Symmetries of free and right-angled Artin groups

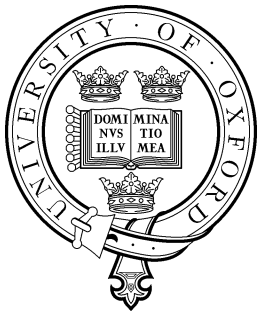

Richard D Wade
Corpus Christi College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity Term, 2012

This thesis would not have been completed without the unending support of my family and friends. I would especially like to thank my supervisor for sharing his knowledge, ideas, and limitless enthusiasm during my time as a graduate student.

Parts of this thesis have appeared previously in papers [17], [71], [72], and [73]. In particular, the material in the first half of Chapter 5 comes from the paper [17] coauthored with Martin Bridson.

Abstract

The objects of study in this thesis are automorphism groups of free and right-angled Artin groups. Right-angled Artin groups are defined by a presentation where the only relations are commutators of the generating elements. When there are no relations the right-angled-Artin group is a free group and if we take all possible relations we have a free abelian group.

We show that if no finite index subgroup of a group G contains a normal subgroup that maps onto $\mathbb{Z}$, then every homomorphism from G to the outer automorphism group of a free group has finite image. The above criterion is satisfied by $\mathrm{SL}_m(\mathbb{Z})$ for $m \geq 3$ and, more generally, all irreducible lattices in higher-rank, semisimple Lie groups with finite centre.

Given a right-angled Artin group A_Γ we find an integer n, which may be easily read off from the presentation of A_Γ, such that if $m \geq 3$ then $\mathrm{SL}_m(\mathbb{Z})$ is a subgroup of the outer automorphism group of A_Γ if and only if $m \leq n$. More generally, we find criteria to prevent a group from having a homomorphism to the outer automorphism group of A_Γ with infinite image, and apply this to a large number of irreducible lattices as above.

We study the subgroup $IA(A_\Gamma)$ of $\mathrm{Aut}(A_\Gamma)$ that acts trivially on the abelianisation of A_Γ. We show that $IA(A_\Gamma)$ is residually torsion-free nilpotent and describe its abelianisation. This is complemented by a survey of previous results concerning the lower central series of A_Γ.

One of the commonly used generating sets of $\mathrm{Aut}(F_n)$ is the set of Whitehead automorphisms. We describe a geometric method for decomposing an element of $\mathrm{Aut}(F_n)$ as a product of Whitehead automorphisms via Stallings' folds. We finish with a brief discussion of the action of $\mathrm{Out}(F_n)$ on Culler and Vogtmann's Outer Space. In particular we describe translation lengths of elements with regards to the 'non-symmetric Lipschitz metric' on Outer Space.

Contents

1 Introduction **1**

I Rigidity, IA automorphisms, and central series **8**

2 The lower central series of a right-angled Artin group **9**
2.1 Associative algebras, Lie algebras, and central filtrations 12
2.2 Lie algebras from central filtrations 13
2.3 The cast 17
2.3.1 The monoid M_Γ and algebra U_Γ 17
2.3.2 U^∞, an ideal X, and the group of units U^* 19
2.4 The Magnus map 19
2.5 Lyndon elements of M 25
2.5.1 Lyndon words 27
2.5.2 Lyndon elements 28
2.5.3 The standard factorisation of a Lyndon element 30
2.5.4 A basis theorem for the algebra L_Γ 30
2.6 An isomorphism between L_Γ and the LCS algebra of A_Γ 38
2.7 More information on the structure of the LCS algebra 40

3 Foundations for studying $Out(A_\Gamma)$ **41**
3.1 A generating set of $\mathrm{Aut}(A_\Gamma)$ 41
3.2 Ordering the vertices of Γ 43
3.2.1 The standard order on $V(\Gamma)$ 43
3.2.2 G–ordering vertices 44
3.2.3 Relation of G–orderings to the action on $H_1(A_\Gamma)$ 45
3.3 Restriction, exclusion, and projection homomorphisms. 46
3.4 SL–dimension for subgroups of $\mathrm{Out}(A_\Gamma)$. 48
3.5 An example 50

4 IA_n and $\mathrm{IA}(A_\Gamma)$ 52
4.1 Finitely generated subgroups of IA_n and $\mathrm{IA}(A_\Gamma)$ . . . 54
4.1.1 $\mathrm{Fix}_c(\{x_{m+1}, \ldots, x_n\}) \cap \mathrm{IA}_n$ . . . 54
4.1.2 A generating set of $\mathrm{IA}(A_\Gamma)$ . . . 57
4.2 The Andreadakis–Johnson Filtration of $\mathrm{IA}(A_\Gamma)$ . . . 61
4.2.1 A central filtration of $\mathrm{IA}(A_\Gamma)$ . . . 62
4.2.2 The image of the filtration in $\overline{\mathrm{IA}}(A_\Gamma)$ . . . 64
4.2.3 The structure of $H_1(\mathrm{IA}(A_\Gamma))$ . . . 66

5 Homomorphisms to $\mathrm{Out}(F_n)$ and $\mathrm{Out}(A_\Gamma)$ 68
5.1 Homomorphisms to $\mathrm{Out}(F_n)$ . . . 71
5.1.1 Controlling the action of Λ on homology . . . 71
5.1.2 Alternative Hypotheses . . . 72
5.2 Homomorphisms to $\mathrm{Out}(A_\Gamma)$ . . . 74
5.2.1 Γ is disconnected. . . . 75
5.2.2 Γ is connected and $Z(A_\Gamma)$ is trivial. . . . 76
5.2.3 Γ is connected and $Z(A_\Gamma)$ is nontrivial. . . . 76
5.3 Consequences of Theorem 5.10 . . . 77

II Folding and Outer Space 80

6 Folding free-group automorphisms 81
6.1 Graphs, folds, and associated automorphisms . . . 82
6.1.1 The fundamental group of a graph . . . 83
6.1.2 Folding maps of graphs . . . 85
6.1.3 Branded graphs and their associated automorphisms . . . 86
6.2 An algorithm . . . 87
6.2.1 Folding edges contained in T . . . 88
6.2.2 Swapping edges into a tree . . . 90
6.3 Fixing generators . . . 91

7 Displacement functions on Outer Space 94
7.1 Outer space and the Lipschitz metric . . . 96
7.2 Train tracks and displacement functions . . . 99

A $\langle \mathcal{M}_\Gamma \rangle$ is a normal subgroup of $\mathrm{SAut}^0(A_\Gamma)$. 103

A.1 Conjugates of the form $\rho_{ij} K_{kl} \rho_{ij}^{-1}$ and $\rho_{ij}^{-1} K_{kl} \rho_{ij}$. 104

A.2 Conjugates of the form $\rho_{ij} K_{klm} \rho_{ij}^{-1}$ and of the form $\rho_{ij}^{-1} K_{klm} \rho_{ij}$ 105

Bibliography 108

Chapter 1

Introduction

A motivating theme in this thesis is the principle that we may reduce questions about $\mathrm{Aut}(F_n)$ to questions about $\mathrm{GL}_n(\mathbb{Z})$ and the group IA_n, the subgroup of $\mathrm{Aut}(F_n)$ that acts trivially on the abelianisation of F_n. This is because of the exact sequence:

$$1 \to \mathrm{IA}_n \to \mathrm{Aut}(F_n) \to \mathrm{GL}_n(\mathbb{Z}) \to 1.$$

The fact that $\mathrm{Aut}(F_n)$ has a torsion-free finite-index subgroup follows from fact that IA_n is torsion-free and most congruence subgroups of $\mathrm{GL}_n(\mathbb{Z})$ are torsion free. Dyer and Formanek's proof [33] that $\mathrm{Aut}(\mathrm{Aut}(F_n)) = \mathrm{Aut}(F_n)$ also heavily involves the structure of IA_n, in particular Magnus' generating set [56] of this group. However, a look through the literature will reveal that this mantra is utilised less often than one would expect. One reason could be that besides Magnus' 1935 result, not much progress on IA_n has been made in the intervening 70 years. Indeed, although $\mathrm{IA}_2 \cong F_2$, and IA_3 is not finitely presented [49], it is not known if there exists a finite presentation for IA_n when $n \geq 4$. We do not claim to transform this situation in this thesis, but we will use the above principal productively. In fact, we use it to reinforce a second, more fruitful, philosophy, that "$\mathrm{Out}(F_n)$ behaves like a lattice in a higher-rank semisimple Lie group."

Roughly speaking, Margulis' superrigidity tells us that the only maps between such lattices are the obvious ones. In particular:

Theorem 1.1 (Margulis)**.** *Let Λ be an irreducible lattice in a real semisimple Lie group G that has finite centre, no compact factors, and real rank* $\mathrm{rank}_{\mathbb{R}} G \geq k$*. Then every homomorphism $f : \Lambda \to \mathrm{SL}_k(\mathbb{Z})$ has finite image.*

As there aren't any obvious maps from a higher-rank lattice to $\mathrm{Out}(F_n)$, at least with infinite image, the above aphorism leads one to expect that there shouldn't be any maps at all with infinite image. We show that this indeed the case:

Theorem 5.2 (Bridson–W, [17]). *If G is a connected, semisimple Lie group of real rank at least 2 that has finite centre, and Λ is an irreducible lattice in G, then every homomorphism from Λ to the outer automorphism group of a finitely generated free group has finite image.*

In Theorem 5.2, our input from the technology that Margulis developed is his *normal subgroup theorem*, which implies that no finite index subgroup of such a lattice contains a normal subgroup that maps onto $\mathbb{Z}$. Theorem 5.2 is then deduced from the following:

Theorem 5.1. *Let Λ be a group. Suppose that no subgroup of finite index in Λ has a normal subgroup that maps surjectively to $\mathbb{Z}$. Then every homomorphism from Λ to the outer automorphism group of a finitely generated free group has finite image.*

We say that a group Λ satisfying the hypothesis of Theorem 5.1 is $\mathbb{Z}$-*averse*. The proof of Theorem 5.1 roughly goes as follows: we use some deep geometric results by Bestvina and Feighn [11], Dahmani, Guirardel and Osin [25], and Handel and Mosher [43] to analyse the action of the image of a $\mathbb{Z}$-averse group on the abelianisation of F_n (that is, on $H_1(F_n)$). We show that the group must act trivially on $H_1(F_n)$, and therefore its image lies in $\overline{\mathrm{IA}}_n$, the projection of IA_n to $\mathrm{Out}(F_n)$. One can show that if N is a torsion-free nilpotent group then every subgroup of N maps onto $\mathbb{Z}$ and therefore any homomorphism from a $\mathbb{Z}$-averse group to N is trivial. We complete our proof by appealing to the following theorem of Bass and Lubotzky:

Theorem 1.2 ([5], Theorem 10.4). *$\overline{\mathrm{IA}}_n$ is residually torsion-free nilpotent.*

After comparing Theorems 1.1 and 5.2, we were drawn naturally to right-angled Artin groups. Such groups have become of central importance in geometric group theory in recent years, most notably through their association with special cube complexes [42] and Bestvina–Brady groups [9]. A right-angled Artin group A_Γ, or RAAG for short, is determined by a graph Γ, and depending on your choice of Γ, the group A_Γ can exhibit both free and free-abelian behaviour. Indeed, if Γ is a discrete graph then A_Γ is a free group, and if Γ is a complete graph then A_Γ is a free-abelian group. One therefore expects traits shared by both $\mathbb{Z}^n$ and F_n to be shared by an arbitrary RAAG. Similarly, one optimistically hopes that properties shared by both $\mathrm{GL}_n(\mathbb{Z})$ and $\mathrm{Out}(F_n)$ will also by shared by $\mathrm{Out}(A_\Gamma)$ for an arbitrary right-angled Artin group. For instance, there is a Nielsen-type generating set of $\mathrm{Out}(A_\Gamma)$ given by the work of Laurence [52] and Servatius [68], and $\mathrm{Out}(A_\Gamma)$ has finite virtual

cohomological dimension [20]. $\mathrm{Out}(A_\Gamma)$ is residually finite [21, 61], and for a large class of graphs, $\mathrm{Out}(A_\Gamma)$ satisfies the Tits alternative [21].

Unlike $\mathrm{Out}(F_n)$, there are maps from higher-rank lattices to the outer automorphism groups of some right-angled Artin groups. Suppose that Γ' is a k–vertex subgraph of Γ which is a clique (i.e. any two vertices of Γ' are connected by an edge). Then $A_{\Gamma'} \cong \mathbb{Z}^k$ is a subgroup of A_Γ and it is tempting to look for a copy of $\mathrm{SL}_k(\mathbb{Z})$ in $\mathrm{Out}(A_\Gamma)$ supported on $A_{\Gamma'}$. In general, such a copy of $\mathrm{SL}_k(\mathbb{Z})$ does not exist, but if every vertex in v has the same star (the set of vertices at distance less than or equal to one from v) then the embedding $\mathbb{Z}^k \to A_\Gamma$ induces an injection $\mathrm{SL}_k(\mathbb{Z}) \to \mathrm{Out}(A_\Gamma)$. Define the SL–dimension of $\mathrm{Out}(A_\Gamma)$, written $d_{SL}(\mathrm{Out}(A_\Gamma))$, to be the number of vertices in the largest subgraph Γ' satisfying this condition. This can reasonably be thought of as a notion of rank for $\mathrm{Out}(A_\Gamma)$. We generalise Margulis' theorem and Theorem 5.2 like so:

Theorem 5.16. *Let G be a real semisimple Lie group with finite centre, no compact factors, and $\mathrm{rank}_{\mathbb{R}} G \geq 2$. Let Λ be an irreducible lattice in G. If $\mathrm{rank}_{\mathbb{R}} G \geq d_{SL}(\mathrm{Out}(A_\Gamma))$, then every homomorphism $f : \Lambda \to \mathrm{Out}(A_\Gamma)$ has finite image.*

Note that $d_{SL}(\mathrm{GL}_n(\mathbb{Z})) = n$ and $d_{SL}(\mathrm{Out}(F_n)) = 1$. Our previous observation that $\mathrm{Out}(A_\Gamma)$ contains a copy of $SL_m(\mathbb{Z})$ for $m = d_{SL}(\mathrm{Out}(A_\Gamma))$ tells us that the bound on $\mathrm{rank}_{\mathbb{R}} G$ given in Theorem 5.16 is the best that one can provide. As with Theorem 5.2, the above result is deduced from a more algebraic criterion:

Theorem 5.10. *Suppose that $d_{SL}(\mathrm{Out}(A_\Gamma)) = m$. Let*

$$F(\Gamma') = \max\{|V(\Gamma')| : \Gamma' \subset \Gamma \text{ and } A_{\Gamma'} \leq A_\Gamma \text{ is a free group}\}.$$

Let Λ be a group. Suppose that for each finite index subgroup $\Lambda' \leq \Lambda$, we have:

- *Every homomorphism $\Lambda' \to \mathrm{SL}_m(\mathbb{Z})$ has finite image,*
- *For all $N \leq F(\Gamma)$, every homomorphism $\Lambda' \to \mathrm{Out}(F_N)$ has finite image.*

Then every homomorphism $f : \Lambda \to \mathrm{Out}(A_\Gamma)$ has finite image.

The constant $F(\Gamma)$ can be viewed as the maximal rank of a free subgroup of A_Γ that arises from a subgraph of Γ. Again, the proof of this theorem is split into two steps: we first analyse the action of such a group Λ on $H_1(A_\Gamma)$, using the *projection homomorphisms* for $\mathrm{Out}(A_\Gamma)$ developed by Charney and Vogtmann [20, 21]. These allow one to reduce questions about $\mathrm{Out}(A_\Gamma)$ to questions about outer automorphism

groups of RAAGs associated to subgraphs of Γ, thus allowing for an inductive argument on the size of Γ. We show that the image of Λ must lie in $\overline{\mathrm{IA}}(A_\Gamma)$, the subgroup of $\mathrm{Out}(A_\Gamma)$ that acts trivially on $H_1(A_\Gamma)$. We then prove the following analogue of the theorem of Bass and Lubotzky:

Theorem 4.22. *For any graph Γ, the group $\overline{\mathrm{IA}}(A_\Gamma)$ is residually torsion-free nilpotent.*

This theorem was discovered independently by Toinet [70]. The groups $\mathrm{IA}(A_\Gamma)$ and $\overline{\mathrm{IA}}(A_\Gamma)$ are less well understood than IA_n and $\overline{\mathrm{IA}}_n$, so we will spend time developing the machinery behind the proof of Theorem 4.22.

One would ideally like to understand the lower central series of $\mathrm{IA}(A_\Gamma)$ in detail. However a direct method for studying this lower central series has remained elusive. We therefore replace the lower central series with a coarser central series that we call the Andreadakis–Johnson filtration. Let $\gamma_1(A_\Gamma)$, $\gamma_2(A_\Gamma)$, $\gamma_3(A_\Gamma), \ldots$ be the lower central series of A_Γ. As each term in the lower central series is a characteristic subgroup of A_Γ, for each c there is a homomorphism $\mathrm{Aut}(A_\Gamma) \to \mathrm{Aut}(A_\Gamma/\gamma_{c+1}(A_\Gamma))$. Let G_c be the kernel of this map. Then $G_0 = \mathrm{Aut}(A_\Gamma)$ and $G_1 = \mathrm{IA}(A_\Gamma)$. The Andreadakis–Johnson filtration of $\mathrm{IA}(A_\Gamma)$ is the central series given by $G_1, G_2, G_3, \ldots$. This was first defined in the case of the free group by Andreadakis [3]. This central series can be studied with the following homomorphisms found by Andreadakis for use with $\mathrm{Aut}(F_n)$, and used by Johnson [45] in an analogous situation with surface homeomorphisms:

Proposition 4.13. *Let $\phi \in G_c$, where $c \geq 1$. Define a map*

$$\tau_c(\phi) : H_1(A_\Gamma) \to \gamma_{c+1}(A_\Gamma)/\gamma_{c+2}(A_\Gamma)$$

by $\tau_c(\phi)(\bar{g}) = g^{-1}\phi(g).\gamma_{c+2}(A_\Gamma)$. Then τ_c is a homomorphism

$$\tau_c \colon G_c \to \mathrm{Hom}(H_1(A_\Gamma), \gamma_{c+1}(A_\Gamma)/\gamma_{c+2}(A_\Gamma))$$

with $\ker(\tau_c) = G_{c+1}$.

We say that τ_c is the *cth Johnson homomorphism*. A good understanding of the lower central series of A_Γ will help us to understand the Andreadakis–Johnson filtration. In Chapter 2 we give a description of the work of Duchamp, Krob and Lalonde [31, 32, 50, 51, 48] in this direction. The information we require is contained in the *lower central series algebra* $L_\mathcal{C} = \sum_{c=1}^{\infty} \gamma_c(A_\Gamma)/\gamma_{c+1}(A_\Gamma)$ associated to A_Γ. We show that each term $\gamma_c(A_\Gamma)/\gamma_{c+1}(A_\Gamma)$ is a free abelian group; combined with

Proposition 4.13 this implies each consecutive quotient G_c/G_{c+1} is also a free abelian group. In turn, this implies $\mathrm{IA}(A_\Gamma)$ is residually torsion-free nilpotent. However, we are interested in $\overline{\mathrm{IA}}(A_\Gamma)$ rather than $\mathrm{IA}(A_\Gamma)$, and this is a more subtle object. In Chapter 4 we use machinery developed by Bass and Lubotzky to analyse the image of $G_1, G_2, G_3, \ldots$ in $\overline{\mathrm{IA}}(A_\Gamma)$. The key step involving the lower central series algebra is the following theorem, which we prove in Chapter 2. We use use $Z(_)$ to denote the centre of a Lie algebra.

Theorem 2.56. *Suppose that $Z(A_\Gamma) = 1$. Then $Z(L_\mathcal{C}) = Z((\mathbb{Z}/p\mathbb{Z}) \otimes_\mathbb{Z} L_\mathcal{C}) = 0$ and $Z(L_\mathcal{C}/\oplus_{i>c} L_{\mathcal{C},i})$ is the image of $L_{\mathcal{C},c}$ under the quotient map $L_\mathcal{C} \to L_\mathcal{C}/\oplus_{i>c} L_{\mathcal{C},i}$.*

This information allows us to complete the proof of Theorem 4.22. Even in the free group case, it is not known if the Andreadakis–Johnson filtration coincides with the lower central series of IA_n or if there is an equivalent of Witt's formula ([57], Theorem 5.11) that gives the rank of G_c/G_{c+1} in general. In the free group case, Pettet and Satoh give descriptions of G_1/G_2, G_2/G_3 and G_3/G_4 [64, 66], and it is known that $G_2 = \gamma_2(\mathrm{IA}_n)$ and G_3 is finite index in $\gamma_3(\mathrm{IA}_n)$ (see [64]). The proof that $\gamma_2(\mathrm{IA}_n) = G_2$ goes as follows: one takes Magnus' generating set for IA_n and shows that the generators map to linearly independent elements in the free abelian group that is the image of the first Johnson homomorphism. It follows that the commutator subgroup of IA_n contains G_2, the kernel of this map. As the terms in any central series contain the respective terms of the lower central series, we know that G_2 is contained in $[\mathrm{IA}_n, \mathit{ian}]$, and equality follows. We reprove a result of Day [27] that gives a finite generating set $\mathcal{M}_\Gamma$ of $\mathrm{IA}(A_\Gamma)$, and show that $\gamma_2(\mathrm{IA}(A_\Gamma)) = G_2$ for a general RAAG.

Theorem 4.23. *The first Johnson homomorphism τ_1 maps $\mathcal{M}_\Gamma$ to a free generating set of a subgroup of* $\mathrm{Hom}(H_1(A_\Gamma), \gamma_2(A_\Gamma)/\gamma_3(A_\Gamma))$. *The abelianisation of* $\mathrm{IA}(A_\Gamma)$ *is isomorphic to the free abelian group on the set $\mathcal{M}_\Gamma$, and G_2 is the commutator subgroup of* $\mathrm{IA}(A_\Gamma)$.

The first part of this thesis uses mostly combinatorial methods (although, as previously stated, we do use algebraic results whose proofs arise from techniques in *geometric* group theory). In Part 2, we consider two more modern approaches to automorphisms of free groups. Chapter 6 has a topological flavour, looking at how Stallings' folding algorithm gives a method of decomposing an element of $\mathrm{Aut}(F_n)$ as a product of Whitehead automorphisms, and hence Nielsen automorphisms. This gives a pictorial description of how such a decomposition occurs, via a sequence of folding

operations on a graph with fundamental group F_n. This is in contrast with the more combinatorial methods of Nielsen's original proof that $\mathrm{Aut}(F_n)$ is finitely generated. We use this algorithm to show that if Y is a subset of a generating set of F_n then the subgroup $\mathrm{Fix}(Y)$ of $\mathrm{Aut}(F_n)$ that fixes each element of Y is finitely generated. So too is the subgroup $\mathrm{Fix}_c(Y)$ that fixes each element of Y up to conjugacy. (This can also be proved using methods of McCool [59], though we feel that the proof using the folding algorithm is more informative and intuitive.)

In Chapter 7 we look at the *non-symmetric Lipschitz metric* on Culler and Vogtmann's Outer Space (written CV_n.) Outer Space (introduced in [24]) is a finite dimensional, contractible space on which $\mathrm{Out}(F_n)$ acts properly discontinuously with finite stabilisers. This makes it a central tool in the study of $\mathrm{Out}(F_n)$. Previously, much of the work on Outer Space has focused on its topological properties, but attention has recently shifted to its metric properties.

Given a metric d on a space X, and an action of a group G on X by isometries, a natural problem is to find which elements of G are *elliptic*, *hyperbolic* and *parabolic* with respect to this action. These terms are defined as follows: we start with the displacement function $D : G \to \mathbb{R}_{\geq 0}$ given by

$$D(g) = \inf\{d(x, gx) : x \in X\}.$$

An element $g \in G$ is said to be *elliptic* if $D(g) = 0$ and the infimum is realised, so that g fixes a point in X. We say that g is *hyperbolic* if $D(g) > 0$ and there exists $x \in X$ with $d(x, gx) = D(g)$. We say that g is *parabolic* if $D(g)$ is not realised at a point in X, so that g is neither elliptic nor hyperbolic. We study the action of $\mathrm{Out}(F_n)$ on CV_n with the Lipschitz metric. We do not succeed in completely classifying the elements of $\mathrm{Out}(F_n)$ with respect to this action: we do not yet have a pleasing description of when an element with positive displacement is hyperbolic or parabolic. However, we do determine the value of the displacement function for an arbitrary element $\Phi \in \mathrm{Out}(F_n)$. We show that $D(\Phi)$ is equal to three related concepts: the asymptotic translation length of Φ, the exponential growth rate of Φ, and the maximal Perron–Frobenius eigenvalue of a relative train track representative of Φ. Definitions of these concepts are given in Chapter 7.

With this final result, we reach the end of a path that has taken us through selected parts of a century of study on automorphisms of free groups, from the combinatorial methods championed by Nielsen and Magnus, though the topological ideas introduced by Whitehead and Stallings in the second half of the twentieth century, and ending at the geometric approach to Outer Space that is still in development. We take pleasure

in seeing a similar toolbox form for automorphisms of right-angled Artin groups: we have a Nielsen type generating set thanks to Laurence [52] and Servatius [68], an approach to Whitehead's theory of peak reduction by Day [26], and for certain graphs Charney, Crisp and Vogtmann [19] have defined a topological space that satisfies many of the properties we'd like from an equivalent of Outer Space for $\mathrm{Out}(A_\Gamma)$. We hope to convince the reader that combining these three lines of thought will lead to new results, and exciting new mathematics.

Part I

Rigidity, IA automorphisms, and central series

Chapter 2

The lower central series of a right-angled Artin group

One can often translate problems concerning Lie groups to the world of Lie algebras. When we linearise a problem our life becomes much easier: we understand vector spaces and their endomorphisms very well, and we may use our knowledge here to give us information about the underlying Lie group. By analogy, given a discrete group G, one may form a Lie $\mathbb{Z}$–algebra by taking the direct sum $\sum_{i=1}^{\infty} \gamma_i(G)/\gamma_{i+1}(G)$, where $\gamma_i(G)$ is the ith term in the lower central series; the bracket operation is given by taking commutators in G. In general, this Lie algebra may not give much information about G, but if G is a free group the picture is very nice indeed. The Lie algebra one attains is a *free Lie algebra*, and the structure theory of free Lie algebras not only allows us to obtain information about free groups, but $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$ also. This correspondence is well-known, and is covered in detail by Magnus in Chapter 5 of [57]. The aim of this chapter is to give a description of the analogous theory for right-angled Artin groups. These results are not new, however we feel that a unified summary of key parts of the papers of Duchamp, Krob, and Lalonde [31, 32, 50, 51, 48] will make a useful reference.

We have attempted to make this work as self contained as possible. In particular, we do not assume any results concerning free Lie algebras, which allows the theory of free Lie algebras and the *partially-commutative free Lie algebras* studied here to be developed in parallel. This comes at the cost of assuming certain facts about the combinatorics of words in RAAGs. We feel that this is a reasonable trade-off.

This algebraic approach to the study of the lower central series of a RAAG has much wider implications than the methods in this chapter might suggest. In Chapter 4 we shall use the structure of the lower central series to study $\mathrm{IA}(A_\Gamma)$, and Linnell, Okun, and Schick used the fact the RAAGs are residually torsion-free nilpotent (also

shown by this theory) as part of their proof of the strong Atiyah conjecture for RAAGs [53].

Let Γ be a graph with vertex set V and edge set E. Let ι and τ be the maps that send an edge to its initial and terminal vertices respectively. We will work with a fixed right-angled Artin group determined by Γ and defined like so:

$$A_\Gamma = \langle v \in V | [\iota(e), \tau(e)] : e \in E \rangle.$$

We will assume Γ is finite with vertex set $v_1, \ldots, v_n$.

This chapter is set out as follows: we start with a basic introduction to associative algebras and Lie algebras and recount how, given a *central series* $\mathcal{G} = \{G_i\}_{i=1}^{\infty}$ of a group, one can build a Lie algebra $L_\mathcal{G}$. This is a generalisation of the construction of the Lie algebra associated to the lower central series mentioned before. It is functorial in the sense that if you have two central filtrations $\mathcal{G} = \{G_i\}$ and $\mathcal{H} = \{H_i\}$ of groups G and H respectively, and $\phi : G \to H$ is a homomorphism such that $\phi(G_i) \subset H_i$ for all i, then there is an induced algebra homomorphism $L_\mathcal{G} \to L_\mathcal{H}$.

In Section 2.3 we build up a host of *partially-commutative* objects associated to a right-angled Artin group. Of central importance is the *free partially-commutative monoid* M, which may be viewed as the monoid of positive elements in A_Γ. We define U to be the free $\mathbb{Z}$–module on M. The module U inherits a graded algebra structure, with the grading coming from word length in A_Γ, and multiplication induced by multiplication in A_Γ. One can extend U to an algebra U^∞ by allowing infinitely many coefficients in a sequence of elements of M to be nonzero. U^∞ behaves very much like an algebra of formal power series. For instance, $1 + \mathbf{v_i}$ is a unit in U^∞, with inverse

$$(1 + \mathbf{v_i})^{-1} = 1 - \mathbf{v_i} + \mathbf{v_i}^2 - \mathbf{v_i}^3 + \cdots$$

and if we define U^* to be the group of units of U^∞, the mapping $v_i \mapsto 1 + \mathbf{v_i}$ gives an embedding

$$\mu : A_\Gamma \to U^*,$$

called the *Magnus map*. We define a sequence of subsets $\mathcal{D} = \{D_i\}_{i=1}^{\infty}$ of A_Γ by saying that $g \in D_i$ if and only if $\mu(g)$ is of the form:

$$\mu(g) = 1 + (\text{elements of } U \text{ of degree} \geq i).$$

As U^∞ can be treated like an algebra of formal power series, this allows us to show that the sequence of subsets $\mathcal{D}$ has a particularly nice stucture. This is encapsulated in the following theorem:

Theorem 2.20. *For all k, the set D_k is a subgroup of A_Γ. These subgroups satisfy:*

1. *$\mathcal{D}$ is a central filtration of A_Γ.*
2. *$D_{k+1} \trianglelefteq D_k$ and D_k/D_{k+1} is a finitely generated free abelian group.*
3. *$\gamma_k(A_\Gamma) \subset D_k$.*

As μ is injective, we have $\cap_{k=1}^{\infty} D_k = \{1\}$, and this fact combined with properties (1) and (2) imply that a right-angled Artin group is residually torsion-free nilpotent. If $\mathcal{C}$ is the central filtration given by the lower central series, then property (3) implies that we have a Lie algebra homomorphism $L_\mathcal{C} \to L_\mathcal{D}$. We finish our study of the Magnus map by using it to give a new proof of the normal form theorem for words in right-angled Artin groups.

The free $\mathbb{Z}$–module U has an associated Lie algebra $\mathcal{L}(U)$ consisting of the elements of U and bracket operation $[a, b] = ab - ba$. In Section 2.5, we study the Lie subalgebra L_Γ of $\mathcal{L}(U)$ generated by the set $V = \{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$ by looking at Lalonde's description of the *partially commutative free Lie algebra* determined by the graph Γ [50, 51]. His construction goes as follows. One first defines a subset $LE(M) \subset M$ known as the set of *Lyndon elements* of M. These have a very rigid combinatorial structure. In particular there is a way of assigning a bracketing to each Lyndon element; given a subset $X = \{x_1, \ldots, x_n\}$ of a Lie algebra L, this bracketing induces a $\mathbb{Z}$–module homomorphism $\phi_X : \mathbb{Z}[LE(M)] \to L$. When $X = V$, the induced map $\phi_V : \mathbb{Z}[LE(M)] \to L_\Gamma$ is an isomorphism. This gives a basis of L_Γ as a free $\mathbb{Z}$–module, and allows us to give a universal defining property of L_Γ:

Theorem 2.49. *Let Γ be a graph with vertices $v_1, \ldots, v_n$. Let L be a Lie algebra, and suppose that $X = \{x_1, \ldots, x_n\}$ is a subset of L that satisfies:*

$$[x_i, x_j] = 0 \text{ when } v_i \text{ and } v_j \text{ are connected by an edge in } \Gamma.$$

Then there is a unique algebra homomorphism $\psi_X : L_\Gamma \to L$ such that

$$\psi_X(\mathbf{v_i}) = x_i \text{ for } 1 \leq i \leq r.$$

We use this in Section 2.6 to construct a chain of algebra homomorphisms

$$L_\Gamma \to L_\mathcal{C} \to L_\mathcal{D} \to L_\Gamma$$

and show that the composition of the three maps is the identity on L_Γ. In fact:

Theorem 2.52. *L_Γ, $L_\mathcal{C}$, and $L_\mathcal{D}$ are isomorphic as graded Lie algebras. Furthermore, the central filtrations $\mathcal{C}$ and $\mathcal{D}$ are equal, so that $\gamma_k(A_\Gamma) = D_k$ for all $k \geq 1$.*

We are now able to use Lyndon elements and L_Γ to describe the lower central series of A_Γ in more detail. For instance, Proposition 2.20 now implies:

Theorem 2.53. *If $k \in \mathbb{N}$, then $\gamma_k(A_\Gamma)/\gamma_{k+1}(A_\Gamma)$ is free-abelian, and $A_\Gamma/\gamma_k(A_\Gamma)$ is torsion-free nilpotent.*

A wonderful aspect of Magnus' approach to the study of free groups is how nicely the overall structure of his work translates to right-angled Artin groups. An avid reader is encouraged to compare Section 2.4 of this chapter with Section 5.5 of [57]. The statements contained in this chapter are adapted to deal with the more general setting of RAAGs, however very little work needs to be done in ensuring the proofs then follow through as well.

2.1 Associative algebras, Lie algebras, and central filtrations

An *R–algebra* is a (left) R–module A equipped with a bilinear map:

$$_ \cdot _ : A \times A \to A,$$

which we call multiplication. An R–algebra A is *associative* if multiplication is associative: for all $a, b, c \in A$,

$$(a \cdot b) \cdot c = a \cdot (b \cdot c).$$

In this case we may ignore the brackets when multiplying elements. Furthermore we will often write $a \cdot b$ as simply ab. We say that A is a *Lie algebra* if it satisfies the following pair of identities:

$$a \cdot a = 0 \tag{2.1}$$

$$a \cdot (b \cdot c) + b \cdot (c \cdot a) + c \cdot (a \cdot b) = 0 \tag{2.2}$$

Equation (2.2) is called the Jacobi identity. In this case, we usually call multiplication the bracket operation, and write $[a, b]$ rather than $a \cdot b$. Given any associative algebra A we may form a Lie algebra $\mathcal{L}(A)$ consisting of the elements of A with the bracket operation

$$[a, b] = a \cdot b - b \cdot a.$$

Given a subset $S \subset A$ we define $\langle S \rangle$ to be the smallest subalgebra of A containing S. We say that $\langle S \rangle$ is the *subalgebra of A generated by S*, and if $\langle S \rangle = A$ we say that S *generates* A. We say that S is a *basis* of A if A is free as an R–module on the set S. A map $\phi : A \to B$ is an algebra homomorphism if ϕ respects multiplication and the R–module structure of A and B. If S generates A then ϕ is uniquely determined by where it sends S. If A and B are associative, then a morphism $\phi : A \to B$ induces a morphism $\mathcal{L}(\phi) : \mathcal{L}(A) \to \mathcal{L}(B)$ ($\mathcal{L}$ is a functor from the category of associative algebras to the category of Lie algebras). If U and V are two subalgebras of A, define the algebra $U.V$ to be the linear span of the set $\{u.v : u \in U, v \in V\}$. We recursively define $U^k = U.U^{k-1}$. We say that an algebra A is *graded* if there exist subspaces A_i of A indexed by $\mathbb{N}$ such that $A = \oplus_{i=0}^{\infty} A_i$ and $A_i.A_j \subset A_{i+j}$ for all i, j. If $\oplus_{i=0}^{\infty} A_i$ is a grading of an associative algebra A, then $\oplus_{i=0}^{\infty} A_i$ is also a grading of the associated Lie algebra $\mathcal{L}(A)$. A homomorphism $\phi : A \to B$ of algebras is *graded* if A and B are graded and $\phi(A_i) \subset B_i$ for all $i \in \mathbb{N}$.

2.2 Lie algebras from central filtrations

Let G be a group. Let $\mathcal{G} = \{G_k\}_{k \geq 1}$ be a sequence of subgroups of G such that for all k, l:

$$\begin{aligned} &\text{(F1)} \qquad G_1 = G, \\ &\text{(F2)} \qquad G_{k+1} \leq G_k, \\ &\text{(F3)} \qquad [G_k, G_l] \subset G_{k+l}. \end{aligned}$$

We say that $\mathcal{G}$ is a *central filtration*, or a *central series* of G. The above conditions imply that $G_k \trianglelefteq G$ and $G_{k+1} \trianglelefteq G_k$ for all k. The main construction in this section is a Lie algebra $L_{\mathcal{G}}$ built out of the consecutive quotients G_k/G_{k+1} of $\mathcal{G}$.

One example of a central filtration is $\gamma(G) = \{\gamma_k(G)\}_{k \geq 1}$, the lower central series of G. This is defined recursively by $\gamma_1(G) = G$ and $\gamma_{k+1}(G) = [G, \gamma_k(G)]$. Where it is clear which group we are using, we shall simply write γ_k (or γ) rather than $\gamma_k(G)$ (or $\gamma(G)$). The lower central series is contained in all central filtrations of G:

Proposition 2.1. *Let $\mathcal{G} = \{G_k\}$ be a central filtration of G. Then $\gamma_k \subset G_k$ for all k.*

Proof. We use induction on k. When $k = 1$ we have $G = G_1 = \gamma_1$. Thereafter, if $\gamma_{k-1} \subset G_{k-1}$ then $\gamma_k = [\gamma_1, \gamma_{k-1}] \subset [G_1, G_{k-1}] \subset G_k$ by (F3). □

Central filtrations tell us about residual properties of groups. We say that a central filtration $\mathcal{G}$ is separating if $\cap_{k=1}^{\infty} G_k = \{1\}$. The following proposition encapsulates a key observation that we shall use throughout this thesis:

Proposition 2.2. *Suppose that $\mathcal{G}$ is a central filtration of G that is separating. Furthermore, suppose that each consecutive quotient G_k/G_{k+1} is free-abelian. Then:*

(1) G_k is a normal subgroup of G.

(2) For all k the group G/G_k is torsion-free nilpotent.

