Javier de la Nuez González

On expansions of non-abelian free groups by cosets of a finite index subgroup

2016

Mathematik

On expansions of non-abelian free groups by cosets of a finite index subgroup

Inaugural-Dissertation
zur Erlangung des Doktorgrades
der Naturwissenschaften im Fachbereich
Mathematik und Informatik
der Mathematisch-Naturwissenschaftlichen Fakultät
der Westfälischen Wilhelms-Universität Münster

vorgelegt von
Javier de la Nuez González
aus Madrid
2016

Dekan:	Prof. Dr. Martin Stein
Erster Gutachter:	Prof. Dr. Dr. Katrin Tent
Zweiter Gutachter:	Prof. Dr. Zlil Sela
Tag der Disputation:	24/11/2016
Tag der Promotion:	

Abstract

Let $\mathbb{F}$ be a finitely generated non-abelian free group and Q a finite quotient. Denote by $\mathcal{L}^Q$ the language obtained by adding unary predicates $\{P_q\}_{q\in Q}$ to the language of groups.

By generalizing some of the techniques involved in Zlil Sela's solution to Tarski's question on the elementary theory of non-abelian free groups, we provide a few basic results on the validity of first order sentences in the $\mathcal{L}^Q$-expansion of $\mathbb{F}$ in which every P_q is interpreted as the preimage of q in $\mathbb{F}$.

In particular, we prove an analogous result to Sela's generalization of Merzlyakov's theorem for $\forall\exists$-sentence and show that the positive theory of such a structure depends only on Q and not on the rank of $\mathbb{F}$ nor the particular quotient map.

Acknowledgements

First and foremost, I would never have come this far without the unconditional love and support of my family, who were always there in the highs and lows and had to sacrifice more than I did during these years. This is partly theirs as well.

It would also have been impossible without the support of my advisor Prof. Katrin Tent, the many opportunities for exchange with which she provided me.

I am very grateful for Prof. Zlil Sela for inviting me to visit to Jerusalem and his openness for discussion, often around what for him must have been fairly trite and boring technical details of his wonderful work. His patient and unassuming attitude is something to appreciate.

It is impossible for me to forget the truly exceptional hospitality of Montserrat Casals and Ilya Kazachkov, who were always ready to lend me an ear and helped me to gain focus and confidence I need of it.

I must also thank Chloe Perin and Rizos Sklinos, who allowed this nobody to collaborate with them and taught some invaluable lessons on how to put the book aside and start doing math.

I would also like to mention Prof. Gilbert Levitt, who I had the chance to visit for a week in Caen and with whom I came to experience the beauty of geometric group theory for the first time. Also Martin Bays, who helped me chose this dissertation's title and Franziska Janke and Daniel Palacín, who shared with me very valuable insights on the process of writing a thesis.

Finally, life in the math building would not have been the same without the warmth of my fellow graduate students in Münster: Alex (who saved my life several times), Isabel and Zaniar.

Contents

Introduction

Some history Model theory concerns itself with the interaction between syntactic objects, first order formulas and consistent collections thereof, and the mathematical structures which interpret them. A fundamental notion is that of the (first order) theory $Th(\mathcal{S})$ of a given structure $\mathcal{S}$.

Free groups are a very natural and simple to describe mathematical object that plays a role in several areas outside of algebra, such as topological group theory. Interest in the model theory of free groups can be traced back to the forties by one of the fathers of modern logic, Alfred Tarski, who proposed the following question:

Question. *Do $\mathbb{F}_m$ and $\mathbb{F}_n$ have the same order theory for $2 \leqslant m < n$?*

The free group of rank m, $\mathbb{F}_m$ can be conceived as the collection of all finite words in m-letters and their iverses, up to the cancellation of consecutive pairs of dual letters, so that each element can be represented by a unique reduced such word. This point of view, within the larger context of what is known as "combinatorial group theory", is the one privileged by classical reference works like [MKS07] and [SL77] dominated mathematicians initial understanding of free groups [SL77].

Some early successes Tarski's problem was achieved using such methods by Merzlyakov [Mer66], who proved that if a generates a free group $\mathbb{F}$, then the validity of any sentence of the form $\forall x \exists y\ \phi(x, y, a)$ true in $\mathbb{F}$, where $\phi(x, y, a)$ is a system of equations over a, is witnessed by what is known as a formal solution: some algebraic expression of the y's in terms of x and a, rendering the equations in the system trivial. Generalizing this result, he was able to prove that all non-abelian free groups satisfy the same positive sentences.

Another early milestone was the work of Makanin (see [Mak83]), who provided an algorithm for deciding whether system of equations with parameters in the free group admits a solution. Later on by refining Makanin's technique Razborov (see [Raz85]) managed provide a description of the set of solutions to such systems.

Meanwhile P. Serre [SB77], established a duality between algebraic decompositions of a group as a graph of groups and particular actions on such groups on simplicial trees, providing a new approach to free groups and other notions in combinatorial group theory. Later on there was an increased interest in the possibility of generalizing Serre's work to so called $\mathbb{R}$-trees, metric trees where the set of branching points is allows to be non-discrete. An important problem in its day was that of characterizing those finitely generated groups admitting a free action on an $\mathbb{R}$-tree (for simplicial trees the answer is precisely free groups). It was E. Rips' outstanding contribution to the subject was to realized that the combinatorial iterative procedure behind Makanin and Razborov's work could be understood in in a geometrical sense and hence generalized to such actions, providing a solution of the problem above. More general developments of the idea appeared later in works as [BF95] or [Gui08] and its predecessors, which show how certain classes of actions always decompose into simpler pieces of a certain type.

Another important idea, going back to the work Gromov, is that there is endow a class of metric spaces with a topology, that some form of compactness holds under certain circumstances and that one certain large scale properties of a metric space can be distilled from the limit of a sequence of rescalings of the space by smaller and smaller positive constants. Later Paulin [Pau88] and Bestvina generalized this idea to isometric actions of a group G on certain metric spaces.

These are some of the main ingredients in Z. Sela remarkable positive answer to Tarski's question, which came at the turn of the century in the shape of a series of six papers (culminating in [Sel06]). Tarski's problem is a consequence of the stronger result that the embedding of free group as a free factor of another one is an elementary embedding. A solution to Tarski's problem was reached independently by Kharlampovich and Myasnikov [KM06] using combinatorial techniques.

This thesis This work provides a small illustration of the versatility of Sela's geometric approach. Fix a finite group Q and consider the language $\mathcal{L}^Q$ obtained by adding unary predicates $\{P_q\}_{q\in Q}$ to the language of groups. Given a group G, together with an epimorphism $\pi : G \to Q$. This is what we call a π-group. By interpreting P_q as the preimage of q in G by π we can promote G to an $\mathcal{L}^Q$ structure.

This sort of expansion is fairly common move in model theory. Sometimes a mathematical object can be seen both as an $\mathcal{L}_1$-structure and as an $\mathcal{L}_2$-structure, where the theory is 'well-behaved' in both cases in a certain sense, while the $\mathcal{L}_1 \cup \mathcal{L}_2$ structure exhibits an extremely wild behaviour. In our case, the evidence collected so far points in the opposite direction.

We are interested in the case in which $\mathbb{F}$ is a finitely generated free group and π an epimorphism. What is the general picture? We conjecture the following generalizations of Sela's result:

Conjecture. *Given a finite group Q, non-abelian free groups F_1, F_2 and epimorphisms π_i from F_i to Q for $i = 1, 2$ the pairs (F_1, π_1) and (F_2, π_2) are elementary equivalent as $\mathcal{L}^Q$-structures.*

Unfortunately due to the complexity of the techniques involved, are not able to offer an aswer to the question in this work. We dapt several of his tools to the study of the language $\mathcal{L}^Q$.

A solution of a system of equations (for simplicity without constants) $\Sigma(x) = 1$ in a given free group $\mathbb{F}$ can be corresponds to a homomorphism from the group G_Σ with presentation $\langle x | \Sigma(x) = 1\rangle$. Makanin showed the existence of finitely many finite chains of proper epimorphisms, terminating in a free group:

$$G_\Sigma \xrightarrow{\eta_0} L_1 \xrightarrow{\eta_1} L_2 \xrightarrow{\eta_2} \cdots L_m = F_k$$

Such that any $f \in Hom(G_\Sigma, \mathbb{F})$ can be written as $\eta_0 \circ \tau_1 \circ \eta_1 \circ \tau_2 \cdots \eta_{m-1} \circ h$ for automorphisms τ_j of L_j and a homomorphism h from F_k to $\mathbb{F}$. The groups L_i happen to be so called limit groups: finitely generated models of the universal theory of a non-abelian free group. This is also one of the initial steps in Sela's work (see [Sel01]), whose proof is based on the Bestvina-Paulin limiting method and Rips' analysis of axions of groups on real trees, which makes it easily generalizable to hyperbolic groups (see [RW]).

In the language $\mathcal{L}^Q$ the basic quantifier free formula that needs to be considered is the conjunction of a system of equations together with conditions of the form $P_q(x_j)$ (this is a special case of the notion of an equation with rational constraints, explored in works such as [DGH05] under a different angle, more akin to computer science).

We work in the category finitely generated π-groups as a category, a morphism from (G, π) to (G', π') being any homorphism $f : G \to G'$ for which $\pi' \circ f = \pi$. Solutions to any condition as above can be identified with morphisms from a finitely generated π-group (G, π_G) to $(\mathbb{F}, \pi)$.

In chapter 4 we provide a description of the family of such morphisms by generalizing Sela's version of the Makanin-Razborov procedure and one of its refinements (namely that of a taut-Makanin Razborov diagram) to our new category.

The above mentioned morphism-classifying sequences of epimorphisms are what one usually refers to as a resolution. Our approach to resolutions is formally slightly different from the one found in Sela in the sense that we use a rooted tree to index what would be the free factors of a resolution in a classical sense. For the scope of this work this probably makes little difference, but we think it might be useful in later stages of the analysis of $\mathcal{L}^Q$-sentences.

The concept analogous to limit groups is that of what we call π-limit groups: finite models of the universal theory of our reference structure $(\mathbb{F}, \pi)$ or, equivalently, those finitely generated π-groups which are discriminated by morphisms to $(\mathbb{F}, \pi)$. As in the standard case, this notion can be eventually be shown to be independent from the particular reference structure. Not every homomorphism $\pi : L \to Q$ from a limit group to Q makes it into a π-limit group, but we will formulate a sufficient (and necessary) condition for this to be the case, akin to that of a constructible limit group (see [BF09]).

In chapter 5 we present Sela's notions of a tower (a group together with a certain algebraic structure) and a test sequence (a sequence of homomorphisms from the tower to the free group), suitably adapted to our setting. Test sequences are a generalization of small cancellation sequences. They play a crucial role in Sela's generalization of Merzlyakov's theorem for $\forall\exists$-sentence, where universal variables to be constrained by certain algebraic relations.

Our exposition differs from Sela's in that we isolate those properties of an individual sequence must satisfy from those that a collection of sequences has to satisfy in order for the theorem to hold. In the final section of the chapter we give a proof of the existence of such families.

In chapter 6 we adapt the aforementioned theorem to our setting, together with the following generalization of the other result by Merzlyakov mentioned before:

Theorem 6.2.6. *Let A be a free factor of non-abelian free groups F_1 and F_2 and for $i = 1, 2$ let π_i be a surjective homomorphism from F_i to the finite group Q.*
Then $Th^+_A(F_1, \pi_1) = Th^+_A(F_2, \pi_2)$.

Here $Th^+_A(M)$ stands for the positive A-theory of M, that is, the collection of all sentences built without use of the negation symbol with constants in $A \subset M$.

The results above seem to suggest there is little interaction between the newly introduced predicates and the old family of definable sets. We hope to be able to explore this line of thought in future work.

Chapter 1

Preliminaries

1.1 Reminder: the free group

We start with a quick reminder of what free groups and free products are; see [LS15] for a thorough introduction. Given a set of elements X a group F containing X is free over X if any map from X to a group G extends uniquely to some homomorphism from F to G. This object is unique up to isomorphism relative to X and we denote it by $\mathbb{F}(X)$. In this situation any element of G can be uniquely written as a word in $\{x, \bar{x}\}_{x\in X}$, where $\bar{x}$ is a new letter (formal inverse) dual to x. The group $\mathbb{F}(X)$ can be constructed by taking the quotient of the set of such terms (where inverses are a priori formal) by the relationship generated by equivalences of the form $uxx^{-1}v \sim uv$, where $x \in X^{-1}$. Any element can be represented by a reduced word, i.e., one in which a consecutive pair xx^{-1} does not appear. Such a word is cyclically reduced if in addition the first and last elements are not mutually inverse. A group F which is free over a certain set X is called free and any such X is called a base of F. All bases of a given free group have the same cardinality which, if finite, is denoted as the rank of the free group in question. The free group of rank k is often denoted as $\mathbb{F}_k$.

Given a collection $\{G_i\}_{i\in I}$ of groups, its free product, denoted by an expression such as $*_{i\in I}G_i$ or $G_1 * G_2 \cdots G_k$ is a group containing each G_i and satisfying a similar universal property. Namely, given any group H any collection of homomorphisms $\{f_i : G_i \to H\}_{i\in I}$ extends to a unique f, which we will denote as $\amalg_{i\in I} f_i$. Any non-trivial element of G can be written in a unique way as a product $a_1a_2 \cdots a_m$, where a_l and a_{l+1} belong to A_{i_l} and $A_{i_{l+1}}$ for any $1 \leqslant l \leqslant m$.

A free decomposition is an isomorphism (often thought as an identity) between a group and some non-trivial free product.

1.2 A few words about first order logic

Let us introduce here some basic notions of first order logic with which to formulate some of the results in this work. We will gloss over many formalities. A complete yet swift introduction to the matter can be found in [TZ12]; for a more fleshed-out the reader can take for example [Sho67].

A language $\mathcal{L}$ is a collection of symbols including constants, functions and relation symbols. Additionally, each function symbol and relation symbols are associated with a natural number $n \geqslant 1$, referred to as the arity of the symbol. When naming the constituents of $\mathcal{L}$ we will often include the n-arity of the symbols in a subscript between parenthesis. For most purposes, constants can be regarded as functions of arity 0. In order to do anything meaningful with $\mathcal{L}$ one needs some additional logical symbols, namely:

- A set of infinitely many variables: $\{x_n\}_{n\in\mathbb{N}}$
- The equality symbol $=$,
- Negation $\neg$, conjunction $\wedge$ and disjunction $\vee$ symbols
- An existential quantifier $\exists$ and a universal quantifier $\forall$

In addition to this we will also use commas and parenthesis as auxiliary symbols (in a somewhat loose way in the case of parenthesis), although they are not strictly necessary. Likely, one loses no expressive power by foregoing one of the two quantifier symbols, or one among $\{\wedge, \vee\}$. Expressions involving them can be safely regarded as abbreviations of expressions without them. In our case we will restrict that treatment to the logical connectors $\leftarrow$,$\rightarrow$ and $\leftrightarrow$. $\mathcal{L}$-terms and $\mathcal{L}$-formulas are sets of finite strings of symbols of $\mathcal{L}$ and logical symbols $\mathcal{L}$-structure which are 'meaningful' or 'sound' in an intuitive way. An $\mathcal{L}$-terms is any such string that can be generated in finitely many steps using the following rules:

- The string containing a single variable or constant is an $\mathcal{L}$-term
- Given a function symbol $f^{(m)} \in \mathcal{L}$, and terms $t_1, t_2, \cdots t_m$, the string $f(t_1, t_2, \cdots t_m)$ is an $\mathcal{L}$-term.

for $\mathcal{L}$-formulas, the generating rules are:

- If t_1 and t_2 are terms, then $t_1 = t_2$ is a formula.
- If $R^{(m)} \in \mathcal{L}$ is a relation symbol, and $t_1, t_2, \cdots t_m$ are terms, then $R(t_1, t_2, \cdots t_m)$ is a formula.
- If ψ is an $\mathcal{L}$-formula, then $\neg\psi$ is one as well.
- If ψ_1 and ψ_2 are $\mathcal{L}$-formulas, then $(\psi_1 \vee \psi_2)$ and $(\psi_1 \wedge \psi_2)$ are one as well.
- If ψ is a formula then $\exists x\ \psi$ and $\forall x\ \psi$ are one as well.

Those formulas as in the first two points are called *atomic*. For any occurrence in a formula of a quantifier Qx arising from an application of the third role, ones refers to the corresponding formula ψ as the scope of the quantifier and to any occurrence of x in it as a bounded occurrence of x. A variable with at least one non-bounded occurrence in a certain formula is called a free variable. A formula without free variables is called a sentence.

A formula with no occurrence of the negation symbol $\neg$ is called a *positive formula* and one with no instances of quantifiers is called a *quantifier free*. Typically when introducing a term t one writes $t(x_1, x_2, \cdots x_k)$, where $x_1, x_2, \cdots x_k$ are (distinct) variables containing all those appearing in t. Likewise $\phi(x_1, x_2 \cdots x_k)$ will stand for a formula with free variables among $x_1, x_2 \cdots x_k$. The notation $\pi(x_1, x_2 \cdots x_k)$ might also indicate a set of formulas whose free variables are among $x_1, x_2 \cdots x_k$.

For $\mathcal{L}$ as above, a $\mathcal{L}$-structure $\mathfrak{U}$ is given by a set U, referred to as the universe of the structure, together with a tuple an interpretation $Z^{\mathfrak{U}}$ for each of the symbols $Z \in \mathcal{L}$, which is:

- an element of U in case Z is a constant
- a function from A^m to A, if Z is a function symbol of arity m
- some subset of A^m if Z is a function symbol of arity m

In the future we might incur in a certain lack of precision and refer to $\mathfrak{U}$ simply as U. Given an $\mathcal{L}$-structure $\mathfrak{U}$, one can interpret any term $t(x_1, x_2, \cdots x_k)$ as a function $t^{\mathfrak{U}} : U^k \rightarrow U$. Likewise, a formula $\phi(x_1, x_2, \cdots x_k)$ for $k \geqslant 1$ can be interpreted as a predicate in $\phi(\mathfrak{U}) \subset U^k$. The formal definition merely reflects the intuitive way such expressions are used in common mathematical practice; the interested reader can take a look at the references mentioned above.

In the case of a $\mathcal{L}$-sentence ϕ one can ask whether ϕ holds (is true) in $\mathcal{M}$, denoted by the notation $\mathcal{M} \models \phi$. Again, the concept the formal definition merely reflects the informal commonsense notion of satisfaction.

Given a term $t(x_1, x_2, \cdots x_k)$ and any tuple $s_1, s_2, \cdots s_k$ of terms (possibly other variables) the expression $t(s_1, \cdots s_k)$ denotes the term obtained by replacing each variable x_i by s_i. One can also 'substitute' $t_1, t_2, \cdots t_k$ in place of free variables of a formula, but this is slightly more subtle.

As it is common practice in model theory, we might use a single letter, such as x to denote a whole tuple of $k = |x|$ many variables. In this case, given $Q \in \{\forall, \exists\}$, an expression like $Q \in \{\forall, \exists\}$ then Qx will abbreviate the string $Qx_1\, Qx_2\, \cdots Qx_k$. A formula of the form $\forall x\, \phi(x)$, where ϕ is quantifier free is called a universal formula. One of the form $\forall x\, \exists y\ \phi(x, y)$ with $\phi(x, y)$ quantifier free is called an $\forall\exists$-formula.

Given an $\mathcal{L}$-structure $\mathfrak{U}$ and a subset $A \subset U$, it is formally useful to consider the language $\mathcal{L}_A$ obtained by adding a constant c_a to the language for any element A; in practice we will use the same letter $a, b \cdots$ denote both the element in A and the associated constant. There is an obvious way one can extend $\mathfrak{U}$ to an $\mathcal{L}_A$-structure $\mathfrak{U}_A$.

The (positive) elementary theory of an $\mathcal{L}$-structure $\mathfrak{U}$, denoted by $Th(\mathfrak{U})$ ($Th^+(\mathfrak{U})$), is the collection of all (positive) $\mathcal{L}$-sentences valid in $\mathfrak{U}$. We say that two $\mathcal{L}$-structures $\mathfrak{U}$ and $\mathfrak{U}'$ are (positively) elementarily equivalent, denoted by $\mathfrak{U} \equiv \mathfrak{U}'$ ($\mathfrak{U} \equiv^+ \mathfrak{U}'$) if $Th(\mathfrak{U}) = Th(\mathfrak{U}')$ ($Th^+(\mathfrak{U}) = Th^+(\mathfrak{U}')$).

Given $\mathcal{L}$-structures $\mathfrak{U}$ and $\mathfrak{U}'$ such that their respective universes contain some common subset A, we say that they are elementarily equivalent over A, denoted by $\mathfrak{U}' \equiv_A \mathfrak{U}$, if $\mathfrak{U}'_A \equiv \mathfrak{U}_A$. The positive theory of $\mathfrak{U}$, denoted by $Th^+(\mathcal{U})$ is the collection of all positive sentences valid in $\mathfrak{U}$. We denote the corresponding weak elementarily equivalent relationships by $\equiv^+$ and $\equiv^+_A$. Considering only universal sentences instead of positive ones we obtain the notion of the positive theory and positive elementary equivalence respectively.

An $\mathcal{L}$-theory T implies an $\mathcal{L}$-sentence ϕ, something we denote as $T \vdash \phi$, if $\mathcal{M} \models \phi$ for any model $\mathcal{M}$ of T. We say that two $\mathcal{L}$-formulas $\phi(x)$ and $\psi(x)$ are T-equivalent if $T \models \forall x\, (\phi(x) \Leftrightarrow \psi(x))$. Any formula $\phi(x)$ is $\varnothing$-equivalent to one of the form:

$$\forall y_1\ \exists y_2\ \cdots \forall y_{2m-1} \exists y_{2m} (\bigvee_{j=1}^{m} (\bigwedge_{i=1}^{r_i} \psi_i^j))$$

where each ψ_i^j is either an atomic formula or the negation of one; we will refer to this as a normal form (for $\phi(x)$). If the formula we started with is positive, then the ψ_i^j can be all taken to be atomic.

1.3 Introducing our setting and some basic definitions

Groups can be seen as $\mathcal{L}_{gp}$ structures, where $\mathcal{L}_{gp} = \{1, \cdot, (-)^{-1}\}$, containing symbols for the identity, multiplication and inverse respectively.

Fix a finite group Q and let $\mathcal{L}^Q$ be the language resulting from adding an unary predicate P_q to $\mathcal{L}_{gp}$ for any $q \in Q$.

By a π-*group* we intend a pair (G, π), where $\pi : G \to Q$ is a homomorphism. Such an object can be seen as an $\mathcal{L}^Q$-structure in a natural way by interpreting P_q as the preimage of q by π. The epression as $T_\pi = Th(G, \pi)$ denotes the corresponding $\mathcal{L}^Q$-theory.

We are interested in the particular case of a π-group $(\mathbb{F}, \pi)$, fixed for most of this work, where $\mathbb{F}$ is a non-abelian free group and π a surjection. This structure will be usually abbreviated as F.

A *morphism* between π-groups (G, π) and (G', π') is a homomorphism $\phi : G \to G'$ satisfying $\pi_{G'} \circ \phi = \pi_G$. To help lighten the notation we will (for now) adopt the notation G to refer to a π-group (G, π_G), $\mathcal{H}$ one of the form (H, π_H) and so forth. Likewise, F will stand for our fixed pair $(\mathbb{F}, \pi)$.

A fundamental notion in the study of the theory $T_{\mathbb{F}}$ is that of a limit group. Many basic questions about the theory $T_{\mathbb{F}}$ can be restated in terms of limit groups and the homomorphisms between them. This notion and all of the main related results can be easily adapted to the theory T_π.

A sequence $(f_n)_n$ of homomorphisms from a countable group G to a group H is said *convergent* if and only if for any $g \in H$ eventually either $f_n(g) = 1$ or $f_n(g) \neq 1$. By the limit kernel of such a sequence $(f_n)_n$ we intend $lker_n\ f_n = \{w \in G \mid \#\{n \mid f_n(w) = 1\} = \infty\}$. It is a trivial matter to check that the limit kernel is a subgroup of G. By the *limit quotient* of $(f_n)_n$ we intend the quotient map from G to $G' = G/lker_n\ f_n$. Often we will use the term to refer to the quotient itself. If the f_n are morphisms from some π-group G to H, then we will be referring to the π-group G' obtained by pushing forward the π-structure of G onto G'.

Given a countable set X, we can identify $\mathcal{P}(X)$ with 2^X, which is a compact separable metric space for the product topology. Convergence in the sense above is then equivalent to the topological convergence of $Ker(f_n)$ in 2^G . In particular:

Observation 1.3.1. Any sequence of homomorphisms between two groups contains a convergent subsequence.

Definition 1.3.1.1. By a π-limit group we intend a π-group L which is the limit quotient of a finitely generated π-group G by a converging sequence of morphisms from G to F.

Letting $Q = \{1\}$ one recovers the standard notion of a limit group. We would like to remark that the notion of a π-limit group is strictly stronger than that of a limit group equipped with a morphism to Q. An obvious restriction is the fact that given a π-limit group (L, π_L) and any $x \in L$ the image by π of the centralizer $Z_L(x)$ is necessarily cyclic.

A π-subgroup of a π-group G is simply a subgroup of G endowed with a π-structure making the inclusion map a π-morphism.

By a *(A-)restricted* π-group we intend a pair $(\mathsf{G}, \mathsf{A}, \iota)$, where G is a π-group and ι an injective morphism from a finitely generated π-subgroup A of F to G.

Since we will never be interested in comparing restricted groups differing only in the particular isomorphism ι and not in the underlying groups, we will often call such (G, A, ι) by the name G_A and possibly even identify A with its image.

Definition 1.3.1.2. By a morphism from an A-restricted π-group (G, A, ι) to another one (G', A, ι') we intend a morphism $\phi : \mathsf{G} \to \mathsf{G}'$ such that $\iota' \circ \phi = \iota$.

Of course, given a restricted π-group G_A and a morphism $\phi : G \to H$ injective on A, there is a unique way to push the A-restricted structure of G to one on H, with respect to which ϕ is a restricted π-morphism.

Definition 1.3.1.3. A restricted π-limit group is a restricted Q-group L_A where L is the quotient of a convergent sequence of morphisms ϕ_n from a finitely generated restricted π-group G_A to F_A and $\phi_n \circ \iota^A_L = \iota^A_F$ for all m.

Definition 1.3.1.4. A graded π-group is given by a tuple $(\mathsf{G}_A, (P_j)_{i \in I})$, comprising:

i) A restricted π-group G_A

ii) A family $\{P_j\}_{j \in I}$ of finitely generated subgroups of G, which we call parameter subgroups.

We will often refer to such a structure by an expression of the form $\mathsf{G}_{A,(P_i)_{i\in I}}$. A graded π-limit group is just a graded π-group for which the underlying restricted π-group is a restricted limit group.

By a marking of a group G we intend a map $\iota : x \to G$ where $x = (x_i)_{=1}^{k}$ is a finite set of syntactic variables . A marked group is a pair (G, ι) where G is a group and $\iota : x \to G$ a marking of G for which $\langle \iota(x) \rangle = G$. One should think of ι as an interpretation of the tuple x in G. In accordance, we will always refer to $\iota(x)$, or its image in a quotient of G as simply x. If instead of a group we have a π-group, restricted π-group we will talk of marked π-groups, marked restricted π-groups and so forth. The expression $x(G)$ will stand for $\iota(x)$ and $\iota(x_i)$ for $x_i(G)$. A marked restricted group G is a pair (G, ι) where ι is a marking such that $\langle A, \iota(x) \rangle = G$. A marked graded group A consists of a tuple $(G, (\iota_i)_{i\in I\setminus\{0\}}, \iota_0)$, together with markings $\iota_i : p_i \to P_i$ and $\iota : x \to G$ such that $\pi(\langle A, \iota(x), \iota_i(p_i) \rangle) = G$. A marked restricted group G is a pair (G, ι) where ι an x-marking such that $\langle A, \iota(x) \rangle = G$. A marked graded group A consists of a tuple $(G, (\iota_i)_{i\in I\setminus\{0\}}, \iota_0)$, together with markings $\iota_i : p_i \to P_i$ and $\iota : x \to G$ such that $\langle A, \iota(x), \iota_i(p_i) \rangle$. The notation might replace A by some finite tuple a of generators. Notice that given an epimorphism of (restricted , graded) groups, a marking of the domain pushes forward to a marking of the target univocally.

Definition 1.3.1.5. A group G is called CSA if and only if every maximal abelian subgroup A oef G is malnormal, i.e. $g \in A$ whenever $A \cap A^g \neq \{1\}$. It is called commutative transitive if and only if the relationship $R(x, y) \cong [x, y] = 1$ is transitive when restricted to $G\setminus\{\{1\}\}$.

The following is well-known and easy to prove:

Observation 1.3.2. In a commutative transitive group any non-trivial element is contained in a unique maximal abelian subgroup. Any CSA group is commutative transitive.

For the following theorem see for example [CG05]

Theorem 1.3.3. *Limit groups are CSA.*

Remark 1.3.4. Given π-groups G and H we denote the unique extension π_K of π_G and π_H to $K = G * H$ as $\mathsf{G} * \mathsf{H}$. If G can be promoted to an A-restricted π-group, the same injection map makes K into one.

1.4 Preorders

We will sometimes prove properties of a certain class of strucutres $\mathcal{C}$ by induction on some well-founded partial order of $\mathcal{C}$. Given a partially ordered set (poset) $(J, \leqslant)$, we denote by $\equiv$, as usual $\leqslant \cap \geqslant$. A strict poset will be one for which $\leqslant$ is irreflexive. Recall a poset $(J, \leqslant)$ is said well-founded if there are no infinite chains $\lambda_1 > \lambda_2 > \cdots \lambda_n \cdots$ of elements of P. For $\lambda \in J$ we let $J\restriction_\lambda = \{\mu \in P \mid \mu \geqslant \lambda\}$. Any poset determines a successor relationship, $\succ$, where $\lambda \prec \mu$ if and only if $\lambda < \mu$ and there is no element between λ and μ. By a rooted simplicial tree we will intend a poset $(J, \geqslant, r)$ together with a distinguished element r, the root such that:

i) If $i, j \in J$ have a common lower bound, they are comparable.

ii) Any $i, j \in J$ have a least common upper bound.

iii) For any $\lambda \in J$ there is a unique sequence $r = \lambda_0 \succ \lambda_1 \succ \cdots \lambda_n = \lambda$ for some $n \in \mathbb{N}$.

By a branch of it we mean a sequence $\lambda_0 \succ \lambda_1 \succ \cdots \lambda_n$ or $\lambda_0 \succ \lambda_1 \succ \cdots \lambda_n \succ \cdots$. Whenever $\lambda \prec \mu$ in this context we will refer to λ as a child of μ and to μ as the (unique) parent of λ, or $p(\lambda)$. Let $Ch(\lambda)$ stand for the set of all the children of λ. Given $\lambda \succcurlyeq \mu$, we will refer to λ as an ancestor of μ and to μ as a descendant of λ. We call *leaves* those elements of J which are minimal and denote the set of all of them by $lv(J)$, while $\hat{J}$ will stand for the complement of the latter in J. The following is a consequence of König's lemma.

Lemma 1.4.1. *Every finitely branching well-founded rooted simplicial tree is finite.*

Given preorders $\leqslant_1$ and $\leqslant_2$ on a set P, we denote by $\leqslant_0 = \leqslant_1 \times \leqslant_2$ the relationship defined by $p \leqslant_0 q \Leftrightarrow p \leqslant_1 q \vee (p \equiv_1 q \wedge p \leqslant_2 q)$. It is well-known and easy to check this is a preorder and and that it is well-founded whenever both the $\leqslant_i$ are

Notation. Given a class $\mathcal{C}$ of structures and a function $f : \mathcal{C} \to \mathbb{N}$ we let $\leqslant_f$ be the preorder given by $\leqslant_f = f^{-1}(\leqslant_{\mathbb{N}})$.

We say that a rooted tree (S, r) extends some subtree (T, r) in case each $v \in S \backslash T$ is a descendant of some $w \in lv(T)$. The l-th level of a rooted tree (T, r), denoted by $[T]_l$ consists of those nodes of T at distance l from r. Let also $[T]_{\leqslant l} = \bigcup_{k \leqslant l} [T]_k$.

Given a well founded preorder $(P, \leqslant)$ let $Tr(P, \leqslant)$ be the set of pairs $((T, \leqslant, r), p)$ where $(T, \leqslant, r)$ is a finite rooted tree and p a monotonous map from T to P taking root to root and whose restriction to $T \backslash \{r\}$ is strictly monotonous.

We define a partial order $\leqslant_{Tr}$ on $Tr(P, \leqslant)$ as follows: we say that

$$((T, \leqslant, r), p) <^0_{Tr} ((T', \leqslant, r'), p')$$

if and only if there is a subtree $S \subset T$ and an isomorphism f between S and some subtree S' of T' such that $f(r) = r$ and:

i) T extends S and T' extends S'

ii) Either $S = T$ or $p'(f(u)) < p(u)$ for some $u \in S$

iii) $p'(f(u)) < p(u)$ if $u \in lv(S) \backslash lv(T)$

iv) Each $v \in T' \backslash f(S)$ lies below $f(u)$ for some $u \in lv(S)$ for which *ii*) holds (in particular, T' extends $f(S)$).

It is easy to check that $<_{Tr}$ is anti-reflexive and transitive.

Lemma 1.4.2. *The preorder $\leqslant_{Tr}$ is well-founded.*

Proof. Suppose there is some infinite descending chain $(((T_k, \leqslant, r), p_k))_{k \in \mathbb{N}}$ in $Tr(P, \leqslant)$. For each $k \in \mathbb{N}$ take $S_k \subset T_k$ and a map f_k witnessing $((T_k, \leqslant, r), p_k) > ((T_{k+1}, \leqslant, r), p_{k+1})$. For each $m \leqslant k$, let $g^m_k = f_k \circ f_{k-1} \circ \cdots \circ f_m$ (its domain of definition might be smaller than that of the f_m). Let $U_k = \{u \in S_k \mid p_{k+1}(f_k(u)) < p_k(u)\}$, as in *ii*).

Claim 1.4.3. *For each $L \in \mathbb{N}$ there is $N_L \in\in \mathbb{N}$ such that for all $k \geqslant N_L$ $S_k \supset [T_k]_L$, f_k restricts to an isomorphism between $[T_k]_{\leqslant L}$ and $[T_k]_{\leqslant L}$ and $U_k \cap [T_k]_{\leqslant L-1} = \varnothing$.*

Proof. We will prove the existence of N_L by induction on L. Take $L = 0$ as a base case. Then $N_{-1} = 0$ does the job. The first property is satisfied by all f_k, as they take roots to roots. The second one follows from the well-foundedness of $\leqslant$.

We now deal with the induction step. Suppose N_L is given such that for any $k \geqslant N_L$ the map f_k maps $[T_k]_{\leqslant L}$ isomorphically onto $[T_{k+1}]_{\leqslant L}$ and $lv(U_k) \cap [T_k]_{\leqslant L} = \varnothing$. Now, for any

$\lambda \in [T_{N_L}]_L$ either its image v_k in T_k belongs to $lv(T_k)$ for any $k \geqslant N_L$ or eventually this image has children.

Suppose this occurs for T_M. Then property (iii) implies that $Ch(v_k) \subset S_k$ and therefore f_k restricts to a bijection between $Ch(v_k)$ and $Ch(v_{k+1})$ from that point on (it is surjective because T_k extends S_k as opposed to merely containing it as a subtree). We conclude that f_k restricts to an isomorphism between $[T_k]_{\leqslant N+1}$ and $[T_{k+1}]_{\leqslant N+1}$ for all k greater or equal than some M_L. The rest of the induction hypothesis for $L+1$ clearly follows from the well-foundedness of the base partial order. □

This, together with property (ii) of the definition of the order $<$ implies that $[T_k]_{\geqslant L+1} \neq \varnothing$ for $k \geqslant N_L$. Now consider the direct limit T_∞ of the restriction of the maps $g^{N_L}_{N_{L+1}}$ to $[T_{N_L}]_L$ with $[T_{N_{L+1}}]_L$. Let p_∞ be the map from T_∞ to $(P, \leqslant)$ which restricts to p_{N_L} on the image of $[T_{N_L}]_L$. It follows from the properties of the N_L, that p_∞ is strictly monotonous. Since the other hand T_∞ is clearly finite-branching, so by König's lemma it has to contain some infinite branch. This contradicts well-foundedness of $(P, \leqslant)$. □

Strictly speaking, this result needs only be applied to linear trees in $Tr(P, \leqslant)$ (which is simpler), but we believe it can be of use in proving the iterative procedure for analyzing general ∀∃ formulas, which unfortunately goes beyond the scope of this work.

Chapter 2

Actions on trees

2.1 $\mathbb{R}$-trees

2.1.1 Basic definitions

So called $\mathbb{R}$-trees are class of metric spaces which generalize the notion of a simplicial tree which play a crucial role in the study of the first order theory of free groups. We will give a quick survey of the main results that will be needed later and point the reader to the main source for the material of this chapter, [Chi01], for a more in detail and general account (see also [Bes02]).

Given a metric space (X, d), by an oriented segment in X we denote an isometric embedding i from a bounded interval $[a, b] \subset \mathbb{R}$, into X; we say that such i is a segment from $i(a)$ to $i(b)$. If we are only interested in i up to precomposition with an isometry of $[a, b]$, then we talk about an unoriented segment. In this case we will often blur the distinction between ι and its image, as this hardly generates any ambiguity. For example, we might say that the union of two segments is a segment, when what we mean is that formally the union of their images is the image of another segment. If instead of a bounded interval we take the entire real line $\mathbb{R}$ (and we allow ourselves to precompose by translations of $\mathbb{R}$) we obtained a geodesic line; if we take a semi-infinite interval $[a, \infty)$, a geodesic ray. A simple path in X is a continuous embedding of an interval $[a, b]$ into X, defined up to precomposition by a homeomorphism of $[a, b]$ relative to a, b.

Definition 2.1.0.1. A metric space is said geodesic if for any $x, y \in X$ there is a segment $\iota : [a, b] \to X$, with $\iota(a) = x$ and $\iota(b) = y$.

Definition 2.1.0.2 (Gromov)**.** Given a metric space (X, d), and $x \in X$, for $y, z \in X$ let $(y, z)_x = 1/2(d(x, y) + d(x, z) - d(z, y))$. Given a positive constant δ, a metric space is called δ-hyperbolic if for some (equivalently for each) $x \in X$, and for all $y, z, w \in X$:

$$(y, w)_x \geqslant \min\{(y, z)_x, (w, z)_x\} - \delta$$

We collect some background results.

Lemma 2.1.1 ([Chi01, p.43])**.** *Let (X, d) be a metric space. The following are equivalent:*

a) Any two distinct points of X are joined by unique simple path in X, represented by a segment.

b) (X, d) is geodesic and contains no subspace homeomorphic to the circle.

i) *(X, d) is geodesic*

ii) *If a point is a common endpoint of two points and their only intersection point, then their union is a segment.*

iii) *The intersection of two segments with a common endpoint is again a segment.*

We refer to any space satisfying the equivalent conditions above as an $\mathbb{R}$-tree. We will often use the expression $[x, y]$ to denote the unique oriented segment from x to y.

Lemma 2.1.2. *Let (X, d) be a tree and $x, y, z \in X$. Then either*

i) *Exactly one of $\{x, y, z\}$ belongs to the segment between the other two.*

ii) *There is a unique point w in the interior of each of the segments $[x, y]$, $[x, z]$, $[y, z]$, so that $[x, z] \cap [y, z] = [w, z]$, $[y, x] \cap [z, x] = [w, x]$, $[z, y] \cap [x, y] = [w, y]$.*

In both cases $[x, y] \cap [x, z] \cap [z, y]$ contains a unique point, which we denoted by $Y(x, y, z)$.

Call a subset of an $\mathbb{R}$-tree convex if it contains the segment between any of its points. The following is well-known:

Lemma 2.1.3. *Given a closed convex $C \subset X$ and a point $x \in X \backslash C$ there is a unique point $x' \in C$ such that $d(x, C) = d(x, x')$.*

Proof. Given any segment $[x, y]$, where $y \in C$, there is clearly a point $x_0 \in [x, y] \cap C$ closest to x. All we need to show is that x_0 is contained in $[x, z]$ for any $z \in C$. So consider any $z \in C \backslash \{y\}$ and let $u = Y(x, y, z)$. Clearly $u \in [y, z] \subset C$, so $c \notin [y, u]$. Therefore $c \in [x, u] \subset [x, z]$. □

For the following result, see the discussion in chapter 3 of [Chi01], starting on p.79.

Lemma 2.1.4. *Let g be an isometry of a $\mathbb{R}$-tree X. Then either:*

- *g fixes a point.*
- *There is a unique line invariant under g, on which g acts by proper translations. .*

In the first case we will say that g acts *elliptically* on X. In the second case we say that g acts hyperbolically on X and we refer to the aforementioned line as the *axis* of g, or $Ax(g)$ and to through constant distance $d(x, g \cdot x)$ for $x \in Ax(g)$ as the translation length of g, or $tl(g)$.

Lemma 2.1.5. *Given an isometry g of an $\mathbb{R}$-tree Y and $x \in Y$ we have $l_x(g) = tl(g) + 2d(x, Ax(g))$.*

We denote by $Fix(g)$ the set of points fixed by g. Note that this is a convex set.

2.1.2 Group actions on $\mathbb{R}$-trees by isometries

An action λ of a group G on a real tree X is called *minimal* if X contains no proper non-empty G-invariant sub-trees. It is called faithful if no $g \in G \backslash \{1\}$ fixes X point-wise. It is called *trivial* if G fixes some $x \in X$. We say that X is *minimal* if it contains no proper invariant subtree. A subgroup $H \geqslant G$ is called *elliptic* with respect to λ if the restriction of λ to H is trivial. In the absence of ambiguity as to what the intended action is we will tend to refer to the tree itself as minimal. Likewise, we might refer to elements of G itself as hyperbolic or elliptic. By the *kernel* of λ we intend the subgroup of all those elements fixing the whole space X.

An action is relative to a family $\mathcal{A}$ of subgroups of G if and only if every $A \in \mathcal{A}$ is elliptic.

We denote by $tl^{\lambda} : G \to \mathbb{R}$ the map which to each $g \in G$ associates the translation length of the isometry $\lambda(g)$. Likewise, given $x \in X$ we can consider the function $l_x : G \to \mathbb{R}_{\geqslant 0}$

defined by $l_x^\lambda(g) = d(\lambda(g)\cdot x, x)$. In both cases λ might be replaced by X or even fall from the notation altogether when the action is unambiguous. Recall that in general, given actions of a group G on sets S and S' a map $S : S' \to$ is equivariant if $f(g\cdot s) = g\cdot s$ for any $s \in S$ and $g \in G$. Given isometric actions λ and λ' of a group G on real trees X and X' we will say that X and X' are said *equivariantly isometric* if some equivariant isometry between the two exists.

Lemma 2.1.6. *Suppose we are given a finite set S of generators of a group G which acts by isometries on a $\mathbb{R}$-tree T. If x and $x\cdot y$ are elliptic for each $x, y \in S$, then the action is trivial.*

Proof. This follows from the fact that $Fix(gh) = Fix(g) \cap Fix(h)$ for any two elliptic elements g and h together with the well-known fact that the intersection of a finite set of convex subsets of an $\mathbb{R}$-tree with pair-wise non-empty intersection is itself non-empty (Helly's thoerem). □

As a consequence, a subgroup $H \leqslant G$ is elliptic with respect to λ if and only if all of its elements are. We now summarize the fundamental classifications of isometric actions λ of a group G on a real tree X.

To begin with, an action λ as above is called *abelian* if the translation length map $tr^\lambda : G \to \mathbb{R}$ is a homomorphism.

Lemma 2.1.7. *The action λ is abelian if and only if one of the following possibilities occurs:*

i) It is trivial.

ii) It restricts to an action by translations on some invariant axis.

iii) There is some infinite ray ρ of X such that $\rho\backslash Fix(g)$ is a bounded segment for any $g \in G$, in which case G is not finitely generated.(1)

Corollary 2.1.8. *An isometric action of a finitely generated abelian group A on a tree X is either trivial, or restricts to an action by translations on a geodesic line of X.*

The following lemma characterizes all remaining actions.

Proposition 2.1.9 ([Chi01][Proposition 3.7])**.** *Let G be a group acting by isometries on an $\mathbb{R}$-tree X. The following are equivalent:*

i) There are hyperbolic elements $g, h \in G$ such that $tr(gh) > tr(g) + tr(h)$

ii) There are hyperbolic elements $g, h \in G$ such that $tl([g, h]) \neq 0$.

iii) There are hyperbolic elements $g, h \in G$ such that $Ax(g) \cap Ax(h)$ is a segment of length strictly less than $tr(g) + tr(h)$.

iv) G contains a free subgroup of rank 2 which acts freely, without inversions and properly discontinously on X.

One refers to any action satisfying the properties above *irreducible*.

Lemma 2.1.10 (see [Chi01, Theorem 4.1])**.** *Suppose we are given an isometric action λ of a group G on a $\mathbb{R}$-tree X which is either irreducible or restricts to a non-trivial action by translations on a line of X. Then the union of the axis of hyperbolic elements of G is a minimal G-invariant subtree $X_{min} \subset X$. Given any other action λ' of G on an $\mathbb{R}$-tree X', if the functions tl^λ and $tl^{\lambda'}$ coincide, then the trees X_{min} and X'_{min} are equivariantly isometric.*

(1) One usually describes this by saying that G is not elliptic in X but fixes and end of X.

In the context of an action of a supergroup H of G on X we will usually denote the minimal tree of G in X as X_G.

The following is well-known.

Lemma 2.1.11. *Suppose that a group G acts by isometries on an $\mathbb{R}$-tree in such a way that some subgroup H of G of finite index fixes a point y of Y. Then the action is trivial.*

sketch. Let y_1 and y_2 be two points in the orbit of y by G at maximum distance from each other and m the midpoint of the segment between them. Given $g \in G$ of course $g \cdot y_j \in G \cdot y$. From the maximality of $d(y_1, y_2)$ one can easily see that $[y_1, y_2] \cap [g \cdot y_1, g \cdot y_2] \neq \emptyset$ and in fact $g \cdot m = m$. □

We say that a real tree on which G acts by isometries is K-acylindrical, for $K > 0$, if the point-wise stabilizer of any segment of length greater than K is trivial.

2.2 Simplicial trees: Bass-Serre theory

Bass-Serre theory provides a correspondence between each from a very general class of actions of a group G on a simplicial trees and certain types of presentations of G in terms of small building blocks (finitely many in case G is finitely generated). It is the natural setting for many essential tools in the study of the first order theory of the free group. We will roughly follow the approach found in [DD89, Chapter 2]. By a *graph* Y we intend a structure consisting of two disjoint sets. One $EY = E$ of (oriented) edges and another one $VY = V$ of vertices, together with a function $\alpha : E \to V$ and an involution $\bar{\cdot} : E \to E$ without fix points. We will use $\omega(e)$ as an abbreviation of $\alpha(\bar{e})$. If $(\alpha(e), \omega(e)) = (u, v)$ we will say that e originates at u and terminates at v, or that it is an edge from u to v and we will refer to u and v as adjacent.

A *path* in Y from a vertex u to a vertex v is a finite sequence $u = u_0, e_0, u_1 \cdots u_m = v$ ($m \geqslant 0$), where $\alpha(e_i) = u_i$ and $\omega(e_i) = u_{i+1}$ for $0 \leqslant i \leqslant m-1$. The path is closed if $u = v$ and simple if $u_i = u_j$ and $i < j$ implies $i = 0, j = m$. A graph is called connected when for any two distinct vertices u, w there is a path between them. A *simplicial tree* (or simply a tree) is a connected graph without simple closed paths. In particular, given two adjacent vertices u, v of a simplicial tree, there is a unique edge between them, which we will often denote as (u, v). For edges e, f in a tree we say that $e \succ f$ if and only if $\alpha(f) \in [\alpha(e), \omega(f)]$.

Given graphs X and Y, a cellular map $f : X \to Y$ is a map $VX \cup EY \to VY \cup EY$ preserving VX and VY and commuting with α and $\bar{\cdot}$. In case f is an inclusion, we refer to X as a subgraph of Y. Given a graph (V, E) and $v \in V$, by the star around v we intend the subgraph consisting of v, all edges at distance 1 from v, together with all the edges between any two of two said vertices.

By an action without inversions of a group G on a graph $X = (V, E)$ we intend an action of G on $E \cup V$ respecting the sets V and E, commuting with α and $\bar{\cdot}$ and such that $g \cdot e \neq \bar{e}$ for no $g \in G$ and $e \in E$. Given such an action, α and $\overline{}$ push to the origin and inverse function of a graph structure with edge st $G\backslash E$ and vertex group $G\backslash V$, the quotient graph, or $G\backslash Y$.

We will refer to a simplicial tree T endowed with an action without inversions of a group G as a G-tree. This is called *trivial* if there is some vertex of T fixed by the entire G. It is called *reduced* if the stabilizer of a vertex is never contained in that of an adjacent edge. It is called minimal if T does not contain a proper G-invariant subtree.

Fact 2.2.1. *If G is finitely generated and T minimal, then $G\backslash T$ is finite for any minimal G-tree T.*

A Bass-Serre presentation compresses in a certain sense all the information contained in a G-tree into a smaller object, finite whenever G is finitely generated.

Definition 2.2.1.1. A *Bass-Serre presentation* of a G-tree T consists of a triple $(Y_0, Y, (g_e)_{e\in Y})$, where $Y_0 \subset Y \subset T$ and Y_0 is a subtree of T satisfying the following conditions:

- p_G restricts to a bijection between EY and $G\backslash ET$ and to a graph isomorphism between Y and a maximal subtree of $G\backslash T$
- $\alpha(e) \in VY_0$ for every edge $e \in Y$
- The edge $g_e \cdot \overline{e}$ belongs to Y and $g_f = g_e^{-1}$.
- $g_e = 1$ for $e \in Y_0$

The elements g_e in the definition above are usually referred to as Bass-Serre elements. Given $x \in G\backslash T$ denote by $\tilde{x}$ be the unique preimage of x in Y.

Definition 2.2.1.2. A *graph of groups* consists of a connected graph (E, V), together with the following assignments:

- A group Γ_v for each $v \in V$,
- For all $e \in E$ a subgroup $\Gamma_e \leqslant \Gamma_{\alpha(v)}$ and an isomorphism $i_e : \Gamma_e \cong \Gamma_{\overline{e}}$, such that $i_{\overline{e}} = (i_e)^{-1}$.

Definition 2.2.1.3. Let Γ be a graph of groups and Z a maximal subtree of $|\Gamma|$. By the fundamental group of Γ with respect to Z, denoted by $\pi(\Gamma, Z)$ we intend the group defined by the following relative presentation:

$$\pi(\Gamma, Z) = \langle \bigcup_{v\in V} \Gamma_v \cup \{t_e\}_{e\in E} | \{t_e = 1\}_{e\in Z} \cup \{i_e(g) = g^{t_e},\ t_e = t_{\overline{e}}^{-1}\}_{e\in E, g\in\Gamma_e} \rangle$$

It can be shown that if G is finitely generated, then Δ_v is generated by finitely many elements together with Δ_e for all e with $\alpha(e) = v$. The natural homomorphism from each G_v to $\pi(\Gamma, Z)$ can be shown to be injective and we will think of it from now on as an inclusion. One can prove that given two maximal trees Z, Z' there is an isomorphism $\theta : \pi(\Gamma, Z) \equiv \pi(\Gamma, Z')$ compatible with the embeddings of G_v into $\pi(\Gamma, Z)$ and $\pi(\Gamma, Z')$, up to inner automorphism of one of the fundamental groups. Let $P = (Y_0, Y, (t_e)_{e\in EY})$ a Bass-Serre presentation the action $\cdot$ of a group G on a tree $T = (V, E)$. To this presentation we can associate a graph of groups $\Gamma(P)$ with underlying graph $X = G\backslash T$ and assigments:

- $\Gamma_v = Stab_G(s(v))$ for $v \in V$
- $\Gamma_e = Stab_G(\tilde{e})$ for $e \in G\backslash E$.
- For $e \in S$ the map i_e is the restriction of the inner automorphism inn_{g_e} to the subgroup G_e (in particular, the identity when $e \in EY$).

Let $Z_0 = p_G(Y_0)$. This is a maximal subtree of $G\backslash T$. Consider the map from $\{t_e\}_{e\in EX\backslash Z_0} \cup (\bigcup_{v\in VX} \Gamma_x)$ which sends each $t_e \in \pi(\Gamma, p_G(Y_0))$ to $g_e \in G$ and restricts to the inclusion on each Γ_v. One can easily check that all relations in the presentation of $\pi(\Gamma, p_G(Y_0))$ hold for their images in G, so it extends to a homomorphism $\phi_P : \pi(\Gamma, Z_0) \to G$. By a *graph of groups decomposition* of a group G we intend a triple (ϕ, Δ, Z) where ϕ is an isomorphism between G and $\pi_1(\Delta, Z)$.

The fundamental theorem of Bass-Serre theory establishes a strong correspondence between graph of groups decompositions of a group and Bass-Serre presentations of actions of the group on a tree.

Theorem 2.2.2. *For any pair* (Γ, Z) *where* Γ *is a graph of groups and* Z *a maximal subtree of its underlying graph* X *we can construct an action* $\lambda(\Gamma, Z)$ *of* $\pi(\Gamma, Z)$ *on a tree* $\widetilde{(\Gamma, Z)}$ *and a Bass-Serre presentations* $P(\Gamma, Z)$ *of it so that the following properties are satisfied:*

- *The graph of groups associated to the presentation* $P(\Gamma, Z)$ *in the way described coincides with* Γ.[(2)]
- *Suppose one is given an action* λ *of a group* G *on a tree* T *and a Bass-Serre presentation* $P = (Y_0, Y, (t_e)_{e\in EY})$ *of it and let* $Z = p_G(Y_0)$. *Then the map* $\phi_P : \pi(\Gamma(\lambda, P), Z) \to G$ *is an isomorphism and there is a graph isomorphism* $f : T(\Gamma(\lambda, P), Z) \to T$ *such that* $\phi(h) \cdot v = f(h \cdot v)$ *and pair* (f, ϕ) *sends* $P(\Gamma, Z)$ *to* P.

The two most basics examples of fundamental groups of graph of groups are so called amalgamated products and HNN extensions. Given two groups A, B and i an embedding of some subgroup C of A into B (usually one tends to think of C as a common subgroup of A and B), by the *amalgamated product* of A and B over C, denoted by $A *_B C$ we intend the quotient of $A * B$ by the normal subgroup generated by all the elements of the form $i(c)c^{-1}$ for $c \in C$. This is merely the fundamental group of a graph of groups with a single pair of mutually iverse edges and vertex groups A and B. The fundamental group of a graph of groups with a single pair of mutually inverse edges and one single vertex, with associated group A, is called an *HNN extension* of A. It is the quotient of $A * \langle t\rangle$ by the normal subgroup generated by all the elements of the form $i(c)^{-1}c^t$ for some $C \leqslant A$ and some embedding i of C into A.

The defining property of the free group $\mathbb{F}(x)$ can be restated as saying that the $Cayl(\mathbb{F}, x)$ is a tree. In fact any group which acts freely on a tree is a free and the collection of Bass-Serre elements in any presentation of this action a basis for it. More generally, any G-simplicial tree with trivial edge stabilizers is associated to a decomposition of G as a free product

$$(*_{i=1}^m G_i) * F$$

where each G_i is the stabilizer of some vertex v_i and F is a group acting freely on T. We say that a graded π-group $\mathsf{G}_{A,(P_i)_{i\in I}}$ is freely indecomposable if G admits no non-trivial action without inversions on a simplicial tree with trivial edge stabilizers relative to $\{A, P_i\}_{i\in I}$.

A metric realization of a connected graph $X = (V, E)$ is the metric space obtained by gluing together a segment of a certain positive length for each pair of mutually inverse edges according to the maps α and ω. The distance between two points is given by the shortest total length of a path between those points.

In case of a G-tree T, we require the assigned length to be the same for edges in the same G-orbit. The resulting space $|T|$ is easily seen to be an $\mathbb{R}$-tree on which G acts by isometries.

Each of the notions used to describe the action of G or particular elements of G on T defined in the previous section have discrete counterparts which hold if and only if the original holds for the geometric realization.

Let us only remark that the fact that the action of G on T is without inversions guarantees that an element of subgroup of G is elliptic in T if and only if it fixes a vertex of T. In the standard geometric realization every edge is assigned length 1.

