Between buildings and free factor complexes

Benjamin Brück

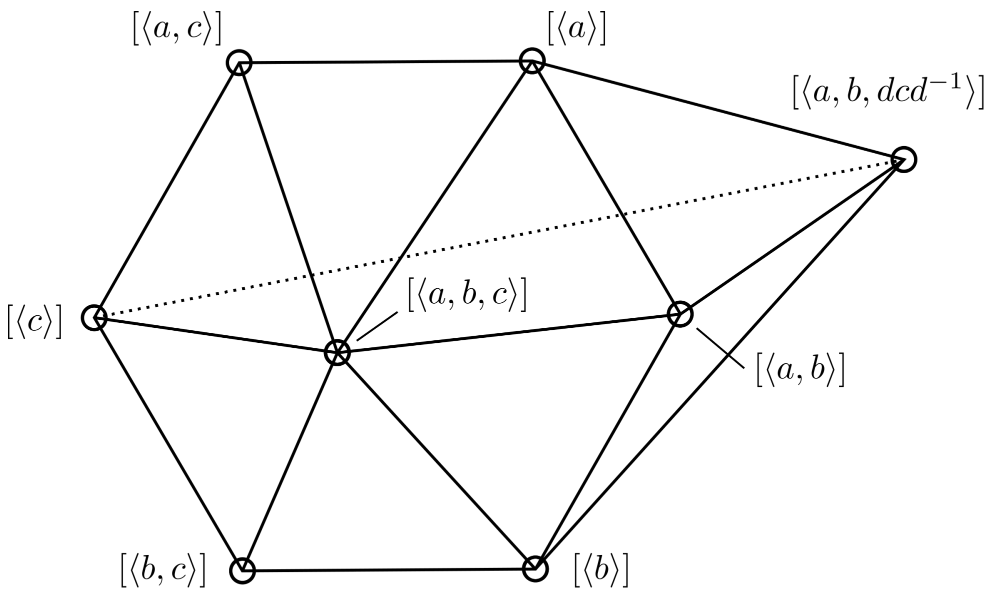

Fakultät für Mathematik
Universität Bielefeld
Dissertation
September 2019

Abstract

The main achievement of this thesis is the construction of a new family of simplicial complexes interpolating between Tits buildings and free factor complexes. For every finite graph Γ, we obtain a simplicial complex $\mathcal{CC}$ associated to the outer automorphism group of the right-angled Artin group A_Γ. These complexes are defined using the intersection patterns of cosets of parabolic subgroups. Each of them is homotopy Cohen–Macaulay and in particular homotopy equivalent to a wedge of d-spheres. The dimension d can be read off from the defining graph Γ and provides a new invariant for the automorphism group of A_Γ.

In order to deduce this and further properties of $\mathcal{CC}$, we introduce new methods for studying the topology of coset complexes and coset posets, refine the decomposition sequence for automorphism groups of right-angled Artin groups established by Day–Wade and study the asymptotic geometry of Culler–Vogtmann Outer space. In particular, we show that the simplicial boundary of the Outer space of the free group F_n can be described in terms of complexes of free factors of F_n and study the connectivity properties of these complexes.

Figure on title page: A part of the free factor complex $\mathcal{F}_4$ associated to the free group $F_4 = \langle a, b, c, d \rangle$.

Für Lara

Acknowledgements

This work would not have been possible without the help of all those who contributed to it, directly or indirectly, over the last few years. I am grateful to my advisor Kai-Uwe Bux for his guidance during this and other projects and for being receptive to my own plans and suggestions. Many thanks are also due to Dawid Kielak for helpful advice on mathematics and everything that comes along with it and to Yuri Santos Rego for many discussions on mathematics and on life. I am indebted to Radhika Gupta for letting me enter into what was originally her project alone and which has since become an important part of this thesis. I have thoroughly enjoyed our work together. I would like to warmly thank Ric Wade for many helpful comments at various stages of this project and for hosting me for a two-month visit in Oxford in Autumn 2018. Thanks are also due to Herbert Abels and Stephan Holz for interesting conversations about higher generation and coset complexes.

This work profited greatly from the good atmosphere, both mathematical and social, that I encountered in Bielefeld. I would like to thank all who made my time here so enjoyable.

I am grateful to Radhika Gupta, Christoph Hilmes, Dawid Kielak, Yuri Santos Rego and in particular Alastair Litterick for reading first drafts of this text.

Lastly, I would like to thank my family for always supporting me and my interests and Lara for always being there for me and for listening (at least sometimes) to my explanations of my mathematical problems and discoveries.

Contents

List of Figures

Chapter 1

Introduction

Automorphism groups of free groups have been studied since the early twentieth century. Right from the beginning, a combination of algebraic, combinatorial and topological methods have been used in their study (see the work of Nielsen [Nie17], [Nie24] and Whitehead [Whi36]). A lot of the modern treatment of these groups is motivated by similarities with arithmetic groups. These similarities come from the observation that $\mathrm{GL}_n(\mathbb{Z})$, the general linear group over the integers, can also be seen as the automorphism group of the free *abelian* group $\mathbb{Z}^n$. Over the last decades, many ideas and concepts from the rich, well-developed theory of arithmetic groups have been adapted and generalised to the setting of automorphisms of free groups. Most notably, Culler and Vogtmann defined an analogue of symmetric spaces called *Outer space* [CV86]. Since its introduction in the eighties, this space has played an important role in the investigation of homological and geometric properties of automorphism groups of free groups.

Given the similarities between these groups and arithmetic groups, it is natural to consider classes of groups which sit in between them—such a class is formed by the automorphism groups of right-angled Artin groups (RAAGs). Given a simplicial graph Γ, the associated RAAG is the group A_Γ generated by the vertex set of Γ subject to the relations $[v, w] = 1$ whenever v and w are adjacent. If Γ is a discrete graph (no edges) then A_Γ is a free group, whereas if Γ is complete then the corresponding RAAG is a free abelian group; for arbitrary graphs, A_Γ is seen as an interpolation between these two extremal cases. Accordingly, its automorphism group is often seen as interpolating between the arithmetic group $\mathrm{GL}_n(\mathbb{Z})$ and the automorphism group of the free group. RAAGs attracted a lot of interest in recent years and, through their connection to special cube complexes, played an important role in the solution of the virtual Haken conjecture (see the work of Haglund–Wise [HW08], [HW12] and Agol [Ago13]). Over the last years, viewing automorphism groups of RAAGs as an interpolating class between automorphism groups of free groups and arithmetic groups has served both as a motivation for studying this class and as a source of techniques for improving our understanding of it (see e.g. the work of Charney–Crisp–Vogtmann [CCV07], Hensel–Kielak [HK18], Charney–Stambaugh–Vogtmann [CSV17], Guirardel–Sale [GS18] and Day–Wade [DW]).

The present thesis contributes to this programme by providing a new structure which generalises well-studied complexes associated to arithmetic groups and automorphism groups of free groups. On the arithmetic side, we have the Tits building associated to $\mathrm{GL}_n(\mathbb{Q})$. This can be defined as the order complex of the partially ordered set (poset) of proper subspaces of $\mathbb{Q}^n$, ordered by inclusion, and is homotopy equivalent to a wedge of $(n-2)$-spheres (this is the Solomon–Tits Theorem [Sol69]). On the side of $\mathrm{Aut}(F_n)$, the automorphism group of the free group on n generators, there is the free factor complex. This is defined as the order complex of the poset of proper free factors of F_n, ordered by inclusion (a free factor being a subgroup A such that F_n can be written as a free product $F_n = A * B$). This complex was defined by Hatcher and Vogtmann [HV98b] who also showed that it is homotopy equivalent to a wedge of $(n-2)$-spheres. The aim and main achievement of this thesis is to construct for every graph Γ a simplicial complex $\mathcal{CC}$ associated to the automorphism group of A_Γ, which interpolates between these two structures.

For our construction, we use the concept of *coset complexes* which are defined as follows: Let G be a group and $\mathcal{H}$ a family of subgroups. The coset complex of G with respect to $\mathcal{H}$, denoted by $\mathrm{CC}(G, \mathcal{H})$, is the simplicial complex whose vertices are the cosets gH, with $g \in G$ and $H \in \mathcal{H}$, and where a collection of such cosets forms a simplex if and only if its intersection is non-empty. Left-multiplication induces a natural action of G on this complex. Both buildings and the free factor complex can be described as coset complexes with respect to families of so-called *parabolic subgroups*. The complex $\mathcal{CC}$ we define is a coset complex related to the automorphism group of a RAAG.

Let $\mathrm{Out}(A_\Gamma)$ be the outer automorphism group of A_Γ, i.e. the quotient of its automorphism group by the group of conjugations. Instead of looking at $\mathrm{Out}(A_\Gamma)$ itself, we will work with its finite index subgroup $O := \mathrm{Out}^0(A_\Gamma)$, called the *pure outer automorphism group*. These groups coincide if A_Γ is free or free abelian. For further details, see Section 5.1.1. Given O, we will define a family of *maximal standard parabolic subgroups* $\mathcal{P}(O)$. The complex $\mathcal{CC}$ we consider now is the coset complex $\mathrm{CC}(O, \mathcal{P}(O))$. Our main result about its structure is the following:

Theorem A. *The complex $\mathcal{CC} := \mathrm{CC}(O, \mathcal{P}(O))$ is homotopy equivalent to a wedge of spheres of dimension $|\mathcal{P}(O)| - 1$.*

In fact, we will see that $\mathcal{CC}$ is even *Cohen–Macaulay*, which provides further information about its local homology. We define the *rank* of the group O by $\mathrm{rk}(O) := |\mathcal{P}(O)|$. It seems to be an interesting invariant of O, which has, to the best of the author's knowledge, not been studied in the literature so far.

The main ingredient for the proof of Theorem A is an inductive procedure first used by Charney–Crisp–Vogtmann [CCV07] and further developed by Day–Wade [DW]. This procedure allows one to decompose O using short exact sequences into basic building blocks which consist of free abelian groups, $\mathrm{GL}_n(\mathbb{Z})$ and so-called Fouxe-Rabinovitch groups, which are groups of certain automorphisms of free products. The aim of applying this induction to $\mathcal{CC}$ and ultimately proving Theorem A is the guiding line of this thesis. The additional results that we obtain are woven around this central thread. We now give a brief outline of each chapter of this thesis, highlighting main results and how they contribute towards proving Theorem A.

Chapter 2: Preliminaries on (poset) topology

In this work, various methods are used in order to deduce connectivity properties of simplicial complexes and determine their homotopy types. However, the point of view that is predominant is to describe them as the order complexes of appropriate posets and to use tools from *poset topology* in order to understand their homology and homotopy groups. Chapter 2 sets up the notation and collects the necessary basics that are needed for this. It also contains some general topological background, in particular on Cohen–Macaulay complexes. Most of this can be found in several other texts. An exception to this is the Quillen-type fibre Lemma 2.3. Although its proof is fairly standard, it has to the best of the author's knowledge not appeared in the literature before.

Chapter 3: Coset complexes

Abels and Holz in [AH93] defined the notion of "higher generation" of a group G by a family of subgroups $\mathcal{H}$. We say that $\mathcal{H}$ is *n-generating* for G if $\mathrm{CC}(G, \mathcal{H})$ is $(n-1)$-connected. Higher generation can be interpreted as an answer to the question "How much information about G is contained in $\mathcal{H}$?" In Chapter 3, we study coset complexes and higher generation for arbitrary pairs $(G, \mathcal{H})$. Our first result here is Theorem 3.11, which gives a criterion for obtaining higher generating families from group actions on Cohen–Macaulay complexes. As an application of this, we show that the family of Levi subgroups in groups with a BN-pair is highly generating (see Corollary 3.18).

We also observe that a theorem due to Walker (Theorem 2.15) implies the following characterisation of coset complexes which satisfy the Cohen–Macaulay property:

Theorem B. *Let G be a group and $\mathcal{H}$ a finite family of subgroups of G. Then $\mathrm{CC}(G, \mathcal{H})$ is homotopy Cohen–Macaulay if and only if every $\mathcal{H}' \subseteq \mathcal{H}$ is $(|\mathcal{H}'| - 1)$-generating for G.*

We then generalise a theorem of Brown [Bro00] regarding the behaviour of coset complexes under short exact sequences to the following:

Theorem C. *Let G be a group, $\mathcal{H}$ a family of subgroups of G and $N \triangleleft G$ a normal subgroup. If $\mathcal{H}$ is strongly divided by N, there is a homotopy equivalence*

$$\mathrm{CC}(G, \mathcal{H}) \simeq \mathrm{CC}(G/N, \overline{\mathcal{H}}) * \mathrm{CC}(N, \mathcal{H} \cap N).$$

Here, $*$ denotes the join on geometric realisations, $\overline{\mathcal{H}}$ and $\mathcal{H} \cap N$ are certain families of subgroups of G/N and N, respectively, and being *strongly divided by* N is a compatibility condition on the family $\mathcal{H}$. (For the definitions, see Section 3.3; for an explicitly stated special case of Theorem C, see Corollary 3.30.)

Note that if two spaces X and Y are homotopy equivalent to wedges of spheres, then so is their join $X * Y$. Thus, combining Theorem C with the decomposition sequence of Day–Wade, we are able to reduce Theorem A to the cases where O is either isomorphic to $\mathrm{GL}_n(\mathbb{Z})$ or a Fouxe-Rabinovitch group. In the former case, the result follows from the Solomon–Tits Theorem. In the latter case, we are lead to study relative versions of free factor complexes; this is done in Chapter 4.

Chapter 4: Homotopy type of the complex of free factors

As mentioned above, Hatcher–Vogtmann [HV98b] defined the free factor complex as the order complex of the poset of proper free factors of F_n and showed that it is homotopy equivalent to a wedge of $(n-2)$-spheres. This complex is equipped with a natural action of $\mathrm{Aut}(F_n)$. Following a general trend towards the study of $\mathrm{Out}(F_n)$ (see Remark 4.1), the following version of this complex has become more popular in recent years: Let $\mathcal{F}_n$ be the poset of all *conjugacy classes* of proper free factors of F_n, ordered by inclusion of representatives. The order complex of $\mathcal{F}_n$ comes with an action of $\mathrm{Out}(F_n)$ and will also be called "free factor complex" in what follows. The geometry of $\mathcal{F}_n$ has been studied very well in recent years and it has been used to improve the understanding of $\mathrm{Out}(F_n)$. Most notably, Bestvina and Feighn in [BF14] showed that $\mathcal{F}_n$ is Gromov-hyperbolic, in analogy to Masur–Minsky's hyperbolicity result for the curve complex of a surface [MM99]. A question that remained open, however, is whether the topology of $\mathcal{F}_n$ looks similar to the one of its $\mathrm{Aut}(F_n)$ counterpart. In Chapter 4, we answer this affirmatively by showing:

Theorem D. *The free factor complex $\mathcal{F}_n$ is homotopy equivalent to a wedge of spheres of dimension $n-2$.*

In fact, we prove a more general result which also applies to relative versions of $\mathcal{F}_n$ associated to Fouxe-Rabinovitch groups (see Definition 4.3); these appear as base cases of the induction we use in order to establish Theorem A.

A key idea in the proof of Theorem D is to show that $\mathcal{F}_n$ can be seen as a substructure of the simplicial boundary $\partial_s \mathcal{CV}_n$ of Culler–Vogtmann Outer space (for the definitions, see Section 4.1.4). This leads us to study this boundary in general. The results here can be summarised as follows.

Theorem E. *The free factor complex $\mathcal{F}_n$ is homotopy equivalent to a subspace $b\mathcal{FS}^1$ of the simplicial boundary $\partial_s \mathcal{CV}_n$ of Outer space. The entire boundary $\partial_s \mathcal{CV}_n$ is homotopy equivalent to the complex $\mathcal{FF}_n$ of free factor systems of F_n. Furthermore, $\mathcal{FF}_n$ is $(n-2)$-connected.*

The complex $\mathcal{FF}_n$ of free factor systems was defined by Handel and Mosher in [HM] (see Definition 4.2). In fact, $\mathcal{F}_n$ and $\mathcal{FF}_n$ are quasi-isometric to each other [HM, Proposition 6.3]. However, Theorem D and Theorem E imply that they are not homotopy equivalent.

The main technical step towards establishing the connectivity results for $\mathcal{F}_n$ and $\mathcal{FF}_n$ is to show that relative versions of the free splitting complex are contractible. The free splitting complex $\mathcal{FS}$ was (in its non-relative version) introduced and shown to be contractible by Hatcher [Hat95]. Relative versions of it (see Section 4.1.1) were defined by Handel–Mosher who showed that these complexes are non-empty, connected and hyperbolic [HM]. The third main achievement of Chapter 4 is the following extension of their results:

Theorem F. *For every finitely generated group A and every free factor system $\mathcal{A}$ of A, the relative free splitting complex $\mathcal{FS}(A, \mathcal{A})$ is contractible.*

Chapter 5: A Cohen–Macaulay complex for Out(RAAGs)

In Chapter 5, we finally turn towards automorphism groups of RAAGs. After the necessary background on these groups, this chapter contains a description and several refinements of the inductive procedure of Charney–Crisp–Vogtmann and Day–Wade. Following this induction, one is led to study groups of the form $O = \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$, which are relative versions of $\mathrm{Out}^0(A_\Gamma)$ (for the definitions, see Section 5.1). This is why we prove all of our results (and in particular Theorem A) in this general setting. To do this, we first define the set $\mathcal{P}(O)$ of maximal standard parabolic subgroups and the rank $\mathrm{rk}(O)$. We then combine Theorem C and Theorem D in order to inductively prove Theorem A. Using Theorem B, we deduce that $\mathcal{CC} = \mathrm{CC}(O, \mathcal{P}(O))$ is Cohen–Macaulay. As a consequence of these results, we obtain higher generating families of subgroups for O and are able to give presentations of O in terms of the parabolic subgroups (Corollary 5.32 and Corollary 5.33).

Furthermore, we show that $\mathcal{CC}$ has the following properties which indicate that it is a reasonable analogue of Tits buildings and free factor complexes:

Properties of $\mathcal{CC}$

- *Building.* If $O = \mathrm{GL}_n(\mathbb{Z})$, the complex $\mathcal{CC}$ is isomorphic to the building associated to $\mathrm{GL}_n(\mathbb{Q})$. (Proposition 3.15)
- *Free factor complex.* If $O = \mathrm{Out}(F_n)$, the complex $\mathcal{CC}$ is isomorphic to the free factor complex $\mathcal{F}_n$. (Proposition 4.7)
- *Cohen–Macaulayness.* $\mathcal{CC}$ is homotopy Cohen–Macaulay and in particular a chamber complex. (Theorem 5.29)
- *Facet-transitivity.* Any maximal simplex of $\mathcal{CC}$ forms a fundamental domain for the action of O. (Section 3.1.3)
- *Stabilisers.* The vertex stabilisers of this action are exactly the conjugates of the elements of $\mathcal{P}(O)$. Stabilisers of higher-dimensional simplices are given by the intersections of such conjugates and can be seen as parabolic subgroups of lower rank. (Section 5.4.2)
- *Parabolics as relative automorphism groups.* Every maximal standard parabolic $P \in \mathcal{P}(O)$ is itself a relative automorphism group of the form $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ and $\mathrm{rk}(P) = |\mathcal{P}(P)| = \mathrm{rk}(O) - 1$. (Proposition 5.31)
- *Rank via Weyl group.* Similar to a group with BN-pair, the rank $\mathrm{rk}(O)$ is equal to the rank of a naturally defined Coxeter subgroup $\mathrm{Aut}^0(\Gamma) \leq O$. (Corollary 5.36)
- *Direct and free products.* The construction is well-behaved under taking direct and free products of the underlying RAAGs, i.e. under passing from $\mathrm{Out}^0(A_\Gamma)$ to $\mathrm{Out}^0(A_\Gamma \times A_{\Gamma'})$ or to $\mathrm{Out}^0(A_\Gamma * A_{\Gamma'})$. (Section 5.3.2)

Published parts of this work

This thesis essentially consists of the material of three articles written by the author. For the present text, they have been slightly modified in order to improve exposition. The first one is called "Higher generating subgroups and Cohen–Macaulay complexes" [Brüb] and is accepted for publication in the Proceedings of the Edinburgh Mathematical Society. Its content is given by Lemma 2.13 and

Section 3.2 of this text. The second article is joint work with Radhika Gupta and has the title "Homotopy type of the complex of free factors of a free group" [BG]. It is available on the arXiv and submitted to a journal but has as of today (September 2019) not been accepted for publication. This work is presented in Chapter 4. The third article "Between buildings and free factor complexes: A Cohen–Macaulay complex for Out(RAAGs)" [Brüa] is also available on the arXiv but has not been published in a journal yet. It forms Section 3.3 and Chapter 5 of this thesis.

Chapter 2

Preliminaries on (poset) topology

This chapter collects topological background material which is used throughout this thesis. Section 2.1 and Section 2.2 introduce notation and preliminaries on partially ordered sets and their topology. In particular, several Quillen-type fibre theorems are mentioned; most of them are well-known, one (Lemma 2.3) is original to this work. Section 2.3 and Section 2.4 contain the Nerve Theorem and the theorems of Whitehead and Hurewicz. In Section 2.5, Cohen–Macaulay complexes are defined and some consequences of the definition are derived.

2.1 Posets and their realisations

Let $P = (P, \leq)$ be a poset (partially ordered set). If $x \in P$, the sets $P_{\leq x}$ and $P_{\geq x}$ are defined by

$$P_{\leq x} := \{y \in P \mid y \leq x\}, \qquad P_{\geq x} := \{y \in P \mid y \geq x\}.$$

Similarly, one defines $P_{<x}$ and $P_{>x}$. For $x, y \in P$, the *open interval* between x and y is defined as

$$(x, y) := \{z \in P \mid x < z < y\}.$$

A *chain of length* l in P is a totally ordered subset $x_0 < x_1 < \ldots < x_l$. For each poset $P = (P, \leq)$, one has an associated simplicial complex $\Delta(P)$ called the *order complex* of P. Its vertices are the elements of P and higher dimensional simplices are given by the chains of P. When we speak about the *realisation of the poset* P, we mean the geometric realisation of its order complex and denote this space by $\|P\| := \|\Delta(P)\|$. By an abuse of notation, we will attribute topological properties (e.g. homotopy groups and connectivity properties) to a poset when we mean that its realisation has these properties. We will also often identify a simplicial complex X and its geometric realisation $\|X\|$ if what is meant is clear from the context.

The *join* of two posets P and Q, denoted $P * Q$, is the poset whose elements are given by the disjoint union of P and Q equipped with the ordering extending the orders on P and Q and such that $p < q$ for all $p \in P$, $q \in Q$. The geometric realisation of the join of P and Q is homeomorphic to the topological join of their geometric realisations:

$$\|P * Q\| \cong \|P\| * \|Q\|$$

The *direct product* $P \times Q$ of two posets P and Q is the poset whose underlying set is the Cartesian product $\{(p,q) \mid p \in P,\ q \in Q\}$ and whose order relation is given by

$$(p,q) \leq_{P\times Q} (p',q') \text{ if } p \leq_P p' \text{ and } q \leq_Q q'.$$

A map $f\colon P \to Q$ between two posets is called a *poset map* if $x \leq y$ implies $f(x) \leq f(y)$. Such a poset map induces a simplicial map from $\Delta(P)$ to $\Delta(Q)$ and hence a continuous map on the realisations of the posets. It will be denoted by $\|f\|$ or just by f if what is meant is clear from the context.

2.2 Tools from poset topology

2.2.1 Fibre theorems

An important tool to study the topology of posets is given by so-called fibre lemmas comparing the connectivity properties of posets P and Q by analysing the fibres of a poset map between them. These can be seen as poset versions of the Vietoris–Begle Theorem, see [GV09, Corollary 2.4]. The first such fibre theorem appeared in [Qui73, Theorem A] and is known as Quillen's fibre lemma:

Lemma 2.1 ([Qui78, Proposition 1.6])**.** *Let $f\colon P \to Q$ be a poset map such that the fibre $f^{-1}(Q_{\leq x})$ is contractible for all $x \in Q$. Then f induces a homotopy equivalence on geometric realisations.*

The following result shows that if one is given a poset map f such that the fibres have only vanishing homotopy groups up to a certain degree, one can also transfer connectivity results between the domain and the image of f. Recall that for $n \in \mathbb{N}$, a space X is *n-connected* if $\pi_i(X) = \{1\}$ for all $i \leq n$ and X is (-1)-connected if it is non-empty.

Lemma 2.2 ([Qui78, Proposition 7.6])**.** *Let $f\colon P \to Q$ be a poset map such that the fibre $f^{-1}(Q_{\leq x})$ is n-connected for all $x \in Q$. Then P is n-connected if and only if Q is n-connected.*

For a poset $P = (P, \leq)$, let $P^{op} = (P, \leq_{op})$ be the poset that is defined by $x \leq_{op} y \Leftrightarrow y \leq x$. Using the natural identification $\Delta(P) \cong \Delta(P^{op})$, one can draw the same conclusion as in the previous lemmas if one shows that $f^{-1}(Q_{\geq x})$ is contractible or n-connected, respectively, for all $x \in Q$.

Usually, the connectivity results one can obtain using fibre lemmas is bounded above by the degree of connectivity of the fibre. The following lemma gives a sufficient condition for obtaining a slightly better degree of connectivity. It appears as Lemma 2.3 in [BG].

Lemma 2.3. *Let $f\colon P \to Q$ be a poset map where Q is $(k+1)$-connected. Assume that for all $q \in Q$, the fibre $f^{-1}(Q_{\leq q})$ is k-connected and the map $g_*\colon \pi_{k+1}(f^{-1}(Q_{\leq q})) \to \pi_{k+1}(P)$ induced by the inclusion $g\colon f^{-1}(Q_{\leq q}) \hookrightarrow P$ is trivial. Then P is $(k+1)$-connected.*

Proof. Applying Lemma 2.2, one sees that P is k-connected.

We now show that $\pi_{k+1}(P)$ also vanishes, which implies that P is in fact $(k+1)$-connected. Consider a map $i\colon S^{k+1} \to \|P\|$ from the $(k+1)$-sphere to P. Using simplicial approximation [Spa66, Chapter 3.4] we can (after possibly precomposing with a homotopy) assume that i is simplicial with respect to a simplicial structure τ on S^{k+1}. We wish to show that i extends to a map $\hat{\imath}\colon B^{k+2} \to \|P\|$, where B^{k+2} is the $(k+2)$-ball and $\hat{\imath}|_{\partial B^{k+2}} = i$.

Consider the simplicial map $h := f \circ i\colon S^{k+1} \to \|Q\|$. Since Q is $(k+1)$-connected, it extends to a map $\hat{h}\colon B^{k+2} \to \|Q\|$ such that $\hat{h}|_{\partial B^{k+2}} = h$. Simplicial approximation applied to the pair (B^{k+2}, S^{k+1}) allows us to assume that $\hat{h}$ is simplicial with respect to a simplicial structure τ' on B^{k+2} such that τ' agrees with τ on $\partial B^{k+2} = S^{k+1}$. For this, we might need to do barycentric subdivision and replace i by a homotopic map again. We now show that $\hat{h}$ lifts to a map $\tilde{h}\colon B^{k+2} \to \|P\|$ such that $\tilde{h}|_{\partial B^{k+2}} = i$ by defining $\hat{h}$ inductively on the simplices of τ'.

To start, let v be a vertex of τ'. If $v \in \tau$, then $\tilde{h}(v) := i(v)$; otherwise set $\tilde{h}(v)$ to be any vertex in $f^{-1}(\hat{h}(v))$. Now assume that for $m \leq k+1$, the map $\tilde{h}$ has been defined on every $(m-1)$-simplex σ_{m-1} in τ' such that

$$\tilde{h}(\sigma_{m-1}) \subseteq \left\| f^{-1}(Q_{\leq q_{m-1}}) \right\|,$$

where q_{m-1} is the largest vertex in $\hat{h}(\sigma_{m-1})$ and $\tilde{h}$ restricts to i on τ. Let σ_m be an m-simplex of τ'. Then $\tilde{h}(\partial\sigma_m) \subseteq f^{-1}(Q_{\leq q_m})$ and $f^{-1}(Q_{\leq q_m})$ is k-connected. Thus $\tilde{h}$ extends to σ_m such that $\tilde{h}(\sigma_m) \subseteq f^{-1}(Q_{\leq q_m})$. Now for a $(k+2)$-simplex σ and a corresponding $q \in Q$, we have $\tilde{h}(\partial\sigma) \subseteq f^{-1}(Q_{\leq q})$. Since the image $g_*(\pi_{k+1}(f^{-1}(Q_{\leq q})))$ in $\pi_{k+1}(P)$ is trivial, the map $\tilde{h}$ extends to σ. Thus we have shown that P is $(k+1)$-connected. □

2.2.2 Homotopic poset maps and monotonicity

Another standard tool which is helpful for studying the topology of posets is:

Lemma 2.4 ([Qui78, 1.3])**.** *If two poset maps $f, g\colon P \to Q$ satisfy $f(x) \leq g(x)$ for all $x \in P$, then they induce homotopic maps on geometric realisations.*

A poset map $f\colon P \to Q$ is called *monotone* if $f(x) \leq x$ for all $x \in P$ or $f(x) \geq x$ for all $x \in P$. Later on, we will mostly use the following consequence of the preceding lemma.

Corollary 2.5. *Let Q be a subposet of P and $f\colon P \to Q$ a poset map such that $f|_Q = \mathrm{id}_Q$. Then if f is monotone, it defines a deformation retraction $\|P\| \to \|Q\|$.*

Proof. Without loss of generality, assume that $f(x) \leq x$ for all $x \in P$. Let $i\colon Q \hookrightarrow P$ denote the inclusion map. Then for all $x \in P$, we have $i \circ f(x) \leq x$, so by Lemma 2.4, this composition is homotopic to the identity. As $f \circ i = \mathrm{id}_Q$, the inclusion i is a homotopy equivalence and the claim follows from [Hat02, Proposition 0.19]. □

2.2.3 Alexander duality for posets

Alexander duality allows one to compute homology groups of compact subspaces of spheres by looking at the homology of their complement. We will need the following poset version of it which is due to Stanley.

Lemma 2.6 ([Sta82], [Wac07, Theorem 5.1.1]). *Let P be a poset such that $\|P\|$ is homeomorphic to an n-sphere and let $Q \subset P$ be a subposet. Then for all i, one has*

$$\tilde{H}_i(\|Q\|\,;\mathbb{Z}) \cong \tilde{H}^{n-i-1}(\|P \setminus Q\|\,;\mathbb{Z}).$$

2.3 The nerve of a covering

The *nerve* of a family of sets $(X_i)_{i\in I}$ is the simplicial complex $\mathcal{N}(X_i)_{i\in I}$ that has vertex set I and where a finite subset $\sigma \subseteq I$ forms a simplex if and only if $\bigcap_{i\in\sigma} X_i \neq \emptyset$. The *Nerve Theorem* is another standard tool which exists in various versions. For simplicial complexes, it can be stated as follows:

Lemma 2.7 ([Bjö95, Theorem 10.6]). *Let X be a simplicial complex and $(X_i)_{i\in I}$ a family of subcomplexes such that $X = \bigcup_{i\in I} X_i$. Suppose that every non-empty finite intersection $X_{i_1} \cap \ldots \cap X_{i_k}$ is contractible. Then X is homotopy equivalent to the nerve $\mathcal{N}((X_i)_{i\in I})$.*

2.4 Spherical complexes

A topological space is *n-spherical* if it is homotopy equivalent to a wedge of n-spheres; as a convention, we consider a contractible space to be homotopy equivalent to a (trivial) wedge of n-spheres for all n and the empty set to be (-1)-spherical.

In order to deduce information about the homotopy type of a space from its homology groups, we need a corollary of the theorems of Hurewicz and Whitehead.

Theorem 2.8 (Hurewicz, [Hat02, Theorem 4.32]). *If a space X is $(n-1)$-connected, $n \geq 2$, then $\tilde{H}_i(X) = 0$ for all $0 < i < n$ and $\pi_n(X)$ is isomorphic to $H_n(X)$.*

Theorem 2.9 (Whitehead, [Hat02, Corollary 4.33]). *Let $f\colon X \to Y$ be a map between simply-connected CW-complexes such that $f_*\colon H_k(X) \to H_k(Y)$ is an isomorphism for each k. Then f is a homotopy equivalence.*

Corollary 2.10. *Let X be a simply-connected CW-complex such that*

$$\tilde{H}_i(X) = \begin{cases} \mathbb{Z}^\lambda & , i = n, \\ 0 & , \text{otherwise}. \end{cases}$$

Then X is homotopy equivalent to a wedge of λ spheres of dimension n.

Proof. By the Hurewicz Theorem, X is $(n-1)$-connected and $\pi_n(X) \cong \tilde{H}_n(X) = \mathbb{Z}^\lambda$. Now take a disjoint union $\bigsqcup_{\mu\leq\lambda} S_\mu$ of n-spheres. For each $\mu \leq \lambda$, choose a generator $S_\mu \to X$ of the μ-th summand of $\pi_n(X)$. This gives rise to a map

$f : Y \to X$ where Y is the space obtained by wedging together the S_μ along their base points. This induces an isomorphism f_* on all homology groups, so the claim follows from the Whitehead Theorem. □

In particular, it follows that an n-dimensional CW-complex is n-spherical if and only if it is $(n-1)$-connected. Sphericity is preserved under taking joins:

Lemma 2.11. *Let X and Y be CW-complexes such that X is n-spherical and Y is m-spherical. Then the join $X * Y$ is $(n+m+1)$-spherical.*

2.5 The Cohen–Macaulay property

For the remainder of this section, let $\mathbf{k}$ be a field or the ring of integers $\mathbb{Z}$.

Definition 2.12. Let X be a simplicial complex of dimension $d < \infty$. Then X is *Cohen–Macaulay over* $\mathbf{k}$ if it is $(d-1)$-acyclic over $\mathbf{k}$, i.e. $\tilde{H}_i(X, \mathbf{k}) = \{0\}$ for all $i < d$, and the link of every s-simplex is $(d-s-2)$-acyclic over $\mathbf{k}$.

X is *homotopy Cohen–Macaulay* if it is $(d-1)$-connected and the link of every s-simplex is $(d-s-2)$-connected.

The notion of Cohen–Macaulayness over $\mathbf{k}$ was introduced in the mid-70s and came up in the study of finite simplicial complexes via their Stanley-Reisner rings, see [Sta96]. The homotopical version was introduced by Quillen in [Qui78]. While it can be shown that "being Cohen–Macaulay over $\mathbf{k}$" only depends on the geometric realisation $\|X\|$ and not on its specific triangulation, the homotopical version is not a topological invariant but a property of the simplicial complex X itself. One has implications

$$\text{homotopy CM} \quad \Rightarrow \quad \text{CM over } \mathbb{Z} \quad \Rightarrow \quad \text{CM over any field } \mathbf{k},$$

which are all strict. For more details on Cohen–Macaulayness and its connections to other combinatorial properties of simplicial complexes, see [Bjö95].

An advantage of a complex that is Cohen–Macaulay over one that is merely spherical is that it allows for inductive methods using its local structure. This is also what we will make use of in the proof of the following lemma; it follows the proof of [AH93, Theorem 3.3] and appears as Lemma 2.8 in [Brüb].

Lemma 2.13. *Let X be a d-dimensional complex and let $X_s := \|X\| \setminus \|X^{(s)}\|$ denote the complement of the s-skeleton of $\|X\|$. The following holds true:*

1. *If X is Cohen–Macaulay over $\mathbf{k}$, the homology with $\mathbf{k}$-coefficients of X_s is concentrated in dimension $d-s-1$, i.e. $\tilde{H}_i(X_s, \mathbf{k})$ is trivial if $i \neq d-s-1$.*

2. *If X is homotopy Cohen–Macaulay, X_s is $(d-s-1)$-spherical.*

Proof. We prove the two statements in parallel, proceeding by induction on s. Setting $X_{-1} := \|X\|$, the statements hold for $s = -1$ as $\|X\|$ itself is assumed to be $(d-1)$-acyclic or $(d-1)$-connected, respectively. For all s, the space X_{s-1} is the union of X_s and the open s-simplices of $\|X\|$, so we will successively adjoin these simplices to X_s while keeping track of the homotopy type. Assume that we have already constructed X' as the union of X_s and a set of open s-simplices of $\|X\|$. Then for every s-simplex σ in $\|X\|$ that is not contained in X', there is an open contractible neighbourhood U of the interior of σ in $X'' := X' \cup \mathring{\sigma}$ such that $U \cap X' = U \setminus \mathring{\sigma}$ is homotopy equivalent to the link of σ in X. As X is Cohen–Macaulay, this link is $(d-s-2)$-acyclic in the homological and $(d-s-2)$-connected in the homotopical setting. This means that X'' can be constructed by gluing together X' and U, which is contractible, along the open subset $U \setminus \mathring{\sigma}$, which is $(d-s-2)$-acyclic or $(d-s-2)$-connected. Hence, the inclusion $X' \hookrightarrow X''$ induces for all $i \leq d-s-2$ an isomorphism on homology groups $\tilde{H}_i(\cdot, \mathbf{k})$ or homotopy groups $\pi_i(\cdot)$, respectively.

By induction, we can conclude that if X is Cohen–Macaulay over $\mathbf{k}$, we have $\{0\} = \tilde{H}_i(X_{s-1}, \mathbf{k}) \cong \tilde{H}_i(X_s, \mathbf{k})$ and if it is homotopy Cohen–Macaulay, we have $\{1\} = \pi_i(X_{s-1}) \cong \pi_i(X_s)$ for $i \leq d-s-2$. Noting that the complement of the s-skeleton of any simplicial complex of dimension d is homotopy equivalent to a complex of dimension $(d-s-1)$ (contract all the simplices of dimension $(s+1)$ to their barycentres), the result follows. □

A simplicial complex X is called *pure* if all of its facets, i.e. its maximal faces, have the same dimension. Such a complex is called a *chamber complex* (or *strongly connected*) if every pair of facets $\sigma, \tau \in X$ can be connected by a sequence of facets $\sigma = \tau_1, \ldots, \tau_k = \tau$ such that for all $1 \leq i \leq k$, the intersection of τ_i and τ_{i+1} is a face of codimension one. The facets of a chamber complex are also called *chambers*.