(3) G is residually torsion-free nilpotent.

Proof. Part (1) is a general fact about central filtrations. By (F2) and (F3) if $g \in G$ and $h \in G_k$ then $ghg^{-1}h^{-1} \in G_{k+1} \subset G_k$, therefore $ghg^{-1} \in G_k$ and G_k is a normal subgroup of G. Let $\mathcal{H} = \{H_1, H_2, \ldots\}$ be the image of $\mathcal{G}$ in $H = G/G_k$. Then $\mathcal{H}$ is a central filtration of H such that $H_l = \{1\}$ for $l \geq k$. By Proposition 2.1 we have $\gamma_k(H) \subset H_k$, so $\gamma_k(H) = \{1\}$ and H is nilpotent. Let gG_k be a nontrivial element of H. There exists $l < k$ such that $g \in G_l$ but $g \notin G_{l+1}$. As G_l/G_{l+1} is free-abelian, g has infinite order in G_l/G_{l+1}, and therefore has infinite order in G_l/G_k, which can naturally be viewed as a subgroup of H. Hence H is torsion free. As $\mathcal{G}$ is separating, for any nontrivial element $g \in G$ we may find a k such that $g \notin G_k$, whence g is nontrivial in G/G_k, and G is residually torsion-free nilpotent. □

We shall provide a proof that γ is a central filtration in Proposition 2.5. We first have to take a short detour to look at some commutator identities. We use the convention that for $x, y \in G$ we have $[x, y] = xyx^{-1}y^{-1}$, and for conjugation we write ${}^{y}x = yxy^{-1}$.

Lemma 2.3. *Let x, y, z be elements of G. Then the following identities hold:*

$$ {}^{x}y = [x, y].y \tag{2.3} $$

$$ [xy, z] = {}^{x}[y, z].[x, z] = [x, [y, z]].[y, z].[x, z], \tag{2.4} $$

$$ [x, yz] = [x, y].{}^{y}[x, z] = [x, y].[y, [x, z]].[x, z], \tag{2.5} $$

As well as the Witt–Hall identity:

$$ [[x, y], {}^{y}z].[[y, z], {}^{z}x].[[z, x], {}^{x}y] = 1. $$

The reader should be aware that the above equations are different to those that occur in many group theory text books; the commutation and conjugation conventions we use are set up for left, rather than right, actions. The Witt–Hall identity implies the following '3 subgroup' theorem:

Theorem 2.4 (Hall, 1933)**.** *Let X, Y and Z be three normal subgroups of G. Then*

$$[[X,Y],Z] \subset [[Y,Z],X].[[Z,X],Y]$$

Proof. If $x \in X, y \in Y$ and $z \in Z$, then as X and Y are normal, we have ${}^{z}x \in X$ and ${}^{x}y \in Y$. The Witt–Hall identity implies that $[[x,y],{}^{y}z] \in [[Y,Z],X].[[Z,X],Y]$. As Z is a normal subgroup of G we may replace z with $y^{-1}zy$ to show that $[[x,y],z] \in [[Y,Z],X].[[Z,X],Y]$, and the result follows. □

We may now prove the previously promised result:

Proposition 2.5. *The lower central series is a central filtration of G.*

Proof. (F1) holds by definition of the lower central series. For (F2), a simple induction argument shows that $\gamma_k \trianglelefteq G$, therefore if $x \in G$ and $y \in \gamma_k$ we have $[x,y] = {}^{x}y.y^{-1} \in \gamma_k$. As γ_{k+1} is generated by elements of this form, $\gamma_{k+1} \leq \gamma_k$, and (F2) also holds. We are left to show that

$$[\gamma_k, \gamma_l] \subset \gamma_{k+l}$$

for all k, l. We allow l to vary and proceed by induction on k. Note that $[\gamma_1, \gamma_l] = \gamma_{l+1}$ by definition. For the inductive step, suppose that $[\gamma_{k-1}, \gamma_l] = [\gamma_l, \gamma_{k-1}] \subset \gamma_{k+l-1}$ for all l. Then by Theorem 2.4:

$$\begin{aligned}
[\gamma_k, \gamma_l] &= [[\gamma_1, \gamma_{k-1}], \gamma_l] \\
&\subset [[\gamma_{k-1}, \gamma_l], \gamma_1].[[\gamma_l, \gamma_1], \gamma_{k-1}] \\
&\subset [\gamma_{k+l-1}, \gamma_1].[\gamma_{l+1}, \gamma_{k-1}] \\
&\subset \gamma_{k+l}.
\end{aligned}$$

Hence the lower central series satisfies (F1), (F2), and (F3) and is a central filtration of G. □

Now let $\mathcal{G} = \{G_i\}_{i \geq 1}$ be any central filtration of G. Let $L_{\mathcal{G},i} = G_i/G_{i+1}$. As $[G_i, G_i] \subset G_{2i} \subset G_{i+1}$ each $L_{\mathcal{G},i}$ is an abelian group, therefore we can form a $\mathbb{Z}$–module $L_{\mathcal{G}} = \oplus_{i=1}^{\infty} L_{\mathcal{G},i}$. Any element in $L_{\mathcal{G}}$ is of the form $\sum_i x_i G_{i+1}$, where each $x_i \in G_i$ and only finitely many x_i are not equal to the identity. As we are in the

abelian setting, when there is no danger of confusion we will switch between additive and multiplicative notation. For instance, $kx_iG_{i+1} = x_i^k G_{i+1}$, and $x_iG_{i+1} - y_iG_{i+1} = x_iy_i^{-1}G_{i+1}$ in $L_{\mathcal{G}}$. We write

$$x = y \bmod G_i$$

if $xG_i = yG_i$.

Proposition 2.6. *The bracket operation*

$$[\sum_i x_iG_{i+1}, \sum_j y_jG_{j+1}] = \sum_{i,j}[x_i, y_j]G_{i+j+1}$$

gives $L_{\mathcal{G}}$ the structure of a graded Lie $\mathbb{Z}$–algebra.

Proof. We need to show that the bracket operation is well-defined, bilinear, and satisfies the Lie algebra axioms (2.1) and (2.2) from Section 2.1. First note that if $x \in \gamma_i$, $y \in G_j$ and $z \in G_{i+1}$ then $[x, [z, y]] \in G_{2i+j+1}$ and $[z, y] \in G_{i+j+1}$. Therefore by Equation (2.4) of Lemma 2.3, we have

$$\begin{aligned}[xz, y] &= [x, [z, y]].[z, y].[x, y] \\ &= [x, y] \bmod G_{i+j+1}.\end{aligned}$$

Hence the choice of coset representatives in the left hand side of the product does not affect the bracket operation. Similarly, Equation (2.5) shows that the choice of coset representatives in the right hand side of the bracket does not affect the bracket operation, so the bracket operation is well-defined. We may also use Equation (2.4) to show that if $x, y \in G_i$ and $z \in G_j$ then

$$[xy, z] = [x, z] + [y, z] \bmod G_{i+j+1},$$

which gives linearity of the bracket operation in the left hand side. Linearity in the right hand side follows in the same way. For the first of the Lie algebra axioms:

$$\begin{aligned}[\sum_i x_iG_{i+1}, \sum_j x_jG_{j+1}] &= \sum_{i,j}[x_i, x_j]G_{i+j+1} \\ &= \sum_i [x_i, x_i]G_{2i+1} + \sum_{i \neq j}([x_i, x_j] + [x_j, x_i])G_{i+j+1} \\ &= 0,\end{aligned}$$

as $[x_i, x_i] = 1$ and $[x_i, x_j] = [x_j, x_i]^{-1}$ for all i, j. To prove the Jacobi identity, we first note that if $x \in G_i$, $y \in G_j$ and $z \in G_k$ then

$$\begin{aligned}[[x, y], {}^y z] &= [[x, y], [y, z]z] && \text{by (2.3)} \\ &= [[x, y], [y, z]].[[y, z], [[x, y], z]].[[x, y], z] && \text{by (2.5)} \\ &= [[x, y], z] \bmod G_{i+j+k+1},\end{aligned}$$

and similarly we have:

$$[[y,z],{}^{z}x] = [[y,z],x] \bmod G_{i+j+k+1}$$
$$[[z,x],{}^{x}y] = [[z,x],y] \bmod G_{i+j+k+1}.$$

The Witt–Hall identity of Lemma 2.3 then implies that

$$[[x,y],z] + [[y,z],x] + [[z,x],y] = 0 \bmod G_{i+j+k+1},$$

and the Jacobi identity for a general triple of elements in $L_{\mathcal{G}}$ follows. □

The identities (2.4) and (2.5) imply that if G has a generating set $\{x_1, \ldots, x_n\}$ then any consecutive quotient $\gamma_k(G)/\gamma_{k+1}(G)$ of terms in the lower central series is generated by elements of the form $[x_{i_1}, [x_{i_2}, [\cdots [x_{i_{k-1}}, x_{i_k}] \cdots]]].\gamma_{k+1}(G)$. In particular:

Proposition 2.7. *If G is generated by $\{x_1, \ldots, x_n\}$ then $L_{\gamma(G)}$ is generated by the set $\{x_1\gamma_1(G), \ldots, x_n\gamma_1(G)\}$.*

We finish this section with a useful observation:

Proposition 2.8. *Let $\mathcal{G} = \{G_i\}$ and $\mathcal{H} = \{H_i\}$ be central filtrations of groups G and H respectively. Let $\phi : G \to H$ be a homomorphism such that $\phi(G_i) \subset \phi(H_i)$ for all $i \in \mathbb{N}$. Then ϕ induces a graded Lie algebra homomorphism $\Phi : L_{\mathcal{G}} \to L_{\mathcal{H}}$.*

Proof. Define $\Phi(\sum_i x_i G_{i+1}) = \sum_i \phi(x_i) H_{i+1}$. As $\phi(G_i) \subset H_i$, this map is well-defined, and the fact that ϕ is a homomorphism ensures that the map Φ is a graded algebra homomorphism. □

2.3 The cast

In this section we introduce a host of partially-commutative structures associated with A_Γ.

2.3.1 The monoid M_Γ and algebra U_Γ

Let $W(V)$ be the set of (positive) words in $\{v_1, \ldots, v_n\}$. The empty word is denoted by $\emptyset$ or 1. We write $|w|$ to denote the length of a word in $W(V)$. We define $\|w\|$, the *multidegree* of a word $w = v_{p_1}^{e_1} \cdots v_{p_k}^{e_k}$ to be the element of $\mathbb{N}^r$ with ith coordinate given by

$$\sum_{p_j = i} e_j.$$

If $w, w' \in W(V)$, we write $w \leftrightarrow w'$ if there exist $w_1, w_2 \in W(V)$ and $v_i, v_j \in A_\Gamma$ such that $[v_i, v_j] = 1$ and

$$w = w_1 v_i v_j w_2$$
$$w' = w_1 v_j v_i w_2.$$

We then define an equivalence relation on $W(V)$ by saying that $w \sim w'$ if there exist $w_1, \ldots, w_n \in W(V)$ such that

$$w = w_1 \leftrightarrow w_2 \leftrightarrow \cdots \leftrightarrow w_n = w'.$$

Let $M_\Gamma = W(V)/\sim$. Let $\overline{w}$ be the equivalence class of w under the equivalence relation $\sim$. If $w_1 \sim w_1'$ and $w_2 \sim w_2'$ then $w_1 w_2 \sim w_1' w_2'$, therefore multiplication of words in $W(V)$ descends to a multiplication operation on M_Γ, with an identity element given by the equivalence class of the empty word. Similarly, if $w \sim w'$ then $|w| = |w'|$ and $\|w\| = \|w'\|$, so we may define the length and multidegree of an element $m \in M_\Gamma$ to be the respective length and multidegree of a word in $W(V)$ representing m. Length and multidegree are additive with respect to multiplication, so that if $m_1, m_2 \in M_\Gamma$ we have:

$$|m_1.m_2| = |m_1| + |m_2|$$
$$\|m_1.m_2\| = \|m_1\| + \|m_2\|$$

This gives the free $\mathbb{Z}$–module on M_Γ a graded algebra structure in the following way:

Proposition 2.9. *Let U_Γ be the free $\mathbb{Z}$–module with a basis given by elements of M_Γ. Let $U_{\Gamma,i}$ be the submodule of U_Γ spanned by the elements of M_Γ of length i. Then $U_\Gamma = \oplus_{i=0}^{\infty} U_{\Gamma,i}$, and multiplication in M_Γ gives U_Γ the structure of a graded associative $\mathbb{Z}$–algebra.*

As we will be keeping the graph Γ fixed throughout this chapter, we will refer to M_Γ as M, U_Γ simply as U, and a graded piece as U_i rather than $U_{\Gamma,i}$. We will distinguish elements of U from A_Γ by writing positive words in $\{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$ rather than $\{v_1, \ldots, v_n\}$.

2.3.2 U^∞, an ideal X, and the group of units U^*

Let U^∞ be the algebra extending U by allowing infinitely many coefficients of a sequence of positive elements to be non-zero. Any element of U^∞ can be written uniquely as a power series $a = \sum_{i=0}^{\infty} a_i$, where a_i is an element of U_i. We say that a_i is the *homogeneous part* of a of degree i, and a_0 is the *constant term* of a. Each a_i is a linear sum of elements of M_i, so is of the form $a_i = \sum_{m \in M_i} \lambda_m m$, where $\lambda_m \in \mathbb{Z}$. If $a = \sum_{i=0}^{\infty} a_i$ and $b = \sum_{i=0}^{\infty} b_i$ then the homogeneous part of $a.b$ of degree i is

$$c_i = \sum_{j=0}^{i} a_j b_{i-j}.$$

If $a^{(0)}, a^{(1)}, a^{(2)}, \ldots$ is a sequence of elements of U^∞, then the sum $\sum_{j=0}^{\infty} a^{(j)}$ does not always make sense. However, if the set

$$S_i = \{j : a_i^{(j)} \neq 0\}$$

is finite for all i we define $\sum_{j=0}^{\infty} a^{(j)}$ to be the element of U^∞ with homogeneous part of degree i equal to

$$\sum_{j \in S_i} a_i^{(j)}.$$

Let X be the ideal of U^∞ generated by $\mathbf{v_1}, \ldots, \mathbf{v_n}$. Alternatively, X is the set of elements of U^∞ with a trivial constant term. In a similar fashion, X^k is the ideal of U^∞ containing all elements $a \in U^\infty$ such that $a_i = 0$ for all $i < k$.

Let U^* be the group of units of U^∞. One can show (cf. Proposition 2.11) that $a \in U^*$ if and only if $a = \pm 1 + x$ for some $x \in X$. Note that this is much larger than the group of units of U: there is an embedding of A_Γ into U^* called the *Magnus morphism*, or *Magnus map* (Proposition 2.13).

2.4 The Magnus map

To make U^∞ easier to work with, we would like to treat it as a (noncommutative) polynomial algebra. Specifically, we would like to have an idea of 'substitution' of elements of U^∞ 'into other elements of U^∞'. For instance, given a positive word $w = v_{p_1} \ldots v_{p_k}$ in $W(V)$ and $Q_1, \ldots, Q_n$ in U^∞ we may define $w(Q_1, \ldots, Q_n) = Q_{p_1} \ldots Q_{p_k} \in U^\infty$. Suppose that $Q_1, \ldots, Q_n$ satisfy

$$Q_i Q_j = Q_j Q_i \text{ for all } i, j \text{ such that } [v_i, v_j] = 1. \tag{2.6}$$

If w and w' are words such that $w \leftrightarrow w'$ then

$$w(Q_1, \ldots, Q_n) = w'(Q_1, \ldots, Q_n).$$

It follows that if w and w' represent the same element of M, the above equality also holds. Therefore we may define $m(Q_1, \ldots, Q_n) = w(Q_1, \ldots, Q_n)$, where w is any word in the equivalence class m. This definition respects multiplication in M, so that for $m_1, m_2 \in M$ we have:

$$m_1(Q_1, \ldots, Q_n)m_2(Q_1, \ldots, Q_n) = m_1m_2(Q_1, \ldots, Q_n). \tag{2.7}$$

We can't quite substitute variables in any element of U^∞ with this level of generality; for example it is not possible to set $x = 1$ in

$$1 + x + x^2 + x^3 + \cdots .$$

However, as long as $Q_1, \ldots, Q_n$ have a trivial constant part (in other words they all lie in the ideal X) this problem does not occur.

Proposition 2.10. *Let $Q_1, \ldots, Q_n$ be elements of X which satisfy condition* (2.6)*. Then the mapping*

$$\mathbf{v_i} \mapsto Q_i$$

may be extended to an algebra morphism $\phi : U^\infty \to U^\infty$.

Proof. Let $a = \sum_{i=0}^{\infty} a_i$, with $a_i = \sum_{m \in M_i} \lambda_m m$. We define:

$$\phi(a_i) = \sum_{m \in M_i} \lambda_m m(Q_1, \ldots, Q_n).$$

If $|m| = i$ then as $Q_j \in X$ for all j, it follows that $m(Q_1, \ldots, Q_n)$ lies in X^i. Therefore the smallest nonzero homogeneous part of $\phi(a_i)$ is of degree at least i. Hence the sum $\phi(a) = \sum_{i=1}^{\infty} \phi(a_i)$ is well defined. It is clear from the definition that ϕ is well-behaved under addition and scalar multiplication. Equation (2.7) tells us that ϕ also behaves well under multiplication, and is an algebra homomorphism. □

Such substitutions make our life much easier while working in U^∞; this is exemplified by the following three propositions:

Proposition 2.11. *If a is of the form $a = 1 + \sum_{i=1}^{\infty} a_i$, then $a \in U^*$ and*

$$a^{-1} = 1 - (a_1 + a_2 + \cdots) + (a_1 + a_2 + \cdots)^2 - \ldots = 1 + \sum_{i=1}^{\infty} c_i.$$

Here $c_1 = -a_1$ and $c_i = -\sum_{j=0}^{i-1} c_j a_{i-j} = -\sum_{j=1}^{i} a_j c_{i-j}$ recursively.

Proof. One first checks that if $a = 1 + \mathbf{v_i}$ then the element $a^{-1} = 1 - \mathbf{v_i} + \mathbf{v_i}^2 - \cdots$ satisfies $a.a^{-1} = a^{-1}.a = 1$. We then attain the general formula for an element of the form $a = 1 + x$ with $x \in X$ by applying the algebra homomorphism given by Proposition 2.10 under the mapping $\mathbf{v_i} \mapsto x$ for all i. The recursive formula is obtained by equating homogeneous parts in the equation $a^{-1}.a = a.a^{-1} = 1$. □

Proposition 2.12. *Let $x, y \in X$. Then the following formulas hold:*

$$(1+x)(1+y)(1+x)^{-1} = 1 + y + (xy - yx)\sum_{i=0}^{\infty}(-1)^i x^i, \tag{2.8}$$

$$(1+x)(1+y)(1+x)^{-1}(1+y)^{-1} = 1 + (xy - yx)\sum_{i,j=0}^{\infty}(-1)^{i+j} x^i y^j. \tag{2.9}$$

Proof. As in the proof of Proposition 2.11, we first note that these identities hold for $x = \mathbf{v_i}$ and $y = \mathbf{v_j}$ for any i and j. For the general case, we wish to apply Proposition 2.10. If $xy = yx$ then we may pick any i and j and study the algebra homomorphism induced by the mappings $\mathbf{v_i} \mapsto x$, $\mathbf{v_j} \mapsto y$, and $\mathbf{v_k} \mapsto 0$ when $k \neq i, j$. If $xy \neq yx$ then in particular A_Γ is not abelian: in this case pick vertices v_i and v_j such that $[v_i, v_j] \neq 1$, and use the same map as above. □

Proposition 2.13. *The mapping $v_i \mapsto 1 + \mathbf{v_i}$ induces a homomorphism $\mu : A_\Gamma \to U^*$.*

Proof. The mapping $v_i \mapsto 1 + \mathbf{v_i}$ induces a homomorphism $\overline{\mu} : F(V) \to U^*$ from the free group on the set V. If $[v_i, v_j] = 1$ in A_Γ then $\mathbf{v_i v_j} - \mathbf{v_j v_i} = 0$ in U^∞, therefore by Equation (2.9), relations in the standard presentation of A_Γ are sent to the identity in U^*, and $\overline{\mu}$ descends to a homomorphism $\mu : A_\Gamma \to U^*$. □

The homomorphism μ is called the *Magnus map*, and is the central object of study in this section. Our first task is to gain some understanding of the image of a generic element of A_Γ under μ.

Definition 2.14. We say that an element $m \in M$ is *square-free* if for all words $w \in W(V)$ representing m there exists no element $v \in V(\Gamma)$ such that vv occurs as a subword of w.

We will now relate square-free elements of M to reduced words representing elements of A_Γ. (Note that our words representing elements of A_Γ are in $W(V \cup V^{-1})$ rather than just $W(V)$).

Definition 2.15. Let $g \in A_\Gamma$ and suppose that $w = v_{p_1}^{e_1} \cdots v_{p_k}^{e_k}$ is a word representing g with $e_i \in \mathbb{Z}$. We say that w is *fully reduced* if $e_i \neq 0$ for all i and for all i, j such that $v_{p_i} = v_{p_j}$ there exists $i < l < j$ such that $[v_{p_i}, v_{p_l}] \neq 1$.

We define three moves on the set a words of the form $w = v_{p_1}^{e_1} \cdots v_{p_k}^{e_k}$:

(M1) Remove $v_{p_i}^{e_i}$ if $e_i = 0$.

(M2) Replace the subword $v_{p_i}^{e_i} v_{p_{i+1}}^{e_{i+1}}$ with $v_{p_i}^{e_i+e_{i+1}}$ if $p_i = p_{i+1}$.

(M3) Replace the subword $v_{p_i}^{e_i} v_{p_{i+1}}^{e_{i+1}}$ with $v_{p_{i+1}}^{e_{i+1}} v_{p_i}^{e_i}$ if $[v_{p_i}, v_{p_{i+1}}] = 1$.

Given any word w representing g we may find a fully reduced representative of g by applying a sequence of moves of the form (M1), (M2), and (M3). Moves of type (M3) are called *swaps*. If $w = v_{p_1}^{e_1} v_{p_2}^{e_2} \cdots v_{p_k}^{e_k}$ is fully reduced then $\mathbf{v_{p_1} v_{p_2} \cdots v_{p_k}}$ is square-free. The following key lemma shows that we can find this square-free form in the kth homogeneous part of $\mu(g)$. We will use $\mu(g)_i$ to denote the ith homogeneous part of $\mu(g)$.

Lemma 2.16. *Let g be a nontrivial element of A_Γ. There exists $k \in \mathbb{N}$ such that k is the largest integer such that there is a square-free element $m \in M_k$ with nonzero coefficient λ_m in the decomposition of $\mu(g)_k$. This element is unique. Furthermore, if $v_{p_1}^{e_1} v_{p_2}^{e_2} \cdots v_{p_l}^{e_l}$ is a fully reduced representative of g then $l = k$, $\mathbf{v_{p_1} \cdots v_{p_l}} = m$, and $e_1 \cdots e_l = \lambda_m$.*

Proof. By an induction argument on e_i , we have

$$\mu(v_{p_i}^{e_i}) = 1 + e_i \mathbf{v_{p_i}} + \mathbf{v_{p_i}^2} u_i$$

for some $u_i \in U^*$. Therefore if $v_{p_1}^{e_1} v_{p_2}^{e_2} \cdots v_{p_k}^{e_k}$ is a fully reduced representative of g, we have:

$$\begin{aligned} \mu(g) &= \mu(v_{p_1}^{e_1}) \mu(v_{p_2}^{e_2}) \cdots \mu(v_{p_k}^{e_k}) \\ &= (1 + e_1 \mathbf{v_{p_1}} + \mathbf{v_{p_1}^2} u_1)(1 + e_2 \mathbf{v_{p_2}} + \mathbf{v_{p_2}^2} u_2) \cdots (1 + e_k \mathbf{v_{p_k}} + \mathbf{v_{p_k}^2} u_k). \end{aligned}$$

In this expansion we see that any positive element occurring with length greater than k must contain $\mathbf{v_{p_i}^2}$ as a subword for some i, and the only element of length k without such a subword is $m = \mathbf{v_{p_1} \cdots v_{p_k}}$, with coefficient $\lambda_m = e_1 \cdots e_k$. As $\mu(g)$ is independent of the choice of fully reduced representative of g, every fully reduced representative $v_{q_1}^{f_1} \ldots v_{q_l}^{f_l}$ must satisfy $l = k$, with $\mathbf{v_{q_1} \cdots v_{q_l}} = m$ and $f_1 \cdots f_l = \lambda_m$. □

We have shown that for every nontrivial $g \in A_\Gamma$ there exists $k > 0$ such that $\mu(g)_k$ is nontrivial.

Corollary 2.17. *The homomorphism $\mu : A_\Gamma \to U^*$ is injective.*

We may now use μ to study the lower central series of A_Γ.

Definition 2.18. Let $g \in A_\Gamma$. We define the *derivation* $\delta(g)$ of g to be equal to $\mu(g)_k$, where k is the smallest integer ≥ 1 such that $\mu(g)_k \neq 0$. If no such k exists, then $g = 1$ and we define $\delta(g) = 0$.

The derivation $\delta : A_\Gamma \to U$ satisfies the following properties:

Lemma 2.19. *Let $g, h \in A_\Gamma$ and suppose that $\delta(g) = \mu(g)_k$ and $\delta(h) = \mu(h)_l$.*

1. *For all integers N, $\delta(g^N) = N\mu(g)_k$.*
2. *If $k < l$ then $\delta(gh) = \delta(hg) = \mu(g)_k$.*
3. *If $k = l$ and $\mu(g)_k + \mu(h)_l \neq 0$ then*
$$\delta(gh) = \delta(hg) = \mu(g)_k + \mu(h)_l.$$
4. *If $k = l$ and $\mu(g)_k + \mu(h)_l = 0$ then either*
$$gh = 1 \text{ or } \delta(gh) \in X^{k+1}.$$
5. *If $\mu(g)_k\mu(h)_l - \mu(h)_l\mu(g)_k \neq 0$ then*
$$\delta([g,h]) = \mu(g)_k\mu(h)_l - \mu(h)_l\mu(g)_k.$$
6. *If $\mu(g)_k\mu(h)_l - \mu(h)_l\mu(g)_k = 0$ then either*
$$[g,h] = 0 \text{ or } \delta([g,h]) \in X^{k+l+1}.$$

Proof. Parts (2) , (3) and (4) follow from standard properties of multiplication in U^∞. Part (1) follows from part (3), an induction argument on $N > 0$, and induction on $N < 0$. Parts (5) and (6) follow from Equation (2.9) in Proposition 2.12. □

Let $D_k = \{g \in A_\Gamma : \mu(g)_l = 0 \text{ if } 0 < l < k\}$. Alternatively, D_k is the set of elements $g \in A_\Gamma$ such that either $g = 1$ or $\delta(g) \in X^k$.

Proposition 2.20. *For all k, the set D_k is a subgroup of A_Γ and these subgroups satisfy:*

1. *$\mathcal{D} = \{D_i\}_{i=1}^\infty$ is a central filtration of A_Γ.*
2. *D_k/D_{k+1} is a finitely generated free abelian group.*

3. $\gamma_k(A_\Gamma) \subset D_k$.

Proof. Parts (2)–(4) of Lemma 2.19 imply that D_k is a subgroup of A_Γ. By definition, $D_1 = A_\Gamma$ and $D_{k+1} \leq D_k$ for all k. Also, if $g \in D_k$ and $h \in D_l$, then $[g, h] \in D_{k+l}$ by parts (5) and (6) of Lemma 2.19. Therefore $\mathcal{D} = \{D_i\}$ satisfies the requirements (F1), (F2) and (F3) given in Section 2.2 and is a central filtration of A_Γ. For part (2), we define the map $\phi : D_k \to U_k$ by defining $\phi(g) = \mu(g)_k$. Equivalently:

$$\phi(g) = \begin{cases} \delta(g) & \text{if } \delta(g) = \mu(g)_k \\ 0 & \text{otherwise, when } \delta(g) \in X^{k+1}. \end{cases}$$

Parts (2)–(4) of Lemma 2.19 imply that ϕ is a homomorphism to U_k, with kernel D_{k+1}. Therefore the quotient group D_k/D_{k+1} is isomorphic to a subgroup of U_k. As U_k is a finitely generated free abelian group, so is D_k/D_{k+1}. Part (3) is satisfied for all central filtrations of A_Γ by Proposition 2.1. □

As $\mathcal{D}$ is a central filtration of A_Γ, we have $\gamma_i(A_\Gamma) \subset D_i$ for all i, and as the Magnus map is injective, $\cap_{i=1}^{\infty} D_i = \{1\}$. Hence we may apply Proposition 2.2 to the central filtration $\mathcal{D}$ to obtain:

Theorem 2.21. *The intersection* $\cap_{i=1}^{\infty} \gamma_i(A_\Gamma) = \{1\}$ *and* A_Γ *is residually torsion-free nilpotent.*

We finish this section with a proof of a normal form theorem for elements of A_Γ. This is reasonably well-known; Green's thesis [39] contains a combinatorial proof involving case-by-case analysis. Green's work also extends more generally to graph products of groups. We give a proof for RAAGs using the Magnus map. The first step is an immediate consequence of Lemma 2.16:

Proposition 2.22. *Let* $g \in A_\Gamma$. *Let* $w = v_{p_1}^{e_1} \cdots v_{p_k}^{e_k}$ *and* $w' = v_{q_1}^{f_1} \cdots v_{q_l}^{f_l}$ *be two fully reduced representatives of* g. *Then* $k = l$.

In fact, we can prove something much more detailed:

Theorem 2.23. *Let* $g \in A_\Gamma$. *Let* $w = v_{p_1}^{e_1} \cdots v_{p_k}^{e_k}$ *and* $w' = v_{q_1}^{f_1} \cdots v_{q_k}^{f_k}$ *be two fully reduced representatives of* g. *Then we may obtain* w *from* w' *by a sequence of swaps (moves of the form* $v_{p_i}^{e_i} v_{p_{i+1}}^{e_{i+1}} \mapsto v_{p_{i+1}}^{e_{i+1}} v_{p_i}^{e_i}$ *when* $[v_{p_i}, v_{p_{i+1}}] = 1$*).*

Proof. We proceed by induction on k. We first look at the element $v_{p_1}^{-e_1} g \in A_\Gamma$. Note that $v_{p_2}^{e_2} \cdots v_{p_k}^{e_k}$ and $v_{p_1}^{-e_1} v_{q_1}^{f_1} \cdots v_{q_k}^{f_k}$ are two representatives of $v_{p_1}^{-e_1} g$, and the former

representative is fully reduced. By Proposition 2.22 the latter cannot be fully reduced, so there exists l such that $q_l = p_1$ and $[v_{p_1}, v_{q_i}] = 1$ for $i \leq l$. If $f_l \neq e_1$, then

$$v_{q_1}^{f_1} \cdots v_{q_l}^{f_l - e_1} \cdots v_{q_k}^{f_k}$$

is a fully reduced representative of $v_{p_1}^{-e_1} g$, however this also contradicts Proposition 2.22. Therefore $e_1 = f_l$, and after applying a sequence of swaps to w' we may assume that $v_{p_1} = v_{q_1}$ and $e_1 = f_1$. By induction, $v_{p_2}^{e_2} \cdots v_{p_k}^{e_k}$ may be obtained from $v_{q_2}^{f_2} \cdots v_{q_k}^{f_k}$ by a sequence of swaps, therefore w may be obtained from w' by a sequence of swaps. □

Given $g \in A_\Gamma$, let $init(g)$ (respectively $term(g)$) be the set of vertices of Γ that can occur as the initial (respectively terminal) letter of a fully reduced word representing g. We say that g is *positive* if $g = 1$ or g can be written as a product $v_1^{e_1} \cdots v_k^{e_k}$ with $e_i > 0$ for all i. As any two fully reduced representatives may be obtained from each other by a sequence of swaps, we have the following immediate corollaries:

Corollary 2.24. *For any $g \in A_\Gamma \setminus \{1\}$, the sets $init(g)$ and $term(g)$ form cliques in Γ: any pair of vertices in $init(g)$ or $term(g)$ commute.*

Corollary 2.25. *The monoid M is isomorphic to the set of positive elements of A_Γ under multiplication.*

2.5 Lyndon elements of M

We will now study the Lie subalgebra of $\mathcal{L}(U)$ generated by the set $\{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$. We call this subalgebra L_Γ. The approach is as follows: we first introduce a subset of M called the set of *Lyndon elements*, $LE(M)$. We describe a method for supplying each Lyndon element with a bracketing. If L is a Lie algebra and $X = \{x_1, \ldots, x_n\} \subset L$ then this bracketing induces a homomorphism (as $\mathbb{Z}$–modules) $\phi_X : \mathbb{Z}[LE(M)] \to L$. In the case that $X = \{\mathbf{v_1}, \ldots, \mathbf{v_n}\} \subset L_\Gamma$ we call this induced homomorphism ℓ, and show that ℓ is bijective. Thus we obtain a basis of L_Γ in terms of bracketed Lyndon elements. In general, if $X \subset L$ satisfies

$$[x_i, x_j] = 0 \text{ if } [v_i, v_j] = 1$$

then we will show that $\phi_X \ell^{-1} : L_\Gamma \to L$ is an algebra homomorphism taking $\mathbf{v_i}$ to x_i. This property will then be used in the next section to show that L_Γ and the lower central series algebra of A_Γ are isomorphic.

We deviate here from the approach in Magnus, and instead follow the paper of Lalonde [50]. The analogous free group version is contained in Chapter 5 of [54], and we must start in this world. We first define a lexicographic order on $W(V)$:

Definition 2.26. The *lexicographic ordering* on $W(V)$ is the unique total order $<$ on $W(V)$ that satisfies the following:

1. For any nonempty word w, we have $\emptyset < w$.
2. If w_1 and w_2 are distinct nonempty words and $x, y \in W(V)$ such that $w_1 = v_i x$ and $w_2 = v_j y$, then $w_1 < w_2$ if either
   (a) $i < j$ or:
   (b) $i = j$ and $x < y$.

In particular, $\emptyset < v_1 < v_2 < \ldots < v_n$. We state two basic properties of this order:

Lemma 2.27. *Let $x, y, z \in W(V)$.*

- *if $y < z$ then $xy < xz$*
- *if $|x| \geq |y|$ and $x < y$ then $xz < yz$*

The above lemma remains valid if we replace all occurrences of strong inequalities with weak inequalities. The natural projection $\pi : W(V) \to M$, when coupled with the ordering of $W(V)$, gives us a way of choosing a representative in $W(V)$ for each $m \in M$:

Definition 2.28. Let $m \in M$. Then we define $std(m) \in W(V)$, the *standard representative* of m to be the largest element of $\pi^{-1}\{m\}$ with respect to the lexicographic order.

We then define a total order on M as follows: if $a, b \in M$ we say

$$a < b \text{ if and only if } std(a) < std(b).$$

In view of Lemma 2.27, the following is true:

Lemma 2.29. *Let $a, b, c \in M$*

- *$std(ab) \geq std(a)std(b) \geq std(a)$*
- *If $b < c$ then $std(a)std(b) < std(a)std(c)$*
- *If $|a| \geq |b|$ and $a < b$, then $std(a)std(c) < std(b)std(c)$*

2.5.1 Lyndon words

We now describe the notion of *Lyndon words*. These were first introduced by Chen, Fox, and Lyndon in [23]. In this paper, the authors show that in the free group case, the groups D_k introduced in the last section are equal to the terms of the lower central series of F_n, and they give an algorithm to determine a presentation of a consecutive quotient γ_k/γ_{k+1} of the lower central series for any finitely presented group. This algorithm is quite complicated, however we shall use the notion of *Lyndon elements* in M, introduced by Lalonde in [50], to give a simple algorithm to describe γ_k/γ_{k+1} in an arbitrary right-angled Artin group. Chen, Fox, and Lyndon also relate coefficients of elements in $\mu(g)$ to *Fox derivatives*. Unfortunately these have no natural analogue in the partially commutative setting.

We say that w_1 and w_2 are *conjugate* in $W(V)$ if there exist $x, y \in W(V)$ such that $w_1 = xy$ and $w_2 = yx$. Alternatively, w_1 and w_2 are conjugate if they are conjugates in the free group F_n in the usual sense, where $W(V)$ is viewed as a subset of F_n. The *conjugacy class* of w in $W(V)$ is the set of all elements conjugate to w in $W(V)$. A word w is *primitive* if there does not exist $x, y \in W(V) \setminus \{\emptyset\}$ such that $w = xy = yx$.

Definition 2.30. $w \in W(V)$ is a *Lyndon word* if it is nontrivial, primitive and minimal with respect to the lexicographic ordering in its conjugacy class.

Example 2.31. If $V = \{v_1, v_2, v_3, v_4\}$ then v_i is a Lyndon word for all i, and $v_1v_2v_1v_3$ and $v_1v_1v_2$ are Lyndon words. v_1v_1 is not a Lyndon word as it is not primitive, and $v_1v_3v_1v_2$ is not a Lyndon word as it is not minimal in its conjugacy class ($v_1v_2v_1v_3$ is).

There is an assortment of equivalent definitions of Lyndon words.

Theorem 2.32 ([23], Theorem 1.4)**.** *Let $w \in W(V)$. The following are equivalent:*

1. *w is a Lyndon word.*
2. *For all $x, y \in W(V) \setminus \{\emptyset\}$ such that $w = xy$, $w < y$.*
3. *Either $w = v_i$ for some i or there exist Lyndon words x and y with $x < y$ such that $w = xy$.*

The third of these characterisations is particularly appealing, as it allows one to build up a list of Lyndon words recursively.

Example 2.33. If $V = \{v_1, v_2, v_3\}$, then the Lyndon words of length less than or equal to 3 are:

$$v_1,\ v_2,\ v_3,$$

$$v_1v_2,\ v_1v_3,\ v_2v_3,$$

$$v_1v_1v_2,\ v_1v_1v_3,\ v_1v_2v_3,\ v_2v_2v_3,\ v_1v_2v_2,\ v_1v_3v_3,\ v_2v_3v_3.$$

Note that the decomposition of a Lyndon word of length > 1 as a product of two smaller Lyndon words assured to us by part (3) of Theorem 2.32 is not always unique. In this example $v_1v_2v_3$ may be decomposed as $v_1.v_2v_3$ and $v_1v_2.v_3$.