Normal forms The following appears within the proof of the fundamental theorem of Bass-Serre theory. Let T be a G-tree, $(Z_0, Z, (t_e)_{e\in EZ})$ a Bass-Serre presentation for it and v_0 a vertex in Z_0. Let g be any element of G and consider the path from u_0 to $g \cdot u_0$:

$$u_0, f_0, u_1 \cdots u_m = g \cdot u_0$$

[(2)]In other words, $\pi(\Gamma, Z)\backslash\widetilde{(\Gamma, Z)}$ is identical to X and Γ_v coincides with the stabilizer of the lift of x in the maximal subtree of $P(\Gamma, Z)$.

then we can write g as

$$g = g_0 t_{e_0} g_1 \cdots t_{e_{m-1}} g_m$$

where $g_i \in Stab(v_i)$ for the unique translate of u_i in v_i and e_i is a translate of f_i (in the language of graph of groups the same can be expressed as saying that $g_i \in \Gamma_{v_i}$ for some closed path $u_0, f_0, u_1 \cdots u_m$ in the underlying graph). This word is reduced in the sense that for $1 \leqslant i \leqslant m-1$ we have $g_i \notin Stab(e_i)$ whenever $e_{i+1} \neq t_{e_i} \cdot \overline{e_i}$ (equivalently $g_i \notin \Delta_{\overline{e_i}}$ whenever $\overline{e_i} = e_{i+1}$)

Expansions of groups acting on trees by elliptic elements The following lemma should be well-known, but we have failed to find an adequate reference.

Lemma 2.2.3. *Suppose we are given groups $G \leqslant H$, as well as a simplicial G-tree T and a subtree $S \leqslant T$ invariant under the action of G, whose translates by elements of H span T. Assume moreover that H is generated by G and a family $\mathcal{E}$ of elliptic elements of G, invariant under conjugation by elements of G. Then the following holds:*

- *i) $\bigcup_{h \in H} h \cdot S = T$ or, equivalently, the quotient $q : G\backslash S \to H\backslash T$ induced by inclusion is onto.*
- *ii) For any vertex $v \in S$ denote by K_v the subgroup generated by $Stab_G(v)$ and $\mathcal{E} \cap Stab_H(v)$, let $\bar{\mathcal{E}} = \bigcup_{v \in VS} K_v$ and assume that for any $v \in VS$ and $e, f \in ES$ originating at v one has $e \in K_v \cdot f$ if and only if $e \in Stab_G(v) \cdot f$. Then:*
  - *a) $Stab_H(v)$ is generated by $Stab_G(v) \cup (\mathcal{E} \cap Stab_H(v))$ For any $v \in VS$.*
  - *b) q is an isomorphism.*

Proof. Clearly $\bigcup\limits_{h \in H} h \cdot S$ is connected. Given a finite product $g_1 g_2 \cdots g_m$, where each $g_j \in \mathcal{E}$, let $\bar{g}_j = g_1 g_2 \cdots g_j$ for $1 \leqslant j \leqslant m$, clearly $\bar{g}_j \cdot S \cap \bar{g}_{j+1} \cdot S \neq \varnothing$ for $1 \leqslant j \leqslant m-1$. We have shown (i), since we assume T is minimal. We now assume that two edges adjacent to $v \in VS$ are in the same $Stab_G(v)$-orbit if and only if they are in the same K_v-orbit.

Lemma 2.2.4. *Suppose that $h \cdot S \cap S \neq \varnothing$ for some $h \in H$. Then $h = kg$, for some $k \in \bar{\mathcal{E}}$ and $g \in G$.*

Proof. Clearly, since $\mathcal{E}$ is invariant under the action of G by conjugation, the subgroup N generated by $\bar{\mathcal{E}}$ is normal in H, so one can write $h = ng$ for $n \in N$ and some $g \in G$. All we have to show is that $n \in \bar{\mathcal{E}}$. So write it as $a_1 a_2 \cdots a_r$, where $a_j \in \bar{\mathcal{E}}$; we can assume that r is minimal among all such possible representations of n. We deal with case $k = 2$ first. For the sake of contradiction we assume that $a_1 a_2 \notin \bar{\mathcal{E}}G$. Let F_j be the fixed point set of a_j for $j = 1, 2$. Take $u_1, u_2 \in S$ be such that $a_1^{-1} \cdot u_1 = a_2 \cdot u_2 := v$ and $L(a_1, a_2)$ the minimal value of $d(u_1, F_1) + d(u_2, F_2)$ for any such u_1, u_2. We can assume that $L(a_1, a_2)$ is minimal among all a_1, a_2 such that $a_1 a_2 \in h \cdot G$. Denote by w be the vertex of S closest to v. For each $j = 1, 2$ one of the following two possibilities occurs:

- $w \in F_j$
- $w \notin F_j$ and there are two consecutive edges $e, f \in S$ in the path from u_j to v such that $a_j^{-1^j} \cdot e = \overline{f}$. In particular, a_j fixes $\omega(u_1) = u_2$.

The first possibility cannot be the case for both values of j, since then $F_1 \cap F_2 \neq \varnothing$, which implies that $a_1 a_2 \in \bar{\mathcal{E}}$. By symmetry we can assume the second holds for $j = 1$ (otherwise work on $a_2^{-1} a_1^{-1}$ and then use the G-invariance of $\bar{\mathcal{E}}$). Our condition on K_v implies the existence of some $g_0 \in Stab(u)$ sending $\overline{f}$ to e, which in turn implies the inequality $d(g_0 a_1^{-1} \cdot u_1, u_1) = d(g_0 \cdot v, u_1) < d(v, u_1)$.

On the other hand $d(a_2^{g_0^{-1}} \cdot (g_0 \cdot u_2)) = d(g_0 a_2^{-1} \cdot u_2, g_0 \cdot u_2) = d(v, u_2)$. This means that $L(a_1 g_0^{-1}, a_2^{g_0^{-1}}) < L(a_1, a_2)$, contradicting the minimality assumption, since $(a_1 g_0^{-1}) a_2^{g^{-1}} = a_1 a_2 g^{-1}$.

Suppose now that $k > 2$ and let $\bar{a}_j = a_1 \cdots a_j$ (so $\bar{a}_0 = 1$), $S_j := \bar{a}_j \cdot S$ and $d_l = d(S_l, S)$ for $0 \leqslant j \leqslant m$. Pick any $1 < j_0 < m$ maximizing d_{j_0} and $D = d_{j_0}$. In view of the previous case we can assume that $D > 0$. We claim that $S_{j_0-1} \cap S_{j_0+1} \neq \emptyset$. Indeed, both intersect S_{j_0} and therefore they must contain the unique ($D > 0$) point of S_{j_0} at minimal distance from S. □

Using this the following additional claim, from which properties (i) and (iib) are easy to prove.

Lemma 2.2.5. *Let $h \in H$ and suppose we are given vertices $u, v \in S$ such that $h \cdot u = v$. Then $h = kg$ for some $k \in K_v$ and $g \in G$ such that $g \cdot u = v$.*

Proof. Given any such h consider a pair $(k_0, g_0) \in \bar{\mathcal{E}} \times G$ such that $k_0 g_0 = h$ and $d = d(g_0 \cdot u, v)$ is minimal for this property. If $d = 0$ we are done, by the previous case, so assume this is not the case. Let $v = g_0 \cdot u = n_0 \cdot u$. Just as in the proof of the previous lemma, the segment from $g \cdot u$ to v must contain two consecutive edges e, f such that $\omega(e) = \alpha(e) := w$ and $n_0 \cdot e = \overline{f}$ and some $g_1 \in Stab_w(G)$ such that $g_1 \cdot e = \overline{f}$ as well. It is clear then that $d(g_0 g \cdot u, v) < d(g \cdot u, v)$. Since k fixes w as well, $k g_0^{-1} \in \bar{\mathcal{E}}$ which contradicts the minimality of d, since of course $g_0 g \in G$. □

□

2.3 Operations on trees

2.3.1 Collapses and blow-ups

Given a graph X and $F = \overline{F} \subset EX$ let $p_F : X \to C_F(X)$ be the cellular map obtained by collapsing the edges in F. More precisely, let $X_F = F \cup \alpha(F)$ and $\mathcal{C}$ the collection of its connected components. We let

$$V(C_F(X)) = V(X)/\{u \sim v | u, v \in V(C),\, C \in \mathcal{C}\}$$
$$E(C_F(X)) = E(X) \backslash \bigcup_{C \in \mathcal{C}} EC$$

where the new incidence function is obtained by post-composing the previous one with the quotient map and the inverse operation the adequate restriction of the old one. The cellular map p_F is the obvious quotient mapping each component C to a point.

For simplicity, in this situation above let $Y^F = C^F(Y)$ for $F \subset EX$ and Y a subgraph of X. Assume now that X is a G-tree and F a G-invariant family. Clearly X^F inherits a G-graph structure from X and there is an isomorphism $(G\backslash X)^{(G\backslash F)} \equiv (G\backslash X^F)$ compatible with p^F and $p_{(G\backslash F)}$. Suppose we are given a G-tree T, as well as some G-invariant $F \subset ET$ take any Bass-Serre presentation $(Z_0, Z, (t_e)_{e \in EZ})$ of it such that whenever $e \in (G\backslash F)$ for some edge $e \in Z\backslash Z_0$ we have $f \in (G\backslash F)$ for each f in the unique path in Z between the endpoints of e. Then $(G\backslash Y^F)$ is a maximal subtree of $(G\backslash X^F)$, and $P' = (Z_0^F, Z^F, \{g_e\}_{e \in EZ^F})$ a Bass-Serre presentation of the G-tree T^F. It is easy to see that at the level of graph sof groups we can describe the operation as replacing each subgraph in the image of a connected component of F by the group generated by the corresponding generators.

The inverse operation is what one calls a *blow-up*, or refinement. The pertinent data in this case is some $v_0 \in VT$ and some $Stab(v_0)$-tree S in which $Stab(e)$ is elliptic for each edge e incident to v_0. Using this one can construct G-tree in the following way. First, for each $[g] \in G/Stab(v_0)$ take a copy $S^{[g]}$ of S. We can assume that $S^1 = S$. Pick any set of representatives $\mathcal{R}$ of $G/Stab(v_0)$ and for each $r \in \mathcal{R}$ an isomorphism $\theta^r : S^{[1]} \to S^{[r]}$. There is a unique way of extending this assignment to a collection of graph isomorphisms $\theta^g_{[h]}$ between $S^{[h]}$ and $S^{[gh]}$ for any $[g] \in G/Stab(v_0)$ and any $g \in G$ in such a way that $r^g_{[1]} = \lambda(g)$ for any $g \in Stab(v_0)$ and $r^{gh}_{[1]} = r^g_{[h]} \circ r^h_{[1]}$. Take any system of representatives $\mathcal{ER}$ of orbits by G of edges of T in the family $\mathcal{E}_{v_0}$ of all the edges originating at a at a translate of v_0. Any assignment of a vertex $p_e \in S^{\alpha(e)}$ to each $e \in \mathcal{ER}$ can be extended uniquely to the whole $\mathcal{E}_{v_0}$ in a way that the equality $p_{h \cdot e} = \theta^h_{[g]}$ holds for any $e \in \mathcal{E}$ with $\alpha(e) = g \cdot v$ and $h \in G$. The tree $T_{S,(p_e)_{e \in \mathcal{E}}}$ (often simply T_S, when context allows) is defined by:

$$VT_{S,(p_e)_{e \in \mathcal{E}}} = (VT \backslash (G \cdot \{v_0\})) \cup \bigcup_{g \in G} VS^{[g]}$$

$$ET_{S,(p_e)_{e \in \mathcal{E}}} = \{\bar{e} \mid e \in ET\} \cup (\bigcup_{g \in G} ES^{[g]})$$

where for each $e \in ET$ we have $\bar{\bar{e}} = \bar{\bar{e}}$ and the copy $\bar{e}$ of e originates either at $v = \alpha(e)$ in case $\alpha(e) \neq G \cdot v_0$ or else at p_e. Note that although there is no canonical bijection between the vertex and edge sets of S and $S^{[h]}$, the maps θ do induce a canonical isomorphism between $Stab(v_0)\backslash S$ and $Stab(v_0)\backslash S^{[h]}$.

2.3.2 Lifting decompositions through blow-up and collapse

Clearly, if we collapse all edges inherited from S in the blow-up T_S described in the previous subsections we recover the original tree S. One can also collapse those edges inherited from T instead of from S. This operation has very nice properties under certain circumstances, a fact that we will be useful later.

Lemma 2.3.1. *Suppose we are given a non-trivial simplicial H-tree T, $G \leqslant H$, a G-invariant subtree U of T and for some $v_0 \in V$ a simplicial $Stab_H(v_0)$-tree S. Assume furthermore that:*

a) Two edges of U are in the same H-orbit only if they are already in the same G-orbit

b) $Stab_H(e)$ is elliptic in S for any edge e adjacent to v_0.

c) For any $h \in H$ such that $h \cdot v_0 \in T_G$ (equivalently, for all h in some system of representatives of $G\backslash H$) if we let $\mathcal{E}_v$ be the set of edges of T_G originating at v then the subgroup generated by

$$Stab_G(h \cdot v_0)^h \cup \{Stab_H(e)^h\}_{e \in \mathcal{E}_{h \cdot v_0}}$$

fixes a vertex in some $Stab(v_0)$-invariant family $\mathcal{O} \subset VS$.

Then a simplicial H-tree S' exists together with an equivariant embedding ψ of S into T with the following properties:

i) The union of translates of $\psi(S)$ cover S'.

ii) For each $v \in VT\backslash H \cdot \{v\}$ the group $Stab(v)$ fixes some vertex in $H \cdot \psi(\mathcal{O})$.

iii) $Stab(\psi(e)) = Stab(e)$ *for any* $e \in ES$

iv) $Stab(v) = Stab(\psi(v))$ *for any* $v \in VS$

v) G *fixes some vertex* v_G *in* $T \backslash H \cdot \psi(\mathcal{O})$. *Let* $H_0 = Sb(v_G)$

vi) If $e, e' \in ET$ *are in distinct* H_0*-orbits, then* $\psi(e), \psi(e')$ *are in distinct* H*-orbits.*

vii) Any $v, v' \in VT \backslash \mathcal{O}$ *in distinct* H_0*-orbits are also in distinct* H*-orbits.*

Furthermore, for some $v_G \in \psi(\mathcal{O})$ *fixed by* G *if we let* H_0 *be the stabilizer of* v_G *and* U_0 *be the convex closure in* T *of the union of translates of* U *by elements of* H_0*. The intersection with* U_0 *of either an* H*-orbit of edges or an* H*-orbit of vertices not containing* v_0 *is either empty or an* H_0*-orbit.*

Proof. We first construct an H-tree $T_S := T_{S,(p_e)_e}$ by equivariantly blowing up all vertices in $H \cdot v_0$ to a copy of S. Condition (b) guarantees this can be done (the result is generally non-unique). Afterwards, we collapse all edges of T_S inherited from T. As ψ is simply the isomorphism between S and $S^{[1]} \subset T_S$. The resulting H-tree T_S^c together with the obvious map from S into T_S^c clearly satisfies properties (i),(v), (iii), (iv) and the last part of (vii) by construction. Let $\mathcal{V}^G$ be a set of representatives of $G \backslash (H \cdot v_0)$. For any $u \in \mathcal{V}^G$ chose a set of representatives $\mathcal{ER}_r^G$ of the orbits by $Stab(v_0)$ of the family of edges of T originating at u. Our first hypothesis implies that for any $u \in \mathcal{V}^G$ and any distinct $u, u' \in \mathcal{V}^G$ and $h \in H$ the sets $h \cdot \mathcal{ER}_u^G$ and $\mathcal{ER}_{u'}^G$ are disjoint.

In particular, we can assume that $\bigcup_{u \in \mathcal{U}^G} \mathcal{ER}_u^G$ is a subset of the transversal $\mathcal{ER}_{v_0}^H$ of the set of edges originating in $H \cdot v_0$ used in the blow-up construction. In virtue of the third hypothesis, for any $r \in \mathcal{R}$ we can arrange so that all $e \in \mathcal{ER}_u^G$ have the same attaching point $p_r = p_e \in \mathcal{O}_u$, where $\mathcal{O}_u$ is the image of $\mathcal{O}$ in the copy of S at u. This p_r is fixed by $Stab_G(r \cdot v_0)$.

Clearly, every pair of edges (e, f) of U originating at a common translate of v_0 is an H-translate of a pair of the form $(e_0, k \cdot f_0)$, where $k \in Stab_G(u)$, $e, f \in \mathcal{ER}_u^G$ and $u \in \mathcal{V}^G$, so $p_e = p_f$ in that case as well. This implies that subtree spanned by the images of the edges of U in $T_{S,(p_e)_e}$ is connected, and hence maps to a unique point $v_G \in T_S^c$, fixed by G. The preimage C of v_G in T_S has to contain an attaching point (since it does not fill the whole T_S), so $v_G \in H \cdot \psi(\mathcal{O})$. The same argument applies to $Stab(v)$ for $v \in VT \backslash H \cdot v_0$. On the other hand C is H_0-equivariantly isomorphic to soem subtree of T containing U, since the collapse map to T does not touch any of its edges. If two vertices $v, v' \in C$ inherited from T or two edges $e, e' \in C$ differ by translation by some $h \in H$ then $h \cdot C \cap C \neq \varnothing$, which implies $C = h \cdot C$, since C is a connected component of the complement of the union of the collection of all copies of S in T_S we conclude that $h \in Stab(v_G)$, which proves the last claim. □

2.3.3 Foldings and slides

Given a graph $X = (V, E)$, and $e, f \in E$ such that $\alpha(a) = \alpha(f)$ by the *fold* of e and f we intend the quotient map $F_{e,f} : X \to X'$, where $X' = (V', E')$, V' obtained from V by identifying $\omega(e)$ and $\omega(f)$ and E' is obtained from E by identifying e and f. Given a family $\mathcal{P}$ of such pairs we can consider the (inverse) limit graph obtained by quotiening each of those pairs one by one (in successive steps we use the image of the initial pairs in the current quotient tree). It can be shown this does not depend on the order in which the different foldings are performed. There is a natural cellular map from X onto the resulting graph preserving the set of edges.

If X is a tree, the result of folding $\mathcal{P}$ is a tree as well. If we are dealing with a simplicial G-tree the operation is to be interpreted equivariantly. Namely, we not only fold (e, f) but also any image of the pair by the action of G. The result is again a $\mathcal{G}$ tree. See [Sta91] or [Dun98] for more details. The effect of an equivariant folding at the level of graph of groups has a varied typology. We will restrict ourselves to the case relevant to us.

Lemma 2.3.2. *Let T be a simplicial G-tree and suppose we are given $v \in VT$ and $e_1, \cdots e_m$ origiating at v and in distinct orbits by the action of G and suppose that for $1 \leqslant j \leqslant m$ we are given $Stab(e_j) < H_j \geqslant Stab(\omega(e_j))$. Assume that for any $1 \leqslant j \leqslant m$ the vertex $\omega(e_j)$ is not in the orbit of v. Let $f_{\bar{H}} : T \to T^f_{\bar{H}}$ the equivariant fold of all the pairs $(\overline{e_j}, h \cdot \overline{e_j})$, where $h \in H_j$. Then $Stab(f_{\bar{H}}(v))$ is equal to the amalgamated product of $Stab(v)$ with H_j over $Stab(e_j)$ and the stabilizer of $f_{\bar{H}}(e_j)$ is equal to H_j.*

Proof. That the stabilizer of the image of e_i is equal to H_i is immediate from the fact that no fold is performed at $\epsilon(e_i)$.

The stabilizer H of $f_{\bar{H}}(v)$ coincides with the set-wise stabilizer of its preimage $\mathcal{P}$ in T, which is precisely the set of all vertices v' in the orbit of v that can be joined to v by some path of the form $v_0 = v, f_0, u_0, f_0', v_1, e_1 \cdots v_m = v'$, where for each $0 \leqslant j \leqslant m-1$ edges f_j and $\overline{f_j'}$ are folded together (we are using the assumption that none of the ends of an edge which is folded is in the orbit of v). The union of all those paths is a subtree S whose intersection with $G \cdot v$ is precisely $\mathcal{P}$, so that the set-wise stabilizer of S is equal to H. We claim that the action of H on S is dual to the graph of groups decomposition described above. It is easy to check that $\mathcal{P}$ is a single orbit by the action of H. This implies that the orbits by the action of H of edges in S are in bijective correspondence with $e_1, \cdots e_m$. Since the star around $\omega(e_j)$ in S contains only edges in the orbit of e_j for $1 \leqslant i < j \leqslant m$ vertices $\omega(e_j)$ and $\omega(e_i)$ must be in distinct orbits by the action of H. This implies the action of H on G is of the required form. ☐

Note that the situation above $f_{\bar{H}}$ induces an isomorphism from $G\backslash T$ to $G\backslash T^f_{\bar{H}}$.

Another equivariant move that can be performed on a tree is a *slide*. It moves the origin of of some edge e with $\alpha(e) = v$ with v' (as well as that of all other edges in the orbit of e, in an equivariant way) on the condition that there is some edge $f = (v, v')$ such that $Stab(e) \leqslant Stab(f)$.

2.4 Trees and surfaces

Let Σ be a non-simply connected compact surface with boundary. We quickly recall some notions on the topology of compact surfaces. For more details the reader is referred to [FM12]. We will refer to those compact surfaces with (possibly empty) boundary, Σ, for which either $\xi(\Sigma) \leqslant -2$ or Σ is a torus minus an open disk as *big*. We say that a surface is *closed* if contains no boundary. By a simple closed curve (abbreviated as s.c.c.) in a surface Σ we intend an homeomorphic embedding from a circle into Σ. We will usually be interested in the class $\alpha = [a]$ of a curve a modulo homotopy and reparametrization (so not an oriented curve), rather than a itself.

There are two mutually excluding possibilities. Either α has a neighbourhood homeomorphic to an annulus or one homeomorphic to a Möbius band; in the first case we refer to the curve as two-sided. A s.c.c. will be called *essential* if it is two-sided and is not homotopic to a boundary component of Σ. A subsurface $\Sigma' \subset \Sigma$ will be called essential if none of its boundary components are nullhomotopic in Σ. By an essential system of curves of Σ, δ, we intend a collection of distinct homotopy classes of essential s.c.c. in Σ. Sometimes it is convenient to think of it as the union of disjoint representatives of such classes.

Recall that for any path connected subspace $X \subset \Sigma'$ of Σ, $*' \in X$, any homotopy class of continuous paths $[p]$ from $*$ to $*'$ defines a homomorphism $\iota_{[p]}$ from $\pi_1(X, *')$ into $\pi_1(\Sigma, *)$.

The effect of chosing a different p is that of postcomposing $\iota_{[p]}$ by some inner automorphism of $\pi_1(\Sigma, *)S$ and unless $* = *'$, there is no canonical one, but in most of the applications we are interested in we will be usually able to ignore this issue. The homomorphism ι_* is injective in case of an incompressible subsurface (one none of whose boundary connected compoenents is null-homotopic in S). We refer to the cyclic group in one of the conjugacy classes associated with the boundary of Σ as a boundary subgroup and to any of its members as a boundary element. The collection $Homeo_\partial(\Sigma)$ of all homeomorphisms of Σ fixing its boundary components, quotiened by the equivalence relation any two such homeomorphisms which are related by an isotopy preserving the boundary components set-wise. We will refer to the resulting group as the modular group of Σ, or $Mod(\Sigma)$.(3) It follows from by the Dehn-Nielsen Baer theorem that the group, which we denote by $Mod(\pi_1(\Sigma))$ of all automorphisms of $\pi_1(\Sigma)$ induced by homeomorphisms of the surface relative to the boundary has finite index in the group of all the automorphisms of $\pi_1(\Sigma)$ which restrict to inner automorphisms on peripheral subgroups. In truth, the classic result (see [FM12, Theorem 8.1]) is only valid for orientable surfaces; an appropriate generalization to non-orientable surfaces can be found in [Fuj02, p.278].

The action of $Homeo_\partial(\Sigma)$ on Σ descends to a well-defined action of $Mod(\Sigma)$ on the set of homotopy classes of essential simple closed curves in Σ and on the set of essential systems of curves on Σ. It is possible to show the latter has finitely many orbits. Given an essential system of curves δ, we denote by $[\Sigma]_\delta$ the collection of pieces in which δ cuts Σ, i.e., the (closures of) the connected components of $\Sigma\backslash\delta$. Of course, these are only well defined up to isotopy of the surface. The system δ dertermines an action λ_δ without inversions of S on a simplicial tree T_δ whose associated graph of groups Δ has the following properties:

i) There is a bijective correspondence by which each $\alpha \in \delta$ corresponds to some orbit $S \cdot e$ of edges of T_δ, so that the family $\{Stab(s \cdot e)\}_{s\in S}$ coincides with that of cyclic groups associated to α.

ii) There is a bijective correspondence associating every $\Sigma_0 \in [\Sigma]_\delta$ with an orbit $S \cdot v$ of vertices of T_δ, so that the family $\{Stab(s \cdot v)\}_{s\in S}$ coincides with all the mutually conjugate images of $\pi_1(\Sigma_0)$ in S by the embeddings associated to the inclusion map.

Alternatively, consider the preimage $\mathcal{D}$ of δ by the universal covering map $p : \widetilde{\Sigma} \to \Sigma$. Necessarily $\widetilde{\Sigma}$ is homeomorphic to an infinite plane. The family $\mathcal{D}$ is invariant under the action of S and consists of infinitely many disjoint bi-infinite curves. Each one of them, say $\tilde{\alpha}$, projecting to α, separates D_2 into two connected components and is stabilized by some $Z \in C^{\Sigma}_{\alpha}$. The tree T_δ can be obtained by taking a vertex V_U for each connected component U of $D\backslash(\cup\mathcal{D})$ and a pair of edges $e_{\tilde{\alpha}}, \overline{e_{\tilde{\alpha}}}$ for each $\tilde{\alpha} \in \mathcal{D}$ joining the two connected components meeting at $\tilde{\alpha}$.

Clearly S acts isomorphically on T, with edge stabilizers in C^{Σ}_{α}. That T is a tree comes from the fact that the $\tilde{\alpha}$ separate D. The hypothesis that α is one-sided guarantees that the action is without inversions.

Alternatively, the dual graph of groups decomposition can be obtained using Van Kampen's theorem, which provides the required presentation for $\pi_1(\Sigma)$.

Given two homotopy classes of essential s.c.c., α, β, the intersection number of α and β, denoted by $i(\alpha, \beta)$ is the minimum $n \in \mathbb{N}$ such that there are representatives $a \in \alpha$ and $b \in \beta$ which intersect in exactly n distinct points.

(3) In the literature this is usually defined as the modular group of the corresponding punctured surface (in which a point instead of a closed disk is removed from the closed surface). Likewise, when referred to an orientable surface, in absence of punctures the modular group stands for the set of classes of orientation preserving homeomorphism exclusively. If we include non-orientable homeomorphisms as well (in the empty boundary is empty) we are left with what is usually called the extended mapping class group.

Fact 2.4.1. *The intersection number $i(\alpha,\beta)$ coincides with the translation length of a generator of a subgroup associated to α in T_β.*

We will refer to those compact surfaces with (possibly empty) boundary, Σ, for which either $\xi(\Sigma) \leqslant -2$ or Σ is a torus minus an open disk as big.

Part of the important role played by surface groups in the theory of simplicial actions of trees has to do with the following property (see [MS84, Theorem III.2.6] or the proof of [GL10, Lemma 9.4]):

Lemma 2.4.2. *Let S be the fundamental group of a compact surface with boundary Σ supporting a hyperbolic metric (i.e. neither a torus nor the connected sum of one or two projective planes) and T a minimal non-trivial $\pi_1(\Sigma,*)$-tree in which all peripheral subgroups of S are elliptic. Then there is an equivariant map f from T to T_δ for some essential system of curves δ on Σ. If all stabilizers of edge groups in T are cyclic then f can be taken to be an isomorphism.*

2.4.1 Geometric abelian decompositions

We will borrow the notion of a geometric abelian decomposition from [BF09], with only slight modifications. Let G be a finitely generated group and $\mathcal{R}$ a family of subgroups of G. By a geometric-abelian G-tree we intend a simplicial G-tree T with abelian edge stabilizers together with a partition of the set of vertices V of T, into G-invariant sets V_a,V_s and V_r, the set of abelian, surface and rigid type vertices respectively, such that the following properties are satisfied:

i) For $v \in V_a$ the group $Stab_G(v)$ is abelian. Additionally, we distinguished what we will call the generalized peripheral subgroup of $Stab(v)$, or $Per^*(v)$. This is a free summand of $Stab(v)$ containing the peripheral subgroup $Per(v)$, which is the smallest summand containing the stabilizer of each of the edges incident at E.

ii) Each $v \in V_s$ is endowed with an isomorphism θ_v between $Stab(v)$ and the fundamental group of a compact surface with boundary (possibly empty) Σ_v, which admits essential simple closed curves [(4)] . Each edge stabilizer of an edge incident to v is mapped by θ_v to some finite index subgroup of the peripheral subgroup of $\pi_1(\Sigma)$ and the map thus induced from the set of orbits of edges incident to v to the family of boundary components of Σ is injective.

We say that such T is relative to $\mathcal{R}$ in case for any $R \in \mathcal{R}$:

- R is elliptic in T.
- For any $v \in V_s$ the intersection of R with $Stab(v)$ is contained in a boundary subgroup .

In this context, if the preimage of each boundary subgroup of $\pi_1(\Sigma_v)$ contains a finite index subgroup which is either conjugate into $\mathcal{R}$ or the stabilizer of an edge incident to v we say that all boundaries are used. A sufficient condition for all boundaries to be used is for G to be freely indecomposable relative to $\mathcal{R}$ [(5)] . A graph of groups decomposition Δ dual to T, together with the induced partition of the vertices of the underlying graph and peripheral subgroups of the abelian type vertex groups is called a geometric abelian decomposition of G. In this context we will often talk about rigid/abelian/surface vertex stabilizers or vertex groups.

---

(4) In other words,

(5) The fundamental group S of a surface with boundary is a free group and given any boundary component any collection $\mathcal{B}$ of generators of each of the remaining boundary components can be extended to a base of S. In particular S admits a free decomposition relative to $\mathcal{B}$.

By a trivial geometric abelian tree we intend one in which G stabilizes a vertex of rigid type. The surface Σ_v can have empty boundary, in which case the definition implies the tree T consists of a single point and $\mathcal{R}$ can only contain the trivial group.

We will say that a geometric abelian tree (or the corresponding geometric abelian decomposition) is commutative transitive if every stabilizer is its own centralizer in the stabilizer of at least one of the two endpoints.

2.4.2 Pinching

Let Δ be a geometric abelian decomposition of a group finitely generated group G. By an essential system of curves in Δ we will intend simply the union of an essential system of curves Σ_v for each surface type vertex v of the underlying graph (for that purpose, we regard surfaces corresponding to distinct vertices as distinct and disjoint). Denote by Δ_δ the graph of groups decomposition obtained by blowing each surface type vertex v in the tree dual to Δ to a copy of $T_{\Sigma_{[v]}\cap\delta}$ (since edge groups incident to v fix a unique point in $T_{\Sigma_{[v]}\cap\delta}$, the result is unique). This can be seen as a geometric abelian decomposition as well, whose surfaces are the pieces in which δ cuts each of the surfaces in Δ. Let $pinch_\delta : G \to G^p_\delta$ be the quotient by the normal subgroup generated by the cyclic groups corresponding to simple closed curves in δ. This amounts to collapsing to a point each curve in some disjoint set of representatives of the classes in δ making the corresponding edge groups. Denote by Δ^p_δ the resulting graph of groups decomposition of G^p_δ and by $Pinch(\Delta,\delta)$ the union of all surfaces resulting from those of $Cut(\Sigma_v, \delta\cap\Sigma_v)$ by collapsing the curves in δ to a point. We distinguish between the set $Pinch^{in}(\Sigma,\Delta)$ of all of those which contain some boundary component of the original Σ_v and that $Pinch^{ex}(\Sigma,\Delta)$ of those which do not. Let Δ^c_δ be the free decomposition obtained by collapsing the non-trivially stabilized edges of Δ_δ and X^c_δ its underlying graph. This is of the form:

$$(*_{i=1}^m A_i) * (*_{\Pi\in Pinch^{ex}(\Delta,\delta)}\pi_1(\Pi)) * \mathbb{F}_{b_1(X^c_\delta)} \tag{2.1}$$

where the last factor is the free group generated by the Bass-Serre elements and $b_1(X)$ stands for the first betti number of the graph X, while $A_1, A_2\cdots A_m$ are what we call the internal pinching factors, that is, the images of the fundamental groups of the finitely many graph of groups into which Δ^p_δ decomposes after removing the edges associated with the curves in δ. By the exterior rank of Δ^p_δ we intend the sum of the maximum rank of a free images of the remaining factors, what we call external factors.

Lemma 2.4.3. *Suppose we are given a closed compact surface Σ which is not a sphere. Then any homomorphism f from $\pi_1(\Sigma)$ to a free group $\mathbb{F}$ pinches some maximal (in the set theoretical sense) essential system of curves in Σ.*

sketch. In case $\xi(\Sigma) = 0$, i.e., when Σ is either a klein bottle or a torus, it follows immediately from the presentation of $\pi_1(\Sigma)$ and the fact that free groups are CSA with cyclic maximal abelian groups that any such morphism must have a cyclic image. The sought essential system of curves consists of a single curve that be then be found by hand.

As for the other cases, a straightforward induction argument implies that it is enough to show the existence of some essential simple closed curve in Σ which is pinched by f. This can be deduced from 2.4.2, applied to the action of $\pi_1(\Sigma)$ on the Cayley graph of $\mathbb{F}$ induced by f (see also [Rez97, Theorem 5.2]). □

2.4.3 Seifert type actions of surface groups on real trees

The notion of a singular measured foliation was introduce by Thurston in his groundbreaking work on the classification of homeomorphisms of surfaces (see [T+88] and [FLP12]). It generalizes that of an essential system of curves on a surface. Discrete versions of that notion are often useful: this is the approach found in [LP97], which uses simplicial complexes, or [BF95], which uses band complexes (a finite graph to which foliated squares are attached along the sides perpendicular to the foliation, plus some additional) are used instead.

What follows is only an intuitive sketch. Roughly speaking, given a surface with boundary Σ, a singular foliation on Σ is a partition $\mathcal{F} = \{F_\lambda\}_{\lambda\in\Lambda}$ of Σ into path-connected sets called leafs, admitting a differential atlas $\mathcal{A}$ of Σ, such that for some finite set $x_1, x_2 \cdots x_m$ of points of the surface (possibly on its boundary) the following properties are satisfied. First of all, any boundary component is contained in a leaf. For any chart $\phi : U \to \phi(U) \subset \mathbb{R}^2$ in $\mathcal{A}$ such that $U \cap \{x_i\}_{i=1}^m = \varnothing$ the intersection of each F_λ with U is equal to $\phi^{-1}(\bigcup_{i\in I}\{(x,y)\,|\,y = c_i\})$ for some countable collection $\{c_i\}_{i\in I}$ of constants. For any other chart $\phi : U \to \phi(U) \subset \mathbb{R}^2$ of $\mathcal{A}$ the set U contains a unique point x_{i_0}, sent to $(0,0)$ by ϕ and any F_λ is the pullback of a similar set by the map $g_k \circ \phi$, where $g_k : \mathbb{R}^2 \to \mathbb{R}^2$, is some branch of the complex exponentiation $z \mapsto z^{\frac{k}{2}}$, for $k = 1, 3, 4\cdots$ up to identifying $\mathbb{R}^2$ with $\mathcal{C}$. A leaf that contains one of the x_i is called singular. An arc α in Σ is said *transverse* to $\mathcal{F}$ if it can be covered by a union of domains of charts $\phi : U \to \phi(U)$ in $\mathcal{A}$ with the property that for any of them $\phi(\alpha \cap U)$ does not intersect any of the level sets $\{y = c\}$ twice (resp. the preimage of such a set by f_k). A measured foliation on Σ is given by a pair $(\mathcal{F}, m)$, together with a non-negative real valued function m on the set of transverse paths between points of Σ such that $m(\alpha) = m(\alpha')$ for paths α and α' whenever there is a homotopy between α and α' keeping any point in the same leaf.

Any measured foliation lifts to one, $(\tilde{\mathcal{F}}, \tilde{m})$ on the fundamental cover $\tilde{\Sigma}$ of Σ invariant under the action of $\pi_1(\Sigma)$ on $\tilde{\Sigma}$. One can define a pseudo-metric on the set $\tilde{\mathcal{F}}$ by letting the distance between to leaves be equal to the infimum of the measure of the paths between them and that the corresponding quotient metric space $X = \tilde{\Sigma}/\mathcal{F}$ is an $\mathcal{R}$-tree. It is clear that if no leaf of $\mathcal{F}$ contains a path homotopic to an essential simple closed curves the only elements of $\pi_1(\Sigma)$ fixing a $\tilde{\mathcal{F}}$ are the boundary subgroups. Under certain additional assumptions the quotient becomes nice enough, so that an element $\pi_1(\Sigma)$ fixes a point of $Y_{\mathcal{F}}$ if and only if it is contained in a boundary subgroup. One can also assume each boundary component for contains a singularity.

The resulting action of $\pi_1(\Sigma)$ on X is what is known as a Seifert type action. Its main properties are the fact that no non-trivial element fixes a non-degenerate segment and that distinct boundary subgroups fix distinct points of X.

2.5 Simplicial trees and group automorphisms

Let $(Y_0, Y, (t_e)_{e\in EY})$ be a Bass-Serre presentation for an action without inversion of a finitely generated group G on a simplicial tree T. Fix $e \in Y$ and $c \in Z_{Stab(e)}(Stab(\omega(e)))$ and for each $v \in VY$ let c_v be equal to 1 if e is in the unique path from $\alpha(e)$ to v in Y_0 and c otherwise. Likewise, let c_f be equal to c if $e \succeq f$ and 1 otherwise. Consider the map which sends:

- $g \in Stab(v)$ to g^{c_v} for $v \in VY$.
- t_f to $c_f^{-1} t_f c_f$ for $f \in E\backslash Z$

It is easy to check that this induces a homomorphism from G to itself, $\tau_{e,c}$, which we refer to as the Dehn twist over e by c. It is also easy to check that $\tau_{e,c} \circ \tau_{e,c^{-1}} = \tau_{e,c^{-1}} \circ \tau_{e,c} = Id$, so that $\tau_{e,c}$ is in fact an automorphism of G. We will refer to e as the base of $\tau_{e,c}$, denoted by $B(\tau_{e,c})$.

Now suppose we are given $v_0 \in VY$ and an automorphism σ of $Stab(v_0)$ compatible with the incidence structure, by which we mean that $\sigma\restriction_{Stab(e)}$ restricts to conjugation by some element $c_e \in Stab(v_0)$ for each $e \in E$ with $\alpha(e) = v_0$.

Then there is a natural way of extending σ to an endomorphism $\bar{\sigma} = \bar{\sigma}_{(c_e)_e}$ of G such that:

i) $\bar{\sigma}$ restricts to σ on $Stab(v_0)$

ii) $\bar{\sigma}$ restricts to inn_{c_v} on $Stab(v)$, for $v \in VY$, where $c_v = c_{in(e)}$ for $in(e)$ the first edge in $[v_0, v]$

iii) $\bar{\sigma}(g_e) = c_e^{-1} t_e c_{\bar{e}}$ for $e \in E$ where $c_e = c(u)$ for the unique vertex $u \in Z_0 \cap G \cdot \omega(e)$.

One can check, just as before, that this induces an isomorphism of G with inverse $\overline{\sigma^{-1}}_{(c_e^{-1})_e}$, to which we will refer to as a natural extension of a vertex group automorphism. In this case we let $\{v\}$ be the base $B(\sigma)$ of $\bar{\sigma}$. We will refer to the elements c_e above as twisting elements.

We will refer to Dehn twists and extensions of vertex group automorphisms as elementary automorphisms of G.

Suppose we are given a graded π-group $\mathsf{G}_{A,(P_i)_{i\in I}}$. By geometric abelian tree (resp. geometric abelian decomposition) of it we intend a geometric abelian tree (resp. geometric abelian decomposition) T of G relative to $\{A\} \cup \{P_i\}_{i\in I}$. By a geometric abelian tree of G_A we intend one relative to $\{A\}$. Fix a Bass-Serre presentation $\mathcal{P} = (Y, S, (t_e)_{e\in ES})$ for the action of T. By $Mod(\mathsf{G}_A, T)$ or the *modular group* of G_A with respect to T we intend the subgroup of $Aut(G)$ generated by all those elementary automorphisms σ of the form:

i) A Dehn twist over an edge of Y.

ii) A vertex group automorphism $\bar{\sigma}_{(c_e)_{\substack{\alpha(e)=v,\\ e\in S}}}$ of some $v \in VY$, such that:

  i) σ is equal to $Id_{Stab(v)}$ in case $v \in V_r$.

  ii) σ is induced by a homeomorphism of Σ_v in case $v \in V_s$

  iii) σ fixes $Per^*(v)$ for any $v \in V_a$

iii) An inner automorphism inn_c of G (conjugation by c).

and satisfying properties:

i) $\pi\restriction_{Stab(v)} \circ \sigma = \pi\restriction_{Stab(v)}$ and $\pi(c_e), \pi(c) \in Z(Q)$

ii) $\sigma\restriction_A = Id\restriction_A$

Denote by $MGen(\mathsf{G}_A, \mathcal{P})$ the set of generators listed above; we call them elementary automorphisms. It can be shown that the subgroup they generate does not depend on the presentation $\mathcal{P}$, justifying the use of the alternative notation $MGen^{\pi}(\mathsf{G}_A, T)$ or $MGen^{\pi}(\mathsf{G}_{A,(,i)_{i\in\Delta}})$ for any associated geometric abelian decomposition . We denote the subgroup generated by $MGen^{\pi}(\mathsf{G}_{A,(,i)_{i\in\Delta}})$ as $Mod(\mathsf{G}_A, \Delta)$ (or dually $Mod(\mathsf{G}_A, T)$, where. Alternatively, we might use the notation $Mod^{\pi}_A(\Delta)$.

Ignoring the π-structure yields the standard modular group $Mod_A(G, T)$. The definition extends to restricted π-groups and π-groups by regarding them as graded π-groups in the obvious way. Note that in particular, σ will fix $Per^*(v)$ for any $v \in V_a$. Given v, we will let $Mod^{\pi}(v)$ stand for the subgroup of all possible $\sigma \in Aut(Stab(v))$ satisfying the restrictions above.

If the given T is merely a pure G-tree over abelian edge stabilizers and not a geometric abelian tree then in an expression such as $Mod(G, T)$ we will be implicitly promoting T to a geometric abelian tree by declaring all its vertices as rigid (so that only Dehn twists are taken into account).

Observation 2.5.1. $PM := Mod(\mathsf{G}_A, T)$ is a subgroup of finite index in $M := Mod_A(G, T)$.

Proof. The group M contains the group $I := Inn(Z_G(A))$ of inner automorphisms of G by elements of $Z_A(G)$ as a normal subgroup. Clearly PM intersects I in a subgroup of finite index. On the other hand, given a transversal $\mathcal{R}$ of $G\backslash VT$, for any $v \in \mathcal{V}$ and $\tau \in M$ there is c such that $inn_c \circ \tau$ preserves $Stab(v)$ and it is possible to check that its projection in $Mod(G_v)/Inn(G_v)$ is well defined. This determines a map $\eta_v : M/I \to Mod(G_v)$. It is well known that the collection of all of them induces an isomorphism between M/I and $\bigoplus_{v\in\mathcal{R}} Mod(Mod(G_v)/Inn(G_v)$. Indeed, this is well-known to be true in case $A = 1$ (see [Lev05]), and it is easy to see that the inclusion map induces an isomorphism between M/I and $Mod(G,T)/Inn(G)$. Clearly the image of $Mod\mathsf{G}_A, T$ intersects each of the direct summands in a finite index subgroup. Of course, given any group G, $N \trianglelefteq G$ and $H \leqslant G$, if both $[N, N \cap H]$ and $[G/K, NH/K]$ are finite, then $[G, H]$ is finite as well. □

The following can be easily checked by inspection.

Observation 2.5.2. If e is a vertex from u to v, and $Stab(v)$ is abelian, then $\tau_{e,c} = \tau_{e_1,c} \cdots \tau_{e_k,c}$, where e_i, $i = 1, 2 \cdots k$ enumerates the edges $f \neq \bar{e}$ with $\alpha(f) = v$. Notice also that $inn_{c^{-1}}\tau_{e,c} = \tau_{\bar{e},c^{-1}}$ and for every $v \in VT$ and vertex group automorphism $\bar{\sigma}_{(c_e)_{\alpha(e)=v}}$ and $g \in Stab(v)$ one has $inn_g\bar{\sigma}_{(c_e)_{\alpha(e)=v}} = \bar{inn}_g\sigma_{(c_e g)_{\alpha(e)=v}}$. In particular, for any e_0 with $\alpha(e_0) = v$ we have $e_0 \notin Supp(inn_{c_{e_0}^{-1}} \circ \bar{\sigma}_{(c_e)_{\alpha(e)=v}})$.

Definition 2.5.2.1. The *π-modular group* of a freely indecomposable $\mathsf{G}_{A,(P_i)_{i\in I}}$, denoted by $Mod(\mathsf{G}_{A,(P_i)_{i\in I}})$ is the subgroup of $Aut(G)$ generated by $Mod(\mathsf{G}_A, \Delta)$ where Δ ranges among all geometric abelian decomposition of $\mathsf{G}_{A,(P_i)_{i\in I}}$. The modular group of G_A is simply the modular group of the associated graded π-group (with trivial grading).

Lemma 2.5.3. *Suppose we are given a geometric abelian decomposition Δ of a CSA graded π-group $\mathsf{G}_{A,(P_i)_{i\in I}}$. Assume furthermore that any maximal abelian subgroup of G is finitely generated. Then any $\sigma \in Mod(\mathsf{G}_A, \Delta)$ is contained in $\sigma \in Mod(\mathsf{G}_A, \Delta')$, for a geometric abelian decomposition Δ' of $\mathsf{G}_{A,(P_i)_{i\in I}}$ whose underlying tree is relative to the family of all abelian non-cyclic subgroups of G.*

Proof. Take any non-cyclic maximal abelian subrgroup M of G which is not elliptic. Since M is finitely generated, in virtue of 2.1.8, it acts by translations on some line L of $T = (\tilde{\Delta, Z})$.

Since M is non-cyclic, the group $K = Fix(L) \cap M$ is non-trivial. Let T' be the result of collapsing all edges in L. If σ is based on a vertex outside L (in particular, if this vertex is of surface type), then it is clear it belongs to $Mod(\mathsf{G}_A, T')$ as well. If not, observe that by 3.1.16 we can decompose G as an amalgamated product of M over the edge group K. By inspection one can easily check that when σ is a Dehn twist over an edge in L or the extension of an automorphism of the stabilizer of an abelian type subgroup of L the latter decomposition does the job. □

Let G be a finitely generated CSA group and let T be a geometric abelian tree for G. We will say that T is *normalized* if:

i) It is relative to the family of all non-cyclic abelian subgroups of G.

ii) Exactly one of the endpoints of each edge $e \in T$ is contained in an abelian type vertex, the other endpoint being non-abelian.

iii) The intersection of $Stab(e)$ and $Stab(e')$ for each rigid type vertex $v \in T$ and any two distinct edges e, e' originating at v.

Observation 2.5.4. A normalized geometric abelian tree in which the stabilizer of every rigid type vertex is CSA is 2-acylindrical.

2.6 Grushko and JSJ decompositions

In this section we give a quick overview of JSJ theory, as laid out in [GL09] and [GL10]. Or main goal is to show the following:

Proposition 2.6.1. *Any graded π-limit group $\mathsf{L}_{A,(P_i)_{i\in I}}$ admits some geometric abelian decomposition Δ_{JSJ} such that any π-modular automorphism of $\mathsf{L}_{A,(P_i)_{i\in I}}$ is contained $Mod(\mathsf{L}_A, \Delta_{JSJ})$.*

Let G be a fixed finitely generated group and $\mathcal{R}$ some family of subgroups of G. The following result was first proven by Grushko in case $\mathcal{R} = \varnothing$ (a nice topological proof due to Stallings can be found in [Sta65]).

Theorem 2.6.2. *There is a free decomposition relative to $\mathcal{R}$ of the form:*

$$G = (\ast_{i=1}^{m} G_i) \ast F$$

Where $l, m \geqslant 0$, F is a (possibly trivial) free group and each G_i is freely indecomposable relative to $\mathcal{R} \cap G_i$. The decomposition is finest in the sense that given any other decomposition of the same form $(\ast_{i=1}^{m'} G'_i) \ast F'$ then every G'_i is the free product of a free factor of a conjugate of F and some conjugates of groups G_j, while F' is a free factor of a conjugate of F.

As a matter of fact Bass-Serre presentations are pretty much irrelevant to the discussion, which will thus take place in the language of trees rather than in that of graph of groups decompositions. Given a conjugacy invariant family $\mathcal{E}$ of subgroups an $(\mathcal{E}, \mathcal{R})$-tree is a simplicial G-tree relative to $\mathcal{R}$ with edge stabilizers in $\mathcal{E}$. The notion of an $(\mathcal{E}, \mathcal{R})$-JSJ generalizes the notion of a Grushko decomposition to the class of $(\mathcal{E}, \mathcal{R})$-trees. A subgroup $H \leqslant G$ will be called $(\mathcal{E}, \mathcal{R})$-universally elliptic in case it is elliptic in any $(\mathcal{E}, \mathcal{R})$-tree. A simple example of an universally elliptic subgroup is that of some subgroup $H \leqslant G$ containing some $R \in \mathcal{R}$ as a finite index subgroup. Given G-trees S and T, one says that S dominates T if and only if $Ell(S) \subset Ell(T)$, where $Ell(T)$ stands for the familiy of all elliptic subgroups of G elliptic in T. It is easy to check this is equivalent to the existence of an equivariant map from VS to VT.

A deformation space is the class of all $(\mathcal{E}, \mathcal{R})$-trees T for which $Ell(T)$ has a particular value. A $(\mathcal{E}, \mathcal{R})$-universally elliptic tree is an $(\mathcal{E}, \mathcal{R})$-tree all of whose edge stabilizers are $(\mathcal{E}, \mathcal{R})$-universally elliptic. We say that an universally elliptic tree T is an $(\mathcal{E}, \mathcal{R})$-JSJ tree if T is $(\mathcal{E}, \mathcal{R})$-universally elliptic and $Ell(T)$ is minimal among universally elliptic trees. The following follows from [GL09, 3.2,3.8] (the second of those results is formulated there only in case $\mathcal{R} = \varnothing$ but the difference is inessential).

Lemma 2.6.3. *Let T be an universally elliptic $(\mathcal{E}, \mathcal{R})$-tree and S an $(\mathcal{E}, \mathcal{R})$-tree. Then there is a refinement $\hat{T}$ of T such that $Ell(\hat{T}) = Ell(T) \cap Ell(S)$. If S is a universal $(\mathcal{E}, \mathcal{R})$tree then $\hat{T}$ is an universal $(\mathcal{E}, \mathcal{R})$tree as well.*

As a consequence, if the class of $(\mathcal{E}, \mathcal{R})$-JSJ trees is non-empty, it is a deformation space, referred to as the $(\mathcal{E}, \mathcal{R})$-JSJ deformation space, and every $(\mathcal{E}, \mathcal{R})$-universally elliptic T tree has a $(\mathcal{E}, \mathcal{R})$-JSJ refinement. Using this it is possible to give an outline of the basic strategy used in proving existence. The following is lemma 3.9 in [GL09]

Lemma 2.6.4. *Let $\{T_i\}_{i\in I}$ be any family of G-trees, where G is finitely generated. There exists a countable subset $J \subset I$ such that if T is elliptic with respect to every T_i for $i \in I$ and dominates every T_j for $j \in J$ then T dominates T_i for all $i \in I$.*

Let $\{T_n\}_{n\in\mathbb{N}}$ be the countable subfamily extracted from the family of all $(\mathcal{E}, \mathcal{R})$-universally elliptic trees using the previous lemma. Iterative application of 2.6.3 yields a sequence of $(T'_n)_n$

of $(\mathcal{E}, \mathcal{R})$-universally elliptic trees, where T'_n is a collapse of T'_{n+1} and dominates every T_m for $m \leqslant n$. All it is then left to show is that the deformation space of any chain of refinements of $(\mathcal{E}, \mathcal{R})$-trees eventually stabilizes. This is what is commonly referred to as an accessibility result.

The following is a relative version of [GL09, Proposition 4.4], which was originally proved in [Dun85]

Lemma 2.6.5. *Suppose G is finitely presented and let $(T_n)_n$ be a sequence of $(\mathcal{E}, \mathcal{R})$-trees, where T_{n+1} is a refinement of T_n for any $n \in \mathbb{N}$. Then there is some G-tree S dominating each of the T_n.*

If one takes for granted, as we do, that limit groups are finitely presented, this is good enough for our purposes. Without this assumption one needs to apply acylindrical accesibility. See [Sel97] and [Wei02] for the basic combinatorial result and [RW] for its application to limit groups.

Given an $(\mathcal{E}, \mathcal{R})$-JSJ tree T_{JSJ} one distinguishes two types of vertices (vertex stabilizers) of T. We say that v (resp. $Stab(v)$) is *rigid* if $Stab(v)$ is $(\mathcal{E}, \mathcal{R})$-universally elliptic. Otherwise we will refer to it as *flexible* The following lemma is easy to prove using the blow-up construction.

Lemma 2.6.6. *If $Stab(v)$ is rigid then there is no non-trivial $Stab(v)$-tree S over edge stabilizers in $\mathcal{E}$ in which the stabilizer of any e incident to v in T and the intersection with $Stab(v)$ of any conjugate of a subgroup from $\mathcal{R}$ are elliptic.*

Let us now restrict to the case in which $\mathcal{E}$ contains only abelian subgroups. Take any $(\mathcal{E}, \mathcal{R})$-tree T and suppose v is of surface type, with $Stab(v) \cong \pi_1(\Sigma)$. We claim that for any minimal non-trivial $Stab(v)$-tree S relative to the boundary subgroups over edge stabilizers in $\mathcal{E}$ and any edge $e \in S$ the stabilizer of e is not universally elliptic. Indeed, (up to subdivision) we know S to be dual to some essential system of curves δ on Σ. On the other hand there is some essential simple closed curve α such that $i(\alpha, \beta) \neq 0$ for any $\beta \in \delta$. The result of blowing up v to T_α in S witnesses the fact $Stab(e)$ is not $(\mathcal{E}, \mathcal{R})$-universally elliptic.

Under certain weak conditions $Stab(v)$ can be proved to be elliptic in any $(\mathcal{E}, \mathcal{R})$-universally elliptic tree (see [GL09, Proposition 7.13]). Note that this is obvious if we assume T to be $(\mathcal{E}, \mathcal{R})$-universally elliptic to begin with, since in that case it can be refined to a $(\mathcal{E}, \mathcal{R})$-JSJ tree.

The crux of JSJ theory is the fact that under certain assumptions all flexible vertex groups of a JSJ-tree are of this form (see [GL09, theorem 7.7]. With a different formulation this was proven for example in [RS97] using a finer version of the technique used to construct test sequences presented in this work. In truth, in the general case the relevant notion is not that of a surface type vertex group but the broader one of a quadratically hanging (QH) vertex group (see definition 8.17 in [GL09]). Both are equivalent, however, in case G is a torsion free CSA (see [GL09, 8.12]).

An $(\mathcal{E}, \mathcal{R})$-tree T is called universally compatible if T and T' admit a common refinement $\hat{T}$ for any $(\mathcal{E}, \mathcal{R})$tree T'. This is more restrictive than being an $(\mathcal{E}, \mathcal{R})$-JSJ tree and more convenient (but not essential) when dealing with modular automorphisms. The existence of a universally compatible $(\mathcal{E}, \mathcal{R})$-JSJ tree is proven in [GL10, 11.1], under certain hypothesis, in particular in case of a limit group which is freely indecomposable relative to $\mathcal{R}$ and $\mathcal{E}$ the family of all abelian subgroups of G. Let us collect this result and the above discussion.

Proposition 2.6.7. *Let G be a limit group which is freely indecomposable with respect to some family $\mathcal{R}$ of subgroups. Let also $\mathcal{A}^{nz}$ be the family of all the non-cyclic abelian subgroups of G. Then G admits a (necessarily unique) universally compatible $(\mathcal{E}, \mathcal{R} \cup \mathcal{A}^{nz})$-JSJ tree T_{JSJ} with the following properties:*

- *T_{JSJ} is normalized*
- *A vertex group of T_{JSJ} is flexible if and only if it is of surface type (relative to $\mathcal{R}$).*
- *Up to edge subdivision, any minimal reduced geometric abelian tree S of G relative to $\mathcal{R} \cup \mathcal{A}^{nz}$ can be obtained from T_{JSJ} by equivariantly blowing up some of the flexible vertices of T_{JSJ} according to a essential system of curves δ and collapsing (equivariantly) some of the edges inherited from T_{JSJ}.*
- *Given a tree S as above the stabilizer of any surface type vertex of S is the fundamental group of some of the pieces in which the associated to vertices of T.*

Proof. Flexible vertices of T_{JSJ} are the only ones that can be affected by the blow-up (as we assume S is reduced). Regarding the last claim, notice that for any vertex stabilizer $Stab(v)$ of S which is not of the form specified above either:

- $Stab(v)$ equals some rigid vertex stabilizer of T_{JSJ}
- $Stab(v)$ splits relative to stabilizers of incident edges over universally elliptic vertex stabilizers.