Remark 2.14. Every Cohen–Macaulay complex is pure and a chamber complex [Bjö95, Proposition 11.7]. The preceding lemma is a generalisation of this well-known fact in the following sense: Let X be pure of dimension $d \geq 1$. Define a graph Γ whose vertices are given by the facets of X and where two vertices are joined by an edge if and only if the corresponding facets intersect in a face of codimension one. The graph Γ, which is also called the *chamber graph of* X, is homotopy equivalent to the complement of the $(d-2)$-skeleton of X. Furthermore, X is a chamber complex if and only if Γ is connected, which is equivalent to $\tilde{H}_0(\Gamma) = \{0\}$. So if we assume that X is Cohen–Macaulay, Lemma 2.13 implies that it is a chamber complex.

A pure simplicial complex X of dimension d is called *coloured* (or *completely balanced*) if there is a map $c : X^{(0)} \to \{0, \ldots, d\}$ restricting to a bijection on each facet. In this setting, for each $J \subseteq \{0, \ldots, d\}$, let X_J be the induced subcomplex of X with vertex set $c^{-1}(J)$. Colourings of simplicial complexes naturally appear in the context of coset complexes, which we will study later on. For coloured simplicial complexes, the Cohen–Macaulay property has an equivalent formulation using these colour-sorted subcomplexes as the following theorem shows. As stated below, this result is due to Walker.

Theorem 2.15 ([BWW09, Theorem 5.2], [Bjö95, Theorem 11.14])**.** *Let X be a pure d-dimensional coloured complex. Then X is Cohen–Macaulay over* $\mathbf{k}$ *if and only if X_J is $(|J|-2)$-acyclic over* $\mathbf{k}$ *for every $J \subseteq \{0, \ldots, d\}$. It is homotopy Cohen–Macaulay if and only if X_J is $(|J|-2)$-connected for every $J \subseteq \{0, \ldots, d\}$.*

Chapter 3

Coset complexes

In this chapter, we define and study coset complexes and families of higher generating subgroups. Section 3.1 contains the definitions, some basic properties due to Abels and Holz and a characterisation of coset complexes due to Zaremsky. In Section 3.2, we specialise to the setting of coset complexes which satisfy the Cohen–Macaulay property. In this context, a general result on higher generation is derived (Theorem 3.11), examples related to the theory of Tits buildings are presented (Section 3.2.1 and Section 3.2.2) and a characterisation of coset complexes with this property is given (Theorem 3.20). Section 3.3 is devoted to the proof of Theorem C, which describes the behaviour of coset complexes under short exact sequences.

3.1 Definitions and basic properties

Standing assumptions Throughout this section, let G be a group, let $\mathcal{H}$ be a family of proper subgroups of G and let $\mathcal{U} := \{gH \mid g \in G,\, H \in \mathcal{H}\}$ be the collection of cosets of the subgroups from $\mathcal{H}$.

3.1.1 Coset complex and coset poset

Definition 3.1. The *coset complex* $\mathrm{CC}(G, \mathcal{H})$ is defined as the nerve $\mathcal{N}(\mathcal{U})$, i.e. the simplicial complex that has vertex set $\mathcal{U}$, and where $U_0, \ldots, U_k \in \mathcal{U}$ form a simplex if and only if $U_0 \cap \ldots \cap U_k \neq \emptyset$.

Observe that the cosets $g_0H_0, \ldots, g_kH_k$ with $g_i \in G$ and $H_i \in \mathcal{H}$ intersect non-trivially if and only if there is $g \in G$ such that

$$g_0H_0 \cap \ldots \cap g_kH_k = g(H_0 \cap \ldots \cap H_k).$$

Hence, the set of k-simplices of $\mathrm{CC}(G, \mathcal{H})$ is in bijection with the set

$$\{(g \cdot \cap\mathcal{H}', \mathcal{H}') \mid g \in G,\, \mathcal{H}' \subseteq \mathcal{H},\, |\mathcal{H}'| = k+1\}.$$

In particular, if for all $\mathcal{H}_1, \mathcal{H}_2 \subseteq \mathcal{H}$ with $\mathcal{H}_1 \neq \mathcal{H}_2$, one has $\cap\mathcal{H}_1 \neq \cap\mathcal{H}_2$, then the set of k-simplices of $\mathrm{CC}(G, \mathcal{H})$ is in bijection with

$$\{g(H_0 \cap \ldots \cap H_k) \mid g \in G,\, H_i \in \mathcal{H},\, H_i \neq H_j \text{ for } i \neq j\}.$$

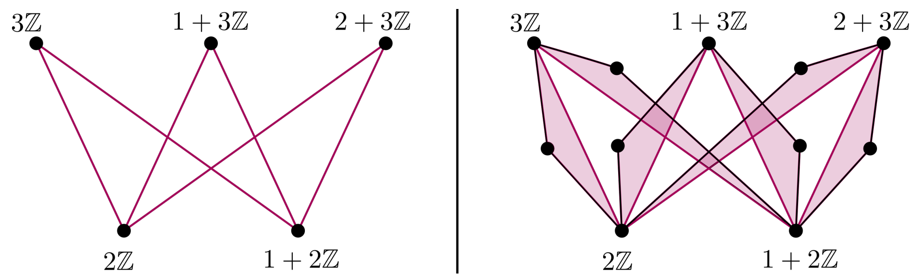

Figure 3.1: The left hand side shows $\mathrm{CC}(\mathbb{Z}, \mathcal{H})$ and $\mathrm{CP}(\mathbb{Z}, \mathcal{H})$, the right hand side shows $\mathrm{CC}(\mathbb{Z}, \widetilde{\mathcal{H}})$ and $\mathrm{CP}(\mathbb{Z}, \widetilde{\mathcal{H}})$, for the families of subgroups $\mathcal{H} = \{2\mathbb{Z}, 3\mathbb{Z}\}$ and $\widetilde{\mathcal{H}} = \{2\mathbb{Z}, 3\mathbb{Z}, 6\mathbb{Z}\}$. In both pictures, the coset poset is drawn in black and the coset complex is obtained from it by adding the magenta parts.

In this form, coset complexes were introduced by Abels and Holz in [AH93] but they appear with different names in several branches of group theory. The main motivation of Abels–Holz was to study finiteness properties of groups. Recent developments in this direction can be found in the work of Bux–Fluch–Marschler–Witzel–Zaremsky [BFM$^+$16] and Santos Rego [SR]. In [MMV98], Meier–Meinert–VanWyk used these complexes to study the BNS invariants of right-angled Artin groups. They also appear in Griffin's thesis [Gri11] on automorphisms of free products and in the thesis of Welsch [Wel18]. However, the examples that are most important to the present work are given by Tits buildings and free factor complexes (see Section 3.2.1 and Section 4.2.1).

Closely related to these complexes is the following:

Definition 3.2. The *coset poset* $\mathrm{CP}(G, \mathcal{H}) := (\mathcal{U}, \subseteq)$ is the partially ordered set consisting of the elements of $\mathcal{U}$, ordered by inclusion.

Well-known examples of coset posets are given by Coxeter and Deligne complexes [CD95]. Brown [Bro00] studied the coset poset of all subgroups of a finite group and its connection to zeta functions. Generalisations of his work can be found in the articles of Ramras [Ram05] and Shareshian–Woodroofe [SW16].

The order complex of the coset poset $\mathrm{CP}(G, \mathcal{H})$ has the same vertices as the coset complex $\mathrm{CC}(G, \mathcal{H})$ but the higher-dimensional simplices do not have to agree (see Fig. 3.1). However, if we assume that $\mathcal{H}$ is closed under finite intersections, the topology of these complexes is the same:

Lemma 3.3. *Suppose that $H_1, H_2 \in \mathcal{H}$ implies $H_1 \cap H_2 \in \mathcal{H}$. Then there is a homotopy equivalence*

$$\mathrm{CP}(G, \mathcal{H}) \simeq \mathrm{CC}(G, \mathcal{H}).$$

Proof. As $\mathcal{H}$ is closed under intersections, the intersection of two cosets from $\mathcal{U}$ is either empty or also an element of $\mathcal{U}$. The result now follows from [AH93, Theorem 1.4 (b)]. □

Let $\widetilde{\mathcal{H}}$ denote the *family consisting of all finite intersections of elements from* $\mathcal{H}$. The following was proved by Holz in his thesis [Hol85].

Lemma 3.4.

1. *Let $\mathcal{H}'$ be a family of subgroups of G with $\mathcal{H} \subseteq \mathcal{H}'$ and such that for all $H' \in \mathcal{H}'$, there is $H \in \mathcal{H}$ with $H' \subseteq H$. Then there is a homotopy equivalence* $\mathrm{CC}(G, \mathcal{H}) \simeq \mathrm{CC}(G, \mathcal{H}')$.

2. *There is a homotopy equivalence* $\mathrm{CC}(G, \mathcal{H}) \simeq \mathrm{CC}(G, \widetilde{\mathcal{H}})$.

Remark 3.5. The preceding lemmas imply that for any family $\mathcal{H}$ of subgroups of G, we have

$$\mathrm{CC}(G, \mathcal{H}) \simeq \mathrm{CC}(G, \widetilde{\mathcal{H}}) \simeq \mathrm{CP}(G, \widetilde{\mathcal{H}}).$$

It follows that we can always replace a coset complex by a coset poset. The advantage of this is that it allows us to apply the tools of poset topology, e.g. the Quillen fibre lemma, to study the topology of these complexes. However, the trade-off is that we have to increase the size of our family of subgroups.

3.1.2 Higher generation

We now turn our attention to coset complexes.

Definition 3.6. The *free product of $\mathcal{H}$ amalgamated along its intersections* is the group given by the presentation $\langle X \mid R \rangle$ where $X = \{x_g \mid g \in \bigcup \mathcal{H}\}$ and $R = \left\{ x_g x_h x_{gh}^{-1} \,\middle|\, \exists H \in \mathcal{H} : g, h \in H \right\}$.

Definition 3.7. We say that $\mathcal{H}$ is *n-generating* for G if $\mathrm{CC}(G, \mathcal{H})$ is $(n-1)$-connected, i.e. $\pi_i(\mathrm{CC}(G, \mathcal{H})) = \{1\}$ for all $i < n$.

The term "higher generating subgroups" was coined by Holz in [Hol85] and is motivated by the following:

Theorem 3.8 ([AH93, Theorem 2.4])**.**

1. *$\mathcal{H}$ is 1-generating if and only if $\bigcup \mathcal{H}$ generates G.*

2. *$\mathcal{H}$ is 2-generating if and only if G is the free product of $\mathcal{H}$ amalgamated along its intersections.*

Roughly speaking, the latter means that the union of the subgroups in $\mathcal{H}$ generates G and that all relations that hold in G follow from relations in these subgroups. The concept of 3-generation has a similar interpretation using identities among relations, see [AH93, 2.8].

3.1.3 Group actions and detecting coset complexes

Coset complexes are endowed with a natural action of G given by left multiplication. These complexes are highly symmetric in the sense that this action is *facet transitive*: Assume that $\mathcal{H}$ is a finite family of subgroups of G. Then $\mathrm{CC}(G, \mathcal{H})$ has dimension $|\mathcal{H}| - 1$ and $\mathcal{H}$ itself is the vertex set of a *facet*, i.e. a maximal simplex, of the coset complex. We will write this facet as $C_{\mathcal{H}}$. This (and hence any other) facet is a *fundamental domain* for the action of G; this means that for all $0 \leq k \leq |\mathcal{H}| - 1$, the set of k-faces of $C_{\mathcal{H}}$ contains exactly one element of each G-orbit of k-simplices of $\mathrm{CC}(G, \mathcal{H})$. The following converse of this observation is due to Zaremsky.

Proposition 3.9 (see [BFM+16, Proposition A.5]). *Let G be a group acting by simplicial automorphisms on a simplicial complex X, with a single facet C as a fundamental domain and assume that for all vertices $v \neq v'$ of C, we have* $\operatorname{Stab}_G(v) \neq \operatorname{Stab}_G(v')$. *Let*

$$\mathcal{P} \coloneqq \{\operatorname{Stab}_G(v) \mid v \text{ is a vertex of } C\}.$$

Then the map

$$\begin{aligned} \psi\colon \operatorname{CC}(G, \mathcal{P}) &\to X \\ g \operatorname{Stab}_G(v) &\mapsto g.v \end{aligned}$$

is an isomorphism of simplicial G-complexes.

Note that in the original statement of Zaremsky, the technical condition $\operatorname{Stab}_G(v) \neq \operatorname{Stab}_G(v')$ for $v \neq v'$ is not mentioned, but it clearly is necessary for the statement to be true in this form (and for Zaremsky's proof to work).

3.2 Higher generation and Cohen–Macaulay complexes

We now want to use Proposition 3.9 in order to obtain higher generating families of subgroups for groups acting on Cohen–Macaulay complexes. The content of this section was published in [Brüb]. We start with a rather technical observation.

Lemma 3.10. *Let G be a group acting by simplicial automorphisms on a Cohen–Macaulay complex X, with a single facet $C = \{v_0, \dots, v_d\}$ as fundamental domain. Then the following holds true: For all $0 \leq k \leq d$ and all k-dimensional faces σ, σ' of C,*

$$\operatorname{Stab}_G(\sigma) = \operatorname{Stab}_G(\sigma') \neq G \tag{3.1}$$

implies $\sigma = \sigma'$.

Proof. Assume that for σ and σ' as above, Eq. (3.1) holds. The action of G determines a colouring $c\colon X^{(0)} \to \{0, \dots, d\}$ by assigning to an element $g.v_i$ in the orbit of v_i the colour i. Let $J, J' \subseteq \{0, \dots, d\}$ be the sets of colours that are present in σ and σ', respectively. By Theorem 2.15, the colour-sorted subcomplex $X_{J \cup J'}$ is again Cohen–Macaulay of dimension $|J \cup J'| - 1$. In particular, it is a chamber complex [Bjö95, Proposition 11.7]. The group G acts facet transitively on $X_{J \cup J'}$ and there are at least two distinct facets because $\operatorname{Stab}_G(\sigma) = \operatorname{Stab}_G(\sigma \cup \sigma') \neq G$. Hence, there must be a facet $\tau \neq \sigma \cup \sigma'$ of $X_{J \cup J'}$ that intersects $\sigma \cup \sigma'$ in a face of dimension $|J \cup J'| - 2$. Then either $\sigma = \sigma'$ or τ contains one of them. Without loss of generality, assume that $\sigma \subseteq \tau$. Because G acts facet transitively, we have $\tau = g.(\sigma \cup \sigma')$ for some $g \in G$. But as σ is contained in τ, this implies $g \in \operatorname{Stab}_G(\sigma) = \operatorname{Stab}_G(\sigma \cup \sigma')$, so it follows that $\tau = \sigma \cup \sigma'$, which is a contradiction. □

The main result of this section is as follows.

Theorem 3.11. *Let G be a group acting by simplicial automorphisms on a simplicial complex X, with a single facet C as fundamental domain and such that for every vertex $v \in C$, the stabiliser $\mathrm{Stab}_G(v)$ is a proper subgroup of G. If X is homotopy Cohen–Macaulay and has dimension d, the set*

$$\mathcal{P}_k := \{\mathrm{Stab}_G(\sigma) \mid \sigma \text{ is a } k\text{-dimensional face of } C\}$$

is $(d-k)$-generating for all $0 \leq k \leq d$. Furthermore, the corresponding coset complex $\mathrm{CC}(G, \mathcal{P}_k)$ is $(d-k)$-spherical.

Proof. By Lemma 3.10, the assumption $\mathrm{Stab}_G(v) < G$ for all $v \in C$ implies that distinct vertices of C have distinct stabilisers. Hence by Proposition 3.9, we can identify X with the coset complex $\mathrm{CC}(G, \mathcal{P}_0)$.

As C is a fundamental domain for the action of G, the stabiliser of a k-face σ of C is equal to the intersection of the stabilisers of all the vertices of σ. Hence, the elements of $\mathcal{P}_k$ are given by all the intersections of $(k+1)$ pairwise distinct elements from $\mathcal{P}_0$.

The comments after Definition 3.1 and the preceding Lemma 3.10 imply that the vertices of $\mathrm{CC}(G, \mathcal{P}_k)$ are in one-to-one correspondence with the k-simplices of $\mathrm{CC}(G, \mathcal{P}_0) \cong X$. Moreover, a set of vertices in $\mathrm{CC}(G, \mathcal{P}_k)$ forms a simplex if and only if the corresponding k-simplices in X are all faces of one common facet. It follows that the geometric realisation $\|\mathrm{CC}(G, \mathcal{P}_k)\|$ is homotopy equivalent to $\|Y\|$, where Y is the induced subcomplex of the barycentric subdivision $\mathcal{B}(X)$ whose vertices are the barycentres of all simplices of X that have dimension greater or equal to k.

The complex $\|Y\|$ is homotopy equivalent to the complement of the $(k-1)$-skeleton of $\|X\|$. As X is Cohen–Macaulay, we can use Lemma 2.13 to conclude that $\mathrm{CC}(G, \mathcal{P}_k)$ is $(d-k)$-spherical. This finishes the proof. □

In what follows, we give two rather immediate applications of Theorem 3.11. Both of them come from the theory of buildings. Definitions and background material needed for these subsections can be found in [AB08]. A third application will be given in Section 5.4.2 where we will use Theorem 3.11 to find higher generating families of subgroups for automorphism groups of right-angled Artin groups; this in particular includes the case of $\mathrm{Out}(F_n)$, the outer automorphism group of the free group (see Remark 4.42).

3.2.1 Parabolic subgroups and buildings

Our first application recovers [AH93, Theorem 3.3] of Abels and Holz. We will be brief here and refer to their text for further details. The following result is originally due to Solomon [Sol69].

Theorem 3.12 (Solomon–Tits)**.** *Let Δ be a building of rank r. Then Δ is $(r-1)$-spherical. It is contractible if and only if its Weyl group is infinite.*

The link of a k-simplex in Δ is again a building of rank $r-k-1$, so the Solomon–Tits Theorem already implies that Δ *is homotopy Cohen–Macaulay.*

Now let G be a group with a BN-pair, denote by Δ the corresponding building and by $Ch(\Delta)$ the set of its chambers. The action of G is transitive on the chambers of Δ, so we can apply Theorem 3.11 to deduce that for any choice of chamber $C \in Ch(\Delta)$, the family $\mathcal{P}_k$ of stabilisers of the k-dimensional faces of

C is $(r-1-k)$-generating for G. If we take C to be the "fundamental" chamber associated to the Borel subgroup B, these stabilisers are exactly the *standard parabolic subgroups of rank* $r-k-1$. Hence we get:

Theorem 3.13 ([AH93, Theorem 3.3])**.** *The family of rank-m standard parabolic subgroups is m-generating for G.*

Here, 2-generation was also already shown by Tits [Tit74, Section 13].

The building associated to $\mathrm{GL}_n(\mathbb{Z})$

A special case that will become important later on is the following: Let $G = \mathrm{GL}_n(\mathbb{Q})$. A BN-pair in this group is given by choosing $B \leq G$ as the group of upper-triangular matrices and $N \leq G$ as the monomial group, i.e. the group of all matrices in G with exactly one non-zero entry in every row and column. In what follows, we will call the corresponding building Δ the *building associated to* $\mathrm{GL}_n(\mathbb{Q})$. The associated Weyl group is the symmetric group $\mathrm{Sym}(n)$, which is finite and has rank $r = n-1$. Thus, by the Solomon–Tits Theorem, the building Δ is homotopy equivalent to a (non-trivial) wedge of $(n-2)$-spheres.

This building can be described as the order complex of the poset $\mathcal{Q}$ of proper (i.e. non-trivial and not equal to $\mathbb{Q}^n$) subspaces of $\mathbb{Q}^n$, ordered by inclusion, see e.g. [AB08, Chapter 6.5]. It is well-known that Δ can equivalently be seen as the coset complex of $\mathrm{GL}_n(\mathbb{Q})$ with respect to the family of maximal standard parabolic subgroups. We will now show that it can also be described as a coset complex of $\mathrm{GL}_n(\mathbb{Z})$.

A subgroup $A \leq \mathbb{Z}^n$ is called a *direct summand* if there is $B \leq \mathbb{Z}^n$ such that $\mathbb{Z}^n = A \oplus B$. We say that a direct summand A is *proper* if it is neither trivial nor equal to $\mathbb{Z}^n$. Let $\mathcal{Z}$ be the poset of all proper direct summands of $\mathbb{Z}^n$, ordered by inclusion. The group $\mathrm{GL}_n(\mathbb{Z})$ acts naturally on $\mathcal{Z}$.

Fix a basis $\{e_1, \ldots, e_n\}$ of $\mathbb{Z}^n$ and for all $1 \leq i \leq n-1$, set $S_i \coloneqq \langle e_1, \ldots, e_i \rangle$. Note that one has $S_i \in \mathcal{Z}$ for all i and define

$$P_i \coloneqq \mathrm{Stab}_{\mathrm{GL}_n(\mathbb{Z})}(S_i)$$

to be the stabiliser of S_i under the action of $\mathrm{GL}_n(\mathbb{Z})$ on $\mathcal{Z}$. We define the set of *maximal standard parabolic subgroups of* $\mathrm{GL}_n(\mathbb{Z})$ as

$$\mathcal{P} = \mathcal{P}(\mathrm{GL}_n(\mathbb{Z})) \coloneqq \{P_i \mid 1 \leq i \leq n-1\}.$$

Remark 3.14. We called the elements of $\mathcal{P}$ the maximal *standard* parabolic subgroups of $\mathrm{GL}_n(\mathbb{Z})$ to match the usual convention where an arbitrary parabolic subgroup is defined as the conjugate of a standard one. We will, however, not work with non-standard parabolic subgroups in this text, thus we leave out this adjective from now on.

In terms of matrices, the maximal parabolic subgroups can be written in the form

$$P_i = \begin{pmatrix} \mathrm{GL}_i(\mathbb{Z}) & M_{i,n-i}(\mathbb{Z}) \\ 0 & \mathrm{GL}_{n-i}(\mathbb{Z}) \end{pmatrix} \leq \mathrm{GL}_n(\mathbb{Z}).$$

Proposition 3.15. *The building associated to* $\mathrm{GL}_n(\mathbb{Q})$ *is* $\mathrm{GL}_n(\mathbb{Z})$*-equivariantly isomorphic to the coset complex* $\mathrm{CC}(\mathrm{GL}_n(\mathbb{Z}), \mathcal{P})$.

Proof. Each $A \in \mathcal{Z}$ is isomorphic to $\mathbb{Z}^i$ for an integer $i := \mathrm{rk}(A) \in \{1, \ldots, n-1\}$, the rank of A. Furthermore, if $A \leq B$ in $\mathcal{Z}$, we have $\mathrm{rk}(A) \leq \mathrm{rk}(B)$ with equality if and only if A and B are equal. It follows that the maximal simplices of $\Delta(\mathcal{Z})$ are given by chains $A_1 \leq \ldots \leq A_{n-1}$, where $\mathrm{rk}(A_i) = i$. The group $\mathrm{GL}_n(\mathbb{Z})$ acts transitively on the set of all such chains and preserves the rank of each summand. Hence, the facet $S_1 \leq \ldots \leq S_{n-1}$ is a fundamental domain for this action and Proposition 3.9 implies that the order complex of $\mathcal{Z}$ is $\mathrm{GL}_n(\mathbb{Z})$-equivariantly isomorphic to $\mathrm{CC}(\mathrm{GL}_n(\mathbb{Z}), \mathcal{P})$.

On the other hand, for every subspace V of $\mathbb{Q}^n$, the intersection $V \cap \mathbb{Z}^n$ is a direct summand of $\mathbb{Z}^n$. Hence, sending V to $V \cap \mathbb{Z}^n$ defines a poset map $f\colon \mathcal{Q} \to \mathcal{Z}$ and, using the fact that every subspace of $\mathbb{Q}^n$ has a basis in $\mathbb{Z}^n$, one has $\mathrm{rk}(f(V)) = \dim(V)$. In the opposite direction, we have a poset map $g\colon \mathcal{Z} \to \mathcal{Q}$ defined by sending $A \leq \mathbb{Z}^n$ to its $\mathbb{Q}$-span $\langle A \rangle_{\mathbb{Q}}$. Again using standard linear algebra, one sees that $\dim(g(A)) = \mathrm{rk}(A)$. Now f and g are isomorphisms that are inverse to one another: Indeed, it is clear that $g \circ f(V) = V$ because V has a basis in $\mathbb{Z}^n$ and it is also clear that $f \circ g(A)$ contains A. As both A and $f \circ g(A)$ are direct summands of the same rank, the claim follows. It is easy to see that these isomorphisms are $\mathrm{GL}_n(\mathbb{Z})$-equivariant. □

3.2.2 Levi subgroups and the opposition complex

To show that the families of standard parabolic subgroups in a group G with a BN-pair are higher generating, we only needed to use chamber-transitivity of the action of G on the associated building. However, this action is known to satisfy stronger transitivity conditions; we will exploit them to find other families of higher generating subgroups in this subsection.

Let Δ be a spherical building. The chamber distance $d(-,-)$ induces an *opposition relation* op between chambers of Δ which is defined by

$$C \,\mathrm{op}\, C' :\Leftrightarrow d(C, C') = \max\left\{ d(C_1, C_2) \mid C_1, C_2 \in Ch(\Delta) \right\}.$$

This opposition relation can be extended to arbitrary simplices $\sigma, \sigma' \in \Delta$ of equal dimension by saying that σ is opposite to σ' if and only if the following holds true:

> *For every chamber $C \geq \sigma$ in Δ, there is a chamber $C' \geq \sigma'$ such that $C \,\mathrm{op}\, C'$ and for every chamber $C' \geq \sigma'$, there is a chamber $C \geq \sigma$ such that $C \,\mathrm{op}\, C'$.*

Using this opposition relation, one can define a new complex from Δ as follows:

Definition 3.16. The *opposition complex* $\mathrm{Opp}(\Delta)$ is the simplicial complex whose simplices are of the form (σ, σ') with $\sigma, \sigma' \in \Delta$, $\sigma \,\mathrm{op}\, \sigma'$ and where the face relation is given by

$$(\tau, \tau') \leq (\sigma, \sigma') :\Leftrightarrow \tau \leq \sigma \text{ and } \tau' \leq \sigma'.$$

$\mathrm{Opp}(\Delta)$ has the same dimension as Δ and it was shown to be homotopy Cohen–Macaulay by von Heydebreck in [vH03]. The complex is pure and its facets are given by pairs (C, C') of opposite chambers $C, C' \in Ch(\Delta)$.

Every building Δ comes with a map

$$\delta : Ch(\Delta) \times Ch(\Delta) \to W,$$

where W is the Weyl group of Δ. This function is called the *Weyl distance function* (of Δ) and it is related to the gallery distance as follows:

$$d(C, C') = l_S(\delta(C, C')),$$

where l_S denotes the Coxeter length function on W. If a group acts by type-preserving automorphisms on Δ, we say that the action is *Weyl transitive* if for each $w \in W$, the action is transitive on the set of order pairs of chambers (C, C') with $\delta(C, C') = w$.

Theorem 3.17. *Let G be a group acting Weyl transitively by type-preserving automorphisms on a spherical building Δ of dimension d. Choose any pair (C, C') of opposite chambers $C, C' \in Ch(\Delta)$. Then the set*

$$\mathcal{P}_k := \{\mathrm{Stab}_G(\sigma) \cap \mathrm{Stab}_G(\sigma') \mid \sigma, \sigma' \text{ } k\text{-dimensional faces of } C, C';\ \sigma \operatorname{op} \sigma'\}$$

is $(d - k)$-generating for G.

Proof. As the action of G on Δ preserves distances and adjacency relations, it induces a simplicial action on $\mathrm{Opp}(\Delta)$ given by

$$g.(\sigma, \sigma') := (g.\sigma, g.\sigma').$$

We claim that the simplex $(C, C') \in \mathrm{Opp}(\Delta)$ is a fundamental domain for this action of G. Because Δ is spherical, its Weyl group W is finite and has a unique element w_S of maximal length. Hence, two chambers $D, D' \in Ch(\Delta)$ are opposite to each other if and only if $\delta(D, D') = w_S$ and by Weyl transitivity, G acts transitively on such pairs of opposite chambers. This implies that the set of vertices of (C, C') contains a representative of each G-orbit of vertices in $\mathrm{Opp}(\Delta)$. Furthermore, the type of any vertex of the chamber C is preserved by all the elements of G. Hence, no two distinct vertices of (C, C') lie in the same G-orbit which proves that this facet is indeed a fundamental domain.

As a consequence, Theorem 3.11 shows that the set $\mathcal{P}_k$ of stabilisers of k-simplices in $\mathrm{Opp}(\Delta)$ is $(d - k)$-generating. Since a k-simplex in $\mathrm{Opp}(\Delta)$ is a pair (σ, σ') of k-simplices $\sigma, \sigma' \in \Delta$, this finishes the proof. □

In particular, the conditions of the preceding theorem are fulfilled in the following situation: If G is a group having a BN-pair of rank r with finite Weyl group $W = \langle S \rangle$, it acts Weyl transitively on the associated spherical building. The chambers associated to B and $B^- = w_S B w_S$ are opposite to each other and after setting $C := B$ and $C' := B^-$, the family $\mathcal{P}_k$ defined in Theorem 3.17 is the set of *standard rank-$(r - k - 1)$ Levi subgroups.* We state this as follows:

Corollary 3.18. *Let (G, B, N, S) be a Tits system with finite Weyl group. Then the family of standard rank-m Levi subgroups is m-generating for G.*

Example 3.19. As an illustration, we spell out the following special case of this result: If Δ is the flag complex of proper subspaces of the vector space $\mathbf{k}^n$, i.e.

a building of type $\mathtt{A}_{n-1}$, the opposition complex $\mathrm{Opp}(\Delta)$ is the complex with vertex set

$$\{(U,U') \mid U,\, U' \text{ are proper subspaces of } \mathbf{k}^n \text{ and } U \oplus U' = \mathbf{k}^n\}$$

in which $(U_0, U_0'), \ldots, (U_k, U_k')$ form a simplex if and only if (possibly after reordering), one has $U_0 < U_1 < \ldots < U_k$ and $U_0' > U_1' > \ldots > U_k'$.

Let $\{e_1, \ldots, e_n\}$ be the standard basis of $\mathbf{k}^n$. The flags

$$\begin{aligned} C &:= \langle e_1 \rangle < \langle e_1, e_2 \rangle < \ldots < \langle e_1, \ldots, e_{n-1} \rangle \text{ and} \\ C' &:= \langle e_2, \ldots, e_n \rangle > \langle e_3, \ldots, e_n \rangle > \ldots > \langle e_n \rangle \end{aligned}$$

form opposite chambers of Δ. The building Δ has dimension $n-2$ and $\mathrm{GL}_n(\mathbf{k})$ acts Weyl transitively on it. The corresponding family of stabilisers $\mathcal{P}_k$ with $0 \le k \le n-3$ consists of all subgroups of the form

$$\begin{pmatrix} \mathrm{GL}_{n_1}(\mathbf{k}) & 0 & \cdots & 0 \\ 0 & \mathrm{GL}_{n_2}(\mathbf{k}) & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & \mathrm{GL}_{n_{k+2}}(\mathbf{k}) \end{pmatrix} \le \mathrm{GL}_n(\mathbf{k}).$$

So the number of blocks in the corresponding matrices is $k+2$ and the n_i are natural numbers such that $\sum_{i=1}^{k+2} n_i = n$. These are exactly the standard rank-$(n-2-k)$ Levi subgroups of $\mathrm{GL}_n(\mathbf{k})$ and by Theorem 3.17, this family is $(n-2-k)$-generating.

3.2.3 Characterisation of CM coset complexes

In this section, we characterise the class of pairs $(G, \mathcal{H})$ which can be obtained using Theorem 3.11: By Proposition 3.9, the conditions of Theorem 3.11 are fulfilled if and only if $\mathrm{CC}(G,\, \mathcal{P}_0)$ is homotopy Cohen–Macaulay. We will give an alternative characterisation of this condition for coset complexes.

Every finite-dimensional coset complex is a pure simplicial complex which can be given a colouring

$$c : \mathrm{CC}(G,\, \{H_0, \ldots, H_d\}) \to \{0, \ldots, d\}$$

by setting $c(gH_i) := i$. Hence, we obtain Theorem B as an immediate consequence of Theorem 2.15:

Theorem 3.20. *Let G be a group and $\mathcal{H}$ be a finite family of subgroups of G.*

1. $\mathrm{CC}(G,\, \mathcal{H})$ *is Cohen–Macaulay over* $\mathbf{k}$ *if and only if for all $\mathcal{H}' \subseteq \mathcal{H}$, the coset complex* $\mathrm{CC}(G,\, \mathcal{H}')$ *is $(|\mathcal{H}'|-2)$-acyclic over* $\mathbf{k}$.
2. $\mathrm{CC}(G,\, \mathcal{H})$ *is homotopy Cohen–Macaulay if and only if every $\mathcal{H}' \subseteq \mathcal{H}$ is $(|\mathcal{H}'|-1)$-generating for G.*

Being a coset complex imposes rather strong restrictions: In addition to being coloured, every such complex is endowed with a facet transitive group action. One might ask whether in this setting, Cohen–Macaulayness implies already stronger combinatorial conditions like shellability. A finite complex is *shellable* if and only if the set of its facets admits a sufficiently nice ordering, called a *shelling*; for the precise definition, see [Bjö95, Section 11.2]. In general, being shellable is strictly stronger than being homotopy Cohen–Macaulay. Buildings form a class of coset complexes which are shellable see [Bjö84]. However, the following example shows that there are also coset complexes which are Cohen–Macaulay over $\mathbb{Z}$, but are not homotopy Cohen–Macaulay and so in particular not shellable.

Let Alt_5 be the alternating group on the set $\{1, 2, 3, 4, 5\}$ and consider the following subgroups:

$$\begin{aligned} H_1 &:= \mathrm{Stab}_{\mathrm{Alt}_5}(\{2\}), \\ H_2 &:= N_{\mathrm{Alt}_5}(\,\langle(1,2,3,4,5)\rangle\,), \\ H_3 &:= N_{\mathrm{Alt}_5}(\,\langle(1,3,5)\rangle\,), \end{aligned}$$

where $\mathrm{Stab}_{\mathrm{Alt}_5}$ and N_{Alt_5} denote stabiliser and normaliser in Alt_5. The group H_1 is isomorphic to Alt_4 and H_2 and H_3 are isomorphic to the dihedral groups D_5 and D_3, respectively. Let $\mathcal{H} := \{H_1, H_2, H_3\}$. The coset complex $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})$ has dimension two and consists of 21 vertices, 80 edges and 60 two-simplices. This complex was first found by R. Oliver, an explicit description of it as a coset complex can be found in [Seg93]. For further details and a picture, see [Lut99, Section 7.3]; note that $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})$ is isomorphic to the complex N_0 in [Lut99].

Lemma 3.21. *The coset complex* $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})$ *is Cohen–Macaulay over* $\mathbb{Z}$*, but is not homotopy Cohen–Macaulay.*

Proof. In [Lut99], Lutz shows that $\|\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})\|$ is homeomorphic to a cell complex Q obtained by taking the boundary of a dodecahedron and identifying opposite pentagons by a coherent twist of $\pi/5$. The complex Q arises in triangulations of the Poincaré homology 3-sphere Σ^3. It is $\mathbb{Z}$-acyclic and one has $\pi_1(Q) \cong \pi_1(\Sigma^3)$, see [Bre72, p. 57]. As this fundamental group is non-trivial, Q and therefore $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})$ cannot be homotopy Cohen–Macaulay.

It remains to show that $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})$ is Cohen–Macaulay over $\mathbb{Z}$. By Theorem 3.20, it suffices to show that for all $\mathcal{H}' \subseteq \mathcal{H}$, the complex $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H}')$ is $(|\mathcal{H}'| - 2)$-acyclic. For $\mathcal{H}' = \mathcal{H}$, this is true as Q is $\mathbb{Z}$-acyclic and for $|\mathcal{H}'| = 1$, there is nothing to show. Hence, one only needs to check that for all two-element subsets $\mathcal{H}'$ of $\mathcal{H}$, the corresponding subcomplex of $\mathrm{CC}(\mathrm{Alt}_5, \mathcal{H})$ is connected. This can easily be verified, e.g. by using Figure 7.5 of [Lut99]. □

A further question in the same direction which might be interesting to consider is whether every coset complex that is *homotopy* Cohen–Macaulay is already shellable. A counterexample to that (if existent) would have to be a pure, completely balanced simplicial complex with a facet-transitive group action that is homotopy Cohen–Macaulay but not shellable. It seems likely that such a complex exists but the author is not aware of any examples.