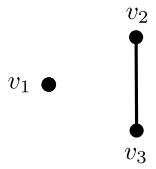

Figure 2.1: A small example graph Γ.

2.5.2 Lyndon elements

Lyndon elements are the natural generalisations of Lyndon words to the partially commutative setting. Defining conjugation here is more tricky. We first say that two elements m_1, m_2 of M are *transposed* if there exist $x, y \in M$ such that $m_1 = xy$ and $m_2 = yx$. Unfortunately transposition is not an equivalence relation; if Γ is the graph shown in Figure 2.1, then

$$v_2v_1v_3 \leftrightarrow_{trans.} v_1v_3v_2 = v_1v_2v_3 \leftrightarrow_{trans.} v_3v_1v_2,$$

however $v_3v_1v_2$ cannot be obtained from $v_2v_1v_3$ by a single transposition. We therefore say two elements of M are *conjugate* if one can be obtained from the other by a sequence of transpositions. Equivalently, two elements are conjugate in M if and only if they are conjugate in A_Γ in the group theoretic sense (when M is viewed as a subset of A_Γ). The set of all elements in M conjugate to m is its *conjugacy class*. We say that m is primitive if there do not exist nontrivial x and y in M such that $m = xy = yx$.

Definition 2.34. $m \in M$ is a *Lyndon element* if it is nontrivial, primitive, and minimal with respect to the ordering of M in its conjugacy class.

Given $g \in A_\Gamma$, we remind the reader that $init(g)$ is the set of vertices that can appear as the initial letter in reduced words representing g.

Proposition 2.35 ([48], Corollary 3.2)**.** *If m is a Lyndon element, then $init(m)$ is a single vertex.*

Given $m \in M$, we say that $v_i \in \zeta(m)$ if either $v_i \in supp(m)$ or there exists $v_j \in supp(m)$ such that $[v_i, v_j] \neq 1$. Equivalently $v_i \in \zeta(m)$ if and only if either $v_i \in supp(m)$ or $v_i m \neq m v_i$. In a similar fashion to Lyndon words, there is a selection of equivalent definitions of Lyndon elements.

Theorem 2.36 ([48], Propositions 3.5, 3.6, and 3.7)**.** *Let $m \in M$. The following are equivalent.*

1. *m is a Lyndon element.*
2. *For all $x, y \in M \setminus \{1\}$ such that $m = xy$, $m < y$.*
3. *Either $|m| = 1$ or there exist Lyndon elements x, y such that $x < y$, $init(y) \in \zeta(x)$ and $m = xy$.*
4. *$std(m)$ is a Lyndon word.*

Once again, the third part of the classification gives a simple recursive process for writing down Lyndon elements.

Example 2.37. If Γ is the small example graph of Figure 2.1, then the Lyndon elements of length ≤ 3 are:

$$v_1,\ v_2,\ v_3$$

$$v_1v_2,\ v_1v_3$$

$$v_1v_1v_2,\ v_1v_1v_3,\ v_1v_2v_2,\ v_1v_2v_3,\ v_1v_3v_3$$

The words given here are a subset of the set of Lyndon words on $\{v_1, v_2, v_3\}$. So for example, the element v_2v_3 does not appear in this list as $v_3 \notin \zeta(v_2) = \{v_1, v_2\}$. As with Lyndon words, the decomposition of a Lyndon elements of length > 1 as a product of Lyndon elements is not necessarily unique. In this example $v_1v_2v_3$ has two possible decompositions as $v_1v_2.v_3$ and $v_1v_3.v_2$.

2.5.3 The standard factorisation of a Lyndon element

We now give each Lyndon element a unique 'bracketing'. If m is a Lyndon element of length greater than 1, there may exist many pairs of Lyndon elements x and y such that $m = xy$. If y is minimal in the lexicographic ordering out of all such pairs, we say that $S(m) = (x, y)$ is the *standard factorisation* of m. The standard factorisation behaves well with respect to standard decompositions:

Theorem 2.38 ([51], Proposition 2.1.10)**.** *If $S(a) = (x, y)$ is the standard factorisation of a, then $std(a) = std(x)std(y)$.*

Note that if x, y are any two elements of M with $x < y$ then $std(x)std(y)$ is a Lyndon word, and is strictly less than its nontrivial conjugates, hence

$$std(x)std(y) < std(y)std(x). \tag{2.10}$$

We shall use this trick repeatedly in the work that follows. There is one final combinatorial fact we need before we can move on:

Theorem 2.39 ([51], Proposition 2.3.9)**.** *If a and b are Lyndon elements with $a < b$ and $init(b) \in \zeta(a)$, then $S(ab) = (a, b)$ if and only if $|a| = 1$ or $S(a) = (x, y)$ and $y \geq b$.*

Example 2.40. We now have a recursive way of giving a bracketing to any Lyndon element. Given $m \in M$, take its standard factorisation $S(m) = (a, b)$, and define the bracketing on m to be equal to $[[a], [b]]$, where $[_]$ denotes the bracketing on a and b respectively. In our small example graph, the only interesting case is $std(v_1v_2v_3) = v_1v_3v_2 = std(v_1v_3)std(v_2)$. We then obtain the following bracketing on Lyndon elements of length 3:

$$[v_1, [v_1, v_2]],\ [v_1, [v_1, v_3]],\ [[v_1, v_2], v_2],\ [[v_1, v_3], v_2],\ [[v_1, v_3], v_3].$$

2.5.4 A basis theorem for the algebra L_Γ

Let $LE(M)$ be the set of Lyndon elements of M. Let $\mathbb{Z}[LE(M)]$ be the free $\mathbb{Z}$–module with basis $LE(M)$.

Definition 2.41. Let L be a Lie algebra, and let $X = \{x_1, x_2, \ldots, x_n\}$ be a subset of L. Let $\phi_X : \mathbb{Z}[LE(M)] \to L$ be the $\mathbb{Z}$–module homomorphism defined recursively as follows:

$$\begin{aligned} \phi_X(\mathbf{v_i}) &= x_i && \text{for all } i \\ \phi_X(a) &= [\phi_X(x), \phi_X(y)] && \text{if } |a| > 1 \text{ and } S(a) = (x, y). \end{aligned}$$

Example 2.42. Let L_Γ be the Lie subalgebra of $\mathcal{L}(U)$ generated by the set $V = \{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$. We attain a $\mathbb{Z}$–module homomorphism $\phi_V : \mathbb{Z}[LE(M)] \to L_\Gamma$. We write $\phi_V = \ell$. The map ℓ can be thought of as the bracketing procedure for Lyndon elements described above.

The following technical lemma gives us a way of understanding the bracket operation in $\mathcal{L}(U)$.

Lemma 2.43. *Suppose that $f = \sum_{b\in I} \alpha_b b$ and $g = \sum_{c\in J} \beta_c c$ are two homogeneous elements in U^∞, so that $|b| = |b'|$ for all $b, b' \in I$ and $|c| = |c'|$ for all $c \in J$. Let x be the minimal element in I with α_x nonzero and y be the minimal element in J with β_y nonzero. Suppose that x and y are Lyndon elements, $x < y$ and $init(y) \in \zeta(x)$, so that xy is a Lyndon element. Then*

- *$[f, g]$ is a homogeneous element of U^∞ of degree $|xy|$;*
- *xy is the minimal element of M with nonzero coefficient in $[f, g]$;*
- *The coefficient of xy in $[f, g]$ is $\alpha_x\beta_y$.*

Furthermore, if f and g are homogeneous with respect to multidegree, so that $\|b\| = \|b'\|$ for all $b, b' \in I$ and $\|c\| = \|c'\|$ for all $c, c' \in J$, then $[f, g]$ is homogeneous with respect to the multidegree $\|xy\|$.

Proof. We have:

$$[f, g] = \sum_{b\in I} \sum_{c\in J} \alpha_b\beta_c(bc - cb), \tag{2.11}$$

where we may assume that $b \geq x$ and $c \geq y$, and $|bc| = |cb| = |xy|$. If either $b > x$ or $c > y$ then by Lemma 2.29:

$$\begin{aligned} std(bc) &\geq std(b)std(c) \\ &> std(x)std(y) \\ &= std(a). \end{aligned}$$

By the identities in Lemma 2.29 and the identity (2.10) we also have:

$$\begin{aligned} std(cb) &\geq std(c)std(b) \\ &\geq std(y)std(x) \\ &> std(x)std(y) \\ &= std(a). \end{aligned}$$

Hence $cb > xy$ for all $b \in I, c \in J$ and $bc \geq xy$ with equality if and only if $b = x$ and $c = y$, so the coefficient of a in the above sum is $\alpha_x \beta_y$. The final remark about homogeneity with respect to multidegree follows as if f and g are homogeneous with respect to multidegree then $\|bc\| = \|cb\| = \|xy\|$ for all $b \in I$ and $c \in J$. □

Proposition 2.44. *For each $a \in LE(M)$, there exists a subset $I \subset M$ and a set of nonzero integers $\{\alpha_b\}_{b \in I}$ indexed by I such that*

$$\ell(a) = \sum_{b \in I} \alpha_b b.$$

Furthermore, $a \in I$ with $\alpha_a = 1$, and for all $b \in I$ we have $\|b\| = \|a\|$ and $b \geq a$.

Proof. We proceed by induction on $|a|$. If $|a| = 1$ then $\ell(a) = a$ and we are done. Suppose that $|a| > 1$. Let $S(a) = (x, y)$ be the standard decomposition of a. By our inductive hypothesis we may write

$$\ell(x) = \sum_{b \in I} \alpha_b b \text{ and } \ell(y) = \sum_{c \in J} \beta_c c$$

with $b \geq x$, $c \geq y$ and $\|b\| = \|x\|$, $\|c\| = \|y\|$ for all $b \in I$ and $c \in J$. Furthermore we may assume $\alpha_x = \beta_y = 1$. As $\ell(a) = [\ell(x), \ell(y)]$ the result follows from Lemma 2.43. □

A consequence of the above theorem is that the image of $LE(M)$ under ℓ forms a linearly independent set.

Corollary 2.45. *The map $\ell : \mathbb{Z}[LE(M)] \to L_\Gamma$ is injective.*

We now go back to the more general situation.

Lemma 2.46. *Let L be a Lie algebra, and suppose that $X = \{x_1, \ldots, x_n\}$ is a subset of L that satisfies*

$$[x_i, x_j] = 0 \text{ when } [v_i, v_j] = 1.$$

Suppose that a is a Lyndon element of M, and $v_i \in V$ such that $[\mathbf{v_i}, a] = 0$ in U. If ϕ_X is defined as in Definition 2.41, then

$$[\phi_X(a), \phi_X(\mathbf{v_i})] = 0.$$

Proof. We induct on the length of a. If $a = \mathbf{v_j}$ for some j then $[v_i, v_j] = 1$. Therefore $[\phi_X(a), \phi_X(v_i)] = [x_j, x_i] = 0$. If $|a| > 1$ then $S(a) = (x, y)$ for some $x, y \in LE(M)$ such that $[x, \mathbf{v_j}] = [y, \mathbf{v_j}] = 0$. Therefore by induction $[\phi_X(\mathbf{v_i}), \phi_X(x)] = [\phi_X(y), \phi_X(\mathbf{v_i})] = 0$, and by the Jacobi identity in L:

$$\begin{aligned}
[\phi_X(a), \phi_X(\mathbf{v_i})] &= [[\phi_X(x), \phi_X(y)], \phi_X(\mathbf{v_i})] \\
&= -[[\phi_X(\mathbf{v_i}), \phi_X(x)], \phi_X(y)] - [[\phi_X(y), \phi_X(\mathbf{v_i})], \phi_X(x)] \\
&= -[0, \phi_X(y)] - [0, \phi_X(x)] \\
&= 0.
\end{aligned}$$

□

What follows is the main technical theorem of this section, which will allow us to extend the $\mathbb{Z}$–module homomorphism ϕ_X to something that behaves well with respect to brackets also.

Proposition 2.47. *Let L be a Lie algebra, and suppose that $X = \{x_1, \ldots, x_n\}$ is a subset of L that satisfies*

$$[x_i, x_j] = 0 \text{ if } [v_i, v_j] = 1.$$

Let ϕ_X be the homomorphism defined in Definition 2.41. Let $a, b \in LE(M)$ be such that $a < b$. Then there exists a subset $I_{a,b} \subset LE(M)$ and a set of integers $\{\alpha_c\}_{c \in I_{a,b}}$ indexed by $I_{a,b}$ such that

$$[\phi_X(a), \phi_X(b)] = \sum_{c \in I_{a,b}} \alpha_c \phi_X(c).$$

Furthermore, each $c \in I_{a,b}$ satisfies the following:

(B1) $c < b$,

(B2) $std(c) \geq std(a)std(b)$,

(B3) $\|c\| = \|ab\|$,

and the sets $I_{a,b}$ and $\{\alpha_c\}_{c \in I_{a,b}}$ are independent of L and X.

Proof. The first step is to define an order $\prec$ on the set of pairs $(a, b) \in LE(M) \times LE(M)$ satisfying $a < b$. We say $(a, b) \prec (a', b')$ if

- $|ab| < |a'b'|$, or
- $|ab| = |a'b'|$ and $std(a)std(b) > std(a')std(b')$, or
- $std(a)std(b) = std(a')std(b')$ and $b < b'$.

Note that the second criterion is possibly the reverse of what one might expect. We shall prove Proposition 2.47 by using induction on the order given by $\prec$. We drop the subscript of ϕ_X for the remainder of this proof. The base case is when $(a, b) = (v_{n-1}, v_n)$ and is trivial. The inductive step splits into two cases.

Case 1. $init(b) \in \zeta(a)$.

If $|a| = 1$, then Theorem 2.39 tells us $S(ab) = (a, b)$, and $[\phi(a), \phi(b)] = \phi(ab)$ by definition. Also, $ab < b$ by part 2 of Theorem 2.36, and $std(ab) \geq std(a)std(b)$.

If $|a| > 1$, let $S(a) = (x, y)$. This now splits into two subcases.

Subcase 1. $y \geq b$. By Theorem 2.39, we have $S(ab) = (a, b)$, and we are in exactly the same situation as case 1.

Subcase 2. $y < b$ We use the Jacobi identity in L:

$$\begin{aligned}
[\phi(a), \phi(b)] &= [[\phi(x), \phi(y)], \phi(b)] \\
&= -[[\phi(b), \phi(x)], \phi(y)] - [[\phi(y), \phi(b)], \phi(x)] \\
&= [[\phi(x), \phi(b)], \phi(y)] + [\phi(x), [\phi(y), \phi(b)]]
\end{aligned}$$

We look at the two parts of this sum separately.

The $[[\phi(x), \phi(b)], \phi(y)]$ part:

Note that $x < a < b$, and $|xb| < |ab|$, so we have $(x, b) \prec (a, b)$. Therefore by induction there exists a decomposition:

$$[\phi(x), \phi(b)] = \sum_{c \in I_{x,b}} \alpha_c \phi(c)$$

with each c satisfying (B1)–(B3) with respect to (x, b). Then for each c, if $y < c$ then

$$\begin{aligned}
std(y)std(c) &\geq std(y)std(x)std(b) && \text{by (B3)} \\
&> std(x)std(y)std(b) && \text{by (2.10)} \\
&= std(a)std(b),
\end{aligned}$$

so that $(y, c) \prec (a, b)$. If $y = c$ then $[\phi(y), \phi(c)] = 0$. If $c < y$ then as $std(c) \geq std(x)std(b)$ and $std(y) < std(b)$ we have:

$$\begin{aligned}
std(c)std(y) &\geq std(x)std(b)std(y) \\
&> std(x)std(y)std(b) \\
&= std(a)std(b),
\end{aligned}$$

so that $(c, y) \prec (a, b)$. In any case, by induction there exists a decomposition:

$$[\phi(c), \phi(y)] = \sum_{d \in I_{c,y}} \beta_d \phi(d)$$

with each d satisfying (B1)–(B3) with respect to either (y, c) or (c, y). As the c here satisfies (B1)–(B3) with respect to (x, b) one can check that each d also satisfies (B1)–(B3) with respect to (a, b) and we have the required decomposition:

$$[[\phi(x), \phi(b)], \phi(y)] = \sum_{c \in I_{x,b}} \sum_{d \in I_{c,y}} \alpha_c \beta_d \phi(d).$$

The $[\phi(x), [\phi(y), \phi(b)]]$ part:

Since $y < b$ and $|yb| < |ab|$ there exists a decomposition $[\phi(y), \phi(b)] = \sum_{c \in I_{y,b}} \alpha_c c$ with each c satisfying (B1)–(B3) with respect to (y, b). Also for each c we have

$$\begin{aligned} std(c) &\geq std(y)std(b) \\ &\geq std(y) \\ &> std(x), \end{aligned}$$

so that $x < c$ and

$$\begin{aligned} std(x)std(c) &\geq std(x)std(y)std(b) \\ &= std(a)std(b). \end{aligned}$$

Hence $(x, c) \prec (a, b)$, and by induction we have the decomposition

$$[\phi(x), \phi(c)] = \sum_{d \in I_{x,c}} \beta_d \phi(d)$$

with each d satisfying (B1)–(B3) with respect to (x, c). As $c < b$ and $std(x)std(c) \geq std(a)std(b)$ each d also satisfies (B1)–(B3) with respect to (a, b). This gives our required decomposition

$$[\phi(x), [\phi(y), \phi(b)]] = \sum_{c \in I_{y,b}} \sum_{d \in I_{x,c}} \alpha_c \beta_d \phi(d)$$

Adding the above two parts gives the required decomposition of $[\phi(a), \phi(b)]$, and finishes the inductive step in this first case.

Case 2. $init(b) \notin \zeta(a)$.

If $|b| = 1$ then $[\phi(a), \phi(b)] = 0$ by Lemma 2.46, and we are done. If $|b| > 1$, then we write $S(b) = (x, y)$. By the Jacobi identity in L:

$$\begin{aligned}[\phi(a), \phi(b)] &= [\phi(a), [\phi(x), \phi(y)]] \\ &= -[\phi(x), [\phi(a), \phi(y)]] - [\phi(y), [\phi(x), \phi(a)]] \\ &= [[\phi(a), \phi(y)], \phi(x)] - [[\phi(a), \phi(x)], \phi(y)].\end{aligned}$$

Again we look at the two separate parts in this sum. First, $[[\phi(a), \phi(y)], \phi(x)]$. As $(a, y) \prec (a, b)$ by induction there exists a decomposition

$$[\phi(a), \phi(y)] = \sum_{c \in I_{a,y}} \alpha_c \phi(c),$$

with each c satisfying (B1)–(B3) with respect to (a, y). We would like to show that $c < x$ and $(c, x) \prec (a, b)$. Note that the smallest letter (with respect to the ordering $v_1 < v_2 < \cdots < v_n$) of any Lyndon word must be its initial letter, otherwise there would be a conjugate of that word that is smaller with respect to the ordering of M. Let $inf(g)$ denote the smallest letter in $supp(g)$ for any $g \in M$. As $\|c\| = \|ay\|$, we have:

$$init(c) = inf(c) = inf(ay) \leq inf(a) = init(a) < init(b) = init(x).$$

The strict inequality holds in the above as $a < b$ and $init(a) \neq init(b)$ because $init(b) \notin \zeta(a)$. Hence $c < x$, and

$$\begin{aligned}std(c)std(x) &\geq std(a)std(y)std(x) \\ &> std(a)std(x)std(y) \\ &= std(a)std(b)\end{aligned}$$

Therefore $(c, x) \prec (a, b)$, and there is a decomposition

$$[\phi(c), \phi(x)] = \sum_{d \in I_{c,x}} \beta_d \phi(d),$$

with each d satisfying the required (B1)–(B3) with respect to (c, x). Once again it is not hard to check that d also satisfies (B1)–(B3) with respect to (a, b). For $[[\phi(a), \phi(x)], \phi(y)]$ the same methods apply as before and we will spare the reader any further details.

This completes the induction proof. The only part we have not covered is the fact that the sets $I_{a,b}$ and $\{\alpha_c\}_{c \in I_{a,b}}$ are independent of X and L, however this is clear as we did not need use our choice of L or X at any point in the proof. $\square$

Proposition 2.47 implies that that the image of ℓ in L_Γ is closed under the bracket operation, so is a subalgebra of L_Γ. As L_Γ is the smallest subalgebra of $\mathcal{L}(U)$ containing $\{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$, and this set is in the image of ℓ, this means that ℓ is surjective. We have shown in Corollary 2.45 that ℓ is also injective.

Corollary 2.48. *The map $\ell : \mathbb{Z}[LE(M)] \to L_\Gamma$ is bijective.*

For our toils, we can now show that L_Γ satisfies the following universal property:

Theorem 2.49. *Let Γ be a graph with vertices $v_1, \ldots, v_n$. Let L be a Lie algebra, and suppose that $X = \{x_1, \ldots, x_n\}$ is a subset of L that satisfies:*

$$[x_i, x_j] = 0 \text{ when } v_i \text{ and } v_j \text{ are connected by an edge in } \Gamma.$$

Then there is a unique algebra homomorphism $\psi_X : L_\Gamma \to L$ such that

$$\psi_X(\mathbf{v_i}) = x_i \text{ for } 1 \leq i \leq r.$$

Proof. As L_Γ is generated by V, if such a map exists then it is unique. Let $\psi_X = \phi_X \ell^{-1}$. As ψ_X is a $\mathbb{Z}$–module morphism, we only need to check the bracket operation on the basis $\ell(LE(M))$ of L_Γ. Let $a, b \in LE(M)$ and without loss of generality suppose that $a < b$. By Proposition 2.47 there exists $I \subset LE(M)$ and a set of integers $\{\alpha_c\}_{c \in I}$ such that

$$\begin{aligned} [\ell(a), \ell(b)] &= \sum_{c \in I} \alpha_c \ell(c) \\ \text{and } [\phi_X(a), \phi_X(b)] &= \sum_{c \in I} \alpha_c \phi_X(c). \end{aligned}$$

Therefore

$$\begin{aligned} \psi_X([\ell(a), \ell(b)]) &= \psi_X(\sum_{c \in I} \alpha_c \ell(c)) \\ &= \sum_{c \in I} \alpha_c \psi_X \ell(c) \\ &= \sum_{c \in I} \alpha_c \phi_X(c) \\ &= [\phi_X(a), \phi_X(b)] \\ &= [\psi_X(\ell(a)), \psi_X(\ell(b))]. \end{aligned}$$

□

2.6 An isomorphism between L_Γ and the LCS algebra of A_Γ

The algebra L_Γ inherits a grading from $\mathcal{L}(U)$ by letting $L_{\Gamma,i} = L_\Gamma \cap \mathcal{L}(U)_i$. We note that

$$L_{\Gamma,i} = \langle \ell(a) : a \in LE(M), |a| = i \rangle.$$

Previously we defined $\mathcal{C}$ and $\mathcal{D}$ to be the linear filtrations of A_Γ given by the lower central series, and the central series $\{D_i\}$ given in section 2.4 respectively.

Lemma 2.50. *Let $X = \{v_i\gamma_1(A_\Gamma) : 1 \leq i \leq n\} \subset L_\mathcal{C}$. The algebra homomorphism $\psi_X : L_\Gamma \to L_\mathcal{C}$ given by Theorem 2.49 respects the gradings of L_Γ and $L_\mathcal{C}$.*

Proof. We show that $\psi_X(L_{\Gamma,k}) \subset L_{\mathcal{C},k}$ by induction on k. As $\psi_X(\mathbf{v_i}) = v_i\gamma_1(A_\Gamma)$, and $L_{\Gamma,1}$ is spanned by $\{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$, the case $k = 1$ holds. For the inductive step, pick $a \in LE(M)$ such that $|a| = k$. Let $S(a) = (b, c)$ be the standard decomposition of a, with $|b| = i$, $|c| = j$, and $i + j = k$. Then by induction $\psi_X(\ell(b)) \in L_{\mathcal{C},i}$ and $\psi_X(c) \in L_{\mathcal{C},j}$, hence

$$\psi_X(\ell(a)) = [\psi_X(\ell(b)), \psi_X(\ell(c))] \in L_{\mathcal{C},i+j} = L_{\mathcal{C},k}. \qquad \square$$

By Proposition 2.20 we know that $\gamma_k(A_\Gamma) \subset D_k$ for all k. Hence by Proposition 2.8 the identity map $A_\Gamma \to A_\Gamma$ induces a graded algebra homomorphism $\alpha : L_\mathcal{C} \to L_\mathcal{D}$.

Lemma 2.51. *The mapping $gD_{k+1} \mapsto \mu(g)_k$ induces a graded algebra homomorphism $\beta : L_\mathcal{D} \to \mathcal{L}(U)$.*

Proof. The group D_{k+1} is the kernel of the homomorphism $D_k \to U_k$ given by $g \mapsto \mu(g)_k$. Therefore the induced map $\beta : L_\mathcal{D} \to \mathcal{L}(U)$ is well-defined. As $\mu(g)_k \in \mathcal{L}(U)_k$, this map also respects gradings. The fact that β is a homomorphism is implied by parts (1), (5) and (6) of Lemma 2.19. $\square$

We now have a chain of graded algebra homomorphisms

$$L_\Gamma \xrightarrow{\psi_X} L_\mathcal{C} \xrightarrow{\alpha} L_\mathcal{D} \xrightarrow{\beta} \mathcal{L}(U),$$

which allows us to prove the main theorem of this chapter.

Theorem 2.52. *L_Γ, $L_\mathcal{C}$, and $L_\mathcal{D}$ are isomorphic as graded Lie algebras. Furthermore, the central filtrations $\mathcal{C}$ and $\mathcal{D}$ are equal, so that $\gamma_k(A_\Gamma) = D_k$ for all $k \geq 1$.*

Proof. We start by calculating the image of $\{\mathbf{v_1}, \ldots, \mathbf{v_n}\}$ under $\beta\alpha\psi_X$. We have

$$\begin{aligned}\beta\alpha\psi_X(\mathbf{v_i}) &= \beta\alpha(v_i\gamma_1(A_\Gamma)) \\ &= \beta(v_i D_1) \\ &= \mu(v_i)_1 \\ &= \mathbf{v_i}.\end{aligned}$$

Therefore the image of $\beta\alpha\psi_X$ is L_Γ, and as $\beta\alpha\psi_X$ takes the generators to themselves, it is the identity map on L_Γ. In particular, ψ_X must be injective. By Proposition 2.7, the algebra $L_\mathcal{C}$ is generated by the set X, hence ψ_X is also surjective, and is an isomorphism. We now know that $L_\mathcal{C}$ and L_Γ are isomorphic as graded Lie algebras. Then $\beta\alpha$ maps $L_\mathcal{C}$ isomorphically onto L_Γ, so the map α is also injective. Looking at each graded piece, each homomorphism

$$\gamma_k(A_\Gamma)/\gamma_{k+1}(A_\Gamma) \xrightarrow{\alpha_k} D_k/D_{k+1}$$

is injective. We shall use this to show that $\gamma_k(A_\Gamma) = D_k$ by induction on k, and this will complete the proof of the main theorem. Note that $D_1 = \gamma_1(A_\Gamma)$ by definition. Suppose that $\gamma_k(A_\Gamma) = D_k$. Then α_k is also surjective, so is an isomorphism. If $g \in D_k = \gamma_k(A_\Gamma)$, then

$$\begin{aligned} g \in D_{k+1} &\iff gD_{k+1} = 1 && \text{in } D_k/D_{k+1} \\ &\iff \alpha_k^{-1}(gD_{k+1}) = 1 && \text{in } \gamma_k(A_\Gamma)/\gamma_{k+1}(A_\Gamma) \\ &\iff g\gamma_{k+1}(A_\Gamma) = 1 && \text{in } \gamma_k(A_\Gamma)/\gamma_{k+1}(A_\Gamma) \\ &\iff g \in \gamma_{k+1}(A_\Gamma). \end{aligned}$$

Hence $\gamma_{k+1}(A_\Gamma) = D_{k+1}$. □

We conclude with an important consequence of Theorem 2.52 and Proposition 2.20:

Theorem 2.53. *If $k \in \mathbb{N}$, then $\gamma_k(A_\Gamma)/\gamma_{k+1}(A_\Gamma)$ is free-abelian, and $A_\Gamma/\gamma_k(A_\Gamma)$ is torsion-free nilpotent.*

Example 2.54. Let Γ be the small example graph given in Figure 1. We have already worked out the bracketing of Lyndon elements of length 3 in example 2.40. The isomorphism given in Theorem 2.52 tells us that $\gamma_3(A_\Gamma)/\gamma_4(A_\Gamma)$ is freely generated by $[v_1,[v_1,v_2]]\gamma_4(A_\Gamma)$, $[v_1,[v_1,v_3]]\gamma_4(A_\Gamma)$, $[[v_1,v_2],v_2]\gamma_4(A_\Gamma)$, $[[v_1,v_3],v_2]\gamma_4(A_\Gamma)$, and $[[v_1,v_3],v_3]\gamma_4(A_\Gamma)$.

2.7 More information on the structure of the LCS algebra

In Chapter 4 we shall use the structure of $L_{\mathcal{C}}$ to study $\mathrm{IA}(A_\Gamma)$. The relevant proofs are contained in this section.

Proposition 2.55. *The free abelian group $L_{\mathcal{C},2} = \gamma_2(A_\Gamma)/\gamma_3(A_\Gamma)$ has a basis given by the set $S = \{[v_i, v_j]\gamma_3(A_\Gamma) : i < j, [v_i, v_j] \neq 0\}$.*

Proof. By Theorem 2.52 we know that $L_{\mathcal{C},2}$ is freely generated by the image of the elements of A_Γ obtained by taking the Lyndon elements of M of length 2 with their unique bracketing. The Lyndon elements of M of length 2 are exactly the elements of the form $v_i v_j$ with $i < j$ and $[v_i, v_j] \neq 1$ and have a unique bracketing $[v_i, v_j]$. □

We shall use Proposition 2.55 in Chapter 4 to describe the abelianisation of $\mathrm{IA}(A_\Gamma)$. We also want to understand the centre of $L_{\mathcal{C}}$ and some associated structures. We use $Z(_)$ to denote the centre of a Lie algebra. The following theorem is similar to Exercise 3.3 in Chapter 2 of [16].

Theorem 2.56. *Suppose that $Z(A_\Gamma) = 1$. Let p and c be positive integers. Then $Z(L_{\mathcal{C}}) = Z((\mathbb{Z}/p\mathbb{Z}) \otimes_{\mathbb{Z}} L_{\mathcal{C}}) = 0$ and $Z(L_{\mathcal{C}}/\oplus_{i>c} L_{\mathcal{C},i})$ is the image of $L_{\mathcal{C},c}$ under the quotient map $L_{\mathcal{C}} \to L_{\mathcal{C}}/\oplus_{i>c} L_{\mathcal{C},i}$.*

Proof. As $L_{\mathcal{C}}$ and L_Γ are isomorphic as graded algebras, we may work in L_Γ. Let $f = \sum_{a_i \in U_i} a_i$ be a nontrivial element of L_Γ. Let j be the smallest integer such that $a_j \neq 0$, and suppose that $a_j = \sum_{b \in I} \alpha_b b$, where $I \subset M_j$. Let x be the minimal element of I with nonzero coefficient. As $f \in L_\Gamma$, we know that a_j is a sum of bracketed Lyndon elements, and therefore by Proposition 2.44 we know that x is a Lyndon element. As $Z(A_\Gamma) \neq 0$, there exists a vertex v such that $[v, init(x)] \neq 1$, so that $v \in \zeta(x)$ and $init(x) \in \zeta(v)$. Either $v < x$ or $x < v$. Without loss of generality, assume the former. Then by Lemma 2.43 vx is the minimal element in the decomposition of $[v, a_j]$ with coefficient α_x. As the bracket operation respects degrees, vx has coefficient α_x in the decomposition of $[v, f]$. Hence $[v, f] \neq 0$ and $Z(L_\Gamma) = 0$. If f is nonzero in $Z((\mathbb{Z}/p\mathbb{Z}) \otimes_{\mathbb{Z}} L_\Gamma)$ we may assume that p does not divide α_x so that $[v, f]$ is nonzero in $(\mathbb{Z}/p\mathbb{Z}) \otimes_{\mathbb{Z}} L_\Gamma$. Hence $Z(\mathbb{Z}/p\mathbb{Z} \otimes_{\mathbb{Z}} L_\Gamma) = 0$. Finally, if f nonzero in $L_\Gamma/\oplus_{k>c} L_{\Gamma,k}$ and not in the image of $L_{\Gamma,c}$ then $j < c$ and $|vx| \leq c$. Hence f is not in $Z(L_\Gamma/\oplus_{k>c} L_{\Gamma,k})$ and therefore $Z(L_\Gamma/\oplus_{k>c} L_{\Gamma,k})$ is the image of $L_{\Gamma,c}$ under the quotient map $L_\Gamma \to L_\Gamma/\oplus_{k>c} L_{\Gamma,k}$. □

Chapter 3

Foundations for studying $Out(A_\Gamma)$

In this chapter we shall describe some background material on $Out(A_\Gamma)$, where A_Γ is the right-angled Artin group determined by the graph Γ. The key tools are the restriction and projection homomorphisms developed by Charney and Vogtmann in [20, 21]. These allow one to study the automorphism group of a right-angled Artin group in terms of automorphism groups of RAAGs with strictly smaller defining graphs. Our contribution to this machinery is to slightly extend the standard generating set of $Out(A_\Gamma)$ (described below) so that this generating set is preserved under the restriction and projection homomorphisms. This allows one to study all subgroups of $Out(A_\Gamma)$ generated by subsets of the generating set. This is important as the restriction and projection homomorphisms are usually not surjective, and in our applications (and we suspect this will be useful in future work) it is important that we only study the image of the homomorphism, rather than the whole automorphism group that the image is contained in. We study how such subgroups act on the abelianisation of A_Γ. At the end of the chapter, we define a notion of rank, called SL–dimension, for a subgroup of $G < \mathrm{Out}(A_\Gamma)$. We show that if G is generated by a subset of our generating set, then the rank of G does not increase under restriction and projection homomorphisms.

3.1 A generating set of $\mathrm{Aut}(A_\Gamma)$

Given a vertex v, the *link of* v is the set of vertices of Γ adjacent to v. The star of v is the union of v and the link of v. We write $lk(v)$ for the link of v, and $st(v)$ for the star of v.

Laurence [52] proved a conjecture of Servatius [68] that $\mathrm{Aut}(A_\Gamma)$ has a finite generating set consisting of the following automorphisms:

Graph symmetries If a permutation of the vertices comes from a self-isomorphism of the graph, then this permutation induces an automorphism of A_Γ. These automorphisms form a finite subgroup of $\mathrm{Aut}(A_\Gamma)$ called $\mathrm{Sym}(A_\Gamma)$.

Inversions These are automorphisms that come from inverting one of the generators of A_Γ, so that:

$$s_i(v_k) = \begin{cases} v_i^{-1} & i = k \\ v_k & i \neq k. \end{cases}$$

Partial conjugations Suppose $[v_i, v_j] \neq 0$. Let Γ_{ij} be the connected component of $\Gamma - st(v_j)$ containing v_i. Then the partial conjugation K_{ij} conjugates every vertex of Γ_{ij} by v_j, and fixes the remaining vertices, so that:

$$K_{ij}(v_k) = \begin{cases} v_j v_k v_j^{-1} & v_k \in \Gamma_{ij} \\ v_k & v_k \notin \Gamma_{ij}. \end{cases}$$

Note that if $lk(v_i) \subset st(v_j)$ then $\Gamma_{ij} = \{v_i\}$, so in this case K_{ij} fixes every basis element except v_i.

Transvections If $lk(v_i) \subset st(v_j)$, then there is an automorphism ρ_{ij} which acts on the generators of A_Γ as follows:

$$\rho_{ij}(v_k) = \begin{cases} v_i v_j & i = k \\ v_k & i \neq k. \end{cases}$$

There are two important finite index normal subgroups of $\mathrm{Aut}(A_\Gamma)$ that we obtain from this classification. The first is the subgroup generated by inversions, partial conjugations, and transvections and is denoted $\mathrm{Aut}^0(A_\Gamma)$. The second is the smaller subgroup generated by only partial conjugations and transvections. Denote this group $\mathrm{SAut}^0(A_\Gamma)$. In some cases we will need to look at groups generated by (outer) automorphisms that conjugate more than one component of $\Gamma - st(v_j)$ by v_j.

Definition 3.1. Let T be a subset of $\Gamma - st(v_j)$ such that no two vertices of T lie in the same connected component of $\Gamma - st(v_j)$. An *extended partial conjugation* is defined to be an automorphism of the form $\prod_{t \in T} K_{tj}$.

We will abuse notation by describing the images of the above elements in $\mathrm{Out}(A_\Gamma)$ by the same names, so that the groups $\mathrm{Out}^0(A_\Gamma)$ and $\mathrm{SOut}^0(A_\Gamma)$ are defined in the same manner. If $\phi \in \mathrm{Aut}(A_\Gamma)$, we use $[\phi]$ to denote the equivalence class of ϕ in $\mathrm{Out}(A_\Gamma)$.

Definition 3.2. Let $\mathcal{S}_\Gamma$ be the enlarged generating set of $\mathrm{Out}(A_\Gamma)$ given by graph symmetries, inversions, extended partial conjugations, and transvections.

We shall be studying subgroups of $\mathrm{Out}(A_\Gamma)$ generated by subsets of $\mathcal{S}_\Gamma$, however some of these groups are not generated by subsets of the standard generating set. This is because under the restriction, exclusion and projection maps defined in Section 3.3, partial conjugations are not always mapped to partial conjugations, but are always mapped to extended partial conjugations. Throughout we will assume $\mathrm{Aut}(A_\Gamma)$ and $\mathrm{Out}(A_\Gamma)$ act on A_Γ on the left.

3.2 Ordering the vertices of Γ

We look at two methods of ordering the vertices of Γ. These orderings allow us to describe the action on $\mathrm{Aut}(A_\Gamma)$ and $\mathrm{Out}(A_\Gamma)$ on the abelianisation of A_Γ and, later on, define the aforementioned restriction and projection homomorphisms.

3.2.1 The standard order on $V(\Gamma)$

Extending the definition of the link and star of a vertex, given any full subgraph Γ' of Γ, the subgraph $lk(\Gamma')$ is defined to be the intersection of the links of the vertices of Γ', and we define $st(\Gamma') = \Gamma' \cup lk(\Gamma')$. Given any full subgraph Γ', the right-angled Artin group $A_{\Gamma'}$ injects into A_Γ, so can be viewed as a subgroup. As we shall only be interested in full subgraphs of Γ we shall often blur the distinction between a subset of the vertex set and the full subgraph of Γ spanned by this vertex set.