Both are impossible for surface type vertex stabilizers. □

We can regard T_{JSJ} itself as a geometric abelian tree relative to $\mathcal{R}$ by declaring all its flexible vertices to be of surface type and all those which are abelian and as being of abelian type, with extended peripheral subgroup the smallest summand containing the peripheral subgroup and all the conjugates of $\mathcal{R}$ in $Stab(v)$. All the remaining ones are considered to be of rigid type. Proposition 2.6.1 can be proven by inspection using 2.6.7 (see [PS]) .

Remark 2.6.8. In the propostion above normality comes from the fact that T_{JSJ} can be taken to be a 'tree of cylinders' in the sense of Guirardel and Levitt.

Given a graded π-group we will refer to the JSJ of G relative to $\{A\} \cup \{P_i\}_{i \in I}$ as above as its abelian* JSJ tree .

Chapter 3

Limits of trees and the shortening argument

3.1 Limits of actions on real trees: the Bestvina Paulin method

3.1.1 Definitions

In [Pau89] Paulin adapted Gromov's notion of convergence among metric spaces to the context of isometric actions of a group G on a hyperbolic space. In our case we are mostly concerned with actions on real trees that arise as limits of sequences of actions of a group on the Cayley graph of $\mathbb{F}$ induced by homomorphisms to $\mathbb{F}$. In this section we give an overview of the techniques and state the main compactness result needed later. Most of the definitions and results come from [Bes02].

Definition 3.1.0.1. An ϵ-approximation between metric spaces (X_1, d_1) and (X_2, d_2) is a relation $R \subseteq X_1 \times X_2$ such that:

- For all $x_1, y_1 \in X_1$ and $x_2, y_2 \in X_2$ we have:

$$R(x_1, x_2) \cap R(y_1, y_2) \Rightarrow |d_1(x_1, y_1) - d_2(x_2, y_2)| < \epsilon$$

- $R(X_1) = X_2$ and $R^{-1}(X_2) = X_1$

If X_1 and X_2 are equipped with an action of a group G by isometries, for any $S \subset G$ we say that the approximation is S-equivariant if, in addition for any (x_1, x_2) and $s \in S$, if $s \cdot x_i \in X_i$ then $(s \cdot x_1, s \cdot x_2) \in R$.

Let G be a group and suppose we are given $\mathcal{X}$ a set of tuples (X, d, λ), where (X, d) is a metric space, λ an action by isometries of G on X and $\mathcal{Y}$ a set of tuples of the form $(X, d, \lambda, \cdot)$ with (X, d, λ) as before and $\bar{x} \in X^m$ a tuple of m points of X (a marking). In both cases we can assume that λ an action by isometries of G on X.

For any $S \in \mathcal{P}^{<\omega}(G)$ and $K_0 \subset S$ a finite subspace of Y_0, we define $N_{K,S,\epsilon}(X)$ as the set of all $(Y, \bar{y}, \cdot)\mathcal{Y}$ such that there exists a finite subspace $K \subseteq Y$ and a S-equivariant ϵ-approximation R between $K_0 \cup \bar{y}$ and $K \cup \bar{y}$ such that $(y_{0,j}, y_j) \in R$ for all $1 \leqslant j \leqslant m$. Let $N_{K,S,\epsilon}(X)$ be the analogous for $\mathcal{X}$, without the condition on the marking. Both sets are well defined in the sense that they only depend on the equivalence class up to equivariant isometry.

The (marked) equivariant Gromov topology on $\mathcal{X}$ ($\mathcal{Y}$) is the topology defined by taking as a neighbourhood system of $\mathcal{X}$ ($\mathcal{Y}$) the family of all $N_{K,S,\epsilon}(X)$, where S ranges over all $S\mathcal{P}^{<\omega}(G)$ and K over all finite $K \subseteq X$. When the marking consists of a single point we use the adjective *pointed* instead.

Observe that in case the metric spaces involved are $\mathbb{R}$-trees any ϵ-approximation between $K \subset X \in \mathcal{X}$ and $K' \subset X' \in \mathcal{X}$ can be extended to an ϵ-approximation between the convex envelopes $conv(K)$ and $conv(K')$, which are compact. So in the definition of the Gromov Hausdorff topology we might as well have used 'compact' instead of 'finite'. Let $\mathcal{T}(G)$ be the space of non-trivial minimal actions of G on $\mathbb{R}$-trees which are either linear or irreducible, up to equivariant isometry, equipped with the equivariant Gromov topology. Let $\mathcal{PA}(G)$ be the space of non-trivial minimal actions of G on pointed real trees which are spanned by the basepoint, equipped with the pointed equivariant Gromov topology .

Definition 3.1.0.2. A function $l : G \to \mathbb{R}_{\leqslant 0}$ is called a Lyndon length function if it satisfies the following conditions:

i) $l(1) = 0$

ii) $\forall g \in G \; l(g) = l(g^{-1})$

iii) For all $g, h, k \in G$, the inequality $c(g,h) \geqslant \min\{c(g,k), c(h,k)\}$ holds, where $c(g,h) = \frac{1}{2}(l(g) + l(h) - l(g^{-1}h))$

We say that a point $y \in Y$ *spans* Y in case Y is the convex closure of the orbit of y by the action of G.

Theorem 3.1.1 ([Chi01], 4.6 Chapter 2)**.** *Given any length function l, there is an action λ of G by isometries on a $\mathbb{R}$-tree Y and $y \in Y$ such that l is equal to the length function l^{λ}_{y}. The tuple (Y, d, λ, y) is unique up to basepoint preserving equivariant isometry.*

The proof of 3.1.1 yields also the stronger result:

Proposition 3.1.2. *The map* $\boldsymbol{l} : \mathcal{PA}(G) \to \mathbb{R}^{G}_{\geqslant 0}$ *taking any* $(X, d, \cdot, x)$ *to* $(l^{\cdot}_{x}(g))_{g \in G}$ *is a homeomorphism between* $\mathcal{PA}(G)$ *and its image.*

Assume that G is countable. Then in particular the pointed equivariant Gromov topology is separable. In fact a similar non-pointed result is true:

Proposition 3.1.3. *The map* $\boldsymbol{tl} : \mathcal{PA}(G) \to \mathbb{R}^{G}_{\geqslant 0}$ *taking any* $(X, d, \cdot)$ *to* $(tl^{\cdot}(g))_{g \in G}$ *is a homeomorphism between the subset of* $\mathcal{T}(G)$ *of all those actions which are either linear or irreducible and its image.*

Suppose we are given a sequence $((Y_n, d_n, y_n, \rho_n))_n \subset \mathcal{PA}(G)$ converging to some $(Y, d, y, \rho) \neq \mathcal{PA}(G)$ in the equivariant Gromov topology. We say that a sequence of points $(y_n)_n \in Y_n$ converges to a point $y \in Y$ if the sequence $(Y_n, d_n, \rho_n, (x_n, y_n))$ converges to $(Y, d, \rho, (x, y))$ in the marked equivariant Gromov topology. This behaves as expected:

Lemma 3.1.4 (see [RW])**.** *In the situation above, any point of y is the limit of some convergent sequence of points. If the sequence of points $(y_n)_n$ is convergent, then for any other sequence $(y'_n)_n$ of points if $(d(y_n, y'_n))_n$ converges to 0 then $(y'_n)_n$ is convergent and has the same limit as $(y_n)_n$.*

3.1.2 Compactness, rescaling and apriori basepoints

The last result above is enough to prove compactness of the projectivization of $\mathcal{T}(G)$ in the cases we are interested in. Otherwise formulated, any sequence $(\rho_n)_n$ of minimal nontrivial actions of our finitely generated G on real trees contains a subsequence which, after rescaling converges to another non-trivial action.

For the shortening argument one needs however to be able to make an apriori choice of a point $*^\rho \in Y$ for every $(Y, d, \rho) \in \mathcal{PA}(G)$ in such a way that when extracting a subsequence we can guarantee the points $*^{\rho_n}$ converge as well.

Definition 3.1.4.1. Given an action λ by isometries of a group G on an $\mathbb{R}$-tree Y and $U \subset G$ a finite set we define the following functions:

$$ml^\lambda_y(U) = \max_{u\in U} l^\lambda_y(u)$$
$$sl^\lambda_y(U) = \sum_{u\in U} l_y(U)$$
$$md^\lambda(U) = \min_{y\in Y} ml_y(U)$$

.

Lemma 3.1.5. *Let s be a finite tuple of generators of a group G and suppose that for each $n \in \mathbb{N}$ we are given a minimal non-trivial action (Y_n, d_n, ρ_n) of G on a real tree and some y_n for which $ml^{\rho_n}_{y_n}(s) = md^{\rho_n}(s) =: \mu_n$. Then some subsequence of $((Y_n, \frac{1}{\mu_n}d_n, y_n, \rho_n))_n$ converges in the equivariant pointed Gromov topology to a minimal non-trivial pointed action (Y, d, y, ρ) of G on a real tree Y.*

Proof. First of all, for any $g \in G$ the length functions $l^{1/\mu_n Y_n}_{y_n}(g)$ are uniformly bounded, since in general $l_x(gh) \leqslant l_x(g) + l_x(h)$. This implies some subsequence of $(l^{1/\mu_{n_k} Y_k}_{y_{n_k}})_{k\in\mathbb{N}}$ converges to some non-constant function l, which is necessarily the length function associated to some tuple (Y, d, y, ρ) to which $((Y_{n_k}, \frac{1}{\mu_{n_k}}d_{n_k}, y_{n_k}, \rho_{n_k}))_{n_k}$ converges in the pointed equivariant Gromov topology . We claim this action is non-trivial. Indeed, suppose that $z \in Y$ is fixed by G and let (z_n) a sequence of points in Y_n converging to z. Since $\lim_k l^{(1/\mu_{n_k} d_{n_k}, \rho_{n_k})}_{z_{n_k}}(g) = l^{d,\rho}_z(g)$ for any $g \in G$, eventually $l^{(d_{n_k}, \rho_{n_k})}_{z_{n_k}}(s) < \mu_{n_k}$, contradicting the choice of z_{n_k}. We claim that the action is minimal. Indeed, in case it is not, let Z be a proper G-invariant subtree of Y. Since y spans Y, clearly $y \notin Z$. But then, given any $g \in s$ which moves y, it is clear that there must be some $z \in Z$ such that $l_z(g) < l_y(g)$. Any sequence of approximating z will eventually contradict the choice of y_n. □

Definition 3.1.5.1. From now on, given a finitely generated G we fix some finite tuple s of generators. For any ρ as above let $\mu^\rho = md^\rho(s)$ and fix some $*^\rho_{min}$ for which $ml^\rho_{*^\rho_{min}}(s) = \mu^\rho$. If Y is the geometric realization of a simplicial tree, one can assume that $*^\rho_{min}$ corresponds to a simplicial vertex.

In certain situations, in which G is equipped with a distinguished subgroup H the point $*^\rho_{min}$ is not good enough. In order to be able to apply the shortening argument we will need to be able to make an a priori choice of points that will converge (perhaps after subsequence extraction) to a point in the minimal tree or in the fixed point set of H in the limit tree. A small complication arises from the fact that we cannot tell in advance whether H will be elliptic or not in the limiting action. In the final writing of this work it turns out the case in which H is not elliptic in the limiting tree can be ignored, but we will present the other case in case it might be useful in other circumstances.

Given non-empty convex subsets X, Z of a real tree Y and $y \in Y$, let $pr_X(y)$ be the point in X closest to Y, and $Lpr_{X,Z}(y)$ the points in $\{pr_X(y), pr_Z(y)\}$ whose distance to y is minimal. Let $Lpr_{X,Z}(y)$ be either the unique point in $Lpr_{X,Z}(y)$ or $pr_X(y)$ in case of a draw. Given $h \in G$, we let $Ax^{\rho}(h)$ stand for the set of points minimally displaced by h with respect to ρ.

Observation 3.1.6. Suppoe a group G endowed with a finite set of generators s and a sequence $((Y_n, d_n, y_n, \rho_n))_n$ of minimal pointed isometric actions of G on real trees, converging to some (Y, d, y, ρ) in the pointed equivariant Gromov topology is given. Fix some $h \in G$. For each n let $A_n = Ax^{\rho_n}(h)$, and $A = Ax^{\rho}(h)$. Then $pr_{A_n}(y_n)$ converges to $pr_A(y)$.

If we are also given $h' \in G$ and we let $A'_n = Ax^{\rho_n}(h')$ and $A' = Ax^{\rho}(h')$, then some subsequence of $Lpr_{A_n,A'_n}(y_n)$ converges to a point in $Lpr_A(y)$.

Proof. Let $z_n = pr_{A_n}(y)$ and $z = pr_A(y)$. Let $(w_n)_n$ be a sequence of points converging to w. Now, since both $l^{\rho_n}_{w_n}(h)$ and $tl^{\rho_n}(h)$ tend to $tl^{\rho}(h)$, the sequence of distances $d(w_n, A_n) = {}^{l^{\rho_n}_{w_n}(h) - tl^{\rho_n}(h)}/_2$ converges to 0, and thus we can assume $w_n \in A_n$.

Now, on one hand $d(y_n, w_n)$ converges to $d(y, z)$ by definition. On the other $d(y_n, z_n) = d(y_n, A_n) = {}^{l^{\rho_n}_{y_n}(h) - tl^{\rho_n}(h)}/_2$ must also converge to $d(y, z) = d(y, A) = {}^{l^{\rho}_{y}(h) - tl^{\rho}(h)}/_2$. But then $|d(y_n, w_n) - d(y_n, z_n)|$ tends to 0 (regardless of whether $y \in A$), so $(z_n)_n$ converges to z.

The second statement follows almost immediately from this. □

Lemma 3.1.7. *Let ρ be an isometric action of a group G on a real tree Y be given, together with $h_1, h_2 \in H \leqslant G$ and $y \in Y$. Let $A_i = Ax^{\rho}(h_i)$. Assume that if both h_1 and h_2 are elliptic then $A_1 \cap A_2$ has at most one point. Then $Lpr_{A_1,A_2}(y)$ belongs to the fix point set of H if H acts elliptically, or to its minimal tree otherwise.*

Proof. Let $w := lpr_{A_1,A_2}(y) = pr_{A_{i_0}}(y)$. Of course, each A_i contains some point a_i in Y_H. If either h_1 or h_2 is hyperbolic, then $w \in Y_H$, since then w either belongs to the axis of a hyperbolic element in H or belongs to the path between some $a_i \in A_i$ and some axis contained in Y_H.

If both h_1 and h_2 are elliptic, then w is closest to A_{3-i_0} among the points in A_{i_0} and hence must belong to $[a_1, a_2]$, since A_1 and A_2 intersect in at most one point in that case. □

Observe that if the A^1_n and A^2_n in the lemma intersect in a single point w_n, this coincides with z^i_n.

Definition 3.1.7.1. Now, suppose we are given some finitely generated limit group G and $1 \neq H \lneq G$. For each action ρ of G on some real tree X we assign a point $*^{\rho}_H$ as follows:

i) If H is trivial, take as $* = *^{\rho}_{min}$

ii) If H is abelian, pick some $h \in H \backslash \{1\}$. For any ρ as above we let $*^{\rho}_H = pr_{Ax^{\rho}(h)}(*^{\rho}_{min})$.

iii) If H is non-abelian, chose $h_1, h_2 \in H$ such that $[h_1, h_2] \neq 1$. For any ρ as above we let $*^{\rho,s}_H = lpr_{Ax^{\rho}(h_1),Ax^{\rho}(h_2)}(*^{\rho,s}_{min})$.

As before, when the tree X is a geometric realization of some simplicial tree, we take the basepoints to correspond to simplicial vertices.

3.1.3 Limit actions of increasingly acylindrical sequences of actions

Remark 3.1.8. Suppose we are given a sequence $((Y_n, d_n, \rho_n))_n$ of minimal isometric actions of a finitely generated group G on real trees and two different sequences of rescalings $((Y_n, \lambda_n d_n, \rho_n))_n$ and $((Y_n, \lambda'_n, d_n, \rho_n))_n$ $(\lambda_n, \lambda'_n \in \mathbb{R})$ converging to minimal actions (Y, d, ρ) and (Y', d', ρ') respectively in the equivariant Gromov topology . If both Y and Y' are non-trivial, then there is an equivariant homotety between (Y, d) and (Y', d').

One says that an action by isometries of a group G on a tree Y is *superstable* if and only if for any non-degenerate segments $I \subsetneq J$ either $Fix(I) = \{1\}$ or $Fix(I) = Fix(J)$. The following is a well now result (see [RS94] or [Wil09]) that can be generalized to sequences of action on hyperbolic spaces (see [Per11] or [RW]).

Lemma 3.1.9. *Let G be a torsion free finitely generated group and $((Y_n, d_n, \rho_n))_n$ a sequence of minimal actions by isometries of G on real trees, converging to some action (Y, d, ρ) in the equivariant Gromov topology. Assume that the stabilizer of any segment of length greater than $\geqslant \epsilon_n$ is contained in $K_n := \ker \rho_n$, for some sequence $(\epsilon_n)_n$ of positive constants converging to 0, and that for any $g \in G \backslash \{1\}$ we have $n \notin K_n$ for n big enough. Then:*

i) Point-wise stabilizers of non-degenerate segments of Y are free abelian and each of their non-trivial elements is hyperbolic in T_n for n big enough.

ii) Fix point stabilizers of non-degenerate tripods are trivial.

iii) The action of G on Y is super-stable.

iv) Either Y_n is linear for n big enough (so that Y is as well) or the kernel of ρ is trivial.

If we assume no subgroup of G acts on a line of Y_n dihedrally, then the same is true for the action of G on Y.

3.1.4 Limit actions induced by a sequence of homomorphisms

Suppose finitely generated groups G, and K are given, as well as an acylindrical action ρ of K on an $\mathbb{R}$-tree (T, d). Given a homomorphism $f : G \to K$, let ρ^f denote the action of G by the pullback of ρ by f on its minimal tree $Y^f = Y_{f(G)}$ in T.

Suppose that a sequence $(f_n)_n$ of homomorphisms from G to K with trivial limit kernel is given. For each $n \in \mathbb{N}$ let $Y_n := Y^{f_n}$, $\rho_n = \rho^{f_n} \restriction_{Y_n}$, $*_n = *^{\rho_n}_{min} \in Y$ the minimum displacement point and $\mu_n = \mu^{\rho_n}$ the corresponding maximal displacement by some fixed generating set, as in subsection 3.1.2. Then we know that some subsequence of rescaled pointed actions $((Y_{n_k}, \frac{1}{\mu_{n_k}} d \restriction_{Y_{n_k}}), *_n, \rho_{n_k})_k$ will converge to an action of G on a pointed real tree $(Y, d_\infty *)$.

If the sequence $(\mu_{n_k})_{n_k}$ tends to infinity, we say that $(f_{n_k})_{n_k}$ is unbounded with respect to the action of K on Y. Otherwise we say that it is *bounded*. We refer to any (pointed) tree obtained from a subsequence of $(f_n)_n$ in this way as the one as a *limit tree* for $(f_n)_n$ and the action ρ. In the first case, the sequence of acylindricity constants of the sequence of rescaled actions tends to 0, so that the hypothesis of theorem 3.1.9 hold. In particular, point-wise stabilizers of non-degenerate segments are abelian and hence, by 3.1.7, given $H \leqslant G$, the sequence consisting of each of the points $*^{\rho_{n_k}}_H$ chosen as specified in 3.1.7.1 converges to a point of Y which is either in the minimal tree of H, in case H does not fix any point of Y, or it is otherwise fixed by H. For simplicity, in case the action ρ to which we are referring is clear we will $*^{\rho_f}_H$ simply by $*^f_H$. In the second case, the limit action is a simplicial tree.

We will be mainly concerned with the particular case in which H is our reference free group $\mathbb{F}$, and ρ the action of $\mathbb{F}$ on its Cayley graph with respect to a fixed base. In this situation we might refer to Y simply as a limit tree for the sequence $(f_n)_n$.

For $h, k \in F$ we have that $l_k(h) = l_1(h^k)$. Hence in this case μ_n can be expressed as $\min_h \max_{g \in s} |inn_k \circ f_n(g)|$, for $k \in H$. Since for any positive constant N there are finitely many homomorphisms from G to H for which $\max_{g \in s} |f(g)| < N$ is bounded, any sequence $(f_n)_n \subset Mor(G_A, \mathbb{F}_A)$ contains an unbounded subsequence if and only if it contains infinitely many distinct members up to postcomposition by an inner automorphism of $\mathbb{F}$.

Definition 3.1.9.1. Let G_1 and G_2 be finitely generated groups and consider sequences $(f^i_n)_n \subset Hom(G^i, \mathbb{F})$ for $i = 1, 2$. We say that $(f^2_n)_n$ grows faster than $(f^1_n)_n$ if for any finitely generated K containing G_1, G_2 and any sequence $(h_n)_n \subset Hom(K, \mathbb{F})$ such that $h_n \restriction_{G_i} = inn_{c_n} \circ f^i_n$ for some $c_n \in \mathbb{F}$ the group G_1 fixes a point in any limit action associated to $(h_n)_n$. Alternatively, given a single sequence $(f_n)_n \subset Hom(K, \mathbb{F})$, where $G_1, G_2 \leqslant K$, we say that $(f_n)_n$ makes G_2 grow faster than G_1 if $(f_n \restriction_{G_1})_n$ grows faster than $(f_n \restriction_{G_2})_n$.

It is perhaps not a bad idea to formally state the following (well-known) lemma.

Observation 3.1.10. Let G_1 and G_2 be finitely generated groups and consider sequences $(f^i_n)_n \subset Hom(G^i, \mathbb{F})$ for $i = 1, 2$. The following are equivalent:

- $(f^2_n)_n$ grows faster than $(f^1_n)_n$
- There is some $g \in G_2$ and some tuple x of generators of G_1 such that the sequence $tl^\lambda(f_n(x_i x_j))$ tends to infinity with n for any $x_i \neq x_j$, where λ is the action of $\mathbb{F}$ on its Cayley graph.
- Some $g_2 \in G_2$ is not killed by f^2_n for n big enough and for every $g_1 \in G_1$ the quotient $\frac{tl^\lambda(f^1_n(g_1))}{tl^\lambda(f^2_n(g_2))}$ tends to 0 with n.

Proof. Suppose we are given a group G which acts by isometries on a real tree Y and a set $\mathcal{X} \subset G$, the group generated by $\mathcal{X}$ fixes a point of Y if and only if for each $x, y \in \mathcal{X}$ both elements x, xy fix a point of Y. The equivalence between 3.1.10 and 3.1.10 follows easily from continuity of translation length with respect to the equivariant Gromov topology and its invariance by conjugation. Using continuity it is also easy to show that 3.1.10 implies 3.1.10, while the direction from 3.1.10 to 3.1.10 is trivial. □

3.1.5 Trees of actions

We borrow some tools from [Gui08], the term "graph of actions" instead of "tree of actions":

Definition 3.1.10.1. A *tree of actions* can be given as a tuple $(T^G, Y_{(-)}, p_{(-)}, \cdot)$, comprising:

i) A simplicial G-tree T^G, which we call the skeleton of the tree of actions.

ii) A map assigning to each $v \in VT$ a real tree Y_v.

iii) For each $e \in ET$ a point $p_e \in Y_{\alpha(e)}$.

iv) An action of G on $\amalg_{v \in VT} Y_v$. So that for each $g \in G$ and $v \in VT$ the action of g restricts to an isometry between Y_v and $Y_{g \cdot v}$.

In particular the data makes each Y_v into a $Stab(v)$-tree.

Given a tree of actions as above, consider the quotient Y of $\amalg_{v \in VT} Y_v$ by the equivalence relation generated by $\{(p_{\alpha(e)}, p_{\alpha(\bar{(e)})})\}_{e \in ET}$. It can be shown there is a metric structure on Y making it an $\mathbb{R}$-tree with respect to which the inclusion map of Y_v into Y is an isometric embeddings for every v. We denote the resulting real tree by $\mathcal{D}((S^G, Y_{(-)}, p_{(-)}, \cdot))$.

Definition 3.1.10.2. Let Y be an $\mathbb{R}$-tree on which a group G acts by isometries. A decomposition of the action as a tree of actions comprises:

- A tree of actions $(S^G, Y_{(-)}, p_{(-)}, \cdot)$
- An equivariant isomorphism ϕ between $\mathcal{D}((S^G, Y_{(-)}, p_{(-)}, \cdot))$ and Y.

It is easy to see that S is minimal whenever Y is. We will usually identify the vertex trees Y_v with their preimage in Y and by an abuse of notation refer to the decomposition as $(T^G, Y_{(-)}, p_{(-)}, \cdot)$. We say that a family $\{Z_j\}_{j\in J}$ of subtrees of Y is fundamental if $Z_i \cap h \cdot Z_j$ is non-degenerate only in case $i = j$. We say that a subtree Z is fundamental if $\{Z\}$ is.

i) If $Y, Z \in \mathcal{Y}$ intersect in more than two points, then $Y = Z$.

ii) Every arc of X is covered by finitely many members of $\mathcal{Y}$.

Observation 3.1.11. There is a transverse covering $\mathcal{Z}$ of Y containing $\{Z_j\}_{j\in J}$, with Z_i and Z_j in different orbits for $i \neq j$ if and only if $\{Z_j\}_{j\in J}$ is fundamental.

Given a decomposition of an action of G on Y as a tree of actions $(S, Y_{(-)}, p_{(-)})$, the family of images in Y of non-degenerate vertex trees Y_v constitutes a transverse covering of Y. Viceversa: given a transverse covering $\mathcal{Y}$ of a Y on which a group G acts by isometries, let $V = \mathcal{Y} \amalg \mathcal{I}$, where $\mathcal{I}$ is the family of all points of Y belonging to more than one member of Y; notice how G acts on both sets. Put an edge (p, Z) between $p \in \mathcal{I}$ and $Z \in \mathcal{Y}$ whenever $p \in Z$. Clearly G acts without inversion in the resulting graph S. For each $Z \in \mathcal{Y}$ there is an associated action of the set-wise stabilizer of $\mathcal{Y}$ on Y. For each p we can consider the trivial action of $Stab(p)$ on $\{p\}$. It is easy to see that the dual of the corresponding tree of actions is isomorphic to Y. We call decompositions as trees of actions of this form normalize.

Observation 3.1.12. Let $(S, Y_{(-)}, p_{(-)})$ a normalized decomposition as a tree of actions of the action of G on Y, as above. Then any $H \leqslant G$ fixing a point of Y fixes a vertex of S.

Observation 3.1.13. Suppose we are given a decomposition as a tree of actions $(S, Y_{(-)}, p_{(-)})$ of an action of a group G on a $\mathbb{R}$-tree Y. Let T a subtree of S which is invariant under the action of some $H \leqslant G$, and such that for any $v \in VT$ such that Y_v is not degenerate condition $G \cdot v \cap T = H \cdot v$ holds and $Stab(v) \leqslant H$. Then $Y_T := \bigcup_{v\in VT} Y_v \subset Y$ is fundamental and $Stab(T) = H$, ie., $g \cdot Y_T \cap Y_T$ is non-degenerate for $g \in G$ only if $g \in H$.

Proof. Given $g \in G$ such that $g \cdot Y_T \cap Y_T$ is non-degenerate, our assumption on the action of H on S implies $g \cdot Y_v \cap h \cdot Y_v$ is non-degenerate for some $v \in VS$. Since Y_v is itself fundamental, $g^{-1}h \in Stab(v) \subset Stab(Y_v) \subset H$. But $h \in H$, so this proves the result. □

We note in passing the following:

Observation 3.1.14. Suppose that we are given a normalized decomposition as a tree of actions of an isometric action of a group G on a real tree Y, $(S, Y_{(-)}, p_{(-)})$, where the action on S is non-trivial. If two distinct degenerate components Y_u, Y_v intersect in a point p with $Fix_{Stab(v)}(p) = \{1\}$ then S contains trivially stabilized edges and, therefore G is freely decomposable relative to $\{Stab(v) \mid v \in VS\}$.

Corollary 3.1.15. *Given G, Y and $(S, Y_{(-)}, p_{(-)})$ as above, suppose that G is freely indecomposable relative to $\{Stab(v) \mid v \in VS\}$. Suppose that some non-degenerate component Y_v, with stabilizer G_v and $p \in Y_v$ such that $Stab_{G_v}(p) = \{1\}$. Then $Stab_G(p) = \{1\}$ and $G \cdot p \cap Y_v = G_v \cdot p$.*

One property of decompositions of the type above worth mentioning is the following:

Lemma 3.1.16. *Let G be a CSA group equipped with an action by isometries on a real tree Y in which segment stabilizers are abelian. Let Z be a line of Y and suppose that $M \leqslant G$ consist of those elements preserving Z set-wise, with its orientation. Let also $K \leqslant M$ be the kernel of the action of M on Z. If $M \neq K$, then M is a maximal abelian subgroup of G. If moreover G is torsion-free, then M coincides with the set-wise stabilizer of Z. If $K \neq \{1\}$ as well, then Z is fundamental and G decomposes as an amalgamated product of the form $G' *_K M$, where every subgroup of G elliptic in Y can be conjugated into G'.*

Proof. First of all we claim that M is abelian. Indeed, first of all clearly M/K is free abelian. On the other hand K is abelian by our assumption on the action on Y. If $K \neq \{1\}$, since also $K \trianglelefteq M$, the CSA property implies that M is abelian.

Now, M is maximal as well. Indeed, pick any $g \in G$ commuting with M, let $Z' = a \cdot Z$ and suppose that $Z' \neq Z$. Any $m \in M \backslash K$ has to act as a proper translation on both Z and Z', which is impossible in any action by isometries on a real tree. So $a \cdot Z = Z$. It is clear as well that g has to preserve the orientation of Z, since otherwise $m^g = m^{-1}$ so $g \in M$. Finally, assume that G is torsion free and $s \in G$ stabilize Z set-wise. Then $s^2 \in M$. Since $s^2 \neq 1$, commutative transitivity implies that $s \in M$.

We now add the assumption that $K \neq \{1\}$. We claim that Z is fundamental in Y. Indeed, pick any $k \in K \backslash \{1\}$. Given any $g \in G$ for which $J := g^{-1} \cdot Z \cap Z$, k^g fixes J which implies that $k^g \in M$, since segment stabilizers are abelian. The CSA property implies that $g \in M$ as well, so that $g \cdot Z = Z$. This implies that there is a decomposition of the action on Y as a normalized tree of actions $(S, Y_{(-)}, p_{(-)})$, where $Z = Y_v$ for some $v \in VS$. And observe that the edge group associated to any edge e incident on v is precisely K. This implies that G admits the required decomposition. Any subgroup of G elliptic in Y is elliptic in Y. If it is contained in a conjugate of M, then it must be contained in a conjugate of K. This proves the last claim. ☐

3.1.6 Rips' decomposition

The combinatorial core of the following result goes back to the work of Makanin and Razborov on solutions to equations in the free group. Rips' contribution was the realization that their ideas could be read geometrically and hold in a much wider context. The initial goal was to study free actions on $\mathbb{R}$ trees but it turned out it could be applied to a larger class of actions (see [BF95]). We now state the version found as theorem 4.1 of [Gui08], with some extra detail borrowed from its proof. See theorem 5.1 in that same paper for the most general result of this kind presented there.

Theorem 3.1.17. *Any minimal superstable action by isometries of a group G on an $\mathbb{R}$-tree Y, where G is finitely generated relative to the family of its elliptic subgroups admits a decomposition as a tree of actions $(S, Y_{(-)}, p_{(-)})$ where either:*

a) There is $v \in V$ such that Y_v contains an infinite tripod and some $e \in S$ incident to v such that $Stab(e)$ coincides with the kernel K of the vertex action on Y_v. In particular there is a splitting of G over K relative to the stabilizer of any point in Y.

b) Each vertex action is either:

i) Simplicial, i.e., isomorphic to the geometric realization of a simplicial tree.

ii) Of Seifert type. If we denote by K the kernel of the action of $Stab(Z)$ on Y_v, then $Stab(Z)/K$ is dual to the fundamental group of an orbifold with boundary and the action on Y_v dual to an arational measured foliation on it.

iii) Of axial type: Z is a line and the image of $Stab(Z)$ in $Isom(Z)$ is a non-discrete group of isometries of Z.

One can assume that the decomposition as a tree of actions above is normalized. ALso, after further decomposing any simplicial action, that in the second alternative of the theorem all simplicial vertex trees are either a point or a non-trivially stabilized segment with its endpoints are the only attaching points p_e. The first alternative in the theorem can be understood in terms of the existence of exotic components. See for example [BF95]. If G is torsion free, then in case (bii) the orbifold is actually a surface and the action of Seifert type as defined earlier.

The result above can be slightly refined:

Corollary 3.1.18. *Consider an isometric action of a finitely generated group G on a $\mathbb{R}$-tree Y, which is superstable and has trivial tripod stabilizers and let $\{Z_j\}_{j\in J}$ be a fundamental family of subtrees of Y. Then either:*

i) G is freely decomposable relative to $\{Stab(Z_1), Stab(Z_2), \cdots Stab(Z_m)\} \cup \mathcal{E}$ where $\mathcal{E}$ is the family of all the subgroups of G elliptic in Y.

ii) The action decomposes as a tree of actions $(S, Y_{(-)}, p_{(-)})$, where $Y_{v_j} = Z_j$ for some $v_j \in VS$ and for any $w \in VS\backslash(G \cdot \{v_j \mid j \in J\})$ the vertex action on Y_v is either simplicial with non-trivially stabilized edges, of axial type or of Seifert type. Given a non-degenerate component Y_v, all its points of intersection with other non-degenerate components have non-trivial stabilizer in $Stab(v)$.

Proof. Let $(S, Y_{(-)}, p_{(-)})$ a graph of actions, where each $Z_j = Y_{v_j}$ for some v_j in a different orbit. For all non-trivial $v \in VS\backslash G \cdot \{Z_j\}_{j\in J}$ we can apply 3.1.17 to Y_v, which yields a decomposition of the action of $Stab(v)$ as a normalized tree of actions $(S^v, Y^v_{(-)}, p_{(-)})$, with associated transverse covering $\mathcal{Y}_v$. Denote by Y'_v the minimal tree of Y_v. If one of Y'_v falls into first alternative of theorem 3.1.17, then $Stab(v)$ splits over a tripod stabilizer relative to its elliptic subgroups; in particular, relative to the stabilizer of edges of S incident to v. Hence this free splitting extends to a free splitting of G relative to $\{H_j\}_{j=1}^m$. If the second alternative, observe that all the components of $Y_v\backslash Y'_v$ are segments. Indeed, a segment with and endpoint in Y'_v and another in $\partial(Y\backslash Y_v)$ cannot be trivially stabilized, since otherwise S would contain trivially stabilized edges. Superstability tells us that its point-wise stabilizer coincides with that any of its non-degenerate initial subsegments. The fact non-degenerate tripods have trivial stabilizer then implies any two such segments can only intersect in their initial points.

We conclude that Y_v itself admits a transverse covering $\mathcal{S}_v$ as in the previous lemma. The covering $\{Y_{g\cdot v_j}\}_{\substack{g\in G\\ j\in J}} \cup \bigcup_{\mathcal{Y}_v} v \in VS\backslash(G \cdot \{Z_j\}_{j\in J})$ of Y satisfies the required conditions. The final claim follows from observation 3.1.14. □

3.2 The shortening argument

The following theorem contains the geometric engine of the shortening argument. Except for the constraint on the modular group, the slightly looser hypothesis and the additional clause regarding the limit of the shortening argument. This is usually stated in the more general context of actions on hyperbolic spaces (see [RS94] or [Per08]). For the sake of completeness we provide an almost self contained exposition of the more restrictive version used in this work. The original result could have been enough for most of the applications in this thesis, had we used a different approach but in view of future work, we have chosen to state the result involving a finite index subgroup of the modular group.

Proposition 3.2.1. *Let G be a group and $(Y_n, d_n, y_n, \lambda_n)$ a sequence of minimal pointed actions by isometries of L on real trees, converging to some faithful pointed action (Y, d, y, λ) in the pointed equivariant Gromov topology , such that any segment stabilizer is abelian. Let U be a finite tuple of elements from G. Suppose that λ admits a decomposition as a graph of actions $(S, Y_{(-)}, p_{(-)})$ and $\mathcal{O} \subset VS$ an equivariant set such that for each $v \in \mathcal{O}$ the vertex action of $Stab(v)$ on Y_v is either simplicial with non-trivially stabilized edges or of Seifert or axial type.*

Let S' be the tree obtained by collapsing all those edges of S none of whose ends belong to $\mathcal{O}$ and blowing up into Y_v every $v \in VS$ for which Y_v is simplicial.

Regard S as a geometric abelian tree by declaring of surface type all vertices associated to actions of Seifert type, of abelian type those associated with mixing components and regarding all the rest as rigid. Suppose that a normal finite index subgroup $\mathscr{P} \trianglelefteq Mod(\Delta)$ is given.

Then there is $\tau \in \mathscr{P}$ such that for infinitely many $n \in \mathbb{N}$:

- *i) $\tau(g) = g$ for any $g \in G$ such that $[y, s \cdot y]$ does not intersect Y_v for $v \in \mathcal{O}$ in a non-degenerate segment.*
- *ii) $l_{y_n}^{\lambda_n \circ \tau}(g) < l_{y_n}^{\lambda_n}(g)$ for any $g \in U$ for which $[y, g \cdot y]$ intersects Y_v in a non-degenerate segment for some $v \in \mathcal{O}$.*

Proof. We can assume that the decomposition is normalized, so that $VS = V_0 \amalg V_1$, with Y_v a single point in case $v \in V_0$ and a non-degenerate tree in case $v \in V_1$, as well that any attaching points on an action of simplicial type correspond to simplicial vertices. For convenience, we will assume that $x \in Y_{v_0}$ for some $v_0 \in S \backslash \mathcal{O}$ (the other case does not represent any fundamental complication). We will start with some preliminary considerations, valid for any decomposition of Y as a graph of actions. Fix some presentation $P = (W_0, W, (t_e)_{e \in EW})$ of the action of G on S for which $v_0 \in Z$. Let g be any element of g; it has a loop normal form $g_0 t_{e_0} \cdots t_{e_{m-1}} g_m$ with respect to this presentation, as in 2.2, where $g_i \in Stab(v_i)$ and $v_0, e_0, v_1 \cdots e_{m-1}, v_m \subset W_0 \cup W$ projects to a closed loop in $G \backslash S$. For each $0 \leqslant j \leqslant m-1$ we let f_j be the unique edge in W whose inverse is in the orbit of e_j. Suppose we are given $v \in V_1 \cap W_0$ and a natural extension $\bar{\sigma} = \bar{\sigma}_{(c_e)_{\alpha(e)=v}}$ of some $\sigma \in Aut(Stab(v))$ to G. with $c_{v_0} = 1$ (recall that $\bar{\sigma}$ restricts to conjugation by c_u on $Stab(v)$ for $u \in W_0 \backslash \{v\}$. Then the element $\bar{\sigma}(g)$ has a normal form $g_0^{\bar{\sigma}} t_{e_0} \cdots t_{e_{m-1}} g_m^{\bar{\sigma}}$, where

- $g_i^{\bar{\sigma}} = g_i$ in case $v_i \neq v$
- $g_i^{\bar{\sigma}} = c_{e_{i-1}} \sigma(g_i) c_{e_i}^{-1}$ in case $v_i = v$

Recall that a vertex $v \in \mathcal{O} \cap W_0$ appears in the loop representation of g if and only if $[g \cdot v_0, v_0]$ intersects the orbit of w. Since we assume our decomposition $(S, Y_{(-)}, p_{(-)})$ to be normalized, this happens precisely when $[y, g \cdot y]$ intersects some translate of Y_{v_0} in a non-degenerate segment. Let $N(v)$ be the subgroup of all those $g \in G$ for which that is not the case. Clearly $\bar{\sigma}$ fixes $N(v)$. Likewise, observe that the axis of g intersects the orbit of v in S if and only if its axis in Y intersects some translate of Y_v in a non-degenerate segment. The automorphism $\bar{\sigma}$ restricts to an inner automorphism of G on any subgroup $H \leqslant G$ which is either elliptic in S or whose minimal tree does not contain the vertex v. By the previous assertion, the latter is equivalent to the fact that none of the axis of hyperbolic elements in H intersects Y_{v_0} in a non-degenerate segment.

We define:

$$l_y^v(g) = |[y, g \cdot y] \cap \bigcup_{u \in [v]} Y_u| = \sum_{\substack{0 \leqslant j \leqslant m \\ v_j = v}} |J_j|$$

More precisely, the segment $[y, g \cdot y]$ can be decomposed into mutually non-overlapping subsegments of the form $I_j = h_j \cdot J_j$ for $0 \leqslant j \leqslant m$, where $J_j \subset Y_{v_j}$ and h_j are given by:

- $h_0, \ J_0 = [y, p_{e_0}]$
- $h_m = g, \ J_m = [p_{f_{m-1}}, x]$
- $h_j = g_0 t_{e_0} g_1 \cdots t_{e_{j-1}}$ and $J_j = [p_{f_{j-1}}, g_j \cdot p_{e_j}]$ for any $1 \leqslant j \leqslant m-1$.

Let $\bar{\sigma} = \bar{\sigma}_{(c_e)_{\alpha(e)=v}}$ be a natural extension of some $\sigma \in Aut(Stab(v))$ to G with $c_{v_0} = 1$, as before. Let $h_j^{\bar{\sigma}}$ be defined in the same way as h_j, but with $g_j^{\bar{\sigma}}$ in place of g_j. Exactly as before,

$[y, \bar{\sigma}(g) \cdot y]$ is the union of the non-overlapping segments $h_j^{\bar{\sigma}} \cdot J_j^{\bar{\sigma}}$, where $J_j^{\bar{\sigma}} = [p_{f_{j-1}}, g_j^{\bar{\sigma}} \cdot p_{e_j}]$ for $1 \leqslant j \leqslant m-1$. In particular, for $\bar{\sigma}$ as above $l_y^w(\bar{\sigma}(g)) = l_y^w(\bar{\sigma}(g))$ for $v \neq w \in W_0$.

The proof of our proposition reduces to the following two lemmas:

Lemma 3.2.2 (Mixing case). *Let $v \in W_0$ be associated to an action of Seifert or axial type, U a finite subset of G and $\epsilon > 0$. Then there is a natural extension $\bar{\sigma} = \bar{\sigma}_{(c_e)} \in \mathscr{P}$ of an automorphism σ of $Stab(v)$ such that for all $g \in U \backslash N(v)$ we have:*

$$l_y^{[v]}(\bar{\sigma}(g)) < \epsilon$$

As we know, we can assume any simplicial component is an edge containing no attaching points in its interior. Given any such I, one consider the simplicial tree T_I which has a vertex v_C in correspondence to each connected component of $Y \backslash \bigcup_{g \in G} g \cdot I$ and a simplicial edge between v_{C_1} and v_{C_2} whenever C_1 and C_2 intersect some common translate of I.

Lemma 3.2.3 (Simplicial case). *Let $v \in W_0$ and suppose Y_v is a non-trivially stabilized edge $[p, q]$ and e the corresponding simplicial edge in $T^{[p,q]}$. Then there is a Dehn twist $\tau_{e,c} \in \mathscr{P}$ such that for some diverging sequence $(k_n)_n \subset \mathbb{N}$, and all $g \in G \cap L(v)$ we have:*

$$l_{y_{n_k}}(\tau(g)) < l_{y_{n_k}}(g)$$

for k big enough (the particular bound dependent on g)

Let us see how to prove the proposition using the lemmas above. Denote by M be the set of all those $v \in W_0$ associated to an axion of Seifert or axial type. For each $v \in M$ let $\bar{\sigma}^v$ be the automorphism obtained by applying the first lemma to $\epsilon_v = \min_{g \in U} l_y^v(g)$. We know that $\bar{\sigma}^v(g) = g$ for $g \in N(v)$, while in case of $g \in U \backslash N(v)$ we have

$$\begin{aligned} l_y(\bar{\sigma}^v(g)) &= l_y^{[v]}(\bar{\sigma}^v(g)) + l_y^{[v]^c}(\bar{\sigma}^v(g)) = \\ &= l_y^{[v]}(\bar{\sigma}^v(g)) + l_y^{[v]^c}(g) < l_y^{[v]}(g) + l_y^{[v]^c}(g) = l_y(g) \end{aligned}$$

If we let θ be the product of all the $\bar{\sigma}^v$ as we let v range among all vertices in W_0, (the order is irrelevant) then for any $v \in U \backslash \bigcup_{v \in M} N(v)$ we have $l_y(g) < l_y(\theta(g))$. Which implies that $l_{y_n}^{\lambda_n}(\tau(g)) < l_{y_n}^{\lambda_n}(\tau(g))$ for n large enough, since of course $(Y_n, d_n, y_n, \lambda_n)$ has to converge to $\lambda \circ \sigma$ in the pointed equivariant Gromov topology . The proof can be completed by applying the second lemma to each of the orbits of simplicial edges in Y, starting with the sequence of actions $\lambda_n \circ \sigma$. What remains is dedicated to the proof of the aforementioned lemmas. For simplicity, denote $N_v \cap U$ by U_v.

Mixing cases

Let now $v \in W_0$ be a vertex associated with an action of Seifert or axial type. In view of the discussion above, it will be enough for us to find, given $\epsilon > 0$ and any finite $H \subset Stab(v)$, some extension $\bar{\sigma}_{(c_e)} \in \mathscr{P}$ such that for all $h \in H$:

$$d(c_e^{-1} \cdot p_e, \sigma(h) c_f^{-1} \cdot p_f) < \epsilon$$

In general, given an action by isometries of a group H on a $\mathbb{R}$-tree, and $A \subset Aut(H)$, let us say that A has the shortening property if for any $V \subset H$, any $z \in G$ and any $\epsilon > 0$ some $\sigma \in \mathcal{A}$ exists for which $l_y(\sigma(g)) < 0$ for all $g \in H$. Ww claim that if H has the shortening property and $K \leqslant H$ has finite index in H then K has the shortening property as well. Let $\mathcal{R}$ be a finite set of representative of H/K. Given ϵ, V and z as above, chose

$\theta \in Mod(\Sigma)$ such that $l_y(\theta(g)) < \epsilon$ for all $h \in \bigcup_{\tau \in \mathcal{R}} \tau^{-1}(v)$. Then for some $\tau \in \mathcal{R}$ the composition $\sigma = \theta \circ \tau^{-1} \in B$ witnesses shortening for (ϵ, V). We set $z = p_{e*}$. Denote by $Mod(Stab(v))$ the group of automorphisms of $Stab(v)$ which coincide with an inner automorphism on each of the stabilizers of an edge incident to e.

Seifert case Here $Stab(v) \cong \pi_1(\Sigma)$, where Σ is a compact surface with boundary. For a proof of the following result see [Per08, Lemma 5.13] or [RS94, p.347-348]

Proposition 3.2.4. *In the situation above, given $v \in W_0$ associated to an action of surface type $Mod(Stab(v)) \leqslant Aut(Stab(v))$ has the shortening property.*

Let $M_{\mathscr{P}}$ be the subgroup consisting of all those elements of $Mod(\Sigma)$ extending to $\mathscr{P}$; this clearly has finite index in $Mod(\Sigma)$, hence has the shortening property as well. Let $\mathcal{E} = \{e \in W \,|\, e \neq E^*, \alpha(e) = v\}$ and for each $e \in \mathcal{E}$ let s_e generate $Stab(e)$.

Let $V = \{s_e\}_{\substack{\alpha(e)=v \\ e \in W}} \cup U_v$. Pick $\sigma \in M_H$ such that $l_x(h) < \frac{\epsilon}{6}$ for all $h \in V$. If it restricts to conjugation by some c_e on each $\langle s_e \rangle$, then $\sigma(\langle s_e \rangle)$ stabilizes $c_e^{-1} p_e$ and no other point. Hence $d(z, c_e^{-1} \cdot p_e) = l_z(\sigma(s_e)) < \frac{\epsilon}{5}$. Given any other $f \in \mathcal{E}$ we get:

$$d(hc_f^{-1} \cdot p_f, z) \leqslant d(hc_f^{-1} \cdot p_f, c_f^{-1} \cdot p_f) + d(c_f^{-1} \cdot p_f, z) \leqslant$$
$$\leqslant (l_z(h) + 2d(c_f^{-1} p_f, z)) + l_z(\sigma(s_f)) \leqslant \frac{4\epsilon}{6}$$

It follows that $d(c_e^{-1} \cdot p_e, hc_f^{-1} \cdot p_f) \leqslant \frac{5\epsilon}{6} < \epsilon$ and we are done.

Axial case

Fact 3.2.5. *The action of $Stab(v)$ on Y_v satisfies the following properties:*

i) Orbits are dense.

ii) $Mod(Stab(v))$ has the shortening property.

First of all, we claim it was possible for us to chose our transversal W, containing v_0, for the action of G on S in such a way that $d(p_e, z) < \frac{\epsilon}{3}$ for any $e \in W$ originating at v. Indeed, if W already satisfies the property for all v' in the path from v_0 to v associated to an axial type action, let $W^1, W^2, \cdots W^m$ consist of the connected components of $W \backslash \{v\}$ which do not contain v_0. Using (i), we can guarantee that the property holds for v, by replacing each of the W^i by its translate by an appropriate element of $Stab(v)$, an operation which preserves the condition on those vertex on which it has already been arranged for. Let $M_{\mathscr{P}}$ be the subgroup of $\sigma \in Aut_{\partial}(Stab(v))$ such $\bar{\sigma}_{\bar{1}} \in \mathscr{P}$. Being of finite index in $Stab(v)$, $M_{\mathscr{P}}$ has the shortening property, so $\sigma \in M_{\mathscr{P}}$ exists such that $l_z(\sigma(h)) < \frac{\epsilon}{3}$ for any $h \in U_v$. Now, for any $e, f \in \mathcal{E}$ we have $c_e = c_f = 1$ while

$$d(p_e, \sigma(h)p_f) \leqslant d(p_e, z) + d(p_f, z) + tl(\sigma(h)) < \epsilon$$

Simplicial case

Let $[p, q]$ be our simplicial edge and $(p_n)_n$ and $(q_n)_n$ approximating sequences for p and q respectively. We can assume that p is closest to y. Set $l_n = d(p_n, q_n)$ and $l = d(p, q)$, $t_n = tl^{\lambda_n}(c)$ and $\zeta_n = \max\{d(p_n, h \cdot p_n), d(q_n, h \cdot q_n)\}$. By convergence both $(t_n)_n$ and $tl^{\lambda_n}(h) \neq 0$ tend to 0 with n. Choose $c \in Stab(v) = Fix([p, q]) \backslash \{1\}$ hyperbolical for the action λ_n. The acylindricity condition implies that for n big enough c has to act hyperbolically for λ_n.

Indeed there are points p'_n and q'_n in $Fix_{\lambda_n(c)}$ at distance strictly less than $\frac{\zeta_n}{2}$ from p_n and q_n respectively, so that the diameter of $Fix(c)$ is greater or equal than $l_n - \zeta_n$. Since this value tends to $d(p,q) \neq 0$ for n big enough this cannot be the case.

Up to restricting to a subsequence, we can assume that c always translates $Ax(c)$ in the direction from p'_n to q'_n and up to replacing h with some power we can assume that $\tau = \tau_{e,c} \in \mathscr{P}$.

Fix now $g \in G$ and suppose that $[y, g \cdot y]$ intersects some translate of $[p,q]$ non-trivially. Let $g_0 t_{e_0} \cdots t_{e_{m-1}} g_m$ be a loop representation of g with respect to some presentation of the tree $T_{[p,q]}$ described above, associated to some sequence $u_0, e_0 \cdots u_m \subset T_{[p,q]}$. Notice that in this case e_j is either the edge $e \in T_{[p,q]}$ associated to $[p,q]$ or $t_e^{-1} \cdot \bar{e}$ for all i (possibly $t_e = 1$). For each $0 \leqslant j \leqslant m-1$ let $k_j = g_0 t_{e_0} g_{e_1} \cdots t_{e_{i-1}} g_j$. For any $0 \leqslant j \leqslant m$ let $p^j = p$, $q^j = q$ if $e^j = e$ and $p^j = q$, $q^j = p$ otherwise.

Observe that we encounter the sequence of points $p^0, q_0, k_1 \cdot p^1, k_1 \cdot q^1 \cdots k_m \cdot q^m$ in that same order along the path between y and $g \cdot y$. By convergence in the equivariant Gromov topology,

$$lim_n\, d(p_n, q_n) + \sum_{m-1=0}^{L} d(t_{e_j^{-1} \cdot q_n^j}, g_{j+1} p_n^{j+1}) + d(y_n, p_n^0) + d(g_m \cdot y_n, t_{e_m}^{-1} q_n^{0,m}) = l_y(g)$$

Now, let $k_0 = 1$ and for $0 \leqslant j \leqslant m-1$ let $k_{j+1} = k$. $k_m = \tau(g)$. This implies the existence of sequences of points $(a_n)_n, (b_n)_n \subset [p'_n, q'_n]$ with $(a_n)_n$ converging to p and $(b_n)_n$ converging to q respectively such that if we let : $a_n^j = a_n$ and $b_n^j = b_n$ in case $e_j = e$ and $a_n^j = b_n$ and $b_n^j = a_n$ otherwise, then the sequence of points $w_n^0, z_n^0, w_n^1 \cdots z_n^m$ defined by $w_n^j = k_j a_n^j$ and $z_n^j = k_j b_n^j$ appears in that order along the segment $[y_n, g \cdot y_n]$ for n big enough.

Let the pair $(\hat{w}_n^j, \hat{z}_n^j)$ be equal to $(k')_n^j \cdot (a_n, c^{-1} \cdot b_n)$ if $e_j = e$ and $(b_n, c \cdot a_n)$ otherwise. Since clearly $d(\hat{w}_n^j, \hat{z}_n^j) = d(w_n^j, z_n^j) - t_n$, in order to show that $l_{y_n}(\tau(g)) = l_{y_n}(\tau(g)) - t_n m$ for n big enough it is enough to prove:

- The segments $[z_n^j, w_n^j]$ and $[\hat{z}_n^j, \hat{w}_n^j]$ are translates of each other for $1 \leqslant j \leqslant m-1$
- $\hat{w}_n^0 \in [y_n, \hat{z}_n^0]$
- $\hat{z}_n^m \in [\hat{w}_n^m \tau(g) \cdot y_n]$
- $\hat{z}_n^j \in [\hat{w}_n^j, \hat{w}_n^{j+1}]$, $\hat{w}_n^{j+1} \in [\hat{z}_n^j, \hat{z}_n^{j+1}]$ for each $1 \leqslant j \leqslant m-1$

First of all, for each $1 \leqslant j \leqslant m$ we have $\hat{z}_n^j = k'_j c_{e_j}^{-1} \cdot b_n^j$, while $\hat{w}_n^{j+1} = k'_j \tau(t_{e_j}) g_{j+1} a_n^{j+1} = k'_j c_{e_j}^{-1} t_{e_j} g_{j+1} a_n^{j+1}$. It is clear from this that $(\hat{z}_n^j, \hat{w}_n^{j+1}) = \mu_j \cdot (z_n^j w_n^{j+1})$, where $\mu_j = (h'_j)^{-1} c_{e_j}^{-1}$.

Likewise, the triple $(\hat{w}_n^j, \hat{z}_n^j, \hat{w}_n^{j+1})$ is a translate of $(a_n^j, c_{e_j}^{-1} \cdot b_n^j, g_{j+1} \cdot a_n^{j+1})$. Now, $b_n^j \in [a_n^j, g_{j+1} \cdot a_n^{j+1}]$, since $z_n^j \in [w_n^j, a_n^{j+1}]$. On the other hand $[a_n^j, b_n^j] \subset Ax(c)$, so as soon as the translation length of c with respect to λ_n is smaller than $d(a_n^j, b_n^j)$, we have $\hat{z}_n^j \in [\hat{w}_n^j, \hat{w}_n^{j+1}]$. The argument showing $\hat{w}_n^{j+1} \in [\hat{z}_n^{j+1}, \hat{z}_n^j]$ is entirely symmetrical. □

3.2.1 Proper quotients through shortening

It is time to draw some consequences of a more algebraic nature from the previous theorem.

Let $\mathsf{L}_{A,(P_i)_{i \in I}}(x)$ be a marked non-cyclic graded π-limit group, and Δ an graded geometric abelian decomposition of $\mathsf{G}_{A,(P_i)_{i \in I}}$. If $A \neq \{1\}$, fix some $h \in A \backslash \{1\}$. Let H_A be π-group, endowed with an isometric action without inversion λ on a K-acylindrical $\mathbb{R}$-tree (Y,d) and f a morphism from L_A to H_A.

We say that f is Δ-*short* with respect to x-marking if and only if $sl^{\lambda \circ \tau}_{*_A^{\lambda_f}}(x) \geqslant sl^{T}_{*_A^{\lambda_f}}(x)$, for any $\tau \in Mod(\mathsf{G}_A, \Delta)$, where $*_A^{\lambda_f}$ is the point as defined in 3.1.2.

