UNIVERSITY OF SOUTHAMPTON

Faculty of Social Sciences
School of Mathematics

Barking Up the Right Tree: Group Actions on Trees, Automorphisms and Separability

by

Naomi Grace Andrew

A thesis for the degree of
Doctor of Philosophy

July 2021

University of Southampton

Abstract

Faculty of Social Sciences
School of Mathematics

Doctor of Philosophy

Barking Up the Right Tree: Group Actions on Trees, Automorphisms and Separability

by Naomi Grace Andrew

Actions on trees are powerful tools for understanding the structure of a group. In this thesis, we use them to understand separability and automorphisms of free products, and automorphisms of free-by-cyclic groups.

This is a three paper thesis; the main body of the work is contained in the following papers:

[1] Naomi Andrew. A Bass–Serre theoretic proof of a theorem of Burns and Romanovskii. Preprint, July 2021, available at arXiv:2107.02548.

[2] Naomi Andrew. Serre's property (FA) for automorphism groups of free products. *J. Group Theory*, 24(2):385–414, 2021.

[3] Naomi Andrew and Armando Martino. Free-by-cyclic groups, automorphisms and actions on nearly canonical trees. Preprint, June 2021, available at arXiv:2106.02541.

In [1], we use properties of actions on trees – described in the combinatorial language of graphs of groups, due to Bass and Serre – to re-prove that free products of subgroup separable groups are themselves subgroup separable.

In [2], we suppose G is a free product of groups and investigate when $\mathrm{Aut}(G)$ admits actions on trees. Under the assumption that the factor groups are freely indecomposable and not $\mathbb{Z}$ (for example, if they are finite) this depends only on a count of the isomorphism classes appearing in the decomposition. To build actions on trees we use length functions and a theorem of Culler and Morgan; to rule them out we use commutation relations within the automorphism groups.

In [3], we investigate the outer automorphisms of free-by-cyclic groups, and in some cases prove that they are finitely generated. To do this we introduce the notion of a "nearly canonical" action on a tree, construct such an action for certain free-by-cyclic groups, and use these actions to understand the outer automorphisms.

Contents

List of Figures **vii**

Declaration of Authorship **ix**

Acknowledgements **xi**

Definitions and Abbreviations **xiii**

Introduction **1**

1 Groups and Presentations 2
2 Actions on Trees 5
3 Residual Finiteness and Subgroup Separability 24
4 Automorphisms of Groups 28
5 Deformation Spaces and Canonical Actions 32
6 Out(F_n) and Outer Space 36

References **41**

1 A Bass–Serre theoretic proof of a theorem of Burns and Romanovskii **47**

1 Introduction 47
2 Graphs of Groups 48
3 Kurosh rank and finite index subgroups of free products 53
4 Subgroup separability for free products 56
References 61

2 Serre's Property (FA) for automorphism groups of free products **63**

1 Introduction 63
2 Background 65
2.1 Actions on trees 65
2.2 Automorphisms of free products 67
2.3 Bass-Serre theory 69
2.4 Translation length 70
3 Sufficient conditions 71
4 Necessary conditions 78
References 87
A A presentation of Out(G) 88

3 Free-by-cyclic groups, automorphisms and actions on nearly canonical trees **93**

1 Introduction 93

1.1 Free-by-cyclic groups . . . 93
2 Background . . . 96
2.1 Notation, Actions on trees and Bass-Serre Theory . . . 96
2.2 Length Functions and Twisting Actions by Automorphisms . . . 98
2.3 Trees of cylinders . . . 100
2.4 Automorphisms of free groups . . . 101
3 Extending actions to the automorphism group . . . 104
3.1 Canonical actions and nearly canonical actions . . . 104
3.2 Automorphisms which preserve a splitting, and a theorem of Bass–Jiang . . . 107
4 Exponential growth . . . 110
4.1 Relative Hyperbolicity . . . 110
5 Linear growth . . . 115
5.1 Strategy . . . 115
5.2 Constructing a tree . . . 115
5.3 Reducing to free-by-finite groups . . . 118
5.4 McCool groups for free-by-finite groups . . . 122
6 Quadratic growth . . . 127
6.1 Strategy . . . 127
6.2 Normal forms and a tree to act on . . . 127
6.3 Invariance of the tree . . . 129
6.4 Calculating the tree of cylinders, T_c . . . 130
References . . . 136

List of Figures

1 The Cayley graph for $F_2 = \langle a, b : \rangle$, showing the hyperbolic axis of a in blue. 7
2 Constructing $\text{axis}(g)$ when g does not fix a point 9
3 A schematic of the situation described in the proof of Lemma 2.11 10
4 The segment $[u, v]$ is translated and contained in $\text{axis}(g)$; the segment $[v', u]$ is not. 11
5 The image under h of the axis of g is the axis of $h^{-1}gh$. 12
6 If the axes of g and h do not intersect, then the axis of gh (shown in blue) intersects both. 12
7 Example of graphs of groups 17
8 Calculating translation length for $g = php^{-1}$ 25
9 Three graphs of groups for $G = A * B * C$. 32
10 A graph of groups for $G = A * B * C$ with an extra edge group. 34
2.1 The graph T' described in Lemma 3.6(2) 74
2.2 Graphs of groups realising each G in Proposition 4.8. 82
2.3 The graphs of groups at each stage of Proposition 4.10 84
3.1 T_0, as described in Corollary 6.2.2 129

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

(1) This work was done wholly or mainly while in candidature for a research degree at this University;

(2) Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

(3) Where I have consulted the published work of others, this is always clearly attributed;

(4) Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

(5) I have acknowledged all main sources of help;

(6) Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

(7) Parts of this work have been published as:

[1] Naomi Andrew. A Bass–Serre theoretic proof of a theorem of Burns and Romanovskii. Preprint, July 2021, available at `arXiv:2107.02548`.

[2] Naomi Andrew. Serre's property (FA) for automorphism groups of free products. *J. Group Theory*, 24(2):385–414, 2021.

[3] Naomi Andrew and Armando Martino. Free-by-cyclic groups, automorphisms and actions on nearly canonical trees. Preprint, June 2021, available at `arXiv:2106.02541`.

Signed:.. Date:..................

Acknowledgements

First, thank you Armando. Thanks for introducing me to this corner of maths, your willingness to talk, for encouraging me, and for persuading me to do things I was ready for even if I didn't believe it beforehand.

Thank you Ashot and Gilbert, for your careful reading, detailed comments and for an interesting and useful discussion.

Southampton has been a wonderfully supportive environment as I entered mathematics: thanks to everyone (students, postdocs, academics, admin staff) for making it so. Mathematically I've grown through conversations, seminars, reading groups (thanks especially to everyone involved in the $\mathrm{Out}(F_n)$ group this year) and with your encouragement and feedback. But equally important has been the coffee breaks, crosswords, pub trips, and board game sessions. This last year and a bit has been very peculiar: thank you all so much for persevering with finding ways to keep a sense of community while we've been working away from the office and Southampton.

Also, thank you to everyone who has welcomed me (in person at first, and virtually more recently) to your own institutions, and all those who've made conferences enjoyable (for reasons including but by no means limited to the mathematics).

This would have been impossible without the support of so many people outside maths too. Thanks to my housemates (and road-mates); to Concert Band for giving me something else to work towards; and to the (extended) Connaught rangers (the free groups are doing pretty well this week!). Thanks too to everyone at Portswood Church, especially my home group: for the time, walks, meals and faithful encouragement. Finally thanks to my family for all your support through the years, without which I would not be who or where I am today.

Definitions and Abbreviations

$\mathbb{Z}$	The integers
$\mathbb{R}$	The real numbers
S_n	The symmetric group on n points
F_n	The free group on n generators
$\mathrm{GL}, \mathrm{SL}, \mathrm{PGL}$	The general, special and projective general linear groups (usually taken over $\mathbb{Z}$)
$G * H$	The free product of G and H
$N \rtimes H$	The semidirect product with kernel N and quotient H
$H \wr S_n$	The permutational wreath product $H^n \rtimes S_n$
$\mathrm{Aut}(G)$	The automorphism group of G
$\mathrm{Inn}(G)$	The inner automorphisms of G
$\mathrm{Out}(G)$	The outer automorphism group of G, $\mathrm{Aut}(G) / \mathrm{Inn}(G)$
$\mathrm{Out}(G; \{G_i\})$	The subgroup of $\mathrm{Out}(G)$ preserving the conjugacy classes of the subgroups $\{G_i\}$
$\mathrm{Mc}(G; \{G_i\})$	The subgroup of $\mathrm{Out}(G)$ restricting to conjugation on each of the subgroups $\{G_i\}$
$\mathrm{Aut}_T(G)$ or $\mathrm{Aut}^T(G)$	The subgroup of $\mathrm{Aut}(G)$ preserving the action on a given G-tree T
$\mathrm{Out}^T(G)$	The image of $\mathrm{Aut}^T(G)$ in $\mathrm{Out}(G)$
$\mathrm{FR}(G)$	When G is a free product, the subgroup of $\mathrm{Aut}(G)$ generated by partial conjugations
$\mathrm{Fact}(G)$	When G is a free product, the subgroup of $\mathrm{Aut}(G)$ generated by automorphisms of its free factors
$\mathrm{Perm}(G)$	When G is a free product, the subgroup of $\mathrm{Aut}(G)$ corresponding to (fixed) permutations of its free factors
G_x	The subgroup of G stabilising x
$\mathrm{Fix}(H)$	The fixed subgroup of H
$\\|g\\|_T$	The translation length of g in an action on T
$N_G(H)$	The normaliser of H in G
$C_G(H)$	The centraliser of H in G
$Z(G)$	The centre of G
$g \sim h$	g is conjugate to h

$g \sim_K h$	g is conjugate to h by an element of K
$[G : H]$	The index of the subgroup H of G

Introduction

In this introduction we provide background material and context for the three papers that form the main body of the thesis. Note that the ordering of the papers within the thesis does not reflect publication or posting dates but rather the order I worked on them during my PhD.

Papers 1 and 2 are single author papers; Paper 3 is a joint paper with Armando Martino (my supervisor).

The statements and proofs in Section 4 were largely my work; the statements and proofs in Sections 5 and 6 were formulated and improved together over the course of several meetings. (For example, Armando suggested using limiting trees in the rank 3 linear case, while I spotted the connection to trees of cylinders and the extension to the general linear case.)

I wrote the majority of the first draft; we both reviewed and revised the manuscript before finalising it.

All three papers rely substantially on Bass–Serre theory, which studies groups with actions on trees via those actions. How they do so varies: all three consider actions for specific classes of groups (free products, automorphisms of free products, or free-by-cyclic groups and their automorphisms), and take differing perspectives on these actions. Paper 1 is more combinatorial, using covers and immersions from [2] to construct subgroups with desired properties. Paper 2 is concerned with when an action exists, using the geometry of trees to construct large fixed subgroups, and Paper 3 involves constructing (nearly) canonical actions, in order to understand properties of the automorphism group.

For this reason, this introduction is not divided according to the papers: however, it is not necessary to read the whole introduction for each paper. All three papers rely on material in Sections 1 and 2; for the other chapters, the dependencies are:

Paper 1 Section 3;

Paper 2 Section 4;

Paper 3 Sections 4-6.

In particular, Sections 4-6 are independent of Section 3. Material on free groups is included throughout, since they make for good examples. However, they are strictly speaking only necessary for the final introductory section (on $\mathrm{Out}(F_n)$) and Paper 3.

1 Groups and Presentations

We begin with a brief review of free groups and group presentations, taken largely from [54, Sections I.1 and II.1-2]

Given a set X, the free group on X, $F(X)$ is defined as follows.

Consider words $x_1 x_2 \ldots x_n$, where each x_i is an element of $X \cup X^{-1}$ (an element of X^{-1} is a "formal inverse" x^{-1} of some element x of X). An elementary reduction consists of removing a pair yy^{-1} or $y^{-1}y$ from a word, and we impose the equivalence relation generated by all elementary reductions: this sets equivalent any words which can be transformed into each other by a sequence of elementary reductions, and insertions of pairs yy^{-1} or $y^{-1}y$. Every word is equivalent to a unique reduced word under this relation. We also allow the empty word.

The free group $F(X)$ is the set of reduced words under concatenation (followed by reduction, if necessary). It satisfies a universal property: every map from X to a group G factors through the inclusion $X \hookrightarrow F(X)$ (as words of length one) and a unique homomorphism $F(X) \to G$:

$$\begin{array}{ccc} & & F(X) \\ & \nearrow & \downarrow \\ X & \longrightarrow & G \end{array}$$

In fact, this universal property may be taken as a definition; the definition given above serving to verify that objects exist which satisfy this property. Two free groups $F(X)$ and $F(Y)$ are isomorphic if and only if the sets X and Y have the same cardinality; when $|X| = n$, we write F_n.

Given any group G and a subset $X \subseteq G$, say that X *generates* G if for every element g of G we can write $g = x_1 x_2 \ldots x_n$ as a product of elements of X (the trivial element is the empty product, in this context). If there is some finite set X which generates G then G is said to be *finitely generated*; sometimes written $G = \langle X \rangle$.

If X generates G there is a homomorphism from $F(X)$ to G by the universal property for free groups; so to describe G we need only to give a generating set X and describe the kernel of this map, N. Sufficient information is a "normal generating set" of this subgroup: that is a set of elements R where every element of N is a product of conjugates of elements of N. The set R is said to be a set of *relators*. Then G has a presentation $\langle X : R\rangle$.

Sometimes it is more instructive to give elements of R as *relations* of the form $u = v$ (where u and v are elements of the free group $F(X)$). The information is always equivalent: a relator r can be viewed as a relation $r = 1$, and so a relation $u = v$ implies uv^{-1} is a relator.

A group G is said to be *finitely presented* if it has a presentation $\langle X : R\rangle$ where both X and R are finite sets.

For example, $\langle X :\rangle$ is a presentation of $F(X)$, and it is a finite presentation whenever X is a finite set. Another example is $\langle x : x^n\rangle$, which is a presentation of the finite cyclic group with order n. We also have that $\mathbb{Z}^2 = \langle x, y : x^{-1}y^{-1}xy\rangle$ (giving R as a relator) or equivalently $\langle x, y : xy = yx\rangle$ (giving R as a relation).

There are many different presentations for every group; and even many finite presentations for any finitely presented group. One way to find more is by using Tietze transformations:

Definition 1.1. Tietze transformations are the following two kinds of operation on a group presentation:

- Introducing or removing a relation implied by the other relations (that is, any relation implying a relator already contained in the normal subgroup N of $F(X)$).
- Introducing a new generator and a relation setting it equal to some element of the group, or removing such a generator and relation (substituting all its occurrences in the other relations).

The first kind alters just the relations; the second kind the relations as well as the generators. In some sense, these transformations describe all the ways to change a presentation:

Theorem 1.2 ([54, Proposition II.2.1])**.** *Two finite presentations $\langle X_1 : R_1\rangle$ and $\langle X_2 : R_2\rangle$ define isomorphic groups if and only if there is a finite sequence of Tietze transformations taking one to the other.*

Note that the theorem does not explicitly construct the transformations required, so it is less useful than it might at first appear.

New Groups From Old

We introduce a few ways of combining groups. First,

Definition 1.3. Suppose two groups are defined via presentations as $G_1 = \langle X_1 : R_1 \rangle$ and $G_2 = X_2 : R_2$. Their *free product* is defined as $G_1 * G_2 = \langle X_1 \cup X_2 : R_1 \cup R_2 \rangle$.

Free products satisfy a universal property similar to free groups: the inclusions of G_1 and G_2 into $G_1 * G_2$ are such that that given homomorphisms $G_1 \to H$ and $G_2 \to H$, they factor through the inclusions and a unique homomorphism $G_1 * G_2 \to H$.

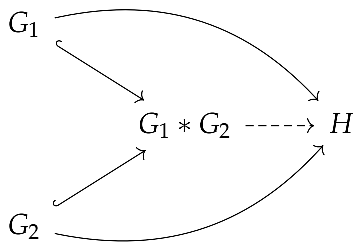

We can also define a free product of any set of groups, by combining the presentations in the same way. (The free product of no groups is the trivial group.) The same property (and diagram) can be constructed for general free products; in fact they are the coproduct for groups.

We also note here two other constructions:

Definition 1.4. Let $N = \langle X_1 : R_1 \rangle$ and $H = \langle X_2 : R_2 \rangle$ be groups and Ψ a map $H \to \mathrm{Aut}(N)$. The semidirect product of N by H with respect to Ψ is

$$N \rtimes_\Psi H = \langle X_1, X_2 : R_1, R_2, h^{-1}nh = \Psi(h)(n) \quad \forall n \in N, h \in H \rangle.$$

That is, conjugation by elements of H acts as a specified automorphism on N. A particular kind of semidirect product is known as a wreath product:

Definition 1.5. A permutational wreath product $H \wr S_n$ is a semidirect product $H^n \rtimes S_n$, with the action of S_n being to permute the factors of H^n.

(This definition works in more generality, with the right hand group equipped with an action on some set X and taking the kernel of the semidirect product to be $H^{|X|}$.)

Group Actions

Throughout the introduction (and the papers) we consider actions of groups on objects: often but not always trees. A group action of G on an object X is a homomorphism from G to $\mathrm{Aut}(X)$, the transformations of X; equivalently it is a map $X \times G \to X$ satisfying $x \cdot 1 = x$ and $x \cdot (gh) = (x \cdot g) \cdot h)$. Throughout, actions are generally taken to be on the right. Whenever G acts on a object X, X is called a G-object; $\mathrm{Fix}(g)$ refers to those elements $x \in X$ where $xg = x$, and G_x to the stabiliser of x, $\{g \in G | xg = x\}$.

A map f between two G-objects X, Y is *equivariant* if for every element x of X and g of G we have that $f(xg) = (f(x))g$. There is also a more general notion: given a homomorphism $\varphi : G \to H$, a G-object X and an H-object Y, a map $f : X \to Y$ is *φ-equivariant* if for every element x of X and g of G, $f(xg) = (f(x))(\varphi(g))$. (If we take φ to be the identity homomorphism, this recovers the simpler notion.)

2 Actions on Trees

Trees are reasonably simple combinatorial objects, which makes them powerful tools for studying groups which act on them. We largely follow [63] and [2]; there are other expositions in for example [26]. We begin by exploring the geometry of trees and some properties of group actions on trees, before moving onto the framework of Bass–Serre theory used to study groups with these actions. Note that most sources put the actions on the left; here they will be on the right.

Definition 2.1 (Serre)**.** A graph Γ consists of a set of vertices $V\Gamma$ and a set of edges $E\Gamma$, together with two maps: $\iota : E\Gamma \to V\Gamma$; and an involution $E\Gamma \to E\Gamma, e \mapsto \bar{e}$. We also define $\tau : E\Gamma \to V\Gamma, \tau(e) = \iota(\bar{e})$. An *orientation* of Γ is a choice of one edge from each pair $\{e, \bar{e}\}$.

For an edge e the vertices $\iota(e)$ and $\tau(e)$ are referred to as the *initial* and *terminal* vertices of e respectively. Since $\tau(e) = \iota(\bar{e})$, the initial vertex of e is the terminal vertex of $\bar{e}$, and conversely.

These graphs are combinatorial objects, but we can – and frequently do – think of them geometrically: as CW complexes, with vertices corresponding to zero cells and edge-pairs corresponding to one cells. The attaching maps are then $\iota(e)$ and $\tau(e)$. They are also metric spaces, by giving each geometric edge length 1, identifying it with the interval $[0, 1]$, and taking the length of the shortest path between any two vertices. This is the *edge path metric*. This definition allows loops and double edges, so these graphs are not simplicial. However, by subdividing twice, any graph can be made into a simplicial graph.

The formalism of edge-pairs (with twice as many edges as geometrically required) is not strictly necessary, but we do need a way to distinguish direction of travel along a geometric edge, and this is one such way: traversing e and $\bar{e}$ correspond to travelling in different directions.

Definition 2.2. A *path* of length n is a sequence of edges $e_1 e_2 \ldots e_n$ such that $\tau(e_i) = \iota(e_{i+1})$ for all $0 < i < n$. A *cycle* is a path where in addition $\tau(e_n) = \iota(e_1)$. A path or a cycle is called *reduced* if no pair e_ie_{i+1} is of the form $e\bar{e}$, and a cycle is *cyclically reduced* if it is reduced and in addition e_ne_1 is not of the form $e\bar{e}$.

We can extend the notion of "reversing" to paths as a whole: if $p = e_1e_2 \dots e_n$, then we say $\overline{p} = \overline{e}_n \dots \overline{e}_1$.

Definition 2.3. A graph Γ is *connected* if for every pair of vertices $u \neq v$ in Γ there is a path with $\iota(e_1) = u$ and $\tau(e_n) = v$.

In this case, we will also have that $\overline{p}$ is a path from v to u.

A tree is a particularly simple kind of graph:

Definition 2.4. A *forest* is a graph with no cycles; a *tree* is a connected forest.

Alternatively, we can express the property of "being a tree" in terms of paths:

Lemma 2.5. *A graph T is a tree if and only if there is a unique reduced path between any two vertices.*

Proof. A graph is connected precisely when there are paths joining any two vertices, so it remains to check that the uniqueness of these paths is equivalent to there being no cycles.

Suppose T is a tree. If there were two distinct reduced paths p, q joining u and v, then $p\overline{q}$ must (after reduction, if necessary) be a non-trivial cycle. So in fact the path joining u and v must be unique.

Conversely, suppose there is a reduced cycle $e_1 \dots e_n$ in the graph. Then for any $0 < i < n$, the paths $e_1 \dots e_i$ and $\overline{e}_n \dots \overline{e}_{i+1}$ both join $\iota(e_1)$ to $\tau(e_i)$, and they are reduced (since the cycle was) and distinct for the same reason: since the cycle is reduced, $e_i \neq \overline{e}_{i+1}$. □

Definition 2.6. A *graph map* $f : \Gamma_1 \to \Gamma_2$ consists of maps $E(\Gamma_1) \to E(\Gamma_2)$ and $V(\Gamma_1) \to V(\Gamma_2)$ satisfying $f(\overline{e}) = \overline{f(e)}$ and $f(\iota(e)) = \iota(f(e))$.

This is a fairly restrictive definition of a graph map, but it is the most useful for group actions on graphs or trees. Other notions – allowing edges to collapse, or to be sent to paths with length greater than 1 – are possible, and are used for example in the contexts of deformation spaces (see Section 5).

We say a group acts on a graph if it does so by invertible graph maps: edges are sent to edges, vertices to vertices, preserving adjacency. We also demand that the actions are without inversions: that $eg \neq \overline{e}$ for all edges e and group elements g. This ensures that the the quotient object is itself a graph: its vertices and edges are the orbits of vertices and edges (respectively) under the action. It also means that an edge stabiliser G_e is

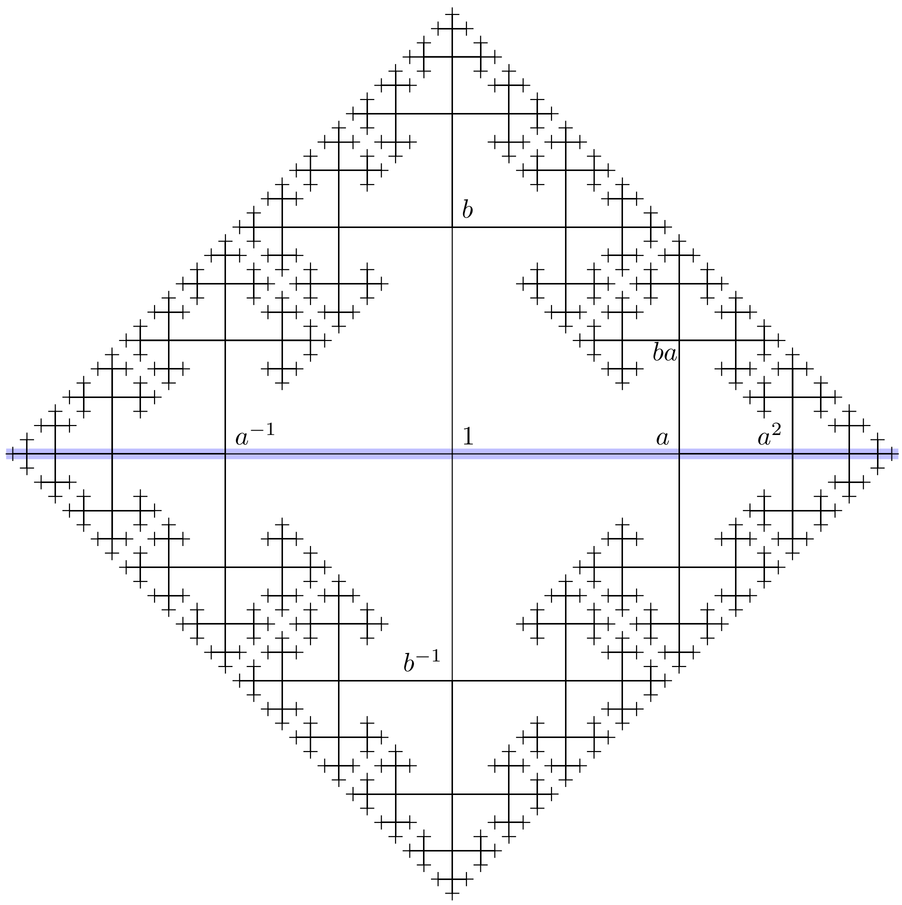

FIGURE 1: The Cayley graph for $F_2 = \langle a, b : \rangle$, showing the hyperbolic axis of a in blue.

contained in the adjacent vertex stabilisers. This does not actually restrict the actions we can consider, since we can always achieve it by subdividing every edge.

Viewing the graph as a metric space, this corresponds to an action by isometries: the shortest path between two vertices must be sent to the shortest path between their images. We largely consider groups acting on trees. Within a tree, use $[v, w]$ to denote the unique shortest path (or *geodesic*) joining the vertices v and w.

Example 2.7. The free group F_2 acts freely on the infinite four-regular tree.

We can construct this tree T as follows: fix a basis $\{a, b\}$ for F_2. The vertex set is the set of elements of F_2, and there is an edge between two vertices u and v if $u = sv$ with $s \in \{a, b, a^{-1}, b^{-1}\}$. This tree is shown in Figure 1. (This is a Cayley graph for this presentation $F_2 = \langle a, b : \rangle$, although the multiplication is taking place on the left to facilitate a right action.)

We need to check that T is a tree: it is connected since $\{a, b\}$ generates F_2. Suppose there is a reduced cycle in the graph. Then (picking any vertex v in the cycle) we must have that $s_1 \dots s_n v = v$, where $s_i \in \{a, b, a^{-1}, b^{-1}\}$ are the generators (or inverses) defining each edge. But these are elements of F_2, and so equality only holds if $s_1 \dots s_n = 1$.

Since the cycle was reduced this word must be too: but since there is only one reduced word in each equivalence class, it must also be empty. So there cannot be a non-trivial reduced cycle in T.

There is a right action of F_2 on T: right multiplication permutes the vertices, since vg is still an element of F_2. We need to check adjacency: suppose we have an edge $[v, sv]$: then under g these vertices become vg and svg, and since these differ by s (on the left) there is another edge $[vg, svg]$. If there was an inversion, it would imply (amongst other things) that for some $v, g \in F_2$, $vg^2 = v$ and so that $g^2 = 1$: but there are no torsion elements in F_2 so this cannot happen.

By the same argument, vertex (and hence edge) stabilisers are trivial. The quotient graph is a "rose with two petals": there is one orbit of vertices, and two of edges. One edge orbit consists of edges of the form $[v, av]$ and the other of edges $[v, bv]$.

Note that the fundamental group of this graph is F_2 again: this is a simple example of the fact that the fundamental group of a space acts freely on its universal cover (when it exists) [43, Proposition 1.39]. The same arguments hold for F_n, giving a regular tree with valence $2n$ and a quotient graph with one vertex and n edges. In fact (see for example [63, Theorem I.4]) a group acts freely on a tree if and only if it is a free group: the other direction can be seen topologically, by an argument about fundamental groups and universal covers, or more combinatorially, by equivariantly collapsing a maximal tree in the quotient graph in order to recover a tree and action realising the Cayley graph again.

Stallings [65] uses this idea (expressed purely combinatorially) to express statements about free groups in terms of coverings of graphs. See for example the proof of Theorem 3.6, which follows Stallings' proof.

Of course most actions are not free, and the stabilisers can complicate the situation. Here are some general facts about group actions on trees.

Lemma 2.8. *If a group G acts on a tree T, then each element g either fixes a point or preserves and translates along a line.*

Proof. Suppose g does not fix a point, and let v be any vertex. If $vg = vg^{-1}$ then the path $[v, vg]$ is inverted by g, in which case the midpoint is a fixed vertex (if it has even length) or the middle edge is inverted (if it has odd length). Neither is possible, so vg and vg^{-1} are distinct. The paths $[v, vg^{-1}]$ and $[v, vg]$ intersect in a segment $[v, u]$, which is strictly shorter than $[v, vg^{-1}]$, although possibly has length 0, if $v = u$.

Consider the vertices ug and ug^{-1}. Again, these must be distinct (and distinct from u), otherwise there is a fixed point or inverted edge. Say there is a segment $[x, y_1, \ldots, y_n, z]$ if the points $y_1, \ldots, y_n$ lie on the geodesic $[x, z]$ in the given order.

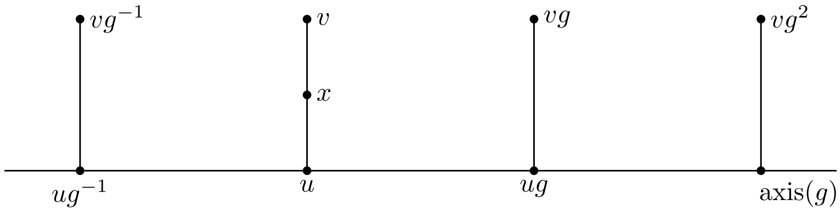

FIGURE 2: Constructing axis(g) when g does not fix a point

Since there is a segment $[vg^{-1}, u, v]$, there is also a segment $[v, ug, vg]$. Similarly there is a segment $[vg^{-1}, ug^{-1}, v]$ because there is a segment $[v, u, vg]$. Since there are segments $[v, ug, vg]$ and $[v, u, vg]$, there must either be a segment $[v, u, ug, vg]$ or $[v, ug, u, vg]$. (As u and ug are distinct, only one of these occurs.) We will see that it must be $[v, u, ug, vg]$.

Suppose not, so instead there is a segment $[v, ug, u, vg]$. (For instance, ug might be at the point labelled x in Figure 2.) Then, applying g^{-1}, there is a segment $[vg^{-1}, u, ug^{-1}, v]$. That is, ug^{-1} also lies between u and v.

But then ug and ug^{-1} are the same point, since they both lie in $[u, v]$ and are the same distance from u, contradicting the fact they are distinct.

So there are segments $[v, u, ug, vg]$ and $[vg^{-1}, ug^{-1}, u, v]$ (by applying g^{-1}). These intersect only along $[u, v]$, so $[ug^{-1}, u]$ and $[u, ug]$ intersect only at u. Therefore the concatenation of $[ug^{-1}, u]$ and $[u, ug]$ cannot contain a backtrack and so is exactly the segment $[ug^{-1}, u, ug]$.

Similarly, $[ug^{n-1}, ug^n]$ and $[ug^n, ug^{n+1}]$ concatenate to give $[ug^{n-1}, ug^n, ug^{n+1}]$. So the images of $[u, ug]$ under g^n as n varies form a line, preserved by g, and on which g acts by a translation of length $d(u, ug)$. □

The point u constructed in this proof is sometimes called the Y-*point* of vg^{-1}, v, and vg, and does not depend on the order in which vg^{-1}, v, and vg are taken [15, Lemma 2.1.2].

Definition 2.9. An element preserving a point is called *elliptic*. An element translating along a line is called *hyperbolic*, and that line its *hyperbolic axis* In addition, if a subgroup H of G has a common fixed point, the subgroup is again called elliptic.

Lemma 2.10. *Suppose G acts on a tree T and $H \leq G$ has non-empty fixed point set. Then* $\mathrm{Fix}(H)$ *is a subtree of T.*

Proof. Since edge stabilisers are contained in vertex stabilisers, $\mathrm{Fix}(H)$ will be a subgraph and hence a subforest of T. We need to show it is connected. To do this, suppose u and v are fixed by H, and let p be the unique reduced path joining them. Then for all $h \in H$, we have that $uh = u$, $vh = v$, and ph is the unique reduced path joining them. But since it is unique, this implies that $ph = p$, so all edges and vertices in p are again

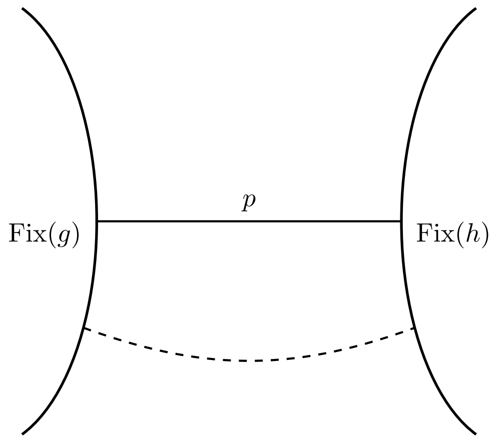

FIGURE 3: A schematic of the situation described in the proof of Lemma 2.11

fixed by h: any pair of points fixed by h are joined by a path also fixed by h. This was true for all elements of H, so $\mathrm{Fix}(H)$ is connected. □

This strategy of arguing about a unique path joining two vertices can be very fruitful. Here is another example:

Lemma 2.11. *Suppose a group G acts on a tree, and has elements g, h which are elliptic and commute. Then they have a common fixed point.*

Proof. Let v be fixed by g, and consider its image under h: $vh = vgh = vhg$, so $\mathrm{Fix}(g)$ is preserved by h. Similarly, $\mathrm{Fix}(h)$ is preserved by g. Either these fixed point sets intersect, or there is a unique path p joining them. (See Figure 3: Uniqueness follows since if there were two such paths joining them, there would be a cycle consisting of those paths, and paths contained in $\mathrm{Fix}(g)$ and $\mathrm{Fix}(h)$ joining the endpoints.) But since both elements preserve both fixed point sets, they must both fix p, and there was a common fixed point after all. □

The same argument shows that two commuting elliptic subgroups have a common fixed point.

There is a useful characterisation of the hyperbolic axis of an element in terms of the way an edge is moved by the element:

Lemma 2.12. *Suppose T is a G-tree, g an element of G, and u a vertex of T that is not fixed by g. Let v be the first vertex (after u) on the path $[u, ug]$. Then vg is outside $[u, ug]$ if and only if g is hyperbolic and u and v are contained in* $\mathrm{axis}(g)$.

Proof. Certainly if vg is contained in $[u, ug]$ then g is not translating along a line including that segment. So suppose vg is outside $[u, ug]$. Then the intersection of $[u, ug]$ and $[ug, ug^2]$ is only the vertex ug, and as in the proof of Lemma 2.8 the union of the images of the segment $[u, ug]$ under g form the axis of g. □

FIGURE 4: The segment $[u, v]$ is translated and contained in $\text{axis}(g)$; the segment $[v', u]$ is not.

The situation is shown in Figure 4. In this case we say the edge $[u, v]$ is *translated* by g. (Other terminology is possible; for example Serre uses "coherent under g".) If this edge is not translated by g, then the axis or fixed point set of g includes the midpoint of $[u, ug]$.

This argument is valid even if $v = ug$; alternatively, one can equivariantly subdivide the first edge to guarantee a vertex between them.

A group can act on multiple trees in multiple ways, but some of these differences are more interesting than others. For example, given an action on a tree, we can add a new edge and then a leaf vertex at every vertex: the new edges and vertices will have to share the stabiliser of the vertex to which they are attached. Many more "new" actions can be constructed in similar ways, by adding subforests and extending the original action. The tree we started with is a subtree of the tree we have constructed, and the action can be restricted back to this subtree – we didn't change the existing orbits, only added new ones. Often we might want to consider an actions which cannot be restricted to an invariant subtree, which leads to the following definition.

Definition 2.13. An action on a tree is called *minimal* if it has no proper invariant subtrees.

It turns out that – in most situations – we can pass to a minimal subtree.

Proposition 2.14. *If a group G acting on a tree T has hyperbolic elements, there is a unique minimal invariant subtree of T which is the union of the hyperbolic axes of elements of G.*

Proof. First, the collection of hyperbolic axes is G-invariant, as $\text{axis}(g) \cdot h = \text{axis}(h^{-1}gh)$. To see this, consider the situation as shown in Figure 5; the four colinear points on $\text{axis}(g)$ must be sent by h to four colinear points in the same order. Then Lemma 2.12 implies that they lie on $\text{axis}(h^{-1}gh)$. The same is true for any points along the axis of g, so the image must be exactly $\text{axis}(h^{-1}gh)$.

To show that it is a subtree, we need to show it is connected. Suppose g and h are hyperbolic elements, and consider their axes. If they do not intersect, we consider $\text{axis}(gh)$. See Figure 6: that the blue shaded section is part of $\text{axis}(gh)$ can be checked by using Lemma 2.12 with the vertices u and v. In particular, gh is hyperbolic, and its axis intersects the axes of both g and h.

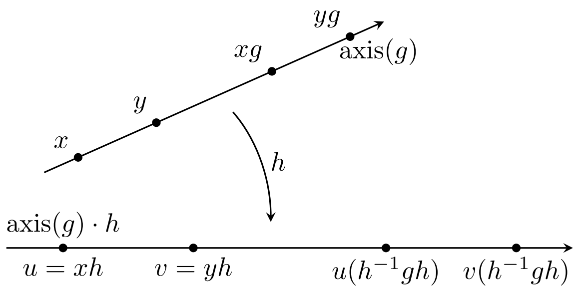

FIGURE 5: The image under h of the axis of g is the axis of $h^{-1}gh$.

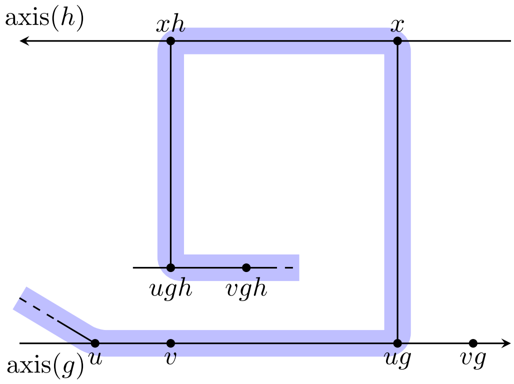

FIGURE 6: If the axes of g and h do not intersect, then the axis of gh (shown in blue) intersects both.

Finally we show that any invariant tree contains every axis. To see this, consider any vertex v in an invariant tree, and its image vg under any hyperbolic element. The invariant tree must contain the segment $[v, vg]$ joining them, and this contains part of the axis (as in the proof of Lemma 2.8). Let u be the nearest point to v in axis(g): then ug is the nearest point to vg. So the invariant tree contains $[u, ug]$, and iterating g and g^{-1} will translate this segment across the whole axis, which must again be contained in the invariant tree. □

If there is a global fixed point, then depending on the situation we may want to consider one point (though this is usually not unique) or all (though this is usually not minimal) of the fixed point set in similar situations. There are actions on trees where despite every element being elliptic there is no global fixed point (see the proof of Theorem 2.32), but these cannot be minimal.

Graphs of groups and Bass–Serre theory

The idea of a graph of groups is to combinatorially encode sufficient information in a (usually finite) graph to describe a group action on a tree. It turns out (see Theorem 2.24) that what is required is the quotient graph and the vertex and edge stabilisers.

This is not the case for general actions on graphs – for example, taking Γ to be a cycle with n vertices acted on by the cyclic group C_n, then the action is always free, and the quotient always has one vertex and one edge. So we could not hope to recover the original group or action from this data. But the (topologically) simpler structure of trees means that there is enough information to reconstruct the action.

The definitions here are closest to those used in [2], although it is worth noting that here they are changed so the action on the Bass–Serre tree will be on the right.

Definition 2.15. A *graph of groups* $\mathcal{G}$ consists of a graph Γ together with groups G_v for every vertex and $G_e = G_{\overline{e}}$ for every (oriented) edge, and monomorphisms $\alpha_e : G_e \to G_{\tau(e)}$ for every (oriented) edge.

The fundamental group of a graph of groups can be defined in two ways, with respect to a maximum tree of the graph, and by considering loops in the graph of groups. We take the second route, which simplifies some subsequent calculations.

Definition 2.16 (Paths)**.** Let $F(\mathcal{G})$ be the group generated by all the vertex groups and all the edges of $\mathcal{G}$, subject to the additional relations $e\alpha_e(g)\overline{e} = \alpha_{\overline{e}}(g)$ for $g \in G_e$. Note that taking $g = 1$ this gives that $e^{-1} = \overline{e}$, as expected.

Define a *path* (of length n) in $F(\mathcal{G})$ to be a sequence $g_0e_1g_1 \ldots e_ng_n$, where each e_i has $\iota(e_i) = v_{i-1}$ and $\tau(e_i) = v_i$ for some vertices v_i (so there is a path in the graph), and each $g_i \in G_{v_i}$. A *loop at* v_0 is a path where $v_0 = v_n$.

Some authors treat the set of all paths in $F(\mathcal{G})$ as a *groupoid*, denoted $\pi(\mathcal{G})$ in [2], and called the *fundamental groupoid*; the multiplication is inherited from $F(\mathcal{G})$, with the product of two paths defined precisely when their endpoints match.

Definition 2.17 (Reduced paths)**.** A path is *reduced* if it contains no subpath of the form $e\alpha_e(g)\overline{e}$ (for $g \in G_e$). A loop is *cyclically reduced* if, in addition to being reduced, $e_n(g_ng_0)e_1$ is not of the form $e\alpha_e(g)\overline{e}$.

Every path is equivalent (by the relations for $F(\mathcal{G})$) to a reduced path, and similarly every loop is equivalent to both a reduced loop and a cyclically reduced loop. In general these reduced representations are not unique, although all equivalent reduced paths (and cyclically reduced loops) will have the same edge structure:

Theorem 2.18 (see [63, Theorem I.11] or [2, Theorem 1.8, Corollary 1.10])**.**

(1) A reduced path represents the trivial element of $F(\mathcal{G})$ if and only if it has length 0 and $g_0 = 1$.

(2) Two reduced paths representing the same element of $F(\mathcal{G})$ have the same edge structure.

Remark 2.19. It is a consequence of this theorem that vertex groups embed in $F(\mathcal{G})$, by the map taking g (as an element of G_v) to g, the path of length 0 at v. With only one edge, the first statement reduces to Britton's lemma (see [54, Theorem VI.2.1] and the surrounding section), and in fact by adding an edge at a time inductively this more general statement is a consequence of Britton's lemma. A proof of the second statement proceeds by induction, given the first: the base case is paths of length 0, then for the inductive step one observes that if p and q are reduced and represent the same element then pq^{-1} is not reduced, but the only possible location of s segment $e\alpha_e(s)\bar{e}$ is between the last edge of p and the first of q^{-1}.

Note that a cyclically reduced loop is only defined up to a choice of initial vertex, and might not include the initial (and terminal) vertex v_0 of the original loop.

Definition 2.20. The *fundamental group of* $\mathcal{G}$ (at a vertex v) is the set of loops at v in $F(\mathcal{G})$, and is denoted $\pi_1(\mathcal{G}, v)$. The multiplication is that of $F(\mathcal{G})$, restricted to these loops. (Equivalently, this is the multiplication from the fundamental groupoid, since in the groupoid the product of two loops at the same vertex is always defined.)

Note that loops are counted equal if they are equal under the relations for $F(\mathcal{G})$, so (for example) two equivalent reduced loops are "the same" for this definition. The isomorphism class of this group does not depend on the vertex chosen. (In fact, the two groups obtained by choosing different base vertices are conjugate in the groupoid.)