3.3 Coset complexes and short exact sequences

The present section appeared as Section 3.2 in [Brüa]. We will later on study coset complexes in the setting where $G = \mathrm{Out}(A_\Gamma)$, the outer automorphism group of a right-angled Artin group. For this, we want to use the decomposition sequences of $\mathrm{Out}(A_\Gamma)$ developed in [DW]. In order to do so, we need to study the following question: If G fits into a short exact sequence, can the coset complex $\mathrm{CC}(G, \mathcal{H})$ be decomposed into "simpler" complexes related to the image and kernel of the sequence? There is a special case where this question can easily be answered:

Coset complexes and direct products Assume that we have a group factoring as a direct product $G = G_1 \times G_2$ and let $\mathcal{H}$ be a family of subgroups such that each $H \in \mathcal{H}$ contains either $\{1\} \times G_2$ or $G_1 \times \{1\}$; denote the set of those elements of $\mathcal{H}$ satisfying the former by $\mathcal{H}_1$ and the set of those satisfying the latter by $\mathcal{H}_2$. Now given $H_i, H_i' \in \mathcal{H}_i$, we have

$$\begin{aligned} &(g_1, g_2) \cdot H_i \cap (g_1', g_2') \cdot H_i' \neq \emptyset \\ \Leftrightarrow\ &(g_1, 1) \cdot H_i \cap (g_1', 1) \cdot H_i' \neq \emptyset \\ \Leftrightarrow\ &g_1 \cdot p_i(H_i) \cap g_1' \cdot p_i(H_i') \neq \emptyset, \end{aligned}$$

where p_i is the projection map $G \to G_i$. On the other hand, if we take $H_1 \in \mathcal{H}_1$ and $H_2 \in \mathcal{H}_2$, all of their cosets intersect non-trivially because

$$(g_1, g_2) \cdot H_1 = (g_1, g_2') \cdot H_1 \qquad \text{and} \qquad (g_1', g_2') \cdot H_2 = (g_1, g_2') \cdot H_2.$$

It follows that the coset complex $\mathrm{CC}(G, \mathcal{H})$ decomposes as a join

$$\mathrm{CC}(G, \mathcal{H}) \cong \mathrm{CC}(G_1, p_1(\mathcal{H}_1)) * \mathrm{CC}(G_2, p_2(\mathcal{H}_2)).$$

However, the situation becomes more complicated if we consider semi-direct products or general short exact sequences

$$1 \to N \to G \to Q \to 1.$$

[Hol85, Proposition 5.17], [Wel18, Theorem 7.3] and [Bro00, Proposition 10] contain results in this direction for the cases where every $H \in \mathcal{H}$ is a complement of N, every $H \in \mathcal{H}$ contains N and where G is a finite group and $\mathcal{H}$ is the set of *all* subgroups of G, respectively. The work in this section provides a common generalisation of all three of these results (see Theorem 3.29).

Notation and standing assumptions From now on, we will fix a normal subgroup $N \triangleleft G$ and assume that $\mathcal{H}$ is a set of proper subgroups of G. In this situation, we can write $\mathcal{H}$ as a disjoint union $\mathcal{H} = \mathcal{H}_N \sqcup \mathcal{H}^N$, where

$$\mathcal{H}_N := \{H \in \mathcal{H} \mid HN \neq G\} \text{ and } \mathcal{H}^N := \{K \in \mathcal{H} \mid KN = G\}.$$

For elements $g \in G$ and subgroups $H \leq G$ of G, let $\bar{g}$ and $\overline{H}$ denote the image of g and H in the quotient G/N, respectively.

The family $\mathcal{H}_N$ gives rise to a family of proper subgroups of G/N, denoted by

$$\overline{\mathcal{H}} := \left\{\overline{H} \mid H \in \mathcal{H}_N\right\}.$$

Similarly, $\mathcal{H}^N$ gives rise to a family of proper subgroups of N, denoted by

$$\mathcal{H} \cap N := \left\{K \cap N \mid K \in \mathcal{H}^N\right\}.$$

3.3.1 Coset posets

We start by considering the behaviour of coset posets under short exact sequences.

Definition 3.22. The family $\mathcal{H}$ of proper subgroups of G is *divided by* N if the following holds true:

1. For all $H \in \mathcal{H}_N$, one has $HN \in \mathcal{H}$.
2. For all $H \in \mathcal{H}_N$ and $K \in \mathcal{H}^N$, one has $HN \cap K \in \mathcal{H}$.

In what follows, we will use the following elementary observation several times:

Lemma 3.23. *Let $H, K \leq G$ be two subgroups of G and assume that $KN = G$. Then one has $(HN \cap K) \cdot N = HN$.*

Proof. Obviously, $(HN \cap K) \cdot N$ is contained in HN. We claim that in fact, these sets are equal. Indeed, as $KN = G$, each $hn \in HN$ can be written as $hn = kn'$ with $k \in K$ and $n' \in N$. As $k = hnn'^{-1}$, it is contained in $HN \cap K$. Hence, $hn = kn' \in (HN \cap K) \cdot N$. □

The next proposition is a generalisation of [Bro00, Proposition 10]. Our proof closely follows the ideas of Brown.

Proposition 3.24. *If $\mathcal{H}$ is divided by N, then there is a homotopy equivalence*

$$\mathrm{CP}(G, \mathcal{H}) \simeq \mathrm{CP}(G/N, \overline{\mathcal{H}}) * \mathrm{CP}(G, \mathcal{H}^N).$$

Proof. Set $C := \mathrm{CP}(G, \mathcal{H})$, $C_N := \mathrm{CP}(G, \mathcal{H}_N)$ and $C^N := \mathrm{CP}(G, \mathcal{H}^N)$. We define a map

$$f \colon C \to \mathrm{CP}(G/N, \overline{\mathcal{H}}) * C^N$$

such that f restricts to the identity on C^N and $f(gH) = \bar{g}\overline{H}$ for all $gH \in C_N$. As no coset from C^N can be contained in a coset from C_N, this map is order-preserving, i.e. a poset map. We claim that it is actually a homotopy equivalence.

For $K \in \mathrm{CP}(G/N, \overline{\mathcal{H}}) * C^N$, define

$$F := f^{-1}(\,(\mathrm{CP}(G/N, \overline{\mathcal{H}}) * C^N\,)_{\leq K})$$

to be the fibre of K with respect to f. By Lemma 2.1, it suffices to show that F is contractible.

If $K \in \mathrm{CP}(G/N, \overline{\mathcal{H}})$, this is clear: Write $K = \bar{g}\overline{H}$ such that $g \in G$, $H \in \mathcal{H}_N$. As N divides $\mathcal{H}$, the subgroup HN is contained in $\mathcal{H}$ and $g \cdot HN$ is the unique maximal element of F. This immediately implies contractibility of F.

Now assume $K \in C^N$. Using the natural action of G on these posets, we can assume that $K \in \mathcal{H}^N$. By definition of the join, the fibre F can be written as

$$F = C_N \cup C_{\leq K}.$$

The intersection $C' := C_N \cap C_{\leq K}$ is equal to $(C_N)_{\leq K}$, i.e. it is given by all cosets $gH \subseteq K$ such that $HN \neq G$. As we noted earlier, no coset from C^N can be contained in a coset from C_N; also if $gH \in C_N$ is contained in some $g'H' \in C_{\leq K}$ we have $gH \in C'$. It follows that on the level of geometric realisations, one has

$$\|F\| = \|C_N\| \cup_{\|C'\|} \|C_{\leq K}\| .$$

We want to show that $\|C'\|$ is a strong deformation retract of $\|C_N\|$.

To prove this, we first claim that for $gH \in C_N$, the intersection $(g \cdot HN) \cap K$ is an element of C'. Indeed, as $G = KN$, we can write $g = kn$ with $n \in N$ and $k \in K$. Then the intersection

$$(g \cdot HN) \cap K = (kn \cdot HN) \cap K = (k \cdot HN) \cap K \tag{3.2}$$

contains k, so it is equal to $k \cdot (HN \cap K)$. We know that $H \in \mathcal{H}_N$ and $K \in \mathcal{H}^N$, so as $\mathcal{H}$ is divided by N, the subgroup $HN \cap K$ is contained in $\mathcal{H}$ as well. Furthermore, we have $(HN \cap K) \cdot N = HN \neq G$ by Lemma 3.23 and obviously $(g \cdot HN) \cap K$ is contained in K. Hence, the claim follows.

This allows us to define poset maps

$$\begin{aligned} \phi\colon C_N &\to C' \\ gH &\mapsto (g \cdot HN) \cap K \end{aligned} \qquad \text{and} \qquad \begin{aligned} \psi\colon C' &\to C_N \\ gH &\mapsto g \cdot HN. \end{aligned}$$

For $gH \in C'$, we have $gH \subseteq K$, hence

$$\phi \circ \psi(gH) = (g \cdot HN) \cap K \supseteq gH \cap K = gH.$$

If on the other hand $gH \in C_N$, one has by Eq. (3.2)

$$\begin{aligned} \psi \circ \phi(gH) &= ((g \cdot HN) \cap K) \cdot N \\ &= k \cdot (HN \cap K) \cdot N \end{aligned}$$

for some $k \in g \cdot HN \cap K$. By Lemma 3.23, we have $(HN \cap K) \cdot N = HN$, so it follows that $\psi \circ \phi(gH) = k \cdot HN \supseteq gH$.

Lemma 2.4 now implies that ϕ and ψ are homotopy equivalences which are inverse to each other. Furthermore, we have $gH \subseteq \psi(gH)$ for all $gH \in C'$, so again by Lemma 2.4, the map ψ is homotopic to the inclusion $C' \hookrightarrow C_N$ which must hence be a homotopy equivalence as well. It follows that $\|C'\|$ is a strong deformation retract of $\|C_N\|$.

This implies that F is homotopy equivalent to $C_{\leq K}$, which is contractible as it has K as unique maximal element. □

3.3.2 Coset complexes

We will now translate the results obtained in the last section to coset complexes. The following observation follows from elementary group theory.

Lemma 3.25. *Let $K_1 \neq K_2$ be subgroups of G such that $G = (K_1 \cap K_2)N$. Then one has $K_1 \cap N \neq K_2 \cap N$.*

We obtain the following relation between $\mathrm{CC}(G, \mathcal{H}^N)$ and $\mathrm{CC}(N, \mathcal{H} \cap N)$:

Lemma 3.26. *Assume that for every finite collection $K_1, \ldots, K_m \in \mathcal{H}^N$, one has $(K_1 \cap \ldots \cap K_m)N = G$. Then there is an isomorphism*

$$\mathrm{CC}(G, \mathcal{H}^N) \cong \mathrm{CC}(N, \mathcal{H} \cap N).$$

Proof. As $G = KN = NK$ for all $K \in \mathcal{H}^N$, each vertex of $\mathrm{CC}(G, \mathcal{H}^N)$ can be written as nK with $n \in N$. Use this to define the map

$$\begin{aligned} \psi\colon \mathrm{CC}(G, \mathcal{H}^N) &\to \mathrm{CC}(N, \mathcal{H} \cap N) \\ nK &\mapsto n \cdot K \cap N \end{aligned}$$

which we claim is an isomorphism of simplicial complexes.

As $n \in N$, this map is well-defined on vertices. It also clearly is surjective on vertices. Now assume that for $n_1, n_2 \in N$ and $K_1, K_2 \in \mathcal{H}^N$, one has $n_1 \cdot K_1 \cap N = n_2 \cdot K_2 \cap N$. As the two cosets coincide, so do the subgroups $K_1 \cap N = K_2 \cap N$. By Lemma 3.25, this implies that $K_1 = K_2$. It follows in particular that $n_1 K_1 = n_2 K_2$ which shows that ψ defines a bijection between the vertex sets of the two coset complexes.

To see that ψ is a simplicial map which defines a bijection between the set of simplices of the two complexes, take $n_1, \ldots, n_m \in N$ and $K_1, \ldots, K_m \in \mathcal{H}^N$ and consider the following chain of equivalences:

$$\begin{aligned} &\bigcap_i n_i K_i \neq \emptyset \\ \Leftrightarrow\ &\exists\, g \in G : \bigcap_i n_i K_i = \bigcap_i g K_i = g \bigcap_i K_i \\ \overset{*}{\Leftrightarrow}\ &\exists\, n \in N : \bigcap_i n_i K_i = n \bigcap_i K_i \\ \Leftrightarrow\ &\emptyset \neq \left(\bigcap_i n_i K_i \right) \cap N = \bigcap_i n_i (K_i \cap N), \end{aligned}$$

where $*$ follows because $G = N(K_1 \cap \ldots \cap K_m)$. □

This motivates the following definition:

Definition 3.27. The family $\mathcal{H}$ of proper subgroups of G is *strongly divided by* N if the following holds true:

1. For all $H \in \mathcal{H}_N$, one has $N \subseteq H$.
2. For all $K_1, \ldots, K_m \in \mathcal{H}^N$, one has $(K_1 \cap \ldots \cap K_m)N = G$.

Using Lemma 3.23, it is easy to see that every family of subgroups which is strongly divided by N is also divided by N. On top of that, given a family which is strongly divided, we can even produce a family which is closed under intersections and still divided by N as the following lemma shows. Recall that $\widetilde{\mathcal{H}}$ denotes the family of all finite intersections of elements from $\mathcal{H}$.

Lemma 3.28. *If $\mathcal{H}$ is strongly divided by N, the family $\widetilde{\mathcal{H}}$ is divided by N. Furthermore, we have*

1. *$\widetilde{\mathcal{H}}^N$ is equal to the family of all finite intersections of elements from $\mathcal{H}^N$, i.e.*
$$\widetilde{\mathcal{H}}^N = \widetilde{\mathcal{H}^N}.$$
2. *The image of $\widetilde{\mathcal{H}}_N$ in G/N is equal to the family of finite intersections of elements from $\overline{\mathcal{H}}$, i.e.*
$$\overline{\widetilde{\mathcal{H}}} = \widetilde{\overline{\mathcal{H}}}.$$

Proof. Every $\widetilde{H} \in \widetilde{\mathcal{H}}$ can be written as

$$\widetilde{H} = H_1 \cap \ldots \cap H_n \cap K_1 \cap \ldots \cap K_m$$

where for all i and j, one has $N \subseteq H_i$ and $K_j \in \mathcal{H}^N$.

If $\widetilde{H} \in \widetilde{\mathcal{H}}^N$, it follows that $n = 0$, i.e. $\widetilde{H} = K_1 \cap \ldots \cap K_m$ is a finite intersection of elements from $\mathcal{H}^N$. On the other hand, every such finite intersection forms an element of $\widetilde{\mathcal{H}}^N$ because one has $(K_1 \cap \ldots \cap K_m)N = G$, which proves Item 1.

This also implies that if $\widetilde{H} \in \widetilde{\mathcal{H}}_N$, we have $n \geq 1$. It follows from Lemma 3.23 that $\widetilde{H}N$ is equal to $H_1 \cap \ldots \cap H_n$. This is a finite intersection of elements from $\mathcal{H}_N$ and hence contained in $\widetilde{\mathcal{H}}$. Furthermore, this implies that the image $\overline{H}$ of $\widetilde{H}$ in G/N is equal to $\overline{H} = \overline{H}_1 \cap \ldots \cap \overline{H}_n$, showing Item 2.

The last thing that remains to be checked is that $\widetilde{H}$ is divided by N, i.e. that for all $\widetilde{H} \in \widetilde{\mathcal{H}}_N$ and $\widetilde{K} \in \widetilde{\mathcal{H}}^N$, one has $\widetilde{H}N \cap \widetilde{K} \in \widetilde{\mathcal{H}}$. However, we already know that $\widetilde{H}N = H_1 \cap \ldots \cap H_n$, so $\widetilde{H}N \cap \widetilde{K}$ is itself a finite intersection of elements from $\mathcal{H}$. □

We are now ready to prove Theorem C, which we restate as:

Theorem 3.29. *If $\mathcal{H}$ is strongly divided by N, there is a homotopy equivalence*

$$\mathrm{CC}(G, \mathcal{H}) \simeq \mathrm{CC}(G/N, \overline{\mathcal{H}}) * \mathrm{CC}(N, \mathcal{H} \cap N).$$

Proof. It follows from Lemma 3.3 and Lemma 3.4 that $\mathrm{CC}(G, \mathcal{H})$ is homotopy equivalent to $\mathrm{CP}(G, \widetilde{\mathcal{H}})$. Furthermore, Lemma 3.28 tells us that $\widetilde{\mathcal{H}}$ is divided by N. Hence, we can apply Proposition 3.24 to see that there is a homotopy equivalence

$$\mathrm{CP}(G, \widetilde{\mathcal{H}}) \simeq \mathrm{CP}(G/N, \overline{\widetilde{\mathcal{H}}}) * \mathrm{CP}(G, \widetilde{\mathcal{H}}^N).$$

By Lemma 3.28, we have $\overline{\widetilde{\mathcal{H}}} = \widetilde{\overline{\mathcal{H}}}$. Hence, using Lemma 3.3 and Lemma 3.4 again,

$$\mathrm{CP}(G/N, \overline{\widetilde{\mathcal{H}}}) \simeq \mathrm{CC}(G/N, \widetilde{\overline{\mathcal{H}}}) \simeq \mathrm{CC}(G/N, \overline{\mathcal{H}}).$$

On the other hand, Lemma 3.28 also tells us that $\widetilde{\mathcal{H}}^N$ consists of all finite intersections of elements from $\mathcal{H}^N$. It follows that

$$\mathrm{CP}(G, \widetilde{\mathcal{H}}^N) \simeq \mathrm{CC}(G, \widetilde{\mathcal{H}}^N) \simeq \mathrm{CC}(G, \mathcal{H}^N).$$

As $\mathcal{H}$ is strongly divided by N, we can finally apply Lemma 3.26 and get that $\mathrm{CC}(G, \mathcal{H}^N) \cong \mathrm{CC}(N, \mathcal{H} \cap N)$. □

3.3.3 Summary

We summarise the results of this section in the form that we will use later on:

Corollary 3.30. *Let G be a group and assume we have a short exact sequence*

$$1 \to N \to G \xrightarrow{q} Q \to 1.$$

Let S be a set of generators for $G = \langle S \rangle$ and let $\mathcal{P}$ be a family of proper subgroups. Furthermore, assume that for all $P \in \mathcal{P}$, one of the following holds:

1. *Either* P *contains the kernel* $N = \ker q$*, or*

2. P *contains* $S \setminus N$.

Then there is a homotopy equivalence

$$\mathrm{CC}(G, \mathcal{P}) \simeq \mathrm{CC}(Q, \overline{\mathcal{P}}) * \mathrm{CC}(N, \mathcal{P} \cap N),$$

where $\overline{\mathcal{P}} = \{q(P) \mid P \in \mathcal{P}, N \subseteq P\}$ *and* $\mathcal{P} \cap N = \{P \cap N \mid P \in \mathcal{P}, S \setminus N \subseteq P\}$.

Proof. We stick with the notation defined on page 23. If $P \in \mathcal{P}^N$, it cannot contain N. Hence, all such P must contain the set $S \setminus N$ of elements from S that are not contained in the kernel. It follows that for any $P_1, \ldots, P_m \in \mathcal{P}^N$, one has $(P_1 \cap \ldots \cap P_m)N = G$. On the other hand, for every $P \in \mathcal{P}_N$, our assumption implies that $N \subseteq P$. Hence, $\mathcal{P}$ is strongly divided by N and the claim follows from Theorem 3.29. $\square$

Chapter 4

Homotopy type of the complex of free factors

The topics of this chapter are automorphism groups of free products and complexes associated to them. After providing background material in Section 4.1, we in Section 4.2 define the free factor complex $\mathcal{F}$ associated to a Fouxe-Rabinovitch group and show that it is isomorphic to a coset complex of this group. Section 4.3 has a different flavour than the rest of this chapter: Using tools from poset topology, we study (finite) posets of graphs and determine their homotopy types. The aim of Section 4.4 is to establish Theorem F, which states that relative versions of the free splitting complex are contractible. In Section 4.5, this is used to prove that the free factor complex $\mathcal{F}$ and the complex of free factor systems $\mathcal{FF}$ are homotopy equivalent to certain complexes of free splittings. We show that these complexes are highly connected in Section 4.6. This uses the results on posets of subgraphs from Section 4.3. As a consequence, we obtain sphericity of the free factor complex $\mathcal{F}$, establishing Theorem D. The chapter ends in Section 4.7 with comments about the asymptotic geometry of Outer space, which put the results into context.

Standing assumption Throughout this chapter, let A denote a finitely generated group. If we do not explicitly state something different, we assume that A is *freely decomposable*, i.e. it can be written as a non-trivial free product of non-trivial groups. We write F_n for the free group of rank $n \in \mathbb{N}$.

The content of the present chapter mainly consists of the article [BG] published by Radhika Gupta and the author of this thesis. In that article, we worked throughout in the setting where $A = F_n$. Here, most of the material is presented in the more general case where A is any freely decomposable group, as was done in [Brüa, Section 4.2]. However, it should be mentioned that this is only a minor improvement. It changes some technical details but none of the main ideas of [BG].

4.1 Preliminaries: Free products and their automorphisms

The *outer automorphism group* $\mathrm{Out}(A)$ is the quotient $\mathrm{Aut}(A)/\mathrm{Inn}(A)$, where $\mathrm{Inn}(A)$ denotes the subgroup of the automorphism group $\mathrm{Aut}(A)$ consisting of all conjugations.

The modern study of $\mathrm{Out}(F_n)$ proceeds largely by considering the actions of this group on spaces related to free factorisations or free splittings of F_n. The most important space used for this is Outer space defined by Culler and Vogtmann in [CV86]. A lot of this machinery has by now been generalised for the study of $\mathrm{Out}(A)$, where A is an arbitrary freely decomposable group and to "relative" settings defined below. In what follows, we collect definitions and basic results in this field which will be used later on. This is done rather briefly. Further details on these preliminaries can be found in the following literature: The original definitions of Outer space and its relative versions are due to Culler–Vogtmann [CV86] and Guirardel–Levitt [GL07], a short introduction to this space and the motivation for its definition are given by Vogtmann in [Vog08]; for more details on free splittings, the reader is referred to the text [GL17] of Guirardel–Levitt; the setting of relative complexes of free splittings and free factor (systems) is described in detail in the article [HM] by Handel–Mosher.

Remark 4.1. The group $\mathrm{Out}(F_n)$ has become more common than $\mathrm{Aut}(F_n)$ in recent years because of its similarities with the mapping class group $\mathrm{MCG}(S)$ of a surface S, which is isomorphic to $\mathrm{Out}(\pi_1(S))$. Conceptually, whether one studies $\mathrm{Aut}(F_n)$ or $\mathrm{Out}(F_n)$ is mostly a matter of taste and there are $\mathrm{Aut}(F_n)$-versions of the free factor complex [HV98b], the free splitting complex [Hat95] and Outer space [HV98a]. For the sake of using similarities with arithmetic groups, this difference is irrelevant as well because we have $\mathrm{GL}_n(\mathbb{Z}) = \mathrm{Aut}(\mathbb{Z}^n) = \mathrm{Out}(\mathbb{Z}^n)$. In practice however, passing from one to the other is not always a trivial task. As many of the techniques that we use here are better developed for outer automorphism groups, we will thus mostly work with these in this text.

4.1.1 Free factors and free splittings

Free factors and free factor systems

A *free factor* of A is a subgroup $B \leq A$ such that A splits as a free product $A = B * C$. Let $[\cdot]$ denote the conjugacy class of a subgroup of A. There is a natural partial order on the set of conjugacy classes of free factors of A given by $[B_1] \leq [B_2]$ when a conjugate of B_1 is contained in B_2.

A *free factor system* of A is a finite collection $\mathcal{A} = \{[A_1], \ldots, [A_k]\}$, where each A_i is a non-trivial subgroup of A, such that there exists a free factorisation

$$A = F_n * A_1 * \cdots * A_k.$$

We call the conjugacy classes $[A_i]$ the *components* of $\mathcal{A}$. Of course, each A_i then is a free factor of A. A free factor system $\mathcal{A}$ is called *proper* if $\mathcal{A} \neq \{[A]\}$. Note that we allow $\mathcal{A} = \emptyset$, but this can only occur if $A = F_n$.

The partial order on the set of conjugacy classes of free factors of A extends to a partial order on the set of free factor systems: We write $\mathcal{A} \sqsubseteq \mathcal{A}'$ if for every $[A_i] \in \mathcal{A}$ there exists $[A'_j] \in \mathcal{A}'$ such that $[A_i] \leq [A'_j]$. In this case, if we want to emphasise that $\mathcal{A} \neq \mathcal{A}'$, we also write $\mathcal{A} \sqsubset \mathcal{A}'$.

To simplify notation, we will occasionally identify free factor systems having only one component with the corresponding conjugacy class of free factors and write $[A_1]$ for the system $\{[A_1]\}$.

Free splittings

An action of A on a tree T is *minimal* if T does not contain any proper A-invariant subtree. A *free splitting* S of A is a non-trivial, minimal, simplicial A-tree with finitely many edge orbits and trivial edge stabilisers. The *vertex group system* of a free splitting S is the (finite) set of conjugacy classes of its vertex stabilisers. Two free splittings S and S' are equivalent if they are equivariantly isomorphic. We say that S' collapses to S if there is a *collapse map* $S' \to S$ which collapses an A-invariant set of edges.

Let $\mathcal{A}$ be a proper free factor system of A. The *poset of free splittings of A relative to $\mathcal{A}$*, denoted $\mathcal{FS}(A, \mathcal{A})$, is given by the set of all equivalence classes of free splitting S of A such that $\mathcal{A} \sqsubseteq \mathcal{V}(S)$ and where $S \leq S'$ if S' collapses to S. Its order complex $\Delta(\mathcal{FS}(A, \mathcal{A}))$ is the *relative free splitting complex* studied in [HM], where the authors showed that it is non-empty, connected and hyperbolic.

4.1.2 Relative automorphism groups

The main interest for studying free factors and free splittings of A is to get a better understanding of its (outer) automorphism group.

We will often use capital letters for elements from the outer automorphism group of A and lower-case letters for the corresponding representatives from the automorphism group of A. In particular, for $\Phi \in \mathrm{Out}(A)$, we write $\Phi = [\phi]$ where $\phi \in \mathrm{Aut}(A)$. Let Φ be an outer automorphism of a group A and $H \leq A$ a subgroup. Then Φ *stabilises* H or H *is invariant under* Φ if there exists a representative $\phi \in \Phi$ such that $\phi(H) = H$. We say that Φ *acts trivially* on H if there is $\phi \in \Phi$ restricting to the identity on H.

If $\mathcal{G}$ and $\mathcal{H}$ are families of subgroups of A, the *relative outer automorphism group* $\mathrm{Out}(A; \mathcal{G}, \mathcal{H}^t)$ is the subgroup of $\mathrm{Out}(A)$ consisting of all elements stabilising each $H \in \mathcal{G}$ and acting trivially on each $H \in \mathcal{H}$. If $\mathcal{G}$ or $\mathcal{H}$ are given by the empty set, we also write $\mathrm{Out}(A; \mathcal{H}^t)$ or $\mathrm{Out}(A; \mathcal{G})$ for this group.

If $O \leq \mathrm{Out}(A)$ is a subgroup of the outer automorphism group of A and $G \leq A$, we also write

$$\mathrm{Stab}_O(G)$$

for the subgroup of O consisting of all elements that stabilise G. In the case where O is equal to $\mathrm{Out}(A; \mathcal{G}, \mathcal{H}^t)$, we have $\mathrm{Stab}_O(G) = \mathrm{Out}(A; \mathcal{G} \cup \{G\}, \mathcal{H}^t)$.

In the case where A is a freely decomposable group and $\mathcal{A}$ is a free factor system of A, the group $\mathrm{Out}(A; \mathcal{A}^t)$ is also called a *Fouxe-Rabinovitch group* because of the work of Fouxe-Rabinovitch on automorphism groups of free products [FR40]. Note that we defined $\mathrm{Out}(A; \mathcal{G}, \mathcal{H}^t)$ for families of subgroups $\mathcal{G}$ and $\mathcal{H}$; writing $\mathrm{Out}(A; \mathcal{A}^t)$ for a free factor system $\mathcal{A}$ is justified because the definition only depended on the conjugacy classes of the elements in $\mathcal{G}$ and $\mathcal{H}$.

4.1.3 Relative Outer space

Fix a proper free factor system $\mathcal{A} = \{[A_1], \ldots, [A_k]\}$ of A. In [GL07], Guirardel and Levitt define a topological space called *relative Outer space* used to study the group $O \coloneqq \mathrm{Out}(A; \mathcal{A}^t)$. Subgroups of A that are conjugate into one of the A_i are called *peripheral* subgroups. An $(A, \mathcal{A})$-tree is an $\mathbb{R}$-tree with an isometric action of A, in which every peripheral subgroup fixes a unique point. Two $(A, \mathcal{A})$-trees are equivalent if there exists an A-equivariant isometry between them. A *Grushko* $(A, \mathcal{A})$-graph is the quotient by A of a minimal, simplicial metric $(A, \mathcal{A})$-tree, whose set of point stabilisers is the free factor system $\mathcal{A}$ and edge stabilisers are trivial. *Relative Outer space* $\mathcal{O} = \mathcal{O}(A, \mathcal{A})$ is the space of equivalence classes of Grushko $(A, \mathcal{A})$-graphs with volume, i.e. sum of all edge lengths, equal to one.

The *spine* $L = L(A, \mathcal{A})$ of $\mathcal{O}$ is the subposet of $\mathcal{FS}(A, \mathcal{A})$ consisting of all free splittings S such that $\mathcal{A} = \mathcal{V}(S)$. The group O acts cocompactly on L: Taking the quotient by the action of A, each free splitting $S \in L$ can equivalently be seen as a marked graph of groups $\mathbb{G}$. The edge groups of $\mathbb{G}$ are trivial and for every $1 \leq i \leq k$, there is exactly one vertex group which is conjugate to A_i. All the other vertex groups are trivial. The *marking* is an isomorphism $\mathfrak{m} : A \to \pi_1(\mathbb{G})$ which is well-defined up to composition with inner automorphisms. The underlying graph G is finite, has fundamental group of rank n and all of its vertices with valence less than or equal to two have non-trivial vertex group. Using this description, the action of O on L is given by changing the marking.

The marking $\mathfrak{m}$ is determined by the following information: Choose a maximal forest F_G in G and a base point contained in F_H. After choosing an orientation for all the edges of G which are not contained in F_G, each such edge determines an element $g \in \pi_1(\mathbb{G})$. The marking identifies this element with $\mathfrak{m}^{-1}(g) \in A$ which we assign as a label to the corresponding edge. This assignment of an edge labelling completely determines $\mathfrak{m}$. For details, see [CV86, Section 1.2]. Let e be an edge of G and let G_o and G_t be the vertex groups of the initial and terminal end point of e, respectively. Then if we change the label of e by left-multiplication with an element of $\mathfrak{m}^{-1}(G_o)$ or right-multiplication by an element of $\mathfrak{m}^{-1}(G_t)$, the marked graph of groups we obtain still determines the same element $S \in L$. This follows from the choices involved when passing from a graph of groups to the corresponding Bass–Serre tree, see [Ser03, Section I.4.5] for the original definitions and [Bog08, Chapter 2.18] for a slightly more comprehensive introduction.

We can interpret $\|L\|$ as a subspace of $\mathcal{O}$. It consists of all Grushko $(A, \mathcal{A})$-graphs satisfying the following property: The subgraph spanned by the set of all edges not having maximal length forms a forest. The poset L is a contractible deformation retract of $\mathcal{O}$.

4.1.4 The case $A = F_n$: Culler–Vogtmann Outer space

If $A = F_n$, the empty set forms a free factor system $\mathcal{A} = \emptyset$ of A. In this case, we recover the more classical setting of Outer space as defined by Culler and Vogtmann [CV86]. In what follows, we spell out some of the definitions given above for this situation.

We can identify F_n with $\pi_1(\mathcal{R}, *)$ where $\mathcal{R}$ is a rose with n petals, i.e. the graph with one vertex and n edges which form loops attached at this vertex. A *marked graph* G is a graph of rank n equipped with a homotopy equivalence $\mathfrak{m} : \mathcal{R} \to G$ called a *marking*. The marking determines an identification of F_n with $\pi_1(G, \mathfrak{m}(*))$. (Unreduced) Culler–Vogtmann *Outer space* $\mathcal{CV}_n$ is the space

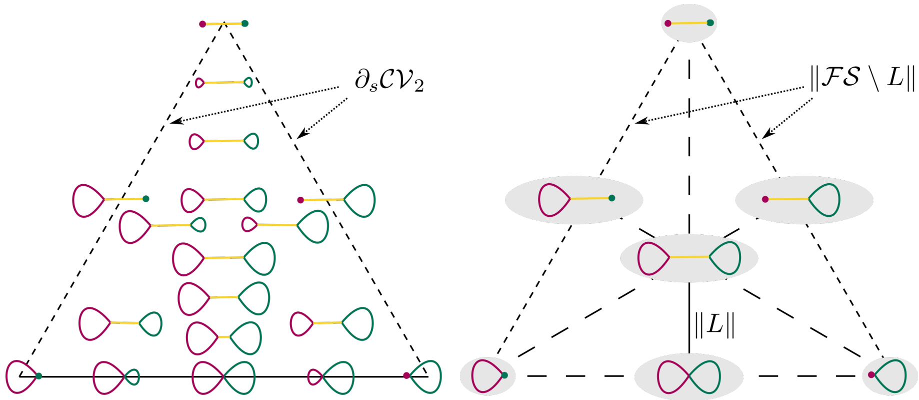

Figure 4.1: A simplex in $\mathcal{CV}_2$. The left-hand side shows points inside an open 2-simplex, corresponding to the dumbbell graph in its barycentre. The left vertex corresponds to the red, the right vertex to the green and the upper vertex to the yellow coordinate. The lower edge of the triangle is a face that is contained in $\mathcal{CV}_2$. The other faces (including the vertices) are in $\partial_s\mathcal{CV}_2$. The right-hand side shows the corresponding part of the order complex of $\mathcal{FS}$. Markings have been omitted to improve visibility.

of equivalence classes of marked metric graphs G of volume one such that every vertex of G has valence at least three. It is equal to the relative Outer space $\mathcal{O}(F_n, \emptyset)$ defined in Section 4.1.3.

Simplicial completion and boundary of Outer space

$\mathcal{CV}_n$ can be decomposed into a disjoint union of open simplices: Every marked graph determines a subspace of $\mathcal{CV}_n$ which can be parametrised as an open simplex σ_G of dimension $|E(G)| - 1$, where $|E(G)|$ denotes the number of edges of G: Assigning a length $0 < l < 1$ to each edge of G determines $|E(G)|$-many numbers which sum up to one, the volume of G. These determine a point in σ_G, expressed in its barycentric coordinates. Letting the length of the edges of a subgraph H of G go to zero results in changing the marked graph by collapsing H. If H is a forest, the resulting graph G/H is a marked graph with fundamental group F_n again; this corresponds to passing from σ_G to its face $\sigma_{G/H}$. If on the other hand H has non-trivial fundamental group, we have $\pi_1(G/H) \neq F_n$, so the corresponding ideal face $\sigma_{G/H}$ is not contained in $\mathcal{CV}_n$. These missing faces are thought of as “sitting at infinity”. There is a simplicial complex called the *simplicial completion of* $\mathcal{CV}_n$, which is obtained by adding the missing faces at infinity. The subspace of this completion consisting of all the open faces sitting at infinity is called the *simplicial boundary* $\partial_s\mathcal{CV}_n$ of Outer space. The points of $\partial_s\mathcal{CV}_n$ can be seen as marked metric graphs of groups, i.e. elements in the (geometric realisation of the) free splitting complex $\mathcal{FS} = \mathcal{FS}(F_n, \emptyset)$. In this way, Outer space embeds naturally as a subspace of $\|\mathcal{FS}\|$ and the free splitting complex is identified with the barycentric subdivision of the simplicial completion of $\mathcal{CV}_n$ (see Fig. 4.1). Using this translation, the spine $L = L(F_n, \emptyset)$

of $\mathcal{CV}_n$ is given by the subposet of $\mathcal{FS}$ consisting of all free splittings that have trivial vertex stabilisers. By the definitions above, we have a homeomorphism

$$\partial_s \mathcal{CV}_n \cong \|\mathcal{FS} \setminus L\| .$$

These notions generalise in a straightforward way to the relative Outer space $\mathcal{O}$ defined in Section 4.1.3.

4.2 Complexes of free factors

Standing assumptions In this section, we fix a proper free factor system $\mathcal{A} = \{[A_1], \ldots, [A_k]\}$ of A and let $O := \mathrm{Out}(A; \mathcal{A}^t)$.

The following definition is due to Handel and Mosher [HM].

Definition 4.2. Let $\mathcal{FF} = \mathcal{FF}(A, \mathcal{A})$ denote the poset of all free factor systems $\mathcal{B}$ of A such that $\mathcal{A} \sqsubset \mathcal{B} \sqsubset \{[A]\}$. Its order complex is called the *complex of free factor systems of A relative to $\mathcal{A}$*.

By [HM, Proposition 6.1], this complex has dimension $2n + k - 3$ and any of its simplices is contained in a simplex of maximal dimension.

Definition 4.3. Let $\mathcal{F} = \mathcal{F}(A, \mathcal{A})$ denote the poset of all conjugacy classes of proper free factors $B \subset A$ such that there is a free factor B' of A with $\mathcal{A} \sqsubseteq \{[B']\}$ and B' is a proper subgroup of B. We call the order complex of $\mathcal{F}$ the *free factor complex of A relative to $\mathcal{A}$*.

An equivalent definition of $\mathcal{F}$ will be given in Remark 4.6. Note that $\mathcal{F}$ can be seen as a subposet of $\mathcal{FF}$. However, it is generally not the subposet of $\mathcal{FF}$ given by all free factor systems which have only a single component: If $k > 1$, we have $[A_1 * \cdots * A_k] \in \mathcal{FF}$, but $[A_1 * \cdots * A_k] \notin \mathcal{F}$.

Remark 4.4. In the setting where $A = F_n$ and $\mathcal{A} = \emptyset$, the complexes $\mathcal{F}$ and $\mathcal{FF}$ are quasi-isometric to each other by [HM, Proposition 6.3]. However, we will see that they are not homotopy equivalent.

Remark 4.5. If $k = 0$, the poset $\mathcal{F}$ consists of *all* conjugacy classes of proper free factors of F_n, so we recover the $\mathrm{Out}(F_n)$-version of the free factor complex defined by Hatcher–Vogtmann [HV98b] as used e.g. in [BF14]. However, our definition differs from the one of the relative free factor complex used by Guirardel–Horbez [GH]; in the notation of Definition 4.3, they do not require that $\mathcal{A} \sqsubseteq \{[B]\}$.

4.2.1 The free factor complex as a coset complex

Both the free factor complex $\mathcal{F}$ and the complex of free factor systems $\mathcal{FF}$ carry a natural, simplicial action of O. We will use this action to show that $\mathcal{F}$ can be expressed as a coset complex of O. By assumption, A can be written as a free product $A = F_n * A_1 * \cdots * A_k$. In what follows, we will assume that $n \geq 2$ and fix a basis $\{x_1, \ldots, x_n\}$ of F_n.