Given a property $\mathcal{P}$ (and an L-invariant family $\mathcal{A}$ of subgroups of L) we say that f is Δ-short (relative to $\mathcal{A}$) with respect to the x-marking among those satisfying property $\mathcal{P}$ if it satisfies $\mathcal{P}$ and $sl^{\lambda\circ\tau}_{*^{\lambda_f}_A}(x) \geqslant sl^{T}_{*}(x)$ for any $\tau \in Mod(\mathsf{L}_{A,(P_i)_{i\in I}}, \Delta)$ for which $f \circ \tau$ also satisfies $\mathcal{P}$.

Corollary 3.2.6. *Let $\mathsf{G}_{A,(P_i)_{i\in I}}(x)$ be a marked graded π-limit group and H_A a restricted limit group which is equipped with an acylindrical isometric action λ on a real tree. Suppose we are given a system of inequalities $\Psi(x,a) \neq 1$ and $(f_n)_n$ an unbounded sequence of morphisms from $\mathsf{G}_A(x)$ to H_A which are short among those preserving $\Psi(x,a) \neq 1$. Let Y a limiting tree for the sequence $(f_n)_n$ and λ, as described in 3.1.9. If each $\mathsf{G}_{A,(P_i)_{i\in I}}$ is freely indecomposable and each of the P_i is elliptic in Y, then $lker_n\ f \neq \{1\}$. In particular the latter is the case whenever the domain is a freely indecomposable restricted group (trivial grading case).*

Proof. Suppose for the sake of contradiction that $lker_n\ f_n = \{1\}$. Up to replacing f_n with a subsequence we can assume that the sequence of rescalings of the actions induced by the f_n converges to a limiting action ρ on a real tree Y. In virtue of 3.1.9, the action of G on Y is superstable with trivial tripod stabilizers. Notice that since the sequence is unbounded, the subgroup A is elliptic in Y. In particular it fixes the limit $*$ of the sequence $(*_n)_{n\in\mathbb{N}}$ given by $*_n = *^{\lambda_{f_n}}_A$.

Notice $[h\cdot *, *] = \{*\}$ for any $h \in A$. At the same time, $[*, y_j \cdot *]$ is necessarily non-degenerate for some $y_j \in x$. Taking $*$ as a basepoint in 3.2.1, we conclude that an automorphism $\sigma \in Mod(\mathsf{G}_A, S)$ exists such that $sl_{*_n}(f_n \circ \sigma(x)) < sl_{*_n}(f_n(x))$ for n big enough. Since $lker_n\ f_n \circ \sigma = \{1\}$, for n big enough the map $f_n \circ \sigma$ has to preserve the system $\Psi(x,a) \neq 1$, contradicting the assumption on f_n. □

Let $\mathsf{L}_{A,(P_i)_{i\in I}}(x)$ be a marked graded π-limit group and Δ a graded geometric abelian decomposition of it. We say that $(\mathsf{L}_{A,(P_i)_{i\in I}}(x), \Delta)$ is a solid pair if and only if a sequence of Δ-short restricted morphisms $f_n : \mathsf{G}_A \to \mathsf{K}_A$ with trivial limit kernel exists. We say that a freely indecomposable graded π-limit group $\mathsf{L}_{A,(P_i)_{i\in I}}(x)$ is *solid* if the pair $(\mathsf{L}_{A,(P_i)_{i\in,}}\Delta_{JSJ})$ is solid, where Δ_{JSJ} is the abelian* JSJ tree of $\mathsf{L}_{A,(P_i)_{i\in I}}$.

Suppose that for each $i \in I$ some finite tuple of generators p^i of P^i is given. A *flexible sequence* of morphisms from $\mathsf{L}_{A,(P_i)_{i\in I}}(x)$ to K_A is a sequence $(f_n : \mathsf{L}_A \to \mathsf{K}_A)_{n\in\mathbb{N}}$ of short morphisms satisfying the condition:

$$tl(f_n(p^i_k p^i_l)) \leqslant n \cdot \max_{1\leqslant k,l\leqslant |x|} tl(f_n(x_{l_0}x_{k_0}))$$

For any $i \in I$ and $1 \leqslant k, l \leqslant |p^i|$. By a *flexible π-quotient* of $\mathsf{G}_{A,(P_i)_{i\in I}}$ we will intend the limit π-quotient of some flexible sequence. The fact that parameter subgroups must be elliptic in any limiting tree associated to such a sequence implies, by the results above, that its limit kernel must be non-trivial. Notice that the collection of all flexible sequences is closed under subsequences and diagonal sequences.

A morphism f from $\mathsf{L}_A(x)$ to K_A is said Δ-*flexible* if it factors through some flexible π-quotient after precomposing by some $\tau \in Mod^{\pi}_A(\Delta)$. Otherwise we call it a Δ-solid morphism. By a Δ-solid π-quotient of $\mathsf{L}_{A,(P_i)_{i\in(}}x)$ we intend the limit π-quotient of a sequence of solid morphisms. In absence of an explicit reference to Δ we will tacitly assume it coincides with abelian* JSJ tree of $\mathsf{G}_{A,(P_i)_{i\in I}}$. It can be proven that the distinction between solid and flexible morphisms is independent from the choice of x and p_i, but we will not prove it here. In the degenerate case of a restricted π-group one obtains the definition above has a simple interpretation.

Observation 3.2.7.

- L_A is solid if and only if it injects into F_A
- A morphisms from a freely indecomposable restricted π-limit group L_A to some free K_A is solid if and only if it is injective.

- If $A \neq \{1\}$ or L_A is not cyclic, then up to precomposition with an element of $Mod(\mathsf{L}_A)$, there are only finitely many distinct injective morphisms from a solid L_A into F_A.

Proof. For the first point, suppose a converging sequence of short morphisms $(f_n)_n$ from L_A to K_A exists such that $lker_n\ f_n = \{1\}$. By 3.2.6 this sequence has to be bounded, which implies that all morphisms in some subsequence $(f_{n_m})_m$ can be obtained from each other by postcomposing with an inner automorphism of $\mathbb{F}$, implying they all have trivial kernel. But then $\ker f_{n_m} = lker_m\ f_{n_m} = \{1\}$. Any $f \in Mor(\mathsf{G}_A, \mathsf{F}_A)$ which is non-injective is flexible; here the condition on the growth of parameter groups is empty, so if $f \circ \sigma$ is short for some $\sigma \in Mod(\mathsf{G}_A)$ the constant sequence $f_n = f \circ \sigma$. □

In case A is a free factor of $\mathbb{F}$, the situation above can only be the case if $L = A$, since otherwise the pullback of a free splitting of $\mathbb{F}$ containing A as a factor is a free splitting of L relative to A. This becomes crucial later.

Chapter 4

Makanin-Razborov diagrams

4.1 Basic notions

By a system of equations in the tuple of variables x and the tuple of parameters a we mean a condition of the form:

$$\bigwedge_{i\in I} w_i(x,a) = 1$$

where for $w_i(x,a)$ is a term in the language of groups. This is mostly abbreviated as $\Sigma(x,a) = 1$. By a system of inequations we mean a condition of the form:

$$\bigwedge_{i\in I} w_i(x,a) \neq 1$$

mostly abbreviate by an expression such as $\Sigma(x,a) \neq 1$. We can think of Σ as a subset of elements of any particular $\langle a\rangle$-restricted group. Any quantifier-free definable formula is equivalent to one of the form $\bigvee_{i=1}^{m} \Sigma((x,a)) = 1 \wedge \Pi(x,a) \neq 1$. When passing on to the language $\mathcal{L}^Q$ the same property holds for the intersection of the two types of formulas above with those of the form $\bigcap_{i=1}^{k} P_{q_i}(x_i)$, for $k = |x|$ and $q = (q_i)_{i=1}^{k} \in Q^k$, which we will often abbreviate the latter as $x \in q$. This will be be referred to as π-systems of equations and π-systems of inequations respectively and denoted by an expression of the form $\Sigma^q(x,a) = 1$ (resp. $\Sigma^q(x,a) \neq 1$).

Definition 4.1.0.1. A group G is said to be equationally noetherian if any system of equations with parameters in G is equivalent to some finite subsystem, i.e., they both have the same set of solutions in G.

The property holds for free groups, see [Gub86]. Given a group K, $A \leqslant K$ and a marked restricted group $G_A(x)$ the variety $Var(G,A,x)[K]$ is defined as the set of tuples $x \in K^{\boldsymbol{x}}$ which are the image of $x(G)$ by some morphism from G_A to $\mathbb{F}_A$. Let $Var(G_A,x)$ stand for $Var(G,A,x)[\mathbb{F}]$.

For any surjective resolution homomorphism $\phi : G_A \to H_A$ and compatible x-markings on both G_A and H_A, clearly $Var(H_A,x) \subset Var(G_A,x)$ and it can be easily seen this inclusion is proper in case that G is residually free, that is, whenever for any $g \in G$ there is some homomorphism (morphism) from G (G) to A. Observe that equational noetherianity implies that for every $A \leqslant K$ and any chain of epimorphisms between finitely generated marked groups:

$$G^1_A(x) \to G^2_A(x) \to \cdots G^m_A(x) \to \cdots$$

the sequence

$$Var(G^1_A; x)[K] \supset Var(G^2_A; x)[K] \supset \cdots Var(G^m_A; x)[K] \cdots$$

contains finitely many proper inclusions. Another consequence is that all the homomorphisms in some convergent sequence with $\mathbb{F}$ as target eventually factor through the limit quotient of the sequence.

Corollary 4.1.1. *Every chain of epimorphisms between residually free groups (hence, in particular, between limit groups) contains only finitely many which are proper.*

If we are given a marked restricted π-group, $\mathsf{G}_A(x)$ instead, we can define a π-*variety* $Var(\mathsf{G}_A, x)$ as the set of tuples in $\mathbb{F}^x$ which are the image of $x(G)$ by some morphism from G_A to K_A. We might refer to the points in the π-variety as solutions of $\mathsf{G}_A(x)$.

Notice that the correspondence $f \mapsto f(x)$ is a bijection between $Hom_A(G, \mathbb{F})$ ($Mor(\mathsf{G}_A, \mathsf{F}_A)$) and $Var(G_A, x)$ ($Var(\mathsf{G}_A, x)$). Given a system of equations $\Sigma(x, a) = 1$, let G^{Σ}_A be the A-restricted marked group with presentation $\langle x, A | \Sigma(x, a) = 1\rangle$ (relative to A), where $A = \langle a \rangle$. The set of morphisms from it to F_A determines in turn the same variety defined by Σ. Given a π-system of equations we can expand the previous group to a π-group by sending each generator in the marking to its specified value in Q. We denote the resulting restricted π-group as $\mathsf{G}^{\Sigma}{}_A(x)$, where $A = \langle a \rangle$. The same comment applies.

Given a family $\mathcal{C}$ of π-limit groups or one $\mathcal{C}_A$ of A-restricted π-limit groups, define the partial order $\leqslant_Z$ on $\mathcal{C}$ by saying that $\mathsf{G} \leqslant \mathsf{H}$ ($\mathsf{G}_A \leqslant \mathsf{H}_A$) if and only if there is a surjective morphism from H to G (G_A to H_A). The discussion above implies $\leqslant_Z$ is is well-founded.

Equational noetherianity implies the following characterization of π-limit groups, entirely analogous to that of limit groups.

Lemma 4.1.2. *Let G_A be any finitely generated restricted π-group. The following are equivalent:*

i) *G_A is a restricted π-limit group.*

ii) *G_A is π-discriminated by F_A, that is: for any finite set $S \subset G$ there is a morphism $f : \mathsf{L}_A \to \mathsf{F}_A$ which is injective on S.*

iii) *G_A is a model of the universal part of $Th(\mathsf{F}_A)$.*

Proof. The implication from (i) to (ii) follows from equational noetherianity. The homomorphisms in a ny convergent sequence targetting must factor through the limit quotient after a certain point.

Let us look now at the implication from (ii) to (iii). Take a restricted π-group G_A which is π-discriminated by F_A. All one has to show is that for any tuple b of elements from G and any atomic formula $\phi(x, a)$ ($|x| = |b|$) with constants a in A such that $\phi(b, a)$ holds in L_A some tuple c in F exists such that $\phi(c, a)$ holds in F_A. Observe that it is enough to verify this for a formula of the form:

$$(\bigwedge_{i=1}^{|x|} P_{q_i}(x_i)) \wedge \Sigma(x, a) = 1 \wedge \Delta(x, a) \neq 1$$

since any atomic formula with free variables in x is the disjunction of finitely many formulas of that same form. Property (ii) implies the existence of a morphism f from G_A to F_A which does not kill the words in $\Delta(x,a) \neq 1$. Of course, since this is a homorphism fixing A any element which can be written as $u(b,a)$ for some word $u(x,y)$ is sent by such f to $u(f(b),a)$, so that $\Delta(f(b),a) \neq 1$ holds as well. That $\phi(f(b),a)$ holds is now obvious, the fact that a morphism preserves the positive part of such a formula being equally easy to prove.

The implication from (iii) to (i) follows the inverse path. Let s be a tuple of generators of a given finitely generated model G_A of the universal theory of F_A and u a finite tuple of elements of G. Let H_A the π-group with underlying group $A \circledast \mathbb{F}(s)$, together with the obvious homomorphism to Q. There is a natural morphism p from H_A to G_A. Let $\Sigma_1(x,a) = 1, \Sigma_2(x,a) = 1 \cdots$ an increasing chain of systems of equations whose union contains the collection of all the equations with parameters in a satisfied by the tuple (s,a). Let $\Delta_1(x,a) \neq 1, \Delta_2(x,a) \neq 1 \cdots$ a similar exhaustion, but with inequations in place of equations.

Denote by $\phi_n(x,a)$ be the conjunction of $\Sigma_n(x,a) = 1$, $\Delta_n(x,a) \neq 1$ and the condition $P_{\pi(s_i)}(x_i)$. Since s satisfies ϕ in G_A and G_A is a model of the universal theory of F_A, some tupe s' must also satisfy ϕ in F_A. The map sending s_i to s'_i determines a morphism f_n from K_A to F_A. The choice of f_n implies that the limit kernel of the sequence $(f_n)_n$ is the same as the kernel of the projection p.

□

While equational noetherianity guarantees the termination of the Makanin-Razborov procedure, the next result takes care of the finite branching of the process. We provide a topological interpretation of the diagonalization argument found in [Sel01].

4.1.1 Finite width

Let $\mathcal{F}$ be the family of homomorphisms from a group G to an equationally noetherian group H and let $\mathcal{K}$ a subset of the closure of $\{ker(f) : f \in \mathcal{F}\}$ in 2^G. Recall that the product topology on 2^G is compact and metrizable, hence sequentially compact. One can easily check that in this case $\mathcal{K}$ consists of limit kernels of convergent sequences of homomorphisms in $\mathcal{F}$. Equationally noetherianity conditions the way $\mathcal{K}$ sits in 2^G, namely:

$$\text{For each } K_0 \in 2^G \text{ there is a neighbourhood } V \text{ of } K_0 \text{ such that } \forall K \in V \cap \mathcal{F}\ K_0 \subseteq K$$

Observe that $\subset$ is a closed relationship in 2^G, so one can strengthen the previous statement to:

$$\text{For each } K_0 \in 2^G \text{ there is a neighbourhood } V \text{ of } K_0 \text{ such that } \forall K \in V \cap \mathcal{K}\ K_0 \subseteq K \qquad (*)$$

The intersection of the elements of a chain $\mathfrak{C} = \{\mathcal{K}_i\}_{i \in I}$ of $(\overline{\mathcal{K}}, \subseteq)$ belongs to $\overline{\mathfrak{C}} \subset \overline{\mathcal{K}}$, since the restriction of $\overline{\mathcal{K}}$ to the elements in any ball of the Cayley graph of G eventually stabilizes. Hence $\bigcap_{i \in I} K_i \in \overline{\mathcal{K}}$. The observation above, and Zorn's lemma imply that for any element $K \in \overline{\mathcal{K}}$ there is some $L \in \mathcal{L}$ such that $L \neq K$.

Claim 4.1.3. *If $\mathcal{K}$ is closed then $\mathcal{L}$ is finite.*

Proof. Since the space 2^G is sequentially compact, it is enough to show that $\mathcal{L}$ has no accumulation points. Suppose there was a sequence $(L_n)_{n \in \mathbb{N}} \subset \mathcal{L}$ and $L \in \overline{\mathcal{K}}$, such that $L \neq L_n$ for all n and $(L_n)_n$ converges to L. Then (*) implies that eventually $L \subsetneq L_n$, contradicting the fact that $L_n \in \mathcal{L}$. □

If $\mathcal{K}$ is the set of limit kernels of a family of sequences of homomorphisms from G to H, there is a simple criterion on $\mathcal{S}$ ensuring that $\mathcal{K}$ is compact.

We say that a family $\mathcal{S}$ of sequences is closed under diagonal sequences if for any sequence $((f_n^k)_n)_k \subset \mathcal{F}$ the sequence $(f_n^n)_n$ belongs again to $\mathcal{S}$.

Lemma 4.1.4. *If $\mathcal{S}$ is closed under subsequences and diagonal sequences, then $\mathcal{K}$ is closed in 2^G.*

Proof. This follows easily from the fact that 2^G is metrizable. For any $K_n = lker_n f_m^n$ in a sequence $(K_n)_n$ converging to K, take some $f_{k_n}^n \in \mathcal{F}$ in the corresponding sequence $(f_m^n)_n$ such that $Ker\, f_{k_n}^n$ is as close to K_n as K_n is to K. □

We will refer to a family satisfying the property above as being closed under subsequences. Here are some first consequences of the above discussion of importance to us:

Lemma 4.1.5. *Let G_A a finitely generated restricted π-group and $\mathcal{F}$ a family of morphisms from finitely generated restricted π-group G_A to F_A. Then any $f \in \mathcal{F}$ factors through one of finitely many π-quotients which are limits of sequences of elements of $\mathcal{F}$.*

In particular we have:

Corollary 4.1.6. *For any restricted group G_A there is a finite family $\mathcal{MQ}(\mathsf{G}_A)$ of restricted π-quotient maps from L_A onto an A-restricted π-limit group such that any morphism from G_A to a given π-group H_A factors through some $q \in \mathcal{MQ}(\mathsf{G}_A)$.*

4.1.2 π-Resolutions

One of the most important notions in the work of Sela is that of a resolution. One of their uses is that of parametrizing the set of solutions of a systems of equations. They generalize straightforwardly to our context.

Definition 4.1.6.1. A π-resolution $\mathcal{R} = \mathcal{R}^{(J,r)}$ consists of a finite rooted tree (J, r) and a series of assignments:

- To each $\lambda \in J$:
  - i) A non-trivial π-group R^λ, which is a π-limit group in case $\lambda \neq r$ or $(\Delta\mathcal{R})^\lambda$ is non-trivial.
  - ii) A geometric abelian decomposition $(\Delta\cdot)^\lambda$ of R^λ relative to the family of all non-cyclic abelian subgroups.
- To each $\lambda \in \hat{J}$, a surjective π-morphism η_R^λ from R^λ to the free product:

$$*_{\lambda>\mu}\mathsf{R}^\mu$$

We will refer to R^λ as the top of the resolution $\mathcal{R}$ and to $\mathcal{R}$ as a resolution of G.

Definition 4.1.6.2. By a restricted π-resolution $\mathcal{R}_A = \mathcal{R}_A^{(J,r)}$ we intend a resolution of the underlying π-group G in which for some $A \leqslant \mathbb{F}$ each of the node groups appearing along some branch $\mathcal{B}$ of J are all endowed with an additional restricted structure: A_{R^λ}. Moreover:

- The group of constants A is elliptic in the geometric abelian decomposition $(\Delta\cdot)^\lambda$ for any $\lambda \in \mathcal{B}$.
- For $\lambda \in \hat{\mathcal{B}}$ the map η^λ is a morphism of restricted π-groups onto the free product $*_{\lambda>\mu}\mathsf{R}^\lambda$.

Formally we will consider the node group at any node λ as an A_λ-restricted π-limit groups, where $A_\lambda = A$ in case $\lambda \in \mathcal{B}$ and $\{1\}$ if not. Again we will refer to R^λ_A as the top of the π-resolution and to it as a restricted π-resolution of R^λ_A.

Given a π-resolution $\mathcal{R}^{(J,r)}$, we say that λ is of free product type in case $(\Delta\mathcal{R})^\lambda$ is trivial and $\eta^{\lambda,R}_{\lambda,R}$ is injective. A (restricted or graded) π-resolution will be called proper in case η^λ_λ is non-injective whenever μ is the only child of λ. Given $\mathcal{R}^J$, and $\lambda \in J$, by $\mathcal{R}^J\restriction_\lambda$ we intend its restriction of $\mathcal{R}$ to $J\restriction_\lambda$.

Given a π-resolution $\mathcal{R}^{(J,r)}$, a concretion of it to a π-group H is a pair $(f_\lambda, \tau_\mu)_{\lambda\in J,\mu\in J}$, where the $f_\lambda : \mathsf{R}^\lambda \to \mathsf{H}$ is a morphism and $\tau_\lambda \in Mod(\mathsf{R}^\lambda, \Delta\mathcal{R}^\lambda)$ are such that for any $\lambda \in (\hat{J})$:

$$f_\lambda = (\amalg_{\lambda>\mu} f_\mu) \circ \eta^\lambda_\lambda \circ \tau_\lambda$$

An obvious induction argument shows that any choice of morphisms $\{f_\lambda\}_{\lambda\in lv(J)}$ and of automorphisms $\{\tau_\lambda\}_{\lambda\in(\hat{J})}$ determines a unique concretion of $\mathcal{R}^{(J,r)}$.

The morphism f_r, will be said to factor through $\mathcal{R}$ and we will refer to f_λ as a factorizationfor f at λ. Given a π-resolution $\mathcal{R}^{(J,r)}$, we let $Fct(\mathcal{R})$ be the set of morphisms from G to F which factor through it.

Given a restricted π-resolution , $\mathcal{R}^{(J,r)}_A$, a concretion of $\mathcal{R}$ to a restricted π-group H_A is a concretion $(\{f_\lambda\}_{\lambda\in J}, \{\tau_\lambda\}_{\lambda\in(\hat{J})})$, of the underlying π-resolution , to H such that the f_λ are morphisms from A_{R^λ} to H_A and $\tau_\lambda \in Mod(\mathsf{R}^\lambda, \Delta\mathcal{R}^\lambda)$. We will say that a morphism $f : \mathsf{G}_A \to \mathsf{F}_A$ factors through $\mathcal{R}_A$ if there is a restricted concretion $(f_\lambda)_{\lambda\in J}$ such that $f = f_r$. Once more, we denote by $Fct(\mathcal{R}_A)$ the set of morphisms from G_A to F_A factoring through $\mathcal{R}_A$.

Definition 4.1.6.3. A graded π-resolution $\mathcal{R} = \mathcal{R}_{A;(P_i)_{i\in I}}$ of a graded π-group $\mathsf{G}_{A,(P_i)_{i\in I}}$ is a restricted resolution of G_A together with an additional graded π-structure ${\mathsf{R}^\lambda}_{A,(P_i)_{i\in I}}$ on each of its node groups in such a way that the following conditions are satisfied:

i) For each $\lambda \in \hat{J}$ the index set I_λ is equal to $\bigcup\limits_{\mu<\lambda} I_\lambda$.

ii) For each $j \in I_\lambda$, the parameter group $P^\lambda_j \leqslant \mathsf{R}^\lambda$ is mapped by $\eta^\lambda_{\mathbb{Q}RR}$ onto a conjugate in $*_{\mu<\lambda}\mathsf{R}^\mu$ of P^μ_j for some $\mu < \lambda$ and $j \in I_\mu$. This defines a bijective map from I_λ to $\bigcup_{\mu<\lambda} I_\mu$.

iii) For each $\lambda \in \hat{J}$ the geometric abelian decomposition $(\Delta\mathcal{R})^\lambda$ is relative to $\{P^\lambda_j\}_{j\in I_\lambda}$.

We say that the node group at λ carries parameter groups in case $I_\lambda \neq \varnothing$.

We use the notation $\mathcal{R}$ indistinctively for graded ,restricted or simple π-resolution as long as the context makes clear the kind of structure we are referring to. To abbreviate, given a graded π-resolution $\mathcal{R}_A$ as above and a node $\lambda \in J$ denote the group $Mod(\mathsf{A}_{R^\lambda}, (\Delta\mathcal{R})^\lambda)$ by $Mod(\mathcal{R}^\lambda_A)$. Same comment applies. We say that a graded π-resolution $\mathcal{R}$ is closed if for any $\lambda \in lv(J)$ either:

i) R^λ carries the constants or some parameter subgroup, in which case ${\mathsf{R}^\lambda}_{A,(P_i)_{i\in I}}$ is freely indecomposable and solid

ii) R^λ is free and carries no parameters nor constants.

In the last case we refer to λ as a solid leaf. Recall that in the case of a merely restricted π-resolution G_A if the group A of constants is a free factor of $\mathbb{F}$ then A occupies a whole group at one of the leafs.

By a solid concretion of a closed graded π-resolution (in particular of a restricted π-resolution) we mean a concretion $(f_\lambda, \tau_\mu)_{\lambda\in J,\mu\in\hat{J}}$ of it such that for any $\lambda \in lv(J)$ not carrying any constants for parameters the morphism f_λ is solid. In the case of a restricted π-resolutions this means the groups carrying the constants are injected. In this case we say that f_r factors solidly through $\mathcal{R}_A$ and to $Fct^{sld}(\mathcal{R})$ as the set of morphism factoring solidly through $\mathcal{R}_A$.

We will entertain for some time the the possibility that any π-group along a (graded or restricted) π-resolution might not be a (restricted , graded) π-limit group. We will refer to such an object as a 'weak' (restricted , graded) π-resolution . It turns out that if a weak resolution is closed and strict in the sense below then it is also in fact a π-resolution in the standard sense, as we will show later.

4.2 The Makanin-Razborov procedure

We say that a π-resolution $\mathcal{R}^{(J,r)}$ is *cautious* if $(\Delta\mathcal{R})^r$ is trivial. This will be the case by default for us whenever R^λ is not a limit group. As we saw, understanding the set of solutions of a π-system of equations over some $A \subset \mathbb{F}$ amounts to that of the family of all morphisms from a certain A-restricted π-group G_A to F_A. It will be in fact convenient to analyze the family of all morphisms from G_A to any π-group $\mathsf{F'}_A$ of the form $\mathsf{F}_A * \mathsf{H}$. For a fixed L_A we might as well restrict to a single H.

Claim 4.2.1. *Given a finitely generated* G_A*, there is some* π*-group* H_0 *such that the image of any morphism* $f : \mathsf{G}_A \to \mathsf{F'}_A$ *for* $\mathsf{F'}_A$ *as above factors through* $\mathsf{F}_A * \mathsf{H}_0$*.*

Proof. Given any free π-group $\mathsf{F'} = \mathsf{F}_A * \mathsf{H}$, consider the $\mathbb{F}'$-tree T with trivial edges stabilizers and a single orbit of vertices, stabilized by conjugates of $\mathbb{F}$. Now, suppose that $f : \mathsf{G}_A \to \mathsf{F}_A * \mathsf{H}$ is also given. The map f induces an action of $Q = G/ker(f)$ on T with trivial edge stabilizers. Consider the action of G' on its minimal tree $S \subset T$. There cannot be more than $rk(G') \leqslant rk(G)$ distinct orbits of non-trivially stabilized vertices of S (by the action of Q). Since they are all contained in the same $\mathbb{F}'$-orbit of VT, by adding at most $rk(G)-1$ elements of $\mathbb{F}'$ to Q, we can get $Q \leqslant Q' \leqslant \mathbb{F}'$ with $rk(Q') \leqslant 2rk(G) - 1$ acting on its minimal tree S' with a single non-trivially stabilized orbit. Now, fix some maximal spanning tree of the quotient graph $Q'\backslash S'$ and collapse all the edges in its preimage in S'. The tree so obtained is dual to that of a free decomposition of G' of the form $K * H'$ where $K \subset \mathbb{F}$ and H' is free. Now, $rk(H') \leqslant rk(G') \leqslant 2rk(G) - 1$. Let $\mathsf{H}^1 \cdots \mathsf{H}^m$ be all the non isomorphic π-structures on a free group of rank $\leqslant 2rk(G) - 1$. Hence $\mathsf{H}^0 = *_{i=1}^m \mathsf{H}^i$ satisfies the properties we need. □

The idea behind this is that of preserving some more information on the set of solutions in F, rather than studying the set of solutions in a larger model. We are now in a position to state a very basic Makanin-Razborov type result:

Proposition 4.2.2. *For any finitely generated graded restricted group* $\mathsf{G}_{A,(P_i)_{i\in I}}$ *there is a finite family* $\mathcal{MR}(\mathcal{G})$ *of closed graded* π*-resolutions such that*

$$\mathcal{G} \subset \bigcup_{\mathcal{R}\in\mathcal{MR}} Fct(\mathcal{R})$$

Proof. The proof is by induction on the well-founded partial order $\leqslant^{\mathbf{r}} \cong \leqslant_{rk(G)} \times \leqslant_Z$. If $\mathsf{G}_{A,(P_i)_{i\in I}}$ is freely decomposable, let

$$G = \left(*_{l=0}^k G_l\right) * \mathbb{F}_l$$

be a Grushko decomposition of G relative to $\{A, P_i\}_{i\in I}$, where $A \leqslant G_0$ and no P_i is conjugate into the free factor $\mathbb{F}_l$. Let $A(l)$ be A in case $l = 0$ and $\{1\}$ otherwise and $I(l)$ the set of those

i for which P_i is conjugate into $A(i)$. We can assume that $P_j \leqslant G_l$ for any $j \in I(l)$. For each $0 \leqslant l \leqslant k$ there is by induction a finite family of graded π-resolution $\{(\mathcal{R}^j_l)^{(J^j_l, r^j_l)}_{A(l)}\}_{j\in\Lambda_l}$ such that

$$Mor(\mathsf{G}_{\mathsf{l}A(l)}, \mathsf{K}_{A(l)}) = \bigcup_{j\in\Lambda_i} Fct^{sld}((\mathcal{R}^j_l)_{A(l)})$$

Let $(\mathsf{G}^{\mathsf{j}}_{\mathsf{l}})_{A(l),(P^l_i)_{i\in I_j}}$ be the top graded π-limit group appearing in $\mathcal{R}^{\mathsf{l}}$.

For each choice of a tuple $\bar{j} = (j_l)^k_{i=0} \subset \Lambda_0 \times \cdots \Lambda_k$ we can construct a graded π-resolution $(\mathcal{R}^{\bar{j}})^{(J^{\bar{j}}, r^{\bar{j}})}_A$ of $\mathsf{G}_{A,(P_i)_{i\in I}}$, where $\Delta\mathcal{R}^{r^{\bar{j}}}$ is trivial, the children of $r^{\bar{j}}$ contain the groups $\{G^{j_i}_l, \lambda\}^k_{l=1}$ and F_l with the appropriate additional structure, the map $\eta^{R^{\bar{j}}}_{R^{\bar{j}}}, r^{\bar{j}}$ is the obvious quotient, $\mathcal{R}^{\bar{j}}_A\restriction_{r^{j_l}_l} = (\mathcal{R}^{j_l}_l)^{(J_l, r^{j_l}_l)}_{A_l}$ and F_l is terminal. Clearly any morphism from G_A to K_A factors solidly through some of the π-resolutions constructed this way.

Suppose now $\mathsf{G}_{A,(P_i)_{i\in I}}$ is freely indecomposable. If $\mathsf{G}_{A,(P_i)_{i\in I}}$ is solid add the graded resolution $\mathcal{R}^{sol}$ whose single node contains the pair $(\mathsf{G}_{A,(P_i)_{i\in I}}, \Delta_{JSJ})$, where Δ_{JSJ} is an abelian* JSJ decomposition of $\mathsf{G}_{A,(P_i)_{i\in I}}$. Chose a marking x of G_A and let $\mathcal{F}$ be the finite family of maximal flexible limit quotients of sequences from $\mathsf{G}_{A,(P_i)_{i\in I}}$ to K_A with respect to x and the action of K on its Cayley graph.

Each $q : \mathsf{G}_A \to \mathsf{L^q}_A \in \mathcal{F}$ is proper, so we can apply the induction hypothesis to $\mathsf{L^q}_{A,(P_i)_{i\in I}}$, which results in a finite set $\mathcal{MR}^q$ of resolutions of L_A such that $Mor(\mathsf{L^q}_A, \mathsf{K}_A) = \bigcup_{q\in\mathcal{M}} Fc^{sol}(R^q)$.

For each $\mathcal{R}^{(J,r)}_A \in \mathcal{MR}^q$ construct a graded π-resolution $\mathcal{R}'^{(J',r')}_A$ of $\mathsf{G^q}API$ by attaching a new root r' associated to L_A and letting $(\Delta\mathcal{R}')^{r'} = \Delta_{JSJ}$ and taking q as $\eta^{r'}_{\mathcal{R}'}$. Given any flexible $f \in Mor(\mathsf{G}_{A,(P_i)_{i\in I}}, \mathsf{K}_A)$, there is by definition some $\sigma_r \in Mod(\mathsf{G}_{A,(P_i)_{i\in I}})$ such that $f \circ \sigma_r = f' \circ q$ for some $q \in \mathcal{M}$.

It is obvious how to extend some essential system of curves of f' through $\mathcal{R} \in \mathcal{MR}^q$ to some essential system of curves of f through the corresponding $\mathcal{R}'$. In case G_A is not a restricted π-limit group we know that there is some finite set $\mathcal{MQ}(\mathsf{G}_A)$ of quotients of L_A to A-restricted limit groups such that every morphism from G_A to K_A factors through one of them. This clearly reduces this case to the previous one. □

4.3 Strict π-resolutions

Definition 4.3.0.1. Suppose we are given a restricted π-limit groups G_A and $\mathsf{H}A$, as well as a geometric abelian decomposition Δ of G_A, we say that a morphism $g : \mathsf{G} \to \mathsf{H}$ is *strict* with respect to G, Δ if for any sequence $(h_n)_n$ of morphisms from G_A to F_A there is a sequence of automorphisms $\sigma_n \in Mod(\mathsf{G}_A)$ such that $(h_n \circ g \circ \sigma_n)_n$ has trivial limit kernel as well.

Remark 4.3.1. The former is equivalent to the fact that for any (some) sequence $(h_n : \mathsf{G} \to \mathsf{F})_n$ with trivial limit kernel there are $\sigma_n \in Mod(\mathsf{G}_A, \Delta)$ such that $lker_n\ h_n \circ h \circ \sigma_n = \{1\}$.

Given a geometric abelian tree T of G we let T^f be the G-tree obtained by folding together for each abelian type vertex v all the edges incident to v which be long to the same $Per^*(v)$-orbit. Denote by e^f and v^f the image in T^f of an edge e and a rigid type vertex v respectively. All those edges and vertices of T not involved in the folding can be thought of as identical to their image in T^f.

Given a rigid vertex v of T, $Stab(v^f)$, is the subgroup of G generated by $Stab(v)$ and $Per(v)$ for all $w \in VT_a$ of abelian type adjacent to v, what one usually refers to as the envelope of v.

Definition 4.3.1.1. We say that a homomorphism between groups G and H is formally strict with respect to a geometric abelian tree T of G if:

i) For any rigid type vertex v of T^f, f is injective on $Stab(v)$.

ii) For any $e \in E$, f is injective on $Stab(e)$

iii) For any surface type vertex v of T^f the image $f(Stab(v))$ is non-abelian.

The following fact is well-known in the standard case (see for example [BF09] and [Wil09])).

Lemma 4.3.2. *A morphism f from a restricted π-group G_A and a restricted π-limit group L_A which is formally strict is also strict.*

Proof. The implication from right to left is immediate. Given any of the standard generators of $\tau \in Mod(\mathsf{R}_A, \Delta)$ it is easy to check by inspection that τ restricts to an inner automorphism on:

i) $Stab(v)$ for any rigid type vertex $v \in T^f$.

ii) $Stab(e)$ for any $e \in E$

Hence the kernel of the restriction of f and $f \circ \tau$ to any of those groups coincide. Moreover, for any surface type vertex $v \in T^f$, τ takes the image of $Stab(v)$ to a conjugate in G, hence $f \circ \tau$ sends it to an abelian group if and only if f does.

The converse can be shown using a generalization of the method fleshed out in [Wil09]. See section 5.6 for more details. □

In particular, if such a formally strict morphism exists G_A is also a restricted π-limit group.

Definition 4.3.2.1. A restricted π-resolution $\mathcal{R}^{(J,r)}_A$ will be called strict if and only if each map η^λ_λ is strict with respect to $(\Delta\mathcal{R})^\lambda$. It will be called quasi-strict if the property holds for any $\lambda \neq r$ and the root r has a single child and is associated to a trivial decomposition.

Lemma 4.3.3. *For any graded π-group $\mathsf{G}_{A,(P_i)_{i\in I}}$ there is a finite family $\mathscr{S}$ of closed quasi-strict graded π-resolution of $\mathsf{G}_{A,(P_i)_{i\in I}}$ such that any morphism from G_A to K_A factors solidly through one of them.*

Proof. Start with each $\mathcal{R}^{\in}_{(J,r)}\mathcal{MR}(\mathsf{G}_{A,(P_i)_{i\in I}}, \mathcal{F})$ of $\mathsf{L}_{A,(P_i)_{i\in I}}$ which is not strict. For each $\lambda \in (J,r)$ let $\mathcal{H}_\lambda = Fct(\mathcal{R}\restriction_\lambda) \subset Mor(\mathsf{A}_{R^\lambda}, \mathsf{F}_{A_\lambda})$. The fact that the resolution is not strict implies that for some $\lambda \in J$ any limit quotient of morphisms in $\mathcal{H}_\lambda$ is proper; let λ_0 be some maximal index in J for which this is the case and $\mathcal{MQ}_\lambda$ the associated finite set of maximal quotients.

For each $p : \mathsf{A}_{R^\lambda} \to \mathsf{L}_{A_\lambda} \in \mathcal{MQ}_\lambda$ we consider all possible resolutions obtained by replacing $\mathcal{R}\restriction_\lambda$ with some π-resolution $\mathcal{R}'^{(J',r')} \in \mathcal{MR}(\mathsf{L^p}_{A_\lambda})$ in the obvious sense. Let $Rep_\lambda(\mathcal{R})$ be the finite family of all the π-resolutions obtained as we let p range within $\mathcal{MQ}_\lambda$. Clearly any morphism factoring through $\mathcal{R}^\lambda_A$ factors through some $\mathcal{R}'_A \in Rep_\lambda(\mathcal{R})$.

Iterate the operation in parallel on each $\mathcal{R} \in Rep_\lambda(\mathcal{R})$ as long as it is possible. We can represent the course of this substitution procedure using a rooted tree $(\mathfrak{T}, \rho)$ each of whose nodes τ indexes some π-resolution $\mathcal{R}^\tau$. Each pair $\tau \prec \tau'$ corresponds to the derivation of $\mathcal{R}^\tau$ from $\mathcal{R}^{\tau'}$ by performing a substitution as described above.

If we manage to prove that $\mathfrak{T}$ is finite, then we are done. Any resolution appearing at a leaf of $\mathfrak{T}$ will be strict. Moreover, in that case the union of the π-resolutions which appear at the leaves of $\mathfrak{T}$ at the starting π-resolution ranges within $\mathcal{MR}(\mathsf{G}_{A,(P_i)_{i\in I}})$ will meet our requirements. In virtue of König's lemma this amounts to showing that $\mathfrak{T}$ contains no infinite branches. A π-resolution resolution can be seen as an element of the partial order $Tr := Tr(PL, \leqslant_{rk} \times \leqslant_Z)$ described in subsection 1.4, so the result follows from the somewhat more general fact that $(Tr, \leqslant_{Tr})$ is well-founded (the map $\tau \mapsto \mathcal{R}^\tau \in Tr$ is strictly monotonous by construction). □

Although in general it is not possible to capture all morphisms from L_A to a free K_A using finitely many strict restricted π-resolutions , there is always at least one.

Lemma 4.3.4. *Let L_A be a freely indecomposable restricted limit group and Δ_{JSJ} its JSJ decomposition. Then L_A admits a non-injective Δ_{JSJ}-strict morphism $\phi : \mathsf{L}_A \to \mathsf{K}_A$. As a consequence, any such L_A admits a strict π-resolution $\mathcal{R}_A^{(J,r)}$.*

Proof. Fix a tuple of generators of L. For each $n \in \mathbb{N}$ we know a morphism $f_n : \mathsf{L}_A \to \mathsf{F}_A$ which is inyective on the ball of radius n in the Cayley graph. Now take $\tau_n \in Mod(\mathsf{L}_A, \Delta_{JSJ})$ such that $f_n \circ \tau_n$ is short in the sense of 3.2.1. Any convergent subsequence of $(f_n)_n$ will have non-trivial limit kernel. The second statement can be proven by induction using equational noetherianity, as usual. □

4.4 Well-separated π-resolutions

A well-separated graded π-resolution is a special type of quasi-strict graded π-resolution $\mathcal{R}^{(J,r)}$ to which is possible to assign an completion (see section 5.3). More precisely, for any $\lambda \in \hat{J}$ not of free product type there is a distinguished set

$$DFt(\lambda) \subset Ch(\lambda) \cap lv(J)$$

so that the group at $\mu \in DFt(\lambda)$ is free and the following conditions hold:

i) For every vertex v of non-surface type the image of $(\Delta\mathcal{R})_v^\lambda$ by η_λ^λ is contained in R^μ for some $\mu_v \prec \lambda$.[(1)]

ii) For $\mu \prec \lambda$, if R^μ does not intersect the image of a non-surface type vertex $\mu \in lv((J))$. Moreover, R^μ is a free factor of the image of the subgroup generated by some unique surface type vertex group $(\Delta\mathcal{R})_v^\lambda$ and the Bass-Serre elements associated to edges incident to v.

iii) If $(\Delta\mathcal{R})^\lambda$ consists of a single vertex of surface type or non-cyclic abelian type, then $Ch(\lambda) = DFt(\lambda)$. .

iv) If R^λ is free and does not carry a constant or parameter group, then $\lambda \in lv(J)$.

We will refer to $\lambda \in DFt(\lambda)$ as a dropping node at λ and to R^λ as a factor dropping at λ. Let also $nDFT(\lambda) = Ch(\lambda) \backslash DFt(\lambda)$.

Observe that the second condition implies that no $\lambda \in DFt(\lambda)$ can carry any parameters or constants. We will let $DFt(T) = \bigcup_{\lambda \in \hat{J}} DFt(\lambda)$ and $DFT = (J) \backslash DFt(T)$.

Note that leaves in $DFt(\lambda)$ are not necessarily the only ones for which T^μ is free. Also that the fact that the resolution is strict implies $\eta_{\mathcal{R}}^\lambda$ is injective on $P = Per^*(v)$, whose image is necessarily elliptic in the free decomposition $\ast_{\lambda \succ \mu} \mathsf{R}^\mu$, as otherwise the set of its elliptic elements is the kernel of a map onto $\mathbb{Z}$ containing the finite index subgroup generated by all incident edge groups (see 2.1.8). If $Per^*(v, (\Delta\mathcal{R})^\lambda)$ is cyclic, the fact that $\mathcal{R}$ is well-separated implies that the image of some finite index subgroup of P is elliptic, implying that the same holds for P itself. This leads to the following observation:

(1) It is important that it is contained as opposed to merely conjugate into it.

Remark 4.4.1. Given any path $[u_1, u_2, \cdots u_m]$ of the underlying graph (V, E) of $(\Delta\mathcal{R})^\lambda$ containing only abelian an rigid subgroups, all Δ_{u_j} for $u_j \in V_s$ or $Per(u_j)$ for $u_j \in V_a$ project into the same factor R^μ. Moreover, the Bass-Serre element associated to the edge e_i from u_i to u_{i+1} maps into R^μ as well.

Define $DFt(R, \lambda)$ as the set of μ such that R^μ does not contain the image of a non-surface type vertex of $(\Delta\mathcal{R})^\lambda$, and $nDFT(\lambda)$ as $Ch(\lambda)\backslash DFt(R, \lambda)$.

By the dropping component at λ we intend $DCt(R, \lambda) := *_{\mu \in DFt(R,\lambda)}\mathsf{R}^\mu$ and by the non-dropping component at λ we intend $NDCt(R, \lambda) := *_{\mu \in nDFT(\lambda)}\mathsf{R}^\mu$.

4.5 Taut π-resolutions

The fact that a morphism factors through a resolution provides plenty of information about it. Unfortunately, not enough for the verification process of $\forall\exists$ formulas, which relies on a fine control of both completion and shortening. This is what motivates the introduction of taut π-resolutions; properly speaking they consist of a resolution together with some additional information on each of its nodes, which determines a restricted notion of factoring. Suppose we are given a group G, a geometric abelian decomposition Δ of G, an essential system of curves δ in Δ. Recall that in this situation the image G^p_δ admits a free decomposition of the form:

$$(*_{i=1}^m A_i) * (*_{\Pi \in Pinch^{ex}(\Delta,\delta)}\pi_1(\Pi)) * \mathbb{F}_{b_1(X^c_\delta)} \tag{4.1}$$

Given a homomorphism $f : G \to H$ we say that f is taut with respect to δ if:

- f pinches δ
- $\bar{f}$ does not pinch any essential simple closed curve in any of the surfaces in $Pinch^{in}(\Delta, \delta)$
- The external rank of Δ^p_δ is maximal among all possible choices of δ for which the first two properties are satisfied.

Suppose now that we are also given a π-structure G on G. For any surface type vertex v the group $Mod(\mathsf{G}, \Delta)$ restricts to some finite index subgroup of $Out^\delta(\pi_1(\Sigma_v))$, corresponding to a finite index subgroup of $Mod(\Sigma_v)$. For each such v chose a set $\mathcal{X}_v$ of representatives of the action of the latter on the family of all essential systems of curves on Σ_v; this is finite. Denote by $Ess(\mathsf{G}, \Delta)$ (the result does not depend on A) the family of all the essential systems of curves on Δ obtained by picking some $\delta_v \in \mathcal{X}_v$ for each v of surface type. The discussion above implies that this is finite. By a taut structure on Δ we intend an essential system of curves in Δ such that $\Sigma_v \cap \delta \in Ess(\Sigma_v)$ for all vertices v.

Definition 4.5.0.1. A taut restricted π-resolution of a restricted π-group G_A is given by $(\mathcal{R}^{(J,r)}_A, \delta^\lambda)_{\lambda \in (J)}$, comprising a well-separated restricted π-resolution $\mathcal{R}^{(J,r)}_A$ and for each $\lambda \in J$ a taut structure on δ^λ on Δ such that for $\lambda \in \hat{J}$:

- η^R_λ is taut with respect to δ^λ
- $\{\mathsf{T}^\mu \mid \mu \in nDFT(\lambda)\}$ are precisely the images of those free factors in the first term of 4.1.
- $\{\mathsf{T}^\mu \mid \mu \in DFt(\lambda)\}$ are each one a maximum rank free image of the remaining factors.

We will refer to $(\Theta^\lambda)_\lambda$ as a taut structure on $\mathcal{R}_A$.

By a concretion of $\mathscr{R}_A = \mathcal{R}_A[\Theta]$ we intend a concretion $(f_\lambda, \tau_\mu)_{\lambda \in (J), \mu \in (\hat{J})}$ of $\mathcal{R}_A$ such that for all $\lambda \in (J)$ $f^\lambda \circ \sigma^\lambda$ is taut with respect to δ for some $\sigma^\lambda \in Mod^\pi_{A_\lambda}((\Delta\mathcal{R})^\lambda)$, where $\sigma^\lambda = (\tau^\lambda)^{-1}$ in case $\lambda \in \hat{J}$. A morphism from G_A to F_A factors through $\mathcal{R}_A[\Theta]$ if and only if it belongs to some taut concretion of $\mathcal{R}_A[\Theta]$. If we are additionally given an additional graded structure on the underlying π-resolution talk about a taut graded π-resolution .

Recall we are allowed to have one non-strict map at the top of the π-resolution , provided the geometric abelian decomposition at the top is trivial. We will say that a taut graded π-resolution $\mathcal{R}[\Theta]$ is closed if $\mathcal{R}$ is closed as a graded π-resolution and $\delta_\lambda = \varnothing$ at any leaf λ.

Lemma 4.5.1. *Suppose we are given a graded π-limit group $\mathsf{G}_{A,(P_i)_{i\in I}}$ and a family $\mathcal{G}$ of morphisms from G_A to some free K_A. Then there is a finite family $\mathcal{TMR}(\mathcal{F})$ of closed taut resolutions with the following properties:*

i) Any morphism in $\mathcal{G}$ factors solidly through one of them.

ii) For each $\mathcal{R}^{(J,r)}$ and any δ^λ the restriction of $\eta^\lambda_{\mathcal{R}}$ to each of the inner factors of Δ^c_δ is the limit quotient of a sequence of factorizations at λ of morphisms in $\mathcal{G}$. (2)

Proof. To begin with, define the notion of a pseudo-taut graded π-resolution in the same way as a taut one, by replacement the requirement that the π-resolution is strict with the weaker one that the groups assigned to those edges of $(\Delta\mathcal{R})^\lambda$ adjacent to a surface type vertex group to be mapped injectively by $\eta^\lambda_{\mathcal{R}}$. We modify the definition of a concretion accordingly: the map f^λ is not required to be injective on boundary subgroups of surface type vertex groups of $(\Delta\mathcal{R})^\lambda$.

Claim 4.5.2. *Given any $\mathsf{G}_{A,(P_i)_{i\in I}}$ and $\mathcal{G}$ as above, there is a finite family $\mathcal{MR}^{qt}(\mathcal{G})$ of closed quasi-taut graded π-resolutions such that every morphism in $\mathcal{G}$ factors solidly through one of them and property (ii) holds.*

Proof. The proof is by induction on the $\leqslant_{rk} \times \leqslant_Z$ of G, using a very similar construction to that used in 4.2.2. In view of this, we will merely underscore the differences between both constructions. By the arguments employed there we can assume that $\mathsf{G}_{A,(P_i)_{i\in I}}$ is a freely indecomposable graded π-limit group, with abelian* JSJ decomposition Δ^{JSJ}.

First of all, for any $e \in \Delta^{JSJ}$ one of whose endpoints is of surface type let $\mathcal{G}_e$ the family of all those morphism in e which kill Δ_e. Let $\mathcal{MQ}_e$ the finite family of maximal limit quotients of sequences of morphisms in $\mathcal{G}_e$. Applying the induction hypothesis to each of them we can easily obtain a family $\mathcal{MR}^{qt}(\mathcal{G}_e)$ satisfying the required conditions with respect to the family $\mathcal{G}_e$. Let $\mathcal{G}_{ndeg}$ be the family of all those $f \in \mathcal{G}$ which do not belong to any of the $\mathcal{G}_e$ above. We can write $\mathcal{G}_{ndeg} = \bigcup_{\delta\in Ess} \mathcal{G}_\delta$, where $\mathcal{G}_\delta$ is the family of all those $f \in \mathcal{G}$ such that $f \circ \tau$ pinches $\delta \in Ess := Ess(\mathsf{G}, \Delta^{JSJ})$ for some $\delta \in Ess$ and the external rank associated to δ is maximal for that property. If $\delta \neq \varnothing$, then consider the inner factors of $G_1 \cdots G_k$ of G^p_δ. We can assume that $A \leqslant G_1$ and $P_i \leqslant G_{j_i}$ for any $i \in I$ and some $1 \leqslant j_i \leqslant k$. For each $f \in \mathcal{G}$ denote by $\bar{f}$ the unique map such that $f = \bar{f} \circ p^\delta$ and for any $1 \leqslant j \leqslant k$ by $\mathcal{G}_{\delta,j}$ the collection $\{\bar{f}\upharpoonright_{A_j} \mid f \in \mathcal{G}\}$. Let $\mathcal{MR}_j$ the finite collection of graded π-resolutions of $\mathsf{G}_{A_i,(P_i)_{i\in I_i}}$ obtained applying the induction hypothesis to the $\mathcal{G}_{\delta,j}$, where $A_i = A \cap G_i$ and $I_i = \{j \mid i_j = i\}$. Given any tuple $(\mathscr{R})_j \in \mathcal{MR}^{qt}_1 \times \cdots \mathcal{MR}^{qt}_k$ for all $1 \leqslant j \leqslant k$ we can construct a pseudo-taut π-resolution of $\mathsf{G}_{A,(P_i)_{i\in I}}$ with $(\mathsf{G}_{A,(P_i)_{i\in I}}, \Delta_{JSJ}, \delta)$ at its root r, where $\eta^r_{\mathcal{R}} = (q_1 \amalg q_2 \cdots q_k) \circ p^\delta$ and q_j is the quotient map of G_j onto $\mathcal{R}_j$. Denote by $\mathcal{MR}^{qt}_\delta$ the collection of all of them. If $\delta = \varnothing$ then we proceed exactly as in the proof of 4.2.2. If $\mathsf{G}_{A,(P_i)_{i\in I}}$ is solid, we have to take both the trivial π-resolution where $\mathsf{G}_{A,(P_i)_{i\in I}}$ is the only vertex and those coming from applying the induction hypothesis to the maximal flexible $\mathcal{G}$-quotients. By the latter we mean the limit quotients of flexible sequences whose members are of the form $f \circ \tau$, where $f \in \mathcal{G}$ and $\tau \in Mod(\mathsf{G}_{A,(P_i)_{i\in I}})$.

It is easy to check that the union $(\bigcup_{e\in E\Delta^{JSJ}} \mathcal{MR}^{qt}_e) \cup (\bigcup_{e\in Ess} \mathcal{MR}^{qt}_\delta)$ meets all of our demands. □

(2) This is in fact stronger than what will be needed in the analysis of $\forall\exists$-sentences.

One can use the same argument as in the proof of 4.3.3 to replace each $\mathcal{R}$ in $\mathcal{MR}^{qt}$ by finitely many taut π-resolution capturing all the morphisms which factor through it. This time each of the nodes in the rooted tree $(\mathfrak{T}, \leqslant, r)$ is assigned some pseudo-taut $\mathscr{R}_{\alpha}^{J_{\alpha}}$, satisfying property (i) of the statement of the lemma. Denote by $\mathcal{G}_{\alpha}$ the family of all the morphisms of $\mathcal{G}$ which factor through it.

There is essentially one single obstruction to the termination of the process at any given node $\alpha \in \mathfrak{T}$, namely, that for some $\lambda \in J_{\alpha}$ the map $\eta_{\mathcal{R}}^{\lambda}$ is not the limit of a sequence of factorizations at λ of members of $\mathcal{G}_{\alpha}$; note that this is necessarily the case if $\mathcal{R}_{\alpha}$ fails to be strict [(3)].

Chose some maximal λ witnessing this failure. Let $\Delta = (\Delta\mathcal{R})^{\lambda}$ and $\bar{G}_{\alpha}^{\lambda}$ be the set of all factorizations at λ of morphisms from $\mathcal{G}_{\alpha}$. This implies for some child μ of λ which is the image of an inner factor G_{μ} of $\Delta_{\delta_{\lambda}^{\alpha}}^{p}$ all $f \in \bar{G}_{\alpha}^{\lambda}$ factor through $q \circ \eta_{\mathcal{R}_{\alpha}}^{\lambda}$ for q in some finite family $\mathcal{MQ}_{\alpha}^{\lambda}$ of proper quotients of ${\mathsf{R}_{\alpha}}^{\lambda}$. At each children of α we place the taut π-resolution obtained by replacing (the construction is described in more detail in the proof of 4.3.3) $\mathcal{R}_{\alpha} \restriction_{\lambda}$ by some member of $\mathcal{MR}^{qt}(\bar{\mathcal{G}}_{\mu})$ (composing $\eta_{\mathcal{R}_{\alpha}}^{\lambda}$ with the top quotient map of the latter). Just as in 4.3.3, the process must eventually stop. □

(3) We can always assume the root has a trivial decomposition and a single child π-resolution .

Chapter 5

π-Towers

5.1 Definition

Although π-limit groups admit a fairly regular description in terms of resolutions, as seen in the previous section, for many purposes one must restrict oneself to the smaller class of those having a π-tower structure, which we now go on to define.

Definition 5.1.0.1. By a π-tower structure we will intend a (a priory) weak well-separated πresolution $\mathcal{T} = \mathcal{T}^{(J,r)}$ where for each $\lambda \in (\hat{J})$ not of free product type the following conditions are satisfied:

a) The graph $X = (V, E)$ underlying the decomposition $(\Delta\mathcal{T})^{\lambda}$ is a bipartite graph, $V = V_1 \coprod V_2$, where $V_1 \neq \varnothing$ and V_2 consist precisely of those vertices of rigid type.

b) Each $v \in V_2$ of abelian type has incidence number at most 1 (incidence number 0 implies the whole group is abelian).

c) There is a bijective correspondence between V_1 and the set $nDFT(\lambda)$ assigning to each $v \in V_1$ some $\mu_v \in nDFT(\lambda)$ for which $(\Delta\mathcal{T})^{\lambda}_{v} = \mathsf{T}^{\mu_v}$.

d) The map $\eta^{\lambda}_{\mathcal{R}}$ restricts to the identity on $(\Delta\mathcal{T})^{\lambda}_{v}$ for any $v \in V_1$. Observe that the collection $(\Delta\mathcal{T})^{\lambda}_{v}$ for $v \in V_1$ generate a group isomorphic to their free product.