We take the corresponding definition of the Bass-Serre tree:

Definition 2.21 (Bass-Serre Tree)**.** Let T be the graph formed as follows: the vertex set consists of 'cosets' $G_w p$, where p is a path in $F(\mathcal{G})$ from w to v. There is an edge(-pair) joining two vertices $G_{w_1}p_1$ and $G_{w_2}p_2$ if $p_1 = eg_{w_2}p_2$ or $p_2 = eg_{w_1}p_1$ (where $g_w \in G_w$).

Anticipating the next proposition, the graph T is usually called the *Bass–Serre tree* (or universal cover) for $\mathcal{G}$.

Proposition 2.22. *The graph T defined in Definition 2.21 is a tree, and it has a natural right action of* $\pi_1(\mathcal{G}, v)$ *which is without edge inversions.*

Proof. First we see that T is a tree: by following the defining path p we can connect each vertex to the vertex G_v. If it contains a reduced cycle there is some $G_w p = G_w$ with p reduced. But as cosets of G_w in $F(\mathcal{G})$ this implies that $p \in G_w$, and in particular has length 0: so there are no non-trivial reduced cycles.

Since loops at v both start and finish at v, $\pi_1(\mathcal{G}, v)$ acts on the right on the set of vertices, preserving adjacency. To see it is without edge inversions, we begin by investigating vertices and showing that $g \in \pi_1(\mathcal{G}, v)$ fixes a vertex if and only if it cyclically reduces until it has length 0.

Suppose g stabilises some vertex $G_u p$. Then $G_u pg = G_u p$, which implies that g is an element of $p^{-1}G_u p$. (Conversely any such element fixes $G_u p$; in particular the stabiliser is isomorphic to G_u.) Note that this implies that any vertex stabiliser has cyclically reduced length 0. The converse is true too: suppose g has cyclically reduced length 0, then we can write $g = p^{-1}hp$ where h is an element of G_w, with w the initial vertex of p. In particular, g fixes the vertex $G_w p$.

Now note that if g has cyclically reduced length n, g^k (for $k \geq 1$) has cyclically reduced length nk: write $g = php^{-1}$ where php^{-1} is reduced as written (this condition ensures that p has the shortest edge structure possible) and h is cyclically reduced; then $g^2 = ph^2p^{-1}$ and h^2 must be reduced (in fact cyclically reduced). So g^2 has cyclically reduced length $2n$; proceeding inductively gives the result for general k.

If g inverts an edge e, g^2 fixes the vertices at each end of e, and so has cyclically reduced length 0; but this can only happen if g had cyclically reduced length 0, in which case it fixes a vertex. But g fixes the midpoint of the alleged inverted edge and not its endpoints; in particular it cannot fix the path between this midpoint and the stabilised vertex, so the fixed point set would not be connected. So no edge inversions are possible. □

In fact, the action is precisely characterised by the initial graph of groups.

Definition 2.23 (Quotient graph of groups)**.** Given a group G acting on a tree T, there is a *quotient graph of groups* formed by taking the quotient graph from the action and assigning edge and vertex groups as the stabilisers of a representative of each orbit from a fundamental domain. Edge monomorphisms are then the inclusions, after conjugating appropriately if incompatible representatives were chosen.

Theorem 2.24 (Structure theorem, see [2, Corollary 3.7],[63, Theorem I.13])**.** *Up to isomorphism of the structures concerned, the processes of constructing the quotient graph of groups, and of constructing the fundamental group and Bass-Serre tree are mutually inverse.*

Remark 2.25. The appropriate notion of equivalence at the graph of groups level is characterised by isomorphism of the underlying graphs, together with isomorphisms of corresponding edge and vertex groups (and those isomorphisms must be compatible with the edge monomorphisms).

The appropriate notion of isomorphism for groups acting on trees is an isomorphism of groups and an equivariant isometry of trees (that is, $f(x \cdot g) = f(x) \cdot (g\varphi)$).

Careful choice of fundamental domain makes these isomorphisms more obvious: for example, most times there is a choice we will be wise to choose the fundamental domain so it includes G_v.

It is worth noting that there is a more general notion of morphism of graphs of groups, corresponding to group homomorphisms and equivariant maps between trees the groups act on. The details are complicated in general; however they are spelled out in sufficient detail for the case with trivial edge groups in Paper 1.

Example 2.26.

(1) The simplest example is one we have already seen: a graph can be viewed as a graph of groups taking every edge and vertex group to be trivial. The fundamental group is the usual fundamental group of the graph, a free group, and the Bass–Serre tree is the universal cover.

(2) Free products (see Definition 1.3) provide another example: if every edge group is trivial, the fundamental group is isomorphic to the free product of the vertex groups and a free group which is the fundamental group of the underlying graph.

(3) A free product with amalgamation (sometimes an amalgam) corresponds to a graph of groups with two vertices and one edge, as shown in Figure 7a. Sometimes this group is written $G_1 *_H G_2$.

(4) For example, $\mathrm{SL}_2(\mathbb{Z})$ has a presentation $\langle x, y : x^4 = 1, y^6 = 1, x^2 = y^3 \rangle$ which corresponds to the amalgamated free product $C_4 *_{C_2} C_6$ and the graph of groups in Figure 7b.

(5) an HNN extension corresponds to a graph of groups with one vertex and one edge (see Figure 7c. The edge inclusions specify two embeddings φ and ψ of the subgroup C into A, and given a presentation $\langle X : R \rangle$ for A, the fundamental group may be written $\langle X, t : R, t^{-1}\varphi(h)t = \psi(h) \quad \forall h \in C \rangle$. The new generator t arises from the loop, and is sometimes called the *stable letter*. Sometimes an HNN extension is written $G*_C$.

(6) A specific example of a HNN extension is the fundamental group of a Klein bottle, which has presentation $\langle a, t : t^{-1}at = a^{-1} \rangle$, expressing it as an HNN extension of $\mathbb{Z}$ (generated by a) with the edge group another infinite cyclic group (generated by c, say): one embedding sends c to a, while the other sends c to a^{-1}, as in Figure 7d.

The subscript notation for amalgamated free products and HNN extensions is not very precise: to understand the groups $A *_C B$ or $A*_C$ we need to know the two embeddings of the subgroup C, not just its isomorphism class. The last example gives the fundamental group of a Klein bottle as $\mathbb{Z}*_{\mathbb{Z}}$, but the same would be true of any presentation $\langle a, t : t^{-1}a^m t = a^n \rangle$ with $m, n \in \mathbb{Z} \setminus 0$. This is a presentation for the *Baumslag–Solitar group* $\mathrm{BS}(m, n)$; two such groups are isomorphic if and only if $\{m_1, n_1\} = \{m_2, n_2\}$ or $\{m_1, n_1\} = \{-m_2, -n_2\}$ [57].

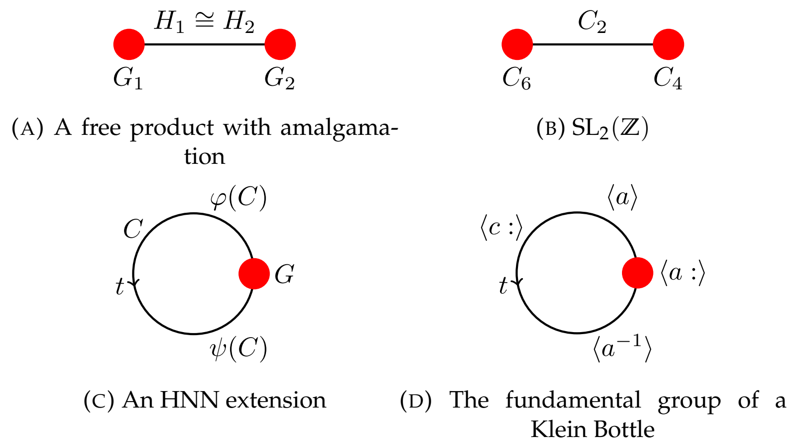

(A) A free product with amalgamation

(B) $SL_2(\mathbb{Z})$

(C) An HNN extension

(D) The fundamental group of a Klein Bottle

FIGURE 7: Example of graphs of groups

We can view the fundamental group of a graph of groups as iterating these two constructions (amalgamated free products and HNN extensions) over each edge of the graph:

Lemma 2.27. *Let e be any edge in a graph of groups $\mathcal{G}$. Then*

(1) *if e is separating, $\pi_1(\mathcal{G}) \cong \pi_1(\mathcal{G}_1) *_{G_e} \pi_1(\mathcal{G}_2)$, where $\mathcal{G}_1$ and $\mathcal{G}_2$ are the graphs of groups defined on the connected components of Γ after removing e.*

(2) *if e is non-separating, $\pi_1(\mathcal{G}) \cong \pi_1(\mathcal{G}_1)*_{G_e}$, where $\mathcal{G}_1$ is the graph of groups obtained by removing e.*

Proof. First we pick basepoints for $\mathcal{G}$, $\mathcal{G}_1$ and $\mathcal{G}_2$: let $v = \iota(e)$, and choose v as the basepoint for $\mathcal{G}$ and $\mathcal{G}_1$. If e is separating, let $w = \tau(e)$ be the basepoint for $\mathcal{G}_2$. If e is not separating, fix a (reduced) path p in $F(\mathcal{G}_1)$ joining v and $w = \tau(e)$.

Let $\mathcal{H}$ be the graph of groups with one edge (identified with e, and with edge group G_e) and in the separating case two vertices with vertex groups $\pi_1(\mathcal{G}_1, v)$ and $\pi_1(\mathcal{G}_2, w)$. In the non-separating case take one vertex with vertex group $\pi_1(\mathcal{G}_1, v)$. Take the vertex x corresponding to $\mathcal{G}_1$ to be the basepoint for $\mathcal{H}$.

Note that in both cases the embeddings $\alpha_{\bar{e}}$ of G_e into $\pi_1(\mathcal{G}_1, v)$ are simply the inclusions into G_v. Similarly, in the separating case, α_e is the original inclusion into G_w. (In all these situations we abuse notation a little by writing α for the inclusion into a vertex group of $\mathcal{H}$ as well as the original inclusion into a vertex group of $\mathcal{G}$.) However in the non-separating case, we take the inclusion $\alpha'_e : G_e \to \pi_1(\mathcal{G}_1, v)$ to be the map $s \mapsto p\alpha_e(s)p^{-1}$, reserving α_e for the original map $G_e \to G_w$.

Now we are in a position to define our map. The idea is to take a reduced path representing an element of $\pi_1(\mathcal{G}, v)$, split it up every time there is an occurrence of e or $\bar{e}$ and replace the intermediate sections with elements of $\pi_1(\mathcal{G}_1, v)$ or $\pi_1(\mathcal{G}_2, w)$. Recall that by Theorem 2.18 the edge structure of any reduced path is uniquely determined.

The remaining segments are paths in $\mathcal{G}_1$ or $\mathcal{G}_2$. In the separating case they are loops at v or w and therefore elements of $\pi_1(\mathcal{G}_1, v)$ or $\pi_1(\mathcal{G}_2, w)$ as written. In the non-separating case they may be loops at v or w, or paths from v to w or vice versa. Use the fixed path p to rewrite all the cases as loops at v: a loop g at w should be taken to pgp^{-1}, a path q from v to w as qp^{-1} and a path r from w to v as pr.

This is a homomorphism. The edge relations not involving e are still satisfied in the relevant subgraph; the edge relation over e is satisfied in the separating case since the inclusions are the original inclusions here. It is also satisfied in the non-separating case since $e\alpha_e(s)\bar{e} \mapsto ep\alpha_e(s)p^{-1}\bar{e} = e\alpha'_e(s)\bar{e} = \alpha_{\bar{e}}(s)$. (Note that the final $\alpha_{\bar{e}}$ includes s into $\pi_1(\mathcal{G}_1, v)$; this is the image under our map of the inclusion of s into G_v.

The kernel of this map is trivial: we use Theorem 2.18. If a reduced cycle is sent to a path representing the trivial element and did not cross e then it represents the trivial element of $\pi_1(\mathcal{G}_1, v)$, and so is also trivial in $\pi_1(\mathcal{G}, v)$. So suppose it crosses e, in which case there must be a reduction available. But since the edge group is just the edge group in $\mathcal{G}$, this implies the original cycle was not reduced.

It is also surjective; given an element of $\pi_1(\mathcal{H}, x)$ we can "expand" every vertex group element to a loop in $\pi_1(\mathcal{G}_1, v)$ or $\pi_1(\mathcal{G}_2, w)$. In the separating case this immediately gives an element of $\pi_1(\mathcal{G}, v)$; in the separating case we also need to insert p^{-1} after each use of e and p before each use of $\bar{e}$. In both cases the element of $\pi_1(\mathcal{G}, v)$ we have constructed maps to the given element of $\pi_1(\mathcal{H}, x)$ (the p and p^{-1} cancel with those in the homomorphism). □

The equivalent procedure in the Bass–Serre tree is "equivariantly collapsing a subforest" comprising the orbits of all edges other than the lifts of e.

Given a graph of groups $\mathcal{G}$, we can define another graph of groups $\mathcal{G}'$ by taking a subgraph Γ' of the underlying graph together with the edge and vertex groups and inclusions corresponding to Γ'. This is known as a *subgraph of groups*. A graph of groups is called *minimal* if no proper subgraph of groups has isomorphic fundamental group.

Lemma 2.28 ([2, Proposition 7.12]). *An action of a group G on a tree T is minimal if and only if the quotient graph of groups is minimal.*

Proof. If there is a proper invariant subtree T', then consider the quotient graph of groups with respect to just this subtree. By Theorem 2.24 the fundamental group is isomorphic to G. This corresponds to a (connected) subgraph of the quotient graph arising from T; edge and vertex stabilisers are the same so this is a subgraph of groups.

Conversely, suppose $\mathcal{G}'$ is a subgraph of groups with isomorphic fundamental group. Then (fixing a basepoint in $\mathcal{G}'$) every reduced cycle stays in $\mathcal{G}'$, since otherwise the fundamental groups would not be isomorphic. In particular, there is no reduced path to a vertex of $\mathcal{G}'$ that goes outside $\mathcal{G}'$, since composing with any reduced path inside $\mathcal{G}'$ would give such a reduced cycle. So the Bass–Serre tree corresponding to $\mathcal{G}'$ can be identified with a proper, $\pi_1(\mathcal{G})$-invariant subtree of the Bass–Serre tree corresponding to $\mathcal{G}$. □

Bass–Serre theory simplifies many proofs by allowing us to work with the geometry arising from the action on the tree. As an example, consider the Kurosh subgroup theorem on subgroups of free products.

Theorem 2.29 (Kurosh subgroup theorem [48])**.** *Suppose G is a free product $* G_i$ (over some index set I) and H is a subgroup of G. Then $H \cong (* H_j) * F$ where each H_j is isomorphic to an intersection $H \cap G_i^{k_i}$ of H with a conjugate of some G_i Further, the set $\{H_j\}$ is unique up to conjugation and reindexing, and the rank of F is uniquely determined.*

Proof. We may view G as acting on a tree T, with vertex stabilisers corresponding to the conjugates of the free factors G_i and trivial edge stabilisers. (We can take T as the Bass–Serre tree of any graph of groups with underlying graph a tree, trivial edge groups and vertex groups corresponding to the G_i.) Any subgroup H of G acts on the same tree: consider the quotient graph of groups $\mathcal{H}$ for this action. By the Structure Theorem 2.24, H is isomorphic to the fundamental group of this quotient graph of groups. Since the edge groups must still be trivial, this expresses H as a free product of the vertex groups of $\mathcal{H}$ and the fundamental group of the underlying graph of $\mathcal{H}$: the vertex groups are isomorphic to $H \cap G_i^{k_i}$ for some element k_i of G, and the fundamental group of the graph provides the free group. Choosing different representatives for each vertex orbit has the effect of conjugating the vertex groups, so the H_j are uniquely determined up to conjugation; since the graph is determined by the action the rank of the free group F is too. □

We also record here a result about decomposing groups as free products:

Theorem 2.30 (Grushko decomposition)**.** *Any finitely generated group G can be decomposed as a free product $G = G_1 * \ldots G_k * F_r$, where the G_i are non-trivial, freely indecomposable and not infinite cyclic, and F_r is a free group of rank r. Further, the G_i are unique up to conjugacy, and the rank of F_r is unique.*

This decomposition theorem is a well known consequence of Grushko's theorem on the rank of free products [35, 59] and the Kurosh subgroup theorem; see for example [64]. As Stallings notes, the existence of such a decomposition follows from Grushko's theorem , and uniqueness follows from the Kurosh subgroup theorem.

Serre's Property (FA)

Since actions on trees provide are useful for understanding a group, it is also useful to know when to stop looking for one. Serre [63, Section I.6] defines and investigates the property of "not having such an action":

Definition 2.31. A group G has *Property (FA)* if whenever it acts on a tree it does so with a global fixed point.

Serre provided an algebraic characterisation of this property.

Theorem 2.32 ([63, Theorem 15])**.** *A countable group G has Property (FA) if and only if*

(1) it is not a non-trivial free product with amalgamation,

(2) it has no quotient isomorphic to $\mathbb{Z}$, and

(3) it is not finitely generated.

A trivial free product with amalgamation would be of the form $A *_A B$ where the edge group is equal to one of the vertex groups. This proof is essentially Serre's, adapted slightly for the way we have set up graphs of groups above.

Proof. First we show that if any of the conditions hold then G admits a non-trivial action on a tree.

If G is a free product with amalgamation, then we can represent G as the fundamental group of a graph of groups with two vertices and one edge, and G acts on the Bass–Serre tree for this graph of groups. Assuming both inclusions of the edge group are proper, there are elements which do not stabilise a point in this tree. The infinite cyclic group $\mathbb{Z}$ acts by translation on the tree consisting of a single line, and so any group admitting a quotient to $\mathbb{Z}$ does too.

Now suppose G is not finitely generated. In this case G is the union of a strictly increasing sequence of subgroups $G_1 < G_2 < G_3 < \ldots$: one way to construct such a sequence is to enumerate an infinite generating set $\{g_1, g_2, \ldots\}$, and add one generator at a time. Perhaps the generating set was not minimal, so adding some generator g_i will not lead to a new subgroup. In this case we should discard that step; but the process cannot terminate, since if it did there would be a finite subset of generators that generated G.

We can now build a graph where the vertices correspond to cosets of subgroups in this sequence, and edges correspond to inclusions "up a single level". So for example, with the vertices corresponding to $G_1 < G_2 < G_3$ there are edges corresponding to the inclusions of G_1 into G_2 and G_2 into G_3, but not to the inclusion of G_1 into G_3. The group G acts on this graph by permuting the cosets of the subgroups; we claim that in fact the graph is a tree.

To see this, suppose there is a reduced cycle, and consider the vertices in this cycle. Consider a "lowest" vertex in the cycle: that is a vertex corresponding to a coset $G_n h$, with all other vertices in the cycle corresponding to cosets of G_i with i at least n. Since this vertex is lowest, the cycle has no edges going down; there are no edges between vertices on the same level, and so both edges at this vertex must be going up. But there is only a single upwards edge, representing the inclusion of $G_n h$ into $G_{n+1} h$: so this is a backtrack, and in fact the cycle was not reduced.

Now suppose that a group satisfies all three conditions and acts on a tree. The quotient graph Γ must be a tree, since if it contains any non-trivial loops there is a surjection to a free group (by sending vertex group elements to the identity and preserving elements corresponding to the edges) and then onwards to $\mathbb{Z}$. Also, the finite generation assumption allows us to pass to a subtree with a finite quotient graph: consider taking the subgroups corresponding to fundamental groups of increasing finite subgraphs of Γ, and observe that after finitely many steps this will contain a generating set of G, and therefore the fundamental group will already be all of G. Since this is realised by a finite subgraph, we can consider one that is minimal. If such a minimal graph is a point, the original tree had an invariant point, which must be a global fixed point.

Otherwise, consider any edge e. Since the quotient graph is a tree, all edges are separating, so by Lemma 2.27 we have that $G \cong \pi_1(\Gamma_1) *_{G_e} \pi_1(\Gamma_2)$, where Γ_1 and Γ_2 are the two connected components of $\Gamma \setminus \{e\}$. But we assumed that G was not an amalgamated free product, so in fact we must have that $G_e = \pi_1(\Gamma_i)$. But then there are no reduced loops crossing e, and so Γ was not minimal. So in fact the minimal tree is just a point, which is to say that G has a global fixed point for the action. □

The condition used in the proof was not so much that G was not finitely generated, but that this implies G is the union of a properly increasing sequence of subgroups. We can replace the third condition with "is not the union of a properly increasing sequence of subgroups" even among uncountable groups, and there are uncountable groups which satisfy this and the first two conditions, and so do have Property (FA), despite being very far from finitely generated [46].

Example 2.33. Finite groups all have Property (FA): since they are finitely generated it is enough to show that all elements are elliptic. This is the case since if g is hyperbolic, then so is g^n for all $n \geq 1$ (it translates along the same axis, but the translation length is multiplied by n): but in a finite group for every element there is some n so that g^n is trivial.

The same argument goes through for finitely generated torsion groups (whose existence was proved by Golod–Shafarevich [34], but as soon as they are infinitely generated of course they cannot have Property (FA).

Property (FA) is not inherited by finite index subgroups, but if a group has the stronger Property (T) (see [4]) then all its finite index subgroups have Property (FA) [66]. So the existence of a finite index subgroup with a non-trivial action on a tree serves as an obstruction to Property (T). Paper 2 investigates when the automorphism group of a free product of groups has Property (FA).

Translation Length Functions and $\mathbb{R}$-trees

Given an action of a group G on a tree (or indeed any metric space), we may define a translation length function for the action

$$\|g\|_T = \inf\{d(x, xg) : x \in T\}.$$

If the action is assumed to be without inversions, the infimum can be taken to range over vertices only. (If inversions were a possibility it must include points on the edges too.)

This infimum is attained, since if g does not fix a point, it is hyperbolic by Lemma 2.8. It is a consequence of those arguments that for all $x \in T$, $d(x, xg) = 2d(x, \text{axis}(g)) + \ell$, where ℓ is the distance g translates along its axis, so the minimum is realised by any point on the hyperbolic axis.

Translation lengths are constant on conjugacy classes: if g fixes x then $h^{-1}gh$ fixes xh, and if g is hyperbolic then so is $h^{-1}gh$, with $\text{axis}(g)h = \text{axis}(h^{-1}gh)$. Since h induces an isometry of T, the translation lengths must be the same. (See Figure 5.)

In [22] Culler and Morgan investigate actions on $\mathbb{R}$-trees, where branching can occur at any distance; they are characterised as metric spaces where there is a unique arc (a geodesic) joining any two elements.

We can view simplicial trees as $\mathbb{R}$-trees, with the usual edge path metric - all vertices are at integer distances. We can also vary the edge lengths to obtain more examples – these are known as *metric simplicial trees*. It is important to note that not all $\mathbb{R}$-trees are metric simplicial trees, and in particular there are actions on $\mathbb{R}$-trees that do not correspond to actions on simplicial trees: for example most surface groups act freely on $\mathbb{R}$-trees [58], but not on simplicial trees. There are even groups with actions on $\mathbb{R}$-trees but which have Property (FA) – so there are no non-trivial actions on simplicial trees [56].

A *branch point* in an $\mathbb{R}$-tree is a point with no neighbourhood homeomorphic to an interval; in a metric simplicial tree the branch points correspond to the vertices with valence at least 3. An $\mathbb{R}$-tree is a metric simplicial tree if and only if the set of branch points is discrete.

We can define a translation length function for an action on an $\mathbb{R}$-tree too: it can assign g any value in $\mathbb{R}_{\geq 0}$ (for an action on a simplicial tree with the edge path metric and no inversions, it must take values in $\mathbb{Z}$). By the same arguments (the proofs used in the lemmas apply equally to $\mathbb{R}$-trees) the infima in the definition are realised, and the translation length function is invariant under conjugation. Culler and Morgan proved that many actions on $\mathbb{R}$-trees are uniquely characterised by their translation length function, in the following sense:

Theorem 2.34 (Culler–Morgan, [22, Theorem 3.7])**.** *Suppose G acts minimally and with no fixed ends on two $\mathbb{R}$-trees T_1 and T_2, with the same translation length function. Then there is a unique G-equivariant isometry $T_1 \to T_2$.*

The condition that there are no fixed ends is formulated by Culler and Morgan as "semi-simple and not a shift" and by some other authors as "non-abelian". This is partly because the actions we exclude are those with translation length functions which are homomorphisms to $\mathbb{R}$, and therefore factor through the abelianisation of G (see [22, Corollary 2.3]).

Whatever the terminology, the actions being excluded either translate along a single line, or have several hyperbolic axes but they share a point at infinity. Perhaps the simplest example of the second case is that of an ascending HNN-extension, such as $\mathrm{BS}(1,n) = \langle a,t : t^{-1}at = a^n \rangle$. (Note that a point by definition has no ends, so in particular cannot have fixed ends. Also recall that an action where every element is elliptic but there is no global fixed point cannot be minimal.)

The consequences of this theorem are explored and used in Papers 2 and 3, and it plays a key role in the study of deformation spaces, defined in Section 5.

Culler and Morgan proposed, and Parry [60] proved that the functions $G \to \mathbb{R}$ which are translation length functions of an action on an $\mathbb{R}$ tree can be characterised as those which satisfy certain axioms:

Theorem 2.35 (Parry, [60])**.** *Let G be a group, and f a function $G \to \mathbb{R}_{\geq 0}$. Then f is the translation length function for the action on G on an $\mathbb{R}$-tree if and only if, for every $g, h \in G$,*

(1) $f(g^{-1}hg) = f(h)$;

(2) either $f(gh) = f(gh^{-1})$, *or* $\max\{f(gh), f(gh^{-1}\} \leq f(g) + f(h)$;

(3) if $f(g), f(h) > 0$, *then either* $f(gh) = f(gh^{-1}) > f(g) + f(h)$ *or* $\max\{f(gh), f(gh^{-1}\} = f(g) + f(h)$

In fact, Parry's result holds for a more general notion of Λ-trees, which have a kind of geometry where the vertices may be at "distances" taking values in any ordered abelian group Λ [1, 15]. In particular, we can use this to say that the length function of an action on a simplicial tree (with the standard edge path metric) requires a function taking values in $\mathbb{Z}$, and satisfying an additional condition (which rules out inversions; this is not a concern in $\mathbb{R}$-trees since "vertices" are distinguished points which can occur at any distance).

There is another notion, due to Lyndon [53], of a *based length function*: where the length assigned to g is the distance it moves some base point v (which is chosen once and for all). Similar theorems hold for these based length funtions [14]; in fact the proofs of results about translation length functions often involve constructing base points and then using the results for based length functions.

For an action on a simplicial tree, by Theorem 2.24 there is a graph of groups representing the action. The translation length can be determined from this graph of groups. (The same is true for metric simplicial trees, with the modification that the edges of the graph of groups must be assigned the lengths of the edges in the orbit they represent.)

Proposition 2.36. *Suppose T is a metric simplicial G-tree, and let $\mathcal{G}$ be its quotient graph of groups. Then for every element g of G, $\|g\|_T$ is the length after cyclic reduction of any loop in $\mathcal{G}$ representing g*

Proof. We work with the graph of groups and Bass–Serre tree; by Theorem 2.24 this is equivariantly isomorphic to the original tree. Fixing a base point v and viewing g as an element of $\pi_1(\mathcal{G}, v)$, write $g = php^{-1}$ reduced as written and with h cyclically reduced. Let $\tau(p) = w$. If h has length 0, then g stabilises $G_w p^{-1}$ in the Bass–Serre tree and is elliptic. So suppose h does not have length 0. Let tes be the terminal segment of h, so $h = h'tes$, reduced as written. Consider the adjacent vertices $G_w p^{-1}$ and $G_u esp^{-1}$ and their images under g: $G_w p^{-1}(php^{-1}) = G_w hp^{-1}$ and $G_u esp^{-1}(php^{-1} = G_u eshp^{-1})$ (the situation is shown in Figure 8); since everything was reduced as written all four vertices lie on the segment $[G_v, G_u eshp^{-1}]$. So g translates $[G_w p^{-1}, G_u esp^{-1}]$, and so Lemma 2.12 gives that this is contained in axis(g). This translation is exactly by the length of h, so this is the translation length, as required. □

3 Residual Finiteness and Subgroup Separability

Given a group, we can ask "how easy" it is to tell its elements apart. One way to try to do this is to look at finite quotients of the group:

Definition 3.1. A group G is *residually finite* if for every element $g \neq 1$ of G, there is a homomorphism $G \to K$ where K is a finite group and g lies outside the kernel.

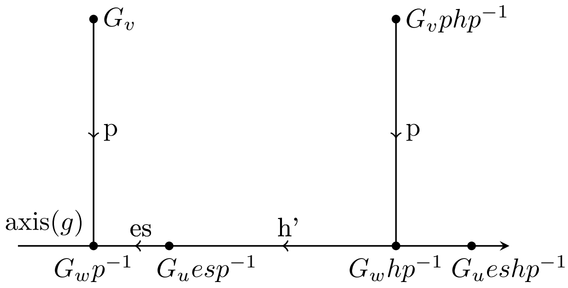

FIGURE 8: Calculating translation length for $g = php^{-1}$

The idea is that elements can be distinguished (from the identity, and so by considering gh^{-1} from each other) by looking at some finite quotient of the group. There are several equivalent definitions of residual finiteness, one of which first needs the definition of a topology on any group G (see [42, 61]):

Definition 3.2. For a group G, the *profinite topology* on G is the topology with basis given by all cosets of finite index subgroups of G.

The cosets forming this basis are closed as well as open: they are the complement of the (finite) union of the remaining cosets of the same subgroup. Also notice that this is a true basis and not just a subbasis: the intersection of two finite index subgroups is again finite index and the intersection of cosets, if non-empty, is a coset of the intersection of the underlying groups.

We now give several equivalent characterisations of residual finiteness:

Lemma 3.3. *For a group G, the following are equivalent:*

(1) G is residually finite;

(2) for every non-identity element g of G, there is a finite index normal subgroup N of G that does not contain g;

(3) for every non-identity element g of G, there is a finite index subgroup H of G that does not contain g;

(4) the trivial subgroup of G is closed in the profinite topology.

The last condition implies that the profinite topology on G is Hausdorff: since the topology respects the multiplication, either all or no singletons are closed.

Proof. Assertions (1) and (2) are equivalent since the kernel of a map to a finite group is a normal finite index subgroup.

Assertion (2) implies assertion (3); to see the converse suppose H is a finite index subgroup, and consider the action of G on the cosets gH. This acts to permute the cosets, so it defines a map to S_n, where n is the index of H. We let N be the kernel of this action; it is a finite index normal subgroup, and it is contained in H (which stabilises itself in this action) so does not contain g.

To see that (4) is equivalent to these conditions, consider the collection of all finite index subgroups of G. They are closed in the profinite topology, and therefore so is their intersection. If G is residually finite, every non-identity element G lies outside some finite index subgroup, and therefore the intersection (the *finite residual*) is the trivial group.

Conversely, if the trivial subgroup is closed, its complement is covered by a subset of the basis. That is, for every non-identity element g there is a coset gH of a finite index subgroup that contains g and not 1. This implies that $gH \neq H$: that is, g lies outside H. □

Given a subroup H, the right coset Hg is the same set as the left coset $gH^g = g(g^{-1}Hg)$, so the basis given for the profinite topology does not need this to be specified. Also, by passing to the *normal cores* (the largest normal subgroup N contained in a subgroup H; in fact this agrees with the subgroup N constructed in the proof that (2) implies (3) above) taking the basis to be cosets of finite index normal subgroups defines the same topology.

The last equivalent condition of Lemma 3.3 suggests a generalisation: given a group, we can ask which subgroups are closed in the profinite topology. This leads to the following definitions.

Definition 3.4. A subgroup H of a group G is called *separable* if it is closed in the profinite topology. The group G is *subgroup separable* if all finitely generated subgroups are separable.

So a group is residually finite if the trivial subgroup is separable. As with residual finiteness, these definitions have an interpretation involving finite index subgroups.

Lemma 3.5. *A subgroup H of a group G is separable if and only if for every element g there is a finite index subgroup K containing H but not g.*

Proof. As for residual finiteness, we may consider the intersection of all finite index subgroups containing H. If every g outside H is also outside a finite index subgroup containing H, this intersection is exactly H which is therefore closed.

Now suppose H is closed in the profinite topology: so there is some finite index subgroup K such that gK and H are disjoint. By passing to the normal core, we can assume without losing generality that K is normal. Now HK is a subgroup (since K is normal) and finite index (since it contains K, which is finite index). Lastly, g cannot be an element of HK: if so, we have $g = hk$ with h in H and k in K. But then $h = gk^{-1} \in gK$, and H is not disjoint from gK. □

A subgroup separable group is sometimes said to be LERF, for "Locally Extended Residually Finite".

One notable class of groups which are subgroup separable are free groups, a result originally due to Marshall Hall.

Theorem 3.6 (Marshall Hall Jr. [41])**.** *Let F_n be the free group on n generators, H be a finitely generated subgroup of F_n and $g \in F_n$ be outside of H. Then there is a finite index subgroup K of F_n, which contains H as a free factor and does not contain g. In particular, finitely generated free groups are subgroup separable.*

The proof given here is an adaptation of the one due to Stallings [65] .

Proof. We view F_n as the fundamental group of a rose R_n – a graph with one vertex and n edges. The universal cover is a $2n$ regular tree T, on which F_n acts freely. Fix a basis of F_n corresponding to the edge set of R_n; then the axes of the basis elements all intersect, at a "preferred lift" v of the vertex of R_n.

The subgroup H can be represented by a locally injective map f (an *immersion*) of a finite graph Γ into R. Stallings constructs such a graph by taking a rose representing a generating set for H and *folding* until it is locally injective; one can also consider the quotient graph of the action of H on its minimal invariant subtree of T, extending equivariantly to include the preferred lift of the base point. By labelling each edge of Γ with its image in R, and choosing the base point corresponding to the orbit of the preferred lift, we can read loops in Γ as elements of F_n. (Choosing any other base point, this reading would give a conjugate of the intended subgroup.)

We can further extend Γ to "read" the element g, by adding a line (at the base point) whose edges read g, and then folding so there is once again an immersion from Γ to R. Since g was given as outside of H, it must not be a loop, even after this process. (From the perspective of the action on the tree, this corresponds to equivariantly extending the invariant tree to include the path from v to vg.)

Our aim now is to extend Γ and the immersion to construct a finite cover of R corresponding to a subgroup K as described in the statement. The preimages of each edge of R give a partial bijection between the vertices of Γ: since f is an immersion, there cannot be two edges with the same image starting or ending at the same vertex.

Adding the edges corresponding to some extension of this partial bijection to a full bijection, and repeating for each edge of R, we will construct extend Γ and f to a cover of

R. The subgroup K corresponding to this cover contains Γ as a subgraph, and therefore contains H but does not contain the element g, since the path representing g still cannot be a loop.

The vertices of this covering graph correspond to the cosets of K (since two vertices vg_1 and vg_2 are in the same K-orbit if and only if $g_1g_2^{-1}$ is in K). Since it is finite, K is finite index. □

Burns [13] and Romanovskii [62] proved that free products are subgroup separable if the factor groups are; Paper 1 consists of a new proof of this theorem in the language of Bass–Serre theory.

4 Automorphisms of Groups

Definition 4.1. Let G be a group. Its automorphism group, $\mathrm{Aut}(G)$ is the set of self-isomorphisms under composition.

We identify one notable kind of automorphism.

Definition 4.2. Conjugation by $g \in G$ induces an *inner automorphism* of G. The *inner automorphism group* denoted $\mathrm{Inn}(G)$ is the collection of such automorphisms.

These automorphisms are "easier" to understand – at least no harder than the group itself:

Lemma 4.3. *The group* $\mathrm{Inn}(G)$ *is a normal subgroup of* $\mathrm{Aut}(G)$*, and is isomorphic to* $G/Z(G)$*.*

Proof. Suppose γ is an inner automorphism, induced by conjugating by $g \in G$, and φ is any element of $\mathrm{Aut}(G)$. Then applied to some $h \in G$, we see that

$$\begin{aligned} h(\varphi^{-1}\gamma\varphi) &= (g^{-1}(h\varphi^{-1})g)\varphi \\ &= (g^{-1}\varphi)h(g\varphi) \end{aligned}$$

That is, we have that the image of γ under this conjugation is the inner automorphism induced by conjugating by $g\varphi$.

The image of the map from G to $\mathrm{Aut}(G)$ taking elements to the inner automorphism they induce is exactly $\mathrm{Inn}(G)$, so we need to check the kernel. If g is in the kernel then $g^{-1}hg = h$ for all elements h of G, so this is exactly the centre $Z(G)$. □

We often will want to work "up to an inner automorphism" and consider the quotient by $\mathrm{Inn}(G)$.

Definition 4.4. The quotient $\mathrm{Aut}(G)/\mathrm{Inn}(G)$ is known as the *outer automorphism group*, or $\mathrm{Out}(G)$.

There is also a topological justification for studying the outer automorphisms: a base point preserving self homotopy equivalence induces an automorphism of the fundamental group of a space. If we do not insist that the homotopy equivalence preserves the base point then the resulting automorphism is only defined up to picking a path from the base point to its image. In particular, when two such paths are not homotopic (relative to the endpoints) then they form a loop, and the automorphisms will differ by conjugation by the element of the fundamental group that loop represents.

Example 4.5. There are some examples where $\mathrm{Inn}(G)$ trivial, so $\mathrm{Aut}(G) \cong \mathrm{Out}(G)$:

(1) If G is a finite cyclic group, $\mathrm{Aut}(G)$ is isomorphic to its group of units (when G is viewed as a ring): we have to send a generator to another element with the correct order.

(2) If G is infinite cyclic (that is, $\mathbb{Z}$) then we have $\mathrm{Aut}(\mathbb{Z}) = \mathbb{Z}/2\mathbb{Z}$: we must either fix or invert the generator.

(3) If G is finitely generated free abelian, we have $\mathrm{Aut}(\mathbb{Z}^n) = \mathrm{GL}_n(\mathbb{Z})$: invertible integer valued matrices with integer valued inverses (equivalently, integer valued matrices with determinant ± 1).

These groups are all abelian, and in particular equal to their centres: there are no non-trivial inner automorphisms, and so we have $\mathrm{Out}(G) \cong \mathrm{Aut}(G)$. In contrast,

(4) With $n \geq 7$ the symmetric group S_n has $\mathrm{Aut}(S_n) \cong \mathrm{Inn}(S_n) \cong S_n$: every automorphism is inner, and there is no centre.

Of course there are many examples where both the inner and outer automorphism groups are non-trivial:

(5) With $n \geq 7$ the alternating group A_n has $\mathrm{Aut}(A_n) \cong S_n$ (A_n is a subgroup of S_n that is preserved by all automorphisms, and in fact the induced map to $\mathrm{Aut}(A_n)$ is surjective.) There is no centre, so $\mathrm{Out}(A_n) \cong S_n/A_n \cong \mathbb{Z}/2\mathbb{Z}$: a representative of the non-trivial outer automorphism can be found by taking the conjugation by any odd element of S_n.

Automorphisms of free groups and free products

The automorphism groups of free groups and free products have been studied in some form for many years. There are presentations (of both cases combined: a free group is a free product of copies of $\mathbb{Z}$) due to Fouxe-Rabinovitch in [28] and [29] and later by Gilbert in [32]. Gilbert's presentation can be transformed into a finite presentation provided that the factor groups and their automorphism groups are finitely presented [32, Theorem 3.3].

The following generating set is taken from Gilbert's paper, under the assumption that $G = G_1 * \cdots * G_n * F_r$ is given as a Grushko decomposition: so each G_i is freely indecomposable and not isomorphic to $\mathbb{Z}$.

There are three kinds of automorphism, which generate the whole automorphism group. Fixing a basis X for F_r (which gives a decomposition of F_r as a free product of infinite cyclic groups), they are:

- Factor automorphisms, which are automorphisms of just one free factor and do not affect the rest (including replacing a single element of X with its inverse);
- Permutation automorphisms, which permute isomorphic free factors according to a fixed, compatible set of isomorphisms (including permuting the elements of X);
- Whitehead automorphisms, of two kinds:
  - partial conjugations sending a free factor G_i to G_i^a or an element x of X to x^a
  - for an element x of X, transvections sending x to ax.

  In the case of a partial conjugation of G_i, a must be drawn from some G_j, $j \neq i$ or from X. In the case of a partial conjugation or transvection of an element of X, a must be drawn from some G_i or from $X \setminus x$.

In the case where $r = 0$ – so this is a free product of groups which are not $\mathbb{Z}$ – there are only Whitehead automorphisms of the first kind. If G is given as a free product but not in a Grushko decomposition (for example, some G_i is freely decomposable) then all the generators above still define automorphisms of G, but together they no longer generate $\mathrm{Aut}(G)$.

The automorphism group or outer automorphism group of some group G does not inherit the properties of G. For example, free groups act freely on trees, but $\mathrm{Aut}(F_n)$ (and therefore $\mathrm{Out}(F_n)$ too) has Property (FA) for n at least 3 [10, 24].

The properties of the automorphism group $\mathrm{Aut}(G)$ do not usually depend in obvious ways on a presentation of G: for example the Baumslag–Solitar groups $\mathrm{BS}(2,3) = \langle a, t : t^{-1}a^2t = a^3 \rangle$ and $\mathrm{BS}(2,4) = \langle a, t : t^{-1}a^2t = a^4 \rangle$ differ only by a power of a in their presentations, but $\mathrm{Out}(\mathrm{BS}(2,3))$ is finite (in fact it has order 2) [19, 33], while $\mathrm{Out}(\mathrm{BS}(2,4))$ is not finitely generated [20].

Paper 2 studies the automorphisms (and outer automorphisms) of free products (without infinite cyclic factors); Paper 3 seeks to describe the outer automorphism groups of free-by-cyclic groups, which requires information about $\mathrm{Out}(F_n)$.

Characteristic Subgroups

A normal subgroup is preserved by conjugation, and therefore it is preserved by all inner automorphisms. We can also consider subgroups preserved by all automorphisms.

Definition 4.6. A subgroup H of a group G is *characteristic* if $H\varphi = H$ for all automorphisms φ of G.

Characteristic subgroups have different inheritance properties to normal subgroups: if $H \leq K \leq G$ and both inclusions are characteristic then in fact H is characteristic in G: an automorphism of G preserves and therefore induces an automorphism of H, which in turn preserves K. In fact if H is characteristic in K which is only normal in G, the same argument shows that H is normal in G.

However, it is not true that a characteristic subgroup is necessarily characteristic (though it will be normal) in all intermediate subgroups:

Example 4.7. Let K be $\langle a, t : t^{-1}at = a^{-1}\rangle$. This has an index two subgroup isomorphic to $\mathbb{Z}^2$, $\langle a, t^2\rangle$. The centre of K is $\langle t^2\rangle$, and it is characteristic in K but not in this copy of $\mathbb{Z}^2$.

Characteristic subgroups allow us to pass automorphisms to the quotient group:

Proposition 4.8. *If N is a characteristic subgroup of G, then there is a homomorphism* $\mathrm{Aut}(G) \to \mathrm{Aut}(G/N)$ *given by* $\varphi \mapsto (Ng \mapsto N(g\varphi))$.

We need N to be characteristic to ensure this definition is well defined; then composition and inverses follow.

In the same way as the normal core of a finite index subgroup is another finite index subgroup, we can find finite index characteristic subgroups given some finite index subgroup.

Proposition 4.9. *Suppose G is finitely generated and H is a finite index subgroup of G. Then*

(i) G has only finitely many subgroups of index $n = [G : H]$;

(ii) there is a finite index subgroup of H that is characteristic in G.