Definition 3.3 (Standard order)**.** We define the binary relation $\leq$ on $V(\Gamma)$ to be such that $u \leq v$ if and only if $lk(u) \subset st(v)$.

This was introduced in [47], where it was shown that the relation is transitive as well as reflexive, so defines a pre-order on the vertices. As the structure of this ordering is important in all of our work on $\mathrm{Out}(A_\Gamma)$ we will supply a proof. Let d_Γ be the edge metric on Γ. We do not suppose that Γ is connected, so allow distances to be infinite.

Proposition 3.4 (cf. [20], page 95)**.** *Let u, v, and w be distinct vertices of Γ such that $u \leq v \leq w$. Then $u \leq w$ and $\leq$ is a pre-order on $V(\Gamma)$. Furthermore, either:*

(i) u is an isolated vertex, i.e. $lk(u) = \emptyset$,

(ii) $d_\Gamma(u, v) = d_\Gamma(v, w) = d_\Gamma(w, u) = 1$, or

(iii) $d_\Gamma(u, v) = 2$.

If $[v]$ *is an equivalence class of the equivalence relation induced by this pre-order then* $A_{[v]}$ *is either free or free-abelian.*

Proof. If u is isolated, then $lk(u) = \emptyset$ and $u \leq w$. Suppose that u is not isolated. Then there exists a vertex $u' \in lk(u)$ and $u' \in st(v)$, so $d_\Gamma(u, v) \leq 2$.

Suppose that $d_\Gamma(u, v) = 1$. Then $u \in lk(v) \subset st(w)$, and therefore $d_\Gamma(u, w) = 1$. Also $w \in lk(u) \subset st(v)$, so $d_\Gamma(v, w) = 1$. If $u' \in lk(u)$ then either $u' \neq v$, so $u' \in lk(v) \subset st(w)$, or $u' = v$, which also lies in $st(w)$. Hence $lk(u) \subset st(w)$ and $u \leq w$.

Suppose that $d_\Gamma(u, v) = 2$. Then $v \notin lk(u)$, so $lk(u) \subset lk(v) \subset st(w)$ and $v \leq w$.

We have shown that $\leq$ is transitive. As $\leq$ is also reflexive, it is a pre-order, and so induces an equivalence relation by saying that $u \sim v$ if and only if $u \leq v$ and $v \leq u$. Let $[v]$ be an equivalence class of $\sim$. If no pair of vertices in $[v]$ are adjacent then $A_{[v]}$ is a free group. If any two vertices in $[v]$ are adjacent, it follows from point (ii) that all vertices in $[v]$ are adjacent, and $A_{[v]}$ is a free abelian group. □

We say that the equivalence class $[v]$ is *abelian* if $A_{[v]}$ is a free-abelian group, and $[v]$ is *non-abelian* if $A_{[v]}$ is a non-abelian free group. The pre-order descends to a partial order of the equivalence classes. We say that $[v]$ is *maximal* if it is maximal with respect to this ordering. Suppose that there are r equivalence classes of vertices in Γ. We may choose an enumeration of the vertices so that there exists $1 = m_1 < m_2 < \ldots < m_r < n$ such that the equivalence classes are the sets $\{v_1 = v_{m_1}, \ldots, v_{m_2-1}\}, \ldots, \{v_{m_r}, \ldots, v_n\}$. With further rearrangement we may assume that $v_{m_i} \leq v_{m_j}$ only if $i \leq j$. We formally define $m_{r+1} = n + 1$ so that for all i, the equivalence class of $[v_{m_i}]$ contains $m_{i+1} - m_i$ vertices.

3.2.2 G–ordering vertices

Given a subgroup $G \leq \mathrm{Out}(A_\Gamma)$, we may define the relation $\leq_G$ on the vertices of Γ by letting $v_i \leq_G v_j$ if either $i = j$ or the element $[\rho_{ij}]$ lies in G. Note that this is a subset of the previous relation $\leq$ defined on the vertices. Also, $\leq_G$ is reflexive by definition and transitive as $\rho_{il} = \rho_{jl}^{-1}\rho_{ij}^{-1}\rho_{jl}\rho_{ij}$, so $\leq_G$ is a pre-order, induces an equivalence relation $\sim_G$ on the vertices, and induces a partial ordering of the equivalence classes of $\sim_G$. Let $[v_i]_G$ be the equivalence class of the vertex v_i. Each equivalence class $[v_i]_G$ is a subset of the equivalence class $[v_i]$. In particular the subgroup $A_{[v_i]_G}$ is either free abelian or free and non-abelian, so $[v_i]_G$ may also be described as *abelian*

or *non-abelian*. Suppose that there are $r' \geq r$ equivalence classes of vertices in $\sim_G$. We may further refine the enumeration of the vertices given previously so that there exists $1 = l_1 < l_2 < \ldots l_{r'} < n$ such that the equivalence classes of $\sim_G$ are the sets $\{v_1 = v_{l_1}, \ldots, v_{l_2-1}\}, \ldots, \{v_{l_{r'}}, \ldots, v_n\}$, and $v_{l_i} \leq_G v_{l_j}$ only if $i \leq j$. Define $l_{r'+1} = n+1$ so that for all i, the equivalence class of $[v_{l_i}]_G$ contains $l_{i+1} - l_i$ vertices.

3.2.3 Relation of G–orderings to the action on $H_1(A_\Gamma)$

Let G be a subgroup of $\mathrm{Out}(A_\Gamma)$ generated by a subset $T \subset \mathcal{S}_\Gamma$. We shall assume that T is maximal, so that $T = \mathcal{S}_\Gamma \cap G$. We are going to study the action of G on $H_1(A_\Gamma)$, and to simplify matters we want to avoid graph symmetries and inversions.

Definition 3.5. Let G^0 be the subgroup of G generated by the extended partial conjugations, inversions, and transvections in T and let SG^0 be the subgroup of G generated solely by the extended partial conjugations and transvections in T.

Proposition 3.6. *Let G be a subgroup of $\mathrm{Out}(A_\Gamma)$ generated by a subset T of $\mathcal{S}_\Gamma$. Then G^0 and SG^0 are finite index normal subgroups of G.*

Proof. Suppose α is a graph symmetry that moves the vertices according to the permutation σ. We find that $\alpha K_{ij}\alpha^{-1} = K_{\sigma(i)\sigma(j)}$, $\alpha\rho_{ij}\alpha^{-1} = \rho_{\sigma(i)\sigma(j)}$ and $\alpha s_i \alpha^{-1} = \alpha_{\sigma(i)}$. We may assume that T is maximal, so that if $[\phi]$ and $[\alpha]$ belong to T, then so does $[\alpha\phi\alpha^{-1}]$. Therefore if W is a word in $T \cup T^{-1}$ one may shuffle graph symmetries along so that they all occur at the beginning of a word W' representing the same element as W. As the group of graph symmetries is finite, this shows that G^0 is finite index in G, and the above computations verify that G^0 is normal in G. Similarly, with inversions one verifies that:

$$\rho_{kl}s_i = \begin{cases} s_i\rho_{kl} & i \neq k, l \\ s_i\rho_{ki}^{-1} & i = l \\ s_iK_{il}^{-1}\rho_{il}^{-1} & i = k, [v_i, v_l] \neq 0 \\ s_i\rho_{il}^{-1} & i = k, [v_i, v_l] = 0 \end{cases} \text{and} \quad K_{kl}s_i = \begin{cases} s_iK_{kl} & i \neq l \\ s_iK_{kl}^{-1} & i = l \end{cases}.$$

These show that SG^0 is normal in G^0, and we may write any element of G^0 in the form $[s_1^{\epsilon_1} \ldots s_n^{\epsilon_n}\phi']$, where $\epsilon_i \in \{0, 1\}$ and ϕ' is a product of extended partial conjugations and transvections. Therefore SG^0 is of index at most 2^n in G^0. □

Now let us look at the generators of $\mathrm{Aut}(A_\Gamma)$ (respectively $\mathrm{Out}(A_\Gamma)$) under the map $\Phi : \mathrm{Aut}(A_\Gamma) \to \mathrm{GL}_n(\mathbb{Z})$ (respectively $\overline{\Phi} : \mathrm{Out}(A_\Gamma) \to \mathrm{GL}_n(\mathbb{Z})$). If α is a graph symmetry, then $\Phi(\alpha)$ is the appropriate permutation matrix corresponding to the

permutation α induces on the vertices. For a partial conjugation K_{ij} we see that $\Phi(K_{ij})$ is the identity matrix I; Φ sends the inversion s_i to the matrix S_i which has 1 everywhere on the diagonal except for -1 at the $(i,i)^{th}$ position, and zeroes everywhere else, and Φ sends the transvection ρ_{ij} to the matrix $T_{ji} = I + E_{ji}$, where E_{ji} is the elementary matrix with 1 in the $(i,j)^{th}$ position, and zeroes everywhere else. The swapping between ρ_{ij} and T_{ji} may seem a little unnatural, but occurs as a choice of having $\mathrm{Aut}(A_\Gamma)$ act on the left. It follows that the image of $\mathrm{SAut}^0(A_\Gamma)$ under Φ is the subgroup of $\mathrm{GL}_n(\mathbb{Z})$ generated by matrices of the form T_{ji}, where $v_i \leq v_j$. If we order the vertices as in Section 3.2.1, then $v_i \leq v_j$ only if either v_i and v_j are in the same equivalence class of vertices, or $i \leq j$. It follows that a matrix in the image of $\overline{\Phi}|_{\mathrm{Out}^0(A_\Gamma)}$ has a block decomposition of the form:

$$M = \begin{pmatrix} M_1 & 0 & \dots & 0 \\ * & M_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ * & * & \dots & M_r \end{pmatrix},$$

where the $*$ in the $(i,j)^{th}$ entry in the block decomposition may be nonzero if $[v_{m_j}] \leq [v_{m_i}]$, but zero otherwise.

Similarly, given a subgroup $G \leq \mathrm{Out}(A_\Gamma)$ generated by a subset of $\mathcal{S}_\Gamma$, if we order the vertices by the method given in Section 3.2.2, a matrix in the image of $\overline{\Phi}|_{G^0}$ has a block decomposition of the form:

$$M = \begin{pmatrix} N_1 & 0 & \dots & 0 \\ * & N_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ * & * & \dots & N_{r'} \end{pmatrix},$$

where the $*$ in the $(i,j)^{th}$ entry in the block decomposition may be nonzero if $[v_{l_j}]_G \leq [v_{l_i}]_G$, but is zero otherwise.

3.3 Restriction, exclusion, and projection homomorphisms.

Suppose that Γ' is a full subgraph of Γ, so that $A_{\Gamma'}$ can be viewed as a subgroup of A_Γ in the natural way. Suppose that the conjugacy class of $A_{\Gamma'}$ is preserved by some $G < \mathrm{Out}(A_\Gamma)$. Then there is a natural *restriction map* $R_{\Gamma'} \colon G \to \mathrm{Out}(A_{\Gamma'})$ obtained by taking a representative of an element $[\phi] \in G$ that preserves $A_{\Gamma'}$. This is well defined because the normaliser of $A_{\Gamma'}$ in A_Γ is generated by $A_{\Gamma'}$ and the

centraliser $C_{A_\Gamma}(A_{\Gamma'}) = A_{lk(\Gamma')}$. (One can see this via the normal form theorem for elements of RAAGs: if $v \in \Gamma'$ and $gvg^{-1} \in A_{\Gamma'}$, then $supp(g)$ must be contained in $st(\Gamma)$.) Similarly, if the normal subgroup of A_Γ generated by $A_{\Gamma'}$ is preserved by some $G < \mathrm{Out}(A_\Gamma)$, then there is a natural *exclusion map* $E_{\Gamma'}\colon G \to \mathrm{Out}(A_\Gamma/\langle\langle A_{\Gamma'}\rangle\rangle) \cong \mathrm{Out}(A_{\Gamma - \Gamma'})$. There are two key examples:

Example 3.7. If Γ is connected and v is a maximal vertex then the conjugacy classes of $A_{[v]}$ and $A_{st[v]}$ are preserved by $\mathrm{Out}^0(A_\Gamma)$ ([20], Proposition 3.2). Therefore there is a restriction map

$$R_v : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}^0(A_{st[v]}),$$

an exclusion map

$$E_v : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}^0(A_{\Gamma - [v]}),$$

and a projection map

$$P_v : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}^0(A_{lk[v]})$$

obtained by taking a representative of an element that preserves both $A_{st[v]}$ and the normal closure of $A_{[v]}$ and taking the induced automorphism on $A_{st[v]}/\langle\langle A_{[v]}\rangle\rangle \cong A_{lk[v]}$. We can take the direct sum of these projection maps over all maximal equivalence classes $[v]$ to obtain the *amalgamated projection homomorphism*:

$$P : \mathrm{Out}^0(A_\Gamma) \to \bigoplus_{[v] \text{ maximal}} \mathrm{Out}^0(A_{lk[v]})$$

Example 3.8. If Γ is not connected, suppose that Γ has m isolated vertices and $\{\Gamma_i\}_{i=1}^k$ is the set of connected components of Γ containing at least two vertices. Then $A_\Gamma \cong F_m * *_{i=1}^k A_{\Gamma_i}$. By looking at the action of the generating set of $\mathrm{Out}^0(A_\Gamma)$, we find that the conjugacy class of A_{Γ_i} is fixed by $\mathrm{Out}^0(A_\Gamma)$, therefore for each i we obtain a restriction map

$$R_i : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}^0(A_{\Gamma_i}),$$

and as the normal subgroup generated by $*_{i\in I} A_{\Gamma_i}$ is preserved by $\mathrm{Out}(A_\Gamma)$ there is an exclusion map

$$E : \mathrm{Out}(A_\Gamma) \to \mathrm{Out}(F_m).$$

Charney and Vogtmann have shown that when Γ is connected, the maps in Example 3.7 describe $\mathrm{Out}(A_\Gamma)$ almost completely. There are two cases: when the centre of A_Γ, which we write as $Z(A_\Gamma)$, is trivial, and when $Z(A_\Gamma)$ is nontrivial. In the first case, they show the following:

Theorem 3.9 ([20], Theorem 4.2)**.** *If* Γ *is connected and* $Z(A_\Gamma)$ *is trivial, then* $\ker P$ *is a finitely generated free abelian group.*

In [21], Theorem 3.9 is extended by giving an explicit generating set of $\ker P$. However we will not need this description in the work that follows. When $Z(A_\Gamma)$ is nontrivial there is a unique maximal abelian equivalence class $[v]$ consisting of the vertices in $Z(A_\Gamma)$, and we are in the following situation:

Proposition 3.10 ([20], Proposition 4.4)**.** *If* $Z(A_\Gamma) = A_{[v]}$ *is nontrivial, then*

$$\mathrm{Out}(A_\Gamma) \cong Tr \rtimes (GL(A_{[v]}) \times \mathrm{Out}(A_{lk[v]})),$$

where Tr is the free abelian group generated by the transvections $[\rho_{ij}]$ *such that* $v_i \in lk[v]$ *and* $v_j \in [v]$*. The map to* $GL(A_{[v]})$ *is given by the restriction map* R_v*, and the map to* $\mathrm{Out}(A_{lk[v]})$ *is given by the projection map* P_v*. The subgroup Tr is the kernel of the product map* $R_v \times P_v$*.*

In the above proposition we do not need to restrict R_v and P_v to $\mathrm{Out}^0(A_\Gamma)$, as every automorphism of A_Γ preserves $Z(A_\Gamma) = A_{[v]}$. When Γ is disconnected, the restriction and exclusion maps of Example 3.8 give us less information. As above, we may amalgamate the restriction maps R_i and the exclusion map E, however in this situation the kernel of the amalgamated map is much more complicated.

3.4 SL–dimension for subgroups of $\mathrm{Out}(A_\Gamma)$.

In the introduction we defined the SL–dimension of $\mathrm{Out}(A_\Gamma)$ to be the size of a largest abelian equivalence class of vertices in Γ. Unfortunately this definition can behave badly under projection, restriction, and exclusion homomorphisms. In particular, if v is a maximal vertex in a connected graph Γ, then it is not always true that $d_{SL}(\mathrm{Out}(A_{lk[v]})) \leq d_{SL}(\mathrm{Out}(A_\Gamma))$. To get round this problem, we will extend the definition of SL–dimension to arbitrary subgroups of $\mathrm{Out}(A_\Gamma)$, and show that if instead we look at the image of $\mathrm{Out}(A_\Gamma)$ under such homomorphisms, then the SL–dimension will not increase (for instance, it will always be the case that $d_{SL}(P_v(\mathrm{Out}(A_\Gamma))) \leq d_{SL}(\mathrm{Out}(A_\Gamma))$). Once again we fix a subgroup G of $\mathrm{Out}(A_\Gamma)$ generated by a subset $T \subset \mathcal{S}_\Gamma$. We look at the G–ordering on $V(\Gamma)$. If $[v]_G$ is an abelian equivalence class of vertices, then G contains a copy $\mathrm{SL}(A_{[v]_G})$ generated by the $[\rho_{ij}]$ with $v_i, v_j \in [v]_G$. This fact lends itself to the following definition:

Definition 3.11. For any subgroup $G \leq \mathrm{Out}(A_\Gamma)$, the SL–dimension of G, $d_{SL}(G)$, is defined to be the size of the largest abelian equivalence class under $\sim_G$.

Roughly speaking, $d_{SL}(G)$ is the largest integer such that G contains an obvious copy of $SL_{d_{SL}(G)}(\mathbb{Z})$. Note that $d_{SL}(\mathrm{Out}(A_\Gamma))$ is simply the size of the largest abelian equivalence class under the relation $\leq$ defined on the vertices, so this generalises our previous definition. As each abelian equivalence class of vertices is a clique in Γ, the SL–dimension of $\mathrm{Out}(A_\Gamma)$ is less than or equal to the dimension of A_Γ. We can now look at how G and its SL–dimension behave under restriction, exclusion, and projection maps.

Lemma 3.12. *Let G be a subgroup of $\mathrm{Out}(A_\Gamma)$ generated by a subset $T \subset \mathcal{S}_\Gamma$. Suppose that Γ' is a full subgraph of Γ and the conjugacy class of $A_{\Gamma'}$ in A_Γ is preserved by G. Then under the restriction map $R_{\Gamma'}$, the group $R_{\Gamma'}(G)$ is generated by a subset of $\mathcal{S}_{\Gamma'}$, and $d_{SL}(R_{\Gamma'}(G)) \leq d_{SL}(G)$.*

Proof. One first checks that for an element $[\phi] \in T$, either $R_{\Gamma'}([\phi])$ is trivial or $R_{\Gamma'}([\phi]) \in \mathcal{S}_{\Gamma'}$. This is obvious in the case of graph symmetries, inversions, and transvections. In the case of partial conjugations if v_j is not in Γ', or if $\Gamma_{ij} \cap \Gamma' = \emptyset$, then $R_{\Gamma'}([K_{ij}])$ is trivial. Otherwise, $\Gamma_{ij} \cap \Gamma'$ is a union of connected components of $\Gamma' - st(v_j)$, so that $R_{\Gamma'}([K_{ij}])$ is an extended partial conjugation of $A_{\Gamma'}$. This proves the first part of the lemma. To prove the second part of the lemma, we first give an alternate definition of $d_{SL}(G)$. Elements in the image of G^0 under $\overline{\Phi}$ are of the form:

$$M = \begin{pmatrix} N_1 & 0 & \dots & 0 \\ * & N_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ * & * & \dots & N_{r'} \end{pmatrix}, \tag{3.1}$$

where each N_i is an invertible matrix of size $l_{i+1} - l_i$. Each of the blocks is associated to either an abelian or non-abelian equivalence class in $\sim_G$, so $d_{SL}(G)$ is the size of the largest diagonal block in this decomposition associated to an abelian equivalence class. Each abelian equivalence class with at least 2 vertices in $\sim_{R_{\Gamma'}(G)}$ is the image of an abelian equivalence class of $\sim_G$, and the action of $R_{\Gamma'}(G)^0$ on $A^{ab}_{\Gamma'}$ is obtained by removing rows and columns from the decomposition given in Equation (3.1). Therefore, the largest diagonal block in the action of $R_{\Gamma'}(G)^0$ on $A^{ab}_{\Gamma'}$ associated to an abelian equivalence class will be of size less than or equal to $d_{SL}(G)$. Therefore $d_{SL}(R_{\Gamma'}(G)) \leq d_{SL}(G)$. □

The following lemma is shown in the same way:

Lemma 3.13. *Let G be a subgroup of $\mathrm{Out}(A_\Gamma)$ generated by a subset $T \subset \mathcal{S}_\Gamma$. Let Γ' be a full subgraph of Γ. Suppose that the normal subgroup generated by $A_{\Gamma'}$ in A_Γ*

is preserved by G. Then under the exclusion map $E_{\Gamma'}$, the group $E_{\Gamma'}(G)$ is generated by a subset of $\mathcal{S}_{\Gamma'}$, and $d_{SL}(E_{\Gamma'}(G)) \leq d_{SL}(G)$.

As projection maps are obtained by the concatenation of a restriction and an exclusion map, combining the previous two lemmas gives the following:

Proposition 3.14. *Let G be a subgroup of $\mathrm{Out}(A_\Gamma)$ generated by a subset $T \subset \mathcal{S}_\Gamma$. Suppose that Γ is connected and v is a maximal vertex of Γ. Under the projection homomorphism P_v of Example 3.7, the group $P_v(G^0)$ is generated by a subset of $\mathcal{S}_{lk[v]}$ and $d_{SL}(P_v(G^0)) \leq d_{SL}(G) = d_{SL}(G^0)$.*

3.5 An example

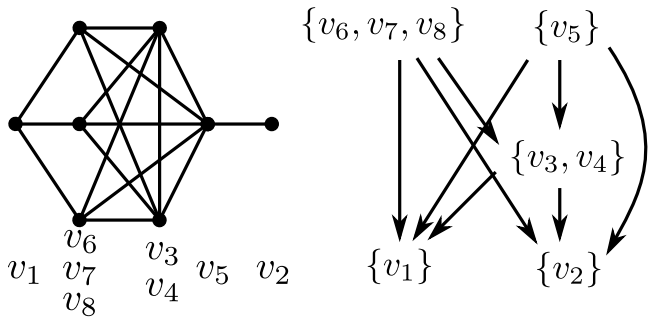

Figure 3.1: Our example graph Γ is pictured on the left. On the right is a diagram indicating the partial order on the equivalence classes, with the maximal equivalence classes at the top.

Let Γ be the graph given in Figure 3.1. The equivalence classes given by the standard order on the vertices are $\{v_1\}$, $\{v_2\}$, $\{v_3, v_4\}$, $\{v_5\}$ and $\{v_6, v_7, v_8\}$. All are abelian except for $\{v_6, v_7, v_8\}$, therefore $d_{SL}(\mathrm{Out}(A_\Gamma)) = 2$. The partial ordering is indicated by the diagram on the right in Figure 3.1. The maximal equivalence classes are $\{v_5\}$ and $\{v_6, v_7, v_8\}$. The link of this latter equivalence class is $\Gamma' = lk(\{v_6, v_7, v_8\})$, shown in Figure 3.2 below.

There is a projection map

$$P_{v_6} : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}(\Gamma').$$

Let $G = P_{v_6}(\mathrm{Out}^0(A_\Gamma))$. The transvections in G are exactly the images of transvections in $\mathrm{Out}^0(A_\Gamma)$, so are $\rho_{13}, \rho_{14}, \rho_{15}, \rho_{34}, \rho_{43}, \rho_{53}$ and ρ_{54}. The G–ordering on Γ' is then given by restriction of the standard order on Γ, so has equivalence classes $\{v_1\}$, $\{v_3, v_4\}$ and $\{v_5\}$. Note here that $d_{SL}(\mathrm{Out}(A_{\Gamma'})) = 3$, but $d_{SL}(G) = 2$.

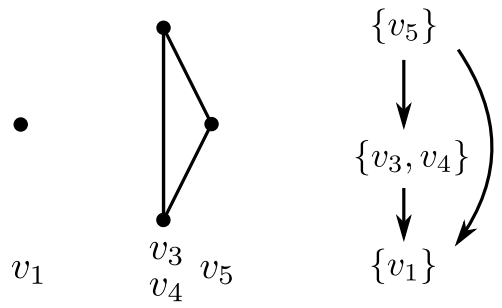

Figure 3.2: The graph $\Gamma' = lk(\{v_6, v_7, v_8\})$ is shown on the left, and on the right is a diagram of the equivalence classes given by the G–ordering of Γ'.

Chapter 4

IA_n and $\mathrm{IA}(A_\Gamma)$

As the commutator subgroup of a group is invariant under automorphisms, there is a homomorphism:

$$\Phi : \mathrm{Aut}(G) \to \mathrm{Aut}(G/[G,G]) \cong \mathrm{Aut}(H_1(G)).$$

We call the kernel of this map $\mathrm{IA}(G)$. The I stands for 'identity' and the A for 'abelianisation.' We have previously used the name *Torelli subgroup*, as this is the name for $\mathrm{IA}(G)$ when G is a surface group, however this isn't very respectful of the history of such groups. Lyndon and Schupp [55, p. 25] tell us that the IA notation was coined by Bachmuth, and we feel that it is more fitting in the general situation.

Inner automorphisms act trivially on $H_1(G)$, therefore there is an induced map

$$\overline{\Phi} : \mathrm{Out}(G) \to \mathrm{Aut}(H_1(G)),$$

and we call the kernel of this map $\overline{\mathrm{IA}}(G)$. When G is a free group we write $\mathrm{IA}(F_n) = \mathrm{IA}_n$. The study of this group goes back to Magnus [56], who found a finite generating set for IA_n. Krstić and McCool [49] have shown that IA_3 is not finitely presentable. Bestvina, Bux, and Margalit [10] have strengthened this result to show that $\overline{\mathrm{IA}}_n$ has cohomological dimension $2n-4$ and $H_{2n-4}(\overline{\mathrm{IA}}_n, \mathbb{Z})$ is not finitely generated. The question of whether IA_n admits a finite presentation for $n \geq 4$ is still a large open problem.

Magnus' generating set of IA_n is given by elements of the form:

$$K_{ij}(x_l) = \begin{cases} x_j x_i x_j^{-1} & i = l \\ x_l & i \neq l \end{cases}$$

$$K_{ijk}(x_l) = \begin{cases} x_i[x_j, x_k] & i = l \\ x_l & i \neq l, \end{cases}$$

where i, j, and k are distinct. Now let Y be a subset of F_n, and define $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$ to be the subgroups of $\mathrm{Aut}(F_n)$ consisting of the automorphisms that respectively fix each element of Y, and fix each element of Y up to conjugacy. Chein [22] showed that when $n = 3$ and Y is a subset of our chosen basis for F_3 then $\mathrm{Fix}(Y) \cap \mathrm{IA}_3$ is generated by the elements of Magnus' generating set that lie $\mathrm{Fix}(Y)$. He also proved the same result for $\mathrm{Fix}_c(Y) \cap \mathrm{IA}_3$ when $Y = \{x_1, x_2, x_3\}$. In his paper, he states: "Some, although not all, of our results can be obtained for $n > 3$ by the same methods." We show that in the $\mathrm{Fix}_c(Y)$ case, this is indeed true:

Theorem 4.4. $\mathrm{IA}_n \cap \mathrm{Fix}_c(\{x_{m+1}, \ldots, x_n\})$ *is generated by the set*

$$\mathcal{C}_m = \{K_{ij} : 1 \leq i \leq n,\ 1 \leq j \leq n\} \cup \{K_{ijk} : 1 \leq i \leq m,\ 1 \leq j, k \leq n\}.$$

We do not give any results for the Fix case, although we compare the above to results of Day and Putman [28] that suggest the analogue of Chein's result may not be true for $n > 3$.

The group $\mathrm{IA}(A_\Gamma)$ has only recently become an object of interest, this coming with the rise of interest in $\mathrm{Aut}(A_\Gamma)$. Day [27] has found a generating set $\mathcal{M}_\Gamma$ of $\mathrm{IA}(A_\Gamma)$ similar to Magnus' generating set of IA_n. In Section 4.1.2 we give a description of $\mathcal{M}_\Gamma$ and a proof of Day's result that $\mathcal{M}_\Gamma$ does indeed generate $\mathrm{IA}(A_\Gamma)$.

In Section 4.2 we move on to studying the Andreadakis–Johnson filtration of $\mathrm{IA}(A_\Gamma)$. This is a central series $\mathcal{G} = \{G_1, G_2, G_3, \ldots\}$ of $\mathrm{IA}(A_\Gamma)$ that is *separating* in the sense that $\cap_{c=1}^{\infty} G_c = \{1\}$. Furthermore each consecutive quotient G_c/G_{c+1} is free-abelian. This latter fact is proved by constructing a *Johnson homomorphism* τ_c from each G_c to an free-abelian group with kernel G_{c+1}. A consequence is that $\mathrm{IA}(A_\Gamma)$ is residually torsion-free nilpotent (and in particular is torsion-free). For a general group G, Bass and Lubotzky [5] give a general method for deciding when the image $\mathcal{H} = \{H_1, H_2, H_3, \ldots\}$ of $\mathcal{G}$ in $\overline{\mathrm{IA}}(G)$ satisfies the same properties. The work on the lower central series of A_Γ discussed in Chapter 2, along with results of Toinet [70] and Minasyan [61] combine to show that this criteria is fulfilled by $\overline{\mathrm{IA}}(A_\Gamma)$. In particular:

Theorem 4.22. *For any graph* Γ*, the group* $\overline{\mathrm{IA}}(A_\Gamma)$ *is residually torsion-free nilpotent.*

This was discovered independently by Toinet [70]. Our methods differ from Toinet's because the full force of the Bass–Lubotzky machinery is not required in the RAAG situation, and we give a simplification of their methods to assemble a proof of Theorem 4.22.

In Section 4.2.3 we study $H_1(\mathrm{IA}(A_\Gamma))$. We show the following:

Theorem 4.23. *The first Johnson homomorphism τ_1 maps $\mathcal{M}_\Gamma$ to a free generating set of a subgroup of* $\mathrm{Hom}(H_1(A_\Gamma), \gamma_2(A_\Gamma)/\gamma_3(A_\Gamma))$. *The abelianisation of* $\mathrm{IA}(A_\Gamma)$ *is isomorphic to the free abelian group on the set $\mathcal{M}_\Gamma$, and G_2 is the commutator subgroup of* $\mathrm{IA}(A_\Gamma)$.

This mimics an analogous result in $\mathrm{IA}(F_n)$, and has the following corollary:

Corollary 4.24. *$\mathcal{M}_\Gamma$ is a minimal generating set of* $\mathrm{IA}(A_\Gamma)$.

We finish this chapter by looking at these results through a specific example: when Γ is a pentagon graph.

4.1 Finitely generated subgroups of IA_n and $\mathrm{IA}(A_\Gamma)$

Magnus' proof that IA_n is finitely generated comes from the following general procedure:

Procedure 4.1. Let G be a group, H a normal subgroup of G and $\overline{G} = G/H$. Let A be a generating set of G, let $\overline{A}$ be the image of A in $\overline{G}$, and let R be a set of words in G such that $\overline{G}$ has the presentation $\overline{G} = \langle \overline{A} | \overline{R} \rangle$. Then H is the subgroup of G normally generated by the elements of R.

If B is a subset of H such that B generates a normal subgroup of G and $\langle B \rangle$ contains R, then B is a generating set of H.

Hence to find a finite generating set for the kernel of a map, one method is to find a presentation for the quotient group and then hope we are lucky enough that a subset B of H satisfying the above criteria is easy to find.

We apply this in two situations. In the first case we show that when Y is a subset of our chosen basis basis for F_n, the intersection $\mathrm{Fix}_c(Y) \cap \mathrm{IA}_n$ is finitely generated, and in the second we describe a theorem of Day [27] that gives a Magnus-esque generating set of $\mathrm{IA}(A_\Gamma)$.

4.1.1 $\mathbf{Fix}_c(\{x_{m+1}, \ldots, x_n\}) \cap \mathrm{IA}_n$

Let $Y = \{x_{m+1}, \ldots, x_n\}$ to be a subset of a fixed basis for F_n. As mentioned previously, $\mathrm{Fix}_c(Y)$ is the subgroup of $\mathrm{Aut}(F_n)$ given by automorphisms that fix each element of $\{x_{m+1}, \ldots, x_n\}$ up to conjugacy. We shall use an adaptation of Magnus' proof (also used by Chein [22]) to show that $\mathrm{Fix}_c(Y) \cap \mathrm{IA}_n$ is generated by Magnus' generators that lie in $\mathrm{Fix}_c(Y)$. (This includes Magnus' theorem in the case $Y = \emptyset$.)

Proposition 4.2. $\mathrm{Fix}_c(Y) = \mathrm{Fix}_c(\{x_{m+1}, \ldots, x_n\})$ *is generated by the set*

$$\mathcal{B}_m = \{K_{ij} : 1 \leq i, j \leq n\} \cup \{s_i, \rho_{ij} : i \leq m, 1 \leq j \leq n\}$$

This can be proved by using peak reduction, and we shall give an alternative proof via folding in Chapter 6. To show that $\mathrm{Fix}_c(Y) \cap \mathrm{IA}_n$ is generated by Magnus' generators that lie in $\mathrm{Fix}_c(Y)$ we shall proceed as follows: we first find a presentation for the group

$$G_m = \left\{ \begin{pmatrix} A & 0 \\ B & I \end{pmatrix} : A \in \mathrm{GL}_m(\mathbb{Z}), B \in M_{n-m,m}(\mathbb{Z}) \right\} \leq \mathrm{GL}_n(\mathbb{Z})$$

in Proposition 4.3. The group G_m is the image of $\mathrm{Fix}_c(Y)$ under the homomorphism $\Phi : \mathrm{Aut}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$. Hence the kernel of this restricted map is $\mathrm{Fix}_c(Y) \cap \mathrm{IA}_n$. It only remains to check that all the relations lie in the subgroup generated by our chosen set, and that this set generates a normal subgroup of $\mathrm{Fix}_c(Y)$.

Let M_{ij} be the matrix taking the value 1 in the (i,j)th entry, and zeroes everywhere else. When $i \neq j$ let $E_{ij} = I + M_{ij}$, and let $T_i = I - 2M_{ii}$, the matrix that takes the value -1 in the (i,i)th entry, 1 in the other diagonal entries, and zero everywhere else. The group G_m is isomorphic to the semidirect product $\mathbb{Z}^{(n-m)m} \rtimes \mathrm{GL}_m(\mathbb{Z})$, where

$$\mathbb{Z}^{(n-m)m} \cong \left\{ \begin{pmatrix} I & 0 \\ B & I \end{pmatrix} \in G_m \right\}$$
$$\mathrm{GL}_m(\mathbb{Z}) \cong \left\{ \begin{pmatrix} A & 0 \\ 0 & I \end{pmatrix} \in G_m \right\},$$

therefore to find a presentation of G_m it is sufficient to find presentations for $\mathbb{Z}^{(n-m)m}$ and $\mathrm{GL}_m(\mathbb{Z})$, and relations that describe the action of $\mathrm{GL}_m(\mathbb{Z})$ on $\mathbb{Z}^{(n-m)m}$ by conjugation. The $\mathbb{Z}^{(n-m)m}$ part of G_m has the obvious presentation $\langle E_{ij} \mid R_{1,m} \rangle$, where $m+1 \leq i \leq n$, $1 \leq j \leq m$ and $R_{1,m}$ contains the commutators of these elements. The $\mathrm{GL}_m(\mathbb{Z})$ part of G_m has a presentation $\langle T_1, E_{ij} \mid R_{2,m} \rangle$, where $1 \leq i, j \leq m$ and

$$R_{2,m} = \left\{ \begin{array}{ll} T_1^2 & \\ (E_{12}E_{21}^{-1}E_{12})^4 & \\ E_{12}E_{21}^{-1}E_{12}E_{21}E_{12}^{-1}E_{21} & \\ [E_{ij}, E_{kl}] & i \neq k, j \neq l \\ [E_{ij}, E_{jk}]E_{ik}^{-1} & i, j, k \text{ distinct} \\ [T_1, E_{ij}] & i \neq 1, j \neq 1 \\ T_1 E_{ij} T_1 E_{ij} & 1 \in \{i, j\} \end{array} \right\}.$$

This is easily deduced from the Steinberg presentation of $\mathrm{SL}_n(\mathbb{Z})$, which can be found in [60, pages 81–82], and the decomposition $\mathrm{GL}_n(\mathbb{Z}) = \mathrm{SL}_n(\mathbb{Z}) \rtimes \langle T_1 \rangle$. There is an exception for $m = 1$, which has the much simpler presentation $\langle T_1 \,|\, T_1^2 \rangle$. The relations that occur from the action of $\mathrm{GL}_m(\mathbb{Z})$ on $\mathbb{Z}^{(n-m)m}$ by conjugation are of the form:

$$R_{3,m} = \left\{ \begin{array}{ll} E_{ij}E_{kl}E_{ij}^{-1} = E_{kl} & i \neq k \\ E_{ij}E_{kl}E_{ij}^{-1} = E_{kj}^{-1}E_{kl} & i = l \text{ and } i,j,k \text{ are distinct} \\ T_1 E_{kl} T_1 = E_{kl} & k,l \neq 1 \\ T_1 E_{kl} T_1 = E_{kl}^{-1} & 1 \in \{k,l\} \end{array} \right\}$$

where E_{ij} is taken over elements in our copy of $\mathrm{GL}_m(\mathbb{Z})$ and E_{kl} is taken over elements in our copy of $\mathbb{Z}^{(n-m)m}$. Summarising:

Proposition 4.3.

$$\langle T_1, E_{ij} \quad 1 \leq i \leq n,\ 1 \leq j \leq m \,|R_{1,m} \cup R_{2,m} \cup R_{3,m} \rangle$$

is a presentation of G_m.

