}^{-1}), \varphi_n(x_j)) \geq \frac{\eta}{4\ell_i} \|\varphi_n(x_i)\|.$$

According to the fourth condition of Definition III.2.11, the tuple $(\varphi_n(x_1), \ldots, \varphi_n(x_p))$ satisfies the ε_n-small cancellation condition for some sequence of positive real numbers $(\varepsilon_n)_{n \in \mathbb{N}}$ converging to 0; particularly, ω-almost surely, ε_n is smaller than $\eta/(4\ell_i)$. We deduce that i and j are equal, and $\varphi_n(\overline{s})$ is contained in $\Lambda(\varphi_n(x_i)) = E(G) \rtimes \langle \varphi_n(x_i) \rangle$.

Lastly, since s belongs to Y, the translation length of $\varphi_n(\overline{s})$ is negligible compared to that of $\varphi_n(x_i)$ ω-almost surely, so $s = \varphi_\infty(\overline{s})$ belongs to $F = \varphi_\infty(K)$ (where H is the subgroup from Lemma III.2.16). ☐

Claim 2: *for every $1 \leq i \leq p$,*

- *the pointwise stabilizer of the axis $A(x_i)$ of x_i is the finite group F,*
- *the setwise stabilizer of $A(x_i)$ (denoted by Z_i) is equal to $\langle x_i, F \rangle$.*

Moreover, $A(x_i)$ is transverse to its translates, which means that for every $\gamma \in \Gamma \setminus Z_i$, the intersection of $A(x_i)$ and $A(\gamma x_i \gamma^{-1})$ is either empty or a single point.

Proof of Claim 2. Let $\gamma \in \Gamma$ be such that the intersection of $A(x_i)$ and $A(\gamma x_i \gamma^{-1})$ contains two distinct points v and w. Let $\overline{\gamma}$ be a preimage of γ in G_Σ. As in the proof of Claim 1, $\varphi_n(\overline{\gamma})$ is contained in $\Lambda(\varphi_n(x_i)) = E(G) \rtimes \langle \varphi_n(x_i) \rangle$. Hence, for every n, there exists an integer p_n such that $\varphi_n(\overline{\gamma} x_i^{p_n})$ belongs to $E(G)$ ω-almost surely. Note that the sequence $(p_n)_{n \in \mathbb{N}}$ is bounded ω-almost surely, since $\|\varphi_n(x_i)\|/\|\varphi_n\|$ is bounded away from 0 ω-almost surely. In particular, p_n is constant ω-almost surely, equal to some $p \in \mathbb{N}$. It follows that γx_i^p fixes the basepoint o. In other words, γ is equal to sx_i^{-p} for some element $s \in Y$. Now, observe that one has $\gamma x_i \gamma^{-1} = s x_i s^{-1}$. Since the intersection of $A(x_i)$ with $A(\gamma x_i \gamma^{-1})$ is neither empty nor a single point, Claim 1 implies that s belongs to F. This shows that $Z_i = \langle x_i, F \rangle$ and that $A(x_i)$ is transverse to its translates, which completes the proof of Claim 2. □

We are now ready to prove that T_Γ is a simplicial tree. First, we prove that the number of orbits of directions at branch points in T_Γ is finite by using Theorem III.2.5. Note that Γ is countable since it is a subgroup of the finitely generated group L (which is a quotient of G_Σ). We need to check that the two assumptions of Theorem III.2.5 are satisfied, namely that the action of Γ on T_Γ is K-superstable for some constant K (i.e. that every arc whose stabilizer has order greater than K is stable) and that arc stabilizers are finitely generated. In fact, we will prove the following stronger fact: arc stabilizers are contained in the finite group F.

Let $I \subset T_\Gamma$ be an arc. Observe that there exists a subarc $J \subset I$ that is contained in the axis of a hyperbolic element of Γ. Indeed, let p_1, p_2 be two distinct points of I, and let $\gamma_1, \gamma_2 \in \Gamma$ be two hyperbolic elements such that $p_1 \in A(\gamma_1)$ and $p_2 \in A(\gamma_2)$. Since T_Γ is a tree, if $A(\gamma_1)$ and $A(\gamma_2)$ are not disjoint then the arc $[p_1, p_2]$ is contained in the union $A(\gamma_1) \cup A(\gamma_2)$; thus, there exists a subarc $J \subset [p_1, p_2]$ that is contained in $A(\gamma_1)$ or in $A(\gamma_2)$. If $A(\gamma_1)$ and $A(\gamma_2)$ are disjoint, then $\gamma_1\gamma_2$ is hyperbolic and $[p_1, p_2]$ is contained in $A(\gamma_1\gamma_2)$. In this case, one can take $J = [p_1, p_2]$. In either case, there exists a subarc $J \subset I$ that is contained in the axis of a hyperbolic element $\gamma \in \Gamma$.

We next observe that the axis of γ is obtained by concatenating subarcs of the axes of some conjugates of $x_1, \ldots, x_p$, and therefore J contains a subarc J' that is contained in the axis of a conjugate $h x_i h^{-1}$ of x_i for some $1 \leq i \leq p$ and $h \in \Gamma$. Now, let g be an element that fixes I pointwise. In particular g fixes the subarc J' and it follows that J' is contained in $A(h x_i h^{-1})$ and in $gA(h x_i h^{-1})$. Since $A(x_i)$ is transverse to its translates, g belongs to F. Hence, arc stabilizers are contained in the finite group F, and so the assumptions of Theorem III.2.5 are satisfied. We deduce that the number of orbits of directions at branch points in T_Γ is finite.

We next seek to bound the number of branch points in arcs in T_Γ. More specifically, let I be an arc of T_Γ; we will prove that there are only finitely many branch points on I in T_Γ. Without loss of generality one can assume that the length of I is smaller than the translation length ℓ_i of x_i for every $1 \leq i \leq p$. Assume towards a contradiction that there are infinitely many branch points on I. Then, by III.2.5, there exist necessarily two non-degenerate subsegments J_1 and J_2 of I such that

- $J_1 \cap J_2 = \varnothing$,

- there exists some $\gamma \in \Gamma$ such that $\gamma J_1 = J_2$.

As above, up to taking subarcs of J_1 and J_2, one can assume that J_1 and J_2 are contained in the axes of two hyperbolic elements γ_1 and γ_2 respectively. Since the axes of γ_1 and γ_2 are obtained by concatenating subarcs of the axes of some conjugates of $x_1, \dots, x_p$, one can assume (again, up to taking subarcs of J_1 and J_2) that J_1 and J_2 are contained in the axes of two conjugates of x_{i_1} and x_{i_2} for some $1 \leq i_1, i_2 \leq p$. By Claim 1, one has that $i_1 = i_2$ and that γ belongs to $\langle F, x_{i_1} \rangle$ (up to conjugacy). This is a contradiction, since F fixes the axis of x_{i_1} pointwise, and since the intersection of I with $x_{i_1} I$ is empty or reduced to a point (because the length of I is smaller than the translation length ℓ_i of x_i for every $1 \leq i \leq p$). Therefore, there are only finitely many branch points on I, which proves that the tree T_Γ is simplicial.

To finish, we give a precise description of the graph of groups T_Γ / Γ. Since the pointwise stabilizer of the axis of x_i is F for every $1 \leq i \leq p$, and since $\boldsymbol{x}$ and Y generate Γ by definition, it follows that the underlying graph of T_Γ / Γ is a rose in which the unique vertex group is Y, every edge group is equal to F, and the stable letters are $x_1, \dots, x_p$. Therefore, Γ can be decomposed as an amalgamated product as follows:

$$\Gamma = Y \ast_F \langle \boldsymbol{x}, F \mid \mathrm{ad}(x_i)|_F = \alpha_i, \forall i \in \{1, \dots, p\} \rangle,$$

where α_i denotes the automorphism of F induced by the action of x_i on F by conjugation. ☐

Corollary III.2.18. *With the same notations and the same hypotheses as in Lemma III.2.17 above, the tree T_Γ is transverse to its translates, i.e. for every element $h \in L \smallsetminus \Gamma$, the intersection $hT_\Gamma \cap T_\Gamma$ is at most one point. In addition, if e is an edge of T_Γ, there are only finitely many branch points on e in T_L.*

Proof. Let h be an element of L such that $hT_\Gamma \cap T_\Gamma$ is non-degenerate. As a consequence of the description of T_Γ above, we can find two elements $u, v \in \Gamma$ such that the axes of ux_iu^{-1} and $h(vx_jv^{-1})h^{-1}$ have a non-trivial overlap in the limiting tree T_L, for some $1 \leq i, j \leq p$; note that it could be that $i = j$. Let $\overline{u}, \overline{v}$ and $\overline{h}$ be three preimages of u, v, h respectively in G_Σ. As in the proof of Lemma III.2.17, we have that

$$\Delta(\varphi_n(x_i), \varphi_n(\overline{u}^{-1}\overline{h}\overline{v})\varphi_n(x_j)\varphi_n(\overline{u}^{-1}\overline{h}\overline{v})^{-1}) \geq \varepsilon_n \min(\|\varphi_n(x_i)\|, \|\varphi_n(x_j)\|)$$

ω-almost surely. Hence, $i = j$ and $\varphi_n(\overline{u}^{-1}\overline{h}\overline{v})$ belongs to $\Lambda(\varphi_n(x_i)) = E(G) \times \langle \varphi_n(x_i) \rangle$. Thus, for every n, there exists p_n such that $\varphi_n(\overline{u}^{-1}\overline{h}\overline{v}x_i^{p_n})$ lies in $E(G)$. On the other hand, as in the proof of Lemma III.2.17, since x_i acts hyperbolically on T_Γ, p_n is bounded by a constant that does not depend on n ω-almost surely. Otherwise,

$$\lim_{\omega}\left(\frac{\|\varphi_n(x_i)\|}{\|\varphi_n(\overline{u}^{-1}\overline{h}\overline{v})\|}\right) = 0.$$

Since $\varphi_n(\overline{u}^{-1}\overline{h}\overline{v})/\|\varphi_n\|$ is bounded, $\|\varphi_n(x_i)\|/\|\varphi_n\|$ ω-tends to 0, contradicting that x_i is hyperbolic. As a consequence, we can assume that for some p, $p_n = p$ for all n. Thus, $\varphi_n(\overline{u}^{-1}\overline{h}\overline{v}x_i^p)$ belongs to $E(G)$ ω-almost surely, that is, $\overline{u}^{-1}\overline{h}\overline{v}x_i^p$ lies in H and $u^{-1}hvx_i^p$ belongs to $\varphi_\infty(H) = F$. We conclude that there is an element $f \in F$ such that $u^{-1}hvx_i^p = f$, or in other words $h = ufx_i^{-p}v^{-1}$. Note that $h \in \Gamma$ since u, v, f and x_i belong to Γ.

Lastly, let e be an edge of T_Γ. Assume without loss of generality that e is adjacent to the basepoint o (since T_Γ is simplicial). We will prove that there are only finitely many branch points on e in T_L. First, note that the stabilizer of e in L is contained in the finite group F. Indeed, if an element $g \in L$ fixes the edge e, then it fixes the basepoint o, whose stabilizer is $Y \subset \Gamma$. Hence g belongs to the stabilizer of e in Γ, which is equal to F by Lemma III.2.17. This shows that the hypotheses of Theorem III.2.5 are satisfied.

It follows that the number of orbits of directions at each branch point of T_L is finite. Now, assume towards a contradiction that there are infinitely many branch points on e. Then, by III.2.5 there exist necessarily two non-degenerate subsegments I and J of e, with $I \cap J = \varnothing$, and an element $g \in L$ such that $gI = J$. But we just proved that T_Γ is transverse to its translates, so g belongs to Γ. This is a contradiction, since T_Γ is a simplicial tree by Lemma III.2.17. □

III.3 Merzlyakov's theorem and an outline of the proof of Theorem A

The main theorem proved in this chapter is the following generalisation of Merzlyakov's theorem. In this section we state the general version of Merzlyakov's theorem, and subtly reduce it to obtain a slightly easier statement. Note that Theorem E stated in the introduction corresponds to the case where $\ell = 1$ in the result below. We also provide a rough outline of the proofs of Theorems E and A in preparation for the following sections.

Theorem III.3.1 (cf. Theorem E)**.** *Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G (called constants). Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let*

$$\theta(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) : \bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) \neq 1)$$

be a finite disjunction of finite systems of equations and inequations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are two tuples of variables. For every $1 \leq k \leq \ell$, let G_{Σ_k} be the following group, which is finitely presented relative to $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$:

$$G_{\Sigma_k} = \langle \boldsymbol{x},\boldsymbol{y},\boldsymbol{a} \mid R(\boldsymbol{a}) = 1, \ \Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1\rangle.$$

Let $p = |\boldsymbol{x}|$ be the arity of $\boldsymbol{x}$, and let x_i denote the i-th component of $\boldsymbol{x}$. Suppose that G satisfies the following first-order sentence:

$$\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) \neq 1).$$

Then, for every p-tuple $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_p) \in \mathrm{Aut}_G(E(G))^p$, there exist an integer $1 \leq k \leq \ell$ and a morphism

$$\pi_{\boldsymbol{\sigma}} : G_{\Sigma_k} \to G_{\boldsymbol{\sigma}} = G *_{E(G)} \left\langle \boldsymbol{x}, E(G) \mid \mathrm{ad}(x_i)|_{E(G)} = \sigma_i, \ \forall i \in \{1, \ldots, p\}\right\rangle$$

such that the following hold:

1. $\pi_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \boldsymbol{x}$,
2. $\pi_{\boldsymbol{\sigma}}(\boldsymbol{a}) = \boldsymbol{a}$,
3. $\Psi(\boldsymbol{x}, \pi_{\boldsymbol{\sigma}}(\boldsymbol{y}), \boldsymbol{a}) \neq 1$.

Moreover, the image of $\pi_{\boldsymbol{\sigma}}$ is a subgroup of $G_{\boldsymbol{\sigma}}$ of the form

$$\langle \boldsymbol{g}, \boldsymbol{a}\rangle *_{E(G)} \left\langle \boldsymbol{x}, E(G) \mid \mathrm{ad}(x_i)|_{E(G)} = \sigma_i, \ \forall i \in \{1, \ldots, p\}\right\rangle$$

for some tuple $\boldsymbol{g}$ of elements of G.

In fact, as Lemma III.3.3 will show, it is enough to prove the following version of Merzlyakov's theorem, which is *a priori* weaker than Theorem III.3.1 since the group $E(G)$ is replaced with a subgroup $E \subset E(G)$ that may be proper.

Theorem III.3.2. *Keep the assumptions of Theorem III.3.1. Then for every p-tuple* $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_p) \in \mathrm{Aut}_G(E(G))^p$*, there exist an integer* $1 \leq k \leq \ell$*, a finite subgroup* E *of* $E(G)$*, and a morphism*

$$\pi_{\boldsymbol{\sigma}} : G_{\Sigma_k} \to G_{\boldsymbol{\sigma}} = \langle G, \boldsymbol{x} \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E, \ \forall i \in \{1, \ldots, p\} \rangle$$

such that the following hold:

1. $\pi_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \boldsymbol{x}$,
2. $\pi_{\boldsymbol{\sigma}}(\boldsymbol{a}) = \boldsymbol{a}$,
3. $\Psi(\boldsymbol{x}, \pi_{\boldsymbol{\sigma}}(\boldsymbol{y}), \boldsymbol{a}) \neq 1$.

Moreover, the image of $\pi_{\boldsymbol{\sigma}}$ *is a subgroup of* $G_{\boldsymbol{\sigma}}$ *of the form*

$$\langle \boldsymbol{g}, \boldsymbol{a} \rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E, \ \forall i \in \{1, \ldots, p\} \rangle$$

for some tuple $\boldsymbol{g}$ *of elements of* G*.*

Lemma III.3.3. *Theorems III.3.1 and III.3.2 are equivalent.*

Proof. Theorem III.3.2 follows immediately from Theorem III.3.1. Let us prove the converse. The proof consists in slightly modifying the first-order formula

$$\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \theta(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}).$$

Let $\boldsymbol{b}$ be the tuple of elements of G consisting of $\boldsymbol{a}$ and $E(G)$, and let $\boldsymbol{\sigma}_1, \ldots, \boldsymbol{\sigma}_N$ be an enumeration of the elements of $\mathrm{Aut}_G(E(G))^p$. For $1 \leq i \leq N$, let $\mu_i(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{b})$ be the quantifier-free formula which states that "$\theta(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ is true and $\boldsymbol{x}$ acts on $E(G)$ as $\boldsymbol{\sigma}_i$". Since $\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \theta(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ holds in G, the following first-order sentence holds in G as well:

$$\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \bigvee_{i=1}^{N} \mu_i(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{b}).$$

Theorem III.3.1 follows from Theorem III.3.2 applied to this new first-order sentence. □

III.3.1 An outline of the proof of Theorem A

In order to illustrate the main ideas appearing in the following sections, we sketch a proof of Theorem A in the particular (and easier) case where the maximal normal finite subgroup $E(G)$ is trivial.

Suppose that G satisfies a first-order sentence

$$\theta : \forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \Sigma(\boldsymbol{x}, \boldsymbol{y}) = 1 \wedge \Psi(\boldsymbol{x}, \boldsymbol{y}) \neq 1.$$

Let $\Gamma = G * \langle t \rangle \simeq G * \mathbb{Z}$. Observe that the following two assertions are equivalent, where p denotes the arity of $\boldsymbol{x}$.

- Γ satisfies the sentence θ.
- For every p-tuple $\boldsymbol{\gamma}$ of elements from Γ, there exists a retraction r from $\Gamma_{\Sigma,\boldsymbol{\gamma}} = \langle \Gamma, \boldsymbol{y} \mid \Sigma(\boldsymbol{\gamma}, \boldsymbol{y}) = 1 \rangle$ onto Γ such that no $\psi(\boldsymbol{\gamma}, \boldsymbol{y})$ which appears in the conjunction of inequations $\Psi(\boldsymbol{\gamma}, \boldsymbol{y})$ is killed by r, i.e. the inequations remain valid in the image of r.

In order to prove that Γ satisfies the sentence θ, we will construct such a retraction $r : \Gamma_{\Sigma,\boldsymbol{\gamma}} \twoheadrightarrow \Gamma$ for any $\boldsymbol{\gamma} \in \Gamma^p$. The very first step of the construction of this retraction relies on the existence of a quasi-convex free subgroup $F(a,b) \subset G$ (see [39, Theorem 6.14] combined with [5, Lemma 3.1]). Using a sequence of elements $(w_n(a,b))_{n \in \mathbb{N}} \in F(a,b)^{\mathbb{N}}$ satsifying certain small cancellation conditions in the free group $F(a,b)$, we construct a test sequence $(\varphi_n : \Gamma \twoheadrightarrow G)_{n \in \mathbb{N}}$ by setting $\varphi_n|_G = \mathrm{id}_G$ and $\varphi_n(t) = w_n(a,b)$. Since, by assumption, the sentence θ is true in the group G, each morphism φ_n extends to a morphism $\psi_n : \Gamma_{\Sigma,\boldsymbol{\gamma}} \twoheadrightarrow G$ mapping $\boldsymbol{y}$ to a tuple $\boldsymbol{g}_n$ such that $\Sigma(\psi_n(\boldsymbol{\gamma}), \boldsymbol{g}_n) = 1$ and $\Psi(\psi_n(\boldsymbol{\gamma}), \boldsymbol{g}_n) \neq 1$.

The fact that $F(a,b)$ is quasi-isometrically embedded in G enables us to prove that the sequence of elements $(\psi_n(t) = w_n(a,b))_{n \in \mathbb{N}}$ satisfies nice geometric conditions in G, which, in some sense, encapsulate the first-order sentence

$$\exists \boldsymbol{y} \ \Sigma(\boldsymbol{\gamma}, \boldsymbol{y}) = 1 \wedge \Psi(\boldsymbol{\gamma}, \boldsymbol{y}) \neq 1.$$

The sequence $(\psi_n)_{n \in \mathbb{N}}$ yields a divergent limit group L acting on the corresponding tree T. We analyse the action of L on T using the Rips machine, which converts the action $L \curvearrowright T$ into an action of L on a simplicial tree. Using the version of the shortening argument proved in Section III.1.4, we can assure that some properties of the test sequence $(\psi_n)_{n \in \mathbb{N}}$ are reflected in this splitting of L, and the rest of the proof consists of constructing a retraction from L onto Γ, using this splitting.

III.4 Proof of Merzlyakov's theorem III.3.2 in a particular case

In this section, we deal with the case where $(\sigma_1, \ldots, \sigma_p) = (\mathrm{id}_{E(G)}, \ldots, \mathrm{id}_{E(G)})$. More precisely, we prove the following result, which is stated here again for the reader's convenience, and which is a restrictive version of Theorem III.3.2.

Theorem III.4.1. *Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G (called constants). Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let*

$$\bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) \neq 1)$$

be a finite disjunction of finite system of equations and inequations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are two tuples of variables. For every $1 \leq k \leq \ell$, let G_{Σ_k} denote the following group, finitely presented relative to $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$:

$$\langle \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a} \mid R(\boldsymbol{a}) = 1, \ \Sigma_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1\rangle.$$

Let $p = |\boldsymbol{x}|$ be the arity of $\boldsymbol{x}$, and let x_i denote the i-th component of $\boldsymbol{x}$. Suppose that G satisfies the following first-order sentence:

$$\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) \neq 1).$$

Then there exist an integer $1 \leq k \leq \ell$, a finite subgroup E of $E(G)$, and a morphism

$$\pi_{\boldsymbol{\sigma}} : G_{\Sigma_k} \to G_{\boldsymbol{\sigma}} = \langle G, \boldsymbol{x} \mid \mathrm{ad}(x_i)|_E = \mathrm{id}_E, \ \forall i \in \{1, \ldots, p\}\rangle$$

such that the following hold:

- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \boldsymbol{x}$,
- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{a}) = \boldsymbol{a}$,
- $\Psi(\boldsymbol{x}, \pi_{\boldsymbol{\sigma}}(\boldsymbol{y}), \boldsymbol{a}) \neq 1$.

Moreover, the image of $\pi_{\boldsymbol{\sigma}}$ is a subgroup of $G_{\boldsymbol{\sigma}}$ of the form

$$\langle \boldsymbol{g}, \boldsymbol{a}\rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \mathrm{id}_E, \ \forall i \in \{1, \ldots, p\}\rangle$$

for some tuple $\boldsymbol{g}$ of elements of G.

Recall that an $(\mathrm{id}_{E(G)}, \ldots, \mathrm{id}_{E(G)})$-test sequence is simply called a *test sequence*. We commence by constructing a test sequence enjoying two additional special properties.

III.4.1 Construction of a test sequence

The construction relies crucially on the existence of a quasi-isometrically embedded subgroup of G of the form $F(a,b) \times E(G)$ (provided by [39, Theorem 6.14] together with [5, Lemma 3.1]), which will enable us to use small cancellation within $F(a,b)$.

Proposition III.4.2. *Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G. Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let*

$$\bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) \neq 1)$$

be a finite disjunction of finite system of equations and inequations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are two tuples of variables. For every $1 \leq k \leq \ell$, denote

$$G_{\Sigma_k} = \langle \boldsymbol{x},\boldsymbol{y},\boldsymbol{a} \mid R(\boldsymbol{a}) = 1, \ \Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1\rangle.$$

Suppose that G satisfies the following first-order sentence:

$$\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}) \neq 1).$$

Then, there exists $1 \leq k \leq \ell$ and a test sequence $(\varphi_n : G_{\Sigma_k} \to G)_{n\in\mathbb{N}}$ satisfying the following two conditions ω-almost surely:

1. *no $\psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$ which appears in the system of inequations $\Psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$ is killed by φ_n, and*
2. *the morphism φ_n maps $\boldsymbol{a}$ to $\boldsymbol{a}$ (and not only to a conjugate of $\boldsymbol{a}$).*

Proof. By Theorem 6.14 in [39], there exists a hyperbolically embedded subgroup $H \hookrightarrow_h G$ (see [39, Definition 2.1]) such that $H = F(a,b) \times E(G)$ (where $F(a,b)$ denotes the free group on two generators a and b) and the elements a and b are loxodromic in G.

By Theorem I.2.7, there exists a (possibly infinite) generating set S of G such that the Cayley graph of G with respect to S is hyperbolic, and such that the natural action of G on this Cayley graph is non-elementary and acylindrical. By [93, Lemma 5.1], the conclusion of Theorem I.2.7 is still satisfied for $S \cup \{a,b\}$ instead of S. Therefore, we can assume without loss of generality that a and b belong to S. Denote by (X,d) the Cayley graph of G with respect to this enlarged set S.

Let d' denote the metric in the free group $\langle a,b\rangle$ for the generating set $\{a,b\}$. By [5, Lemma 3.1], there exist two constants q and r such that

$$d'(1,h) \leq qd(1,h) + r \tag{III.1}$$

for every $h \in \langle a,b\rangle$.

We next choose a collection of elements in $F(a,b)$ that satisfy a small cancellation condition. For any $1 \leq i \leq p$ (where p is the arity of $\boldsymbol{x}$) and $n \geq 0$, we define

$$g_{i,n} = a^{(i-1)n+1}ba^{(i-1)n+2}b\cdots a^{in}b.$$

Let $\boldsymbol{g}_n$ be the p-tuple $(g_{1,n},\dots,g_{p,n})$.

There exists $1 \leq k \leq \ell$ such that, for infinitely many integers n, there is a tuple $\boldsymbol{h}_n$ of elements of G for which

$$G \vDash \Sigma_k(\boldsymbol{g}_n,\boldsymbol{h}_n,\boldsymbol{a}) = 1 \wedge \Psi_k(\boldsymbol{g}_n,\boldsymbol{h}_n,\boldsymbol{a}) \neq 1.$$

By passing to a subsequence we assume without loss of generality that this system of equalities and inequalities holds for all integers n. We will prove that the sequence $(\varphi_n)_{n\in\mathbb{N}}$ defined by the three equalities below is a test sequence,

$$\varphi_n(\boldsymbol{x}) = \boldsymbol{g}_n, \quad \varphi_n(\boldsymbol{y}) = \boldsymbol{h}_n \quad \text{and} \quad \varphi_n(\boldsymbol{a}) = \boldsymbol{a}.$$

For $1 \leq i \leq p$ and $n \geq 0$, let $\tau_{i,n}$ be the path in X connecting 1 to $g_{i,n}$; assume that $\tau_{i,n}$ is labeled with the word $g_{i,n}$ in a and b. Consider the bi-infinite path $\overline{\tau}_{i,n} = \bigcup_{k\in\mathbb{Z}} g_{i,n}^k\tau_{i,n}$. The path $\overline{\tau}_{i,n}$ is a geodesic in the Cayley graph of $F(a,b)$ (equipped with the metric d'). By the inequality (III.1), this graph is quasi-isometrically embedded into (X,d) (for some constants that do not depend on n). Therefore, $\overline{\tau}_{i,n}$ can also be seen as a quasi-geodesic in (X,d). Consequently, $\overline{\tau}_{i,n}$ lies in the λ-neighbourhood of the quasi-axis $A(g_{i,n})$ of $g_{i,n}$ for some constant $\lambda \geq 0$ (which is independent of n). Similarly, let α be the edge of X that connects 1 to a, let $\overline{\alpha}$ denote the quasi-geodesic $\overline{\alpha} = \bigcup_{k\in\mathbb{Z}} a^k\alpha$ and let μ be a constant such that $\overline{\alpha}$ lies in the μ-neighbourhood of $A(a)$.

Since $g_{i,n}$ is cyclically reduced in $\langle a,b\rangle$, an easy calculation shows that $d'(1,g_{i,n}) \sim (i-1/2)n^2$. Thus, the second and third conditions of Definition III.2.11 hold. In addition, the inequality (III.1) implies that there exists a constant $R > 0$ such that $||g_{i,n}|| \geq Rn^2$ for every n that is sufficiently large and every $i \in \{1,\dots,p\}$.

It remains to prove the fourth condition of Definition III.2.11, namely that each tuple $\varphi_n(\mathbf{x})$ satisfies a suitable geometric small cancellation condition. Since a is loxodromic, there exists $N \geq 0$ such that, for every $g \in G$, if $\Delta(a, gag^{-1}) \geq N$, then g belongs to $\Lambda(a) = \langle a \rangle \times E(G)$ (see Subsection III.2.3). Let n_0 be such that

$$16 \cdot q \cdot R \cdot n_0 \gg N' = N + 204 \cdot \delta + 2 \cdot \lambda + 2 \cdot \mu.$$

We remark that N' was chosen to cater for the use of [37, Lemma 2.13] which gives a bound on the diameter of the intersection of two quasiconvex sets in a hyperbolic space. We remind the reader that

- q is the constant involved in the inequality (III.1),
- R satisfies $\|g_{i,n}\| \geq Rn^2$ for every $1 \leq i \leq p$ ω-almost surely,
- λ was chosen such that $\overline{\tau}_{i,n}$ lies in the λ-neighbourhood of the quasi-axis $A(g_{i,n})$ of $g_{i,n}$, and
- μ was chosen such that $\overline{\alpha}$ lies in the μ-neighbourhood of $A(a)$.

We will show that for every $n \geq n_0$, the tuple $\varphi_n(\boldsymbol{x}) = \boldsymbol{g}_n$ satisfies the $(16q/n)$-small cancellation condition (see Definition III.2.9). Let $n > n_0$ and let $g \in G$ be such that

$$\Delta(g_{i,n}, g \cdot g_{j,n} \cdot g^{-1}) \geq \frac{16q \min(\|g_{i,n}\|, \|g_{j,n}\|)}{n} \tag{III.2}$$

for some $i, j \in \{1, \ldots, p\}$. We need to show that $i = j$ and that g belongs to the subgroup $\langle g_{i,n} \rangle \times E(G)$. For convenience, we assume without loss of generality that $j \geq i$. Thus, ω-almost surely, one has that $\min(\|g_{i,n}\|, \|g_{j,n}\|) = \|g_{i,n}\|$. We first show that g lies in the subgroup $\langle a, b \rangle \times E(G)$. Since

$$\Delta(g_{i,n}, g \cdot g_{j,n} \cdot g^{-1}) \geq \frac{16q \cdot \|g_{i,n}\|}{n} \geq 16 \cdot q \cdot R \cdot n \geq 16 \cdot q \cdot R \cdot n_0 \gg N', \tag{III.3}$$

we can choose two subpaths $\mu_{i,n}$ and $\mu_{j,n}$ of $\overline{\tau}_{i,n}$ and $g\overline{\tau}_{j,n}$ respectively, such that

1. each of $\mu_{i,n}$ and $\mu_{j,n}$ is labeled by $a^{N'}$ (and is of length N'), and
2. $\operatorname{diam}((\mu_{i,n})^{+(100\delta+\lambda)} \cap (\mu_{j,n})^{+(100\delta+\lambda)}) \geq N'$.

Denote the initial points of $\mu_{i,n}$ and $\mu_{j,n}$ by $o_{i,n}$ and $o_{j,n}$ respectively. We have that

$$\operatorname{diam}(o_{i,n}\overline{\alpha}^{+(100\delta+\lambda)} \cap o_{j,n}\overline{\alpha}^{+(100\delta+\lambda)}) \geq N'.$$

It follows that

$$\operatorname{diam}(A(a)^{+(100\delta+\lambda+\mu)} \cap o_{i,n}^{-1} \cdot o_{j,n} \cdot A(a)^{+(100\delta+\lambda+\mu)}) \geq N'.$$

By [37, Lemma 2.13] we have that

$$\begin{aligned}\Delta(a,(o_{i,n}^{-1}\cdot o_{j,n})\cdot a\cdot(o_{i,n}^{-1}\cdot o_{j,n})^{-1}) &\geq \operatorname{diam}\left(A(a)^{+(100\delta+\lambda+\mu)}\cap o_{i,n}^{-1}\cdot o_{j,n}A(a)^{+(100\delta+\lambda+\mu)}\right)\\ &\quad -(204\delta+2\lambda+2\mu)\\ &\geq N'-(204\delta+2\lambda+2\mu)\\ &= N.\end{aligned}$$

It follows that $o_{i,n}^{-1}\cdot o_{j,n}$ belongs to $\Lambda(a)=\langle a\rangle\times E(G)$. Next, observe that since $o_{i,n}$ lies in $\langle a,b\rangle$, it is on the quasi-geodesic $\overline{\tau}_{i,n}$. Similarly, $o_{j,n}$ can be written as $o_{j,n}=gw_{j,n}$ with $w_{j,n}\in\angle a,b\rangle$. It follows that g belongs to the subgroup $\langle a,b\rangle\times E(G)$.