A restricted (graded) π-tower is simply a restricted (graded) π-resolution whose underlying π-resolution is a π-tower.

If a restricted π-tower $\mathcal{T}_A$ is closed, we will refer to it as a π-tower structure for $\mathsf{A}_{S^{\lambda}}$. If a collection $\Lambda \subset lv(J)\backslash DFt(\mathcal{T})$ exists such that declaring $\mathsf{A}_{S^{\lambda}}$ for $\lambda \in \Lambda$ as a parameters makes $\mathcal{T}_A$ into a closed graded resolution, we refer to it as a tower structure for T^{r}_{A} relative to $\{\mathsf{T}^{\lambda}\}_{\lambda\in\Lambda}$. Conversly, we will refer to T^{r}_{A} as the top of the π-tower.

The way towers are described in most of the literature, the fundamental group of a closed surface must belong to the very bottom of the tower, while the existence of a homomorphism with non-abelian image from this fundamental group to the free group with is expressed in terms of the surface having characteristic $\leqslant -2$.

Let now $\mathcal{T}_A$ be a restricted π-tower relative to the collection $\{\mathsf{T}^{\lambda}\}_{\lambda\in\Lambda}$. Observe that if in the situation above $\lambda \in nDFT(\lambda)$, and $\mu > \lambda$, then T^{λ} is a rigid vertex of $(\Delta\mathcal{R})^{\mu}$ and hence is contained in T^{μ}. Iterating the argument one can easily see that in fact $\mathsf{T}^{\lambda} \leqslant \mathsf{T}^{r}$ in this case. Denote by $tlf(J)$ the set of minimal nodes in $nDFT(\mathcal{T})$. The group generated by $\{\mathsf{T}^{\lambda}\}_{\lambda\in tlf(J)}$ is isomorphic to $*_{\lambda\in tlf(J)}\mathsf{T}^{\lambda}$.

As a convention, we denote the π-group on top of a closed π-tower with the same letter, so unless otherwise specified, by T_A we will always mean T^r_A and so on. We will also sometimes improperly refer to T_A itself as a π-tower instead of as a group admitting a π-tower structure.

5.2 Closures

Let $\mathcal{T}_A^{(J,r)}$ be a restricted π-tower structure and M a maximal non-cyclic abelian subgroup of T. Clearly the set of those $\lambda \in nDFT$ for which some conjugate of M intersects R^λ non-trivially is a descending branch $r = \lambda_0 > \lambda_1 > \lambda_2 > \cdots \lambda_m$ in J. Up to replacing M by some of its conjugates, we can assume that $M_{\lambda_j} = M \cap \mathsf{T}^{\lambda_j}$ is non-trivial. Given λ, let $r_j = rk(M_{\lambda_j})$. We claim that $r_j > r_{j+1}$ for some $0 \leqslant i \leqslant m-1$ if and only if M is conjugate to $Z_{G(A)}$, where A is some abelian vertex group of $(\Delta\mathcal{T})^{\lambda_j}$ attached to the rigid vertex group $\mathsf{T}^{\lambda_{j+1}}$ of $(\Delta\mathcal{T})^{\lambda_j}$. It is easy to see that there can be no growth in the case λ_j is not of free-product type, so we might as well assume it is of floor type. All non-cyclic abelian subgroups of T^{λ_j} are elliptic in $(\Delta\mathcal{T})^{\lambda_j}$.

Now, since $M \cap \mathsf{T}^{\lambda_{j+1}} \neq \{1\}$, if M did not intersect any abelian vertex group of $(\Delta\mathcal{T})^{\lambda}$, then the rigid vertex group T^{λ_j} would be the only one into which one might conjugate M_{λ_j}, contradicting the growth in rank. On the other hand, centralizers in surface groups are cyclic, so necessarily M contains a conjugate of an abelian vertex group, as desired. Let us note in passing that this shows that:

Lemma 5.2.1. *Any maximal non-cyclic abelian subgroup of T^r is either conjugate into some of the abelian vertex groups in the decomposition associated to some node of the tower or can be conjugated into L_i for some $1 \leqslant i \leqslant m$.*

Consider first the case in which λ_m neither belongs to Λ nor is of isolated abelian type (observe this includes the possibility that T^{λ_m} is cyclic). Let us give a closer look at T^{λ_m}. In case λ_m is of free-product type M_{λ_m} has to be hyperbolic in the corresponding free product, hence cyclic. In the floor case the minimality of λ_m implies M_{λ_m} cannot intersect a conjugate of a rigid type vertex non trivially, hence (M is maximal) neither a conjugate of an abelian type vertex. On the other hand, maximal abelian subgroups of surface groups and free groups are cyclic, so M_{λ_m} has to be cyclic in every other case as well.

Whenever M_{λ_m} is cyclic we will refer to M as a pegged abelian group and to any generator of the conjugate of M_{λ_m} in M as a peg of M. Let $\mathcal{PM}$ denote some system of representatives of the conjugacy classes of pegged non-cyclic maximal abelian subgroups of T.

In case M_{λ_m} is an isolated abelian vertex group, we will say that M is *free hanging non-cyclic abelian* and refer to the unique conjugate of M_{λ_m} contained in M as the bottom of M, denoted by $bot(M)$. Let $\mathcal{FHM}(T)$ stand for the collection of representatives of conjugacy classes of free hanging non-cyclic abelian groups. We remark that if $\Lambda = \emptyset$, any non-cyclic maximal abelian group is conjugate to an element of $\mathcal{PM}(T) \cup \mathcal{FHM}(T)$.

Definition 5.2.1.1. Suppose we are given a closed restricted -π-tower $\mathcal{T}^J_A$. By a *closure* of $\mathcal{T}$ we intend a pair $(\kappa, \mathcal{S}^J_A)$ where $\mathcal{S}^J_A$ is a π-tower with the same index set and dropped nodes as $\mathcal{T}_A$ and κ an injective morphism

$$\kappa : \mathsf{T}_r * (*_{\lambda \in DFt(T)} \mathsf{T}^\lambda) \to \mathsf{S}_r * (*_{\lambda \in DFt(S)} \mathsf{S}^\lambda)$$

compatible with the quotient maps in the two resolutions and such that:

i) λ is of free product type in $\mathcal{S}$ if and only if it is in $\mathcal{T}$

ii) If $\lambda \in lv(J)$ then κ restricts to an isomorphism between T^λ and S^λ, unless the first is the image of a free hanging non-cyclic abelian group, in which case the codomain of the inclusion is abelian as well and contains the image as a finite index group.

iii) For any other $\lambda \in J$ there is a vertex type preserving isomorphism ϕ between the underlying graphs of the decompositions $\Delta = (\Delta\mathcal{T})^{\lambda}$ and $\Delta^{S} = \Delta\mathsf{S}^{\lambda}$ such that $\kappa(t_e) = t_{\phi(e)}$, κ restricts to embeddings of Δ^T_e into $\Delta^S_{\phi(e)}$ and of Δ^T_v into $\Delta^S_{\phi(v)}$. Moreover, $\kappa(\Delta^T_v)$ is equal to Δ^S_v if v is of surface type and a finite index subgroup if it is of abelian type. The image in Q of Δ^S_v and Δ^T_v are the same.

Incurring in a certain abuse of language we sometimes refer as the underlying group S_A itself as a closure of T_A.

Notice that given a system of representatives $\mathcal{M}$ of the conjugacy classes of non-cyclic maximal abelian subgroups of T_A and $\mathcal{S}_A$ a closure of $\mathcal{T}_A$, the group S is the fundamental group of a star shaped graph of groups, with T in its center, and a free abelian group $\bar{M}$ amalgamated to T over the finite index subgroup M at each of the other vertices.

Observation 5.2.2. In case M is pegged, its peg is a peg of $\bar{M}$.

Proof. All we have to show is that any peg p in T has no proper roots in S either. If some non-trivial element of $\bar{M}$ could be conjugated to a rigid vertex group of the decomposition in $(\Delta\mathcal{T})^{\lambda}$ in which p is hyperbolic, then the same would be true for some finite power in M of said element, contradicting choice of p. In general, suppose we are given a reduced simplicial G-tree T in which no non-trivial element fixes an infinite axis and $H \leqslant G$, with minimal tree $T_H \subset T$, so that the natural map from $H\backslash T_H$ to $G\backslash T$ is injective on edges. Then any hyperbolic element $h \in H$ lacking a proper root in H also lacks one in G. Indeed, suppose this is not the case, so that there some $g \in G$ and $n \in \mathbb{N}$ such that $g^n = h$.

We claim that g preserves T_H. By assumption action of g preserves the orbits of edges of $A = Ax(h)$ by the action of G. If we pick $e \in A$ this implies the existence of some $k \in H$ such that $k \cdot e = g \cdot e$. Since both e and $g \cdot e$ point in the same direction along A, necessarily $Ax(k) = A$. It follows that $k^{-1}g$ fixes A, which by assumption can only be the case if $g = k$. □

We say that a morphism f which factors through $\mathcal{T}_A$ is primitive if $f(M) = Z_{\mathbb{F}}(M)$ for any $M \in \mathcal{MA}$. We say that f extends to some closure $\mathcal{S}_A$ of $\mathcal{T}_A$ if it extends to a morphism from S_A to F_A factoring through $\mathcal{S}_A$.

Definition 5.2.2.1. By a covering system of closures of a restricted π-tower $\mathcal{T}^J_A$ of [(1)] we intend a finite family $\mathcal{CL}$ of closures of $\mathcal{T}^J_A$ with the property that every morphism $f : \mathsf{T}_A \to \mathsf{F}_A$ factoring through $\mathcal{T}_A$ extends to some $\bar{f} : \mathsf{S}^r_A \to \mathsf{F}_A$ factoring through $\mathcal{S}_A$ for some $\mathcal{S}_A \in \mathcal{CL}$.

This extends to a homomorphism $\bar{f} : \mathsf{S}_A \to \mathsf{H}_A$ if and only if for any $M \in \mathcal{FHM}$ the restriction $f_M := \hat{f}\restriction_{\mathsf{M}}$ extends to $f_{\bar{M}} : \hat{\bar{\mathsf{M}}} \to \mathsf{H}$ (notice that these, and thus $\bar{f}$, are unique). The necessary and sufficient condition for this extension to factor through $\mathcal{S}_A$ is precisely that:

i) $f_{\bar{M}}(\bar{M}) = f_{\bar{M}}(pg_M)$ for any $M \in \mathcal{MA}$

ii) $f_{\bar{M}}(\bar{M}) \leqslant f_{\bar{M}}(bot(\bar{M}))$ for any $M \in \mathcal{FHM}$

An obvious but important observation that can be drawn from this is that if some primitive f as above extends to $\mathcal{S}_A$ then any f' such that the maps $f_M : M \to f(M)$ and $f'_M : M \to f'(M)$ differ by a π-automorphism of their images will also extend to $\mathcal{S}_A$. In particular, only primitive morphisms are needed to test whether a system of closures is covering or not.

It can come in handy to modify $\mathcal{T}$ by turning all maximal non-cyclic abelian subgroups into pegged ones, as described below. Let $\hat{T}$ be the group obtained by amalgamating $\hat{M} := M \oplus Z_M$ over M to T for each $M \in \mathcal{FHM}$, where Z_M is isomorphic to $\mathbb{Z}$. For each $M \in \mathcal{MA} :=$

(1) A covering closure in Sela's terminology.

$\mathcal{PM} \cup \mathcal{FHM}$ let pg_M be either a peg of M in $\mathcal{T}$ in case $M \in \mathcal{PM}$ or a fixed generator of Z_M otherwise. Denote by $pgq(T)$ the collection all $q \in Q^{\mathcal{FHM}}$ with the property $\pi(Z_M) = \pi(M)$ for any $M \in \mathcal{FHM}$. For any $pgq(T)$ denote by $\mathsf{T^q}_A$ the restricted π-group with underlying group $\hat{T}$ in which $\pi(pg_M) = q_M$ and the rest of the structure stays the same. There is an obvious restricted π-tower structure $(\mathcal{T}^q)_A^{(J,r)}$, obtained by enlarging the appropriate vertex groups of $(\Delta\mathcal{T})^\lambda$ for $\lambda \in J$. At the bottom of the new tower, for each $\lambda \in lv(J)$ for which $\mathsf{T}^\lambda = bot(M)$ for some $M \in \mathcal{FHM}$ we find a retraction $\eta_\lambda^{T^q}$ from $Z_M \oplus bot(M)$ to $\mathsf{Z_M}$ instead. Here Z_M is the unique vertex group of an obvious (degenerate) decomposition, making pg_M a peg of $\hat{M}$. We will refer to any such $\mathsf{T^q}_A$ as a pegging of T_A. For any morphism f which factors through $\mathcal{T}$ there is some $q \in pgq(M)$ such that f admits a 'good' extension $\hat{f}$ to $\mathsf{T^q}_A$, i.e., one for which $\hat{f} : \mathsf{T^q}_A \to \mathsf{G}_A$ factors through $\mathsf{T^q}_A$ and with the additional property that $\hat{f}(M) = \hat{f}(\langle pg_M \rangle)$ in case $M \in \mathcal{FHM}$.

Definition 5.2.2.2. Let K be a positive integer divided by $|\pi(M)|$ for any $M \in \mathcal{MA}$. An atomic K-congruence condition on $\mathcal{T}_A$ is given by a tuple $\mathcal{C} = (q, (\xi_{PM})_{PM \in \mathcal{MA}})$, where $q \in pgq(T)$ and for each $M \in \mathcal{MA}$ the map ξ_M is a homomorphism from $\hat{M}$ to $\mathbb{Z}/N\mathbb{Z}$. We say that a primitive f factoring through $\mathcal{T}_A$ satisfies $\mathcal{C}$ if there is an extension $\hat{f} : \mathsf{T^q}_A \to \mathsf{H}_A$ of f such that $\hat{f}(M) = \hat{f}(\langle pg_M \rangle)$ for all $M \in \mathcal{FHM}$ and for all $M \in \mathcal{MA}$ and $m \in M$ equality $\hat{f}(m) = f(pg_M)^l$ holds for some $l \in \xi(m) \in \mathbb{Z}/K\mathbb{Z}$. A K-congruency condition is a union of K-congruency conditions.

The intersection of a (q, K) and a (q, K') atomic congruency conditions is a (possibly empty) $lcm(K, K')$ atomic congruency condition and that the complement of a congruency condition is again a congruency condition. Note that the condition might be empty even being satisfied by some non-primitive f. Our goal now is to show the following:

Lemma 5.2.3. *Given a finite family $\mathcal{CL}$ of closures of a restricted π-tower $\mathcal{T}_A$, a primitive morphism $f : \mathsf{T}_A \to \mathsf{H}_A$ factoring through $\mathcal{T}_A$ extends to one of the resolutions in $\mathcal{CL}$ precisely in case it satisfies certain congruence condition.*

Suppose we are given a family $\mathcal{F}$ of sequences of primitive maps factoring through $\mathcal{T}_A$ We say that $\mathcal{F}$ is congruence complete if and only if for any non-empty congruency condition $\mathcal{C}$ there is $(f_n)_n \in \mathcal{F}$ such that f_n satisfying $\mathcal{C}$ for n big enough.

Corollary 5.2.4. *Let $\mathcal{T}_A$ be a restricted π-tower and $\mathcal{F}$ a congruency complete family of sequences of morphisms from T_A to F_A factoring through $\mathcal{T}_A$. Then any finite family $\mathcal{CL}$ of closures of $\mathcal{T}_A$ with the property that f_n eventually extends to some $\mathcal{S}_A$ in $\mathcal{CL}$ for any $(f_n)_n \in \mathcal{F}$ is a covering system of closures.*

Proof. (of 5.2.3) Fix $\mathcal{S}_A \in \mathcal{CL}$ and let f be primitive factoring through $\mathcal{T}_A$ and conforming to some $q \in Q$. Recall the discussion after definition 5.2.2.1. Since M has finite index in $\bar{M}$ and pg_M is primitive in $\bar{M}$ in case $M \in \mathcal{PM}$, for $M \in \mathcal{MA}$, the basic theory $\mathbb{Z}$-modules tells us that there are bases $\{pg_M, m_j\}_{j=1}^r$ of $\hat{M}$ and $\{pg_M, \bar{m}_j\}_{j=1}^r$ of $\hat{\bar{M}}$ and positive integers k_j for $1 \leqslant j \leqslant r$ such that $m_j = \bar{m}_j^{k_j}$. Let $K_M = lcm\{k_j|\pi(M)| \mid 1 \leqslant j \leqslant r\}$. Consider $M \in \mathcal{PM}$ first. For each $1 \leqslant j \leqslant r$, let e_j the unique integer such that $f(m_j) = z^{e_j}$, where $z = f(pg_M)$. Clearly, whether f_M has an extension to $\bar{M}$ satisfying (i) above depends only on whether each $f(m_j)$ has a k_j-th root in $\langle z \rangle$, hence only on the tuple $([e_j])_{j=1}^r \in (\mathbb{Z}/N\mathbb{Z})^r$. The same is true fir the question of whether this extension is in fact a π-morphism, i.e., whether the k_j-th root of $f(m_j)$ is mapped to $\pi(\bar{m}_j)$ by π, since $|\pi(M)||K$.

In the case of $M \in \mathcal{FHM}$, it is a priory not enough for us that $f(m_j)$ has a k_j-th root, in order for (ii) to be satisfied we need that $f(m_j) = f(n_j)$ for some $n_j \in bot(M)$ which itself has a k_j-th root in $\bar{M}$. This follows from the fact that $Z(f_M(M)) = \bar{f}_M(\bar{M}) = f_M(M)$. $\square$

5.3 Completions

Let $\mathcal{R}_A^J$ be a graded well-separated π-resolution for which the π-group at any leaf is a restricted π-limit group. We will assume that all the $(\Delta\mathcal{R})^\lambda$ are normalized, as defined earlier (if not, the construction can be done using their normalization instead). By adding edges to groups of surface type (possibly making the tree non-minimal) we can guarantee that

In this section we describe the construction of a so called completion of aa well-separated π-resolution $\mathcal{R}$: this will be a restricted (graded) π-tower $\mathcal{S}_A^{(J,r)}$ over the same index set, together with certain compatibility data. The standard construction for plain groups can be found in [Sel03]. First of all, for each $\lambda \in J$ we have an injective morphism $\iota_\lambda : \mathsf{R}_A^\lambda \to \mathsf{S}_A^\lambda$. This is an isomorphism in case $\lambda \in lv(J)$ or $(\Delta\mathcal{R})^\lambda$ consists of a single non-rigid type vertex (so the associated group is a closed surface group or a free non-cyclic abelian group.

If λ is not of free product type and $(\Delta\mathcal{R})^\lambda$ has at least one edge, Let $T_\lambda = \widetilde{(\Delta\mathcal{R})^\lambda}$ and $\hat{T}_\lambda^*$ the R^λ-tree obtained by normalizing $\widetilde{(\Delta\mathcal{S})^\lambda}$ sliding all edges between a rigid and a surface type vertex. Then there is also an isomorphism of geometric abelian trees $\Theta_\lambda\ T_\lambda^f$ (as described in the previous subsection) and some subtree of $\hat{T}_\lambda^*$ such that:

i) $\eta_{\mathcal{S}}^\lambda \circ \iota_\lambda = *_{\mu<\lambda}\iota_\mu \circ \eta_{\mathcal{R}}^\lambda$.

ii) For each concretion $(f_\lambda, \tau_\mu)_{\lambda\in J,\mu\in\hat{J}}$ of $\mathcal{R}_A$, there is a concretion $(g_\lambda, \tau_\mu)_{\lambda\in J,\mu\in\hat{J}}$ such that $f_\lambda = g_\lambda \circ \iota_\lambda$ for each $\lambda \in J$.

iii) Θ_λ is equivariant with respect to ι_λ, i.e., for each $g \in Sl$ and $x \in T_\lambda^f$ we have $\Theta_\lambda(g\cdot x) = \iota_\lambda(g) \cdot \Theta_\lambda(x)$.

iv) Each abelian vertex group of T_λ is mapped by Θ into an abelian group of $\hat{T}_\lambda^*$.

v) The star around each surface type vertex of T_λ is mapped isomorphically onto the star of a surface type vertex of $\hat{T}_\lambda^*$.

We construct $QRS{\restriction}_\lambda$ and the embedding ι_λ by induction on $\leqslant$. For $\lambda \in lv(J)$ let $Sl = Rl$ and ι_λ the idenity. Now, suppose that $QRS{\restriction}_\lambda$ is defined for $\mu > \lambda$ and that same is true for the embeddings $\iota(\mu)$. There are two cases to consider. If λ is of free product type, then we simply let S^λ be $*_{\lambda>\mu}\mathsf{S}^\mu$ and $\eta_{\mathcal{S}}^\lambda$ the map $\amalg_{\mu<\lambda}\iota_\mu$. In the next section we will deal with the case in which $(\Delta\mathcal{R})^\lambda$ is non-tivial and $\eta_{\mathcal{R}}^\lambda$ non-injective. All along the way we will use the fact that any towers having restricted π-limit groups at its leaves is itself a limit group.

5.3.1 Floor case

In order to ease the notation, we will denote $\amalg_{\mu<\lambda}\iota_\mu$ by κ, $\theta = \kappa \circ \phi$ and by ι the map ι_λ to be constructed. In the same fashion, we will rename η_R^λ as ϕ, the new map $\eta_{\lambda,S}^{\lambda,S}$ to be constructed as η and Θ_λ as Θ. Let also $\mathsf{G} := \mathsf{R}^\lambda$ and $\hat{\mathsf{G}}$ the π-group G^λ. The geometric abelian decomposition $(\Delta\mathcal{R})^\lambda$ puts the group R^λ into isomorphism with some $\pi_1(\Delta, Z)$. If this decomposition consists of a single isolated vertex of surface of abelian type we then can just set $\mathsf{S}^\lambda = \mathsf{R}^\lambda$.

Let $X = (V, E)$ be the underlying graph of Δ. Set also $\mathsf{H}_\mu = \mathsf{S}^\mu$ and $\mathsf{H} = *_{\lambda>\mu}\mathsf{S}^\mu$. We now describe the construction of a geometric abelian decomposition $\hat{\Delta}$, with underlying graph $\hat{X} = (\hat{V}, \hat{E})$. Its set of rigid vertex groups wil be precisely $\{H_\mu\}_{\mu\in nDFT(\lambda)}$. We also define a π-structure to each vertex group a $\hat{\Delta}_v$. Later in the subsection we specify a maximal tree $\tilde{Z}$ in $\hat{X}$ and show that the partial π-structures on vertex groups extend to a π-structure on $\hat{G} := \pi_1(\hat{\Delta}, \hat{Z})$, which we take as Sl, and a morphism from it to H. As $(\Delta\mathcal{S})^\lambda$ we take the decomposition as the graph of groups Δ above. We start the construction of a graph of groups Λ with underlying graph

$$(W,F) = (\{v_\mu, \hat{u}, \hat{w}\}_{\mu \in NDCt(\lambda), u \in V_a, w \in V_s}, \{e_v, \overline{e_v}, \hat{f}, \overline{\hat{f}}\}_{v \in V_a, \alpha(f) \in V_s})$$

To begin with, set $\Lambda_\mu = \mathsf{S}^\mu$. For any $w \in V_a$ let $\mathsf{P_v} = (Per(v), \pi\restriction)$ and chose some $\mathsf{K_v}$ such that $(\Delta_w, \pi\restriction) = \mathsf{P_v} \oplus \mathsf{K_v}$. Since ι is injective on P_v the morphism η injects P into $\mathsf{H}_{\mu_\mathsf{w}}$ for some rigid vertex $w \in V$. In correspondence, the new graph includes vertex $\hat{w}$ and a single edge e_w from $\hat{w}$ to v_{μ_w}. Let $O(w)$ denote the set of those edges e from w to a vertex distinct from the first edge in the unique path in Z from w to v_0. Let $Q_w = C_{H_{\mu_w}}(\theta(\Delta_w))$. As $\Lambda_{\hat{w}}$ we take the group $Q_w \oplus K'_w \oplus (\langle \bar{t}_e \rangle_{O(w)})_{ab}$, where $\mathsf{K'_w}$ is a copy of K_w isomorphic to it via some α_w. We endow $\Lambda_{\hat{w}}$ with a π-structure restricting to $\pi_{H_{\mu_w}}$ and to $\pi_{K'_w}$ respectively on the first two direct summands in the expression above and sending each of the generators $\bar{t}_e$ to the identity. As Λ_{e_w} we take the leftmost summand in the previous decomposition and as i_{e_w} simply the identity map. Our ζ_v will be the π-isomorphism $\theta\restriction_{K_w} \oplus \alpha_v$ between $(\Delta_w \pi\restriction_{\Delta_w})$ and the subgroup of $\mathsf{\Lambda_w}$ given by $\mathsf{P_w} \oplus \mathsf{K'_w}$.

To continue, we associate to each $u \in V_s$ a new vertex $\hat{u}$ and as $\mathsf{\Lambda_{\hat{u}}}$ we take some copy of $\mathsf{\Delta_u}$, isomorphic to it via some isomorphism ζ_u. For each e originating at u there is $\mu_e \prec \lambda$ such that $\eta(\Delta_{\bar{e}}) \subset H_{\mu_e}$. To e we associate some $\hat{e}$, originating at u, with $\alpha\overline{\hat{e}} = v_{\mu_u}$. We let $\Delta_{\hat{e}} = \zeta_u(\Delta_e)$. As $i_{\hat{e}}$ we take the map $\theta \circ i_e \circ \zeta_u^{-1}$ to Δ_e. The fact that Δ is a commutative transitive graph of groups does not guarantee that Λ is one as well, as the images of two elements in non-conjugate maximal abelian subgroups might as well commute. By induction on the size of J, one can assume, however, that the H_μ are limit groups and, therefore, commutative transitive.

Dealing with this issue involves three steps. First of all, for each $w \in V_a$ we chose $\delta_v \in H_{\mu_w}$ so that $Q_w^{\delta_w} = Q_{w'}^{\delta_{w'}}$ whenever Q_w and $Q_{w'}$ are conjugate (which implies $\mu_w = \mu_{w'}$). Then replace the vertex group Λ_w by its formal conjugate $\Lambda'_w = \Lambda_w^{\delta_w} \subset \pi_1(\Lambda\restriction_{v_{\mu_w}, e_w, \hat{w}}, \{e_w\})$. By this we mean that Λ'_w inherits the π-structure $\pi(\delta_w)$-conjugate to that on Λ_w as well. The new left summand is still a subset of H_{μ_w} and as i_{e_w} we can take the inclusion again. The operation does not affect the partial πisomorphism type of the fundamental group. Accordingly, we can see ζ_w as an injective π-morphism of Δ_w into $(\Lambda'_w)^{\delta_w^{-1}}$, coinciding with θ on K_w.

After the operation is completed, we can assume that for any two $w, w' \in V_a$ if two conjugates of the images of Λ_{e_w} and $\Lambda_{e_{w'}}$ in H_μ for $\mu = \mu_w = \mu_{w'}$ commute, then they coincide, in which case we say that w and w' are equivalent. For each μ_w fold together all the edges from v_μ ending in the same equivalence class $[w]$. The vertex group associated to any of the new vertices $\widehat{[w]}$ is the amalgamated product of $\Lambda_{w'}$ for $w' \sim w$ over the common peripheral subgroup. At this point we replace any of those groups by their abelianization:

$$P_{[w]} \oplus (\bigoplus_{w \sim w'} K_{w'}) \oplus (\bigoplus_{\substack{e \in O(w'),\\ w' \sim w}} s_e)$$

Here s_e stands for the image of the generator $\bar{t}_e^{\delta_w}$ of Λ_w. We keep the notation $\bar{t}_e$ for $s_e^{\delta_w^{-1}}$. Notice that since peripheral subgroups are mapped injectively by this quotient, the what we are left with is still a graph of groups. The fact that the quotient group inherits a π-structure from the amalgamated free product follows from the following:

Claim 5.3.1. *Let $w, w' \in V_a$ such that $w \sim w'$. Then $[\pi(\Delta_w)^{\pi(\delta_w)}, \pi(\Delta_{w'})^{\phi(\delta_w)}] = 1$.*

Proof. Let $\phi = \iota \circ \eta$. Observe that $\pi(\Delta_u) = (\pi \circ \theta)(\Delta_u) \subset C_{H_{\mu_w}}(P_u)$ for $u \in \{w, w'\}$. Now, we know that H_{μ_w} is a π-limit group, hence CSA. Since $P_w^{\delta_w}$ and $P_{w'}^{\delta_{w'}}$ are contained in some maximal abelian subgroup, the same is true for $\theta(\Delta_w^{\delta_w})$ and $\theta(\Delta_{w'}^{\delta_{w'}})$. Hence

$$1 = [\pi(\phi(\Delta_{w'})^{\delta_{w'}}), \pi(\phi(\Delta_w)^{\delta_w})] = [\pi(\phi(\Delta_{w'}))^{\pi(\delta_{w'})}, \pi(\phi(\Delta_w))^{\pi(\delta_w)}] = [\pi(\Delta_{w'})^{\pi(\delta_{w'})}, \pi(\Delta_w)^{\pi(\delta_w)}]$$

□

The previous embedding ζ_w corresponds now to one of Δ_w into $\hat{\Delta}_{[\hat{w}]}^{\delta_w^{-1}}$, for which we employ the same term. Extend the notation by letting $\zeta_v = \theta\restriction_{\Delta_v} : \Delta_v \to \hat{\Delta}_{\mu_v}$ for $v \in V_r$.

Remark 5.3.2. Let $e \in E$ and $v = \alpha(e)$. Then $\zeta_v\restriction_{\Delta_e}$ is equal to:

i) $\theta\restriction_{\Delta_e}$ in case $v \in V_r \cup V_a$

ii) $inn_{t_{\hat{e}}^{-1}} \circ \theta \circ inn_{t_e}\restriction_{\Delta_e} = inn_{\theta(t_e)t_{\hat{e}}^{-1}} \circ \theta\restriction_{\Delta_e}$ in case $v \in V_s$

We first specify the restriction of η to each vertex group of $\hat{\Delta}$. This is the identity on every $\hat{\Delta}_{v_{\mu_u}}$, as required by the definition of a tower. On $\hat{\Delta}_{[\hat{w}]}$ with $w \in V_a$, we let η kill any of the generators of the form $\tilde{t_e}$ for $w' \sim w$ and $e \in O(w)$. On the direct summand which is the image of $\mathsf{K}_{\mathsf{w}'}$ in $\hat{\Delta}_{[\hat{w}]}$ we let the restriction of η coincide with $inn_{\delta_v} \circ \theta \circ \zeta_v^{-1} \circ inn_{\delta_v^{-1}}$. The map $\eta\restriction_{\Delta_{[\hat{w}]}}$ is a π-morphism since both θ and ζ_v are. In case $v \in V_s$, we let $\eta\restriction_{\hat{\Delta}_{\hat{v}}} = \theta \circ \zeta_v^{-1}$. This, again, preserves the π-structure. The identity $\eta \circ \zeta_v = \theta$ clearly holds for every $v \in V$.

Notice that removing the set $\mathcal{O}$ of all those $\hat{e}$ associated with edges $e \in E\backslash Z$ does not disconnect $\hat{X}$. This means it is possible to chose a maximal subtree $\hat{Z}$ of $\hat{X}$ which does not intersect $\mathcal{O}$. For each $v \in V$ let $\Phi(v) \in \hat{V}$ be either:

- The unique v_μ such that $\eta(v) \subset G_\mu$
- $\hat{[v]}$ in case $v \in V_a$
- $\hat{v}$ in case $v \in V_s$

Given edges $e \in E\backslash Z$ and $f \in \hat{E}\backslash\hat{Z}$ denote by $t_e \in \pi_1(\Delta, Z)$ and $\hat{t}_f$ the Bass-Serre elements associated to e and f in the respective graph of groups.

We now have to check that both the map η and the partial π-structure defined on vertex groups extend to the entire $\hat{G}$. In view of the fact that $\eta\restriction_H = Id\restriction_H$ and that every edge in $\hat{E}$ has an endpoint in some of the v_μ, this is equivalent to the existence for all $e \in \hat{E}$ of $q_e = \eta(t_e) \in \hat{G}$ such that the equality

$$i_e = inn_{q_e} \circ \eta\restriction_{\hat{\Delta}_e} : \hat{\Delta}_e \to H_{\mu_e}$$

is satisfied for all $e \in \hat{X}\backslash\hat{Z}$. Note that any such edge is adjacent to some surface type vertex. So suppose that $e = \hat{f}$ for some f originating at $w \in V_s$. Then:

$$i_e = \theta \circ i_f \circ \zeta_w^{-1} = \theta \circ inn_{t_f} \circ \zeta_w^{-1} = inn_{\theta(t_f)} \circ \theta \circ \zeta_w^{-1}$$

We also have $\eta\restriction_{\Delta_e} = \theta \circ \zeta_w^{-1}$. So letting $q_e = \theta(t_f)$ does the job. Note that in virtue of the way in which we chose $\hat{Z}$, in case $t_e = 1$ we have $q_e = 1$, while $t_f = 1$ in this case by virtue of the choice of $\hat{Z}$. In the same fashion, setting $\pi(t_e) = \pi(\theta(t_f))$ extends the π-structure on vertex groups to a global one compatible with the previous collapse map.

5.3.2 The morphism ι

Chose a reference vertex $v_0 \in V_r \cup V_a$. If the group A of constants is not-trivial, then we can assume that $A \subset \Delta_{v_0}$. The construction depends on the assignment, to each $v \in V$ of a conjugating element $\gamma_v \in \pi_1(\hat{\Delta}, \hat{Z})$. With this information one defines a morphism ι_v from each Δ_v into $\hat{G}$, given by $inn_{\gamma_v} \circ \zeta_v$.

The choice of γ_v is completed with the assignment of an element $r_e \in \hat{G}$ for each $e \in E$ in such a way that the map sending t_e to r_e and restricting to ζ_v on Δ_v extends to a morphism ι from G_A to $\hat{\mathsf{G}}_A$ with the properties sought for. In particular $r_e = 1$ for $e \in Z$. The condition for an extension which is a group homomorphism to exist is that the equation

$$inn_{r_e} \circ inn_{\gamma_v} \circ \zeta_v \restriction_{\Delta_e} = inn_{\gamma_w} \circ \zeta_w inn_{t_e} \restriction_{\Delta_e}$$

holds for all $e \in E$ from v to w (recall that t_e acts like i_e on Δ_e seen as a subgroup of G). Equivalently:

$$inn_{\gamma_v r_e \gamma_w^{-1}} \zeta_v \restriction_{\Delta_e} = \zeta_w \circ inn_{t_e} \restriction_{\Delta_e}$$

In order to define γ_v, consider the unique simple path $v_0, e_0, v_1, e_1, \cdots v_k, e_k, v_{k+1} = v$ contained in Z. We let $\gamma_v = g_k g_{k-1} \cdots g_0$, where g_j is equal to:

i) $\bar{t}_{e_j} \in \hat{\Delta}_{\hat{v_j}}^{\delta_j^{-1}}$ in case $v_j \in V_a$ and $v_{j+1} \in V_r$

ii) $t_{\hat{e_j}}^{-1}$ in case $V_s \cap \{v_j, v_{j+1}\} \neq \varnothing$ (possibly 1)

iii) 1 in case $v_j \in V_r$ and $v_{j+1} \in V_a$

As for r_e, it is enough to deal with the case in which $v = \alpha(e) \in V_a$. If $w = \omega(e)$, then we let r_e be equal to:

i) $\gamma_v^{-1} \bar{t}_e \theta(t_e) \gamma_w$ in case $w \in V_r$

ii) $\gamma_v^{-1} t_{\hat{e}} \gamma_w$ in case $w \in V_s$

For $e \in E$ from a vertex v to a vertex w we distinguish several cases:

i) None of w, v belongs to V_s. In this case we have:

$$\begin{aligned} \zeta_w \circ inn_{t_e} \restriction_{\Delta_e} &= \theta \circ inn_{t_e} \restriction_{\Delta_e} = \\ = inn_{\theta(t_e)} \circ \theta \restriction_{\Delta_e} &= inn_{\gamma_v r_e \gamma_w^{-1}} \circ \zeta_v \restriction_{\Delta_e} \end{aligned}$$

ii) One of w, v belongs to V_s, let's say w does. Then $\zeta_v \restriction_D = \theta \restriction_D$ while $\zeta_w \restriction_{\Delta_{\overline{e}}} = inn_{\theta(t_e^{-1} t_{\hat{e}})} \theta \restriction_{\Delta_{\overline{e}}}$ and:

$$\begin{aligned} \zeta_w \circ inn_{t_e} \restriction_{\Delta_e} &= inn_{\theta(t_e^{-1}) t_{\hat{e}}} \circ \theta inn_{t_e} \restriction_{\Delta_e} = \\ = inn_{\theta(t_e^{-1}) t_{\hat{e}}} \circ inn_{\theta(t_e)} \theta \restriction_{\Delta_e} &= inn_{t_{\hat{e}}} \circ \theta = \\ &= inn_{\gamma_v r_e \gamma_w^{-1}} \circ \theta \restriction_{\Delta_e} = \zeta_w \restriction_{\Delta_e} \end{aligned}$$

We note a couple of properties of the choice of γ_v and their implications for ι.

i) $\gamma_{v_0} = 1$ imples that ι is a restricted morphism,

ii) $\pi(\gamma_v) = 1$ for any v (recall that $\pi(t_{\hat{e}}) = 1$ in case $e \in Z$ but $\hat{e} \notin Z$), which implies the restriction of ι to any vertex group is, just as ζ_v is, a π-morphism.

iii) $\eta(\gamma_v) = 1$ for all $v \in V$, as $\eta(t_{\hat{e_j}}) = 1$ always in case (ii) of the definition of g_j. This implies that $\eta \circ \iota \restriction_{\Delta_v} = \theta \restriction_{\Delta_v}$. Since also $\eta(\iota(t_e)) = \theta(t_e)$ for each $e \in E$ (as $\theta(t_{\hat{e}}) = t_e$) we deduce that $\eta \circ \iota = \theta$.

5.3.3 The embedding Ψ, injectvity of ι

The following will use the assumption that $\hat{G}$ is a limit group so, in particular, CSA. Recall Z lifts to a subtree $\tilde{Z}$ of the covering G-tree $T = (\tilde{\Delta, Z})$ and that if we let $\tilde{v}$ and $\tilde{e}$ stand for the images of v and e respectively, then $Stab(T; \tilde{v}) = \Delta_v$ and $Stab(T; \tilde{e}) = \Delta_e$ for all $v \in V$ and $e \in E$. Additionally, any $e \in E \backslash Z$ can be lifted to $\tilde{e} \in ET$ in such a way that if we let the tuple $(\hat{Z}, \hat{Z} \cup \{\tilde{e}\}_{e \in E}, (g_{\tilde{e}})_{e \in E \backslash E})$ is a presentation of the action of G on T, where $g_{\tilde{e}} = t_e$ (in particular $\alpha(\tilde{e})\tilde{Z}$ for such e). The graph of groups $\hat{\Delta}$ is not normalized, as vertices of surface type are connected to vertices of rigid type. However, for each $e \in E$, ending in some $v \in V_s$, However, for any edge $e \in \hat{T}$ originating in a vertex of surface type there is a unique edge f with $\alpha(\overline{f})$ of abelian type such that $Stab(e) \subset Stab(f)$. Sliding each such e over the corresponding f yields a $\hat{G}$-tree $\hat{T}^*$ which is normalized. We denote by e^* the edge of $\hat{T}^*$ resulting from the slide of the edge e of $\hat{T}$. Any other edge or vertex of the tree $\hat{T}^*$ can be thought of as identical to the corresponding ones in $\hat{T}$. We will define a map Ψ from VT to $\hat{T}^*$ and then prove it extends to a cellular map from T^f to $\hat{T}*$ for which we keep the same name mapping the star of any vertex v injectively into the one around $\Psi(v)$, bijectively in case v is a surface-type vertex. At that point one can apply the following easy fact:

Fact 5.3.3. *Any locally cellular map $f : T \to T'$ between G-trees injective on stars of vertices is an isomorphic embedding.*

From this we will prove the map Θ to be a graph embedding.

Corollary 5.3.4. *The morphism ι is injective.*

Proof. If $g \in G \backslash \{1\}$ does not fix T^f, the same is true for $\iota(g)$ and $\hat{T}*$, so $\iota(g) \neq 1$. There might be non-trivial elements $g \in G$ acting trivially on T_v. Since the restriction of ι to any stabilizer of T^f is injective, though, in this case $\iota(g) \neq 1$ as well. □

of the existence of Ψ. We first define Ψ on $\tilde{V} = V$ by letting $\Psi(\tilde{v}) = \gamma_v^{-1} \cdot \tilde{\Phi(v)}$ for $v \in V$. After that we extend Ψ equivariantly to the whole VT. That is, we send any $v \in VT$ of the form $g \cdot v_0$ for $v_0 \in \tilde{Z}$ to $\iota(g) \cdot \Psi(v_0)$. This is well defined since $\iota(\Delta_v)$ stabilizes $\Psi(v)$ for $v \in \tilde{V}$.

Claim 5.3.5. *Given $e \in T$ the vertices $\Psi(\alpha(e))$ and $\Psi(\omega(e))$ are adjacent in $\hat{T}^*$*

Proof. Let $v = \alpha(e)$ and $w = \omega(e)$. The map Ψ preserves vertex type, hence $v' := \Psi(v)$ and $w' := \Psi(w)$ are two vertices of distinct type (one of them abelian). The fact that $\{1\} \neq \iota(Stab_T(v) \cap Stab_T(w)) \subset Stab_{\hat{G}}(v') \cap Stab_{\hat{G}}(w')$ implies v' and w' are adjacent. Indeed, otherwise some diameter 2 segment of $\hat{T}^*$ centered at a vertex of non-abelian type would be fixed by a non-trivial element, which is impossible, since $\hat{T}^*$ is in normal form. □

For any $e \in ET$ we let $\Psi(e)$ be the edge given by the previous claim. One needs to show that the map Ψ folds exactly as much as the quotient map from T onto T^f, namely:

Lemma 5.3.6. *Any two edges e, f of T with $\alpha(e) = \alpha(f) = v$ have the same image by Ψ only if:*

- *v is of non-rigid type and $e = f$*
- *v is of abelian type and e and f are in the same orbit under the action of the peripheral subgroup of $Stab(v)$*

Proof. Since the stabilizers of rigid vertices of $\hat{T}^*$ are CSA, the proof amounts to showing that for $u \in VT$ of non-rigid type:

i) For any $e \in ET$ the intersection $Stab_G(v) \cap Stab(\Psi(e), G)$ is equal to:

- $Per(T;v)$ if e is adjacent to $u \in VT_r$ and $v \in VT_a$
- $Stab_G(e)$ otherwise

ii) For any edges e, f of T originating at a common non-rigid vertex v and which are in different orbits by the action of $Stab(v)$, their images $\Psi(e)$ and $\Psi(f)$ lie in different orbits by the action of $Stab_G(v)$.

Local injectivity at rigid type vertex v depends on the fact that the geometric abelian decomposition Δ is normalized. Edges incident at v in different orbits will have mutually stabilizers non-conjugate in $Stab_G(v)$, so the same holds for their images by Ψ with respect to $\iota(Stab(v))$.

As usual, it is enough if we check the desired properties for the lifts of vertices in V. Start by considering $v \in V_a$. Then $\Psi(\tilde{v}) = g \cdot x$, where $x = [\hat{\tilde{v}}]$ and $g = \gamma_v^{-1}\delta_v$. Of course for any two w, w' adjacent to $\tilde{v}$ which are of different type their images are of different type as well, so they lie in different orbits; the same is true in case w, w' are both of surface type.

Suppose we are given an edge e from v to some vertex of surface type w. The definition of the attaching maps in $\hat{\Delta}$ reformulates as saying that $Stab(\hat{\tilde{e}}) = \theta(Stab_G(e)) \subset H_{\mu_v}$. Now, $\Psi(\tilde{e})$ is a translate of $\tilde{\hat{e}}^*$, so the group $Stab(\Psi(\tilde{e}), \hat{G})$ is a conjugate in $\hat{G}$ of $\theta(\Delta_e)$. Since $\iota(Stab(\tilde{e})) = \iota(\Delta_e)$ is of that form as well and $\hat{G}$ is CSA[2] . As ι is injective on Δ_v, this proves (i) for those edges ending in a surface type vertex.

A more careful analysis is needed to deal with adjacent rigid type vertices. To begin with, let RE_v be the set of all $e \in E$ from v to a vertex of rigid type and $RV_v = \{\omega(e)\}_{e \in RE_v}$. For each $u \in RV_v$ we have $\Phi(u) = v_{\mu_v}$, lifting to some $y_u \in Z$. For any $e = (v,u) \in E$, with $u \in RV_v$ the lift $\tilde{e} \in ET$ ends in the vertex $t_e\tilde{u}$. Now, $\Psi(\tilde{u}) = \gamma_u^{-1} \cdot y_u$, so by equivariance we get:

$$\Psi(\omega(\tilde{e})) = \iota(t_e)\gamma_u^{-1} \cdot y_u \tag{5.1}$$

Remember that $\Psi(\tilde{v}) = \gamma_v^{-1}\delta_v \cdot x$. To each $e = (v,u) \in RE_v$ we know there is an element h_e which is either 1 or of the form $\bar{t}_e^{\pm} \in \Delta_{[\tilde{v}]}^{\delta_v^{-1}}$ and:

- If $e \in Z$, then $\gamma_u = h_e\gamma_v$. So, as $t_e = 1$, the equation (5.1) yields $\Psi(\omega(\tilde{e})) = \gamma_v^{-1}h_e^{-1} \cdot y$
- If $e \notin Z$, then $\iota(t_e) = \gamma_v^{-1}h_e^{-1}\theta(t_e)\gamma_u$, so in this case we get

$$\Psi(\omega(\tilde{e})) = \gamma_v^{-1}h_e^{-1}\theta(t_e) \cdot y = \gamma_v^{-1}h_e^{-1} \cdot y_u$$

recall that $\theta(t_e)$ belongs H_{μ_v}.

Now, for any distinct $e, e' \in RE_v$ the element $h_e h_{e'}^{-1} \notin \iota(\Delta_v)^{\gamma_v^{-1}} = \zeta_v(\Delta_v)^{\delta_v^{-1}} \subseteq \Delta_{[\hat{w}]}^{\delta_v^{-1}}$, and therefore $\Psi(\tilde{e})$ and $\Psi(\tilde{e'})$ lie in different orbits with respect to the action of $\iota(Stab(\tilde{v})) = \iota(\Delta_v)$. Remember, moreover, that $\Delta_v = P_v \oplus K_v$, where P_v is $Per(v)$, and the map ζ_v has been constructed in such a way that it sends P_v to $\hat{\Delta}_{e_{[v]}}^{\delta_w^{-1}} \subset \hat{\Delta}_{[\hat{v}]}^{\delta_w^{-1}}$ and K_v into some complementary factor $\hat{\Delta}_{[\hat{v}]}^{\delta_w^{-1}}$. This clearly implies the property (i) for $\tilde{v}$. □

□

5.3.4 Factoring

We shall now prove that for any morphism $f : \mathsf{G}_A \to \mathsf{F}_A$ factoring through the π-resolution $\mathcal{R}_A^{(J,r)}$ there is one $\hat{f} : \mathsf{R}_A^{\lambda} \to \mathsf{G}_A$ factoring through $\mathcal{S}_A^{(J,r)}$ such that $f = \hat{f} \circ \iota$. We claim the statement reduces to the following lemma:

[2] It admits a commutative transitive geometric abelian decomposition with CSA rigid type vertex groups.

Lemma 5.3.7. *For each $\lambda \in \hat{J}$ (of non-free product type) and each $\sigma \in Mod((\Delta\mathcal{R})^{\lambda})$ there is $\hat{\sigma} \in Mod(PMod(GSla; A))$ such that*

$$\hat{\sigma} \circ \iota_{\lambda} = \iota_{\lambda} \circ \sigma \tag{5.2}$$

Indeed, in that case given a concretion $(f_{\lambda}, \sigma_{\mu})_{\lambda \in lv(J), \mu \in \hat{J}}$ of $\mathcal{R}_A$ let $(\hat{f}_{\lambda}, \hat{\sigma}_{\mu})_{\lambda \in J, \mu \in \hat{J}}$ be a concretion of $\mathcal{S}_A$ such that $\hat{f}_{\lambda} = f_{\lambda}$ for $\lambda \in lv(J)$. Given $\lambda \in J$ it is straightforward to prove by induction that $\hat{f}_{\lambda} \circ \iota_{\lambda} = f_{\lambda}$. This is trivial in case $j \in lv(J)$. Otherwise we have:

$$\begin{aligned} \hat{f}_{\lambda} \circ \iota_{\lambda} &= \amalg_{\mu<\lambda} \hat{f}_{\mu} \circ \eta^{\lambda,S}_{\lambda,S} \circ \hat{\sigma}_{\lambda} \circ \iota_{\lambda} = \\ = \amalg_{\mu<\lambda} \hat{f}_{\mu} \circ \eta^{\lambda,S}_{\lambda,S} \circ \iota_{\lambda} \circ \sigma_{\lambda} &= \amalg_{\mu<\lambda} \hat{f}_{\mu} \circ \amalg_{\mu<\lambda} \iota_{\mu} \circ \eta^{\lambda,R}_{\lambda,R} \circ \sigma_{\lambda} = \\ = \amalg_{\mu<\lambda} \hat{f}_{\mu} \circ \iota_{\mu} \circ \eta^{\lambda,R}_{\lambda,R} \circ \sigma_{\lambda} &= \amalg_{\mu<\lambda} f_{\mu} \circ \eta^{\lambda,R}_{\lambda,R} \circ \sigma_{\lambda} = f_{\lambda} \end{aligned}$$

Where the first equality follows from the properties satisfied any concretion of $\mathcal{S}_A$, the second from equality 5.2, the third from the properties of the family $(\iota_{\lambda})_{\lambda \in J}$, the fifth by the induction hypothesis and the last one by the equalities valid for any concretion of RRA.

of 5.3.7. First of all, notice that it is enough to prove the property for a set of generators of $PMod(\Delta)$. Indeed, suppose that extensions $\hat{\sigma}$ and $\hat{\tau}$ exist for $\sigma, \tau \in PMod(\Delta)$. Then $\iota \circ \sigma \circ \tau = \hat{\sigma} \circ \iota \circ \tau = \hat{\sigma} \circ \hat{\tau} \circ \iota$, so $\widehat{\sigma \circ \tau} = \hat{\sigma} \circ \hat{\tau}$ satisfies the condition Obviously, for inn_c we can take $\widehat{inn}_c = inn_{\iota(c)}$. Contrary to the definition, to make the notation lighter when denoting the generators of the π-modular group we will refer to edges and vertices of V rather than to their lifts in the fundamental trees T and $\hat{T}$. In virtue of 2.5.2 it is then enough to prove the lemma for σ of the following form:

i) A Dehn twist automorphism $\tau_{e,c}$ where $e \in E \backslash Z$ and $\tilde{e} = (\tilde{v}, \tilde{v'})$ points away from $\tilde{v_0}$, where $c \in \Delta_v$, one among $\{v', v\}$ is of rigid type and the other of abelian type.

ii) A natural extension $\bar{\rho}_{(c_e)_{\substack{\alpha(e)=v,\\ e \in E}}}$ of some $\rho \in Mod(\Delta_v)$ for some $v \in Z$ of surface type.

iii) A vertex group automorphism $\bar{\rho}_{(1_e)_{\substack{\alpha(e)=v,\\ e \in E}}}$ extending some $\rho \in Aut(\Delta_v)$ which fixes $Per(T; v)$, for some $v \in Z$ of abelian type.