Theorem 4.4. $\mathrm{IA}_n \cap \mathrm{Fix}_c(\{x_{m+1}, \ldots, x_n\})$ *is generated by the set*

$$\mathcal{C}_m = \{K_{ij} : 1 \leq i \leq n,\ 1 \leq j \leq n\} \cup \{K_{ijk} : 1 \leq i \leq m,\ 1 \leq j,k \leq n\}.$$

Proof. We can remove the elements $S_2, \ldots, S_m$ from the generating set $\mathcal{B}_m$ of the group $\mathrm{Fix}_c(Y)$, as $S_i = S_1\rho_{1i}\rho_{i1}^{-1}S_1\rho_{1i}^{-1}S_1\rho_{1i}S_1\rho_{i1}\rho_{1i}^{-1}S_1$, to make a smaller generating set $\mathcal{B}'_m$. Then $\mathcal{B}'_m$ maps onto the generating set of G_m given in Proposition 4.3 by taking $\rho_{ij} \to E_{ji}$, $S_1 \to T_1$. The elements K_{ij} are taken to the identity matrix. Using Procedure 4.1, it suffices to show that $\langle \mathcal{C}_m \rangle$ is a normal subgroup of $\mathrm{Fix}_c(Y)$ that contains the lift of each element of $R_{1,m} \cup R_{2,m} \cup R_{3,m}$ obtained by swapping E_{ij} with ρ_{ji} and T_1 with S_1. It is not hard to check that the lift of each relation to $\mathrm{Aut}(F_n)$ lies in $\langle \mathcal{C}_m \rangle$. To prove normality it is sufficient to show that the conjugate of every element of $\mathcal{C}_m$ by each element of $\mathcal{B}'_m \cup \mathcal{B}'^{-1}_m$ lies in $\mathit{langle}\mathcal{C}_m\rangle$. Most of these computations are simple, except in the case of $\rho_{pk}K_{kpq}\rho_{pk}^{-1}$ and $\rho_{pk}^{-1}K_{kpq}\rho_{pk}$, which we write as products of elements of $\mathcal{C}_m$ below:

$$\rho_{pk}K_{kpq}\rho_{pk}^{-1} = K_{qk}K_{qp}K_{pq}K_{pqk}K_{kp}K_{kpq}K_{kq}^{-1}K_{kp}^{-1}K_{qp}^{-1}K_{qk}^{-1},$$
$$\rho_{pk}^{-1}K_{kpq}\rho_{pk} = K_{qk}^{-1}K_{qp}K_{pq}^{-1}K_{qp}^{-1}K_{kpq}K_{pqk}K_{qk}K_{kq}. \qquad \square$$

The two long identities in this proof are the main difficulty in showing that $\mathrm{Fix}(Y) \cap \mathrm{IA}_n$ is finitely generated. As with the fixing-up-to-conjugacy case, one can show that the set $\mathcal{C}_m$ normally generates $\mathrm{Fix}(Y) \cap \mathrm{IA}_n$, however we cannot prove that $\mathcal{C}_m$ generates a normal subgroup of $\mathrm{Fix}(Y)$: the element K_{qk} appears in our decomposition of $\rho_{pk} K_{kpq} \rho_{pk}^{-1}$, and if $x_q \in Y$, this causes trouble. We would conjecture that $\mathrm{Fix}(Y) \cap \mathrm{IA}_n$ is not finitely generated in general. Our reasoning behind this is a recent paper of Day and Putman [28] who prove the following:

Theorem 4.5. *[28, Theorem A, Theorem E] Let $Y = \{x_{l+1}, \ldots, x_{m+1}, \ldots, x_n\}$ with $l \leq m < n$ and $m \geq 2$. Let*

$$\theta_1 : \mathrm{Fix}_c(Y) \to \mathrm{Aut}(F_m)$$
$$\theta_2 : \mathrm{Fix}(Y) \to \mathrm{Aut}(F_m)$$

be the maps induced by the homomorphism $F_n \to F_m$ taking $x_i \mapsto x_i$ if $i \leq m$ and $x_i \mapsto 1$ if $i > m$. Then $\ker(\theta_1)$ is finitely generated. In contrast, $H_1(\ker(\theta_2), \mathbb{Q})$ is not finitely generated. In particular, $\ker(\theta_2)$ is not finitely generated.

Day and Putman give an explicit description of an infinite set of independent cycles in $H_1(\ker(\theta_2), \mathbb{Q})$. It would be interesting to see if they could be adapted to study $H_1(\mathrm{Fix}(Y) \cap \mathrm{IA}_n, \mathbb{Q})$.

4.1.2 A generating set of $\mathrm{IA}(A_\Gamma)$

We shall now describe an analogous generating set of $\mathrm{IA}(A_\Gamma)$. For each i, j we take K_{ij} to be the partial conjugation described in Section 3.1. We'd also like generators of the form K_{ijk}. The map

$$K_{ijk}(v_l) = \begin{cases} v_i[v_j, v_k] & i = l \\ v_l & i \neq l, \end{cases}$$

induces a nontrivial automorphism of A_Γ if $[v_j, v_k] \neq 1$ and any basis element of A_Γ that commutes with v_i also commutes with v_j and v_k. Hence we are lead to the following definition:

Definition 4.6. Let $\mathcal{M}_\Gamma$ be the subset of $\mathrm{Aut}(A_\Gamma)$ consisting of:

1. Partial conjugations.

2. Elements of the form K_{ijk}, where $[v_j, v_k] \neq 1$, $lk(v_i) \subset st(v_j) \cap st(v_k)$, and $j < k$.

We add the restriction that $j < k$ for elements of the form K_{ijk} as $K_{ijk}^{-1} = K_{ikj}$. Magnus' result may be generalised as follows:

Theorem 4.7. *$\mathcal{M}_\Gamma$ is a finite generating set of* $\mathrm{IA}(A_\Gamma)$.

This was first proved by Day [27]. We shall give our own proof in the work that follows. It is slightly more direct, but our methods do not differ vastly from his. We again set up a situation where we may apply Procedure 4.1. As the image of $\mathrm{Aut}(A_\Gamma)$ in $\mathrm{GL}_n(\mathbb{Z})$ can be quite complicated, our first step is to show that $\mathrm{IA}(A_\Gamma)$ is contained in the finite index subgroup $\mathrm{SAut}^0(A_\Gamma)$. We then find a presentation of $\Phi(\mathrm{SAut}^0(A_\Gamma))$, find lifts of all the relations in terms of our chosen generating set, and show that this generating set generates a normal subgroup of $\mathrm{SAut}^0(A_\Gamma)$.

In the case that the right-angled Artin group is two-dimensional, parts (1) and (2) of the following proposition are shown in Corollary 3.3 of [19]. Below, we alter the proof to work for all RAAGs.

Proposition 4.8. *Let* $\mathrm{Sym}^0(A_\Gamma) = \mathrm{Sym}(A_\Gamma) \cap \mathrm{Aut}^0(A_\Gamma)$. *Let* $\mathcal{G} = \Phi(\mathrm{Aut}^0(A_\Gamma))$, *and let* $Q(\Gamma) = \mathrm{Aut}(A_\Gamma)/\mathrm{Aut}^0(A_\Gamma) \cong \mathrm{Sym}(A_\Gamma)/\mathrm{Sym}^0(A_\Gamma)$.

1. *The group* $\mathrm{Sym}(A_\Gamma)$ *acts on the set of equivalences classes of vertices, and a graph symmetry* α *lies in* $\mathrm{Sym}^0(A_\Gamma)$ *if and only if* α *acts trivially. (i.e.* $[\alpha(v)] = [v]$ *for all* $v \in V(\Gamma)$.*)*

2. *The quotient maps from* $\mathrm{Aut}(A_\Gamma)$ *and* $\mathrm{Sym}(A_\Gamma)$ *to* $Q(\Gamma)$ *split.*

3. $\Phi^{-1}(\mathcal{G}) = \mathrm{Aut}^0(A_\Gamma)$.

Proof. If $\alpha \in \mathrm{Sym}(A_\Gamma)$ then $lk(v_i) \subset st(v_j)$ if and only if $lk(\alpha(v_i)) \subset st(\alpha(v_j))$. Hence α preserves the partial order on $V(\Gamma)$ and so permutes the induced equivalence classes. As two elements v_i, v_j in the same equivalence class can be transposed by the automorphism $s_i\rho_{ij}^{-1}s_is_j\rho_{ji}\rho_{ij}^{-1}$, any permutation that preserves equivalence classes lies in $\mathrm{Sym}^0(A_\Gamma)$. In Section 3.2.3 we saw that matrices in $\mathcal{G}$ are of the form:

$$M = \begin{pmatrix} M_1 & 0 & \dots & 0 \\ * & M_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ * & * & \dots & M_r \end{pmatrix} \tag{4.1}$$

(when there are r equivalence classes of vertices in $V(\Gamma)$), therefore permutation matrices that lie in $\mathcal{G}$ are of the form

$$M = \begin{pmatrix} P_1 & 0 & \dots & 0 \\ 0 & P_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & P_r \end{pmatrix}.$$

Hence, if $\alpha \in \mathrm{Sym}(A_\Gamma)$ does not fix equivalence classes it does not lie in $\mathrm{Aut}^0(A_\Gamma)$ and the coset of $\alpha \in \mathrm{Sym}(A_\Gamma)$ in $Q(\Gamma)$ is determined by how α permutes equivalence classes. We can then define a splitting $Q(\Gamma) \to \mathrm{Sym}(A_\Gamma)$ by choosing the representative from each coset that preserves the numerical ordering of the vertices in each equivalence class. The injection $\mathrm{Sym}(A_\Gamma) \to \mathrm{Aut}(A_\Gamma)$ then gives a splitting $Q(\Gamma) \to \mathrm{Aut}(A_\Gamma)$. Hence, if $\phi \in \mathrm{Aut}(A_\Gamma) - \mathrm{Aut}^0(A_\Gamma)$ we can find a decomposition $\phi = \alpha\phi'$, where $\alpha \in \mathrm{Sym}(A_\Gamma) - \mathrm{Sym}^0(A_\Gamma)$ and $\phi' \in \mathrm{Aut}^0(A_\Gamma)$. As $\Phi(\phi')$ lies in $\mathcal{G}$ and $\Phi(\alpha)$ does not, this tells us that $\Phi(\phi) \notin \mathcal{G}$, so $\Phi^{-1}(\mathcal{G}) = \mathrm{Aut}^0(A_\Gamma)$. □

Proposition 4.9. *According to the notation of Equation* (4.1)*, let* $f\colon \mathcal{G} \to (\pm 1)^r$ *be the homomorphism defined by*

$$f(M) = (\det(M_1), \det(M_2), \ldots, \det(M_r)).$$

Let $\mathcal{H} = \ker(f)$*. Then* $\Phi^{-1}(\mathcal{H}) = \mathrm{SAut}^0(A_\Gamma)$*.*

Proof. If $[v_i] = [v_j]$ then we can write $S_i = S_j\rho_{ji}\rho_{ij}^{-1}S_j\rho_{ji}^{-1}S_j\rho_{ji}S_j\rho_{ij}\rho_{ji}^{-1}S_j$, so we may replace the generators $s_1, \ldots, s_n$ in the generating set of $\mathrm{Aut}^0(A_\Gamma)$ by the smaller set $s_{m_1}, \ldots, s_{m_r}$. Suppose that $\Phi(\phi) \in \mathcal{G}$. By Proposition 4.8, we have $\phi \in \mathrm{Aut}^0(A_\Gamma)$. By the shuffling argument in the proof of Proposition 3.6 we may then shuffle generators to write any element ϕ of $\mathrm{Aut}^0(A_\Gamma)$ in the form $s_{m_1}^{\epsilon_1}s_{m_2}^{\epsilon_2}\cdots s_{m_r}^{\epsilon_r}\phi'$, where $\epsilon_i \in \{0, 1\}$ and $\phi' \in \mathrm{SAut}^0(A_\Gamma)$. As $\Phi(\mathrm{SAut}^0(A_\Gamma))$ is generated by elementary matrices, $\Phi(\mathrm{SAut}^0(A_\Gamma)) \leq \ker f$, and

$$\begin{aligned} f\Phi(\phi) &= f\Phi(s_{m_1}^{\epsilon_1}s_{m_2}^{\epsilon_2}\cdots s_{m_r}^{\epsilon_r}\phi') \\ &= f\Phi(s_{m_2}^{\epsilon_2}\cdots s_{m_r}^{\epsilon_r}) \\ &= ((-1)^{\epsilon_1}, \ldots, (-1)^{\epsilon_r}). \end{aligned}$$

Hence $f\Phi(\phi) = 0$ if and only if $\phi \in \mathrm{SAut}^0(A_\Gamma)$. □

In particular, the above proposition shows that if $\Phi(\phi) = 1$ then $\phi \in \mathrm{SAut}^0(A_\Gamma)$:

Corollary 4.10. $\mathrm{IA}(A_\Gamma)$ *is a subgroup of* $\mathrm{SAut}^0(A_\Gamma)$*.*

The final ingredient we need is a presentation for the group $\mathcal{H}$. Given a presentation of $\mathrm{SL}_n(\mathbb{Z})$, this is an exercise in finding presentations of semidirect products of groups.

Proposition 4.11. *The following four types of relations are sufficient to give a presentation of* $\mathcal{H}$ *with respect to the generating set of elementary matrices* $X = \{E_{ij} : v_j \leq v_i\}$*:*

Type I $[E_{ij}, E_{kl}]$, *when* $i \neq k$ *and* $j \neq l$.

Type II $[E_{ij}, E_{jk}]E_{ik}^{-1}$ *for* i, j, k *distinct.*

Type III $(E_{m_i(m_i+1)}E_{(m_i+1)m_i}^{-1}E_{m_i(m_i+1)})^4$ *if* $m_{i+1} - m_i \geq 2$.

Type IV $E_{m_i(m_i+1)}E_{(m_i+1)m_i}^{-1}E_{m_i(m_i+1)}E_{(m_i+1)m_i}E_{m_i(m_i+1)}^{-1}E_{(m_i+1)m_i}$ *if* $m_{i+1} - m_i = 2$.

Proof. Let $\mathcal{H}_a$ be the subgroup of $\mathcal{H}$ of matrices that are trivial outside the top left $a \times a$ block in the decomposition given in Equation (4.1). This is the subgroup generated by all E_{ij} in X such that $i, j < m_{a+1}$. As $\mathcal{H}_1 \cong \mathrm{SL}_{m_2-m_1}(\mathbb{Z})$ the above relations in the generators of $\mathcal{H}_1$ are those of the Steinberg presentation of $\mathrm{SL}_{m_2-m_1}(\mathbb{Z})$ given on pages 81–82 of [60]. Suppose that the above relations in the generators of $\mathcal{H}_a$ give a presentation of $\mathcal{H}_a$. We can decompose $\mathcal{H}_{a+1}$ as the semidirect product $\mathcal{H}_{a+1} = (\mathcal{H}_a \times \mathcal{H}'_a) \ltimes A$, where

$$\mathcal{H}'_a = \langle E_{ij} : m_{a+1} \leq i, j < m_{a+2} \rangle$$
$$A = \langle E_{ij} : m_{a+1} \leq i < m_{a+2}, 1 \leq j < m_{a+1} \rangle.$$

As $\mathcal{H}'_a \cong \mathrm{SL}_{m_{a+2}-m_{a+1}}(\mathbb{Z})$ and A is free abelian the relations listed above are sufficient to give a presentation for these subgroups. We attain a presentation for $\mathcal{H}_a \times \mathcal{H}'_a$ by taking relations of type I to show that pairs of generators from distinct groups in the product commute, and to get a presentation for the semidirect product we need to encode the conjugation action of $\mathcal{H}_a \times \mathcal{H}'_a$ on A, and this is given by relations of type I and type II. Hence the relations listed above are sufficient to provide a presentation of $\mathcal{H}_{a+1}$. By induction we have provided a presentation of $\mathcal{H}$. □

Proof of Theorem 4.7. By Corollary 4.10, the group $\mathrm{IA}(A_\Gamma)$ is the kernel of the restriction of Φ to $\mathrm{SAut}^0(A_\Gamma)$. The group $\mathrm{SAut}^0(A_\Gamma)$ is generated by partial conjugations and transvections, and these elements are mapped to the identity matrix and elementary matrices respectively. As we have a presentation of $\Phi(\mathrm{SAut}^0(A_\Gamma)) = \mathcal{H}$ where the generators are the images of our chosen generators of $\mathrm{SAut}^0(A_\Gamma)$, any element of $\mathrm{IA}(A_\Gamma)$ can be written as a product of conjugates of the lifts of the relations given in Proposition 4.11. Therefore we need to show that every such lift can be written as a product of elements in $\mathcal{M}_\Gamma$, and that $\langle \mathcal{M}_\Gamma \rangle$ is normal in $\mathrm{SAut}^0(A_\Gamma)$. Firstly, let's go through the relations. As $\rho_{ij} \mapsto E_{ji}$ we have the following to check:

Type I $[\rho_{ji}, \rho_{lk}] = \begin{cases} 1 & \text{if } j \neq l \text{ or } [v_i, v_k] = 0 \\ K_{jik} & \text{if } j = l \text{ and } [v_i, v_k] \neq 0, \end{cases}$

Type II $[\rho_{ji}, \rho_{kj}]\rho_{ki}^{-1} = \begin{cases} 1 & [v_j, v_i] = 0 \\ K_{kji} & [v_j, v_i] \neq 0 \end{cases}$

Type III $(\rho_{(m_i+1)m_i}\rho_{m_i(m_i+1)}^{-1}\rho_{(m_i+1)m_i})^4$ is trivial when $[v_{m_i}, v_{m_i+1}] = 0$ and is equal to $K_{m_i(m_i+1)}K_{(m_i+1)m_i}K_{m_i(m_i+1)}^{-1}K_{(m_i+1)m_i}^{-1}$ when $[v_{m_i}, v_{m_i+1}] \neq 0$.

Type IV $\rho_{(m_i+1)m_i}\rho_{m_i(m_i+1)}^{-1}\rho_{(m_i+1)m_i}\rho_{m_i(m_i+1)}\rho_{(m_i+1)m_i}^{-1}\rho_{m_i(m_i+1)}$ is trivial in $\mathrm{Aut}(A_\Gamma)$.

The final part of the proof — showing that $\langle \mathcal{M}_\Gamma \rangle$ is a normal subgroup of $\mathrm{SAut}^0(A_\Gamma)$ is rather technical, and rather than give the details here we will provide them in Appendix A. □

4.2 The Andreadakis–Johnson Filtration of $\mathrm{IA}(A_\Gamma)$

In this section we follow the methods of Bass and Lubotzky [5] to extend the notion of *higher Johnson homomorphisms* from the free group setting to general right-angled Artin groups. These allow us to describe the abelianisation of $\mathrm{IA}(A_\Gamma)$, and show that $\mathrm{IA}(A_\Gamma)$ has a separating central series $G_1, G_2, G_3, \ldots$ where each quotient G_i/G_{i+1} is a finitely generated free abelian group. This was first studied in the case of free groups by Andreadakis. We show that the image of this series in $\mathrm{Out}(A_\Gamma)$ satisfies the same results.

We shall be using results on the lower central series algebra of A_Γ studied in Chapter 2. We denote this algebra L. This simplified notation should not cause any trouble, as the three algebras we studied in Chapter 2 (the lower central series algebra $L_{\mathcal{C}}$, the central series algebra coming from the Magnus map $L_{\mathcal{D}}$, and algebra of Lyndon elements L_Γ) were shown to be isomorphic. We write $\gamma_c(A_\Gamma)$ as simply γ_c and label each graded piece as $L_c = \gamma_c/\gamma_{c+1}$. In particular $L_1 = H_1(A_\Gamma)$. Let $(L)_p$ be the Lie algebra obtained by taking the tensor product of $\mathbb{Z}/p\mathbb{Z}$ with L and L^c be the quotient algebra $L/\oplus_{i>c} L_i$. We will use the following results:

Theorem 2.53. *If $k \in \mathbb{N}$, then $\gamma_k(A_\Gamma)/\gamma_{k+1}(A_\Gamma)$ is free-abelian, and $A_\Gamma/\gamma_k(A_\Gamma)$ is torsion-free nilpotent.*

Theorem 2.56. *If $Z(A_\Gamma) = \{1\}$ then $Z(L) = Z((L)_p) = 0$ and $Z(L^c)$ is the image of L_c under the quotient map $L \to L^c$.*

4.2.1 A central filtration of $\mathrm{IA}(A_\Gamma)$

As γ_c is characteristic, there is a natural map $\mathrm{Aut}(A_\Gamma) \to \mathrm{Aut}(A_\Gamma/\gamma_c)$. Let G_{c-1} be the kernel of this map. Then $G_0 = \mathrm{Aut}(A_\Gamma)$ and $G_1 = \mathrm{IA}(A_\Gamma)$. If $g \in A_\Gamma$ and $\phi \in G_c$ we write g_ϕ for the unique element of γ_{c+1} such that $\phi(g) = gg_\phi$.

Lemma 4.12. *Let $\phi, \psi \in G_c$ and $g, h \in \gamma_d$. Then:*

(1). $g_\phi \in \gamma_{c+d}$, so that $\phi(g) = g \mod \gamma_{c+d}$.

(2). $(g^{-1})_\phi = g_\phi^{-1} \mod \gamma_{c+2d}$.

(3). $(gh)_\phi = g_\phi h_\phi \mod \gamma_{c+2d}$.

(4). $(g)_{\phi\psi} = g_\phi g_\psi \mod \gamma_{2c+d}$.

Proof. We prove (1) by induction on d. This holds for $d = 1$ from the definition of G_c. Now suppose that (1) holds for $d-1$. Let $g \in \gamma_1$ and $h \in \gamma_{d-1}$. Then $g_\phi \in \gamma_{c+1}$, $hh_\phi = \phi(h) \in \gamma_{d-1}$, and by induction $h_\phi \in \gamma_{c+d-1}$. We can apply the commutator identities $[xy, z] = {}^x[y, z].[x, z]$ and $[x, yz] = [x, y].{}^y[x, z]$ from Lemma 2.3 to show that

$$\begin{aligned}
\phi([g,h]) &= [\phi(g), \phi(h)] \\
&= [gg_\phi, hh_\phi] \\
&= {}^g[g_\phi, hh_\phi].[g, hh_\phi] \\
&= {}^g[g_\phi, hh_\phi].[g,h].{}^h[g, h_\phi] \\
&= [g,h] && \mod \gamma_{c+d}.
\end{aligned}$$

As γ_d is generated by such commutators, the proof of (1) follows. As $g_\phi = g^{-1}\phi(g)$ we calculate:

$$\begin{aligned}
(g^{-1})_\phi &= g\phi(g^{-1}) \\
&= gg_\phi^{-1}g^{-1} \\
&= [g, g_\phi^{-1}]g_\phi^{-1}
\end{aligned}
\qquad
\begin{aligned}
(gh)_\phi &= h^{-1}g^{-1}\phi(gh) \\
&= h^{-1}g^{-1}gg_\phi hh_\phi \\
&= h^{-1}g_\phi hh_\phi \\
&= [h^{-1}, g_\phi]g_\phi h_\phi
\end{aligned}
\qquad
\begin{aligned}
(gh)_{\phi\psi} &= g^{-1}\phi\psi(g) \\
&= g^{-1}\phi(gg_\psi) \\
&= g^{-1}gg_\phi g_\psi (g_\psi)_\phi \\
&= g_\phi g_\psi (g_\psi)_\phi
\end{aligned}$$

Parts (2), (3) and (4) follow from these calculations and part (1).

□

We may now define the notion of *higher Johnson homomorphisms*. The word 'higher' here arises as in Johnson's work on Torelli subgroups of mapping class groups, he looked at the first homomorphism τ_1 in the sequence $\tau_1, \tau_2, \tau_3, \ldots$ which we define below. (Of course, we replace the fundamental group of a surface with A_Γ.)

Proposition 4.13. *Let $\phi \in G_c$, where $c \geq 1$. The mapping $g.\gamma_2 \mapsto g_\phi.\gamma_{c+2}$ induces a homomorphism*

$$\tau_c \colon G_c \to \mathrm{Hom}(L_1, L_{c+1})$$

such that $\ker(\tau_c) = G_{c+1}$. *We say that τ_c is the* cth Johnson homomorphism.

Proof. Suppose that g_1 and g_2 are elements of A_Γ representing the same element of L_1. Then there exists $h \in \gamma_2$ such that $g_1 = g_2 h$. By part (1) of Lemma 4.12 we have $h_\phi \in \gamma_{c+2}$, therefore by part (3) of Lemma 4.12:

$$(g_1)_\phi = (g_2 h)_\phi = (g_2)_\phi h_\phi = (g_2)_\phi \mod \gamma_{c+2}.$$

Hence the map $g.\gamma_2 \mapsto g_\phi.\gamma_{c+2}$ is well-defined. Part (3) of Lemma 4.12 also implies that the map $g.\gamma_2 \mapsto g_\phi.\gamma_{c+2}$ is a homomorphism, so that τ_c is well defined. Finally, part (4) of Lemma 4.12 tells us that τ_c is a homomorphism. An automorphism ϕ lies in $\ker(\tau_c)$ if and only if $g_\phi \in \gamma_{c+2}$ for all $g \in A_\Gamma$, and this happens if and only if $\phi \in G_{c+1}$. Hence $\ker(\tau_c) = G_{c+1}$. □

As $\cap_{c=1}^{\infty} \gamma_c = \{1\}$, it follows that $\cap_{c=1}^{\infty} G_c = \{1\}$. As L_1 and L_{c+1} are free abelian, $\mathrm{Hom}(L_1, L_{c+1})$ is free abelian. As G_c/G_{c+1} is isomorphic to a subgroup of a free abelian group, it is also free abelian.

Proposition 4.14. *$G_1, G_2, G_3, \ldots$ is a central series of* $\mathrm{IA}(A_\Gamma)$.

Proof. $G_1 = \mathrm{IA}(A_\Gamma)$ and $G_{c+1} < G_c$ for all c, therefore we only need to check that $[G_c, G_d] \subset G_{c+d}$ for all $c, d \geq 1$. Suppose that $\phi \in G_c$ and $\psi \in G_d$. If $g \in A_\Gamma$, then $g_\phi \in \gamma_{c+1}$ and $g_\psi \in \gamma_{d+1}$ and we may repeatedly apply parts 1– 3 of Lemma 4.12 to show that:

$$\begin{aligned}
[\phi, \psi](g) &= \phi\psi\phi^{-1}\psi^{-1}(g) \\
&= g g_\phi g_\psi g_\phi^{-1} g_\psi^{-1} && \mod \gamma_{c+d+1} \\
&= g && \mod \gamma_{c+d+1}
\end{aligned}$$

Hence $[\phi, \psi] \in G_{c+d}$. □

4.2.2 The image of the filtration in $\overline{\mathrm{IA}}(A_\Gamma)$

We'd now like to move our attention to the images of $G_1, G_2, \ldots$ in $\mathrm{Out}(A_\Gamma)$, which we will label $H_1, H_2, \ldots$ respectively. Let π be the natural projection $\mathrm{Aut}(A_\Gamma) \to \mathrm{Out}(A_\Gamma)$. The action of an element of A_Γ on itself by conjugation induces a homomorphism $ad\colon A_\Gamma \to \mathrm{Aut}(A_\Gamma)$. For each $g \in A_\Gamma \setminus \{1\}$ there exists a unique integer d such that $g \in \gamma_d$ and $g \notin \gamma_{d+1}$. We identify g with the element $g\gamma_{d+1}$ in the submodule L_d of the Lie algebra L. We use this to study the map ad in the following lemma:

Lemma 4.15. *If $Z(A_\Gamma) = \{1\}$, then $g \in \gamma_c$ if and only if $ad(g) \in G_c$.*

Proof. If $g \in \gamma_c$ then $ghg^{-1} = h \mod \gamma_{c+1}$, for all $h \in A_\Gamma$. Hence $ad(g) \in G_c$. Conversely, suppose that $g \notin \gamma_c$. Then $g \in L_d$ for some $d < c$, and by Theorem 2.56, the image of $g\gamma_{d+1}$ under the quotient map $L \to L^c$ is not central. As L is generated by $v_1, \ldots, v_n$, there exists v_i such that the image of $[g, v_i]$ is nonzero in L^c. Hence $[g, v_i] \neq 1 \mod \gamma_{c+1}$, so $gv_ig^{-1} \neq v_i \mod \gamma_{c+1}$, and $ad(g) \notin G_{c+1}$. □

Proposition 4.16. *When $Z(A_\Gamma) = \{1\}$ we have an exact sequence of abelian groups:*

$$0 \to \gamma_c/\gamma_{c+1} \xrightarrow{\alpha} G_c/G_{c+1} \xrightarrow{\beta} H_c/H_{c+1} \to 0,$$

where α and β are induced by ad and π respectively.

Proof. By Lemma 4.15, the map α is injective. The sequence $H_1, H_2, \ldots$, is defined to be the image of $G_1, G_2, \ldots$ in $\mathrm{Out}(A_\Gamma)$, so the map β is surjective. Furthermore, $\pi(ad(A_\Gamma)) = \{1\}$, so that $\mathrm{im}(\alpha) \subset \ker(\beta)$. It remains to check that $\ker(\beta) \subset \mathrm{im}(\alpha)$. Suppose that $\phi \in G_c$ and $[\phi] \in H_{c+1}$. Then there exists $g \in A_\Gamma$ and $\psi \in G_{c+1}$ such that $\phi = ad(g)\psi$. Then $ad(g) = \phi\psi^{-1} \in G_c$, so by Lemma 4.15 we have $g \in \gamma_c$. Hence $\phi G_{c+1} = ad(g)G_{c+1}$ and ϕG_{c+1} is in the image of γ_c/γ_{c+1}. □

Theorem 4.17. *If $Z(A_\Gamma) = \{1\}$ and $c \geq 1$ then the group H_c/H_{c+1} is free abelian.*

Proof. By the exact sequence above,

$$H_c/H_{c+1} \cong (G_c/G_{c+1})/(ad(\gamma_c)G_{c+1}/G_{c+1}).$$

As G_c/G_{c+1} is free abelian, it is sufficient to show that if there exists $\phi \in G_c$, $g \in \gamma_c$, and a prime p such that $ad(g)G_{c+1} = \phi^p G_{c+1}$, then there exists $h \in \gamma_c$ such that $ad(h)G_{c+1} = \phi G_{c+1}$.

Suppose that ϕ, g, and p exist as above. As $\phi \in G_c$, for every generator v of A_Γ we have $v^{-1}\phi(v) = v_\phi \in \gamma_{c+1}$. By part (4) of Lemma 4.12:

$$\phi^p(v) = vv_\phi^p \mod \gamma_{c+2},$$

therefore

$$gvg^{-1} = vv_\phi^p \mod \gamma_{c+2}.$$

Hence $[g,v]\gamma_{c+2} = v_\phi^p\gamma_{c+2}$ and $[g,v]$ is zero in $(L)_p$. As $((L)_p)$ is generated by such elements, the image of g lies in $Z((L)_p)$. However, Theorem 2.56 tells us that $Z((L)_p) = 0$, therefore g must lie in the kernel of the map $L \to (L)_p$. Hence there exists $h \in \gamma_c$ such that $g = h^p \mod \gamma_{c+1}$. Then $ad(h^p)G_{c+1} = ad(g)G_{c+1} = \phi^p G_{c+1}$. As G_c/G_{c+1} is torsion-free it has unique roots. Hence $ad(h)G_{c+1} = \phi G_{c+1}$. □

We'd now like to show that the filtration $H_1, H_2, H_3, \ldots$ is separating (has trivial intersection). We first state two theorems without proof:

Theorem 4.18 (Toinet, [70])**.** *Right-angled Artin groups are conjugacy separable in their nilpotent quotients: if $g, h \in A_\Gamma$ are conjugate in A_Γ/γ_c for all c, then g and h are conjugate in A_Γ.*

Proposition 4.19 (Minasyan, [61], Proposition 6.9)**.** *Let $\phi \in \mathrm{Aut}(A_\Gamma)$. If $\phi(g)$ is conjugate to g for all $g \in A_\Gamma$ then ϕ is an inner automorphism, i.e. $\phi = ad(h)$ for some $h \in A_\Gamma$.*

These combine to give:

Proposition 4.20. *The intersection $\cap_{c=1}^\infty H_c$ is trivial.*

Proof. Let $\phi \in \mathrm{Aut}(A_\Gamma)$, and suppose that $[\phi] \in H_c$ for all c. Let $g \in A_\Gamma$. We know that $\phi(g)$ is conjugate to g in A_Γ/γ_c for all c. Therefore by Toinet's theorem, g is conjugate to $\phi(g)$. As this applies to every element of A_Γ, by the proposition of Minasyan, ϕ itself is an inner automorphism. Hence $\cap_{c=1}^\infty H_c = \{1\}$. □

Therefore if $Z(A_\Gamma)$ is trivial then $H_1, H_2, H_3 \ldots$ is central series of $\overline{\mathrm{IA}}(A_\Gamma)$, with trivial intersection, such that the consecutive quotients H_c/H_{c+1} are free abelian.

Corollary 4.21. *If $Z(A_\Gamma)$ is trivial then $\overline{\mathrm{IA}}(A_\Gamma)$ is residually torsion-free nilpotent.*

Suppose that $Z(A_\Gamma)$ is nontrivial. Let $[v]$ be the unique maximal equivalence class of vertices in Γ. By Proposition 3.10, there is a restriction map R_v and a projection map P_v like so:

$$R_v\colon \mathrm{Out}(A_\Gamma) \to \mathrm{GL}(Z(A_\Gamma)) \cong \mathrm{GL}(A_{[v]})$$
$$P_v\colon \mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_\Gamma/Z(A_\Gamma)) \cong \mathrm{Out}(A_{lk[v]}).$$

Also by Proposition 3.10, the kernel of the map $R_v \times P_v$ is the free abelian subgroup Tr generated by the transvections $[\rho_{ij}]$ such that $v_i \in lk[v]$ and $v_j \in [v]$. Elements of $\mathrm{Out}(A_\Gamma)$ that lie in Tr act non-trivially on $H_1(A_\Gamma)$, as do elements that are nontrivial under R_v. It follows that $\overline{\mathrm{IA}}(A_\Gamma)$ is mapped isomorphically under P_v onto $\overline{\mathrm{IA}}(A_{lk[v]})$. As the centre of $lk[v]$ is trivial, this lets us promote the above work to any right-angled Artin group:

Theorem 4.22. *For any graph Γ, the group $\mathrm{IA}(A_\Gamma)$ is residually torsion-free nilpotent.*

4.2.3 The structure of $H_1(\mathrm{IA}(A_\Gamma))$

The rank of each L_c has been calculated in [31], although more work is needed to calculate the ranks of the quotients G_c/G_{c+1}. In the free group case G_1/G_2, G_2/G_3 and G_3/G_4 are known [64, 66] but as yet there is no general formula. We restrict ourselves to studying the abelianisation of $\mathrm{IA}(A_\Gamma)$, using the generating set $\mathcal{M}_\Gamma$ given in Definition 4.6.

Theorem 4.23. *The first Johnson homomorphism τ_1 maps $\mathcal{M}_\Gamma$ to a free generating set of a subgroup of $\mathrm{Hom}(L_1, L_2)$. The abelianisation of $\mathrm{IA}(A_\Gamma)$ is isomorphic to the free abelian group on the set $\mathcal{M}_\Gamma$, and G_2 is the commutator subgroup of $\mathrm{IA}(A_\Gamma)$.*

Proof. We found the following basis for L_2 in proposition 2.55:

$$S = \{[v_i, v_j]\gamma_3(A_\Gamma) : i < j, [v_i, v_j] \neq 0\}.$$

This allows us to obtain an explicit description of the images of elements of $\mathcal{M}_\Gamma$:

$$\tau_1(K_{ij})(v_l) = \begin{cases} 1\gamma_3(A_\Gamma) & \text{if } v_l \notin \Gamma_{ij} \\ [v_j, v_l]\gamma_3(A_\Gamma) & \text{if } v_l \in \Gamma_{ij} \end{cases}$$

$$\tau_1(K_{ijk})(v_l) = \begin{cases} 1\gamma_3(A_\Gamma) & \text{if } l \neq i \\ [v_j, v_k]\gamma_3(A_\Gamma) & \text{if } l = i \end{cases}$$

These elements are linearly independent in $\mathrm{Hom}(L_1, L_2)$. The second statement follows immediately, and the third follows as $G_2 = \ker(\tau_1)$. □

We have show that the rank of $H_1(\mathrm{IA}(A_\Gamma))$ is equal to the size of $\mathcal{M}_\Gamma$.

Corollary 4.24. *$\mathcal{M}_\Gamma$ is a minimal generating set of $\mathrm{IA}(A_\Gamma)$.*

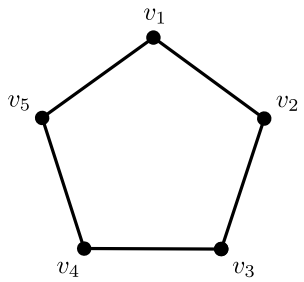

Figure 4.1: The pentagon graph

Example 4.25 (The pentagon graph). Suppose that Γ is the pentagon shown in Figure 4.1. In this case, if $v_i \leq v_j$ then $v_i = v_j$, therefore no elements of the form K_{ijk} exist in $\mathrm{IA}(A_\Gamma)$. Removing $st(v_i)$ from Γ leaves exactly one connected component consisting of the two vertices opposite v_i, therefore

$$\mathcal{M}_\Gamma = \{K_{13}, K_{24}, K_{35}, K_{41}, K_{52}\}.$$

Hence $H_1(\mathrm{IA}(A_\Gamma)) = G_1/G_2 \cong \mathbb{Z}^5$. Also, $\{[v_1, v_3], [v_1, v_4], [v_2, v_4], [v_2, v_5], [v_3, v_5]\}$ is a set of coset representatives of $\gamma_2(A_\Gamma)$ in $\gamma_3(A_\Gamma)$, therefore $\mathrm{Hom}(L_1, L_2) \cong \mathbb{Z}^{25}$. In particular τ_1 is not surjective (in contrast to the free group situation — see [64]).

Chapter 5

Homomorphisms to $\mathrm{Out}(F_n)$ and $\mathrm{Out}(A_\Gamma)$

In this chapter we give results that place restrictions on groups that can map to $\mathrm{Out}(F_n)$ and $\mathrm{Out}(A_\Gamma)$ with infinite image. We start with $\mathrm{Out}(F_n)$. We prove:

Theorem 5.1. *Let Λ be a group. Suppose that no subgroup of finite index in Λ has a normal subgroup that maps surjectively to $\mathbb{Z}$. Then every homomorphism from Λ to the outer automorphism group of a finitely generated free group has finite image.*

Note that in Theorem 5.1 and indeed in any of the results in this chapter, we do not assume that Λ is finitely generated.