Up to replacing g with gc for some $c\in E(G)$, we can assume that $g\in\langle a,b\rangle$; this does not affect the inequality in (III.3). Indeed, $g\cdot c\cdot g_{j,n}(g\cdot c)^{-1}$ is equal to $g\cdot g_{j,n}g^{-1}$ since $g_{j,n}$ centralizes $E(G)$ as an element of $\langle a,b\rangle$.

Let Y denote the Cayley graph of the free group $\langle a,b\rangle$ (equipped with the metric d'). The following inequality can be easily deduced from the inequalities (III.1) and (III.2):

$$\operatorname{diam}\left((\overline{\tau}_{i,n})^{+(q(100\delta+r)+1)}\cap(g\overline{\tau}_{j,n})^{+(q(100\delta+r)+1)}\right)\geq\frac{16\cdot q\cdot d'(1,g_{i,n})}{2\cdot q\cdot n}=8\cdot\frac{d'(1,g_{i,n})}{n}.$$

Since Y is a tree, this inequality implies that the axes of $g_{i,n}$ and $g\cdot g_{j,n}g^{-1}$ have an overlap of length at least $8d'(1,g_{i,n})/n$. Recall that $d'(1,g_{i,n})\sim(i-1/2)n^2$. Thus, $8d'(1,g_{i,n})/n$ is asymptotically equivalent to $8n(i-1/2)$. Therefore, ω-almost surely, the axes of $g_{i,n}$ and $g\cdot g_{j,n}\cdot g^{-1}$ have an overlap of length at least $4n\cdot(i-1/2)$ in Y.

To conclude, note that $4n(i-1/2)>2ni-2$ and that two distinct cyclic conjugates of $g_{i,n}$ and $g_{j,n}$ can have a mutual initial subword of length at most $2ni-2$ (here we use the assumption that $j>i$). Thus, if the axes of $g_{i,n}$ and $g\cdot g_{j,n}\cdot g^{-1}$ have a common subsegment of length strictly larger than $2ni-2$ in Y, then $g_{i,n}$ and $g\cdot g_{j,n}\cdot g^{-1}$ have the same axis. It follows that $i=j$, and that $g_{i,n}$ and g have a common root. Since $g_{i,n}$ is not a power, we must have that g is a power of $g_{i,n}$, which finishes the proof. □

III.4.2 Proof of Theorem III.4.1 (partial version of Merzlyakov's Theorem)

Theorem III.4.1 (Merzlyakov's theorem where $\sigma=(\mathrm{id}_{E(G)},\ldots,\mathrm{id}_{E(G)})$) is an immediate consequence of Proposition III.4.2 (which proved the existence of a test sequence) and Proposition III.4.3 below (which is essentially Theorem E under the additional assumption that a suitable $(\sigma_1, \dots, \sigma_p)$-test sequence exists), applied with $(\sigma_1, \dots, \sigma_p) = (\mathrm{id}_{E(G)}, \dots, \mathrm{id}_{E(G)})$.

Proposition III.4.3. *Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G. Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1 \rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let*

$$\Sigma(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1 \ \wedge \ \Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) \neq 1$$

be a finite system of equations and inequations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are tuples of variables. Suppose that there exists a $(\sigma_1, \dots, \sigma_p)$-test sequence $(\varphi_n : G_\Sigma \to G)_{n \in \mathbb{N}}$ satisfying the following two conditions ω-almost surely:

1. *no no $\psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ which appears in the system of inequations the system of inequations $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ is killed by φ_n, and*
2. *the morphism φ_n maps $\boldsymbol{a}$ to $\boldsymbol{a}$ (not only to a conjugate).*

Then there exist a finite subgroup E of $E(G)$ and a morphism

$$\pi_{\boldsymbol{\sigma}} : G_\Sigma \to G_{\boldsymbol{\sigma}} = \langle G, \boldsymbol{x} \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E, \ \forall i \in \{1, \dots, p\} \rangle$$

such that the following hold:

- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \boldsymbol{x}$,
- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{a}) = \boldsymbol{a}$,
- *no $\psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ which appears in the system of inequations $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ is killed by π_σ.*

Moreover, the image of $\pi_{\boldsymbol{\sigma}}$ is a subgroup of $G_{\boldsymbol{\sigma}}$ of the form

$$\langle \boldsymbol{g}, \boldsymbol{a} \rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E, \ \forall i \in \{1, \dots, p\} \rangle$$

for some tuple $\boldsymbol{g}$ of elements of G.

We remark that the proof of Proposition III.4.3 is quite intricate, and will therefore be interrupted by supplementary lemmas.

Proof. By assumption, there exists a $(\sigma_1, \dots, \sigma_p)$-test sequence $(\varphi_n : G_\Sigma \to G)_{n \in \mathbb{N}}$ that satisfies the following two conditions ω-almost surely:

1. for each $\psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ which appears in the system of inequations $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$, $\varphi_n(\psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}))$ is non-trivial, and

2. the morphism φ_n maps $\boldsymbol{a}$ to $\boldsymbol{a}$ (not only to a conjugate).

Let U be the subgroup of G_Σ generated by $\boldsymbol{x}$ and $\boldsymbol{a}$. Since the Cayley graph of G with respect to S (on which G acts acylindrically and non-elementarily) is discrete, the length of any morphism $G_\Sigma \to G$ is a natural number. As a consequence, there exists a sequence of morphisms $(\theta_n : G_\Sigma \to G)_{n\in\mathbb{N}}$ that satisfies simultaneously the following three conditions ω-almost surely:

1. θ_n coincides with φ_n on U up to conjugation,
2. for each $\psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$ that appears in the system of inequations $\Psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$, $\theta_n(\psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a}))$ is non-trivial, and
3. there is no morphism that satisfies simultaneously the conditions 1. and 2. above, and that is stricly shorter than θ_n.

Note that by Remark III.2.13, the sequence $(\theta_n)_{n\in\mathbb{N}}$ is a $(\sigma'_1,\ldots,\sigma'_p)$-test sequence for some $(\sigma'_1,\ldots,\sigma'_p) \in \mathrm{Aut}_G(E(G))^p$; therefore, the sequence $(\theta_n)_{n\in\mathbb{N}}$ is divergent (by Remark III.2.12). However, one cannot guarantee that σ'_i coincides with σ_i. Moreover, θ_n maps $\boldsymbol{a}$ to a conjugate of $\boldsymbol{a}$, and not necessarily to $\boldsymbol{a}$ itself.

Let $L = G_\Sigma/\underleftarrow{\ker}_\omega((\theta_n)_{n\in\mathbb{N}})$ and let $\theta_\infty : G_\Sigma \twoheadrightarrow L$ be the corresponding limiting map. Observe that θ_∞ is injective on $U = \langle \boldsymbol{a}, x_1, \ldots, x_p\rangle$: indeed, by Lemma III.2.17, the limiting tree associated with the sequence $(\theta_n|_U)_{n\in\mathbb{N}}$ coincides with the Bass-Serre tree of the decomposition of U obtained in Lemma III.2.17. Recall that the underlying graph of the graph of groups decomposition of U is a rose, that the unique vertex group of this decomposition is $\langle \boldsymbol{a}\rangle$ and that the edges groups are all equal to F, with stable letters $x_1,\ldots,x_p$. It follows that the sequence $(\theta_n|_U)_{n\in\mathbb{N}}$ is discriminating (that is, for every $u \in U$, $\theta_n(u) \neq 1$ ω-almost surely) for the following two reasons:

1. θ_n is injective on the vertex group $\langle \boldsymbol{a}\rangle$, and
2. every element that does not belong to a conjugate of $\langle \boldsymbol{a}\rangle$ is hyperbolic in the limiting tree.

Therefore θ_∞ is injective on U. Consequently, in the proof below, we abuse notation and use U to denote its image in the successive quotients of G_Σ involved in the construction of the formal solution $\pi_{\boldsymbol{\sigma}}$.

Standing Assumption III.4.4. *For the rest of the proof, C denotes the constant defined in the Stability Lemma III.1.4.*

A particular case. For presentational purposes, we first prove Proposition III.4.3 in the particular case where L does not split non-trivially over a finite group of order less than C. Under this assumption, if one assumes (towards a contradiction) that U is elliptic in the limiting tree associated to the test sequence $(\theta_n : G_\Sigma \to G)_{n\in\mathbb{N}}$, then the relative Shortening Argument (Theorem III.1.26) implies the existence of a sequence of homomorphisms $(\rho_n : G_\Sigma \to G)_{n\in\mathbb{N}}$ satisfying the following three conditions ω-almost surely:

1. ρ_n coincides with θ_n (and therefore with φ_n) on U up to conjugation,
2. ρ_n kills no $\psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$ which appears in the tuple $\Psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$,
3. and ρ_n is stricly shorter than θ_n relative to H.

This contradicts the definition of θ_n as the shortest morphism satisfying both conditions (1) and (2) ω-almost surely. Hence, U is not elliptic in the limiting tree of the test sequence $(\theta_n : G_\Sigma \to G)_{n\in\mathbb{N}}$. The conclusion now follows from the following technical lemma (see Lemma III.4.8 for a more general version). We postpone the proof (of both lemmas) to the end of this section, so as not to disrupt the proof of Proposition III.4.3.

Lemma III.4.5. *Let F be the finite subgroup of L defined in Lemma III.2.16. If U is not elliptic in the limiting tree, then the group L admits a splitting $\mathbb{S}_L$ with two vertex groups, $\langle \boldsymbol{\ell},\boldsymbol{a}\rangle$ (for some tuple $\boldsymbol{\ell}$ of elements of L) and $\langle \boldsymbol{x},F\rangle$, and one edge group F. Let A be an $\mathbb{S}_L$-approximation of L as in Proposition III.1.11. There exist a finite subgroup E of $E(G)$ and an epimorphism r from A onto a group of the form*

$$\langle \boldsymbol{g},\boldsymbol{a}\rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E,\ \forall i\in\{1,\ldots,p\}\rangle,$$

where $\boldsymbol{g}$ denotes a tuple of elements of G, and such that $r(\boldsymbol{x})=\boldsymbol{x}$, $r(\boldsymbol{a})=\boldsymbol{a}$ and r does not kill any element appearing in the image of the tuple $\Psi(\boldsymbol{x},\boldsymbol{y},\boldsymbol{a})$ in A.

Lastly, one defines the formal solution $\pi_{\boldsymbol{\sigma}} : G_\Sigma \to G$ by $\pi_{\boldsymbol{\sigma}} = r\circ q$ where q denotes the natural epimorphism from G_Σ onto A. This concludes the proof of Proposition III.4.3 in the particular case where L does not split non-trivially over a finite group of order less than C. In general, however, this hypothesis is not satisfied and one has to deal with the complications that arise from splittings over finite subgroups. In particular, one needs a strengthened version of the relative Shortening Argument Theorem III.1.26, namely Theorem III.1.29.

General case. Since we are going to describe an iterative process, let us rename θ_n to θ_n^0, and L to L_0. For any G-limit group L_i that appears in the proof, denote by L_i^U the vertex group containing U in a reduced JSJ decomposition $\mathbb{J}_i$ of L_i, relative to U, over finite groups of order less than C.

The proof of Proposition III.4.3 consists of constructing the following diagrams (indexed by n), illustrated in Figure III.1, ω-almost surely.

Remark III.4.6. **A word of caution**: not all parts of these diagrams commute as it might appear. That is, θ_n^i is not equal to $\theta_n^{i+1} \circ q_{i+1}$, but there are other maps $\rho_n^i : A_i \to G$ defined below such that $\theta_n^i = \rho_n^{i+1} \circ q_{i+1}$.

$$\begin{tikzcd}
A_0 = G_\Sigma \arrow[rrrrr, "\theta_n^0"] \arrow[dr, "q_1"] \arrow[ddddd, two heads, "\theta_\infty^0"] & & & & & G \\
& A_1 \arrow[urrrr, "\theta_n^1"] \arrow[dr, "q_2"] \arrow[ddddl, two heads, "\pi_0"] \arrow[dddd, two heads, "\theta_\infty^1"] & & & & \\
& & A_2 \arrow[uurrr, "\theta_n^2"] \arrow[dddl, two heads, "\pi_1"] \arrow[dr, dashed, "q_i \circ \cdots \circ q_3"] & & & \\
& & & A_i \arrow[uuurr, "\theta_n^i"] \arrow[dr, "q_{i+1}"] \arrow[dd, two heads, "\theta_\infty^i"] & & \\
& & & & A_{i+1} \arrow[uuuur, "\theta_n^{i+1}"] \arrow[dl, "\pi_i"] \arrow[d, two heads, "\theta_\infty^{i+1}"] & \\
L_0 & L_1 & & L_i & L_{i+1} &
\end{tikzcd}$$

Figure III.1: Constructing the maps θ_n^i by iteratively shortening approximations of splittings of the resulting limit groups. The process eventually terminates with U being non-elliptic in the corresponding limiting tree, yielding the desired formal solution.

These diagrams are built iteratively, as follows: given the divergent sequence $(\theta_n^i : A_i \to G)_{n\in\mathbb{N}}$, one defines L_i by $L_i = A_i / \underleftarrow{\ker}_\omega((\theta_n^i)_{n\in\mathbb{N}})$. Let $\mathbb{J}_i$ be a reduced JSJ spltting of L_i over finite groups of order less than C, let $\mathbb{R}_i$ be the splitting of the vertex group L_i^U as a graph of actions outputted by the Rips machine and let $\mathbb{RJ}_i$ be the splitting of L obtained from $\mathbb{J}_i$ by replacing the vertex fixed by L_i^U with the graph of groups $\mathbb{R}_i$. Let A_{i+1} be an $\mathbb{RJ}_i$-approximation of L_i given by Proposition III.1.11

and Corollary III.1.15, and let $\rho_n^{i+1} : A_{i+1} \to G$ be the factorization of $\theta_n^i : A_i \to G$ through the natural epimorphism $q_{i+1} : A_i \twoheadrightarrow A_{i+1}$.

Since A_{i+1} is an $\mathbb{R}\mathbb{J}_i$-approximation of L_i, it is also a $\mathbb{J}_i$-approximation of L_i (indeed, one can collapse the subgraph corresponding to $\mathbb{R}_i$ to a point). We denote by A_{i+1}^U the vertex group of the splitting of A_{i+1} corresponding to L_i^U. Note that A_{i+1}^U, unlike L_i^U, may split non-trivially relative to U over finite subgroups of order less than C. It remains to define the divergent sequence $(\theta_n^{i+1} : A_{i+1} \to G)_{n\in\mathbb{N}}$.

If U is elliptic in the limiting tree of the sequence $(\rho_n^{i+1})_{n\in\mathbb{N}}$, then by the generalised relative Shortening Argument (Theorem III.1.29) there exists a sequence of homomorphisms $(\theta_n^{i+1} : A_{i+1} \to G)_{n\in\mathbb{N}}$ satisfying the following three conditions ω-almost surely:

1. θ_n^{i+1} coincides with ρ_n^{i+1} (and therefore with φ_n) on U up to conjugation,
2. θ_n^{i+1} kills no element appearing in the image of the tuple $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A_{i+1}, and
3. the restriction $\theta_n^{i+1}|_{A_{i+1}^U}$ is stricly shorter than the restriction $\rho_n^{i+1}|_{A_{i+1}^U}$, relative to U.

In addition, since the length of θ_n^{i+1} is a natural number, one can assume without loss of generality that θ_n^{i+1} is the shortest morphism from A_{i+1} to G that satisfies the first two conditions above ω-almost surely. Note that the sequence $(\theta_n^{i+1})_{n\in\mathbb{N}}$ is divergent since it is a test sequence (see Remarks III.2.12 and III.2.13).

To show that this iteration eventually terminates, we have to prove that there exists some $i \in \mathbb{N}$ such that U is not elliptic in the limiting tree of the sequence $(\rho_n^{i+1})_{n\in\mathbb{N}}$. In other words, we shall prove the following claim:

Claim III.4.7. *There exists an integer i such that $q_{i+1}(A_i^U) = A_{i+1}^U$.*

Before proving Claim III.4.7 (whose proof, as one might expect, relies on accessibility arguments), we first explain how to complete the proof of Proposition III.4.3. First, note that if $q_{i+1}(A_i^U)$ is equal to A_{i+1}^U, then shortening the restriction of θ_n^{i+1} to A_{i+1}^U automatically shortens the restriction of ρ_n^i to A_i^U. This is not possible by the definition of ρ_n^i. As a consequence, if $q_{i+1}(A_i^U) = A_{i+1}^U$, then U cannot be elliptic in the limiting tree of the sequence $(\theta_n^{i+1})_{n\in\mathbb{N}}$ by the generalised relative Shortening Argument (Theorem III.1.29). In order to construct the formal solution, we will use the following lemma (the generalised version of Lemma III.4.5) (We remind the reader that the proof of Lemma III.4.8 is postponed to the end of this section)

Lemma III.4.8. *Let F be the finite subgroup of L_{i+1} defined in Lemma III.2.16. If U is non-elliptic in the limiting tree of the sequence $(\theta_n^{i+1})_{n\in\mathbb{N}}$, then the group L_{i+1} admits a splitting $\mathbb{S}$ with two vertex groups, $\langle \boldsymbol{\ell}, \boldsymbol{a}\rangle$ (for some tuple $\boldsymbol{\ell}$ of elements of L_{i+1}) and $\langle \boldsymbol{x}, F\rangle$, and one edge group F. Let A be an $\mathbb{S}$-approximation of L_{i+1} given by Proposition III.1.11. There exists a finite subgroup E of $E(G)$ and an epimorphism r from A onto a group of the form*

$$\langle \boldsymbol{g}, \boldsymbol{a}\rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E,\ \forall i \in \{1,\ldots,p\}\rangle,$$

where $\boldsymbol{g}$ denotes a tuple of elements of G, and such that $r(\boldsymbol{x}) = \boldsymbol{x}$, $r(\boldsymbol{a}) = \boldsymbol{a}$ and r kills no element appearing in the image of the tuple $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A.

As a consequence of this lemma, if U is not elliptic in the limiting tree of the sequence $(\rho_n^{i+1})_{n\in\mathbb{N}}$, one can define the formal solution $\pi_{\boldsymbol{\sigma}} : G_\Sigma \to G$ by

$$\pi_{\boldsymbol{\sigma}} = r \circ q_{i+1} \circ q_i \circ \cdots \circ q_0.$$

Therefore, in order to conclude the proof of Proposition III.4.3, it remains to prove Claim III.4.7 (according to which there exists an integer i such that $q_{i+1}(A_i^U) = A_{i+1}^U$). Denote by η_i the number of edges in a reduced JSJ splitting $\mathbb{J}_i$ of L_i over finite groups of order less than C, relative to U. Let $E(\mathbb{J}_i)$ be the set of edges of $\mathbb{J}_i$. We first prove the following simple lemma, required for the proof of Claim III.4.7:

Lemma III.4.9. *Let G and H be two groups, equipped with two splittings $\mathbb{S}_G$ and $\mathbb{S}_H$ over finite groups. Let T_G and T_H be the Bass-Serre trees corresponding to these splittings. Suppose that there exists an epimorphism $\theta : G \twoheadrightarrow H$ and a θ-equivariant bijection $f : T_G \to T_H$ such that θ is injective on finite vertex groups and maps infinite vertex groups onto infinite vertex groups. Then, the following implication holds: if T_G is reduced, then T_H is reduced.*

Proof of Lemma III.4.9. Suppose that T_G is reduced. Let $\varepsilon = [v, w]$ be an edge of T_H such that $H_v = H_\varepsilon = H_w$. We need to prove that w is a translate of v, i.e. that there exists an element $h \in H$ such that $w = hv$. Let $e = [x, y]$ be a preimage of ε by f. Since H_ε is finite, H_v and H_w are finite. Therefore, G_x and G_y are finite (indeed, we assume that θ maps infinite vertex groups onto infinite vertex groups). Moreover, since θ is injective on finite vertex groups, we have that $G_x = G_e = G_y$. It follows that $y = gx$ for some $g \in G$. In addition, $f(y) = w$, and, since f is θ-equivariant, $f(gx) = \theta(g)f(x) = \theta(g)v$. Hence, $w = \theta(g)v$. □

Before proving Claim III.4.7, we also make the following observation:

Observation: Note that L_i is finitely generated as a quotient of the finitely generated group G_Σ. Using a folding sequence argument, Dunwoody proved in [43] that the sum

$$\sum_{e \in E(\mathbb{J}_i)} \frac{1}{|L_{ie}|},$$

where L_{ie} denotes the edge group of e, is smaller than the rank of L_i (that is, the minimal number of elements required to generate L_i). Therefore, for every $i \in \mathbb{N}$, one has that $\eta_i \leq C\mathrm{rank}(L_i)$. In addition, one has that $\mathrm{rank}(L_i) \leq \mathrm{rank}(G_\Sigma)$ since L_i is a quotient of G_Σ. Putting these inequalities together, we obtain that η_i is bounded from above by $C \cdot \mathrm{rank}(G_\Sigma)$.

Proof of Claim III.4.7. We begin by showing that $\eta_{i+1} \geq \eta_i$, with equality if and only if $q_{i+1}(A_i^U) = A_{i+1}^U$. Since A_{i+1} is a $\mathbb{J}_i$-approximation of L_i, there exists a splitting $\mathbb{J}_i'$ of A_{i+1} with the same underlying graph as $\mathbb{J}_i$, and whose edge groups have the same order as the corresponding edge groups in $\mathbb{J}_i$. In particular, $\mathbb{J}_i'$ is a splitting over finite groups of order less than C, whose underlying graph has η_i edges. Moreover, by Proposition III.1.11 and Lemma III.4.9, the splitting $\mathbb{J}_i'$ is reduced since $\mathbb{J}_i$ is reduced.

In order to establish that $\eta_{i+1} \geq \eta_i$, we take a closer look at the defining sequence $(\theta_n^{i+1} : A_{i+1} \to G)_{n \in \mathbb{N}}$ of $L_{i+1} = A_{i+1}/\underleftarrow{\ker}_\omega((\rho_n^{i+1})_{n \in \mathbb{N}}$. In the proof of Theorem III.1.29, each morphism θ_n^{i+1} is obtained by precomposing ρ_n^{i+1} by an automorphism α of A_{i+1} (note that α is independent of n). The restriction of α to A_{i+1}^U is a modular automorphism (or, more precisely, a lift of a modular automorphism of L_i^U to A_{i+1}^U); furthermore, the restriction of α to any other vertex group of $\mathbb{J}_i'$ is a conjugation. We remark that α is a lift of a natural extension of a modular automorphism of L_i^U (see Lemma III.1.20), and this extension exists because modular automorphisms restrict to conjugation on finite subgroups. As a consequence, L_{i+1} admits a splitting $\mathbb{J}_i''$ over finite groups of order less than C with η_i edge groups. This splitting obtained from the splitting $\mathbb{J}_i'$ of A_{i+1} by replacing each vertex group by its image under the limiting map θ_∞^{i+1}. This shows that a reduced JSJ splitting of L_{i+1} has at least η_i edges, or in other words, that $\eta_{i+1} \geq \eta_i$.

Next, suppose that $\eta_i = \eta_{i+1}$. We will prove that $\mathbb{J}_i''$ is a reduced JSJ splitting of L_{i+1} over finite groups of order less than C. Since we already know that $\mathbb{J}_i''$ is a splitting of L_{i+1} over finite groups of order less than C with η_{i+1} edges, we only need to show that $\mathbb{J}_i''$ is reduced. To do so, it suffices to verify that the conditions of Lemma III.4.9 are satisfied. The manner in which $\mathbb{J}_i''$ was defined implies the limiting map

$\theta_\infty^{i+1} : A_{i+1} \twoheadrightarrow L_{i+1}$ maps each vertex group of $\mathbb{J}_i'$ onto the corresponding vertex group of $\mathbb{J}_i''$. We will verify that the following two hold:

1. θ_∞^{i+1} is injective on finite vertex groups, and
2. θ_∞^{i+1} maps infinite vertex groups onto infinite vertex groups.

Consider the following diagram, where π_i denotes the natural epimorphism from A_{i+1} onto L_i (in other words, π_i is the limiting map ρ_∞^{i+1}):

$$\begin{tikzcd}
 & A_{i+1} \arrow[dl, two heads, "\pi_i"'] \arrow[dr, two heads, "\theta_\infty^{i+1}"] & \\
L_i & & L_{i+1}.
\end{tikzcd}$$

Note that for each vertex group $V \leq A_{i+1}$ of $\mathbb{J}_i'$ that does not contain U, the kernel of the restriction of π_i to V coincides with the kernel of the restriction of θ_∞^{i+1} to V. Indeed, recall that θ_n^{i+1} is obtained by precomposing ρ_n^{i+1} by an automorphism α of A_{i+1} whose restriction to V is a conjugation. As a consequence, the vertex groups $\pi_i(V)$ and $\theta_\infty^{i+1}(V)$ are isomorphic. We know that $\pi_i(V)$ is infinite if and only if V is infinite, and that $\pi_i(V) \simeq V$ if V is finite; this follows from the construction of $\mathbb{J}_i'$ and A_{i+1} (see also Proposition III.1.11 and Corollary III.1.14). Therefore, $\theta_\infty^{i+1}(V)$ is infinite if and only if V is infinite. In addition, $\theta_\infty^{i+1}(V)$ and V are isomorphic whenever V is finite. Finally, note that the image of the vertex group of $\mathbb{J}_i'$ containing U under θ_∞^{i+1} is infinite since U is infinite (and since θ_∞^{i+1} is injective on U). Hence, the conditions 1. and 2. above are satisfied, Lemma III.4.9 applies and $\mathbb{J}_i''$ is reduced, as required.

Hence, if $\eta_i = \eta_{i+1}$, then $\mathbb{J}_i''$ is a reduced JSJ splitting of L_{i+1} over finite groups of order less than C. It follows that the image of the vertex group A_{i+1}^U in L_{i+1} coincides with L_{i+1}^U and $A_{i+1}^U = q_{i+1}(A_i^U)$. ☐

It remains to prove Lemma III.4.8, whose statement is recalled below (for the sake of readability, the index $i+1$ is replaced with i).

Lemma. *Let F be the finite subgroup of L_i defined in Lemma III.2.16. If U is non-elliptic in the limiting tree of the sequence $(\theta_n^i)_{n\in\mathbb{N}}$, then the group L_i admits a splitting $\mathbb{S}$ with two vertex groups, $\langle \boldsymbol{\ell}, \boldsymbol{a}\rangle$ (for some tuple $\boldsymbol{\ell}$ of elements of L_i) and $\langle \boldsymbol{x}, F\rangle$, and one edge group F. Let A be an $\mathbb{S}$-approximation of L_i given by Proposition III.1.11.*

There exists a finite subgroup E of $E(G)$ and an epimorphism r from A onto a group of the form

$$\langle \boldsymbol{g}, \boldsymbol{a} \rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E, \ \forall i \in \{1, \ldots, p\} \rangle,$$

where $\boldsymbol{g}$ denotes a tuple of elements of G, and such that $r(\boldsymbol{x}) = \boldsymbol{x}$, $r(\boldsymbol{a}) = \boldsymbol{a}$ and r kills no element appearing in the image of the tuple $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A.

Proof. By assumption, the group U is non-elliptic in the limiting tree $T := T_{L_i^U}$ associated with the divergent sequence $(\theta_n^i|_{A_i^U})_{n\in\mathbb{N}}$. First, we construct a splitting $\mathbb{S}$ of L_i with exactly two vertex groups $\langle \boldsymbol{x}, F \rangle$ and $\langle \boldsymbol{\ell}, \boldsymbol{a} \rangle$ (for some tuple $\boldsymbol{\ell}$ of elements of L_i) and one edge group F. Let Y be the stabilizer of the base point o in T. Let Γ be the subgroup $\langle U, Y \rangle$ of L_i^U. Since U is contained in Γ, this group is non-elliptic in the limiting tree T. Let $T_\Gamma \subset T$ be the minimal invariant subtree of Γ. By Lemma III.2.17, the tree T_Γ is simplicial and Γ admits the following splitting:

$$\Gamma = Y *_F \langle \boldsymbol{x}, F \mid \mathrm{ad}(x_i)|_F = \alpha_i, \forall i \in \{1, \ldots, p\} \rangle,$$

where F denotes the finite subgroup of L defined in Lemma III.2.16 and α_i denotes the automorphism of F induced by the action of x_i.

Define an equivalence relation $\sim$ on T by setting $x \sim y$ if $[x, y] \cap uT_\Gamma$ contains at most one point for every $u \in L_i^U$. Let $(Y_j)_{j\in J}$ denote the equivalence classes of $\sim$ that are not a single point. Each Y_j is a subtree of T. We next show that $(Y_j)_{j\in J} \cup \{uT_\Gamma \mid u \in L_i^U/\Gamma\}$ is a transverse covering of T (recall Definition III.2.1):

- *Transverse intersection.* For every $i \neq j$, the intersection $Y_i \cap Y_j$ is clearly empty. For every i and $u \in L_i^U$, $Y_i \cap uT_\Gamma$ contains at most one point by definition. For every $u, u' \in L_i^U$ such that $u'u^{-1} \notin \Gamma$, $|uT_\Gamma \cap u'T_\Gamma| \leq 1$ by to Lemma III.2.18.
- *Finiteness condition.* Let x and y be two points of T. By Lemma III.2.18, there exists a constant $\varepsilon > 0$ such that, for every $u \in U$, if the intersection $[x, y] \cap uT_\Gamma$ is non-degenerate, then the length of $[x, y] \cap uT_\Gamma$ is bounded from below by ε. Consequently, the arc $[x, y]$ is covered by at most $\lfloor d(x, y)/\varepsilon \rfloor$ translates of T_Γ and at most $\lfloor d(x, y)/\varepsilon \rfloor + 1$ distinct subtrees Y_j.

Hence, the collection $(Y_j)_{j\in J} \cup \{uT_\Gamma \mid u \in L_i^U\}$ is a transverse covering of T. We next consider the skeleton T_c associated to the transverse covering $(Y_j)_{j\in J} \cup \{uT_\Gamma \mid u \in L_i^U\}$ (see Definition III.2.2). Since the action of L_i^U on T is minimal (from the definition of T), [61, Lemma 4.9] implies that the same holds for the action of L_i^U on T_c.

We will now focus our attention on understanding the decomposition $\Delta_c = T_c/L_i^U$ of L_i^U as a graph of groups. We begin by describing the stabilizer (in L_i^U) of an edge e of T_Γ. Let u be an element of L_i^U that fixes e; e is contained in $T_\Gamma \cap uT_\Gamma$, so u belongs to Γ by Lemma III.2.18. It follows that u belongs to F, because the stabilizer of e in Γ is contained in F (indeed, recall that by Lemma III.2.17, T_Γ is isometric to the Bass-Serre tree of the rose-like splitting of Γ in which every edge group is F). Thus, the stabilizer of e in L_i^U is equal to F.

We next prove that if one of the subtrees in the transverse covering $(Y_j)_{j\in J} \cup \{uT_\Gamma \mid u \in L_i^U\}$ (other than T_Γ) intersects T_Γ in a point, then this point is necessarily one of the terminal vertices of a translate of the edge $e \in T_\Gamma$. Assume towards a contradiction that Y_j or uT_Γ with $u \notin \Gamma$ intersects T_Γ in a point x that is not one of the extremities of e. Then the skeleton T_c contains an edge $\varepsilon = (x, T_\Gamma)$ whose stabilizer is $\mathrm{Stab}(x) \cap \Gamma$ (where $\mathrm{Stab}(x)$ denotes the stabilizer of x in L_i^U) By the previous paragraph, $\mathrm{Stab}(x) \cap \Gamma = F$. Recall that φ_n maps F into $E(G)$ ω-almost surely, and therefore $|F| \leq C$. It follows that the splitting Δ_c of L_i^U is a non-trivial splitting over finite groups of order $\leq C$ relative to Γ. This is impossible since L_i^U does not split over a finite subgroup of order $\leq C$ non-trivially relative to Γ. Hence, if $Y_j \cap T_\Gamma = \{x\}$ or $uT_\Gamma \cap T_\Gamma = \{x\}$ (with $u \notin \Gamma$), then the point x is a terminal vertex of e in T_Γ. As a consequence, $\mathrm{Stab}(x)$ is a conjugate of Y (the stabilizer of the basepoint o) in Γ; furthermore, every edge adjacent to T_Γ in T_c admits the form $(\gamma x, T_\Gamma) = \gamma\varepsilon$ with $\varepsilon = (x, T_\Gamma)$.