Moreover, in the last two cases we can assume that $c_{e_0} = 1$ for the unique $e_0 \in \tilde{Z}$ originating at v and pointing towards v_0. Start with case (i) and let $x = [\hat{v}]$. We will assume v is of abelian type, the other one being idetical. Remember that $\hat{\Delta}_{[\hat{v}]}$ can be written as $P_{[\hat{v}]} \oplus \bigoplus_{u \sim v} L_u \oplus \bigoplus_{f \in O(v)} \langle \bar{t}_f \rangle$, where $L_u = (K'_v)^{\delta_u}$ for $\mathsf{K}'_\mathsf{v} \subset \hat{\mathsf{G}}$ isomorphic via ζ_u to K_v. Consider the automorphism ξ of $\hat{Delta}_x$ fixing the two leftmost direct summands in the expression above and sending $\bar{t}_e$ to $\zeta_v(c)^{\delta_v} \bar{t}_e$ and thus $\bar{t}_e$ to $\zeta_v(c) \bar{t}_e$. Let $\bar{\xi}$ be its extension to $\hat{G}$ with trivial twisting elements. Let us check that $\hat{\sigma} = \bar{\xi}$ serves as the appropriate lift. We first check the behaviour of $\hat{\sigma}$ on a vertex group Δ_u of Δ. Notice that in any case $\bar{\xi} \circ \zeta_u = \zeta_u$, since x has a single incident edge. There are two cases:

- If $\tilde{e} \notin [\tilde{v_0}, \tilde{u}]$, so that $\sigma \restriction_{\Delta_u} = Id_{\Delta_u}$, then γ_u is the product of elements of the form $\bar{t}_f$ for $f \neq e$ or of the form $t_{\hat{f}}$ for $f \in E \backslash Z$, hence $\bar{\xi}(\gamma_u) = \gamma_u$.

$$\begin{aligned} \bar{\xi} \circ \iota \restriction_{\Delta_u} &= \bar{\xi} \circ inn_{\gamma_u} \zeta_u = \\ = inn \bar{\xi}(\gamma_u) \circ \bar{\xi} \circ \zeta_u \circ \sigma \restriction_{\Delta_u} &= \iota \circ \sigma \restriction_{\Delta_u} \end{aligned}$$

- If $\tilde{e} \in [\tilde{v_0}, \tilde{u}]$, then $\sigma\restriction_{\Delta_u} = inn_c\restriction_{\Delta_u}$. Going back to the definition of γ_v, we can write $\gamma_v = g\bar{t}_e\gamma_v$ where γ_v and g are products of Bass-Serre elements of $\hat{\Delta}$ and elements of Δ_w for $w \neq u$ or elements of Δ_x in a direct summand of Δ_u fixed by σ. Hence

$$\bar{\xi}(\gamma_u) = g\zeta_v(c)\bar{t}_e\gamma_v = g\bar{t}_e\gamma_v\zeta_v(c)^{\zeta_v} = \gamma_u\iota(c)$$

  So that: .

$$\begin{aligned} \bar{\xi} \circ \iota\restriction_{\Delta_u} &= inn_{\gamma_u\iota(c)} \circ \bar{\xi} \circ \zeta_u = \\ = inn_{\iota(c)}inn_{\gamma_u} \circ \zeta_u &= inn_{\iota(c)} \circ \iota\restriction_{\Delta_u} = \\ &= \iota \circ inn_c\restriction_{\Delta_u} = \iota \circ \sigma\restriction_{\Delta_u} \end{aligned}$$

Let us now check the equality on $t_f \neq 1$, for $f \in E$ from an edge u to an edge w. If $e \neq f$, then $\sigma(t_e) = c_u^{-1}t_f c_w$, where the assignment $x \mapsto c_x \in \{1, c\}$ is as in the definition of a Dehn twist. For any $w \in V$ we know from the previous case that $\bar{\xi}(\gamma_w) = \gamma_w\iota(c_w)$. On the other hand $\iota(t_f)$ was defined as $\gamma_u^{-1}\epsilon_f\gamma_w$, where $\epsilon_f = \bar{t}_f\theta(t_f)$ in case $w \in V_r$, and $t_{\hat{f}}$ in case $w \in V_s$. Note that $\bar{\xi}(\epsilon_f) = \epsilon_f$.

$$\begin{aligned} \bar{\xi}(\iota(t_f)) &= \bar{\xi}(\gamma_u^{-1}\epsilon_f\gamma_w) = \\ &= \bar{\xi}(\gamma_u)^{-1}\bar{\xi}(\epsilon_f)\bar{\xi}(\gamma_w) = \\ &= \iota(c_u)^{-1}\gamma_u^{-1}\epsilon_f\gamma_w\iota(c_w) = \\ &= \iota(c_u)^{-1}\iota(t_f)\iota(c_w) = \\ &= \iota(c_u^{-1}t_f c_w) = \iota(\sigma(t_f)) \end{aligned}$$

as desired. If $e = f$, then $\sigma(t_e) = c^{-1}t_e$ and $u = v$. Also $c_u = c_w = 1$. Now $\iota(t_e) = \gamma_v^{-1}\bar{t}_e\theta(t_e)\gamma_w$, so that:

$$\begin{aligned} \bar{\xi}(\iota(t_e)) &== \bar{\xi}(\gamma_v^{-1}\bar{t}_e\theta(t_e)\gamma_w) = \\ &= \gamma_v^{-1}\bar{t}_e\theta(c)^{-1}\theta(t_e)\gamma_w = \\ &= \gamma_v^{-1}\bar{t}_e\zeta_v(c^{-1}t_e)\gamma_w = \\ &= \iota(c^{-1}t_e) = \iota(\sigma(t_e)) \end{aligned}$$

We now deal with case (ii). Here ρ extends an automorphism of Δ_v for some $v \in V$ of surface type. Let Or be the set of all $e \in E$ with $\alpha(e) = v$ and $\hat{Or} = \{\hat{e} : e \in Or\}$. As $\hat{\sigma}$ we take $\hat{\bar{\xi}} = \hat{\bar{\xi}}_{(\hat{c}_e)}$, where $\xi = \zeta_v \circ \rho \circ \zeta_v^{-1}$ and $\hat{c}_{\hat{e}} = \iota(c_e)^{\gamma_v^{-1}}$ for each $e \in \hat{Or}$. Extend the assignment $e \mapsto \hat{c}_e$ to $\hat{V}$ and the whole of E as described in definition 2.5.2 (by lifting to the fundamental domain). We start by showings the following:

Lemma 5.3.8. *The following equalities hold for $u \neq v$:*

$$\begin{aligned} \hat{c}_{\Phi(u)}\hat{\bar{\xi}}_{(\hat{c}_e)}(\gamma_u) &= \gamma_u\iota(c_u) \\ \hat{\bar{\xi}}_{(\hat{c}_e)}(\gamma_v) &= \gamma_v \end{aligned}$$

Proof. For $x \in V_a$ observe that $\hat{\sigma}(\bar{t}_e) = \bar{t}_e^{c_{[\hat{v}]}}$ for any of the generators $\bar{t}_e \in \hat{\Delta}_{[\hat{x}]}$, as δ_x belongs to a vertex v_μ for which $\hat{c}_{v_\mu} = c_{[\hat{x}]}$. Let $v_0 = undefined, u_m = u$ be the unique simple path in $\mathbb{Z}$ from v_0 to u. There are two possibilities:

i) In case $v \neq u_i$ for all i, let $0 \leqslant i_1 < i_2 < \cdots i_l \leqslant m$ enumerate the (possibly empty) set of those indices i for which e_i is incident to a rigid type vertex and $\hat{e_i} \notin Z$; extend the notation by letting $i_0 = 0$ and $i_{l+1} = m$. Then, for $0 \leqslant j \leqslant l$, the value $\hat{c}_{\Phi(u_i)}$ is constant, say equal to c_j for all $i \in (i_j, i_{j+1}]$. If we let $f_j = \hat{e_{i_j}}$, then of course

$$\bar{\hat{\xi}}_{(\hat{c}_e)}(t_{f_j}) = \hat{c}^{-1}_{\alpha(f_j)} t_{f_j} \hat{c}_{\omega(f_j)} = c_j^{-1} t_{f_j} c_{j+1}$$

Now, we know that $\gamma_u = a_l t_{f_l} a_{l-1} \cdots t_{f_1} a_0$, where a_j is the product of either elements in some Δ_x and $c_{\Phi(x)} = c_j$, or of the form $t_{\hat{f}}$, where $\Phi(\alpha(f)) = \Phi(\omega(f)) = c_j$. We conclude that $\bar{\hat{\xi}}_{(\hat{c}_e)}(a_j) = a_j^{c_j}$ for all j. Since $c_u = 1$ and $c_{v_0} = 1$ we have:

$$\begin{aligned}\hat{c}_{\Phi(u)} \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u) &= c_l (a_l^{c_l} (c_l^{-1} t_{f_l} c_{l-1}) a_{l-1}^{c_{l-1}} \cdots a_0^{c_0}) = \\ &= a_l t_{f_l} a_{l-1} \cdots t_{f_1} a_0 = \gamma_u = \gamma_u \iota(c_u)\end{aligned}$$

ii) In case $v = u_{i*}$ for some $i*$, this index unique. Observe also that for any $e \in E$ originating at v and such that $\hat{e} \notin \hat{Z}$, $\bar{\hat{\xi}}_{(\hat{c}_e)}(t_{\hat{e}}) = \hat{c}_{\hat{e}}^{-1} t_{\hat{e}}$. We can write γ_u as $g\gamma_v$, where g corresponds to the subpath $u_i, e_i \cdots u_m$. Now, $\bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_v) = \gamma_v$, as $\hat{c}_{\hat{e_i}} = \hat{c}_{u_i} = e_{i*-1} = 1$ for $i < i*$. Similarly, $\bar{\hat{\xi}}_{(\hat{c}_e)}(g) = \hat{c}^{-1}_{\Phi(u)} g \hat{c}_{\hat{e_u}}$, where e_u is the first edge in the path from v to u in Z. Since $\hat{c}_{\hat{e_u}} = \iota(c_{e_u})$, we get

$$\hat{c}_{\Phi(u)} \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u) = g \hat{c}_{\hat{e_u}} \gamma_u = g \iota(c_u)^{\gamma_v^{-1}} \gamma_v = g \gamma_v \iota(c_u) = \gamma_u \iota(c_u)$$

□

We can now check the equality on Δ_u, for any $v \neq u \in V$. Note, first of all, that $\sigma \circ \iota\restriction_{\Delta_u} = inn_{\iota(c_v)} \circ \iota\restriction_{\Delta_u}$. On the other hand, since $\zeta_u(\Delta_u) \subset \hat{\Delta}_{\Phi(\hat{u})}$, we get

$$\begin{aligned}\bar{\hat{\xi}}_{(\hat{c}_e)} \circ \iota\restriction_{\Delta_u} &= inn_{\bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u)} \circ \bar{\hat{\xi}}_{(\hat{c}_e)} \circ \zeta_u = \\ = inn_{\bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u)} \circ \bar{\hat{\xi}}_{(\hat{c}_e)} \circ \zeta_u &= inn_{\bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u)} \circ \hat{c}_{\Phi(u)} \circ \zeta_u = \\ &= inn_{\hat{c}_{\Phi(u)} \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u)} \circ \zeta_u = \\ &= inn_{\hat{c}_{\gamma_u} \iota(c_u)} \circ \zeta_u = \\ &= inn_{\iota(c_u)} \circ inn_{\hat{c}_{\gamma_u}} \circ \zeta_u = \\ &= inn_{\iota(c_u)} \circ \iota\restriction_{\Delta_u} = \iota \circ inn_{c_u}\restriction_{\Delta_u} = \\ &= \sigma\restriction_{\Delta_u}\end{aligned}$$

As for v itself, $\bar{\hat{\xi}}_{(\hat{c}_e)} \circ \zeta_v = \zeta_v \circ \xi\restriction_{\Delta_v}$ by definition, hence:

$$\begin{aligned}\bar{\hat{\xi}}_{(\hat{c}_e)} \circ \iota\restriction_{\Delta_u} &= inn_{\hat{c}_{\bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_v)}} \circ \zeta_u \circ \xi\restriction_{\Delta_v} = \\ &= inn_{\gamma_v} \circ \zeta_v \circ \xi\restriction_{\Delta_v} = \iota \circ \sigma\restriction_{\Delta_v}\end{aligned}$$

All is left to check is that $\bar{\hat{\xi}}_{(\hat{c}_e)}(\iota(t_f) = \iota(\sigma(t_f))$. Let us first check the case in which $f \in E \backslash Z$ originating at a vertex $u \neq v$ and ending at a vertex w. We have $\sigma(t_f) = (c_u)^{-1} t_f c_v$, hence:

$$\begin{aligned}\iota(\sigma(t_f)) &= \iota(c_u)^{-1} \iota(t_f) \iota(c_w) = \\ = (\gamma_u^{-1} \hat{c}_{\Phi(u)} \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u))^{-1} (\gamma_u^{-1} \epsilon_f \gamma_v) \gamma_w^{-1} \hat{c}_{\Phi(u)} \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_w) &= \\ = \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u)^{-1} (\hat{c}^{-1}_{\Phi(u)} \epsilon_f \hat{c}_{\Phi(w)}) \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_w) &= \\ = \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u^{-1} \epsilon_f \gamma_w) = \bar{\hat{\xi}}_{(\hat{c}_e)}(\iota(t_f))\end{aligned}$$

Here ϵ_f is equal to either $\epsilon_f = \bar{t}_f\theta(t_f)$ or $t_{\hat{f}}$. Up to taking inverses, the only case left case is that in which f originates at v. Then $\bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_v) = \gamma_v$ and using 5.3.8 again we get:

$$\begin{aligned}\iota(\sigma(t_f)) = \iota(t_f c_f) = \gamma_v^{-1} t_{\hat{f}} \gamma_w \iota(c_w) = \\ = \gamma_v^{-1} t_{\hat{f}} \hat{c}_{\Phi(w)} \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_u) = \\ = \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_v^{-1}) \bar{\hat{\xi}}_{(\hat{c}_e)}(t_{\hat{f}}) = \\ = \bar{\hat{\xi}}_{(\hat{c}_e)}(\gamma_v^{-1} t_{\hat{f}} \gamma_w) = \bar{\hat{\xi}}_{(\hat{c}_e)}(\iota(t_f))\end{aligned}$$

In case (iii) one can take as $\hat{\sigma}$ a vertex group automorphism extending an automorphism of $\hat{\Delta}_{[\hat{v}]}$ supported in the direct summand $K_v'^{c_w}$ corresponding to v. This case lacks the minor subtleties of the previous two and is left for the reader.

In all these cases, it is clear that if σ is an π-modular automorphisms and the twisting element c_e maps to the center of Q then the lift $\hat{\sigma}$ constructed above is a π-modular automorphism of $\hat{\Delta}$. Likewise, if σ fixes A the same is true of $\hat{\sigma}$. This concludes the proof. □

Remark 5.3.9. It follows from the proof that if $\mathcal{R}$ is well-separated (taut) the same is true for its completion. In the latter case the extension of a morphism factoring in a taut fashion does as well.

5.4 π-test sequences

In [Sel03], a certain class of sequences of homomorphisms from a group endowed with a tower structure to the free group, called test sequences, are defined in terms of a certain (long) list of combinatorial conditions. Roughly speaking, the goal of the definition is, to ensure that the image of a fixed tuple of generators by the test sequence (often itself denoted by the term 'test sequence') eventually escapes any diophantine condition not already witnessed within the tower. This will be made more precise in the last chapter.

For our purposes all we need is test sequences which are morphisms of π-groups, what we call π-test sequences. In an attempt to make the reading easier for the reader, our presentation differs slightly from that in [Sel03]. The definition of a π-test sequence provided below is (regardless of π) is in fact weaker than the version presented there, as it regards the geometric properties of the sequence in isolation from which we immediately draw a couple of simple consequences. The next section is devoted to the proof of the existence of a family of test sequences with certain properties. In Sela's terms this would be phrased as the existence of test sequences, without further qualifiers.

Given a pegging $\mathsf{T^q}_A$ of $\mathcal{T}_A$, it will be convenient to replace the π-tower structure inherited from $\mathcal{T}_A$ by what we will call a 'refinement' of $\mathcal{T}_A$. The underlying index tree J^*, node groups and graph of groups decomposition T^λ will be again essentially independent of q. being the for each T. Decompose the complement of the peripheral group of each abelian vertex group appearing in $\mathcal{T}$ into a direct sums of cyclic groups; let $\mathcal{CS}$ the collection of all the resulting cyclic groups and $\mathcal{SG}$ be that of all the surface type vertex groups appearing along $\mathcal{R}$. The properties we need $\mathcal{T}_A^q$ to have are as follows. Firstly for each $\lambda \in (\hat{J})$ either:

a) λ is of free product type, with exactly two children

b) $(\Delta\mathcal{T})^\lambda$ has exactly one vertex v of non-rigid type and either $(\Delta\mathcal{T})_v^\lambda$ is conjugated to a member of $\mathcal{SG}$ or else v is of abelian type, $(\Delta\mathcal{T})^\lambda$ contains at least a rigid vertex and $(\Delta\mathcal{T})_v^\lambda$ is the direct sum of its peripheral group with a cyclic group conjugated to one in $\mathcal{CS}$.

Secondly, any group in $\mathcal{SG}$ or in $\mathcal{CS}$ appears at some node of floor type in the way described above.

We start by placing $\mathsf{T^q}_A$ at the root r (same as before) of our tower. Assume that for some $\lambda \in J$, the group T^λ already appears as T^λ for some $\lambda^* \in J^*$ which is a leaf of the partial resolution. We extend our partially constructed tower $\mathcal{T}$ by attaching at λ a resolution at whose leaves we find precisely all the $\bar{\mathsf{T}}^\mu$ for $\mu \prec \lambda, \mu \in lv(J)$, as described now. Assign an order $v_1, v_2, \cdots v_{m_S}$ to the non-rigid surface groups appearing in $\Delta = (\Delta\mathcal{T})^\lambda$. In case v_i is of abelian type, pick up an ordered base $(m_i^j)_{j=1}^{n_i}$ of the complement of the peripheral subgroup of Δ_{u_i} and let $M_i^j = \langle m_i^k \rangle_{k=1}^{n_i}$.

If removing v_1 from $|\Delta|$ does not disconnect the graph, after renumbering the v_i, we can assume it does not disconnect our reference spanning tree either. At the nodes $nDFT(\lambda)$ we take the fundamental groups of each of the resulting connected components with respect to the restriction of the spanning tree. We proceed by removing the vertices of non rigid type of those components in the order specified by the given numbering, working with each of the components in parallel. If the vertex v_i is of abelian type, we add a chain $\lambda^* = \mu_0 \succ \mu_1 \succ \cdots \mu_{n_1}$, where T^{μ_l} is the amalgamated product of $M_i^{n_i - l} *_{Per(M_i)} H$, for H is the fundamental group of the graph of groups obtained removing v_i. If $\lambda \in J$ is of free product type, we can dislodge free factors one by one in a series of nodes of free product type with only two factors.

In addition to this, we will impose a linear order, denoted by $\lhd$, on J^*. On the set of children of any fixed $\lambda \in J^*$, we define $\lhd$ arbitrary, except for the constraint that the child carrying the constants must be the minimum element. Finally, given any pair of distinct nodes $\lambda_1, \lambda_2 \in J^*$, consider their lowest upper bound ν and suppose that $\lambda_i \leqslant \mu_i$ for $\mu_i \prec \lambda$ for $i \in \{1, 2\}$. We let $\lambda_i \lhd \lambda_j$ if and only if $\mu_i \lhd \mu_j$, where $\{i, j\} = 1, 2$. This is a standard construction: it corresponds to a depth-first search of the nodes of the tree.

Definition 5.4.0.1. Let $\mathcal{T}_A^{(J,r)}$ be a restricted π-tower structure of a group G_A over π-limit groups $\{\mathsf{T}^\lambda\}_{\lambda\in\Lambda}$ and suppose that for each $\lambda \in \Lambda$ we are given a sequence $(g_n^\lambda)_n$ of morphisms from A_{S^λ} to F_A with trivial limit kernel. We will say that a sequence $(f_n)_n$ of morphisms factoring through $\mathcal{T}_A$ is a *π-test sequence* relative to the family $\{(g_n^\lambda)_n\}_{\lambda\in\Lambda}$ if and only if it has trivial limit kernel, $(f_n\restriction_{\mathsf{T}^\lambda})_n$ is equal to some subsequence $(g_n^\lambda)_n$ up to postcomposition by an inner automorhisms and the following properties are satisfied.

For any $\lambda \in J\backslash\Lambda^*$, where $\Lambda^* \subset lv(J^*)$ is the set of leaves corresponding to those in Λ, we require the sequence f_n^λ to be unbounded and metrically convergent and the action of T^λ on a limiting tree Y associated to the sequence $(f_n^\lambda)_n$, where $f_n^\lambda := f_n\restriction_{(T*)^\lambda}$ to be as follows:

(FG) If $\lambda \in lv(J\backslash\Lambda^*)$, (which implies T^λ is unrestricted) Y is a geometric realization of a simplicial T^λ-tree S, with a single orbit of trivially stabilized vertices and as many orbits of trivially stabilized edges as the rank of the free group.

(FP) If T^λ is of free product type, with children μ_1 and μ_2, then Y is equivariantly isomorphic to the geometric realization of the simplicial action associated with the free product $\mathsf{T}^{\mu_1} * \mathsf{T}^{\mu_2}$.

(Fl) If T^λ is of floor type then the action of T^λ on Y admits a decomposition $(S, Y_{(-)}, p_{(-)})$ as a graph of actions. in which S is the tree dual to Γ. There are two cases, according to the type of the unique non-rigid vertex u of $(\Delta\mathcal{T})^\lambda$:

- **(S)** If u is of surface type, then $\Gamma = (\Delta\mathcal{T})^\lambda$ and for any surface type vertex $u \in VS$, the action of $Stab_S(u)$ on Y_u is dual to an arational measured foliation on Σ_u and is not the pullback of any similar action by the map induced by a covering map $q : \Sigma \to \Sigma'$, while the action of the stabilizers of rigid vertices is trivial.

(A) If u is of abelian type, in which case Y is isomorphic to the geometric realization of the tree dual to an HNN extension, where the vertex group is that of the single rigid type vertex of $(\Delta\mathcal{T})^{\lambda}$ and the Bass-Serre a generator of a non-peripheral direct summand of $Stab(u)$.

In cases (FG) and (FP) above we also impose that for any finitely generated restricted π-limit group H_{A_λ} containing $(\mathcal{T}^*_A)^{\lambda}$ and any diverging subsequence of extensions $(h_{k_n})_n$ of $(f_{k_n})_n$ to H_A with trivial limit kernel and limiting tree Z either T^{λ} is elliptic in Z or its minimal tree $Z_{\mathsf{T}^{\lambda}}$ satisfies:

$$\text{For any } h \in H, \text{ if } Z \cap h \cdot Z \text{ is non-degenerate, then } h \in \mathsf{T}^{\lambda}. \qquad (\dagger)$$

We also require:

$$\mathsf{T}^{\mu} \text{ grows faster than } \mathsf{T}^{\nu} \text{ under the sequence } (f_n)_n \text{ for } \mu \lhd \nu. \qquad (\dagger\dagger)$$

As we will see in the course of the construction, the condition that the sequence $(f_n)_n$ has limit kernel is to a certain extent redundant, and can be pulled up from the bottom of the tower even if one starts with a slightly weaker (but somehow cumbersome) definition. We note in passing the following:

Observation 5.4.1. Let $\mathcal{T}_A^{(J,r)}$ be a π-tower structure of a group G and H a subgroup of G. There is a unique minimal $\lambda \in J_0$ such that H can be conjugated into T^{λ}.

Proof. It follows easily from the following two facts:

i) Given a retraction r of a group G onto some $H \leqslant G$, two subgroups $K_1, K_2 \leqslant H$ are conjugate in K if and only if they are conjugate in H.

ii) Given a free product $G = *_{i=1}^{m} G_i$, a subgroup $H \leqslant G$ can be conjugated into G_i for at most on value of i.

□

Lemma 5.4.2. *Fix $\lambda \in J^*$ and let G be T^{λ}, H finitely generated containing G and g_n an extension of f_n^{λ} to some finitely generated H for each $n \in \mathbb{N}$ such that $lker_n\ g = \{1\}$, with limting tree Y.*

Assume that H is freely indecomposable relative to $\{\mathsf{G}\} \cup \mathcal{A}$, where $\mathcal{A}$ is the family of all the subgroups of H elliptic in Y. Then the action of H on Y admits a decomposition as a tree of actions $(S, Y_{(-)}, p_{(-)})$ in which each vertex action is either of Seifert (surface), axial or simplicial type. If G is elliptic in Y, then all the the simplicial components are non-trivially stabilized. If G is not elliptic in Y, denote by L either the non-rigid type vertex group of $(\Delta\mathcal{T})^{\lambda}$ in cases (S) and (A) or G in case (FG). Then there is a unique orbit of non-degenerate components and the minimal tree Y_L is equal to one of them Y_{v_0}.

- *In case (FG) we have $\mathsf{G} = H$.*
- *In cases (S) and (FP) the group $Stab(v_0)$ coincides with L and all the points of intersection betweeen Y_L and other non-degenerate components are non-trivially stabilized in L.*

Proof. First, observe that if the group T^{λ} is itself elliptic in Y, then H cannot be freely indecomposable relative to the family of all point stabilizers of Y. If it is not elliptic, the same is true in cases (S) and (A) of the definition of test sequences, since then T^{λ} is freely indecomposable relative to its rigid vertex groups, which are elliptic in Y in virtue of the properties of test sequences. This means that the first alternative of 3.1.17 does not hold, and we can decompose Y as a normalized tree of actions $(S, Y_{(-)}, p_{(-)})$ of the desired type, with no trivially stabilized simplicial edges.

So assume now that T^λ is not elliptic. Consider (S) first. Let the action on Z be dual to an arational measured folitation $\mathcal{F}_0$ on Σ_0. Observe that since the action of the vertex surface group is indecomposable and non-linear $Z \subset Y_{v_0}$ for some component Y_{v_0} of Seifert type. Of course, this implies $L \leqslant Stab(Y_{v_0})$. Now, boundary subgroups of L fix points in Y_{v_0}, hence they are contained in the ones associated to boundary components of Σ. In virtue of the results of [Sco78], the embedding of L into $\pi_1(\Sigma)$ is induced by some covering map $q : \Sigma_0 \to \Sigma$. The defining properties of test sequences imply q has to be a homeomorphism, hence $L = Stab(v_0)$. In (A), the minimal tree $Z = Y_L$ is a line which according to 3.1.16 has to be fundamental in Y and the conclusion follows from corollary 3.1.16. In (FG) and (FP) and corollary 3.1.18 yields a decomposition $(S, Y_{(-)}, p_{(-)})$ of the required type which includes Z as a vertex tree Y_{v_0}. The full strength of (†) implies that $Stab(v_0) = L$. □

Corollary 5.4.3. *In the situation above, if T^λ is not elliptic in Y, there is a minimal simplicial H-tree U in which T^λ is not elliptic and a T^λ-subtree (sub-geometric abelian tree) W of Y such that:*

- *If λ is as in case (FG), then $\mathsf{T}^\lambda = H$.*
- *If λ is as in (FP) U has trivially stabilized edges and W is equivariantly isomorphic to the tree dual to the decomposition $\mathsf{T}^\lambda = (\mathsf{T}^{\mu_1})^{\gamma_\lambda} * (\mathsf{T}^{\mu_2})^{\gamma_\lambda}$, where $\{\mu_1, \mu_2\} = Ch(\lambda)$.*
- *In case λ is of floor type then W is equivariantly isomorphic as a geometric abelian tree to the dual tree of the decomposition $(\Delta\mathcal{T})^\lambda$. In case (S), $Stab_{\mathsf{T}^\lambda}(v) = Stab_H(v)$ holds for any non-rigid vertex v of W.*

Any subgroup of H elliptic in Y is elliptic in U and stabilizes a rigid vertex of U in case λ is of floor type.

Proof. Case (FG) is already dealt with in the lemma. In case (A) the conclusion follows directly from lemma 3.1.16. In case (FP), we can apply lemma 2.3.1 to the skeleton S and the simplicial tree underlying the action of T^λ on Y_{v_0}. Recall that the stabilizers of edges of S incident to v_0 can only stabilize points of Y_{v_0} which are non-trivially stabilized in T^λ and is therefore elliptic in the associated simplicial tree. In case (S) the decomposition can be obtained by collapsing all the edges of S which are not incident to a translate of v_0. Each edge of S incident to v_0 corresponds to a point of Y_L stabilized by a boundary subgroup of L and have that same stabilizer in H. This implies the union of the translates by elements of T^λ of the star around v_0 is isomorphic to the tree dual to $(\Delta\mathcal{T})^\lambda$. A subgroup of H elliptic in Y is elliptic in S. On the other hand any subgroup of $Stab(v_0)$ elliptic in Y_{v_0} is contained in the stabilizer of an incident edge group. This proves the last claim for cases (FG) and (S). In case (A) the property was already contained in lemma 3.1.16. □

5.4.1 Two basic properties of test sequences

Lemma 5.4.4. *Assume $\mathcal{T}_A^{(J,r)}$ is a closed tower structure of G_A and that all members of a given π-test sequence $(f_n)_n : \mathsf{T}_A^q \to \mathsf{F}_A$ restrict to the same embedding κ of the node group H_A which carries the constants into F_A. Assume that for $c \in G$ either $Z_G(c) = \langle c \rangle$ or a peg of $Z_{G(c)}$. Then either of the following happens:*

i) Some conjugate c^g of c belongs to the terminal group of $\mathcal{R}_A$ which carries the constants.

ii) For n big enough $f_n(c)$ is primitive in $\mathbb{F}$.

Proof. Assume for the sake of contradiction none of the two options were true and let $\lambda \in J\backslash DFt(\mathcal{T})$ be minimal such that c is conjugate to some $d \in \mathsf{T}^{\lambda} = \mathsf{L}_{A_\lambda}$. Let K the amalgamated product of $M = \langle d, u\rangle_{ab}$ and $L = \mathsf{T}^{\lambda}$ over $\langle d\rangle$. Notice that this is a CSA group, since d generates its own centralizer in L.

Then a metrically convergent sequence $(g_n)_n$ of extensions of $(f_n)_n$ to K exists so that $g_n(u)^{k_n} = f_n(d)$ for some $k_n \notin \{-1, 0, 1\}$. Let $P = K/lker_{g_n}\ n$. The first observation is that $Z_Q(d) \neq \langle z\rangle$ as for no $n, m \in \mathbb{N}$ the element $d^m u^{-1}$ in the kernel of g_n, and hence that L does not map onto P. Clearly P is freely indecomposable relative to H, since free factors are centralizer closed. Consider now the action of P on a limiting tree for the sequence L.

Observe that L is not elliptic in Y. Indeed, suppose for a second it is elliptic we claim that then M must necessarily be elliptic as well. For M can only act by translations on its minimal tree (abelianity of an action is clearly preserved by taking limits in the equivariant Gromov topology). On the other hand $tl^Z(u) \leqslant \lim_n \frac{1}{k_n} tl^Z(d)$, so M in this case is generated by elliptic elements. As a result there are points x, stabilized by M and y whose stabilizer contains L such that Y is equivariantly isomorphic to a geometric realization of the tree dual to the amalgamated product of $Stab(x)$ and $Stab(y)$ over the set-wise stabilizer of $[x, y]$ (since by superstability if an element of $Stab(x) \cup Stab(y)$ fixes a non-degenerate segment of $[x, y]$ then it must fix the entire $[x, y]$). But this is impossible, since implies that either d stabilizes a non-degenerate tripod (if the index of the set-wise stabilizer of $[x, y]$ in $\langle u, d\rangle$ is greater than 2) or that some element of M acts as a reflection on some infinite line (if the index is equal to 2).

So G is not elliptic in Y. It follows from minimality and our assumptions on d that λ falls in one of cases (FP), (FG) or (S) and d is not elliptic in Y. On the other hand, consider the decomposition of the action of P on Y as a tree of actions $(S, Y_{(-)}, p_{(-)})$ provided by corollary 5.4.2. In case (S) let $S_0 \subset S$ be the miminal tree of L in S and $Z = \bigcup_{v\in S_0} Y_v$. In the other ones let $Z = Y_v$ for the unique v stabilized by L. Observation 3.1.13 implies that Z is fundamental and $Stab(Z) = L$ in both. This means that $P = L$, since $Ax(d) \subset \bigcup_{v\in VS_0} Y_v$ fixed by u as well (in other words ug belongs to the limit kernel of $(g_n)_n$ for some $g \in G$). □

As a warm up before the proof of Merzlyakov's theorem, let us give one of the simplest and most natural examples of the idea behind test sequences.

Lemma 5.4.5. *Let $\mathcal{T}_A^{(J,r)}$ be a closed tower structure where A is full and malnormal in $\mathbb{F}$, $u, u_2 \in \hat{T}$ non-trivial elements of $\hat{T}$ and for some π-test sequence $(f_n)_n$ relative to a fixed embedding κ of the level group carrying the constants into $\mathbb{F}$ and n big enough $f_n(u_1) = f_n(u_2)^{t_n}$ for some $t_n \in \mathbb{F}$. Let κ be image of u_1 and u_2 are conjugate in $\hat{T}$.*

Proof. Assume without loss of generality that u_j belongs to T^{λ_j} for the minimal $\lambda_j \in J^*$ such that u_j can be conjugated to T^{λ} (this is well defined in virtue of 5.4.3). Consider first the case in which $\lambda_1 \neq \lambda_2$ and assume without loss of generality that $\lambda_1 \lhd \lambda_2$. For simplicity let $H_i = \mathsf{T}^{\lambda_i}$ for $i = 1, 2$ up to replacing $(f_n)_n$ by a subsequence, a geometrically convergent sequence $(g_n)_n$ of homomorphisms from $H_1 * H_2$ to $\mathbb{F}$ exist with the properties:

i) $g_n\restriction_{H_1} = f_n\restriction_{H_1}$

ii) $g_n\restriction_{H_2} = inn_{t_n} \circ f_n\restriction_{H_2}$ for some $t_n \in \mathbb{F}$ satisfying the property $f_n(u_1) = f_n(u_2)^{t_n}$ for which $sl_{*_n^1}(f_n(x)^{t_n s})$ is minimal, where x is a fixed tuple of generators of H_2.

Let $q : H_1 * H_2 \to Q$ be the limit quotient of $(g_n)_n$ and Y the associated limiting tree. We can think of H_1, H_2 as a subgroups of Q, which necessarily satisfies equality $u_1 = u_2^t = u_1^{ts}$. Property (††) implies $\mathsf{T}^{\lambda_1} \leqslant G$ (hence u_1) fixes some $*_1 \in Y$ which is the limit of the sequence $(*_n^1)_n$. We claim that H_2 is elliptic as well. Indeed, otherwise u_2, not being conjugate to a non-rigid vertex group of $(\Delta\mathcal{T})^{\lambda}$, would have to act hyperbolically on Y, something clearly incompatible with being conjugate to u_1. In view of this, let $*_2 \in Y$ be fixed by H_2. Since the action of H on Y is non-trivial, necessarily $*_1 \neq *_2$. This implies that the action of $\langle H_1, H_2^{ts}\rangle$ on its minimal tree in Y must be isomorphic to some geometric realization of the tree dual to an amalgamated free product of the form $H_1 *_E H_2^{ts}$, where E is the abelian segment stabilizer of $I = [*_1, *_2]$, since superstability of the action implies for any $g \in H_1$ either $g \in E$ or $g \cdot I \cap I = \{*_1\}$. But then 3.2.1 implies there is $c \in E$ and n big enough $sl_{*_n^1}(g_n(x)^{g_n(c)}) < sl_{*_n^1}(g_n(x))$, and the element $t'_n = t_n g_n(c)$ contradicts the minimality of t_n.

If $\lambda_1 = \lambda = \lambda_2$, then if T^λ carries the constants the claim is clear. If it does not, let $H = H_1$ again, up to replacing $(f_n)_n$ by a subsequence, a geometric convergent $(g_n)_n$ of homomorphisms form $H * \langle t\rangle$ to $\mathbb{F}$ exist such that:

i) $g_n\restriction_{H_1} = f_n\restriction_{H_1}$

ii) $g_n(u_2)^{g_n(t)} = g_n(u_1)$ and $t_n := g_n(t)$ minimizes $l_{*_n^1}(t_n)$ among those with that property.

Let $q : H * \langle t\rangle \to Q$ be the limit quotient of $(g_n)_n$ and Y the associated limiting tree. We can identify H with $q(H)$, so that $u_1 = u_2^{q(t)}$; this equation implies that Q is freely indecomposable relative to H. If H is elliptic in Y, then it fixes the limit $*_1$ of the sequence $(*_n^1)$ and an argument entirely analogous to the one used above shows that Y is isomorphic to some geometric realization of an HNN extension with vertex group H and Bass-Serre element t. An application of 3.2.1 contradicts the minimality condition for t_n in this case. If H is not elliptic in Y then neither are u_1 and u_2 and the intersection $Y_H \cap q(t)^{-1} \cdot Y_H$ must contain the whole axis of u_1. If λ falls in cases (FG),(S) or (FP) this implies right away that $q(t) \in u_1$ and we are done. If λ is as in case (A), we know that there is a decomposition of Q obtained from $(\Delta\mathcal{T})^\lambda$ by enlarging the two vertex groups (the non-rigid one stays abelian), dual to some tree S. If one among u_1, u_2 fixes an abelian-type vertex of S, then the other must fix that same vertex as well and by the CSA property $[q(t), u_2] = 1$, implying that $u_1 = u_2$.

If one of them is hyperbolic in S, then the other one must be as well. By inspecting their normal form one can easily check that the only possible case in which they are not conjugate in H is that in which, up to conjugation in Q (by different elements) we have $u_1 = r_1a_1$, $u_2 = r_2a_2$, where $a_i \in A$ and the two elements $r_1, r_2 \in R$ are not mutually conjugate in R but $r_1 = r_2^b$ for some $b \in B\backslash H$ (boolean '\'). This can be ruled out by induction hypothesis, so we are done.

□

5.5 Existence of π-test sequences

Fix now a restricted π-tower structure $\mathcal{T}_A^{(J,r)}$ relative to the the groups at the nodes of $\Lambda \neq (J)$ and suppose that for each $\lambda \in \Lambda$ a sequence $(g_n^\lambda)_n$ of morphisms from A_{S^λ} to F_{A_λ} with trivial limit kernel is given. For the remaining of the section, let $\mathcal{T}_A$ be a restricted π-tower which has already been refined, as described at the beginning of section 5.4 The goal of this section is to prove the following result:

Proposition 5.5.1. *For any $q \in pgq(T)$ there is a family $\mathcal{F}$ of test sequences of $\mathcal{T}_A$ relative to $((g_n^\lambda)_n)_{\lambda\in\Lambda}$ such that:*

- *$\mathcal{F}$ is closed under taking subsequences*
- *$\mathcal{F}$ is closed under diagonal subsequences*
- *$\mathcal{F}$ is congruency complete*

5.5.1 Proof of proposition 5.5.1

The proof is by induction on the size of J; the induction step involve showing the existence of an adequate family of π-test sequences of $\mathcal{T}_A^{(J,r)}$ provided such a family exists for $\mathcal{T}_A\restriction_\lambda$ for any $\lambda \in Ch(r)$. We will refer to a family of divergent sequences satisfying the first two points above as a CSD family. It is not hard to see that a CSD family $\mathcal{F}$ corresponds to a descending chain

$$Hom(G,\mathbb{F}) \supseteq [\mathcal{F}]_0 \supset [\mathcal{F}]_1 \cdots$$

with trivial intersection, from which $\mathcal{F}$ is recoverable as the set of all sequences $(f_n)_n$ such that $f_n \in [\mathcal{F}]_n$. Clearly all sequences in such an $\mathcal{F}$ have trivial limit kernel if and only if for any $g \in G\backslash\{1\}$ there is some $n \in \mathbb{N}$ such that $f_m(g) \neq 1$ for all $(f_n)_n \in \mathcal{F}$ and $m \geqslant n$. If all sequences in the family share a common limiting tree we say that the family is convergent. Note that in this case there is some $g \in G$ which grows uniformly fast, by which we mean that for any $L > 0$ some n exists such that $tl^{\lambda^{f_m}}(g) > L$ for any $m \geqslant n$ and $(f_m)_m \in \mathcal{F}$.

Observation 5.5.2. Suppose we are given groups $G_1, G_2 \cdots m$ and for $1 \leqslant i \leqslant m$ a CSD family $\mathcal{F}_i$ of convergent sequences in $Hom(G_i,\mathbb{F})$. Then there is a CSD family $\mathcal{F}$ of sequences of homomorphisms from $G_1 * G_2 \cdots G_m$ to $\mathbb{F}$ such that:

- $(f_n\restriction_{G_i}) \in \mathcal{F}_i$ for $1 \leqslant i \leqslant m$
- Given any tuple $((f_n^j)_n)_{1 \leqslant j \leqslant m} \in \mathcal{F}_1 \times \mathcal{F}_2 \cdots \mathcal{F}_m$ there is some $(f_n)_n$ in $\mathcal{F}$ such that $(f_n\restriction_{G_i})_n$ is a subsequence of $(f_n^i)_n$ for $1 \leqslant i \leqslant m$.
- $(f_n)_n$ makes G_i grow faster than G_j for $i < j$

Proof. Simply chose a finite set of generators $\mathcal{X}_i$ of G_i for $1 \leqslant i \leqslant m-1$ and take as $\mathcal{F}$ the collection of all sequences $(g_n)_n$ of maps from $G_1 * G_2 \cdots G_m$ such that $(g_n\restriction_{G_j})_n \in \mathcal{F}_i$ and that $tl^{\lambda^{f_n^j}}(xy) \leqslant n \cdot tl^{\lambda^{f_n^{j+1}}}(g_{j+1})$ for all $n > 0$, $1 \leqslant j \leqslant m-1$, $x, y \in \mathcal{X}_j$ and some $g_{j+1} \in G_{j+1}$. □

Free groups, free products and small cancellation sequences

Fix a base $\mathcal{B}$ of a free group $\mathbb{F}$. We say that $x \in \mathbb{F}$ is cyclically reduced if the translation length of its normal form is minimal among those of all of its conjugates. Alternatively, if the the length of its normal form coincides with the translation length of its action on the Cayley graph of $\mathbb{F}$, i.e. if the orbit of $1 \in Cayl(\mathbb{F})$ by powers of x is contained in the axis of x. Any element $y \in \mathbb{F}$ can be expressed in a unique way as a product without cancellation of the form $u \cdot x \cdot u^{-1}$ so that $x = y^u$ is cyclically reduced (see [LS15] for more details. Given elements x, y (possibly equal) of a free group $\mathbb{F}$, a common piece of x and y is a word in $\mathcal{B}^{\pm}$ appearing as a subword of the normal form of cyclically reduced conjugates x^g and y^h of x and y. If $x = y$, we additionally require that $gh^{-1} \notin Z(x)$. A tuple $x = (x_j)_{j=1}^n$ of cyclically reduced elements of a free group $\mathbb{F}$ is said to have the small cancellation property $C'(k)$ for $k \in \mathbb{N}$ if any common piece of x_i and $x_j^{\pm}$ for $1 \leqslant i, j \leqslant n$ has length smaller than $\min\{|x_i|, |x_j|\}$.

Lemma 5.5.3. *Let (F^1, π^1) and (F^2, π^2) be free π-groups, where π^i and chose a base $(x_1, \cdots x_m)$ for F^1. Then for each $m \in \mathbb{N}$ there is a morphism ϕ^m from (F^1, π^1) to (F^2, π^2) such that if we let $y_j^m = \phi^m\bar{(}x_j)$ and $z_{i,m}$ be a cyclically reduced conjugate of $y_{i,m}$ then:*

- *The tuple $(z_1^m, z_2^m, \cdots z_n^m)$ has property $C'(m)$*
- *$|y_{i,m}| > m$ for all $1 \leqslant i \leqslant n$*
- *$\frac{|y_{i,m}|}{|y_{j,m}|} < 1 + \frac{1}{m}$ for all $i, j \in \{1, 2 \cdots m\}$*
- *$|z_{i,m}| \geqslant \frac{(m-1)}{m}|y_{i,m}|$.*

Proof. Let $\mathcal{R} \subset \mathbb{F}$ map bijectively onto Q via π and take $L = \max_{r\in\mathcal{R}} |r|$. Let n be the rank of n. It is a well known fact that for any m a tuple $\bar{y}'_m = (y'_{1,m} \cdots y'_{n,m})$ cyclically reduced elements of F^2 of length $\geqslant m$ satisfying the three conditions above exists. The only thing left is to ensure the compatibility of those tuples with the π-structure on both sides.

In order to do so, for each $1 \leqslant i \leqslant n$, let $r_{i,m} \in \mathcal{R}$ be such that $\pi^1(x_i) = \pi^2(y_{i,m})$, where $y_{i,m} = y'_{i,m} r_{i,m}$. Let u be such that $y'_{i,m} = w_{i,m} \cdot u$, $r_{i,m} = u^{-1} s_{i,m}$ and $y_{i,m} = w_{i,m} \cdot s_{i,m}$. Here the notation $u \cdot v$ indicates the fact that the product is reduced, i.e., no cancellation between a sub-word of the normal form for u and that for v takes place. As in the statement, let $z_{i,m}$ be a cyclically reduced conjugate of $y_{i,m}$. Since $\bar{y}_m$ satisfies the small cancellation property, if $y_{i,m}$ starts with the inverse of some final subword w of $z_{i,m}$, then $|w| \leqslant |s_{i,m}| + \frac{1}{m}|y_{i,m}|$. In short, $|z_{i,j}| \geqslant \frac{m-2}{m}|y_{i,j}| - L$. Furthermore, any piece of length C contained in some $z_{i,m}$ has to contain another one $y_{i,m}$ of length $\geqslant C - L$. Hence if we let $P(\bar{u})$ stand for the length of the biggest piece of $\bar{z}_m$ not intersecting any of the $u_{i,m}$ then:

$$\frac{P(z,\bar{m})}{\min_{1\leqslant i\leqslant n}|z_{i,m}|} \leqslant \frac{P(y,\bar{m}) + L}{\frac{m-1}{m}\min_{1\leqslant i\leqslant n}|z_{i,m}| - L} \leqslant \frac{\min_{1\leqslant i\leqslant n}|z_{i,m}|}{\min_{1\leqslant i\leqslant n}|z_{i,m}| - L}\left(\frac{1}{m} + \frac{L}{m}\right)$$

Clearly this term goes to 0 as m goes to infinity, so we are done. ☐

Observation 5.5.4. If a tuple $(x_1, x_2 \cdots x_n)$ of cyclically reduced elements of a free group satisfies $C'(3)$, it is automatically a base of the subgroup of the free group it generates, as in particular it satisfies Nielsen property (see [LS15] for a definition).

We will call a sequence of morphisms f_m satisfying the first two properties listed in the statement a small cancellation sequence with respect to the base $(x_1, \cdots x_m)$. If it also satisfies the third we will say it is balanced. Note that given a sequence of balanced small cancellation sequences $((f^m_n)_n)_{m\in\mathbb{N}}$ with respect to a certain base x of F_1, the sequence $(f^n_n)_{n\in\mathbb{N}}$ is also a balanced small cancellation sequence.

Lemma 5.5.5. *Let H be a free group, x a base of H and $(f_n)_n$ a balanced small cancellation sequence of morphisms from H to some other free group F with respect to x. Suppose we are given a finitely generated extension L of H and a sequence $(h_n)_n$ of maps extending f_n with trivial limit kernel and Y a limiting tree for $(h_n)_n$. Suppose that H is not elliptic in Y and let Y_H be the minimal tree of H. Then $h \cdot Y_H \cap Y_H$ can only be a non-degenerate segment in case $h \in H$ and the action of H on Y_H is a geometric realization of the action of H on $Cayl_x(S)$ which assigns the same length to all the edges of $Cayl(S, x)$.*

Proof. It all comes down to the following easy observation:

Observation 5.5.6. Let $(x_1, \cdots x_n)$ be a tuple of cyclically reduced elements of the free group $\mathbb{F}$. Let l_i be the translation length of x_i for its action on the Cayley graph of $\mathbb{F}$. The tuple has property $C'(n)$ if for any $h \in \mathbb{F}$ whenever $Ax(x_i)$ and $h \cdot Ax(x_j)$ overlap on a segment of length greater than $\frac{1}{n}\min\limits_{1\leqslant i\leqslant m} tl^{Cayl(\mathbb{F})}(x_i)$ necessarily $i = j$ and $h \in Z(x_i)$.

Assume that H does not fix a point in Y. Let $\nu_n = \min\limits_{1\leqslant i\leqslant |x|} f_n(x_i)$. Small cancellation implies that 1 cannot be at distance greater than $\frac{1}{n}\nu_n$ from the axis of any $f(x_i)$ for $1 \leqslant i \leqslant m$. On the other hand, for any $h \in H$, we have $\frac{n-4}{n}|h|\nu_n \leqslant |f_n(h)| \leqslant \frac{m+1}{m}|h|\nu_n$, the left inequality following from that given a reduced word $x^{\epsilon_1}_{i_1} x^{\epsilon_2}_{i_2} \cdots x^{\epsilon_r}_{i_r}$ at most $\frac{4}{n}$ of the reduced word for $f_n(x^{\epsilon_j}_{i_j})$ cancels out in the word $f_n(x^{\epsilon_1}_{i_1}) f_n(x^{\epsilon_2}_{i_2}) \cdots f_n(x^{\epsilon_r}_{i_r})$. This implies that the functions $\frac{1}{\nu_n} l_1^{Cayl(\mathbb{F})}(f_n(-))$ converge point-wise to $l_1^{Cayl_x(H)}(-)$ and the action of H on Y_H is exactly as desired. Suppose now that for some $g \in L$ the intersection $Y_H \cap g \cdot Y_H$ is a non-degenerate segment.

Since Y_H is covered by translates of the axis of the x_i, this implies $Ax(x_i)$ and $h \cdot Ax(x_j)$ overlap in a non-degenerate segment for $1 \leqslant i, j \leqslant m$. This in turn implies that for some $\alpha > 0$ and any n in some infinite subsequence $|Ax(f_n(x_i)) \cap f_n(h) \cdot Ax(f_n(x_j))| > \alpha tl(f_n(x_i))$. The small cancellation property together with the observation above implies that $i = j$ and $f_n(h) = f_n(x_i)^{r_n}$ for some $r_n \in \mathbb{Z}$ and any such n. Now, $r_n = \frac{tl(f_n(h))}{tl(f_n(h))}$ converges to $r_\infty = \frac{tl^Y(x_j)}{tl^Y(x_i)}$, so it must be constantly equal to $r_\infty \in \mathbb{Z}$ after a certain point, which implies that $h \in \langle x_i \rangle$. □

Corollary 5.5.7. *Suppose we are given a group $L = K * H$ and a $(g_n)_n$ a sequence of homomorphisms from L to $\mathbb{F}$ such that $g_n\restriction_H = f_{k_n}$ for some balanced small cancellation sequence f_n and $\frac{l_1(g_n(y_i y_j))}{k_j}$ tends to 0 for some finite tuple y of generators of L and $1 \leqslant i, j \leqslant |y|$. Then $lker_n\, g_n = 1$. and the limiting tree is a geometric realization of the tree dual to a graph of groups decomposition with a single vertex group K and a one-edge trivially stabilized loop for each of the given generators of H.*

Proof. Let $q : L \to Q$ be the limit quotient of the sequence $(f_n)_n$. Consider the action of $Q = L/lker_n\, h$ on any limiting tree Y for $(h_n)_n$. The last property in the definition of small cancellation sequences implies that $q(K)$ fixes a point in the minimal tree Y_H(3) . We know that $Y_H \cap g \cdot Y_H$ can only be non-degenerate for $g \in q(H)$ which then, if non-trivial, cannot fix any point of Y. It follows that the limit kernel of $(g_n)_n$ is trivial and Y of the desired form. □

Corollary 5.5.8. *Any π-tower without constants and a single node of type (FG) admits a CSD family of π-test sequences.*

Corollary 5.5.9. *Suppose that the root r is as in case (FP) and that for each $\mu \in Ch(r)$ a congruency complete CSD family of π-test sequences of $\mathcal{T}_A\restriction_\mu$ exists. Then a familiy $\mathcal{TS}$ with such properties exists for $\mathcal{T}_A$ as well.*

Proof. Let $Ch(r) = \{\mu_1, \mu_2\}$ and assume that μ_1 carries the constants. Let $\mathcal{TS}_0$ be the family of sequences of morphisms from T to $\mathbb{F}$ obtained by applying 5.5.2 to the families $\mathcal{TS}^{\mu_1}$ and $\mathcal{TS}^{\mu_2}$. Let Z be cyclic, with $\pi(Z) = \{1\}$ and $(g_n)_n$ a π-test sequence. Take $L = T * Z$, Let ι be the embedding of T into L given by $\iota\restriction_{\mathsf{T}_{\mu_1}} = Id$ and $\mathsf{T}_{\mu_2}^{=} inn_t$. Fix a tuple of generators x of T. As $\mathcal{TS}$ we take the set of all sequences of the form $(g_n \circ \iota) \subset Mor(\mathsf{L}_A, \mathsf{F}_A)$, where:

- $(g_n\restriction_T)_n \in \mathcal{TS}_0$
- $g_n\restriction_Z = h_{k_n}$ for some π-test sequence $(h_n)_n$ of Z and $k_n \in \mathbb{N}$ is such that $tl(g_n(x_i x_j)) \leqslant k_n$ for any $1 \leqslant i, j \leqslant |x|$.

It is easy to check that $\mathcal{TS}$ is CSD sequence of morphisms factoring through $\mathcal{T}_A$. The action of T on any limiting tree for $(f_n)_n \in \mathcal{TS}$ is simply the minimal tree of the pullback by ι of the action of L on the limiting tree for $(g_n)_n$. The result we are looking for is immediate, given what the we know about the latter from the previous lemma. It is also true, although not needed for this, that $lker_n\, g_n$. □

Approximating actions by twisting

Lemma 5.5.10. *Suppose we are given a group G, $\mathcal{F}$ a finite subset of G, $(Z_0, Z, (t_e)_{e \in EZ})$ a presentation of an action of G without inversions on a simplicial tree S with abelian edge stabilizers. Then there is a finite subset $\mathcal{H} \subset \bigcup_{v \in VS} Stab(v)$ such that given any $f : G \to H$ and any action λ of H by isometries on a real tree (X, d), if the following properties are satisfied:*

(3)The ratio between the distance between the axis of the small cancellation elements in $\mathbb{F}$ from the identity and the rescaling factor tends to 0.

i) λ *is* D *acylindrical for some* $D > 0$

ii) $\mathcal{H} \cap ker\, f = \varnothing$

iii) $f(Stab(e))$ *acts hyperbolically on* X *for all* $e \in S$

iv) Given e, f *incident at* $v \in S$*, if* $Stab(e) \cap Stab(f)^g \neq \{1\}$ *then either* [4]*:*

- $e = f$
- $e = t\overline{f}$ *for some Bass-Serre element* t *commuting with* $Stab(e)$

Then for each $x \in X$*,* $v_0 \in S$ *and any geometric realization* $(|S|, d')$ *of* S *there is a sequence* $(\tau_n)_n \subset \mathcal{A}$ *of products of Dehn twists over the edges of* e *such that the sequence* $\frac{1}{n} l_x^{\lambda \circ f \circ \tau_n}(g)$ *converges to* $l_{v_0}^{(\rho, d')}(g)$ *for any* $g \in \mathcal{F}$*, where* ρ *is the obvious action of* G *on* $|S|$*.*

In particular, if we take H to be free and as λ the action of H on its Cayley graph we deduce:

Corollary 5.5.11. *Suppose we are given a finitely generated* π*-group* G *a simplicial* G*-tree* S *with abelian edge stabilizers satisfying condition (iv) or of the lemma above and a sequence* $(f_n)_n$ *of morphisms from* G *to* $\mathbb{F}$ *such that for any* $v \in S$ *the sequence* $(f_n{\restriction}_{Stab(v)})_n$ *has trivial limit kernel. Then there is a sequence* $(\tau_n)_n \subset Mod^{\pi}(\mathsf{G}, S)$ *such that* $(f_n \circ \tau_n)_n$ *has trivial limit kernel.*

Example 5.5.12. The corollary is not true if one simply removes condition (iv). Consider the amalgamated product of a free group $\mathbb{F}(a, b, c)$ with the free abelian group with base $\{c, d\}$ over their common subgroup $\langle c \rangle$. Let $F' = \mathbb{F}(x, y)$ and let $(f_n)_n$ be the sequence of maps given by:

$$f_n(a) = y, \ \ f_n(b) = y^x, \ \ f_n(c) = x^{n+3}, \ \ f_n(d) = x^{n+2}$$

It is easy to check that the restriction of f_n to both vertex groups has trivial limit kernel, as the triple $(f_n(a), f_n(b), f_n(c))$ is Nielsen reduced. Now, any g_n obtained form f_n by precomposing by a Dehn twist of the decomposition above is equal to f_n up to post-composing by an automorphism of $\mathbb{F}$ (some power of the one fixing x and sending y to y^x), so in particular $ker\, f_n = ker\, g_n$. But $lker_n\ f_n \neq \{1\}$, as the element $b^{-1}a^{cd^{-1}} \neq 1$ is sent by the identity by all the f_n.

The hypothesis can be slightly weakened, though:

Corollary 5.5.13. *Suppose we are given a finitely generated* π*-group* G *and* S *a a simplicial* G*-tree with non-trivial abelian stabilzers such that* $VS = V_0 \amalg V_1$ *and*

- *Each vertex of* V_0 *is adjacent only to vertices of* V_1
- *There are no non-trivial elements fixing two distinct edges adjacent to a vertex of* V_1*.*
- *The stabilizer of any edge of* S *coincides with its centralizer in the stabilizer of any of its endpoints.*

[4] Equivalently, at the level of graph of groups:

- No two conjugates of the groups $\Delta_{[e]}$ and $\Delta_{[f]}$ intersect non-trivially for distinct and non mutually inverse $[e], [f]$ originating at $[v]$.
- The two edge groups associated to a one-edge loop coincide and the corresponding Bass-Serre generator commutes with them.
- $\Delta_{[e]}$ is malnormal in $\Delta_{[v]}$ for any edge $[e] \in G\backslash S$

If G is CSA the last bullet is equivalent to $\Delta_{[e]}$ being a maximal abelian subgroup of $\Delta_{[v]}$.

Assume we are also given a sequence $(f_n)_n$ of homomorphisms from G to $\mathbb{F}$ such that its restriction to the stabilizer of any vertex of S has trivial limit kernel. Then there is a sequence $(\tau_n)_n \subset Mod^{\pi}(\mathsf{G}, S)$ such that $lker_n\ f_n \circ \tau_n$ is trivial.

Proof. First of all notice that the conditions listed above imply that S is 2-acylindrical and $Stab(e)$ is self-centralizerd in G for any edge e in S [(5)] . Our proof is by induction on the number of edges of the associated decomposition Δ. Pick any edge $e \in S$ and let S^* be the tree resulting from the collapse of any edge outside the orbit of e. The induction hypothesis implies the existence of a sequence $(\sigma_n)_n \subset \mathcal{A}$ such that $lker_n\ f_n \circ \sigma_n\restriction_{Stab(u)} = \{1\}$ for any vertex $u \in S^*$.