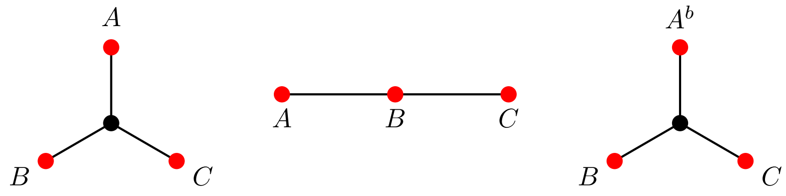

(A) (B) (C)

FIGURE 9: Three graphs of groups for $G = A * B * C$.

Proof. As in the construction of the normal core, we consider the action of G on the cosets G/H, where H is its own stabiliser. This action gives a map from G to S_n, with $n = [G : H]$, and H is mapped to the subgroup stabilising $\{1\}$. These maps correspond to a choice of elements of S_n for each generator and if G is finitely generated there are only finitely many such maps. However, as we vary H across subgroups of index n, we get different maps, since different subgroups are mapped to the stabiliser of $\{1\}$. So there can be only finitely many subgroups of a given finite index n.

To find a characteristic subgroup, observe that the index of a subgroup is preserved by an automorphism. So the finite set of index n subgroups is also preserved. Take the intersection of these subgroups: it is finite index (as the intersection of finitely many subgroups) and characteristic, since being an element in every index n subgroup is invariant under automorphisms. □

Note that automorphisms do not in general fix the elements of a characteristic subgroup in the same way that conjugations do not in general fix the elements of a normal subgroup.

Example 4.10. Characteristic subgroups include the commutator subgroup, the centre of a group, the subgroup generated by all finite order elements, and (in a free product) the normal subgroup generated by all free factors of the same isomorphism class (provided the group is written as a Grushko decomposition, and that isomorphism class is not $\mathbb{Z}$).

5 Deformation Spaces and Canonical Actions

There may be many graphs of groups suggested by the same presentation of a group G. For example, if $G = A * B * C$, two candidates are a "star" and a "line", as shown in Figure 9a-b.

These graphs of groups give very different actions: for example the products ab and ac (with $a \in A, b \in B$ and both non-trivial) are hyperbolic in all cases. Taking all edge lengths to be 1, in Figure 9a the translation length $\|ab\|_{T_1}$ is 4, whereas in Figure 9b we have $\|ab\|_{T_2} = 2$. However, $\|ac\|_{T_1} = \|ac\|_{T_2} = 4$, showing that the difference is not due to a scaling.

If we allow the edge lengths to vary, we can deform the star into the line, by collapsing the edge with B until it has length zero. We could also get lines where a different vertex group was in the middle, by collapsing a different edge.

Another variation is to swap one of the vertices for a conjugate, as in Figure 9c this has the same fundamental group, but (letting $b \in B$ be the conjugating element) the element $cb^{-1}ab$ will have translation length 4 in T_3 but 8 in T_1, while bc has translation length 4 in both cases.

None of these graphs of groups is obviously "better" than the others: some are more symmetric, but they have vertices that seem perhaps unnecessary since they have trivial stabiliser. Similarly, the choice of A over A^b was arbitrary: but when we swapped them the actions changed. In some situations it is better to consider not one action, but the whole family that can be obtained in these (and similar) ways. This leads us to deformation spaces, introduced by Forester [27]; the definitions here follow Guirardel and Levitt [36].

Definition 5.1. A deformation space of G-trees consists of all actions of G on metric simplicial trees having the same set of elliptic subgroups. That is, T_1 and T_2 are G-trees in the same deformation space, then any $H \leq G$ fixes a point of T_1 if and only if it fixes a point of T_2.

Any action of G on a tree T defines a deformation space, by considering all simplicial actions of G on a tree with the same elliptic subgroups. Note that there can still be vertices with stabilisers that are not conjugate to a stabiliser in the original action. (In the example above for $G = A * B * C$, some trees have an orbit of trivially stabilised vertices, and others do not, but they are all elements of the same deformation space.)

Definition 5.2. Given G-trees T_1 and T_2, say that T_1 *dominates* T_2 if there is a G-equivariant map $T_1 \to T_2$.

Trees are in the same deformation space if and only if they *dominate* each other; the maps are built by matching up vertices according to their stabilisers, and then mapping the paths between them appropriately. Note that the maps allowed in this definition are far more general than the graph maps defined earlier: they send vertices to vertices, but may collapse edges to vertices or send them to edge paths.

There are three "kinds" of equivalence relation on G-trees we may want to consider; each has benefits and drawbacks depending on the context. We will assume all actions are minimal, and usually that they are irreducible. (Equivalently that the minimal action is on a tree with at least 3 ends, none of which are fixed; many of the results also

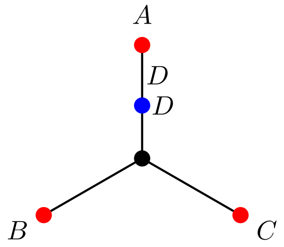

FIGURE 10: A graph of groups for $G = A * B * C$ with an extra edge group.

hold for the "dihedral actions" on a line without preserving its orientation, but do not for actions with fixed ends.)

First, we will always want to use the translation length functions to distinguish actions. So, following Theorem 2.34 we will always consider actions up to G-equivariant isometry.

We can also "projectivise", and consider actions that are distinct up to equivariant *homothety* (or scaling: a map such that for all points x, y we have $d(f(x), f(y)) = Cd(x, y)$ for some global constant $C > 0$): this corresponds to a scaling of the length function. (Note that in the trivial action every element is elliptic, so this translation length is not a scaling of the length function of any action with hyperbolic elements.) We can choose a representative of all classes simultaneously by fixing the *covolume*, the sum of all edge lengths in the quotient graph.

Finally we may want to ignore variations in edge lengths, counting G-trees as equivalent if they are equivariantly homeomorphic. This reflects the underlying simplicial structure, taking the vertices to be the set of branch points together with any "inversion points", any $x \in [v_1, v_2]$ stabilised by an element swapping the segments $[x, v_1]$ and $[x, v_2]$ (this means the combinatorial object will have no "removable" vertices with only two incident edges, and the same stabiliser as those edges): in fact the homeomorphism can be chosen to be linear on edges.

Rather than considering all trees in a deformation space, we can pass to a "restricted" deformation spaces, where we only allow trees with edge groups in some family of subgroups $\mathcal{F}$, which should be closed under conjugation and taking subgroups. This provides a way to remove examples such as Figure 10, which is another graph of groups for $G = A * B * C$, but with a (not particularly necessary or interesting) edge that is not trivially stabilised. In general [36, Definition 4.9] there is a way to pass from a deformation space to a subset which is a restricted deformation space, where the family restricted to comprises subgroups fixing an edge in every point of the deformation space.

Example 5.3. (1) The Culler–Vogtmann Outer Space CV_n [23] for a free group consists of minimal free actions of F_n on metric simplicial trees: this is the deformation space where the only elliptic subgroup is the trivial group. Usually Culler–Vogtmann space is taken to be projectivised, in practice often by restricting to actions with covolume 1. The *spine* of CV_n is constructed by putting a partial order on the homeomorphism classes (defined by collapse maps), and realising this partial order as a simplicial complex.

(2) The examples considered at the start of the section for $G = A * B * C$ are points of the deformation space with elliptic subgroups conjugates of A, B, and C. Taking the restricted deformation space with trivial edge groups (which includes all the points shown in Figure 9, but not that in Figure 10) gives the space studied by McCullough and Miller in [55]; and when the factors are freely indecomposable and not $\mathbb{Z}$, this is sometimes called the *Grushko decomposition space* (since it realises the Grushko decomposition of a free product), and studied in [37].

(3) JSJ decompositions over a family of subgroups [40] are actions on trees where edge groups are in the given family and elliptic in every such splitting, and which dominate any other tree with this property. They are generally not unique, and so the correct object to study is a deformation space. JSJ decompositions over the trivial group recover the Grushko decomposition space; another example is the JSJ deformation space of a hyperbolic group over virtually cyclic groups, or a relatively hyperbolic group over elementary (parabolic and virtually) cyclic groups. (These have particularly nice properties when they are one-ended [11, 39].)

Actions on Deformation Spaces

Given any G-tree T, we can produce another by *twisting* the action by an automorphism: that is, for every automorphism φ of G there is a new action $*_\varphi$ defined as $x *_\varphi g = x \cdot (g\varphi)$. The length function of the new action is given by $\|g\varphi\|$. This gives an action of $\mathrm{Aut}(G)$ on the set of G-trees; this action respects the three equivalences (equivariant isometry, homothety or homeomorphism) given above.

If the elliptic subgroups defining a deformation space are $\mathrm{Aut}(G)$-invariant – as happens for example in outer space CV_n (since only the trivial group is elliptic) or the Grushko decomposition space for a free product – then this action preserves the deformation space.

By Theorem 2.35, an inner automorphism preserves all length functions and so must be in the kernel of this action (working up to equivariant isometry, of course): so we can consider it an action of the outer automorphism group.

This perspective has been very fruitful. For example, careful inspection of the deformation space for BS(2, 4) in [16] provides a proof that Out(BS(2, 4)) is not finitely generated via Bass–Serre theory.

In general this action of Out(G) does not have global fixed points. (Sometimes very far from this: in CV_n, point stabilisers are finite [21, 45].) If there is a global fixed point – equivalently a length function which is preserved by all automorphisms – then this is called a *canonical tree*. In this situation we can study the (outer) automorphism group via Bass–Serre theory: as is discussed in Paper 2, when the action satisfies the hypotheses of Theorem 2.34 there is an action of Aut(G) on the same tree. Bass and Jiang in [3] provide a filtration of Out(G) in this situation, which is used in Paper 3.

One way of constructing canonical trees is via a *tree of cylinders*, as defined by Guirardel and Levitt in [38]. This is described in Paper 3; the input is any tree in the deformation space, and an equivalence relation on its edge groups; the output is a tree where the induced splitting is preserved by all outer automorphisms which preserve the deformation space. The equivalence relation is required to satisfy certain axioms; one of their effects is that it can be extended to an equivalence relation on edges in such a way that the equivalence classes form subtrees, which are known as cylinders. From this a new tree T_c is constructed by replacing each cylinder with the cone on its boundary; the axioms on the original equivalence class also guarantee that this is equivariant, in the sense that the new tree inherits an action of G.

The tree of cylinders T_c depends only on the deformation space of T, in the sense that given two minimal, non-trivial trees T, T' in the same deformation space, there is a canonical equivariant isomorphism between T_c and T'_c [38, Corollary 4.10]. In particular this means that this tree of cylinders is fixed by any automorphism which preserves the deformation space, and so can be used to study these (outer) automorphisms.

It is always true that T dominates T_c, but cylinder stabilisers were not necessarily elliptic in T and so T_c might not be an element of the original deformation space. In some cases the failure is spectacular: the tree of cylinders for CV_n is a single point with the trivial action. The deformation space of the tree of cylinders depends on the size of the cylinders: if all cylinders are bounded, or equivalently contain no hyperbolic axis, then the cylinder stabilisers are elliptic in T and so T_c lies in the same deformation space, and conversely [38, Proposition 5.2].

6 Out(F_n) and Outer Space

Let $n \geq 2$, and consider the outer automorphisms Out(F_n). Given two automorphisms φ and ψ representing the same outer automorphism, say they are *isogredient* if there is an inner automorphism γ conjugating one to the other (so $\gamma^{-1}\varphi\gamma = \psi$).

The introduction of train tracks by Bestvina and Handel in [8] proved a useful tool for studying automorphisms of free groups. One of the results of that paper concerns the ranks of fixed subgroups of automorphisms:

Theorem 6.1 (Bestvina–Handel, [8])**.** *Let $\Phi \in \mathrm{Out}(F_n)$. Then,*

$$\sum \max\{\mathrm{rank}(\mathrm{Fix}(\varphi)) - 1, 0\} \leq n - 1,$$

where the sum is taken over representatives, ϕ, of isogredience classes in Φ.

In particular the fixed subgroup of any automorphism has rank bounded by n, so this result resolved the Scott conjecture.

Fix a basis X, and consider lengths of words in this basis. In order to work with outer automorphisms, we consider conjugacy length, setting $\|g\|$ to be the length of the shortest word in its conjugacy class. As we iterate an (outer) automorphism, we can consider how this length changes. By taking λ to be the maximum, $\|x\Phi\|$ over elements x of the basis X, every element satisfies

$$\frac{\|\Phi^k(g)\|}{\|g\|} \leq \lambda^k,$$

In fact, Bestvina–Handel proved that for any $\Phi \in \mathrm{Out}(F_n)$ and $g \in F_n$, we either have that

$$\mu^k \leq \frac{\|\Phi^k(g)\|}{\|g\|} \leq \lambda^k,$$

for some $\mu \in (1, \lambda)$, or

$$Ak^d \leq \frac{\|\Phi^k(g)\|}{\|g\|} \leq Bk^d,$$

with $d \in \{0, 1, \ldots, n-1\}$, and $A < B$ positive constants. Levitt in [51, Theorem 6.2] gives a more precise characterisation, showing that each conjugacy class grows as a polynomial times an exponential under iteration.

In the first case g is *exponentially growing*, and if any element is exponentially growing then Φ is also called exponentially growing. If no elements grow exponentially then Φ is said to be *polynomially growing of degree d*, where d is the smallest degree bounding the growth of every element.

Example 6.2. Consider the following automorphisms defined on basis elements:

$$\begin{array}{ccc} & \varphi & \\ a & \longrightarrow & a \\ b & \longrightarrow & ba \\ c & \longrightarrow & cb \end{array} \qquad \begin{array}{ccc} & \psi & \\ a & \longrightarrow & b \\ b & \longrightarrow & ab \end{array}$$

The automorphism φ is quadratically growing; this is realised by c. Its restriction to $\langle a, b\rangle$ is linearly growing, and the fixed subgroup $\mathrm{Fix}(\varphi)$ is $\langle a, bab^{-1}\rangle$.

The automorphism ψ is exponentially growing; nevertheless the commutator $aba^{-1}b^{-1}$ is periodic: it is fixed by ψ^2.

Following Bestvina and Handel [7], a stronger condition we may put on elements of $\mathrm{Out}(F_n)$ is being *unipotent polyomially growing* or UPG: an element has this property if it is polynomially growing and its image in $\mathrm{GL}_n(\mathbb{Z})$ is unipotent (that is, conjugate to a matrix with all diagonal entries equal to 1). Intuitively, this is excluding "periodic behaviour" – so every element fixed by a power of φ is already fixed by φ, for example, although periodic behaviour can also occur among elements that are growing. By [7, Corollary 5.7.6], every polynomially growing automorphism has a power that is UPG.

Outer Space and its boundary

As mentioned above, projectivised Outer Space CV_n is the space of minimal free F_n actions on metric simplicial trees, working up to equivariant homothety. It is equipped with an action of $\mathrm{Out}(F_n)$. Usually representatives of homothety classes are chosen by requiring covolume 1. The equivariant homeomorphism classes define simplices, as we vary the lengths of the edges in the quotient graph; however not all combinations of edges can be collapsed while the action remains free, so some faces are missing.

Outer space CV_n may be given a topology by identifying an action with its length function (free actions of F_n are irreducible for n at least 2, so Theorem 2.34 applies), giving a subset of $\mathbb{R}^{|F_n|} \setminus \{0\} / \sim$ (the equivalence relation $\sim$ is scaling). (There are other topologies one can put on a deformation space of trees; since the trees in CV_n are locally finite these are all equivalent [36, Section 5].)

Outer Space is not closed; missing faces provide examples of limit points outside CV_n. There are limit points arising in other ways too: for example some have non-trivial edge stabilisers (see Theorem 6.5). It turns out that the closure can be characterised as as space of "very small" actions on real trees:

Definition 6.3. An action of G on a real tree T is *very small* if

(i) arc stabilisers do not contain F_2 (it is small);

(ii) for every elliptic element g and n such that $g^n \neq 1$, $\mathrm{Fix}(g) = \mathrm{Fix}(g^n)$;

(iii) for every ellipic element g, $\mathrm{Fix}(g)$ does not contain a tripod.

We are concerned with free groups, in which case the first condition implies arc stabilisers are trivial or cyclic, and the second condition must hold for all $n \neq 1$.

Theorem 6.4. *The closure of CV_n is precisely the space of minimal, irreducible, very small actions of F_n on real trees.*

That the closure consists only of very small actions is due to Cohen and Lustig [17]; that every very small action arises in this way is due to Horbez [44]; see also Bestvina and Feighn's preprint [6]. Note that some of these actions are *not* on metric simplicial trees, unlike CV_n itself.

One way to construct explicit limit points is to produce a "limiting tree" for an automorphism, by taking an element of CV_n, twisting it by an automorphism, and rescaling the result appropriately. In some cases there are strong existence and uniqueness theorems for these trees. For example,

Theorem 6.5 (Parabolic Orbits Theorem (see [17] and [18])**.** *Let $\Phi \in \mathrm{Out}(F_n)$ be linear and unipotent. Then for any $X \in CV_n$, $\lim_{k\to\infty} \Phi^k(X) = T \in \overline{CV_n}$ exists, is a simplicial tree and lies in the same simplex (that is, any two such limit trees are equivariantly homeomorphic) independently of X. Moreover, T is a simplicial F_n-tree with the following properties.*

(i) Edge stabilisers are maximal infinite cyclic

(ii) Vertex stabilisers are precisely the subgroups $\mathrm{Fix}(\varphi)$*, where $\varphi \in \Phi$ has a fixed subgroup of rank at least 2.*

Note that the theorem is usually stated for *Dehn Twist* automorphisms of F_n: these are exactly the linear unipotent automorphisms (see [18], [47], [8] and [7]).

Another example comes from automorphisms of F_n where no power preserves a non-trivial free splitting of F_n (known as *irreducible with irreducible powers*). These fix two points on the boundary of CV_n: an "attracting tree" and a "repelling tree" (neither are simplicial). Every point in the closure of CV_n, except the repelling tree, tends to the attracting tree under iterating the automorphism [5, 52].

These limiting trees have been used to investigate the conjugacy problem in Out(F_n) (for example, [18]). In Paper 3 the parabolic orbits theorem is used to construct actions of free-by-cyclic groups (with linearly growing defining automorphism) on trees.

Free-by-cyclic groups

Paper 3 studies free-by-cyclic groups. These are semidirect product of F_n by $\mathbb{Z}$, defined by an element φ of Aut(F_n) as

$$F_n \rtimes_\varphi \mathbb{Z} = \langle x_1, \ldots x_n, t : t^{-1}x_i t = x_i\varphi \rangle.$$

By introducing $t' = t^{-1}$ or $t' = tg$ to the presentation and then removing t it is apparent that choosing different representatives of the same outer automorphism or its inverse do not change the isomorphism type of the free-by-cyclic group. In fact more is true; it is an invariant up to conjugacy and inverses within $\mathrm{Out}(F_n)$ [9]. So properties of free-by-cyclic groups may be expected to correspond to properties of their defining outer automorphism.

For example, periodic outer automorphisms give rise to free-by-cyclic groups which are act on trees with all edge and vertex groups isomorphic to $\mathbb{Z}$ (known as generalised Baumslag–Solitar groups); exponentially growing automorphisms yield relatively hyperbolic free-by-cyclic groups [30, 31, 25]; and automorphisms with no periodic conjugacy classes (*atoroidal* automorphisms) give hyperbolic free-by-cyclic groups [12].

In the periodic [50] and atoroidal [49] cases this leads to descriptions of $\mathrm{Out}(G)$. Additionally, the outer automorphisms of all groups isomorphic to $F_2 \rtimes \mathbb{Z}$ were characterised (up to a finite index subgroup) in [9].

Paper 3 proves finite generation for $\mathrm{Out}(G)$ when G is isomorphic to $F_n \rtimes_\varphi \mathbb{Z}$ with φ linearly growing, or $F_3 \rtimes_\varphi \mathbb{Z}$ in general. The techniques used are those from Bass–Serre theory and the study of deformation spaces: in particular we define "nearly canonical" trees that the groups act on, and use these to characterise $\mathrm{Out}(G)$.

References

[1] Roger Alperin and Hyman Bass. Length functions of group actions on Λ-trees. Combinatorial group theory and topology, Sel. Pap. Conf., Alta/Utah 1984, Ann. Math. Stud. 111, 265-378 (1987)., 1987.

[2] Hyman Bass. Covering theory for graphs of groups. *J. Pure Appl. Algebra*, 89(1-2):3–47, 1993.

[3] Hyman Bass and Renfang Jiang. Automorphism groups of tree actions and of graphs of groups. *J. Pure Appl. Algebra*, 112(2):109–155, 1996.

[4] Bachir Bekka, Pierre de la Harpe, and Alain Valette. *Kazhdan's property (T)*, volume 11 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2008.

[5] M. Bestvina, M. Feighn, and M. Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geom. Funct. Anal.*, 7(2):215–244, 1997. (Erratum 7(6):1143, 1997).

[6] Mladen Bestvina and Mark Feighn. Outer limits. Preprint, 1994. available at https://www.math.utah.edu/~bestvina/eprints/bestvina.feighn..outer_limits.pdf.

[7] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $\mathrm{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms. *Annals of Mathematics. Second Series*, 151(2):517–623, 2000.

[8] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[9] O. Bogopolski, A. Martino, and E. Ventura. The automorphism group of a free-by-cyclic group in rank 2. *Comm. Algebra*, 35(5):1675–1690, 2007.

[10] O. V. Bogopolski. Arboreal decomposability of groups of automorphisms of free groups. *Algebra i Logika*, 26(2):131–149, 271, 1987.

[11] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.

[12] P. Brinkmann. Hyperbolic automorphisms of free groups. *Geom. Funct. Anal.*, 10(5):1071–1089, 2000.

[13] R. G. Burns. On finitely generated subgroups of free products. *J. Austral. Math. Soc.*, 12:358–364, 1971.

[14] I. M. Chiswell. Abstract length functions in groups. *Math. Proc. Camb. Philos. Soc.*, 80:451–463, 1976.

[15] Ian Chiswell. *Introduction to Λ-trees*. Singapore: World Scientific, 2001.

[16] Matt Clay. Deformation spaces of G-trees and automorphisms of Baumslag-Solitar groups. *Groups Geom. Dyn.*, 3(1):39–69, 2009.

[17] Marshall M. Cohen and Martin Lustig. Very small group actions on $\mathbf{R}$-trees and Dehn twist automorphisms. *Topology. An International Journal of Mathematics*, 34(3):575–617, 1995.

[18] Marshall M. Cohen and Martin Lustig. The conjugacy problem for Dehn twist automorphisms of free groups. *Commentarii Mathematici Helvetici*, 74(2):179–200, 1999.

[19] Donald J. Collins. The automorphism towers of some one-relator groups. *Proc. London Math. Soc. (3)*, 36(3):480–493, 1978.

[20] Donald J. Collins and Frank Levin. Automorphisms and Hopficity of certain Baumslag-Solitar groups. *Arch. Math. (Basel)*, 40(5):385–400, 1983.

[21] Marc Culler. Finite groups of outer automorphisms of a free group. Contributions to group theory, Contemp. Math. 33, 197-207 (1984)., 1984.

[22] Marc Culler and John W. Morgan. Group actions on $\mathbb{R}$-trees. *Proc. London Math. Soc. (3)*, 55(3):571–604, 1987.

[23] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84:91–119, 1986.

[24] Marc Culler and Karen Vogtmann. A group-theoretic criterion for property FA. *Proc. Amer. Math. Soc.*, 124(3):677–683, 1996.

[25] François Dahmani and Ruoyu Li. Relative hyperbolicity for automorphisms of free products. Preprint, January 2019. arXiv:1901.06760.

[26] Warren Dicks and M. J. Dunwoody. *Groups acting on graphs*, volume 17 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989.

[27] Max Forester. Deformation and rigidity of simplicial group actions on trees. *Geom. Topol.*, 6:219–267, 2002.

[28] D. I. Fouxe-Rabinovitch. Über die Automorphismengruppen der freien Produkte. I. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):265–276, 1940.

[29] D. I. Fouxe-Rabinovitch. Über die Automorphismengruppen der freien Produkte. II. *Rec. Math. [Mat. Sbornik] N. S.*, 9 (51):183–220, 1941.

[30] Francois Gautero and Martin Lustig. The mapping-torus of a free group automorphism is hyperbolic relative to the canonical subgroups of polynomial growth. Preprint, July 2007. `arXiv:0707.0822`.

[31] Pritam Ghosh. Relative hyperbolicity of free-by-cyclic extensions. Preprint, February 2018. `arXiv:1802.08570`.

[32] N. D. Gilbert. Presentations of the automorphism group of a free product. *Proc. London Math. Soc. (3)*, 54(1):115–140, 1987.

[33] N. D. Gilbert, J. Howie, V. Metaftsis, and E. Raptis. Tree actions of automorphism groups. *J. Group Theory*, 3(2):213–223, 2000.

[34] E. S. Golod. On nil-algebras and finitely approximable p-groups. *Izv. Akad. Nauk SSSR Ser. Mat.*, 28:273–276, 1964.

[35] I. Gruschko. Über die Basen eines freien Produktes von Gruppen. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):169–182, 1940.

[36] Vincent Guirardel and Gilbert Levitt. Deformation spaces of trees. *Groups Geom. Dyn.*, 1(2):135–181, 2007.

[37] Vincent Guirardel and Gilbert Levitt. The outer space of a free product. *Proc. Lond. Math. Soc. (3)*, 94(3):695–714, 2007.

[38] Vincent Guirardel and Gilbert Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

[39] Vincent Guirardel and Gilbert Levitt. Splittings and automorphisms of relatively hyperbolic groups. *Groups Geom. Dyn.*, 9(2):599–663, 2015.

[40] Vincent Guirardel and Gilbert Levitt. *JSJ decompositions of groups*, volume 395. Paris: Société Mathématique de France (SMF), 2017.

[41] Marshall Hall, Jr. Coset representations in free groups. *Trans. Amer. Math. Soc.*, 67:421–432, 1949.

[42] Marshall Hall, Jr. A topology for free groups and related groups. *Ann. of Math. (2)*, 52:127–139, 1950.

[43] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[44] Camille Horbez. The boundary of the outer space of a free product. *Isr. J. Math.*, 221(1):179–234, 2017.

[45] D. G. Khramtsov. Finite subgroups of groups of outer automorphisms of free groups. *Algebra and Logic*, 26(3):217–228, 1987.

[46] Sabine Koppelberg and Jacques Tits. Une propriété des produits directs infinis de groupes finis isomorphes. *C. R. Acad. Sci. Paris Sér. A*, 279:583–585, 1974.

[47] Sava Krstić, Martin Lustig, and Karen Vogtmann. An equivariant Whitehead algorithm and conjugacy for roots of Dehn twist automorphisms. *Proceedings of the Edinburgh Mathematical Society. Series II*, 44(1):117–141, 2001.

[48] Alexander Kurosch. Die Untergruppen der freien Produkte von beliebigen Gruppen. *Math. Ann.*, 109(1):647–660, 1934.

[49] Gilbert Levitt. Automorphisms of hyperbolic groups and graphs of groups. *Geom. Dedicata*, 114:49–70, 2005.

[50] Gilbert Levitt. On the automorphism group of generalized Baumslag-Solitar groups. *Geom. Topol.*, 11:473–515, 2007.

[51] Gilbert Levitt. Counting growth types of automorphisms of free groups. *Geom. Funct. Anal.*, 19(4):1119–1146, 2009.

[52] Gilbert Levitt and Martin Lustig. Irreducible automorphisms of F_n have north-south dynamics on compactified outer space. *J. Inst. Math. Jussieu*, 2(1):59–72, 2003.

[53] R. C. Lyndon. Length functions in groups. *Math. Scand.*, 12:209–234, 1963.

[54] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Classics in Mathematics. Springer-Verlag, Berlin, 2001. Reprint of the 1977 edition.

[55] Darryl McCullough and Andy Miller. *Symmetric automorphisms of free products*, volume 582. Providence, RI: American Mathematical Society (AMS), 1996.

[56] Ashot Minasyan. New examples of groups acting on real trees. *J. Topol.*, 9(1):192–214, 2016.

[57] D. I. Moldavanskii. Isomorphism of the Baumslag-Solitar groups. *Ukr. Math. J.*, 43(12):1569–1571, 1991.

[58] John W. Morgan and Peter B. Shalen. Free actions of surface groups on $\mathbb{R}$-trees. *Topology*, 30(2):143–154, 1991.

[59] B. H. Neumann. On the number of generators of a free product. *J. London Math. Soc.*, 18:12–20, 1943.

[60] Walter Parry. Axioms for translation length functions. In *Arboreal group theory. Proceedings of a workshop, held September 13-16, 1988, in Berkeley, CA (USA)*, pages 295–330. New York etc.: Springer-Verlag, 1991.

[61] Luis Ribes and Pavel Zalesskii. *Profinite groups. 2nd ed*, volume 40. Berlin: Springer, 2nd ed. edition, 2010.

[62] N. S. Romanovskiĭ. On the residual finiteness of free products with respect to subgroups. *Izv. Akad. Nauk SSSR Ser. Mat.*, 33:1324–1329, 1969.

[63] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[64] John Stallings. Coherence of 3-manifold fundamental groups. In *Séminaire Bourbaki, Vol. 1975/76, 28 ème année, Exp. No. 481*, pages 167–173. Lecture Notes in Math., Vol. 567. 1977.

[65] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[66] Yasuo Watatani. Property T of Kazhdan implies property FA of Serre. *Math. Japon.*, 27(1):97–103, 1982.

Paper 1: **A Bass–Serre theoretic proof of a theorem of Burns and Romanovskii**

Naomi Andrew

ABSTRACT. A well known theorem of Burns and Romanovskii states that a free product of subgroup separable groups is itself subgroup separable. We provide a proof using the language of immersions and coverings of graphs of groups, due to Bass.

1 Introduction

Subgroup separability is a strengthening of residual finiteness. It has many equivalent definitions; we will use the following:

Definition 1.1. Let G be a group and H be a subgroup of G. Say H is *separable in* G if for every element $g \in G \setminus H$, there is a finite index subgroup K of G containing H and not g.

If every finitely generated subgroup of G is separable in G, say that G is *subgroup separable*.

This note is concerned with the following theorem, giving that subgroup separability is closed under finite free products:

Main Theorem. *Suppose G is a finite free product of subgroup separable groups. Then G is itself subgroup separable.*

This theorem is originally due (independently) to Romanovskii [9] and Burns [4], and there are subsequent proofs due to Gitik [5] and Wilton [12].

The object of this paper is to provide a new proof of this theorem, generalising the proof Stallings gives in [11] that free groups are subgroup separable (a theorem originally due to Hall [7]; see also [3]) to graphs of groups with trivial edge groups. Previous proofs have worked with other objects associated to the free product, such as graphs of spaces (viewed topologically) or relative Cayley graphs; ours uses the graph of groups structure more immediately.

The notions of immersions and coverings of graphs of groups, due to Bass, are rather more technical than those used by Stallings for graphs. So we begin by covering the necessary definitions for graphs of groups, then the notion of *Kurosh rank* for a subgroup of a free product (given an action on a tree). Given a group acting on any set (or in particular a tree) we provide a way to calculate the index of a subgroup from its action in Lemma 3.4. Finally in Section 4 we combine these results to show how to complete an immersion of graphs of groups to a cover, and how this implies the Burns–Romanovskii theorem.

2 Graphs of Groups

We follow Bass' exposition [2], although we change some notation. There are other sources covering the same material, such as [10]. Unlike many expositions, we put the action on the right.

Graphs of groups are a combinatorial tool encoding group actions on trees: they consist of a graph corresponding to the quotient together with edge and vertex groups corresponding to stabilisers.

Definition 2.1. A graph Γ consists of a set of vertices $V\Gamma$ and a set of edges $E\Gamma$, together with two maps: $\iota : E\Gamma \to V\Gamma$; and an involution $E\Gamma \to E\Gamma, e \to \bar{e}$. We also define $\tau : E\Gamma \to V\Gamma, \tau(e) = \iota(\bar{e})$. An *orientation* of Γ is a choice of one edge from each pair $\{e, \bar{e}\}$.

Definition 2.2. A *Graph of Groups*, $\mathcal{G}$, consists of

- a connected graph $\Gamma_{\mathcal{G}}$;
- for each vertex v of $\Gamma_{\mathcal{G}}$, a group G_v;
- for each edge e of $\Gamma_{\mathcal{G}}$, a group G_e such that $G_e = G_{\bar{e}}$ and there is a monomorphism $\alpha_e : G_e \to G_{\tau(e)}$.

Where the graph of groups is clear, we may just refer to Γ for the underlying graph.

There are two main ways of defining the fundamental group and universal cover of a graph of groups: by a maximal tree, and by considering loops at a base point. We follow Bass, and consider paths and loops in the graph of groups.

Definition 2.3 (Paths)**.** Let $F(\mathcal{G})$ be the group generated by all the vertex groups and all the edges of $\mathcal{G}$, subject to relations $e\alpha_e(g)\bar{e} = \alpha_{\bar{e}}(g)$ for $g \in G_e$. Note that taking $g = 1$ this gives that $e^{-1} = \bar{e}$, as expected.

Define a *path* (of length n) in $F(\mathcal{G})$ to be a sequence $g_0e_1g_1 \ldots e_ng_n$, where each e_i has $\iota(e_i) = v_{i-1}$ and $\tau(e_i) = v_i$ for some vertices v_i (so there is a path in the graph), and each $g_i \in G_{v_i}$. A *loop* is a path where $v_0 = v_n$.

The set of all paths in $F(\mathcal{G})$ forms a groupoid (sometimes called the fundamental groupoid of $\mathcal{G}$).

Definition 2.4 (Reduced paths)**.** A path is *reduced* if it contains no subpath of the form $e\alpha_e(g)\bar{e}$ (for $g \in G_e$). A loop is *cyclically reduced* if, in addition to being reduced, $e_n(g_n g_0)e_1$ is not of the form $e\alpha_e(g)\bar{e}$.

Every path is equivalent (by the relations for $F(\mathcal{G})$) to a reduced path, and similarly every loop is equivalent to both a reduced loop and a cyclically reduced loop. In general these reduced representations are not unique, although all equivalent (cyclically) reduced paths (or loops) will have the same edge structure. Note that a cyclically reduced loop might not be at the same vertex as the original loop.

Definition 2.5. The *fundamental group of* $\mathcal{G}$ at a vertex v is the set of loops in $F(\mathcal{G})$ at v, and is denoted $\pi_1(\mathcal{G}, v)$.

The isomorphism class of this group does not depend on the vertex chosen. (In fact, the two groups obtained by choosing different base vertices are conjugate in the groupoid.)

We define the Bass–Serre tree (or universal cover) in the corresponding way:

Definition 2.6 (Bass–Serre Tree)**.** Let T be the graph formed as follows: the vertex set consists of 'cosets' $G_w p$, where p is a path in $F(\mathcal{G})$ from w to v. There is an edge(-pair) joining two vertices $G_{w_1}p_1$ and $G_{w_2}p_2$ if $p_1 = eg_{w_2}p_2$ or $p_2 = eg_{w_1}p_1$ (with $g_w \in G_w$).

This graph is a tree, and there is a right action of $\pi_1(\mathcal{G}, v)$ on the vertex set, since this multiplication is possible in the groupoid, and the paths will still start at v. This action preserves adjacency and is without inversions, and so $\pi_1(\mathcal{G}, v)$ acts on T.

There is another construction, that takes a group action on a tree and returns a graph of groups:

Definition 2.7 (Quotient graph of groups)**.** Suppose a group G acts on a tree T. Form a graph of groups whose underlying graph to be the quotient graph of the action, with edge and vertex groups are assigned as follows: choose subtrees $T^v \subseteq T^e$ such that T^v contains exactly one representative of each vertex orbit (that is, a lift of a maximal tree in the quotient), and T^e exactly one representative of each edge orbit, in such a way that at least one end of every edge is in T^v. We abuse notation a little by identifying vertices in T^v and edges in T^e with their orbits (that is, their image in the quotient graph). Set the vertex and edge groups to be the stabilisers G_v and G_e. To define the monomorphisms, we choose elements $g_v \in G$ which act to bring each vertex of T^e into T^v: if $v \in T^v$ then set $g_v = 1$, and otherwise choose any element with this property. Now we may set the monomorphisms α_e to be the composition of the inclusion with conjugation by our chosen elements (so $s \mapsto g_{\tau(e)}^{-1} s g_{\tau(e)}$).

In many cases, the full complexity of this definition is unnecessary: we can consider the stabiliser of *any* orbit representative and assert that the injection is the composition of an inclusion and the relevant conjugation.

Theorem 2.8. *Up to isomorphism of the structures concerned, the processes of constructing the quotient graph of groups, and of constructing the fundamental group and Bass–Serre tree are mutually inverse.*

From the perspective of groups acting on trees, the isomorphisms required are an isomorphism between the original group and the fundamental group, and an equivariant isometry between the original tree and the Bass–Serre tree. From the perspective of graphs of groups, they are an isomorphism of underlying graphs, together with isomorphisms of corresponding edge and vertex groups (and these must respect the edge monomorphisms). This is the fundamental result linking actions on trees with splittings of groups.

Morphisms, immersions and covers

If G is the fundamental group of a graph of groups, then any subgroup of G will also act on the Bass–Serre tree, and this action will give a quotient graph of groups carrying that subgroup. In the case of a free action (where G must be a free group), then we know that the quotient graph is a cover of the original graph - in fact, there is a correspondence between covers and subgroups. This point of view has been fruitful for investigating free groups, and is the main tool of Stallings' paper [11]. The aim of Bass' definitions of morphisms and covers (and immersions) is to recover the same correspondence for graphs of groups.

There is a lot of structure, and so any definition of a morphism must feature a graph map and several group homomorphisms. It turns out that slightly more data is needed as well, in the form of group elements attached to each edge and vertex.

The definitions we give here are specialised to the case of free products - that is, when all G_e are trivial. In general the elements δ_e defined below must satisfy conditions involving the edge group inclusions, but these are automatically satisfied for any choices with trivial edge groups.

Definition 2.9. Suppose $\mathcal{H}$ and $\mathcal{G}$ are graphs of groups with all edge groups trivial. A *morphism* of graphs of groups $\Phi : \mathcal{H} \to \mathcal{G}$ consists of:

- a graph morphism $\varphi : \Gamma_{\mathcal{H}} \to \Gamma_{\mathcal{G}}$;
- a group homomorphism $\phi_v : H_v \to G_{\varphi(v)}$ for every vertex $v \in \Gamma_{\mathcal{H}}$;
- an element λ_v in $\pi_1(\mathcal{G}, \varphi(v))$ for every vertex v of $\Gamma_{\mathcal{H}}$;
- an element $\delta_e \in G_{\varphi(\iota(e))}$ for every edge of $\Gamma_{\mathcal{H}}$.

Such a morphism induces maps on the structures that can be defined from a graph of groups, as follows:

- A homomorphism (of groups) $F(\mathcal{H}) \to F(\mathcal{G})$ by $s \mapsto \lambda_v^{-1}\phi_v(s)\lambda_v$ for $s \in H_v$ and $e \mapsto \lambda_{\iota(e)}^{-1}\delta_e^{-1}e\delta_{\bar{e}}\lambda_{\tau(e)}$.
- A homomorphism Φ_P of fundamental groupoids, by restricting that map to the paths in $F(\mathcal{H})$. Note that, in this case, for each edge e, the extra elements introduced at e, $\iota(e)$ and $\tau(e)$ will cancel to leave δ_e and $\delta_{\bar{e}}$ which are elements of the vertex groups at either end.
- A homomorphism $\Phi_v : \pi_1(\mathcal{H}, v) \to \pi_1(\mathcal{G}, \varphi(v))$ of the fundamental groups, by further restricting the above map to loops at v.
- An equivariant graph map $\tilde{\Phi}$ on the Bass–Serre trees, defined on vertices by $H_w p \mapsto G_{\varphi(w)}\lambda_w\Phi_P(p)$.

Additionally, we can define a 'local map' at each edge of $\mathcal{G}$. (Requiring edge stabilisers to be trivial simplifies this considerably compared to Bass' general definition.)

Given a vertex v and an edge with $\tau(e) = v$, the lifts of e at a single vertex in the Bass–Serre tree correspond to the elements of G_v, by identifying the edge $[G_v p, G_w esp]$ with the element s.

Given a morphism $\Phi : \mathcal{H} \to \mathcal{G}$, let v be a vertex of $\Gamma_\mathcal{H}$ and f be an edge of $\Gamma_\mathcal{G}$ with $\tau(f) = \varphi(v)$. Define a map

$$\Phi_{v/f} : \coprod_{e \in \varphi^{-1}(f), \tau(e) = v} H_v \to G_{\varphi(v)}$$

by

$$h \mapsto \delta_{\bar{e}}\phi_v(h).$$

Alternatively, we can view $\Phi_{v/f}$ as a map $H_v \times \{e \in E(\mathcal{H}) : \iota(e) = v, \varphi(e) = f\} \to G_{\varphi(v)}$ taking $(H_v, e) \mapsto \delta_e\phi_v(H_v)$. These maps are useful for "locally" understanding the image of the Bass–Serre tree under a morphism: see Proposition 2.12.

Given two group actions on trees, and an equivariant map between the trees we can induce a graph of groups morphism between the quotient graphs of groups. We continue to assume that the actions are free on edges.

Proposition 2.10. *Suppose S is an H-tree, T a G-tree, $\psi : H \to G$ is a homomorphism and $f : S \to T$ is a ψ-invariant graph map. (That is, f sends vertices to vertices, edges to edges, preserves adjacency, and $vhf = (vf)(h\psi)$.) Let $\mathcal{H}$ and $\mathcal{G}$ be the quotient graphs of groups corresponding to the actions of H on S and G on T respectively. Then ψ, f induce a graph of* *groups morphism $\mathcal{H} \to \mathcal{G}$, which (after the isomorphisms required by Theorem 2.8) recovers ψ and f as maps of fundamental groups and Bass–Serre trees.*

For details, and full proofs, see [2, Section 4]. Here we give sufficient details to explain how the induced morphism is constructed. First, since f was ψ-equivariant it induces a graph map φ on the quotients $S/H \to T/G$: this is our map between the underlying graphs. Let S^v, S^e, T^v and T^e be the subtrees of S and T used in Definition 2.7, and let h_u and g_v be the elements given there which bring vertices of S^e and T^e into S^v and T^v respectively.

For vertices v in S^v, choose a k_v in G, so that $f(v)k_v$ is in T^v; similarly for an edge e let k_e be an element of G with $f(e)k_e$ in T^e, and $k_e = k_{\bar{e}}$. Since f is ψ-equivariant, ψ takes stabilisers to stabilisers, though not necessarily of the preferred representative of each orbit. Define $\phi_v : H_v \to G_{f(v)k_v}$ by $s \mapsto k_v^{-1}\psi(s)k_v$.

To define a morphism also requires elements λ_v and δ_e. For an edge e in S^e, let $v = \iota(e)$, $x = f(v)k_e$ and $y = vh_v$. Then let $\delta_e = g_x^{-1}k_e^{-1}\psi(h_v)k_y$. To see that this is indeed an element of $G_{\varphi(v)}$, observe that both $\psi(h_v)k_y$ and $k_e g_x$ act to bring $f(v)$ into T^v, and the vertex group $G_{\varphi(v)}$ (of $\mathcal{G}$) is defined as the stabiliser of the vertex of T^v in the same orbit as $f(v)$.

We will want to let $\lambda_v = k_v^{-1}$; however to be in the right group we must first apply the isomorphism (of Theorem 2.8 from $G \to \pi_1(\mathcal{G}, \varphi(v))$. (This can be thought of as "reading" the path between v and vk_v in T.)