Corank [HM, Lemma 2.11] implies that the elements of $\mathcal{F}$ are conjugacy classes of groups of the form

$$B = F * A_1^{x_1} * \ldots * A_k^{x_k},$$

where $x_j \in A$ and F is a free group with $1 \leq \mathrm{rk}(F) \leq n-1$. Furthermore, we can write A as a free product $A = B * C$, where C is a free group of rank $n - \mathrm{rk}(F)$. The rank of C is an invariant of the conjugacy class $[B]$, see [HM, Section 2.3]. It is called the *corank* of $[B]$ and will be denoted by $\mathrm{crk}[B]$.

Remark 4.6. Using the corank, the definition of $\mathcal{F}$ can be rephrased as follows: The free factor complex $\mathcal{F} = \mathcal{F}(A, \mathcal{A})$ of A relative to $\mathcal{A}$ is the poset of all conjugacy classes of proper free factors $B \subset A$ such that $\mathrm{crk}[B] > 0$.

Let

$$S_i \coloneqq \langle x_1, \ldots, x_i \rangle * A_1 * \ldots * A_k.$$

Every S_i is a free factor of A because for all i, we have $A = S_i * \langle x_{i+1}, \ldots, x_n \rangle$. We set $P_i \coloneqq \mathrm{Stab}_O(S_i)$ and define the set of *maximal standard parabolic subgroups of O* as

$$\mathcal{P} = \mathcal{P}(O) \coloneqq \{P_i \mid 1 \leq i \leq n-1\}.$$

As in the case of $\mathrm{GL}_n(\mathbb{Z})$, we will usually leave out the adjective "standard" (see Remark 3.14).

Proposition 4.7. *The free factor complex $\mathcal{F}$ of A relative to $\mathcal{A}$ is O-equivariantly isomorphic to the coset complex $\mathrm{CC}(O, \mathcal{P})$. In particular, it has dimension $n-2$.*

Proof. If $[B_1] \leq [B_2]$, we know from [HM, Proposition 2.10] that the corank of $[B_2]$ is smaller than or equal to the corank of $[B_1]$ and that equality holds if and only if $[B_1] = [B_2]$. Consequently, the simplices of $\Delta(\mathcal{F})$ are given by chains of the form

$$[B_1] \leq [B_2] \leq \ldots \leq [B_m]$$

with $\mathrm{crk}[B_1] < \mathrm{crk}[B_2] < \ldots < \mathrm{crk}[B_m]$. Let $i_j \coloneqq \mathrm{crk}[B_j]$.

We claim that for each such chain, there exists $\Phi \in O$ with $[\phi(S_{i_j})] = [B_j]$ for all j. To see this, first observe that sending each A_i to a conjugate of itself and fixing all the other generators defines an automorphism of A that represents an element in O. Hence, we can assume that $A_1 * \ldots * A_k \leq B_1$. Now choose representatives such that $B_j \leq B_{j+1}$ for all j. In order to use induction, assume that there is $\Phi' \in O$ such that for some l, we have $\phi'(S_{i_j}) = B_j$ for all $0 \leq j \leq l$—this is true for $l = 0$ where we define $i_0 = 0$ and $B_0 = S_0 = A_1 * \ldots * A_k$. By assumption, $\phi'(S_{i_l}) = B_l \leq B_{l+1}$, so [HM, Lemma 2.11] implies that

$$A = \phi'(S_{i_l}) * C * D, \qquad \text{where} \qquad B_{l+1} = \phi'(S_{i_l}) * C$$

and C and D are free groups of rank $(i_{l+1} - i_l)$ and $(n - i_{l+1})$, respectively. On the other hand, the group A also decomposes as a free product

$$A = S_{i_l} * \langle x_{i_l+1}, \ldots, x_{i_{l+1}} \rangle * \langle x_{i_{l+1}+1}, \ldots, x_n \rangle.$$

This allows us to define an automorphism ϕ of A which agrees with ϕ' on S_{i_l}, maps $\langle x_{i_l+1}, \ldots, x_{i_{l+1}} \rangle$ isomorphically to C and $\langle x_{i_{l+1}+1}, \ldots, x_n \rangle$ to D. As ϕ agrees with ϕ' on S_{i_l}, we know that $[\phi(S_{i_j})] = [B_j]$ for all $j \leq l$ and that ϕ acts by conjugation on each A_i, i.e. $[\phi] \in O$. Furthermore, we have

$$\begin{aligned}\phi(S_{i_{l+1}}) &= \phi(S_{i_l}) * \phi(\langle x_{i_l+1}, \ldots, x_{i_{l+1}} \rangle) \\ &= \phi'(S_{i_l}) * C \\ &= B_{l+1}.\end{aligned}$$

By induction, this proves the claim.

On the other hand, for each $[\phi] \in O$, the chain

$$[\phi(S_1)] \leq [\phi(S_2)] \leq \ldots \leq [\phi(S_{n-1})]$$

forms a facet in $\Delta(\mathcal{F})$. Hence, every facet of $\Delta(\mathcal{F})$ can be written in this form. It follows that the natural action of O on $\Delta(\mathcal{F})$ has a fundamental domain given by the simplex

$$[S_1] \leq [S_2] \leq \ldots \leq [S_{n-1}]$$

The result now follows from Proposition 3.9. ☐

Remark 4.8. In contrast to $\mathcal{F}$, the complex $\mathcal{FF}$ of free factor systems is in general not isomorphic to a coset complex: If this were the case, the group O would have to act facet transitively on $\mathcal{FF}$ (see Section 3.1.3). However, this does not even hold in the case where $A = F_3$ and $\mathcal{A} = \emptyset$. The chains

$$\{[\langle x_1 \rangle]\} \sqsubseteq \{[\langle x_1 \rangle], [\langle x_2 \rangle]\} \sqsubseteq \{[\langle x_1, x_2 \rangle]\} \sqsubseteq \{[\langle x_1, x_2 \rangle], [x_3]\} \qquad \text{and}$$
$$\{[\langle x_1 \rangle]\} \sqsubseteq \{[\langle x_1 \rangle], [\langle x_2 \rangle]\} \sqsubseteq \{[\langle x_1 \rangle], [\langle x_2 \rangle], [\langle x_3 \rangle]\} \sqsubseteq \{[\langle x_1, x_2 \rangle], [x_3]\}$$

both form facets in $\mathcal{FF}(F_3, \emptyset)$, but there is no element of O which maps one of them to the other.

Remark 4.9. The $\mathrm{Aut}(F_n)$-version of the free factor complex defined by Hatcher–Vogtmann [HV98b]—i.e. the poset of proper free factors of F_n, ordered by inclusion—is isomorphic to a coset complex of $\mathrm{Aut}(F_n)$ with respect to a family of "parabolic" subgroups as well. Details on this can be found in [Brüb, Section 3.3].

4.3 Posets of graphs

In this section, we study (finite) posets of subgraphs of a given graph G. For the combinatorial arguments we use, let us set up the following notation:

In what follows, all graphs are allowed to have loops and multiple edges. For a graph G, we denote the set of its vertices by $V(G)$ and the set of its edges by $E(G)$. If $e \in E(G)$ is an edge, then $G - e$ is defined to be the graph obtained from G by removing e and G/e is obtained by collapsing e and identifying its two endpoints to a new vertex v_e. A graph is called a *tree* if it is contractible. It is called a *forest* if it is a disjoint union of trees. An edge $e \in E(G)$ is called *separating* if removing it from G results in a disconnected graph.

Throughout this section, we will only care about *edge-induced subgraphs*, i.e. when we talk about a "subgraph H of G", we will always assume that H is possibly disconnected but does not contain any isolated vertices. Hence, we can interpret any subgraph of G as a subset of $E(G)$.

Definition 4.10. A *labelled graph* is a pair (G, l) consisting of a graph G and a map $l\colon \{1, \ldots, k\} \to V(G)$. We call the image of l the *labelled vertices* of G.

If one has a labelled graph (G, l) and an edge $e \in E(G)$, there are canonical labellings $\{1, \ldots, k\} \to G - e$ and $\{1, \ldots, k\} \to G/e$ that will be denoted by l as well.

Definition 4.11. A labelled graph is called a *core graph* if every vertex of valence one is labelled and each connected component of the graph either has non-trivial fundamental group or contains a labelled vertex. Every labelled graph (G, l) contains a unique maximal core subgraph. It is a labelled graph as well and contains $\mathrm{im}(l)$. We will refer to this as the *core of* G, denoted by $(\mathring{G}, l)$.

The reason why we study (posets of) labelled graphs is the following: As remarked above, every free splitting S in the spine $L(A, \mathcal{A})$ can be seen as a marked graph of groups whose system of (conjugacy classes of) vertex groups is given by $\mathcal{A} := \mathcal{V}(S) = \{[A_1], \ldots, [A_k]\}$. Let G be the underlying graph of this graph of groups. There is a natural labelling $l\colon \{1, \ldots, k\} \to V(G)$ of G given by defining $l(i)$ as the vertex with vertex group conjugate to A_i. From the definition of free splittings (see Section 4.1.1), it follows that (G, l) is in fact a core graph. Labelled graphs and their core graphs were also studied by Bestvina and Feighn in [BF00]. However, note that in contrast to the definition given above, their core graphs are not allowed to have separating edges. This difference is due to the fact that Bestvina–Feighn use *reduced* Outer space for their definitions, while we use its unreduced version (for comments on this, see Section 4.7.1).

4.3.1 The poset of core subgraphs

In this subsection, we restrict ourselves to the setting of unlabelled graphs. Every unlabelled graph G can equivalently be viewed as a labelled graph (G, l), where $l\colon \{1, \ldots, k\} \to V(G)$ is trivial (i.e. $k = 0$). In this case, a core graph is simply a graph such that each of its connected components has non-trivial fundamental group and such that no vertex has valence one. For our purposes, unlabelled graphs arise from splittings S of $A = F_n$ that have trivial vertex stabilisers $\mathcal{V}(S) = \emptyset$; these splittings are exactly the elements of the spine $L = L(F_n, \emptyset)$ of Culler–Vogtmann Outer space $\mathcal{CV}_n$.

Definition 4.12. Let G be a graph. We define the following posets of subgraphs of G; all of them are ordered by inclusion:

1. $\mathrm{Sub}(G)$ is the poset of all proper subgraphs of G that are non-empty. Equivalently, $\mathrm{Sub}(G)$ can be seen as the poset of all proper, non-empty subsets of $E(G)$.

2. $\mathrm{For}(G)$ denotes the poset of all proper, non-empty subgraphs of G that are forests.

3. $\mathrm{X}(G)$ is defined to be the poset of proper subgraphs of G that are non-empty and where at least one connected component has non-trivial fundamental group.

4. $\mathrm{C}(G)$ is the poset of all proper core subgraphs of G.

Clearly one has:

$$\mathrm{C}(G) \subseteq \mathrm{X}(G) \subseteq \mathrm{Sub}(G)$$

and

$$\mathrm{X}(G) = \mathrm{Sub}(G) \setminus \mathrm{For}(G).$$

Examples of the realisation of $\mathrm{X}(G)$ can be found in the Appendix, see Figure A.3.

Lemma 4.13. $\mathrm{X}(G)$ *deformation retracts to* $\mathrm{C}(G)$.

Proof. Every subgraph $H \in \mathrm{X}(G)$ contains a unique maximal core subgraph $\mathring{H}$ and if $H_1 \subseteq H_2$, one has $\mathring{H}_1 \subseteq \mathring{H}_2$. Hence, sending each H to this core subgraph $\mathring{H}$ defines a monotone poset map $f : \mathrm{X}(G) \to \mathrm{C}(G)$ restricting to the identity on $\mathrm{C}(G)$. The claim now follows from Corollary 2.5. □

Proposition 4.14. *Let G be a finite connected graph whose fundamental group has rank $n \geq 2$ and assume that every vertex of G has valence at least three. Then* $\mathrm{X}(G)$ *is $(n-2)$-spherical. It is contractible if and only if G has a separating edge.*

Proof. Note that $\mathrm{Sub}(G)$ can be seen as the poset of all proper faces of a simplex with vertex set $E(G)$. Hence, its realisation $\|\mathrm{Sub}(G)\|$ is homeomorphic to a sphere of dimension $|E(G)| - 2$. By [Vog90, Proposition 2.2], the poset $\mathrm{For}(G)$ is contractible if and only if G has a separating edge and is homotopy equivalent to a wedge of $(|V(G)| - 2)$-spheres if it does not contain a separating edge. We want to use Alexander duality as stated in Lemma 2.6 to describe the homology groups of $\mathrm{X}(G) = \mathrm{Sub}(G) \setminus \mathrm{For}(G)$.

If G has a separating edge, it immediately follows from Alexander duality that all reduced homology groups of $\mathrm{X}(G)$ vanish. If on the other hand G does not have a separating edge, then the only non-trivial homology group of $\mathrm{X}(G)$ appears in dimension

$$(|E(G)| - 2) - 1 - (|V(G)| - 2) = n - 2$$

where it is given by a direct sum of copies of $\mathbb{Z}$.

We next want to show that for $n \geq 4$, the realisation of $\mathrm{X}(G)$ is simply-connected in order to apply the Whitehead theorem.

Denote by $\mathrm{Sub}(G)^{(k)}$ the subposet of $\mathrm{Sub}(G)$ given by those subgraphs having precisely $(|E(G)| - k)$ edges. As $n \geq 4$, removing at most three edges from G results in a graph with non-trivial fundamental group. Hence, we have $\mathrm{Sub}(G)^{(k)} \subset \mathrm{X}(G)$ for $k = 1, 2, 3$. The realisation of

$$\mathrm{Sub}(G)^{(\leq 3)} := \mathrm{Sub}(G)^{(1)} \cup \mathrm{Sub}(G)^{(2)} \cup \mathrm{Sub}(G)^{(3)}$$

forms a subspace of $\|\mathrm{X}(G)\|$ that is homeomorphic to the 2-skeleton of an $(|E(G)| - 2)$-simplex. In particular, it is simply-connected.

Now let ρ be a closed edge path in $\|\mathrm{X}(G)\|$ given by the sequence of vertices $(H = H_1, H_2, ..., H_k = H)$. We want to show that it can be homotoped to a path in $\left\|\mathrm{Sub}(G)^{(1)} \cup \mathrm{Sub}(G)^{(2)}\right\|$. Whenever we have an edge $(H_{i-1} \subset H_i)$ such that H_i has at least two edges fewer than G, there is a subgraph $H_i' \in \mathrm{Sub}(G)^{(1)}$

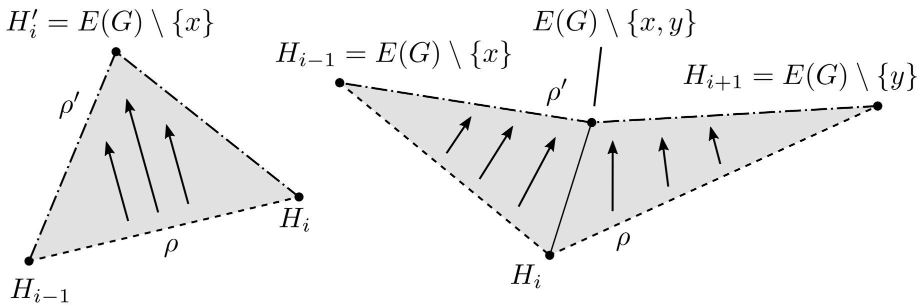

Figure 4.2: Simple connectedness of $\mathrm{X}(G)$.

containing H_i. As the chain $(H_{i-1} \subset H_i \subset H'_i)$ forms a simplex in $\mathrm{X}(G)$, we can replace the segment (H_{i-1}, H_i) by (H_{i-1}, H'_i, H_i) and hence assume that every second vertex crossed by ρ lies in $\mathrm{Sub}(G)^{(1)}$ (see the left hand side of Figure 4.2). Next take a segment $(H_{i-1} \supset H_i \subset H_{i+1})$ where $H_{i-1} = E(G) \setminus \{x\}$ and $H_{i+1} = E(G) \setminus \{y\}$ lie in $\mathrm{Sub}(G)^{(1)}$. In this situation, the two chains $(H_i \subseteq E(G) \setminus \{x, y\} \subset H_{i-1})$ and $(H_i \subseteq E(G) \setminus \{x, y\} \subset H_{i+1})$ form simplices contained in $\mathrm{X}(G)$. It follows that we can perform a homotopy in order to replace $(H_{i-1} \supset H_i \subset H_{i+1})$ by $(H_{i-1} \supset E(G) \setminus \{x, y\} \subset H_{i+1})$.

This argument shows that every closed path can be homotoped to a path that lies in $\left\|\mathrm{Sub}(G)^{(\leq 3)}\right\|$. As this is a simply-connected subset of $\|\mathrm{X}(G)\|$, it follows that $\mathrm{X}(G)$ itself is simply-connected for $n \geq 4$. Applying Corollary 2.10 yields the result.

The only cases that remain are those where $n = 2$ or 3. However, as we assumed that every vertex of G has valence at least three, there are only finitely many such graphs. Using Lemma 4.13, it is not hard to verify the claim using a case-by-case analysis. For completeness, the proof for $n = 3$ can be found in the Appendix A. □

Remark 4.15. Assuming that each vertex of G has valence at least 3 does not impose any restrictions for the considerations in this text as every graph in Outer space satisfies this condition. However, note that we only used this assumption in the case where $n = 2$ or 3 and there it only shortened the argument and could easily be dropped.

4.3.2 The poset of connected core subgraphs

We turn to the more general setting of (non-trivially) labelled graphs.

Definition 4.16. Let (G, l) be a labelled graph.

1. $\mathrm{cX}(G, l)$ is the poset of all connected subgraphs of G which are not trees, contain all the labelled vertices and whose fundamental group is strictly contained in $\pi_1(G)$.

2. $\mathrm{cC}(G, l)$ is the poset of all proper connected core subgraphs of (G, l) which have non-trivial fundamental group.

In both cases, the partial order is given by inclusion of subgraphs.

If (G, l) is a labelled graph such that exactly one component of G has non-trivial fundamental group and this component contains $\mathrm{im}(l)$, then its core $(\mathring{G}, l)$ is connected (here, we use that our core subgraphs are allowed to have separating edges). We use this observation in the proof of the following lemma:

Lemma 4.17. $\mathrm{cX}(G, l)$ *deformation retracts to* $\mathrm{cC}(G, l)$.

Proof. By restricting the labelling, every $H \in \mathrm{cX}(G, l)$ can be seen as a labelled graph (H, l) and, as noted above, its core $(\mathring{H}, l)$ is connected. Also, if $H_1 \leq H_2$ in $\mathrm{cX}(G, l)$, one has $(\mathring{H}_1, l) \subseteq (\mathring{H}_2, l)$. Hence, sending (H, l) to its core $(\mathring{H}, l)$ defines a poset map $f\colon \mathrm{cX}(G, l) \to \mathrm{cC}(G, l)$ which restricts to the identity on $\mathrm{cC}(G, l)$. The claim now follows from Corollary 2.5. ☐

Lemma 4.18. *Let (G, l) be a labelled graph where G is finite and connected. Let $v \in V(G)$ be a vertex of valence one and e the edge adjacent to v. Then* $\mathrm{cX}(G, l) \simeq \mathrm{cX}(G/e, l)$.

Proof. Whenever $H \in \mathrm{cX}(G, l)$, the set $H \setminus \{e\}$ can be seen as a connected subgraph of G/e. It contains all labelled vertices of $(G/e, l)$ and has non-trivial fundamental group. Using this, we can define poset maps $f\colon \mathrm{cX}(G, l) \to \mathrm{cX}(G/e, l)$ and $g\colon \mathrm{cX}(G/e, l) \to \mathrm{cX}(G, l)$ by setting $f(H) \coloneqq H \setminus \{e\}$ and

$$g(K) \coloneqq \begin{cases} K \cup \{e\} & , v \in \mathrm{im}(l), \\ K & , \text{else.} \end{cases}$$

To see that f is well-defined, note that for $H \in \mathrm{cX}(G, l)$, one has

$$1 < \pi_1(H) = \pi_1(H \setminus \{e\}) < \pi_1(G) = \pi_1(G/e).$$

On the other hand, the case distinction in the definition of g ensures that for each $K \in \mathrm{cX}(G/e, l)$, the image $g(K)$ contains all labelled vertices of (G, l). Also, if v is in the image of l, then the vertex v_e to which e is collapsed is a labelled vertex of $(G/e, l)$ whence we know that it is contained in any $K \in \mathrm{cX}(G/e, l)$. It follows that for all such K, the subgraph $K \cup \{e\}$ of G is connected. This implies that g is well-defined.

For all $K \in \mathrm{cX}(G/e, l)$, one obviously has $f \circ g(K) = K$. On the other hand, if v is in the image of l, the edge e must be contained in any graph $H \in \mathrm{cX}(G, l)$ which implies that $g \circ f(H) = H$, so f and g form bijections which are inverse to each other. If v is not labelled, this might not be true, but in this case, one still has $g \circ f(H) \subseteq H$. Hence by Lemma 2.4, f and g induce inverse homotopy equivalences on geometric realisations. ☐

To prove the following result, we apply an argument similar to the one used in [Vog90, Proposition 2.2].

Proposition 4.19. *Let (G, l) be a labelled graph where G is finite, connected and has fundamental group of rank $n \geq 2$. Then $\mathrm{cX}(G, l)$ is $(n-2)$-spherical.*

Proof. By Lemma 4.18 we can assume that every vertex of G has valence at least two.

We proceed by induction on n and start with the case $n = 2$. By Lemma 4.17, it suffices to show that $\mathrm{cC}(G, l)$ is homotopy equivalent to a wedge of 0-spheres, i.e. a disjoint union of points. Let $H \in \mathrm{cC}(G, l)$. As $1 < \pi_1(H) < \pi_1(G)$, the fundamental group of H is infinite cyclic. Let $e \in H$ be an edge of H. We distinguish between the two cases where e is non-separating or separating in H. As we assumed that there is no vertex of valence one, e being non-separating implies that $H \setminus \{e\}$ has trivial fundamental group while if e is separating, $H \setminus \{e\}$ has two connected components both of which either have non-trivial fundamental group or contain at least one labelled vertex. In both cases, no $K \in \mathrm{cC}(G, l)$ can be contained in $H \setminus \{e\}$. Hence, the order complex of $\mathrm{cC}(G, l)$ does not contain any simplex of dimension greater than zero which proves the claim.

Now let $n > 2$. If every edge of G is a loop, G is a rose with n petals and every proper non-empty subset of $E(G)$ is an element of $\mathrm{cX}(G, l)$. In this case, the order complex of $\mathrm{cX}(G, l)$ is given by the set of all proper faces of a simplex of dimension $n - 1$ whose vertices are in one-to-one correspondence with the edges of G. This is homotopy equivalent to an $(n - 2)$-sphere.

So assume that G has an edge e which is not a loop. We want to define poset maps between $\mathrm{cX}(G, l)$ and $\mathrm{cX}(G/e, l)$, just as in the proof of Lemma 4.18. If $H \in \mathrm{cX}(G, l)$ is not equal to $G - e$, then $H \setminus \{e\}$ is a proper subgraph of G/e which, as we assumed that G has no vertex of valence one, implies that $\pi_1(H \setminus \{e\}) < \pi_1(G/e)$. Consequently, we get a poset map

$$\begin{aligned} f\colon\ \mathrm{cX}(G, l) \setminus \{G - e\} &\to \mathrm{cX}(G/e, l) \\ H &\mapsto H \setminus \{e\}\,. \end{aligned}$$

On the other hand, if $K \in \mathrm{cX}(G/e, l)$ contains the vertex v_e to which e was collapsed, it is easy to see that $K \cup \{e\}$ is an element of $\mathrm{cX}(G, l) \setminus \{G - e\}$. This allows us to define a poset map

$$\begin{aligned} g\colon\ \mathrm{cX}(G/e, l) &\to \mathrm{cX}(G, l) \setminus \{G - e\} \\ K &\mapsto \begin{cases} K \cup \{e\} & , v_e \in V(K), \\ K & , \text{else}. \end{cases} \end{aligned}$$

One has $g \circ f(H) \supseteq H$ and $f \circ g(K) = K$, so using Lemma 2.4, these two posets are homotopy equivalent.

If e is separating, the graph $G - e$ is not connected so in particular not an element of $\mathrm{cX}(G, l)$. It follows that $\mathrm{cX}(G, l)$ is homotopy equivalent to $\mathrm{cX}(G/e, l)$. As G/e has one edge less than G, we can apply induction.

If on the other hand e is not a separating edge, $G - e$ is a connected graph having the same number of vertices as G and one edge less. This implies that $\mathrm{rk}(\pi_1(G - e)) = n - 1$. Collapsing edges adjacent to valence-one vertices and using Lemma 4.18, we see that $\mathrm{cX}(G - e, l) \simeq \mathrm{cX}(G', l)$ where G' has the same rank as $G - e$, at most as many edges and such that every vertex in G' has valence at least two. Hence, $\mathrm{cX}(G - e, l) \simeq \mathrm{cX}(G', l)$ is by induction homotopy equivalent to a wedge of $(n - 3)$-spheres.

$\|\mathrm{cX}(G, l)\|$ is obtained from $\|\mathrm{cX}(G, l) \setminus \{G - e\}\|$ by attaching the star of $G - e$ along its link. The link of $G - e$ in $\|\mathrm{cX}(G, l)\|$ is isomorphic to $\|\mathrm{cX}(G - e, l)\|$ and its star is contractible. Gluing a contractible set to an $(n - 2)$-spherical complex along an $(n - 3)$-spherical subcomplex results in an $(n - 2)$-spherical complex, so the claim follows. □

4.4 Contractibility of relative free splitting complexes

Standing assumptions Throughout this section, let $\mathcal{A} = \{[A_1], \ldots, [A_k]\}$ be a proper free factor system of A. We write $\mathcal{FS} := \mathcal{FS}(A, \mathcal{A})$ and $L := L(A, \mathcal{A})$ for the associated free splitting complex and the spine of the corresponding relative Outer space.

The aim of this section is to show that the relative free splitting complex $\mathcal{FS} = \mathcal{FS}(A, \mathcal{A})$ is contractible (Theorem 4.29). In order to prove this, we need several relativisations of the complexes in question and need to introduce some notation in order to keep track of them. To make the proof more accessible, we start with an informal outline before spelling out the technicalities in greater detail.

4.4.1 Outline of the proof

The poset of all free splittings of A having vertex group system equal to $\mathcal{A}$ is nothing but the spine L of the Outer space of A relative to $\mathcal{A}$, so we know that it is contractible. In this way, we can see $\mathcal{FS}$ as being assembled from the contractible pieces $L(A, \mathcal{B})$ where $\mathcal{B}$ ranges over all free factor systems with $\mathcal{A} \sqsubseteq \mathcal{B}$. In order to understand the (order) relation between these different pieces, we need a way of organising them. The natural choice is to use the ordering of free factor systems in $\mathcal{FF} = \mathcal{FF}(A, \mathcal{A})$.

Roughly speaking, the order "$\sqsubseteq$" of $\mathcal{FF}$ is coarser than the one on $\mathcal{FS}$: It is true that if S is greater or equal to S' in $\mathcal{FS}$, i.e. if there is a collapse map $S \to S'$, one has $\mathcal{V}(S) \sqsubseteq \mathcal{V}(S')$. Hence, the map $\mathcal{FS} \to \mathcal{FF}$, sending S to $\mathcal{V}(S)$ is an order-inverting poset map. However, if one is given $S \in \mathcal{FS}$ such that $\mathcal{A} \sqsubseteq \mathcal{V}(S)$, it is not necessarily true that there is a collapse map $S' \to S$ such that S' has vertex group system $\mathcal{V}(S') = \mathcal{A}$.

Recall that in the order complex of $\mathcal{FS}$, an edge between S and S' is added if and only if one collapses to the other. So in order to understand how the spines of the different relative Outer spaces are glued together to form the relative free splitting complex, we need to understand the following situation: If $\mathcal{A} \sqsubseteq \mathcal{B}$, which elements $S \in L(A, \mathcal{A})$ collapse to some free splitting $S' \in L(A, \mathcal{B})$? Adopting the graph of groups point of view, one can intuitively see S' as being obtained from S by collapsing a subgraph of groups. This is why in this case, we will say that S has a "subgraph with fundamental group $\mathcal{B}$" and denote the poset of all such $S \in L(A, \mathcal{A})$ by $X(\mathcal{A} : \mathcal{B})$ (see below for the precise definitions). These are the posets whose connectivity properties we want to understand.

Eventually, our inductive argument requires us to consider intersections and unions of such $X(\mathcal{A} : \mathcal{B})$ as well, so we are led to consider slightly more general versions of these posets and need to show that they all are contractible.

4.4.2 Subgraphs of groups

As mentioned above, each $S \in L$ can be seen as a marked graph of groups $\mathbb{G}$ with system of (conjugacy classes of) vertex groups given by $\mathcal{A}$. Let G be the underlying graph of $\mathbb{G}$ and let H be a connected subgraph of G that contains all the vertices with non-trivial vertex group. Then there is an induced structure

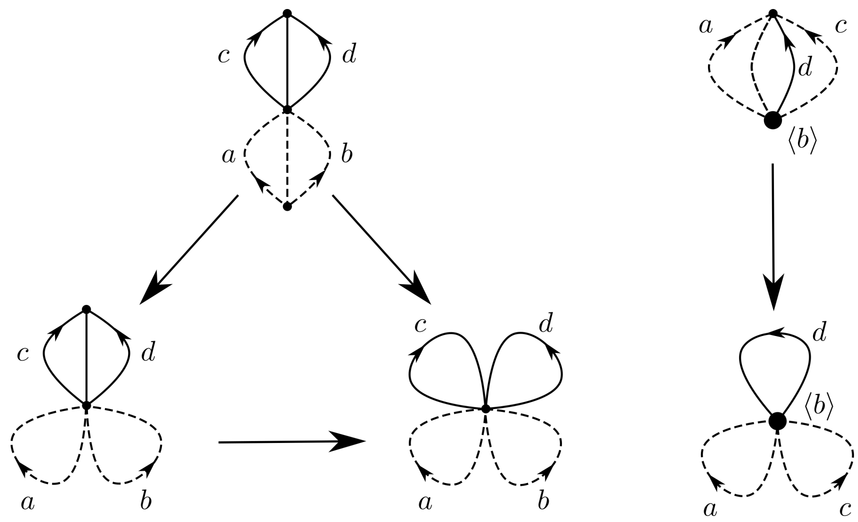

Figure 4.3: Let $A = F_4 = \langle a, b, c, d\rangle$. The three graphs on the left are elements of $X(\emptyset : [\langle a, b\rangle])$ and the two on the right are in $X([\langle b\rangle] : [\langle a, b, c\rangle])$. The dashed parts are the corresponding subgraphs with fundamental group $[\langle a, b\rangle]$ and $[\langle a, b, c\rangle]$, respectively.

of a marked graph of groups on H. We define the *fundamental group* $\pi_{\mathbb{G}}(H)$ as the fundamental group of this graph of groups. It is a subgroup of A that is well-defined up to conjugacy and has the form

$$\pi_{\mathbb{G}}(H) = F * A_1^{x_1} * \ldots * A_k^{x_k},$$

where $x_i \in A$ and F is a free group with rank equal to the rank of $\pi_1(H)$.

Definition 4.20. Let $S \in L$, let $\mathbb{G}$ be the associated graph of groups and (G, l) the underlying labelled graph. Let $\mathcal{B} = \{B_1, \ldots, B_l\}$ be a free factor system in A such that $\mathcal{A} \sqsubseteq \mathcal{B}$. We say that S *has a subgraph with fundamental group* $\mathcal{B}$ if there are disjoint subgraphs $H_1, \ldots, H_l$ of G such that $[\pi_{\mathbb{G}}(H_i)] = [B_i]$.

If such subgraphs exists, there is also a unique core subgraph H of (G, l) which has connected components $H_1, \ldots, H_l$ and such that $[\pi_{\mathbb{G}}(H_i)] = [B_i]$. We will denote this core subgraph by $\mathcal{B}|S$. We then also say that $\mathcal{B}|S$ *is a subgraph of* S.

The notation $\mathcal{B}|S$ is borrowed from Bestvina–Feighn who use it for the setting where $A = F_n$ in [BF14].

Notation To simplify notation, we will from now on not distinguish between a free splitting S and the corresponding graph of groups. For example, we will talk about "(core) subgraphs of S" and mean (core) subgraphs of the corresponding labelled graph (G, l). Instead we will use the letter G for elements in $L = L(A, \mathcal{A})$ and the letter S for free splittings that have vertex group system different than $\mathcal{A}$. If $G \in L$ and H is a subgraph, let G/H denote the free splitting obtained by collapsing H.

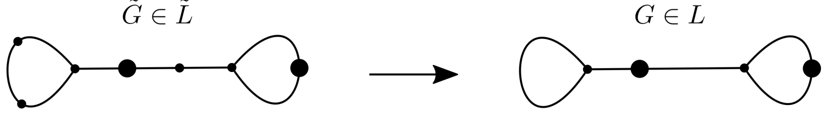

Figure 4.4: Two graphs of groups (each having two non-trivial vertex groups) that differ by a sequence of three valence-two-homotopies.

Definition 4.21. Let $\mathcal{A}_0 \sqsubset \ldots \sqsubset \mathcal{A}_l \sqsubset \mathcal{B}_0 \sqsubset \ldots \sqsubset \mathcal{B}_m$ be a chain of free factor systems of A. We define

$$X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)$$

as the poset of all free splittings S of A such that one has $\mathcal{V}(S) \in \{\mathcal{A}_0, \ldots, \mathcal{A}_l\}$ and $\mathcal{B}_i | S$ is a subgraph of S for every $0 \leq i \leq m$.

By definition, we have $X(\mathcal{A} : \{[A]\}) = L(A, \mathcal{A})$ and more generally, the poset $X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)$ is contained in the union of $L(A, \mathcal{A}_0), \ldots, L(A, \mathcal{A}_l)$. See Fig. 4.3 for examples of these posets.

The main technical result we prove in order to show contractibility of relative free splitting complexes is that for all chains of free factor systems, the poset $X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)$ is contractible (see Proposition 4.27). This is done in the next three subsections.

4.4.3 Blow-up construction

We start by defining a poset map Ψ from $X(\mathcal{A} : \mathcal{B})$ to a contractible poset and claim that the fibres of this map are contractible.

Let G be a finite connected graph and $v \in V(G)$ be a vertex of valence two with adjacent edges $e_1 \neq e_2$. We define G^v to be the graph obtained from G by replacing the segment $e_1 v e_2$ by a new edge e_v; i.e.

$$V(G^v) = V(G) \setminus \{v\}, \qquad E(G^v) = \{e_v\} \cup E(G) \setminus \{e_1, e_2\}$$

and e_v connects the endpoints of e_1 and e_2 that are not equal to v. Replacing G by G^v is called a *valence-two-homotopy* (see Fig. 4.4). Whenever we have a marked graph $(G, \mathfrak{m}) \in L$, every valence-two-homotopy of the combinatorial graph G induces a valence-two-homotopy of marked graphs, changing $(G, \mathfrak{m})$ to $(G^v, \mathfrak{m}')$ for a marking $\mathfrak{m}'$, well-defined up to equivalence of marked graphs, see [BH92, p. 13].

Let $\tilde{L}$ denote the poset of all marked graphs of groups $\tilde{G}$ such that after applying a finite number of valence-two-homotopies to $\tilde{G}$, we obtain an element $G \in L$. In other words, the elements of $\tilde{L}$ are obtained by applying the inverse of a valence-two-homotopy to $G \in L$, introducing new midpoints to edges of G. In contrast to the elements in L, they can thus contain vertices of valence two that have trivial vertex group—this is not true for the elements of L. The partial order on $\tilde{L}$ is again given by the collapse relation, extending the partial order on L. Observe that for every $\tilde{G} \in \tilde{L}$, there is a unique maximal $G \in L$ such that $\tilde{G}$ collapses to G, i.e. $\tilde{G} \geq G$ in $\tilde{L}$; it is obtained by removing all the valence-two vertices with trivial vertex group. This allows us to show:

Lemma 4.22. *There is a deformation retraction $\tilde{L} \to L$.*

Proof. The map $f\colon \tilde{L} \to L$ sending $\tilde{G} \in \tilde{L}$ to the unique maximal element $G \in L$ such that $\tilde{G} \geq G$ is a monotone poset map restricting to the identity on L. Hence, the claim follows from Corollary 2.5. $\square$

Now let $\mathcal{B} = \{[B_1], \ldots, [B_{k'}]\}$ be a free factor system in A such that $\mathcal{A} \sqsubseteq \mathcal{B}$. Let $G \in L = L(A, \mathcal{A})$ such that $H = \mathcal{B}|G$ is a subgraph of G and let $H_1, \ldots, H_{k'}$ be the connected components of H. Each H_i can be seen as an element of the spine of another relative Outer space: For this, choose a presentation of G as a graph of groups that assigns to H_i the fundamental group B_i. The non-trivial vertex groups of this graph of groups are given by $A_1^{x_1}, \ldots, A_k^{x_k}$ for some $x_1, \ldots, x_k \in A$. We obtain a free factor system of B_i by setting

$$\mathcal{A}'_i \coloneqq \left\{[A_j^{x_j}] \,\middle|\, \text{the vertex corresponding to } A_j^{x_j} \text{ is contained in } H_i\right\},$$

where $[\cdot]$ denotes the conjugacy class in B_i. Now H_i can be seen as an element of $\tilde{L}(B_i, \mathcal{A}'_i)$. After applying the deformation retraction from Lemma 4.22, we thus obtain a poset map

$$\begin{aligned} \Psi\colon X(\mathcal{A} : \mathcal{B}) &\to L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A}) \\ G &\mapsto (G/H, H_1, \ldots, H_{k'}), \end{aligned}$$

where $L(\mathcal{B}, \mathcal{A}) \coloneqq L(B_1, \mathcal{A}'_1) \times \cdots \times L(B_{k'}, \mathcal{A}'_{k'})$.