The desired result will follow as soon as we manage to prove that S^* satisfies condition (iv). Since $Stab(e)$ is self-centralized and $Stab(u)$ is CSA, the only thing left to show is that $Stab(e) \cap Stab(f) = \{1\}$ for e, f originating at u in distinct orbits, this is clear, since the lifts $\tilde{e}$, $\tilde{f}$ in S of e and f respectively must span a segment of diameter at least 3. $\square$

Proof. (of lemma 5.5.10) Let f be any homomorphism from G to H. Chose a Bass-Serre presentation $(Z_0, Z, (t_e)_{e\in EZ})$ of the action of G on S for which $v_0 \in Z_0$. For any edge $e \in Z$ mapping onto a loop in $G\backslash S$, chose a Bass-Serre element t_e commuting with $Stab(e)$, in accordance to (iv). Assign to every edge $e \in S$ some $c_e \in Stab(e)$ in such a way that:

- $c_{\bar{e}}^{-1} = c_e$
- $c_{g^{-1}\cdot e} = c_e^g$
- c_e is hyperbolical with respect to λ^f for each $e \in E$; let tl_e be the translation length of $f(c_e)$ with respect to λ

For each $e \in S$ let α_e be the length of the segment corresponding to e in $(|S|, d')$ and tl_e the translation length of $f(c_e)$ with respect to λ. For each $n \in \mathbb{N}$ chose an integer $exp_e(n)$ closest to $\left\lfloor \frac{n\alpha}{tl_e} \right\rfloor$, with the property that the Dehn twist over e by $c_e^{exp_e(n)} \in \mathcal{A}$. Clearly there is a constant C independent of e and n such that $exp_e(n)$ and $\frac{n\alpha}{tl_e}$ differer at mostby C regardless of n. As τ_n we take the composition of the Dehn twists $\tau_{e,c_e}^{exp_e(n)}$, for e ranging among all edges of Z pointing away from v_0. Represent any $g \in \mathcal{G}$ as a product of the form $g_0 t_{e_0} \cdots t_{e_{m-1}} g_m$, associated to a sequence $v_0, e_0, v_1, e_1 \cdots e_{m-1}, v_m = v_0 \subset Z$ projecting to a closed path in $G\backslash T$, where $\alpha(e_i) \in V$ for each i. Let R be some constant bigger than:

i) Any of the tl_e.

ii) $l_x^{\lambda}(g_i)$ for $0 \leqslant i \leqslant m$

iii) The distance between x and $Ax(c_e)$ for any $e \in S$.

iv) $(C+1)(\sum_{e\in E} tl^{\lambda}(c_e))$

Then $\tau_n(g)$ can be written as $g_0 s_0 t_{e_0} g_1 s_1 t_{e_1} g_2 \cdots s_{m-1} g_m$ where $s_i = c_{e_i}^{-exp_{e_i,n}}$. We want to estimate the distance between the points x and $y = \lambda(\tau(g)) \cdot x$. Let: $I_0 = [x, g_0 \cdot x]$, $J_j = \bar{g}_j \cdot [x, s_j \cdot x]$ for $0 \leqslant j \leqslant m$, $I_j = \bar{h}_j \cdot [x, t_{e_{j-1}} g_j \cdot x]$ for $1 \leqslant j \leqslant m$, where $\bar{g}_j = g_0 s_0 t_{e_0} \cdots g_j$ and $\bar{h}_j = g_0 s_0 t_{e_0} g_1 \cdots s_{j-1}$ (the reader is reminded that although we have decided to leave the subscripts out, all of this is dependent of n). Notice that for $0 \leqslant j \leqslant m-1$ the right endpoint of the oriented segment I_j is the left endpoint of J_j, while the right endpoint of J_j is the left endpoint of I_{j+1} and the right endpoint of I_m is y. Now, clearly $|I_j|, |K| \leqslant 2R$ for $0 \leqslant j \leqslant m$. On the other hand $|J_j - n| \leqslant R$ and the subsegment $J'_j \subset J_j$ consisting of those

[(5)] That elliptic elements commuting with $Stab(e)$ are in $Stab(e)$ is easy to check. The axis of an hyperbolic element commuting with $Stab(e)$ has to be fixed by any non-trivial element of $Stab(e)$, against acylindricity.

points at distance $\geqslant R$ from ∂J_j is contained in $\bar{h}_j Ax(s_j) = \bar{h}_j Ax(c_{e_j})$, since x is at distance $\leqslant R$ from $Ax(c_{e_j})$. For each $e \in S$, let $l_e(g)$ be the number translates of e or $\bar{e}$ contained in the path $[v_0, g \cdot v_0]$. We deduce that:

$$\frac{1}{n} l_x^{\lambda}(\tau_n(g)) \leqslant \frac{1}{n}(\sum_{e \in E} \alpha_e |\{0 \leqslant j \leqslant m \mid [e_j] = e\}| + (2m+1)R) = l_\alpha(g) + \frac{2m+1}{n}$$

Assume for a moment that the subsegment J''_j of J_j consisting of those points at distance greater or equal than $3R + D + 1$ from ∂J_j intersects some I_k or J_l for $j \neq l$. Since $|I_j|, |I_{j+1}| \leqslant R$, for n big enough some subsegment of J_j of length greater or equal than $2R + D + 1$ is contained in J_{j-1} or J_{j+1} with the opposite orientation, let us say J_{j-1}. This implies J'_{j-1} and J'_j share a subsegment of length $\geqslant R + D + 1$ and hence the same is true for $(t_{e_{j-1}} g_j s_j)^{-1} \cdot Ax(c_{e_{j-1}}) = Ax(c_{e_{j-1}}^{t_{e_{j-1}} g_j s_j})$ and $Ax(c_{e_j})$ (in the sense of λ^f). Now, $c_{e_{j-1}}^{t_{e_{j-1}} g_j s_j} = (c_{t_{e_{j-1}}}^{-1})^{g_j s_j}) \in Stab(v_j)$. The following fact is well known and can be easily proven by inspection:

Fact 5.5.14. *Given an action of a group G on a real tree, $L > 0$ and two hyperbolical elements $h, g \in G$ such that $|Ax(g) \cap Ax(h)| \geqslant tl(g) + tl(h) + L$, the commutator $[h, g]$ fixes a segment of length at least L of $Ax(g) \cap Ax(h)$.*

It follows from this, together with the D-acylindricity of Y, that $f(([c_{t_{e_{j-1}}}^{-1})^{g_j s_j}, c_{e_j}]) = 1$, which implies that $f([(c_{t_{e_{j-1}}}^{-1})^{g_j}, c_{e_j}]) = 1$. We claim that $[(c_{t_{e_{j-1}}}^{-1})^{g_j}, c_{e_j}]$ is non-trivial in G. Indeed, assumption iv together with the fact that the expression $g_0 t_{e_0} \cdots t_{e_{m-1}} g_m$ is reduced imply that the only way it can be trivial is if $[g_j, c_{e_j}] = 1$ and $e_{j-1} = e_j$. This implies the existence of two inversely oriented subsegments of $Ax(c_j)$ both of which are oriented in the same direction in which c_{e_j} acts by translations, which is obviously an absurdity. We deduce from this that there is a finite collection of non-trivial elements of G such that if f does not kill any of them for each n big enough the segment $[x, y]$ will contain the disjoint union of the segments J''_j, which in turn implies $\frac{1}{n} l_x^Y(\tau_n(g)) \geqslant {}^{n l_\alpha(g) - m(6R+3D+D))}/_{n}$. Both this and the lower bound above tend to l_α with n, giving us the desired result. □

Let Σ be a big compact surface with boundary. We say that δ cuts the surface into small pieces if and only if any $\Sigma_0 \in [\Sigma]_\delta$ is either a a sphere minus three discs (a.k.a. 'pair of pants') or a projective plane with two disks removed.

Remark 5.5.15. Observe that such a essential system of curves has the property that given any homotopy class α of essential s.c.c. either $\alpha \in \delta$ or else any of the cyclic subgroups associated with α are hyperbolic in T_δ.

We say that two essential systems of curves δ and ϵ fill Σ if for suitable representatives of δ, ϵ, each connected component of $\Sigma \backslash (\delta \cup \epsilon)$ is either simply connected or a boundary-parallel annulus. In particular, for any $\alpha \in \gamma$ there is some $\beta \in \delta$ such that $i(\alpha, \beta) \neq 0$ and viceversa.

Remark 5.5.16. This is equivalent to the fact that any non-trivial element of $\pi_1(\Sigma, *)$ must be hyperbolic in either T_γ or T_γ.

Given a essential system of curves γ on Σ let Δ_γ be a graph of groups decomposition corresponding to the tree T_γ, as defined in 2.4, whose vertex groups are the image in $\pi_1(\Sigma, *)$ of the fundamental groups of the subsurfaces in $[\Sigma]_\delta$ (this image is well defined up to conjugation).

Fact 5.5.17. *Given a essential system of curves γ in a compact bounded surface Σ there is another essential system of curves δ such that γ and δ fill Σ.*

The curve graph of a big surface Σ, which we denote by $\mathcal{C}^1(\Sigma)$ is the graph whose vertices are homotopy classes of essential simple closed curves in Σ, where an edge is added between α and β in case $i(\alpha, \beta)$ is minimal. Clearly a essential system of curves is a subset $\mathcal{C}^1(\Sigma)$ of diameter $\leqslant 1$. On the other hand, any two curves at distance $\geqslant 3$ in $\mathcal{C}^1(\Sigma)$ fill Σ, as there is no essential simple closed curve disjoint with both. The statement above follows therefore from the very well-known fact that $\mathcal{C}^1(\Sigma)$ has infinite diameter. A proof of the following can be extracted from [Wil09, 4.17]

Lemma 5.5.18. *Let S be the fundamental group of a big bounded compact surface Σ . Suppose we are given a homomorphism f from S to some limit group G which injects the groups corresponding to boundary components of S and has non-abelian image. Then there is a essential system of curves δ on Σ cutting Σ into small pieces such that for each $\alpha \in \delta$ such that the fundamental group of any $\Sigma_0 \in [\Sigma]_\delta$ is mapped injectively by f into G.*

Surface type vertex groups

As a warm up, let us start by proving the following weaker and well-understood result:

Lemma 5.5.19. *Suppose we are given a π-tower structure $\mathcal{T}_A^{(J,r)}$ and $\lambda \in J$ of floor type, where $(\Delta\mathcal{T})^\lambda$ contains a unique vertex of non-rigid type, which is of surface type, as well as a sequence of morphisms $(f_n)_n : *_{\mu<\lambda}\mathsf{A}_{S^\lambda} \to \mathsf{F}_A$ with trivial limit kernel. Let δ be a essential system of curves on the corresponding surface Σ which cuts it into small pieces.*

Proof. Let δ be the essential system of curves provided by lemma 5.5.18, applied to the unique surface group of $\Delta = (\Delta\mathcal{T})^\lambda$, isomorphic to $\pi_1(\Sigma, *)$, and $\eta := \eta^\lambda_{\mathcal{R}}$. Let S be the tree dual to $G := (\Delta\mathcal{T})^\lambda$ and S_δ the one obtained by blowing up those vertices in the unique orbit $G \cdot v_0$ of S to a copy of the action on T_δ. Let S'_δ be the tree obtained from S_δ by folding together $(e, g \cdot e)$ for any e originating at a rigid type vertex v and any g in the centralizer of $Stab(e)$ in $Stab(v)$. The resulting decomposition is identical to that associated to S_δ, except that at each of the vertices where we saw $H := \pi_1(\Sigma_0)$ for some $\Sigma_0 \in [\Sigma]_\delta$ we now see the amalgamated product of H with the centralizers of groups corresponding to each of the one or two boundary components of Σ_0. It is clear now that the tree S'_δ satisfies the list of conditions of corollary 5.5.13 (with the set of rigid vertices as V_0).

We need to check that the restriction of $(f_n)_n$ to the stabilizer of a vertex u of S'_δ not already in S_δ has trivial limit kernel. Observe that in that case $Stab(u) = A * B$, with A and B are the centralizers of incident edges.

Claim 5.5.20. *Let $G = A * B$ for free abelian groups A and B and let $(f_n)_n$ be a sequence of homomorphisms from G to a free group such that $(f_n\restriction_A)_n$ and $(f_n\restriction_B)_n$ both have trivial limit kernel and $f_n(G)$ is non-abelian for any n. Then $(f_n)_n$ has trivial limit kernel.*

Proof. This follows easily from the well-known fact that any pair of two non-commuting elements of the free groups form a basis of the subgroup of the free group they generate; in particular here $f_n(G) \cong f_n(A) * f_n(B)$ for any n. For any $g \in G$ with normal form (let's say) $a_1b_1a_2b_2\cdots b_n$, where $a_i, b_i \neq 1$, for n big enough both $f_n(a_i)$ and $f_n(b_i)$ are non-trivial for all i, which together with the fact above clearly implies $f_n(g) \neq 1$. □

□

Lemma 5.5.21. *Let G be a group admitting an action on a geometric abelian tree S, such that $G\backslash S$ is bipartite, with edges joining a single vertex $[v_0]$ of surface type with (finitely many) vertices of rigid type. Let Σ be the surface associated to v_0, so that $Q_0 := Stab(v_0) \cong \pi_1(\Sigma)$. Suppose we are given some finite index subgroup $R \leqslant Mod(\Delta) \subset Aut(G)$. and essential* *system of curves δ on Σ cutting it into small pieces. Let W be the tree obtain by blowing up $v_0 \in S$ to the $Stab(v_0)$-tree T_δ dual to the essential system of curves δ*[6] *and then collapsing all edges inherited from S and let $(|W|, d, \rho)$ a geometric realization of it. Then there is a sequence $(\theta_n)_n \in R$ and a diverging sequence $(\mu_n)_n$ of positive constants such that if we let W_0 be the minimal tree of W the sequence $(|W_0|, \frac{1}{\mu_n}d, \rho \circ \theta_n)$ converges in the equivariant Gromov topology to some action on a real tree (Y, d_{lim}, ρ_{lim}) of G on a real tree which can be decomposed as a tree of actions $(S, Y_{(-)}, p_{(-)})$, where*

- *Y_u is a point for $u \in G \cdot v_0$*
- *Y_{v_0} is dual to an arational measured foliation on Σ and for any covering map $q : \pi_1(\Sigma, *) \to \pi_1(\Sigma', *)$ the action of $Stab(v_0)$ on its minimal tree does not extend via q_v* to any action of $\pi_1(\Sigma', *)$ on a real tree.*

Proof. Let γ be a essential system of curves such that γ and δ fill Σ. Take any collection $\{\phi_n\}_{n \in \mathbb{N}}$ of non-trivial Deck transformations of Σ such that (up to an homeomorphism isotopic to the identity) any proper covering map with domain Σ admits ϕ_n as a Deck transformation for some n. Let us first show how to reduce the lemma to the following claim:

Claim 5.5.22. *There are sequences $(\theta_n)_n \subset R$ restricting to an automorphisms of Q_0 and $(\mu_n)_n \subset \mathbb{R}^+$ such that:*

i) For any $g \in G$ which does not stabilize a rigid type vertex of W and n big enough $\theta_n(g)$ is hyperbolic in W.

ii) $|ln(\frac{tl^{(\rho,d)}(\theta_n(g))}{tl^{(\rho,d)}(\theta(h))})|$ is convergent for h, g none of which stabilizes a rigid type vertex of W.

iii) $tl^{(\rho,d)}(g)$ diverges for some $g_0 \in Q_0$

iv) For any two elements $h, g \in G$ which do not stabilize a rigid type vertex of W $lim_n|ln(\frac{tl^{(\rho,d)}(\theta_n(g))}{tl^{(\rho,d)}(\theta(h))})| < \infty$

v) For any n there is some $a_n \in Q_0$ such that $lim_m|ln(\frac{tl^{(\rho,d)}(\theta_n(a_n))}{tl^{(\rho,d)}(\theta(a'_n))})| > 0$, where a'_n stands for $(\phi_n)_*(a_n)$ by a'_n [7]*.*

Indeed, this clearly implies some divergent (by iii) sequence $(\mu_n)_n$ of positive constants exists such that the sequence $\frac{1}{\mu_n}tl^{(d,\rho\circ\theta_n)}(--)$ converges pointwise to $tl^{(d_{lim},\rho_{lim})}(--)$ for some isometric action ρ_{lim} of G on real tree Y_{lim}. As we know, this implies convergence in the equivariant Gromov topology. Observe the following:

a) For any u of rigid type any $g \in Stab(u)\backslash\{1\}$ is elliptic with respect to each of the actions $\rho \circ \theta_n$ and therefore the point it fixes in the limit action is unique.

b) For such u the group $Stab(u)$ itself is elliptic. let x_u be the unique point it fixes.

c) Observe that G is freely indecomposable with respect to the family of all its subgroups which are elliptic in Y. [8]

[6] In this particular case there is a unique way of performing such a blow-up.

[7] Note that $(\phi_n)_*(a_n)$ being defined up to conjugacy is enough for the expression above to be well defined.

[8] Consider the action of G on any tree U dual a free decomposition that would contradict this. Stabilizers of rigid vertices of T must act elliptically on U, as their image is elliptic in Y and all stabilizers of vertices in T inject in Q. Since surface groups are freely indecomposable relative to the family of their boundary subgroups, this means that $Stab(u)$ fixes a unique vertex of U for any $u \in VT$ so that the action of G on U is trivial (free factors are malnormal).

d) The tree W is 2-acylindrical.

Let Z stand for either singleton containing the point fixed by $Stab(v_0)$ (unique since the group is not abelian) or else the minimal tree of $Stab(v_0)$. For any segment $v_0, u, g \cdot v_0$ of length 2 the sets Z and $g \cdot Z$ have to intersect in the unique point fixed by $Stab(u)$. If $Z = \{x\} = Fix(Stab(v_0))$ it follows easily by induction that $Stab(v_1) = Stab(v_0)^{g^{-1}}$ fixes x for any $v_1 = g \cdot v_0$ of surface type, so that in fact every $g \in G$ does, contradicting the non-triviality of the action. Therefore $Stab(v_0)$ is not elliptic in Y. In virtue of the results of section 3.1 we know that the action of Y decomposes as a normalized tree of actions $(S, Y_{(-)}, p_{(-)})$, where each non-degenerate Y_v is either of Seifert, axial or simplicial type.

Assume for a moment that the action of $H := Stab(v_0)$ on Y were not elliptic in S. Since the peripheral subgroups of $Stab(v_0)$ are elliptic in Y, they are elliptic in the minimal tree $S_{Stab(v_0)}$ and hence there is a non-empty essential system of curves δ and an equivariant map from the tree T_δ dual to δ to S_H. In particular the conjugacy class corresponding to any of the curves in δ is mapped into an edge stabilizer of S and it must hence be elliptic in Y, contradicting hypothesis (iv) listed above, which implies $tl^{(d_{lim}, \rho_{lim})}(g) \neq 0$ for any $g \in G$. We conclude that $Stab(v_0)$ is contained in $Stab(u)$ for some component Y_w, which cannot be other than of Seifert type, the other ones being either lines or segments.

Assume that $Stab(w) \cong \pi_1(\Sigma')$ with subgroups of $Stab(w)$ elliptic in S in correspondence with peripheral subgroups of Σ'. Since boundary subgroups of $Stab(w)$ are sent to peripheral groups of $Stab(w)$, a result due to Scott [Sco78] that the inclusion of $Stab(v_0)$ into $Stab(w)$ is induced by a finite covering map $q : \Sigma \to \Sigma'$. That this is in fact a homeomorphism can be proven directly(9), but it also follows from the last property in the statement of the lemma, which is all is now left to show. If the action of $Stab(v_0)$ on its minimal tree is the pullback of an action of $\pi_1(\Sigma', *)$ on a real tree for any covering map $q : \Sigma \to \Sigma'$, pick a deck transformation ϕ_{n_0} associated to the covering map q. The fact that for any loop α in Σ we have $q \circ \phi_{n_0} \circ \alpha = q \circ \alpha$ implies that a and $(\phi_n)_*(a)$ are conjugate in $Stab(w)$ for any element $a \in Stab(v)$. In particular, this implies that $tl^{(\rho,d)}(a_n) = tl^{(\rho,d)}((\phi_n)_*(a_n))$, contradicting property (v) listed above and the continuity of translation length with respect to the Gromov topology.

Proof. (of the claim) For each $\alpha \in \delta$ chose $w_\alpha \in (0,1)$ such that $w_\alpha \neq w_\beta$ for any distinct $\alpha\beta \in \delta$ and let d^w be the distance on $|S_\delta|$ obtained by assigning length w_α to each of the edges in the orbit corresponding to α, $(|T_\gamma|, d^1)$ the standard geometric realization of T_δ and $\epsilon = 2 \min_{\alpha,\beta \in \gamma \alpha \neq \beta} |ln \frac{w_\alpha}{w_\beta}|$.

Given an automorphism σ of S we let $tl^{\lambda_\gamma^w}(g)$ stand for the translation length of $\sigma(g)$ with respect to the action of g on (T_δ, d^w). Denote by $Ell(\lambda)$ the set of elements which are elliptic for an action λ and by $Hyp(\lambda)$ the set of those which are hyperbolic. For each R let B_R stand for the ball of radius R around the identity in the Cayley graph of S with respect to some fixed finite tuple of generators. In the course of the construction, we generate a sequence $(A_n)_n$ of finite subsets of S, where $B_n \subset A_n$. Assume we have successfully constructed θ_n with the property that $A_n \cap Ell(\lambda_\delta \circ \theta_n) = \{1\}$.

Let $\bar{\theta_n}$ be some homeomorphism of Σ (fixing the basepoint $*$) which induces θ_n. Consider the essential system of curves $\delta' = \bar{\theta_n}^{-1}$. Since δ' splits Σ into small pieces and ϕ_n is not homotopic to the identity, it might behave in two different ways with respect to the deck transformation ϕ_n. The first possibility is that $i(\phi_n(\alpha), \delta') \neq 0$ for some $\alpha \in \delta'$. Given a non-trivial element a_n in the conjugacy class associated to α, this is the same as to say that $a_n \in Ell(\lambda_\delta \circ \theta_n)$ but $\phi_n(a_n) \in Hyp(\lambda_\delta \circ \theta_n)$. The other one is that $\phi_n(\delta') = \delta'$. In that

(9) As $Stab(w)$ contains $Stab(v_0)$ as a subgroup of finite index, it must be elliptic in the tree S as well. Since v_0 is the unique vertex fixed by $Stab(v_0)$ in fact $Stab(v_0) = Stab(w)$

case some $\alpha \in \delta'$ is sent by ϕ to some other $\beta \in \delta'$. Chose as a_n in this case an element with translation length 2 in $T_{\gamma'}$ (the original simplicial tree) and whose axis consists entirely of translates of e_α of $\overline{e_\alpha}$. Clearly the same property holds for $\phi_n(a_n)$ only with respect to e_β instead of e_α. In virtue of lemma 5.5.10 there is a power $\tau_n^\delta \in R$ of Dehn twists over edges of T_δ satisfying the following inequalities for any $h, g \in (Hyp(\lambda_\delta \circ \theta_n) \cap B_{n+1}) \cup A_n$:

$$|ln(\frac{(tl^{\lambda_\gamma \circ \tau_n^\delta \circ \theta_n, d^1)}(g) tl^{(\lambda_\delta \circ \theta_n, d^w)}(h)}{tl^{(\lambda_\delta \circ \theta_n, d^w)}(g) tl^{(\lambda_\gamma \circ \tau_n^\delta \circ \theta_n, d^1)}(h)})| < \frac{\epsilon}{2^{n+3}}$$

$$tl^{\lambda_\gamma \circ \tau_n^\delta \circ \theta_n}(g) > n$$

For $\eta_n = \tau_n^\delta \circ \theta_n$ this implies in particular that $Hyp(\lambda_{\delta'}) \cap A_n = Hyp(\lambda_\delta \circ \theta_n) \cap A_n \subset Hyp(\lambda_\gamma \circ \eta_n)$. On the other hand if $g \in Ell(\lambda_\delta \circ \theta_n) = Ell(\lambda_\delta \circ \eta_n)$ then $g \in Hyp(\lambda_\gamma \circ \eta_n)$, since $\delta' = \bar{\eta_n}^{-1}(\delta)$ and $\bar{\eta_n}^{-1}(\gamma)$ fill Σ. If ϕ_n is as in the first of the two possibilities discussed above, we can also take:

$$|ln(\frac{tl^{(\lambda_\gamma \circ \eta_n, d^1)}((\phi)_*(a_n))}{tl^{(\lambda_\gamma \circ \eta_n, d^1)}(a_n)})| > \epsilon$$

If the second one is the case, we can assume that

$$|ln(\frac{tl^{(\lambda_\gamma \circ \eta_n, d^1)}(a_n)}{tl^{(\lambda_\gamma \circ \eta_n, d^1)}(\phi_n)_*(a_n))})| > \epsilon$$

for $g \in \{a_n, (\phi_n)_*(a_n)\}$ as well. Let $A_{n+1} = A_n \cup B_n \cup \{a_n, (\phi_n)_*(a_n)\}$. Another application of lemma 5.5.10 yields the existence of a product τ_n^γ of Dehn twists over edges of T_γ such that if we let $\theta_{n+1} := \tau_n^\gamma \circ \eta_n$ then $Ell(\lambda_\delta \circ \theta_{n+1}) \cap A_{n+1} = \{1\}$ and:

$$|ln(\frac{tl^{(\lambda_\delta \circ \tau_n^\gamma \circ \eta_n, d^w)}(g)}{tl^{(\lambda_\delta \circ \tau_n^\gamma \circ \eta_n, d^w)}(h)})| < \frac{1}{2^{n+3}}$$

for any $h, g \in A_{n+1}$. At this point we can let $\theta_{n+1} = \tau_n^\gamma \circ \tau_n^\delta \circ \theta_n$ and iterate the process. Given any $h, g \in B_{n_0}$ for $n \geqslant n_0$ obviously $h, g \in Hyp(\lambda_\delta \circ \theta_n)$ and:

$$\sum_{n \geqslant n_0} |ln(\frac{tl^{(\lambda_\delta \circ \theta_n, d^w)}(g) tl^{(\lambda_\delta \circ \theta_{n+1}, d^w)}(h)}{tl^{(\lambda_\delta \circ \theta_n, d^w)}(h) tl^{(\lambda_\delta \circ \theta_{n+1}, d^w)}(g))|} \leqslant \sum_{n \geqslant n_0} \frac{\eta_n}{2^{n+2}} \leqslant \frac{\epsilon}{2}$$

Using a standard diagonal argument, one can obtain a subsequence $(\theta_{n_k})_k$ and any h, g not contained in a boundary subgroup of S the sequence of values $(|ln(\frac{tl^{(\lambda_\delta \circ \theta_{n_k}, d^w)}(g)}{tl^{(\lambda_\delta \circ \theta_{n_k}, d^w)}(h)})|)_n$ converges to some limit, which is greater than $\frac{\epsilon}{2}$ in case $(g, h) = (a_n, (\phi_n)_*(a_n))$. There is a small catch: we claimed to be able to take as $(|S_\delta|, d^w)$ any geometric realization of S_δ however this restriction can be easily removed by postcomposing each θ_n with a sufficiently big power of Dehn twists over δ and a big powers of Dehn twists over γ. □

This concludes the proof of 5.5.21. □

Another useful lemma is the following:

Lemma 5.5.23. *Suppose we are given a (restricted) π-tower $\mathcal{T}^{(J,r)}$, a test sequence $(f_n)_n$ of $\mathcal{T}$, a limit group L and a convergent sequence $(g_n)_n$ of homomorphisms from $G := L * T$ to $\mathbb{F}$ such that:*

- $g_n \restriction_T = f_n$

- *$lker_n\ g_n \cap G = \{1\}$*
- *T^{λ} grows faster than G under $(f_n)_n$ for any $\lambda \in J$*

Then $lker_n\ g_n = \{1\}$.

Proof. This follows from the arguments used in the proof of Merzlyakov's theorem, so we won't exercise excesive rigor in the proof. Let $N := L/lker_n\ g_n$. Let Y be the limiting tree for the sequence $(g_n)_n$ and Y_T the minimal tree of T in Y. If no non-trivial element of G fixes a point of Y_T with non-trivial stabilizer in T, then it is easy to show that N is isomorphic to the free product of G and T (see the proof of 6.1.9). Otherwise N is isomorphic to the amalgamated product G and some group B_{λ} obtained from T by enlarging T^{λ} for some $\lambda \prec r$. By induction we can write B_{λ} as the free product of $G \cap B_{\lambda}$ and T^{λ}. This easily implies that N is the free product of T and G. □

Corollary 5.5.24. *Suppose we are given a π-tower structure $\mathcal{T}_A^{(J,r)}$, where r is of floor type and $(\Delta\mathcal{T})^r$ contains a unique vertex of non-rigid type, which is of surface and that for any $\mu \prec r$ we are given a family $\mathcal{TS}^{\mu}$ of (relative) π-test sequences $\mathcal{T}_A\upharpoonright_{\mu}$, closed under taking subsequences and diagonal subsequences and such that $\mathsf{A}_{S^{\mu}}$ is elliptic in any limit tree of a sequence extending both $(f_n^1)_n$ and $(f_n^{\mu_2}\upharpoonright_{\mathsf{T}^{\nu}})$ for $\mu_1 \lhd \mu_2$, $(f_n^j)_n \in \mathcal{TS}^{\mu_j}$ and $\nu \leqslant \mu$. Then, given any tuple $((f_n^{\mu})_n)_{\mu} \in \times_{\mu \prec r}\mathcal{TS}^{\mu}$ there is a π-test sequence $(f_n)_n$ of f_n extending $\amalg_{f_n^{\mu}} \mu \prec r$ and such that the collection $\mathcal{TS}$ of all of them is closed under taking subsequences and diagonal sequences. Given any fixed increasing sequence $(\kappa_n)_n$ of positive constants we can require that $\frac{tl^r(f_n(g))}{\kappa_n}$ tends to infinity with n for any $(f_n)_n \in \mathcal{TS}$ and some $g \in G$.*

Proof. Let δ be the essential system of curves provided by lemma 5.5.18, applied to the unique surface group of $\Delta = (\Delta\mathcal{T})^r$, isomorphic to $\pi_1(\Sigma)$, and $\eta := \eta^r_{\mathcal{R}}$. Take γ a essential system of curves on Σ which together with δ fills Σ. Let $(\theta_n)_n \subset Mod^{\pi}(\mathsf{G}, S)$ and $(Y_{lim}, d_{lim}, *_{lim}, \rho_{lim})$ as provided by the last lemma. Let $\mathcal{TS}_{\prime}$ be the family of sequences obtained by applying 5.5.2 to $(\mathcal{TS}^{\mu})_{\mu \in Ch(r)}$, in a way that makes T^{μ} grow faster than T^r in case $\mu \prec r$. In virtue of 5.5.23, we know that $lker_n\ g_n$ for all $(g_n)_n \in \mathcal{TS}_0$. Let $\mathcal{UTS}_0$ be the collection of all the $f \in Hom(G, \mathbb{F})$ appearing as f_n for some $(f_n)_n \in \mathcal{TS}_0$. For any $f \in \mathcal{UTS}_0$ denote by N_f the maximum value of n such that $f = inn_c \circ f_n$ for for some $c \in \mathbb{F}$ and $(f_n)_n \in \mathcal{F}$. Let $A_0 \subset A_1 \subset A_2 \cdots A_n \cdots$ some exhaustion of G. For each n let $B_n \subset G\backslash\{1\}$ be the set of S-elliptic elements obtained by applying lemma 5.5.10 to $\sigma_n(A_n)$ and k_n a positive integer such that $ker\, f_m \cap B_n = \varnothing$ for any $m \geqslant k_n$ and $(f_n)_n \in \mathcal{F}$. Pick any basepoint $*'$ in the Cayley graph of $\mathbb{F}$. Given $f \in \mathcal{UTS}_0$ let $u_f := max\,\{n \in \mathbb{N} \,|\, k_n \leqslant N_f\}$, so that $ker\, f \cap B_{u_f} = \varnothing$. In virtue of lemma 5.5.10 we know that there is some $\tau_f \in Mod^{\pi}(\mathsf{G}, S)$ and $\nu_f \in \mathbb{R}^+$ such that

$$|\frac{1}{\nu_f \mu_{u_f}} l_{*'}^{r\circ f\circ\tau_f}(\sigma_{u_f}(g)) - \frac{1}{\mu_{u_f}} l_*^{\rho}(\sigma_{u_f}(g))| \leqslant |l_{*_{lim}}^{\rho_{lim}}(g) - \frac{1}{\mu_{u_f}} l_*^{\rho}(\sigma_{u_f}(g))|$$

for every $g \in B_{u_f}$, where r is the action of $\mathbb{F}$ on its Cayley graph. We claim that the family $\mathcal{TS}$ consisting of all the sequences of the form $(g_n \circ \tau_{g_n})_n$, where $(g_n)_n \in \mathcal{TS}_0$ satisfies the desired properties. First of all, the family clearly inherits being closed under subsequences and taking diagonal sequences from $\mathcal{TS}_0$. Now, given some $(f_n)_n \in \mathcal{TS}_0$ and $g \in B_m$ for any $n \geqslant k_m$ we have

$$|\frac{1}{\nu_{f_n} \mu_{u_{f_n}}} l_{*'}^{r\circ f_n\circ\tau_{f_n}\circ\sigma_{u_{f_n}}}(g) - l_{*_{lim}}^{\rho_{lim}}(g)| \leqslant 2|l_{*_{lim}}^{\rho_{lim}}(g) - \frac{1}{\mu_{u_{f_n}}} l_*^{\rho_{lim}}(\sigma_{u_{f_n}}(g))|$$

since u_{f_n} tends to infinity with n, the right hand side tends to 0, which implies the sequence $(f_n \circ \tau_{f_n} \circ \sigma_{f_n})_n$ has trivial limit kernel and admits Y_{lim} as a limiting tree, which implies it satisfies the test sequence condition at r. □

Note that in the theorem above, if each $\mathcal{TS}^{\mu}$ is congruency complete, the same is true for $\mathcal{TS}$.

5.5.2 The induction step: abelian vertex groups

Lemma 5.5.25. *Let G be a group of the form $L *_M (M \oplus \langle t \rangle)$, where M is a non-trivial free abelian group and suppose we are given a sequence $(f_n)_n \neq Hom(G, \mathbb{F})$ with trivial limit kernel. Let x be some finite set of generators of G. Suppose that for each n we are given an extension $g_n : G \to \mathbb{F}$ of f_n such that $tl(g_n(t)) \geqslant n \cdot tl(g_n(x_i x_j))$ for any $1 \leqslant i < j \leqslant |x|$. Then then any limiting tree for the sequence $(g_n)_n$ is isomorphic to a geometric realization of the simplicial tree S dual to an HNN extension with vertex group L and Bass-Serre element t.*

Proof. This can be proven directly without any advanced technology using the same techniques of the proof of 5.5.10. But it also follows straightforwardly from the fact that L is elliptic and t (its image in the limit quotient of the sequence g_n) hyperbolic in any limiting tree Y, together with all what we know about the action on Y. In particular, the point-wise stabilizer of $Ax(t)$ coincides with the centralizer of t in G. □

Corollary 5.5.26. *Let $\mathcal{T}_A^{(J,r)}$ a refined π-tower where r falls in case (A). If there is a congruency complete CSD family $\mathcal{TS}_\lambda$ of π-test sequences of $\mathcal{T}_A \restriction_\lambda$, where λ is the unique child of r, then there is a congruency complete CSD family of π-test sequences of $\mathcal{T}_A$.*

Proof. Let $M \oplus \langle t \rangle$ be the abelian type vertex group of $(\Delta\mathcal{R})^r$, where $M \leqslant \mathsf{T}^\lambda =: L$ and x some finite tuple of generators of L. As $\mathcal{TS}$ we take the family of sequences $(g_n)_n$ of morphisms factoring through $\mathcal{T}_A$ such that $g_n \restriction_L \in \mathcal{TS}$ and $tl(g_n(t)) \geqslant n \cdot tl(f_n(x_i x_j))$ for $1 \leqslant i < j \leqslant |x|$. Any such sequence is a π-test sequence, by the previous lemma, and the resulting family is clearly congruency complete. □

5.6 Characterization of restricted π-limit groups

The most basic examples of limit groups, other than free groups, are finitely generated abelian groups and closed surface groups. It turns out that the family of all π-structures on such a group making it into π-limit groups admits a simple description. The following is a consequence of observation 5.5.4 below.

Lemma 5.6.1. *Let H, H' be π-groups and suppose that H, H' are non-abelian free groups. Then there is an injective morphism from H to H'. Let $\mathsf{L}_A = \mathsf{A} * \mathsf{H}$, where H is free. Then L_A is a restricted π-limit group.*

Corollary 5.6.2. *The notion of a π-limit group depends only on Q and not on the particular epimorphism $\pi : \mathbb{F} \to Q$.*

Lemma 5.6.3. *Let M a π-group whose underlying group M is free abelian. Then M is a restricted π-group if and only if $\pi(M) \leqslant C$ for some cyclic $C \leqslant \mathbb{F}$.*

Proof. The implication from left to right is clear. For the opposite direction, take $x \in M$ such that $\pi(x)$ generates $\pi(M)$ in Q. We can assume $Z = \langle x \rangle$ is root closed in M. Write M as $Z \oplus N$ and let $y_1, \cdots y_m$ a base for N. Pick $K \in \mathbb{N}$ dividing the order of $\pi(C)$ in Q. Fix some tuple $\bar{r} = (r_i)_{i=1}^m \in \{0 \cdots K\}$ of distinct integers such that $\pi(y_i) = x^{r_i}$. For each $n \in \mathbb{N}$ consider the π-retraction f_n of M onto Z sending y_i to $x^{e_{i,n}}$, where $e_{i,n} = K(n+i)! + r_i$. It is easy to check that given $0 \leqslant i_1 < i_2 \cdots < i_r \leqslant m$ and non-zero integers $k_1, k_2 \cdots k_r$ the last term of the sum $k_1 e_{i_1,n} + k_2 e_{i_2,n} \cdots k_r e_{i_r,n}$ grows in absolute value faster with n than all others and that therefore the total number is non-trivial for n big enough. This easily implies taht the limit kernel of the f_n is trivial. Since Z is itself a π-limit group we are done. □

Lemma 5.6.4. *Let* S *be a* π*-group whose underlying group* S *is isomorphic the fundamental group of a closed surface* Σ *of Euler characteristic* $\geqslant -2$*. Then* S *is a* π*-limit group if and only if for some maximal essential system of curves* δ *in* Σ *each of the root closure of the conjugacy classes in* S *in correspondence with the curves is killed by the homomorphism* π*.*

Proof. For the 'only if' all is needed is the fact that $\pi_S = \pi \circ h$ for some homomorphism h from S to $\mathbb{F}$. In virtue of 2.4.3, h can be written as $h' \circ p$, where $p = q \circ pinch_\delta$ for an essential system of curves δ in Σ and q kills the fundamental group of any projective plane in $Pinch(\Sigma, \delta)$.

Notice that the image of p is the fundamental group of an finite graph, hence free. For the opposite direction notice that viceversa, the fact that some δ satisfies the property of the statement implies π_S factors through the quotient p above.

Now, by virtue of the discussion in 2.4.2 the image of p is a free group F and the assumption that $\xi(\Sigma) \leqslant -2$ implies it is non-abelian. Push the π-structure on S forward onto F. The result follows from 5.6.2 together with the existence of π-test sequences (the π-tower in this case consists of a unique vertex group of surface type). ☐

Lemma 5.6.5. *Let* G_A *be a restricted* π*-group,* S *a geometric abelian decomposition* Δ *of* G_A*. Then any morphism* f *from* G_A *to a restricted* π*-limit group* H_A *which is formally strict with respect to* S*. is strict (in particular* G_A *is also a restricted* π*-limit group).*

Proof. Let $(h_n)_n$ be a sequence of morphisms from H_A to F_A with trivial limit kernel. To begin with, let G' be the retract of G obtained by replacing any abelian type vertex group of Δ by its extended peripheral subgroup. Corollary 5.5.13, together with the fact that f is formally strict with respect to S implies the existence of a sequence of automorphisms $(\tau_n)_n \supset Mod^\pi(\mathsf{G}, S)$ such that the limit kernel of $(h_n \circ f \circ \tau_n)$ has trivial intersection with G' (up to subsequence extraction we can assume that this sequence is convergent).

Now, up to the fact that the image of an abelian vertex group of Δ by h_n might be larger than that of its extended peripheral subgroup, the result is now a consequence of regarding G_A as a π-tower over $\mathsf{G'}_A$. A quick look at the proof of case (A) of the construction of π-test suffices to check that this is not really an obstruction. ☐

Corollary 5.6.6. *A restricted* π*-group* G_A *is am* A*-restricted* π*-limit group if and only if it admits a strict closed weak restricted* π*-resolution .*

The latter can be also seen as a direct consequence of proposition 5.5.1 together with the embedding constructed in section 5.3.

Chapter 6

Merzlyakov-type theorems

6.1 π-towers and formal solutions

In this section we provide a generalization of the following theorem, originally due to Merzlyakov:

Theorem 6.1.1. *Let $\mathbb{F}$ be a free group and $\Sigma(x, y, a) = 1$ a finite system of equations in the tuple x over a tuple a of constants from $\mathbb{F}$. Assume that the group generated by a in $\mathbb{F}$ is a free factor of $\mathbb{F}$ and suppose that:*

$$\mathbb{F} \models \forall x \exists y \;\; \Sigma(x, a) = 1$$

Then there is a word $w(x, a)$ such that the equality $\Sigma(x, w(x, a), a) = 1$ holds in $\langle x \rangle * \langle a \rangle$.

We would like to remark that the result does not hold in case $\langle a \rangle$ is not a free factor. For a trivial example of how this can fail if A take $\mathbb{F} = \mathbb{F}(\alpha, \beta)$, $a = (\alpha, \alpha^{\beta})$, x and y single variables and our system of equations simply:

$$\begin{cases} x = 1 \\ a_1^y = a_2 \end{cases}$$

The corresponding $\forall\exists$ sentence is clearly true, but there is no way that some word $w(x, a)$ as above can exist (in this case really $w(a)$) since a_1 and a_2 are not conjugate in $\langle a_1, a_2 \rangle$. One refers to the word w as a *formal solution* for the sentence above. Observe that the existence of such a word is equivalent to that of a retraction from the group G_Σ to the subgroup generated by x and a.

Merzlyakov theorem can be generalized to finite disjunctions of systems of equations. If one lets systems of inequations enter the hypothesis, though, the picture becomes much more complex. Indeed, given a sentence with constants of the form:

$$\forall x \exists y \, \Sigma(x, y, a) = 1 \wedge \Pi(x, y, a) \neq 1$$

which is valid in $\mathbb{F}$, in general there is no way of finding an algebraic expression $w(x, a)$ such that $\Sigma(x, w(x, a), a) = 1 \wedge \Pi(x, w(x, a), a) \neq 1$ holds for any possible value of x. In case that $\mathbb{F}(x) * \langle a \rangle$ is non abelian one can guarantee however that $\Pi(x_0, w(x_0, a), a) \neq 1$ for at least some value x_0 of x, in oder words, that the system $\Sigma(x, a) = 1$ does not imply $\Pi(x, w(x, a), a) = 1$ in $\mathbb{F}$ and hence the set of those values of x for which the condition fails is a proper variety of $\mathbb{F}^{|x|}$.

The natural way to proceed is trying to construct an ad-hoc formal solution that works in the subvariety of the points where first formal solution fails, except for those points in a yet smaller proper variety. Equational noetherianity would then imply that the process eventually stops, yielding a piecewise algebraic function that witnesses the validity of the sentence.

The first, but inessential, complication lies in the fact that in this case there is no single formal solution that can cover all values of x in the specified variety. The main obstacle is however, that these more specialized formal solutions cannot be found in terms of x and a alone. They necessarily involve some new tuple t of variables, which are linked to x and a by certain relationships. And in that case, knowing the failure of the formal solution of a particular (x_0, t_0) might fail to give any new information on x_0 in isolation, breaking the previous argument. Roughly speaking, the way out of this is to show that the relationships linking x are t are tight enough for a certain notion of complexity to properly decrease from one step to the next of the procedure, when adequately refined.

Before being able to tackle such difficulties, a subject that unfortunately exceeds the scope of this dissertation, one needs to adapt the generalization of Merzlyakov theorem from [Sel03]), which handles $\forall\exists$ formulas with algebraic conditions on universal variables to our particular context.

The class of sentences dealt with in the next theorem are almost general $\forall\exists$ sentences with constants, as any $\forall\exists$ sentence can be written as a finite conjunction of formulas of that type.

Theorem 6.1.2. *Suppose that* F *satisfies the sentence with constants:*

$$\forall x\,((x \in q^{q^x} \wedge \Delta(x,a) = 1) \to \exists y \in q^y \wedge (\bigvee_{j=1}^{r} (y \in q^y(j) \wedge \Sigma(x,y,a) = 1 \wedge \Pi(x,y,a) \neq 1)))$$

Then there is a tuple of variables s*, some* $m \geqslant 1$ *and for any* $1 \leqslant l \leqslant m$*:*

i) A system of equations $\Lambda_l(x,t,a) = 1$

ii) Some $q^s(l) \in Q^{|s|}$.

iii) A $|y|$*-tuple* $w^l(x,s,a)$ *of words in* x, s, a

iv) Some $k_l \in \{1, \cdots r\}$

so that the following list of conditions are satisfied:

i) The following holds in any group containing A*:*

$$\forall x\,\forall s\,(\Delta(x,a) = 1 \wedge \Lambda_l(x,s,a) = 1) \to \Sigma_{k_l}(x, w_l(x,s,a), a) = 1))$$

ii) The following equality holds in Q *for* $1 \leqslant l \leqslant m$*:*

$$w^l(q^x, q^s(l), q^a) = q^y(k_l)$$

iii) The following holds in F*:*

$$\forall x\,\exists s\,(\bigvee_{i=1}^{m} s \in q^s(l) \wedge \Lambda_l(x,s,a) = 1)$$

iv) For all $1 \leqslant l \leqslant s$ *there is* $(x_0, s_0) \in \mathbb{F}^{|s|+|x_0|}$*:*

$$x_0 \in q^x \wedge s_0 \in q^s(l) \wedge \Delta(x_0,a) = 1 \wedge \Lambda_l(x_0,s,a) = 1 \wedge \Pi(x_0, w^l(x_0,s_0),a) \neq 1$$

If $\Delta(x,a) = 1$ is empty and $\langle \pi(a), q^x \rangle = Q$, then all conditions except possibly the last one can be attained in the case in which s the empty tuple. If in addition to this the free product of $\mathbb{F}(x) * \langle a \rangle$ is non-abelian, then the last condition can be achieved for $s = \varnothing$ as well.

This result illustrates neatly what the obstruction to a full generalization of Merzlyakov's theorem for sentences with algebraically constrained universal variables is, but as it stands it is not precise enough to be of use for the ultimate purpose of analyzing all $\forall\exists$ sentences with parameters. In order to achieve this goal one needs to keep track of a greater amount on the information provided by the proof. This will be phrased in the language of π-groups morphisms and resolutions introduced in chapter 4.

Let $\mathsf{G} = \mathsf{G}^{\Delta}$. We know that there is a finite family $\mathcal{MR}^{ws}$ of strict well-separated restricted π-resolution and for each $\mathcal{R}_A \in \mathcal{MR}^{ws}$ some surjective morphism q^R from G_A to its top group L_A, so that each $f \in Mor(\mathsf{L}_A, \mathsf{F}_A)$ can be written as $\bar{f} \circ q^R$, for some $\mathcal{R}_A \in \mathcal{MR}^{ws}$, and some $\bar{f}$ which factors through $\mathcal{R}_A$. To each $\mathcal{R}_A^{(J,r)} \in \mathcal{MR}^{ws}$ we can associate a completion $\mathcal{T}_A^{(J,r)} = Compl(\mathcal{R})$, as constructed in 5.3 and a morphism which embeds R_A^r into T_A^r.

Let H be a free π-group such that $\langle \pi(\mathsf{G}_A), \pi(H) \rangle = Q$ and t be a tuple of variables, such that for each of the completions T_A obtained above the group $T \ast H$ admits a generating t-marking. We can assume that $x \subset t$.

We claim that the theorem above follows from the following proposition:

Proposition 6.1.3. *Let $\mathcal{T}_A$ be a closed restricted π-tower structure for a group $\mathsf{T}_A(t)$, where A is a free factor of $\mathbb{F}$. Suppose that any solution t_0 of $\mathsf{T}_A(t)$ satisfies:*

$$\exists y \bigvee_{j=1}^{s} (y \in q(j) \wedge \Sigma_j(t_0, y, a) = 1 \wedge \Pi_j(t_0, y, a) \neq 1)$$

for a system of equalities $\Sigma_j(t,a) = 1$ and a system of inequalities $\Pi_j(x,a) = 1$ and tuples $q(j) \in Q^{|y|}$. Assume moreover that for all $1 \leqslant j \leqslant s$ the tuple $q(j)$ is contained in $\pi(T) \leqslant Q$.

Then there is a covering system of closures $\mathcal{CL}$ of $\mathcal{T}_A$ and for each $\mathcal{S}_A \in \mathcal{CL}$ some retraction:

$$r^S : \mathsf{S}_A \ast \mathbb{F}(y) \to \mathsf{S}_A \ast \mathsf{H}$$

and some $1 \leqslant j_S \leqslant s$ such that $r^S(y) = q(j_S)$ and $\Sigma_{j_S}(x, r^S(y), a) = 1$.

If in addition T is non-abelian then r_j can be chosen to preserve the validity of the system $\Pi_{j_S}(x,y,a) \neq 1$.

Indeed, we can always take a free π-group H mapping onto Q and replace the completion of each of the resolutions in $\mathcal{MR}^{ws}$ by its product with H, which can be given a π-tower structure extending that of T in an obvious way. Apply the proposition to them and take a tuple s of variables such that each of the resulting closures admits a (generating) s marking. Any solution of $\mathsf{G}_A(x)$ extends to one which factors through at least one of the completions above, and hence to one factoring through one of the closures [(1)], a property which translates into property (iii) of the theorem; each Δ_l appearing in the statement corresponds to the presentation of one of those closures and each $q(l)$ to the image in Q of the associated tuple of generators by the corresponding π-map. If we let $w_l(s,a)$ be the expression of $r^S(y)$ in terms of s, a, items (i) and (ii) follow immediately.

For the last item, observe that since each of the closures S_A are themselves restricted π-limit groups, if the system of inequalities $\Pi_l(x, w(x,t,a), a) \neq 1$ holds in S_A, then necessarily $\Pi_l(x_0, w(x_0,t_0,a), a) \neq 1$ for some particular solution (x_0, t_0) of S_A, since the latter is a restricted π-limit group.

[(1)] It is important that it factors through the π-tower, not any morphism from the π-tower group T_A to F_A extends to a closure.

6.1.1 Formal Makanin-Razborov diagrams: proving propostion 6.1.3

For each $1 \leqslant l \leqslant r$ and $q \in pgq(T)$ let $(\mathsf{K}^{\mathsf{q}}_{\mathsf{l}})_A$ be the quotient of $\mathsf{T}_A * \mathbb{F}(y)$ by the relation given by the system $\Sigma_l(t, y, a) = 1$, where $\pi(y) = q(l)$.

We can assume for any l this assignment induces an extension of the map $\pi_{\hat{T}}$, since otherwise there are not any tuples (t_0, y_0) in $\mathbb{F}$ satisfying both $\Sigma_l(t_0, y_0, a) = 1$ and $y \in q(l)$, and we can remove the corresponding term from the disjunction. Likewise, we can assume that T injects in the quotient obtained this way.

By a (q, Σ_l)-formal π-group we intend a free factor containing $\hat{T}$ of a π-quotient of $(\mathsf{K}^{\mathsf{q}}_{\mathsf{l}})_A$ onto which T maps injectively and by.

Fix some family $\mathcal{TS}$ of test sequences of $\mathcal{R}_A$. By a formal sequence of morphisms from $\mathsf{L}_A(z)$ a sequence $(g_n)_n$ of morphisms from L_A to F such that $(g_n\restriction_T)_n \in \mathcal{TS}$. We refer to it as merely 'formal' if when not wanting to specify the pair (q, Σ_l).

Given a π-resolution $\mathcal{R}^{(J,r)}$ of a marked restricted π-group $\mathsf{G}_A(x)$ and a system of inequalities $\Sigma(x, a) \neq 1$, with $a \subset A$, we say that $\mathcal{R}$ preserves the validity of $\Sigma(x, a) \neq 1$ if the map from G_A to $*_{\lambda \in lv(J)} \mathsf{R}^{\lambda}$ induced by the quotient maps along the resolution does.

Lemma 6.1.4. *For any π-limit group L there is some free subgroup $H \leqslant L$ such that*

$$\pi(H) = \pi(G)$$

Proof. We will use the following well-known fact:

Claim 6.1.5. *Given an isometric action of a group G on a real tree Y and elements $g_0, g_1, \cdots g_m$ of G acting hyperbolically on Y and such that $Ax(g_i) \cap Ax(g_i)$ has bounded diameter for $0 \leqslant i < j < m$, there is a positive integers N such that for any choice of $k_i \geqslant N$ for $0 \leqslant i \leqslant m$ the set $\{g_i^{k_i}\}_{0 \leqslant i \leqslant m}$ generates a free group, of which it is a base.*

Proof. (sketch) The boundary of Y, denoted by ∂Y is a Hausdorff space whose underlying set is the family of all infinite geodesic rays in Y, quotiented by the equivalence relationship that identifies two such rays with infinite intersection. The action of G on ∂Y induces an action of G on δY. Any bi-infinite geodesic α is associated with a pair of endpoints $\partial \alpha \subset \delta Y$. Given any hyperbolic element h of G the pair $\partial Ax(h)$ is of the form $\{e^+, e^-\}$, with the property that given any neighbourhoods U^+ and U^- of e^+ and e^- respectively there is some positive $N > 0$ such that $h^k \cdot (\delta Y \backslash U^-) \subset U^+$ for any $k \geqslant N$ and $h^k \cdot (\delta Y \backslash U^+) \subset U^-$ for any $k \leqslant -N$.

Now, the fact that $Ax(h_i) \cap Ax(h_j)$ is bounded for $i \neq j$ implies that the intersection of $\partial Ax(h_i)$ and $\partial Ax(h_j)$ is empty.

For each $0 \leqslant i \leqslant m$ one can chose neighbourhoods U_i^+, U_i^- of the endpoints of $Ax(h_i)$, as above, in a way that $\{U_i^{\pm}\}_{0 \leqslant i \leqslant m}$ are mutually disjoint. If we let $A_i = U_i^+ \cup U_i^-$ then there is $N > 0$ such that for any choice of $k_j \geqslant N$, and $l \in \mathbb{Z} \backslash \{0\}$ we have $h_i^{k_i l} \cdot A_j \subset A_i$ for any $j \neq i, 0 \leqslant j \leqslant m$. An application of the ping-pong lemma yields the desired conclusion. □

Let Y be the limiting tree for some sequence of morphisms from L to F with trivial limit kernel and let $\mathcal{X}$ be any finite set of generators of L. Perform the following two operations until none of them is possible:

- Suppose $x \in \mathcal{X}$ is elliptic and let h be hyperbolic. If xh is elliptic then $Fix(x)$ must intrsect the axis of h (indeed, consider the segment $I = [p, xh^k \cdot p]$ for some p fixed by xh). Now, $Fix(x)$ is linear; since our action is irreducible, this implies the existence of some element h hyperbolic in Y such that $Fix(x) \cap Ax(y) = \varnothing$, so that xh is hyperbolic. Up to replacing h with a proper power we can assume $\pi(h) = 1$. Replace x by xh for some such h.
- If two elements $x, y \in \mathcal{X}$ commute, replace $\{x, y\}$ by some $z \in Z_L(x)$ such that $\pi(z)$ generates $\pi(Z_L(x))$.

We end up with some set $\mathcal{X} \subset L$ of hyperbolic elements such that $\langle \pi(\mathcal{X}) \rangle = \pi(L)$ and $[x, y] \neq 1$ for any distinct $x, y \in \mathcal{X}$, which implies that $Ax(x) \cap Ax(y)$ has bounded diameter (see 3.1.16). On the other hand any element of L has arbitrarily big powers with the same image in Q, so the sought for result follows from the fact above. □

Proposition 6.1.6. *Suppose that $\mathcal{TS}$ is congruency complete and closed under subsequences and diagonal subsequences.*

Then for each $1 \leqslant j \leqslant s$ and $q \in pgq(T)$ a set $\mathcal{FMR}_l^q$ of finitely many closed graded π-resolution of $(\mathsf{K}_\mathsf{l}^\mathsf{q})_A$ with respect to the parameter group $\hat{T}$ exists such that for any formal sequence $(g_n)_n$ of morphisms from $(\mathsf{K}_\mathsf{l}^\mathsf{q})_A$ to F_A with the following property.

Given any formal sequence $(g_n)_n$, if $t_n = g_n(t)$ and $y_n = g_n(t)$ satisfy

$$y_n \in q(l) \wedge \Sigma_l(t_n, y_n, a) = 1 \wedge \Pi_l(t_n, y_n, a) \neq 1$$

for n big enough, then all the members of some infinite subsequence of $(g_n)_n$ factor through some π-resolution in $\mathcal{FMR}_l^q$.

Moreover, for any $\mathcal{T}^{(J,r)} \in \mathcal{FMR}_l^q$, for any $\lambda \in \tilde{J}$ either R^λ is free or a closure of $\mathcal{T}_A^q$ and the morphism $\psi : \mathsf{K}_{\mathsf{l}\,A}^\mathsf{q} \to *_{\lambda \in lv(J)} \mathsf{R}^\lambda$ induced by the projections along $\mathcal{R}$ preserves the validity of $\Pi_l(t, y, a) \neq 1$.

Proof. (of 6.1.3 from proposition 6.1.6) Let $\mathcal{R}_J \in \mathcal{FMR}_l^q$ and λ_0 the leaf of J carrying $\hat{T}$.

Since $\pi(K_l^q) = \pi(T)$ and R^λ for is free for $\lambda \in lv(J) \backslash \{\lambda_0\}$, there must be some π-retraction $\bar{r} : \mathsf{S} * (*_{\lambda \in lv(J) \lambda \neq \lambda_0} \mathsf{R}^\lambda) \to \mathsf{S}$ sending R^λ into T for $\lambda \in lv(J)$ into If in addiiton T is non-abelian, then in virtue of lemma 6.1.4 there is some $H \leqslant \hat{T}$ such that $\pi(H) = \pi(\hat{T})$; it follows from lemma 5.6.1 that there is some π-retraction from $\mathsf{R}_A^{\lambda_0} * (*_{\lambda \in lv(J) \lambda \neq \lambda_0} \mathsf{R}^\lambda)$ to $\mathsf{R}_A^{\lambda_0}$ preserving the validity of $\Psi(t, \psi(y), a) \neq 1$ and sending $*_{\lambda \in lv(J) \backslash \{\lambda_0\}} \mathsf{R}^\lambda$ to T.

Now, $\mathsf{U}_A := \mathsf{R}_A^{\lambda_0}$ has the structure of a closure of $\mathcal{T}_A^q$, which is the result of adding a peg to the original tower π-resolution $\mathcal{T}_A$. Given $M \in \mathcal{FHM}(T)$, contained in $\tilde{M} \in \mathcal{FHM}(S)$, we can write $\tilde{M}$ as $\bar{M} \oplus Z'_M$, where $\bar{M}$ is the root closure of $Z \oplus M$.