We say that a group satisfying the hypothesis of Theorem 5.1 is $\mathbb{Z}$*–averse*. The Normal Subgroup Theorem of Kazhdan and Margulis [75] tells us that irreducible lattices in connected, higher-rank, semisimple Lie groups with finite centre have no infinite normal subgroups of infinite index. Since such lattices are not virtually cyclic, it follows that they are $\mathbb{Z}$–averse. Hence we have the following theorem:

Theorem 5.2. *If G is a connected, semisimple Lie group of real rank at least 2 that has finite centre, and Λ is an irreducible lattice in G, then every homomorphism from Λ to the outer automorphism group of a finitely generated free group has finite image.*

An additional argument allows one to remove the hypothesis that G has finite centre (see Remark 5.7). Further examples of $\mathbb{Z}$–averse groups come from Bader and Shalom's recent work on the Normal Subgroup Theorem [4]. If Λ is a hereditarily just infinite group (i.e. every finite index subgroup of Λ has no non-trivial, infinite quotients) is not virtually cyclic then it is $\mathbb{Z}$–averse; examples are described in [40].

Theorem 5.2 has implications for the Zimmer programme [76], broadly speaking, the aim of which is to understand the actions of such lattices on manifolds. Specifically, it allows one to extend Farb and Shalen's theorem about actions of higher-rank

lattices on 3–manifolds to the general case, removing the non-uniform hypothesis from Theorem II of [34] and recasting their Theorem III as follows:

Theorem 5.3. *Let Λ be an irreducible uniform lattice in a semisimple Lie group of real rank at least 2, and let M be any closed, orientable, connected 3–manifold. Then for every action $\Lambda \to \mathrm{Homeo}(M)$, the image of Λ in $\mathrm{Aut}(H_*(M, \mathbb{Q}))$ is finite.*

Our proof of Theorem 5.1 relies heavily on recent results of Bestvina and Feighn [11], Dahmani, Guirardel and Osin [25], and Handel and Mosher [43]. The work of Bestvina and Feighn was inspired in part by the desire to prove Theorem 5.2, following the lines of the proof given in the case of mapping class groups by Bestvina and Fujiwara [13], which invokes Burger and Monod's theorem that irreducible lattices in higher-rank Lie groups have trivial bounded cohomology [18]. One can replace this use of bounded cohomology with an argument of Dahmani, Guirardel and Osin that applies small cancellation theory to the study of purely pseudo-Anosov subgroups; this is used in [25] to prove an analogue of Theorem 5.1 for homomorphisms to mapping class groups. We also give a version of Theorem 5.1 in terms of bounded cohomology:

Theorem 5.8. *Let Λ be a group. Suppose that for every finite index subgroup $\Lambda' \subset \Lambda$ the second bounded cohomology of Λ' is finite dimensional and $\mathrm{Hom}(\Lambda, \mathbb{R}) = 0$. Then every homomorphism $\phi : \Lambda \to \mathrm{Out}(F_n)$ has finite image.*

Whether one uses bounded cohomology or the alternative endgame from [25], the key step in the Bestvina–Feighn–Fujiwara approach is to get a finitely generated subgroup of $\mathrm{Out}(F_n)$ to act in a suitable way on a hyperbolic metric space. Bestvina and Feighn [11] construct such actions for subgroups of $\mathrm{Out}(F_n)$ that contain a fully irreducible automorphism, and hence deduce that a higher-rank lattice cannot map onto such a subgroup. If one could construct suitable actions for more general subgroups of $\mathrm{Out}(F_n)$, then Theorem 5.1 would follow. A significant step in this direction was taken recently by Handel and Mosher [43], who proved that if a subgroup $H < \mathrm{Out}(F_n)$ does not contain the class of a fully irreducible automorphism, then H has a subgroup of finite index that leaves the conjugacy class of a proper free factor of F invariant. Handel and Mosher also indicate that they hope to extend their work so as to prove Theorem 5.2 along the lines sketched above.

Our proof of Theorem 5.1 proceeds as follows. In Proposition 5.4 we shall use the results of Handel–Mosher and Dahmani–Guirardel–Osin to see that if Λ is $\mathbb{Z}$–averse, then the image of every homomorphism $\Lambda \to \mathrm{Out}(F_n)$ will have a subgroup of finite index that lies in the kernel $\overline{\mathrm{IA}}_n$ of the map $\mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$ given by the action

of $\mathrm{Out}(F_n)$ on the first homology of F_n. In Theorem 4.22 of the previous chapter we showed that $\overline{\mathrm{IA}}_n$ is residually nilpotent, and therefore every nontrivial subgroup of $\overline{\mathrm{IA}}_n$ maps onto $\mathbb{Z}$. As no finite index subgroup of Λ maps onto $\mathbb{Z}$, this completes the proof of Theorem 5.1.

We then move on to $\mathrm{Out}(A_\Gamma)$. We look at subgroups $G \leq \mathrm{Out}(A_\Gamma)$ generated by subsets of our generating set $\mathcal{S}_\Gamma$ (described in Section 3.1). In Section 3.2.2 we defined a partial order on $V(\Gamma)$ given by letting $v_i <_G v_j$ if either $i = j$ or $[\rho_{ij}] \in G$. This partial order induces an equivalence relation $\sim_G$ on $V(\Gamma)$, and each equivalence class generates either a free or free-abelian subgroup of A_Γ. We defined the SL–dimension of G to be equal to the size of a largest abelian equivalence class in $\sim_G$. The name SL–dimension comes from the fact that if $[v]_G$ is an abelian equivalence class then the subgroup of G generated by the $[\rho_{ij}]$ with $v_i, v_j \in [v]_G$ is isomorphic to $\mathrm{SL}_{|[v]_G|}(\mathbb{Z})$. We dedicate Section 5.2 to proving the following key result:

Theorem 5.10. *Suppose that G is a subgroup of $\mathrm{Out}(A_\Gamma)$ generated by a subset $T \subset \mathcal{S}_\Gamma$ and $d_{SL}(G) \leq m$. Let $F(\Gamma)$ be the size of a maximal, discrete, full subgraph of A_Γ. Let Λ be a group. Suppose that for each finite index subgroup $\Lambda' \leq \Lambda$, we have:*

- *Every homomorphism $\Lambda' \to \mathrm{SL}_m(\mathbb{Z})$ has finite image,*
- *For all $N \leq F(\Gamma)$, every homomorphism $\Lambda' \to \mathrm{Out}(F_N)$ has finite image.*
- $\mathrm{Hom}(\Lambda', \mathbb{Z}) = 0$

Then every homomorphism $f : \Lambda \to G$ has finite image.

The proof follows a similar structure to that of Theorem 5.1. Suppose Λ satisfies the hypothesis of the above theorem. In Chapter 3 we described restriction and projection homomorphisms from $\mathrm{Out}^0(A_\Gamma)$ to the outer automorphism groups of RAAGs associated to smaller defining graphs, and showed that the SL–dimension of the images of such homomorphisms does not increase. We use this to study the action of Λ on $H_1(A_\Gamma)$, and show that, up to finite index, this action must be trivial. We then use our result from Chapter 4 that $\overline{\mathrm{IA}}(A_\Gamma)$ is residually torsion-free nilpotent to deduce that the image of Λ in G must be finite. In Section 5.3 we combine the above theorem with our result in the $\mathrm{Out}(F_n)$ case to restrict the maps from $\mathbb{Z}$-averse groups and higher-rank lattices to $\mathrm{Out}(A_\Gamma)$. In particular, we justify our definition of SL–dimension in the following sense:

Corollary 5.17. *Let $k \geq 3$. Then $\mathrm{Out}(A_\Gamma)$ contains a subgroup isomorphic to $\mathrm{SL}_k(\mathbb{Z})$ if and only if $k \leq d_{SL}(\mathrm{Out}(A_\Gamma))$.*

5.1 Homomorphisms to $\mathrm{Out}(F_n)$

We fix a $\mathbb{Z}$–averse group Λ. We do not assume that Λ is finitely generated.

5.1.1 Controlling the action of Λ on homology

Proposition 5.4. *For every subgroup of finite index $\Lambda' \subset \Lambda$ and every homomorphism $\phi : \Lambda' \to \mathrm{Out}(F_n)$, the intersection $\phi(\Lambda') \cap \overline{\mathrm{IA}}_n$ has finite index in $\phi(\Lambda')$.*

Proof. The proof is by induction on n. The case $n = 1$ is trivial. $\mathrm{Out}(F_2)$ has a free subgroup of finite index and no subgroup of finite index in Λ can map onto a free group, so every homomorphism $\Lambda' \to \mathrm{Out}(F_2)$ has finite image.

Suppose $n \geq 3$. Recall that $\psi \in \mathrm{Aut}(F_n)$ (and its image in $\mathrm{Out}(F_n)$) is said to be *fully irreducible* if no power of ψ sends a proper free factor of F_n to a conjugate of itself. Let $[\psi]$ denote the image of ψ in $\mathrm{Out}(F_n)$. Using the actions constructed in [11] and drawing on the approach to small cancellation theory developed in [29], Dahmani, Guirardel and Osin [25] prove that if ψ is fully irreducible then for some positive integer N, the normal closure of $[\psi]^N$ is a free group. It follows that any subgroup of $\mathrm{Out}(F_n)$ that contains a fully irreducible automorphism also contains an infinite normal subgroup that is free. In particular, $\phi(\Lambda')$ cannot contain a fully irreducible automorphism.

According to [43], if $\phi(\Lambda')$ does not contain a fully irreducible automorphism then a subgroup of finite index $H \subset \phi(\Lambda')$ leaves a free factor of F_n invariant up to conjugacy; say $F_n = L * L'$, where $\psi(L) = g_\psi^{-1} L g_\psi$ for all $[\psi] \in H$. Note that the image in $\mathrm{Out}(L)$ of $x \mapsto g_\psi \psi(x) g_\psi^{-1}$, which we denote $[\psi]_L$, depends only on the image of ψ in $\mathrm{Out}(F_n)$, and that $[\psi] \mapsto [\psi]_L$ defines a homomorphism from H to $\mathrm{Out}(L)$. Likewise, the action on the quotient $F_n / \langle\langle L \rangle\rangle$ induces a homomorphism $H \to \mathrm{Out}(L')$. By induction, we know that the induced action of H on the abelianisation of both L and L' factors through a finite group. Thus the action of H on the abelianisation of $F_n = L * L'$ lies in a block triangular subgroup (with respect to a basis that is the union of bases for L and L')

$$\begin{pmatrix} G & 0 \\ * & G' \end{pmatrix} \leq \mathrm{GL}_n(\mathbb{Z})$$

where G and G' are finite. This matrix group is finitely generated and virtually abelian, whereas Λ, and therefore H, does not have a subgroup of finite index that maps onto $\mathbb{Z}$. Thus the action of H on the homology of F_n factors through a finite group, and hence that of $\phi(\Lambda')$ does too, i.e. $\phi(\Lambda') \cap \overline{\mathrm{IA}}_n$ has finite index in $\phi(\Lambda')$. This completes the induction. □

It follows that if Λ' is a non-trivial subgroup of $\overline{\mathrm{IA}}_n = H_1$ then there exists $c \geq 1$ such that $\Lambda' \leq H_c$ and $\Lambda' \not\leq H_{c+1}$. In Theorem 4.17 we saw that for $c \geq 1$ the quotient H_c/H_{c+1} is a finitely generated free abelian group.

Corollary 5.5. *Every non-trivial subgroup of $\overline{\mathrm{IA}}_n$ maps onto $\mathbb{Z}$.*

As Λ is $\mathbb{Z}$–averse, this completes the proof of Theorem 5.1.

5.1.2 Alternative Hypotheses

It emerges from the proofs in the previous section that one can weaken the hypotheses of Theorem 5.1 as follows. Note that since we have not assumed Λ to be finitely generated, condition (1) is not equivalent to assuming that every finite index subgroup of Λ has finite abelianisation.

Theorem 5.6. *Let Λ be a group. Suppose that each finite-index subgroup $\Lambda' \subset \Lambda$ satisfies the following conditions:*

1. *Λ' does not map surjectively to $\mathbb{Z}$;*
2. *Λ' does not have a quotient containing a non-abelian, normal, free subgroup.*

Then every homomorphism $\Lambda \to \mathrm{Out}(F_n)$ has finite image.

Proof. The only additional argument that is needed concerns the normal closure I of $[\psi]^N$ in $\mathrm{Out}(F_n)$, as considered in the second paragraph of the proof of Proposition 5.4. We must exclude the possibility that the intersection of I with the image of $\phi : \Lambda' \to \mathrm{Out}(F_n)$ is cyclic, generated by $[\psi]^m$ say. But if this were the case, $\langle [\psi]^m \rangle$ would be normal in $\phi(\Lambda')$. Since the normaliser in $\mathrm{Out}(F_n)$ of the subgroup generated by any fully irreducible element is virtually cyclic [6], it would follow that $\phi(\Lambda')$ itself was virtually cyclic, contradicting the fact that no finite-index subgroup of Λ' maps onto $\mathbb{Z}$. □

Remark 5.7. The class of groups that satisfy the hypotheses of Theorem 5.6 is closed under certain extension operations. For example, if $1 \to A \to \hat{\Lambda} \to \Lambda \to 1$ is a short exact sequence and, in the notation of Theorem 5.6, we suppose that

- every finite-index subgroup of A satisfies (2),
- every finite-index subgroup of $\hat{\Lambda}$ satisfies (1), and
- every finite-index subgroup of Λ satisfies both (1) and (2),

then an elementary argument shows that every finite-index subgroup of $\hat{\Lambda}$ satisfies (2). (Hence every homomorphism $\hat{\Lambda} \to \mathrm{Out}(F_n)$ has finite image.)

Let $\hat{G}$ be an arbitrary semisimple Lie group of real rank at least two, with centre $Z(\hat{G})$. Let $\hat{\Lambda} < \hat{G}$ be an irreducible lattice, let $A = \hat{\Lambda} \cap Z(\hat{G})$ and let $\Lambda = \hat{\Lambda}/A$. The abelianisation of any subgroup of finite index in $\hat{\Lambda}$ is finite ([58], page 333) and Λ is an irreducible lattice in the centreless semisimple Lie group $G = \hat{G}/Z(\hat{G})$. Thus the above remark allows us to remove from Theorem 5.2 the hypothesis that the centre of G is finite.

In Theorem 5.6, condition (2) is used only to exclude the possibility that a homomorphic image of Λ in $\mathrm{Out}(F_n)$ might contain a fully irreducible element. An alternative way of ruling out such images is to use bounded cohomology, as in [11]. We briefly review some notation. A map $f : \Lambda \to \mathbb{R}$ is a *quasi-homomorphism* if the function

$$(g, h) \mapsto |f(g) + f(h) - f(gh)|$$

is bounded on $\Lambda \times \Lambda$. Let $V(\Lambda)$ be the vector space of all quasi-homomorphisms from Λ to $\mathbb{R}$. Two natural subspaces of $V(\Lambda)$ are $B(\Lambda)$, the vector space of bounded maps from Λ to $\mathbb{R}$, and $\mathrm{Hom}(\Lambda; \mathbb{R})$, the vector space of genuine homomorphisms. Define $\widetilde{\mathrm{QH}}(\Lambda) = V(\Lambda)/(B(\Lambda) + \mathrm{Hom}(\Lambda; \mathbb{R}))$.

Theorem 5.8. *Let Λ be a group. Suppose that for every finite index subgroup $\Lambda' \subset \Lambda$ the second bounded cohomology of Λ' is finite dimensional and $\mathrm{Hom}(\Lambda, \mathbb{R}) = 0$. Then every homomorphism $\phi : \Lambda \to \mathrm{Out}(F_n)$ has finite image.*

Proof. Bestvina and Feighn [11] show that if $H < \mathrm{Out}(F_n)$ contains a fully irreducible automorphism then either H is virtually cyclic or $\widetilde{\mathrm{QH}}(H)$ is infinite dimensional. If $\mathrm{Hom}(\Lambda; \mathbb{R}) = 0$ then a surjective map $\Lambda \to H$ induces an injection $\widetilde{\mathrm{QH}}(H) \to \widetilde{\mathrm{QH}}(\Lambda)$. The vector space $\widetilde{\mathrm{QH}}(\Lambda)$ injects into the second bounded cohomology of Λ (see [62]). Therefore, for all finite index subgroups $\Lambda' \subset \Lambda$ and integers m the image of a homomorphism $\Lambda' \to \mathrm{Out}(F_m)$ cannot contain a fully irreducible automorphism. It follows from Corollary 5.5 and the arguments in Proposition 5.4 that $\phi(\Lambda)$ is finite. □

As the lattices of interest in this chapter have trivial second bounded cohomology [18], this alternative to Theorem 5.1 also implies Theorem 5.2.

5.2 Homomorphisms to $\mathrm{Out}(A_\Gamma)$

By Theorem 4.22 we know that $\overline{\mathrm{IA}}(A_\Gamma)$ is residually torsion-free nilpotent. Hence:

Proposition 5.9. *Suppose that* $\mathrm{Hom}(\Lambda, \mathbb{Z}) = 0$ *and* $f : \Lambda \to \overline{\mathrm{IA}}(A_\Gamma)$ *is a homomorphism. Then* f *is trivial.*

The overriding theme of this chapter is that we may build homomorphism rigidity results from weaker criteria by carefully studying a group's subgroups and quotients. This is very much the flavour of our main theorem:

Theorem 5.10. *Suppose that* G *is a subgroup of* $\mathrm{Out}(A_\Gamma)$ *generated by a subset* $T \subset \mathcal{S}_\Gamma$ *and* $d_{SL}(G) \leq m$*. Let* $F(\Gamma)$ *be the size of a maximal, discrete, full subgraph of* A_Γ*. Let* Λ *be a group. Suppose that for each finite index subgroup* $\Lambda' \leq \Lambda$*, we have:*

- *Every homomorphism* $\Lambda' \to \mathrm{SL}_m(\mathbb{Z})$ *has finite image,*
- *For all* $N \leq F(\Gamma)$*, every homomorphism* $\Lambda' \to \mathrm{Out}(F_N)$ *has finite image.*
- $\mathrm{Hom}(\Lambda', \mathbb{Z}) = 0$

Then every homomorphism $f : \Lambda \to G$ *has finite image.*

This section is dedicated to a proof of Theorem 5.10. We proceed by induction on the number of vertices in Γ. If Γ contains only one vertex, then $\mathrm{Out}(A_\Gamma) \cong \mathbb{Z}/2\mathbb{Z}$, so there is no work to do. As the conditions on Λ are also satisfied by finite index subgroups, we shall allow ourselves to pass to such subgroups without further comment.

Remark 5.11. If either $m \geq 2$ or $F(\Gamma) \geq 2$, then as there exist no homomorphisms from Λ' to $\mathrm{SL}_m(\mathbb{Z})$ or $\mathrm{Out}(F_{F(\Gamma)})$ with infinite image, it follows that $\mathrm{Hom}(\Lambda', \mathbb{Z}) = 0$ also. This is always the case when $G = \mathrm{Out}(A_\Gamma)$ and $|V(\Gamma)| \geq 2$. Hence the above statement of Theorem 5.10 is a strengthening of the version given in the introduction to this thesis.

Let $f : \Lambda \to G$ be such a homomorphism. There are three cases to consider: either the defining graph Γ is disconnected, or the defining graph is connected and $Z(A_\Gamma)$ is trivial, or the defining graph is connected and $Z(A_\Gamma)$ is non-trivial.

5.2.1 Γ is disconnected.

In this case $A_\Gamma \cong F_N *_{i=1}^{k} A_{\Gamma_i}$, where each Γ_i is a connected graph containing at least two vertices. Let $\Lambda' = f^{-1}(\mathrm{Out}^0(A_\Gamma))$. As $\mathrm{Out}^0(A_\Gamma)$ is finite index in $\mathrm{Out}(A_\Gamma)$, this means Λ' is finite index in Λ. We showed in Example 3.8 that for each Γ_i there is a restriction homomorphism:

$$R_i : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}^0(A_{\Gamma_i}).$$

By Lemma 3.12, $R_i(G)$ is generated by a subset $T_i \subset \mathcal{S}_{\Gamma_i}$, and $d_{SL}(R_i(G)) \leq d_{SL}(G)$. As Γ_i is a proper, full subgraph of Γ, we have $F(\Gamma_i) \leq F(\Gamma)$ and $|V(\Gamma_i)| < |V(\Gamma)|$. Hence, by induction $R_i f(\Lambda')$ is finite for each i, and there exists a finite index subgroup Λ_i of Λ' such that $R_i f(\Lambda_i)$ is trivial. We may also consider the exclusion homomorphism:

$$E : \mathrm{Out}(A_\Gamma) \to \mathrm{Out}(F_N).$$

As $N \leq F(\Gamma)$, the group $\ker(Ef)$ is a finite index subgroup of Λ. Let

$$\Lambda'' = \cap_{i=1}^{k} \Lambda_i \cap \ker(Ef).$$

As Λ'' is the intersection of a finite number of finite index subgroups, it is also finite index in Λ. We now study the action of Λ'' on $H_1(A_\Gamma)$. The transvection $[\rho_{ij}]$ belongs to $\mathrm{Out}^0(A_\Gamma)$ only if either v_i and v_j belong to the same connected component of Γ, or if v_i is an isolated vertex of Γ. Therefore the action of $\mathrm{Out}^0(A_\Gamma)$ on $H_1(A_\Gamma)$ has a block decomposition of the following form:

$$\begin{pmatrix} M_1 & 0 & \dots & 0 & * \\ 0 & M_2 & \dots & 0 & * \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & M_k & * \\ 0 & 0 & \dots & 0 & M_{k+1} \end{pmatrix},$$

where M_i (for $i \leq k$) corresponds to the action on A_{Γ_i}, and M_{k+1} corresponds to the action on F_N. However, as $R_i f(\Lambda'')$ is trivial for each i, and $Ef(\Lambda'')$ is trivial, the action of Λ'' on $H_1(A_\Gamma)$ is of the form:

$$\begin{pmatrix} I & 0 & \dots & 0 & * \\ 0 & I & \dots & 0 & * \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & I & * \\ 0 & 0 & \dots & 0 & I \end{pmatrix}.$$

This means there is a homomorphism from Λ'' to an abelian subgroup of $\mathrm{GL}_n(\mathbb{Z})$. As $\mathrm{Hom}(\Lambda'', \mathbb{Z}) = 0$, this homomorphism must be trivial. Hence $f(\Lambda'') \subset \overline{\mathrm{IA}}(A_\Gamma)$. By Proposition 5.9, this shows that $f(\Lambda'')$ is trivial. Hence $f(\Lambda)$ is finite.

5.2.2 Γ is connected and $Z(A_\Gamma)$ is trivial.

Let $\Lambda' = f^{-1}(\mathrm{Out}^0(A_\Gamma)) = f^{-1}(G^0)$. For each maximal vertex v of Γ we have a projection homomorphism:

$$P_v : \mathrm{Out}^0(A_\Gamma) \to \mathrm{Out}^0(A_{lk[v]}).$$

By Proposition 3.14, $P_v(G^0)$ is generated by a subset $T_v \subset \mathcal{S}_{lk[v]}$ and $d_{SL}(P_v(G^0)) \leq d_{SL}(G^0) = d_{SL}(G)$. As $lk[v]$ is a proper subgraph of Γ, we have $F(lk[v]) \leq F(\Gamma)$ and $|V(lk[v])| < |V(\Gamma)|$. Therefore by induction $P_v f(\Lambda')$ is finite. Let

$$\Lambda'' = \bigcap_{[v] \text{ max.}} \ker(P_v f).$$

Then Λ'' is a finite index subgroup of Λ and its image lies in the kernel of the amalgamated projection homomorphism:

$$P : \mathrm{Out}^0(A_\Gamma) \to \bigoplus_{[v] \text{ max.}} \mathrm{Out}^0(A_{lk[v]}).$$

By Theorem 3.9, $\ker P$ is a finitely generated free-abelian group. As $\mathrm{Hom}(\Lambda'', \mathbb{Z}) = 0$, a homomorphism from Λ'' to $\ker P$ must be trivial. Therefore $f(\Lambda'')$ is trivial and $f(\Lambda)$ is finite.

5.2.3 Γ is connected and $Z(A_\Gamma)$ is nontrivial.

Suppose that $Z(A_\Gamma)$ is nontrivial. Let $[v]$ be the unique maximal equivalence class in Γ, so that $Z(A_\Gamma) = A_{[v]}$. Let P_v and R_v be the restriction and projection maps given in Proposition 3.10 so that:

$$R_v : \mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_{[v]}) \cong \mathrm{GL}(A_{[v]})$$
$$P_v : \mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_{lk[v]})$$

If $[v]$ is not equal to the whole of Γ then by induction $P_v f(\Lambda)$ and $R_v f(\Lambda)$ are both finite, and there exists a finite index subgroup Λ' of Λ such that $f(\Lambda')$ is contained in the kernel of $P_v \times R_v$. By Proposition 3.10 this is a free abelian group Tr, so the image of Λ' in Tr is trivial, and $f(\Lambda)$ is finite.

Therefore we may assume that $\Gamma = [v]$. We now look at the $\sim_G$ equivalence classes in Γ. As A_Γ is free abelian, each $[v_i]_G \subset [v]$ is abelian, and as $d_{SL}(G) \leq m$, every such $[v_i]_G$ contains at most m vertices. Therefore matrices in (the image of) G^0 (under $\overline{\Phi}$) are of the form:

$$M = \begin{pmatrix} N_1 & 0 & \dots & 0 \\ * & N_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ * & * & \dots & N_{r'} \end{pmatrix},$$

where the $*$ in the $(i,j)^{th}$ block is possibly nonzero if $[v_{l_j}] \leq [v_{l_i}]$. For each i, we can look at the projection $M \mapsto N_i$ to obtain a homomorphism $g_i : SG^0 \to \mathrm{SL}_{l_{i+1}-l_i}(\mathbb{Z})$. As $l_{i+1} - l_i \leq m$, our hypothesis on Λ implies that $g_i f(f^{-1}(SG^0))$ is finite for all i. Let Λ_i be the kernel of each map $g_i f$ restricted to $f^{-1}(SG^0)$. Each Λ_i is finite index in Λ. Let $\Lambda' = \cap_{i=1}^{k} \Lambda_i$. Then matrices in the image of Λ' under f are of the form:

$$M = \begin{pmatrix} I & 0 & \dots & 0 \\ * & I & \dots & 0 \\ \dots & \dots & \dots & \dots \\ * & * & \dots & I \end{pmatrix},$$

therefore $f(\Lambda')$ is a torsion-free nilpotent group. As $\mathrm{Hom}(\Lambda', \mathbb{Z}) = 0$, this implies that $f(\Lambda')$ is trivial. Hence $f(\Lambda)$ is finite, and this finishes the final case of the theorem.

5.3 Consequences of Theorem 5.10

As there are no homomorphisms from a $\mathbb{Z}$–averse group to $\mathrm{SL}_2(\mathbb{Z})$ with infinite image (as $\mathrm{SL}_2(\mathbb{Z})$ is virtually free), combining our result for $\mathrm{Out}(F_n)$ (Theorem 5.1) with Theorem 5.10 we obtain:

Corollary 5.12. *If Λ is a $\mathbb{Z}$–averse group, and Γ is a finite graph that satisfies $d_{SL}(\mathrm{Out}(A_\Gamma)) \leq 2$, then every homomorphism $f : \Lambda \to \mathrm{Out}(A_\Gamma)$ has finite image.*

We would like to apply Theorem 5.10 to higher-rank lattices in Lie groups. For the remainder of this section Λ will be an irreducible lattice in a semisimple real Lie group G with real rank $\mathrm{rank}_{\mathbb{R}} G \geq 2$, finite centre, and no compact factors. As previously stated, such lattices are $\mathbb{Z}$-averse by Margulis' normal subgroup theorem, and the work of Margulis also lets us restrict the linear representations of such lattices:

Proposition 5.13. *If $\mathrm{rank}_{\mathbb{R}} G \geq k$ then every homomorphism $f : \Lambda \to \mathrm{SL}_k(\mathbb{Z})$ has finite image.*

To prove this we appeal to Margulis superrigidity. The following two theorems follow from [58], Chapter IX, Theorems 6.15 and 6.16 and the remarks in 6.17:

Theorem 5.14. *Let H be a real algebraic group and $f : \Lambda \to H$ a homomorphism. The Zariski closure of the image of f, denoted $\overline{f(\Lambda)}$, is semisimple.*

Theorem 5.15 (Margulis' Superrigidity Theorem)**.** *Let H be a connected, semisimple, real algebraic group and $f : \Lambda \to H$ a homomorphism. If*

- *H is adjoint (equivalently $Z(H) = 1$) and has no compact factors, and*
- *$f(\Lambda)$ is Zariski dense in H,*

then f extends uniquely to a continuous homomorphism $\tilde{f} : G \to H$. Furthermore, if $Z(G) = 1$ and $f(\Lambda)$ is nontrivial and discrete, then $\tilde{f}$ is an isomorphism.

We may combine these to prove Proposition 5.13:

Proof of Proposition 5.13. Let $f : \Lambda \to \mathrm{SL}_k(\mathbb{Z})$ be a homomorphism. By Theorem 5.14, the Zariski closure of the image $\overline{f(\Lambda)} \subset \mathrm{SL}_k(\mathbb{R})$ is semisimple. Also, $\overline{f(\Lambda)}$ has finitely many connected components — let $\overline{f(\Lambda)}_0$ be the connected component containing the identity. Decompose $\overline{f(\Lambda)}_0 = H_1 \times K$, where K is a maximal compact factor. Then H_1 is a connected semisimple real algebraic group with no compact factors. We look at the finite index subgroup $\Lambda_1 = f^{-1}(H_1)$ of Λ, so that $\overline{f(\Lambda_1)} = H_1$. As the centre of a subgroup of an algebraic group is contained in the centre of its Zariski closure, $f(Z(\Lambda_1)) \subset Z(H_1)$. This allows us to factor out centres in the groups involved. Let $G_2 = G/Z(G)$, $\Lambda_2 = \Lambda_1/Z(\Lambda_1) = \Lambda_1/(\Lambda_1 \cap Z(G))$ and $H_2 = H_1/Z(H_1)$. Then there is an induced map $f_2 : \Lambda_2 \to H_2$ satisfying the conditions of Theorem 5.15. Therefore if $f_2(\Lambda_2) \neq 1$ there is an isomorphism $\tilde{f}_2 : G_2 \to H_2$. However

$$\mathrm{rank}_{\mathbb{R}} G_2 = \mathrm{rank}_{\mathbb{R}} G \geq k$$
$$\mathrm{rank}_{\mathbb{R}} H_2 = \mathrm{rank}_{\mathbb{R}} H_1 \leq \mathrm{rank}_{\mathbb{R}} \mathrm{SL}_k(\mathbb{R}) = k - 1.$$

This contradicts the isomorphism between H_2 and G_2. Therefore $f_2(\Lambda_2) = 1$. As $Z(\Lambda_1)$ is finite, and Λ_1 is finite index in Λ, this show that the image of Λ under f is finite. □

Combining Proposition 5.13 with Theorems 5.1 and 5.10, this gives:

Theorem 5.16. *Let G be a real semisimple Lie group with finite centre, no compact factors, and $\mathrm{rank}_{\mathbb{R}} G \geq 2$. Let Λ be an irreducible lattice in G. If $\mathrm{rank}_{\mathbb{R}} G \geq d_{SL}(\mathrm{Out}(A_\Gamma))$, then every homomorphism $f : \Lambda \to \mathrm{Out}(A_\Gamma)$ has finite image.*

The following corollary justifies our definition of SL–dimension, and shows that you can't hide any larger copies of SL_n inside $\mathrm{Out}(A_\Gamma)$:

Corollary 5.17. *Let $k \geq 3$. Then $\mathrm{Out}(A_\Gamma)$ contains a subgroup isomorphic to $\mathrm{SL}_k(\mathbb{Z})$ if and only if $k \leq d_{SL}(\mathrm{Out}(A_\Gamma))$.*

This corollary notably excludes the case $k = 2$ (It also excludes $k = 1$, but $\mathrm{SL}_1(\mathbb{Z})$ is trivial!). As $\mathrm{SL}_2(\mathbb{Z})$ is virtually free, it is much easier to embed into other groups (in particular it is of index 2 in $\mathrm{Out}(F_2) \cong \mathrm{GL}_2(\mathbb{Z})$), so we cannot expect such a result to hold.

Part II

Folding and Outer Space

Chapter 6

Folding free-group automorphisms

The idea of controlling cancellation between words in a group can be traced along a line of thought spanning the twentieth century, from Nielsen's 1921 paper [63] showing that a finitely generated subgroup of a free group is free, through to the combinatorial and geometric methods in small cancellation theory now prevalent in the study of group actions on CAT(0) and hyperbolic complexes. In the free group, Nielsen's method of *reduction* was extended and given a topological flavour by Whitehead, who looked at sphere systems in connected sums of copies of $S^1 \times S^2$ [74]. Whitehead's idea of *peak reduction* was refined and recast in combinatorial language by Rapaport [65], Higgins and Lyndon [44], and McCool [59]. There is a good description of this viewpoint in Lyndon and Schupp's book on combinatorial group theory [55].

Peak reduction is very powerful. Given a finite set Y of elements in F_n, McCool [59] gives an algorithm to obtain finite presentations of $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$, the subgroups of $\mathrm{Aut}(F_n)$ that fix Y pointwise, and fix each element of Y up to conjugacy, respectively. Culler and Vogtmann's work on Outer Space shows that such subgroups also satisfy higher finiteness properties [24].

The generating sets for $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$ are built up out of *Whitehead Automorphisms*. These are automorphisms of two types. The first consists of the group W_n of automorphisms that permute and possibly invert elements of a fixed basis. So if F_n is generated by $X = \{x_1, \ldots, x_n\}$, then for each $\phi \in W_n$ there exists $\sigma \in S_n$ and $\epsilon_1, \ldots, \epsilon_n \in \{-1, 1\}$ such that $\phi(x_i) = x_{\sigma(i)}^{\epsilon_i}$. For the second type, we pick an element $a \in X \cup X^{-1}$ and for each basis element, we either pre-multiply by a, post-multiply by a^{-1}, or do both of these things. Traditionally this is defined by taking a subset $A \subset X \cup X^{-1}$ such that $a \in A$ and $a^{-1} \notin A$, and defining $(A, a) \in \mathrm{Aut}(F_n)$ by

$$(A, a)(x_j) = \begin{cases} x_j & \text{if } x_j = a^{\pm 1} \\ a^{\alpha_j} x_j a^{-\beta_j} & \text{if } x_j \neq a^{\pm 1} \end{cases},$$

where $\alpha_j = \chi_A(x_j)$ and $\beta_j = \chi_A(x_j^{-1})$ (here χ_A is the the indicator function of A, so that $\chi_A(y) = 1$ if $y \in A$ and otherwise $\chi_A(y) = 0$).

Beyond the work of Nielsen and Whitehead, a third approach to reduction in free groups comes from Stallings [69], who cast Nielsen reduction in terms of folds on graphs. Since 'Topology of finite graphs' appeared in 1983, folding has become a key tool in geometric group theory, notably in its applications to graphs of groups and their deformation spaces [41, 46, 36], and to the dynamics of free group automorphisms (and endomorphisms) [7, 35, 30]. In this chapter we give an account of how folding gives an algorithm to decompose an automorphism as a product of Whitehead automorphisms. This algorithm is hinted at by Stallings [69, Comment 8.2], and will be familiar to many authors who have used his techniques, but no explicit account appears in the literature.

The chief advantage of folding over peak reduction is the ease of application: folding a graph is less complicated than searching through a list of possible Whitehead automorphisms (a list that grows exponentially with n). Moreover, folding gives an intuitive, pictorial way of looking at the decomposition. The proofs here are geared towards making it easy to produce such calculations by hand or with a computer.

Finite generation of some subgroups of the form $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$ also follows very naturally from this description. In Section 6.3 we show that if Y is a subset of our preferred basis for F_n then the folding algorithm implies that $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$ are generated by the Whitehead automorphisms that lie in $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$, respectively (see Figure 6.4 for a quick idea of how this is done.) We used this result in Chapter 4 to show that when Y is a subset of a basis the intersection of $\mathrm{Fix}_c(Y)$ with IA_n is finitely generated.

6.1 Graphs, folds, and associated automorphisms

The fundamental group of a graph gives a pleasant pictorial description of the free group, and can be thought of as both a topological and a combinatorial construction. We will focus on the latter approach, borrowing most of our notation from Serre's book [67]. Proofs in this first section will either be sketched or omitted.

6.1.1 The fundamental group of a graph

Definition 6.1. A *graph* G consists of a tuple $(\mathrm{E}G, \mathrm{V}G, inv, \iota, \tau)$ where $\mathrm{E}G$ and $\mathrm{V}G$ are sets and $inv : \mathrm{E}G \to \mathrm{E}G$, $\iota, \tau : \mathrm{E}G \to \mathrm{V}G$ are maps which satisfy

$$\begin{aligned} inv(e) &\neq e \\ inv(inv(e)) &= e \\ \iota(inv(e)) &= \tau(e). \end{aligned}$$

$\mathrm{E}G$ is said to be the *edge set* of G and $\mathrm{V}G$ the *vertex set* of G. For an edge $e \in \mathrm{E}G$ we write $inv(e) = \bar{e}$, and say that $\iota(e), \tau(e)$ are the *initial* and *terminal* vertices of e respectively.

A *path* p in G is either a sequence of edges $e_1, \dots e_k$ such that $\iota(e_{i+1}) = \tau(e_i)$, or a single vertex v. Let $\mathrm{P}G$ be the set all paths. The functions inv, ι and τ extend to $\mathrm{P}G$; in the case where p is a sequence of edges we define $\iota(p) = \iota(e_1)$, $\tau(p) = \tau(e_k)$ and $\bar{p} = \bar{e}_k, \dots \bar{e}_1$, and evaluate these functions at v if p is a single vertex v. We say that G is connected if for any two vertices v, w there exists a path p such that $\iota(p) = v$ and $\tau(p) = w$. If $\tau(p_1) = \iota(p_2)$ we define $p_1.p_2$ to be the concatenation of the two sequences. We define an equivalence relation $\sim$ on $\mathrm{P}G$ by saying two paths p_1, p_2 are equivalent if and only if one can be obtained from the other by insertion and deletion of a sequence of pairs of edges of the form $(e, \bar{e})$. We say that a path p is *reduced* if there are no consecutive edges of the form $(e, \bar{e})$ in p.