There are two cases which deserve additional comments because they are not covered by the definition of the spine given in Section 4.1: If $\mathcal{A}'_i = \{[B_i]\}$, then H_i is a graph consisting of a single vertex with vertex group B_i. In this case, we interpret $L(B_i, \mathcal{A}'_i)$ as the singleton consisting of this graph. If B_i is infinite cyclic and $\mathcal{A}'_i = \emptyset$, we interpret $L(B_i, \emptyset)$ as the singleton consisting of a graph with one vertex and one edge whose fundamental group is identified with $B_i \cong \mathbb{Z}$.

Note that for any element $(S, H) \in L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A})$, we can obtain any $G \in \Psi^{-1}(S, H)$ by "blowing up" the vertex v_i stabilised by B_i into the graph H_i. Thus, the map Ψ is surjective. However, it is in general not injective, as when trying to reconstruct G from the pair (S, H), one faces two ambiguities: The first one occurs because for each i, one can choose where to attach the edges of S that are adjacent to v_i to the graph H_i. The second ambiguity arises because before blowing up v_i to H_i, one can change the marking of each of the adjacent edges of S by an element of $\pi_1(H_i) = B_i$. These two choices can be used to parametrise the fibre $\Psi^{-1}(S, H)$.

We now use this point of view to show that the fibres of the map Ψ are contractible. Then because of the contractibility of $L(A, \mathcal{B})$ and $L(\mathcal{B}, \mathcal{A})$, it will follow that $X(\mathcal{A} : \mathcal{B})$ is contractible.

Proposition 4.23. *(Blow-up construction) For $(S, H) \in L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A})$, the fibre $\Psi^{-1}(S, H)$ is contractible.*

4.4.4 Proof of Proposition 4.23

While reading the following proof, the reader might find it helpful to have in parallel a look at the last subsection of Section 4.4.4 (p. 48ff.). It presents in detail an example of the constructions that follow and might help in getting a better understanding of the ideas behind the rather formal definitions.

One component

For simplicity, let us first assume that $\mathcal{B}$ consists of a single free factor B that is not infinite cyclic.

Subdivide all loops incident at the vertex v of S that is stabilised by B into two edges. Let m be the number of edges of S incident at v and let $E_1, \ldots, E_m$ denote the outgoing edges incident at v. Let $\tilde{H}$ be the Bass–Serre tree corresponding to H. The space $\Psi^{-1}(S, H)$ is the subcomplex of $X(\mathcal{A} : \mathcal{B})$ spanned by marked graphs obtained by blowing up the vertex v of S to the subgraph H. Its geometric realisation $\left\|\Psi^{-1}(S, H)\right\|$ can be naturally seen as a subspace of the Outer space of A relative to $\mathcal{A}$. We claim that it is homeomorphic to the m-fold product of $\tilde{H}$, a contractible space. We now construct a map $f\colon \tilde{H}^m \to \left\|\Psi^{-1}(S, H)\right\|$ and show that it is a homeomorphism. The idea for defining the map is the following: An m-tuple of points in H gives us an attaching point for the initial end points of the edges E_i of S and varying the tuple continuously corresponds to sliding the edges E_i along the graph H. If we slide one of these attaching points along a loop in H, the marking of the underlying graph changes. This is encoded by considering points in the Bass–Serre $\tilde{H}$ instead of points in H itself.

We first set up some notation and choose a lift of H to $\tilde{H}$. Let $e_1, e_2, \ldots, e_q$ be the collection of edges of H. Choose a base point v_1 in H and a maximal forest F_H. Using the marking of H, the edges of H not contained in F_H are labelled and oriented (see Section 4.1.3; this is not to be confused with the labelling of the vertices of H as defined in Section 4.3). Also choose orientations for the edges in F_H. Let o_j and t_j denote the initial and terminal end points of the edge e_j, respectively. Denote the label of e_j by $\alpha_j \in B$, where α_j is trivial if $e_j \in F_H$. Consider a metric on H where each edge has length one. Choose a lift $\tilde{v}_1$ of v_1 in $\tilde{H}$. Let $\tilde{F}_H$ be the lift of F_H that contains $\tilde{v}_1$. Let $v_1, \ldots, v_l$ denote the vertices of H and $\tilde{v}_1, \ldots, \tilde{v}_l$ the respective lifts that are contained in $\tilde{F}_H$. Let $\tilde{e}_j$ be the lift of e_j such that $\tilde{o}_j \in \tilde{F}_H$. The tree $\tilde{H}$ gets the lifted metric from H.

Definition of $\boldsymbol{f(P)}$ Consider a point $P = (p_1, \ldots, p_m) \in \tilde{H}^m$. If p_i is a vertex of $\tilde{H}$, then there exists $h_i \in \pi_1(H, v_1)$ and $1 \leq j \leq l$ such that $p_i = h_i\tilde{v}_j$; the element h_i is well-defined up to right multiplication by an element of the vertex stabiliser $\mathrm{Stab}_A(\tilde{v}_j)$. If p_i is in the interior of an edge of $\tilde{H}$, then there is a j such that p_i is specified by the pair $(h_i\tilde{e}_j, l_j(i))$ for $h_i \in \pi_1(H, v_1)$ and $l_j(i) \in (0, 1)$. Given P, we will first construct a marked metric graph of groups $(G, \mathfrak{m}, \ell)$ in the Outer space of A relative to $\mathcal{A}$. Then we will show that this marked metric graph is actually contained in $\|L(A, \mathcal{A})\|$ seen as a subspace of the Outer space.

The graph of groups $\boldsymbol{G}$ For each $1 \leq j \leq q$, order the numbers $0 < l_j(i_1) \leq \cdots \leq l_j(i_j) < 1$. Now subdivide the edge e_j of H (which has length one) according to the numbers $l_j(i_r)$, $i_1 \leq i_r \leq i_j$ and denote the vertices by $u_j(i_r)$. It is possible that a vertex has multiple labels. Let H' be the graph (of groups) obtained from H by this subdivision. From this, G is obtained as follows: if $p_i = (h_i\tilde{e}_j, l_j(i))$ or $h_i\tilde{v}_j$, attach the initial end point of E_i at the vertex $u_j(i)$ of H'. Then remove the vertices of valence two which were introduced when we subdivided the loops of S.

The marking $\mathfrak{m}$ For each $1 \leq j \leq q$, if there exists an $l_j(i_r) \in (0,1)$, then e_j gets subdivided. Define a marking of H' as follows: for every edge e_j that gets subdivided, label the edge $[u_j(i_j), t_j]$ of H', where t_j denotes the terminal end point of e_j, by α_j. For the edges that did not get subdivided, keep the same label as in H. For $p_i = (h_i\tilde{e}_j, l_j(i))$ or $h_i\tilde{v}_j$, label the edge E_i by multiplying the label it inherits from S from the left by h_i^{-1}. The remaining edges of G retain the labelling from S. As mentioned above, the element h_i is only well-defined up to right multiplication by elements g of the adjacent vertex group. Replacing it by $h_i g$ leads to left-multiplying the label by g^{-1}. This does not change the marking, see Section 4.1.3. A way of verifying that this assignment indeed defines a marking of G is the following: Collapse the edges of H' that are unlabelled (α_j being trivial is considered a labelling). Then for an edge E_i with prefix h_i^{-1}, the last letter of h_i coincides with a label of an edge incident at the initial vertex of E_i and hence can be folded away (using Stalling's folds, see e.g. [Wad14]). Continue inductively.

The metric $\boldsymbol{\ell}$ We say an edge of H' is an *edgelet* if at least one of its end points has valence two in H'. Each edgelet e of H' has a length $l'(e)$ induced by the metric on H. For each edge e_j of H which has been subdivided, let $l_j(max)$ be the length of a longest edgelet contained in e_j. Define $l_j(max)$ to be 1 if e_j did not get subdivided. Let $M = \sum_{j=1}^{q} l_j(max)$. Now assign lengths to the edges of H' as follows: For each edgelet e that was part of $e_j \in E(H)$, set the length of e to be $\frac{M}{l_j(max)} \cdot l'(e)$. Edges of S are assigned the length M. Thus we get a metric on G. Now normalise the metric on G to volume one.

Set $f(P) = (G, \mathfrak{m}, \ell)$. We claim that $f(P) \in \left\|\Psi^{-1}(S,H)\right\|$. Indeed, the set $\mathcal{C}$ of edges of G of non-maximal length is precisely the set of non-maximal edgelets of H'. These form a forest, so $f(P)$ is an element of $\|L(A,\mathcal{A})\|$, seen as a subspace of the relative Outer space. Furthermore, $f(P)$ by construction contains a subgraph that differs from H only by valence-two-homotopies and collapsing this subgraph yields the marked graph S.

***f* is a homeomorphism** The image $f(P)$ depends continuously on P: The definition of f is set up in a way such that moving the point $P \in \widetilde{H}^m$ inside the product of these Bass–Serre trees corresponds to sliding the "feet" of the edges E_i, i.e. the points to which these edges are attached, along the graph H. Assume one has $P \in \widetilde{H}^m$ and slightly perturbs it such that: none of its coordinates crosses a vertex of $\tilde{H}$; and for each edge e_j of H, the order $l_j(i_1) \leq l_j(i_2) \leq \cdots \leq l_j(i_j)$ given by the positions of the attaching points of the E_i on e_j does not change, i.e. no foot overtakes another one. Then neither the combinatorial graph nor the marking of $f(P)$ change. Hence, $f(P)$ moves inside an open simplex of the relative Outer space. This movement is specified by the metric of $f(P)$ and it is not hard to see that it depends continuously on P. If on the other hand one of the coordinates passes a vertex of $\widetilde{H}$ or the order of the attaching points on some edge of H changes, the image $f(P)$ continuously moves to an adjacent simplex in the relative Outer space passing through a face of smaller dimension.

Next we claim that f is a homeomorphism. Indeed, the map f is injective because for two different points in $\widetilde{H}^m$ their images will, by construction, differ either in the combinatorial graph, or the marking or the metric. To see that f

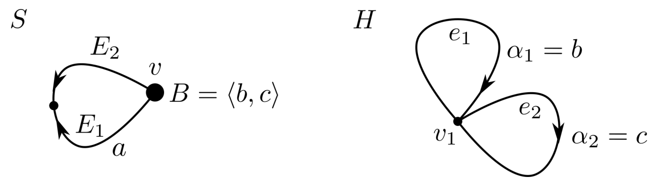

Figure 4.5: $(S, H) \in L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A})$.

is also surjective, take a point $(G, \mathfrak{m}, \ell) \in \left\|\Psi^{-1}(S, H)\right\|$. Then the data from G and $\mathfrak{m}$ provides the information $h_i\tilde{e}_j$ or $h_i\tilde{v}_j$ and the metric ℓ allows us to solve for the lengths $l_j(i_r)$ to give the precise gluing points.

Infinite cyclic components If B is infinite cyclic, a blow-up of S is invariant under conjugation by elements of B. Thus fixing the attaching point for (any) one edge incident at v, we get that $\left\|\Psi^{-1}(S, H)\right\|$ is homeomorphic to $\widetilde{H}^{m-1}$. Since $\widetilde{H}$ is contractible, we get the desired result.

Several components

The general case where $\mathcal{B} = \{B_1, \ldots, B_{k'}\}$ and $(S, H) = (S, H_1, \ldots, H_{k'})$ follows easily from the considerations above: Let m_i be the number of edges of S (after subdividing the loops) incident at the vertex of S that has vertex group B_i. With this notation $\left\|\Psi^{-1}(S, H)\right\|$ is homeomorphic to $\tilde{H_1}^{m_1} \times \cdots \times \tilde{H_k}^{m_{k'}}$ (if B_i is infinite cyclic, then the exponent changes to $m_i - 1$). To define the homeomorphism, take $P \in \tilde{H_1}^{m_1} \times \cdots \times \tilde{H_k}^{m_{k'}}$ and apply the blow-up construction described above at each vertex v_i independently to get a (combinatorial) marked graph $(G, \mathfrak{m})$. The graph G now has a (disconnected) subgraph H' whose components are subdivisions of $H_1, \ldots, H_{k'}$. The metrics on the H_i determine a length function on the edges of H'. This allows us just as above to define a metric ℓ on G giving the same length to all maximal edgelets of H' and the edges coming from S. □

An example

Let $A = F_3 = \langle a, b, c\rangle$, $\mathcal{A} = \emptyset$ and $B = \langle b, c\rangle$. Let $S \in L(A, \{[B]\})$ be the graph of groups consisting of a single vertex v with vertex group B. Subdivide this loop into two edges E_1 and E_2 starting at v and assume that E_1 is labelled by a. We have $m = 2$. Let $H \in L(B, \emptyset)$ be a rose with two petals e_1 and e_2, such that e_1 is labelled by $\alpha_1 = b$ and e_2 is labelled by $\alpha_2 = c$. The maximal forest F_H here is equal to the base point v_1, which is defined as the unique vertex of H (see Fig. 4.5). We have $l = 1$.

The corresponding Bass–Serre tree $\tilde{H}$ is simply the Cayley graph of B with respect to the generating set $\{b, c\}$, which is depicted in Fig. 4.6. Its vertices are identified with the elements of B. As a lift $\tilde{v}_1 = \tilde{F}_H$ of v_1, we choose the vertex corresponding to the identity $\mathrm{id} \in B$. Accordingly, the initial end points $\tilde{o}_1$ and $\tilde{o}_2$ of the lifts $\tilde{e}_1$ and $\tilde{e}_2$ are both equal to $\tilde{v}_1$ and their terminal end points are the vertices corresponding to b and c. We want to construct a homeomorphism

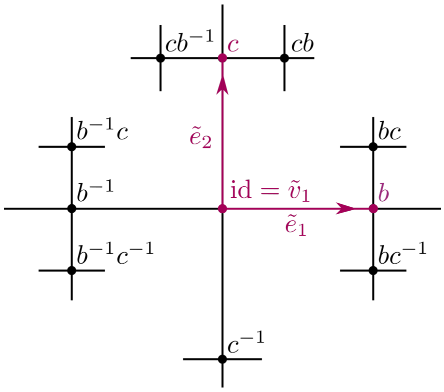

Figure 4.6: A part of the Bass–Serre tree $\tilde{H}$; the chosen lift of H is marked in magenta.

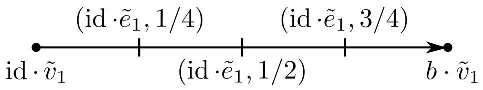

Figure 4.7: Parametrisation of points along $\tilde{e}_1$.

$f\colon \tilde{H}^2 \to \left\|\Psi^{-1}(S, H)\right\|$. In what follows, we describe the image of the square $\tilde{e}_1^2 \subseteq \tilde{H}^2$.

Fig. 4.7 shows how we parametrise points $p_1 \in \tilde{H}$ along the edge $\tilde{e}_1$. Let $p_1 = (\mathrm{id} \cdot \tilde{e}_1, 1/4)$ and $p_2 = (\mathrm{id} \cdot \tilde{e}_1, 1/2)$. We will now describe the image of $P = (p_1, p_2)$. Following the notation above, we have $h_1 = h_2 = \mathrm{id} \in B$ and $l_1(1) = 1/4 \leq l_1(2) = 1/2$. To obtain the graph H', subdivide the edge e_1 in H twice and denote the new vertices by $u_1(1/4)$ and $u_1(1/2)$. From this, G is obtained by attaching E_1 at $u_1(1/4)$ and E_2 at $u_1(1/2)$. The marking of this graph is given by assigning the label b to the edge $s_3 := [u_1(1/2), v_1]$ in H' (this is the terminal segment of what used to be the edge e_1 in H). The edges $s_1 := [v_1, u_1(1/4)]$ and $s_2 := [u_1(1/4), u_1(1/2)]$ remain unlabelled and the remaining edges keep their labels from H (see Fig. 4.8). Finally, we assign a metric to the edges of G. Its edgelets are given by s_1, s_2 and s_3. We have $l'(s_1) = 1/4 = l'(s_2)$, $l'(s_3) = l_1(max) = 1/2$ and $l_2(max) = 1$. Hence, $M = 3/2$ and the length of the edges of H' are given by $l(s_1) = 3/4 = l(s_2)$ and $l(s_3) = 3/2 = l(e_2)$.

The images of other points in $\tilde{e}_1^2$ are obtained similarly, see Fig. 4.9 for a schematic picture.

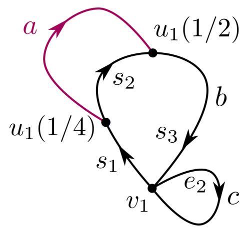

Figure 4.8: $G = f(p_1, p_2)$ with the correct marking and metric. The graph H' is coloured in black.

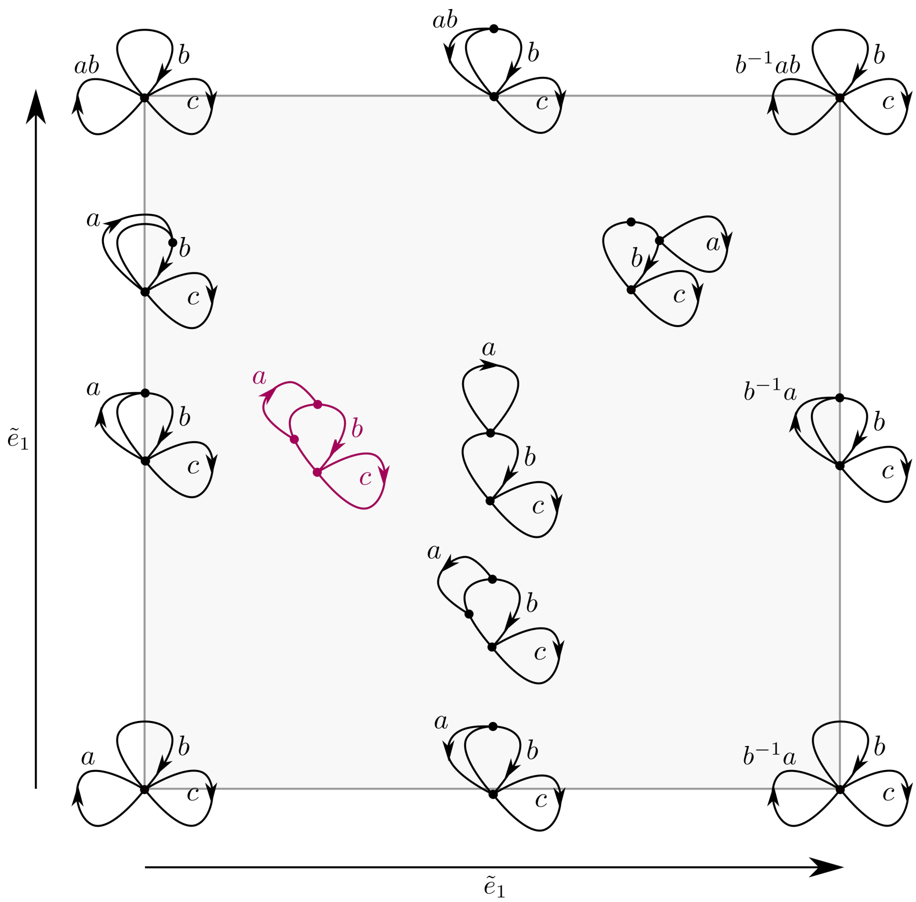

Figure 4.9: Some points in the image of $\tilde{e}_1^2 \subseteq \tilde{H}^2$ under f. The horizontal direction corresponds to the p_1- and the vertical to the p_2-coordinate. The graph $f((\mathrm{id}\cdot\tilde{e}_1, 1/4), (\mathrm{id}\cdot\tilde{e}_1, 1/2))$ is marked in magenta. The metrics of the graphs are not exactly those given by the definition of f but have been adapted for better visibility.

4.4.5 Contractibility of $\mathbf{X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)}$

Using Proposition 4.23 we now show that $X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)$ is contractible.

Lemma 4.24. *Let $(S, H) \in L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A})$ and set*

$$Y := (L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A}))_{\geq (S,H)}$$

Then there is a monotone poset map $f\colon \Psi^{-1}(Y) \to \Psi^{-1}(S, H)$ which restricts to the identity on $\Psi^{-1}(S, H)$. In particular, the fibre $\Psi^{-1}(Y)$ is homotopy equivalent to $\Psi^{-1}(S, H)$.

Proof. Let $(S', H') \in Y$ where $H' = \{H'_1, \ldots, H'_{k'}\}$. Then S and H are obtained from S' and H' by collapsing unique forests $F_{S'}$ and $F_{H'_1}, \ldots, F_{H'_{k'}}$, respectively. Each element of the fibre $\Psi^{-1}((S', H'))$ is a marked graph $G' \in X(\mathcal{A} : \mathcal{B})$ obtained by blowing up the vertex $v'_i \in S'$, stabilised by B_i, to the graph H'_i. Up to valence-two-homotopies, we can view $H' = \bigcup_{i=1}^{k'} H'_i$ as a subgraph of G'. In this way, the union $F' := F_{S'} \cup F_{H'_1} \cup \cdots \cup F_{H'_k}$ can be seen as a subgraph of G'.

We claim that this subgraph is a forest. Suppose not. Then there is a loop $l \in F'$ crossing each edge at most once. As $F_{H'} := \bigcup_{i=1}^{k'} F_{H'_i} \subset H'$ is a forest, we know that l cannot be completely contained in $H' \subset G'$. Hence, collapsing H' maps l to a non-trivial loop in $F_{S'} \subset S'$ which is a contradiction.

Collapsing $F' \subset G'$, we obtain a graph G which lies in $\Psi^{-1}(S, H)$. This defines a map

$$\begin{aligned} f : \Psi^{-1}(Y) &\to \Psi^{-1}(S, H) \\ G' &\mapsto G \end{aligned}$$

restricting to the identity on $\Psi^{-1}(S, H) \subseteq \Psi^{-1}(Y)$. If $G' \geq G'' \in \Psi^{-1}(Y)$, we have $\Psi(G') \geq \Psi(G'') \geq (S, H)$, so the collapse map $\Psi(G') \to (S, H)$ can be written as a concatenation of collapses $\Psi(G') \to \Psi(G'') \to (S, H)$. Whence we know that the forest $F'' \subset G''$ is obtained from $F' \subset G'$ by applying the collapse $G' \to G''$, so $f(G') \geq f(G'')$. This means that f is a monotone poset map; the second claim follows from Corollary 2.5. ☐

Lemma 4.25. *For all free factor systems $\mathcal{A} \sqsubseteq \mathcal{B}$ of A, the poset $X(\mathcal{A} : \mathcal{B})$ is contractible.*

Proof. The image of the map Ψ is the contractible poset $L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A})$. Using Lemma 2.1, the claim hence is an immediate consequence of Proposition 4.23 and Lemma 4.24. ☐

We now generalise Lemma 4.25 to arbitrary chains of free factor systems. We start by proving the following:

Lemma 4.26. *For all chains of free factor systems $\mathcal{A} \sqsubset \mathcal{B}_0 \sqsubset \ldots \sqsubset \mathcal{B}_m$ in A, the poset $X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)$ is contractible.*

Proof. For this proof, we write $X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)$ as $X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)[A]$ to emphasise that this poset contains free splittings of A. Our proof is by induction on m, the hypothesis being that for all finitely generated groups A and all chains of free factor systems $\mathcal{A} \sqsubset \mathcal{B}_0 \sqsubset \ldots \sqsubset \mathcal{B}_m$ of A, the poset $X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)[A]$ is contractible. The base case where $m = 0$ was proved in Lemma 4.25. Now fix m and assume that the claim holds up to $m - 1$.

Let $\mathcal{B}_m = \{[B_1], \ldots, [B_{k'}]\}$. By definition, the poset $X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)[A]$ is contained in $X(\mathcal{A} : \mathcal{B}_m)[A]$. Let

$$\begin{aligned} \Psi \colon X(\mathcal{A} : \mathcal{B}_m)[A] &\to L(A, \mathcal{B}_m) \times L(B_1, \mathcal{A}'_1) \times \cdots \times L(B_{k'}, \mathcal{A}'_{k'}) \\ G &\mapsto (G/H, H_1, \ldots, H_{k'}), \end{aligned} \tag{4.1}$$

be the poset map defined in Section 4.4.3 and let Φ be the restriction of Ψ to $X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)[A]$.

In contrast to Ψ itself, the restriction Φ is not surjective: Let $\mathcal{B}^i_0, \ldots, \mathcal{B}^i_{m-1}$ be the free factor systems of B_i given by intersecting the free factor systems $\mathcal{B}_0, \ldots, \mathcal{B}_{m-1}$ with B_i—this is well-defined because we have $\mathcal{B}_0 \sqsubset \ldots \sqsubset \mathcal{B}_{m-1} \sqsubset \mathcal{B}_m$. Now let $G \in X(\mathcal{A} : \mathcal{B}_0, \ldots, \mathcal{B}_m)[A]$ and let H_i be the $(i+1)$st coordinate of $\Phi(G)$ as in Eq. 4.1. As G has a subgraph with fundamental group $\mathcal{B}_j$, it follows that $H_i \in L(B_i, \mathcal{A}'_i)$ must have a subgraph with fundamental group $\mathcal{B}^j_i$ for all $1 \leq j \leq m - 1$. Hence, the image of Φ is given by

$$L(A, \mathcal{B}_m) \times X(\mathcal{A}'_1 : \mathcal{B}^1_0, \ldots, \mathcal{B}^1_{m-1})[B_1] \times \cdots \times X(\mathcal{A}'_{k'} : \mathcal{B}^{k'}_0, \ldots, \mathcal{B}^{k'}_{m-1})[B_{k'}].$$

By our inductive hypothesis, this product forms a contractible poset, hence it suffices to show that all fibres of Φ are contractible. To see this, note that for every $(S, H) \in \mathrm{im}(\Phi)$, we have $\Phi^{-1}((S, H)) = \Psi^{-1}((S, H))$. This is contractible by Proposition 4.23. It remains to check that the preimage of $\mathrm{im}(\Phi)_{\geq (S,H)}$ is contractible as well. This follows from Corollary 2.5: By Lemma 4.24, there is a monotone poset map

$$f \colon \Psi^{-1}\left(\mathrm{im}(\Psi)_{\geq (S,H)}\right) \to \Psi^{-1}(S, H)$$

which restricts to the identity on $\Psi^{-1}(S, H) = \Phi^{-1}(S, H)$. Restricting f, we obtain a poset map $\Phi^{-1}(\mathrm{im}(\Phi)_{\geq (S,H)}) \to \Phi^{-1}(S, H)$ with the same properties. □

Proposition 4.27. *For every chain $\mathcal{A}_0 \sqsubset \ldots \sqsubset \mathcal{A}_l \sqsubset \mathcal{B}_0 \sqsubset \ldots \sqsubset \mathcal{B}_m$ of free factor systems of A, the poset $X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)$ is contractible.*

Proof. The proof is by induction on l. By Lemma 4.26, the claim holds true for all m if $l = 0$. Now assume that it holds true up to $l - 1$.

Then in particular, the posets

$$\begin{gathered} X_{l-1} \coloneqq X(\mathcal{A}_0, \ldots, \mathcal{A}_{l-1} : \mathcal{B}_0, \ldots, \mathcal{B}_m), \qquad X_l \coloneqq X(\mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m) \\ \text{and} \qquad X_{l-1,l} \coloneqq X(\mathcal{A}_0, \ldots, \mathcal{A}_{l-1} : \mathcal{A}_l, \mathcal{B}_0, \ldots, \mathcal{B}_m) \end{gathered}$$

are contractible. By definition $X_{l-1,l}$ is the subposet of X_{l-1} consisting of all those $S \in X_{l-1}$ that collapse to some free splitting in X_l. For each such $S \in X_{l-1}$, there is a unique maximal splitting $S' \in X_l$ on which S collapses, namely $S' = S/(\mathcal{A}_l|S)$. Hence, the map

$$\begin{aligned} X_l \cup X_{l-1,l} &\to X_l \\ S &\mapsto \begin{cases} S' & , S \in X_{l-1,l}, \\ S & , S \in X_l, \end{cases} \end{aligned}$$

induces a deformation retraction $\|X_{l-1,l} \cup X_l\| \to \|X_l\|$.

It follows that $\|X(\mathcal{A}_0, \ldots, \mathcal{A}_l : \mathcal{B}_0, \ldots, \mathcal{B}_m)\| = \|X_{l-1} \cup X_l\|$ is obtained by gluing together $\|X_{l-1}\|$ and $\|X_l\|$ along $\|X_{l-1,l}\|$. Now $\|X_{l-1}\|, \|X_{l-1,l}\|$ and $\|X_l\|$ are contractible by assumption, whence the claim follows. □

4.4.6 Proof of contractibility of free splitting complexes

We are now ready to prove the contractibility of relative free splitting complexes and of the following poset which will occur in the study of the free factor complex $\mathcal{F}(A, \mathcal{A})$ later on:

Definition 4.28. Let $\mathcal{FS}^1(A; \mathcal{A})$ be the poset of all free splittings S of A that have *exactly one* conjugacy class $\mathcal{V}(S)$ of non-trivial vertex stabilisers and such that $\mathcal{A} \sqsubseteq \mathcal{V}(S) \sqsubset \{[A]\}$.

Collecting the results of this section, we prove Theorem F, which is the first item of the following:

Theorem 4.29. *Let A be a finitely generated group and $\mathcal{A}$ a proper free factor system in A.*

1. *The poset of free splittings $\mathcal{FS} = \mathcal{FS}(A, \mathcal{A})$ is contractible.*
2. *If $\mathcal{A}$ has only one component, $\mathcal{A} = \{[A']\}$, the poset $\mathcal{FS}^1(A; \mathcal{A})$ is contractible.*

Proof. We start by proving the first claim: Each simplex σ in the order complex $\Delta(\mathcal{FS})$ is of the form $S_0 \to \ldots \to S_k$ where each S_i is a free splitting of A collapsing to S_{i+1}. Furthermore, the vertex group systems of these free splittings form a chain $\mathcal{V}(S_0) \sqsubseteq \ldots \sqsubseteq \mathcal{V}(S_k)$ of free factor systems such that $\mathcal{A} \sqsubseteq \mathcal{V}(S_i)$ for all i. It follows that σ is contained in $\Delta(X(\mathcal{A}, \mathcal{V}(S_0), \ldots, \mathcal{V}(S_k) : [A]))$. Hence the realisation $\|\mathcal{FS}\|$ can be written as a union

$$\|\mathcal{FS}\| = \bigcup_{\mathcal{A} \sqsubset \mathcal{A}_1 \sqsubset \ldots \sqsubset \mathcal{A}_l} \|X(\mathcal{A}, \mathcal{A}_1, \ldots, \mathcal{A}_l : [A])\|.$$

By Proposition 4.27, each $\|X(\mathcal{A}, \mathcal{A}_1, \ldots, \mathcal{A}_l : [A])\|$ is contractible. Furthermore, one has

$$\|X(\mathcal{A}, \mathcal{A}_1, \ldots, \mathcal{A}_l : [A])\| \cap \|X(\mathcal{A}, \mathcal{B}_1, \ldots, \mathcal{B}_m : [A])\| = \|X(\mathcal{A}, \mathcal{C}_1, \ldots, \mathcal{C}_j : [A])\|$$

where $\mathcal{A} \sqsubset \mathcal{C}_1 \sqsubset \ldots \sqsubset \mathcal{C}_k$ is the longest common subchain of $\mathcal{A} \sqsubset \mathcal{A}_1 \sqsubset \ldots \sqsubset \mathcal{A}_l$ and $\mathcal{A} \sqsubset \mathcal{B}_1 \sqsubset \ldots \sqsubset \mathcal{B}_m$. Consequently, all intersections of these sets are contractible and Lemma 2.7 implies that $\|\mathcal{FS}\|$ is homotopy equivalent to the nerve of this covering. However, as all of these sets contain $\|X(\mathcal{A} : [A])\|$, they intersect non-trivially, so this nerve complex is contractible.

By the same arguments, $\mathcal{FS}^1(A; \{[A']\})$ is homotopy equivalent to the nerve of its covering given by all the sets $\|X([A'], [A_1], \ldots, [A_l] : [A])\|$ where each A_i is a free factor and $[A'] \leq [A_i]$. Again, the intersection of all of these sets contains $\|X([A'] : [A])\|$ and hence is non-empty, so the second claim follows. □

4.5 Factor complexes at infinity

Standing assumptions As before, let $\mathcal{A} = \{[A_1], \dots, [A_k]\}$ be a proper free factor system of A and write $\mathcal{FS} := \mathcal{FS}(A, \mathcal{A})$, $\mathcal{FS}^1 := \mathcal{FS}^1(A, \mathcal{A})$ and $L := L(A, \mathcal{A})$. Let $\mathcal{F} := \mathcal{F}(A, \mathcal{A})$ and $\mathcal{FF} := \mathcal{FF}(A, \mathcal{A})$ be the corresponding complexes of free factors and free factor systems.

In this short section, we connect the factor complexes defined in Section 4.2 to subposets of $\mathcal{FS}$. Let

$$b\mathcal{FS} := \{S \in \mathcal{FS} \mid \mathcal{V}(S) \in \mathcal{FF}\} = \mathcal{FS} \setminus L \text{ and}$$
$$b\mathcal{FS}^1 := \{S \in \mathcal{FS} \mid \mathcal{V}(S) \in \mathcal{F}\}.$$

Note that we have $b\mathcal{FS}^1 \subseteq b\mathcal{FS}$. These posets should be regarded as the posets of all splittings in $\mathcal{FS}$ and $\mathcal{FS}^1$ that have a non-minimal vertex group system: For $b\mathcal{FS}$ this is true on the nose; it consists of all those splittings $S \in \mathcal{FS} = \mathcal{FS}(A, \mathcal{A})$ such that $\mathcal{A} \sqsubset \mathcal{V}(S)$. For $b\mathcal{FS}^1$ this is true if we add the extra assumption that $k = 1$, i.e. $\mathcal{A}$ has only one component; in this case, we have $b\mathcal{FS}^1 = \mathcal{FS}^1 \setminus L$. These posets can be seen as boundary structure of the corresponding Outer space $\mathcal{O}(A, \mathcal{A})$. This was explained for the case where $A = F_n$ in Section 4.1.4.

The next proposition follows almost immediately from the contractibility of the relative free splitting complexes established in the preceding section. It implies the first two sentences of Theorem E.

Proposition 4.30. *The following hold true:*

1. *$b\mathcal{FS}$ is homotopy equivalent to $\mathcal{FF}$.*
2. *$b\mathcal{FS}^1$ is homotopy equivalent to $\mathcal{F}$.*

Proof. Assigning to each splitting $S \in b\mathcal{FS}$ the free factor system $\mathcal{V}(S)$ given by its non-trivial vertex stabilisers defines a poset map $f \colon b\mathcal{FS} \to \mathcal{FF}^{op}$. As there is a natural isomorphism of the order complexes $\Delta(\mathcal{FF}^{op}) \cong \Delta(\mathcal{FF})$, we will interpret f as an order-inverting map $f \colon b\mathcal{FS} \to \mathcal{FF}$.

For any free factor system $\mathcal{B} \in \mathcal{FF}$, the fibre $f^{-1}((\mathcal{FF})_{\geq \mathcal{B}})$ is equal to the poset $\mathcal{FS}(A, \mathcal{B})$ of free splittings relative to $\mathcal{B}$. This poset is contractible by the first point of Theorem 4.29.

The image $f(b\mathcal{FS}^1)$ is equal to $\mathcal{F}$, so restricting f provides us with a poset map $g : b\mathcal{FS}^1 \to \mathcal{F}$. For any $B \in \mathcal{F}$, the fibre $g^{-1}(\mathcal{F}_{\geq B})$ is equal to the poset $\mathcal{FS}^1(A, \{[B]\})$, so the second point of Theorem 4.29 finishes the proof. □

Remark 4.31. The map $f \colon b\mathcal{FS} \to \mathcal{FF}$ defined in the proof of Proposition 4.30 has already been used to study the geometry of the complexes in question:

In [HM, Section 6.2], the authors define "projection maps" $\pi : \mathcal{FS} \to \mathcal{FF}$ and show that these maps are Lipschitz with respect to the metrics on the 1-skeleta of $\mathcal{FS}$ and $\mathcal{FF}$ assigning length one to each edge. The map f can be seen as the restriction of such a projection map to $b\mathcal{FS}$ and hence is Lipschitz as well.

In the "absolute" setting where $A = F_n$ and $\mathcal{A} = \emptyset$, Hilion and Horbez in [HH17, Section 8] consider the poset $\mathcal{FS}^c \subset b\mathcal{FS}^1$ of all free splittings of F_n whose corresponding graph of groups is a rose with non-trivial vertex group, i.e.

those free splittings of $b\mathcal{FS}^1$ having only one orbit of vertices. They show that the inclusion $\mathcal{FS}^c \subset b\mathcal{FS}^1$ defines a quasi-isometry of the 1-skeleta and that the restriction $f : \mathcal{FS}^c \to \mathcal{F}$ has quasi-convex fibres. This is used to deduce hyperbolicity of $\mathcal{F}$. They formulate their results using the language of sphere systems.

4.6 Higher connectivity of factor complexes

Standing assumptions We keep the assumptions of Section 4.5. In particular, the group A can be written as $A = F_n * A_1 * \dots * A_k$. In this section, we assume that $n \geq 2$. This implies that $\mathcal{F}$ and $\mathcal{FF}$ are non-empty.

Recall our convention that when we talk about "(core) subgraphs" of a free splitting of A, we mean (core) subgraphs of the corresponding *labelled* graphs (see Section 4.4.2).

Let $\mathcal{Z}$ be the subposet of $L \times b\mathcal{FS}$ given by all pairs (G, S) such that $G \in L$ and $S = G/H$ is obtained by collapsing a proper core subgraph $H \subset G$. Let $p_1 : \mathcal{Z} \to L$ and $p_2 : \mathcal{Z} \to b\mathcal{FS}$ be the natural projection maps.