There is usually some choice as to the subtrees used to construct the quotient graph of groups. In particular, if we arrange for $f(S^v)$ to (maximally) intersect T_v, we may choose several λ_v to be 1, simplifying the morphism and allowing choices of basepoint (in $\mathcal{H}$) so that the map on fundamental groups is "as written" – meaning it does not involve a conjugation by a non-trivial λ_v.

We are most interested in studying subgroups H of a group G with an action on a tree T, so usually ψ is an inclusion, and f is either the inclusion $T_H \to T$ (sometimes a slightly larger H-invariant tree) or the identity $T \to T$. In this case, we should expect the induced morphism to have good properties, since the map on trees makes no identifications. These good properties are characterised by the morphism being a *cover* or *immersion*.

In the context of a graph (with no groups) a covering map corresponds to the usual topological definition, and an immersion relaxes "locally bijective" to "locally injective". This allows the universal cover of the immersed graph to be strictly contained in the original universal cover.

The Bass–Serre tree gives the "universal cover" in this world: of course it is not a true cover, since an edge may have many (even infinitely many) preimages at each vertex.

Similarly, our covers and immersions might have several preimages of an edge at a vertex:

Definition 2.11 ([2, Definition 2.6]). A morphism $\Phi : \mathcal{H} \to \mathcal{G}$ is an *immersion* if

(1) each $\phi_v : H_v \to G_{\phi(v)}$ is injective and

(2) each $\Phi_{v/e}$ is injective

And a *covering* if the second condition is replaced by

(2′) each $\Phi_{v/e}$ is bijective.

Bass proves that these properties exactly characterise the situation of a subgroup acting on a subtree:

Proposition 2.12 ([2, Proposition 2.7]). *A morphism Φ is an immersion if and only if Φ_{v_0} (on fundamental groups), and $\tilde{\Phi}$ (on Bass–Serre trees) are injective. Furthermore, it is a covering if and only if Φ_{v_0} (on fundamental groups) is injective and $\tilde{\Phi}$ (on Bass–Serre trees) is bijective.*

Viewing $\Phi_{v/f}$ as a map $H_v \times \{e \in E(\mathcal{H}) : \iota(e) = v, \varphi(e) = f\} \to G_{\varphi(v)}$ taking $(H_v, e) \mapsto \delta_e \phi_v(H_v)$, we have that it will be injective if and only if ϕ_v is, and the δ_e represent different right cosets of $G_{\varphi(v)}/\phi_v(H_v)$.

One way to construct immersions is by taking a *subgroup of subgroups*: restrict to a subgraph of the underlying graph, and take subgroups of each vertex group. Then letting Φ consist of the graph and group inclusions, and all λ_v and δ_e trivial, this is an immersion.

3 Kurosh rank and finite index subgroups of free products

Recall Kurosh's theorem about subgroups of free products:

Theorem 3.1 ([8]). *Suppose G is a free product $* G_i$ (over some index set I) and H is a subgroup of G. Then $H \cong (* H_j) * F$ where each H_j is isomorphic to an intersection $H \cap G_i^{k_i}$ of H with a conjugate of some G_i. Further, the set $\{H_j\}$ is unique up to conjugation and reindexing, and the rank of F is uniquely determined.*

The idea of Kurosh rank is inspired by this theorem, and aims to measure the "complexity" of such a subgroup in terms of its free factors. Here we take the approach of [1] of defining it with respect to an action on a tree:

Definition 3.2. For a group G and a G-tree T with trivial edge stabilisers, the *Kurosh rank* (relative to T) of a subgroup H is

$$\kappa_T(H) = \text{rank}(H\backslash T) + |\{Hv \in H\backslash T : H_v \neq 1\}|$$

where $\text{rank}(H\backslash T)$ is the number of edges outside a maximal tree.

The reduced Kurosh rank of a subgroup is $\overline{\kappa}_T(H) = \max\{\kappa_T(H) - 1, 0\}$.

Note that $\text{rank}(H\backslash T)$ is the rank of the fundamental group of the graph $H\backslash T$. In particular, in the case of a free group acting freely on a tree, (reduced) Kurosh rank reduces to the usual definition of (reduced) rank.

Proposition 3.3 ([1, Proposition 2.3])**.** *Let H be a subgroup of G, and T a G-tree with trivial edge stabilisers. Then*

(1) The Kurosh rank with respect to any H-invariant subtree T' of T, $\kappa_{T'}(H)$ is equal to the Kurosh rank with respect to the minimal H-invariant subtree T_H, $\kappa_{T_H}(H)$.

(2) The Kurosh rank $\kappa_T(H)$ is finite if and only if the quotient $H\backslash T_H$ is finite.

If T is clear from context then we will often just write $\overline{\kappa}$, without subscripts, even if we are reasoning with some H-subtree.

By Grushko's theorem [6], the Kurosh rank of a subgroup is bounded above by its true rank. (Each vertex group adds 1 to the Kurosh rank while adding at least 1 to the true rank of the subgroup.)

We need to calculate the index of a subgroup from its covering graph of groups, for which we use the following lemma:

Lemma 3.4. *Suppose a group G acts transitively on a set X, and H is a subgroup of G. Let X_0 be a set of orbit representatives for the action of H on X. Then*

$$[G : H] = \sum_{x \in X_0} [G_x : H_x].$$

In particular, $[G : H]$ is finite if and only if X_0 is finite as is every index $[G_x : H_x]$.

Proof. We will exhibit a bijection (though it is very far from canonical) from $\coprod_{x\in X_0} H_x\backslash G_x$ to $H\backslash G$. To define this, fix a base point $x_0 \in X$, and for every element x of X a $g[x] \in G$ such that $g[x]x_0 = x$.

Now define a function sending a coset $H_x g$ to the coset $Hgg[x]$. This is well defined in the sense that if g_1 and g_2 are in the same H_x-coset of G_x, then $(g_1 g[x])(g_2 g[x])^{-1} = g_1 g_2^{-1}$ which is an element of H_x and in particular H. To see it is injective, suppose there are orbit representatives x, y and elements $g_1 \in G_x, g_2 \in G_y$ such that $Hg_1g[x] = Hg_2g[y]$. That is, $g_1g[x]g[y]^{-1}g_2^{-1}$ is an element of H. But this element moves y to x, and so these must represent the same orbit. So we may assume $x = y$, and this reduces to considering $g_1g_2^{-1}$. This is an element of H, and also of G_x, and therefore of H_x, so g_1 and g_2 represent the same H_x-coset.

To see surjectivity, we will write an arbitrary $g \in G$ as a product $h\hat{g}g[x]$, where $\hat{g} \in G_x$. To do this, let y be the element gx_0, and observe that we may write y as $hx = hg[x]x_0$ for some orbit representative x. The element $hg[x]$ is in the same G_y-coset as g, so we may write $g = \tilde{g}hg[x]$ with $\tilde{g} \in G_y$. But recall that $G_y = hG_xh^{-1}$, so we have that $\tilde{g} = h\hat{g}h^{-1}$, with $\hat{g}$ an element of G_x. In particular, we have that $g = (h\hat{g}h^{-1})hg[x] = h\hat{g}g[x]$, as required. □

Given any (not necessarily transitive) action, we can calculate the index of a subgroup by restricting our attention to one orbit and using Lemma 3.4. In the context of an action on a tree, this would mean looking at the orbit of a single edge or vertex and considering the edge or vertex groups that arise in the covering graph of groups. In particular, for a free product, counting the occurences of any edge in the same G-orbit will give the the index. If the original action gave rise to a finite quotient graph, so will the action of a finite index subgroup. In particular, if the original action was minimal, the minimal invariant subtree for the subgroup will be the whole tree again.

This gives a criterion for a subgroup of a graph of groups to be finite index: this happens if and only if the covering graph of groups has finite underlying graph, and every vertex (or indeed edge) group is finite index in the relevant vertex group of the original graph of groups.

Note that this lemma provides a link between the index and Kurosh rank of a finite index subgroup of a free product, though this is complicated somewhat by the possibility that a vertex has non-trivial stabiliser under G but is trivially stabilised by the subgroup H. However, if we reduce to the case that the vertex groups are infinite these complications disappear. (Restricting to normal subgroups provides a different simplification.)

Corollary 3.5. *Suppose a group G is expressed as a free product of infinite groups, and H is a finite index subgroup of G. Then $\overline{\kappa}(H) = [G : H]\overline{\kappa}(G)$.*

Proof. Represent G as the fundamental group of a graph of groups $\mathcal{G}$ with trivial edge stabilisers, and let $\mathcal{H}$ be the covering graph of groups corresponding to the action of H on the Bass–Serre tree of $\mathcal{G}$. Write $\overline{\kappa}(H) = |E\mathcal{H}| - |\{v \in V\mathcal{H} : H_v = 1\}|$. If a vertex has trivial stabiliser under H, it will also have trivial stabiliser under G, since otherwise G_v

would be a finite subgroup. So use Lemma 3.4 to rewrite as follows:

$$\begin{aligned}\overline{\kappa}(H) &= |E\mathcal{H}| - |\{v \in V\mathcal{H} : H_v = 1\}| \\ &= [G : H]|E\mathcal{G}| - [G : H]|\{v \in V\mathcal{G} : G_v = 1\}| \\ &= [G : H]\overline{\kappa}(G). \end{aligned}$$

□

For a free group acting either freely or as a free product of infinite cyclic groups, the Kurosh rank and the true rank agree, so this recovers Schreier's formula for free groups:

$$\text{rank}(H) - 1 = [F : H](\text{rank}(F) - 1).$$

4 Subgroup separability for free products

In this section we provide the Bass–Serre theoretic proof of the theorem that free products of subgroup separable groups are themselves subgroup separable. The proof is in three steps, dealing in turn with "completing" a graph of groups immersion, enlarging the vertex groups, and then doing this in general for all finitely generated subgroups of a free product.

Stallings' proof (in the free group case) begins with a (finite) labelled graph where the labelling provides an immersion to a rose. The lifts of any edge in a cover would provide a bijection from the vertex set to itself, and the lifts present in the immersion give a partial bijection. Thus any way of completing the partial bijection to a full bijection is admissible in a cover. (There are only finitely many options, and all of them will work.) The condition on a cover can be checked one edge of the rose at a time, so we may do this separately to each edge and the final graph will be a covering graph.

There are several obstacles in extending this to graphs of groups. First, the graph (of groups) we are covering may have more than one vertex. In the free group case this is easily surmountable, although a little care is needed: we must first make sure that all vertices have equal numbers of lifts. We can achieve this by adding isolated vertices to the immersed graph, to take each vertex up to the maximum. Notice that after adding edges (which now correspond to bijections between different kinds of vertex, in general), every vertex will be connected to at least one vertex of every kind. Since at least one vertex had no extra preimages added, all of these are connected. Thus the constructed graph will have only one component.

Another obstacle is that the notion of "local injectivity" is different, and we will need to assign δ_e values to any new edges in a way which preserves the immersion. Finally, we will have to alter the vertex groups so that all of them are finite index, otherwise (by Lemma 3.4) the subgroup cannot be.

We deal in turn with the obstacles presented in generalising Stallings' proof. First, Theorem 4.1 gives a way to complete an immersion to a cover when the vertex groups have finite index image in the target graph of groups; then Theorem 4.2 gives sufficient conditions for enlarging the vertex groups so this condition is met. Finally we prove the Burns–Romanovskii theorem by combining these to produce a covering graph of groups containing a given finite index subgroup but excluding any given element outside it.

Theorem 4.1. *Suppose G is a free product, expressed as the fundamental group of a graph of groups $\mathcal{G}$ where every edge group is trivial. Suppose H is a subgroup of G, corresponding to an immersion $\Phi : \mathcal{H} \to \mathcal{G}$, where $\Gamma_{\mathcal{H}}$ is finite and each H_v is mapped to a finite index subgroup of $G_{\varphi(v)}$.*

Then there is a finite index subgroup M of G containing H as a free factor.

Proof. By Lemma 3.4, in a cover the index of the subgroup can be calculated by looking at the preimages of any edge or vertex and their stabilisers. So we need to ensure these are equal. To this end, for each vertex u of $\mathcal{G}$ calculate

$$d_u = \sum_{v \in \varphi^{-1}(u)} [G_u : H_v].$$

Each d_u is finite: the sum is over finitely many vertices (since $\mathcal{H}$ is finite), and each $[G_u : H_v]$ is finite by assumption. Since $\mathcal{G}$ is also finite, there is a maximum among the d_u, say d. This will be the degree of the cover. For each vertex in $\mathcal{G}$ add $d - d_u$ isolated vertices to $\mathcal{H}$, declaring them to be in the pre-image of u, and assigning each the full subgroup G_u. (Recalculating d_u after doing this, all are equal to d.)

Though it is disconnected, we can still extend Φ to the new vertices: each v is in the pre-image of some vertex u of $\mathcal{G}$, and set each new ϕ_v to be the identity map. We now need to add edges, further extending the morphism Φ by assigning $\varphi(e)$ and δ_e as we do.

We have the local maps on cosets, $\Phi_{v/f}$ and we may extend these to the new vertices (note that, where there are no edges in the pre-image of f at v, this is a map from the empty set). These maps are all injections since we began with an immersion and maps out of the empty set must be injective. Our goal is that they should all be bijections: we will need to add more edges, choosing values for δ_e to achieve this.

For each edge f of $\mathcal{G}$, there should be d pre-images in $\mathcal{H}$. Consider a vertex v of $\mathcal{H}$, and the pre-images e of f at v. The values δ_e form a partial system of coset representatives for $G_{\varphi(v)}/\phi_v(H_v)$, since $\Phi_{v/f}$ is injective.

Suppose f has initial vertex u and terminal vertex y. The edges in the pre-image of f provide a partial bijection

$$\coprod_{v\in\varphi^{-1}(u)} G_{\varphi(v)}/\phi_v(H_v) \to \coprod_{x\in\varphi^{-1}(y)} G_{\varphi(x)}/\phi_x(H_x),$$

by

$$\phi_{\iota(e)}(H_{\iota(e)})\delta_e \mapsto \phi_{\tau(e)}(H_{\tau(e)})\delta_{\bar{e}}.$$

Both these disjoint unions have cardinality d, so this can be completed to a bijection. Add new edges (in the pre-image of f) and coset representatives δ_e and $\delta_{\bar{e}}$ according to this bijection. (The choice of bijection will usually change the subgroup we construct, but not its index.)

Let $\mathcal{M}$ be the graph of groups constructed by repeating this for each edge in $\mathcal{G}$, and Φ be the extension of the original morphism to all of $\mathcal{M}$. This process added finitely many edges to $\mathcal{H}$. Every connected component of $\mathcal{M}$ contains at least one pre-image of each vertex of $\mathcal{G}$, and since at least one vertex of $\mathcal{G}$ had no pre-images added, and this means the underlying graph of $\mathcal{M}$ will be connected. Also, each $\Phi_{v/f}$ is now bijective, so the morphism Φ has been extended to a cover.

Picking a base point for $\mathcal{M}$ (in the pre-image of a chosen base point for $\mathcal{G}$) we recover a subgroup M of G, which has index d since $\mathcal{M}$ by Lemma 3.4. □

Just as in the free group case, restricting $\mathcal{M}$ to the edges and vertices of $\mathcal{H}$ recovers $\mathcal{H}$: so we may view H as a free factor of $M = \pi_1(\mathcal{M})$.

For a general subgroup H of G, the vertex groups of $\mathcal{H}$ are not finite index subgroups of the corresponding vertex groups of $\mathcal{G}$ so the process used proving Theorem 4.1 will not terminate – in fact, any cover must have infinite degree by Lemma 3.4, so there will be infinitely many edges in each pre-image.

So to say anything for general H, we must first replace each vertex group H_v with a group mapping to a finite index subgroup of $G_{\varphi(v)}$. Done carelessly, this enlarging of vertex groups is likely to cause some δ_e values to represent the same coset, and we will no longer have an immersion. So care – and separability assumptions – will be needed as we do this.

Theorem 4.2. *Suppose G is a free product, expressed as the fundamental group of a graph of groups $\mathcal{G}$ where every edge group is trivial. Suppose H is a subgroup of G, corresponding to an immersion $\Phi : \mathcal{H} \to \mathcal{G}$ with Γ_H finite. If each $\phi_v(H_v)$ is separable in $G_{\varphi(v)}$, then there is a finite index subgroup K of G corresponding to a cover $\mathcal{K}$ that contains $\mathcal{H}$ as a subgraph of subgroups.*

Proof. Our first goal is to alter $\mathcal{H}$ and Φ, so each vertex group maps to a finite index subgroup of the relevant G_v, while keeping Φ an immersion. In order to achieve this, we must ensure that the elements δ_e continue to represent different cosets $G_{\varphi(v)}/\phi_v(H_v)$.

For each vertex v of $\mathcal{H}$ and edge f with $\iota(f) = \varphi(v)$, let $X_{v/f}$ be the finite set of elements $\delta_{e_i}^{-1}\delta_{e_j}$ where e_i and e_j are distinct edges with $\iota(e_i) = \iota(e_j) = v$ and $\varphi(e_i) = \varphi(e_j) = f$. Let X_v be the disjoint union of the $X_{v/f}$ over edges f at $\varphi(v)$.

Since each $\phi_v(H_v)$ was assumed separable in $G_{\varphi(v)}$, there is a finite index subgroup of $G_{\varphi(v)}$ that contains $\phi_v(H_v)$ but no elements of X_v. Let K_v be isomorphic to this subgroup, contain H_v, and extend ϕ_v to K_v so its image is this subgroup.

This remains an immersion: the vertex maps are still injective, so we have to check the local coset maps $\Phi_{v/f}$. This amounts to checking that for each edge f at $\varphi(v)$ the elements δ_e with e in the preimage of f represent different right cosets of $\phi_v(K_v)$, or equivalently that $\delta_{e_i}^{-1}\delta_{e_j}$ are outside $\phi_v(K_v)$ for all pairs e_i, e_j of edges.

But this is exactly what we ensured by requiring $\phi_v(K_v)$ to exclude X_v, so this condition is satisfied, and Φ remains an immersion.

Now we are able to apply Theorem 4.1: we have an immersion where the vertex groups correspond to finite index subgroups. This immersion can be completed to a cover corresponding to a finite index subgroup. Since the procedure to do this does not identify any edges or vertices, we can recover $\mathcal{H}$ (and the original immersion) by restricting to a subgraph and subgroups of the vertex groups. □

Remark 4.3. Note that the hypothesis that each vertex group embedding is separable was stronger than needed for the conclusion: all that was used was the fact that each $\phi_v(H_v)$ was contained in a finite index subgroup excluding X_v in order to keep the cosets separate. This condition is necessary as well as sufficient, since otherwise there is no way to enlarge H_v to a finite index subgroup where the δ_{e_i} represent different cosets.

Now we have the tools to prove the main theorem,

Main Theorem. *Suppose G is a finite free product of subgroup separable groups. Then G is itself subgroup separable.*

Proof. The proof is essentially an application of Theorem 4.2 (and therefore also of Theorem 4.1), but a little care is needed in the set up to be sure that we can exclude any element.

Let $\mathcal{G}$ be a finite graph of groups (with trivial edge groups) representing G as $\pi_1(\mathcal{G}, v_0)$, and T be the Bass–Serre tree for $\mathcal{G}$. Let v be the vertex in T represented by G_{v_0} – the "preferred lift" of v_0 to T.

Let H be any finitely generated subgroup of G and let g be any element of G outside of H. Define a subtree T_1 of T as the smallest subtree which

- is H-invariant (that is, contains T_H);
- contains v;
- contains vg.

This subtree is the union of T_H and the orbits of paths from v and vg to T_H; since these paths are finite the quotient T_1/H is again finite.

Let u_0 be the vertex of $\mathcal{H}$ corresponding to the H-orbit of v: this will be our basepoint for $\mathcal{H}$. Let $\mathcal{H}$ be the quotient graph of groups obtained from the action of H on T_1, choosing the subtree S^v to contain v.

The group and graph inclusions $H \to G$ and $T_1 \to T$ induce an immersion $\Phi : \mathcal{H} \to \mathcal{G}$, as described in Proposition 2.10. This realises the embedding of H into G as $\Phi_{u_0} : \pi_1(\mathcal{H}, u_0) \to \pi_1(\mathcal{G}, v_0)$; careful choice of subtrees avoids any conjuations appearing in this map, so a cyclically reduced loop at u_0 will be mapped to a cyclically reduced loop at v_0.

Since H is finitely generated, it also has finite Kurosh rank as this is bounded above by the true rank. So Proposition 3.3 implies that $\mathcal{H}$ is a finite graph of groups, since in constructing it we used a tree which differed from T_H only by the addition of the orbits of two finite paths. Notice that since the Kurosh rank must not change when we add these paths, there are no non-trivial vertex stabilisers: the "interest" will be captured by the elements δ_e of the immersion Φ.

The path p from vg to v in T and T_1 is the lift of g as a loop in $\mathcal{G}$, and quotients to a path beginning at u_0 in $\mathcal{H}$. Since g is not an element of H, either this path is not a loop, or it is a loop but the final group element s at G_{v_0} is not the image of an element in H_{u_0}.

In the first case, we can apply Theorem 4.2 directly. Since G was a free product of subgroup separable groups, all the $\phi_v(H_v)$ are separable in $G_{\varphi(v)}$. Since no identifications are made between edges or vertices in this process, the path p will still not be a loop in $\mathcal{K}$, and therefore $\pi_1(\mathcal{K}, u_0)$ will not be an element of K.

If, on the other hand, p is a loop, we must slightly adapt the proof. When the sets X_v are constructed (containing the elements that must be excluded from each K_v), we simply add s to X_{u_0}. Then s will still not be the image of an element of K_{u_0}, so g is not an element of K. □

References

[1] Yago Antolín, Armando Martino, and Inga Schwabrow. Kurosh rank of intersections of subgroups of free products of right-orderable groups. *Math. Res. Lett.*, 21(4):649–661, 2014.

[2] Hyman Bass. Covering theory for graphs of groups. *J. Pure Appl. Algebra*, 89(1-2):3–47, 1993.

[3] R. G. Burns. A note on free groups. *Proc. Amer. Math. Soc.*, 23:14–17, 1969.

[4] R. G. Burns. On finitely generated subgroups of free products. *J. Austral. Math. Soc.*, 12:358–364, 1971.

[5] Rita Gitik. Graphs and separability properties of groups. *J. Algebra*, 188(1):125–143, 1997.

[6] I. Gruschko. Über die Basen eines freien Produktes von Gruppen. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):169–182, 1940.

[7] Marshall Hall, Jr. Coset representations in free groups. *Trans. Amer. Math. Soc.*, 67:421–432, 1949.

[8] Alexander Kurosch. Die Untergruppen der freien Produkte von beliebigen Gruppen. *Math. Ann.*, 109(1):647–660, 1934.

[9] N. S. Romanovskiĭ. On the residual finiteness of free products with respect to subgroups. *Izv. Akad. Nauk SSSR Ser. Mat.*, 33:1324–1329, 1969.

[10] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[11] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[12] Henry Wilton. Elementarily free groups are subgroup separable. *Proc. Lond. Math. Soc. (3)*, 95(2):473–496, 2007.

Paper 2: **Serre's Property (FA) for automorphism groups of free products**

Naomi Andrew

ABSTRACT. We provide some necessary and some sufficient conditions for the automorphism group of a free product of (freely indecomposable, not infinite cyclic) groups to have Property (FA). The additional sufficient conditions are all met by finite groups, and so this case is fully characterised. Therefore this paper generalises the work of Leder in [18] for finite cyclic groups, as well as resolving the open case of that paper.

1 Introduction

Serre introduced Property (FA) in [22] as a 'near opposite' to a group splitting as a free product with amalgamation or an HNN extension. A group G has Property (FA) if every action of G on a tree has a fixed point.

Serre proves ([22, Theorem 15]) that Property (FA) is equivalent to the following conditions

(1) G is not a (non-trivial) amalgamated free product

(2) G has no quotient isomorphic to $\mathbb{Z}$

(3) G is not the union of a strictly increasing sequence of subgroups.

If G is countable, then the third condition is equivalent to finite generation; there are uncountable groups satisfying Property (FA) ([17]). Examples of groups with Property (FA) include finitely generated torsion groups and $\mathrm{SL}(n, \mathbb{Z})$ for $n \geq 3$ (both due to Serre in [22]); $\mathrm{Aut}(F_n)$ for $n \geq 3$ (due to Bogopolski in [4], with an alternative proof in [8]) and the automorphism group of a free product of at least four copies of $\mathbb{Z}/n\mathbb{Z}$ (due to Leder in [18]).

In fact, Leder shows (in most cases) that for free products of finite cyclic groups, whether the automorphism group has Property (FA) depends only on the number of times each isomorphism class appears. Our results give the following generalisation and completion of Leder's work:

Corollary 1.1. *Let G be a free product of finite groups. Then $\mathrm{Aut}(G)$ has Property (FA) if and only if all but possibly one factor appear at least four times (up to isomorphism), and the remaining factor (if present) appears only once.*

This is a consequence of our main results which are, in the positive direction:

Theorem 3.1. *Let G be a finite free product of groups satisfying one of the following two conditions.*

1. *Each free factor satisfies:*
   (a) *its isomorphism class appears at least four times in the decomposition,*
   (b) *it has Property (FA), and*
   (c) *its automorphism group has finite abelianisation and cannot be expressed as the union of a properly increasing sequence of subgroups;*
2. *There is a free factor appearing exactly once that has Property (FA) and its automorphism group has Property (FA), and all other free factors are as in (1).*

Then $\mathrm{Aut}(G)$ *has Property (FA).*

And in the opposite direction, we have:

Theorem 4.1. *Let G be a free product of freely indecomposable groups, with no infinite cyclic factors. Suppose G satisfies any of the following three conditions:*

1. *At least one free factor appears exactly two or three times, or at least two free factors appear exactly once.*
2. *The automorphism group of any factor appearing exactly once does not have Property (FA).*
3. *The automorphism group of any factor appearing more than once does not have finite abelianisation or can be expressed as a union of a properly increasing sequence of subgroups.*

Then $\mathrm{Aut}(G)$ *does not have Property (FA).*

These imply Corollary 1.1 since finite groups and their automorphism groups have Property (FA), and so the extra conditions of Theorem 3.1 are always satisfied.

Comparing Theorems 3.1 and 4.1, most of the sufficient conditions in Theorem 3.1 are also necessary. The exception is the requirement that each factor has Property (FA): since the structure of the automorphism group places significant restrictions on the possible trees it could act on, it seems plausible that there are examples of groups that act on trees but not in a way that extends to the automorphism group of their free product.

The cases with infinite cyclic factors are in general still open although some cases of Theorem 4.1 go through allowing free rank 1 or 2, and (as observed above) in the opposite direction $\mathrm{Aut}(F_n)$ has Property (FA) for $n \geq 3$.

All of the groups considered have finite index subgroups that do act on trees, which will be shown as Proposition 4.12, and so we obtain

Corollary 4.13. *Suppose G is a finite and non-trivial free product where each factor is freely indecomposable and not infinite cyclic. Then* $\mathrm{Aut}(G)$ *does not have Kazhdan's Property (T).*

This is in contrast to the situation for free groups, where it has recently been shown that $\mathrm{Aut}(F_n)$ has Property (T) for n at least 4. (See [16, 15, 21].)

Remark 1.2. In view of Remark 1.10 of [6], Theorem 3.1(1) is true for Property (F$\mathbb{R}$), as is Theorem 3.1(2) with the extra hypothesis that the free factor appearing once only is finitely generated.

Acknowledgements. I am grateful to my supervisor, Armando Martino, for all his guidance and encouragement. I am also grateful to Ric Wade for pointing out the observation of Proposition 4.12, as well as to Ashot Minasyan for helpful comments on this manuscript. I would also like to thank the anonymous reviewer for their helpful comments and suggestions.

2 Background

2.1 Actions on trees

First, we collect some lemmas about trees, subtrees and fixed point sets of elliptic subgroups of groups acting on trees, that are needed at various points in the later arguments. Many of the statements and proofs hold for both real and simplicial trees, but unless otherwise specified, all trees are simplicial trees equipped with the edge-path metric. All actions are on the right.

Lemma 2.1. *Let X_i be a family of subtrees of a tree T with non-empty intersection, and let Y be another subtree. Suppose that for each i, $X_i \cap Y$ is non-empty. Then $(\bigcap X_i) \cap Y$ is also non-empty.*

Proof. Let v be a nearest point in Y to $\bigcap X_i$. Then for each i, X_i contains v, since X_i includes both $\bigcap X_i$ and part of Y. So $v \in \bigcap X_i$, and since it was in Y by definition it is in $(\bigcap X_i) \cap Y$). □

In the finite case, but not in general, we may weaken the hypotheses to give the following lemma. (In fact, it can be proved by using Lemma 2.1 as an induction step.)

Lemma 2.2 (Serre, Lemma 10 of [22]). *Let $X_1, \ldots, X_m$ be subtrees of a tree T. If the X_i meet pairwise, then their intersection is non-empty.*

Definition 2.3. Let G act on a tree T. A subgroup $H \leq G$ is *elliptic* if there is a point $v \in T$ such that $vh = h$ for all elements $h \in H$. In this case, the *fixed point set* for H (denoted $\text{Fix}(H)$) consists of all such points, and forms a subtree of T.

Lemma 2.4. *Suppose H and K are subgroups of a group G acting on a tree. If H and K are elliptic and every element of H commutes with every element of K, then the subgroup they generate is elliptic.*

Proof. Consider some point v in $\text{Fix}(H)$. Since $vkh = vhk = vk$, for all h and k, the point vk is also in $\text{Fix}(H)$. So the geodesic $[v, vk]$ is contained in $\text{Fix}(H)$. Since k is elliptic, the midpoint of this geodesic is fixed by k which puts it in the intersection $\text{Fix}(H) \cap \text{Fix}(k)$ which must be non-empty. Since $\text{Fix}(K)$ is the non-empty intersection of all the $\text{Fix}(k)$, the subtrees $\text{Fix}(k)$ and $\text{Fix}(H)$ satisfy Lemma 2.1, and so $\text{Fix}(H) \cap \text{Fix}(K)$ is non-empty. □

Note that this also gives that the direct product of two groups with Property (FA) itself has Property (FA). The converse is also true, since factors are quotients so if either factor has an action on a tree the direct product will.

Lemma 2.5.

(1) *Suppose a tree has subtrees S_1, S_2, T_1, T_2 such that S_1 has non-empty intersection with T_1 and T_2, and S_2 has non-empty intersection with T_1 and T_2. Then S_1 and S_2 have non-empty intersection, or T_1 and T_2 have non-empty intersection.*

(2) *Suppose a group G acts on a tree, and has subgroups H_1 and H_2 which are elliptic, and an element g such that H_1 has common fixed points with H_1^g and H_2^g, and H_2 has common fixed points with H_1^g and H_2^g. Then H_1 and H_2 have a common fixed point.*

Proof.

(1) Suppose S_1 and S_2 do not intersect. Consider the bridge joining S_1 and S_2. Since T_1 has non-empty intersection with both these subtrees, T_1 contains this bridge. Similarly, T_2 contains this bridge. So $T_1 \cap T_2$ contains the bridge, and must be non-empty.

(2) The fixed point subtrees of the four subgroups satisfy the conditions of part (1), so either H_1 and H_2 or H_1^g and H_2^g have a common fixed point. But since $\text{Fix}(H_1^g) \cap \text{Fix}(H_2^g) = (\text{Fix}(H_1) \cap \text{Fix}(H_2))g$ if one is non-empty both are. So in fact both are non-empty and so H_1 and H_2 have a common fixed point. □

2.2 Automorphisms of free products

Presentations of the automorphism group of a free product were found by Fouxe-Rabinovitch in [11] and [12] and later by Gilbert in [14]. (Gilbert's is a finite presentation, under some reasonable finiteness assumptions on the factor groups and their automorphisms.) They assume that the free product is given as a Grushko decomposition:

Theorem 2.6 (Grushko decomposition)**.** *Any finitely generated group G can be decomposed as a free product $G = G_1 * \dots G_k * F_r$, where the G_i are non-trivial, freely indecomposable and not infinite cyclic, and F_r is a free group of rank r. Further, the G_i are unique up to conjugacy, and the rank of F_r is unique.*

Remark 2.7. This decomposition theorem is well known; see for example [23]. As Stallings notes, the existence of such a decomposition (for a finitely generated group) is provided by Grushko's theorem, and its uniqueness by the Kurosh subgroup theorem. Proofs of these theorems can be found in many textbooks, for example [20].

They distinguish three kinds of automorphism, which generate the whole automorphism group. Fixing a basis X for F_r (which gives a decomposition of F_r as a free product of infinite cyclic groups), they are:

- Factor automorphisms, which are automorphisms of just one free factor and do not affect the rest (including replacing a single element of X with its inverse);
- Permutation automorphisms, which permute isomorphic free factors according to a fixed, compatible set of isomorphisms (including permuting the elements of X);
- Whitehead automorphisms, of two kinds:
  - partial conjugations sending a free factor G_i to G_i^a or an element x of X to x^a
  - for an element x of X, transvections sending x to ax.

  In the case of a partial conjugation of G_i, a must be drawn from some G_j, $j \neq i$ or from X. In the case of a partial conjugation or transvection of an element of X, a must be drawn from some G_i or from $X \setminus x$.

We will refer to the following subgroups of $\mathrm{Aut}(G)$: $\mathrm{Fact}(G)$, which is generated by the factor automorphisms; $\mathrm{Perm}(G)$, which is generated by the permutation automorphisms, and $\mathrm{FR}(G)$, which is generated by partial conjugations. (Usually, we only consider $\mathrm{FR}(G)$ in the case where G does not have any infinite cyclic factors.)

We denote the first kind of Whitehead automorphism by (A, b), indicating that the free factor A is to be conjugated by the element b. In addition, we use (A, B) to denote the subgroup consisting of the partial conjugations (A, b) for all elements b of some other free factor B.

Note that an *inner factor automorphism*, conjugating a free factor by one of its own elements, is not a Whitehead automorphism because of the requirement that a is drawn from a different factor. It lies only in the first class.

Gilbert gives the following characterisation of $\mathrm{Fact}(G)$ and $\mathrm{Perm}(G)$. We restrict to the case with no infinite cyclic factors, since we do not require the general case.

Proposition 2.8 (Gilbert, Proposition 3.1 of [14]). *Let $G = G_1 * \cdots * G_k$ be a free product of freely indecomposable and not infinite cyclic groups. After reordering if necessary, suppose $G_1, \ldots, G_d$ are representatives of all the distinct isomorphism classes of the G_i, and that the isomorphism class represented by G_i occurs n_i times. Then:*

(1) $\mathrm{Fact}(G) \cong \prod_{i=1}^{d} \mathrm{Aut}(G_i)$

(2) $\mathrm{Perm}(G) \cong \prod_{i=1}^{d} S_{n_i}$

(3) $\langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle \cong \prod_{i=1}^{d} (\mathrm{Aut}(G_i) \wr S_{n_i})$

(Here S_n is the symmetric group on n elements, and the wreath products are permutation wreath products on a set of n elements, not $n!$.)

The subgroup $\mathrm{FR}(G)$ has the following presentation given explicitly as Proposition 3.1 of [5] (although it can be deduced from [11] and [14]):

Proposition 2.9. *Suppose G is a free product of freely indecomposable groups with no infinite cyclic factors. Then the subgroup $\mathrm{FR}(G)$ of $\mathrm{Aut}(G)$ is generated by the partial conjugations (A, b) subject to the relations:*

$$(A, b)(A, b') = (A, b'b) \qquad \textit{for } b, b' \in B \tag{1}$$

$$(A, b)(C, d) = (C, d)(A, b) \qquad \textit{for } A \neq C, b \notin C, d \notin A \tag{2}$$

$$[(A, b)(C, b), (A, c)] = 1 \qquad \textit{for } A, B, C \textit{ all different}, b \in B, c \in C \tag{3}$$

If each factor is finitely generated or presented, the same is true of $\mathrm{FR}(G)$, by rewriting each generator (A, b) in terms of a finite generating set for the subgroup $B \ni b$, and eliminating all unnecessary relators.

Finally, a presentation for $\mathrm{Aut}(G)$ is found by adding a set of generators and relations for $\langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$ together with the relations $\varphi^{-1}(A, b)\varphi = (A\varphi, b\varphi)$ for each $\varphi \in \langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$.

Since inner factor automorphisms are not in $\mathrm{FR}(G)$, it intersects $\langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$ trivially. So together with the final relation above giving that it is normal, we have that $\mathrm{Aut}(G) = \mathrm{FR}(G) \rtimes \langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$.

2.3 Bass-Serre theory

We restate some of the definitions and results of Bass-Serre theory, making sure the notation lines up with this paper. In particular, the action will be on the right. (This is closest to the exposition by Bass [2], but other expositions can be found in [22] and [9].)

Definition 2.10. A *graph of groups* $\mathcal{G}$ consists of a graph Γ together with groups G_v for every vertex and $G_e = G_{\bar{e}}$ for every (oriented) edge, and monomorphisms $\alpha_e : G_e \to G_{\tau(e)}$ for every (oriented) edge.

(Here, the graph Γ should be understood as it is defined by Serre [22], with edges in oriented pairs indicated by $\bar{e}$, and maps $\iota(e)$ and $\tau(e)$ from each edge to its initial and terminal vertices.)

The fundamental group of a graph of groups can be defined in two ways, with respect to a maximum tree of the graph, and by considering loops in the graph of groups. We take the second route, which simplifies some subsequent calculations.

Definition 2.11 (Paths)**.** Let $F(\mathcal{G})$ be the group generated by all the vertex groups and all the edges of $\mathcal{G}$, subject to relations $e\alpha_e(g)\bar{e} = \alpha_{\bar{e}}(g)$ for $g \in G_e$. Note that taking $g = 1$ this gives that $e^{-1} = \bar{e}$, as expected.

Define a *path* (of length n) in $F(\mathcal{G})$ to be a sequence $g_0e_1g_1 \ldots e_ng_n$, where each e_i has $\iota(e_i) = v_{i-1}$ and $\tau(e_i) = v_i$ for some vertices v_i (so there is a path in the graph), and each $g_i \in G_{v_i}$. A *loop* is a path where $v_0 = v_n$.

Some authors prefer to treat the set of all paths in $F(\mathcal{G})$ as a *groupoid*, denoted $\pi(\mathcal{G})$ in [2], and called the *fundamental groupoid*; the multiplication is inherited from $F(\mathcal{G})$, with the exception that some products are not defined. In this sense, the product of two paths is defined precisely when their endpoints match.

Definition 2.12 (Reduced paths)**.** A path is *reduced* if it contains no subpath of the form $e\alpha_e(g)\bar{e}$ (for $g \in G_e$). A loop is *cyclically reduced* if, in addition to being reduced, $e_n(g_ng_0)e_1$ is not of the form $e\alpha_e(g)\bar{e}$.

Every path is equivalent (by the relations for $F(\mathcal{G})$) to a reduced path, and similarly every loop is equivalent to both a reduced loop and a cyclically reduced loop. In general these reduced representations are not unique, although all equivalent (cyclically) reduced paths (or loops) will have the same edge structure. Note that a cyclically reduced loop might not be at the same vertex as the original loop.

Definition 2.13. The *fundamental group of* $\mathcal{G}$ (at a vertex v) is the set of loops in $F(\mathcal{G})$ at v, and is denoted $\pi_1(\mathcal{G}, v)$. The multiplication is that of the fundamental groupoid, restricted to these loops. (Equivalently, this is just the multiplication from $F(\mathcal{G})$, since in the groupoid the product of two loops at the same vertex is always defined.)

The isomorphism class of this group does not depend on the vertex chosen. (In fact, the two groups obtained by choosing different base vertices are conjugate in the groupoid.)

We take the corresponding definition of the Bass-Serre tree:

Definition 2.14 (Bass-Serre Tree)**.** Let T be the graph formed as follows: the vertex set consists of 'cosets' $G_w p$, where p is a path in $F(\mathcal{G})$ from w to v. There is an edge(-pair) joining two vertices $G_{w_1} p_1$ and $G_{w_2} p_2$ if $p_1 = e g_{w_2} p_2$ or $p_2 = e g_{w_1} p_1$ (where $g_w \in G_w$).

The graph T is a tree, usually called the Bass-Serre tree (or universal cover) for $\mathcal{G}$. Since loops at v both start and finish at v, $\pi_1(\mathcal{G}, v)$ acts on the right on the set of vertices, preserving adjacency.

Definition 2.15 (Quotient graph of groups)**.** Given a group G acting on a tree T, there is a *quotient graph of groups* formed by taking the quotient graph from the action and assigning edge and vertex groups as the stabilisers of a representative of each orbit. Edge monomorphisms are then the inclusions, after conjugating appropriately if incompatible representatives were chosen.

Theorem 2.16 (Structure theorem)**.** *Up to isomorphism of the structures concerned, the processes of constructing the quotient graph of groups, and of constructing the fundamental group and Bass-Serre tree are mutually inverse.*

2.4 Translation length

The results in Section 4 require some calculations involving the translation length function for an action on a tree. This function was investigated in [7]; Section 1 of that paper proves many of its basic properties.

Definition 2.17 (Translation length function)**.** For a group G acting on an ($\mathbb{R}$-)tree T the *translation length function* is $\|-\| : G \to \mathbb{R}$ with $\|g\| = \inf_{x \in T} d(x, xg)$.

If g stabilises a point, then $\|g\| = 0$, and if g is a hyperbolic element $\|g\|$ is the distance between a point on the axis and its image. Translation length is invariant under conjugation (that is, $\|h^{-1}gh\| = \|g\|$). Also, if T is a simplicial tree (with edge lengths equal to 1), then the translation length function takes only integer values.

For the action of the fundamental group of a graph of groups on its Bass-Serre tree, using the definitions above, the translation length function is easy to calculate:

Proposition 2.18. *Let $\mathcal{G}$ be a graph of groups, with fundamental group G, acting on its Bass-Serre tree T. For each element $g \in G$, the translation length $\|g\|$ is the path length of g after cyclic reduction.*

3 Sufficient conditions

In this section we prove the following sufficient conditions for the automorphisms of a free product to have Property (FA):

Theorem 3.1. *Let G be a finite free product of groups satisfying one of the following two conditions.*

(1) Each free factor satisfies:

(a) its isomorphism class appears at least four times in the decomposition,

(b) it has Property (FA), and

(c) its automorphism group has finite abelianisation and cannot be expressed as the union of a properly increasing sequence of subgroups;

(2) There is a free factor appearing exactly once that has Property (FA) and its automorphism group has Property (FA), and all other free factors are as in (1).

Then $\mathrm{Aut}(G)$ *has Property (FA).*

Since these conditions require each factor to have Property (FA), they are certainly freely indecomposable and not infinite cyclic. So their automorphism group decomposes as $\mathrm{Aut}(G) = \mathrm{FR}(G) \rtimes \langle \mathrm{Fact}(G), \mathrm{Perm}(G)\rangle$ as described in Proposition 2.9. First we will show that the quotient $\langle \mathrm{Fact}(G), \mathrm{Perm}(G)\rangle$ has Property (FA). In [6] Cornulier and Kar characterise the permutational wreath products with Property (FA). Their result is:

Theorem 3.2 (Theorem 1.1 of [6])**.** *Let G be a group that is a permutational wreath product $G = A \wr_X B$ where $A \neq 1, X \neq \emptyset$ and X has finitely many B-orbits each with more than one element. Then G has Property (FA) if and only if B has Property (FA) and A has finite abelianisation and cannot be expressed as the union of a properly increasing sequence of subgroups.*

Since Proposition 2.8 gives us a decomposition of $\langle \mathrm{Fact}(G), \mathrm{Perm}(G)\rangle$ as a direct product of permutational wreath products, we may use this result to investigate this subgroup.