Therefore, ε is the only edge adjacent to T_Γ in the quotient graph Δ_c. Its stabilizer is Y. By collapsing all edges of Δ_c except for ε, one obtains a splitting $\Gamma *_Y H$ of L_i^U (where H is some subgroup of L_i^U). Recall that

$$\Gamma = Y *_F \langle \boldsymbol{x}, F \mid \mathrm{ad}(x_i)|_F = \alpha_i, \forall i \in \{1, \ldots, p\}\rangle,$$

and hence the splitting $\Gamma *_Y H$ of L_i^U yields the following splitting of L_i^U:

$$L_i^U = H *_F \langle \boldsymbol{x}, F \mid \mathrm{ad}(x_i)|_F = \alpha_i, \ \forall i \in \{1, \ldots, p\}\rangle.$$

Since every finite subgroup of L_i^U is conjugate to a finite subgroup of H, the group L_i splits as

$$L_i = K *_F \langle \boldsymbol{x}, F \mid \mathrm{ad}(x_i)|_F = \alpha_i, \ \forall i \in \{1, \ldots, p\}\rangle,$$

for some subgroup K of L_i with $\langle \boldsymbol{a}\rangle \subset Y \subset H \subset K$. Denote this splitting of L_i by $\mathbb{S}$. Note that if $\boldsymbol{a}$ was empty, one could just retract L_i onto the free group $F(\boldsymbol{x})$ on $\boldsymbol{x}$.

But $\boldsymbol{a}$ is not empty in general, which makes the construction of the retraction a little bit more involved.

Let A be an $\mathbb{S}$-approximation of L_i (given by Proposition III.1.11), and let $\mathbb{S}_A$ be the corresponding splitting of A. By Remark III.1.12, we can assume that the elements appearing in the tuple $\Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in L_i and A have the same normal forms when written in $\mathbb{S}$ and $\mathbb{S}_A$. The splitting $\mathbb{S}_A$ has a similar form to that of $\mathbb{S}$; more precisely, we have that

$$A = K' *_F \langle \boldsymbol{x}, F \rangle = \langle K', \boldsymbol{x} \mid \mathrm{ad}(x_i)|_F = \alpha_i, \ \forall i \in \{1, \ldots, p\} \rangle.$$

Note that we abuse notation and use F to denote a preimage of $F \subset L_i$ in A. We also use $\boldsymbol{x}$ and $\boldsymbol{a}$ to denote preimages of $\boldsymbol{x}$ and $\boldsymbol{a}$ in A (we also remark that $\boldsymbol{a}$ is contained in K').

We claim that there exists a subgroup E of $E(G)$ and an epimorphism r from A onto a group of the form

$$\langle \boldsymbol{g}, \boldsymbol{a} \rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i|_E, \ \forall i \in \{1, \ldots, p\} \rangle,$$

where $\boldsymbol{g}$ denotes a tuple of elements of G, and such that $r(\boldsymbol{x}) = \boldsymbol{x}$, $r(\boldsymbol{a}) = \boldsymbol{a}$ and r kills no element appearing in the image of the tuple $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A.

For every n, denote by $\psi_n : A \to G$ the factorization of $\theta_n^i : A_i \to G$ through the natural epimorphism from A_i onto A. $\psi_n : A \to G$ is better described by the following commuting diagram:

$$\begin{tikzcd}
A_i \arrow[rr, "\theta_n^i"] \arrow[rd, two heads] \arrow[rdd, two heads, "\theta_\infty^i"'] & & G \\
& A \arrow[ru, "\psi_n"'] \arrow[d, two heads] & \\
& L_i &
\end{tikzcd}$$

We remind the reader that we denote by $\boldsymbol{x}$ and $\boldsymbol{a}$ the copies of $\boldsymbol{x}$ and $\boldsymbol{a}$ in both A_i and A. The homomorphism ψ_n coincides with θ_i^n on $U = \langle \boldsymbol{x}, \boldsymbol{a} \rangle$. Therefore, up to postcomposing ψ_n with an inner automorphism of G, we can assume without loss of generality that ψ_n coincides with φ_n on $U = \langle \boldsymbol{x}, \boldsymbol{a} \rangle$, and in particular ψ_n restricts to the identity on $\boldsymbol{a}$. This also means that the inner automorphisms $\mathrm{ad}(\psi_n(x_i))$ and $\mathrm{ad}(\varphi_n(x_i))$ induce the same automorphism σ_i of $E(G)$ (recall that $(\varphi_n : G_\Sigma \to G)_{n \in \mathbb{N}}$ is the initial $(\sigma_1, \ldots, \sigma_p)$-test sequence).

For every n, since ψ_n is the identity on $\boldsymbol{a}$, the group $\psi_n(K')$ contains $\boldsymbol{a}$. Since A is finitely presented relative to $U = \langle \boldsymbol{x}, \boldsymbol{a} \rangle$, and since F is a finite group, K' and $\psi_n(K')$ are finitely generated relative to $\boldsymbol{a}$. Therefore, there exists a tuple $\boldsymbol{g}_n$ of elements of G such that

$$\psi_n(K') = \langle \boldsymbol{g}_n, \boldsymbol{a} \rangle.$$

Let $E = \psi_n(F) \subset E(G)$. Consider the following amalgamated product:

$$Q_n = \langle \boldsymbol{g}_n, \boldsymbol{a} \rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i, \ \forall i \in \{1, \ldots, p\} \rangle.$$

For every n, one can define an epimorphism $\pi_n : A \twoheadrightarrow Q_n$ by setting $\pi_n(x_i) = x_i$ and $\pi_n = \psi_n$ on K'. This morphism is well-defined (ω-almost surely). Indeed, for every integer $1 \leq i \leq p$, since x_i normalizes F there exists an automorphism α_i of F such that $x_i f x_i^{-1} = \alpha_i(f)$ for every $f \in F$. The following relation holds ω-almost surely:

$$\psi_n \circ \alpha_i = \sigma_i \circ \psi_n$$

which shows that π_n is well-defined ω-almost surely. In addition, π_n is surjective because its image contains $\boldsymbol{x}$ and $\psi_n(K') = \langle \boldsymbol{g}_n, \boldsymbol{a} \rangle$, which generate the group Q_n. It remains to prove that π_n does not kill any element appearing in the image of the tuple $\Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A ω-almost surely.

Let v be an element appearing in the image of $\Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A. This element can be written in normal form in the splitting $\mathbb{S}_A$ as

$$v = k_0' t_1^{\varepsilon_1} k_1' t_2^{\varepsilon_2} k_2' \cdots t_q^{\varepsilon_q} k_{q+1}'$$

with $k_i' \in K'$ and $t_j \in \{x_1, \ldots, x_p\}$ for every $1 \leq j \leq q$. For every j, if $t_j = t_{j+1} = x_i$ and $\varepsilon_j = -\varepsilon_{j+1}$, then $k_j' \notin F$. By Remark III.1.12, the image of v in L_i can be written in normal form in a similar way, by replacing each k_i' with an element k_i that belongs to the subgroup K of L. Therefore, for every j, if $t_j = t_{j+1} = x_i$ and $\varepsilon_j = -\varepsilon_{j+1}$, then k_j does not belong to F. It follows that $\pi_n(k_j') = \psi_n(k_j')$ does not lie in E ω-almost surely. Otherwise, if $\pi_n(k_j')$ belonged to E ω-almost surely, k_j would lie in F ω-almost surely, contradicting the previous condition. Thus

$$\pi_n(v) = \psi_n(k_0) t_1^{\varepsilon_1} \psi_n(k_1) t_2^{\varepsilon_2} \psi_n(k_2) \cdots t_q^{\varepsilon_q} \psi_n(k_{q+1})$$

is non-trivial ω-almost surely. To finish, simply take $r = \pi_N$ for N such that π_N kills no element appearing in the image of the tuple $\Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ in A. ☐

This seals the proof of Proposition III.4.3. ☐

III.5 Proof of Theorem A

In this section, we prove Theorem A (which follows from the specific case of Theorem E proven in the previous section). First, recall that Theorem A states that every acylindrically hyperbolic group G is $\exists\forall\exists$-elementarily embedded into the HNN extensions $G\dot{*}_{E(G)} = \langle G, t \mid [t,g] = 1, \ \forall g \in E(G)\rangle$. In fact, we just have to prove that G is $\forall\exists$-elementarily embedded into $G\dot{*}_{E(G)}$, in virtue of the following easy and general lemma (which has nothing to do with acylindrical hyperbolicity or groups in general).

Lemma III.5.1. *Let G' be a group, and let G be a subgroup of G. If G is $\forall\exists$-elementarily embedded into G', then G is $\exists\forall\exists$-elementarily embedded into G'.*

Proof. Suppose that G is $\forall\exists$-embedded into G'. Let $\theta(\boldsymbol{t})$ be an $\exists\forall\exists$-formula with m free variables. Suppose that there exists a tuple $\boldsymbol{g} \in G^m$ such that $\theta(\boldsymbol{g})$ holds in G; we will prove that $\theta(\boldsymbol{g})$ holds in G'.

The formula $\theta(\boldsymbol{t})$ can be written as $\exists \boldsymbol{x} \ \mu(\boldsymbol{t},\boldsymbol{x})$, where $\mu(\boldsymbol{t},\boldsymbol{x})$ denotes an $\forall\exists$-formula in $m+n$ variables, where n is the arity of $\boldsymbol{x}$. Since $\theta(\boldsymbol{g})$ holds in G, there exists a tuple $\boldsymbol{h} \in G^n$ such that $\mu(\boldsymbol{g},\boldsymbol{h})$ holds in G. But the formula $\mu(\boldsymbol{t},\boldsymbol{x})$ is an $\forall\exists$ formula, thus $\mu(\boldsymbol{g},\boldsymbol{h})$ holds in G'. This concludes the proof. $\square$

In order to prove Theorem A, it remains to prove that every acylindrically hyperbolic group G is $\forall\exists$-embedded into $G\dot{*}_{E(G)}$. The proof of this result relies on Theorem E.

Theorem III.5.2. *Every acylindrically hyperbolic group G is $\forall\exists$-embedded into $G\dot{*}_{E(G)}$.*

Proof. For brevity, denote $\Gamma = G\dot{*}_{E(G)}$. Let

$$\bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{g}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{g}) \neq 1)$$

be a finite disjunction of systems of equations and inequations in $\boldsymbol{x}$ and $\boldsymbol{y}$. Suppose that G satisfies the following first-order sentence $\mu(\boldsymbol{g})$:

$$\forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{g}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x},\boldsymbol{y},\boldsymbol{g}) \neq 1).$$

Let $\boldsymbol{\gamma}$ be a tuple of elements of Γ of the same arity as $\boldsymbol{x}$. We will prove that there exists a tuple $\boldsymbol{\gamma}'$ of elements of Γ of the same arity as $\boldsymbol{y}$ such that the following holds in Γ:

$$\bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{\gamma},\boldsymbol{\gamma}',\boldsymbol{g}) = 1 \ \wedge \ \Psi_k(\boldsymbol{\gamma},\boldsymbol{\gamma}',\boldsymbol{g}) \neq 1).$$

To this end, we would like to construct a retraction π from the group

$$\langle \Gamma, \boldsymbol{y} \mid \Sigma_k(\boldsymbol{\gamma}, \boldsymbol{y}, \boldsymbol{g}) = 1 \rangle$$

onto Γ, for some $1 \leq k \leq \ell$, such that π kills no $\psi_k(\boldsymbol{\gamma}, \boldsymbol{y}, \boldsymbol{g})$ appearing in the system of inequations $\Psi_k(\boldsymbol{\gamma}, \boldsymbol{y}, \boldsymbol{g}) \neq 1$. Indeed, given such a retraction π, one can simply take $\boldsymbol{\gamma}' = \pi(\boldsymbol{y})$. We could construct this retraction by mimicking the proof of Theorem E, as sketched in Subsection III.3.1. However, in order to avoid unnecessary repetitions, we will appeal to Theorem E. Before applying this result, we need to fix the following problem: Theorem E does not apply directly in the present situation since it only allows us to deal with constants from G, and $\boldsymbol{\gamma}$ is not a tuple of elements of G in general. In order to be able to use Theorem E, we have first to slightly reformulate the problem.

Let $\boldsymbol{s}$ be a (possibly infinite) generating tuple of G. For every $n \geq 1$, let $\boldsymbol{s}_n$ be the n-tuple composed of the first n components of $\boldsymbol{s}$ and let G_n be the subgroup of G generated by $\boldsymbol{s}_n$. For n sufficiently large, the following two conditions are satisfied:

- The subgroup G_n of G contains the finite subgroup $E(G)$. Therefore, there is a finite system of equations $\theta(\boldsymbol{s}_n, t) = 1$ expressing the fact that the stable letter t centralizes $E(G)$.
- The subgroup $\langle G_n, t \rangle$ of Γ contains each component γ_i of $\boldsymbol{\gamma}$. As a consequence, each γ_i can be written as a word $w_i(\boldsymbol{s}_n, t)$.

Let $\boldsymbol{a}$ be the tuple of elements of G obtained by concatenating $\boldsymbol{g}$ and $\boldsymbol{s}_n$. Let $\Sigma'_k(t, \boldsymbol{y}, \boldsymbol{a}) = 1$ denote the finite system of equations

$$(\Sigma_k((w_1(\boldsymbol{s}_n, t), \ldots, w_p(\boldsymbol{s}_n, t)), \boldsymbol{y}, \boldsymbol{g}) = 1) \ \wedge \ (\theta(\boldsymbol{s}_n, t) = 1),$$

and let $\Psi'_k(t, \boldsymbol{y}, \boldsymbol{a}) \neq 1$ denote the finite system of inequations

$$\Psi_k((w_1(\boldsymbol{s}_n, t), \ldots, w_p(\boldsymbol{s}_n, t)), \boldsymbol{y}, \boldsymbol{g}) \neq 1.$$

By assumption, $G \models \mu(\boldsymbol{g})$. Therefore, G satisfies the following first-order sentence $\theta(\boldsymbol{a})$:

$$\forall t \ \exists \boldsymbol{y} \ \bigvee_{k=1}^{\ell} (\Sigma'_k(t, \boldsymbol{y}, \boldsymbol{a}) = 1 \ \wedge \ \Psi'_k(t, \boldsymbol{y}, \boldsymbol{a}_n) \neq 1).$$

By Theorem E, there exist $1 \leq k \leq \ell$, a subgroup G' of G containing $\langle \boldsymbol{a} \rangle$ and an epimorphism

$$\pi : G_{\Sigma'_k} \twoheadrightarrow \Gamma' = (\langle t \rangle \times E(G)) *_{E(G)} G'$$

such that

1. $\pi(t) = t$,
2. $\pi(\boldsymbol{a}) = \boldsymbol{a}$ (in particular $\pi(\boldsymbol{g}) = \boldsymbol{g}$ and $\pi(\boldsymbol{s}_n) = \boldsymbol{s}_n$, and therefore $\pi(\boldsymbol{\gamma}) = \boldsymbol{\gamma}$),
3. no $\psi'_k(t, \boldsymbol{y}, \boldsymbol{a})$ in the system of inequations $\Psi'_k(t, \boldsymbol{y}, \boldsymbol{a}) \neq 1$ is killed by π.

As a consequence, the following system of equations and inequations holds in Γ':

$$\bigvee_{k=1}^{\ell} (\Sigma'_k(t, \pi(\boldsymbol{y}), \boldsymbol{a}) = 1 \ \wedge \ \Psi_k(t, \pi(\boldsymbol{y}), \boldsymbol{a}) \neq 1).$$

It follows that the following system of equations and inequations holds in Γ':

$$\bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{\gamma}, \pi(\boldsymbol{y}), \boldsymbol{g}) = 1 \ \wedge \ \Psi_k(\boldsymbol{\gamma}, \pi(\boldsymbol{y}), \boldsymbol{g}) \neq 1).$$

Since $\Gamma' \leq \Gamma$, this system holds in Γ as well and one can take $\boldsymbol{\gamma}' = \pi(\boldsymbol{y})$. □

III.6 Proof of Merzlyakov's theorem III.3.2 in the general case

We proved Merzlyakov's theorem III.4.3 under the additional assumption that the tuple σ of elements of $\mathrm{Aut}_G(E(G))$ is trivial, that is, each of its components is the identity map of $E(G)$. In this section we eliminate this assumption.

III.6.1 Reduction to an overgroup G_{2p} of G

As above, p denotes the arity of $\boldsymbol{x}$ in the considered first-order sentence. In the proof of Proposition III.4.2, we used the fact that G contains a quasi-convex non-abelian free subgroup F_2 that centralizes $E(G)$ in order to construct a test sequence (that is, a $(\sigma_1, \ldots, \sigma_p)$-test sequence with $\sigma_i = \mathrm{id}_{E(G)}$ for every $1 \leq i \leq p$).

We wish to to circumvent the difficulties involved in adapting the construction of the aforementioned $F_2 \leq G$ to the general setting; that is, we wish to avoid constructing a non-central free subgroup of G that admits a prescribed action by conjugation on $E(G)$. We do so by utilizing Theorem A proved in the previous section. By Theorem A, the inclusion of G into $G *_{E(G)} (E(G) \times \mathbb{Z})$ is an $\exists\forall\exists$-embedding. More generally, the inclusion of G into $G_m := G *_{E(G)} (E(G) \times F_m)$ is an $\exists\forall\exists$-embedding, for any $m \in \mathbb{N}$.

Take $m = 2p$, and let $t_1, ..., t_{2p}$ be a basis of F_{2p}. Let $(\sigma_1, \ldots, \sigma_p)$ be a p-tuple of elements of $\mathrm{Aut}_G(E(G))$. For every $1 \leq i \leq p$, there exists $g_i \in G$ such that

$\sigma_i = \mathrm{ad}(g_i)|_{E(G)}$ (by the definition of $\mathrm{Aut}_G(E(G))$). Note that $\sigma_i = \mathrm{ad}(g_i t_i)|_{E(G)}$ since t_i centralizes $E(G)$. Let α_i be the automorphim of G_{2p} that coincides with $\mathrm{ad}(g_i)$ on G, maps t_i to $g_i t_i$ and maps t_j to t_j for $j \neq i$. The composition $\alpha = \alpha_1 \circ \cdots \circ \alpha_p$ is an automorphism of G_{2p} that satisfies the following conditions:

1. α coincides with conjugation by $g_1 \cdots g_p$ on G,
2. α maps t_i to $g_i t_i$ for $1 \leq i \leq p$, and
3. α fixes t_i for $p+1 \leq i \leq 2p$.

Therefore, up to replacing t_i by $g_i t_i$, we can assume without loss of generality that $\mathrm{ad}(t_i)|_{E(G)} = \sigma_i$ for every $1 \leq i \leq p$. This also implies that G_{2p} has the following presentation:

$$G_{2p} = G *_{E(G)} \left\langle E(G), t_1, \ldots, t_{2p} \;\middle|\; \begin{array}{c} \mathrm{ad}(t_i)|_{E(G)} = \sigma_i \text{ for } 1 \leq i \leq p \\ \mathrm{ad}(t_i)|_{E(G)} = \mathrm{id} \text{ for } p+1 \leq i \leq 2p \end{array} \right\rangle.$$

III.6.2 Construction of a $(\sigma_1, \ldots, \sigma_p)$-test sequence

We now build a $(\sigma_1, \ldots, \sigma_p)$-test sequence from G_{Σ_k} to G_{2p} for any $(\sigma_1, \ldots, \sigma_p)$ in $\mathrm{Aut}_G(E(G))^p$.

Proposition III.6.1. *Let G be an acylindrically hyperbolic group, and let $\boldsymbol{a}$ be a tuple of elements of G. Fix a presentation $\langle \boldsymbol{a} \mid R(\boldsymbol{a}) = 1\rangle$ for the subgroup of G generated by $\boldsymbol{a}$. Let*

$$\bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) \neq 1)$$

be a finite disjunction of finite systems of equations and inequations over G, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are tuples of variables. For every $1 \leq k \leq \ell$, denote

$$G_{\Sigma_k} = \langle \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a} \mid R(\boldsymbol{a}) = 1, \ \Sigma_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1\rangle.$$

Let $p = |\boldsymbol{x}|$ be the arity of $\boldsymbol{x}$, and let $(\sigma_1, \ldots, \sigma_p)$ be a p-tuple of elements of $\mathrm{Aut}_G(E(G))$. *Suppose that G satisfies the following first-order sentence:*

$$\theta : \forall \boldsymbol{x} \ \exists \boldsymbol{y} \ \bigvee_{k=1}^{\ell} (\Sigma_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) = 1 \ \wedge \ \Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}) \neq 1).$$

Then there exist an integer $1 \leq k \leq \ell$ and a $(\sigma_1, \ldots, \sigma_p)$-test sequence $(\varphi_n : G_{\Sigma_k} \to G_{2p})_{n \in \mathbb{N}}$ such that $\varphi_n(\Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}))$ is non-trivial for every n sufficiently large and such that $\varphi_n(\boldsymbol{a}) = \boldsymbol{a}$.

Proof. Recall that G_{2p} has the following presentation:

$$G_{2p} = G *_{E(G)} \left\langle E(G), t_1, \ldots, t_{2p} \;\middle|\; \begin{array}{c} \mathrm{ad}(t_i)|_{E(G)} = \sigma_i \text{ for } 1 \leq i \leq p \\ \\ \mathrm{ad}(t_i)|_{E(G)} = \mathrm{id} \text{ for } p+1 \leq i \leq 2p \end{array} \right\rangle.$$

For every $1 \leq i \leq p$ and for every integer $n \geq 1$, we let

- o_i be the order of σ_i,
- r_n be the remainder of the division of n by o_i, and
- $q_n = o_i + 1 - r_n$.

Note that $2 \leq q_n \leq o_i + 1$ and that $n + q_n = 1 \mod o_i$. We define the element $g_{i,n}$ of G by

$$g_{i,n} = t_{i+p}^n t_i t_{i+p}^{n+1} t_i \cdots t_{i+p}^{2n} t_i^{q_n}.$$

Observe that $\mathrm{ad}(g_{i,n})|_{E(G)} = \sigma_i$ by the choice of q_n.

Since the inclusion of G into G_{2p} is an $\exists\forall\exists$-embedding, the group G_{2p} also satisfies the first-order sentence θ. By the pigeonhole principle, there exists $1 \leq k \leq \ell$ and an infinite set $A \subset \mathbb{N}$ such that for every integer $n \in A$, the group G_{2p} satisfies the following existential sentence:

$$\exists \boldsymbol{y}_n \; \Sigma_k((g_{1,n}, \ldots, g_{p,n}), \boldsymbol{y}_n, \boldsymbol{a}) = 1 \wedge \Psi_k((g_{1,n}, \ldots, g_{p,n}), \boldsymbol{y}_n, \boldsymbol{a}) \neq 1.$$

Define $\varphi_n : G_\Sigma \to G_{2p}$ by $\varphi_n(x_i) = g_{i,n}$, $\varphi_n(\boldsymbol{a}) = \boldsymbol{a}$ and $\varphi_n(\boldsymbol{y}) = \boldsymbol{y}_n$. One can check that the sequence $(\varphi_n : G_{\Sigma_k} \to G_{2p})_{n \in \mathbb{N}}$ is a $(\sigma_1, \ldots, \sigma_p)$-test sequence. ☐

III.6.3 Proof of Merzlyakov's theorem E

By Proposition III.6.1 above, there exist an integer $1 \leq k \leq \ell$ and a $(\sigma_1, \ldots, \sigma_p)$-test sequence $(\varphi_n : G_{\Sigma_k} \to G_{2p})_{n \in \mathbb{N}}$ such that $\varphi_n(\Psi_k(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a}))$ is non-trivial for every n sufficiently large (where G_{2p} is the group defined in the previous subsection). Therefore, Proposition III.4.3 applied to G_{2p} instead of G provides a finite subgroup E of $E(G_{2p}) = E(G)$ and a morphism

$$\pi_{\boldsymbol{\sigma}} : G_{\Sigma_k} \to (G_{2p})_{\boldsymbol{\sigma}} = \langle G_{2p}, \boldsymbol{x} \mid \mathrm{ad}(x_i)|_E = \sigma_i, \;\; \forall i \in \{1, \ldots, p\} \rangle$$

such that the following three conditions hold:

- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \boldsymbol{x}$,

- $\pi_{\boldsymbol{\sigma}}(\boldsymbol{a}) = \boldsymbol{a}$,
- no $\psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ from the tuple $\Psi(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{a})$ is killed by π_{σ}.

Moreover, the image of $\pi_{\boldsymbol{\sigma}}$ is a subgroup of $(G_{2p})_{\boldsymbol{\sigma}}$ of the form

$$\langle \boldsymbol{g}, \boldsymbol{a}\rangle *_E \langle \boldsymbol{x}, E \mid \mathrm{ad}(x_i)|_E = \sigma_i,\ \forall i \in \{1, \ldots, p\}\rangle$$

for some tuple $\boldsymbol{g}$ of elements of G_{2p}. We conclude the proof of Theorem E by composing the morphism $\pi_{\boldsymbol{\sigma}}$ with the retraction r from G_{2p} onto G defined by $r(t_i) = x_i$ for every $1 \leq i \leq p$ and $r(t_i) = 1$ for every $p + 1 \leq i \leq 2p$.

III.7 Trivial positive theory and verbal subgroups

In this section we prove Corollary C, which claims that acylindrically hyperbolic groups have a trivial positive theory (meaning that every positive sentence satisfied by an acylindrically hyperbolic group is satisfied by all groups). We also deduce Corollary D about verbal subgroups of acylindrically hyperbolic groups.

The proof of Corollary C relies on Theorem A: the canonical inclusion of an acylindrically hyperbolic group G into $G \dot{*}_{E(G)}$ is an $\exists\forall\exists$-elementary embedding. In particular, G and $G \dot{*}_{E(G)}$ have the same $\forall\exists$-theory.

Proof of Corollary C. Let G be an acylindrically hyperbolic group. By Theorem 6.3 of [28], if a group satisfies a non-trivial positive sentence, then it also satisfies a non-trivial positive $\forall\exists$-sentence. As a consequence, in order to prove that G has trivial positive theory, it suffices to prove that G has trivial positive $\forall\exists$-theory. Let θ be a positive $\forall\exists$-sentence satisfied by G. Let $E(G)$ denote the maximal finite normal subgroup of G. It follows from Theorem A that the groups G and $\Gamma = \langle G, x, y \mid [x, g] = [y, g] = 1, \forall g \in E(G)\rangle$ have the same $\forall\exists$-theory. As a consequence, θ is satisfied by Γ. Now, observe that Γ maps onto the free group $\langle x, y\rangle \simeq F_2$. Since positive sentences are preserved under epimorphisms, θ is satisfied by F_2. It follows that θ is satisfied by all free groups, and therefore, θ holds in all groups. In other words, θ is trivial. ☐

We now turn to prove Corollary D. Given a word $w \in F_k$, the *verbal subgroup*

$$w(G) = \langle\{w(\boldsymbol{g}),\ \boldsymbol{g} \in G^k\}\rangle$$

is said to have finite width if there exists $m \in \mathbb{N}$ such that any $g \in w(G)$ can be represented as a product of at most m values of w and their inverses. Otherwise, one says that $w(G)$ has infinite width. We also introduce additional notation that will be used in the proof of Lemma III.7.1:

Let $k \geq 1$ be an integer and let w be an element of the free group $F(x_1, \dots, x_k)$. Denote by e_i the sum of the exponents of x_i in w. If $e_1, \dots, e_k$ are all equal to 0, define $d(w) = 0$. Otherwise, let $d(w)$ be their greatest common divisor. Note that $d(w) = 1$ if and only if the image of w in the abelianization of $F(x_1, \dots, x_k)$ is not a power.

Lemma III.7.1. *Let G be a group, Suppose that G has trivial positive theory, then $w(G)$ has infinite width unless w is trivial or $d(w) = 1$ (in which cases the width is equal to 1).*

Proof. First, suppose that w is trivial or $d(w) = 1$; we will show that the width of $w(G)$ is 1. If w is trivial, this is obvious. Next, suppose that w is non-trivial and that $d(w) = 1$. Then there exist $a_1, \dots, a_k \in \mathbb{N}$ such that $a_1 \cdot e_1 + \cdots + a_k \cdot e_k = 1$. It follows that $w(g^{a_1}, \dots, g^{a_k}) = g$ for every $g \in G$. Hence $w(G)$ is equal to G, and its width is equal to 1.

Now, suppose that G has a trivial positive theory, that w is non-trivial and that $d(w) \neq 1$. We will prove that $w(G)$ has infinite width. Assume towards a contradiction that $w(G)$ has finite width $\ell \geq 1$. Then G satisfies the following positive first-order $(\forall\exists)$-sentence ϕ_n for every integer $n \geq 1$: every element of g that can be represented as a product of n elements of $\{w(\boldsymbol{g})^{\pm 1},\ \boldsymbol{g} \in G^k\}$, can also be represented as a product of ℓ elements of $\{w(\boldsymbol{g})^{\pm 1},\ \boldsymbol{g} \in G^k\}$. Since G has trivial positive theory, the sentence ϕ_n is satisfied by all groups. In particular, ϕ_n is true in the free group F_2 (for every n). Thus, $w(F_2)$ has finite width (equal to ℓ). It follows from [104, Lemma 3.1.1 and Theorem 3.1.2] (inspired by [98]) that either w is trivial or $d(w) = 1$, contradicting our assumption. Hence $w(G)$ has infinite width. □

We seal the discussion by noting that Corollary D is an immediate consequence of Corollary C and Lemma III.7.1 above.

III.8 Questions and comments

In [110], Sela asked the following intriguing question:

Question III.8.1. *Which (algebraic, first-order) properties are satisfied by groups G such that G and $G * \mathbb{Z}$ are elementarily equivalent?*

If the answer to The Generalised Tarski Problem is 'Yes', then every acylindrically hyperbolic group G with a trivial finite radical $E(G)$ is elementarily equivalent to $G * \mathbb{Z}$. As far as we are aware, there are no examples of finitely generated groups G such that G and $G*\mathbb{Z}$ are elementarily equivalent, and such that G is not acylindrically hyperbolic. This raises the following question:

Question III.8.2. *Is there a finitely generated group G that is not acylindrically hyperbolic, and such that G and $G * \mathbb{Z}$ are elementarily equivalent (or at least have the same $\forall\exists$-theory)?*

This question is closely related to the following one:

Question III.8.3. *Is acylindrical hyperbolicity preserved under elementary equivalence among finitely generated groups?*

In [3], André proved that the hyperbolicity is preserved under elementary equivalence among finitely generated groups (this result was proved by Sela in [109] for torsion-free groups). Since acylindrically hyperbolic groups are not necessarily finitely generated, Question III.8.3 makes sense without assuming finite generation; however, the answer to this question is negative in general, even among countable groups. This can be seen by utilizing Łoś's theorem (which states that the ultrapower $G^{\mathbb{N}}/\omega$ of a group G is elementarily equivalent to G): the ultrapower of an acylindrically hyperbolic group is never acylindrically hyperbolic because centralisers in $G^{\mathbb{N}}/\omega$ are not virtually cyclic.

We conclude this chapter with the proof of Proposition B, which provides a partial answer to Question III.8.3. The proof relies on Theorem A and on a theorem of Minasyan and Osin that gives a sufficient condition for a group $H = A *_C B$ or $H = A*_C$ to be acylindrically hyperbolic [86, Corollaries 2.2 and 2.3].

Proposition III.8.4. *Let G be a group and let H be a group that admits a non-trivial splitting over a virtually abelian group C. Suppose that G and H are elementarily equivalent (or simply that they have the same $\exists\forall\exists$-theory). If G is acylindrically hyperbolic, then H is acylindrically hyperbolic.*

Recall that a subgroup C of H is called *weakly malnormal* in H if there exists $h \in H$ such that $hCh^{-1} \cap C$ is finite.