Let $\mathcal{Z}^1$ be the subposet of $L \times b\mathcal{FS}^1$ given by all pairs (G, S) such that $S = G/H$ is obtained by collapsing a proper *connected* core subgraph $H \subset G$. Observe that this implies that H, seen as a labelled graph, must have non-trivial fundamental group. Let $q_1 : \mathcal{Z}^1 \to L$ and $q_2 : \mathcal{Z}^1 \to b\mathcal{FS}^1$ be the natural projection maps. The poset $\mathcal{Z}^1$ is a subposet of $\mathcal{Z}$ and q_1 and q_2 are the restrictions of the projection maps p_1 and p_2.

We think of $\mathcal{Z}$ and $\mathcal{Z}^1$ as *thickened versions* of $b\mathcal{FS}$ and $b\mathcal{FS}^1$, respectively and want to use them to study properties of these posets. The projection maps of these products play an important role for this. For an element $G \in L$, the preimage $p_1^{-1}(G)$ is naturally identified with a subset of $b\mathcal{FS}$ consisting of graphs $S = G/H$ to which G collapses. On the other hand, for $S \in b\mathcal{FS}$, the preimage $p_2^{-1}(S)$ is a subset of L consisting of graphs G which collapse to S. Following the explanations given in Section 4.1.4, the picture one should have in mind is the following: The map p_2 is the natural projection from the interior of Outer space, i.e. its spine L, to its boundary, i.e. $\mathcal{FS} \setminus L = b\mathcal{FS}$, and p_1 is the corresponding projection from the boundary of Outer space to its interior. All of this can be understood as taking place in the simplicial completion of Outer space, i.e. the free splitting complex $\mathcal{FS}$. (For a schematic picture, see Fig. 4.10.)

We proceed in two steps: First we show that the projections p_2 and q_2 to the second factors define homotopy equivalences; then we apply the results of Section 4.3 to understand the fibres of the projections p_1 and q_1.

4.6.1 Projection to the second factor

We can deformation retract the fibres of p_2 and q_2 to simpler subposets:

Lemma 4.32.

1. *For all $S \in b\mathcal{FS}$, the fibre $p_2^{-1}(b\mathcal{FS}_{\geq S})$ deformation retracts to $p_2^{-1}(S)$.*
2. *For all $S \in b\mathcal{FS}^1$, the fibre $q_2^{-1}(b\mathcal{FS}^1_{\geq S})$ deformation retracts to $q_2^{-1}(S)$.*

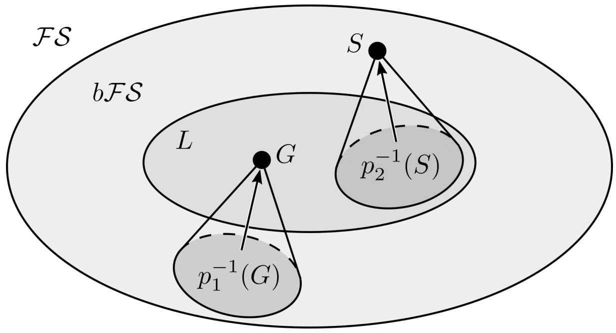

Figure 4.10: The fibres of p_1 and p_2.

Proof. Let $F := p_2^{-1}(b\mathcal{FS}_{\geq S})$. If (G', S') is an element of F, there are collapse maps $G' \to S'$ and $S' \to S$. Concatenating these maps, we see that S is obtained from G' by collapsing a subgraph $H' \subset G'$. As $S \in b\mathcal{FS}$, the graph H' can be written as the union of a (possibly trivial) forest T' and its core $\mathring{H}'$. We define a map

$$\begin{aligned} f\colon F &\to p_2^{-1}(S) \\ (G', S') &\mapsto (G'/T', S). \end{aligned}$$

As $S = (G'/T')/\mathring{H}'$, the image $(G'/T', S)$ is indeed an element of $p_2^{-1}(S)$. Also if $(G'', S'') \geq (G', S')$ in F, we have $c(T'') \supseteq T'$ which implies $G''/T'' \geq G'/T'$. Consequently $f\colon F \to p_2^{-1}(S)$ is a well-defined, monotone poset map restricting to the identity on $p_2^{-1}(S)$. It follows from Corollary 2.5 that this defines a deformation retraction.

For the second point, note that the labelled graph corresponding to a splitting $S \in b\mathcal{FS}^1$ has only a single labelled vertex. This implies that if we have $(G', S') \in q_2^{-1}(b\mathcal{FS}^1_{\geq S})$ and $H' \subset G'$ with $G'/H' = S$ as above, then the core $\mathring{H}'$ must be connected. Using this observation, the map f defined above restricts to a monotone poset map $q_2^{-1}(b\mathcal{FS}^1_{\geq S}) \to q_2^{-1}(S)$. So the second claim follows from Corollary 2.5 as well. □

Hence, instead of studying arbitrary fibres, it suffices to consider the preimages of single vertices. We start by using the results from Section 4.4 to show:

Proposition 4.33. *For all $S \in b\mathcal{FS}$, the preimage $p_2^{-1}(S)$ is contractible.*

Proof. Let $\mathcal{B} := \mathcal{V}(S)$. Every element in $p_2^{-1}(S)$ is given by a pair (G, S) such that $H := \mathcal{B}|G$ is a subgraph of G and $S = G/H$. Forgetting the second coordinate—which is constant S—, we can thus view $p_2^{-1}(S)$ as a subposet of $X(\mathcal{A} : \mathcal{B})$ (see Definition 4.21). Let

$$\begin{aligned} \Psi\colon X(\mathcal{A} : \mathcal{B}) &\to L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A}) \\ G &\mapsto (G/H, H) \end{aligned}$$

be the poset map defined in Section 4.4.3. Restricting Ψ to $p_2^{-1}(S)$ and forgetting the first coordinate of $\mathrm{im}(\Psi)$—it is again constant S—, we obtain a poset map

$\psi\colon p_2^{-1}(S) \to L(\mathcal{B}, \mathcal{A})$. Let $H \in L(\mathcal{B}, \mathcal{A})$. By definition, the poset $\psi^{-1}(H)$ is equal to $\Psi^{-1}(S, H)$, which is a contractible by Proposition 4.23. Furthermore, we have

$$\psi^{-1}\left((L(\mathcal{B}, \mathcal{A}))_{\geq H}\right) = \Psi^{-1}\left(\{(S, H') \mid H' \geq H\}\right).$$

By Lemma 4.24, there is a monotone poset map

$$f\colon \Psi^{-1}\left((L(A, \mathcal{B}) \times L(\mathcal{B}, \mathcal{A}))_{\geq (S,H)}\right) \to \Psi^{-1}(S, H)$$

which restricts to the identity on $\Psi^{-1}(S, H)$. Restricting f, we obtain a poset map $\psi^{-1}((L(\mathcal{B}, \mathcal{A})_{\geq H}) \to \Psi^{-1}(S, H) = \psi^{-1}(H)$ which has the same properties. Using Corollary 2.5 and Lemma 2.1, we see that ψ is a homotopy equivalence. Thus, the claim follows from contractibility of $L(\mathcal{B}, \mathcal{A})$. □

The following shows that Proposition 4.33 also provides us with sufficient information about the fibres of q_2.

Lemma 4.34. *For all $S \in b\mathcal{FS}^1$, one has $q_2^{-1}(S) = p_2^{-1}(S)$.*

Proof. This immediately follows from the definitions: The map q_2 is the restriction of p_2, so it is clear that we have $q_2^{-1}(S) \subseteq p_2^{-1}(S)$. The other inclusion follows because for $G \in L$ and a core subgraph $H \subset G$, the collapse G/H being in $b\mathcal{FS}^1$ implies that H is connected. □

In particular, these fibres are contractible.

Corollary 4.35. *The maps $p_2 : \mathcal{Z} \to b\mathcal{FS}$ and $q_2 : \mathcal{Z}^1 \to b\mathcal{FS}^1$ are homotopy equivalences.*

Proof. Using Quillen's fibre lemma, the claim is an immediate consequence of Lemma 4.32, Proposition 4.33 and Lemma 4.34. □

4.6.2 Projection to the first factor

Corollary 4.35 allows us to replace $b\mathcal{FS}$ by its thickened version $\mathcal{Z}$. This has the advantage that $\mathcal{Z}$ possesses a natural projection map p_1 to the contractible poset L which we will study in this subsection.

Lemma 4.36. *For all $G \in L$, the fibre $p_1^{-1}(L_{\leq G})$ is homotopy equivalent to $\mathrm{C}(G)$, the poset of proper core subgraphs of G.*

Proof. Each element of the fibre $p_1^{-1}(L_{\leq G})$ consists of a pair (G', S') where $G' \leq G$ in L and $S' \in b\mathcal{FS}$ is obtained from G' by collapsing a proper core subgraph H'. As G' is obtained from G by collapsing a forest, there is a unique, proper core subgraph H of G making the following diagram commute:

$$\begin{array}{ccc} H & \hookrightarrow & G \\ \downarrow & & \downarrow \\ H' & \hookrightarrow & G' \end{array}$$

H is equal to the core $\pi_{G'}(H')|G$.

Because this diagram commutes, the collapse $G \to G'$ induces a collapse $G/H \to G'/H' = S'$. Hence, we get a monotone poset map

$$\begin{aligned} f : p_1^{-1}(L_{\leq G}) &\to p_1^{-1}(G) \\ (G', S') &\mapsto (G, G/H) \end{aligned}$$

restricting to the identity on $p_1^{-1}(G) \subseteq p_1^{-1}(L_{\leq G})$. Again Corollary 2.5 implies that f defines a deformation retraction.

If H and H' are proper core subgraphs of G, one has $G/H \geq G/H'$ in $b\mathcal{FS}$ if and only if $H \leq H'$ in $\mathrm{C}(G)$. It follows that $p_1^{-1}(G)$ can be identified with $\mathrm{C}(G)^{op}$. Noting that $\|\mathrm{C}(G)^{op}\| \cong \|\mathrm{C}(G)\|$ finishes the proof. □

Lemma 4.37. *For all $G \in L$, the fibre $q_1^{-1}(L_{\leq G})$ is homotopy equivalent to $\mathrm{cC}(G, l)$, the poset of proper connected core subgraphs of the corresponding labelled graph (G, l) that have non-trivial fundamental group.*

Proof. The proof is literally the same as that of Lemma 4.36 after one makes the following observation: Whenever $G' \leq G$ in L and H' is a proper *connected* core subgraph of G', there is a unique, proper *connected* core subgraph $H \subset G$ making the diagram

$$\begin{array}{ccc} H & \hookrightarrow & G \\ \downarrow & & \downarrow \\ H' & \hookrightarrow & G' \end{array}$$

commute. Furthermore, if H' has non-trivial fundamental group, then so does H. □

4.6.3 Homotopy type of $\mathcal{F}$

Corollary 4.38. *The poset $\mathcal{Z}^1$ is $(n-3)$-connected.*

Proof. The projection $q_1 : \mathcal{Z}^1 \to L$ is a map from $\mathcal{Z}^1$ to the contractible poset L. By Lemma 4.37, the fibres of this map are given by the posets of connected core subgraphs studied in Section 4.3.2. These are $(n-3)$-connected by Lemma 4.17 and Proposition 4.19. Hence, the result follows from Lemma 2.2. □

In order to determine the homotopy type of $\mathcal{F}$ we now only have to collect the work done so far:

Theorem 4.39. *The free factor complex $\mathcal{F}$ of $A = F_n * A_1 * \cdots * A_k$ relative to the free factor system $\mathcal{A} = \{[A_1], \ldots, [A_k]\}$ is $(n-2)$-spherical.*

Proof. By Proposition 4.30 and Corollary 4.35, there is a homotopy equivalence $\mathcal{F} \simeq \mathcal{Z}^1$. By Corollary 4.38, the poset $\mathcal{Z}^1$ is $(n-3)$-connected. As $\mathcal{F}$ is $(n-2)$-dimensional (see Proposition 4.7), the claim follows. □

In the case where $A = F_n$ and $\mathcal{A} = \emptyset$, this is Theorem D from the introduction.

Remark 4.40. The complex $\mathcal{F}$ has non-trivial homology in dimension $n-2$ and is in particular non-contractible: Let $\{x_1, x_2, \ldots, x_n\}$ be a basis of F_n and let Σ be the poset of all conjugacy classes of free factors of the form

$$\langle \{x_i \mid i \in I\} \rangle * A_1 * \cdots * A_k$$

where $I \subset \{1, \ldots, n\}$ is a subset of size $1 \leq |I| \leq n-1$. It is easy to see that $\|\Sigma\|$ is a triangulated $(n-2)$-sphere inside $\|\mathcal{F}\|$. (It is isomorphic to a Coxeter complex of type $\mathtt{A}_{n-1}$ and forms an analogue of an apartment in a Tits-building; see also [HV98b, Section 5].) In particular, this shows that $H_{n-2}(\mathcal{F})$ is non-trivial, so the complex cannot be contractible.

Cohen–Macaulayness

The relative formulations allow us to deduce that $\mathcal{F}$ is even Cohen–Macaulay:

Theorem 4.41. *The free factor complex $\mathcal{F} = \mathcal{F}(A, \mathcal{A})$ is homotopy Cohen–Macaulay.*

Proof. We have to show that the link of every s-simplex $\sigma = [B_0] \leq \ldots \leq [B_s]$ in $\Delta(\mathcal{F})$ is $(n-s-3)$-spherical. However, the link of this simplex is by definition given by the following join of posets

$$\mathrm{lk}(\sigma) = \mathcal{F}_{<[B_0]} * ([B_0], [B_1]) * \ldots * ([B_{s-1}], [B_s]) * \mathcal{F}_{>[B_s]}.$$

As pointed out in Section 4.2.1, each B_i can be written in the form

$$B_i = D_i * A_1^{x_1} * \ldots * A_k^{x_k},$$

where D_i is a free group of rank $n - \mathrm{crk}[B_i]$. Using malnormality of free factors, it follows that two subgroups of a free factor B of A are conjugate in A if and only they are conjugate in B [HM, Lemma 2.1]. It follows that there are isomorphisms

$$\begin{aligned} \mathcal{F}_{<[B_0]} &\cong \mathcal{F}(B_0, \{[A_1^{x_1}], \ldots, [A_k^{x_k}]\}), \\ ([B_i], [B_{i+1}]) &\cong \mathcal{F}(B_{i+1}, \{[B_i]\}), \\ \mathcal{F}_{>[B_s]} &\cong \mathcal{F}(A, \{[B_s]\}). \end{aligned}$$

The result now follows from Lemma 2.11 and Theorem 4.39. ☐

Remark 4.42. Combining the preceding theorem with Proposition 4.7, we see that $\mathcal{F}$ is a Cohen–Macaulay coset complex associated to the Fouxe-Rabinovitch group $\mathrm{Out}(A; \mathcal{A}^t)$. Hence, we can apply the results of Section 3.2. In particular, we obtain higher generating families of subgroups for $\mathrm{Out}(A; \mathcal{A})$ (see Theorem 3.11). We do not spell out the details here because a more general version of this result for the case where A is a right-angled Artin group will be given in Section 5.4.2.

4.6.4 Higher connectivity of $\mathcal{FF}$

We now study the complex of free factor systems $\mathcal{FF}$. Here, we restrict ourselves to the case where $A = F_n$. The methods used for $\mathcal{F}$ do not suffice to determine the homotopy type of $\mathcal{FF}$. Instead, we only obtain a lower bound on its connectivity, which we then slightly improve by invoking the fibre theorem Lemma 2.3.

Proposition 4.43. *The poset $\mathcal{Z}$ is $(n-3)$-connected.*

Proof. Consider the first projection $p_1 : \mathcal{Z} \to L$. By Lemma 4.36, the fibre $p_1^{-1}(L_{\leq G})$ is homotopy equivalent to $\mathrm{C}(G)$ for all $G \in L$. Lemma 4.13 and Proposition 4.14 imply that this poset is at least $(n-3)$-connected. Applying Lemma 2.2 finishes the proof. □

For $\mathcal{Z}^1$, we also showed $(n-3)$-connectivity and this was optimal as $\mathcal{F} \simeq \mathcal{Z}^1$ has non-trivial homology in dimension $n-2$ and hence cannot have a higher degree of connectivity (see Remark 4.40). For $\mathcal{Z}$ however, we can further improve the result of Proposition 4.43 because the following lemma provides us with additional information about the fibres of p_1.

Lemma 4.44. *For $G \in L$, let $f : p_1^{-1}(L_{\leq G}) \to \mathcal{Z}$ be the inclusion map. Then the induced map on homotopy groups $f_* : \pi_{n-2}(p_1^{-1}(L_{\leq G})) \to \pi_{n-2}(\mathcal{Z})$ is trivial.*

Proof. Since $\left\|p_1^{-1}(L_{\leq G})\right\|$, $\left\|p_1^{-1}(G)\right\|$, $\|C(G)\|$ and $\|X(G)\|$ are homotopy equivalent to each other, we blur the distinction between their homology and homotopy groups in the following discussion. For $n \geq 4$, there is an isomorphism $\pi_{n-2}(\|X(G)\|) \cong H_{n-2}(\|X(G)\|)$ by Proposition 4.14 and the Hurewicz Theorem. By Alexander duality, we have

$$\tilde{H}^{v(G)-2}(\|\mathrm{For}(G)\|) \xrightarrow[\cong]{\Psi} \tilde{H}_{n-1}(\|\mathrm{Sub}(G)\|\,, \|X(G)\|) \xrightarrow[\cong]{\partial} \tilde{H}_{n-2}(\|X(G)\|)\ .$$

Here $v(G)$ is the number of vertices in G. We want to find a generating set for $\tilde{H}_{n-2}(\|X(G)\|)$. We start by describing a generating set for $\tilde{H}^{v(G)-2}(\|\mathrm{For}(G)\|)$. Let $\{\sigma^i\}_{i=1}^N$ be the collection of $(v(G)-2)$-simplices of $\|\mathrm{For}(G)\|$ oriented appropriately, such that they form a basis for the free abelian group of $(v(G)-2)$-chains on $\|\mathrm{For}(G)\|$. Let ϕ_i be a co-chain on $\|\mathrm{For}(G)\|$ such that $\phi_i(\sigma^i) = 1$ and $\phi_i(\sigma^j) = 0$ for $j \neq i$. Let $[\phi_i]$ denote the corresponding cohomology class. Then $\{[\phi_i]\}_{i=1}^N$ generates $\tilde{H}^{v(G)-2}(\|\mathrm{For}(G)\|)$. In general it might not be a basis.

The collection $\{\sigma^i\}_{i=1}^N$ of facets of $\|\mathrm{For}(G)\|$ is in bijection with the collection of maximal forests of G, denoted $\{\mathcal{E}_i\}_{i=1}^N$. Under the isomorphism $\partial \circ \Psi$ given by Alexander duality, the dual to $[\phi_i]$ is given by $\|\mathrm{Sub}(G - \mathcal{E}_i)\|$. Since $\mathcal{E}_i$ is a maximal forest of G, we have $\mathrm{Sub}(G - \mathcal{E}_i) \cong X(G/\mathcal{E}_i)$, which is homotopy equivalent to an $(n-2)$-sphere. Thus we conclude that $\{\|X(G/\mathcal{E}_i)\|\}_{i=1}^N$ generates $\tilde{H}_{n-2}(\|X(G)\|)$.

Now for each $G/\mathcal{E}_i$, which is a rose, there exists $G_i \in L$ such that G_i has a separating edge and $G_i > G/\mathcal{E}_i$. Since $p_1^{-1}(L_{\leq G_i})$ is contractible by Proposition 4.14 and $p_1^{-1}(G/\mathcal{E}_i) \subset p_1^{-1}(L_{\leq G_i})$, the preimage $p_1^{-1}(G/\mathcal{E}_i)$ is contractible in $\mathcal{Z}$. Furthermore, $G > G/\mathcal{E}_i$, so there is an inclusion $p_1^{-1}(G/\mathcal{E}_i) \subset p_1^{-1}(L_{\leq G})$. The posets $p_1^{-1}(L_{\leq G})$, $p_1^{-1}(G)$ and $X(G)$ are homotopy equivalent, therefore we can conclude that each generator of $\tilde{H}_{n-2}(\|X(G)\|)$ given by $\|X(G/\mathcal{E}_i)\|$, or equivalently $\left\|p_1^{-1}(G/\mathcal{E}_i)\right\|$, is contractible in $\mathcal{Z}$. For $n = 3$, the lemma follows by an explicit computation. □

Remark 4.45. It is possible that in $\mathcal{Z}$, the preimage $p_1^{-1}(G)$ has multiple contractions. This can give rise to higher dimensional spheres in $\mathcal{Z}$. See Example A.1 in the appendix.

We are now ready to prove:

Theorem 4.46. *For $n \geq 2$, the poset $b\mathcal{FS}$ is $(n-2)$-connected.*

Proof. By Corollary 4.35, the poset $b\mathcal{FS}$ is homotopy equivalent to $\mathcal{Z}$. The spine L of Outer space is contractible and it follows from Lemma 4.36 and Proposition 4.14 that the fibres of $p_1 : \mathcal{Z} \to L$ are either $(n-3)$-connected or contractible. Using Lemma 4.44 and applying Lemma 2.3, one gets that $b\mathcal{FS}$ is $(n-2)$-connected. □

Proposition 4.30 immediately implies the following corollary which completes the proof of Theorem E.

Corollary 4.47. *The complex $\mathcal{FF}$ of free factor systems is $(n-2)$-connected.*

For more comments on the optimality of the result obtained here, see Section 4.7.2.

Remark 4.48. The restriction to the case $A = F_n$ comes from the fact that this is the setting of the article [BG]. Furthermore, connectivity results of $\mathcal{FF}$ are not needed for studying the complex $\mathcal{CC}$ associated to the automorphism group of a right-angled Artin group (see Chapter 5). However, it seems plausible that an analogue of Theorem 4.46 holds in the setting of arbitrary finitely generated groups A. The only thing one would need to prove for this is a version of Proposition 4.14 for non-trivially labelled graphs.

4.7 Boundary structures of Outer space

In the last section of this chapter, we summarise the results about the asymptotic geometry of Outer space that we obtained above and put them into context. For simplicity, we restrict ourselves to the case where $A = F_n$ and $\mathcal{A} = \emptyset$. We will keep track of the rank of the free group by adding a subscript and setting $\mathcal{F}_n := \mathcal{F}(F_n, \emptyset)$, $\mathcal{FF}_n := \mathcal{FF}(F_n, \emptyset)$, $\mathcal{FS}_n := \mathcal{FS}(F_n, \emptyset)$, etc.

4.7.1 Reduced Outer space and Jewel space

Recall that an edge e of a graph G is called *separating* if removing it from G results in a disconnected graph. The subspace of $\mathcal{CV}_n$ consisting of all marked graphs that do not contain separating edges is called *reduced Outer space*, denoted $\mathcal{CV}_n^r$. It is an equivariant deformation retract of $\mathcal{CV}_n$. Similarly to the unreduced case, there is a poset K such that $\mathcal{CV}_n^r$ retracts to $\|K\|$. It is the subposet of $L = L(F_n, \emptyset)$ consisting of all marked graphs having no separating edges and is called the *spine (of reduced Outer space)*. The barycentric subdivision of the simplicial closure of reduced Outer space is given by the order complex of the poset $\mathcal{FS}_n^r$ consisting of all those free splittings $S \in \mathcal{FS}_n$ such that the quotient S/F_n does not have any separating edges. Just as in the unreduced case, we have

$$\partial_s \mathcal{CV}_n^r \cong \|\mathcal{FS}_n^r \setminus K\| .$$

In [BSV18], Bux, Smilie and Vogtmann introduced an equivariant deformation retract of $\mathcal{CV}_n^r$ called *Jewel space*, denoted by $\mathcal{J}_n$. They showed that $\mathcal{J}_n$ is homeomorphic to the bordification $b\mathcal{CV}_n^r$ of Outer space defined by Bestvina and Feighn in [BF00] and asked what the homotopy type of its boundary $\partial\mathcal{J}_n$ is. In so far unpublished work, Vogtmann shows the following:

Complex	Description via free splittings	Best-known degree of connectivity	Connectivity optimal?
$\mathcal{F}_n$	$b\mathcal{FS}^1_n = \mathcal{FS}^1_n$	$n-3$	Yes
$\mathcal{FF}_n$	$b\mathcal{FS}_n = \mathcal{FS}_n \setminus L$	$n-2$	?
$\partial\mathcal{J}_n \cong b\mathcal{CV}^r_n$	$\mathcal{FS}^r_n \setminus K$	$n-3$	?

Table 4.1: Different boundary complexes of Outer space. The results in the first two rows are established in the present work, the third row is due to Vogtmann.

Theorem 4.49 (Vogtmann)**.** *$\mathcal{FS}^r_n \setminus K$ and $\partial\mathcal{J}_n$ are homotopy equivalent. Moreover, for $n \geq 3$, they are $(n-3)$-connected.*

The techniques developed in this chapter can be used to give an alternate proof of the $(n-3)$-connectivity of $\mathcal{FS}^r_n \setminus K$; however, there is no obvious analogue of Lemma 4.44 which would yield $(n-2)$-connectivity for this poset. More details about this can be found in [BG].

4.7.2 Three boundary structures

As we saw above, the free factor complex $\mathcal{F}_n$, the complex of free factor systems $\mathcal{FF}_n$ and the boundary of Jewel space $\partial\mathcal{J}_n$ all form boundary structures of Outer space. All of them can, up to homotopy equivalence, be described as complexes of free splittings and possess connectivity in the order of n. We summarise the presently known results about these complexes in Table 4.1.

The simplicial boundaries of $\mathcal{CV}_2$ and $\mathcal{CV}^r_2$

The difference in the degree of connectivity between the reduced and the unreduced setting might be surprising at first glance, but in fact it can easily be seen when one considers the case where $n = 2$.

Here, reduced Outer space $\mathcal{CV}^r_n$ can be identified with the tesselation of the hyperbolic plane by the Farey graph (an excellent picture of this tesselation can be found in [Vog08]). The triangles of this tessellation correspond to the three-edge "theta graph". Each side of such a triangle is given by graphs that are combinatorially roses with two petals and obtained by collapsing one of the edges of the theta graph; as the rose is a graph of rank two, these edges are contained in the interior of $\mathcal{CV}^r_2$. In contrast to that, the vertices of the triangles correspond to loops obtained by collapsing two edges of the theta graph and hence are points sitting at infinity. Hence, the simplicial boundary of $\mathcal{CV}^r_2$ is homeomorphic to $\mathbb{Q}$, a countable join of 0-spheres.

Starting from reduced Outer space, unreduced $\mathcal{CV}_2$ is obtained by adding "fins" above each edge of the Farey graph. These fins are triangles corresponding to the "dumbbell graph" which consists of two loops connected by a separating edge. Collapsing this separating edge, one obtains the side of the triangle that corresponds to the rose. On the other hand, collapsing one of the two loops of the dumbbell yields a graph of rank one, forcing the other two sides of the triangle to sit at infinity. Inside the simplicial boundary $\partial_s\mathcal{CV}_2$, the concatenation of these

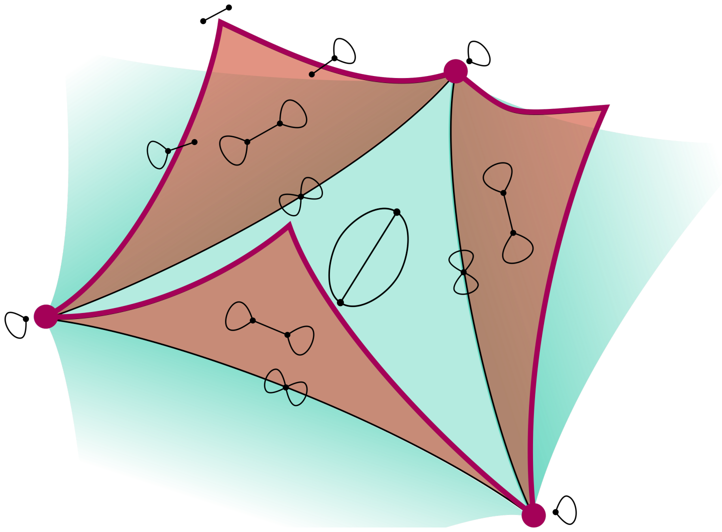

Figure 4.11: A part of $\mathcal{CV}_2$. The turquoise bottom part is reduced Outer space, together with the red fins on top it forms unreduced Outer space. The faces at infinity are coloured in magenta, where the three round vertices are the only points contained in the reduced boundary $\partial_s \mathcal{CV}_n^r$.

sides now connects two vertices of the adjacent theta graph triangles as depicted in Fig. 4.11. It follows that $\partial_s \mathcal{CV}_2$ is isomorphic to the barycentric subdivision of the Farey graph which is in turn homotopy equivalent to a countable wedge of circles.

Questions that remain open

This argument shows that for $n = 2$, the lower bounds for the degree of connectivity of the simplicial boundaries $\partial_s \mathcal{CV}_n^r \simeq \partial \mathcal{J}_n$ and $\partial_s \mathcal{CV}_n \simeq \mathcal{FF}_n$ from Table 4.1 are optimal and furthermore, the homology of these complexes is concentrated in dimension $n-2$ and $n-1$, respectively. For higher rank, however, this is not clear at all as $\partial \mathcal{J}_n$ and $\mathcal{FF}_n$ have dimension $2n-3$. The following question hence remains open:

Question 4.50. What are the homotopy types of $\mathcal{FF}_n \simeq b\mathcal{FS}_n$ and $\partial \mathcal{J}_n \simeq \mathcal{FS}_n^r \setminus K$?

There are two weaker versions of this question which would be interesting to consider. The first one is: Are $\mathcal{FF}_n$ and $\partial \mathcal{J}_n$ homotopy equivalent to simplicial complexes of lower dimension? This is the case for the curve complex $\mathcal{C}(S)$ associated to a surface S, which was defined by Harvey [Har81] as a mapping class group analogue of the Tits building. It was shown to be homotopy equivalent to a wedge of spheres by Harer [Har86] and Ivanov [Iva87]. If S is orientable, closed and has genus g, the curve complex $\mathcal{C}(S)$ is $(3g-4)$-dimensional but nevertheless it is $(2g-2)$-spherical. Unfortunately, attempts to adapt the dimension-reduction argument of [Har86, Section 3] to the setting of $\mathcal{FF}_n$ have so far remained unsuccessful.

The second weakening of Question 4.50 is: Do $\mathcal{FF}_n$ and $\partial\mathcal{J}_n$ possess non-trivial homology in dimension $n-1$ and $n-2$, respectively? In the case of $\partial\mathcal{J}_n$, there are obvious $(n-2)$-spheres one might expect to be non-trivial elements of $\pi_{n-2}(\partial\mathcal{J}_n)$. Namely whenever one has an open $(n-1)$-simplex in $\mathcal{CV}_n^r$ corresponding to a rose with n petals, all of its faces are contained in the simplicial boundary $\partial_s\mathcal{CV}_n^r$. We suspect that the spheres formed by these faces are not contractible inside the boundary but right now we do not see how this could be shown.

As $\mathcal{F}_n$ has non-trivial homology in degree $n-2$ (see 4.40), it cannot be homotopy equivalent to $\mathcal{FF}_n$. However, the following is unclear:

Question 4.51. Is $\partial\mathcal{J}_n$ homotopy equivalent to $\mathcal{F}_n$ or $\mathcal{FF}_n$?

4.7.3 Dualising module of Out(F_n)

The main reason for the interest in these boundary structures of Outer space and their homology is as follows. A group G is called a *duality group* if there is a G-module D, called the *dualising module*, such that for any G-module M and any i, one has

$$H^i(G, M) \cong H_{d-i}(G, D \otimes M).$$

As was shown by Borel and Serre [BS73], arithmetic groups are *virtual duality groups*, i.e. they have finite index subgroups which are duality groups. In this case, the dualising module is given by the top-dimensional homology of the associated rational Tits building (also called the *Steinberg module*). In order to show this, Borel–Serre constructed a bordification of the associated symmetric space and showed that its boundary is given by the rational Tits building. The result then follows from the Solomon–Tits Theorem. This relationship has been successfully extended to the mapping class group of surfaces by Harvey [Har81], Harer [Har86] and Ivanov [Iva87]: The mapping class group MCG(S) is a virtual duality group and the dualising module is given by the top-dimensional reduced homology of the curve complex $\mathcal{C}(S)$. This is shown by interpreting the curve complex as a boundary structure of Teichmüller space and then showing that it is k-spherical for some k depending on S.

Using the bordification $b\mathcal{CV}_n^r$ of Outer space mentioned above, Bestvina–Feighn [BF00] showed that Out(F_n) is a virtual duality group as well. However, they did not obtain an explicit description of the dualising module. The aim of describing a (spherical) complex whose homology realises this module is an important motivation for studying boundary structures like the complexes $\mathcal{F}_n$, $\mathcal{FF}_n$ and $\partial\mathcal{J}_n$.

Chapter 5

A Cohen–Macaulay complex for Out(RAAGs)

This chapter is devoted to automorphisms of right-angled Artin groups and the associated complex $\mathcal{CC}$ mentioned in the introduction. It starts with background material on (relative) automorphism groups of RAAGs in Section 5.1. In particular, this section contains the necessary preliminaries for the decomposition sequence of $\mathrm{Out}^0(A_\Gamma)$ due to Day and Wade. Section 5.2 is in some sense the core of this text: We define (maximal) parabolic subgroups and the complex $\mathcal{CC}$ as a coset complex with respect to these parabolic subgroups. We then combine the results of Section 5.1 and the previous chapters in order to determine the homotopy type of $\mathcal{CC}$ and thus prove Theorem A. Section 5.3 summarises the extent to which our considerations refine the inductive procedure of Day–Wade and provides examples of the constructions for specific graphs Γ. In Section 5.4, we prove Cohen–Macaulayness of $\mathcal{CC}$, define parabolic subgroups of lower rank and show that these form higher generating families. We then explain how the dimension of $\mathcal{CC}$ is related to the rank of a Coxeter subgroup of O (see Corollary 5.36). The chapter closes with comments about the limitations of the construction and open questions in Section 5.5.

The material of this chapter is taken from the article "Between buildings and free factor complexes: A Cohen–Macaulay complex for Out(RAAGs)" [Brüa].

5.1 Relative automorphism groups of RAAGs

In this section, we examine relative automorphism groups of right-angled Artin groups. These groups were studied in detail by Day and Wade [DW] and many of the results here are either taken from their work or build on their ideas. For an overview of other literature on relative automorphism groups, see [DW, Section 6.1]. In this thesis, such relative automorphism groups occur in two ways: On the one hand, they arise as the images and kernels of restriction and projection homomorphisms, which in turn play an important role for the inductive procedure of Day–Wade; on the other hand, the parabolic subgroups we will define in Section 5.2 are themselves relative automorphism groups of RAAGs. For the purpose of this text, the present section mostly serves as a toolbox for the inductive proof of Theorem A in Section 5.2. Its main goals are to collect all the results from [DW] that we will need afterwards, to adapt them to our purposes and, maybe most importantly, to set up the language that will be used later on.

Standing assumption From now on, all graphs that we consider are finite and simplicial, i.e. without loops or multiple edges. To emphasise this difference from Chapter 4, they will be denoted by Greek letters.

5.1.1 RAAGs and their automorphism groups

Subgraphs, links and stars In contrast to Chapter 4, if we talk about a subgraph Δ of a graph Γ, we from now on always mean a *full* subgraph, i.e. if two vertices $v, w \in V(\Delta)$ are connected by an edge in Γ, they are connected in Δ as well. A full subgraph of Γ can also be seen as a subset of the vertex set $V(\Gamma)$; we will often take this point of view, identify Δ with $V(\Delta)$ and write $\Delta \subseteq \Gamma$, or $\Delta \subset \Gamma$ if we want to emphasise that Δ is a proper subgraph of Γ.

Given a vertex $v \in V(\Gamma)$, the *link* $\mathrm{lk}(v)$ of v is the subgraph of Γ consisting of all the vertices that are adjacent to v. The *star* $\mathrm{st}(v)$ of v is the subgraph of Γ with vertex set $\{v\} \cup \mathrm{lk}(v)$. We also write $\mathrm{lk}_\Gamma(v)$ or $\mathrm{st}_\Gamma(v)$ if we want to distinguish between links and stars in different graphs.

RAAGs and special subgroups Given a graph Γ, the associated *right-angled Artin group*—abbreviated as RAAG—A_Γ is defined to be the group generated by the set $V(\Gamma)$ subject to the relations $[v, w] = 1$ for all $v, w \in V(\Gamma)$ which are adjacent to each other.

Given any subgraph $\Delta \subseteq \Gamma$, the inclusion $V(\Delta) \to V(\Gamma)$ induces an injective homomorphism $A_\Delta \hookrightarrow A_\Gamma$. This allows us to interpret A_Δ as a subgroup of A_Γ. Subgroups of this type are called *special subgroups* of A_Γ.

The standard ordering and its equivalence classes There is a so-called *standard ordering* on the vertex set $V(\Gamma)$ that is the partial pre-order given by $v \leq w$ if and only if $\mathrm{lk}(v) \subseteq \mathrm{st}(w)$. The induced equivalence relation of this partial pre-order is denoted by $\sim$, i.e. $v \sim w$ if and only if $v \leq w$ and $w \leq v$. The equivalence class of v is denoted by $[v]$. The standard ordering induces a partial order on the equivalence classes where we say $[v] \leq [w]$ if $v \leq w$ (this does not depend on the choice of representatives). If two equivalent vertices $v \sim w$ are adjacent, it follows that the vertices from their equivalence class $[v]$ form a complete subgraph of Γ. In this case, the special subgroup $A_{[v]}$ is isomorphic to $\mathbb{Z}^{|[v]|}$ and we call $[v]$ an *abelian* equivalence class. If on the other hand $[v]$ does not contain any pair of adjacent vertices, it can be seen as discrete subgraph of Γ. In this case, we call $[v]$ a *free* equivalence class because $A_{[v]}$ is isomorphic to the free group $F_{|[v]|}$. For more details about this ordering and the equivalence relation, see [CV09].