Proposition 3.3. *Letting G be as in Theorem 3.1, the subgroup generated by factor and permutation automorphisms has Property (FA).*

Proof. By Proposition 2.8 this is a direct product of permutation wreath products. If G satisfies part (1) of Theorem 3.1 then each of them satisfies the hypotheses of Theorem 3.2: since each $n_i \geq 4$ the set X is non empty; S_{n_i} acts transitively on it so there is only one orbit; and from the hypotheses of Theorem 3.1 these $\mathrm{Aut}(G_i)$ have finite abelianisation and cannot be expressed as the union of a properly increasing sequence of subgroups, and S_{n_i} is finite so has Property (FA). So each wreath product has Property (FA). If G satisfies part (2) of Theorem 3.1, then the automorphism group of the singleton factor has Property (FA) by assumption, and all others satisfy the hypotheses we need for Theorem 3.2 just as above. So, in either case, we have a direct product of groups with Property (FA). Their direct product must also have Property (FA), by inductively applying the argument of Lemma 2.4. □

Next we show that, whenever $\mathrm{Aut}(G)$ acts on a tree, the subgroup $\mathrm{FR}(G)$ has a fixed point. Most of the arguments are similar, and proceed by finding 'enough commutation' that various elliptic subgroups are forced to have common fixed points, but we write them out in full.

Proposition 3.4. *Let G be as in part (1) of Theorem 3.1. Then any action of* $\mathrm{FR}(G)$ *on a* $\mathbb{Z}$*-tree which extends to an action of* $\mathrm{FR}(G) \rtimes \mathrm{Perm}(G)$ *on the same tree has a global fixed point.*

Proof. The subgroup $\mathrm{FR}(G)$ is generated by finitely many subgroups (A, B). (Recall that this consists of all partial conjugations (A, b), where A is fixed, but b ranges over all of some other factor B.) This is isomorphic (in fact, anti-isomorphic) to B, and so since B has Property (FA), all such subgroups are elliptic. By Lemma 2.2, if their fixed point subtrees intersect pairwise then their intersection is non-empty.

So we check all the possible pairs (A, B) and (C, D): (Different letters always represent different subgroups; some combinations cannot occur due to the fact the inner factor automorphisms are excluded.)

(1) (A, B) and (C, D): These commute (by Relation (2)), and so since they are elliptic there must be a common fixed point by Lemma 2.4.

(2) (A, B) and (C, B): These subgroups commute by Relation (2). Since both are elliptic, they must have a common fixed point by Lemma 2.4.

(3) (A, B) and (A, D): Since there are (at least) four isomorphic copies of each factor group, there is some C' (different to A, B, D) such that $C' \cong A$. Letting τ be the permutation interchanging A and C', then $(A, B), (A, D)$ and τ satisfy the conditions of Lemma 2.5(2): $(A, B)^\tau = (C', B)$ and $(A, D)^\tau = (C', D)$ both commute with both (A, B) and (A, D) (by Relation (2)) and so have common fixed points by Lemma 2.4. So (A, B) and (A, D) have a common fixed point.

(4) (A, B) and (B, D) (and, by symmetry (A, B) and (C, A)): This time take $C' \cong B$ so that A, B, C', D are all different. Let τ swap C' and B, so conjugating by τ gives (A, C') and (C', D). Now (C', D) commutes with both the original elements, and (A, C') commutes with (C, D) (all by relation (2)), and so there are fixed points in common by Lemma 2.4. Also, (A, B) and (A, C') fit the hypotheses of Case (3),

and so they have common fixed point. So $(A, B), (B, D)$ and τ satisfy Lemma 2.5 and there is a common fixed point.

(5) (A, B) and (B, A): Take $C' \cong B$ and $D' \cong A$, so that A, B, C', D' are different factors. Then let τ swap C' with B, and D' with A. The images after conjugating by τ (which are (D', C') and (C', D') respectively) commute (by Relation (2)) and so have common fixed points (by Lemma 2.4) with (A, B) and (B, A), and so (A, B), (B, A) and τ satisfy Lemma 2.5 so these subgroups have a fixed point.

These pairwise intersections satisfy Lemma 2.2, so have a non-empty intersection. This is fixed by every element of $\mathrm{FR}(G)$, and so since these subgroups generate, this intersection is fixed by $\mathrm{FR}(G)$ which must itself be elliptic. □

Before proving the second case, we cover one aspect of the proof in a lemma.

We use H^{*n} to denote the free product $H_1 * \cdots * H_n$ of n copies of the group H. By analogy with a wreath product, we make the following definition.

Definition 3.5. Let H and K be groups, and equip K with an action on a set X. The *wreathed free product* of H and K (with respect to the given action) is the semidirect product $H^{*|X|} \rtimes K$, where the action of K on $H^{*|X|}$ is to permute the free factors according to the action on X.

The symmetries induced by the K-action have the effect of restricting the trees such groups can act on, as we see in the following lemma (restricting to the action of S_n on a set of n elements):

Lemma 3.6. *Suppose the wreathed free product $H^{*n} \rtimes S_n$ acts on a tree T, such that the free factor H_1 fixes a subtree T_1. Then*

(1) Each factor H_i fixes a subtree T_i, and these are permuted by the action of S_n on T.

(2) There are vertices $v_i \in T_i$ such that

- *$d(v_i, v_j) = d(T_i, T_j)$ for all i, j*
- *The vertices v_i are permuted by the action of S_n on T*
- *The geodesics $[v_i, v_j]$ (for $i \neq j$) all have the same midpoint, which we denote w.*

Let T' be the convex hull of the set of all vertices v_i (see Figure 2.1).

(3) There are elements $h_i \in H_i$ such that $\mathrm{Fix}(h_i) \cap T' = v_i$. These elements can be chosen so they are permuted by the action of S_n on H^{*n}.

(4) *Suppose there is a group G, containing $H^{*n} \rtimes S_n$ (with n at least 3) as a subgroup and acting on a tree. Since the subgroup $H^{*n} \rtimes S_n$ also acts on the tree, we may choose elements h_i in accordance with (3). If there is an element g in G which commutes with h_1 and some h_j ($j \neq 2$) and satisfying $g^{-1}h_1h_2g = h_2h_1$, then the $H^{*n} \rtimes S_n$ subgroup is elliptic. (That is, T' is a single point.)*

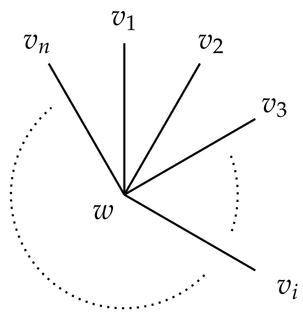

FIGURE 2.1: The graph T' described in Lemma 3.6(2)

Proof. Write the elements of S_n in disjoint cycle notation, but we use the notation σ with a subscript to denote the corresponding isometry of T. That is, $\sigma_{(12)}$ denotes the isometry of T induced by (12).

(1) Since $(1i)^{-1}H_1(1i) = H_i$, we must have that H_i fixes precisely $T_i := T_1\sigma_{(1i)}$, which in particular is non empty.

(2) If $n = 1$ then we may choose any global fixed point for the H-action. If any (and therefore every) pair has a common fixed point, in fact the H^{*n} action must be elliptic. Since the S_n action permutes the T_i, it acts on their (non-empty) intersection. Since it is finite, it does so with a fixed point, which will be a fixed point for the whole action. Let every v_i be this fixed point; then this one point subtree works.

Otherwise, since S_n acts 2-transitively on the set $\{T_i\}$, and is acting by isometries, we have that $d(T_i, T_j) = \lambda(1 - \delta_{ij})$ where λ is a positive constant and δ_{ij} is the Kronecker delta. Let v_{ij} (with $i \neq j$) be the nearest point in T_i to T_j. In fact this is the same point as j varies, since if there were j, k such that v_{ij} and v_{ik} were different, we would have

$$d(T_j, T_k) = d(v_{ji}, v_{ij}) + d(v_{ij}, v_{ik}) + d(v_{ik}, v_{ki}) > d(T_i, T_j) + d(T_i, T_k).$$

Since all three distances are equal, this is not the case. Call this common nearest point v_i; then we have that $d(v_i, v_j) = d(T_i, T_j)$ as we wanted. Since the action is by isometries, $d(T_1, T_j) = d(v_1, T_j) = d(v_1\sigma_{(1i)}, T_j\sigma_{(1i)}) = d(v_1\sigma_{(1i)}, T_{j(1i)})$. As explained above, $d(T_1, T_j) = d(T_i, T_{j(1i)})$, and so since $v_1\sigma_{(1i)}$ lies in T_i it is the nearest point to $T_{j(1i)}$. So $v_1\sigma_{(1i)} = v_i$, and they are permuted by the S_n action.

If $n = 2$ then since $\sigma_{(12)}$ must swap T_1 and T_2, it will invert the geodesic. So (after subdividing if necessary) there must be a fixed vertex w at the same distance from v_1 and v_2.

For $n \geq 3$ consider the vertices v_i, v_j, v_k, and their y-point, w. This tripod contains three geodesics which must all be the same length, and w is therefore the midpoint of the geodesic joining any two. As i, j, k are arbitrary, the point w is the midpoint of the geodesic joining any two of the v_i.

(3) For each $h \in H_1$ consider the intersection $\mathrm{Fix}(h) \cap T'$. This is a collection of closed sets whose intersection must be the point v_1. Since T is a $\mathbb{Z}$-tree, the diameter of each set must be an integer. But then in order for the intersection to consist only of the boundary point v_1, there is some h such that $\mathrm{Fix}(h) \cap T'$ is precisely v_1.

Choose an $h_1 \in H_1$ such that $\mathrm{Fix}(h_1) \cap T' = v_1$, and let $h_i \in H_i$ be $\sigma_{(1i)}^{-1} h_1 \sigma_{(1i)}$ for each i. By definition, the h_i are permuted by the action of S_n, and $\mathrm{Fix}(h_i) \cap T' = \mathrm{Fix}(h_1)\sigma_{(1i)} \cap T' = (\mathrm{Fix}(h_1) \cap T')\sigma_{(1i)} = v_1\sigma_{(1i)} = v_i$ as needed.

(4) Suppose not. So T' is not a single point, h_1 and h_2 have no common fixed point, and so the element h_1h_2 is hyperbolic: its axis includes the geodesic $[v_1, v_2]$, and its translation length is $2d(v_1, v_2)$. The same is true of the element h_2h_1, although it translates the other way along the common segment. Since g commutes with h_1 and h_j, it preserves both fixed point sets, and fixes the geodesic $[v_1, v_j]$ joining them. In particular, it fixes the segment $[v_1, w]$. Consider how the elements $g^{-1}h_1h_2g$ and h_2h_1 move this fixed segment: h_2h_1 must move it past v_1. But $g^{-1}h_1h_2g$ cannot move it past v_1: $g^{-1}h_1h_2$ moves it along $\mathrm{Axis}(h_1h_2)$, and then the nearest point of $\mathrm{Fix}(g)$ is closer than v_1 so the segment must stay the same side. So the equation given cannot hold, giving us a contradiction. So T' is a single point, and the subgroup $H^{*n} \rtimes S_n$ is elliptic in this action. □

Remark 3.7. This lemma is also true for $\mathbb{R}$-trees, with one change: in (3) we must assume the group H is finitely generated, and then we can restrict to a finite generating set and consider the intersection of finitely many $\mathrm{Fix}(h) \cap T'$ to argue the existence of a h_1 such that $\mathrm{Fix}(h_1) \cap T' = v_1$.

Proposition 3.8. *Let G be as in part (2) of Theorem 3.1. Then any action of* $\mathrm{FR}(G)$ *on a* $\mathbb{Z}$*-tree which extends to an action of* $\mathrm{FR}(G) \rtimes \mathrm{Perm}(G)$ *on the same tree has a global fixed point.*

Proof. This is the same idea as for Proposition 3.4, but depends even more on having access to symmetries required by the permutation automorphisms. Recall that the hypotheses provide one factor appearing exactly once, and the others are repeated at least four times. Throughout the proof, we denote the non-repeated factor by K.

For any action of $\mathrm{FR}(G)$ on a tree, extending as required, the following subgroups have global fixed points:

(i) The subgroup generated by all partial conjugations of repeated factors by repeated factors – this is what was proved in Proposition 3.4.

(ii) The subgroup generated by all partial conjugations where the conjugating group is the non-repeated factor K. This is a direct product of several copies of K, and so will have Property (FA) since K does.

(iii) Every subgroup of the form (K, H) where K is the non-repeating factor. We deal with these by isomorphism class of H. Let $H_1, \dots, H_n$ be the $n \geq 4$ repeated copies; then $\langle (K, H_1), \dots, (K, H_n) \rangle$ is isomorphic to H^{*n}. Let S_n be the subgroup of $\mathrm{Perm}(G)$ which permutes the H_i; then $H^{*n} \rtimes S_n$ is a subgroup of $\mathrm{FR}(G) \rtimes \mathrm{Perm}(G)$, and we only consider actions which arise as actions of this group. We will show that this subgroup is elliptic in every action of this kind.

We apply Lemma 3.6, letting $\{(K, h_i)\}$ be the elements $\{h_i\}$ provided by part (3). Consider the element (H_2, h_1). By Relation (2) it commutes with and so has common fixed points with all (K, h_i) with $i \neq 2$. The final commutation relation (3) for $\mathrm{FR}(G)$ gives that $[(K, h_1)(H_2, h_1), (K, h_2)] = 1$. Expanding (and moving some commuting elements past each other) this becomes

$$(K, h_1)(K, h_2) = (H_2, h_1)^{-1}(K, h_2)(K, h_1)(H_2, h_1).$$

Therefore part (4) of the lemma tells us that the $H^{*n} \rtimes S_n$ subgroup is elliptic. Of course, so are all its subgroups, including in particular all subgroups (K, H).

All subgroups of the form (A, B) are contained in one of these subgroups, so must themselves be elliptic. As before, we check that any pair of these subgroups have a common fixed point. Pairs drawn from the same subgroup are already done, so we check the cases where they are drawn from different subgroups. Some cases are by commuting subgroups, others rely on Lemma 2.5 and so are similar to the technique used in the previous result, and others need the use of Relation (3) of Proposition 2.9. Denote by K the factor occuring once, and by $A, B, \dots$ any of the factors that appear at least four times. As before, different letters denote different factors.

(1) (A, B) and (C, K): These commute and therefore have a common fixed point (by Relation (2) and Lemma 2.4).

(2) (A, B) and (A, K): Let τ be the permutation swapping A and some $C' \cong A$ (different to A and B). Conjugating by τ gives (C', B) and (C', K), which both commute with both original subgroups. So by Lemma 2.5 we get that our elements have a common fixed point.

(3) (A, B) and (B, K) Let $C' \cong B$, different to A and B, and let τ swap B and C'. After conjugating, both have common fixed points with the original subgroups: (A, C')

and (A, B) by case (3) of Proposition 3.4 and the rest since they commute. So we satisfy Lemma 2.5 and there is a common fixed point.

(4) (A, B) and (K, C) or (A, B) and (K, B) commute so there will be a common fixed point.

(5) (A, B) and (K, A): In this case we need τ to swap A and a $C' \cong A$, giving (C', B) and (K, C'). These both have common fixed points with both original elements, in three cases because they commute, and in the fourth because (K, A) and (K, C') are subgroups of one of the groups discussed in (iii) above. So we have the common fixed points we need to once again deploy Lemma 2.5 to give us a common fixed point.

(6) (K, A) and (B, K): Consider $[(B, a)(K, a), (B, k)] = 1$ (one of the relations (3) in Proposition 2.9). Since they commute (and are elliptic), $(B, a)(K, a)$ is elliptic. Also, (B, k) is elliptic, and so since these elements commute (and are elliptic) there is a common fixed point for $(B, a)(K, a)$ and (B, k). But this means there is a common fixed point for all three elements, and so in particular for (K, a) and (B, k). We now apply Lemma 2.1 twice: first, fix $a \in A$ and vary $k \in K$ to see that $\text{Fix}(K, a)$ and $\text{Fix}(B, K)$ have non empty intersection for all $a \in A$. Then this gives that $\text{Fix}(K, A)$ and $\text{Fix}(B, K)$ have non-empty intersection, as we wanted.

(7) (A, K) and (K, A): Let $B' \cong A$ (but different), and let τ be the permutation automorphism swaping A and B'. Again, these satisfy Lemma 2.5: the four pairs are (A, K) and (B', K) which commute; (K, A) and (K, B') which are subgroups of one of the groups discussed in (iii); (A, K) and (K, B'), and (K, A) and (B', K) which satisfy the previous case. So this final pair also have a common fixed point.

(8) (K, A) and (K, B) where $A \ncong B$. Consider $[(K, a)(B, a), (K, b)] = 1$: just as above, this gives a common fixed point for (K, a) and (K, b) and then applying Lemma 2.1 gives that $\text{Fix}(K, A)$ and $\text{Fix}(K, B)$ have non-empty intersection. □

Propositions 3.3, 3.4 and 3.8 provide the proof of Theorem 3.1, as follows:

Proof of Theorem 3.1. An action of $\text{Aut}(G)$ on a tree defines an action of $\text{FR}(G)$ on the same tree. Since $\text{FR}(G) \rtimes \text{Perm}(G) \leq \text{Aut}(G)$, this action must extend to the permutation automorphisms. So by Proposition 3.4 or 3.8 this subgroup is elliptic. Now consider $v \in \text{Fix}(\text{FR}(G))$: we have that $vhg = vg'h = vh$ for all $g \in \text{FR}(G), h \in \langle \text{Fact}(G), \text{Perm}(G) \rangle$, where $g' = hgh^{-1} \in \text{FR}(G)$. So $\langle \text{Fact}(G), \text{Perm}(G) \rangle$ acts on the fixed point set of $\text{FR}(G)$. Since it has Property (FA) by Proposition 3.3 that action will have a fixed point, which must be a fixed point for the whole action. So $\text{Aut}(G)$ also has Property (FA). □

4 Necessary conditions

The results in this section, taken together, will prove all parts of Theorem 4.1. First we deal with the (shorter) parts (2) and (3), and then afterwards part (1).

Theorem 4.1. *Let G be a free product of freely indecomposable groups, with no infinite cyclic factors. Suppose G satisfies any of the following three conditions:*

(1) At least one free factor appears exactly two or three times, or at least two free factors appear exactly once.

(2) The automorphism group of any factor appearing exactly once does not have Property (FA).

(3) The automorphism group of any factor appearing more than once does not have finite abelianisation or can be expressed as a union of a properly increasing sequence of subgroups.

Then $\mathrm{Aut}(G)$ *does not have Property (FA).*

Proposition 4.2. *Let G be a free product of freely indecomposable groups, with no infinite cyclic factors. Suppose there is some free factor H whose isomorphism class appears exactly once in the Grushko decomposition, and* $\mathrm{Aut}(H)$ *does not have Property (FA). Then* $\mathrm{Aut}(G)$ *does not have Property (FA).*

Proof. By Proposition 2.9, the group $\langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$ is a quotient of $\mathrm{Aut}(G)$. Then by part (3) of Proposition 2.8, one of the direct summands of this group is $\mathrm{Aut}(H)$. So $\mathrm{Aut}(H)$ is a quotient of $\mathrm{Aut}(G)$. Since $\mathrm{Aut}(H)$ has an action on a tree without global fixed point, the same is true of $\mathrm{Aut}(G)$. □

Proposition 4.3. *Let G be a free product of freely indecomposable groups, with no infinite cyclic factors. Suppose there is a free factor H whose isomorphism class appears at least two times in the decomposition, and* $\mathrm{Aut}(H)$ *does not have finite abelianisation or can be expressed as a union of a properly increasing sequence of subgroups. Then* $\mathrm{Aut}(G)$ *does not have Property (FA).*

Proof. Again, by Proposition 2.9, $\langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$ is a quotient of $\mathrm{Aut}(G)$. Then again using part (3) of Proposition 2.8, one of the direct summands of this group is $\mathrm{Aut}(H) \wr S_n$. By Theorem 3.2 this does not have Property (FA). So since $\mathrm{Aut}(G)$ has a quotient that does not have Property (FA), neither does $\mathrm{Aut}(G)$. □

For part (1) of Theorem 4.1, we will extend the action of G on a Bass-Serre tree to its (outer) automorphisms. It is useful to view an action of a group G by isometries on a tree T as a homomorphism $G \to \mathrm{Isom}(T)$.

We recall a theorem of Culler–Morgan (also due to Alperin–Bass) on $\mathbb{R}$-trees and translation length, starting with the relevant definitions.

Definition 4.4 ([7], page 573)**.** Consider a group, G, acting on an $\mathbb{R}$-tree, T.

(1) The action is called *irreducible* if G does not preserve a point, a line or an end of T. In particular, this means that there are two hyperbolic axes whose intersection has finite length.

(2) The action is called a *shift* if G preserves a line, and the orientation of the line.

(3) The action is called *dihedral* if G preserves a line, but not the orientation of the line.

(4) The action is called *semi-simple* if G either preserves a point, or is irreducible, or a shift or dihedral.

Theorem 4.5 (see [7, Theorem 3.7] and [1, Theorem 7.13(b)])**.** *If a group G acts minimally and irreducibly, or minimally and dihedrally, on $\mathbb{R}$-trees T_1 and T_2 with the same translation length function. Then there is a unique G-equivariant isometry $f : T_1 \to T_2$. That is, f is the unique isometry such that with $f^* : \operatorname{Isom}(T_1) \to \operatorname{Isom}(T_2)$ defined by $\varphi f^* = f^{-1}\varphi f$ the following diagram commutes.*

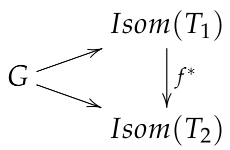

Note that the condition in [7] is that the action is minimal and semi-simple, and for uniqueness that it is not a shift. They also do not give the interpretation as a commutative diagram.

Also note that irreducibility is a property of the length function; that is, if two trees have the same length function then one is irreducible if and only if the other is.

We will consider the following subgroup of automorphisms:

Definition 4.6. Suppose G is a group acting on a tree T. Let $\operatorname{Aut}_T(G)$ be the subgroup of $\operatorname{Aut}(G)$ that preserves the translation length function of the action.

The uniqueness part of Theorem 4.5 leads to the following corollary, for which we give a full proof.

Corollary 4.7. *Given a group G acting minimally and irreducibly, or minimally and dihedrally, on an $\mathbb{R}$-tree T, then $\operatorname{Aut}_T(G)$ acts by isometries on T. Moreover, the following statements hold.*

(1) Denote by $\delta(g)$ the inner automorphism induced by g. Then $\delta(g)$ induces the same isometry as g. So if the original action of G has no fixed points, the same is true for the action of $\operatorname{Aut}_T(G)$.

(2) *The action of* $\mathrm{Aut}_T(G)$ *is compatible with the action of* G*, in the sense that the subgroup* $G \rtimes \mathrm{Aut}_T(G)$ *of the holomorph acts on* T *with the given actions of each factor.*

(3) *If the original tree was a* $\mathbb{Z}$*-tree then the action constructed is also an action on a* $\mathbb{Z}$*-tree, after subdividing if necessary to remove edge inversions.*

Groups acting on trees may have a non-trivial centre, for example in $\mathrm{SL}(2, \mathbb{Z})$ the centre has order 2. However, if the action has at least two axes, or a single axis and an elliptic element that does not preserve its orientation, then the centre of the group must be in the kernel of the action. So since two elements inducing the same inner automorphism must already have the same image in $\mathrm{Isom}(T)$, the isometry described in (1) is unique.

Proof of Corollary 4.7. Given an action $\cdot : G \to \mathrm{Isom}(T)$, and any automorphism φ of G, there is another action defined by $*_\varphi = \varphi \circ \cdot : G \to G \to \mathrm{Isom}(T)$. That is, $t *_\varphi g = t \cdot (g\varphi)$.

For any $\varphi \in \mathrm{Aut}_T(G)$, this action will have the same translation length function as $\cdot$. So we may apply Theorem 4.5 to the actions $\cdot$ and $*_\varphi$ to give a unique isometry f_φ of T (corresponding to the automorphism φ) such that the following diagram commutes:

$$\begin{array}{ccc} G & \xrightarrow{\ \cdot\ } & Isom(T) \\ \varphi \big\downarrow & \searrow^{*_\varphi} & \big\downarrow f_\varphi^* \\ G & \xrightarrow{\ \cdot\ } & Isom(T) \end{array}$$

These isometries do give an action: for $*_1 = \cdot$, the identity map on T is an equivariant isometry of T making the diagram commute. So by uniqueness, $f_1 = \mathrm{Id}$. Now for $\varphi, \psi \in \mathrm{Aut}_T(G)$, consider this diagram:

$$\begin{array}{ccc} G & \xrightarrow{\ \cdot\ } & Isom(T) \\ \varphi \big\downarrow & \searrow^{*_\varphi} & \big\downarrow f_\varphi^* \\ G & \xrightarrow{\ \cdot\ } & Isom(T) \\ \psi \big\downarrow & \searrow^{*_\psi} & \big\downarrow f_\psi^* \\ G & \xrightarrow{\ \cdot\ } & Isom(T) \end{array}$$

Here, the top square shows the equivariant isometry induced by φ, and the bottom square that by ψ. We want to consider the element $\varphi\psi$, which should also induce a unique equivariant isometry. However, from the diagram, the composition $f_\varphi f_\psi$ is just such an equivariant isometry, and so it must be the unique $f_{\varphi\psi}$.

To see (1): let g be some element of G, and consider the inner automorphism that conjugates by g (called $\delta(g)$). In this case, the usual action of g is an equivariant isometry for the conjugation, since we have that $(x \cdot g) *_{\delta(g)} h = (x \cdot g) \cdot (g^{-1}hg) = x \cdot hg = (x \cdot h) \cdot g$. As a commutative diagram (where the right hand arrow is induced by the action of g):

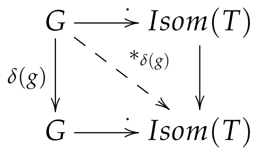

So, by uniqueness, $f_{\delta(g)} = \cdot g$.

To see (2): we need to check that the isometries corresponding to $(g\varphi)$ (as an element of G) and $\varphi^{-1}g\varphi$ are the same for all $g \in G$ and $\varphi \in \mathrm{Aut}_T(G)$. This is immediate from the commutative diagram in the statement of Theorem 4.5 with the actions $\cdot$ and $*_\varphi$, since the downwards arrow is then precisely conjugation by the isometry corresponding to φ.

To see (3): The induced isometries must send branch points to branch points (of the same valence). In a $\mathbb{Z}$-tree, since branch points are vertices and all other vertices are at integer distance, the vertex set must be preserved by the induced isometries. In the case where T is a single line, there must be a vertex stabilised by an orientation reversing element. The induced isometries must send this to another such point, and in a $\mathbb{Z}$-tree these are all vertices. Again, all other vertices are at integer distance, and so the vertex set is preserved. So the action of $\mathrm{Aut}_T(G)$ is still by an action on the $\mathbb{Z}$-tree, as we needed. □

We use this corollary to prove the first part of Theorem 4.1. First, we prove special cases where there are *only* two or three free factors (satisfying the conditions of the theorem) and then use properties of characteristic subgroups to extend these results. In the two factor case we construct an action of the full automorphism group on a Bass-Serre tree for the group; in the three factor case it is an action of the outer automorphism group.

Proposition 4.8. *Suppose $G = H * K$, with H and K freely indecomposable. Then $Aut(G)$ does not have Property (FA).*

This proposition follows from work of Forester [10], generalised by Levitt [19]; we give a proof for this case that does not require the (explicit) use of deformation spaces.

Proof. If neither H nor K are infinite cyclic, realise G as the fundamental group of the graph of groups shown in Figure 2.2a. Consider the action of G on the Bass-Serre tree for this graph of groups. We want to check that the translation length is preserved by every automorphism, which requires us first to calculate it. An elliptic element is a conjugate of an element of a factor group. None of the generating automorphisms given in Subsection 2.2 change this, and so all automorphic images of elliptic elements are themselves elliptic.

Any hyperbolic element is conjugate to a cyclically reduced word $a_1b_1a_2b_2\dots a_nb_n$, where a_1 and b_n are non-trivial. The path length of this cyclically reduced word is

(A) Graph of Groups for $H * K$

(B) Graph of Groups for $H * Z$

FIGURE 2.2: Graphs of groups realising each G in Proposition 4.8.

its length (that is, $2n$): we use Proposition 2.18, noting that for these elements the path length does not depend on the chosen base vertex. Since translation length is conjugacy invariant, it is sufficient to check invariance on words of this form. Factor automorphisms do not change the structure of our word at all, so lie in $\mathrm{Aut}_T(G)$. There is a permutation automorphism (which we write τ) if and only if H and K are isomorphic: applying this and conjugating by $a_1\tau$ we once again have a word of length $2n$ in our preferred structure, and so $\tau \in \mathrm{Aut}_T(G)$. For the partial conjugation (A, b) we have

$$a_1b_1a_2b_2 \dots a_nb_n \mapsto (b^{-1}a_1b)b_1(b^{-1}a_2b)b_2 \dots (b^{-1}a_nb)b_n = (a_1b_1'a_2b_2' \dots a_nb_n')^b$$

where each $b_i' = bb_ib^{-1}$. Once again, we have a cyclically reduced conjugate of length $2n$ in our preferred form. Therefore every partial conjugation lies in $\mathrm{Aut}_T(G)$. (The argument is identical for conjugating the other factor.)

Since all the generators preserve the translation length function, $\mathrm{Aut}_T(G) = \mathrm{Aut}(G)$. By Corollary 4.7 there is an action of the automorphism group on the Bass-Serre tree that is without global fixed points. If H and K were isomorphic, and we have a permutation automorphism, the isometry it induces will invert the edge in the fundamental domain, we need to pass to the barycentric subdivision; otherwise no subdivisions are necessary.

If exactly one factor is $\mathbb{Z}$, we can use the same technique. We write $G = H * \mathbb{Z}$, and use x for the generator of $\mathbb{Z}$. A generating set for this automorphism group consists of the partial conjugations $H \mapsto H^x$ and $x \mapsto h^{-1}xh$ for all $h \in H$, the transvections $x \mapsto hx$ for all $h \in H$, and the factor automorphisms $\mathrm{Aut}(H)$ and $x \mapsto x^{-1}$. Realise G as the fundamental group of the graph of groups in Figure 2.2b and consider the action of G on its Bass-Serre tree. Elliptic elements are in some conjugate of H, and are sent to some other conjugate of H by all of the generators. Hyperbolic elements have a cyclically reduced conjugate of the form $h_1x^{n_1} \dots h_kx^{n_k}$. By Proposition 2.18 the translation length of this element is $\sum |n_i|$. Factor automorphisms of H don't affect this; conjugating H by x is an inner automorphism so can't change the translation length. Replacing x with any of its images, after conjugating by h^{-1} if necessary to return to a cyclically reduced conjugate of the preferred form, has the same absolute exponent sum, and so the translation length is unchanged. So every element of the automorphism group is length preserving, and so it too acts on the Bass-Serre tree by Corollary 4.7. The involution $x \mapsto x^{-1}$ induces an edge inversion, so we must subdivide.

If both factor groups are infinite cyclic, then $\mathrm{Aut}(G)$ is just $\mathrm{Aut}(F_2)$, which does not have Property (FA) (by [4]; it maps onto $\mathrm{GL}_2(\mathbb{Z}) \cong D_4 *_{D_2} D_6$). □

For the case with three isomorphic factors, we will find an action of the outer automorphism group on a tree, similar to that given in Propositions 4.1 and 4.2 of [5] for three non-isomorphic factors, and in [18] for finite cyclic groups.

Proposition 4.9. *Suppose $G = A * B * C$ is a free product of three copies of some freely indecomposable and not infinite cyclic group. Let $\gamma(a)$ denote the inner factor automorphism conjugating A by a; $\varphi = (\varphi_a, \varphi_b, \varphi_c)$ denote an element of* $\mathrm{Fact}(G)$*; and $\sigma_{(12)}, \sigma_{(123)}$ the permutation automorphisms. Then* $\mathrm{Out}(G)$ *has the following presentation.*

Generators: $(A,b), (B,c), (C,a), \mathrm{Aut}(A), \mathrm{Aut}(B), \mathrm{Aut}(C), \sigma_{(123)}, \sigma_{(12)}$

Relations (with terms going over all appropriate generators):

$$(A,b)(A,b') = (A,b'b), \text{ etc.} \tag{1}$$

$$[\mathrm{Aut}(A),\mathrm{Aut}(B)] = [\mathrm{Aut}(A),\mathrm{Aut}(C)] = [\mathrm{Aut}(B),\mathrm{Aut}(C)] = 1 \tag{2}$$

$$\sigma_{(123)}^3 = 1, \sigma_{(12)}^2 = 1, (\sigma_{(123)}\sigma_{(12)})^2 = 1 \tag{3}$$

$$\varphi^{-1}(A,b)\varphi = (A,b\varphi), \text{ etc.} \tag{4}$$

$$\sigma_{(12)}^{-1}(A,b)\sigma_{(12)} = \gamma(a^{-1})(C,a^{-1}) \tag{5}$$

$$\sigma_{(12)}^{-1}(B,c)\sigma_{(12)} = \gamma(c^{-1})(B,c^{-1}) \tag{6}$$

$$\sigma_{(12)}^{-1}(C,a)\sigma_{(12)} = \gamma(b^{-1})(A,b^{-1}) \tag{7}$$

$$\sigma_{(123)}^{-1}(A,b)\sigma_{(123)} = (B,c), \text{ etc.} \tag{8}$$

$$(\varphi_a,\varphi_b,\varphi_c)^{\sigma_{(12)}} = (\varphi_b,\varphi_a,\varphi_c), \quad (\varphi_a,\varphi_b,\varphi_c)^{\sigma_{(123)}} = (\varphi_b,\varphi_c,\varphi_a) \tag{9}$$

Relations (2) are inherited from the direct product structure on $\mathrm{Fact}(G)$*; relations (3) give a presentation of S_3, and relations (9) are a consequence of the wreath product structure on* $\langle \mathrm{Fact}(G), \mathrm{Perm}(G)\rangle$.

A proof of this presentation is given as an appendix, since it closely follows the proof in [5] for three non-isomorphic groups.

This presentation gives a semidirect product decomposition of $\mathrm{Out}(G)$ as $(\hat{G} \rtimes \mathrm{Fact}(G)) \rtimes \mathrm{Perm}(G)$, where $\hat{G}$ is isomorphic to G but generated by the (A,b),(B,c) and (C,a) (denote these factor groups by $\hat{B}$,$\hat{C}$ and $\hat{A}$ respectively) and the actions are the actions from the original semidirect decomposition of $\mathrm{Aut}(G)$. However, whenever the factors are not abelian, the order of evaluation is now important, since $\mathrm{Perm}(G)$ does not normalise $\hat{G}$ in the presence of inner factor automorphisms.

Proposition 4.10. *If $G = A * B * C$ is a free product of three copies of some freely indecomposable group, then* $\mathrm{Out}(G)$*, and therefore* $\mathrm{Aut}(G)$*, acts on a tree without global fixed points.*

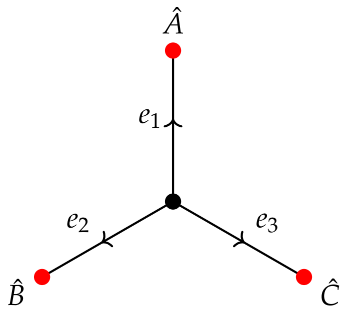

(A) $\hat{G} \leq \mathrm{Out}(G)$

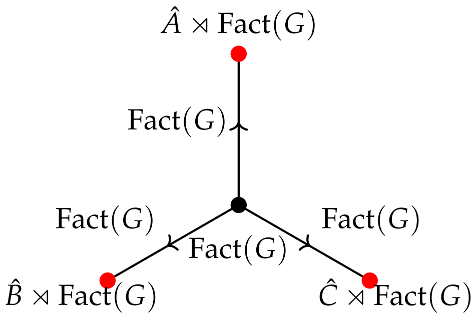

(B) $\hat{G} \rtimes \mathrm{Fact}(G)$

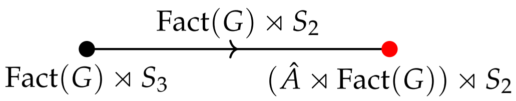

(C) $(\hat{G} \rtimes \mathrm{Fact}(G)) \rtimes S_3$

FIGURE 2.3: The graphs of groups at each stage of Proposition 4.10

Proof. We will construct an action at each stage of the semidirect product decomposition.

Consider the tripod graph of groups for $\hat{G}$ (shown in Figure 2.3a), taking the central vertex to be the base point. Call the Bass-Serre tree for this graph of groups T. Any elliptic word can be cyclically reduced to a single letter - a path of zero length (as expected). The translation length of any hyperbolic word is twice the length of a cyclically reduced conjugate, since every letter will require traversing two edges. The factor automorphisms act by sending (A, b) to $(A, b\varphi)$ (for example), and so they don't change the (cyclically reduced) word length. So the factor automorphisms are translation length preserving, and have an action on T. By Corollary 4.7(2), this action is compatible with the action of $\hat{G}$, and so we have an action of $\hat{G} \rtimes \mathrm{Fact}(G)$ on T.

A quotient graph of groups for this action, taking the same fundamental domain, is shown in Figure 2.3b. The factor automorphisms preserve the subgroups $\hat{A}$, $\hat{B}$, and $\hat{C}$ and so fix the fundamental domain: the equivariance of the induced isometries means fixed points are sent to fixed points, so an automorphism preserving a subgroup will induce an isometry preserving its fixed point set. Since the fixed point sets in an action with trivial edge stabilisers are single vertices, this means that the induced isometry fixes the same vertex. The central vertex must then be fixed in order to preserve adjacency. Also, no orbits are collapsed by this action, so the fundamental domain does remain the same.

A element of $\hat{G} \rtimes \mathrm{Fact}(G)$ may be written uniquely as φw, where φ is an element of $\mathrm{Fact}(G)$ and w is an element of $\hat{G}$. Each conjugacy class has a representative with w cyclically reduced: first write $w = h^{-1}gh$, where g is cyclically reduced, and the last letter of g and the first letter of h are drawn from different factor groups. (So there are no reductions or concatenations to do except possibly at $h^{-1}g$.) We can then conjugate the element as a whole by h^{-1}, giving $\varphi w \sim h\varphi h^{-1}ghh^{-1} = \varphi(h\varphi)h^{-1}g$. The word $(h\varphi)h^{-1}g$ (after cancelling and concatenating as necessary, depending on how many terminal letters of h are fixed by φ) is cyclically reduced, since we chose h to ensure that it (and so also $(h\varphi)$) have a first letter drawn from a different factor group to the last letter of g.

So it is enough to calculate the translation length of φw when w is cyclically reduced. Using the graph of groups in Figure 2.3b, since φ can be picked up at the same vertex as the first non-trivial group element, and w is cyclically reduced, the length of the (cyclically reduced) path for φw is just the same as that for w. So $\|\varphi w\| = \|w\|$.

Now we need to describe the effect of a permutation automorphism on the translation length. We have seen that is enough to understand it in the case where w is cyclically reduced, so we restrict to this case. In general, there are inner factor automorphisms introduced by the permutation, which we will need to move past the rest of the word to get back to our standard form. This can't change the length or structure (in terms of a sequence of factor groups from which the elements have come) of the word, since they either fix each letter or replace it with a different letter from the same factor group. So $(\varphi w)\sigma = \varphi' w'$, where φ' is a (likely different) element of $\mathrm{Fact}(G)$, and w' is the image of w after applying σ and moving any inner factor automorphisms past it. Provided w was cyclically reduced, w' is also cyclically reduced and has the same length. So by the argument above, $\|(\varphi w)\sigma\| = \|w'\| = \|w\| = \|\varphi w\|$, and so the translation length is preserved by the permutation automorphisms.

So the permutation automorphisms are a subgroup of $\mathrm{Aut}_T(\hat{G} \rtimes \mathrm{Fact}(G))$, and so we may further extend the action to the full outer automorphism group $(\hat{G} \rtimes \mathrm{Fact}(G)) \rtimes \mathrm{Perm}(G)$, again by applying Corollary 4.7(2). □

A quotient graph of groups for this action (giving the splitting) is shown in Figure 2.3c. The effect of the permutation automorphisms is to collapse the orbits of the three outer vertices to one, while preserving the orbit of the central vertex. So there are two orbits of vertices and one orbit of edges: the edge is stabilised by $\mathrm{Fact}(G) \rtimes C_2$, and the vertices by $(\hat{A} \rtimes \mathrm{Fact}(G)) \rtimes C_2$ and by $\mathrm{Fact}(G) \rtimes S_3$.

We are now in a position to prove the final part of Theorem 4.1:

Corollary 4.11. *Suppose G is a free product of freely indecomposable groups, such that any of the following occur:*

- *the free rank is exactly* 2*;*
- *the free rank is exactly* 1*, and another free factor appears exactly once*

- *G has no infinite cyclic factors and either a free factor appears exactly two or three times, or any two free factors appear exactly once.*

Then $Aut(G)$ does not have Property FA.

Proof. Let H be the subgroup generated by the free factors matching one of the conditions in the hypotheses. The normal subgroup N generated by all other free factors is characteristic ($N\varphi = N$ for all automorphisms φ of G), since it contains all representatives of these isomorphism classes. So there is a homomorphism from $\mathrm{Aut}(G)$ to $\mathrm{Aut}(H)$. It is onto since $\mathrm{Aut}(G)$ already contains a copy of $\mathrm{Aut}(H)$ (by restricting the presentation to only these generators), and this is sent to itself. We have that $\mathrm{Aut}(H)$ acts on a tree by Proposition 4.8 if H has two free factors, or that the quotient $\mathrm{Out}(H)$ (and therefore the group $\mathrm{Aut}(H)$) does by Proposition 4.10 if there are three. In either case, we have an action (without global fixed points) of a quotient of $Aut(G)$ on a tree, and so $\mathrm{Aut}(G)$ also acts on that tree without global fixed points. So $\mathrm{Aut}(G)$ does not have Property (FA). □

The proof of Theorem 4.1 is just assembling the proofs in this section:

Proof of Theorem 4.1.

(1) This is (3) of Corollary 4.11.

(2) This is Proposition 4.2.

(3) This is Proposition 4.3. □

However, all of these groups (assuming the free product is non-trivial) do have finite index subgroups that admit actions on trees:

Proposition 4.12. *Suppose G is a finite and non-trivial free product, where each factor is freely indecomposable and not infinite cyclic. Then there is a finite index subgroup of* $\mathrm{Aut}(G)$ *that does not have Property (FA).*

This proof uses only methods already described in this paper. A different proof can be established from [13].

Proof. The finite index subgroup we will work with is the group $\mathrm{FR}(G) \rtimes \mathrm{Fact}(G)$, with the finite quotient being $\mathrm{Perm}(G)$. Observe that all the generators of this group preserve the conjugacy class of each free factor, making the normal closure of any collection of free factors 'characteristic for this subgroup'. Let N be the normal closure of all but two factors. There is a map to $\mathrm{Aut}(G/N)$, and all generators (apart from the permutation, if present) are in the image. So we have a quotient isomorphic to (a finite index subgroup of) some $\mathrm{Aut}(H)$, where H is a free product of just two groups. (If all the free factors are isomorphic, then $\mathrm{Aut}(H)$ necessarily contains a permutation automorphism, which is not in the image. However, its index 2 subgroup $\mathrm{FR}(H) \rtimes \mathrm{Fact}(H)$ works just as well for the rest of the argument.) By Proposition 4.8 this admits an action on a tree and so does not have Property (FA). □

Corollary 4.13. *Suppose G is a finite and non-trivial free product where each factor is freely indecomposable and not infinite cyclic. Then* $\mathrm{Aut}(G)$ *does not have Kazhdan's Property (T).*

Proof. Discrete groups with Property (T) are finitely generated [3, Theorem 1.3.1], so if any factor is uncountable they certainly do not have Property (T). If all factors are countable, we may use Watatani's result [24] which gives that if $\mathrm{Aut}(G)$ had Property (T), then every finite index subgroup would have Property (FA). Since Proposition 4.12 gives a finite index subgroup which acts on a tree, and therefore does not have Property (FA), $\mathrm{Aut}(G)$ cannot have Property (T). □

References

[1] Roger Alperin and Hyman Bass. Length functions of group actions on Λ-trees. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 265–378. Princeton Univ. Press, Princeton, NJ, 1987.