Proposition 6.2. *Every element of* $\mathrm{P}G/\sim$ *is represented by a unique reduced path. For* $p \in \mathrm{P}G$ *we let* $[p]$ *denote the reduced path in the equivalence class of* p*.*

The set of reduced paths that begin and end at a vertex v in G form a group that we shall denote $\pi_1(G, v)$, the *fundamental group of* G *based at* v. Multiplication is defined as follows — if p, q are reduced paths, then $p \cdot q = [p.q]$. The identity element is the path consisting of the single vertex v, and the inverse of a reduced path p is the path $\bar{p}$. A path p_{vw} connecting vertices v and w in G induces an isomorphism $[p] \mapsto [p_{vw}.p.\overline{p_{vw}}]$ between $\pi_1(G, w)$ and $\pi_1(G, v)$. A *subgraph* of G is given by subsets of $\mathrm{E}G$ and $\mathrm{V}G$ which are invariant under the operations inv and ι. A connected graph T is called a *tree* if $\pi_1(T, v)$ is trivial for a (equivalently, any) vertex v of T. We say that T is a *maximal tree* in a connected graph G if T is a subgraph of G, T is a tree, and the vertex set of T is $\mathrm{V}G$. Such a tree always exists. Given a base point b in a connected graph G and a maximal tree T, there exists a unique reduced path p_v from b to v. An *orientation* of a subgraph $G' \subset G$ is a set $\mathcal{O}$ that contains

exactly one element of $\{e, \bar{e}\}$ for each element of G'. An *ordered orientation* of G' is an orientation $\mathcal{O}$ of G' with an enumeration of the set $\mathcal{O}$.

Proposition 6.3. *Let T be a maximal tree in a connected graph G with chosen base point b. Then we can define an orientation $\mathcal{O}(T, b)$ of T by saying that $e \in \mathcal{O}(T, b)$ if an only if e occurs as an edge in a path p_v for some v.*

Geometrically, this is the orientation one obtains by drawing arrows on edges 'pointing away from b.' The main use of maximal trees and orientations will be to give a basis for $\pi_1(G, b)$. The following theorem is key to this chapter, so we will give it a name:

Basis Theorem. *Let T be a maximal tree in a connected graph G with chosen base point b. Let $\{e_1, \ldots, e_n\}$ be an ordered orientation of $G \setminus T$. Let*

$$l_i = p_{\iota(e_i)} e_i \overline{p_{\tau(e_i)}}.$$

$\pi_1(G, b)$ is freely generated by $l_1, \ldots, l_n$. Given any loop l based at b, we may write $[l]$ as a product of the generators as follows: remove the edges of l contained in T to obtain a sequence $e_{i_1}^{\epsilon_1}, \ldots, e_{i_k}^{\epsilon_k}$, where $i_j \in \{1, \ldots, n\}$ and $\epsilon_j \in \{1, -1\}$. Then

$$[l] = [l_{i_1}^{\epsilon_1} \cdots l_{i_k}^{\epsilon_k}].$$

Thus, once we have a maximal tree and an ordered orientation of the edges outside of this tree, the Basis Theorem gives us a method for constructing an ordered free generating set of $\pi_1(G, b)$. It also tells us how to write any element of $\pi_1(G, b)$ as a product of these generators. We may determine when a subgraph of G is a maximal tree as follows:

Lemma 6.4. *Let G be a connected graph, T a subgraph of G and b a vertex of G. Then T is a maximal tree if and only if:*

1. *T contains $2(|\mathrm{V}G| - 1)$ edges.*
2. *For each vertex v of G there exists a reduced path p_v from b to v in T.*

6.1.2 Folding maps of graphs

From now on we shall assume that all graphs are connected. A *map of graphs* $f : G \to \Delta$ is a map that takes edges to edges, vertices to vertices and satisfies $f(\bar{e}) = \overline{f(e)}$ and $f(\iota(e)) = \iota(f(e))$ for every edge in G. For a vertex v of G the map f induces a group homomorphism $f_* : \pi_1(G, v) \to \pi_1(\Delta, f(v))$. If f_* is an isomorphism for some (equivalently, any) choice of vertex of G, we say that f is a *homotopy equivalence*. If f is bijective on EG and VG then we say f is a *graph isomorphism*. Stallings' definition of the star of a vertex differs from the one we have used in previous chapters, and we will need his definition in the work that follows:

$$St(v, G) = \{e \in \mathrm{E}G : \iota(e) = v\}.$$

Henceforth we shall only use this second definition, and hope the difference in notation clarifies this abuse of terminology. If f is a map of graphs then for each vertex v in G we obtain a map $f_v : St(v, G) \to St(f(v), \Delta)$ by restricting f to the edges in $St(v, G)$. We say that f is an *immersion* if f_v is injective for each vertex of G, and we say that f is a *covering* if f_v is bijective for each vertex of G. If for some vertex v the map f_v is not injective, Stallings [69] introduced a method called *folding* for improving the map f: take edges e_1 and e_2 in $St(v, G)$ such that $f_v(e_1) = f_v(e_2)$ and form a quotient graph G' by identifying the pairs $\{e_1, e_2\}$, $\{\bar{e}_1, \bar{e}_2\}$ and $\{\tau(e_1), \tau(e_2)\}$ in G to form quotient edges e', $\bar{e}'$ and a quotient vertex v'.

There are then induced maps $q : G \to G'$ and $f' : G' \to \Delta$ such that $f' \cdot q = f$. We call this process *folding f along e_1 and e_2*. If v is a vertex in G the map $q_* : \pi_1(G, v) \to \pi_1(G', q(v))$ is surjective and $f_*(\pi_1(G, v)) = f'_*(\pi_1(G', q(v)))$.

Stallings' Folding Theorem ([69])**.** *Let $f : G \to \Delta$ be a map of graphs, and suppose that G is finite and connected.*

1. *If f is an immersion then f_* is injective.*

2. *If f is not an immersion, there exists a finite sequence of folds $G = G_0 \to G_1 \to G_2 \ldots \to G_n$ and an immersion $G_n \to \Delta$ such that the composition of the above maps is equal to f.*

Sketch proof. If f is an immersion, then reduced paths are sent to reduced paths of the same length. Hence f_* is injective. For the second part, we iterate the folding described above to obtain a sequence of graphs with the required properties. This process must eventually end as G is finite, and folding reduces the number of edges in a graph. □

There are four different types of fold that can occur, which we illustrate in Figure 6.1. If f_* is injective only folds of type 1 or 2 occur. In case 3 the loop $e_1, \bar{e}_2$ is non-trivial in the original graph, but mapped to the trivial element in the quotient, and in 4 the loops e_1 and e_2 are distinct but mapped to homotopic loops in the quotient.

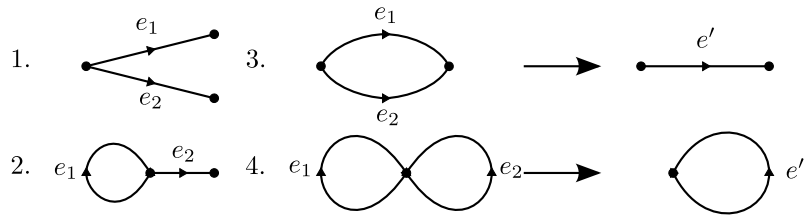

Figure 6.1: Possible folds of a graph

6.1.3 Branded graphs and their associated automorphisms

We may identify F_n with the fundamental group of a fixed graph, R_n:

Definition 6.5. The *rose with n petals*, R_n is defined be the graph with edge set $\mathrm{ER}_n = \{x_1, \ldots, x_n\} \cup \{\bar{x}_1, \ldots, \bar{x}_n\}$, a single vertex b_R with $\iota(e) = \tau(e) = b_R$ for each edge e in ER_n and *inv* taking $x_i \to \bar{x}_i$. We identify F_n with $\pi_1(\mathrm{R}_n, b_R)$ by the map taking each generator x_i of F_n to the path consisting of the single edge with the same name.

Suppose that $f : G \to R_n$ is a homotopy equivalence. Let T be a maximal tree of G, let b be a vertex of G, and let $\{e_1, \ldots, e_n\}$ be an ordered orientation of $G \setminus T$. We call the tuple $\mathcal{G} = (G, f, b, \{e_1, \ldots, e_n\})$ a *branded graph*. If we are given G and f, then we say that a choice of a base point b and an ordered orientation of a complement of a maximal tree in G is a *branding* of (G, f). As b and $\{e_1, \ldots, e_n\}$ determine a choice of basis of $\pi_1(G, b)$, every branded graph has an *associated automorphism* of F_n defined by:

$$\phi_{\mathcal{G}}(x_i) = f_*(l_i),$$

where l_i is the loop $p_{\iota(e_i)}.e_i.\overline{p_{\tau(e_i)}}$ described in Proposition 6.1.1. Topologically, the choice of basepoint b and edges $\{e_1, \ldots, e_n\}$ determines a homotopy equivalence $(R_n, b_R) \xrightarrow{h_{\mathcal{G}}} (G, b)$ given by mapping x_i over l_i. Then $\phi_{\mathcal{G}}$ is the automorphism $f_* h_{\mathcal{G}*}$:

$$\begin{array}{ccc} & & \pi_1(G,b) \\ & \overset{h_{\mathcal{G}*}}{\nearrow} & \downarrow f_* \\ \pi_1(R_n, b_R) & \xrightarrow{\phi_{\mathcal{G}}} & \pi_1(R_n, b_R) \end{array}$$

Example 6.6. If $\phi \in \mathrm{Aut}(F_n)$ and $\phi(x_i) = w_i$ for all i, let G be the graph that is topologically a rose, with the ith loop subdivided into $|w_i|$ edges. Let $f : G \to R_n$ be the homotopy equivalence given by mapping the ith loop to the path given by w_i in R_n. Let b be the vertex in the centre of the rose, and for each i choose an edge e_i in the ith loop oriented in the direction of the word w_i. If $\mathcal{G} = (G, f, b, \{e_1, \ldots, e_n\})$ then l_i is the ith loop, so that $\phi_{\mathcal{G}} = \phi$.

Of particular importance is the situation when f is an immersion:

Lemma 6.7. *Let $f : G \to R_n$ be a homotopy equivalence and an immersion. Then f is an isomorphism, and for any branding $\mathcal{G}$ associated to G, f, we have $\phi_{\mathcal{G}} \in W_n$. ($\phi_{\mathcal{G}}$ acts on the basis of F_n by permuting and possibly inverting basis elements.)*

Proof. If f is an isomorphism of graphs, then $\phi_{\mathcal{G}} \in W_n$ for any branding – each e_i forms a loop in G, so there exists $\sigma \in S_n$ such that each e_i is sent to $x_{\sigma(i)}^{\epsilon_i}$ for some $\epsilon_i \in \{-1, 1\}$ that depends on i. It remains to show that if f is an immersion and a homotopy equivalence then f is an isomorphism. One way to see this is as follows: if f is an immersion, there we can add extra edges to G to build a graph G' containing G and a map $f' : G' \to R_n$ which covers R_n (e.g. [69], Theorem 6.1). However, f_* is surjective, so this cover is degree 1, and $G' = G \cong R_n$. □

6.2 An algorithm

The algorithm for writing an arbitrary element of $\phi \in \mathrm{Aut}(F_n)$ as a product of Whitehead automorphisms proceeds as follows. One first picks a branded graph $\mathcal{G}$ such that $\phi_{\mathcal{G}} = \phi$; to be definite we take the one described in Example 6.6. If f is not an immersion then we may fold f and, since f is a homotopy equivalence, the fold is one of the two types shown in Figure 6.2.

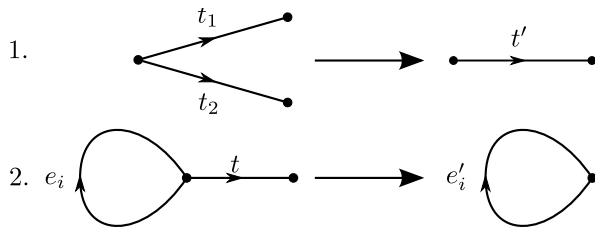

Figure 6.2: Possible folds when f is a homotopy equivalence

If, as the labelling in Figure 6.2 suggests, the folding edges with two distinct endpoints are in the maximal tree T, then we obtain a branding $\mathcal{G}'$ of the folded

graph G' by taking the images of $\{e_1, \ldots, e_n\}$ and b in G'. If the fold is of the first type, then the associated automorphisms $\phi_{\mathcal{G}}$ and $\phi_{\mathcal{G}'}$ are identical (Proposition 6.8). If the fold is of the second type, they differ by a Whitehead automorphism of the form (A, a) that may be read off from the structure of T (Proposition 6.9).

It may happen that one of t_1, t_2, or t does not lie in T. In this case we can swap this edge with an edge already lying in T (see Section 6.2.2), to obtain a new tree T' and a new branding $\mathcal{G}'$. Again, $\phi_{\mathcal{G}}$ and $\phi_{\mathcal{G}'}$ differ by a Whitehead automorphism of the form (A, a) that may be read off from the swap (Proposition 6.10). After at most two such swaps, we can ensure that the folding edges with two distinct endpoints lie in T, and proceed as above (Remark 6.11).

By Stallings' folding theorem, we obtain a finite sequence $\mathcal{G} = \mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_k$ of branded graphs $\mathcal{G}_j = (G_j, f_j, b_j, \{e_1^j, \ldots, e_n^j\})$ such that each f_j is a homotopy equivalence, and f_k is an immersion. By Lemma 6.7 we know that f_k is an isomorphism and $\phi_{\mathcal{G}_k} \in W_n$. Then:

$$\phi = \phi_{\mathcal{G}_1} = \phi_{\mathcal{G}_k}(\phi_{\mathcal{G}_k}^{-1}\phi_{\mathcal{G}_{k-1}}) \cdots (\phi_{\mathcal{G}_3}^{-1}\phi_{\mathcal{G}_2})(\phi_{\mathcal{G}_2}^{-1}\phi_{\mathcal{G}_1})$$

is a decomposition of ϕ as a product of Whitehead automorphisms. We assume that $\mathrm{Aut}(F_n)$ acts on F_n on the left, so that in the above decomposition we apply $\phi_{\mathcal{G}_2}^{-1}\phi_{\mathcal{G}_1}$ first, then $\phi_{\mathcal{G}_3}^{-1}\phi_{\mathcal{G}_2}$, etc.

If we count one step as a (possibly trivial) tree substitution, followed by a fold, then each step reduces the number of combinatorial edges of the graph by two (an e and an $\bar{e}$). If the initial graph has $2m$ edges, then as R_n has $2n$ edges we will obtain a decomposition of ϕ after $m - n$ steps. If $\phi(x_i) = w_i$ and we start with the graph given in Example 6.6, then our algorithm will terminate after $(\sum_{i=1}^{n} |w_i|) - n$ steps.

We will now give a detailed description of the process of folding and exchanging edges in maximal trees.

6.2.1 Folding edges contained in T

Suppose $q : G \to G'$ is a fold from Figure 6.2. The map f factors through q, inducing a homotopy equivalence $f' : G' \to R_n$ such that $f = f' \cdot q$. Let $b', e_1', \ldots, e_n'$ be the images of $b, e_1, \ldots, e_n$ respectively under q. Then $\mathcal{G}' = (G', f', b', \{e_1', \ldots, e_n'\})$ is a branding of G'. The only thing to check is that $T' = G' \setminus \{e_1', \overline{e_1'}, e_2', \overline{e_2'}, \ldots, e_n, \overline{e_n'}\}$ is a maximal tree of G'. The subgraph T' contains $2(|VG'| - 1)$ edges as a fold of type 1 or 2 reduces the number of vertices in a graph by one, and the number of edges in a graph by two. Let v' be a vertex of G'. Take a vertex v of G such that $q(v) = v'$. In the case of a type 1 fold, the path $[q(p_v)]$ is a reduced path from b' to v' lying in T',

and in the case of a type 2 fold, if we remove all occurrences of e'_i from $q(p_v)$, then reduce, we obtain a path from b' to v' lying in T'. Hence by Lemma 6.4, we know that T' is a maximal tree of G'. Let p_v be the unique reduced path from b to v in T and let $l_1, \ldots, l_n$ be the generators of $\pi_1(G, b)$ given by b and $\{e_1, \ldots, e_n\}$. Let $l'_1, \ldots, l'_n$ be the generators of $\pi_1(G', b')$ given by b' and $\{e'_1, \ldots, e'_n\}$. As $f_* = f'_* q_*$, we may find the difference between the automorphisms $\phi_{\mathcal{G}}$ and $\phi_{\mathcal{G}'}$ by finding a decomposition of $q_*(l_i)$ in terms of the l'_i.

Proposition 6.8. *Suppose that q is a fold of type 1, where the folded edges t_1 and t_2 lie in T. Then $\phi_{\mathcal{G}} = \phi_{\mathcal{G}'}$.*

Proof. For each path l_i, the only edge $q(l_i)$ crosses that does not lie in T' is e'_i. By the Basis Theorem, we have $q_*(l_i) = l'_i$. Hence

$$\phi_{\mathcal{G}'}(x_i) = f'_*(l'_i) = f'_*(q_*(l_i)) = f_*(l_i) = \phi_{\mathcal{G}}(x_i). \qquad \square$$

Proposition 6.9. *Suppose that q is a fold of type 2, where an edge t in T is identified with a loop e_i in $G \setminus T$ (and $\bar{t}$ is identified with $\bar{e}_i$). Let $\mathcal{O}(T, b)$ be the orientation of T given by Proposition 6.3. Let*

$$\epsilon = \begin{cases} 1 & \text{if } t \in \mathcal{O}(T, b) \\ -1 & \text{if } \bar{t} \in \mathcal{O}(T, b). \end{cases}$$

Define $A \subset X \cup X^{-1}$ such that $x_i^{\epsilon} \in A$, $x_i^{-\epsilon} \notin A$ and

$$x_j \in A \Leftrightarrow p_{\iota(e_i)} \text{ crosses } t \text{ or } \bar{t}$$
$$x_j^{-1} \in A \Leftrightarrow p_{\tau(e_i)} \text{ crosses } t \text{ or } \bar{t}.$$

Then $\phi_{\mathcal{G}} = \phi_{\mathcal{G}'} \cdot (A, x_i^{\epsilon})$.

Proof. We prove this result for $t \in \mathcal{O}(T, b)$, the other case being similar. If $t \in \mathcal{O}(T, b)$, then t may appear at most once in a path p_v, however $\bar{t}$ may not. Note that:

$$\begin{aligned} q(l_j) &= q(p_{\iota(e_j)} e_j \overline{p_{\tau(e_j)}}) \\ &= q(p_{\iota(e_j)}).e'_j.q(\overline{p_{\tau(e_j)}}). \end{aligned}$$

Removing all the edges of $q(l_j)$ not in T' leaves a sequence of the form (e'_j), (e'_i, e'_j), $(e'_i, e'_j, \overline{e'_i})$ or $(e'_j, \overline{e'_i})$, where e'_i proceeds e'_j if and only if t lies in $p_{\iota(e_j)}$, and $\overline{e'_i}$ follows e'_j if and only if t lies in $p_{\tau(e_j)}$. As e_i is a loop, $p_{\iota(e_i)} = p_{\tau(e_i)}$, and therefore this sequence is either (e'_i) or $(e'_i, e'_i, \overline{e'_i})$. Therefore $q_*(l_i) = l'_i$ and it follows that $\phi_{\mathcal{G}}(x_i) = \phi_{\mathcal{G}'}(x_i)$. If

$j \neq i$ then by the Basis Theorem we have $[q(l_j)] = [l_i']^{\alpha_j}.[l_j'].[l_i']^{-\beta_j}$ where $\alpha_j = \chi_A(x_j)$ and $\beta_j = \chi_A(x_j^{-1})$. Hence

$$\begin{aligned}
\phi_{\mathcal{G}'} \cdot (A, x_i)(x_j) &= \phi_{\mathcal{G}'}(x_i^{\alpha_j} x_j x_i^{-\beta_j}) \\
&= f'_*([l_i']^{\alpha_j}.[l_j'].[l_i']^{-\beta_j}) \\
&= f'_* q_*(l_j) \\
&= f_*(l_j) \\
&= \phi_{\mathcal{G}}(x_j)
\end{aligned}$$

□

6.2.2 Swapping edges into a tree

Suppose that we would like to fold in a branded graph as in Figure 2, but an edge t_1, t_2 or t lies outside the maximal tree. Then either this edge or its inverse is equal to e_i for some i. The edge e_i has distinct endpoints, so $p_{\iota(e_i)} \neq p_{\tau(e_i)}$. Let a be the shared initial segment of these paths. Either $p_{\iota(e_i)} \smallsetminus a$ or $p_{\tau(e_i)} \smallsetminus a$ is non-empty. Choose an edge e_i' such that either $e_i' \in p_{\iota(e_i)} \smallsetminus a$ or $\overline{e_i'} \in p_{\tau(e_i)} \smallsetminus a$. By a similar approach to the one used in Section 6.2.1 one can check that $T' = G \smallsetminus \{e_1, \overline{e_1}, e_2, \overline{e_2}, \ldots, e_i', \overline{e_i'}, \ldots, e_n, \overline{e_n}\}$ is a maximal tree of G, so that $\mathcal{G}' = (G, f, b, \{e_1, \ldots, e_i', \ldots, e_n\})$ is a branding of G.

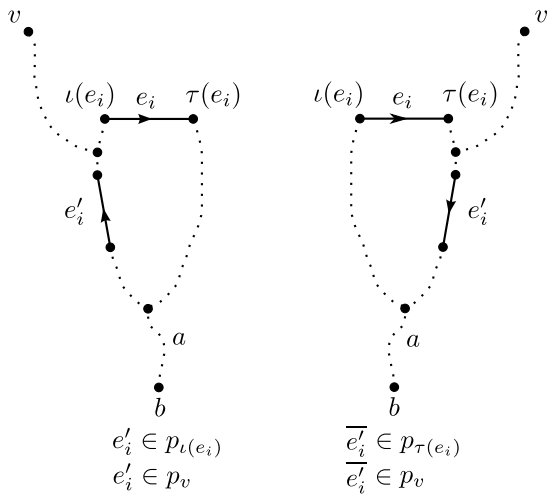

Figure 6.3: Changing maximal trees.

Proposition 6.10. *Let $\mathcal{G}'$ be the branding obtained by swapping an edge as described above and depicted in Figure 6.3. Define*

$$\epsilon = \begin{cases} 1 & \text{if } e_i' \in p_{\iota(e_i)} \\ -1 & \text{if } e_i' \in \overline{p_{\tau(e_i)}}. \end{cases}$$

Now define $A \subset X \cup X^{-1}$ *to be such that* $x_i^{\epsilon} \in A$, $x_i^{-\epsilon} \notin A$ *and*

$$x_j \in A \Leftrightarrow p_{\iota(e_j)} \text{ crosses } e_i' \text{ or } \overline{e_i'}$$
$$x_j^{-1} \in A \Leftrightarrow p_{\tau(e_j)} \text{ crosses } e_i' \text{ or } \overline{e_i'}.$$

Then $\phi_{\mathcal{G}} = \phi_{\mathcal{G}'} \cdot (A, x_i^{\epsilon})$.

Proof. The proof is analogous to the proof of Proposition 6.9. Let $l_1', \ldots, l_n'$ be the new basis of $\pi_1(G, b)$ given by b and $\{e_1, \ldots, e_i', \ldots, e_n\}$. By reading off the edges that lie outside of T' crossed by the paths l_j we find that $l_i = l_i'$ and for $j \neq i$ we have $l_j = [l_i'^{\epsilon\alpha_j}.l_j'.l_i'^{-\epsilon\beta_j}]$, where $\alpha_j = \chi_A(x_j)$ and $\beta_j = \chi_A(x_j)$. It follows that $\phi_{\mathcal{G}} = \phi_{\mathcal{G}'} \cdot (A, x_i^{\epsilon})$. □

Remark 6.11. If we are looking at a fold of the first type in Figure 6.2, we would like both edges t_1 and t_2 to lie in the maximal tree T. If we move one edge t_1 into the maximal tree through the method described above, the edge t_2 may still lie outside the maximal tree. We would like to add it in without removing t_1. We are only unable to do this if t_1 and $\bar{t}_1$ are the only elements of $p_{\iota(t_2)} \setminus a$ and $p_{\tau(t_2)} \setminus a$. This means that $\{p_{\iota(t_2)}, p_{\tau(t_2)}\}$ is either the set $\{a, a.t_1\}$ or the set $\{a, a.\bar{t}_1\}$. These cases would contradict either $\iota(t_1) = \iota(t_2)$ or $\tau(t_1) \neq \tau(t_2)$.

6.3 Fixing generators

The algorithm described in Section 6.2 may be applied to find generating sets of subgroups of $\mathrm{Aut}(F_n)$. Let ρ_{ij}, K_{ij}, and S_i be the elements of $\mathrm{Aut}(F_n)$ defined in the same way as in previous chapters:

$$\rho_{ij}(x_k) = \begin{cases} x_i x_j & \text{if } k = i \\ x_k & \text{if } k \neq i \end{cases},$$
$$K_{ij}(x_k) = \begin{cases} x_i x_k x_i^{-1} & \text{if } k = i \\ x_k & \text{if } k \neq i \end{cases},$$
$$S_i(x_k) = \begin{cases} x_i^{-1} & \text{if } k = i \\ x_k & \text{if } k \neq i \end{cases}.$$

Any Whitehead automorphism can be written as a product of the above elements. Let $\mathrm{Fix}(\{x_{m+1}, \ldots, x_n\})$ be the subgroup of $\mathrm{Aut}(F_n)$ consisting of elements that fix $x_{m+1}, \ldots, x_n$ pointwise, and let $\mathrm{Fix}_c(\{x_{m+1}, \ldots, x_n\})$ be the subgroup of $\mathrm{Aut}(F_n)$ that takes each element of the set $\{x_{m+1}, \ldots, x_n\}$ to a conjugate of itself.

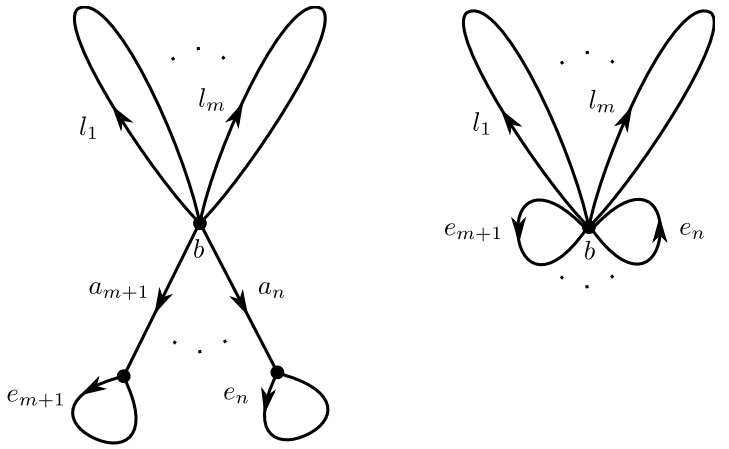

Figure 6.4: The Construction of G in Theorem 6.12

Theorem 6.12. *Let $Y = \{x_{m+1}, \ldots, x_n\}$ be a subset of our preferred basis for F_n. The subgroups $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$ are generated by the Whitehead automorphisms that lie in $\mathrm{Fix}(Y)$ and $\mathrm{Fix}_c(Y)$ respectively. In terms of Nielsen automorphisms, generating sets for $\mathrm{Fix}(Y)$, $\mathrm{Fix}_c(Y)$ are given by*

$$\mathcal{A}_m = \{S_i, \rho_{ij} : 1 \leq i \leq m, 1 \leq j \leq n\},$$
$$\mathcal{B}_m = \mathcal{A}_m \cup \{K_{ij} : m+1 \leq i \leq n, 1 \leq j \leq n\},$$

respectively.

Proof. Let $\phi \in \mathrm{Fix}_c(Y)$ and let G be a graph constructed as follows: take a single vertex b and a loop l_j consisting of $|\phi(x_j)|$ edges about b for $x_1, \ldots, x_m$. When $j > m$ we have $\phi(x_j) = w_j x_j w_j^{-1}$ – add a path a_j containing $|w_j|$ edges to b for each j, and attach an edge loop e_j to the end of each of these paths. We can then define $f : G \to \mathrm{R}_n$ so that when $j \leq m$ the loop l_j is mapped to edge path $\phi(x_j)$ and when $j > m$ each path a_j is sent to the edge path w_j and the edge loops $e_{m+1}, \ldots, e_n$ to the edges $x_{m+1}, \ldots, x_n$ respectively (see Figure 6.4). For $j \leq m$ pick an edge e_j in each l_j oriented in the direction of the word $\phi(x_j)$ being spelt out by l_j. Then ϕ is the automorphism associated to the branded graph $\mathcal{G} = (G, f, b, \{e_1, \ldots, e_n\})$. We apply the algorithm described in Section 6.2 to write ϕ as a product of Whitehead automorphisms. Let $br = \mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_k$ be the sequence of branded graphs $\mathcal{G}_j = (G_j, f_j, b_j, \{e_1^j, \ldots, e_n^j\})$ obtained. Let e_i be an edge in $\{e_{m+1}, \ldots, e_n\}$. Then each e_i^j is a loop, and will never be swapped into a maximal tree, so $e_i^j \to e_i^{j+1}$ at each step in the folding process. As $\iota(e_i^j) = \tau(e_i^j)$, we have $p_{\iota(e_i^j)} = p_{\tau(e_i^j)}$ at each step,

so by Propositions 6.9 and 6.10 the only Whitehead automorphisms of the form (A, a) that occur in the decomposition of $\phi_{\mathcal{G}}$ take x_j to a conjugate. Also, $\phi_{\mathcal{G}_k} \in W_n$ fixes $x_{m+1}, \ldots, x_n$. Hence the Whitehead automorphisms that lie in $\mathrm{Fix}_c(Y)$ generate $\mathrm{Fix}_c(Y)$. In the case where $x_{m+1}, \ldots, x_n$ are completely fixed by ϕ, the loops $e^j_{m+1}, \ldots, e^j_n$ are at the basepoint of each graph in the folding process, therefore Propositions 6.10 and 6.9 tell us every Whitehead automorphism that occurs in the decomposition of ϕ will fix $x_{m+1}, \ldots, x_n$. To obtain the generating sets in terms of Nielsen automorphisms one checks that each Whitehead automorphism that lies in $\mathrm{Fix}(Y)$ may be written as a product of elements of $\mathcal{A}_m$, and that each Whitehead automorphism that lies in $\mathrm{Fix}_c(Y)$ may be written as a product of elements that lie in $\mathcal{B}_m$. □

Chapter 7

Displacement functions on Outer Space

The techniques in the first part of this thesis are mostly combinatorial, and in this sense could be deemed a little dated from the modern viewpoint of geometric group theory. In Chapter 6, with Stallings' methods, we moved into late twentieth century, more topological, techniques. In this final chapter we shall look at metric properties of outer space, an area of research only a few years old. Outer space (denoted CV_n) is a finite dimensional, contractible topological space that admits a proper action of $\mathrm{Out}(F_n)$ with finite stabilisers. Previous results have mostly concerned topological properties of Outer space, however a series of recent papers [38, 37, 2, 1, 8] have turned to study its metric properties. Francaviglia and Martino [38] introduced an $\mathrm{Out}(F_n)$–invariant *non-symmetric Lipschitz metric* $d : \mathrm{CV}_n \times \mathrm{CV}_n \to \mathbb{R}_{\geq 0}$ on Outer space. The map d gets its name as it is positive definite and satisfies the triangle inequality, but is not symmetric. This metric is a new and useful way to study $\mathrm{Out}(F_n)$. Notably, Bestvina and Feighn [12] use the Lipschitz metric as an important part of their proof that the *complex of free factors* is Gromov hyperbolic. Bestvina and Feighn's previous results [11] regarding the action of $\mathrm{Out}(F_n)$ on certain hyperbolic complexes were an essential part of our proofs in Chapter 5.2 regarding homomorphisms to $\mathrm{Out}(F_n)$. These are 'custom builds' in the sense that you input a finite set $\phi_1, \ldots, \phi_k$ of fully irreducible automorphisms and get out a Gromov hyperbolic complex X with an $\mathrm{Out}(F_n)$ action and particularly nice action of $\phi_1, \ldots, \phi_k$. The complex of free factors $\mathcal{F}$ is a major improvement in the sense that it replaces all of these custom builds – all fully irreducible automorphisms act 'nicely' on $\mathcal{F}$. It is possible that $\mathcal{F}$ could also be used as a first step in the proof that $\mathbb{Z}$–averse groups have finite image in $\mathrm{Out}(F_n)$ (Theorem 5.1). Indeed, if similar complexes were developed for subgroups of $\mathrm{Out}(F_n)$ that fix a free factor (up to finite index this is equivalent to not containing a fully

irreducible element [43]), one may be able to give an alternate proof of Theorem 5.1 that does not use the structure of $\overline{\mathrm{IA}}_n$ at all.

An automorphism ϕ is *reducible* if there exists a non-trivial decomposition $F_n = F^1 * F^2 * \ldots * F^k * B$ and $g_1, \ldots, g_k \in F_n$ so that $\phi(F^i) = g_i F^{i+1} g_i^{-1} \mod k$. We say that ϕ is *irreducible* if it is not reducible, and *fully irreducible* if every power of ϕ is irreducible. Note that if we change ϕ by an element of $\mathrm{Inn}(F_n)$ it will stay reducible, irreducible, or fully irreducible, so these definitions extend to elements of $\mathrm{Out}(F_n)$. Bestvina [8] uses the Lipschitz metric to prove the existence of *train track maps* for *irreducible automorphisms*. The first step in Bestvina's proof is to classify the action of an element Φ on CV_n via a *displacement function*. Define

$$D(\Phi) = \inf\{d(\Gamma, \Gamma.\Phi) : \Gamma \in \mathrm{CV}_n\}.$$

If there exists $\Gamma \in \mathrm{CV}_n$ on which $D(\Phi)$ is attained, we say Φ is *semisimple*. In this case Φ is *elliptic* if $D(\Phi) = 0$ and *hyperbolic* if $D(\Phi) > 0$. If $D(\Phi)$ is not attained on any point of CV_n we say that Φ is *parabolic*. Bestvina shows that if Φ is irreducible then Φ is hyperbolic and there is a perturbation of a point on which $D(\Phi)$ is attained that gives a train track map. In the case that Φ is reducible, he uses the metric to find a free factor decomposition of F_n that witnesses this reducibility.

Similar to the displacement function, there are two asymptotic constants that can be assigned to an element of $\mathrm{Out}(F_n)$. The *asymptotic displacement length* again uses the Lipschitz metric and is defined by

$$A(\Phi) = \lim_{n\to\infty} \frac{d(\Gamma, \Gamma.\Phi^n)}{n}.$$

It is easily verified that $A(\Phi)$ is well-defined, independent of the choice of $\Gamma \in \mathrm{CV}_n$ and $A(\Phi) \leq D(\Phi)$. The second constant is the *exponential growth rate* associated to a conjugacy class α of F_n. For $w \in F_n$ define $\|w\|$ to be the cyclically reduced length of w with respect to a word metric on F_n. Then

$$EGR(\Phi, \alpha) = \lim_{n\to\infty} \frac{\log \|\Phi^n(\alpha)\|}{n}.$$

Let $\mathcal{C}$ be the set of conjugacy classes of elements of F_n. We show that these functions are related in the following way:

Theorem 7.7. *Let $\Phi \in \mathrm{Out}(F_n)$. The contsants $\sup_{\alpha\in\mathcal{C}} EGR(\Phi, \alpha)$, $A(\Phi)$, and $D(\Phi)$ are all equal. Furthermore,*

$$D(\Phi) = \begin{cases} 0 & \text{if } \Phi \text{ is NEG} \\ \max\{\log(\lambda) : \lambda \in PF_\Phi\} & \text{if } \Phi \text{ is EG.} \end{cases}$$

Here EG means *exponentially growing*, NEG is *not exponentially growing*, and PF_Φ is a set of Perron–Frobenius eigenvalues associated to a train track representative of Φ. We will give precise definitions of EG, NEG and PF_Φ in Section 7.2. The proof of Theorem 7.7 uses the theory of relative train tracks from [14], which are more general than the train track maps constructed in [8] using the Lipschitz metric. This may appear to be against the philosophy of building up a theory of Outer space completely from the metric viewpoint. However, at the end of [8] Bestvina states that the same ideas on a relative version of Outer space could be used to give a metric proof for the existence of relative train track maps.

7.1 Outer space and the Lipschitz metric

The *topological realisation* of a graph Γ is a cell complex obtained by taking one 0–cell for each vertex of Γ, and one 1–cell for each edge in an orientation $\mathcal{O}$ of Γ. The 1–cells are glued to the 0–cells according to the maps ι and τ. A map $\ell : \mathrm{E}\Gamma \to \mathbb{R}_{>0}$ that satisfies $\ell(e) = \ell(\bar{e})$ induces a metric on the topological realisation by setting the length of each edge to be $\ell(e)$. We shall abuse notation by blurring the distinction between a graph and its (possibly metrized) topological realisation.

We are interested in triples (Γ, f, ℓ), where Γ is a *minimal* graph (every edge of Γ is contained in an immersed loop), $f : R_n \to \Gamma$ is a (topological) homotopy equivalence and $\ell : \mathrm{E}\Gamma \to \mathbb{R}_{>0}$ a map inducing a metric on Γ as above. We define an equivalence relation on this set of triples by saying that two points (Γ, f, ℓ) and (Γ', f', ℓ') are equivalent if there exists an isometry $g : \Gamma \to \Gamma'$ such that gf is freely homotopic to f'. Let $[\Gamma, f, \ell]$ denote the equivalence class of (Γ, f, ℓ) under this relation. Where it does not cause confusion we will sometimes write Γ rather than $[\Gamma, f, \ell]$.

Culler and Vogtmann's *Outer space*, denoted CV_n, is defined to be the set of equivalence classes $[\Gamma, f, \ell]$ such that the metric on Γ has volume one ($\sum_{e \in \mathcal{O}} \ell(e) = 1$ for an orientation $\mathcal{O}$ of Γ). The equivalence class $[\Gamma, f, \ell]$ is called a *marked graph*, and the map f a *marking* of the metric graph Γ. There is a natural action of $\mathrm{Out}(F_n)$ on CV_n; any element $\Phi \in \mathrm{Out}(F_n)$ can be viewed as a homotopy equivalence $\Phi : R_n \to R_n$ from the rose to itself, therefore we can define $[\Gamma, f, \ell].\Phi = [\Gamma, f\Phi, \ell]$.