Example 5.1. Let Γ be the tripod consisting of a central vertex 4 and adjacent vertices $1, 2$ and 3 as depicted in Fig. 5.1. The corresponding RAAG A_Γ is isomorphic to $F_3 \times \mathbb{Z} = \langle 1, 2, 3\rangle \times \langle 4\rangle$. We have $\mathrm{st}(4) = \Gamma$, which implies that $1 \leq 4$, $2 \leq 4$ and $3 \leq 4$. The vertices $1, 2$ and 3 form a (free) equivalence class of size three, i.e. the partial order on the set of equivalence is determined by $[1] = \{1, 2, 3\} \leq [4] = \{4\}$.

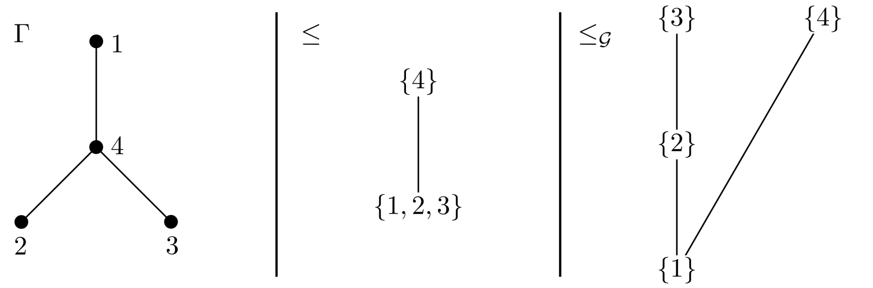

Figure 5.1: The graph Γ, the Hasse diagram of the partial order ≤ on the standard equivalence classes of $V(\Gamma)$ and the Hasse diagram of the partial order $\leq_{\mathcal{G}}$ where $\mathcal{G} = \left\{A_{\{3\}}, A_{\{2,3\}}\right\}$.

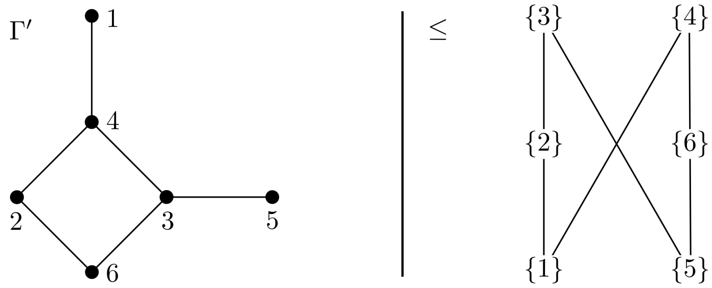

Figure 5.2: The graph Γ′ and the Hasse diagram of the standard ordering.

Fig. 5.2 depicts the standard ordering of a graph Γ′ obtained by attaching two additional vertices 5 and 6 to Γ. In contrast to A_Γ, the RAAG $A_{\Gamma'}$ cannot be obtained from infinite cyclic groups by performing a sequence of direct and free products. The special subgroup $A_{\{1,2,3,4\}} \leq A_{\Gamma'}$ is isomorphic to A_Γ.

Automorphisms of RAAGs Let $\mathrm{Aut}(A_\Gamma)$ and $\mathrm{Out}(A_\Gamma)$ denote the automorphism group and the group of outer automorphisms of A_Γ, respectively. By the work of Servatius [Ser89] and Laurence [Lau95], the group $\mathrm{Aut}(A_\Gamma)$ is generated by the following automorphisms:

- *Graph automorphisms.* Any automorphism of the graph Γ gives rise to an automorphism of A_Γ by permuting the generators of the RAAG.
- *Inversions.* Let $v \in V(\Gamma)$. The map sending v to v^{-1} and fixing all the other generators induces an automorphism of A_Γ. It is called an *inversion* and denoted by ι_v.
- *Transvections.* Let $v, w \in V(\Gamma)$ with $v \leq w$ and $v \neq w$. The *transvection* ρ_v^w is the automorphism of A_Γ induced by sending v to vw and fixing all the other generators. We call w the *acting letter* of ρ_v^w.
- *Partial conjugations.* Let $v \in V(\Gamma)$ and K a union of connected components of $\Gamma \setminus \mathrm{st}(v)$. The map sending every vertex w of K to vwv^{-1} and

fixing the remaining generators induces an automorphism π^v_K of A_Γ and is called a *partial conjugation*. We call v the *acting letter* of π^v_K.

We use the same notation to denote the images of these automorphisms in $\mathrm{Out}(A_\Gamma)$ and call these (outer) automorphisms the *Laurence generators* of $\mathrm{Aut}(A_\Gamma)$ or $\mathrm{Out}(A_\Gamma)$, respectively.

The subgroup of $\mathrm{Out}(A_\Gamma)$ generated by all inversions, transvections and partial conjugations is denoted by $\mathrm{Out}^0(A_\Gamma)$. It is called the *pure outer automorphism group* of A_Γ and has finite index in $\mathrm{Out}(A_\Gamma)$. This group was first defined by Charney, Crisp and Vogtmann in [CCV07] and has since become popular as it avoids certain technical difficulties coming from automorphisms of the graph Γ. If A_Γ is equal to $\mathbb{Z}^n$ or F_n, we have $\mathrm{Out}^0(A_\Gamma) = \mathrm{Out}(A_\Gamma)$.

Example 5.1 (continued). The group $\mathrm{Aut}(A_\Gamma)$ is generated by the following elements: There are four inversions $\iota_1, \iota_2, \iota_3, \iota_4$. The transvections correspond to the relations of the standard ordering $\leq$, so for all $i \neq j \in \{1, 2, 3\}$, there are transvections ρ^4_i and ρ^j_i. As $\Gamma \setminus \mathrm{st}(1) = \{2, 3\}$, we obtain partial conjugations $\pi^1_{\{2\}}$ and $\pi^1_{\{3\}}$. In fact, their product is equal to the inner automorphism conjugating every element of A_Γ by 1, so we have $\pi^1_{\{2\}} = (\pi^1_{\{3\}})^{-1}$ in $\mathrm{Out}(A_\Gamma)$. Analogously, there are the partial conjugations $\pi^2_{\{1\}}, \pi^2_{\{3\}}$ and $\pi^3_{\{1\}}, \pi^3_{\{2\}}$. The group of graph automorphisms of Γ is in an obvious way isomorphic to the group of permutations of the set $\{1, 2, 3\}$. Every such permutation can be written as a product of inversions and transvections. To see this, one verifies that e.g. the transposition exchanging 1 and 2 can be written as the product $\rho^2_1\iota_1\rho^1_2\iota_2\rho^2_1\iota_1$. This implies that $\mathrm{Out}^0(A_\Gamma) = \mathrm{Out}(A_\Gamma)$.

The group $\mathrm{Out}^0(A_{\Gamma'})$ is generated by inversions and the following list of transvections and partial conjugations:

$$\rho^2_1,\ \rho^3_1,\ \rho^4_1,\ \rho^3_2,\ \rho^3_5,\ \rho^4_5,\ \rho^6_5,\ \rho^4_6,$$

$$\pi^2_{\{3,5\}} = \left(\pi^2_{\{1\}}\right)^{-1},\ \pi^3_{\{2\}} = \left(\pi^3_{\{1\}}\right)^{-1},\ \pi^4_{\{5\}} = \left(\pi^4_{\{6\}}\right)^{-1},\ \pi^6_{\{1,4\}} = \left(\pi^6_{\{5\}}\right)^{-1}.$$

The graph Γ' has exactly one non-trivial automorphism, namely the one that swaps the tuples $(1, 2, 4)$ and $(5, 6, 3)$. As we will see in Lemma 5.35, this automorphism cannot be written as a product of inversions, transvections and partial conjugations, so $\mathrm{Out}^0(A_{\Gamma'})$ is a proper subgroup of $\mathrm{Out}(A_\Gamma)$.

5.1.2 Generators of relative automorphism groups

Recall that for a group G and families of subgroups $\mathcal{G}$ and $\mathcal{H}$, the relative automorphism group $\mathrm{Out}(G; \mathcal{G}, \mathcal{H}^t)$ is defined as the subgroup of $\mathrm{Out}(G)$ consisting of all elements that stabilise each $H \in \mathcal{G}$ and that act trivially on each $H \in \mathcal{H}$ (see Section 4.1.2).

Given a pair $(\mathcal{G}, \mathcal{H})$ of families of special subgroups of A_Γ, we define

$$\mathrm{Out}^0(A_\Gamma;\, \mathcal{G}, \mathcal{H}^t) := \mathrm{Out}(A_\Gamma; \mathcal{G}, \mathcal{H}^t) \cap \mathrm{Out}^0(A_\Gamma).$$

Building on the work of Laurence, Day–Wade show:

Theorem 5.2 ([DW, Theorem D])**.** *If $\mathcal{G}$ and $\mathcal{H}$ are families of special subgroups of A_Γ, the group $\mathrm{Out}^0(A_\Gamma;\, \mathcal{G}, \mathcal{H}^t)$ is generated by the set of all inversions, transvections and partial conjugations of $\mathrm{Out}(A_\Gamma)$ it contains.*

In order to prove this, Day and Wade give a description of the Laurence generators contained in such a relative automorphism group. To state it, we first need to set up the terminology developed in their article.

$\mathcal{G}$-components and $\mathcal{G}$-ordering Let $\mathcal{G}$ be a family of proper special subgroups of A_Γ. We say that $v, w \in V(A_\Gamma)$ are *$\mathcal{G}$-adjacent* if v is adjacent to w or if there is some $A_\Delta \in \mathcal{G}$ such that $v, w \in \Delta$. A subgraph $\Delta \subseteq \Gamma$ is *$\mathcal{G}$-connected* if for all $v, w \in \Delta$, there is a sequence of vertices in Δ which starts with v, ends with w and such that each of its vertices is $\mathcal{G}$-adjacent to the next one. A maximal $\mathcal{G}$-connected subgraph of Γ is called a *$\mathcal{G}$-component*.

We define a partial pre-order $\leq_\mathcal{G}$ on $V(\Gamma)$ by saying that $v \leq_\mathcal{G} w$ if and only if $v \leq w$ and for all $A_\Delta \in \mathcal{G}$, if $v \in \Delta$, one has $w \in \Delta$. This is called the *$\mathcal{G}$-ordering* of $V(\Gamma)$. The equivalence relation of this pre-order is denoted by $\sim_\mathcal{G}$, its equivalence classes by $[\cdot]_\mathcal{G}$.

Note that in the case where $\mathcal{G} = \emptyset$, a $\mathcal{G}$-component of Γ is just a connected component and the $\mathcal{G}$-ordering is the standard ordering on $V(\Gamma)$.

Example 5.1 (continued). Let $\mathcal{G} = \{A_{\{3\}}, A_{\{2,3\}}\}$. The vertices 2 and 3 which are not connected by an edge are $\mathcal{G}$-adjacent. Because Γ is connected, it is also $\mathcal{G}$-connected. The $\mathcal{G}$-ordering $\leq_\mathcal{G}$ is given by $1 \leq 4$ and $1 \leq 2 \leq 3$ (see Fig. 5.1). There is no $\mathcal{G}$-equivalence class of size bigger than one.

For $v \in V(\Gamma)$, let $\mathcal{G}^v := \{A_\Delta \in \mathcal{G} \mid v \notin \Delta\}$. It is easy to see that every $\mathcal{G}^v$-component of $\Gamma \setminus \mathrm{st}(v)$ is a union of connected components of $\Gamma \setminus \mathrm{st}(v)$. Suppose that $\mathcal{H}$ is a family of special subgroups of A_Γ. The *power set of $\mathcal{H}$*, denoted by $P(\mathcal{H})$, is defined as the set of all special subgroups $A_\Delta \leq A_\Gamma$ which are contained in some element of $\mathcal{H}$.

Lemma 5.3 ([DW, Proposition 3.9]). *Let $\mathcal{G}$ and $\mathcal{H}$ be families of special subgroups of A_Γ such that $\mathcal{G}$ contains $P(\mathcal{H})$. Let $v, w \in V(\Gamma)$ and let K be a union of connected components of $\Gamma \setminus \mathrm{st}(v)$. Then:*

- *The inversion ι_v is contained in $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ if and only if there is no subgroup $A_\Delta \in \mathcal{H}$ with $v \in \Delta$.*
- *The transvection ρ_v^w is contained in $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ if and only if $v \leq_\mathcal{G} w$.*
- *The partial conjugation π_K^v is contained in $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ if and only if K is a union of $\mathcal{G}^v$-components of $\Gamma \setminus \mathrm{st}(v)$.*

Note that it imposes no great restriction to assume that the power set of $\mathcal{H}$ is contained in $\mathcal{G}$ because for any families $\mathcal{G}$ and $\mathcal{H}$ of special subgroups, one has

$$\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t) = \mathrm{Out}^0(A_\Gamma; \mathcal{G} \cup P(\mathcal{H}), \mathcal{H}^t),$$

see [DW, Lemma 3.8].

Example 5.1 (continued). Let $\mathcal{H} = \emptyset$. Obviously, $\mathcal{G}$ contains $P(\mathcal{H}) = \emptyset$. Among the generators of $\mathrm{Out}^0(A_\Gamma)$ determined above, the following are contained in $\mathrm{Out}^0(A_\Gamma; \mathcal{G})$: All inversions are still there because $\mathcal{H} = \emptyset$. The transvection ρ_2^1, which is an element of $\mathrm{Out}^0(A_\Gamma)$, is not contained in $\mathrm{Out}^0(A_\Gamma; \mathcal{G})$ as it does not stabilise $A_{\{2,3\}}$. This is reflected in the $\mathcal{G}$-ordering $\leq_\mathcal{G}$. According to it, the transvections of $\mathrm{Out}^0(A_\Gamma; \mathcal{G})$ are ρ_1^2, ρ_1^3, ρ_1^4 and ρ_2^3 (see the Hasse diagram in Fig. 5.1). The partial conjugation $\pi_{\{2\}}^1 = (\pi_{\{3\}}^1)^{-1} \in \mathrm{Out}^0(A_\Gamma)$ does not form an element of $\mathrm{Out}^0(A_\Gamma; \mathcal{G})$: We have $\mathcal{G}^1 = \mathcal{G} = \{A_{\{3\}}, A_{\{2,3\}}\}$. Hence, the vertices 2 and 3 are $\mathcal{G}^1$-adjacent in $\Gamma \setminus \mathrm{st}(1) = \{2, 3\}$. All the remaining partial conjugations of $\mathrm{Out}^0(A_\Gamma)$ are contained in $\mathrm{Out}^0(A_\Gamma; \mathcal{G})$.

The next result is the key ingredient for the proof of Lemma 5.3 in [DW]. We include it here because it provides a convenient description of the parabolic subgroups that we will study later on.

Lemma 5.4 ([DW, Lemma 2.2]). *Let A_Δ be a special subgroup of A_Γ. Let $v, w \in V(\Gamma)$ and let K be a union of connected components of $\Gamma \setminus \mathrm{st}(v)$. Then:*

- *The inversion ι_v acts trivially on A_Δ if and only if $v \notin \Delta$; it always stabilises A_Δ.*
- *The transvection ρ_v^w acts trivially on A_Δ if and only if $v \notin \Delta$; it stabilises A_Δ if it acts trivially on it or $w \in \Delta$.*
- *The partial conjugation π_K^v acts trivially on A_Δ if and only if*

$$K \cap \Delta = \emptyset \text{ or } \Delta \setminus \mathrm{st}(x) \subseteq K;$$

it stabilises A_Δ if it acts trivially on it or $w \in \Delta$.

Combining the last two lemmas yields an algorithm which computes all special subgroups of A_Γ that are stabilised by $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$: First determine the generators of this group using Lemma 5.3, then apply Lemma 5.4 to check which special subgroups are stabilised by all of them. This procedure is summarised in [DW, Proposition 3.11] as follows. The subgroup A_Δ is stabilised by $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ if and only if Δ is upwards-closed with respect to $\leq_\mathcal{G}$ and for all $v \in V(\Gamma) \setminus \Delta$, the graph Δ intersects at most one $\mathcal{G}^v$-component of $\Gamma \setminus \mathrm{st}(v)$.

Example 5.1 (continued). The following is a list of all subgraphs $\Delta \subset \Gamma'$ such that A_Δ is stabilised by all of $\mathrm{Out}^0(A_{\Gamma'})$:

$$\begin{gathered}\{3\}, \{4\}, \{2,3\}, \{3,4\}, \{4,6\}, \{2,3,4\}, \{3,4,6\}, \{1,2,3,4\}, \\ \{2,3,4,6\}, \{3,4,5,6\}, \{1,2,3,4,6\}, \{2,3,4,5,6\}.\end{gathered} \tag{5.1}$$

5.1.3 Restriction and projection homomorphisms

Let O be a subgroup of $\mathrm{Out}(A_\Gamma)$. If the special subgroup $A_\Delta \leq A_\Gamma$ is stabilised by O, there is a *restriction homomorphism*

$$R_\Delta \colon O \to \mathrm{Out}(A_\Delta),$$

where $R_\Delta(\Phi)$ is the outer automorphism given by taking a representative $\phi \in \Phi$ that sends A_Δ to itself and restricting it to A_Δ. If the normal subgroup $\langle\langle A_\Delta \rangle\rangle$ generated by A_Δ is stabilised by O, there is a *projection homomorphism*

$$P_{\Gamma \setminus \Delta} \colon O \to \mathrm{Out}(A_{\Gamma \setminus \Delta}),$$

which is induced by the quotient map

$$A_\Gamma \to A_\Gamma / \langle\langle A_\Delta \rangle\rangle \cong A_{\Gamma \setminus \Delta}.$$

Restriction and projection maps were first defined by Charney, Crips and Vogtmann in [CCV07] and have since become an important tool for studying automorphism groups of RAAGs via inductive arguments. These arguments rely on the fact that if Γ is neither a complete nor a discrete graph, there are always proper special subgroups that are stabilised by $\mathrm{Out}^0(A_\Gamma)$. This implies that typically, there are many possible restriction maps one can apply to $\mathrm{Out}^0(A_\Gamma)$.

Generators of image and kernel

Day–Wade obtained a complete description of the image and kernel of restriction homomorphisms. Again let $\mathcal{G}$ and $\mathcal{H}$ be families of special subgroups of A_Γ. We say that $\mathcal{G}$ is *saturated with respect to* $(\mathcal{G}, \mathcal{H})$, if it contains every proper special subgroup stabilised by $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$. As we saw above, there is an algorithm which computes the family $\mathcal{G}'$ of all special subgroups stabilised by $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$. This family is saturated with respect to $(\mathcal{G}', \mathcal{H})$ and $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t) = \mathrm{Out}^0(A_\Gamma; \mathcal{G}', \mathcal{H}^t)$. We also saw that even if we start with $\mathcal{G} = \mathcal{H} = \emptyset$, the family $\mathcal{G}'$ we obtain is a priori non-trivial, see Eq. (5.1).

Given a special subgroup $A_\Delta \leq A_\Gamma$, set

$$\mathcal{G}_\Delta := \{A_{\Delta\cap\Theta} \mid A_\Theta \in \mathcal{G} \text{ and } \Delta \cap \Theta \neq \Delta\}.$$

We define $\mathcal{H}_\Delta$ analogously.

Theorem 5.5 ([DW, Theorem E]). *Let $\mathcal{G}$ be saturated with respect to $(\mathcal{G}, \mathcal{H})$ and let $A_\Delta \in \mathcal{G}$. The restriction homomorphism*

$$R_\Delta \colon \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t) \to \mathrm{Out}(A_\Delta)$$

has image

$$\mathrm{im}\, R_\Delta = \mathrm{Out}^0(A_\Delta; \mathcal{G}_\Delta, \mathcal{H}^t_\Delta)$$

and kernel

$$\ker R_\Delta = \mathrm{Out}^0(A_\Gamma; \mathcal{G}, (\mathcal{H} \cup \{A_\Delta\})^t).$$

It is not hard to see that both restriction and projection maps send each Laurence generator either to the identity or to a Laurence generator of the same type. For the proof of Theorem 5.5, Day–Wade show that for restriction maps, a converse of this is true as well: Every Laurence generator in $\mathrm{im}\, R_\Delta$ is given as the restriction of a Laurence generator of $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$.

Example 5.1 (continued). The special subgroup $A_{\{1,2,3,4\}} \leq A_{\Gamma'}$ is stabilised by $\mathrm{Out}^0(A_{\Gamma'})$, see Eq. (5.1). It is isomorphic to A_Γ. Thus, there is a restriction map

$$R_\Gamma \colon \mathrm{Out}^0(A_{\Gamma'}) \to \mathrm{Out}(A_\Gamma).$$

Let $\mathcal{G}'$ be the family of special subgroups of $A_{\Gamma'}$ corresponding to the list of subgraphs in Eq. (5.1). It is saturated with respect to $(\mathcal{G}', \emptyset)$. The image of R_Γ hence is given by $\mathrm{Out}^0(A_\Gamma; \mathcal{G}'_\Gamma)$, where

$$\mathcal{G}'_\Gamma = \left\{A_{\{3\}}, A_{\{4\}}, A_{\{2,3\}}, A_{\{3,4\}}, A_{\{2,3,4\}}\right\}.$$

It is not hard to see that $\mathrm{Out}^0(A_\Gamma; \mathcal{G}'_\Gamma) = \mathrm{Out}^0(A_\Gamma; \mathcal{G})$ with $\mathcal{G} = \left\{A_{\{3\}}, A_{\{2,3\}}\right\}$ as above. This exemplifies why in general, restriction maps are not surjective: The ambient structure of the graph Γ' (in this case the vertices 5 and 6) excludes some of the automorphisms of the RAAG A_Γ.

The kernel of R_Γ is $\mathrm{Out}^0(A_{\Gamma'};\, \mathcal{G}', \{A_\Gamma\}^t)$, which is equal to $\mathrm{Out}^0(A_{\Gamma'};\, \{A_\Gamma\}^t)$. This group is generated by

$$\iota_1,\, \iota_2,\, \iota_3,\, \iota_4,\, \rho_5^3,\, \rho_5^4,\, \rho_5^6,\, \rho_6^4,$$
$$\pi^4_{\{5\}} = \left(\pi^4_{\{6\}}\right)^{-1} \text{ and } \pi^6_{\{1,4\}} = \left(\pi^6_{\{5\}}\right)^{-1}.$$

Remark 5.6. Theorem 5.5 implies that the class of relative automorphism groups of the form $\mathrm{Out}^0(A_\Gamma;\, \mathcal{G}, \mathcal{H}^t)$ is closed under passing to images and kernels of restriction maps—as we saw in Example 5.1, this is not the case for the class of outer automorphism groups of the form $\mathrm{Out}^0(A_\Gamma)$. A consequence of this is that for our purposes, working with these relative automorphism groups becomes unavoidable, even if one is only interested in $\mathrm{Out}^0(A_\Gamma)$ itself.

In general, images and kernels of projection homomorphisms are more difficult to describe than those of restriction homomorphisms. However, we will only need to consider them in a special case: The centre $Z(A_\Gamma)$ of A_Γ is generated by all vertices $z \in V(\Gamma)$ such that $\mathrm{st}(z) = \Gamma$. If $Z(A_\Gamma)$ is non-trivial, these vertices form an abelian equivalence class $Z := [z]$ and Γ can be written as a join $\Gamma = Z * \Delta$ where $\Delta = \Gamma \setminus Z$. If we have a graph of this form, the centre $Z(A_\Gamma) = A_Z$ is a normal subgroup which is stabilised by all of $\mathrm{Out}(A_\Gamma)$. Hence, there is a projection map

$$P_\Delta\colon\ \mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_\Delta).$$

The image of this projection map can be described similarly as that of a restriction map. In fact, the situation in this special case is even easier as we do not even need to assume any kind of saturation for our families of special subgroups:

Lemma 5.7. *Assume that Γ can be decomposed as a join $\Gamma = Z * \Delta$ where Z is a complete graph. Let $\mathcal{G}$ and $\mathcal{H}$ be any two families of special subgroups of A_Γ and let $A_Z \in \mathcal{G}$. The projection homomorphism*

$$P_\Delta\colon\ \mathrm{Out}^0(A_\Gamma;\, \mathcal{G}, \mathcal{H}^t) \to \mathrm{Out}(A_\Delta)$$

has image equal to

$$\mathrm{im}\, P_\Delta = \mathrm{Out}^0(A_\Delta;\, \mathcal{G}_\Delta, \mathcal{H}^t_\Delta).$$

Proof. The inclusion "$\subseteq$" follows immediately from the definitions.

For the other inclusion, we start by defining $\tilde{\mathcal{G}} := \mathcal{G} \cup P(\mathcal{H})$ as the union of $\mathcal{G}$ and the power set of $\mathcal{H}$. As observed above, we have

$$O := \mathrm{Out}^0(A_\Gamma;\, \mathcal{G}, \mathcal{H}^t) = \mathrm{Out}^0(A_\Gamma;\, \tilde{\mathcal{G}}, \mathcal{H}^t).$$

Furthermore, $\tilde{\mathcal{G}}_\Delta = \mathcal{G}_\Delta \cup P(\mathcal{H}_\Delta)$, so we also have

$$O_\Delta := \mathrm{Out}^0(A_\Delta;\, \mathcal{G}_\Delta, \mathcal{H}^t_\Delta) = \mathrm{Out}^0(A_\Delta;\, \tilde{\mathcal{G}}_\Delta, \mathcal{H}^t_\Delta). \tag{5.2}$$

By Theorem 5.2, we know that O_Δ is generated by the inversions, transvections and partial conjugations it contains. Hence, it suffices to find a preimage under P_Δ for each of those generators. Combining Eq. (5.2) with Lemma 5.3, we have a complete description of the generators in O_Δ. In what follows, we will use this description to construct the preimages for the generators one at a time.

The inversion ι_v is contained in O_Δ if and only if $v \in \Delta$ and there is no $A_{\Delta'} \in \mathcal{H}_\Delta$ such that $v \in \Delta'$. This implies that there is no $A_{\Delta'} \in \mathcal{H}$ with $v \in \Delta'$, so the inversion at v is an element of O. It gets mapped to ι_v under P_Δ.

If one has a transvection $\rho_v^w \in O_\Delta$, Lemma 5.3 implies that $v \leq_{\tilde{\mathcal{G}}_\Delta} w$, i.e. $\mathrm{lk}_\Delta(v) \subseteq \mathrm{st}_\Delta(w)$ and for each $A_{\Delta'} \in \tilde{\mathcal{G}}_\Delta$, one has that $v \in \Delta'$ implies $w \in \Delta'$. We want to show that $v \leq_{\tilde{\mathcal{G}}} w$. As Γ is a join $Z * \Delta$, the link and star of v and w in Γ are of the form

$$\mathrm{lk}_\Gamma(v) = \mathrm{lk}_\Delta(v) \cup Z, \qquad\qquad \mathrm{st}_\Gamma(w) = \mathrm{st}_\Delta(w) \cup Z.$$

In particular, $\mathrm{lk}_\Gamma(v) \subseteq \mathrm{st}_\Gamma(w)$. The vertex v cannot be contained in any Δ' with $A_{\Delta'} \in P(\mathcal{H})$ as this would imply $A_{\{v\}} \in P(\mathcal{H}_\Delta) \subseteq \tilde{\mathcal{G}}_\Delta$, contradicting the assumption that $v \leq_{\tilde{\mathcal{G}}_\Delta} w$. Now take $A_{\Delta'} \in \mathcal{G}$ such that $v \in \Delta'$. If $\Delta \subseteq \Delta'$, both v and w are contained in Δ'. If on the other hand $\Delta \cap \Delta'$ is a proper subset of Δ, one has $A_{\Delta \cap \Delta'} \in \mathcal{G}_\Delta \subseteq \tilde{\mathcal{G}}_\Delta$, so $w \in \Delta'$. It follows that $v \leq_{\tilde{\mathcal{G}}} w$, so the transvection multiplying v by w defines an element of O and is a preimage of ρ_v^w.

Again using Lemma 5.3, the partial conjugation π_K^v is contained in O_Δ if and only if $v \in \Delta$ and K is a union of $\tilde{\mathcal{G}}_\Delta^v$-components of $\Delta \setminus \mathrm{st}(v)$. We claim that every $\tilde{\mathcal{G}}_\Delta^v$-component C of $\Delta \setminus \mathrm{st}_\Delta(v)$ is also a $\tilde{\mathcal{G}}^v$-component of $\Gamma \setminus \mathrm{st}_\Gamma(v)$. To see this, first recall that each element of Z is connected to every vertex of Γ, so $\Gamma \setminus \mathrm{st}_\Gamma(v) = \Delta \setminus \mathrm{st}_\Delta(v)$. Furthermore, it follows right from the definitions that two vertices $x, y \in \Delta \setminus \mathrm{st}_\Delta(v)$ are $\tilde{\mathcal{G}}_\Delta$-adjacent in $\Delta \setminus \mathrm{st}_\Delta(v)$ if and only if they are $\tilde{\mathcal{G}}$-adjacent in $\Gamma \setminus \mathrm{st}_\Gamma(v)$. The claim follows and implies that the partial conjugation of K by v defines an element of O. It is a preimage of π_K^v. $\square$

Remark 5.8. The combinatorial criteria for studying automorphism groups of RAAGs given above allow one to determine things like generators of relative automorphism groups, image and kernel of restriction and projection homomorphisms and families of stabilised special subgroups using a computer. This was done by the author at various occasions in order to produce examples and establish conjectures.

Relative orderings in image and kernel

Standing assumptions and notation From now on and until the end of Section 5.1, let $O := \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ where $\mathcal{G}$ and $\mathcal{H}$ are families of special subgroups of A_Γ such that $\mathcal{G}$ is saturated with respect to $(\mathcal{G}, \mathcal{H})$; note that saturation implies that $P(\mathcal{H}) \subseteq \mathcal{G}$. Set $\preceq := \leq_{\mathcal{G}}$ to be the $\mathcal{G}$-ordering on $V(\Gamma)$.

Remark 5.9. We saw above that, given an arbitrary relative automorphism group, there might be several ways of "representing" this group by families of subgroups that are stabilised or acted trivially upon. This means that we might have

$$\mathrm{Out}^0(A_\Gamma;\, \mathcal{G}_1, \mathcal{H}_1^t) = \mathrm{Out}^0(A_\Gamma;\, \mathcal{G}_2, \mathcal{H}_2^t)$$

with $(\mathcal{G}_1, \mathcal{H}_1) \neq (\mathcal{G}_2, \mathcal{H}_2)$. However, if in this situation, we have both $P(\mathcal{H}_1) \subseteq \mathcal{G}_1$ and $P(\mathcal{H}_2) \subseteq \mathcal{G}_2$, the orderings $\leq_{\mathcal{G}_1}$ and $\leq_{\mathcal{G}_2}$ agree: By Lemma 5.3, for every

$v, w \in V(\Gamma)$, there is a chain of equivalences

$$\begin{aligned} v \leq_{\mathcal{G}_1} w &\Leftrightarrow \rho_v^w \in \mathrm{Out}^0(A_\Gamma; \mathcal{G}_1, \mathcal{H}_1^t) = \mathrm{Out}^0(A_\Gamma; \mathcal{G}_2, \mathcal{H}_2^t) \\ &\Leftrightarrow v \leq_{\mathcal{G}_2} w. \end{aligned}$$

In particular, the ordering $\leq_\mathcal{G}$ of $V(A_\Gamma)$ where $\mathcal{G}$ is saturated with respect to $(\mathcal{G}, \mathcal{H})$ is an invariant of the group $\mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$; it depends on the transvections contained in this group but not on any other choices.

As mentioned above, a restriction homomorphism maps every transvection that is not contained in its kernel to a transvection of the same type. The consequences for the relative ordering in the image and kernel are as follows:

Lemma 5.10. *Let $A_\Delta \in \mathcal{G}$ be a special subgroup that is stabilised by O and let R_Δ denote the corresponding restriction homomorphism. If we write*

$$\operatorname{im} R_\Delta = \mathrm{Out}^0(A_\Delta; \mathcal{G}_{\mathrm{im}}, \mathcal{H}_{\mathrm{im}}^t) \quad \textit{and} \quad \ker R_\Delta = \mathrm{Out}^0(A_\Gamma; \mathcal{G}_{\mathrm{ker}}, \mathcal{H}_{\mathrm{ker}}^t)$$

with $\mathcal{G}_{\mathrm{im}}$ and $\mathcal{G}_{\mathrm{ker}}$ saturated with respect to $(\mathcal{G}_{\mathrm{im}}, \mathcal{H}_{\mathrm{im}})$ and $(\mathcal{G}_{\mathrm{ker}}, \mathcal{H}_{\mathrm{ker}})$, respectively, the following holds true:

1. *For $v, w \in \Delta$, one has $v \leq_{\mathcal{G}_{\mathrm{im}}} w$ if and only if $v \preceq w$.*
2. *For $v \neq w \in V(\Gamma)$, one has $v \leq_{\mathcal{G}_{\mathrm{ker}}} w$ if and only if $v \in V(\Gamma) \setminus \Delta$ and $v \preceq w$.*

Proof. As $\mathcal{G}$ is saturated, we know that $\operatorname{im} R_\Delta = \mathrm{Out}^0(A_\Delta; \mathcal{G}_\Delta, \mathcal{H}_\Delta^t)$. For $v, w \in \Delta$, [DW, Proposition 4.1] shows that $v \leq_{\mathcal{G}_\Delta} w$ if and only if $v \preceq w$. Again because of the saturation of $\mathcal{G}$, we have $P(\mathcal{H}) \subseteq \mathcal{G}$. Hence, $P(\mathcal{H}_\Delta) \subseteq \mathcal{G}_\Delta$. As in Remark 5.9, it follows that $v \leq_{\mathcal{G}_\Delta} w$ if and only if $v \leq_{\mathcal{G}_{\mathrm{im}}} w$ for $\mathcal{G}_{\mathrm{im}}$ saturated with respect to $(\mathcal{G}_{\mathrm{im}}, \mathcal{H}_{\mathrm{im}})$.

For the second point, we have $v \leq_{\mathcal{G}_{\mathrm{ker}}} w$ if and only if $\rho_v^w \in \ker R_\Delta$. This is the case if and only if ρ_v^w is contained in O and acts trivially on A_Δ. The claim now follows from Lemma 5.3 and Lemma 5.4. □

Stabilisers in image and kernel

Theorem 5.5 gives a complete description of the image and kernel of a restriction map $R_\Delta \colon \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t) \to \mathrm{Out}(A_\Delta)$ in the case where $\mathcal{G}$ is saturated with respect to $(\mathcal{G}, \mathcal{H})$. However, if we consider a subgroup of the form

$$\mathrm{Stab}_O(A_\Delta) = \mathrm{Out}^0(A_\Gamma; \mathcal{G} \cup \{A_\Delta\}, \mathcal{H}^t),$$

the family $\mathcal{G} \cup \{A_\Delta\}$ is not necessarily saturated with respect to $(\mathcal{G} \cup \{A_\Delta\}, \mathcal{H})$ and its image under R_Δ is more difficult to describe. The parabolic subgroups we will consider in Section 5.2 are exactly of this form. The next two lemmas show that in special cases, we can describe their images under R_Δ without passing to saturated pairs.

Lemma 5.11. *Assume that O stabilises a special subgroup $A_\Delta \leq A_\Gamma$ and let $R_\Delta \colon O \to \mathrm{Out}(A_\Delta)$ denote the corresponding restriction homomorphism. Take $\Theta \subset \Gamma$. Then:*

1. $\mathrm{Stab}_O(A_\Theta) \cap \ker R_\Delta = \mathrm{Stab}_{\ker R_\Delta}(A_\Theta)$.

2. If $\Theta \subseteq \Delta$, one has $R_\Delta(\mathrm{Stab}_O(A_\Theta)) = \mathrm{Stab}_{\mathrm{im}\, R_\Delta}(A_\Theta)$.

Proof. The first point becomes tautological after spelling out the definitions.

For the second point, the inclusion "$\subseteq$" is clear. On the other hand, each $\Phi \in \mathrm{im}\, R_\Delta$ can by definition be written as $\Phi = [\psi|_{A_\Delta}]$ where $[\psi] \in O$ and $\psi(A_\Delta) = A_\Delta$. If Φ stabilises A_Θ, we know that ψ conjugates A_Θ to a subgroup of A_Δ. Hence, $[\psi] \in \mathrm{Stab}_O(A_\Theta)$ and the second claim follows. □

Lemma 5.12. *Assume that Γ can be decomposed as a join $\Gamma = Z * \Delta$ where Z is a complete graph and $A_Z \in \mathcal{G}$. Let $P_\Delta \colon O \to \mathrm{Out}^0(A_\Delta)$ denote the corresponding projection map. Then for every $\Theta \subset \Gamma$, one has*

$$P_\Delta(\mathrm{Stab}_O(A_\Theta)) = \mathrm{Stab}_{\mathrm{im}\, P_\Delta}(A_{\Theta\cap\Delta}).$$

Proof. The stabiliser $\mathrm{Stab}_O(A_\Theta)$ is the same as the relative automorphism group $\mathrm{Out}^0(A_\Gamma; \mathcal{G} \cup \{A_\Theta\}, \mathcal{H}^t)$. By Lemma 5.7, the image of this group is equal to

$$P_\Delta(\mathrm{Stab}_O(A_\Theta)) = \mathrm{Out}^0(A_\Delta; \mathcal{G}_\Delta \cup \{A_{\Theta\cap\Delta}\}, \mathcal{H}^t_\Delta). \tag{5.3}$$

On the other hand, we have $\mathrm{im}\, P_\Delta = \mathrm{Out}^0(A_\Delta; \mathcal{G}_\Delta, \mathcal{H}^t_\Delta)$, so the right hand side of Eq. (5.3) is also equal to $\mathrm{Stab}_{\mathrm{im}\, P_\Delta}(A_{\Theta\cap\Delta})$ and the claim follows. □

5.1.4 Restrictions to conical subgroups

In this section, we define a family of special subgroups that will play an important role in our inductive arguments later on and study some properties of these special subgroups.