Proof of Proposition B. First, note that H is not virtually cyclic since it has the same first-order theory as G, which contains a non-abelian free subgroup.

We first prove the proposition under the simplifying assumptions that $E(G)$ is trivial and that C is abelian. Fix a non-trivial element c of C. Assume towards a contradiction that H is not acylindrically hyperbolic. By [86, Corollaries 2.2 and 2.3], C is not weakly malnormal in H. Hence, for every $h \in H$, the intersection $hCh^{-1} \cap C$ is infinite. In particular, this intersection contains a non-trivial element z. Since C is abelian, z commutes both with c and hch^{-1}. Therefore, H satisfies the following $\exists\forall\exists$-sentence θ:

$$\exists c \neq 1 \ \forall h \ \exists z \neq 1 \ ([c,z] = 1 \ \wedge \ [hch^{-1}, z] = 1).$$

Since G and H have the same $\exists\forall\exists$-theory, $G \models \theta$. By Theorem A, $G * \mathbb{Z} = G * \langle t \rangle \models \theta$ as well. This is a contradiction, since no non-trivial element of $G * \langle t \rangle$ commutes both with c and tct^{-1}. Indeed, by writing the elements of $G * \langle t \rangle$ in normal form, one easily sees that the centralizer of c in $G * \langle t \rangle$ is contained in G, and that the only element of G that commutes with tct^{-1} is 1.

Suppose now that $E(G)$ is non-trivial and denote its order by $N \geq 2$; suppose furthermore that C contains an abelian subgroup of index d. We slightly modify the sentence θ in order to ensure that $c, z \notin E(G)$ and that c and z lie in the abelian index-d subgroup of C. To do so, we replace the conditions "$\exists c \neq 1$" and "$\exists z \neq 1$" with the conditions

- "there exist $N+1$ pairwise distinct elements $c_1^d, \ldots, c_{N+1}^d$", and
- "there exist $N+1$ pairwise distinct elements $z_1^d, \ldots, z_{N+1}^d$".

□

Remark III.8.5. Note that the sentence θ given in the proof of Corollary B shows in particular that Baumslag-Solitar groups do not satisfy the conclusion of Theorem A: $BS(m,n) = \langle a, t \mid ta^m t^{-1} = a^m \rangle$ is not $\exists\forall\exists$-embedded into $BS(m,n) * \mathbb{Z}$. This observation is interesting because the main result of [28] applies to non-solvable Baumslag-Solitar groups (and shows that these groups have a trivial positive theory). Hence, the techniques used in [28] to deal with positive sentences are not sufficient if one wants to deal with inequations.

IV

Limit groups over coherent right-angled Artin groups are cyclic subgroup separable

RIGHT-ANGLED Artin groups serve as an intermediate ground between free and free abelian groups. In this chapter we explore limit groups over right-angled Artin groups, and prove that they are cyclic subgroup separable. More specifically, we prove that cyclic subgroup separability is preserved under exponential completion for groups that belong to a class that includes all coherent RAAGs and toral relatively hyperbolic groups. We do so by exploiting the structure of these completions as iterated free products with commuting subgroups. From this, we deduce that the cyclic subgroups of limit groups over coherent RAAGs are separable, answering a question of Casals-Ruiz, Duncan and Kazachkov. We also discuss relations between free products with commuting subgroups and the word problem, and recover the fact that limit groups over coherent RAAGs and toral relatively hyperbolic groups have a solvable word problem.

We begin by setting up notation and recollecting a few notions that will be used in this chapter. Given a group G, we denote by $Z(G)$ its center, and by $C_G(g)$ the centraliser of $g \in G$. Throughout this chapter, we also assume that all rings are associative, have a free abelian additive group and a multiplicative identity 1. For such a ring A, the additive group generated by 1 is denoted $\mathrm{char}(A) \cong \mathbb{Z}$ and whenever we refer to $\mathbb{Z} \subset A$ we in fact mean $\mathrm{char}(A)$.

As stated in the Introduction, we will be working with *free product with commuting* *subgroups*. These are defined as follows:

Definition IV.0.1. Let G and H be groups, and suppose that L and M are subgroups of G and H respectively. The *free product of G and H with commuting subgroups L and M* is the quotient of the free product $G * H$ by the normal closure of the set of relations $\{[\ell, m] | \ \ell \in L \text{ and } m \in M\}$. We often abbreviate and refer to this group as $\langle G, H | \ [L, M] = 1\rangle$.

Another key notion in this chapter is that of *cyclic subgroup separability*; to make things perfectly clear:

Definition IV.0.2. A group G is called *cyclic subgroup separable* if each of its cyclic subgroups is separable.

IV.1 The class $\mathcal{C}$ and exponential groups

In this section we discuss the class of groups $\mathcal{C}$ and define exponential groups (namely, groups in which one can raise group elements to a power where the exponent is not an integer). We also illustrate the relation between limit groups over coherent RAAGs and exponential groups.

IV.1.1 Right-angled Artin groups and the class $\mathcal{C}$

We recollect the definition of a *right-angled Artin group*:

Definition IV.1.1. Let Γ be a simple graph; the *right-angled Artin group* (or in short, RAAG) $G(\Gamma)$ is the group with presentation

$$\langle \mathrm{V}\Gamma \mid \ [v, u], \ (v, u) \in \mathrm{E}\Gamma\rangle.$$

We refer to the graph Γ as the *defining graph* of $G(\Gamma)$.

Note that in the definition above we do not restrict ourselves to finite graphs. Recall that a group is called *coherent* if all of its finitely generated subgroups are finitely presented. In [42, Theorem 1], Droms shows that a finitely generated RAAG $G(\Gamma)$ is coherent if and only if its defining graph Γ is *chordal*: every subgraph of Γ that is a cycle of more than 3 vertices admits a chord, i.e., an edge that connects two vertices of the cycle. Note that if Γ is an infinite chordal graph, then every finitely generated subgroup H of $G(\Gamma)$ is a subgroup of a RAAG $G(\Gamma')$, where Γ' is a finite and full subgraph of Γ; Γ' is chordal, which implies that H is finitely presented and therefore $G(\Gamma)$ is coherent. The class of coherent RAAGs includes free groups, free abelian groups and RAAGs which are fundamental groups of 3-manifolds (see [42, Theorem 2]).

In [30], Casals-Ruiz, Duncan and Kazachkov define a new class of groups $\mathcal{C}$ which was crafted specifically to satisfy the following property: the $\mathbb{Z}[t]$-completion (see Subsection IV.1.2) of a group G in the class $\mathcal{C}$ can be built by iterating extensions of centralisers (see Section IV.2), and is fully residually G. The definition of the class $\mathcal{C}$ is long and technical, and is beyond the scope of this dissertation; in fact, we do not use the definition of the class $\mathcal{C}$ directly, but rather use properties of groups in the class $\mathcal{C}$ proven in [30]. We highlight a few of these properties: a group G in the class $\mathcal{C}$ is torsion-free, has unique roots and satisfies the *Big Powers* (BP) property: for every $g_1, \dots, g_k \in G$ such that $[g_i, g_{i+1}] \neq 1$, there exists a positive integer N such that for every $n_1, \dots, n_k > N$,

$$g_1^{n_1} \cdots g_k^{n_k} \neq 1.$$

The BP property is used to show that extending centralisers of G yields a group which is fully residually G. We remind that this was disucssed in Subsection II.1.4, where we mentioned Baumslag's work [10] (Lemma II.1.44) and showed how this condition implies that free extensions of centralisers in a non-abelian free group are fully residually free (Example II.1.50).

Free groups, free abelian groups and more generally coherent RAAGs all lie in the class $\mathcal{C}$; in addition, toral relatively hyperbolic groups (torsion-free groups which are hyperbolic relative to a set of free abelian groups) are also in $\mathcal{C}$. A common property shared by these examples of groups in the class $\mathcal{C}$ is that they are equationally Noetherian. Groves showed in [55, Theorem 5.16] that toral relatively hyperbolic groups are equationally Noetherian. The fact that coherent RAAGs are equationally Noetherian follows from their linearity and is mentioned in [31].

If G is a coherent RAAG or a toral relatively hyperbolic group, then limit groups over G admit yet another description, which was not discussed in Section II.1: they are the finitely generated subgroups of the $\mathbb{Z}[t]$-completion of G (see [30, Corollary 6.12 and Theorem 8.1] and [72, Theorems D. and E.]). We therefore continue by discussing *exponential groups* and *A-completions* of groups in the next subsection. We will further explore this characterisation of limit groups over G as finitely generated subgroups of the $\mathbb{Z}[t]$-completion of G in Section IV.2.

IV.1.2 Exponential groups and A-completions

We carry on with laying further foundations required for the proof of Theorems F and G.

Recall that throughout this chapter we assume that rings are associative, have a free abelian additive group and a multiplicative identity 1; as a consequence, characteristic subrings are always isomorphic to $\mathbb{Z}$. The definitions appearing below are simplified versions of the originals: for every definition which involves a ring A and a subring A_0, we assume that $A_0 = \mathrm{char}(A) \cong \mathbb{Z}$. Further details can be found in [88] and [85]. We begin by axiomatizing what it means to raise a group element to a power, where the exponent is taken from a ring A that is not necessarily isomorphic to $\mathbb{Z}$.

Definition IV.1.2. Let A be a ring. A group G is called an *A-group* if there is a map $G \times A \to G$ which satisfies the following (below, g^a denotes the image of (g, a) under the map $G \times A \to G$):

1. $g^1 = g$, $g^0 = 1$ and $1^a = 1$ for every $g \in G$ and $a \in A$,
2. $g^{a+b} = g^a g^b$ and $(g^a)^b = g^{ab}$ for every $g \in G$ and $a, b \in A$,
3. $(hgh^{-1})^a = hg^a h^{-1}$ for every $g, h \in G$ and $a \in A$, and
4. for every $g, h \in G$ and $a \in A$, if $[g, h] = 1$ then $(gh)^a = g^a h^a$.

We call G a *partial A-group* if there exists $P \subset G \times A$ such that g^a is defined whenever $(g, a) \in P$, and all the properties above hold whenever the arguments belong to P.

A homomorphism $f : G \to H$ where G and H are A-groups is called an *A-homomorphism* if $f(g^a) = (f(g))^a$ for every $g \in G$ and $a \in A$. If $H \leq G$ and G is a partial A-group we say that H is a *full A-subgroup* of G if h^a is defined, and lies in H, for every $h \in H$ and $a \in A$. Within our limited settings, an A-completion of a group G is defined as follows:

Definition IV.1.3. Let G be a group. An A-completion of G is an A-group G^A which satisfies the following two conditions:

1. (*minimality*) there is a homomorphism $\tau : G \to G^A$ such that no proper full A-subgroup of G^A contains $\tau(G)$, and

2. (*universal property*) if $f : G \to H$ is a homomorphism and H is an A-group, then f factors via G^A; in other words, there exists a unique A-homomorphism $\overline{f} : G^A \to H$ such that $f = \overline{f} \circ \tau$.

By [88, Theorems 1 and 2], every group admits an A-completion and this completion is unique (up to A-isomorphism). We also remark that if G is abelian, then G^A is abelian and coincides with $G \otimes_{\mathbb{Z}} A$ as the two groups satisfy the same universal property.

IV.2 Extensions of centralisers and cyclic subgroup separability

Exponential completions of (certain) groups in the class $\mathcal{C}$ exhibit a fairly friendly structure: they can be built, "from the group G up", by iterating *extensions of centralisers.*

Definition IV.2.1. Let G be a group and let $u \in G$. Let H be another group, and let $\varphi : C_G(u) \to H$ be an injective homomorphism such that $\varphi(u) \in Z(H)$. The *extension of the centraliser* $C_G(u)$ *by* H is the group

$$G(u, H) = G *_{C_G(u) = \varphi(C_G(u))} H.$$

If $\varphi(C_G(u))$ is a direct factor of H, then the extension is said to be *direct.* If, furthermore, $H = \varphi(C_G(u)) \times \mathbb{Z}$, the extension is said to be *free.*

Direct extensions of centralisers have a particularly nice structure. If $G(u, C_G(u) \times B)$ is a direct extension of the centraliser $C_G(u)$ by $C = C_G(u) \times B$, then

$$\begin{aligned} G(u, C_G(u) \times B) &= G *_{C_G(u)} (C_G(u) \times B) \\ &= G *_{C_G(u)} \langle C_G(u), B \mid [C_G(u), B] = 1 \rangle \\ &= \langle G, B \mid [C_G(u), B] = 1 \rangle, \end{aligned}$$

or in other words $G(u, C_G(u) \times B)$ is the free product of G and B with commuting subgroups $C_G(u)$ and B.

In [79], Loginova gives a criterion under which free products with commuting subgroups are cyclic subgroup separable. This was later generalized in [112].

Theorem IV.2.2 ([79, Main Theorem], [112, Theorem 2.1])**.** *Let G and H be cyclic subgroup separable groups and let $L \leq G$ and $M \leq H$. If the group with presentation $\langle G, H | \, [L, M] = 1 \rangle$ is residually finite, then it is cyclic subgroup separable.*

This criterion will be used throughout this section; as a warm-up we prove the following lemma which easily follows from Theorem IV.2.2:

Lemma IV.2.3. *Let G be a group in the class $\mathcal{C}$ and let $u \in G$ be such that $C_G(u)$ is abelian. Let B be a free abelian group and write $C = C_G(u) \times B$. If G is cyclic subgroup separable, then so is the direct centraliser extension $G(u, C)$.*

Proof. From the discussion preceding this lemma, $G(u, C)$ is the free product of G and B with commuting subgroups $C_G(u)$ and B. Since B is free abelian, its cyclic subgroups are separable; by our assumption, G is also cyclic subgroup separable. Therefore, by Theorem IV.2.2, it is enough to show that $G(u, C)$ is residually finite.

By [30, Theorem 4.2], $G(u, C)$ is fully residually G. Since G is cyclic subgroup separable it is residually finite, and hence $G(u, C)$ is residually finite. ☐

Remark IV.2.4. It is worth mentioning that in an earlier paper, Loginova shows that a free product with commuting subgroups $\langle G, H | \, [L, M] = 1 \rangle$ is residually finite if and only if G and H are residually finite, L is separable in G and M is separable in H (see [78, Theorem 1]). It follows that if G is a residually finite group in the class $\mathcal{C}$ and $u \in G$ is such that $C_G(u)$ is abelian, then $C_G(u)$ is separable. In particular, abelian centralisers in graph towers over coherent RAAGs (see [30, Section 7]) are separable.

The remainder of this section is devoted to proving Theorems F and G. We do so by analysing the construction of the A-completion of a group G from the class $\mathcal{C}$ in steps, following [30], and proving that each step yields a cyclic subgroup separable group. We also remark that in [30], the authors assume that a group $G \in \mathcal{C}$ satisfies an additional condition, named *condition* R, in order to show that G^A enjoys the structure of an iterated centraliser extension (in the case where G is not abelian) or to show that G is a direct summand of G^A (in the case where G is abelian). We make this assumption too (and state the definition of condition R below).

Definition IV.2.5. A group $G \in \mathcal{C}$ is said to satisfy *condition* R if it is a partial A-group, and for every $u \in G$,

1. if $C_G(u)$ is non-abelian, then the centre $Z(C_G(u))$ of $C_G(u)$ is a full A-subgroup, and
2. if $C_G(u)$ is abelian then for every $g \in C_G(u)$, $a \in A$ and $0 \neq n \in \mathbb{Z}$, if $(g^n)^a$ is defined then so is g^a.

Recall our standing assumption that all rings are associative, have a free abelian additive subgroup and a multiplicative identity 1. We also recall the statement of Theorem F for convenience:

Theorem F. *Let G be a group in the class $\mathcal{C}$ which satisfies condition* R, *and let A be a ring. If G is cyclic subgroup separable, then the A-completion of G, G^A, is cyclic subgroup separable.*

We begin by proving Theorem F under the additional assumption that G is abelian.

Lemma IV.2.6. *Let G be an abelian group in the class $\mathcal{C}$ that satisfies condition* R, *and let A be a ring. If G is cyclic subgroup separable, then the A-completion of G, G^A, is cyclic subgroup separable.*

Proof. By [30, Proposition 6.1], G embeds in G^A, G^A is a torsion-free abelian group and G is a direct summand of G^A; write $G^A = G \times B$. Let $(g_1, g_2), (h_1, h_2) \in G^A$ be such that $(h_1, h_2) \notin \langle (g_1, g_2) \rangle$ and let

$$G' = \langle G, (1, g_2), (1, h_2) \rangle \cong G \times \langle g_2, h_2 \rangle \leq G \times B.$$

Note that $\langle g_2, h_2 \rangle$ is isomorphic to either $\mathbb{Z}^2$ or $\mathbb{Z}$, which implies that G' is cyclic subgroup separable (keeping in line with the other proofs appearing in this Section, one can see this for example by invoking Theorem IV.2.2 since G' is the free product of G and $\langle g_2, h_2 \rangle$ with commuting subgroups G and $\langle g_2, h_2 \rangle$). In particular, one can separate (h_1, h_2) from $\langle (g_1, g_2) \rangle$ in a finite quotient of G'.

To finish, it is enough to show that G^A retracts onto G'. Note that the A-completion of G' coincides with that of G. In addition, one can easily verify that G' satisfies condition R.2 (as it is the direct product of two groups that satisfy this condition). Invoking [30, Proposition 6.1] yet again, we have that G' is a direct summand of G^A, and therefore G^A retracts onto G'. □

In the case where G is not abelian, the strategy behind the construction of the A-completion of G is more complicated but still rather straightforward: we repeatedly extend centralisers in G to obtain a group G^* such that G is a full A-subgroup of G^*; in other words, for every $g \in G$, the action of A on g within G^* is defined. Iterating this construction, we eventually obtain the A-completion of G. We recall the following construction from [85] which allows us to extend multiple centralisers at once:

Definition IV.2.7 ([85, Definition 8], [30, Definition 4.7])**.** Let $\mathfrak{C} = \{C_G(u_i)\}_{i\in I}$ be a set of centralisers in a group G and let $\{\varphi_i : C_G(u_i) \to H_i\}_{i\in I}$ be injective homomorphisms such that $\varphi_i(u_i) \in Z(H_i)$ for every $i \in I$. Let T be the tree whose vertex set is $\{v\} \cup \{v_i\}_{i\in I}$ and whose edge set is $\{e_i = (v, v_i)\}_{i\in I}$. Let T_G be the graph of groups whose underlying graph is T, and whose vertex groups, edge groups and edge maps are as follows:

1. $G_v = G$,
2. $G_{v_i} = H_i$,
3. $G_{e_i} = C_G(u_i)$,
4. the map which maps G_{e_i} into G_v is the inclusion, and
5. the map which maps G_{e_i} into G_{v_i} is φ_i.

The fundamental group of this graph of groups is called a *tree extension of centralisers*. It is denoted by $G(\mathfrak{C}, \mathcal{H}, \Phi)$ where $\mathcal{H} = \{H_i\}_{i\in I}$ and $\Phi = \{\varphi_i\}_{i\in I}$.

We have already seen (Lemma IV.2.3) that direct extensions of abelian centralisers of cyclic subgroup separable groups in $\mathcal{C}$ are cyclic subgroup separable. The following lemma, which relies on [30, Proposition 4.8], shows that the same holds for tree extensions of centralisers:

Lemma IV.2.8. *Let G be a group in the class $\mathcal{C}$ and let $\mathfrak{C} = \{C_G(u_i)\}_{i\in I}$ be a set of abelian centralisers in G such that no two of them are conjugate. For each $i \in I$, let H_i be a free abelian group and let $\varphi_i : C_G(u_i) \to H_i$ be an injective homomorphism such that $H_i = \varphi_i(C_G(u_i)) \times K_i$ for some $K_i \leq H_i$.*
Keeping the notation of Definition IV.2.7, if G is cyclic subgroup separable, then the tree extension of centralisers $G(\mathfrak{C}, \mathcal{H}, \Phi)$ is cyclic subgroup separable.

Proof. Let $<$ be a well-ordering of the set $I \cup \{0\}$ (assuming $0 \notin I$ and $0 < i$ for every $i \in I$). We construct, by recursion, a direct system of groups $\{G_i\}_{i \in I \cup \{0\}}$ over I, along with inclusion maps $f_{i,j} : G_i \to G_j$ (for $i < j$) and retractions $r_{j,i} : G_j \to G_i$ (for $i < j$). For readability, we refer to the maps $r_{j,i}$ as retractions, but formally we mean that $r_{j,i}(f_{i,j}(g)) = g$ for every $g \in G_i$. In addition,

1. $G_0 = G$,
2. for every $j < i \in I$, the centraliser of u_i in G_j coincides with the centraliser of u_i in G, that is $C_{G_j}(f_{0,j}(u_i)) = f_{0,j}(C_G(u_i))$,
3. every G_i is cyclic subgroup separable,
4. $\bigcup_{i \in I} G_i = G(\mathfrak{C}, \mathcal{H}, \Phi)$.

Suppose first that $i \in I$ is a successor ordinal, that is $i = j+1$; suppose in addition that for every $k \leq j$ the groups G_k have been defined, along with the suitable inclusion maps and retractions. Set G_i to be the direct extension of the centraliser $C_{G_j}(f_{0,j}(u_i))$ by H_i. Note that this is well-defined, since we assume that $C_{G_j}(f_{0,j}(u_i)) = f_{0,j}(C_G(u_i))$ so $C_{G_j}(f_{0,j}(u_i))$ is a direct factor of H_i. By Lemma IV.2.3, G_i is cyclic subgroup separable. Let $f_{j,i}$ be the obvious inclusion map $G_j \to G_i$, and for every $k < j$ set $f_{k,i} = f_{j,i} \circ f_{k,j}$. Define the retraction $r_{i,j} : G_i \to G_j$ by mapping every element of K_i to $f_{0,j}(u_i)$. Similarly, for every $k \leq j$ set $r_{i,k} = r_{j,k} \circ r_{i,j}$. In addition, for every $k > i$, the fact that $f_{0,i}(u_k)$ is not a conjugate of any element in $C_{G_j}(f_{0,j}(u_i))$ implies that $C_{G_i}(f_{0,i}(u_k)) = f_{0,i}(C_G(u_k))$.

Suppose now that i is a limit ordinal, and that for every $k < i$ the groups G_k, along with the suitable inclusion maps and retractions, have been defined. Consider the directed system of groups $\{G_j\}_{j<i}$ and let $\overline{G}_i$ be its direct limit. Denote by $\overline{f}_{j,i} : G_j \to \overline{G}_i$ the canonical embedding of G_j in $\overline{G}_i$ for $j < i$. To define retractions $\overline{r}_{i,j} : \overline{G}_i \to G_j$, consider the cofinal system $\{G_k\}_{j<k<i}$; its direct limit is $\overline{G}_i$. For every $j < k \leq \ell < i$ we have the following commuting diagram:

$$\begin{tikzcd}
G_k \arrow[d, "f_{k,\ell}"'] \arrow[dr, "r_{k,j}"] & \\
G_\ell \arrow[r, "r_{\ell,j}"'] & G_j
\end{tikzcd}$$

and by the universal property of direct limits we obtain a map $\overline{r}_{i,j} : \overline{G}_i \to G_j$ along with the following commuting diagrams (for $j < k < i$):

$$\begin{tikzcd}
G_k \arrow[r, "\overline{f}_{k,i}"] \arrow[dr, "r_{k,j}"'] & \overline{G}_i \arrow[d, "\overline{r}_{i,j}"] \\
 & G_j
\end{tikzcd}$$

Note that for $g \in G_j$,

$$\begin{aligned} \overline{r}_{i,j}(\overline{f}_{j,i}(g)) &= \overline{r}_{i,j} \circ \overline{f}_{k,i}(f_{j,k}(g)) \\ &= r_{k,j}(f_{j,k}(g)) \\ &= g \end{aligned}$$

so $\overline{r}_{i,j}$ is a retraction. In addition, the fact that

$$C_{G_j}(f_{0,j}(u_k)) = f_{0,j}(C_G(u_k))$$

for every $j < i$ and $k \geq i$ implies that $C_{\overline{G}_i}(\overline{f}_{0,i}(u_k)) = \overline{f}_{0,i}(C_G(u_k))$ for every $k \geq i$. The existence of these retractions also implies that $\overline{G}_i$ is cyclic subgroup separable. Let $g, h \in \overline{G}_i$ be such that $h \notin \langle g \rangle$; there is $j < i$ and $g', h' \in G_j$ such that $\overline{f}_{j,i}(g') = g$, $\overline{f}_{j,i}(h') = h$ and $h' \notin \langle g' \rangle$. Since G_j is cyclic subgroup separable, there is a map $q : G_j \to Q$ such that $q(h') \notin \langle q(g') \rangle$ and Q is finite. The map $q \circ \overline{r}_{i,j} : \overline{G}_i \to Q$ separates h from $\langle g \rangle$.

We now define G_i to be the direct extension of the centraliser $C_{\overline{G}_i}(\overline{f}_{0,i}(u_i))$ by H_i. By Lemma IV.2.3, G_i is cyclic subgroup separable. We also set $\overline{f}_i : \overline{G}_i \to G_i$ to be the inclusion map, and $\overline{r}_i : G_i \to \overline{G}_i$ to be the retraction which maps K_i to $\overline{f}_{0,i}(u_i)$. Finally, define $f_{j,i} = \overline{f}_i \circ \overline{f}_{j,i}$ and $r_{i,j} = \overline{r}_{i,j} \circ \overline{r}_i$. The desired properties of G_i, the maps $f_{j,i}$ and the retractions $r_{i,j}$ can be verified as in the successor stage.

The cyclic subgroup separability of $\bigcup_{i \in I} G_i = G(\mathfrak{C}, \mathcal{H}, \Phi)$ also follows, as in either the successor or the limit stage, depending on the order type of $\{0\} \cup I$. $\square$

With Lemma IV.2.8 in our arsenal, we are ready to describe the construction of the A-completion of a group G in $\mathcal{C}$ and prove Theorem F. Assume in addition that G satisfies condition R; recall that as mentioned earlier, we first construct a group G^* which contains G, and such that G is a full A-subgroup of G^*. We begin by choosing a set of centralisers $\mathfrak{C}(G) = \{C_G(u_i)\}_{i \in I}$ in G which satisfies the following:

1. every centraliser in $\mathfrak{C}(G)$ is abelian, and *not* a full A-subgroup of G,

2. no two centralisers in $\mathfrak{C}(G)$ are conjugate, and
3. any abelian centraliser in G which is not a full A-subgroup is conjugate to a centraliser in $\mathfrak{C}(G)$.

Note that the existence of a set $\mathfrak{C}(G)$ which satisfies the conditions above is guaranteed by Zorn's Lemma. Recall that as in Subsection IV.1.2, for every $C_G(u_i) \in \mathfrak{C}(G)$ we have that $C_G(u_i)^A = C_G(u_i) \otimes_{\mathbb{Z}} A$; in addition, $C_G(u_i)$ is a direct summand of $C_G(u_i)^A$. Setting

$$\mathcal{H}(G) = \{C_G(u_i)^A\}_{i \in I}$$

and

$$\Phi(G) = \{\varphi_i : C_G(u_i) \to C_G(u_i)^A\}_{i \in I}$$

where each φ_i is the canonical embedding, we define

$$G^* = G(\mathfrak{C}(G), \mathcal{H}(G), \Phi(G)).$$

By [30, Lemma 6.5] G is a full A-subgroup of G^*. In addition, by [30, Lemma 6.6], G^* satisfies condition R and we can iterate this construction. As in [30, Subsection 6.2], we define a directed system of groups

$$G = G^{(0)} < G^{(1)} < \cdots < G^{(n)} < \cdots$$

where

$$G^{(n+1)} = (G^{(n)})^* = G^{(n)}(\mathfrak{C}(G^{(n)}), \mathcal{H}(G^{(n)}), \Phi(G^{(n)}))$$

and the maps $f_{i,j} : G^{(i)} \to G^{(j)}$ are the inclusion maps. The direct limit of this system $\bigcup_{n \in \mathbb{N}} G^{(n)}$ is called an *iterated centraliser extension* of G by A, or in short an ICE of G by A. Note that $\bigcup_{n \in \mathbb{N}} G^{(n)}$ is an A-group, since every $g \in G$ lies in $G^{(n)}$ for some n, and therefore the action of A on g is already defined in $G^{(n+1)}$. As a matter of fact, $\bigcup_{n \in \mathbb{N}} G^{(n)}$ is the A-completion of G as evident in [30, Theorem 6.3]. Theorem F now follows:

Proof of Theorem F. The case where G is abelian was covered in Lemma IV.2.6; we therefore assume that G is not abelian. By Lemma IV.2.8, $G^{(n+1)}$ retracts onto $G^{(n)}$ for every $n \in \mathbb{N}$; composing these retractions we obtain retractions $r_{n,m} : G^{(n)} \to G^{(m)}$ for every $m < n$. As in the proof of Lemma IV.2.8, these retractions imply the existence of retractions from the direct limit $\bigcup_{n \in \mathbb{N}} G^{(n)}$ onto each $G^{(n)}$. In addition, each $G^{(n)}$ is cyclic subgroup separable.

Let $g, h \in \bigcup_{n\in\mathbb{N}} G^{(n)}$ be such that $h \notin \langle g \rangle$; g and h lie in some $G^{(n)}$. Since $G^{(n)}$ is cyclic subgroup separable, there is a homomorphism $q : G^{(n)} \to Q$ such that $q(h) \notin \langle q(g) \rangle$ and Q is finite. The composition $q \circ r_n : \bigcup_{n\in\mathbb{N}} G^{(n)} \to Q$ separates h from $\langle g \rangle$. □

Corollary IV.2.9. *Limit groups over cyclic subgroup separable toral relatively hyperbolic groups are cyclic subgroup separable.*

Proof. Let G be a toral relatively hyperbolic group; in particular G lies in $\mathcal{C}$ and satisfies condition R. By Theorem F, the $\mathbb{Z}[t]$-completion of G is cyclic subgroup separable, and by [72, Theorems D. and E.] limit groups over G are exactly the finitely generated subgroups of $G^{\mathbb{Z}[t]}$. □

With a bit more work, we can also deduce the following:

Theorem G. *Limit groups over coherent RAAGs are cyclic subgroup separable.*

Proof. Let $G(\Gamma)$ be a coherent RAAG. By [30, Corollary 6.12 and Theorem 8.1], limit groups over $G(\Gamma)$ are exactly the finitely generated subgroups of $G(\Gamma, \mathbb{Z}[t])^{\mathbb{Z}[t]}$ (where $G(\Gamma, \mathbb{Z}[t])$ is the graph product whose underlying graph is Γ, and whose vertex groups are all $\mathbb{Z}[t]$). In light of Theorem F, since $G(\Gamma, \mathbb{Z}[t])$ lies in $\mathcal{C}$ and satisfies condition R, it is enough to show that $G(\Gamma, \mathbb{Z}[t])$ is cyclic subgroup separable. Let $g, h \in G(\Gamma, \mathbb{Z}[t])$ be such that $h \notin \langle g \rangle$; there is a finite full subgraph Δ of Γ such that $g, h \in G(\Delta, \mathbb{Z}[t])$. Note that $G(\Gamma, \mathbb{Z}[t])$ retracts onto $G(\Delta, \mathbb{Z}[t])$ by killing each vertex group G_v for $v \notin V\Delta$. Hence it is sufficient to show that $G(\Delta, \mathbb{Z}[t])$ is cyclic subgroup separable for every finite full subgraph Δ of Γ.