Of course M has finite index in $\tilde{M}$, so using lemma 2.3.1 one can easily show that U admits a star shaped graph of groups decomposition with a closure of S_A in its center and for any $M \in \mathcal{FHM}(T)$, a vertex group $\tilde{M}$ amalgamated to it over $\bar{M}$.

We know U_A is a limit of a sequence of morphisms which eventually send p_M to a primitive element of $\mathbb{F}$, so in particular $\pi(\tilde{M}) = \pi(M) = \langle \pi(pg_M) \rangle$, so there is a a π-retraction from U to S; existence of test sequences in case (A) implies we can assume that the latter preserves the validity of the system $\Psi(t, \psi(y), a) \neq 1$.

Composing all the maps above we obtain a homorphism $r^\mathcal{R}$ from $\mathsf{T}_A * \langle y \rangle$ to S restricting to the identity on T. The fact that it can be written as $r' \circ h$, where r' is a morphism from $(\mathsf{K}_l^q)_A$ to S_A implies that $\pi(r^\mathcal{R}(y)) = q(l)$, and $\Sigma_l(x, r^S(y), a) = 1$.

The discussion above implies that the system $\Pi_l(t, r^\mathcal{R}(y), a) \neq 1$ is also valid.

To conclude the proof, notice that the fact that the collection of all the closures $\mathcal{S}_A$ associated to the resolutions in $\bigcup_{\substack{1 \leqslant l \leqslant s \\ q \in pgq(T)}} \mathcal{FMR}_l^q$ constitute a covering system follows from the fact that $\mathcal{TS}$ is congruency complete and each $(f_n)_n \in \mathcal{TS}$ contains an infinite subsequence factoring through $\mathcal{R} \in \mathcal{FMR}_l^q$ for some $1 \leqslant l \leqslant s$ and $q \in pgq(T)$ which automatically implies it extends to the closure associated with $\mathcal{R}$. □

6.1.2 Proof of proposition 6.1.6

Definition 6.1.6.1. Let L_A be a formal restricted π-limit group, $J' := nDFt(J^*)$ and $J_0 \subset J'$. Let $\bar{J}_0$ stand for $J_0 \cup \bigcup_{\lambda \in J_0} Ch(\lambda)$ and let $n(J_0)$ stand for the uppermost node of $J' \backslash J_0$ with respect to the linear order $\lhd$.

By a J_0-tight structure for $\mathcal{R}_A$ we intend a collection:

$$((H_\lambda)_{\lambda\in J_0\cup\{n(J_0)\}}, (S^\kappa)_{\lambda\in J_0}, (\gamma_\lambda)_{\lambda\in\bar{J}_0\backslash J_0}, (v_\lambda)_{\lambda\in\bar{J}_0\backslash\{r\}})$$

where $H_\lambda \leqslant L$, $\gamma_\lambda \in L$ and S^λ is a simplicial H_λ-tree such that the following properties are satisfied, where we extend the notation by letting $\gamma_\lambda = 1$ for $\lambda \in J_0$ and $K_\lambda := (\mathsf{T}^\lambda)^{\gamma_\lambda}$ for $\lambda \in \bar{J}_0$:

i) For any $\mu \in \bar{J}_0$, K_μ is contained in $H_{c(\mu)}$, where $c(\mu) \leqslant \lambda$ is the minimal node of J_0 into which T^μ can be conjugated.

ii) For $\lambda \in \bar{J}_0\backslash\{r\}$ the group K_λ stabilizes a vertex v_λ of $U^{c(\lambda)}$ of rigid type in case $c(\mu)$ is of floor type.

iii) $H_r = L$ and $H_\lambda = Stab_{H_{p(\lambda)}}(v_\lambda)$ for $\lambda \in J_0 \cup \{n(J_0)\}\backslash\{r\}$

iv) For distinct $\mu_1, \mu_2 \in J_0 \cap Ch(\lambda)$ the vertices v_{μ_1} and v_{μ_2} are in distinct orbits by the action of H_λ.

v) If $\lambda \in J_0$ falls in case (FP), the simplicial H_λ-tree S^λ has trivially stabilized edges. The minimal tree U^λ of K_λ in S^λ is equivariantly isomorphic to the tree dual to the free product decomposition $\mathsf{K}_\lambda = (\mathsf{T}^{\mu_1})^{\gamma_\lambda} * (\mathsf{T}^{\mu_2})^{\gamma_\lambda}$ where $\{\mu_1, \mu_2\} = Ch(\lambda)$ and its translates cover S^λ.

vi) For any $\lambda \in J_0$ is of floor type:

   i) The tree dual to $(\Delta\mathcal{T})^\lambda$, on which K_λ acts via conjugation by γ_λ^{-1}, is equivariantly isomorphic to some subtree U^λ of S^λ and the translates of its image cover S^λ.

   ii) The stabilizer in H_λ of the non-rigid type vertices of U^λ is contained in K_λ in case (S). In case (A) given a vertex $v \in U^\lambda$ if we let P be the peripheral subgroup of $Stab_{H_\lambda}(v)$ then $Stab_{K_\lambda}(v) + P$ has finite index in $Stab_{H_\lambda}(v)$.

vii) If $\lambda \in J_0$ falls in case (FG), then $K_\lambda = H_\lambda$

Observation 6.1.7. For $\lambda \in \bar{J}_0\backslash J_0$ the minimality of $c(\lambda)$ implies that v_λ is never a translate of v_μ for any $\mu \in Ch(c(\lambda)) \cap J_0$. In particular $n(J_0) \lhd c(\lambda)$ for any $\lambda \in \bar{J}_0\backslash J_0$, since otherwise $Ch(c(\lambda)) \subset J_0$. This implies that $n(J_0) \succ c(n(J_0))$.

Observation 6.1.8. For a formal π-limit group L_A admitting J'$-$*tight* structure means precisely admitting the structure of a closure of $\mathcal{T}_A^q$.

Proof. All that is left to prove is the existence of a suitable retraction in case (A). This follows right away from lemma 5.4.4 will eventually send each peg of $\hat{T}$ to a primitive element. ☐

Arguing in exactly the same way as it was done in the construction of Makanin-Razborov diagrams, proposition 6.1.6 can be easily reduced to the following:

Lemma 6.1.9. *Let $\mathsf{L}_A(z)$ a formal limit group admitting a J_0-tight structure and $\Psi(z,a) \neq 1$ a system of inequalities true in $\mathsf{L}_A(x)$, where $A = \langle a\rangle$. Suppose that $J_0 \neq J'$. Then there is a finite collection $\mathcal{FMR}^q_{\lambda,L}$ of closed T-graded resolutions of L_A such that from any formal sequence $(g_n)_n$ of L_A preserving the validity of $\Psi(z,a) \neq 1$ one can extract a subsequence factoring through one of $\mathcal{R}_A^{(I,r)} \in \mathcal{FMR}^q_{\lambda_0,L}$. Moreover, for any $\mathcal{R}^J_A \in \mathcal{FMR}^q_{\lambda,L}$:*

- *For $\mathcal{R}^J_A \in \mathcal{FMR}_{\lambda,L}$, the π-limit group at a leaf of J carrying T admits a π-resolution which is $J_0 \cup \{n(J_0)\}$-tight .*
- *$\mathcal{R}$ preserves the validity of the system $\Psi(z,a) \neq 1$*

Proof of 6.1.9: uncovering a node through shortening

Take a J_0-tight structure $((H_\lambda)_\lambda, (\gamma_\mu)_\mu, v_\mu, S^\nu)_{\lambda,\mu,\nu}$ for L_A. Notice that the way the lineaar order $\lhd$ was defined, $n(J_0)$ is either r or the the child of a node in J_0. The result can be proved by yet another application of Rips' and Sela's shortening argument.

Let $\mathcal{FMA}$ be the family of all conjugates in L of non-cyclic maximal abelian subgroups of $\hat{T}$ and $\mathcal{BR} := \{\mu \in \bar{J}_0 \mid v_\mu = v_{n(J_0)}\}$. Given any $M \in \mathcal{FMA}$, let $\bar{M} = Z_L(M)$.

Given $\lambda \in \bar{J}_0$ denote by $\bar{K}_\lambda$ be the subgroup generated by the union of K_λ and $\bar{M} \cap Stab(v_\lambda)$ for those $M \in \mathcal{FMA}$ such that $M \cap K_\lambda \neq \{1\}$. The following lemma will be useful. It consists essentially in an iterative application of lemma 2.3.1.

Lemma 6.1.10. *Let W be a simplicial H-tree $\mathcal{O}$ a minimal invariant family of vertices of W such that each member of the family $\mathcal{A} := \{\bar{K}_{n(J_0)}\} \cup \{\bar{K}_\mu \mid v_\mu = v_{n(J_0)}\}$ fixes a vertex of $\mathcal{O}$.*

Then W equivariantly embeds into a simplicial $\hat{H}$-tree W' with the following properties:

a) *The pegged tower group $\hat{T}$ fixes a vertex v_T in the image of $\mathcal{O}$ such that $Stab(v_T)$ admits a J_0-tight structure.*

b) *The translates of the image of W in W' cover W.*

c) *For any edge e in W the stabilizer of e (in H) is the same as the stabilizer of $\phi(e)$ in H_{λ_j}.*

d) *If two edges $e, e' \in W$ are in distinct orbits, then the same is true for $\phi(e), \phi(e')$.*

e) *Suppose we are given an equivariant family $\mathcal{O}$ of vertices of W such that each member of $\mathcal{A}$ fixes some $u \in \mathcal{O}$. Then for any vertex $v \in W \backslash \mathcal{O}$ we have:*

$$Stab_H(v) = Stab_{H_{\lambda_j}}(\phi(v))$$

f) *Given any two vertices $v, v' \notin \mathcal{O}$ in distinct orbits their images $\phi(v), \phi(v')$ are also in distinct and do not belong to the orbit of any vertex in $\mathcal{O}$.*

Proof. Let us say that $n(J_0) = \lambda_0 \prec \lambda_1 \prec \cdots \lambda_m = r$. Let W be the tree dual to the decompositon Δ. For $1 \leqslant j \leqslant m$ we construct a H_{λ_j}-simplicial tree W_j and an equivariant embedding ϕ_j of W into W_j. Let $\mathcal{O}_j$ the union of translates of $\phi_j(\mathcal{O})$. Additionally, we require the following properties to hold:

- Each subgroup in the family

$$\mathcal{A}_j := \{\bar{K}_{\lambda_j}\} \cup \{Stab(w) \mid w \in VS^{\lambda_j} \backslash (H^{\lambda_j} \cdot v_{\lambda_{j-1}}), \}$$

  fixes some vertex in the orbit $\mathcal{O}_j$.
- The triple (W, ϕ_j, W_j) satisfies properties (c)-(f) of the statement of the lemma. .

We proceed by induction. Suppose that W_l has already been constructed for $0 < l < j$. Our goal is to construct W_j as an application of 2.3.1 to the trees S^{λ_j} and W_{j-1} with $\bar{K}_{\lambda_j}$ in the role of G. Note that there are three possibilities for how some $M \in \mathcal{FMA}$ which intersects K_{λ_j} non-trivially and its centralizer $\bar{M}$ can be placed with respect to S^{λ_j}.

- The first one is that $\bar{M} \cap H_{\lambda_j}$ is hyperbolic in U^{λ_j}. Acylindricity of the tree U^{λ_j} do not allow any non-trivial element elliptic in U^{λ_j} to commute with another one which is hyperbolic. This means the group $M \cap K_{\lambda_j}$ is hyperbolic in U^{λ_j} as well and

$$\bar{M} \cap H_{\lambda_j} = M \cap K_{\lambda_j} = \langle pg_M \rangle$$

- Secondly, it is possible for $\bar{M} \cap H_{\lambda_j}$ to fix a rigid type vertex of U^{λ_j}.
- Finally, it might be that λ_j is of type (A) and $M \cap K_{\lambda_j}$ is the stabilizer of an abelian type vertex of U^{λ_j}. Of course, this can be the case for at most one M.

Thus $\bar{K}_{\lambda_j}$ is generated by K_{λ_j} together with the union $\mathcal{E}$ of $\bar{M} \cap H_{\lambda_j}$ for all the M as in the last two bullets. Let $\bar{U}^{\lambda_j}$ be the tree spanned by the translates of U^{λ_j} by $\bar{K}_{\lambda_j}$ (it clearly contains the minimal tree of $\bar{K}_{\lambda_j}$ in S^{λ_j}).

A first consequence of the property above is that the map $\bar{q} : \bar{K}_{\lambda_j}\backslash\bar{U}^{\lambda_j} \to H_{\lambda_j}\backslash S^{\lambda_j}$ is injective on edges. The analogous quotient $q : K_{\lambda_j}\backslash U^{\lambda_j} \to H_{\lambda_j}\backslash S^{\lambda_j}$ certainly is, so all we need to show is that the quotient $q' : K_{\lambda_j}\backslash U^{\lambda_j} \to \bar{K}_{\lambda_j}\backslash U^{\lambda_j}$ is surjective. This follows from lemma 2.2.3 as $\bar{K}_{\lambda_j}$ is generated by K_{λ_j} and elements of H_{λ_j} fixing vertices of U^{λ_j}. (Since q' is injective on edges, we also know from 2.2.3 is q' an isomorphism.)

Lemma 2.2.3 implies as well that for any $v \in U^{\lambda_j}$ the stabilizer of v in $\bar{K}_{\lambda_j}$ is generated by $Stab_{K_{\lambda_j}}(v) \cup (\mathcal{E} \cap Stab_{H_{\lambda_j}}(v))$. Observe any $M \in \mathcal{FMA}$ can intersect $Stab_{K_{\lambda_j}}(v)$ non-trivially for at most one rigid type vertex v of U^{λ_j}. Also that if w is a vertex of abelian type of U^{λ_j} and $Stab_{H_{\lambda_j}}(w)$ intersects $Stab_{H_{\lambda_j}}(v)$ non-trivially for some $v \in U^{\lambda_j}$ of rigid type, then $Stab_{H_{\lambda_j}}(w) = \bar{M} \cap H_{\lambda_j}$ for some $M \in \mathcal{FMA}$ which intersects $Stab_{K_{\lambda_j}}(v)$ non-trivially. It follows that $Stab_{\bar{K}_{\lambda_j}}(v_{\lambda_{j-1}}) = \bar{K}_{\lambda_{j-1}}$.

We claim that $Stab_{\bar{K}_{\lambda_j}}(v_\mu)$ is conjugate into $\bar{K}_\mu$ in H_{λ_j} for any $\mu \in Ch(\lambda_j)\backslash J_0$ as well. . Indeed, take $\nu \leqslant \lambda_j$, $\nu \in J_0 \cup n(J_0)$ such that $c(\mu) = \nu$. If that was not the case, then for some $\nu \leqslant \kappa \leqslant \lambda_j$ there should be conjugate $M, M' \in \mathcal{FMA}$ such that $K_\mu \subset M$ is non-trivial and M' intersects the stabilizer of some abelian type vertex of S^κ non-trivially. However, this cannot be the case, since it implies the existence of $M_0, M'_0 \in \mathcal{MA}(\hat{T})$ such that $pg_{M_0} \in \mathsf{T}^\mu$ and $pg_M \in \mathsf{T}^{\lambda_j}$. Now, $\mu \lhd \lambda_j$ and μ and λ_j are not $\leqslant$ comparable, so for any n big enough $\gamma_n \in \mathbb{F}$ exists such that the images $f_n(pg_{M_0})$ and $f_n(pg_{M'_0})^{\gamma_n}$ by some test sequence $(f_n)_n$ commute. Since for n big enough these two are primitive elements of different length, this is impossible.

Consider any translate $v = h \cdot v_{\lambda_{j-1}} \in S^{\lambda_j}_{\bar{K}_{\lambda_j}}$ for $h \in H_{\lambda_j}$. For any edge $e \in \bar{U}^{\lambda_j}$ incident at v clearly $Stab_{H_{\lambda_j}}(e) \leqslant \bar{K}_{\lambda_j}$. On the other hand, as a consequence of the last discussion $Stab_{\bar{K}_{\lambda_j}}(v)^h$ is a conjugate in H_{λ_j} of either $\bar{K}_{\lambda_{j-1}}$ or $\bar{K}_\mu$ for some $\mu \in Ch(\lambda_l)\backslash J_0$, where $0 < k < j-1$. Now, we know $\bar{K}_{\lambda_{j-1}}$ to fix a vertex of $\mathcal{O}_{j-1}$, while any such $\bar{K}_\mu$ is fixes one in $\mathcal{O}_l$ for some $0 \leqslant l < j$ (by the initial assumption in case $l = 0$ and by the induction hypothesis otherwise), hence one in $\mathcal{O}_{j-1}$ as well.

Trees S^{λ_j} and W_{j-1} thus fall under the hypothesis of lemma 2.3.1, with $\bar{K}_{\lambda_j}$ in the role of G. The lemma produces a tree, that we take as our W_j, and an equivariant embedding ψ_{j-1} of W_{j-1} into W_j. We defne ϕ_j as $\psi_{j-1} \circ \phi_{j-1}$.

The construction (in view of the previous paragraph) guarantees the ellipticity of the family $\{\bar{K}_{\lambda_j}\} \cup \{Stab(w) \mid w \in VS^{\lambda_j}\backslash(H_{\lambda_j} \cdot v_{\lambda_{j-1}})\}$. Since translates of $\phi_{j-1}(W)$ cover W_{j-1} and translates of $\psi_{j-1}(W_{j-1})$ cover W_j clearly translates of $\phi_j(W)$ cover W_j. Likewise, lemma 2.3.1 guarantees that $\psi_{j-1}(VW_{j-1}\backslash\mathcal{O}_{j-1}) \subset VW_j\backslash\mathcal{O}_j$, and $\phi_{j-1}(VW\backslash\mathcal{O}) \subset VW_{j-1}\backslash\mathcal{O}_{j-1}$ by induction hypothesis, so clearly $\phi_j(VW\backslash\mathcal{O}) \subset VW_j\backslash\mathcal{O}_j$. Transitivity can be easily checked for the remaining properties in order to complete the proof.

Pick some $v_0^T \in \mathcal{O}$ fixed by $K_{n(J_0)}$. Let $v_j^T := \phi_j(v_j^T)$ and for $1 \leqslant j \leqslant m$ let R_{λ_j} the full stabilizer of v_j^T and by P_{λ_j} the tree spanned by all translates of U^λ by elements of R_λ. We claim that the stabilizer of $v_T = v_j^T$ admits a J_0-tight structure. Let us see how to extend this data to a J_0-tight structure

$$((R_\lambda)_{\lambda \in J_0 \cup \{n(J_0)\}}, (P^\lambda)_{\lambda \in J_0}, (\gamma'_\lambda)_{\lambda \in \bar{J}_0\backslash J_0}, (v'_\lambda)_{\lambda \in \bar{J}_0\backslash\{r\}})$$

where $v'_\lambda = v_\lambda$ in case $\lambda \in J_0$. For $1 \leqslant j \leqslant m$, consider the two graph maps induced by inclusion: $q_1 : K_{\lambda_j}\backslash U^{\lambda_j} \to R_{\lambda_j}\backslash P_{\lambda_j}$ and $q_2 : R_{\lambda_j}\backslash P_{\lambda_j} \to H_{\lambda_j}\backslash S^{\lambda_j}$.

Lemma 2.3.1 tells us that q_2 is injective on edges and that the only vertex in its image admitting several preimages is $v_{\lambda_{j-1}}$. $n(J_0) \leqslant \lambda_{j-1} \prec \lambda_j$. On the other hand $q_2 \circ q_1$ is bijective on edges by definition. It follows that q_1 must be surjective as well, so clearly the tree P_λ satisfies the requirements of the definition.

Consider now any $\mu \in Ch(\lambda_j)\backslash\{\lambda_{j-1}\}$. If $\mu \in J_0$ there is no need to modify the given data for $\nu \leqslant \mu$, as $H_\mu \subset R_\lambda$. Suppose now that $\mu \notin J_0$, let u_μ be the vertex stabilized by T^μ and $c'(\mu)$ the minimum $\kappa \in J_0$ such that $K'_\mu := (\mathsf{T}^\mu)^{\gamma'_\lambda} \subset R_{\lambda_j}$ for some $\gamma'_\lambda \in R_{\lambda_j}$. If $c(\mu) = \lambda_j$, then we can take $v'_\mu = u_\mu$. One can show that $\gamma_\mu(\gamma'_\mu)^{-1} \in H_\kappa$ [(2)], so that K'_μ must fix a rigid vertex w of S^μ. If $w \notin P_\mu$, since K'_μ is contained in R_μ, it must fix the vertex in P_μ closest to w. All vertices in the boundary of P_μ in S^μ are rigid, so an adequate v'_μ exists. □

Corollary 6.1.11. *With the notation of the lemma, if W is endowed with a geometric abelian marking relative to $\{\bar{K}_{n(J_0)}\} \cup \{\bar{K}_\mu \,|\, v_\mu = v_{n(J_0)}\}$, then the L-tree W' can be regarded as a geometric abelian tree in such a way ϕ respects vertex type. Every $\sigma \in Mod^\pi_{K_{n(J_0)}}(W)$ extends to some $\tau' \in Mod^\pi_{\hat{T}}(W')$.*

Proof. In this case we can take as $\mathcal{O}$ any minimal invariant family of rigid vertices such that each member of $\mathcal{A}$ stabilizes one of them. Any vertex $v \in W$ not belonging to $\mathcal{O}' := L \cdot \mathcal{O}$ in W' is a translate of $\phi(w)$ for some vertex $w \in W\backslash\mathcal{O}$, whose orbit is uniquely determined. The action of $Stab(w)$ on the star around v is isomorphic to the action of $Stab(w)$ on the star around w, so one can assign to v the type of w (obviously well-defined).

Notice that $\mathcal{O}'$ consists of a unique orbit, as any $v \in \mathcal{O}$ has to become an attaching point at some stage in the iterative construction. The properties of the embedding ϕ imply[(3)] that a geometric abelian decomposition Δ' of L associated to W can be obtained (in a way compatible with the inclusion map) from a geometric abelian decomposition Δ of L by means of:

i) Enlarging Δ_u for some rigid vertex u for which $K_{n(J_0)} \leqslant \Delta_u$.

ii) Removing some of the other rigid vertices and attaching all edges incident to them to u instead. This eequires adding Bass-Serre elements for some of those edges, which have to drop from the maximal tree.

For any of the removed vertices w, the group Δ_u is contained in $(\Delta'_u)^{s_u}$, where s_u is one of the newly introduced Bass-Serre elements. Given any $\tau \in PMod(\Delta)$ in the set of generators given in section 2.5 which fixing Δ_u (either a Dehn twist or an extension of a vertex group automorphisim) one can easily extend τ to some $\tau' \in Mod^\pi(W')$ fixing Δ'_u and sending any of the s_u above to $s_u c_u$, whenever τ restricts to conjugation to c_u on Δ_u. □

Corollary 6.1.12. *Let L_A be a J_0-tight formal π-limit group. Suppose that there is a non-trivial free decomposition of $H_{\lambda_{n(J_0)}}$ relative to $\{K_{n(J_0)}\} \cup \{K_\mu \,|\, v_\mu = v_{n(J_0)}\}$. Then there is a free decomposition of $L = L_0 * L$ such that $\hat{T} \leqslant L_0$ and L_0 admits a J_0-tight structure as well.*

Proof. Notice that any free splitting of $H_{n(J_0)}$ relative to $\{K_{n(J_0)}\} \cup \{K_\mu\}_{c(\mu)=n(J_0)}$ is also relative to $\{\bar{K}_{n(J_0)}\} \cup \{\bar{K}_\mu\}_{v_\mu = v_{n(J_0)}}$, (if an abelian group intersects a free factor non-trivially then it is contained in it) so the previous lemma can be applied to the tree W dual to the said decomposition. □

(2) By repeatedly using the general fact that given $\lambda \in J_0$ an element $h \in H_\lambda$ conjugates some non-trivial element in the stabilizer of a (rigid) vertex of S^λ to another one, then h itself must belong to the stabilizer.

(3) Take a vertex stabilized by $\hat{T}$ in $\phi(\mathcal{O})$. It is easy to find a Bass-Serre presentation $(Z_0, Z, (t_e)_{e \in EZ})$ of W' with $w \in Z_0, Z \subset \phi(W)$, by adding edges to the subtree in a greedy fashion until a transversal is found.

Using the corollary one can reduce the theorem to the case in which $H := H_{n(J_0)}$ is freely indecomposable relative to $\{K_\mu \mid v_\mu = v_{n(J_0)}\} \cup \{K_{n(J_0)}\}$. Suppose it is not; then we know there is a free decomposition of L of the form $L = L_0 * L_1$, where $\hat{T} \leqslant L_0$. Chose a new tuple of variables $u = (u_0, u_1)$ such that each L_i is endowed with a surjective u_i-marking.

By collecting the elements of each L_i appearing in the normal form of the words in Θ one can find systems of inequations $\Theta_0(u^0, a) \neq 1$ and $\Theta_1(u^1) \neq 1$ such that given any restricted homomorphism q_0 with domain L_0 and q_1 with domain L_1, if each q_j preserves the validity of $\Theta_j \neq 1$ then the homomorphism from L to $(*_{j=0}^{1} f_j(L_j))$ induced by them preserves $\Psi(z, a) \neq 1$. It is a routine task to derive the following result given 3.2.6 and the argument following result, just as in 4.2.2.

Lemma 6.1.13. *Given a π-limit group $\mathsf{L}(z)$ and a system of inequalities $\Psi(z, a) \neq 1$ there is a finite family of closed resolutions $\mathcal{MR}_\Psi$ preserving the validity of the system such that any morphism from L to F preserving the validity of $\Psi(z, a) \neq 1$ factors through at least one of $\mathcal{R} \in \mathcal{MR}_\Psi$*

Let $\mathcal{MR}_1$ be the set of π-resolutions which results from applying lemma 6.1.13 to the pair $(\mathsf{L}_1(u^1), \Theta_1(u^1) \neq 1)$. And $\mathcal{MR}_0$ that obtained by applying the induction hypothesis to $(\mathsf{L}_0, \Theta_0(u^0, a) \neq 1)$. Let $\mathcal{MR}^{nd}$ be the result of combining each resolution in the first set with each one in the second in the obvious way into a resolution of L. Each $\mathcal{R} \in \mathcal{MR}^{nd}$ preserves the validity of $\Psi(z, a) \neq 1$. And each formal sequence whose members eventually preserve the validity of $\Theta_0(u^0, a) \neq 1$ and $\Theta_1(u^1) \neq 1$ contains an infinite subsequence factoring through one of them. Now, take any element $v \in \Theta_0 \cup \Theta_1$ and consider the family $\mathcal{FS}_v$ of all formal sequences of L which kill v and preserve $\Psi(z, a) \neq 1$. If $\mathcal{FS}_v = \emptyset$, then just ignore v. If not, since this family is clearly closed under diagonal subsequences, every $(g_n)_n \in \mathcal{FS}_v$ eventually factors (in a strict sense) through one of a collection in a finite collection $\mathcal{Q}_v$ of maximal limit quotients of sequences in $\mathcal{FS}_v$. For each $\mathsf{Q}_A \in \mathcal{Q}_v$ let $\mathcal{MR}^Q$ the finite collection of π-resolutions given by applying the induction hypothesis to the pair $(\mathsf{Q}_A, \Psi(z, a) \neq 1)$. We can complete each one of them into a π-resolution of L_A in an obvious way, by adding the quotient form L_A onto Q_A on top. Let $\mathcal{MR}^d$ be the collection of all the resolutions so obtained, by letting v range in $\Theta_0 \cup \Theta_1$. The family $\mathcal{MR}^{nd} \cup \mathcal{MR}^d$ satisfies all the required conditions.

Let Δ_{JSJ} be a JSJ decomposition of $H_{n(J_0)}$ relative to the family $\{\bar{K}_\mu\}_{v_\mu = v_{n(J_0)}} \cup \{\bar{K}_{n(J_0)}\}$ and let $\mathfrak{M}$ be the group consisting of all the extensions $\tilde{\tau}$ fixing $\hat{T}$, where τ ranges among all elements of $Mod^{\pi}_{K_{n(J_0)}}(\Delta_{JSJ})$. Fix w a tuple of generators for H. Consider the action of $\mathbb{F}$ on $Cayl(\mathbb{F})$. Given a morphism $f : \mathsf{L}_A(z) \to \mathsf{F}_A$ and $\lambda \in J$, let $*_B^{f \upharpoonright H} \in Cayl(\mathbb{F})$ the basepoint chosen as described in 3.1.7.1. We say that f is short among those preserving the validity of $\Psi(z, a) \neq 1$ if for any $\sigma \in \mathfrak{M}$ for which $f \circ \sigma$ preserves the validity of $\Psi(z, a) \neq 1$ the inequality $sl_{*_B^{f \upharpoonright H}}(w) \leqslant sl_{*_B^{f \upharpoonright H}}(\sigma(w))$ holds. By a short formal sequence $(g_n)_n$ for L we intend a formal sequence such that g_n satisfies this condition for any n.

Lemma 6.1.14. *Let L be a formal π-limit group which is freely indecomposable relative to $\mathcal{K} = \{K_\mu\}_{v_\mu = v_{n(J_0)}} \cup \{K_{n(J_0)}\}$ and admits a J_0-tight structure and let $(g_n)_n$ be a short formal sequence consisting entirely of short morphisms in the sense above from L_A to F_A. Then one of the following takes place:*

a) $lker_n\, g_n \neq \{1\}$

b) There is a decomposition of $H_{n(J_0)}$ of the form $H_{n(J_0)} = H'_{n(J_0)} *_{N \cap H'_{n(J_0)}} N$, where $K_{n(J_0)} \leqslant H'$, $N = Z_{H_{n(J_0)}}(M)$ for some $M \in \mathcal{FMA}$ and $N = N \cap H' \oplus E$ for some E.

c) L_A admits a $J_0 \cup n(J_0)$-tight structure.

Proof. Let $J_1 = J_0 \cup \{n(J_0)\}$. Suppose that $lker_n\ g_n = 1$. Let ρ be the action of H on the limiting tree for the sequence $(g_n\restriction_H)_n$. The sequence of points $*_B^{g_n}$ necessarily converges to a point $*$ fixed by B or in the minimal tree of B in the limiting tree Y associated to the restriction of the sequence $(g_n)_n$ to $H_{n(J_0)}$.

Suppose first that $K_{n(J_0)}$ acts elliptically on Y.

Lemma 3.2.6 implies that some $\bar{K}_\mu \in \mathcal{A}$ is not elliptic in Y. In case $\mu \neq n(J_0)$ the properties of test sequences however imply that K_μ does fix some $x_\mu \in Y$. We conclude there is some $M \in \mathcal{FMA}$ intersecting K_μ non-trivially such that $N := Z_M(H_{n(J_0)})$ does not fix x_μ. This implies it does not fix a point $y \neq x_\mu$ either. If not, then notice that in this case M must also fix y, so for some $g \in N$ either the union $[x_\mu, y] \cup g\cdot[x_\mu, y] \cup g^2\cdot[x_\mu, y]$ is a non-degenerate non-trivially stabilized tripod or g acts as a symmetry on a non-degenerate segment $[x_\mu, y] \cup g\cdot[x_\mu, y]$ fixed, both of which cannot occur.

So N acts hyperbolically on Y, while M and $K_{n(J_0)}$ act elliptically, and by 3.1.16 there is a decomposition of $H_{n(J_0)}$ as an amalgamated product of the required form.

We are left with the case in which $K_{n(J_0)}$ does not act elliptically on Y. If the node λ carries the constants then $\mathsf{T}^\lambda = \mathsf{A}$, so this is impossible. In all other cases corollary 5.4.3 implies some $H_{n(J_0)}$-tree $S^{n(J_0)}$ exists with the desired properties.

All is left is to check the existence of appropriate v'_λ and γ'_λ for $\lambda \in \bar{J}_1\backslash J_1$. For $\lambda \in J_1$ we don't effect any change. Consider first some $\lambda \in \bar{J}_0\backslash\{n(J_0)\}$ such that $c(\lambda)$ is the parent of $n(J_0)$. By assumption either v_λ is a translate of $v_{n(J_0)}$ or v_λ and $v_{n(J_0)}$ are not in the same orbit. In the latter case the parent μ of $n(J_0)$ is remains the lowest node of J_1 into which T^λ can be conjugated. In the former one we can replace γ_λ by γ'_λ in such a way that the resulting K'_λ is contained in $Stab(v_{n(J_0)})$. Now, K'_λ is elliptic in the real tree Y, so 5.4.3 implies that it stabilizes a vertex (of rigid type when it applies) of S^λ which we can take as our new v'_λ. The choice of v_λ for $\lambda \in \bar{J}_0\backslash\bar{J}_1 = Ch(n(J_0))$ is clear. All we need to prove in order to verify (iv) is that $v_{n(J_0)}$ is not a translate of v_ν for any $\nu \in Ch(\mu) \cap J_0$, which follows from observation 6.1.7. □

Equipped with this result, let us finish the proof of 6.1.6. Assume first that $H_{n(J_0)}$ admits a decomposition as in b of the previous lemma.

Claim 6.1.15. *Such a decomposition $H' *_{N\cap H'} N$ inducess one of the form $L = L' *_{\bar{M}\cap L'} \bar{M}$, where $H' \leqslant M$, $\bar{M} = Z_L(M)$ and L' admits a J_0-tight structure.*

Proof. The proof is a small variation from that of 6.1.10. We start with the given decomposition, Δ_0, and for $1 \leqslant j \leqslant m$ we construct one Δ_j of the form $H_{\lambda_j} = H'_{\lambda_j} *_{\bar{M}\cap H'_{\lambda_j}} (\bar{M} \cap H_{\lambda_j})$. These are the different possibilities for $0 \leqslant j \leqslant m$:

- λ_j falls in case (A) so that there is a decomposition

$$H_{\lambda_j} = H_{\lambda_{j-1}} *_{\bar{M}\cap H_{\lambda_{j-1}}} ((\bar{M} \cap H_{\lambda_l}) \oplus F)$$

  this means we can write H_{λ_j} as an amalgamated product $H_{\lambda_{j-1}} *_{H_{\lambda_{j-1}}\cap\bar{M}} ((H_{\lambda_{j-1}} \cap \bar{M}) \oplus E \oplus F)$ and we can take $H'_{\lambda_j} := H'_{\lambda_{j-1}} \oplus F$. Clearly the action of H'_{λ_j} on the corresponding subtree of
- If λ falls in any other case, we can use lemma 2.3.1 to lift the amalgamated product from $H_{\lambda_{j-1}}$ to H'_{λ_j}.

At this point one can check that H'_{λ_j} appear in place of H_{λ_j} in some J_0-tight structure for $L' = H'_{\lambda_m}$ in an entirely analogous way as in the proof of 6.1.10. □

In this case the usual argument provides some π-retraction $r : \mathsf{L} \to \mathsf{L}'$ preserving the validity of the sytem $\Psi(z, a) \neq 1$. For any formal sequence $(g_n)_n$ of L, in virtue of 5.4.4 eventually $g_n(N) = g_n(N \cap L')$, and g_n can be written in the form $g'_n \circ r \circ \sigma$, where σ is a π-modular automorphism of the decomposition $L' *_{N \cap L'} N$ (taking N as an abelian type vertex group) fixing L'. We are done by extending each member of the family of π-resolutions provided by the induction hypothesis for the pair to $(\mathsf{L}', \Psi(z, a) \neq 1)$ in the obvious way.

In the other case, observe that since the family of short formal sequences of L is closed under taking diagonal sequences, the collection of limit quotients by short formal sequences contains finitely many maximal elements $\mathcal{MFQ}$.

By 6.1.14 we might as well assume that they are proper. Given any formal sequence $(g_n)_n$ preserving the validity of $\Psi(z, a) \neq 1$, there if for all n we take $\tau_n \in \mathfrak{M}$ such that $g'_n = g_n \circ \tau_n$ is short then the members of any convergent subsequence of $(g'_n)_n$ must factor through $\phi \in \mathcal{MFQ}$ after a certain point.

The sought family of resolutions are obtained in the usual way by combining the quotient $p : \mathsf{L} \to \mathsf{L}' \in \mathcal{MFQ}$ with any resolution resulting from applying the induction hypothesis to $(\mathsf{Q}, \Psi(z, a) \neq 1)$.

6.2 The positive theory

As before, we add constants for the members of some finite tuple $a \leqslant \mathbb{F}$ generating some free factor A of $\mathbb{F}$. If we let $(T^{\pi}_{gp})_A$ stand for the theory of all A-restricted π-groups, it is clear how any positive $\mathcal{L}^{Q}_{A}$-sentence is $(T^{\pi}_{gp})_A$-equivalent to one of the following form:

$$\forall x^1 \, \exists y^1 \cdots \forall x^m \, \exists y^m \bigvee_{i=1}^{k} ((x \in p(i) \wedge y \in q(i) \wedge \Sigma^i(x, y, a) = 1)$$

Here $p(i) \in Q^{|x|}$ and $q(i) \in Q^{|y|}$, $x = (x^1, x^2, \cdots x^m)$ and $y = (y^1, y^2, \cdots y^m)$.

As a matter of fact, it is more convenient to work with a different class of formulas, which we will refer to as simple constrained positive (SCP) formulas, namely those of the form:

$$\forall x^1 \in p^1 \, \exists y^1 \in q^1 \cdots \forall x^m \in p^m \, \exists y^m \in q^m \; \Sigma(x, y, a) = 1$$

For some system $\Sigma(x, y, a) = 1$ of equations with parameters in a. Here the following abbreviations have been used, for a tuple x of variables and $q \in Q^{|x|}$:

$$(\forall x \in q \; \phi) \equiv \forall x \; (x \in q \to \phi)$$
$$(\exists x \in q \; \phi) \equiv \exists x \; (x \in q \wedge \phi)$$

Given an SCP sentence ϕ as above, by a formal solution of ϕ we intend a tuple $(w^1, w^2 \cdots w^m)$ with the following properties:

- w^j is a $|y_j|$-tuple of words in $x^1, \cdots x^j, a$ for $1 \leqslant j \leqslant m$
- $w^j_l(p^1, p^2, \cdots p^j, \pi(a)) = q^j_l$ for $1 \leqslant j \leqslant m$ and $1 \leqslant l \leqslant |y|$
- For any word $u(x, y, a)$ in the system $\Sigma(x, y, a) = 1$ the term:

$$u(x, w^1(x^1, a), w^2(x^1, x^2, a), \cdots w^m(x^1, x^2 \cdots x^m, a), a)$$

  represents the trivial element in the free group $A * \mathbb{F}(x)$.

Clearly, if an SCP formula admits a formal solution, then it is valid in any A-restricted π-group. So, in particular it is valid in F_A. Later we will prove a strong converse of this result with implications for general positive formulas, under the assumption that A is a free factor.

In order to start, we need to go back to the language of π-groups. Associated with any SCP formula as above there is a restricted π-group $\mathsf{G}_{\phi A}$, whose underlying group is the quotient of $A * \mathbb{F}(x,y)$ by the normal subgroup generated by the words in $\Sigma(x,y,a)=1$ and $\pi=\pi_{G_\phi}$ maps each tuple x^j (which we now see as elements of G_ϕ) to the tuple p^j and y^j to q^j. For any $1 \leqslant j \leqslant m$ we let G^j_ϕ be the subgroup of G_ϕ generated by the tuple $(a, x_1, y_1, \cdots x_j)$ and H^j_ϕ be the one generated by $(a, x_1, y_1, \cdots x_j, y_j)$. We say that a ϕ is free on the on the universal variables if for any $0 \leqslant l \leqslant m-1$ the subgroup $\langle G^l_\phi, x_{l+1}, \cdots x_m\rangle$ is isomorphic to the free product of G^l_ϕ and $\mathbb{F}(x_{l+1}, \cdots x_m)$. In particular, $\langle A, x\rangle \cong A * \mathbb{F}(x)$. Viceversa, of course, any finitely presented restricted π-group G_A endowed with a surjective (x,y)-marking can be seen as a G_ϕ, by taking as $\Sigma(x,y,a)=1$ the collection of relators in some presentation, expressed in terms of x and y.

Observation 6.2.1. A SCP ϕ free on the universal variables admits a formal solution if and only if there is a π-retraction $f: G_\phi \to A * \mathbb{F}(x)$ such that $f(H^j_\phi) \subset A * \mathbb{F}(x_1, \cdots x_j)$ for any $1 \leqslant j \leqslant m$.

Corollary 6.2.2. *Suppose that we are given SCP formulas:*

$$\phi \equiv \forall x^1 \in p^1\, \exists y^1 \in q^1 \cdots \forall x^m \in p^m\, \exists y^m \in q^m \Sigma(x,y,a)=1$$
$$\psi \equiv \forall u^1 \in r^1\, \exists v^1 \in s^1 \cdots \forall u^m \in r^m\, \exists v^m \in s^m \Pi(u,v,a)=1$$

which are free in the universal variables and that some surjective restricted morphism f from $\mathsf{G}_{\phi A}$ to $\mathsf{G}_{\psi A}$ exists such that $f(H^j_\phi) \subset H^j_\psi$ for all $1 \leqslant j \leqslant m$. Then the existence of a formal solution for ψ implies that of one for ϕ.

Definition 6.2.2.1. Given an SCP formula:

$$\phi \equiv \forall x^1 \in p^1\, \exists y^1 \in q^1 \cdots \forall x^m \in p^m\, \exists y^m \in q^m \Sigma(x,y,a)=1$$

by a ϕ-formal sequence we intend a metrically convergent sequence of morphisms

$$(f_n)_n : \mathsf{G}_{\phi A} \to \mathsf{F}_A$$

such that:

i) The restriction of $(f_n)_n$ to $\langle x^j\rangle$ is a π-test sequence for the trivial π-tower with $(\langle x^j\rangle, \pi_{G_\phi}\restriction_{\langle x^j\rangle})$ at its single node (i.e. a small cancellation sequence).

ii) The group $\langle x^j\rangle$ grows faster than H^l_ψ for $1 \leqslant l < j \leqslant m$.

Proposition 6.2.3. *Suppose we are given an SCP formula*

$$\phi \equiv \forall x^1 \in p^1\, \exists y^1 \in q^1 \cdots \forall x^m \in p^m\, \exists y^m \in q^m\ \Sigma(x,y,a)=1$$

where A is a free factor of $\mathbb{F}$ and $\langle A, p^1, p^2, \cdots p^m\rangle = Q$. Then ϕ admits a formal solution if and only if some ϕ-formal sequence exists.

Proof. The only if direction is clear. We prove the opposite direction by induction on the pair (m, G_ϕ) with respect to the partial order $\leqslant \times \leqslant_{rk} \times \leqslant_Z$. Case $m=1$ is a particular case of the main result of the previous subsection.

To begin with, let us discard the case in which the ϕ-formal sequence $(f_n)_n$ has non-trivial limit kernel. In this situation let H_A be its limit quotient. Since H is finitely presented H_A equals $\mathsf{G}_{\phi' A}$ for some other SCP formula ϕ' and some subsequence of $(f_n)_n$ pushes forward to a ϕ'-formal sequence. By induction ϕ' admits a formal solutions and therefore so does ϕ.

We can also assume that G_ϕ is freely indecomposable relative to G^m_ϕ. Indeed, otherwise we can write $G_\phi = L_1 * L_2$, where $L_2 \neq \{1\}$ and $G^m_\phi \leqslant L_1$. Take any morphism h from L to $\mathbb{F}$ and let ι be some morphism from F to $\langle A, x^1, x^2, \cdots x^m\rangle$. This exists due to the fact that the latter group maps onto Q. The map $\iota \circ h$ extends to a π-retraction from G_ϕ onto L_1.

So assume now that $(f_n)_n$ has trivial limit kernel and G_ϕ is freely indecomposable relative to G^j_ϕ. And consider the limiting tree Y for the sequence $(f_n)_n$.

Let λ be the action of $\mathbb{F}$ on its Cayley graph. For any homomorphism $f : G_\phi \to \mathbb{F}$ not killing $\langle a, x^1\rangle$ (there is no harm in assuming none of the f_n does), let $*^{\lambda_f}_{G^m_\phi}$ be the basepoint associated to G^m_ϕ in the minimal tree of G_ϕ in X with respect to the action λ_f, chosen as in section 3.1.2 and $\mu(f) = (sl^{\lambda_f}_{bpf}(y^m))^m_{j=1}$. We can assume that $\mu(f_n) \leqslant \mu(f')$ for any $f' : G_\phi \to \mathbb{F}$ which coincides with f_n on G^m_ϕ (the condition for being a ϕ-formal sequence depends only on $f_n\restriction_{G^m_\phi}$).

We claim that this implies $\langle x^m\rangle$ is not elliptic in Y. Indeed, if it was, then G^m_ϕ would fix the limit of the sequence $(*^{\lambda_f}_{G^m_\phi})_n$ and there would thus be an automorphism $\sigma \in Mod^\pi_{G^j_\phi}(G)$ such that for n large enough $\mu(f_n \circ \sigma) \leqslant \mu(f_n)$, contradicting the choice of $(f_n)_n$.

By 5.4.2 we know then that the minimal tree Y_0 of $\langle x^m\rangle$ in Y is fundamental, while H^{m-1}_ϕ fixes some point z in Y. The free indecomposibility hypothesis implies that z can be chosen inside Y_0. Just as in the proof of Merzlyakov theorem, using observation 3.1.14 one can see that then any intersection of Y_0 and one of its translates must belong to the orbit of z and $G_\phi = G' * \langle x^m\rangle$, where $G' = Stab(z)$. Let u be a finite tuple of variables marking elements of G' such that G' is generated by $(x^1, y^1, \cdots y^{m-1}, u)$. We can see G'_A as G_ψ for some SCP formula:

$$\psi \equiv \forall x^1 \in p^1\ \exists y^1 \in q^1 \cdots x^{m-1}\ \exists y^{m-1} \in q^{m-1} \exists u \in r\ \Pi(x^1, \cdots x^{m-1}, y^1, \cdots y^{m-1}, u, a) = 1$$

with the variables in ψ in agreement with the marking of G_ϕ. Now, clearly $(f_n\restriction_{G'})_n$ is a ψ-formal test sequence, so by the induction hypothesis ϕ' admits a formal solution. ψ admits a formal solution. This implies ϕ does, by 6.2.2 (after adding innermost universally quantified variables to ψ). $\square$

Let us now go back to our general positive formula:

$$\phi \equiv \forall x^1\, \exists y^1 \cdots \forall x^m\, \exists y^m \bigvee_{i=1}^{k} (x \in p(i) \wedge y \in q(i) \wedge \Sigma^i(x, y, a) = 1)$$

For any $p = (p^1 \cdots p^m) \in Q^{|x|}$ denote by $\mathcal{D}_p$ the collection of all the SCP formulas with free universal variables of the form:

$$\forall x^1 \in p^1\, \exists y^1 \in q^1 \cdots \forall x^m \in p^m\, \exists y^m \in q^m\ \Sigma^i(x, y, a) = 1$$

for some $1 \leqslant i \leqslant k$ such that $p(i) = p$.

Lemma 6.2.4. *Suppose that $\mathsf{F}_A \models \phi$. Then for each $p \in Q^{|x|}$ there is some sentence $\psi \in \mathcal{D}_p$ admitting a formal ψ-sequence (in particular, $\mathcal{D}_p$ is not empty).*

Proof. Indeed, given any such p, the construction of test sequences implies the existence of a ϕ-formal sequence $(f_n)_n$ such that $\pi(f_n(x)) = p$ for all $n \in \mathbb{N}$. The sought result follows from the validity of ϕ, together with a straightforward application of the pidgeon hole principle. $\square$

Now, let ϕ^{surj} the axiom stating that each of the predicates P_q is non-empty. Modulo the theory $T^\pi_{gp} \cup \phi^{surj}$, any SCP sentence ϕ is equivalent to the one obtained adding any dummy constrained quantifier expression $\forall x' \in p'$ just before the atomic part of ψ, where the variables of x' are disjoint from those appearing in ψ. In particular, we can always assume that the condition $\langle \pi(x, a) \rangle = Q$ is satisfied. Using the previous lemma, the last observation and proposition 6.2.3 we can deduce:

Corollary 6.2.5. *If in the situation above $\mathsf{F}_A \models \phi$, then there is some finite collection $(\psi_i)_{i=1}^r$ of SCP formulas, each of which admits a formal solutions, such that $T_{gp}^{\pi} \vdash (\bigwedge_{i=1}^{r} \psi_i \to \phi)$.*

From this we one can readily deduce:

Theorem 6.2.6. *Let A be a free factor of non-abelian free groups F_1 and F_2 and for $i = 1, 2$ let π_i be a homomorphism from F_i to the finite group Q. Then*

$$Th_A^+(F_1, \pi_1) = Th_A^+(F_2, \pi_2)$$

Proof. Take any positive A-sentence ϕ such that $(F_i, \pi_i) \models \phi$. By the previous corollary, ϕ is implied in T_{gp}^{π} by some conjunction $\bigwedge_{i=1}^{r} \psi_i$ of SCP formulas admitting a formal solution. This implies the validity of each ψ_i, and therefore that of ϕ, in (F_{3-i}, π_{3-i}). $\square$

Bibliography

[Bes02] Mladen Bestvina. R-trees in topology, geometry, and group theory. *Handbook of geometric topology*, pages 55–91, 2002.

[BF95] Mladen Bestvina and Mark Feighn. Stable actions of groups on real trees. *Inventiones mathematicae*, 121(1):287–321, 1995.

[BF09] Mladen Bestvina and Mark Feighn. Notes on selas work: Limit groups and makanin-razborov diagrams. *Geometric and cohomological methods in group theory*, 358:1–29, 2009.

[CG05] Christophe Champetier and Vincent Guirardel. Limit groups as limits of free groups. *Israel Journal of Mathematics*, 146(1):1–75, 2005.

[Chi01] Ian Chiswell. *Introduction to* Λ*-trees*. World Scientific, 2001.

[DD89] Warren Dicks and Martin John Dunwoody. *Groups acting on graphs*, volume 17. Cambridge University Press, 1989.

[DGH05] Volker Diekert, Claudio Gutierrez, and Christian Hagenah. The existential theory of equations with rational constraints in free groups is pspace-complete. *Information and Computation*, 202(2):105–140, 2005.

[Dun85] Martin J Dunwoody. The accessibility of finitely presented groups. *Inventiones mathematicae*, 81(3):449–457, 1985.

[Dun98] MJ Dunwoody. Folding sequences. *Geometry and Topology monographs*, 1:143–162, 1998.

[FLP12] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston's Work on Surfaces (MN-48)*, volume 48. Princeton University Press, 2012.

[FM12] Benson Farb and Dan Margalit. A primer on mapping class groups, volume 49 of princeton mathematical series, 2012.

[Fuj02] Koji Fujiwara. On the outer automorphism group of a hyperbolic group. *Israel Journal of Mathematics*, 131(1):277–284, 2002.

[GL09] Vincent Guirardel and Gilbert Levitt. Jsj decompositions: definitions, existence, uniqueness. i: The jsj deformation space. *arXiv preprint arXiv:0911.3173*, 2009.

[GL10] Vincent Guirardel and Gilbert Levitt. Jsj decompositions: definitions, existence, uniqueness. ii. compatibility and acylindricity. *arXiv preprint arXiv:1002.4564*, 2010.

[Gub86] Victor Sergeevich Guba. Equivalence of infinite systems of equations in free groups and semigroups to finite subsystems. *Mathematical Notes*, 40(3):688–690, 1986.

[Gui08] Vincent Guirardel. Actions of finitely generated groups on $\mathbb{R}$-trees. In *Annales de l'institut Fourier*, volume 58, pages 159–211, 2008.

[KM06] Olga Kharlampovich and Alexei Myasnikov. Elementary theory of free non-abelian groups. *Journal of Algebra*, 302(2):451–552, 2006.

[Lev05] Gilbert Levitt. Automorphisms of hyperbolic groups and graphs of groups. *Geometriae Dedicata*, 114(1):49–70, 2005.

[LP97] Gilbert Levitt and Frédéric Paulin. Geometric group actions on trees. *American Journal of Mathematics*, pages 83–102, 1997.

[LS15] Roger C Lyndon and Paul E Schupp. *Combinatorial group theory.* Springer, 2015.

[Mak83] Gennady S Makanin. Equations in a free group. *Izvestiya: Mathematics*, 21(3):483–546, 1983.

[Mer66] Yuri I Merzlyakov. Positive formulae on free groups. *Algebra i Logika*, 5(4):25–42, 1966.

[MKS07] W Magnus, A Karras, and D Solitar. Combinatorial group theory: Presentations of groups in terms of generators and relations (interscience, new york, 1966). *Received: June*, 27, 2007.

[MS84] John W Morgan and Peter B Shalen. Valuations, trees, and degenerations of hyperbolic structures, i. *Annals of Mathematics*, 120(3):401–476, 1984.

[Pau88] Frédéric Paulin. Topologie de gromov équivariante, structures hyperboliques et arbres réels. *Inventiones mathematicae*, 94(1):53–80, 1988.

[Pau89] Frédéric Paulin. The gromov topology on r-trees. *Topology and its Applications*, 32(3):197–221, 1989.

[Per08] Chloé Perin. *Elementary embeddings into a torsion-free hyperbolic group.* PhD thesis, PhD thesis, Université de Caen, 2008.

[Per11] Chlo Perin. Elementary embeddings in torsion-free hyperbolic groups. *Annales scientifiques de l'cole Normale Suprieure*, 44(4):631–681, 2011.

[PS] C Perin and R Sklinos. Forking and jsj decompositions in the free group, to appear in j. *Eur. Math. Soc.(JEMS).*

[Raz85] Alexander A Razborov. On systems of equations in a free group. *Mathematics of the USSR-Izvestiya*, 25(1):115, 1985.

[Rez97] Alexander Reznikov. Quadratic equations in groups from the global geometry viewpoint. *Topology*, 36(4):849–865, 1997.

[RS94] Eliyahu Rips and Zlil Sela. Structure and rigidity in hyperbolic groups i. *Geometric & Functional Analysis GAFA*, 4(3):337–371, 1994.

[RS97] Eliyahu Rips and Zlil Sela. Cyclic splittings of finitely presented groups and the canonical jsj decomposition. *Annals of Mathematics*, pages 53–109, 1997.

[RW] Cornelius Reinfeldt and Richard Weidmann. Makanin-razborov diagrams for hyperbolic groups, preprint, 201 0.

[SB77] Jean-Pierre Serre and Hyman Bass. Arbres, amalgames, sl2= asrérisque 46. 1977.

[Sco78] Peter Scott. Subgroups of surface groups are almost geometric. *Journal of the London Mathematical Society*, 2(3):555–565, 1978.

[Sel97] Zlil Sela. Acylindrical accessibility for groups. *Inventiones mathematicae*, 129(3):527–565, 1997.

[Sel01] Zlil Sela. Diophantine geometry over groups i: Makanin-razborov diagrams. *Publications Mathematiques de l'IHES*, 93:31–105, 2001.

[Sel03] Zlil Sela. Diophanting geometry over groups ii: Completions, closures and formal solutions. *Israel Journal of Mathematics*, 134(1):173–254, 2003.

[Sel06] Zlil Sela. Diophantine geometry over groups vi: The elementary theory of a free group. *Geometric & Functional Analysis GAFA*, 16(3):707–730, 2006.

[Sho67] Joseph R. Shoenfield. *Mathematical logic*. Addison Wesley, 1967.

[SL77] RC Lyndon PE Schupp and Roger C Lyndon. Combinatorial group theory. *Ergebnisse der Mathematik und ihrer Grenzgebiete Bd*, 89, 1977.

[Sta65] John R Stallings. A topological proof of grushko's theorem on free products. *Mathematische Zeitschrift*, 90(1):1–8, 1965.

[Sta91] John R Stallings. Foldings of g-trees. In *Arboreal Group Theory*, pages 355–368. Springer, 1991.

[T$^+$88] William P Thurston et al. On the geometry and dynamics of diffeomorphisms of surfaces. *Bulletin (new series) of the american mathematical society*, 19(2):417–431, 1988.

[TZ12] Katrin Tent and Martin Ziegler. *A course in model theory*, volume 40. Cambridge University Press, 2012.

[Wei02] Richard Weidmann. The nielsen method for groups acting on trees. *Proceedings of the London Mathematical Society*, 85(1):93–118, 2002.

[Wil09] Henry Wilton. Solutions to bestvina & feighn's exercises on limit groups. *Geometric and cohomological methods in group theory, London Mathematical Society Lecture Note Series*, 358:30–62, 2009.

Curriculum Vitae

Name:	Javier
First surname:	de la Nuez
Second surname:	González
Date of birth:	28/07/1988
Place of birth:	Madrid, Spain

Sep 11	–		PhD studies at the Institüt für mathematische Logik und Grundlagenforschung, Universität Münster, advisor Prof. Katrin Tent
Sep 09	–	Jun 10	Erasmus stay at the Università degli Studi di Pisa, Pisa (Italy)
Sep 06	–	Jun 11	Diploma studies in mathematics (Licenciatura en Ciencias Matemáticas) at Universidad Complutense, Madrid (Spain)
	–	Jun 06	Highschool studies at Liceo Scientifico Italiano *Enrico Fermi*, Madrid