Any continuous map between two metric graphs can be homotoped to a map that is linear on edges. If $g : \Gamma \to \Gamma'$ is such a map between metric graphs, define $\mathrm{Lip}(e)$ to be the slope of g on the edge e. The *Lipschitz constant* of f, denoted $\mathrm{Lip}(f)$, is the largest value of $\mathrm{Lip}(e)$ as e runs through all the edges of Γ. For two points $[\Gamma, f, \ell]$

and $[\Gamma', f', \ell']$ of CV_n we say that a *difference of markings* is a map $g : \Gamma \to \Gamma'$ that is linear on edges such that fg is freely homotopic to f'. Let

$$d([\Gamma, f, \ell], [\Gamma', f', \ell']) = \inf\{\log(\mathrm{Lip}(g)) : g \text{ is a difference of markings}\}.$$

We will sometimes shorten this to $d(\Gamma, \Gamma')$. We say that g is *optimal* if $\mathrm{Lip}(g)$ is minimal as we run through differences of markings. For any pair of points in CV_n a minimal difference of markings always exists (see Proposition 7.1, below). As any difference of markings g is surjective and both Γ and Γ' have volume one, $\mathrm{Lip}(g) \geq 1$. If $\mathrm{Lip}(g) = 1$ then g is an isometry, and $[\Gamma, f, \ell] = [\Gamma', f', \ell']$ in CV_n. Therefore $d(\Gamma, \Gamma') \geq 0$ and $d(\Gamma, \Gamma') = 0$ if and only if $\Gamma = \Gamma'$ in CV_n. For any two maps $g : \Gamma \to \Gamma'$ and $h : \Gamma' \to \Gamma''$ that are linear on edges, $\mathrm{Lip}(hg) \leq \mathrm{Lip}(h)\mathrm{Lip}(g)$. It follows that $d(\Gamma, \Gamma'') \leq d(\Gamma, \Gamma') + d(\Gamma', \Gamma'')$ for any three points in CV_n. It would be pleasing if d were also symmetric, so that $d(\Gamma, \Gamma') = d(\Gamma', \Gamma)$, however this is not always the case. In Figure 7.1 we look at two metrics on R_2. The first assigns both loops length $\frac{1}{2}$, the second assigns the first length $1 - \frac{1}{n}$ and the second length $\frac{1}{n}$. The maps f and f' are obtained in the obvious way. One can show that the obvious choices for differences of markings g and h are optimal (for instance using Proposition 7.1, below), therefore $d(\Gamma, \Gamma') = \log(2 - \frac{2}{n})$ and $d(\Gamma', \Gamma) = \log(\frac{n}{2})$.

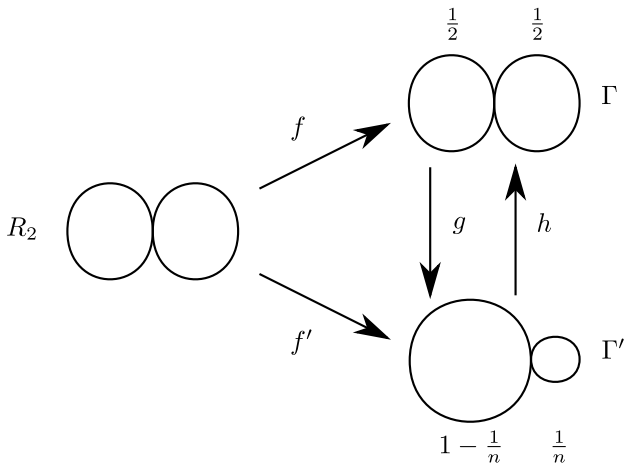

Figure 7.1:

Because of the above properties d is referred to as the *non-symmetric Lipschitz metric* on CV_n, or the Lipschitz metric for short. If α is a conjugacy class in $\mathcal{C}$ and $[\Gamma, f, \ell]$ is a marked graph, there is a unique (up to rotation) shortest loop α_Γ in Γ representing α under the marking given by f. It is obtained by taking a loop in the conjugacy class of α in R_n, mapping it to Γ under f, and then tightening its image

in Γ to an immersion. We use $[\beta]$ to denote a tightened version of a loop β. If β is an edge loop (so does not only partially cross an edge) then $[\beta]$ is obtained by removing any occurrences of $(e, \bar{e})$ from β. Denote the length of α_Γ in Γ by $\ell(\alpha)$. If g is a difference of markings between $[\Gamma, f, \ell]$ and $[\Gamma', f', \ell']$ then $\alpha_{\Gamma'} = [g(\alpha_\Gamma)]$. Where it does not cause confusion we shall refer to α_Γ simply as α. The following proposition is key to our discussion; proofs can be found in [38] and [8].

Proposition 7.1. *Let $[\Gamma, f, \ell]$ and $[\Gamma', f', \ell']$ be marked graphs. Then*

$$\inf\{\mathrm{Lip}(g) : g \text{ is a difference of markings}\} = \sup_{\alpha \in \mathcal{C}} \frac{\ell'([g(\alpha)])}{\ell(\alpha)}.$$

Moreover, both inf and sup are realised.

Note that $\frac{\ell'([g(\alpha)])}{\ell(\alpha)}$ is independent of our choice of g, therefore we do not need to find an optimal g to calculate distance in CV_n. In particular, if g is any difference of markings between (Γ, f, ℓ) and $(\Gamma, f, \ell).\Phi$, we have

$$d(\Gamma, \Gamma.\Phi^n) = \sup_{\alpha \in \mathcal{C}} \frac{\ell([g^n(\alpha)])}{\ell(\alpha)}.$$

Lemma 7.2. *Let X be a basis of F_n and let g be a difference of markings between Γ and $\Gamma.\Phi$. There exists a constant C such that for any conjugacy class α in F_n:*

$$\frac{1}{C}\|\Phi^n(\alpha)\|_X \leq \ell([g^n(\alpha)]) \leq C\|\Phi^n(\alpha)\|_X,$$

and it follows that

$$EGR(\Phi, \alpha) = \lim_{n\to\infty} \frac{\log(\ell([g^n(\alpha)]))}{n} \leq A(\Phi).$$

Proof. As any two bases of F_n are Lipschitz equivalent we may assume that X is obtained by choosing a maximal tree T in Γ, and choosing the generating set determined by T, as described in the previous chapter. We can then find $\|\Phi^n(\alpha)\|_X$ by taking the tightened loop $[g^n(\alpha)]$ and counting number of edges of $\Gamma \setminus T$ that this loop crosses. We can compute upper and lower bounds for this, giving

$$\frac{\ell([g^n(\alpha)])}{2.\mathrm{Diam}(T) + \max\{\ell(e) : e \in \Gamma \setminus T\}} \leq \|\Phi^n(\alpha)\|_X \leq \frac{\ell([g^n(\alpha)])}{\min\{\ell(e) : e \in \Gamma \setminus T\}}.$$

Here Diam(T) is the diameter of the maximal tree T. If follows that such a constant C exists. Then

$$\begin{aligned} EGR(\Phi, \alpha) &= \lim_{n\to\infty} \frac{\log(\|\Phi^n(\alpha)\|_X)}{n} \\ &= \lim_{n\to\infty} \frac{\log(\ell([g^n(\alpha)]))}{n} \\ &= \lim_{n\to\infty} \frac{\log(\frac{\ell([g^n(\alpha)])}{\ell(\alpha)})}{n} \\ &\leq \lim_{n\to\infty} \frac{d(\Gamma, \Gamma.\Phi^n)}{n} \\ &= A(\Phi), \end{aligned}$$

by the work above and Proposition 7.1. □

7.2 Train tracks and displacement functions

We say that a difference of markings g between $[\Gamma, f, \ell]$ and $[\Gamma, f\Phi, \ell]$ is a *topological representative* if it maps vertices to vertices. As a difference of markings is linear, such maps are determined by a consistent assignment of an edge path in Γ to each edge in Γ. Therefore if $e_1, \ldots, e_k$ is an ordering of the (topological) edges of Γ we obtain a $k \times k$ *transition matrix* M, by taking a_{ij} to be the number of times $g(e_j)$ crosses e_i in either direction. We say that M is irreducible if there is no way of permuting the order of the edges so that M has a non-trivial block decomposition $M = \left(\begin{smallmatrix} A & B \\ 0 & D \end{smallmatrix}\right)$. A *filtration* is an increasing sequence of g–invariant subgraphs $\emptyset = \Gamma_0 \subset \Gamma_1 \subset \ldots \Gamma_m = \Gamma$. The edges in $\Gamma_r \setminus \Gamma_{r-1}$ constitute the r*th stratum* of Γ and their union is denoted H_r. Each stratum has an associated $|H_r| \times |H_r|$ submatrix M_r of the transition matrix attained by solely looking at the edges of H_r, and we say that the filtration is *maximal* if each M_r is either irreducible or zero. Where it is 0, H_r is called a *zero-stratum*. If M_r is irreducible the Perron-Frobenius theorem ([15], Appendix A) tells us there is a unique (up to scaling) positive eigenvector of M_r with eigenvalue $\lambda_r \geq 1$. If $\lambda = 1$ then M_r is a permutation matrix, and we say H_r is a polynomially growing (PG) stratum. If $\lambda_r > 1$ then we say that H_r is exponentially growing (EG). If we assign the edges of H_r lengths according to a left positive eigenvector of M_r (which also has eigenvalue λ_r) then

$$\frac{\ell(g(e) \cap H_r)}{\ell(e)} = \lambda_r \tag{7.1}$$

for each edge e in H_r.

To study cancellation of paths under iteration by g we use the notion of *turns*. A topological representative $g : \Gamma \to \Gamma$ induces a map Dg on oriented edges by mapping the edge e to the initial edge of $g(e)$. A *turn* is an unordered pair of oriented edges that have the same initial vertex. The map Dg then induces a map Tg on the set of turns. We say that a turn is *degenerate* if the edges in the pair are identical, and *non-degenerate* if they are distinct. A turn is *legal* if it mapped to a non-degenerate turn. (In some papers a turn is legal if and only if it is mapped to a non-degenerate turn under every iteration of Tg. We only need the weaker definition here.)

We recall the definition of relative train tracks from [14]:

We say that $g : \Gamma \to \Gamma$ is a *relative train track representative* of Φ if g is a topological representative of Φ and the following conditions hold for each exponentially growing stratum H_r:

(RTT-i) Dg maps the set of oriented edges in H_r to itself; in particular all mixed turns in (Γ_r, Γ_{r-1}) are legal.

(RTT-ii) If $\alpha \subset \Gamma_{r-1}$ is a non-trivial path with endpoints in $H_r \cap \Gamma_{r-1}$, then $[g(\alpha)]$ is a non-trivial path with endpoints in $H_r \cap \Gamma_{r-1}$.

(RTT-iii) For each legal path $\beta \subset H_r$, $g(\beta)$ is a path that does not contain any illegal turns in H_r.

We shall make use of the following theorem:

Theorem 7.3 ([14], Theorem 5.12)**.** *For every outer automorphism Φ there exists a relative train track representative $g : \Gamma \to \Gamma$ of Φ.*

Given a relative train track representative $g : \Gamma \to \Gamma$ of Φ, we are free to choose a metric ℓ on Γ as we wish. The set of eigenvalues of exponential strata is independent of the choice of relative train track (see [14]); we therefore let PF_Φ be the set of Perron-Frobenius eigenvalues associated to the exponential strata of a relative train track representative of Φ. We say that Φ is exponentially growing (sometimes shortened to EG) if PF_Φ is nonempty, and Φ is of non-exponential growth (NEG) otherwise. A path (or loop) is said to be *r–legal* if it is contained in Γ_r and each turn in H_r is legal. The relative train track conditions imply that no cancellation occurs in H_r under iteration of an r-legal loop by g.

Proposition 7.4. *Let $g : \Gamma \to \Gamma$ be a relative train track representative of Φ. For every $\epsilon > 0$ there exists a metric ℓ on Γ such that for an edge e in a stratum H_i:*

$$Lip(e) \leq \begin{cases} \lambda_i + \epsilon & \text{if } H_i \text{ is EG} \\ 1 + \epsilon & \text{if } H_i \text{ is PG} \\ \epsilon & \text{if } H_i \text{ is a zero stratum} \end{cases}$$

Proof. We proceed by induction moving up the strata. H_1 is either polynomially growing, in which case we give each edge length 1, or H_1 is exponentially growing and we assign the edges lengths according to a left positive eigenvector of the transition matrix M_1. Then, by equation (7.1) the Lipschitz constant for each edge is λ_1. Now assume we have constructed such a metric ℓ on Γ_{r-1}. As for the base case, if H_r is a polynomial or zero stratum, give each edge length 1, and in the second case assign the edges lengths according to a positive left eigenvector of M_r. Then for any edge e in H_r, we have $\frac{\ell(g(e)\cap H_r)}{\ell(e)}=\lambda_r$ if H_r is EG, $\frac{\ell(g(e)\cap H_r)}{\ell(e)}=1$ if H_r is PG, and $\frac{\ell(g(e)\cap H_r)}{\ell(e)}=0$ if H_r is a zero stratum. Uniformly contracting or expanding the edges of Γ_{r-1} does not change their Lipschitz constants, therefore we may shrink all the edges of Γ_{r-1} so that $\frac{\ell(g(e)\cap\Gamma_{r-1})}{\ell(e)}<\epsilon$. This ensures every edge of Γ_r has Lipschitz constants as required. □

Corollary 7.5.

$$D(\Phi)\leq\begin{cases}0 & \textit{if } \Phi \textit{ is NEG}\\ \max\{\log(\lambda):\lambda\in PF_\Phi\} & \textit{if } \Phi \textit{ is EG}\end{cases}$$

Proof. Let g, Γ be as in Proposition 7.4, and scale the metric given in Proposition 7.4 so that is has volume 1. The map g is a difference of markings between (Γ,f,ℓ) and $(\Gamma,f,\ell).\Phi$, therefore

$$d(\Gamma,\Gamma.\Phi)\leq\begin{cases}\log(1+\epsilon) & \text{if } \Phi \text{ is NEG}\\ \max\{\log(\lambda+\epsilon):\lambda\in PF_\Phi\} & \text{if } \Phi \text{ is EG}\end{cases}$$

The result follows by taking a sequence of graphs where $\epsilon\to 0$. □

Proposition 7.6.

$$\sup_{\alpha\in\mathcal{C}} EGR(\Phi,\alpha)\geq\max\{\log(\lambda):\lambda\in PF_\Phi\}$$

Proof. Let g be a relative train track representative of Φ with a metric given by Proposition 7.4. Let λ_r be the Perron-Frobenius eigenvalue of an EG stratum H_r. Let e be an edge of H_r. The edge e is r-legal, so no cancellation occurs in H_r while iterating e. Therefore there exists k such that $[g^k(e)]$ runs over e at least three times. At least two of these crossings occur with the same orientation, and cutting $[g^k(e)]$ here gives an r–legal loop α such that $e\subset\alpha$. Then by the work in Proposition 7.2

$$EGR(\Phi,\alpha)=\lim_{n\to\infty}\frac{\log(\ell([g^n(\alpha)]))}{n}\geq\lim_{n\to\infty}\frac{\log(\ell(g^n(e)\cap H_r))}{n}=\log(\lambda_r).$$

Repeating this for every EG stratum gives the required result. □

Theorem 7.7. *Let* $\Phi \in \mathrm{Out}(F_n)$. *The contsants* $\sup_{\alpha\in\mathcal{C}} EGR(\Phi,\alpha)$, $A(\Phi)$, *and* $D(\Phi)$ *are all equal. Furthermore,*

$$D(\Phi) = \begin{cases} 0 & \text{if } \Phi \text{ is NEG} \\ \max\{\log(\lambda) : \lambda \in PF_\Phi\} & \text{if } \Phi \text{ is EG.} \end{cases}$$

Proof. We have

$$0 \leq \sup_{\alpha\in\mathcal{C}} EGR(\Phi,\alpha) \leq A(\Phi) \leq D(\Phi) \leq \sup_{\alpha\in\mathcal{C}} EGR(\Phi,\alpha).$$

The second inequality is given by Lemma 7.2, the third follows from the triangle inequality for d. The final inequality follows from Corollary 7.5 and Proposition 7.6. Furthermore, Corollary 7.5 and Proposition 7.6 tell us the value of these functions. □

Appendix A

$\langle \mathcal{M}_\Gamma \rangle$ is a normal subgroup of $\mathrm{SAut}^0(A_\Gamma)$.

We would like to show that the subgroup of $\mathrm{SAut}^0(A_\Gamma)$ generated by elements of $\mathcal{M}_\Gamma$ (see Definition 4.6) is normal. This will complete the proof that $\mathcal{M}_\Gamma$ generates $\mathrm{IA}(A_\Gamma)$ (Theorem 4.7). One would like to directly take the computations made when proving Magnus' generators of IA_n generate a normal subgroup of $\mathrm{Aut}(F_n)$, and indeed this is roughly our strategy. However, K_{ijk} (the automorphism taking v_i to $v_i[v_j, v_k]$) does not exist for all i, j, k. Also K_{ij} conjugates every vertex in Γ_{ij} (the connected component of $\Gamma - st(v_j)$ containing v_i) by v_j, so may conjugate other basis elements of A_Γ in addition to v_i. Therefore more care needs to be taken. We make extensive use of the fact that a transvection ρ_{ij} exists only if $v_i \leq v_j$. Most of the difficulties are dealt with by the following two lemmas.

Lemma A.1. *Suppose that $v_i \leq v_j$. Then the following holds:*

1. *The vertex set of each connected component Γ_{ki} of $\Gamma - st(v_i)$ consists of the vertices of $\Gamma_{ki} \cap st(v_j)$ together with the union of the vertex sets of certain connected components $(\Gamma_{tj})_{t \in T}$ of $\Gamma - st(v_j)$ indexed by a set $T \subset V(\Gamma)$.*

2. *There automorphism $\prod_{t \in T} K_{tj}$ conjugates every element of Γ_{ki} by v_j and fixes all other basis elements of A_Γ.*

3. *Either v_i commutes with v_j or K_{ij} fixes every basis element except v_i.*

4. *Every automorphism of the form K_{lj} where $l \neq i$ fixes v_i.*

Proof. If $[v_i, v_j] = 0$ then $st(v_i) \subset st(v_j)$, hence each connected component of $\Gamma - st(v_j)$ is a subset of a connected component of $\Gamma - st(v_i)$. If $[v_i, v_j] \neq 0$ then $lk(v_i) \subset st(v_j)$, and $\Gamma_{ij} = \{v_i\}$. If two vertices in $\Gamma - st(v_j) - \{v_i\}$ are connected by a path in $\Gamma - st(v_j)$,

then they must also be connected by a path in $\Gamma - st(v_i)$, therefore each connected component of $\Gamma - st(v_j) - \{v_i\}$ is a subset of a connected component of $\Gamma - st(v_i)$. This completes the proof of (1), and (2) is a direct consequence. Parts (3) and (4) follow from the fact that either $v_i \subset st(v_j)$ or $\Gamma_{ij} = \{v_i\}$. □

Lemma A.2. *Suppose that $v_i \leq v_j$, and v_l is a generator distinct from v_i and v_j. Then either*

1. *v_i and v_j both commute with v_l.*
2. *v_j commutes with v_l but v_i does not.*
3. *v_i and v_j do not commute with v_l, and lie in the same component of $\Gamma - st(v_l)$.*
4. *v_j and v_i do not commute with v_l and lie in distinct components of $\Gamma - st(v_l)$. Then $v_i \leq v_l$.*

Proof. If $v_i \in st(v_l)$, then as $lk(v_i) \subset st(v_j)$, this implies that $v_j \in st(v_l)$. Hence if either v_i or v_j commutes with v_l we must be in case 1 or case 2. If v_i and v_j both lie outside of $st(v_l)$, the only way that they could lie in distinct components of $\Gamma - st(v_l)$ would be if $lk(v_i) \subset st(v_l)$, as $lk(v_i) \subset st(v_j)$. This shows that the other possibilities are case 3 and case 4. □

A.1 Conjugates of the form $\rho_{ij}K_{kl}\rho_{ij}^{-1}$ and $\rho_{ij}^{-1}K_{kl}\rho_{ij}$.

Proposition A.3. *For all i, j, k, l such that the relevant automorphisms exist, the elements $\rho_{ij}K_{kl}\rho_{ij}^{-1}$ and $\rho_{ij}^{-1}K_{kl}\rho_{ij}$ lie in the group generated by the set $\mathcal{M}_\Gamma$.*

We work on a case-by-case basis. We first check that if $j = l$ then $\rho_{ij}K_{kl}\rho_{ij}^{-1} = \rho_{ij}K_{kj}\rho_{ij}^{-1} = K_{kj}$, and similarly $\rho_{ij}^{-1}K_{kj}\rho_{ij} = K_{kj}$. When $i = l$, by Lemma A.1 there exists a subset $T \subset V(\Gamma)$ such that the automorphism that conjugates every element of Γ_{ki} by v_j is of the form $\prod_{t\in T} K_{tj}$. We find that

$$\rho_{ij}K_{ki}\rho_{ij}^{-1} = \prod_{t\in T} K_{tj}K_{ki}$$
$$\rho_{ij}^{-1}K_{ki}\rho_{ij} = \prod_{t\in T} K_{tj}^{-1}K_{ki},$$

when $j \notin \Gamma_{ki}$. If $j \in \Gamma_{ki}$, then $K_{ki} = K_{ji}$ and

$$\rho_{ij} K_{ki} \rho_{ij}^{-1} = \rho_{ij} K_{ji} \rho_{ij}^{-1} = \prod_{t \in T} K_{tj} K_{ji} K_{ij}$$
$$\rho_{ij}^{-1} K_{ki} \rho_{ij} = \rho_{ij}^{-1} K_{ji} \rho_{ij} = \prod_{t \in T} K_{tj}^{-1} K_{ji} K_{ij}^{-1}.$$

We are now left with the situation where i and j are both distinct from l, and we can apply Lemma A.2 to leave us with the following cases:

1. $i, j \notin \Gamma_{kl}$;
2. $i \in \Gamma_{kl}$, $[x_j, x_l] = 0$;
3. $i, j \in \Gamma_{kl}$;
4. $i \in \Gamma_{kl}$, $j \notin \Gamma_{kl}, [x_j, x_l] \neq 0, x_i \leq x_l$;
5. $i \notin \Gamma_{kl}$, $j \in \Gamma_{kl}, [x_i, x_l] \neq 0, x_i \leq x_l$.

In the first three sub-cases, ρ_{ij} and K_{kl} commute, giving $\rho_{ij} K_{kl} \rho_{ij}^{-1} = \rho_{ij}^{-1} K_{kl} \rho_{ij} = K_{kl}$. In the remaining two cases, note that as $v_i \leq v_j, v_l$ the automorphisms K_{ijl} and K_{ilj} are well-defined, and Lemma A.1 tells us how conjugating by v_l and v_j behaves. In case 4 we have $K_{kl} = K_{il}$ and

$$\rho_{ij} K_{il} \rho_{ij}^{-1} = K_{il} K_{ilj}$$
$$\rho_{ij}^{-1} K_{il} \rho_{ij} = K_{il} K_{ij}^{-1} K_{ijl} K_{ij}^{-1}.$$

In case 5 we have $K_{kl} = K_{jl}$ and

$$\rho_{ij} K_{jl} \rho_{ij}^{-1} = K_{ijl} K_{jl}$$
$$\rho_{ij}^{-1} K_{jl} \rho_{ij}^{-1} = K_{ij} K_{ilj} K_{ij}^{-1} K_{jl}.$$

A.2 Conjugates of the form $\rho_{ij} K_{klm} \rho_{ij}^{-1}$ and of the form $\rho_{ij}^{-1} K_{klm} \rho_{ij}$

Proposition A.4. *For all i, j, k, l, m such that the relevant automorphisms exist, the elements $\rho_{ij} K_{klm} \rho_{ij}^{-1}$ and $\rho_{ij}^{-1} K_{klm} \rho_{ij}$ lie in the group generated by the set $\mathcal{M}_\Gamma$.*

After noting that $K_{klm} = K_{kml}^{-1}$ we having the following cases to check:

1. $i, j \neq k, l, m$;

2. $i = k$;

3. $i = k$ and $j = l$;

4. $i = l$;

5. $i = l$ and $j = m$;

6. $i = l$ and $j = k$;

7. $j = l$;

8. $j = k$.

When we do not specify a value of i or j we assume that it is not in the set $\{k, l, m\}$. When we are in the situation of 1, 5, or 7 the elements ρ_{ij} and K_{klm} commute. Remembering that $x_j \leq x_k$ implies that K_{kj} fixes every element other than v_k, in case 2 we have:

$$\rho_{kj} K_{klm} \rho_{kj}^{-1} = K_{kj}^{-1} K_{klm} K_{kj}$$
$$\rho_{kj}^{-1} K_{klm} \rho_{kj} = K_{kj} K_{klm} K_{kj}^{-1},$$

unless $[x_j, x_k] = 0$, in which case $[v_j, v_l] = [v_j, v_m] = 0$ and ρ_{kj} and K_{klm} commute. In case 3 we have

$$\rho_{kl} K_{klm} \rho_{kl}^{-1} = K_{kl}^{-1} K_{klm} K_{kl}$$
$$\rho_{kl}^{-1} K_{klm} \rho_{kl} = K_{kl} K_{klm} K_{kl}^{-1}.$$

For case 4, either v_j commutes with v_m, in which case ρ_{lj} and K_{klm} commutes, or since $v_k \leq v_l \leq v_j$ the automorphism K_{kjm} is well-defined. Hence

$$\rho_{lj} K_{klm} \rho_{lj}^{-1} = K_{kl}^{-1} K_{kjm} K_{kl} K_{klm}$$
$$\rho_{lj}^{-1} K_{klm} \rho_{lj} = K_{kl}^{-1} K_{kj} K_{kmj} K_{kj}^{-1} K_{kl} K_{klm}.$$

Case 6 is the trickiest. Firstly, note that as $v_l \leq v_k$ and $v_k \leq v_l$, this implies that $[v_l] = [v_k]$. In particular $st(v_k) = lk(v_l) \cup \{v_k\}$, and we see that the connected components of $\Gamma - st(v_k) - \{v_l\}$ are equal to the components of $\Gamma - st(v_l) - \{v_k\}$. Therefore the automorphisms K_{mk} and K_{ml} conjugate the same basis elements. As $v_k \leq v_m$, this implies that $v_l \leq v_m$, so the automorphism K_{lmk} exists, and furthermore

K_{lm}, K_{kl} and K_{km} fix every basis except v_l, v_k and v_k respectively. These facts allow one to confirm that

$$\rho_{lk} K_{klm} \rho_{lk}^{-1} = K_{mk} K_{ml} K_{lm} K_{lmk} K_{kl} K_{klm} K_{km}^{-1} K_{kl}^{-1} K_{ml}^{-1} K_{mk}^{-1}.$$

For $\rho_{lk}^{-1} K_{klm} \rho_{lk}$ we need to use Lemma A.1 to find a set of connected components $\{\Gamma_{tm}\}_{t \in T}$ such that $\prod_{t \in T} K_{tm}$ conjugates every element of $\Gamma_{mk} = \Gamma_{ml}$ by v_m. (This product may be trivial.) Then

$$\rho_{lk}^{-1} K_{klm} \rho_{lk} = K_{mk}^{-1} K_{ml} K_{lm}^{-1} (\prod_{t \in T} K_{tm}^{-1}) K_{ml}^{-1} (\prod_{t \in T} K_{tm}) K_{klm} K_{lmk} K_{mk} K_{km}.$$

Finally, in case 8, we have $v_i \leq v_k \leq v_l, v_m$, so either $[v_i, v_k] = [v_i, v_l] = [v_i, v_m] = 0$ and conjugation fixes K_{klm}, or the following occurs:

$$\begin{aligned} \rho_{ik} K_{klm} \rho_{ik}^{-1} &= K_{ik}^{-1} K_{iml} K_{ik} K_{klm} \\ \rho_{ik}^{-1} K_{klm} \rho_{ik}^{-1} &= K_{ilm} K_{klm}. \end{aligned}$$

Bibliography

[1] Yael Algom-Kfir. Strongly contracting geodesics in outer space. *arXiv:0812.1555*, 2008.

[2] Yael Algom-Kfir and Mladen Bestvina. Asymmetry of outer space. *arXiv:0910.5408*, 2009.

[3] S. Andreadakis. On the automorphisms of free groups and free nilpotent groups. *Proc. London Math. Soc. (3)*, 15:239–268, 1965.

[4] Uri Bader and Yehuda Shalom. Factor and normal subgroup theorems for lattices in products of groups. *Invent. Math.*, 163(2):415–454, 2006.

[5] Hyman Bass and Alexander Lubotzky. Linear-central filtrations on groups. In *The mathematical legacy of Wilhelm Magnus: groups, geometry and special functions (Brooklyn, NY, 1992)*, volume 169 of *Contemp. Math.*, pages 45–98. Amer. Math. Soc., Providence, RI, 1994.

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

[7] M. Bestvina and M. Handel. Train-tracks for surface homeomorphisms. *Topology*, 34(1):109–140, 1995.

[8] Mladen Bestvina. A Bers-like proof of the existence of train tracks for free group automorphisms. *Fund. Math.*, 214(1):1–12, 2011.

[9] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[10] Mladen Bestvina, Kai-Uwe Bux, and Dan Margalit. Dimension of the Torelli group for Out(F_n). *Invent. Math.*, 170(1):1–32, 2007.

[11] Mladen Bestvina and Mark Feighn. A hyperbolic Out(F_n)-complex. *Groups Geom. Dyn.*, 4(1):31–58, 2010.

[12] Mladen Bestvina and Mark Feighn. Hyperbolicity of the complex of free factors. *arXiv:1107.3308*, 2011.

[13] Mladen Bestvina and Koji Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geom. Topol.*, 6:69–89, 2002.

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

[15] Oleg Bogopolski. *Introduction to group theory*. EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2008.

[16] Nicolas Bourbaki. *Lie groups and Lie algebras. Chapters 1–3*. Elements of Mathematics (Berlin). Springer-Verlag, Berlin, 1998.

[17] Martin R. Bridson and Richard D. Wade. Actions of higher-rank lattices on free groups. *Compos. Math.*, 147(5):1573–1580, 2011.

[18] M. Burger and N. Monod. Bounded cohomology of lattices in higher rank Lie groups. *J. Eur. Math. Soc. (JEMS)*, 1(2):199–235, 1999.

[19] Ruth Charney, John Crisp, and Karen Vogtmann. Automorphisms of 2-dimensional right-angled Artin groups. *Geom. Topol.*, 11:2227–2264, 2007.

[20] Ruth Charney and Karen Vogtmann. Finiteness properties of automorphism groups of right-angled Artin groups. *Bull. Lond. Math. Soc.*, 41(1):94–102, 2009.

[21] Ruth Charney and Karen Vogtmann. Subgroups and quotients of automorphism groups of RAAGs. In *Low-dimensional and symplectic topology*, volume 82 of *Proc. Sympos. Pure Math.*, pages 9–27. Amer. Math. Soc., Providence, RI, 2011.

[22] Orin Chein. Subgroups of IA automorphisms of a free group. *Acta Math.*, 123:1–12, 1969.

[23] K.-T. Chen, R. H. Fox, and R. C. Lyndon. Free differential calculus. IV. The quotient groups of the lower central series. *Ann. of Math. (2)*, 68:81–95, 1958.

[24] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[25] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *arXiv:1111.7048*, 11 2011.

[26] Matthew B. Day. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geom. Topol.*, 13(2):817–855, 2009.

[27] Matthew B. Day. Symplectic structures on right-angled Artin groups: between the mapping class group and the symplectic group. *Geom. Topol.*, 13(2):857–899, 2009.

[28] Matthew B. Day and Andrew Putman. A Birman exact sequence for $Aut(F_n)$. *arXiv:1104.2371*, 2011.

[29] Thomas Delzant and Misha Gromov. Courbure mésoscopique et théorie de la toute petite simplification. *J. Topol.*, 1(4):804–836, 2008.

[30] Warren Dicks and Enric Ventura. *The group fixed by a family of injective endomorphisms of a free group*, volume 195 of *Contemporary Mathematics*. American Mathematical Society, Providence, RI, 1996.

[31] G. Duchamp and D. Krob. The lower central series of the free partially commutative group. *Semigroup Forum*, 45(3):385–394, 1992.

[32] G. Duchamp and D. Krob. Free partially commutative structures. *J. Algebra*, 156(2):318–361, 1993.

[33] Joan L. Dyer and Edward Formanek. The automorphism group of a free group is complete. *J. London Math. Soc. (2)*, 11(2):181–190, 1975.

[34] Benson Farb and Peter Shalen. Lattice actions, 3-manifolds and homology. *Topology*, 39(3):573–587, 2000.

[35] Mark Feighn and Michael Handel. The recognition theorem for $\mathrm{Out}(F_n)$. *Groups Geom. Dyn.*, 5(1):39–106, 2011.

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

[37] Stefano Francaviglia and Armando Martino. The isometry group of outer space. *arXiv:0912.0299*, 2009.

[38] Stefano Francaviglia and Armando Martino. Metric properties of outer space. *Publ. Mat.*, 55(2):433–473, 2011.

[39] Elisabeth R. Green. *Graph products of groups*. PhD thesis, The University of Leeds, 1990.

[40] R. I. Grigorchuk and J. S. Wilson. The conjugacy problem for certain branch groups. *Tr. Mat. Inst. Steklova*, 231(Din. Sist., Avtom. i Beskon. Gruppy):215–230, 2000.

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

[42] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[43] Michael Handel and Lee Mosher. Subgroup classification in $Out(F_n)$. *arXiv:0908.1255*, 2009.

[44] P. J. Higgins and R. C. Lyndon. Equivalence of elements under automorphisms of a free group. *J. London Math. Soc. (2)*, 8:254–258, 1974.

[45] Dennis Johnson. An abelian quotient of the mapping class group $\mathcal{I}_g$. *Math. Ann.*, 249(3):225–242, 1980.

[46] Ilya Kapovich, Richard Weidmann, and Alexei Miasnikov. Foldings, graphs of groups and the membership problem. *Internat. J. Algebra Comput.*, 15(1):95–128, 2005.

[47] Ki Hang Kim, L. Makar-Limanov, Joseph Neggers, and Fred W. Roush. Graph algebras. *J. Algebra*, 64(1):46–51, 1980.

[48] D. Krob and P. Lalonde. Partially commutative Lyndon words. In *STACS 93 (Würzburg, 1993)*, volume 665 of *Lecture Notes in Comput. Sci.*, pages 237–246. Springer, Berlin, 1993.

[49] Sava Krstić and James McCool. The non-finite presentability of $\mathrm{IA}(F_3)$ and $\mathrm{GL}_2(\mathbf{Z}[t, t^{-1}])$. *Invent. Math.*, 129(3):595–606, 1997.

[50] Pierre Lalonde. Bases de Lyndon des algèbres de Lie libres partiellement commutatives. *Theoret. Comput. Sci.*, 117(1-2):217–226, 1993.

[51] Pierre Lalonde. Lyndon heaps: an analogue of Lyndon words in free partially commutative monoids. *Discrete Math.*, 145(1-3):171–189, 1995.

[52] Michael R. Laurence. A generating set for the automorphism group of a graph group. *J. London Math. Soc. (2)*, 52(2):318–334, 1995.

[53] Peter Linnell, Boris Okun, and Thomas Schick. The strong Atiyah conjecture for right-angled Artin and Coxeter groups. *arXiv:1010.0606*, 2010.

[54] M. Lothaire. *Combinatorics on words*. Cambridge Univ Pr, 1997.

[55] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin, 1977.

[56] Wilhelm Magnus. Über n-dimensionale Gittertransformationen. *Acta Math.*, 64(1):353–367, 1935.

[57] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory*. Dover Publications Inc., New York, revised edition, 1976. Presentations of groups in terms of generators and relations.

[58] G. A. Margulis. *Discrete subgroups of semisimple Lie groups*, volume 17 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991.

[59] James McCool. Some finitely presented subgroups of the automorphism group of a free group. *J. Algebra*, 35:205–213, 1975.

[60] John Milnor. *Introduction to algebraic K-theory*. Princeton University Press, Princeton, N.J., 1971. Annals of Mathematics Studies, No. 72.

[61] Ashot Minasyan. Hereditary conjugacy separability of right angled Artin groups and its applications. *Groups. Geom. Dyn.*, 6(2):335–388, 2012.

[62] Nicolas Monod. An invitation to bounded cohomology. In *International Congress of Mathematicians. Vol. II*, pages 1183–1211. Eur. Math. Soc., Zürich, 2006.

[63] J Nielsen. Om Regning med ikke kommutative Faktorer og dens Anvendelse i Gruppetorien. *Mat. Tidsskrift B*, pages 77–94, 1921.

[64] Alexandra Pettet. The Johnson homomorphism and the second cohomology of IA_n. *Algebr. Geom. Topol.*, 5:725–740, 2005.

[65] Elvira Strasser Rapaport. On free groups and their automorphisms. *Acta Math.*, 99:139–163, 1958.

[66] Takao Satoh. New obstructions for the surjectivity of the Johnson homomorphism of the automorphism group of a free group. *J. London Math. Soc. (2)*, 74(2):341–360, 2006.

[67] J.P. Serre. *Trees*. Springer–Verlag, Berlin–New York, 1980.

[68] Herman Servatius. Automorphisms of graph groups. *J. Algebra*, 126(1):34–60, 1989.

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

[70] Emmanuel Toinet. Conjugacy p-separability of right-angled Artin groups and applications. *arXiv:1009.3859*, 2010.

[71] Richard D. Wade. Folding free-group automorphisms. *arXiv:1111.6794*, 2011.

[72] Richard D. Wade. Johnson homomorphisms and actions of higher-rank lattices on right-angled Artin groups. *arXiv:1101.2797*, 2011.

[73] Richard D. Wade. The lower central series of a right-angled Artin group. *arXiv:1109.1722*, 2011.

[74] J. H. C. Whitehead. On equivalent sets of elements in a free group. *Ann. of Math. (2)*, 37(4):782–800, 1936.

[75] Robert J. Zimmer. *Ergodic theory and semisimple groups*, volume 81 of *Monographs in Mathematics*. Birkhäuser Verlag, Basel, 1984.

[76] Robert J. Zimmer. Actions of semisimple groups and discrete subgroups. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Berkeley, Calif., 1986)*, pages 1247–1258, Providence, RI, 1987. Amer. Math. Soc.