[2] Hyman Bass. Covering theory for graphs of groups. *J. Pure Appl. Algebra*, 89(1-2):3–47, 1993.

[3] Bachir Bekka, Pierre de la Harpe, and Alain Valette. *Kazhdan's property (T)*, volume 11 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2008.

[4] O. V. Bogopolski. Arboreal decomposability of groups of automorphisms of free groups. *Algebra i Logika*, 26(2):131–149, 271, 1987.

[5] Donald J. Collins and N. D. Gilbert. Structure and torsion in automorphism groups of free products. *Quart. J. Math. Oxford Ser. (2)*, 41(162):155–178, 1990.

[6] Yves Cornulier and Aditi Kar. On property (FA) for wreath products. *J. Group Theory*, 14(1):165–174, 2011.

[7] Marc Culler and John W. Morgan. Group actions on $\mathbb{R}$-trees. *Proc. London Math. Soc. (3)*, 55(3):571–604, 1987.

[8] Marc Culler and Karen Vogtmann. A group-theoretic criterion for property FA. *Proc. Amer. Math. Soc.*, 124(3):677–683, 1996.

[9] Warren Dicks and M. J. Dunwoody. *Groups acting on graphs*, volume 17 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989.

[10] Max Forester. Deformation and rigidity of simplicial group actions on trees. *Geom. Topol.*, 6:219–267, 2002.

[11] D. I. Fouxe-Rabinovitch. Über die Automorphismengruppen der freien Produkte. I. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):265–276, 1940.

[12] D. I. Fouxe-Rabinovitch. Über die Automorphismengruppen der freien Produkte. II. *Rec. Math. [Mat. Sbornik] N. S.*, 9 (51):183–220, 1941.

[13] Anthony Genevois and Olga Varghese. Conjugating automorphisms of graph products: Kazhdan's property (T) and SQ-universality. *Bull. Aust. Math. Soc.*, 101(2):272–282, 2020.

[14] N. D. Gilbert. Presentations of the automorphism group of a free product. *Proc. London Math. Soc. (3)*, 54(1):115–140, 1987.

[15] Marek Kaluba, Dawid Kielak, and Piotr W. Nowak. On property (T) for $\mathrm{Aut}(F_n)$ and $\mathrm{SL}_n(\mathbb{Z})$. *Ann. of Math. (2)*, 193(2):539–562, 2021.

[16] Marek Kaluba, Piotr W. Nowak, and Narutaka Ozawa. $\mathrm{Aut}(\mathbb{F}_5)$ has property (T). *Math. Ann.*, 375(3-4):1169–1191, 2019.

[17] Sabine Koppelberg and Jacques Tits. Une propriété des produits directs infinis de groupes finis isomorphes. *C. R. Acad. Sci. Paris Sér. A*, 279:583–585, 1974.

[18] Nils Leder. Serre's Property FA for automorphism groups of free products. Preprint, Oct 2018. arXiv:1810.06287.

[19] Gilbert Levitt. Characterizing rigid simplicial actions on trees. In *Geometric methods in group theory*, volume 372 of *Contemp. Math.*, pages 27–33. Amer. Math. Soc., Providence, RI, 2005.

[20] Roger C. Lyndon and Paul E Schupp. *Combinatorial group theory*. Springer-Verlag Berlin and Heidelberg GmbH & Co. K, Berlin, Germany, January 1977.

[21] Martin Nitsche. Computer proofs for Property (T), and SDP duality. Preprint, September 2020. arXiv:2009.05134v2 [math.GR].

[22] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[23] John Stallings. Coherence of 3-manifold fundamental groups. In *Séminaire Bourbaki, Vol. 1975/76, 28 ème année, Exp. No. 481*, pages 167–173. Lecture Notes in Math., Vol. 567. 1977.

[24] Yasuo Watatani. Property T of Kazhdan implies property FA of Serre. *Math. Japon.*, 27(1):97–103, 1982.

A A presentation of $\mathrm{Out}(G)$

This appendix contains a proof of the presentation of $\mathrm{Out}(G)$ given in Section 4:

Proposition 4.9. *Suppose $G = A * B * C$ is a free product of three copies of some freely indecomposable and not infinite cyclic group. Let $\gamma(a)$ denote the inner factor automorphism conjugating A by a; $\varphi = (\varphi_a, \varphi_b, \varphi_c)$ denote an element of $\mathrm{Fact}(G)$; and $\sigma_{(12)}, \sigma_{(123)}$ the permutation automorphisms. Then $\mathrm{Out}(G)$ has the following presentation.*

Generators: $(A,b),(B,c),(C,a),\mathrm{Aut}(A),\mathrm{Aut}(B),\mathrm{Aut}(C),\sigma_{(123)},\sigma_{(12)}$

Relations (with terms going over all appropriate generators):

$$(A,b)(A,b') = (A,b'b), \textit{ etc.} \tag{1}$$

$$[\mathrm{Aut}(A),\mathrm{Aut}(B)] = [\mathrm{Aut}(A),\mathrm{Aut}(C)] = [\mathrm{Aut}(B),\mathrm{Aut}(C)] = 1 \tag{2}$$

$$\sigma_{(123)}^3 = 1, \sigma_{(12)}^2 = 1, (\sigma_{(123)}\sigma_{(12)})^2 = 1 \tag{3}$$

$$\varphi^{-1}(A,b)\varphi = (A,b\varphi), \textit{ etc.} \tag{4}$$

$$\sigma_{(12)}^{-1}(A,b)\sigma_{(12)} = \gamma(a^{-1})(C,a^{-1}) \tag{5}$$

$$\sigma_{(12)}^{-1}(B,c)\sigma_{(12)} = \gamma(c^{-1})(B,c^{-1}) \tag{6}$$

$$\sigma_{(12)}^{-1}(C,a)\sigma_{(12)} = \gamma(b^{-1})(A,b^{-1}) \tag{7}$$

$$\sigma_{(123)}^{-1}(A,b)\sigma_{(123)} = (B,c), \textit{ etc.} \tag{8}$$

$$(\varphi_a,\varphi_b,\varphi_c)^{\sigma_{(12)}} = (\varphi_b,\varphi_a,\varphi_c), \quad (\varphi_a,\varphi_b,\varphi_c)^{\sigma_{(123)}} = (\varphi_b,\varphi_c,\varphi_a) \tag{9}$$

> *Relations (2) are inherited from the direct product structure on* Fact(G)*; relations (3) give a presentation of* S_3*, and relations (9) are a consequence of the wreath product structure on* $\langle \mathrm{Fact}(G), \mathrm{Perm}(G)\rangle$*.*

The proof is largely the same as that given in [5] for three non-isomorphic factors, differing by taking account of the permutation automorphisms which appear when the factor groups are isomorphic.

A presentation of Aut(G) (derived from Propositions 2.9 and 2.8) consists of:

Generators: $(A,b),(A,c),(B,a),(B,c),(C,a),(C,b);\mathrm{Aut}(A),\mathrm{Aut}(B),\mathrm{Aut}(C);\sigma_{(123)},\sigma_{(12)}$
Denote by $\varphi = (\varphi_a,\varphi_b,\varphi_c)$ an arbitrary factor automorphism, and σ some permutation automorphism. Then we have the following relations, which should be taken to range over all appropriate generators:

$$[(A,b),(C,b')] = 1, \text{ etc.} \tag{1}$$

$$[(A,b)(C,b),(A,c)] = 1, \text{ etc.} \tag{2}$$

$$(A,b)(A,b') = (A,b'b), \text{ etc.} \tag{3}$$

$$[\mathrm{Aut}(A),\mathrm{Aut}(B)] = [\mathrm{Aut}(A),\mathrm{Aut}(C)] = [\mathrm{Aut}(B),\mathrm{Aut}(C)] = 1 \tag{4}$$

$$\sigma_{(123)}^3 = 1, \sigma_{(12)}^2 = 1, (\sigma_{(123)}\sigma_{(12)})^2 = 1 \tag{5}$$

$$\varphi^{-1}(A,b)\varphi = (A,b\varphi), \text{ etc.} \tag{6}$$

$$\sigma^{-1}(A,b)\sigma = (A\sigma,b\sigma), \text{ etc.} \tag{7}$$

$$(\varphi_a,\varphi_b,\varphi_c)^{\sigma_{(12)}} = (\varphi_b,\varphi_a,\varphi_c), \quad (\varphi_a,\varphi_b,\varphi_c)^{\sigma_{(123)}} = (\varphi_b,\varphi_c,\varphi_a) \tag{8}$$

Relations (4) are inherited from the direct product structure on $\mathrm{Fact}(G)$; relations (5) give a presentation of S_3, and relations (8) are a consequence of the wreath product structure on $\langle \mathrm{Fact}(G), \mathrm{Perm}(G) \rangle$.

This gives $\mathrm{Aut}(G)$ as the iterated semidirect product $\mathrm{FR}(G) \rtimes \mathrm{Aut}(A)^3 \rtimes S_3$, where $\mathrm{Aut}(A)^3 \rtimes S_3$ is the permutation wreath product in Proposition 2.8. Since S_3 normalises $\mathrm{FR}(G)$, it can be evaluated in either order.

To find a presentation of $\mathrm{Out}(G)$, we add relations to this presentation setting each inner automorphism equal to the identity. To do this, let $\gamma(a)$ be the inner factor automorphism conjugating A by $a \in A$ and fixing the other factor groups, and similarly $\gamma(b)$ and $\gamma(c)$. Then we add relations $\gamma(a)(B,a)(C,a) = 1$, and similarly for elements in B and C.

Use the new relations to rewrite three kinds of generators ((A,c), (B,a) and (C,b)). For example, $(A,c) = \gamma(c^{-1})(B,c^{-1})$. Then we can eliminate both those generators and the new relations. Putting this substitution in the first kind of relation we see that they are implied by the others (and so are unnecessary):

$$
\begin{aligned}
[(A,c),(B,c')] &= [\gamma(c^{-1})(B,c^{-1}),(B,c')] \\
&= (B,c)\gamma(c)(B,c'^{-1})\gamma(c^{-1})(B,c^{-1})(B,c') \\
&= (B,c)(B,cc'^{-1}c^{-1})(B,c^{-1})(B,c') && \text{by (6)} \\
&= (B,c'c^{-1}cc'^{-1}c^{-1}c) && \text{by (3)} \\
&= (B,1) \\
&= 1
\end{aligned}
$$

Similarly for the second kind:

$$
\begin{aligned}
[(A,b)(C,b),(A,c)] &= [\gamma(b^{-1}),\gamma(c^{-1})(B,c^{-1})] \\
&= \gamma(b)(B,c)\gamma(c)\gamma(b^{-1})\gamma(c^{-1})(B,c^{-1}) \\
&= \gamma(b)(B,c)\gamma(b^{-1})(B,c^{-1}) && \text{by (4)} \\
&= (B,c)(B,c^{-1}) && \text{by (6)} \\
&= 1
\end{aligned}
$$

Sometimes only one substitution is required, for example the relation

$$[(A,b)(C,b),(C,a)] = 1$$

.

Relations (3),(4),(5), and (8) are all in terms of only generators we eliminated (in which case they have also been eliminated) or of only generators we still have, so don't need any rewriting. Relations (6) only have the effect of changing the conjugating element for another drawn from the same factor group, so again don't require any rewriting.

However, (7) requires rewriting for transpositions. Taking $\sigma_{(12)}$, to interchange A and B, we have

$$\begin{aligned} \sigma_{(12)}(A,b)\sigma_{(12)} &= (A\sigma_{(12)}, b\sigma_{(12)}) \\ &= (B,a) \\ &= \gamma(a^{-1})(C,a^{-1}) \end{aligned}$$

And similarly:

$$\begin{aligned} \sigma_{(12)}^{-1}(C,a)\sigma_{(12)} &= \gamma(b^{-1})(A,b^{-1}) \\ \sigma_{(12)}^{-1}(B,c)\sigma_{(12)} &= \gamma(c^{-1})(B,c^{-1}) \end{aligned}$$

So with generators $\sigma_{(12)}$ (interchanging A and B) and $\sigma_{(123)}$ (cycling A to B to C to A) we replace (7) above with:

$$\sigma_{(12)}^{-1}(A,b)\sigma_{(12)} = \gamma(a^{-1})(C,a^{-1}) \tag{7a}$$

$$\sigma_{(12)}^{-1}(B,c)\sigma_{(12)} = \gamma(c^{-1})(B,c^{-1}) \tag{7b}$$

$$\sigma_{(12)}^{-1}(C,a)\sigma_{(12)} = \gamma(b^{-1})(A,b^{-1}) \tag{7c}$$

$$\sigma_{(123)}^{-1}(A,b)\sigma_{(123)} = (B,c), \text{ etc.} \tag{7d}$$

After eliminating (1) and (2), and replacing (7) by (7a)-(7d), this gives the presentation of Proposition 4.9.

Note that while this still gives an iterated semidirect product structure, in general the S_3 no longer normalises the first factor and so the order of evaluation is important. If the free factors are abelian then the inner factor automorphisms are trivial and again the semidirect product can be evaluated in either order.

Paper 3: **Free-by-cyclic groups, automorphisms and actions on nearly canonical trees**

Naomi Andrew and Armando Martino

ABSTRACT. We study the automorphism groups of free-by-cyclic groups and show these are finitely generated in the following cases: (i) when defining automorphism has linear growth and (ii) when the rank of the underlying free group has rank at most 3.

The techniques we use are actions on trees, including the trees of cylinders due to Guirardel and Levitt, the relative hyperbolicity of free-by-cyclic groups (due to Gautero and Lustig, Ghosh, and Dahmani and Li) and the filtration of the automorphisms of a group preserving a tree, following Bass and Jiang, and Levitt.

Our general strategy is to produce an invariant tree for the group and study that, usually reducing the initial problem to some sort of McCool problem (the study of an automorphism group fixing some collection of conjugacy classes of subgroups) for a group of lower complexity. The obstruction to pushing these techniques further, inductively, is in finding a suitable invariant tree and in showing that the relevant McCool groups are finitely generated.

1 Introduction

1.1 Free-by-cyclic groups

Given a finite rank free group F_n and an automorphism $\varphi \in \mathrm{Aut}(F_n)$, we can define a free-by-cyclic group $G = F_n \rtimes_\varphi \langle t \rangle = \langle x_1, \dots x_n, t | t^{-1} x_i t = x_i \varphi \rangle$ (so conjugating by the stable letter t acts on F_n as the automorphism φ). The properties of this free-by-cyclic group depend only on the automorphism φ, and in fact only on the conjugacy class of its image in the outer automorphism group, Φ [6, Lemma 2.1].

Various properties of G follow from φ and indeed from Φ: for example, G is hyperbolic if and only if φ is atoroidal (no power of φ fixes the conjugacy class of an element in F_n) [7], and is relatively hyperbolic if and only if the length of some word in F_n grows exponentially under iteration of φ [14, 15, 13]. Both of these properties are properties of the outer class as a whole.

In this paper we study the actions of free-by-cyclic groups on trees, and through this their automorphisms. Even in rank 1 (the two cyclic-by-cyclic groups) it is hard to say anything very general about their automorphisms: for $\mathbb{Z}^2$, the outer automorphism group is $\mathrm{GL}(2, \mathbb{Z})$, whereas for the fundamental group of the Klein bottle it has only four elements.

There are groups which can be expressed as free-by-cyclic groups with more that one possibility for the rank of F_n. However, there are some things which these presentations will have in common: for example, the growth rate of the outer automorphism Φ will be the same [26]. An automorphism is polynomially growing (with degree d) if, as it is iterated, the conjugacy length of a word is bounded by a polynomial (of degree d), and (by a Theorem of [5] – see Subsection 2.4) exponentially growing otherwise. We split our investigation by growth rate.

Levitt's work on Generalised Baumslag-Solitar groups [22] includes (after checking some hypotheses) that if the defining (outer) automorphism is finite order (in which case the free-by-cyclic group G is virtually $F_n \rtimes \mathbb{Z}$) then $\mathrm{Out}(G)$ is VF, and so in particular finitely generated.

We extend finite generation to all cases where the defining outer automorphism has linear growth:

Theorem 1.1.1. *Suppose $G \cong F_n \rtimes_\varphi \mathbb{Z}$, and φ is linearly growing. Then* $\mathrm{Out}(G)$ *is finitely generated.*

Also [6], studies the case when the underlying free group has rank 2. There $\mathrm{Out}(G)$ is calculated up to finite index for all defining automorphisms, and this classification shows that $\mathrm{Out}(G)$ is finitely generated.

We extend the finite generation result to all cases where the underlying free group has rank 3 (in which case the growth is either at most quadratic or exponential):

Theorem 1.1.2. *Suppose $G \cong F_3 \rtimes \mathbb{Z}$. Then* $\mathrm{Out}(G)$ *is finitely generated.*

We understand the automorphism groups through studying certain actions of G on trees. Since they are defined as HNN extensions, all free-by-cyclic groups have a translation action on the real line. But they also admit actions on more complicated trees. These actions are equivalent to alternative presentations which can provide more information about the group. To understand the automorphisms, we use particular trees which are in some sense invariant under all – or sometimes only most – automorphisms.

The details are different in the exponentially growing and polynomially growing cases. With exponential growth, G is one-ended relatively hyperbolic, and so it has a canonical JSJ decomposition by [18]. These decompositions are particularly useful and well understood, and there is a description of the outer automorphism group arising from them. We describe the canonical tree and for the low rank cases carry out the calculations needed for the automorphism group in Section 4.

Using Guirardel and Levitt's tree of cylinders construction [16], we construct canonical trees when the defining automorphism is unipotent polynomially growing (UPG) and either linear or, in low rank, quadratic. These trees arise from fixed points on the boundary of Culler-Vogtmann outer space for the defining (outer) automorphism and restricting the action to F_n will give an action in the same deformation space as such a tree. Every polynomially growing automorphism has a power which is UPG (in fact, the power can be taken to depend only on the rank of the free group – see Definiton 2.4.4). This implies the existence of a normal finite index subgroup which is again free-by-cyclic, this time with a UPG defining automorphism.

Understanding the automorphisms of a finite index subgroup does not necessarily provide insight into those of the larger group: the fundamental group of a Klein bottle (with only four outer automorphisms) contains $\mathbb{Z}^2$ as an index 2 subgroup. A key part of our proof is that we can use the existence of a canonical splitting of a normal finite index subgroup to find a splitting of the larger group which is "nearly canonical" – invariant under at least a finite index subgroup of automorphisms.

The result is:

Proposition 3.1.4. *Let G be a finitely generated group, G_0 a normal finite index subgroup of G, and suppose that T is a canonical G_0-tree. Then*

(i) G acts on T, and this action restricts to the canonical G_0-action.

(ii) With this action, T is nearly canonical as a G-tree.

By this result, we have an action of G on a tree, and we can consider the outer automorphisms preserving this action, which is finite index in the full outer automorphism group. Understanding this group depends on understanding the vertex and edge groups, their automorphisms, and how those automorphisms interact. In particular, we need to calculate "McCool groups", for vertex groups with respect to adjacent edge groups: the outer automorphisms having representatives that restrict to the identity on each of a family of subgroups.

As part of our proof, we carry this calculation out for free-by-cyclic groups defined by a periodic automorphism, with respect to a limited class of subgroups, and when the underlying free group has rank 2.

We note that there appear to be two main obstacles to extending this result further: constructing actions on trees which are (nearly) canonical, and understanding the McCool groups arising from these trees. In the exponential case, the canonical trees exist and the obstruction is only the McCool groups, which are generally required to be with respect to fairly complex subgroups. In the polynomially growing case(s) passing to a (UPG) power should lead to actions arising from limit points of CV_n, and with quadratic growth these are even unique – see [25]. But it is not obvious that the deformation spaces these define are canonical. If canonical trees can be found, the McCool groups are likely to be needed relative only to infinite cyclic subgroups, which may be more manageable.

Acknowledgements. The first author was supported by an EPSRC studentship, and the second author by Leverhulme Trust Grant RPG-2018-058.

2 Background

2.1 Notation, Actions on trees and Bass-Serre Theory

We record here some notation for actions on trees and various subgroups of (outer) automorphisms used throughout the paper.

We recall enough of Bass-Serre theory to set notation; see [32] amongst others for a fuller exposition. Following Serre, the edges of a graph come in pairs denoted e and $\bar{e}$, and $\iota(e)$ and $\tau(e)$ denote the initial and terminal vertices. An orientation, $\mathcal{O}$, is a choice of one edge from each pair $\{e, \bar{e}\}$.

Let a group G act on a tree T. We let G_v and G_e denote the stabiliser of a vertex v or edge e respectively; from the perspective of graphs of groups we use them for vertex and edge groups, and use α_e to denote the monomorphism $G_e \to G_\iota(e)$. Often we simply identify G_e with its image $\alpha_e(G_e)$ in $G_{\iota(e)}$. An action on a tree is called *minimal* if it does not admit a G-invariant subtree; most of our actions will be assumed to be minimal. An action on a tree is *irreducible* if it does not fix a point, line, or end of the tree; to guarantee this it is sufficient that the action has two hyperbolic axes whose intersection is at most finite length.

We use $N_G(H)$, $C_G(H)$ and $Z(H)$ for the normaliser of H in G, the centraliser of H in G, and the centre of H. To save space and subscripts in the context of an action on a tree, we let $N_e = N_{G_{\iota(e)}}(G_e)$ and $C_e = C_{G_{\iota(e)}}(G_e)$.

As usual, $\mathrm{Aut}(G)$ denotes the automorphisms of G, and $\mathrm{Out}(G) = \mathrm{Aut}(G)/\mathrm{Inn}(G)$ the outer automorphisms. We use lower case greek letters (φ) for automorphisms, and upper case (Φ) for outer automorphisms. If the image of φ in $\mathrm{Out}(G)$ is Φ, we say φ *represents* Φ, or $\varphi \in \Phi$ (viewing Φ as a coset of $\mathrm{Inn}(G)$).

Given an automorphism φ of G, we can define the cyclic extension of G by φ as

$$G \rtimes_\varphi \mathbb{Z} = \langle X, t : R, t^{-1}xt = x\varphi \rangle$$

(taking $\langle X : R \rangle$ to be a presentation of G). Automorphisms representing the same outer automorphism define isomorphic extensions, as can be seen by introducing a new generator $t' = tg$. For this reason, we will sometimes use the notation $G \rtimes_\Phi \mathbb{Z}$

to refer to the isomorphism class of cyclic extensions defined by any automorphism representing Φ.

For $g \in G$, we write $\mathrm{Ad}(g)$ for the inner automorphism of G induced by g.

If G normalises H, let $\mathrm{Ad}(G, H)$ denote the automorphisms of H induced by conjugating by elements of G. If H is clear, it may be omitted; in particular (and assuming an action on a tree) $\mathrm{Ad}(N_e)$ always means the automorphisms of G_e induced by conjugating by N_e, its normaliser in $G_{\iota(e)}$. Notice that since N_e contains G_e, the subgroup $\mathrm{Ad}(N_e)$ descends to a subgroup of $\mathrm{Out}(G_e)$.

We identify certain "relative" subgroups of $\mathrm{Out}(G)$:

Definition 2.1.1. Given a family of subgroups $\{G_i\}$ of G, we define

- the subgroup $\mathrm{Out}(G; \{G_i\})$ to be those outer automorphisms of G where for each subgroup G_i there is a representative that restricts to an automorphism of G_i.
- the subgroup $\mathrm{Mc}(G; \{G_i\})$ to be those outer automorphisms of G where for each subgroup G_i there is a representative that restricts to the identity on G_i.

Note that these are subgroups of *outer* automorphisms; any given representative will not usually have the correct restriction for every subgroup G_i.

Throughout the paper we consider actions of $F_n \rtimes_\varphi \mathbb{Z}$ on trees. Dahmani (in Section 2.2 of [12]) gives some useful results about such an action. The following lemma is specialised to free-by-cyclic groups; Dahmani gives it more generally for semidirect products with $\mathbb{Z}$ (suspensions, in the terminology of that paper) of any finitely generated group.

Lemma 2.1.2. *Suppose $G \cong F_n \rtimes \langle t \rangle$ acts minimally and irreducibly on a tree. Then*

(1) F_n acts on the same tree with finite quotient graph

(2) The stabilisers in any action of G on a tree are again free-by-cyclic; the free part is the F_n-stabiliser, and the generator of the cyclic factor has the form $t^k w$.

(3) In particular, all edge stabilisers are at least infinite cyclic, and G is one ended.

(4) If all incident edges at some vertex are cyclically stabilised, its stabiliser cannot be finitely generated and infinitely ended.

The hypotheses given here differ slightly from Dahmani's: we demand an irreducible action while Dahmani uses "reduced". In fact a sufficient condition is that F_n acts non-trivially, which is guaranteed by either of these conditions.

Note that the last point is not immediately obvious: the free part at a vertex could be infinitely generated, and free by cyclic groups of this form can be infinitely ended, and even free. It implies that in any splitting of this kind there cannot be a "quadratically hanging" vertex group, for these have exactly the combination of properties ruled out here (see Section 4).

The crucial observation is that since F_n is a normal subgroup, it also acts minimally on the whole tree. Then finite generation ensures the quotient under this action is finite. To recover the free-by-cyclic structure on the stabilisers, consider the action of $\langle t \rangle$ on the quotient graph, and lift the stabilisers back to the whole group. This kind of argument enables us to analyse the splitting of G by considering the induced splitting of F_n.

To see the last point, in this case, note that this could only occur if the free part of the relevant vertex stabiliser was not finitely generated. Contract all other edges and consider the induced free splitting of F_n: this expresses F_n as a free product where one free factor is not finitely generated, which is impossible.

(More generally, the free part of a vertex may only be infinitely generated if the same is true of at least one incident edge group; control over the edge groups provides some control over the vertex groups.)

2.2 Length Functions and Twisting Actions by Automorphisms

Since we will usually be working with simplicial metric trees, an action of G on a tree T will be equivalent to a map $G \to \text{Isom}(T)$.

Any action of G on a tree T defines a *translation length function* on G, by considering the minimum displacement of points in the tree for each element. That is, given an isometric action of G on T, we can define the function, $l_T : G \to \mathbb{R}$ by $l_T(g) = \min_{x \in T} d_T(x, xg)$ (and this minimum is always realised). Note that l_T is constant on conjugacy classes.

We recall a well-known Theorem of Culler and Morgan:

Theorem 2.2.1 ([11, Theorem 3.7]). *Let G be a finitely generated group and let T_1, T_2 be two $\mathbb{R}$-trees equipped with isometric G actions which are minimal and irreducible. Then $l_{T_1} = l_{T_2}$ if and only if T_1 and T_2 are equivariantly isometric. Moreover, such an equivariant isometry is unique if it exists.*

Remark 2.2.2. This result says that, in many cases, the translation length function determines the action.

The action of G on T defines a deformation space, by considering all simplicial actions of G on a tree with the same elliptic subgroups (this is an equivalence relation on G-trees); the elliptic subgroups are those subgroups of G which fix a point in the tree. Note that there can still be vertices with stabilisers that are not conjugate to a stabiliser in the original action. (For example, consider representing a free product of three groups as a graph of groups where the underlying graph is a line versus a tripod: these are in the same deformation space, despite the extra trivially stabilised vertex.) Trees in the same deformation space *dominate* each other; that is, there are equivariant maps between them.

Definition 2.2.3. Given an isometric action of a group G on a tree, T, a new 'twisted' action of G on T can be defined by pre-composing with any automorphism of G. That is, if $\varphi \in \mathrm{Aut}(G)$, then $x \cdot_\varphi g = x \cdot (g\varphi)$

In terms of length functions, this means that $l_{\varphi T}(g) = l_T(g\varphi)$. (Here φT is the "twisted tree", isometric to T but with the new action defined above.)

Given a deformation space of trees, this defines an action of $\mathrm{Aut}(G)$ on that space.

Remark 2.2.4. Note that there is a switch from left to right; if the automorphisms of G act on elements on the right then the action on trees by pre-composing is on the left and vice versa.

In most cases this changes the length function; we let

$$\mathrm{Aut}^T(G) = \{\varphi \in \mathrm{Aut}(G) \ : \ l_{\varphi T} = l_T\}$$

denote the subgroup of $\mathrm{Aut}(G)$ which leaves it unchanged. Notice that this is true of all inner automorphisms, so these are a subgroup of $\mathrm{Aut}^T(G)$. By Theorem 2.2.1 such an automorphism induces an equivariant isometry of T, and assuming the action is minimal and does not fix an end this is unique and extends to an action of $\mathrm{Aut}^T(G)$. This action is compatible with the original action in at least two senses: $G \rtimes \mathrm{Aut}^T(G)$ (with the usual action of $\mathrm{Aut}^T(G)$ on G) acts on T, restricting to the original action of G and the induced action of $\mathrm{Aut}^T(G)$, and the action of G on T factors through the map sending each element to the inner automorphism it induces.

This is asserting the existence of a commuting diagram (see [1] for how to use Theorem 2.2.1 to produce this diagram):

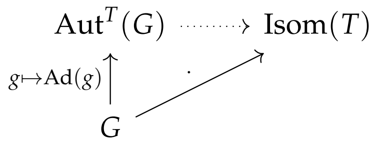

Recall (from Section 2.1) that $\mathrm{Ad}(g)$ is the inner automorphism induced by g.

In fact, such a diagram is also sufficient to recover the definition in terms of length functions since for any $\varphi \in \mathrm{Aut}^T(G)$,

$$l_T(g) = l_T(\mathrm{Ad}(g)) = l_T(\mathrm{Ad}(g)^{\varphi}) = l_T(\mathrm{Ad}(g\varphi)) = l_T(g\varphi),$$

(where we are moving between the G action and the $\mathrm{Aut}^T(G)$ action using the commutative diagram).

We will consider $\mathrm{Out}^T(G) = \mathrm{Aut}^T(G)/\operatorname{Inn}(G)$, the subgroup of outer automorphisms which preserves the length function. By the correspondence theorem, many properties of $\mathrm{Aut}^T(G)$, such as finite index or normality, are inherited by $\mathrm{Out}^T(G)$.

2.3 Trees of cylinders

Guirardel and Levitt in [16] define a tree of cylinders for a deformation space. The input is any tree in the deformation space, and an equivalence relation on the edges; the output is a tree where the induced splitting is preserved by all (outer) automorphisms which preserve the deformation space. They are our main tool for producing trees which allow us to analyse outer automorphisms by considering trees.

We start the construction by defining a family $\mathcal{E}$ of subgroups of G. It should be closed under conjugation, but not under taking subgroups. We then define an *admissible equivalence relation* on $\mathcal{E}$ [16, Definition 3.1]. This must satisfy

(1) if $A \sim B$ then $A^g \sim B^g$ for all $g \in G$

(2) if $A \leq B$ then $A \sim B$

(3) Suppose G acts on a tree with stabilisers in $\mathcal{E}$. If $A \sim B$, $v \in \mathrm{Fix}(A)$ and $w \in \mathrm{Fix}(B)$, then the stabiliser of any edge lying in $[v, w]$ is equivalent to A (and B)

To show (3) it is sufficient to show that $\langle A, B \rangle$ is elliptic [16, Lemma 3.2]

Now suppose that G acts on T with edge stabilisers in $\mathcal{E}$. Define an equivalence relation on the edges of T by saying $e \sim e'$ if $G_e \sim G_{e'}$. A *cylinder* consists of an equivalence class of edges; the conditions on an admissible equivalence relation ensure that cylinders are connected, and that two cylinders may intersect in at most one vertex.

To construct the tree of cylinders T_c, replace each cylinder with the cone on its boundary [16, Definition 4.3]. That is, there is a vertex Y for every cylinder, together with surviving vertices x lying on the boundary of two (or more) cylinders. Edges show inclusion of a boundary vertex x into a cylinder Y. The stabilisers of boundary vertices are unchanged; the stabiliser of a cylinder vertex is the (setwise) stabiliser of the cylinder. Edge stabilisers are the intersection of the relevant vertex stabilisers.

The tree of cylinders T_c depends only on the deformation space of T, in the sense that given two minimal, non-trivial trees T, T' in the same deformation space, there is a canonical equivariant isomorphism between T_c and T'_c [16, Corollary 4.10]. In particular this means that this tree of cylinders is fixed by any automorphism which preserves the deformation space, and so can be used to study these (outer) automorphisms.

It is always true that T dominates T_c, but cylinder stabilisers may not be elliptic in T. The deformation space of the tree of cylinders depends on the size of the cylinders: if all cylinders are bounded, or equivalently contain no hyperbolic axis, then the cylinder stabilisers are elliptic in T and so T_c lies in the same deformation space, and conversely [16, Proposition 5.2].

Edge stabilisers may not be in $\mathcal{E}$; in this case the *collapsed tree of cylinders* T_c^* is defined by collapsing all edges of T_c with stabilisers not in $\mathcal{E}$ [16, Definition 4.5]. Assuming that $\mathcal{E}$ is sandwich closed (if $A \leq B \leq C$ are subgroups of G, and A and C are in $\mathcal{E}$, then so is B), the construction is stable in the sense that $(T_c^*)_c^* = T_c^*$ [16, Corollary 5.8]. If T and T' are in the same deformation space then there is a unique equivariant isometry between T_c^* and $(T')_c^*$ [16, Corollary 5.6] and again this action is canonical.

In general, there may be more restrictions put on the trees: sometimes we require that a certain collection of subgroups is elliptic. In this case the deformation space and tree of cylinders is canonical relative to the automorphisms which preserve this collection.

2.4 Automorphisms of free groups

We recall some of the results we will use about automorphisms of free groups. Here $\text{Aut}(F_n)$ denotes the automorphism group of the free group of rank n and $\text{Out}(F_n) = \text{Aut}(F_n)/\,\text{Inn}(F_n)$, the group of outer automorphisms, which is the quotient by the inner automorphisms. Thus an outer automorphism is a coset of inner automorphisms, and there is an equivalence relation on this set of automorphisms called *isogredience*. Formally,

Definition 2.4.1. Two automorphisms $\varphi, \psi \in \text{Aut}(F_n)$ are said to be *isogredient* if they are conjugate by an inner automorphism. This is an equivalence relation when restricted to any coset of $\text{Inn}(F_n)$, that is an element of $\text{Out}(F_n)$.

Theorem 2.4.2 (Bestvina-Handel Theorem, [5])**.** *Let $\Phi \in \text{Out}(F_n)$. Then,*

$$\sum \max\{\text{rank}(\text{Fix } \varphi) - 1, 0\} \leq n - 1,$$

where the sum is taken over representatives, φ, of isogredience classes in Φ.

Note that isogredient automorphisms have conjugate fixed subgroups, so the ranks of the fixed subgroups do not depend on the representatives chosen.

Growth Rate

If we fix a basis, B, of F_n then we set $\|g\|_B$ to be the length of the shortest conjugacy class of g with respect to B, for any $g \in F_n$. We simply write this as $\|g\|$ if B is understood.

Given a $\Phi \in \mathrm{Out}(F_n)$ it is then clear that there exists a λ such that,

$$\frac{\|\Phi^k(g)\|}{\|g\|} \leq \lambda^k,$$

as we can simply take λ to be the maximum conjugacy length of the image of any element of B. (Also note that since these are conjugacy lengths, we can apply any automorphism in the same outer automorphism class and get the same result. Thus we are effectively applying an outer automorphism.)

One of the results of [5] is that the growth of elements in this sense is either exponential or polynomial. That is, for any $g \in F_n$, we either get that, for some $\mu < \lambda$,

$$\mu^k \leq \frac{\|\Phi^k(g)\|}{\|g\|} \leq \lambda^k,$$

or there exist constants $0 < A < B$ such that

$$Ak^d \leq \frac{\|\Phi^k(g)\|}{\|g\|} \leq Bk^d,$$

where $d \in \{0, 1, \ldots, n-1\}$.

See [23, Theorem 6.2] for a precise description of the growth types of elements of F_n.

Accordingly, we say that

Definition 2.4.3. $\Phi \in \mathrm{Out}(F_n)$ has exponential growth if there is some element g whose conjugacy length grows exponentially. And we say that Φ has polynomial growth of degree d if the conjugacy length of every element grows polynomially and d is the maximum degree of these polynomials.

Note that in our usage "polynomial growth of degree d" implies that d is the smallest degree bounding the growth of every element: so for example for an automorphism of quadratic growth there will be an element whose conjugacy length grows quadratically.

We note that the property of having exponential or polynomial growth (and the degree of polynomial growth) are independent of the basis, B. Also, the growth type (although not the exponential growth rate) of an automorphism is the same as that of its powers. (This includes negative powers, though this is a harder fact to verify).

UPG Automorphisms

We shall look at (outer) automorphisms of polynomial growth and consider a subclass of these, called the *UPG automorphisms*.

Definition 2.4.4. (see [3], Corollary 5.7.6) We say that $\Phi \in \text{Out}(F_n)$ is *Unipotent Polynomially Growing*, or UPG, if it has polynomial growth and it has unipotent image in $\text{GL}_n(\mathbb{Z})$. This is guaranteed if the automorphism induces the trivial map on the homology group of F_n with $\mathbb{Z}_3$ coefficients.

Hence, any polynomially growing automorphism has a power which is UPG. Moreover, this power can be taken to be uniform (given n).

In the subsequent arguments we will have need to refer to a particular type of free group automorphism called a *Dehn Twist*. These are defined in terms of certain maps via a graph of groups. Namely, one takes a splitting of the free group with infinite cyclic edge groups and looks at a map defined by "twisting" along the edges. For our purposes, the following Theorem provides a useful alternative characterisation, and can be taken as a definition.

Theorem 2.4.5. *(see [10], [20], [5] and [3]) Dehn Twist automorphisms of free groups are precisely the linear growth UPG automorphisms.*

Remark 2.4.6. As commented in [20], Theorem 2.4.5 is not proved explicitly in the papers cited, but is well known to experts. The idea is that a UPG automorphism has a 'layered' improved relative train track representative by [3]. The fact that it has linear growth will imply that there are no attracting fixed points on the boundary, and from there is it relatively straightforward to produce a graph-of-groups description in terms of the 'twistors' of [10]. The arguments in [27] show how to go from the relative train track map to the graph of groups description explicitly.

A crucial Theorem about Dehn Twists is the parabolic orbits Theorem, which requires a little notation to set up. The context is Culler-Vogtmann space, CV_n, which is the space of free, simplicial actions of F_n on metric trees. In this formulation, two points – actions on trees – are said to be equivalent if there is an equivariant homothety between them. There is a compactification of this space, $\overline{CV_n}$, which turns out to be the space of *very small actions* of F_n on $\mathbb{R}$-trees. The precise definition is not necessary here, but it is worth noting that the compactification includes points which are actions on trees that are *not* simplicial $\mathbb{R}$-trees.

There is a natural action of $\text{Out}(F_n)$ on CV_n and $\overline{CV_n}$, as in Definition 2.2.3, obtained by pre-composing the action by automorphisms.

Theorem 2.4.7 (Parabolic Orbits Theorem – see [9] and [10]). *Let $\Phi \in \text{Out}(F_n)$ be a Dehn Twist. Then for any $X \in CV_n$, $\lim_{k\to\infty} \Phi^k(X) = T \in \overline{CV_n}$ exists, is a simplicial tree and lies* *in the same simplex – any two such limit trees are equivariantly homeomorphic – independently of X. Moreover, T is a simplicial F_n-tree with the following properties.*

(i) Edge stabilisers are maximal infinite cyclic

(ii) Vertex stabilisers are precisely the subgroups Fix φ*, where $\varphi \in \Phi$ has a fixed subgroup of rank at least 2.*

Since inner automorphisms only fix infinite cyclic groups, and as any vertex stabiliser, H, has rank at least 2, then the corresponding automorphism $\varphi \in \Phi$, such that $H = \text{Fix } \varphi$, is uniquely defined.

Note that if we take two vertices of T in the same orbit, then their stabilisers are conjugate, and the corresponding automorphisms are isogredient. (In general, having conjugate fixed subgroups is not enough to imply isogredience, but it is when the fixed subgroup has rank at least 2). Conversely, if two vertices are in different orbits then the corresponding automorphisms are not isogredient, since edge stabilisers are cyclic.

Given a Dehn Twist, Φ and a $\varphi \in \Phi$, one can construct the free-by-cyclic group, $G = F_n \rtimes_\varphi \langle s \rangle$ (one can do this for any free-by-cyclic group, and the group does not depend on the choice of φ). Using the parabolic orbits Theorem, one gets that G acts on T, with the following properties,

(i) The induced action of s on the quotient of T by F_n is trivial,

(ii) The G edge stabilisers are maximal $\mathbb{Z}^2$,

(iii) The G vertex stabilisers are $F_k \times \mathbb{Z}$, for $k \geq 2$,

(iv) The element sg fixes a vertex of T if and only if $\varphi \, \text{Ad}(g)$ has a fixed subgroup of rank at least 2 (equivalently, if sg has non-abelian centraliser).

3 Extending actions to the automorphism group

3.1 Canonical actions and nearly canonical actions

Since outer automorphisms of free groups often have a power that is better understood (for us, usually a UPG power of a polynomially growing outer automorphism) it can be easier to work with the free-by-cyclic group defined by this power, which is a finite index subgroup of G. However, this means understanding how the automorphisms of a group and a finite index subgroup relate. In general this is hard: recall that the Klein bottle group has a finite outer automorphism group, but contains $\mathbb{Z}^2$ as a finite index subgroup.

Recall that, by Theorem 2.2.1, the action of G on a tree, T, is encoded by its translation length function, l_T. Our strategy is to show that $\mathrm{Aut}^T(G)$, the subgroup of automorphisms preserving the tree (or, equivalently, length function) is finitely generated. Thus we need to find a tree T, such that $\mathrm{Aut}^T(G)$ is either equal to $\mathrm{Aut}(G)$ or is a finite index subgroup of it.

However, the proof of one of our key Lemmas (Lemma 3.1.4) requires us to work with the actions directly rather than via length functions. Therefore we make the following definitions:

Definition 3.1.1. An action of a group G on a tree, T, is called *canonical* if there exists a commuting diagram:

$$\begin{array}{ccc} \mathrm{Aut}(G) & \longrightarrow & \mathrm{Isom}(T) \\ {\scriptstyle g\mapsto \mathrm{Ad}(g)}\uparrow & \nearrow & \\ G. & & \end{array}$$

In the case where G is finitely generated and the G-action is minimal and does not preserve an end, this is equivalent – by Theorem 2.2.1 – to translation length function being preserved by all of $\mathrm{Aut}(G)$, that is $\mathrm{Aut}^T(G) = \mathrm{Aut}(G)$ (and $\mathrm{Out}^T(G) = \mathrm{Out}(G)$).

Definition 3.1.2. We say that an action of G on a tree, T, is called *nearly canonical* if there is a finite index subgroup, $\mathrm{Inn}(G) \leq A \leq \mathrm{Aut}(G)$ such that the following diagram commutes:

$$\begin{array}{ccc} A & \longrightarrow & \mathrm{Isom}(T) \\ {\scriptstyle g\mapsto \mathrm{Ad}(g)}\uparrow & \nearrow & \\ G. & & \end{array}$$

In the case where G is finitely generated and the G-action is minimal and does not preserve an end, this is equivalent to the translation length function being preserved by a finite index subgroup of automorphisms; that is $\mathrm{Aut}^T(G)$ is a finite index subgroup of $\mathrm{Aut}(G)$.