We prove that $G(\Delta, \mathbb{Z}[t])$ is cyclic subgroup separable by induction on the number of vertices of Δ. Write $V\Delta = \{v_1, v_2, \dots, v_n\}$ and for every $i \leq n$ denote by Δ_i the full subgraph of Δ whose vertices are $v_1, v_2, \dots, v_i$; denote by G_i the copy of $\mathbb{Z}[t]$ which corresponds to the vertex v_i of Δ. For $n = 1$, $G(\Delta_1, \mathbb{Z}[t])$ is free abelian and therefore cyclic subgroup separable. Suppose now that $G(\Delta_i, \mathbb{Z}[t])$ is cyclic subgroup separable, let $v_{i_1}, \dots, v_{i_k}$ be the neighbours of v_{i+1} in Δ_i and note that

$$\begin{aligned} G(\Delta_{i+1}, \mathbb{Z}[t]) &= \langle G, G_{i+1} |\, [G_{i_j}, G_{i+1}],\ j = 1, \dots, k \rangle \\ &= \langle G, G_{i+1} |\, [\langle G_{i_1}, \dots, G_{i_k} \rangle, G_{i+1}] \rangle \end{aligned}$$

(if Δ is not connected, the group $\langle G_{i_1}, \dots, G_{i_k} \rangle$ may be trivial). In other words, $G(\Delta_{i+1}, \mathbb{Z}[t])$ is the free product of $G(\Delta_i, \mathbb{Z}[t])$ and G_{i+1} with commuting subgroups

$\langle G_{i_1}, \dots, G_{i_k} \rangle$ and G_{i+1}. Since $G(\Delta_i, \mathbb{Z}[t])$ and G_{i+1} are cyclic subgroup separable and since $G(\Delta_{i+1}, \mathbb{Z}[t])$ is residually finite as the graph product of residually finite groups (see, for example, [52]), it follows that $G(\Delta_{i+1}, \mathbb{Z}[t])$ is cyclic subgroup separable by Theorem IV.2.2. ☐

IV.3 Free products with commuting subgroups and the word problem

It is well-known that an amalgamated product $G *_K H$ admits a solution to the word problem if the word problem is solvable in G and in H and there is a solution to the membership problem for K in both G and H. We refer the reader to [82, Chapter IV, Corollary 2.2] for the similar case of HNN extensions; note that

$$\langle G, H |\, [L, M] = 1 \rangle = (G *_L L \times M) *_M H$$

and therefore, applying the aforementioned result twice, one obtains:

Lemma IV.3.1. *Let G and H be groups with a solvable word problem and let $L \leq G$ and $M \leq H$. Suppose that the membership problem is solvable for L in G and for M in H. Then there is a solution to the word problem in $\langle G, H |\, [L, M] = 1 \rangle$.*

Recall that a *free centraliser extension* is a centraliser extension of the form $G(u, C_G(u) \times \mathbb{Z}) = \langle G, t |\, [C_G(u), t] = 1 \rangle$, where $u \in G$. Using Lemma IV.3.1 above, we obtain:

Proposition H. *Let G be a group in the class $\mathcal{C}$. If G satisfies condition* R *and has a solvable word problem, then every finitely generated subgroup H of G^A has a solvable word problem.*

Proof. The fact that H is finitely generated and embeds in G^A implies that H embeds in a group obtained from G by taking finitely many free extensions of centralisers; that is, there are groups $G_0, G_1, \dots, G_n$ such that $G_0 = G$, $G_{i+1} = \langle G_i, t_i |\, [C_{G_i}(u_i), t_i] = 1 \rangle$ for some $u_i \in G_i$ and $H \leq G_n$.

We prove that G_n, and hence H, has a decidable word problem by induction on n. Suppose that G_i has a solvable word problem. By Lemma IV.3.1, a solution to the membership problem for $C_{G_i}(u_i)$ in G_i would imply that G_{i+1} has a solvable word problem. But checking whether $g \in G_i$ lies in $C_{G_i}(u_i)$ is equivalent to asking whether $[g, u_i] = 1$ in G_i, which is solvable by the induction hypothesis. Hence the word problem in G_{i+1} is solvable, which completes the proof. ☐

Corollary IV.3.2. *Limit groups over coherent RAAGs and toral relatively hyperbolic groups have a solvable word problem.*

Proof. If $G(\Gamma)$ is a coherent RAAG and H is a limit group over $G(\Gamma)$, then by [30, Theorem 8.1] H is a finitely generated subgroup of $G(\Gamma, \mathbb{Z}[t])^{\mathbb{Z}[t]}$. The group $G(\Gamma, \mathbb{Z}[t])$ lies in $\mathcal{C}$, satisfies condition R and admits an algorithm which checks whether a given word in the canonical generators is trivial or not. Similarly, if G is toral relatively hyperbolic and H is a limit group over G, then by [72, Theorems D. and E.] H is a finitely generated subgroup of $G^{\mathbb{Z}[t]}$. G satisfies condition R and by [58, Subsection 2.7] has a solvable word problem. ☐

The fact that limit groups over coherent RAAGs have a decidable word problem was already mentioned in [30], and follows from these groups being finitely presented and residually finite. The following proposition, which we record here for the sake of completeness, proves that there exists a solution to the word problem for limit groups over toral relatively hyperbolic groups (toral relatively hyperbolic groups are finitely generated [in fact, finitely presented] and equationally Noetherian [55, Theorem 5.16]; in addition, every limit group over a toral relatively hyperbolic group G is recursively presented since it embeds in a group obtained from G by taking finitely many free estension of centralisers [72, Theorems D and E]):

Proposition IV.3.3. *Let G be a countable equationally Noetherian group. If G has a solvable word problem, then so does every finitely generated, recursively presented, residually-G group.*

Proof. Let H be a finitely generated, recursively presented, residually-G group. We execute the following two algorithms in parallel: first, since H is recursively presented there is an algorithm which takes a word $g \in H$ as its input, and returns 'yes' if $g = 1$.

Second, let S be a finite generating set of H and let g be a word in the alphabet $S \cup S^{-1}$. The equational Noetherianity of G implies that there is an algorithm which checks, within finite time, whether a map $S \to G$ extends to a homomorphism $f : H \to G$. Given such a homomorphism f, using a solution to the word problem in G the algorithm can further verify whether or not $f(g) \neq 1$. Since H is residually G, the algorithm described will return 'no' whenever $g \neq 1$ in H. ☐

Remark IV.3.4. Note that the proof above just implies the *existence of an algorithm*; this is because it depends on equational Noetherianity, that is, on reducing infinite systems of equations over G to equivalent finite subsystems. Given a group G, if one has a *concrete* algorithm that converts any infinite system of equations over G into an equivalent finite subsystem, then the proof above yields a *concrete* algorithm that solves the word problem in G.

V

Virtual homology of residually free groups and profinite rigidity of direct products

We study limit groups through the homology of their finite index subgroups, and use homological calculations to show that certain groups are profinitely rigid among finitely presented, residually free groups. More specifically, we show that the virtual second Betti number of a finitely generated, residually free group G is finite if and only if G is either free, free abelian or the fundamental group of a closed surface. We continue and calculate the virtual Betti numbers of limit groups in all dimensions. Later, we employ techniques involving rank gradients of pro-p groups in order to recognize direct product decompositions of certain groups from their finite p-quotients. Combining the above ideas, we show that direct products of free and surface groups are profinitely rigid among finitely presented, residually free groups, partially resolving a conjecture of Bridson's.

We amass a few key theorems that will be used in this chapter; many of these were already mentioned in Chapter II, and we state them again for the convenience of the reader.

The first set of theorems that we record revolves around separability properties of HGFC-groups, limit groups and more generally finitely generated, residually free groups:

Theorem V.0.1 ([123], [66], [18], see Theorem II.2.15 and Corollary II.2.18)**.** *Let G*

be an HGFC-group. Then G is subgroup separable. Furthermore, G has a finite-index subgroup which admits local retractions.

Theorem V.0.2 ([116, Theorems A and B], see Theorem II.1.62)**.** *Limit groups are subgroup separable and admit local retractions.*

Theorem V.0.3 ([27, Theorems A and B], see Theorem II.1.74)**.** *Let G be a finitely generated residually free group, and let $H \leq G$. If H is finitely presented, then H is separable in G; if furthermore H is of type $\mathtt{FP}_\infty(\mathbb{Q})$ then G virtually retracts onto H.*

We will utilize subgroup separability and virtual retractions when we construct subgroups of a given group G with certain homological features. The following two lemmas lie at the heart of this strategy:

Lemma V.0.4 (see Lemma I.3.13)**.** *Let G be a subgroup separable group and let H be a finitely generated subgroup of G. Then for every finite-index subgroup $H_0 \leq H$, there is a finite-index subgroup $G_0 \leq G$ such that $G_0 \cap H = H_0$.*

The proof of this lemma is straightforward, and it appears in Subsection I.3. A different lemma of a similar flavour is the following:

Lemma V.0.5. *Let G be a finitely generated group and let H be a virtual retract of G. Then for any n and any field k there is a finite-index subgroup $G_0 \leq G$ such that $b_n^k(G_0) \geq b_n^k(H)$.*

Proof. Let G_0 be a finite-index subgroup of G which retracts onto H; denote the retraction by $r : G_0 \to H$ and denote by $i : H \to G_0$ the inclusion map. Note that $r \circ i : H \to H$ is the identity map. It follows that the same holds for the induced maps on the n-th homology, that is $r_* \circ i_* = \mathrm{Id}_{H_n(H;k)}$. In particular, $i_* : H_n(H;k) \to H_n(G_0;k)$ is injective and $b_n^k(G_0) \geq b_n^k(H)$. $\square$

We also remind the reader that the *virtual i-th Betti number* of a finitely generated group G (with coefficients in k) is given by

$$\mathrm{vb}_n^k(G) = \sup\{\dim_k H_i(H;k) \mid H \text{ is a finite-index subgroup of } G\}$$

(the full definition appears in the Introduction).

Remark V.0.6. Lemma V.0.5 reduces the problem of showing that $\mathrm{vb}_n^k(G) = \infty$ to finding finitely generated subgroups of G with an arbitrarily large n-th Betti number, as long as G admits local retractions.

As a warm-up, we derive the following simple corollary (which follows the plan of action described in Remark V.0.6 above):

Corollary V.0.7. *If G is a finitely generated, residually free group that is not a limit group, then* $\mathrm{vb}_2(G) = \infty$.

Proof. Recall that as mentioned in Remark II.1.7, Baumslag proved that in this case G must contain $F \times \mathbb{Z}$ (where F is a free group of rank 2) as a subgroup [9, Theorems 1 and 3]. It follows that G contains a subgroup that is isomorphic to $F_n \times \mathbb{Z}$ (where F_n is a free group of rank n), and by Künneth formula,

$$H_1(F_n; k) \otimes H_1(\mathbb{Z}; k) = k^n \otimes k \cong k^n$$

embeds in $H_2(F_n \times \mathbb{Z}; k)$. Thus $b_2^k(F_n \times \mathbb{Z}) \geq n$. Note that $F_n \times \mathbb{Z}$ is of type $\mathtt{FP}_\infty(\mathbb{Q})$ (because it is the extension of two $\mathtt{FP}_\infty(\mathbb{Q})$ groups F_n and $\mathbb{Z}$ [90, Proposition 2.7]), and therefore by Theorem II.1.74, G virtually retracts onto $F_n \times \mathbb{Z}$. By Lemma V.0.5, there is a finite-index subgroup $G_0 \leq G$ such that $b_2(G_0) \geq b_2(F_n \times \mathbb{Z}) \geq n$. This shows that $\mathrm{vb}_2(G) = \infty$. □

These easy steps therefore reduce Theorem I to the case where G is either a limit group or an HGFC-group.

Lastly, we also record another simple lemma that has do with cyclic splittings of limit groups. The lemma follows easily from the fact that limit groups are commutative transitive (see Remark II.1.7), and it will aid the proof of Theorem V.1.7:

Lemma V.0.8. *Suppose that a limit group L contains a subgroup that splits as a cyclic amalgamation $G_1 *_{c_1=c_2} G_2$ or as an HNN extension $G*_\phi$ (where ϕ is an isomorphism between two cyclic subgroups $\langle c_1 \rangle$ and $\langle c_2 \rangle$ of G). Then at least one of $\langle c_1 \rangle$ and $\langle c_2 \rangle$ is maximal abelian in its target vertex group.*

Proof. Assume that L contains a subgroup of the form $G_1 *_{c_1=c_2} G_2$; the HNN extension case is similar (and also follows from the proof for cyclic amalgamation: the HNN extension $G*_\phi$ contains a subgroup of index 2 which has a subgroup of the form $G *_{c_1=c_2} G$).

Suppose for a contradiction that both $\langle c_1 \rangle$ and $\langle c_2 \rangle$ are not maximal abelian in G_1 and G_2 respectively (or in G in the case of an HNN extension). Let c_i' be an element that lies outside of $\langle c_i \rangle$ and commutes with c_i. Since limit groups are commutative transitive (see Remark II.1.7), and since both c_1' and c_2' commute with $c_1 = c_2$ in L it follows that c_1' and c_2' commute. However, since both of them do not lie in $\langle c_1 \rangle = \langle c_2 \rangle$, the commutator $c_1' c_2' {c_1'}^{-1} {c_2'}^{-1}$ is a reduced word in the amalgamated product and is therefore non-trivial, which is a contradiction. □

V.1 Classifying residually free groups by their virtual second Betti number

In this section we prove Theorem I:

Theorem V.1.1 (Theorem I)**.** *Let G be a finitely generated residually free group, or a hyperbolic fundamental group of a finite graph of finitely generated free groups with infinite cyclic edge subgroups, and let k be a field. Then*

1. $\mathrm{vb}_2^k(G) = 0$ *if and only if G is free,*
2. $\mathrm{vb}_2^k(G) = 1$ *if and only if $G \cong \pi_1(\Sigma)$ where Σ is a closed, connected surface,*
3. $\mathrm{vb}_2^k(G) = \binom{d}{2}$ *if and only if $L \cong \mathbb{Z}^d$ (for $d > 2$)*
4. $\mathrm{vb}_2^k(G) = \infty$ *otherwise.*

The proof is inspired by Wilton's work on surface subgroups of HGFC-groups, and in particular borrows ideas from [120, Lemmas 5.9 and 5.10]. Recall that by Remark V.0.6, if G is the fundamental group of a graph of spaces X and G admits local retractions, then it suffices to construct precovers of X with arbitrarily large second Betti numbers in order to obtain that $\mathrm{vb}_2^k(G) = \infty$. This will be the strategy carried out in this section. We point out that the definition of a precover (along with a few key techniques) appears in Subsection II.2.1.

V.1.1 Local structures in graphs of spaces

Recall that if G is an HGFC-group with a corresponding graph of spaces decomposition X, then every vertex X_v of X, along with the incident edges, gives rise to a free group equipped with a peripheral structure; we denote the induced pair at v (for $v \in \mathrm{V}(\Xi)$) by $(G_v, [\underline{w}_v])$ and refer to Subsection II.2.3 for further detail.

One-endedness is particularly important to us as it will serve as an indicator for when an HGFC-group is not free (and therefore possibly has non-trivial second homology). We remind that free splittings of HGFC-groups can be detected *locally*: if an HGFC group G admits a free splitting, then for some $v \in \mathrm{V}(\Xi)$ the induced pair at v, $(G_v, [\underline{w}_v])$, is not one-ended. More generally,

Theorem V.1.2 (Relative Shenitzer's Lemma [119, Theorem 18], see Theorem II.2.26)**.** *Let G be a finitely generated group which is the fundamental group of a graph of groups with infinite cyclic edge groups. Then G is one-ended if and only if every vertex group is freely indecomposable relative to the incident edge groups.*

It is worth mentioning at this point that every HGFC-group splits as the free product of one-ended (hyperbolic) graphs of free groups with cyclic edge groups and a free group:

Lemma V.1.3. *Let G be a finitely generated group which splits as a (not necessarily hyperbolic) graph of free groups with cyclic edge groups. Then $G \cong G_1 * \cdots * G_n * F_r$ where each G_i is a one-ended fundamental group of a graph of free groups with cyclic edge groups, and F_r is a free group of rank $r \geq 0$. Furthermore, if G is an HGFC-group, then every G_i is an HGFC-group.*

Proof. By Grushko's theorem we may write $G = G_1 * \cdots * G_n * F_r$ where each G_i is one-ended and F_r (which might not appear in this decomposition, that is we may have $r = 0$) is free. A standard argument using Bass-Serre theory shows that in fact each G_i is the fundamental group of a finite graph of finitely generated free groups with infinite cyclic edge groups: let T be the Bass-Serre tree that corresponds to the cyclic splitting of G and let T_i be a minimal G_i-invariant subtree of T. Taking the core of the quotient of T_i by G_i we obtain a finite graph of groups decomposition $\mathcal{G}(G_i)$ of G_i. Since G_i is freely indecomposable, the edge groups of $\mathcal{G}(G_i)$ are all infinite cyclic; in particular, these edge groups are finitely generated. By Grushko's theorem G_i is finitely generated, which implies that the vertex groups of $\mathcal{G}(G_i)$ are finitely generated. In addition, since the vertex groups of the cyclic splitting of G are free, the vertex groups of $\mathcal{G}(G_i)$ are all free.

Lastly, if G is hyperbolic, then by [18, Theorem D] G is locally quasiconvex, which implies that every G_i is hyperbolic. ☐

In order to prove Theorem I, we want to understand more than just the free splittings of an HGFC-group G; we seek to understand cyclic splittings too, and for this we employ JSJ decompositions (see Section II.3). Utilizing JSJ decompositions will allow us to endow a one-ended HGFC-group G with a refined splitting in which the induced peripheral structures lie in one of two extremes: they are either *flexible* (meaning that they admit many relative cyclic splittings) and come in the shape of a pair of surface type, or *rigid* (meaning that they do not admit any relative cyclic splittings). We will use Cashen's version of a JSJ decomposition, which encompasses all splittings of a free group relative to a finite set of elements.

Theorem V.1.4 ([32, Theorem 4.25], see Theorem II.3.7)**.** *Let F be a free group, let $\underline{w}$ be a multiword in F and let $\mathcal{E}$ be the family of all infinite cyclic subgroups of F. Then there is a canonical* relative *JSJ decomposition $\mathcal{F}$ of F over $\mathcal{E}$ relative to $\underline{w}$*

(that is, a maximal universal splitting with respect to all splittings of F in which the elements of $\underline{w}$ are elliptic) satisfying the following properties:

1. *each vertex group of $\mathcal{F}$ is of one of the following types:*
    - *cyclic, that is $F_v \cong \mathbb{Z}$,*
    - *surface type, for which the induced pair $(F_v, [\underline{w}_v])$ is of surface type,*
    - *rigid, for which the induced pair $(F_v, [\underline{w}_v])$ is rigid.*
2. *the graph is bipartite, and each edge adjoins a cyclic vertex to a non-cyclic vertex.*
3. *if F_v is a non-cyclic vertex group, then the adjacent edge groups map onto maximal cyclic subgroups of F_v that are non-conjugate in F_v.*

Remark V.1.5. We could alternatively use Bowditch's version of JSJ decompositions, which takes a one-ended hyperbolic group G as its input (see Theorem II.3.6). This decomposition would serve us equally well: Bowditch's theorem, in conjunction with the additional structure of an HGFC-group, yields that the rigid vertices in a JSJ dcomposition of an HGFC-group are all free (and hence the corresponding induced pairs are rigid). We choose to use Cashen's version to highlight the interplay between the global properties of an HGFC-group G and the local properties of the induced pairs at the different vertex groups of G.

Wilton showed in [120] that one-ended HGFC-groups contain surface subgroups; since these groups contain finite-index subgroups that admit local retractions, Lemma V.0.5 implies that they have a positive virtual second Betti number. As mentioned in the Introduction, in some cases we will replicate these surface subgroups in precovers, obtaining precovers with large second homology. The following theorem of Wilton will also prove to be extremely useful for replicating surface subgroups in a precover:

Theorem V.1.6 ([119, Theorem 8]). *If $(F, [\underline{w}])$ is rigid then there is a finite-index subgroup $\widehat{F} \leq F$ such that for every finite-index subgroup $F' \leq \widehat{F}$, the pair $(F', [\underline{w}'])$ obtained by pulling back $[\underline{w}]$ to F' admits the following property: for any component w'_i of $\underline{w}'$, the pair $(F', [\underline{w}' - \{w'_i\}])$ is one-ended.*

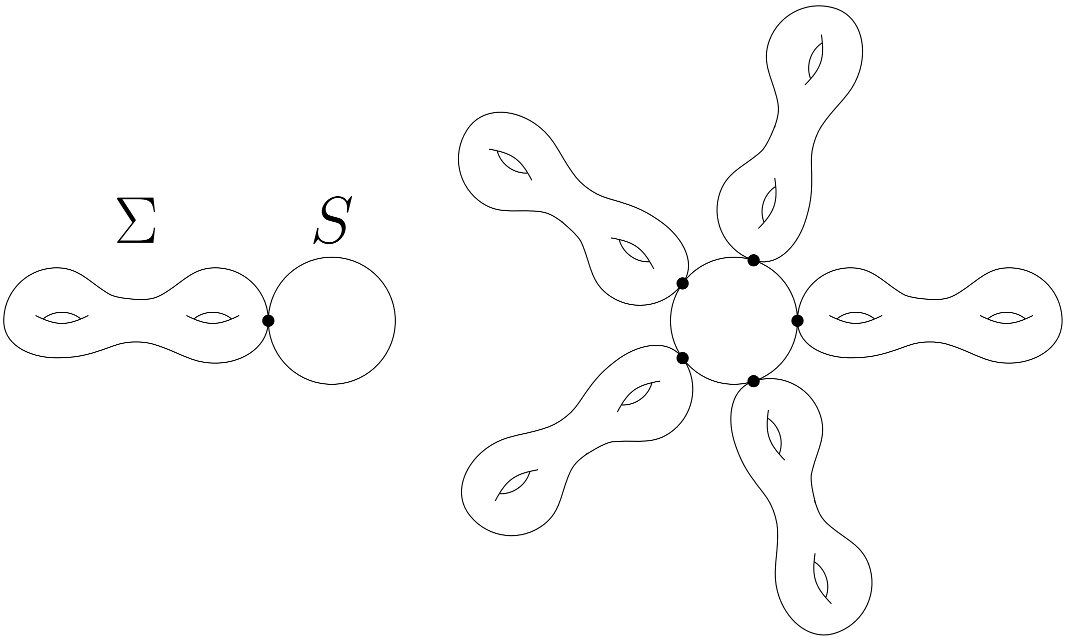

Figure V.1: A wedge of a surface Σ and a circle S on the left and a five-fold covering on the right, corresponding to an index 5 subgroup of $\pi_1(\Sigma) * \mathbb{Z}$ containing 5 surface subgroups.

V.1.2 Precovers with large second homology

The main idea behind the proofs in this section is that if a one-ended HGFC-group G is not a surface subgroup, then it contains a surface subgroup that can be promoted to the existence of many surface subgroups. These surface subgroups will be independent from each other in the second homology of the fundamental group of a precover, producing a subgroup with a large second Betti number. To be precise, we will first prove the following theorem:

Theorem V.1.7. *Let G be a limit group or an HGFC-group. Suppose that G is not a free, a surface or a free abelian group. Then there exists a closed, orientable surface Σ with $\chi(\Sigma) \leq 0$ and an embedding $\pi_1(\Sigma) * \mathbb{Z} \hookrightarrow G$.*

Remark V.1.8. We remark that one can also prove Theorem V.1.7 above, under the assumption that G is hyperbolic, by using a Ping Pong argument. We refer the reader to [6, Theorem 1].

Figure V.1 illustrates how an embedding $\pi_1(\Sigma) * \mathbb{Z} \hookrightarrow G$ implies the existence of many surface subgroups of G.

We divide the proof of Theorem V.1.7 into two cases, depending on whether the relative JSJ decomposition of some induced pair at a vertex of G contains a rigid vertex. For the subsequent two lemmas, G is assumed to be a one-ended HGFC-group. As usual, X will be a graph of spaces whose underlying graph coincides with that of the graph of free groups with cyclic edge groups splitting of G, and whose vertex spaces are either graphs or closed and connected surfaces with boundary (this will depend on our point of view, and we will explicitly describe the vertex spaces of X whenever relevant). Recall that we denote the vertex space corresponding to $v \in \mathrm{V}(\Xi)$ by X_v, and write G_v for $\pi_1 X_v$; fix the same notation for an edge $e \in \mathrm{E}(\Xi)$.

Lemma V.1.9. *If there is a vertex $u \in \mathrm{V}(\Xi)$ such that the relative JSJ decomposition of the induced pair $(G_u, [\underline{w}_u])$ at u has a rigid vertex, then G contains a subgroup H with the following properties:*

1. *H splits as a cyclic amalgamation $H = H_1 *_{c=c'} F$,*
2. *H_1 is a one-ended HGFC-group, and*
3. *F is a non-abelian free group.*

Proof. We begin by "normalizing" X by applying Procedure II.2.27; this will reward us with a graph of spaces that is particularly convenient to work with. We briefly remind the reader how to refine the splitting of G:

Subdivide each edge cylinder of X by adding a cyclic vertex; then, replace each vertex space X_v with the relative JSJ decomposition of the induced pair $(G_v, [\underline{w}]_v)$ at X_v (identifying the cyclic vertices that correspond to the different elements of $[\underline{w}]_v$ with the cyclic vertices of X that are adjacent to X_v). By [123], G is subgroup separable and we may replace X with a finite-sheeted covering $\widehat{X}$ in which all of the attaching maps at cyclic vertices are isomorphisms (see Lemma II.2.28). Folding as in Lemma II.2.29 results in a graph of spaces that satisfies the following properties:

- the graph is bipartite, and edges adjoin cyclic vertices to non-cyclic vertices.
- if F_v is a non-cyclic vertex group, then the adjacent edge groups map onto maximal cyclic subgroups of F_v that are non-conjugate in F_v.

By our assumption, the refined graph of spaces X now contains a vertex space X_u such that the induced pair $(G_u, [\underline{w}_u])$ at u is rigid. We replace $\widehat{X}$ with a finite-sheeted cover, also denoted by $\widehat{X}$, that contains a vertex space $\widehat{X}_{\widehat{u}}$ satisfying the property

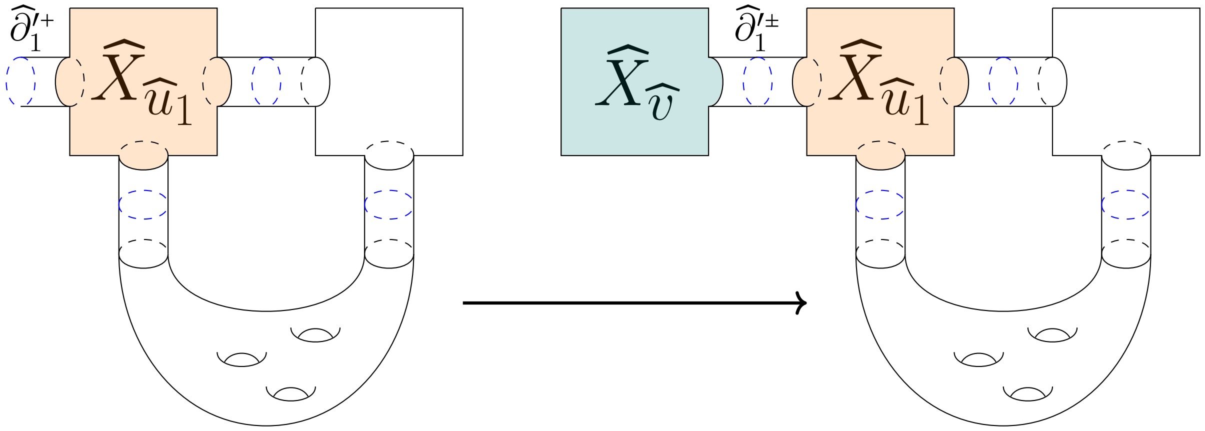

Figure V.2: Constructing X' by attaching an extra copy of $\widehat{X}_{\widehat{v}}$ to X_1' along $\widehat{\partial}_1'^{\pm}$. Cyclic vertices correspond to the blue dashed circles.

described in Theorem V.1.6. In what follows, we will construct a precover X' of X whose fundamental group $H = \pi_1(X')$ satisfies conditions *1.-3.* above.

In a manner similar to [120, Lemma 5.9], we construct a precover $\widehat{X}'$ of $\widehat{X}$ as follows: write $k = \deg(\widehat{u})$ (note that $k > 1$) and let $\widehat{e}_1, \dots, \widehat{e}_k \in \mathrm{E}(\widehat{\Xi})$ be the edges adjacent to $\widehat{u}$. Take k copies of $\widehat{X}_{\widehat{u}}$ and enumerate them $\widehat{X}_{\widehat{u}_1}, \dots, \widehat{X}_{\widehat{u}_k}$; denoting $\pi_1(\widehat{X}_{\widehat{u}_i}) = \widehat{G}_i$ and referring to the peripheral structure on $\widehat{G}_i$ induced by the adjacent edges as $[\underline{\widehat{w}}_i] = \{\widehat{w}_i^1, \dots, \widehat{w}_i^k\}$, the induced pair at $\widehat{X}_{\widehat{u}_i}$ in the resulting precover will be $(\widehat{G}_i, [\underline{\widehat{w}}_i - \{\widehat{w}_i^i\}])$. Take $k-1$ copies of each vertex space $X_{\widehat{v}}$ of $\widehat{X}$ for every $\widehat{v} \neq \widehat{u}$ and enumerate these $\{\widehat{X}_{\widehat{v}_1}, \dots, \widehat{X}_{\widehat{v}_{k-1}}\}$. In order to define the precover $\widehat{X}'$, it suffices to specify which elevations of the attaching maps of X will be attaching maps of $\widehat{X}'$, and to verify that there is a suitable degree-preserving bijection between them. Every elevation of an attaching map to this collection of vertex spaces will be an attaching map of $\widehat{X}'$, except for the elevations that correspond to the edge space of $\widehat{e}_i$ in $\widehat{X}_{\widehat{u}_i}$. Note that if $\partial_e^{\pm}$ is an attaching map in X then every elevation $\widehat{\partial}_{\widehat{e}}^{\pm}$ to $\widehat{X}$ appears $k-1$ times as a hanging elevation in the collection of vertex spaces we have just defined. Therefore one can pick a suitable degree-preserving bijection between these hanging elevations and obtain a precover $\widehat{X}'$ of $\widehat{X}$ (and of X). There are exactly k hanging elevations in this precover, each of them corresponding to the edge space of $\widehat{e}_i$ in $\widehat{X}_{\widehat{u}_i}$ for some $1 \leq i \leq k$.

Let $\widehat{X}_1'$ be the connected component of $\widehat{X}'$ containing $\widehat{X}_{\widehat{u}_1}$ and note that by Theorem V.1.6 the induced pair at every vertex of $\widehat{X}_1'$ is one-ended. It follows from Theorem II.2.26 that $H_1 = \pi_1(\widehat{X}_1')$ is one-ended. Note that H_1 is hyperbolic: by [18, Theorem

D], G is locally quasiconvex, and since H_1 is finitely generated it is a quasiconvex subgroup of a hyperbolic group. Note that $\widehat{X}'_1$ contains at least one hanging elevation, corresponding to the edge space of $\widehat{e}_1$ in $\widehat{X}_{\widehat{u}_1}$; denote this elevation by $\widehat{\partial}'^{+}_1$. Finally, to finish, note that $\widehat{X}'$ contains a cyclic vertex $\widehat{X}_{\widehat{c}}$ and another vertex space $\widehat{X}_{\widehat{v}}$ that may be paired with $\widehat{X}_{\widehat{u}_1}$ along $\widehat{\partial}'^{\pm}_1$; attach these to $\widehat{X}'_1$ as in Figure V.2 to obtain a precover X' of X. Setting $H = \pi_1(X')$ gives the desired result.