Remark 3.1.3. We are not aware of this last definition in the literature, but it is clearly useful. Also, we have avoided calling this *virtually canonical* since it would raise the confusion between what we mean, and canonical for a finite index subgroup of G. (See Proposition 3.1.4).

We are able to extend a canonical action of a normal finite index subgroup to a nearly canonical action of the whole group, as shown in the following proposition.

Proposition 3.1.4. *Let G be a finitely generated group, G_0 a normal finite index subgroup of G, and suppose that T is a canonical G_0-tree. Then*

(i) G acts on T, and this action restricts to the canonical G_0-action.

(ii) With this action, T is nearly canonical as a G-tree.

Proof. The hypotheses tell us that the action of G_0 on T factors through an action of $\mathrm{Aut}(G_0)$ on T:

$$\begin{array}{ccc} \mathrm{Aut}(G_0) & \longrightarrow & \mathrm{Isom}(T) \\ {\scriptstyle g\mapsto \mathrm{Ad}(g)}\uparrow & \nearrow & \\ G_0 & & \end{array}$$

We let A denote the subgroup of $\mathrm{Aut}(G)$ which preserves G_0 setwise. The restriction map (which in general is neither injective nor surjective) gives us a homomorphism from A to $\mathrm{Aut}(G_0)$, and so A acts on T via this map (and the previous action of $\mathrm{Aut}(G_0)$).

$$A \xrightarrow{res} \mathrm{Aut}(G_0) \longrightarrow \mathrm{Isom}(T)$$

Since G_0 is normal, $\mathrm{Inn}(G)$ is a subgroup of A, and so this action defines an action of G on T.

$$\begin{array}{ccccc} A & \longrightarrow & \mathrm{Aut}(G_0) & \longrightarrow & \mathrm{Isom}(T) \\ {\scriptstyle g\mapsto \mathrm{Ad}(g)}\uparrow & & & & \\ G & & & & \end{array}$$

In particular, with respect to this action, $A \leq \mathrm{Aut}^T(G)$. Moreover, since G_0 is finite index in G, A is a finite index subgroup of $\mathrm{Aut}(G)$ and hence the action of G on T is nearly canonical.

Finally, this action of G on T extends the original action of G_0 since the following diagram commutes (the left two maps from G_0 are just the maps sending a group element to the inner automorphism it defines, and the rightmost map is the one given by the original action of G_0):

$$\begin{array}{ccccc} A & \longrightarrow & \mathrm{Aut}(G_0) & \longrightarrow & \mathrm{Isom}(T) \\ \uparrow & \nearrow & & \nearrow & \\ G_0 & & & & \end{array}$$

□

Remark 3.1.5. One can clearly weaken the hypothesis in the Proposition above so that T is only nearly G_0 canonical, and essentially the same proof works. However, the normality of G_0 seems essential to get a G-action. If G_0 were not normal, one could pass to a further finite index subgroup, H, of G_0 which would be normal in G. But then the action of H on T has no reason to be canonical or nearly canonical. The example of $\mathbb{Z}^2$ in the Klein bottle group shows that passing to a finite index subgroup is not a benign process from this point of view.

3.2 Automorphisms which preserve a splitting, and a theorem of Bass–Jiang

Our proof strategy is to use trees of cylinders to produce a tree where enough of the (outer) automorphisms act, and then to analyse that subgroup. (There are some short-cuts when the defining automorphism is exponentially growing, and we do not have to do all the work ourselves.)

There is a thorough discussion of the structure of the group $\mathrm{Out}^T(G)$ of outer automorphisms that preserve an action on a tree in [2].

We recall below the main structural theorem of that paper. Note though that to save on notation we do not state the result in full. (To be precise, their result allows for a centre, although the filtration becomes a step longer. Also, they give a precise description of the quotients at (4) and (5).)

Theorem 3.2.1 ([2, Theorem 8.1]). *Suppose a centreless group G acts on a tree T, minimally and irreducibly. Write Γ for the quotient graph, and $\mathcal{O}$ for a (fixed) choice of orientation of the edges of Γ. Suppose $\mathrm{Out}^T(G)$ is the subgroup of $\mathrm{Out}(G)$ which acts on T – that is, preserves the length function of the action. Then there is a filtration of $\mathrm{Out}^T(G)$,*

$$\mathrm{Out}^T(G) \trianglerighteq \mathrm{Out}_0^T(G) \trianglerighteq \mathcal{T}^+ \trianglerighteq H \trianglerighteq K \trianglerighteq 1$$

The quotients at each stage are as follows:

$$\mathrm{Out}^T(G)/\,\mathrm{Out}_0^T(G) \leq \mathrm{Aut}(\Gamma) \tag{1}$$

$$\mathrm{Out}_0^T(G)/\mathcal{T}^+ \cong {\prod_{v\in V(\Gamma)}}' \mathrm{Out}(G_v; \{G_e\}_{\iota(e)=v}) \tag{2}$$

$$\mathcal{T}^+/H \cong \prod_{e\in\mathcal{O}} \frac{\mathrm{Ad}(N_e)\cap \mathrm{Ad}(N_{\bar{e}})}{\mathrm{Inn}(G_e)} \tag{3}$$

$$H/K \quad \textit{is a quotient of} \prod_{e\in E(\Gamma)} C_{G_{\iota(e)}}(G_e) \tag{4}$$

$$K \quad \textit{is a quotient of} \prod_{e\in\mathcal{O}} Z(G_e) \tag{5}$$

The "prime" on the product at (2) indicates that it is restricted to elements where for every incident edge e_0 the induced outer automorphism of G_{e_0} is also induced by an element in $\mathrm{Out}(G_{\tau(e_0)}; \{G_e\}_{\iota(e)=\tau(e_0)})$.

This property is characterised by the following commutative diagram. Suppose (Θ_v), with v ranging through the vertices of Γ, is an element of the product

$$\prod_{v \in V(\Gamma)} \operatorname{Out}(G_v; \{G_e\}_{\iota(e)=v}).$$

Then (Θ_v) is an element of the restricted product if and only if for every edge e, with $v = \iota(e)$ and $w = \tau(e)$ there are representatives θ_v and θ_w of the relevant outer automorphisms (of G_v and G_w), and an automorphism ψ of G_e so that both squares commute.

$$\begin{array}{ccccc} G_v & \xleftarrow{\alpha_e} & G_e & \xrightarrow{\alpha_{\bar{e}}} & G_w \\ {\scriptstyle\theta_v}\big\downarrow & & \big\downarrow{\scriptstyle\psi} & & \big\downarrow{\scriptstyle\theta_w} \\ G_v & \xleftarrow[\alpha_e]{} & G_e & \xrightarrow[\alpha_{\bar{e}}]{} & G_w \end{array}$$

There is another exposition in [21], from where we have borrowed some notation (for example, $\mathcal{T}^+$ is Levitt's bitwists).

Our common strategy for the polynomially growing case is to construct a canonical tree – possibly only truly canonical for a finite index subgroup – as a tree of cylinders, and then use this theorem to analyse the automorphisms which preserve it.

By Lemma 2.1.2 the quotient graph for the action must be finite, and so the quotient at (1) will be finite in every case. The quotient at (2) contains the McCool groups, which are generally easier to analyse.

The following lemma relates the restricted product at (2) in Theorem 3.2.1 to the McCool groups (see Definition 2.1.1) for the vertex groups with respect to their incident edge groups. It is analogous to part of Proposition 2.3 of [21] which deals with the case where $\operatorname{Out}(G_e)$ is finite.

Lemma 3.2.2. *The product* $\prod_{v \in V(\Gamma)} \operatorname{Mc}(G_v; \{G_e\}_{\iota(e)=v})$ *is a normal subgroup of the restricted product* $\prod_{v \in V(\Gamma)}{}' \operatorname{Out}(G_v; \{G_e\}_{\iota(e)=v})$.

The quotient is isomorphic to a subgroup of $\prod_{e \in E(\Gamma)} A_e / \operatorname{Ad}(N_e)$, *where* A_e *is a subgroup of* $\operatorname{Aut}(G_e)$, *every element of which is induced by an automorphism of* G_v. *Further,* $A_e = A_{\bar{e}}$ *for all edge pairs* $\{e, \bar{e}\}$.

Proof. For each edge at a vertex v there is a map from $\operatorname{Out}(G_v; \{G_e\}_{\iota(e)=v})$ to $\operatorname{Aut}(G_e)/\operatorname{Ad}(N_e)$ (note that this is a quotient of $\operatorname{Out}(G_e)$). Assembling them, we get a map to their product, and the kernel of this map consists of those elements induced by conjugations at every vertex; precisely the McCool group $\operatorname{Mc}(G_v; \{G_e\}_{\iota(e)=v})$. The conditions on the initial restricted product amount to requiring that an element induces the same automorphisms on the stabiliser of an edge and its inverse: that is, the automorphisms A_e of G_e and $A_{\bar{e}}$ of $G_{\bar{e}}$ will be the same. (Though note that the quotient

$\mathrm{Aut}(G_e) / \mathrm{Ad}(N_e)$ depends also on the vertex group, and so there is no reason to expect these will be the same for both an edge and its opposite.) □

Our strategy is to prove that the McCool groups are finitely generated, and that the quotient is too, usually by showing that this is true of every subgroup of this product. The details vary and appear in the relevant case.

In most of our cases, the edge groups are virtually abelian (that is, their free part has rank at most 1). In this case, we can understand the quotient at (3) as well.

Proposition 3.2.3. *Suppose G is free by cyclic, and $H \leq G$ is free-by-cyclic and virtually abelian. Then $N_G(H)$ induces a finite subgroup of* $\mathrm{Out}(H)$. *(That is,* $\mathrm{Ad}(G, H) / \mathrm{Inn}(H)$ *is finite.)*

Proof. The hypotheses give us that H must be trivial, $\mathbb{Z}$, $\mathbb{Z}^2$ or the fundamental group of a Klein bottle. In every case except for $\mathbb{Z}^2$, the outer automorphism group is finite and so there is nothing to prove. So suppose $H = \mathbb{Z}^2$. The kernel of the map to $\mathrm{Out}(H)$ contains H, so it is enough to show that H is finite index in $N_G(H)$. $H \cap F_n$ must be infinite cyclic, and we have that $H \cap F_n \leq N_G(H) \cap F_n \leq N_{F_n}(H \cap F_n)$. (Recall that conjugating cannot change the exponent of the stable letter.) Since the leftmost group is finite index in the rightmost group, it is also finite index in the middle group.

In the quotient, both H and $N_G(H)$ have non-trivial image. So the image of H is finite index in the image of $N_G(H)$. The index of H in $N_G(H)$ is the product of these two indices, and is therefore finite too. □

To show that the quotient at (4) is finitely generated, we will show that the centralisers (and therefore any quotient of their product) are finitely generated. The splittings we define for the polynomial case all have edge and vertex groups with finitely generated free part, so we will use the following lemma.

Lemma 3.2.4. *Suppose $H \leq G$ are (finitely generated free)-by-cyclic. Then $C_G(H)$ is finitely generated.*

Proof. Let F_n be the "free part" of G, the kernel of the given map to $\mathbb{Z}$. If $H \cap F_n$ is rank at least two, then $C_G(H) \cap F_n$ is trivial, and so $C_G(H)$ is either trivial or $\mathbb{Z}$. If $H \cap F_n$ is $\mathbb{Z}$, then so is $C_G(H) \cap F_n$, and $C_G(H)$ may be $\mathbb{Z}$ or $\mathbb{Z}^2$. If $H \cap F_n$ is trivial, then $C_G(H) \cap F_n$ consists of those elements in $G \cap F_n$ which are fixed by conjugating by H. As the fixed subgroup of an automorphism of a free group, this is finitely generated (by Theorem 2.4.2). The full centraliser has an additional generator which is a root of the generator of H. □

The centre of a free-by-cyclic group G is isomorphic to $\mathbb{Z}^2$ if and only if G is; $\mathbb{Z}$ if G is virtually free-times-cyclic and not $\mathbb{Z}^2$, and trivial otherwise. So the group given at (5) is a finitely generated abelian group, as are all its quotients.

In the exponential case, G is a one-ended relatively hyperbolic group. We are able to use previous work in the literature ([18]) on canonical JSJ decompositions and the automorphisms that preserve them.

4 Exponential growth

4.1 Relative Hyperbolicity

In this section, we assume that φ is exponentially growing. Then we have access to a very useful fact: the group $G \cong F_n \rtimes_\varphi \mathbb{Z}$ is relatively hyperbolic (see [14, 15, 13]). Several definitions of relative hyperbolicity, together with proofs of their equivalence, can be found in [19], for instance; we do not include one here since we do not work directly with the definition.

Given a free group outer automorphism Φ, say a subgroup P is *polynomially growing* (for Φ) if there is a power m and a representative α of Φ^m so that $P\alpha = P$ and the restriction of α to P is polynomially growing.

Proposition 4.1.1 ([23, Proposition 1.4])**.** *Every non-trivial polynomially growing subgroup is contained in a unique maximal polynomially growing subgroup. Maximal polynomially growing subgroups have finite rank, are malnormal, and there are only finitely many conjugacy classes of them.*

These maximal polynomially growing subgroups are a key ingredient in the relatively hyperbolic structure of a free-by-cyclic group:

Theorem 4.1.2 ([14, 15, 13])**.** *If φ is an automorphism of F_n with at least one exponentially growing element, the semidirect product $F_n \rtimes_\varphi \mathbb{Z}$ is relatively hyperbolic with respect to subgroups of the form $H \rtimes_{\varphi^m\gamma} \mathbb{Z}$, where H is a maximal polynomially growing subgroup, m is the minimum (positive) power of φ which carries it to a conjugate, and γ is the inner automorphism so $\varphi^m\gamma$ preserves H.*

(This collection is sometimes referred to as the "mapping torus" of the collection of maximal polynomially growing subgroups. For each H, that such an m exists is guaranteed since there are only finitely many conjugacy classes of maximal polynomially growing subgroups, and since φ^m sends H to a conjugate, there is an inner automorphism so that the composition preserves H.)

Recall that Lemma 2.1.2 gives that $F_n \rtimes \mathbb{Z}$ is one ended, and so it is one ended relative to any collection of subgroups.

Now we have access to a wide range of technology used in the study of relatively hyperbolic groups. In [18] there is a careful examination of the subgroup $\mathrm{Out}(G;\mathcal{P})$ for relatively hyperbolic groups which are one ended relative to their parabolic subgroups, using JSJ theory and analysing the subgroup of automorphisms which preserves a splitting. We recall enough of their work to make the statements which follow self contained (although the proofs will not be).

There is a JSJ decomposition space over elementary (parabolic and virtually cyclic) subgroups relative to $\mathcal{P}$, which is invariant under $\mathrm{Out}(G;\mathcal{P})$. It contains a canonical JSJ tree, the tree of cylinders of the deformation space, which again is $\mathrm{Out}(G;\mathcal{P})$-invariant. There are four possibilities for vertex stabilisers:

Maximal loxodromic stabilised by an infinite cyclic group

Maximal parabolic stabilised by a maximal parabolic subgroup

Rigid non-elementary and elliptic in every splitting with elementary edge groups and where $\mathcal{P}$ are elliptic

Flexible QH with finite fiber none of the above, in which case they are "quadratically hanging with finite fiber"

In fact, in our case the last possibility cannot occur. In general, these groups map with finite kernel onto an orbifold group, and the incident edge groups are virtually cyclic (and their images are in boundary subgroups). Since we are considering groups which are torsion free, the structure is actually much simpler here. First, the kernel must be trivial, so the group itself is an obifold group. By [12, Lemma 2.4] this is (virtually) free and hence infinitely ended, and therefore cannot occur as a vertex group with the required (virtually) cyclic incident edge groups by Lemma 2.1.2.

The tree is bipartite: one class of vertices is those stabilised by a maximal elementary group, and the other is the rigid vertices. Edge groups are maximal elementary subgroups of the rigid vertex group they embed in.

As Guirardel and Levitt point out, Lemma 3.2 of [31] tells us that when the groups in $\mathcal{P}$ are not themselves relatively hyperbolic, every automorphism permutes the conjugacy classes of the P_i. This is true in our case. Theorem 4.1.4 below concerns $\mathrm{Out}(G;\mathcal{P})$; since this consists of those (outer) automorphisms which *preserve* each of these conjugacy classes, it is a finite index subgroup of $\mathrm{Out}(G)$.

Before we state the theorem, we define the group of twists, a subgroup of automorphisms of G. (See Section 2 of [21] or Subsection 2.6 of [18].)

Definition 4.1.3. Let e be an edge of a graph of groups, and g an element of C_e (the centraliser of G_e in $G_{\iota(e)}$). Define the *twist by g around e* to be the automorphism that:

- if e is separating, so $G = A *_{G_e} B$, conjugates A by g and fixes B (with B corresponds to the factor containing $G_{\iota(e)}$);
- if e is non-separating, so $G = A*_{G_e}$, fixes A and sends the stable letter t to tg.

The *group of twists*, $\mathcal{T}$, is the group generated by all twists.

The group of twists $\mathcal{T}$ is a quotient of the direct product of all C_e, the centralisers of edge groups in adjacent vertex groups. Also, recall that McCool groups are defined in Definition 2.1.1.

Theorem 4.1.4 ([18, Theorem 4.3])**.** *Let G be hyperbolic relative to $\mathcal{P} = \{P_1, \ldots, P_n\}$, with P_i infinite and finitely generated, and assume that G is one-ended relative to $\mathcal{P}$. Then there is a finite index subgroup* $\mathrm{Out}^1(G;\mathcal{P})$ *of* $\mathrm{Out}(G,\mathcal{P})$ *which fits into the exact sequence*

$$1 \to \mathcal{T} \to \mathrm{Out}^1(G;\mathcal{P}) \to \prod_{i=1}^{p} \mathrm{MCG}^0_{T_{can}}(\Sigma_i) \times \prod_j \mathrm{Mc}(P_j;\mathrm{Inc}(P_j)) \to 1$$

where T_{can} is the canonical JSJ decomposition relative to $\mathcal{P}$, $\mathcal{T}$ is its group of twists; $\mathrm{MCG}^0_{T_{can}}(\Sigma_i)$ *relate to flexible vertex groups; and* $\mathrm{Mc}(P_j;\mathrm{Inc}(P_j))$ *is the McCool group of P_j with respect to the incident edge groups. (The product is taken only over those parabolic subgroups which appear as vertex stabilisers in T_{can})*

(Theorem 4.1.4 is derived from Levitt's discussion in [21], together with some analysis of the bitwists, showing that they are all twists, and extended McCool groups that can appear, to deduce that there is a finite index subgroup fitting into this short exact sequence. Compared to the Bass-Jiang approach, they show that the second normal subgroup is just $\mathcal{T}$ and that the first quotient has a finite index subgroup isomorphic to right hand term above.)

In our case there are no flexible vertex groups, so that term does not appear. We will use this theorem to prove finite generation for $\mathrm{Out}^1(G;\mathcal{P})$, which will give us finite generation of $\mathrm{Out}(G)$. This will follow from showing that the group of twists and the McCool groups which can appear are finitely generated.

Levitt in [23] provides several inequalities relating invariants of an outer automorphism. Theorem 4.1 of that paper concerns the ranks of conjugacy classes of maximal polynomially growing subgroups for an automorphism of F_n and gives that it is at most $n-1$ when the automorphism is exponentially growing (since there is at least one exponentially growing stratum in this case).

Proposition 4.1.5. *The group of twists is finitely generated.*

Proof. The group of twists is a quotient of the direct product of the centralisers of the edge groups in the vertex groups so it is enough to show that all of these are finitely generated. The edge and vertex groups have the structure of a free-by-cyclic group: say the vertex group is $V = F \rtimes \langle t^k g \rangle$, and the edge group is $E = H \rtimes \langle t^\ell h \rangle$, where $H = E \cap F$. (Note that H and F are not necessarily finitely generated but are subgroups of the defining free group, which is.) If H has rank at least two, then its centraliser in F is trivial, and so the centraliser of E in V is at most infinite cyclic. If H is infinite cyclic, then so is its centraliser in F; then the whole centraliser is either $\mathbb{Z}$ or $\mathbb{Z}^2$.

The final case is where H is trivial, so we are interested only in the centraliser of $t^\ell h$. Again, it will be sufficient to show that the centraliser in F is finitely generated, since there is at most one more generator contributed from the "cyclic part" to the full centraliser. The argument is different at rigid and maximal elementary vertex groups.

First consider rigid vertex groups. Since conjugating by $t^\ell h$ induces the automorphism $\varphi^\ell \operatorname{Ad}(h)$, any w in F that commutes with $t^\ell h$ is fixed by $\varphi^\ell \operatorname{Ad}(h)$. In particular, it is polynomially growing for the outer automorphism Φ. This implies that $\langle w, t^\ell h \rangle$ is an elementary subgroup. Since edge groups are maximal elementary in rigid vertex groups, this cannot happen and so there is no such w. (For the same reason, there is no root of $t^\ell h$.)

At maximal elementary vertices, the free part of the centraliser is the fixed subgroup for the automorphism of F induced by conjugating by $t^\ell h$ (again, conjugation induces the automorphism $\varphi^\ell \operatorname{Ad}(h)$, so any element of F that commutes is fixed by this automorphism). Since F is finitely generated (as a maximal polynomially growing subgroup), so is this fixed subgroup (in fact the rank is bounded by the rank of F; Theorem 2.4.2). □

Thus far what we have said is true for any finitely generated free group; but we do not (yet) have the tools to understand McCool groups of free-by-cyclic groups in general. So we specialise to F_3, for the sake of Theorem 1.1.2.

In this case, the bounds on polynomially growing subgroups mean they can have rank at most 2. Here we can analyse the McCool groups, since there is a good classification of the outer automorphism groups for rank 2 in [6], and rank 1 is fairly easy to understand.

Proposition 4.1.6. *Suppose $G = F_n \rtimes \mathbb{Z}$, with $n = 1, 2$. Let $\mathcal{H}$ be a finite collection of finitely generated subgroups of G. Then* $\operatorname{Mc}(G; \mathcal{H})$ *is finitely generated.*

Proof. In rank 1 the outer automorphism groups are $\mathrm{GL}_2(\mathbb{Z})$ or finite, and in both cases this subgroup must be finitely generated. (For $\mathbb{Z}^2$ notice that elements are their own conjugacy classes, and if g is fixed, so is its root, and so after changing basis the only matrices in the subgroup are triangular, and so it is virtually cyclic.)

In rank 2, we refer to [6, Theorem 1.1] for their outer automorphism groups. Most cases are either finite or virtually cyclic: so any subgroup is finitely generated. The remaining cases are $G = F_2 \times \mathbb{Z}$, and $G = F_2 \rtimes_{-I_2} Z$.

In the first of these, we have that $\mathrm{Out}(G) = (\mathbb{Z}^2 \rtimes C_2) \rtimes \mathrm{GL}_2(\mathbb{Z})$ [6, Theorem 1.1(i)]. Since $\mathrm{GL}_2(\mathbb{Z})$ preserves each of the first two factors, we may pass to a finite index subgroup that is $\mathbb{Z}^2 \rtimes \mathrm{GL}_2(\mathbb{Z})$. (An element $u \in \mathbb{Z}^2$ acts by sending $t^k g \to t^{k+u\cdot g_{ab}} g$, and $\mathrm{GL}_2(\mathbb{Z})$ on the free part as you might expect.) Now consider a set of finitely generated subgroups $\mathcal{H}$.

Since t is central, its exponent cannot be changed by inner automorphisms. So any element of the McCool group must fix the t-exponent in each generator: this will give a subgroup of $\mathbb{Z}^2$ (orthogonal to the abelianised free parts of the generators) which is therefore finitely generated. So our McCool group is finitely generated if and only if its intersection with $\mathrm{Out}(F_n)$ is. In fact, this intersection is exactly the McCool group for the free part: since t is central, it cannot identify any conjugacy classes of F_n. These are finitely generated by [29], which completes the proof. Note that McCool proves the result for elements; however in the free group case and more generally for toral relatively hyperbolic groups [17, Corollary 1.6] the McCool group for a finite set of subgroups is equal to the McCool group for some finite set of elements.

For $F_2 \rtimes_{-I_2} \mathbb{Z}$, the outer automorphism group is $\mathrm{PGL}_2(\mathbb{Z}) \times C_2$ [6, Theorem 1.1(ii)]. Again, we can just consider the finite index subgroup $\mathrm{PGL}_2(\mathbb{Z})$, which only acts on the free part. We can consider the McCool group for the free group (as a subgroup of $\mathrm{GL}_2(\mathbb{Z})$). Its image in $\mathrm{PGL}_2(\mathbb{Z})$ is a finite index subgroup of the subgroup we want, which is therefore finitely generated. □

We now summarise this case in a theorem.

Theorem 4.1.7. *Suppose $G \cong F_3 \rtimes_\varphi \mathbb{Z}$, and φ is exponentially growing. Then* $\mathrm{Out}(G)$ *is finitely generated.*

Proof. Use the canonical tree and the analysis of the outer automorphisms derived from it in Theorem 4.1.4. Propositions 4.1.5 and 4.1.6 show that the outside groups in the short exact sequence are finitely generated, and therefore so is $\mathrm{Out}_1(G;\mathcal{P})$ which is a finite index subgroup of $\mathrm{Out}(G)$. □

5 Linear growth

5.1 Strategy

Our strategy for showing that the automorphism group of a free-by-cyclic group, in the case of linear growth, is as follows.

- Start with a free-by-cyclic group, $G = F_n \rtimes_\Phi \mathbb{Z}$, where Φ has linear growth,
- Consider a finite index subgroup, $G_0 = F_n \rtimes_{\Phi^r} \mathbb{Z}$, so that Φ^r is UPG, and hence a Dehn Twist
- Use the parabolic orbits Theorem to find a tree whose deformation space is invariant,
- Deduce that the tree of cylinders, $T = T_c$, of this space is G_0-canonical,
- Use Proposition 3.1.4 to deduce that T is nearly G-canonical
- Show that $\mathrm{Out}^T(G)$ is finitely generated if certain McCool groups for free-by-finite groups are
- Carry out the calculation of the relevant McCool groups, to conclude that $\mathrm{Out}^T(G)$ is finitely generated.

5.2 Constructing a tree

First we record a useful lemma on normalisers in free-by-cyclic groups.

Lemma 5.2.1. *Suppose $F_n \rtimes \langle s \rangle$ is a free-by-cyclic group, and $w \in F_n$ is not a proper power and commutes with s. Then $\langle w, s \rangle$ is its own normaliser.*

Proof. Suppose $s^k g \in F_n \rtimes \langle s \rangle$ so that $\langle w, s \rangle^{s^k g} = \langle w, s \rangle$. This gives that $\langle w^g, s^g \rangle = \langle w, s \rangle$. Taking intersections with F_n, we must have that $w^g \in \langle w \rangle$. But this means that $g \in \langle w \rangle$ so $s^k g \in \langle w, s \rangle$ as required. □

In the following Proposition, we take a Dehn Twist and use the Parabolic Orbits Theorem 2.4.7 to get a tree on which the corresponding free-by-cyclic group acts. We would like, at this stage, to say that the resulting action is canonical for the free-by-cyclic group. Although this seems plausible, our proof goes via the tree of cylinders construction which is guaranteed to be canonical and – as we prove in this case – remains in the same deformation space.

Proposition 5.2.2. *Suppose φ is a UPG and linear automorphism of F_n. Then there is a canonical action of $G_0 = F_n \rtimes_\varphi \mathbb{Z}$ on a tree, where*

(1) Edge stabilisers are maximal $\mathbb{Z}^2$;

(2) Vertex stabilisers are either maximal $\mathbb{Z}^2$, or maximal $F_m \times \mathbb{Z}$ with $n \geq m \geq 2$.

Proof. The initial input for the construction is the Dehn twist, φ. By Theorem 2.4.7, there is a unique simplicial F_n-tree (defining a simplex in the boundary of CV_n) that is preserved by φ. This tree gives a splitting of F_n, where the vertex stabilisers are fixed subgroups (of rank at least two) corresponding to different representatives of the outer automorphism, and the edge groups are maximal infinite cyclic. By Theorem 2.4.2 there are only finitely many conjugacy classes of these subgroups, and their ranks are bounded by n.

Since it is fixed by φ, the same tree provides a splitting for G_0. The vertex groups are now free times cyclic, and the edge groups are maximal $\mathbb{Z}^2$. (They are generated by the original edge group generator g, together with an element sw in either adjacent edge group which commutes with g. They must be maximal since otherwise there would be another element $s^k h$ commuting with g (and sw); writing this element as $(sw)^k h'$ implies that h' commutes with g. Since g generated a maximal infinite cyclic subgroup of F_n, h' is a power of g, and so $(sw)^k h'$ is contained in $\langle g, sw \rangle$.)

This tree defines a deformation space which is preserved by automorphisms, since the vertex stabilisers can be specified algebraically: they are precisely the centralisers of some sw, corresponding to an automorphism in the outer automorphism class of φ with fixed subgroup having rank at least 2. (Equivalently, they are the centralisers that contain a copy of $F_2 \times \mathbb{Z}$.) So they will be permuted by automorphisms of $F_n \rtimes_\varphi \mathbb{Z}$ and the deformation space must be preserved.

We now have most of the tools to start constructing a tree of cylinders for this deformation space: it remains to specify the family $\mathcal{E}$ of allowed edge stabilisers, and the admissible equivalence relation. We will take $\mathcal{E}$ to be maximal $\mathbb{Z}^2$, and the equivalence relation to be equality. (It is easy to check this is admissible, since if $A \leq B$ are both maximal $\mathbb{Z}^2$ then we must have $A = B$).

Now we can calculate the cylinders. First, note that a cylinder may contain at most one edge from each edge orbit. If two edges in the same orbit have the same stabiliser, then there is an element outside the stabiliser which normalises it. However, Lemma 5.2.1 shows that there is no such element.

This also means that a cylinder stabiliser must actually stabilise it pointwise: since it is a subgroup of G_0, it cannot permute edges in different orbits. So cylinder stabilisers are precisely the stabiliser of any (and every) edge in that cylinder. Every vertex is in multiple cylinders, so is also in the tree of cylinders.

Cylinders are finite, and in particular bounded, so the tree of cylinders will lie in the same deformation space. It is already collapsed, since the edge stabilisers are still (maximal) $\mathbb{Z}^2$. □

Remark 5.2.3. Note that an alternative construction of this canonical tree involves subdividing every edge and folding – the effect of constructing the tree of cylinders is to change to original tree so that each vertex has at most one adjacent edge with a given stabiliser. There are examples where the tree of cylinders is not very small – it has tripod stabilisers, so the construction has done something.

However, the (finite index) subgroup of automorphisms which does not permute the underlying graph of groups does act on the original limiting tree, since we can recover it by equivariantly collapsing some edges. This means that in our terminology the action on the limiting tree was itself nearly canonical, though it is not clear how to find a direct proof of this fact.

If a cylinder had only one edge, then it will have been subdivided – allowing (if the endpoints are isomorphic) for the possibility of inversions. (If not, or if the endpoints are not isomorphic, no inversions are possible.)

We now equip ourselves with a nearly canonical action for a general linearly growing automorphism , using this tree of cylinders.

Proposition 5.2.4. *Suppose $G = F_n \rtimes_\Phi \mathbb{Z}$ is a free-by-cyclic group, and Φ is linearly growing. Then G has a nearly canonical action on a tree T, where*

(1) Edge stabilisers are virtually $\mathbb{Z}^2$ (and therefore either $\mathbb{Z}^2$ or the fundamental group of a Klein bottle).

(2) Vertex stabilisers are $F_m \rtimes_\varphi \mathbb{Z}$ where F_m is a subgroup of F_n, the rank m is at most n, and φ is a representative of Φ, which restricts to and is periodic on F_m. (They are virtually free-times-cyclic.)

Proof. Since Φ has a power which is UPG, and therefore a Dehn twist, we pass to the normal finite index subgroup G_0 this suggests and use Proposition 5.2.2 to construct a canonical tree T. We then use Proposition 3.1.4 to extend this action to a nearly canonical action for G. Edge and vertex stabilisers in G will contain edge and vertex stabilisers in G_0 as finite index subgroups, and must themselves be free-by-cyclic by Lemma 2.1.2. Combining these properties gives the conclusions in (i) and (ii). □

5.3 Reducing to free-by-finite groups

We consider the subgroup $\mathrm{Out}^T(G)$ of outer automorphisms which preserves this tree, and apply Theorem 3.2.1 to understand it. The quotients at parts (1) and (3-5) of the theorem are finitely generated by the observations following the theorem; the main difficulty is in understanding the quotient at (2).

First, we reduce to the case where we can consider McCool groups; we will then show that the result we want is implied by a similar result in the free-by-finite group obtained by quotienting by the centre, and in the next section prove the result there. (The arguments involved in the reduction and the following section are easier for the larger groups $\mathrm{Out}(G_v; \{G_e\}_{\iota(e)=v})$ at least when the edge groups all contain the centre of the vertex group as in our case. However, it does not seem possible to take account of the edge compatibility relations through this process, so we do need to pass to McCool groups.)

We begin with a straightforward structural result about free-by-cyclic groups defined by periodic outer automorphisms;

Lemma 5.3.1 ([24, Proposition 4.1])**.** *Suppose G is a free-by-cyclic group which is virtually free-times-cyclic and not virtually $\mathbb{Z}^2$. Then G has an infinite cyclic centre, and is the fundamental group of a graph of groups with all edge and vertex groups isomorphic to $\mathbb{Z}$.*

Such a group is known as a *Generalised Baumslag-Solitar (GBS) group*, and having a non-trivial centre is equivalent to having trivial modulus, in the language of [22]. The free-by-(finite cyclic) groups we will consider are obtained by taking a group of this kind and quotienting by the centre.

We now study the group appearing as a quotient at (2) in Theorem 3.2.1, beginning by considering automorphisms of edge groups that can be induced here.

Lemma 5.3.2. *Suppose G is a free-by-cyclic group that is virtually free-times-cyclic, and H_i is a collection of subgroups isomorphic either to $\mathbb{Z}^2$ or to the fundamental group of a Klein bottle. Then $\mathrm{Out}(G; \{H_i\})$ induces a virtually cyclic subgroup of $\mathrm{Aut}(H_i)$.*

Proof. If H_i is the fundamental group of a Klein bottle, $\mathrm{Out}(H_i)$ is finite, and $\mathrm{Inn}(H_i)$ is virtually cyclic. Therefore $\mathrm{Aut}(H_i)$ is again virtually cyclic, and so is the subgroup induced by $\mathrm{Out}(G; \{H_i\})$.

If H_i is $\mathbb{Z}^2$, it contains a finite index subgroup of the infinite cyclic centre of G. Let δ generate this subgroup. We can choose a basis $\{x_1, x_2\}$ for H_i so that $\delta = x_1^k$ with $k > 0$; roots are unique in $\mathbb{Z}^2$, so x_1 is as uniquely defined as δ: it is unique up to inverses. Any automorphism of G will preserve the centre; in particular it must send δ to itself or its inverse. So any automorphism restricting to G_e will likewise send x_1 to itself or its inverse. Viewing elements of $\mathrm{GL}(2, \mathbb{Z})$ as matrices, this implies that we can only induce automorphisms represented by triangular matrices. This subgroup is virtually cyclic. □

We use this to characterise the subgroup generated when we quotient by the product of McCool groups, which will mean it is sufficient to prove that those are finitely generated.

Proposition 5.3.3. *Suppose G is a free-by-cyclic group where the defining outer automorphism is linearly growing. Let T be the tree constructed in Proposition 5.2.4, with a nearly canonical action of G (where edge stabilisers are virtually $\mathbb{Z}^2$, and vertex stabilisers are either virtually $\mathbb{Z}^2$ or virtually $F_m \times \mathbb{Z}$ with $m \geq 2$). Then the quotient of $\prod_{v \in V(\Gamma)}{}' \operatorname{Out}(G_v; \{G_e\}_{\iota(e)=v})$ (from Theorem 3.2.1 (2)) by $\prod_{v \in V(\Gamma)} \operatorname{Mc}(G_v; \{G_e\}_{\iota(e)=v})$ (as described in Lemma 3.2.2) is finitely generated.*

Proof. We consider the projection to each factor $A_e / \operatorname{Ad}(N_e)$. The subgroup we are interested in is contained in the product of these projections, which we will show is slender, and from there deduce that our subgroup must be finitely generated.

First, we consider the vertices where the stabiliser contains a rank 2 free group. In this case, by Lemma 5.3.2 each of these vertex groups can only induce a virtually cyclic subgroup of automorphisms of each edge group. This is a property closed under subgroups and quotients, so for every edge e with $\iota(e)$ a vertex of this type the projection to $A_e / \operatorname{Ad}(N_e)$ is virtually cyclic.

The remaining vertices arose as cylinders, and their vertex groups are either the fundamental group of a Klein bottle or $\mathbb{Z}^2$ (as are the incident edge groups). If G_v is a Klein bottle, then it has finite outer automorphism group. So $\operatorname{Out}(G_v; \{G_e\}_{\iota(e)=v})$ is finite, and $\operatorname{Ad}(N_e)$ must therefore be finite index in A_e for each edge group. So the projection to $A_e / \operatorname{Ad}(N_e)$ for edges starting at these vertices is finite.

If G_v is $\mathbb{Z}^2$, we need to use the structure of the tree. The quotient graph inherits the bipartite structure of the tree of cylinders constructed in Proposition 5.2.2 – every edge joins a cylinder vertex to a vertex with larger stabiliser. By Lemma 3.2.2 the induced automorphisms A_e and $A_{\bar{e}}$ of the stabilisers of an edge and its inverse are the same. By Lemma 5.3.2 this is virtually cyclic, and so the same is true of the projection to $A_e / \operatorname{Ad}(N_e)$ in this case.

Assembling these projections we get a group that is virtually finitely generated abelian, and in particular is Noetherian. So any subgroup – including the quotient of

$$\prod_{v \in V(\Gamma)}{}' \operatorname{Out}(G_v; \{G_e\}_{\iota(e)=v})$$

by the product of McCool groups – is again finitely generated. □

In the Klein bottle case, the McCool group (as with any subgroup of the outer automorphism group) is finite, and in particular finitely generated. In the $\mathbb{Z}^2$ case the McCool group is trivial since elements of $\mathrm{GL}_2(\mathbb{Z})$ are uniquely characterised by their action on a finite index subgroup of $\mathbb{Z}^2$: as soon as an edge group is fixed, so is the whole vertex group. Therefore the remainder of the work is at the vertices stabilised by some $F_m \rtimes \mathbb{Z}$, with $m \geq 2$.

This reduces the problem to calculating the McCool groups at each vertex. We use Levitt's work in [22] to further reduce the problem to McCool groups of free-by-(finite cyclic) groups.

By Lemma 5.3.1, the larger vertex groups G_v are Generalised Baumslag-Solitar groups with trivial modulus. Levitt proves this theorem, which we use to enable us to understand $\mathrm{Out}(G)$ in terms of the outer automorphisms of a free-by-finite group.

Theorem 5.3.4 (see [22, Theorem 4.4])**.** *Suppose G is a GBS group with trivial modulus, and let H be the quotient of G by its centre. Then there is a finite index subgroup $\mathrm{Out}_0(G)$ of $\mathrm{Out}(G)$ fitting into a split exact sequence*

$$1 \to \mathbb{Z}^k \to \mathrm{Out}_0(G) \to \mathrm{Out}_0(H) \to 1$$

where k is the rank of the underlying graph, and $\mathrm{Out}_0(H)$ is a finite index subgroup of $\mathrm{Out}(H)$. The section of $\mathrm{Out}_0(H)$ fixes the centre of G.

The $\mathbb{Z}^k$ subgroup should be thought of as $\mathrm{Hom}(\pi_1(\Gamma), Z(G))$: it acts by multiplying every "HNN-like generator" by an element of the centre. The subgroup $\mathrm{Out}_0(H)$ consists of (outer classes of) automorphisms which preserve the conjugacy classes of elliptic elements, and the image of a certain map $\overline{\tau}$ to some finite cyclic group.

The map τ is initially defined as a map to $\mathrm{Isom}(\mathbb{R})$, and we then observe that the image is discrete, and so is isomorphic to $\mathbb{Z}$.

This definition does not apply to the "elementary" GBS groups, $\mathbb{Z}^2$ and the fundamental group of a Klein bottle. These are distinguished among free-by-cyclic groups as being virtually $\mathbb{Z}^2$, and this property cannot occur in a free-by-cyclic group with underlying free group having rank at least 2. Since the groups we consider here (corresponding to non-cylinder vertices in the nearly canonical tree) do, this definition (and the following arguments) apply in sufficient generality for our use.

Let δ generate the centre of G, and T be the GBS tree that G acts upon. Recall that if a group acts on a tree without fixing an end, its centre lies in the kernel of the action. In particular, δ is contained in every vertex group. Suppose x_v generates a vertex group, so $x_v^{n_v} = \delta$ for some n_v. Following [22], define τ on x_v as the translation by $1/n_v$. (So δ

is translation by 1.) On generators arising from edges, τ is the identity. Further define $\overline{\tau}$ by taking a quotient by the group generated by $\tau(\delta)$.

In some sense, τ is a "better" map to $\mathbb{Z}$ than the one arising from the presentation of G as a free-by-cyclic group. First consider its kernel:

Lemma 5.3.5. *The kernel of the map τ is a finitely generated free group.*

This follows from the computation of the relevant BNS invariants in [8, Corollary 3.2], but can be proved by more elementary methods as follows:

Proof. Consider the action of the kernel on the GBS tree T. This action is free – non-trivial elements of vertex stabilisers are not in the kernel of τ – and so the kernel is a free group. It remains to show it is finitely generated. To do this consider $\overline{\tau}$, defined by passing to the quotient by $\tau(\delta)$. The kernel of this map is finitely generated (as a finite index subgroup of a finitely generated group), and we claim it is the direct product $\ker(\tau) \times \langle\delta\rangle$. (Notice that δ has non-trivial image under τ, so there is no intersection and this is a direct product.)

Both factors are in the kernel of $\overline{\tau}$, and so too is their product. To show the other inclusion, suppose g lies in $\ker(\overline{\tau})$, so $\tau(g) = \tau(\delta^k)$ for some $k \in \mathbb{Z}$. So $g' = g\delta^{-k}$ lies in $\ker(\tau)$, and we may rewrite $g = g'\delta^k$, as an element of $\ker(\tau) \times \langle\delta\rangle$.

Since $\ker(\tau)$ is a quotient of a finitely generated group, it too is finitely generated. □

This lemma shows that the map τ fibres: it gives us another way to write G as a free-by-cyclic group. Note that the rank of the free group may have changed, but since there is still a centre, the defining outer automorphism must still be periodic. (Sometimes, though not always it becomes periodic as an automorphism – for example, using this construction it becomes apparent that the rank three free-by-cyclic group defined using the automorphism $a \mapsto b^{-1}c, b \mapsto a^{-1}c, c \mapsto c$ is isomorphic to $F_2 \times \mathbb{Z}$.)

By design, this new presentation as a free-by-cyclic group is very well behaved when applying τ: the image under τ of any element is the exponent of the (new) stable letter. This exponent is preserved by conjugation, and (by considering the stable letter as a root of δ) by the section of $\mathrm{Out}_0(H)$. So if an automorphism whose outer class is an element of $\mathrm{Out}_0(G)$ does not preserve the exponent on the stable letter, writing it in the normal form for a semidirect product will involve a non-trivial element of the $\mathbb{Z}^k$ subgroup given in the decomposition of Theorem 5.3.4. Note that since the exponent is preserved by conjugation, this effect is constant across an outer class.

Proposition 5.3.6. *Suppose that G is a free-by-cyclic group that is virtually free-times-cyclic, and $\{G_i\}$ is a family of subgroups. Write H for the quotient of G by its centre, and let H_i be the* *image of the subgroup G_i under this quotient map. Let s be the section of Theorem 5.3.4. Then*

$$\mathrm{Out}_0(G) \cap \mathrm{Mc}(G, \{G_i\}) = (\mathbb{Z}^k \cap \mathrm{Mc}(G, \{G_i\})) \cdot s(\mathrm{Out}_0(H) \cap \mathrm{Mc}(H, \{H_i\})).$$

In particular, it is finitely generated if and only if $(\mathrm{Out}_0(H) \cap \mathrm{Mc}(H, \{H_i\}))$ *is.*

Proof. Consider the split short exact sequence of Theorem 5.3.4. The kernel of the quotient map applied to $\mathrm{Out}_0(G) \cap \mathrm{Mc}(G, \{G_i\})$ will be $(\mathbb{Z}^k \cap \mathrm{Mc}(G, \{G_i\}))$, a finitely generated abelian group. The image of $\mathrm{Out}_0(G) \cap \mathrm{Mc}(G, \{G_i\})$ in $\mathrm{Out}_0(H)$ is contained in $\mathrm{Mc}(H, \{H_i\})$, since subgroups which are conjugate in G have conjugate images in H. To complete the proof, we must show that this accounts for all of $(\mathrm{Out}_0(H) \cap \mathrm{Mc}(H, \{H_i\}))$. To do this consider the section: we need to show that every element lifts to an element of $\mathrm{Out}_0(G) \cap \mathrm{Mc}(G, \{G_i\})$.

For an element of $(\mathrm{Out}_0(H) \cap \mathrm{Mc}(H, \{H_i\}))$ consider the collection of representatives α_i, each fixing the subgroup H_i. Lemma 4.1 of [22] constructs the equivalent section for automorphism groups; one of the properties of the lift $\overline{\alpha}$ of α is that applying $\overline{\alpha}$ first does not alter τ. So if α_i fixes h, and g is any preimage of h, $\overline{\alpha}_i$ must send g to $g\delta^k$. However, since τ must be unaltered, in fact $k = 0$ and g is fixed. So each $\overline{\alpha}_i$ fixes the subgroup G_i. The last thing to check is that they all represent the same outer automorphism. This is the case since inner automorphisms lift to inner automorphisms (by any preimage of the conjugator, as they differ by a central element). So any (indeed every) $\overline{\alpha}_i$ represents an element of $\mathrm{Mc}(G, \{G_i\})$, which is contained in $\mathrm{Out}_0(G)$ since it is the image of the section of Theorem 5.3.4. □

So to show that the McCool groups we are interested in are finitely generated, we need to show the same for the relevant McCool groups of free-by-finite groups. In our situation the edge groups are virtually $\mathbb{Z}^2$, and a power of a generator is central in the vertex group, so in H the image of each edge group becomes virtually infinite cyclic. In this case, we can understand the McCool groups.

5.4 McCool groups for free-by-finite groups

The purpose of this section is to study the groups $\mathrm{Mc}(H, \{H_i\})$, which will complete our proof in the linear growth case.

Proposition 5.4.1. *Suppose H is virtually free and $\{H_i\}$ is a finite collection of virtually infinite cyclic subgroups. Then* $\mathrm{Mc}(H, \{H_i\})$ *is finitely generated.*

First we use a result which allows us to understand the outer automorphisms of the extension by considering the centraliser of the finite cyclic subgroup.

Proposition 5.4.2. *Let H be a group, and F a normal subgroup of H with trivial centre. Let* $\mathrm{Ad}(h)$ *represent the automorphism of F induced by conjugating by h. Let* $\mathrm{Aut}_H(F)$ *be the subgroup of* $\mathrm{Aut}(F)$ *that commutes with* $\mathrm{Ad}(H)$ *up to inner automorphisms. That, is the subgroup defined by*

$$\{\alpha \in \mathrm{Aut}(F) | [\mathrm{Ad}(h), \alpha] \in \mathrm{Inn}(F) \quad \forall h \in H\}.$$

Further, let N be the subgroup of $\mathrm{Aut}(H)$ *which preserves F and all its cosets (that is, it acts trivially on the quotient H/F).*

Then the restriction to F sends N isomorphically to $\mathrm{Aut}_H(F)$.

Proof. First we consider the image of the restriction map. Suppose α is an element of N, and consider its restriction to F. For all $f \in F$, and $h \in H$, $f(\mathrm{Ad}(h)\alpha) = (f^h)\alpha = (f\alpha)^{(h\alpha)} = f(\alpha\,\mathrm{Ad}(h\alpha))$. This gives that, as automorphisms of F, $\mathrm{Ad}(h\alpha) = \alpha^{-1}\,\mathrm{Ad}(h)\alpha$. Since α preserves cosets of F, $h^{-1}(h\alpha) = f$, for some $f \in F$. But then $\mathrm{Ad}(f) = \mathrm{Ad}(h^{-1}(h\alpha)) = \mathrm{Ad}(h)^{-1}\alpha^{-1}\,\mathrm{Ad}(h)\alpha$: the restriction of α to F satisfies the commutator property defining $\mathrm{Aut}_H(F)$, and so the image in $\mathrm{Aut}(F)$ lies in this subgroup.

Next we show that the restriction map is a surjection to $\mathrm{Aut}_H(F)$. To do this, we construct an automorphism of H with a given image in $\mathrm{Aut}_H(F)$. For any $\alpha \in \mathrm{Aut}_H(F)$, we have $\alpha^{-1}\,\mathrm{Ad}(h)\alpha = \mathrm{Ad}(h)\,\mathrm{Ad}(f_{h,\alpha})$ by the defining commutator property, where $f_{h,\alpha}$ is an element of F depending on both h and α. Since F is centreless, it has a unique element inducing any inner automorphism – $f_{h,\alpha}$ is well defined. Extend α to a function $\overline{\alpha}$ defined on all of H by setting $h\overline{\alpha}$ to be $hf_{h,\alpha}$. (On F, since $\alpha^{-1}\,\mathrm{Ad}(h)\alpha = \mathrm{Ad}(h\alpha)$ for inner automorphisms, $h\alpha = hf_{h,\alpha}$, so the restriction is indeed α.) To see $\overline{\alpha}$ is an endomorphism, we need to check that $hkf_{hk,\alpha} = hf_{h,\alpha}kf_{k,\alpha}$.

Consider the following diagram: the squares all commute by the definition of $\mathrm{Aut}_H(F)$, the left hand triangle is a consequence of Ad being a homomorphism, and we are interested in the right hand triangle, whose commutativity follows from chasing the diagram (noting that the top map is an isomorphism). This gives that $\mathrm{Ad}(hkf_{hk,\alpha}) = \mathrm{Ad}(hf_{h,\alpha}kf_{k,\alpha})$, and by normality of F this is equal to $\mathrm{Ad}(hkf_{h,\alpha}^k f_{k,\alpha})$. That is, we have that the unique element of F inducing the correct inner automorphism is $f_{hk,\alpha} = f_{h,\alpha}^k f_k$, with which we get that $hkf_{hk,\alpha} = hf_{h,\alpha}kf_{k,\alpha}$.

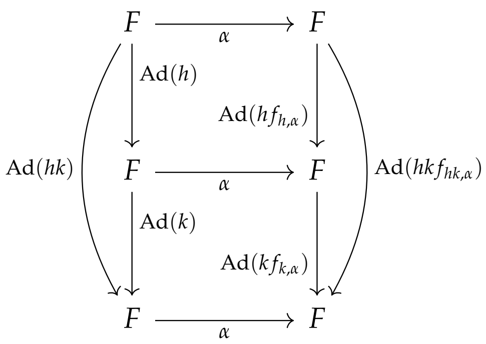

To see $\overline{\alpha}$ is surjective, note that it is surjective on F, and for general h, we have $h = (h(f_{h,\alpha}^{-1}\alpha^{-1}))\overline{\alpha}$. To see injectivity, suppose $h\overline{\alpha} = 1$, so $hf_{h,\alpha} = 1$. In particular, this means h is an element of F; but on F $\overline{\alpha}$ agrees with α, which is an automorphism. So $h = 1$, and $\overline{\alpha}$ is an element of $\mathrm{Aut}(H)$, restricting to α on F as claimed.

Finally, we show that the restriction map $N \to \mathrm{Aut}_H(F)$ is injective. Denote by K the kernel of the map $\mathrm{Ad} : H \to \mathrm{Aut}(F)$. Since F has no centre, $K \cap F$ is trivial.

Suppose α lies in the kernel of the restriction map, so it fixes every element of F. Then for all $f \in F, h \in H$, we have that $f^h \in F$, so $f^{h\alpha} = (f\alpha)^{h\alpha} = (f^h)\alpha = f^h$. So the actions of $h\alpha$ and h on F are the same: that is, $h\alpha$ and h lie in the same K-coset.

So for all elements $h \in H$ we have that $(h\alpha)^{-1}h$ lies in K. Since both automorphisms preserve cosets of F, in fact $(h\alpha)^{-1}h$ lies in $F \cap K$. But these groups intersect trivially, so $h\alpha = h$ for all elements h, α must be the identity, and so the restriction map has trivial kernel. □

We now specialise this general result to our current case of virtually free groups.

Corollary 5.4.3. *Let H be a finitely generated virtually free group, that is not virtually cyclic, and F a normal finite index subgroup (with rank at least 2) of H. Then the subgroup $\mathrm{Aut}_H(F)$ of $\mathrm{Aut}(F)$ is isomorphic to a finite index subgroup of $\mathrm{Aut}(H)$ which preserves F, and where the isomorphism is given by the restriction map.*

(This Corollary is similar in spirit to [28], which deals with centralisers in $\mathrm{Aut}(F)$; ours looks at the preimage of centralisers in $\mathrm{Out}(F)$, and deals simultaneously with the splitting and non-splitting cases.)

Proof. By Proposition 5.4.2, the subgroup $\mathrm{Aut}_H(F)$ of $\mathrm{Aut}(F)$ is isomorphic to the subgroup N of $\mathrm{Aut}(H)$. This subgroup preserves F and all its cosets, and the restriction to F provides the isomorphism, as required. To finish the proof, notice that since H is finitely generated and F is a finite index subgroup, N must be finite index in $\mathrm{Aut}(H)$. □

We want not just the outer automorphism group but the McCool group. The relevant result about $\mathrm{Out}(F_n)$ is the following theorem of Bestvina, Feighn and Handel.

Theorem 5.4.4 ([4, Theorem 1.2(3)]). *Suppose Q is a finite subgroup of $\mathrm{Out}(F_n)$, and $\mathrm{Out}_Q(F_n)$ is its centraliser. Let $K_1, \ldots K_n$ be a collection of conjugacy classes of finitely generated subgroups of F_n. Then the subgroup of $\mathrm{Out}_Q(F_n)$ fixing each K_i is VF (in particular, is finitely presented).*

Note that the conclusion we want is stronger: we want the action on a representative of K_i to be by conjugation, not just sending it to a conjugate. However, as the relevant subgroups are infinite cyclic this is only a matter of passing to a finite index subgroup.

These theorems allow us understand the subgroup of outer automorphisms conjugating an element that lies in the finite index free subgroup; to extend the result to the full subgroups H_i, we need the following lemmas.

Lemma 5.4.5. *Suppose A is a virtually cyclic group. Then $\mathrm{Out}(A)$ is finite.*

See, for instance, [30, Lemma 6.6] for a proof. The key fact from this lemma is that the inner automorphisms are finite index, so 'most' automorphisms of a virtually cyclic group are conjugations.

Lemma 5.4.6. *Suppose H is virtually free (of rank at least 2), and let h be a non-trivial element of the finite index free subgroup F. Then $\langle h \rangle$ has finite index in its normaliser, which in particular is virtually cyclic.*

Proof. First consider the intersection $N_H(\langle h \rangle) \cap F$: this is an infinite cyclic group, generated by the root of h (which we denote $\hat{h}$). This contains $\langle h \rangle$ with finite index. But $N_H(\langle h \rangle) \cap F$ itself is a finite index subgroup of $N_H(\langle h \rangle)$, which must again contain $\langle h \rangle$ with finite index. □

We now combine these results to prove Proposition 5.4.1.

Proof of Proposition 5.4.1. Let F be a finite index normal subgroup of H. By Corollary 5.4.3, it will suffice to show that the subgroup of N (the isomorphic image of $\mathrm{Aut}_H(F)$ in $\mathrm{Aut}(H)$) that acts as conjugation on each subgroup H_i is finitely generated. We use A for this subgroup of N.

Each subgroup H_i is virtually $\mathbb{Z}$; in particular its intersection with F is generated by a single element h_i. This intersection is preserved under conjugation by elements of H_i (since F is a normal subgroup of H): in particular H_i is a subgroup of $N_H(\langle h_i \rangle)$. By Lemma 5.4.6, since it contains h_i, it is finite index in this normaliser.

Let $Q = \mathrm{Ad}(H)/\mathrm{Inn}(F) \cong H/FK$ be the subgroup of $\mathrm{Out}(F)$ induced by H, and denote by $\mathrm{Out}_Q(F)$ the centraliser of Q in $\mathrm{Out}(F)$. This is the projection of $\mathrm{Aut}_H(F)$ to $\mathrm{Out}(F)$. By Theorem 5.4.4, the subgroup $\mathrm{Out}_Q(F)$ preserving the conjugacy class of each $\langle h_i \rangle$ is finitely generated, so this is also true of the subgroup A of N Normalisers must be sent to normalisers, so A sends $N_H(\langle h_i \rangle)$ to a conjugate of itself too.

This normaliser is virtually cyclic, so by Lemma 5.4.5 it has finitely many outer automorphisms. After composing with an inner automorphism we induce an automorphism of $N_H(\langle h_i \rangle)$, and we may restrict to those which induce an inner automorphism. This restriction gives a finite index subgroup of A, which acts as a conjugation on $N_H(\langle h_i \rangle)$, and in particular on the subgroup H_i. Repeating this for each subgroup H_i (there are only finitely many) still defines a finite index subgroup, which is itself finitely generated. □

Remark 5.4.7. Notice that the ad-hoc arguments given in Proposition 4.1.6 for the two cases that are not virtually cyclic can be viewed as a special case of the arguments used here for general periodic automorphisms. (Observe that $\mathrm{PGL}_2(\mathbb{Z}) \cong \mathrm{Out}(C_2 * C_2 * C_2)$, though the Out_0 considered above would be a finite index subgroup isomorphic to $C_2 * C_2 * C_2$.) There the problem can be reduced to understanding McCool groups of free groups, allowing more complicated incident edge groups to appear while leaving the problem tractable.

We are now in a position to prove one of our main theorems.

Theorem 1.1.1. *Suppose $G \cong F_n \rtimes_\varphi \mathbb{Z}$, and φ is linearly growing. Then* $\mathrm{Out}(G)$ *is finitely generated.*

Proof. The defining automorphism φ has a power that is UPG and linearly growing, so it is a Dehn twist. Taking this power to define a normal subgroup G_0 of G, by Proposition 5.2.2, G_0 has a canonical action on a tree T. Then by Proposition 3.1.4, G has a nearly canonical action on the same tree. The vertex stabilisers are free-by-cyclic groups which are virtually free-times-cyclic; edge stabilisers are free-by-cyclic groups that are virtually $\mathbb{Z}^2$ (see Proposition 5.2.4).

Analyse this action using Theorem 3.2.1. The quotient at (1) is finite since by Lemma 2.1.2 the quotient graph is. The quotient at (2) is also finitely generated. By Proposition 5.3.3, this will be finitely generated if and only if the relevant McCool groups are. After passing to a finite index subgroup, Proposition 5.3.6 describes these McCool groups (up to finite index) by considering McCool groups of virtually free groups, arising by quotienting by the centre. Since the edge groups contain the centre of the vertex groups, their image under this quotient map is virtually infinite cyclic. Finally, Proposition 5.4.1 gives finite generation for these McCool groups, completing this part of the proof.

The edge groups are virtually $\mathbb{Z}^2$, and in particular virtually abelian, so by Proposition 3.2.3 the quotient at (3) is finitely generated. Edge and vertex groups are both (finitely generated free)-by-cyclic, so by Lemma 3.2.4 the centralisers are finitely generated groups, and so is their quotient at (4). Finally, the quotient at (5) is a quotient of a finitely generated abelian group, so is itself finitely generated.

Putting this together, we see that $\mathrm{Out}(G)$ admits a finite index subgroup which is finitely generated, and so $\mathrm{Out}(G)$ itself is finitely generated, as claimed. □

6 Quadratic growth

6.1 Strategy

The strategy of the proof of this section is much like the last:

- Start with a free-by-cyclic group, $G = F_3 \rtimes_\Phi \mathbb{Z}$, where Φ has quadratic growth,
- Consider a finite index subgroup, $G_0 = F_3 \rtimes_{\Phi^r} \mathbb{Z}$, so that Φ^r is UPG,
- Find a good basis of F_3 for Φ^r and use this to construct a tree whose deformation space is left invariant by any automorphism of G_0,
- Deduce that the (reduced) tree of cylinders, $T = T_c^*$, of this space is G_0-canonical,
- Use Proposition 3.1.4 to deduce that T is nearly G-canonical
- Show that $\mathrm{Aut}^T(G)$ is finitely generated, using Theorem 3.2.1, and conclude that $\mathrm{Aut}(G)$ is finitely generated.

We establish some notation. Given a group, G, a subgroup H of G and elements g, h of G we set:

(i) We write $g \sim h$ to denote that g and h are conjugate in G, and

(ii) We write $g \sim_H h$ to denote that g and h are conjugate by an element of H (even if g, h might not themselves be elements of H)

6.2 Normal forms and a tree to act on

First we equip ourselves with a useful representative of a UPG automorphism.

Proposition 6.2.1. *Suppose Φ is a UPG element of* $\mathrm{Out}(F_3)$ *of quadratic growth. Then there is a representative $\varphi \in \Phi$ and a basis $\{a, b, c\}$ of F_3 so that*

$$\begin{array}{ccc} & \varphi & \\ a & \longrightarrow & a \\ b & \longrightarrow & ba^{-k} \\ c & \longrightarrow & hcg^{-1}, \end{array}$$

where k is non-zero and h and g are in $\langle a, b\rangle$.

This is close to [8, Proposition 5.9]; we have more control over the images of the first two generators in exchange for less control over the final generator.

Proof. By [3], any UPG automorphism is represented by a homotopy equivalence on a graph, G, such that G consists of edges, $E_1, \ldots, E_k$ and the homotopy equivalence maps E_i to $E_i u_{i-1}$, where the u_{i-1} are closed paths involving only the edges $E_1, \ldots, E_{i-1}$ (u_{i-1} may be the trivial path).

In particular, this implies that any UPG automorphism of F_2 has a representative, such that with respect to some basis, $\{a, b\}$, the automorphism fixes a and sends b to ba^{-k} for some k.

(Briefly, if the top edge, E_k, were separating, then the components on removing this edge would both be homotopic to circles, and then it is easy to see that the map is homotopic to the identity relative to the initial vertex of E_k. If E_k is not separating, then removing E_k leaves a graph, homotopic to a circle, on which the map is homotopic to the identity – giving us the a – and the E_k edge becomes the b basis element. Note that the layered description, which is a consequence of the UPG property, does not allow "inversions" of these various invariant circles.)

Now, if we are given a UPG automorphism, Φ, of F_3, the above description implies that some rank 2 free factor is left invariant, up to conjugacy – again, remove the top edge E_k. Each component of the complement is invariant under the map, and there must be one of rank 2. Moreover, the restriction of Φ to this invariant free factor is also UPG – in fact the restriction of the map has a layered form as above.

This implies that there is a basis, $\{a, b, c\}$ of F_3 and a representative $\varphi \in \Phi$ such that

$$\begin{array}{ccc} & \varphi & \\ a & \longrightarrow & a \\ b & \longrightarrow & ba^{-k} \end{array}$$

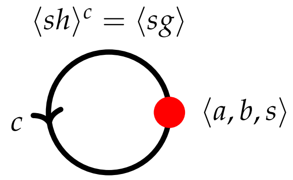

FIGURE 3.1: T_0, as described in Corollary 6.2.2

But now, since the images of a, b, c must also be a basis for F_3, the only possibility for the image of c is $hc^{\pm 1}g^{-1}$ for some $g, h \in \langle a, b\rangle$. The fact that Φ is UPG (or using the description of the map) means that the image must be hcg^{-1}.

Finally, note that if k were to be zero, then Φ would have linear growth, hence we may conclude that $k \neq 0$. □

Corollary 6.2.2. *Let $G = F_3 \rtimes_\Phi \mathbb{Z}$ be a free-by-cyclic group where Φ has quadratic growth. Then G has a normal finite index subgroup $G_0 = F_3 \rtimes_{\Phi^r} \mathbb{Z}$ with presentation:*

$$\begin{aligned} G_0 &= \langle a, b, c, s \;:\; a^s = a, b^s = ba^{-k}, c^s = hcg^{-1}\rangle \\ &= \langle H, c \;:\; (sh)^c = sg\rangle, \text{ where } H = \langle a, b, s\rangle \end{aligned}$$

where $g, h \in \langle a, b\rangle$ and $k \neq 0$.

Moreover, G_0 acts on a tree, T_0, with one orbit of vertices, and one orbit of edges such that the vertex stabilisers are conjugates of $H = \langle a, b, s\rangle$ and edge groups are conjugates of $\langle sg\rangle = \langle sh\rangle^c$. (See Figure 3.1.)

Proof. Every polynomially growing automorphism has a power which is UPG, and Proposition 6.2.1 provides a good generating set and the corresponding presentation. The final relations in both presentations are equivalent, realising G_0 as an HNN extension of $\langle a, b, s\rangle$ with stable letter c, and T_0 is the corresponding Bass-Serre tree. □

6.3 Invariance of the tree

Proposition 6.3.1. *Any automorphism of G_0 fixes the conjugacy class of $\langle a, b, s\rangle$. That is, the deformation space defined by the tree, T_0 from Corollary 6.2.2 is invariant under the automorphisms of G_0.*

Proof. First we will see that s is sent to a conjugate of $s^{\pm 1}$, and then we prove that in the case it is fixed or inverted, the subgroup $\langle a, b, s\rangle$ is preserved.

The subgroup $\langle a, bab^{-1}, s\rangle$ is the centraliser of s and so its image must be the centraliser of some element. Since the free part has rank 2, this element has to be $s^\ell w$ for some ℓ

and w. In fact by Theorem 2.4.2 the only elements which commute with a subgroup of rank at least two are conjugate to s^ℓ. Since s has no roots we must have $\ell = \pm 1$. In particular, this implies $\langle a, bab^{-1}, s\rangle$ is sent to a conjugate.

Therefore, up to composing with an inner automorphism, we may assume that s is fixed or inverted by our automorphism, and we consider the images of a and b. We write u for the image of a and $s^n v$ for the image of b, where u, v are elements of $\langle a, b, c\rangle$. (Notice that the relation $b^s = ba^{-k}$ implies that the image of a lies in this free group.)

The image of bab^{-1} is vuv^{-1}, so $\langle u, vuv^{-1}\rangle = \langle a, bab^{-1}\rangle$. In particular, this shows that u and v (by malnormality of free factors) are contained in $\langle a, b\rangle$.

To see the other inclusion, notice that we also know that $\langle u, v\rangle$ contains $\langle a, bab^{-1}\rangle$. By considering the Stallings graphs (see [33]) of both subgroups, this means it contains either $\langle a, bab^{-1}\rangle$ or $\langle a, b\rangle$ as a free factor.

That is, the subgroup inclusion gives a graph morphism from the Stallings graph of $\langle a, bab^{-1}\rangle$ to that of $\langle u, v\rangle$ with respect to the basis $\{a, b\}$. If this map is injective, then the Stallings graph of $\langle a, bab^{-1}\rangle$ is a subgraph and therefore $\langle a, bab^{-1}\rangle$ is a free factor of $\langle u, v\rangle$. If not, then the two vertices of the Stallings graph of $\langle a, bab^{-1}\rangle$ are identified, and we must get that $\langle a, b\rangle$ is a free factor of, and hence must be equal to, $\langle u, v\rangle$. This is an easy version of the arguments in [34], Theorem 1.7.

Since it has rank 2, this actually says $\langle u, v\rangle$ is equal to either $\langle a, bab^{-1}\rangle$ or $\langle a, b\rangle$; the first is impossible since it would imply that $\langle u, v\rangle = \langle a, bab^{-1}\rangle = \langle u, vuv^{-1}\rangle$, which cannot happen since the last subgroup does not contain v.

Hence, $\langle u, v\rangle = \langle a, b\rangle$ and $\langle u, v, s\rangle = \langle a, b, s\rangle$. □

Corollary 6.3.2. *The (collapsed) tree of cylinders, T_c^*, of T_0 is G_0-canonical and hence nearly G-canonical.*

Proof. The fact that the deformation space of T_0 is invariant, gives us that the (collapsed) tree of cylinders, T_c^* is canonical, see Subsection 2.3.

Then Proposition 3.1.4 gives us the second statement. □

6.4 Calculating the tree of cylinders, T_c

Our goal now is to calculate T_c. In order to do this, we actually modify the basis given by Proposition 6.2.1. The tree, T_0 from Proposition 6.2.2 remains the same, but these modifications aid the calculation.

Throughout this subsection, we are working with the subgroup

$$G_0 = \langle a, b, c, s \ : \ a^s = a, b^s = ba^{-k}, c^s = hcg^{-1} \rangle.$$

First we observe that we can modify the elements g, h from Proposition 6.2.1, and thus in the description of G_0.

Lemma 6.4.1. *The choices of g and h in the statement of Proposition 6.2.1 are not unique. In particular, if $sh \sim_{\langle a,b \rangle} sh'$ and $sg \sim_{\langle a,b \rangle} sg'$, then there exist $x, y \in \langle a, b \rangle$ such that if $c' = x^{-1}cy$, then the image of c' under φ is $h'c'g'^{-1}$.*

Proof. We will work in the corresponding free-by-cyclic group, G_0 from Proposition 6.2.2 and its presentation.

Recall that $G_0 = \langle a, b, c, s \ : \ a^s = a, b^s = ba^{-k}, c^s = hcg^{-1} \rangle$. It will be sufficient to show that $s^{-1}c's = h'c'g'^{-1}$.

Suppose $(sh)^x = sh'$, and $(sg)^y = sg'$, where x and y are elements of $\langle a, b \rangle$. Then put $c' = x^{-1}cy$.

We get that,

$$\begin{aligned} s^{-1}c's &= (x^{-1}cy)^s \\ &= x^{-s}hcg^{-1}y^s \\ &= (x^{-s}hx)c'(y^{-1}g^{-1}y^s) \\ &= s^{-1}(sh)^x c'(sg)^{-y}s \\ &= h'c'g'^{-1}. \end{aligned}$$ □

Note that each of sh and sg normalise $\langle a, b \rangle$. Moreover, they induce the same outer automorphism, and this is a Dehn Twist of $\langle a, b \rangle$. However, while sh and sg are conjugate in G_0 – and so induce isogredient automorphisms of $\langle a, b, c \rangle$ – they might not induce isogredient automorphisms of $\langle a, b \rangle$.

One key point is that:

Lemma 6.4.2. *The following are equivalent:*

(i) *sh and sg induce isogredient automorphisms on $\langle a, b \rangle$*

(ii) $sh \sim_{\langle a,b \rangle} sg$

(iii) $sh \sim_{\langle a,b,s \rangle} sg$.

Proof. The first two are clearly equivalent, and notice that $\langle a, b, sh \rangle = \langle a, b, s \rangle = \langle a, b, sg \rangle$, which makes the second and third equivalent since we can choose the new generator so it centralises the conjugated element. □

We will use the following result, to help us modify g and h as above.

Corollary 6.4.3 ([27, Corollary 3.10])**.** *Let $\Psi \in \mathrm{Out}(F_n)$, $n \geq 2$, be a Dehn Twist outer automorphism fixing a conjugacy class. Then there is a $\psi \in \Psi$ with fixed subgroup of rank at least two fixing an element of that conjugacy class.*

Lemma 6.4.4. *In G_0, the centraliser $C_{\langle a,b \rangle}(sh)$ has rank 0,1 or 2. If the rank is at least 1, then $sh \sim_{\langle a,b \rangle} sh'$ for some $h' \in C_{\langle a,b \rangle}(s) = \langle a, bab^{-1} \rangle$. The same is true for g.*

Moreover, one of $C_{\langle a,b \rangle}(sh)$ and $C_{\langle a,b \rangle}(sg)$ has rank 0 (is the trivial group).

Proof. The first statement follows from the Bestvina-Handel Theorem, Theorem 2.4.2.

For the second statement, we invoke Corollary 6.4.3, to say that if $C_{\langle a,b \rangle}(sh)$ is non-trivial, then there exists a non-trivial $w \in \langle a, b \rangle$ and an $x \in \langle a, b \rangle$ such that:

$$\begin{aligned} (w^x)^{sh} &= w^x \\ w^s &= w. \end{aligned}$$

Here we are using Theorem 2.4.2 to say that since the underlying free group has rank 2, there is exactly one isogredience class with fixed subgroup of rank at least 2, and hence the ψ from Corollary 6.4.3 is, without loss of generality, the automorphism induced by conjugation by s (on $\langle a, b \rangle$). (It is more convenient for the following argument to write w^x for the element fixed by the automorphism induced by sh.)

But these equations imply that,

$$w^{s^{-1}x(sh)x^{-1}} = w^{x(sh)x^{-1}} = (w^x)^{(sh)x^{-1}} = (w^x)^{x^{-1}} = w.$$

Hence, as w is non-trivial and both w and $s^{-1}x(sh)x^{-1}$ are elements of $\langle a, b \rangle$, we get that $s^{-1}x(sh)x^{-1} \in \langle w \rangle \leq C_{\langle a,b \rangle}(s)$, and hence $sh \sim_{\langle a,b \rangle} sw^m$, for some $m \in \mathbb{Z}$ (without loss of generality, we can assume w is not a proper power, and so generates its own centraliser in $\langle a, b \rangle$). The same calculation gives the result for g.

Finally, notice that if both $h, g \in C_{\langle a,b}(s)$, then Φ has linear growth. Thus, via Lemma 6.4.1, we deduce that one of $C_{\langle a,b \rangle}(sh)$ and $C_{\langle a,b \rangle}(sg)$ has rank 0. □

Remark 6.4.5. Given this result, we shall henceforth assume that $C_{\langle a,b \rangle}(sg)$ is the trivial group. (Note that h and g can be interchanged by replacing c with c^{-1} so there is no loss of generality in assuming this.)

We also record that,

Lemma 6.4.6. *Let G be a free-by-cyclic group, with stable letter s. Any subgroup $\langle s^m w\rangle$, with $m \neq 0$, has centraliser and normaliser equal.*

Proof. Notice that conjugation by any element of the normaliser induces an automorphism of $\langle s^m w\rangle$, and in particular either preserves the generator (in which case it is an element of the centraliser) or inverts it. However, conjugating cannot affect the exponent sum of the stable letter s, and so this last case does not arise. □

Since there is only one orbit of edges, we can understand the cylinders by understanding the normaliser of any edge stabiliser. Since the edges are stabilised by infinite cyclic groups of the kind discussed in Lemma 6.4.6, this is equivalent to understanding their centralisers.

Theorem 6.4.7. *Let $G_0 = \langle a, b, c, s \ : \ a^s = a, b^s = ba^{-k}, c^s = hcg^{-1}\rangle$, and T_0 be the Bass-Serre tree on which G_0 acts via the HNN decomposition, $G_0 = \langle H, c \ : \ (sh)^c = sg\rangle$, where $H = \langle a, b, s\rangle$.*

Moreover, assume that $C_{\langle a,b\rangle}(sg)$ is the trivial group, as in Remark 6.4.5.

We form the tree of cylinders, T_c, and collapsed tree of cylinders T_c^* taking maximal infinite cyclic groups to be the family $\mathcal{E}$ and equality to be the admissible equivalence relation.

- *If $sh \not\sim_{\langle a,b\rangle} sg$, then $T_c^* = T_0$, or $T_c = T_c^*$ is simply a subdivision of T_0.*
- *If $sh \sim_{\langle a,b\rangle} sg$, then $T_c = T_c^*$ has one edge orbit, with infinite cyclic stabilisers, conjugates of $\langle sh\rangle$, and two vertex orbits, with stabilisers conjugates of $\langle a, b, s\rangle$ and $C(sh) \cong \mathbb{Z}^2$.*

Proof. Since T_0 has one orbit of edges and one orbit of vertices, the tree of cylinders of T_0 will have two orbits of vertices – one for the cylinders, and one for the T_0-vertices.

Since our relation is equality, edge stabilisers in T_0 are conjugate to sh, and we have Lemma 6.4.6, we deduce that a cylinder is the orbit of an edge under the action of the centraliser of the edge stabiliser (in G_0).

As G_0 acts without inversions on T_0, we may equivariantly orient the edges of T_0. A vertex stabiliser in T_0 acts on the incident edges with two orbits – one orbit for the incoming edges, and one for the outgoing edges.

Choose this orientation so that at the vertex stabilised by $\langle a, b, s\rangle$, the incoming edges have stabiliser conjugate (in $\langle a, b, s\rangle$) to $\langle sg\rangle$ and for the outgoing edges it is conjugate to $\langle sh\rangle$.

The fact that $C_{\langle a,b\rangle}(sg)$ is the trivial group implies that $C_{\langle a,b,s\rangle}(sg) = \langle sg\rangle$ and hence that no cylinder may contain two incoming edges at a vertex.

Suppose $sh \not\sim_{\langle a,b\rangle} sg$:

If a cylinder contained both incoming and outgoing edges at a vertex, then (moving back to the vertex stabilised by $\langle a,b,s\rangle$) we would have $sh \sim_{\langle a,b,s\rangle} sg$, since acting on the edges conjugates the stabilisers. So if $sh \not\sim_{\langle a,b\rangle} sg$ (which is equivalent to $sh \not\sim_{\langle a,b,s\rangle} sg$), then no cylinder may contain both incoming and outgoing edges at a vertex.

Thus if $sh \not\sim_{\langle a,b\rangle} sg$, all cylinders consist of a collection of outgoing edges from a vertex. More concretely, if we take the edge with stabiliser $\langle sh\rangle$, then the corresponding cylinder consists of edges starting from the vertex with stabiliser $\langle a,b,s\rangle$, and are thus all in the same $\langle a,b,s\rangle$-orbit. In particular, this implies that $C(sh) = C_{\langle a,b,s\rangle}(sh) = C_{\langle a,b\rangle}(sh) \times \langle sh\rangle$.

The cylinder stabiliser acts with two orbits on its vertices – the central vertex and all the rest, and hence the tree of cylinders of T_0 has two edge orbits corresponding to these different inclusions. One of these edges has stabilisers equal to the edge stabilisers of the original tree (this is where we have the vertex being one of the 'outside' vertices of the cylinder), whereas the other edge group is equal to the stabiliser of the cylinder, (conjugates of) $C(sh)$.

If $C(sh)$ is not cyclic, then the collapsed tree of cylinders will collapse the corresponding edge, and we will return to the original tree.

If $C(sh)$ is cyclic, then the tree of cylinders is just a subdivision of T_0 – we have subdivided an edge, and given the new vertex the same stabiliser as the edge it is part of.

Suppose $sh \sim_{\langle a,b\rangle} sg$:

If $sh \sim_{\langle a,b\rangle} sg$, then we orient the edges of T_0 as before and now we get that both $C_{\langle a,b\rangle}(sh)$ and $C_{\langle a,b\rangle}(sg)$ are trivial (since they are conjugate). Therefore, $C_{\langle a,b,s\rangle}(sh) = \langle sh\rangle$ and $C_{\langle a,b,s\rangle}(sg) = \langle sg\rangle$.

This means that a cylinder cannot contain either two outgoing or two incoming edges at any vertex. However, each cylinder does contain both an outgoing and incoming edge at each vertex. Hence the cylinder is a line and it is straightforward to verify that $C(sh) \cong \mathbb{Z}^2$. (Since $sh \sim_{\langle a,b\rangle} sg$, we may assume that $h = g$, and in this case, $C(sh) = \langle c, sh\rangle$ – c is acting as a translation, and therefore transitively on the vertices and edges of this line).

In this case, there are again two orbits of vertices in the tree of cylinders – one for the cylinders, one for the vertices of T_0 – with stabilisers (conjugates of) $\langle a,b,s\rangle$ and $C(sh) \cong \mathbb{Z}^2$.

Since the cylinder stabiliser acts transitively on its vertices, there is only one edge, whose stabiliser is (the conjugates of) $\langle sh\rangle$. □

Remark 6.4.8. The tree of cylinders produced by this theorem realises a maximal preserved free factor system for the automorphism induced by s: it is an interesting question if this is true more generally (say, in higher rank or higher polynomial growth).

We now use Theorem 6.4.7 to provide a nearly canonical tree for the general (not just UPG) case.

Corollary 6.4.9. *Let $G \cong F_3 \rtimes_\Phi \mathbb{Z}$, and Φ is quadratically growing. Then G admits an action on a nearly canonical tree, T, such that:*

(i) The action is co-compact (equivalently, co-finite),

(ii) Edge stabilisers are infinite cyclic,

(iii) Vertex stabilisers are of the form $F_r \rtimes \mathbb{Z}$, where $r = 0, 1, 2$.

Proof. We simply apply Proposition 3.1.4 to the collapsed tree of cylinders for G_0 above, Theorem 6.4.7, to get a nearly canonical action on the same tree. The fact that the G action extends the G_0 action tells us about the stabilisers. (For example, edge stabilisers in G must be infinite cyclic since their intersection with F_3 is trivial). □

We now use this to prove the following theorem, which is part of Theorem 1.1.2.

Theorem 6.4.10. *Suppose $G \cong F_3 \rtimes_\varphi \mathbb{Z}$, and φ is quadratically growing. Then* $\mathrm{Out}(G)$ *is finitely generated.*

Proof. We use the tree constructed above, and we calculate the quotients of $\mathrm{Out}^T(G)$ described in Theorem 3.2.1. The quotient graphs are finite, and therefore so is the quotient at (1). For the quotient at (2), the edge groups are all infinite cyclic, and therefore have finite outer automorphism group. So by Lemma 3.2.2, we only need to check the McCool groups of vertex groups. Since vertex groups are free by cyclic groups of rank 0, 1 or 2, these are finitely generated by Proposition 4.1.6.

Since the edge groups are infinite cyclic, we may apply Proposition 3.2.3 to see that the quotient at (3) is finite. The quotient at (4) is finitely generated by Lemma 3.2.4 and that at (5) as a quotient of a finitely generated abelian group. □

Our other main theorem is proved by combining Theorem 1.1.1 (restricted to rank 3) for the linear growth case, Theorem 6.4.10 for the quadratic growth case, Theorem 4.1.7 for the exponential case and [22] for the periodic case.

Theorem 1.1.2. *Suppose $G \cong F_3 \rtimes \mathbb{Z}$. Then* $\mathrm{Out}(G)$ *is finitely generated.*

References

[1] Naomi Andrew. Serre's property (FA) for automorphism groups of free products. *J. Group Theory*, 24(2):385–414, 2021.

[2] Hyman Bass and Renfang Jiang. Automorphism groups of tree actions and of graphs of groups. *J. Pure Appl. Algebra*, 112(2):109–155, 1996.

[3] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $\mathrm{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms. *Annals of Mathematics. Second Series*, 151(2):517–623, 2000.

[4] Mladen Bestvina, Mark Feighn, and Michael Handel. A McCool Whitehead type theorem for finitely generated subgroups of $\mathrm{Out}(F_n)$. Preprint, September 2020. arXiv:2009.10052.

[5] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[6] O. Bogopolski, A. Martino, and E. Ventura. The automorphism group of a free-by-cyclic group in rank 2. *Comm. Algebra*, 35(5):1675–1690, 2007.

[7] P. Brinkmann. Hyperbolic automorphisms of free groups. *Geom. Funct. Anal.*, 10(5):1071–1089, 2000.

[8] Christopher H. Cashen and Gilbert Levitt. Mapping tori of free group automorphisms, and the Bieri-Neumann-Strebel invariant of graphs of groups. *J. Group Theory*, 19(2):191–216, 2016.

[9] Marshall M. Cohen and Martin Lustig. Very small group actions on **R**-trees and Dehn twist automorphisms. *Topology. An International Journal of Mathematics*, 34(3):575–617, 1995.

[10] Marshall M. Cohen and Martin Lustig. The conjugacy problem for Dehn twist automorphisms of free groups. *Commentarii Mathematici Helvetici*, 74(2):179–200, 1999.

[11] Marc Culler and John W. Morgan. Group actions on $\mathbb{R}$-trees. *Proc. London Math. Soc. (3)*, 55(3):571–604, 1987.

[12] François Dahmani. On suspensions and conjugacy of hyperbolic automorphisms. *Trans. Amer. Math. Soc.*, 368(8):5565–5577, 2016.

[13] François Dahmani and Ruoyu Li. Relative hyperbolicity for automorphisms of free products. Preprint, January 2019. arXiv:1901.06760.

[14] Francois Gautero and Martin Lustig. The mapping-torus of a free group automorphism is hyperbolic relative to the canonical subgroups of polynomial growth. Preprint, July 2007. arXiv:0707.0822.

[15] Pritam Ghosh. Relative hyperbolicity of free-by-cyclic extensions. Preprint, February 2018. arXiv:1802.08570.

[16] Vincent Guirardel and Gilbert Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

[17] Vincent Guirardel and Gilbert Levitt. McCool groups of toral relatively hyperbolic groups. *Algebr. Geom. Topol.*, 15(6):3485–3534, 2015.

[18] Vincent Guirardel and Gilbert Levitt. Splittings and automorphisms of relatively hyperbolic groups. *Groups Geom. Dyn.*, 9(2):599–663, 2015.

[19] G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[20] Sava Krstić, Martin Lustig, and Karen Vogtmann. An equivariant Whitehead algorithm and conjugacy for roots of Dehn twist automorphisms. *Proceedings of the Edinburgh Mathematical Society. Series II*, 44(1):117–141, 2001.

[21] Gilbert Levitt. Automorphisms of hyperbolic groups and graphs of groups. *Geom. Dedicata*, 114:49–70, 2005.

[22] Gilbert Levitt. On the automorphism group of generalized Baumslag-Solitar groups. *Geom. Topol.*, 11:473–515, 2007.

[23] Gilbert Levitt. Counting growth types of automorphisms of free groups. *Geom. Funct. Anal.*, 19(4):1119–1146, 2009.

[24] Gilbert Levitt. Generalized Baumslag-Solitar groups: rank and finite index subgroups. *Ann. Inst. Fourier (Grenoble)*, 65(2):725–762, 2015.

[25] Martin Lustig and Kaidi Ye. Normal form and parabolic dynamics for quadratically growing automorphisms of free groups. May 2017. arXiv:1705.04110.

[26] Nataša Macura. Detour functions and quasi-isometries. *Q. J. Math.*, 53(2):207–239, 2002.

[27] Armando Martino. Maximal index automorphisms of free groups with no attracting fixed points on the boundary are Dehn twists. *Algebr. Geom. Topol.*, 2:897–919, 2002.

[28] J. McCool. The automorphism groups of finite extensions of free groups. *Bull. London Math. Soc.*, 20(2):131–135, 1988.

[29] James McCool. Some finitely presented subgroups of the automorphism group of a free group. *Journal of Algebra*, 35:205–213, 1975.

[30] A. Minasyan and D. Osin. Normal automorphisms of relatively hyperbolic groups. *Trans. Amer. Math. Soc.*, 362(11):6079–6103, 2010.

[31] Ashot Minasyan and Denis Osin. Fixed subgroups of automorphisms of relatively hyperbolic groups. *Q. J. Math.*, 63(3):695–712, 2012.

[32] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[33] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[34] Enric Ventura. On fixed subgroups of maximal rank. *Communications in Algebra*, 25(10):3361–3375, 1997.