By Lemma II.2.12, we can complete X'_1 to a cover $\overline{X}$ of X. Let X' be the precover of X which contains the embedded copy of X'_1 in $\overline{X}$, together with two extra vertex spaces: a cyclic vertex space that is attached to the vertex space $\widehat{X}_{\widehat{u}_1}$ of X'_1 along $\widehat{\partial}'^{+}_1$, and a non-cyclic vertex space that is attached to this cyclic vertex. Letting $H = \pi_1(X')$ completes the proof. □

We next deal with the case where none of the relative JSJ decompositions of the induced pairs of X have a rigid vertex. We will make use of the following lemma which determines the finite-sheeted covering spaces of an orientable surface:

Lemma V.1.10 ([89, Lemma 3.2]). *Let Σ be an orientable and connected surface with positive genus and let $\alpha \geq 1$. For each boundary component of Σ, pick a collection of degrees summing to α. Then there is a connected α-sheeted covering $\widehat{\Sigma} \to \Sigma$ such that the connected components of the preimage of each boundary component $\partial\Sigma_i$ of Σ cover $\partial\Sigma_i$ with the prescribed degrees if and only if the number of boundary components of $\widehat{\Sigma}$ has the same parity as $\alpha \cdot \chi(\Sigma)$.*

Lemma V.1.11. *Suppose that for every $v \in \mathrm{V}(\Xi)$ the vertices of the relative JSJ decomposition of the induced pair $(G_v, [\underline{w}_v])$ at v are all of surface type. If G is not the fundamental group of a closed surface, then G contains a subgroup H with the following properties:*

1. *H splits as a cyclic amalgamation $H = \pi_1(\Sigma) *_{c=c'} F$,*
2. *$\pi_1(\Sigma)$ is the fundamental group of a closed and orientable surface, and*
3. *F is a non-abelian free group.*

Proof. Applying Procedure II.2.27 and Lemma II.2.28, we obtain a finite-index subgroup $\widehat{G}$ of G and a graph of spaces decomposition $\widehat{X}$ of G in which the attaching maps at cyclic vertices are homeomorphisms. The construction of $\widehat{X}$ also implies that every non-cyclic vertex is adjacent only to cyclic vertices and vice-versa, and that there is exactly one edge connected to each boundary component of each surface vertex space of $\widehat{X}$. The resulting space $\widehat{X}$ is "almost a surface", with singularities concentrated only at the cyclic vertices of X.

Since G is torsion-free, Nielsen realization implies that $\widehat{G} = \pi_1 \widehat{X}$ is not a surface group. Therefore, there exists $\widehat{c} \in \mathrm{V}(\widehat{\Xi})$ such that $\deg \widehat{c} > 2$ and $\pi_1 \widehat{X}_{\widehat{c}}$ is cyclic. We will use the fact that $\deg \widehat{c} > 2$ in order to construct a precover X' of $\widehat{X}$ whose fundamental group H satisfies properties (1)-(3) above. Let $\widehat{e}_1, \widehat{e}_2, \widehat{e}_3$ be three distinct edges adjacent to $\widehat{c}$ in $\widehat{\Xi}$ and assume that $\widehat{e}_i$ adjoins $\widehat{c}$ to a vertex $\widehat{v}_i \in \widehat{\Xi}$ (it is possible that not all three vertices $\widehat{v}_i$ are distinct). For $1 \le i \le 3$ let $\widehat{X}_i$ be the precover of X obtained by removing the edge cylinder $\widehat{X}_{\widehat{e}_i} \times (-1,1)$ from $\widehat{X}$. $\widehat{X}_i$ has two hanging elevations which we denote by $\widehat{\partial}_i^+ : \widehat{X}_{\widehat{e}_i} \to \widehat{X}_{\widehat{v}_i}$ and $\widehat{\partial}_i^- : \widehat{X}_{\widehat{e}_i} \to \widehat{X}_{\widehat{c}}$. We next construct three precovers X_i' of $\widehat{X}$ (and X) with the property that the geometric realization of X_i' is a connected, orientable surface with two boundary components, each mapped homeomorphically to $\widehat{X}_{\widehat{e}_i} \subset \widehat{X}_{\widehat{v}_i}$ under the natural map $X_i' \to \widehat{X}$.

Take two copies of each surface type vertex space of $\widehat{X}_i$ and take $\deg(\widehat{v})$ copies of each cyclic vertex $\widehat{X}_{\widehat{v}}$ of $\widehat{X}_i$. We may now pair the hanging elevations in this collection of vertex spaces to obtain a precover X_i' of X making sure that

- each non-hanging elevation whose target space is of surface-type in $\widehat{X}_i$ is non-hanging in X_i',
- if $X'_{v'}$ is a cyclic vertex space of X_i' then $\deg(v') = 2$,
- the two copies of $\widehat{\partial}_i^+$ in X_i' are the only hanging elevations in X_i' whose target space is of surface-type.

Since all of the attaching maps of X_i' at cyclic vertices are isomorphisms, each attaching map with target in a surface vertex X_v' identifies the corresponding edge space with a boundary component of X_v' and all of the surface-type vertices are orientable, X_i' is an orientable surface with two boundary components ${S^i}_1$ and ${S^i}_2$ as desired. If X_i' is not connected, replace it with a connected component that has a non-empty boundary. If this connected component has a single boundary component, replace it with a 2-sheeted covering with two boundary components of degree 1; the existence of such a cover is guaranteed by Lemma V.1.10.

To construct X', we use X_1', X_2' and X_3', and take two additional copies $\widehat{X}_{\widehat{c}_1}$ and $\widehat{X}_{\widehat{c}_2}$ of $\widehat{X}_{\widehat{c}}$. We attach $\widehat{X}_{\widehat{c}_1}$ to ${S^1}_1$, ${S^2}_1$ and ${S^3}_1$ by three edge cylinders, and attach $\widehat{X}_{\widehat{c}_2}$ to ${S^1}_2$ and ${S^2}_2$ by two edge cylinders. The geometric realization of X' is described in Figure V.3. The fundamental group H of X' splits as a cyclic amalgamation, where one of the factors is the fundamental group of a closed and orientable surface

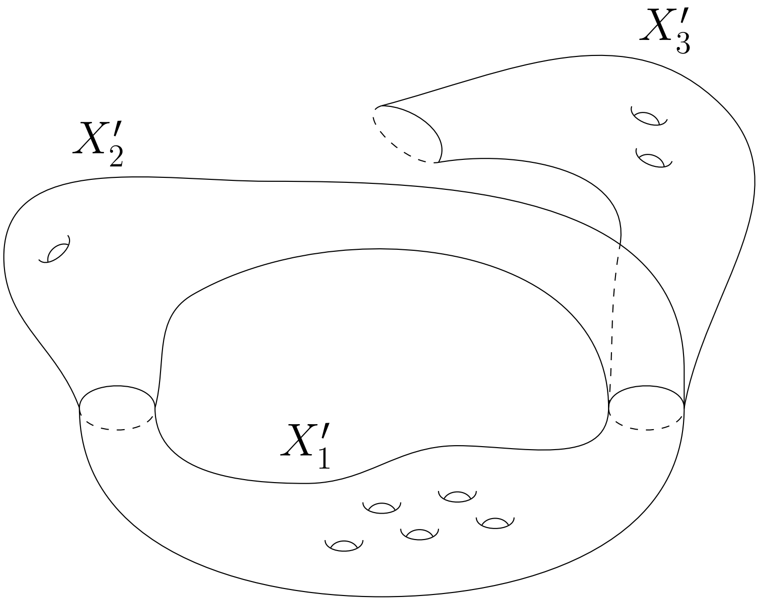

Figure V.3: The precover X' is constructed by assembling a closed surface from X'_1 and X'_2, and then attaching X'_3 along one of its two boundary components.

(obtained from X'_1 and X'_2) and the other factor is $\pi_1(X'_3)$ (which is a non-abelian free group).

$\square$

The following lemma is an easy exercise which will help us turn the subgroups obtained in Lemmas V.1.9 and V.1.11 into subgroups of the form $\pi_1(\Sigma) * \mathbb{Z}$.

Lemma V.1.12. *Let G be a group and let F be a non-abelian free group. Then any cyclic amalgamation $G *_{c=c'} F$ contains a copy of $G * \mathbb{Z}$.*

Proof. Let H be a cyclic amalgamation of the form $G *_{c=c'} F$. We first prove that any $g_1 \cdot h_1 \cdots g_n \cdot h_n \in H$ satisfying the following properties is non-trivial:

1. none of g_i and h_i are trivial (except for, perhaps, g_1 and h_n),
2. each g_i lies in G and each h_i lies in F,
3. none of the h_i lie in the maximal cyclic subgroup of F containing c'.

Induct on the number k of elements g_i such that $g_i \in \langle c \rangle$; if no g_i lies in $\langle c \rangle$ then $g_1 \cdot h_1 \cdots g_n \cdot h_n$ is reduced and hence non-trivial in $G *_{c=c'} F$. Suppose now that the claim holds for $k < \ell$, and let $g_1 \cdot h_1 \cdots g_n \cdot h_n$ be such that ℓ of the elements g_i are in $\langle c \rangle$. If $g_j \in \langle c \rangle$, consider

$$g_1 \cdot h_1 \cdots g_{j-1} \cdot (h_{j-1} \cdot g_j \cdot h_j) \cdots g_n \cdot h_n.$$

Since h_{j-1} and h_j are outside of the maximal cyclic subgroup of F containing c', so is $h_{j-1} \cdot g_j \cdot h_j$. It follows that $\ell - 1$ elements among $g_1, \ldots, g_{j-1}, g_{j+1}, \ldots, g_n$ lie in $\langle c \rangle$, and by the induction hypothesis $g_1 \cdot h_1 \cdots g_{j-1} \cdot (h_{j-1} \cdot g_j \cdot h_j) \cdots g_n \cdot h_n \neq 1$ in H.

To finish, let $g \in F$ be such that $\langle c', h \rangle$ is a free group of rank 2; we claim that the subgroup $\langle G, h \rangle$ of H is isomorphic to $G * \mathbb{Z}$. It suffices to show that every element of $\langle G, h \rangle$ of the form $g_1 \cdot h^{k_1} \cdot g_2 \cdot h^{k_2} \cdots g_n \cdot h^{k_n}$, and in which

1. every g_i is non-trivial (except for, perhaps, g_1),
2. every k_i is different to 0 (except for, perhaps, k_n), and
3. $n \geq 1$ (and if $n = 1$ then either $g_1 \neq 1$ or $k_1 \neq 0$),

is non-trivial. This follows immediately from the first part of the proof. □

We are finally ready to prove Theorem V.1.7:

Theorem V.1.13 (Theorem V.1.7)**.** *Let G be a limit group or an HGFC-group. Suppose that G is not a free abelian, a free or a surface group. Then there exists a closed, orientable surface Σ with $\chi(\Sigma) \leq 0$ and an embedding $\pi_1\Sigma * \mathbb{Z} \hookrightarrow G$.*

Proof of Theorem V.1.7. Suppose first that G is an HGFC-group. If G is one-ended, then Lemmas V.1.9 and V.1.11 imply that G contains a subgroup of the form $H *_{c=c'} F$ where H is a one-ended HGFC-group and F is a non-abelian free group. By Lemma V.1.12, G contains a subgroup isomorphic to $H * \mathbb{Z}$, and if H is not a surface group then by [120, Theorem 6.1] H contains a surface subgroup $\pi_1(\Sigma)$ with $\chi(\Sigma) \leq 0$. The existence of a suitable embedding $\pi_1\Sigma * \mathbb{Z} \hookrightarrow G$ follows.

If G is not one-ended, by Lemma V.1.3 we may write $G = G_1 * \cdots * G_n * F_r$ where each G_i is a one-ended HGFC-group, F_r (which might not appear in this decomposition) is free and there are at least 2 factors in this free product. If some G_i is not a surface group, then the previous paragraph implies that G_i, and hence G, contains a subgroup of the form $\pi_1(\Sigma) * \mathbb{Z}$. Similarly, if some G_i is the fundamental group of a closed, connected hyperbolic surface, then G clearly contains a subgroup of the desired form.

Suppose now that G is a limit group, and let H be a one-ended group appearing in the hierarchy of G with no one-ended groups below it. Recall that as in Corollary

II.1.60, H is either free abelian (of rank at least 2), or an HGFC-group. If H is an HGFC-group, and H is not a surface group, then the discussion above implies that there exists a closed, orientable surface Σ with $\chi(\Sigma) \leq 0$ and an embedding $\pi_1\Sigma * \mathbb{Z} \hookrightarrow H \hookrightarrow G$.

Otherwise, H is either free abelian (of rank at least 2) or the fundamental group of a surface. Since G is not a surface or a free abelian group, there is a group H' lying above H in the hierarchy. If H is a free factor of H', then G contains a subgroup of the form $H * \mathbb{Z}$ as desired (note that a free abelian group of rank 2 is the fundamental group of the torus T^2, and $\chi(T^2) = 0$). If not, then H' contains a free factor K, which admits a cyclic splitting in which H is one of the factors; denote this splitting by $\mathcal{G}_K$. We divide the proof into two cases, depending on whether H is abelian.

Case 1: Suppose that H is non-abelian, and write $H = \pi_1(\Sigma)$ where Σ is a closed, connected surface; up to passing to a finite-index subgroup, we may assume that Σ is orientable. If H is the only vertex group of $\mathcal{G}_K$, realize $\mathcal{G}_K$ as a graph of spaces X_K which consists of the single vertex space Σ accompanied by edge cylinders. Construct a precover $\widehat{X}_K$ of X_K by adjoining two copies of Σ with a suitable edge cylinder. One can see that $\pi_1(\widehat{X}_K) \cong \pi_1(\Sigma) *_{c=c'} \pi_1(\Sigma)$ contains a subgroup of the form $\pi_1(\Sigma) *_{c=c'} F$ where F is a free group of rank 2, and by Lemma V.1.12 there is an embedding $\pi_1\Sigma * \mathbb{Z} \hookrightarrow K \hookrightarrow G$.

If H is not the only vertex group of $\mathcal{G}_K$, then K contains a subgroup that splits as a cyclic amalgamation in which H is one of the factors; denote the other factor by J. If J is non-abelian, then since every two non-commuting elements in a limit group generate a non-abelian free group, K contains a subgroup of the form $H *_{c=c'} F$ (where F is a free group of rank 2). Evoking Lemma V.1.12 completes the proof. If J is abelian (in which case its rank is at least 2), then similarly K contains a subgroup of the form $\mathbb{Z}^2 *_{c=c'} F$ (where F is a free group of rank 2), which yields the desired subgroup of G.

Case 2: Suppose now that H is a free abelian group of rank at least 2. By Lemma V.0.8, H cannot be the only vertex group of $\mathcal{G}_K$, and there is a vertex group J of $\mathcal{G}_K$ that is adjacent to H. Moreover, Lemma V.0.8 implies that J is non-abelian. Since any two non-commuting elements of J must generate a free group of rank 2, K contains a subgroup of the form $H *_{c=c'} F$ (where F is a free group of rank 2). Lemma V.1.12 gives rise to an embedding $\mathbb{Z}^2 * \mathbb{Z} \hookrightarrow G$.

□

Theorem I now follows.

Theorem V.1.14 (Theorem I)**.** *Let G be a finitely generated residually free group or an HGFC-group. Then*

1. $\mathrm{vb}_2^k(G) = 0$ *if and only if G is free,*

2. $\mathrm{vb}_2^k(G) = 1$ *if and only if $G \cong \pi_1(\Sigma)$ where Σ is a closed, connected surface,*

3. $\mathrm{vb}_2^k(G) = \binom{d}{2}$ *if and only if $L \cong \mathbb{Z}^d$ (for $d > 2$)*

4. $\mathrm{vb}_2^k(G) = \infty$ *otherwise.*

Proof. First note that if G is free then $\mathrm{vb}_2^k(G) = 0$, if G is a surface group then $\mathrm{vb}_2^k(G) = 1$ and if $G \cong \mathbb{Z}^d$ then $\mathrm{vb}_2^k(G) = \binom{d}{2}$. Suppose now that G is a limit group or an HGFC-group, and that G is not free, surface or free abelian. By passing to a finite-index subgroup, we may assume that G admits local retractions (since G is torsion-free and not isomorphic to a free, a surface or a free abelian group, it cannot have a finite-index subgroup that is free, free abelian or surface). By Theorem V.1.7, G contains a subgroup isomorphic to $\pi_1(\Sigma) * \mathbb{Z}$; by Lemma V.0.5, it is enough to show that $\pi_1(\Sigma) * \mathbb{Z})$ contains subgroups with arbitrarily large second Betti number. Let $f : \pi_1(\Sigma) * \mathbb{Z} \to \mathbb{Z}/n\mathbb{Z}$ be the map which sends $\pi_1(\Sigma)$ to the trivial element, and the generator of $\mathbb{Z}$ to the generator of $\mathbb{Z}/n\mathbb{Z}$. $\ker f$ is a finite-index subgroup of $\pi_1(\Sigma) * \mathbb{Z}$ that splits as a free product with $n + 1$ factors: one of them is $n\mathbb{Z}$, and the rest are conjugates of $\pi_1(\Sigma)$; $\ker f$ can also be seen as the fundamental group of the cover of a wedge of a circle and Σ illustrated in Figure V.1. A repeated application of Mayer-Vietoris shows that the second Betti number of $\ker f$ is at least $n \cdot \mathrm{b}_2^k(\pi_1(\Sigma)) = n$.

It is left to show that if G is a residually free group that is not a limit group, then $\mathrm{vb}_2^k(G) = \infty$. This was done in Corollary V.0.7. □

V.2 Virtual homology in higher dimensions and residually free manifolds

We now turn to calculating virtual Betti numbers in dimensions $n > 2$. The following technical lemma will easily imply Proposition J.

Lemma V.2.1. *Suppose that a finitely generated group G splits as a finite and connected non-trivial graph of spaces X. Suppose furthermore that there is a vertex group G_v of $\mathcal{G}(X)$ with $\mathrm{b}_n^k(G_v) > 0$ for some $n \in \mathbb{N}$ and a field k. If every edge group G_e of $\mathcal{G}(X)$ satisfies $b_n^k(G_e) = 0$, and at least one of the following holds,*

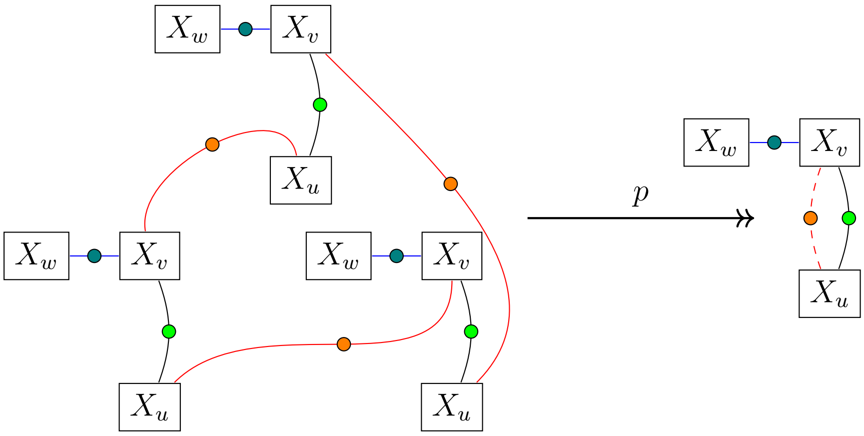

Figure V.4: Splicing together three copies of a precover of X along the red edge.

1. *the underlying graph Ξ of X is not a tree,*

2. *the vertex group G_v is separable in G,*

then for every $\ell \in \mathbb{N}$ there is a finite cover $\widehat{X}$ of X with $\mathrm{b}_n^k(\pi_1(\widehat{X})) \geq \ell$. In particular, $\mathrm{vb}_n^k(G) = \infty$.

Proof. Suppose first that Ξ is not a tree, and therefore the underlying graph Ξ of X contains an edge $e \in \mathrm{E}(\Xi)$ with $\Xi - \{e\}$ a connected graph. Take ℓ copies of X and enumerate them $X_1,\ldots,X_\ell$; let e_i be the copy of e in $\mathrm{E}(\Xi_i)$. Remove $X_{e_i} \times (-1,1)$ from X_i to obtain a precover X_i' of X_i with two hanging elevations $\partial_i^\pm$. Let $\widehat{X}$ be the space obtained by splicing together $X_1',\ldots,X_\ell'$, pairing the elevation ∂_i^+ with ∂_{i+1}^- (and ∂_ℓ^+ with ∂_1^-). The resulting space for $n = 3$ appears in Figure V.4. We remark that the ℓ-fold cover $\widehat{X}$ of X can also be obtained by choosing a loop γ in Ξ that contains e and such that $[\gamma] \in \pi_1(\Xi)$ is a primitive element, mapping G onto $\pi_1(\Xi)$ by killing all of the vertex groups, and retracting $\pi_1(\Xi)$ onto $[\gamma]$. The cover corresponding to the preimage of $\langle [\gamma]^\ell \rangle$ in G is $\widehat{X}$.

Note that $\widehat{X}$ contains ℓ copies of X_v as vertex spaces. Since $\mathrm{b}_n^k(G_v) > 0$ and for every edge $e \in \mathrm{E}(\Xi)$ we have that $\mathrm{b}_n^k(G_e) = 0$, a repeated use of Mayer-Vietoris shows that $b_n^k(\widehat{G}) \geq \sum_{\widehat{v} \in \mathrm{V}(\widehat{\Xi})} b_n^k(\widehat{G}_{\widehat{v}}) \geq \ell$.

We still need to treat the case where Ξ is a tree; in this case, since G_v is separable in G there are finite-index normal subgroups $H_1, H_2, \ldots$ of G such that $G_v = \bigcap_{i \in \mathbb{N}} H_i$. Note that since Ξ is non-trivial, $[G : G_v] = \infty$ and therefore the index of $G_n = \bigcap_{i=1}^n H_i$ in G goes to ∞ as $n \to \infty$. Consider the quotient maps $q_n : G \twoheadrightarrow G/G_n = Q_n$, and note that $[Q_n : q_n(G_v)] = |Q_n| = [G : G_n] \xrightarrow{n \to \infty} \infty$. Let $\widehat{X}_n$ be the cover of X that corresponds to G_n, and note that there are exactly $|Q_n|$ vertices covering G_v in the graph of spaces decomposition that $\widehat{X}_n$ inherits from X. In fact, each of the $|Q_n|$ corresponding vertex groups is isomorphic to G_v since $G_v \leq \pi_1(\widehat{X}_n)$ and $\widehat{X}_n$ is a normal cover of X. Finally, as before, a repeated application of Mayer-Vietoris shows that $b_n^k(G_n) \geq |Q_n| \cdot b_n^k(G_v)$ which goes to infinity as $n \to \infty$. □

We can now easily calculate virtual Betti numbers of limit groups in all dimensions:

Proposition V.2.2 (Proposition J). *Let L be a limit group and let k be a field. Then for any $n \geq 3$, $\mathrm{vb}_n^k(L) < \infty$ if and only if one of the following two holds:*

1. $\mathrm{cd}(L) < n$*, in which case* $\mathrm{vb}_n^k(L) = 0$*, or*

2. *L is free abelian of rank at least n, in which case* $\mathrm{vb}_n^k(L) = \binom{\mathrm{rank}(L)}{2}$.

Proof of Proposition J . If $\mathrm{cd}(L) < n$ then $\mathrm{vb}_n^k(L) = 0$. Suppose that $\mathrm{cd}(L) = m \geq n$, and note that if L is free abelian, then clearly $L \cong \mathbb{Z}^m$ and $\mathrm{vb}_n^k(L) = \binom{m}{2}$. It is left to show that if L is not free abelian then $\mathrm{vb}_n^k(L) = \infty$.

It is a well-known fact, that can be easily proved by employing the hierarchical structure of a limit group, that $\mathrm{cd}(L) = \max(2, m)$ where m is the maximal rank of a free abelian subgroup of L. We may therefore assume that L contains a maximal abelian subgroup M isomorphic to $\mathbb{Z}^m$. If L is freely decomposable, then by [118, Lemma 3.2] M must be contained in one of the factors. It follows that L contains a subgroup isomorphic to $M * \mathbb{Z}$. Given $\ell \in \mathbb{N}$, consider the kernel of the map $f : M \to \mathbb{Z} \to \mathbb{Z}/\ell\mathbb{Z}$ which kills M and maps a generator of $\mathbb{Z}$ to a generator of $\mathbb{Z}/\ell\mathbb{Z}$. $\ker f$ is a finite-index subgroup of $M * \mathbb{Z}$ which admits a free splitting with $\ell + 1$ factors; n of these factors are conjugates of M. Applying Mayer-Vietoris ℓ times yields that $\mathrm{b}_n^k(\ker f) \geq \ell \cdot \mathrm{b}_n^k(M) \geq \ell$ and therefore $\mathrm{vb}_n^k(M * \mathbb{Z}) = \infty$.

If L is freely indecomposable, then Theorem II.1.54 implies that L admits a non-trivial cyclic splitting of one of the following forms:

1. $A *_C B$ where C is infinite cyclic, M is conjugate into one of the two factors, A and B are non-cyclic, and C is maximal abelian in one of the two factors.

2. $A*_C$ where C is an infinite cyclic, maximal abelian subgroup of L, and M is conjugate into A.

Note that by Lemma V.0.5, the n-th Betti number of the vertex group of L that contains a conjugate M is positive. Hence both of the possible splittings of L satisfy the properties required to invoke Lemma V.2.1, and $\mathrm{vb}_n^k(L) = \infty$.

□

Remark V.2.3. Similar arguments show that if $\operatorname{cd}(L) = m \geq 3$ and $L \ncong \mathbb{Z}^n$ then L contains a subgroup isomorphic to $\mathbb{Z}^m * \mathbb{Z}$. This also implies that $\operatorname{vb}_n^k(L) = \infty$ for every $1 \leq n \leq m$ using virtual retractions as in the proof of Theorem I.

A natural class of groups whose n-th virtual Betti number is always finite (and equal to 1) is the class of fundamental groups of aspherical, closed manifolds of dimension n. Using the classification of limit group by their n-th virtual number in Proposition J above, we can deduce:

Corollary V.2.4 (Corollary K). *Let M be an aspherical, closed manifold of dimension n. Then the following two hold:*

1. *if $\pi_1(M)$ is fully residually free and $n \geq 3$ then $M \cong T^n$.*
2. *if $\pi_1(M)$ is residually free and $n \geq 5$ then M has a finite cover that is homeomorphic to the direct product of a torus and finitely many closed surfaces.*

Proof of Corollary K. The manifold M is aspherical, and so are all its finite covers. By Poincaré duality, it follows that $\operatorname{vb}_n(M) = 1$.

If $\pi_1(M)$ is fully residually free, for $n \geq 3$ Proposition J implies that $\pi_1(M) \cong \mathbb{Z}^n$. Note that the classifying map for $\pi_1(M)$, $f : M \to K(\mathbb{Z}^n, 1) = T^n$ is a homotopy equivalence, and in particular M is orientable.

If $n = 3$, then since $\pi_1(M) \cong \mathbb{Z}^n$ is one-ended, the Poincaré Conjecture implies that M is prime, and therefore irreducible. From Waldhausen's homeomorphism theorem [115] we have that $M \cong T^3$. For $n = 4$, since $\mathbb{Z}^4$ is good in the sense of Freedman, [45, Section 11.5] implies that $M \cong \mathbb{Z}^4$. Lastly, for $n \geq 5$, by [7, Theorem A] the Borel Conjecture holds for $\mathbb{Z}^n$, so M is homeomorphic to T^n.

Suppose now that $\pi_1(M)$ is residually free. By [73], M has the homotopy type of a finite CW-complex and therefore $\pi_1(M)$ is of type $\mathrm{FP}(\mathbb{Q})$. By Theorem II.1.72, M has a finite cover $\widehat{M}$ whose fundamental group is a direct product of finitely many limit groups. Since $\operatorname{vb}_n(\pi_1(\widehat{M})) = 1$ and $\operatorname{vb}_\ell(\pi_1(\widehat{M})) = 0$ for $\ell > m$, Proposition J and a direct computation using Künneth formula imply that $\pi_1(\widehat{M})$ is the direct product of finitely many surface groups and a free abelian group. Invoking [7, Theorem A] yields that $\widehat{M}$ is homeomorphic to the direct product of a torus and finitely many closed surfaces, as desired. ☐

Remark V.2.5. The corollary above can also be proved using other cohomological techniques, and in particular by relying on the fact that a Poincaré duality group of dimension n cannot have a Poincaré duality group of dimension n as an infinite-index subgroup.

Remark V.2.6. Passing to a finite cover in (2) of Corollary K above is necessary: let $G = S_1 \times S_2 \times S_3$ where each S_i is the fundamental group of a surface of genus 2, and let H be its index two subgroup which is the kernel of the map $G \to \mathbb{Z}/2\mathbb{Z}$ that maps each of the standard generators of G to the non-trivial element in $\mathbb{Z}/2\mathbb{Z}$. H is the fundamental group of a 6-manifold which is a double cover of the product of three surfaces, and it is residually free as a subgroup of a residually free group. Suppose now that H splits as a direct product $H_1 \times H_2$. Projecting H_1 and H_2 to each of the three factors S_i of G, we get that for every i one of $p_i(H_1)$ and $p_i(H_2)$ is trivial in S_i. By the pigeonhole principle, either H_1 or H_2 is contained in one of the S_i, say H_1 is contained in $S_1 \times \{1\} \times \{1\}$. This implies that H_2 is contained in $\{1\} \times S_2 \times S_3$. Since $[G : H] = 2$ it follows that either $H_1 = S_1 \times \{1\} \times \{1\}$ or $H_2 = \{1\} \times S_2 \times S_3$, which contradicts the fact that H does not contain any of the factors of G.

V.3 On the profinite rigidity of direct products of free and surface groups

In this section we introduce the required tools to prove Theorem N, which asserts that direct products of free, surface and free abelian groups are profinitely rigid among finitely presented groups.

Notation V.3.1. Recall that, as mentioned in Section I, the profinite completion of a group G is denoted by $\widehat{G}$, and the pro-p completion of G is denoted by $G_{\widehat{p}}$. Bold letters such as $\mathbf{A}, \mathbf{B}, \mathbf{G}, \mathbf{H}$ will denote profinite groups. In addition, we write $\mathbf{H} \leq_c \mathbf{G}$ or $\mathbf{H} \leq_o \mathbf{G}$ (and $\mathbf{H} \trianglelefteq_c \mathbf{G}$ or $\mathbf{H} \trianglelefteq_o \mathbf{G}$ respectively) to indicate that the subgroup $\mathbf{H} \leq \mathbf{G}$ (respectively, normal subgroup $\mathbf{H} \trianglelefteq \mathbf{G}$) is closed or open.

The only profinite groups that are finitely generated as abstract groups are finite (since, otherwise, they are not even countable). Therefore, when we say that a profinite group $\mathbf{G}$ is finitely generated, we really mean that it is topologically finitely generated, i.e. that there exists a finite subset $S \subset \mathbf{G}$ that generates a dense abstract subgroup of $\mathbf{G}$. One of the required ingredients for the proof of Theorem N is the following consequence of Theorem V.1.7.

Theorem V.3.2. *Let G be a finitely generated, residually free group. Suppose that for every finite-index subgroups $H \leq G$, there exists a prime p such that $H_{\widehat{p}}$ has a presentation with at most one relator. Then G is either a free or a surface group.*

Proof. Suppose that G is not a free or a surface group; note in addition that under our assumption, G cannot be a free abelian group of rank $n \geq 3$. If G is a limit group, then, by Theorem V.1.7, G contains a subgroup isomorphic to $S * \mathbb{Z}$, where S is the fundamental group of a closed, orientable surface. In particular, G contains $S * S$ as a subgroup. By Theorem II.1.62, G virtually retracts onto $S * S$, so there exists a finite-index subgroup $H \leq G$ and a retraction $r\colon H \longrightarrow S * S$. r induces a retraction at the level of pro-p completions $r\colon H_{\widehat{p}} \longrightarrow (S * S)_{\widehat{p}}$, which yields a retraction of their continuous cohomology groups with trivial $\mathbb{Z}/p$ coefficients $r\colon H^2(H_{\widehat{p}};\mathbb{Z}/p) \longrightarrow H^2((S * S)_{\widehat{p}};\mathbb{Z}/p) \cong (\mathbb{Z}/p)^2$. In particular, the dimension of $H^2(H_{\widehat{p}};\mathbb{Z}/p)$ is at least two and hence $H_{\widehat{p}}$ cannot admit a one-relator presentation for any prime p.

Lastly, if G is not a limit group, then it contains $F \times \mathbb{Z}$, where F is a free group of rank 2. By Theorem II.1.74, $F * \mathbb{Z}$ is a virtual retract of G. As before, we obtain that $\dim_{\mathbb{Z}/p}(H^2(H_{\widehat{p}};\mathbb{Z}/p)) \geq 2$ which completes the proof. □

V.3.1 Nilpotent groups

In this subsection, G denotes a finitely generated nilpotent group and $\mathbf{G}$ denotes a topologically finitely generated nilpotent pro-p group. The following is a classical proposition that encompasses the definition of the *Hirsch length* of a group, along with its properties. For a more detailed account, we refer the reader to Segal's book [103].

Proposition V.3.3. *Let G be a finitely generated nilpotent group. Then the following statements hold:*

(a) There exists a finite subnormal series $1 = G_0 \trianglelefteq G_1 \trianglelefteq \cdots \trianglelefteq G_n = G$ such that each quotient G_{i+1}/G_i is isomorphic to a cyclic group. The number of i's for which G_{i+1}/G_i is infinite is independent of the chosen subnormal series, and is called the Hirsch length *of G. We denote the Hirsch length of G by $h(G)$,*

(b) The Hirsch length is additive in the following sense: for all normal subgroups $N \trianglelefteq G$, we have that $h(G) = h(N) + h(G/N)$,

(c) Let $H \leq G$. Then H has finite index in G if and only if $h(H) = h(G)$.

We next seek to prove a pro-p strengthening of part (c) of Proposition V.3.3. The analogous definition of Hirsch length (along with its properties) carries over to the pro-p setting. More specifically, the classical proof of Proposition V.3.3, that relies on Schreier's Refinement Theorem which states that any two subnormal series of subgroups of a given group have equivalent refinements, can be carried out in the pro-p setting (see for example [75, Lemma 4.3], where the authors prove a version of Schreier's refinement theorem for closed subnormal series of pro-p groups).

Proposition V.3.4. *Let* $\mathbf{G}$ *be a topologically finitely generated nilpotent pro-p group. Then the following statements hold:*

(a) There exists a finite closed subnormal series $1 = \mathbf{G}_0 \trianglelefteq_c \mathbf{G}_1 \trianglelefteq \cdots \trianglelefteq_c \mathbf{G}_n = \mathbf{G}$ *such that each quotient* $\mathbf{G}_{i+1}/\mathbf{G}_i$ *is isomorphic to a pro-p cyclic group. The number of* i*'s for which* $\mathbf{G}_{i+1}/\mathbf{G}_i$ *is infinite is independent of the chosen subnormal series and it is called the Hirsch length of* $\mathbf{G}$*. As in the abstract case, we denote it by* $h(\mathbf{G})$*.*

(b) For every closed normal subgroup $\mathbf{N} \trianglelefteq \mathbf{G}$ *we have that* $h(\mathbf{G}) = h(\mathbf{N}) + h(\mathbf{G}/\mathbf{N})$*.*

(c) Let $\mathbf{H} \leq_c \mathbf{G}$ *be a closed subgroup. Then* $\mathbf{H}$ *has finite index in* $\mathbf{G}$ *if and only if* $h(\mathbf{H}) = h(\mathbf{G})$*.*

The main reason we introduced the Hirsch length is the following proposition, which will help us in detecting when H is a finite-index subgroup of G by looking at finite p-quotients.

Proposition V.3.5. *Let* G *be a finitely generated nilpotent group and let* $H \leq G$ *be a subgroup. Then* H *has a finite index in* G *if and only if the image of the induced map on pro-p completions* $H_{\widehat{p}} \longrightarrow G_{\widehat{p}}$ *has a finite index.*

Proof. The first implication is easy: if the inclusion $\iota\colon H \longrightarrow G$ has finite-index image, then the image of the induced map $\iota_{\widehat{p}}\colon H_{\widehat{p}} \longrightarrow G_{\widehat{p}}$, which is equal to the closure $\overline{H}$ of H in $G_{\widehat{p}}$, has a finite index in $\overline{G} = G_{\widehat{p}}$. To prove the converse, suppose that the induced map on pro-p completions $\iota_{\widehat{p}}\colon H_{\widehat{p}} \longrightarrow G_{\widehat{p}}$ has a finite-index image. Since G' is polycyclic, it admits a finite-index torsion-free subgroup G' [103, Chapter 1]. Consider $H' = H \cap G'$; it is enough to show that H' has finite index in G'. By [87, Theorem C], the induced map $H'_{\widehat{p}} \longrightarrow G'_{\widehat{p}}$ is injective; note that it still has a finite-index image. Hence, by part (c) of Proposition V.3.4, we have that $h(H'_{\widehat{p}}) = h(G'_{\widehat{p}})$. The Hirsch length of a finitely generated torsion-free nilpotent group coincides with that of its pro-p completion (see for example [75, Lemma 5.2] where this is proven in greater generality). Therefore $h(H') = h(H'_{\widehat{p}}) = h(G'_{\widehat{p}}) = h(G')$. Finally, we conclude that $\iota(H) \leq G$ has finite index by part (c) of Proposition V.3.3. □

V.3.2 Pro-p groups with positive rank gradient

It is well-known that an automorphism of a direct product of free groups $F_n \times F_n$ is a direct product of automorphisms of each of the factors (up to permuting the factors). This can be seen by noting that the centraliser of an element $w \in F_n$ is always cyclic, and the centralisers of both $(w,1) \in F_n \times F_n$ and $(1,w) \in F_n \times F_n$ are isomorphic to $\mathbb{Z} \times F_n$. The same argument implies that the same holds for direct products of torsion-free hyperbolic groups.

A different proof can be obtained using Schreier's theorem about normal subgroups of free groups [101], which states that a finitely generated and non-trivial normal subgroup of a free group must be of finite index. Indeed, if $\phi : F_n \times F_n \to F_n \times F_n$ is an automorphism, then the projection of $\phi(F_n \times \{1\})$ onto each of the factors must be either a finite-index subgroup or the trivial group (and the same applies for the projection of $\phi(\{1\} \times F_n)$). One can further verify that $\phi(F_n \times \{1\})$ and $\phi(\{1\} \times F_n)$ must project non-trivially to different factors (otherwise, their intersection in one of the factors would be a finite-index abelian subgroup of F_n) and deduce that $\mathrm{Aut}(F_n \times F_n) \cong (\mathrm{Aut}(F_n) \times \mathrm{Aut}(F_n)) \rtimes (\mathbb{Z}/2\mathbb{Z})$. By employing *rank gradients*, one can prove that a similar phenomenon occurs in the realm of pro-p groups.

Definition V.3.6. Let $\mathbf{G}$ be a pro-p group and let $\mathbf{G} = \mathbf{G}_0 \trianglerighteq \mathbf{G}_1 \trianglerighteq \mathbf{G}_2 \trianglerighteq \cdots$ be a descending chain of normal open subgroups of $\mathbf{G}$ with trivial intersection. The *rank gradient of* $\mathbf{G}$ *relative to* $\{\mathbf{G}_i\}$ is given by

$$\mathrm{RG}(\mathbf{G}; \{\mathbf{G}_i\}) = \lim_{k \to \infty} \frac{d(\mathbf{G}_k) - 1}{|\mathbf{G} : \mathbf{G}_k|}.$$

Remark V.3.7. Note that the limit in Definition V.3.6 above exists since it is the limit of a monotonously decreasing sequence that is bounded from below.

One easily sees that the relative rank gradient of $\mathbf{G}$ does not depend on the choice of a descending chain of open subgroups. We can therefore define:

Definition V.3.8. Let $\mathbf{G}$ be a pro-p group. The *absolute rank gradient* of $\mathbf{G}$ is given by

$$\mathrm{RG}(\mathbf{G}) = \inf_{\mathbf{U} \trianglelefteq_o \mathbf{G}} \frac{d(\mathbf{U}) - 1}{|\mathbf{G} : \mathbf{U}|}.$$

We can extend the notion of a rank gradient to all finitely generated residually-p group G: let $G = G_0 \trianglerighteq G_1 \trianglerighteq G_2 \trianglerighteq \cdots$ be a chain of normal (p-power)-index subgroups such that for every $n \in \mathbb{N}$ there exists N for which $\gamma_n \subseteq G_N$. The *p-gradient* of G, relative to the chain $\{G_i\}$, is given by

$$\mathrm{RG}(G; \{G_i\}) = \lim_{k \to \infty} \frac{\dim_{\mathbb{F}_p} H^1(G_k; \mathbb{F}_p) - 1}{|G : G_k|}.$$

As before, this limit exists and it is independent of the choice of the chain $\{G_i\}$.

The following is a particular instance of [92, Theorem 1.1], which is analogous to the fact that free groups do not have non-trivial infinite-index normal subgroups:

Theorem V.3.9 ([92, Theorem 1.1])**.** *Let* $\mathbf{G}$ *be a finitely generated pro-p group with positive rank gradient and let* $\mathbf{N} \trianglelefteq_c \mathbf{G}$ *be a finitely generated closed normal subgroup. Then* $\mathbf{N}$ *is either finite or open.*

Remark V.3.10. By Lück's approximation [83], if G is a finitely presented residually-p group then $b_1^{(2)}(G) \le \mathrm{RG}(G_{\widehat{p}})$. By [22, Corollary B], if L is a non-abelian limit group then $b_1^{(2)}(L) > 0$, so Theorem V.3.9 above applies to pro-p completions of non-abelian limit groups.

We deduce the following:

Proposition V.3.11. *Let* $\mathbf{G}_i$ *and* $\mathbf{H}_j$ *($1 \le i \le n$ and $1 \le j \le m$) be a collection of finitely generated, torsion-free pro-p groups that have a positive rank gradient. Suppose that there is an isomorphism of pro-p groups*

$$\phi \colon \mathbf{G}_1 \times \cdots \times \mathbf{G}_n \longrightarrow \mathbf{H}_1 \times \cdots \times \mathbf{H}_m.$$

Then $n = m$ *and* ϕ *is a direct product of isomorphisms, that is, there exist a permutation* $\sigma \in \mathrm{Sym}(n)$ *and isomorphisms* $\phi_i \colon (G_i)_{\widehat{p}} \longrightarrow (H_{\sigma(i)})_{\widehat{p}}$ *for* $1 \le i \le n$ *such that*

$$\phi = \left(\phi_{\sigma^{-1}(1)} \circ \pi_{\sigma^{-1}(1)}\right) \times \cdots \times \left(\phi_{\sigma^{-1}(n)} \circ \pi_{\sigma^{-1}(n)}\right).$$

Proof. By assumption, we know that each of the $\mathbf{G}_i$'s and $\mathbf{H}_j$'s is infinite. For both $\mathbf{G}_1 \times \cdots \times \mathbf{G}_n$ and $\mathbf{H}_1 \times \cdots \times \mathbf{H}_m$, we denote by π_i the projection onto the i-th coordinate.

For each $1 \le i \le n$, there exists $1 \le j \le m$ such that $\pi_j(\phi(\mathbf{G}_i))$ is an infinite subgroup of $\mathbf{H}_j$. A priori, there may be multiple choices of such a j for a given i; we choose a single such $j = j(i)$ for every $1 \le i \le n$. Note that $\pi_{j(i)}(\phi(\mathbf{G}_i))$ is normal in $H_{j(i)}$. Therefore, by Theorem V.3.9, $\pi_{j(i)}(\phi(\mathbf{G}_i))$ is open in $\mathbf{H}_{j(i)}$. In particular, it follows that we cannot have that $j(i_1) = j(i_2)$ for different $i_1 \neq i_2$: suppose for a contradiction that $j = j(i_1) = j(i_2)$. Therefore $\mathbf{H}_j$ contain two open commuting subgroups $\pi_{j(i_1)}(\phi(\mathbf{G}_{i_1}))$ and $\pi_{j(i_2)}(\phi(\mathbf{G}_{i_2}))$ that intersect in an open subgroup. It follows that $\mathbf{H}_j$ is infinite and virtually abelian (as it contains the abelian group $\pi_{j(i_1)}(\phi(\mathbf{G}_{i_1})) \cap \pi_{j(i_2)}(\phi(\mathbf{G}_{i_2}))$ as a subgroup of finite index), but such groups have a rank gradient that is equal to 0, which is a contradiction. Therefore $j(i)$ is an injection $\{1, \dots, n\} \longrightarrow \{1, \dots, m\}$ and $n \le m$. Reasoning analogously for the inverse map ϕ^{-1}, we derive that $m \le n$. Hence $m = n$, and so $j(i)$ is a bijection. In fact, for the exact same reasons we must have that $\pi_j(\phi(\mathbf{G}_i)) = 1$ for every $1 \le i \le n$ and $j \ne j(i)$.

Lastly, since ϕ is an isomorphism, $\pi_{j(i)}(\phi(\mathbf{G}_i)) = \mathbf{H}_{j(i)}$. Taking $\sigma(i) = j(i)$ and $\phi_i = \pi_{j(i)} \circ \phi|_{\mathbf{G}_i}$ completes the proof. $\square$

We would like to establish a similar result for direct products of finitely presented residually-p groups (and in particular, non-abelian limit groups). To do so, we begin by proving a duo of lemmas:

Lemma V.3.12. *Let $n \ge 1$ and let $K_1, \dots, K_n$ be finitely generated residually-p groups with a positive p-gradient. Let $K = K_1 \times \dots \times K_n$ and suppose that K' is an intermediate group $K \le K' \le K_{\hat{p}}$ such that the index $|K' : K| < \infty$. Then $K = K'$.*

Proof. We prove the lemma by induction. For $n = 1$, since K' is dense in $K_{\hat{p}}$, $\mathrm{RG}(K') \ge \mathrm{RG}(K)$. In addition, we have that $\mathrm{RG}(K) = |K' : K| \cdot \mathrm{RG}(K')$ and $\mathrm{RG}(K) > 0$ so $|K' : K| = 1$ as desired. Suppose now that $n \ge 2$, and consider the direct product decomposition $K_{\hat{p}} \cong (K_1)_{\hat{p}} \times \dots \times (K_n)_{\hat{p}}$. Denote by $\pi_1 \colon K_{\hat{p}} \longrightarrow (K_1)_{\hat{p}}$ and $\pi^1 \colon K_{\hat{p}} \longrightarrow (K_2)_{\hat{p}} \times \dots \times (K_n)_{\hat{p}}$ the canonical projections. Consider the trio $K_1 \le \pi_1(K' \cap (K_1)_{\hat{p}}) \le (K_1)_{\hat{p}}$ and apply the base step of the induction to obtain that $K' \cap (K_1)_{\hat{p}} = K_1$. Lastly, consider $K_2 \times \dots \times K_n \cong K/K_1 \le K'/K_1 \le K_{\hat{p}}/(K_1)_{\hat{p}} \cong (K/K_1)_{\hat{p}}$; as before we obtain that $K/K_1 = K'/K_1$ and thus $K = K'$. $\square$

We next state a sufficient condition that ensures that direct factors can be recognised by looking at the pro-p completion.

Lemma V.3.13. *Let $n \ge 1$ and let $H_1, \dots, H_n$ be finitely generated residually-p groups with a positive p-gradient. Let H be a finitely generated group with $H_{\hat{p}} \cong (H_1)_{\hat{p}} \times \dots \times (H_n)_{\hat{p}}$. Assume furthermore that the subgroup of H generated by the intersections $H \cap (H_i)_{\hat{p}}$ has a finite index in H. Then $H = (H \cap (H_1)_{\hat{p}}) \times \dots \times (H \cap (H_n)_{\hat{p}})$.*

Proof. We show by induction that $H \cap (H_i)_{\hat{p}}$ is dense in $(H_i)_{\hat{p}}$ for all i and that $H = (H \cap (H_1)_{\hat{p}}) \times \dots \times (H \cap (H_n)_{\hat{p}})$. The base case $n = 1$ is trivial. Let $n \ge 2$ and suppose that the claim holds for $n-1$. By our assumption, the subgroup $\langle H \cap (H_1)_{\hat{p}}, \dots, H \cap (H_n)_{\hat{p}} \rangle$

has a finite index in H and hence $\langle \overline{H \cap (H_1)_{\widehat{p}}}, \dots, \overline{H \cap (H_n)_{\widehat{p}}} \rangle$ is open in $H_{\widehat{p}}$. Hence each $\overline{H \cap (H_i)_{\widehat{p}}}$ is open in $(H_i)_{\widehat{p}}$.

Consider the group $H^1 = H / H \cap (H_1)_{\widehat{p}}$. The natural map

$$H^1 \longrightarrow \Big((H_1)_{\widehat{p}} / \overline{H \cap (H_1)_{\widehat{p}}} \Big) \times (H_2)_{\widehat{p}} \times \cdots \times (H_n)_{\widehat{p}}$$

is the natural injection of H^1 into its pro-p completion. By assumption, the first factor

$$(H_1)_{\widehat{p}} / \overline{H \cap (H_1)_{\widehat{p}}}$$

is a finite p-group and its intersection with H^1 is trivial. In particular, $\mathbf{U} = (H_2)_{\widehat{p}} \times \cdots \times (H_n)_{\widehat{p}}$ is an open subgroup of $(H^1)_{\widehat{p}}$. Suppose that $(H_1)_{\widehat{p}} / \overline{H \cap (H_1)_{\widehat{p}}}$ is of cardinality $k \geq 1$ and let $H^1_0 \leq H^1$ be the finite-index subgroup $H^1_0 = H^1 \cap \mathbf{U}$. Denote by $\pi^1 \colon (H^1)_{\widehat{p}} \longrightarrow (H_2)_{\widehat{p}} \times \cdots \times (H_n)_{\widehat{p}}$ the canonical projection. Note that $|H^1 \colon H^1_0| = k$, and that both H^1 and H^1_0 project injectively to $(H_2)_{\widehat{p}} \times \cdots \times (H_n)_{\widehat{p}} \cong (H^1_0)_{\widehat{p}}$. By the induction hypothesis, we also have that $H^1_0 = K_2 \times \cdots \times K_n$ with $(K_i)_{\widehat{p}} \cong (H_i)_{\widehat{p}}$ (where $K_i = H^1_0 \cap (H_i)_{\widehat{p}}$). So, by Lemma V.3.12, we have that $k = 1$, implying that $H^1 = H^1_0 = (H \cap (H_2)_{\widehat{p}}) \times \cdots \times (H \cap (H_n)_{\widehat{p}})$ and that $H \cap (H_1)_{\widehat{p}}$ is dense in $(H_1)_{\widehat{p}}$. The conclusion follows. □

We seal the discussion with the following proposition that will allow us to recover direct product decompositions of certain finitely generated residually-p groups from their finite p-quotients; this will serve as a key ingredient in the proof of Theorem N.

Proposition V.3.14. *Let $n \geq 1$ and let $H_1, \dots, H_n$ be finitely generated residually-p groups with a positive p-gradient. Suppose that H is a torsion-free and residually-p finitely generated group. Suppose further that $H_{\widehat{p}} \cong (H_1)_{\widehat{p}} \times \cdots \times (H_n)_{\widehat{p}}$ and that there are pairwise-commuting infinite subgroups $N_i \leq H$ such that the natural map $N_1 \times \cdots \times N_n \longrightarrow \langle N_1, \dots, N_n \rangle$ is an isomorphism. If $\langle N_1, \dots, N_n \rangle$ has finite index in H, then there exist subgroups $N_i' \leq H$, each commensurable to N_i respectively, such that $H = N_1' \times \cdots \times N_n'$.*

Proof. Replacing each N_i by one of its finite-index subgroups, we can assume that each N_i is normal in H. For each $1 \leq j \leq n$, denote by $\pi_j \colon H_{\widehat{p}} \longrightarrow (H_j)_{\widehat{p}}$ the canonical projection. Since each N_i is infinite, each of the closures $\overline{N_i}$ is an infinite normal subgroup of $H_{\widehat{p}}$. Hence, for every i there exists $j = j(i)$ such that $\pi_j(\overline{N_i})$ is an infinite normal subgroup of $(H_j)_{\widehat{p}}$. By Theorem V.3.9, $\pi_j(\overline{N_i})$ must have a finite index in $(H_j)_{\widehat{p}}$. Proceeding exactly as in the proof of Proposition V.3.11, we have that for every i, the choice of $j = j(i)$ is unique. For every i and $k \neq j(i)$, the subgroup $\pi_k(\overline{N_i}) \leq (H_k)_{\widehat{p}}$ is finite.

Replacing again each N_i by smaller finite-index subgroups, we can suppose that for every i and $k \neq j(i)$, $\pi_k(\overline{N_i}) = 1$. Therefore $N_i \leq H \cap (H_{j(i)})_{\widehat{p}}$. Since the subgroup generated by all the N_i is still of finite-index in H, we can apply Lemma V.3.13 and obtain that $H = N_1' \times \cdots \times N_n'$ where $N_i' = H \cap (H_{j(i)})_{\widehat{p}}$. ☐

V.3.3 Profinite completions of limit groups

In this subsection we gather a few results about profinite completions of limit groups; these will be used in the proof of Theorem N. Zalesskii and Zapata employ the hierarchical structure of limit groups to show that if L is a non-abelian limit group then $\widehat{L}$ acts faithfully and irreducibly on a profinite tree. They further use this action to establish a number of properties (that were known for abstract limit groups) of profinite completions of limit groups.

Proposition V.3.15 (Corollary 4.4, [124])**.** *Let L be a non-abelian limit group. Then $\widehat{L}$ is centreless.*

The following consequence of Proposition V.3.15 will also be important in the proof of Theorem N.

Proposition V.3.16. *Let L be a non-abelian limit group. Suppose that there are two closed subgroups $\mathbf{A}$ and $\mathbf{B}$ of $\widehat{L}$ with the following two properties:*

- *Every $a \in \mathbf{A}$ and $b \in \mathbf{B}$ commute,*
- *For some prime p, there are embeddings $\mathbb{Z}_p \leq_c \mathbf{A}$ and $\mathbb{Z}_p \leq_c \mathbf{B}$.*

Then the closed subgroup generated by $\mathbf{A}$ and $\mathbf{B}$ in $\widehat{L}$ is not open.

Proof. We argue by contradiction. Suppose that the closed subgroup $\langle \mathbf{A}, \mathbf{B} \rangle$ is open. Then it is isomorphic to $\widehat{L_1}$ for some finite-index subgroup $L_1 \leq_o L$; L_1 must therefore be a non-abelian limit group. Hence we can assume, without loss of generality, that $\langle \mathbf{A}, \mathbf{B} \rangle = \widehat{L}$. By Proposition V.3.15, $\widehat{L}$ is centreless, which implies that $\mathbf{A} \cap \mathbf{B} = \{1\}$. So $\widehat{L} \cong \mathbf{A} \times \mathbf{B}$ and by [99, Proposition 4.2.4] $\widehat{L}$ is projective. However, since $\mathbb{Z}_p$ embeds in both $\mathbf{A}$ and $\mathbf{B}$, $\mathbb{Z}_p^2$ embeds in $\widehat{L}$, which contradicts the projectivity of $\widehat{L}$. ☐

V.3.4 Proof of Theorem N

We reformulate Theorem N to simplify the notation of the proof. The formulations are equivalent because any product of free or surface groups has the form $\Gamma = S_1 \times \cdots \times S_n \times \mathbb{Z}^k$ with $n \geq 1$, $k \geq 0$ and all of the S_i being non-abelian free or hyperbolic surface groups.

Theorem V.3.17 (Theorem N)**.** *Let $n \geq 1$ and $k \geq 0$ be integers and let G be a finitely presented residually free group. Let $S_1, \ldots, S_n$ be non-abelian free or surface groups and let Γ be the direct product $S_1 \times \cdots \times S_n \times \mathbb{Z}^k$. If $\widehat{G} \cong \widehat{\Gamma}$, then $G \cong \Gamma$.*

For the convenience of the reader, we give a brief overview of the different steps of the proof of Theorem N.

(Step 1) Instead of working with G, we work with $G/Z(G)$ which has a trivial centre. This case will be enough as we explain in Claim V.3.25.

(Step 2) We view G as the subdirect product of n non-abelian limit groups $L_1 \times \cdots \times L_n$ (see Corollary II.1.68); a priori, G might be a subdirect product of m limit groups (with $m \neq n$), but we show in Claims V.3.19 and V.3.20 that m must be equal to n.

(Step 3) Using results of Bridson, Howie, Miller and Short regarding the structure of finitely presented subgroups of direct products of limit groups (see Subsection II.1.6), in combination with properties of Hirsch lengths in nilpotent groups (see Section V.3.1), we show in Claim V.3.22 that G must have a finite index in $L_1 \times \cdots \times L_n$.

(Step 4) This clearly shows that G has a finite-index subgroup H that is a direct product of n limit groups. Since each L_i has a positive p-gradient, we can apply the results from Subsection V.3.2 to lift the direct product decomposition structure of H to G. This is done in Claim V.3.23.

(Step 5) We obtain that $G = G_1 \times \cdots \times G_n$, with each G_i a limit group. We show in Claim V.3.24 that each G_i must be a free or surface group using Theorem V.3.2.

As mentioned above, we begin by working with $G_0 = G/Z(G)$ instead of G; this group is indeed centreless:

Lemma V.3.18 ([9], Lemma 4)**.** *Let G be a residually free group. Then $G/Z(G)$ is a residually free group with a trivial centre.*

Note that the profinite completion of G_0 is $\widehat{G_0} \cong \widehat{G}/\overline{Z(G)}$. Since each $\widehat{S}_i$ is centreless, we have that $Z(\widehat{G}) = \widehat{\mathbb{Z}}^k$ and hence $\widehat{G_0} \cong \widehat{S}_1 \times \cdots \times \widehat{S}_n \times \mathbf{A}$, where $\mathbf{A} = Z(\widehat{G})/\overline{Z(G)}$ is an abelian profinite group.

Claim V.3.19. *G_0 is finitely presented and residually free, and its profinite completion is isomorphic to $\widehat{G_0} \cong \widehat{S}_1 \times \cdots \times \widehat{S}_n$.*

Proof. By Lemma V.3.18 and the fact that $Z(G)$ is finitely generated, we have that G_0 is still a finitely presented residually free group. By Corollary II.1.68, G_0 is a subdirect product $G_0 \hookrightarrow L_1 \times \cdots \times L_m$ of m limit groups, for some m. We pick m to be minimal. Hence, in particular, G_0 intersects non-trivially each factor L_i. Since G_0 is centreless, each L_i is non-abelian. Finally, by Proposition V.3.15, $\widehat{L}_i$ is also centreless.

For each $1 \leq j \leq m$, denote by $\pi_j \colon L_1 \times \cdots \times L_m \longrightarrow L_j$ the projection onto the j-th coordinate. By Theorem II.1.74, the induced map of profinite completions $\widehat{G_0} \hookrightarrow \widehat{L_1} \times \cdots \times \widehat{L_m}$ is injective. Since the maps $\widehat{\pi_j} \colon \widehat{G_0} \longrightarrow \widehat{L}_j$ are surjective, and since each $\widehat{L}_j$ is centreless, it follows that $\widehat{\pi_j}(\mathbf{A}) = 1$ for all j, implying that $\mathbf{A} = 1$. ☐

Claim V.3.20. *Keeping the above notation, $n = m$.*

Proof. Since limit groups are torsion-free and $G \cap L_j \neq 1$ for each $1 \leq j \leq m$, there exists an injection $\mathbb{Z}^m \hookrightarrow G_0$. By Theorem II.1.74, this yields an injection in the level of profinite completions $\widehat{\mathbb{Z}}^m \hookrightarrow \widehat{G_0} \cong \widehat{S}_1 \times \cdots \times \widehat{S}_n$, which, in combination with the standard fact that no $\widehat{S}_i$ contains $\mathbb{Z}_p^2$, implies that $m \leq n$.

For the other direction, that is that $n \leq m$, we argue by contradiction. Assume that $n \geq m + 1$. For each $1 \leq i \leq n$, consider a closed infinite pro-p cyclic subgroup $\mathbb{Z}_p \cong \mathbf{Z}_i \leq \widehat{S}_i$. Since $n \geq m + 1$, there must exist two different $\mathbf{Z}_{k_1}$ and $\mathbf{Z}_{k_2}$ such that both of their images $\widehat{\pi_j}(\mathbf{Z}_{k_1})$ and $\widehat{\pi_j}(\mathbf{Z}_{k_2})$ in some $\widehat{L}_j$ are isomorphic to $\mathbb{Z}_p$. However, this would contradict Proposition V.3.16: in this case, the two pairwise commuting closed groups $\widehat{\pi_j}(\widehat{S}_{k_1})$ and $\widehat{\pi_j}(\prod_{i \neq k_1} \widehat{S}_i)$ would both contain a copy of $\mathbb{Z}_p$ and would generate $\widehat{L}_j$. This contradiction shows that $n \leq m$, completing the proof. ☐

We introduce futher notation required for the proof of Claim V.3.22, which states that the injection $G_0 \hookrightarrow L_1 \times \cdots \times L_m$ has a finite-index image. By Theorem II.1.73, there exists $N \in \mathbb{N}$, a finite-index subgroup $E \leq L_1 \times \cdots \times L_m$ and a finite-index subgroup

$G_1 \leq G_0$ such that $\gamma_N E \leq G_1$. Notice that the subgroup $E_1 = (E \cap L_1) \times \cdots \times (E \cap L_n) \leq E$ has finite index in E. We define

$$G_2 = E_1 \cap G_1$$

and note that G_2 has a finite index in G and that $\gamma_N E_1 \leq E_1 \cap G_1 = G_2$. We further denote

$$U_i = E \cap L_i,$$

and each U_i is a limit group. The projection of the direct product $U_1 \times \cdots \times U_n$ onto its i-th coordinate will still be denoted by π_i. The closure of G_2 inside $\widehat{G}_0$ is an open subgroup $\overline{G_2} \leq_o \widehat{G}_0$. Denote

$$\mathbf{K}_i = \overline{G_2} \cap \widehat{S}_i,$$

again, we have that $\mathbf{K} = \mathbf{K}_1 \times \cdots \times \mathbf{K}_n$ is an open subgroup of $\overline{G_2}$.

We define

$$G_3 = \mathbf{K} \cap G_0;$$

one has that $\widehat{G}_3 \cong \mathbf{K}$.

Summarising,

- $G_3 \leq G_0$ has a finite index and $\widehat{G}_3 \cong \mathbf{K}_1 \times \cdots \times \mathbf{K}_n$,
- $E_1 = U_1 \times \cdots \times U_n$, and there is an injection $f\colon G_3 \hookrightarrow E_1$ such that each $\pi_j(G_3)$ has finite index in U_j.

Lastly, since each $\mathbf{K}_i$ is open in $\widehat{S}_i$, they are profinite completions of non-abelian free or surface groups. We denote by $(\mathbf{K}_i)_p$ the maximal pro-p quotient of the profinite group $\mathbf{K}_i$. In fact, $(\mathbf{K}_i)_p$ is the pro-p completion of a non-abelian free or surface group and

$$(G_3)_{\widehat{p}} \cong (\mathbf{K}_1)_p \times \cdots \times (\mathbf{K}_n)_p.$$

Claim V.3.21. *The induced map in the level of pro-p completions,* $f_{\widehat{p}}\colon (G_3)_{\widehat{p}} \longrightarrow (E_1)_{\widehat{p}}$, *has a finite-index image.*

Proof. By the proof of Claim V.3.20, there exists an injection $\mathbb{Z}^n \hookrightarrow G_0$; therefore, there exists an injection $\mathbb{Z}^n \hookrightarrow G_3$. By [87, Theorem 7.12], the induced maps on pro-p completions,

$$\mathbb{Z}_p^n \hookrightarrow (G_3)_{\widehat{p}} \quad \text{and} \quad \mathbb{Z}_p^n \hookrightarrow (E_1)_{\widehat{p}},$$

are injective. Since each $(\mathbf{K}_i)_p$ is the pro-p completion of a non-abelian free or surface group, no $(\mathbf{K}_i)_p$ contains a copy of $\mathbb{Z}_p^2$. This implies that each intersection $\mathbb{Z}_p^n \cap (\mathbf{K}_i)_p$

is infinite. In particular, since the restriction of $f_{\widehat{p}}\colon (G_3)_{\widehat{p}} \longrightarrow (E_1)_{\widehat{p}}$ to $\mathbb{Z}_p^n$ is injective, the restriction of $f_{\widehat{p}}$ to each i-th coordinate $(\mathbf{K}_i)_p$ of $(G_3)_{\widehat{p}}$ has an infinite image.

For each $1 \leq i \leq n$, there exists $j = j(i)$ such that

$$(\pi_j \circ f)_{\widehat{p}}((\mathbf{K}_i)_p) \leq (U_j)_{\widehat{p}}$$

is infinite. For each i there may be, a priori, several j with this property; we fix a choice of $j = j(i)$ for each i. Since H_i is normal in G_3, $(\pi_{j(i)} \circ f)_{\widehat{p}}((\mathbf{K}_i)_p)$ is normal in the finite-index subgroup $(\pi_{j(i)} \circ f)_{\widehat{p}}((G_3)_{\widehat{p}})$ of $(U_{j(i)})_{\widehat{p}}$. By Theorem V.3.9, this implies that $(\pi_{j(i)} \circ f)_{\widehat{p}}((\mathbf{K}_i)_p)$ has finite index in $(U_{j(i)})_{\widehat{p}}$. With this observation we conclude that for $i_1 \neq i_2$, we must have $j(i_1) \neq j(i_2)$: if we had $j = j(i_1) = j(i_2)$ for $i_1 \neq i_2$, then the open commuting subgroups $(\pi_j \circ f)_{\widehat{p}}((\mathbf{K}_{i_1})_p)$ and $(\pi_j \circ f)_{\widehat{p}}((\mathbf{K}_{i_2})_p)$ of $(U_j)_{\widehat{p}}$ would intersect in an open abelian subgroup. This contradicts the fact that $(U_j)_{\widehat{p}}$ is not virtually abelian (because U_j is not virtually abelian). Thus, for each i there exists a unique choice of $j(i)$ with $(\pi_j \circ f)_{\widehat{p}}((\mathbf{K}_i)_p) \leq (U_j)_{\widehat{p}}$ infinite, yielding a bijection $j\colon \{1, \ldots, n\} \longrightarrow \{1, \ldots, n\}$. Furthermore, for each i and $k \neq j(i)$, $(\pi_k \circ f)_{\widehat{p}}((\mathbf{K}_i)_p) \leq (U_k)_{\widehat{p}}$ is finite. Recall that $(\pi_{j(i)} \circ f)_{\widehat{p}}((\mathbf{K}_i)_p)$ has finite index in $(U_{j(i)})_{\widehat{p}}$; this implies that $f_{\widehat{p}}((\mathbf{K}_i)_p) \cap (U_{j(i)})_{\widehat{p}}$ is of finite index, and therefore open, in $(U_{j(i)})_{\widehat{p}}$. We conclude that $f_{\widehat{p}}$ has a finite-index image as desired. □

From this we deduce:

Claim V.3.22. *The injection $G_0 \hookrightarrow L_1 \times \cdots \times L_m$ has a finite-index image.*

Proof. Since $G_3 \leq G_2 \leq G_0$ and $E_1 \leq E$ are finite-index subgroups, it suffices to verify that the injection $G_2 \hookrightarrow E_1$ has a finite-index image. Since $\gamma_N E_1 \leq G_2$, we can instead consider the induced injection of finitely generated nilpotent groups $h\colon G_2/\gamma_N E_1 \hookrightarrow E_1/\gamma_N E_1$. By Claim V.3.21, $f_{\widehat{p}} : (G_3)_{\widehat{p}} \longrightarrow (E_1)_{\widehat{p}}$ has a finite-index image, and so does $h_{\widehat{p}}$. Thus, since $G_2\gamma_N E_1$ and $E_1/\gamma_N E_1$ are finitely generated nilpotent groups, Proposition V.3.5 implies that h has a finite-index image. □

One easily concludes that G_0 admits a direct product decomposition:

Claim V.3.23. *There exist subgroups $G_1, \ldots, G_n \leq G_0$, each with a positive p-gradient, such that $G_0 = G_1 \times \cdots \times G_n$.*

Proof. By Claim V.3.22, the $(G_0 \cap L_1) \times \cdots \times (G_0 \cap L_m)$ has a finite index in G_0 and each $G_0 \cap L_i$ is a non-abelian limit group. By Remark V.3.10, each $G_0 \cap L_i$ has a positive p-gradient. The desired conclusion follows directly from Proposition V.3.14. □

Theorem N now follows. We first deduce it for G_0 (or, equivalently, under the assumption that G is centreless), from which the general case easily follows.

Claim V.3.24 (Theorem N for a centreless group)**.** $G_0 \cong S_1 \times \cdots \times S_n$.

Proof. By Claims V.3.19 and V.3.23 we have that:

- $\widehat{G}_0 \cong \widehat{S}_1 \times \cdots \times \widehat{S}_n$,
- there exist subgroups $G_1, \dots, G_n \leq G_0$, each with a positive p-gradient, such that $G_0 = G_1 \times \cdots \times G_n$.

We know that free and surface groups are distinguished from each other by their pro-p completion; this can be easily seen by, for example, looking at their abelianizations. Combining this observation with Proposition V.3.11, it is enough to show that each G_i is either a free or a surface group. For simplicity, we will verify this for G_1; the same argument yields that every other G_i must also be a free or a surface group.

By Theorem V.3.2, it is enough to show that for every finite-index subgroup $K_1 \leq G_1$, the pro-p completion $(K_1)_{\widehat{p}}$ is the pro-p completion of either a free or a surface group. Let K_1 be a finite-index subgroup of G_1. The subgroup $K_1 \times G_2 \times \cdots \times G_n \leq G$ has finite index in G and so it has the same profinite completion as a finite-index subgroup H of $S_1 \times \cdots \times S_n$. Note that $(H \cap S_1) \times \cdots \times (H \cap S_n)$ has a finite index in H, and we can therefore apply Proposition V.3.14 to H along with $N_i = H \cap S_i$. We obtain subgroups N_i' of H, each commensurable to $H \cap S_i$ respectively, such that

$$H \cong N_1' \times \cdots \times N_n'.$$

From this, we get that

$$(K_1)_{\widehat{p}} \times (G_2)_{\widehat{p}} \times \cdots \times (G_n)_{\widehat{p}} \cong H_{\widehat{p}} \cong (N_1')_{\widehat{p}} \times \cdots \times (N_n')_{\widehat{p}}.$$

Each N_i' is torsion-free, and commensurable to either a free or a surface group; by Nielsen realisation, each N_i' is itself either a free or a surface group. By Proposition V.3.11, there exists $1 \leq i \leq n$ such that $(K_1)_{\widehat{p}} \cong (N_i')_{\widehat{p}}$, as required. □

Claim V.3.25 (Theorem N in the general case)**.** *There is an isomorphism* $G \cong \Gamma$.

Proof. By Claim V.3.24, it is enough to show that $G = G_0 \times \mathbb{Z}^k$. G_0 and G fit into a short exact sequence

$$1 \longrightarrow \mathbb{Z}^k \longrightarrow G \longrightarrow G_0 \longrightarrow 1 \qquad \text{(V.1)}$$

that induces a short exact sequence in the level of profinite completions:

$$1 \longrightarrow \widehat{\mathbb{Z}}^k \longrightarrow \widehat{G} \longrightarrow \widehat{G}_0 \longrightarrow 1.$$

Note that the short exact sequence above splits since

$$\widehat{G} \cong \widehat{\Gamma} \cong \widehat{S_1} \times \cdots \times \widehat{S_n} \times \widehat{\mathbb{Z}}^k \cong \widehat{G}_0 \times \widehat{\mathbb{Z}}^k$$

and since every short exact sequence of profinite groups in which the middle group is isomorphic to the direct product of the other groups splits. By [122, Lemma 8.3] and the goodness of G_0, we have the short exact sequence in (V.1) splits. Hence $G \cong \mathbb{Z}^k \times G_0$ finishing the proof. □

V.4 Other calculations of virtual Betti numbers and further questions

V.4.1 Virtual homology of RAAGs and manifolds

We compute the virtual homology of other classes of groups, including right-angled Artin groups and fundamental groups of closed, hyperbolic 3 manifolds. These calculations are simple and are probably well-known to experts.

Proposition V.4.1. *Let G be a RAAG, let k be a field and let $n \in \mathbb{N}$. Then* $\mathrm{vb}_n(G) < \infty$ *if and only if either* $\mathrm{cd}(G) < n$ *or G is free abelian.*

Notation V.4.2. Recall that we denote the right-angled Artin group associated to the graph Γ by $G(\Gamma)$.

Proof. If $\mathrm{cd}(G) < n$ then $\mathrm{vb}_n^k(G) = 0$. If $G \cong \mathbb{Z}^m$, then $\mathrm{vb}_n^k(G) = \binom{m}{n} < \infty$. So one direction is clear.

Let Γ be the underlying graph of $G = G(\Gamma)$. Let N denote the cohomological dimension of G, which is equal to the size of the biggest clique in Γ. Suppose that $N \geq n$, and that G is not free abelian. We will show that $\mathrm{vb}_n^k(G) = \infty$.

There exists a subgraph Γ_0 with $N+1$ vertices which is not complete but contains a complete subgraph K_n on n vertices. Since $N \geq n$, it follows that Γ_0 contains a further subgraph Γ_1 with $n+1$ vertices that is not complete but contains a complete subgraph K_n on n vertices. The group $G(\Gamma_1$ is a retract of $G = G(\Gamma)$, so $\mathrm{vb}_n^k(G(\Gamma_1)) \leq \mathrm{vb}_n^k(G)$. Hence it suffices to show that $\mathrm{vb}_n^k(G(\Gamma_1)) = \infty$. Note that the group $G(\Gamma_1)$ splits as an HNN extension $\mathbb{Z}^n *_{\mathbb{Z}^m}$ with $m < n$. We apply Lemma V.2.1 and conclude that $\mathrm{vb}_n^k(G) = \infty$. □

Remark V.4.3. Alan Reid pointed out that the fundamental group of an arithmetic hyperbolic n-manifold M of simplest type has $\mathrm{vb}_i^{\mathbb{Z}}(\pi_1(M)) = \infty$ for $1 \leq i < n$. This can be proven by induction as follows: by Agol's work [2, Theorem 9.2], the assertion holds for $n \leq 3$. For $n \geq 4$, we use the fact that M has an immersed totally geodesic submanifold M' of dimension $n - 1$. By the induction hypothesis, $\mathrm{vb}_i^{\mathbb{Z}}(\pi_1(M')) = \infty$ for $1 \leq i < n - 1$. All immersed totally geodesic submanifolds contribute to homology in finite covers by [14], which implies that $\mathrm{vb}_i^{\mathbb{Z}}(\pi_1(M)) = \infty$ for $1 \leq i < n - 1$. Finally, by Poincaré duality, we also have that $\mathrm{vb}_{n-1}^{\mathbb{Z}}(\pi_1(M)) = \infty$.

V.4.2 Further questions and virtual invariants

Theorem I shows that when an HGFC-group G has $\mathrm{vb}_2^k(G) < \infty$, then G is either a free or a surface group. This suggests the following:

Question V.4.4. *Let G be a hyperbolic group with cohomological dimension two. If G has finite virtual second Betti number, does it follow that G is a surface group?*

We remark that this question is closely related to the two most notorious questions about hyperbolic groups (the first asks whether every hyperbolic group is residually finite, and the second asks whether every one-ended hyperbolic group contains a surface subgroup). Therefore this question is particularly ambitious, and an answer is currently far out of reach.

Another property related to virtual homology is *property* (TFab):

Definition V.4.5. A group G has *property* (TFab) if the abelianisation of every finite-index subgroup $H \leq G$ is torsion-free.

Question V.4.6. *Let G be a one-ended limit group with property* (TFab)*, that is, all of the finite-index subgroups of G have a torsion-free abelianisation. Does it follow that G is either a free abelian group or a surface group?*

Remark V.4.7. We remark that computational evidence suggests that the answer to Question V.4.6 is positive.

At this point we have seen many examples of groups, some of which admitting interesting hierarchies, where the boundedness of vb_2 (or higher virtual homology notions) is very exceptional. In particular, most limit groups (with the only exceptions being free, free abelian and surface groups) have unbounded vb_2. We would therefore like to introduce another notion which might aid in obtaining a finer classification of limit groups based on virtual homological invariants.

Definition V.4.8. Let G be a finitely generated group. Define the *virtual i-th cohomology spectrum* of G, with coefficients in a field k, to be the set

$$\mathrm{Spectrum}_i^k(G) = \{\dim_k H^i(H; k) | H \leq G \text{ of finite index}\}.$$

Cohomological goodness implies that if G is a limit group or an HGFC-group,

$$\mathrm{Spectrum}_i^{\mathbb{F}_p}(\widehat{G}) = \mathrm{Spectrum}_i^{\mathbb{F}_p}(G).$$

In other words, the virtual i-th cohomology spectrum with coefficients in $\mathbb{F}_p$ is a profinite invariant of G within good groups. This raises the following question:

Question V.4.9. *Which numbers appear in the virtual i-th cohomology spectrum of a given limit group or a given HGFC-group?*

One can even strengthen the notion of a virtual cohomology spectrum, and define the *filtered virtual i-th cohomology spectrum* of G, with coefficients in a field k, to be

$$\mathcal{F}\mathrm{Spectrum}_i^k(G) = \left\{ \frac{\dim_k H^i(H; k)}{[G : H]} \middle| H \leq G \text{ of finite index} \right\}.$$

This refined version also keeps track of the depth of the finite-sheeted coverings of G in which large i-th homology is exhibited. Again the filtered virtual i-th cohomology spectrum of G (with coefficients in $\mathbb{F}_p$) is a profinite invariant of good groups.

Question V.4.10. *Which limit groups and HGFC-groups are characterized by their filtered second Betti spectrums?*

The proof of Theorem I relies heavily on virtual retractions (see Section II.1.74). Some information is lost in the process of transferring homological data via virtual retractions. In particular, answering questions V.4.9 and V.4.10 should involve methods different to the ones used in this paper. Answering these questions would also, hopefully, produce more examples of groups that are profinitely rigid among limit groups and HGFC-groups.

Bibliography

[1] I. Agol. Virtual betti numbers of symmetric spaces, 2006. arXiv:math/0611828.

[2] I. Agol. The virtual Haken conjecture. *Documenta Mathematica*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[3] S. André. Hyperbolicity and cubulability are preserved under elementary equivalence. *Geometry & Topology*, 24(3):1075–1147, September 2020.

[4] S. André and J. Fruchter. Formal solutions and the first-order theory of acylindrically hyperbolic groups. *Journal of the London Mathematical Society*, 105(2):1012–1072, February 2022.

[5] Y. Antolin, A. Minasyan, and A. Sisto. Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups Geom. Dyn.*, 10(4):1149–1210, 2016.

[6] G. N. Arzhantseva. *Geometriae Dedicata*, 87(1/3):191–208, 2001.

[7] A. Bartels and W. Lück. The Borel Conjecture for hyperbolic and CAT(0)-groups. *Annals of Mathematics*, 175(2):631–689, 2012.

[8] A. Baudisch. On superstable groups. *Journal of the London Mathematical Society*, s2-42(3):452–464, December 1990.

[9] B. Baumslag. Residually free groups. *Proc. London Math. Soc. (3)*, 17:402–418, 1967.

[10] G. Baumslag. On generalised free products. *Math. Z.*, 78:423–438, 1962.

[11] G. Baumslag, B. Fine, and G. Rosenberger. One-relator groups: An overview. In C. M. Campbell, C. W. Parker, M. R. Quick, E. F. Robertson, and C. M.Editors Roney-Dougal, editors, *Groups St Andrews 2017 in Birmingham*, London Mathematical Society Lecture Note Series, page 119–157. Cambridge University Press, 2019.

[12] G. Baumslag, A. Myasnikov, and V. Remeslennikov. Algebraic geometry over groups. I. Algebraic sets and ideal theory. *J. Algebra*, 219(1):16–79, 1999.

[13] G. Baumslag and J. E. Roseblade. Subgroups of direct products of free groups. *Journal of The London Mathematical Society-second Series*, pages 44–52, 1984.

[14] N. Bergeron, F. Haglund, and D. T. Wise. Hyperplane sections in arithmetic hyperbolic manifolds. *Journal of the London Mathematical Society*, 83(2):431–448, February 2011.

[15] M. Bestvina, K. Bromberg, and K. Fujiwara. The verbal width of acylindrically hyperbolic groups. *Algebraic & Geometric Topology*, 19(1):477–489, 2019.

[16] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

[17] Mladen Bestvina. Degenerations of the hyperbolic space. *Duke Mathematical Journal*, 56(1):143–161, February 1988.

[18] H. Bigdely and D. T. Wise. Quasiconvexity and relatively hyperbolic groups that split. *Michigan Mathematical Journal*, 62(2):387 – 406, 2013.

[19] C. Bleak and M. Quick. The infinite simple group v of richard j. thompson: presentations by permutations. *Groups, Geometry, and Dynamics*, 11(4):1401–1436, December 2017.

[20] B. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.

[21] B. Bowditch. Tight geodesics in the curve complex. *Invent. Math.*, 171(2):281–300, 2008.

[22] M. Bridson and D. Kochloukova. The virtual first betti number of soluble groups. *Pacific Journal of Mathematics*, 274(2):497–510, April 2015.

[23] M. Bridson and D. Kochloukova. Volume gradients and homology in towers of residually-free groups. *Mathematische Annalen*, 367(3):1007–1045, April 2017.

[24] M. R. Bridson. Profinite rigidity and free groups. In Bernard Teissier Jean-Michel Morel, editor, *Mathematics Going Forward*, Lecture Notes in Mathematics. Springer Cham, 2023.

[25] M. R. Bridson, J. Howie, C. F. Miller, III, and H. Short. Subgroups of direct products of limit groups. *Annals of Mathematics. Second Series*, 170(3):1447–1467, 2009.

[26] M. R. Bridson, J. Howie, C. F. Miller, III, and H. Short. On the finite presentation of subdirect products and the nature of residually free groups. *American Journal of Mathematics*, 135(4):891–933, 2013.

[27] M. R. Bridson and H. Wilton. Subgroup separability in residually free groups. *Mathematische Zeitschrift*, 260(1):25–30, September 2008.

[28] M. Casals-Ruiz, A. Garreta, and J. de la Nuez González. On the positive theory of groups acting on trees, 2019. arXiv:1910.09000.

[29] M. Casals-Ruiz, A. Garreta, I. Kazachkov, and J. de la Nuez González. Simple groups with infinite verbal width and the same positive theory as free groups, 2019. arXiv:1911.02117.

[30] M. Casals-Ruiz, K. Ilya, and D. Andrew. Limit groups over coherent right-angled Artin groups. *Publicacions Matemàtiques*, 67(1):199–257, feb. 2023.

[31] M. Casals-Ruiz and I. Kazachkov. Limit groups over partially commutative groups and group actions on real cubings. *Geom. Topol.*, 19(2):725–852, 2015.

[32] C. Cashen. Splitting line patterns in free groups. *Algebraic & Geometric Topology*, 16:621–673, 04 2016.

[33] C. Champetier and V. Guirardel. Limit groups as limits of free groups. *Israel J. Math.*, 146:1–75, 2005.

[34] L. Ciobanu, B. Fine, and G. Rosenberger. The surface group conjecture: Cyclically pinched and conjugacy pinched one-relator groups. *Results in Mathematics*, 64(1-2):175–184, February 2013.

[35] D. Cooper, D. Long, and A. W. Reid. On the virtual betti numbers of arithmetic hyperbolic 3–manifolds. *Geometry & Topology*, 11(4):2265–2276, December 2007.

[36] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes*, volume 1441 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1990. Les groupes hyperboliques de Gromov. [Gromov hyperbolic groups], With an English summary.

[37] R. Coulon. Small cancellation theory and Burnside problem. Internat. J. Algebra Comput. 24 (2014), no. 3, 251-345, 2013.

[38] Rémi Coulon. Théorie de la petite simplification : une approche géométrique. *Asterisque*, 380:1–33, 2016.

[39] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.

[40] J. de la Nuez González. On expansions of non-abelian free groups by cosets of a finite index subgroup, 2017. arXiv:1707.03066.

[41] G.-A. Diao and M. Feighn. The grushko decomposition of a finite graph of finite rank free groups: an algorithm. *Geometry & Topology*, 9(4):1835–1880, September 2005.

[42] C. G. Droms. Graph groups, coherence and three-manifolds. *J. Algebra*, 106(2):484–489, 1987.

[43] M. J. Dunwoody. Folding sequences. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 139–158. Geom. Topol. Publ., Coventry, 1998.

[44] B. Fine, O. Kharlampovich, A. Myasnikov, V. Remeslennikov, and G. Rosenberger. On the surface group conjecture. *Scientia, Series A: Mathematical Sciences*, 15:1–15, 2007.

[45] M. H. Freedman and F. Quinn. *Topology of 4-Manifolds (PMS-39).* Princeton University Press, 1990.

[46] J. Fruchter. Limit groups over coherent right-angled artin groups are cyclic subgroup separable. *Michigan Mathematical Journal*, 1(15), January 2022.

[47] J. Fruchter and I. Morales. Virtual homology of residually free groups and profinite rigidity of direct products, 2023. arXiv:2209.14925.

[48] K. Fujiwara. The rates of growth in an acylindrically hyperbolic group, 2023. arXiv:2103.01430.

[49] K. Fujiwara and Z. Sela. The rates of growth in a hyperbolic group. *Inventiones mathematicae*, June 2023.

[50] G. Gardam, D. Kielak, and A. D. Logan. The surface group conjectures for groups with two generators, 2022. arXiv:2202.11093.

[51] R. Gitik. Graphs and separability properties of groups. *Journal of Algebra*, 188(1):125–143, 1997.

[52] E. R. Green. *Graph products of groups.* PhD thesis, University of Leeds, 1990.

[53] R. I. Grigorchuk. Degrees of growth of finitely generated groups, and the theory of invariant means. *Mathematics of the USSR-Izvestiya*, 25(2):259–300, April 1985.

[54] M. Gromov. Groups of polynomial growth and expanding maps. *Publications mathématiques de l'IHÉS*, 53(1):53–78, December 1981.

[55] D. Groves. Limit groups for relatively hyperbolic groups, ii: Makanin-razborov diagrams. *Geometry and Topology*, 9:2319–2358, 2005.

[56] D. Groves. Limit groups for relatively hyperbolic groups. II. Makanin-Razborov diagrams. *Geom. Topol.*, 9:2319–2358, 2005.

[57] D. Groves and M. Hull. Homomorphisms to acylindrically hyperbolic groups i: Equationally noetherian groups and families. *Trans. Amer. Math. Soc.*, 372(10):7141–7190, 2019.

[58] D. Groves and Henry Wilton. The structure of limit groups over hyperbolic groups. *Israel Journal of Mathematics*, 226(1):119–176, Jun 2018.

[59] F. Grunewald, A. Jaikin-Zapirain, and P. A. Zalesskii. Cohomological goodness and the profinite completion of Bianchi groups. *Duke Mathematical Journal*, 144(1):53 – 72, 2008.

[60] V. Guirardel. Bounding the complexity of small actions on $\mathbb{R}$-trees. https://perso.univ-rennes1.fr/vincent.guirardel/papiers/complexity.ps, 2001.

[61] V. Guirardel. Limit groups and groups acting freely on $\mathbb{R}^n$-trees. *Geom. Topol.*, 8:1427–1470, 2004.

[62] V. Guirardel. Actions of finitely generated groups on $\mathbb{R}$-trees. *Ann. Inst. Fourier (Grenoble)*, 58(1):159–211, 2008.

[63] V. Guirardel and G. Levitt. Admissible subtrees. Preliminary version of a preprint, private communication.

[64] V. Guirardel and G. Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

[65] V. Guirardel and G. Levitt. *JSJ Decompositions of Groups*. Asterisque Series. Société Mathématique de France, 2017.

[66] F. Haglund and D. T Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, January 2008.

[67] S. Heil. Test sequences and formal solutions over hyperbolic groups, 2018. arXiv:1811.06430.

[68] T. Hsu and D. T. Wise. Cubulating graphs of free groups with cyclic edge groups. *American Journal of Mathematics*, 132(5):1153–1188, 2010.

[69] M. Hull and D. Osin. Transitivity degrees of countable groups and acylindrical hyperbolicity. *Israel J. Math.*, 216(1):307–353, 2016.

[70] O. Kharlampovich and A. 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.

[71] O. Kharlampovich and A. Myasnikov. Elementary theory of free non-abelian groups. *J. Algebra*, 302(2):451–552, 2006.

[72] O. Kharlampovich and A. Myasnikov. Limits of relatively hyperbolic groups and lyndon's completions. *Journal of the European Mathematical Society*, 014(3):659–680, 2012.

[73] R. C. Kirby and L. C. Siebenmann. On the triangulation of manifolds and the Hauptvermutung. *Bulletin of the American Mathematical Society*, 75(4):742 – 749, 1969.

[74] D. Kochloukova and F.Y. Mokari. Virtual rational betti numbers of abelian-by-polycyclic groups. *Journal of Algebra*, 443:75–98, 2015.

[75] P. Kropholler and K. Lorensen. Virtually torsion-free covers of minimax groups. *Annales scientifiques de l'École normale supérieure*, 53(1):125–171, 2020.

[76] R. Kropholler and G. Wilkes. Profinite properties of RAAGs and special groups. *Bulletin of the London Mathematical Society*, 48(6):1001–1007, September 2016.

[77] P.A. Linnell. On accessibility of groups. *Journal of Pure and Applied Algebra*, 30(1):39–46, 1983.

[78] E.D. Loginova. Residual finiteness of the free product of two groups with commuting subgroups. *Siberian Mathematical Journal*, 40:341 – 350, 1999.

[79] E.D. Loginova. Cyclic subgroup separability of the free product of two groups with commuting subgroups. *Nauch. Tr. Ivan. Gos. Univ. Math.,*, 3:49–55, 2000.

[80] L. Louder. Krull dimension for limit groups. *Geometry & Topology*, 16(1):219–299, January 2012.

[81] R. C. Lyndon. Cohomology theory of groups with a single defining relation. *Annals of Mathematics. Second Series*, 52:650–665, 1950.

[82] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory.* Number 89 in Classics in Mathematics. Springer Berlin Heidelberg, 2001.

[83] W. Lück. Approximating L^2-invariants by their finite-dimensional analogues. *Geometric and Functional Analysis*, pages 455—481, 1994.

[84] J. Merzljakov. Positive formulae on free groups. *Algebra i Logika Sem.*, 5(4):25–42, 1966.

[85] A. Miasnikov and V. Remeslennikov. Exponential groups 2: Extensions of centralizers and tensor completion of csa groups. *International Journal of Algebra and Computation*, 6, 08 1995.

[86] A. Minasyan and D. Osin. Acylindrical hyperbolicity of groups acting on trees. *Math. Ann.*, 362(3-4):1055–1105, 2015.

[87] I. Morales. On the profinite rigidity of free and surface groups, 2022. arXiv:2211.12390.

[88] A. Myasnikov and V. Remeslennikov. Groups with exponents i. fundamentals of the theory and tensor completions. *Siberian Mathematical Journal*, 35(5):986–996, Sep 1994.

[89] W. Neumann. Immersed and virtually embedded π_1–injective surfaces in graph manifolds. *Algebraic and Geometric Topology*, 1:411–426, 07 2001.

[90] W.D. Neumann, R. Bieri, and R. Strebel. A geometric invariant of discrete groups. *Inventiones mathematicae*, 90:451–478, 1987.

[91] G. Niblo and D. Wise. Subgroup separability, knot groups and graph manifolds. *Proceedings of the American Mathematical Society*, 129, 05 2000.

[92] O. Ocaña, A. Garrido, and B. Klopsch. Pro-p groups of positive rank gradient and hausdorff dimension, 01 2019. arXiv:1901.03101.

[93] D. Osin. Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.*, 368(2):851–888, 2016.

[94] F. Paulin. Outer automorphisms of hyperbolic groups and small actions on **R**-trees. In *Arboreal group theory (Berkeley, CA, 1988)*, volume 19 of *Math. Sci. Res. Inst. Publ.*, pages 331–343. Springer, New York, 1991.

[95] C. Perin. *Elementary embeddings in torsion-free hyperbolic groups*. Theses, Université de Caen, October 2008. Thèse rédigée en anglais, avec une introduction détaillée en français.

[96] V. P. Platonov and O. I. Tavgen'. Grothendieck's problem on profinite completions and representations of groups. *K-Theory*, 4(1):89–101, January 1990.

[97] C. Reinfeldt and R. Weidmann. Makanin–razborov diagrams for hyperbolic groups. *Annales mathématiques Blaise Pascal*, 26(2):119–208, 2019.

[98] A. H. Rhemtulla. A problem of bounded expressibility in free products. *Proc. Cambridge Philos. Soc.*, 64:573–584, 1968.

[99] L. Ribes. *Profinite graphs and groups*, volume 66 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]*. Springer, Cham, 2017.

[100] E. Rips and Z. Sela. Structure and rigidity in hyperbolic groups. I. *Geom. Funct. Anal.*, 4(3):337–371, 1994.

[101] Otto Schreier. Die untergruppen der freien gruppen. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 5(1):161–183, December 1927.

[102] P. Scott and T. Wall. *Topological methods in group theory*, page 137–204. London Mathematical Society Lecture Note Series. Cambridge University Press, 1979.

[103] D. Segal. *Polycyclic Groups.* Cambridge Tracts in Mathematics. Cambridge University Press, 1983.

[104] D. Segal. *Words: notes on verbal width in groups*, volume 361. Cambridge University Press, 2009.

[105] Z. Sela. Acylindrical accessibility for groups. *Inventiones Mathematicae*, 129(3):527–565, August 1997.

[106] Z. Sela. Diophantine geometry over groups I: Makanin-razborov diagrams. *Publications Mathématiques de l'IHÉS*, 93:31–105, 2001.

[107] Z. Sela. Diophanting geometry over groups II: Completions, closures and formal solutions. *Israel Journal of Mathematics*, 134(1):173–254, December 2003.

[108] Z. Sela. Diophantine Geometry over Groups. VI. The elementary theory of a free group. *Geom. Funct. Anal.*, 16(3):707–730, 2006.

[109] Z. Sela. Diophantine Geometry over Groups. VII. The elementary theory of a hyperbolic group. *Proc. Lond. Math. Soc. (3)*, 99(1):217–273, 2009.

[110] Z. Sela. Diophantine Geometry over Groups X. The elementary theory of free products of groups, 2010. arXiv:1012.0044.

[111] J-P. Serre. *Galois Cohomology.* Springer Berlin Heidelberg, 1997.

[112] E. V. Sokolov. On the cyclic subgroup separability of the free product of two groups with commuting subgroups. *Int. J. Algebra Comput.*, 24:741–756, 2014.

[113] G.A. Swarup. Decompositions of free groups. *Journal of Pure and Applied Algebra*, 40:99–102, 1986.

[114] T. N. Venkataramana. Virtual betti numbers of compact locally symmetric spaces. *Israel Journal of Mathematics*, 166(1):235–238, August 2008.

[115] F. Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Annals of Mathematics*, 87(1):56–88, 1968.

[116] H. Wilton. Hall's theorem for limit groups. *Geometric and Functional Analysis*, 18(1):271–303, Apr 2008.

[117] H. Wilton. Residualy free 3-manifolds. *Algebraic & Geometric Topology*, 8:2031–2047, 2008.

[118] H. Wilton. Solutions to Bestvina & Feighn's exercises on limit groups. In M. R. Bridson, P. H. Kropholler, and I. J. Leary, editors, *Geometric and Cohomological Methods in Group Theory*, London Mathematical Society Lecture Note Series, page 30–62. Cambridge University Press, 2009.

[119] H. Wilton. One-ended subgroups of graphs of free groups with cyclic edge groups. *Geometry & Topology*, 16(2):665 – 683, 2012.

[120] H. Wilton. Essential surfaces in graph pairs. *Journal of the American Mathematical Society*, 31(4):893–919, June 2018.

[121] H. Wilton. On the profinite rigidity of surface groups and surface words. *Comptes Rendus Mathématique*, 359(2):119–122, 2021.

[122] H. Wilton and P. Zalesskii. Distinguishing geometries using finite quotients. *Geometry & Topology*, 21:345–384, 2017.

[123] D. T. Wise. Subgroup Separability of Graphs of Free Groups with Cyclic Edge Groups. *The Quarterly Journal of Mathematics*, 50(1):107–129, 03 2000.

[124] P. Zalesskii and T. Zapata. Profinite extensions of centralizers and the profinite completion of limit groups. *Revista Matemática Iberoamericana*, 36(1):61–78, 2020.