For a vertex $v \in V(\Gamma)$, define the following subgraphs of Γ:

$$\Gamma_{\succeq v} := \{w \in V(\Gamma) \mid v \preceq w\} \text{ and } \Gamma_{\succ v} := \{w \in V(\Gamma) \mid v \prec w\},$$

where $v \prec w$ if $v \preceq w$ and $w \not\preceq_{\mathcal{G}} v$. We define

$$A_{\succeq v} := A_{\Gamma_{\succeq v}} \text{ and } A_{\succ v} := A_{\Gamma_{\succ v}}$$

as the special subgroups of A_Γ corresponding to these subgraphs. Note that these special subgroups only depend on the $\sim_{\mathcal{G}}$-equivalence class of v, i.e. if $v \sim_{\mathcal{G}} w$, we have $A_{\succeq v} = A_{\succeq w}$.

In the "absolute setting" where $\mathcal{G}$ and $\mathcal{H}$ are trivial and $\preceq$ is equal to the standard ordering of $V(\Gamma)$, these special subgroups appear as *admissible subgroups* in the work of Duncan–Remeslennikov [DR12]. We will also refer to them as *conical subgroups* of A_Γ as they are generated by elements corresponding to an upwards-closed cone in the Hasse diagram of the partial order that $\preceq$ induces on the equivalence classes of $\sim_{\mathcal{G}}$ (see Fig. 5.3).

The elements of $\mathrm{Out}^0(A_\Gamma)$ are characterised among all elements of $\mathrm{Out}(A_\Gamma)$ by the property that they stabilise these special subgroups. Namely, the following holds true:

Lemma 5.13 ([DW, Proposition 3.3])**.** *Let $\mathcal{G}_\geq$ be the set of special subgroups of A_Γ of the form $A_{\geq v}$. Then*

$$\mathrm{Out}^0(A_\Gamma) = \mathrm{Out}(A_\Gamma; \mathcal{G}_\geq).$$

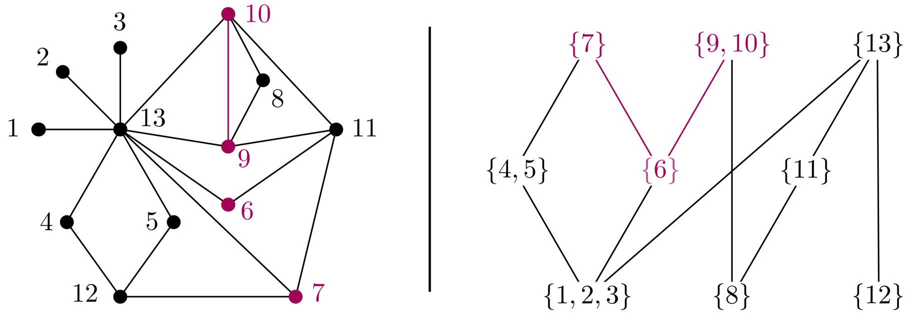

Figure 5.3: A graph Γ with vertex set $V(\Gamma) = \{1, \dots, 13\}$ and the Hasse diagram of the associated partial order $\leq$ of the standard equivalence classes of its vertices. The conical subgroup $\Gamma_{\geq 6}$ is marked in magenta.

In particular, each of these special subgroups is stabilised by all of $\mathrm{Out}^0(A_\Gamma)$. We need the following relative version of this statement.

Lemma 5.14. *Let $v, x \in V(\Gamma)$, let K be a union of $\mathcal{G}^x$-components of $\Gamma \setminus \mathrm{st}(x)$ and let $\pi_K^x \in O$ denote the corresponding partial conjugation. If $v \not\preceq x$, the partial conjugation π_K^x acts trivially on $A_{\succeq v}$.*

Proof. As v is not smaller than x with respect to $\preceq$, there either is an element in $\mathrm{lk}(v)$ which is not contained in $\mathrm{st}(x)$ or there is $A_\Delta \in \mathcal{G}$ such that Δ contains v but does not contain x. We claim that in both cases, $\Gamma_{\succeq v}$ intersects at most one $\mathcal{G}^x$-component of $\Gamma \setminus \mathrm{st}(x)$. Lemma 5.4 then implies that π_K^x acts trivially on $A_{\succeq v}$.

Indeed, if there is $y \in \mathrm{lk}(v) \setminus \mathrm{st}(x)$, one has $y \in \mathrm{st}(w) \setminus \mathrm{st}(x)$ for all $w \in \Gamma_{\succeq v}$. Hence, all elements of $\Gamma_{\succeq v}$ are adjacent to x and $\Gamma_{\succeq v} \setminus \mathrm{st}(x)$ is contained in a single $\mathcal{G}^x$-component of $\Gamma \setminus \mathrm{st}(x)$. If on the other hand for some $A_\Delta \in \mathcal{G}$ one has $v \in \Delta$, then it follows that $w \in \Delta$ for all $w \in \Gamma_{\succeq v}$. If $x \notin \Delta$, this implies that all elements of $\Gamma_{\succeq v}$ are $\mathcal{G}^x$-adjacent. In particular, they lie in the same $\mathcal{G}^x$-component. ☐

Proposition 5.15. *For every vertex $v \in V(\Gamma)$, the special subgroup $A_{\succeq v}$ is stabilised by every element from O.*

Proof. As O is generated by the inversions, transvections and partial conjugation it contains, it suffices to prove the statement for each such element. For this, we again use Lemma 5.3 and Lemma 5.4.

For inversions, there is nothing to show as they always stabilise every special subgroup. If there is a transvection $\rho_x^y \in O$, we have $x \preceq y$. The set $\Gamma_{\succeq v}$ is upwards-closed with respect to $\preceq$, hence $x \in \Gamma_{\succeq v}$ implies $y \in \Gamma_{\succeq v}$. It follows that ρ_x^y stabilises $A_{\succeq v}$. Given a partial conjugation $\pi_K^x \in O$, we either have $v \not\preceq x$, in which case Lemma 5.14 implies that π_K^x even acts trivially on $A_{\succeq v}$, or we have $x \in \Gamma_{\succeq v}$ which implies that π_K^x stabilises $A_{\succeq v}$. ☐

A consequence of this is that for every equivalence class $[v]_\mathcal{G}$ of vertices of Γ, there is a restriction map

$$R_{\succeq v} = R_{A_{\succeq v}} : O \to \mathrm{Out}^0(A_{\succeq v}).$$

We already came across such a restriction map in Example 5.1: There, A_Γ was a conical subgroup of $A_{\Gamma'}$ (one has $\Gamma = \Gamma'_{\geq 1}$, as can be seen in Fig. 5.2) and we saw a detailed description of the image and kernel of R_Γ on page 71. In general, the following holds true:

Lemma 5.16. *Let $v \in V(\Gamma)$ and let $R \coloneqq R_{\succeq v}$ denote the restriction homomorphism to $A_{\succeq v}$. Then:*

1. *For all $\Delta \subseteq \Gamma_{\succeq v}$, one has $\ker R \subseteq \mathrm{Stab}_O(A_\Delta)$.*
2. *For all $w \in V(\Gamma)$, the following holds: If $\Delta \subseteq \Gamma_{\succeq w}$ such that*

$$\Gamma_{\succeq v} \cap \Gamma_{\succeq w} \subseteq \Delta,$$

the stabiliser $\mathrm{Stab}_O(A_\Delta)$ contains all inversions, transvections and partial conjugations of O which are not contained in $\ker R$.

Proof. By Theorem 5.5, the kernel of R consists of all elements from O that act trivially on the special subgroup $A_{\succeq v}$. This immediately implies the first claim.

For the second one, we again use Lemma 5.3 and Lemma 5.4. First note that $\mathrm{Stab}_O(A_\Delta)$ contains all inversions of O.

Next assume we have a transvection $\rho_x^y \in O$. If $x \notin \Gamma_{\succeq v}$, the transvection is contained in $\ker R$. If on the other hand $x \notin \Delta$, the transvection ρ_x^y acts trivially on A_Δ and hence is contained in $\mathrm{Stab}_O(A_\Delta)$. Now observe that the assumption $\Gamma_{\succeq v} \cap \Gamma_{\succeq w} \subseteq \Delta$ implies that $\Delta \cap \Gamma_{\succeq v}$ is equal to $\Gamma_{\succeq v} \cap \Gamma_{\succeq w}$, a set which is upwards-closed with respect to $\preceq$. So if $x \in \Delta \cap \Gamma_{\succeq v}$, we also have $y \in \Delta \cap \Gamma_{\succeq v}$. Again it follows that $\rho_x^y \in \mathrm{Stab}_O(A_\Delta)$.

Lastly, consider a partial conjugation $\pi_K^x \in O$. If $v \npreceq x$, Lemma 5.14 implies that π_K^x is contained in $\ker R$. This lemma also shows that if $w \npreceq x$, the partial conjugation π_K^x acts trivially on $A_{\succeq w}$, and hence is contained in $\mathrm{Stab}_O(A_\Delta)$. The only case that remains is that x is greater than both v and w, i.e. $x \in \Gamma_{\succeq v} \cap \Gamma_{\succeq w}$. As we assumed that $\Gamma_{\succeq v} \cap \Gamma_{\succeq w} \subseteq \Delta$, this implies that $x \in \Delta$, so again $\pi_K^x \in \mathrm{Stab}_O(A_\Delta)$. □

5.2 A spherical complex for $\mathrm{Out}(A_\Gamma)$

In this section, we define maximal parabolic subgroups of $\mathrm{Out}^0(A_\Gamma)$ in the general case. We then prove Theorem A which states that the coset complex associated to these parabolic subgroups is homotopy equivalent to a wedge of spheres.

Notation and standing assumptions As before, let Γ be a graph, $\mathcal{G}$ and $\mathcal{H}$ families of special subgroups of A_Γ such that $\mathcal{G}$ is saturated with respect to $(\mathcal{G}, \mathcal{H})$, define $O \coloneqq \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ and set $\preceq \coloneqq \leq_{\mathcal{G}}$ to be the $\mathcal{G}$-ordering on $V(\Gamma)$. Let $T_{\mathcal{G}}$ denote the set of $\sim_{\mathcal{G}}$-equivalence classes of vertices of Γ.

5.2.1 Rank and maximal parabolic subgroups

Definition 5.17. We define the *rank of O* as

$$\mathrm{rk}(O) \coloneqq \sum_{[v]_{\mathcal{G}} \in T_{\mathcal{G}}} (|[v]_{\mathcal{G}}| - 1) = |V(\Gamma)| - |T_{\mathcal{G}}|.$$

Now fix an ordering $[v]_\mathcal{G} = \{v_1, \ldots, v_n\}$ on each equivalence class $[v]_\mathcal{G} \in T_\mathcal{G}$. For all $[v]_\mathcal{G} \in T_\mathcal{G}$ and $1 \leq j \leq n-1$, let $\Delta_v^j \subset \Gamma$ be the full subgraph of Γ with vertex set $\{v_1, \ldots, v_j\} \cup \Gamma_{\succ v}$.

Lemma 5.18. *For all $[v]_\mathcal{G} \in T_\mathcal{G}$ and $1 \leq j \leq n-1$, the stabiliser $\mathrm{Stab}_O(A_{\Delta_v^j})$ is a proper subgroup of O.*

Proof. Again, we use Lemma 5.3 and Lemma 5.4: As all vertices of $[v]_\mathcal{G}$ are equivalent with respect to $\leq_\mathcal{G}$, the transvection $\rho_{v_1}^{v_n}$ is an element of O. However, this transvection does not stabilise $A_{\Delta_v^j}$ because v_1 is contained in Δ_v^j while v_n is not. □

For $[v]_\mathcal{G} \in T_\mathcal{G}$, let

$$\mathcal{P}_{[v]_\mathcal{G}} := \left\{ \mathrm{Stab}_O(A_{\Delta_v^j}) \,\middle|\, 1 \leq j \leq n-1 \right\},$$

where if $|[v]_\mathcal{G}| = 1$, this is to be understood as $\mathcal{P}_{[v]_\mathcal{G}} = \emptyset$.

Definition 5.19. We define the set of *maximal standard parabolic subgroups of* O as the union

$$\mathcal{P}(O) := \bigcup_{[v]_\mathcal{G} \in T_\mathcal{G}} \mathcal{P}_{[v]_\mathcal{G}}.$$

The reader might at this point want to verify that for the graph Γ depicted in Fig. 5.3 on page 76, one has $|\mathcal{P}(\mathrm{Out}^0(A_\Gamma))| = 4$. The term "maximal" parabolic will become clear in Section 5.4 where we define and study parabolic subgroups of lower rank. As before, we will usually leave out the adjective "standard" (see Remark 3.14).

Remark 5.20. We note the following properties of $\mathcal{P}(O)$ and $\mathrm{rk}(O)$:

1. We have $\mathrm{rk}(O) = |\mathcal{P}(O)|$. We will also see an alternative interpretation of $\mathrm{rk}(O)$ in Section 5.4.3.

2. By Lemma 5.18, every element of $\mathcal{P}(O)$ is a proper subgroup of O.

3. Following Remark 5.9, the definition of parabolic subgroups depends on the ordering chosen for each equivalence class, but not on the pair $(\mathcal{G}, \mathcal{H})$ we chose to represent O.

4. If O is equal to $\mathrm{GL}_n(\mathbb{Z})$ or a Fouxe-Rabinovitch group, we recover the definitions of parabolic subgroups in these groups as defined in Section 3.2.1 and Section 4.2.1. Furthermore, $\mathrm{rk}(\mathrm{GL}_n(\mathbb{Z})) = \mathrm{rk}(\mathrm{Out}(F_n)) = n-1$.

Note that it is possible that there is no $\mathcal{G}$-equivalence class of size bigger than one. In this case, the rank of O is zero and $\mathcal{P}(O)$ is empty (see e.g. the standard equivalence classes of the graph Γ' in Example 5.1) . For further comments on this, see Section 5.5.

5.2.2 The parabolic sieve

In this subsection, we explain the idea of the inductive argument that we will use to show sphericity of the coset complexes $\mathrm{CC}(O, \mathcal{P}(O))$.

Outline of proof Whenever $\Delta \subset \Gamma$ is stabilised by O, the restriction map R_Δ gives rise to a short exact sequence

$$1 \to N \to O \overset{R_\Delta}{\to} Q \to 1$$

and by Theorem 5.5, both N and Q are relative automorphism groups of RAAGs again. Using the considerations of Section 5.1, we will show that for the correct choice of Δ, every $P \in \mathcal{P}(O)$ satisfies the following dichotomy: Either $R_\Delta(P)$ is contained in $\mathcal{P}(Q)$ or $P \cap N$ forms an element of $\mathcal{P}(N)$. Applying a restriction homomorphism hence has the effect of a sieve on $\mathcal{P}(O)$—some of the parabolic subgroups pass through and form parabolics of the quotient Q while others remain in the sieve and form parabolics of the subgroup N. Now using the results of Section 3.3, this allows us to describe the homotopy type of $\mathrm{CC}(O, \mathcal{P}(O))$ in terms of the topology of the lower-dimensional coset complexes $\mathrm{CC}(Q, \mathcal{P}(Q))$ and $\mathrm{CC}(N, \mathcal{P}(N))$. This is used for an inductive argument with two phases: We first apply restriction maps to conical subgroups and then analyse the homotopy type of coset complexes in the conical setting. Concrete examples of this induction are given in Section 5.3.

Conical restrictions

Lemma 5.21 (Induction step)**.** *Let $v \in V(\Gamma)$ and let $R := R_{\succeq v}$ denote the corresponding restriction map to $A_{\succeq v}$. Then there is a homotopy equivalence*

$$\mathrm{CC}(O, \mathcal{P}(O)) \simeq \mathrm{CC}(\operatorname{im} R, \mathcal{P}(\operatorname{im} R)) * \mathrm{CC}(\ker R, \mathcal{P}(\ker R)).$$

Proof. Set $\mathcal{P} := \mathcal{P}(O)$. We want to apply Corollary 3.30, so we have to show that for each $P \in \mathcal{P}$, we either have $\ker R \subseteq P$ or P contains all inversions, transvections and partial conjugations of O which are not contained in $\ker R$.

Take $[w]_{\mathcal{G}} \in T_{\mathcal{G}}$ and $P = \mathrm{Stab}_O(A_\Delta) \in \mathcal{P}_{[w]_{\mathcal{G}}}$ with $\Delta = \Delta_w^j$ as above. If $v \preceq w$, we have $\Delta \subset \Gamma_{\succeq v}$. Hence by the first point of Lemma 5.16, we know that $\ker R \subseteq P$. If on the other hand $w \prec v$, one has $\Gamma_{\succeq v} \cap \Gamma_{\succeq w} = \Gamma_{\succeq v} \subset \Delta$. Similarly if v and w are incomparable, one has $\Gamma_{\succeq v} \cap \Gamma_{\succeq w} \subseteq \Gamma_{\succ w} \subset \Delta$. In both cases, the second point of Lemma 5.16 tells us that P contains all inversions, transvections and partial conjugations of O which are not contained in $\ker R$. From this, it follows that

$$\mathcal{P}_{\ker R} = \left\{ P \in \mathcal{P}_{[w]_{\mathcal{G}}} \mid v \preceq w \right\} \quad \text{and} \quad \mathcal{P}^{\ker R} = \left\{ P \in \mathcal{P}_{[w]_{\mathcal{G}}} \mid v \npreceq w \right\}, \quad (5.4)$$

with notation as defined on page 23. Corollary 3.30 now shows that there is a homotopy equivalence

$$\mathrm{CC}(O, \mathcal{P}) \simeq \mathrm{CC}(\operatorname{im} R, \overline{\mathcal{P}}) * \mathrm{CC}(\ker R, \mathcal{P} \cap \ker R),$$

where $\overline{\mathcal{P}} = \{R(P) \mid P \in \mathcal{P}_{\ker R}\}$ and $\mathcal{P} \cap \ker R = \left\{ P \cap \ker R \mid P \in \mathcal{P}^{\ker R} \right\}$.

If $P \in \mathcal{P}_{\ker R}$, there is $\Delta \subset \Gamma_{\succeq v}$ such that $P = \mathrm{Stab}_O(A_\Delta)$. Using Lemma 5.11, it follows that $R(P) = \mathrm{Stab}_{\operatorname{im} R}(A_\Delta)$. Lemma 5.10 implies that

$$\overline{\mathcal{P}} = \mathcal{P}(\operatorname{im} R).$$

For every $P = \mathrm{Stab}_O(A_\Delta) \in \mathcal{P}^{\ker R}$, we know by Lemma 5.11 that

$$P \cap \ker R = \mathrm{Stab}_{\ker R}(A_\Delta).$$

Write $\ker R = \mathrm{Out}^0(A_\Gamma; \mathcal{G}_{\ker}, \mathcal{H}^t_{\ker})$ where $\mathcal{G}_{\ker}$ is saturated with respect to $(\mathcal{G}_{\ker}, \mathcal{H}_{\ker})$. Then by Lemma 5.10, for $x, y \in V(\Gamma)$, we have $x \leq_{\mathcal{G}_{\ker}} y$ if and only if $v \not\preceq x$ and $x \preceq y$. Combining this with Eq. (5.4), it follows that

$$\mathcal{P} \cap \ker R = \mathcal{P}(\ker R).$$

This finishes the proof. □

For the first phase of our induction, we now use this iteratively in order to obtain:

Proposition 5.22. *There is a homotopy equivalence*

$$\mathrm{CC}(O, \mathcal{P}(O)) \simeq *_{[v]_\mathcal{G} \in T_\mathcal{G}} \mathrm{CC}(O_v, \mathcal{P}(O_v)),$$

where for all $[v]_\mathcal{G} \in T_\mathcal{G}$, one has $O_v = \mathrm{Out}^0(A_{\succeq v}; \mathcal{G}_v, \mathcal{H}^t_v)$ such that:

1. *$\mathcal{G}_v$ is saturated with respect to $(\mathcal{G}_v, \mathcal{H}^t_v)$,*
2. *$\mathcal{H}_v = \mathcal{H}_{\Gamma_{\succeq v}} \cup \{A_{\succeq w} \mid v \prec w\}$,*
3. *for $x \neq y \in \Gamma_{\succeq v}$, one has $x \leq_{\mathcal{G}_v} y$ if and only if $x \in [v]_\mathcal{G}$ and $x \preceq y$.*

Proof. We want to inductively use the restriction maps $R_{\succeq w}$. In order to do this, assume that we have shown that $\mathrm{CC}(\mathrm{Out}^0(A_\Gamma), \mathcal{P}(O))$ is homotopy equivalent to a join of coset complexes of the form

$$\mathrm{CC}(U, \mathcal{P}(U)),$$

where $U = \mathrm{Out}^0(A_\Theta; \mathcal{E}, \mathcal{F}^t)$ with $\Theta = \Gamma_{\succeq v}$ and such that the following hypotheses hold:

1. $\mathcal{E}$ is saturated with respect to $(\mathcal{E}, \mathcal{F}^t)$,
2. $\mathcal{F} = \mathcal{H}_\Theta \cup \mathcal{F}'$ where $\mathcal{F}' \subseteq \{A_{\succeq w} \mid v \prec w\}$,
3. for all $w \in \Theta$, either U acts trivially on $A_{\succeq w}$ or $\Theta_{\geq_\mathcal{E} w} = \Theta_{\succeq w}$.

Note that a priori, there is a slight ambiguity in writing $A_{\succeq w}$ without specifying the ambient graph. Here, however, we can ignore this issue because for all $w \succ v$, we have $\Gamma_{\succeq w} = \Theta_{\succeq w}$. For technical reasons we allow v to be a formal element 0 with $\Gamma_{\succeq 0} := \Gamma$. These three hypotheses hold in particular for the initial case where $v = 0$, $\mathcal{E} = \mathcal{G}$ and $\mathcal{F} = \mathcal{H}$.

Now assume that there is $w \in \Theta$ such that U does not act trivially on $A_{\succeq w}$. In this case, we have $\Theta_{\geq_\mathcal{E} w} = \Theta_{\succeq w}$, so by Proposition 5.15, the special subgroup $A_{\succeq w} \leq A_\Theta$ is stabilised by U and we can consider the restriction map $R\colon U \to \mathrm{Out}^0(A_{\succeq w})$. By Lemma 5.21, this yields a homotopy equivalence

$$\mathrm{CC}(U, \mathcal{P}(U)) \simeq \mathrm{CC}(\mathrm{im}\, R, \mathcal{P}(\mathrm{im}\, R)) * \mathrm{CC}(\ker R, \mathcal{P}(\ker R)).$$

By Theorem 5.5, we can write

$$\ker R = \mathrm{Out}^0(A_\Theta; \mathcal{E}_{\ker}, (\mathcal{F} \cup \{A_{\succeq w}\})^t),$$

where $\mathcal{E}_{\ker}$ is saturated with respect to the pair $(\mathcal{E}_{\ker}, \mathcal{F} \cup \{A_{\succeq w}\})$. Furthermore, Lemma 5.10 together with the third hypothesis of our induction implies that for all $w \in \Theta$, either $\ker R$ acts trivially on $A_{\succeq w}$ or $\Theta_{\geq_{\mathcal{E}_{\ker}} w} = \Theta_{\succeq w}$. It follows that $\mathrm{CC}(\ker R, \mathcal{P}(\ker R))$ satisfies the induction hypotheses.

Again using Theorem 5.5, the image of R can be written as

$$\operatorname{im} R = \mathrm{Out}^0(A_{\succeq w}; \mathcal{E}_{\mathrm{im}}, \mathcal{F}^t_{\mathrm{im}}),$$

where $\mathcal{E}_{\mathrm{im}}$ is saturated with respect to $(\mathcal{E}_{\mathrm{im}}, \mathcal{F}_{\mathrm{im}})$ and the elements of $\mathcal{F}_{\mathrm{im}}$ are the special subgroups generated by the vertices of $\Delta \cap \Gamma_{\succeq w}$ for some $\Delta \in \mathcal{F}$. This implies that $\mathrm{CC}(\operatorname{im} R, \mathcal{P}(\operatorname{im} R))$ satisfies the first hypotheses of our induction. The third one is an immediate consequence of Lemma 5.10.

Now apply induction to these coset complexes. This process ends if we arrive at a case where for all $w \succ v$, the group U acts trivially on $A_{\succeq w}$. But then we can set $\mathcal{F} = \mathcal{H}_\Theta \cup \{A_{\succeq w} \mid v \prec w\}$ and for $x \neq y$, $x \leq_{\mathcal{E}} y$ is only possible if $x \in [v]_{\mathcal{G}}$. If $v = 0$, this means that the relative ordering on $\Gamma_{\preceq 0} = \Gamma$ is trivial, so $\mathcal{P}(U) = \emptyset$. If $v \neq 0$, the group $O_v := U$ satisfies all conditions of the claim. $\square$

Coset complexes of conical RAAGs

We now want to further investigate the coset complexes $\mathrm{CC}(O_v, \mathcal{P}(O_v))$ of conical RAAGs that we obtained in Proposition 5.22. This is why in this subsection, we impose the following assumptions.

Standing assumptions Until the end of Section 5.2.2, we assume:

1. There is a vertex $v \in V(\Gamma)$ such that $\Gamma = \Gamma_{\succeq v}$, i.e. every vertex of Γ is greater than or equal to v with respect to $\preceq$.

2. For all $w \succ v$, the group O acts trivially on $A_{\{w\}} \leq A_\Gamma$.

Observe that Item 1 implies that

$$\text{for all } A_\Delta \in \mathcal{G}, \text{ we have } \Delta \subseteq \Gamma_{\succ v}. \tag{5.5}$$

By Lemma 5.3 and Lemma 5.4, every $\Delta \subseteq \Gamma$ such that O stabilises A_Δ must be upwards-closed with respect to $\preceq$. Hence, if Δ intersects $[v]_{\mathcal{G}}$ non-trivially, it follows that $\Delta = \Gamma$. Furthermore, Item 2 implies that for all $w \succ v$, the equivalence class $[w]_{\mathcal{G}}$ is a singleton. Combining this, we obtain

$$\mathcal{P}(O) = \mathcal{P}_{[v]_{\mathcal{G}}}. \tag{5.6}$$

In this situation, let

$$Z := \{w \in V(\Gamma) \mid v \prec w \text{ and } w \text{ is adjacent to } v\}.$$

We define the *group of twists by elements in* $\Gamma_{\succ v}$ as the subgroup $T \leq O$ generated by the transvections ρ^z_x with $x \in [v]_{\mathcal{G}}$ and $z \in Z$.

Lemma 5.23. *T is a free abelian group. Furthermore, Γ can be decomposed as a join $\Gamma = Z * \Delta$ and there is a short exact sequence*

$$1 \to T \to O \xrightarrow{P_\Delta} \mathrm{Out}^0(A_\Delta; \mathcal{G}_\Delta, \mathcal{H}^t_\Delta) \to 1.$$

Proof. If $Z = \emptyset$, the statement is trivial, so we can assume that Z contains at least one element. By definition, we have $Z \subseteq \mathrm{lk}(v) \setminus [v]_{\mathcal{G}}$. As every vertex of Γ is greater than or equal to v with respect to $\preceq$, this implies that Z is a complete graph and we can write $\Gamma = Z * \Delta$.

Using the assumption that O acts trivially on $A_{\{w\}}$ for all $w \succ v$, Lemma 5.3 and Lemma 5.4 imply that O acts trivially on the normal subgroup $A_Z \triangleleft A_\Gamma$. Consequently, we have a well-defined projection map $P_\Delta \colon O \to \mathrm{Out}(A_\Delta)$. By Lemma 5.7, the image of this map is equal to $\mathrm{Out}^0(A_\Delta;\, \mathcal{G}_\Delta, \mathcal{H}^t_\Delta)$.

The description of the kernel $\ker P_\Delta$ as the free abelian group T generated by the transvections ρ^z_x with $x \in [v]_{\mathcal{G}}$ and $z \in Z$ follows from [CV09, Proposition 4.4] because O acts trivially on A_Z (see also [DW, 5.1.4]). □

Lemma 5.24. *Let $\Delta := \Gamma \setminus Z$ and let P_Δ denote the corresponding projection map. There is a homotopy equivalence*

$$\mathrm{CC}(O,\, \mathcal{P}(O)) \simeq \mathrm{CC}(\mathrm{im}\, P_\Delta,\, \mathcal{P}(\mathrm{im}\, P_\Delta)).$$

Proof. By Lemma 5.23, we have a short exact sequence

$$1 \to T \to O \xrightarrow{P_\Delta} \mathrm{im}\, P_\Delta \to 1,$$

where $\mathrm{im}\, P_\Delta = \mathrm{Out}^0(A_\Delta;\, \mathcal{G}_\Delta, \mathcal{H}^t_\Delta)$. We first claim that every parabolic subgroup $P \in \mathcal{P}(O)$ contains T. Indeed, we observed above that $\mathcal{P}(O) = \mathcal{P}_{[v]_{\mathcal{G}}}$ (see Eq. (5.6)). By definition, every $P \in \mathcal{P}_{[v]_{\mathcal{G}}}$ is of the form $P = \mathrm{Stab}_O(A_\Theta)$ for some Θ containing $\Gamma_{\succ v}$. The claim now follows from Lemma 5.4.

Hence, Corollary 3.30 yields a homotopy equivalence

$$\mathrm{CC}(O,\, \mathcal{P}(O)) \simeq \mathrm{CC}(\mathrm{im}\, P_\Delta,\, \overline{\mathcal{P}(O)}) * \emptyset = \mathrm{CC}(\mathrm{im}\, P_\Delta,\, \overline{\mathcal{P}(O)})$$

The ordering $\leq_{\mathcal{G}_\Delta}$ is just the restriction of $\preceq$ to Δ, so Lemma 5.12 implies that $\overline{\mathcal{P}(O)} = \mathcal{P}(\mathrm{im}\, P_\Delta)$. □

We now distinguish between the case where $[v]_{\mathcal{G}}$ is an abelian and the case where it is a free equivalence class.

Lemma 5.25. *Assume that $\Gamma = \Gamma_{\succeq v}$, where $[v]_{\mathcal{G}}$ is an abelian equivalence class of size $n := |[v]_{\mathcal{G}}| \geq 2$. Then the coset complex*

$$\mathrm{CC}(O,\, \mathcal{P}(O))$$

is homotopy equivalent to the Tits building associated to $\mathrm{GL}_n(\mathbb{Q})$.

Proof. By Lemma 5.24, we have a homotopy equivalence

$$\mathrm{CC}(O,\, \mathcal{P}(O)) \simeq \mathrm{CC}(\mathrm{im}\, P_\Delta,\, \mathcal{P}(\mathrm{im}\, P_\Delta)),$$

where $\Delta = \Gamma \setminus Z$ and $\mathrm{im}\, P_\Delta = \mathrm{Out}^0(A_\Delta;\, \mathcal{G}_\Delta, \mathcal{H}^t_\Delta)$.

By assumption, the abelian equivalence class $[v]_{\mathcal{G}}$ contains at least two elements that are adjacent to each other. As every vertex of Γ is greater than or equal to v with respect to $\preceq$, this implies that every vertex of $\Gamma_{\succ v}$ must be adjacent to v. Hence, $Z = \Gamma_{\succ v}$ and $\Delta = [v]_{\mathcal{G}}$. As observed above (see Eq. (5.5)), every $\Theta \subseteq \Gamma$ with $A_\Theta \in \mathcal{G}$ is entirely contained in $\Gamma_{\succ v}$. Consequently, we have $\mathcal{G}_\Delta = \mathcal{H}_\Delta = \emptyset$ and

$$\mathrm{Out}^0(A_\Delta;\, \mathcal{G}_\Delta, \mathcal{H}^t_\Delta) = \mathrm{GL}_n(\mathbb{Z}).$$

This means that $\mathrm{CC}(O, \mathcal{P}(O)) \simeq \mathrm{CC}(\mathrm{GL}_n(\mathbb{Z}), \mathcal{P}(\mathrm{GL}_n(\mathbb{Z})))$ and this coset complex is isomorphic to the Tits building associated to $\mathrm{GL}_n(\mathbb{Q})$ by Proposition 3.15. □

In the setting of a free equivalence class, the situation is slightly more complicated: As before, we start by projecting away from Z, but we then might have to apply further restriction maps.

Lemma 5.26. *Assume that $\Gamma = \Gamma_{\succeq v}$ where $[v]_{\mathcal{G}}$ is a free equivalence class of size $n := |[v]_{\mathcal{G}}| \geq 2$. Then there is a special subgroup $A \leq A_\Gamma$ such that*

$$A = F * A_1 * \cdots * A_k$$

where F is the free group of rank n generated by $[v]_{\mathcal{G}}$ and the coset complex

$$\mathrm{CC}(O, \mathcal{P}(O))$$

is homotopy equivalent to the free factor complex $\mathcal{F}(A, \mathcal{A})$ of A relative to the free factor system $\mathcal{A} := \{[A_1], \ldots, [A_k]\}$.

Proof. Again by Lemma 5.24, we have a homotopy equivalence

$$\mathrm{CC}(O, \mathcal{P}(O)) \simeq \mathrm{CC}(\operatorname{im} P_\Delta, \mathcal{P}(\operatorname{im} P_\Delta)),$$

where $\Delta = \Gamma \setminus Z$ and $\operatorname{im} P_\Delta = \mathrm{Out}^0(A_\Delta; \mathcal{G}_\Delta, \mathcal{H}^t_\Delta)$. As noted above, the $\mathcal{G}_\Delta$-ordering on Δ is just the restriction of $\preceq$ to Δ; in particular we have $[v]_{\mathcal{G}_\Delta} = [v]_{\mathcal{G}}$ and $\Delta = \Delta_{\succeq v}$.

As no two vertices from $[v]_{\mathcal{G}}$ are adjacent to each other, the link $\mathrm{lk}_\Gamma(v)$ is entirely contained in Z, so every element of $[v]_{\mathcal{G}}$ forms an isolated vertex of Δ. This implies that Δ decomposes as a disjoint union $\Delta = [v]_{\mathcal{G}} \sqcup \bigsqcup \Delta_i$, where each Δ_i is a $\mathcal{G}_\Delta$-component of Δ. In particular, we have

$$A_\Delta = A_{[v]_{\mathcal{G}}} * A_{\Delta_1} * \ldots * A_{\Delta_k}.$$

Moreover, for all i, the group $\operatorname{im} P_\Delta$ stabilises A_{Δ_i}: If Δ_i contains at least two vertices, this is [DW, Lemma 3.13.1] and if Δ_i is a singleton, the action on A_{Δ_i} is trivial by the standing assumptions.

If there is an i such that $\operatorname{im} P_\Delta$ acts non-trivially on Δ_i, then there is a non-trivial restriction map $R\colon \operatorname{im} P_\Delta \to \mathrm{Out}(A_{\Delta_i})$. Its kernel can be written as

$$\ker R = \mathrm{Out}^0(A_\Delta; \mathcal{G}_{\ker}, (\mathcal{H}_\Delta \cup \{A_{\Delta_i}\})^t),$$

where $\mathcal{G}_{\ker}$ is saturated with respect to $(\mathcal{G}_{\ker}, \mathcal{H}_\Delta \cup \{A_{\Delta_i}\})$. One can easily check that each $P \in \mathcal{P}(\operatorname{im} P_\Delta)$ contains all the inversions, transvections and partial conjugations not contained in $\ker R$: The kernel contains all inversions and transvections from $\operatorname{im} P_\Delta$ as well as the partial conjugations that have acting letter contained in $[v]_{\mathcal{G}}$. The remaining partial conjugations are contained in all of the parabolic subgroups.

Hence, by Corollary 3.30, there is a homotopy equivalence

$$\mathrm{CC}(\operatorname{im} P_\Delta, \mathcal{P}(\operatorname{im} P_\Delta)) \simeq \emptyset * \mathrm{CC}(\ker R, \mathcal{P} \cap \ker R) = \mathrm{CC}(\ker R, \mathcal{P} \cap \ker R).$$

Lemma 5.10 implies that $\leq_{\mathcal{G}_{\ker}}$ agrees with $\preceq$ on Δ, so $\mathcal{P}(\ker R) = \mathcal{P} \cap \ker R$ by Lemma 5.11. Every A_{Δ_i} is stabilised by $\ker R$. Hence, we can use induction and apply restriction maps until we reach the group $\mathrm{Out}^0(A_\Delta; \{A_{\Delta_i}\}^t_i)$. This group is equal to $\mathrm{Out}(A_\Delta, \{A_{\Delta_i}\}^t_i)$, a Fouxe-Rabinovitch group. The claim now follows from Proposition 4.7. □

Using the results of Section 3.2.1 and Chapter 4, the last two lemmas can be summarised as:

Corollary 5.27. *Assume that* $\Gamma = \Gamma_{\succeq v}$*, where* $[v]_{\mathcal{G}}$ *is an equivalence class of size* $n := |[v]_{\mathcal{G}}|$*. Then* $\mathrm{CC}(O, \mathcal{P}(O))$ *is* $(n-2)$*-spherical.*

Proof. If $n = 1$, the statement is trivial as in this case, the set $\mathcal{P}(O) = \mathcal{P}_{[v]_{\mathcal{G}}}$ is empty and the complex $\mathrm{CC}(O, \mathcal{P}(O))$ is the empty set, which we consider to be (-1)-spherical (see Section 2.4). Now let $n \geq 2$. If $[v]_{\mathcal{G}}$ is abelian, Lemma 5.25 implies that the coset complex is homotopy equivalent to the Tits building associated to $\mathrm{GL}_n(\mathbb{Q})$, which is $(n-2)$-spherical by the Solomon–Tits Theorem 3.12. If on the other hand $[v]_{\mathcal{G}}$ is free, it is by Lemma 5.26 homotopy equivalent to a relative free factor complex, which is by Theorem 4.39 $(n-2)$-spherical as well. □

5.2.3 Proof of Theorem A

We return to the general situation where Γ is any graph and $\mathcal{G}$ and $\mathcal{H}$ are any families of special subgroups of A_Γ such that $\mathcal{G}$ is saturated with respect to $(\mathcal{G}, \mathcal{H})$. Recall that $\preceq$ denotes the $\mathcal{G}$-ordering of $V(\Gamma)$ and $T_{\mathcal{G}}$ denotes the set of associated $\sim_{\mathcal{G}}$-equivalence classes.

The only thing that is left to be done for the proof of Theorem A, which we restate below, is to collect the results obtained in Section 5.2.2.

Theorem 5.28. *Let* $O := \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$*. The coset complex* $\mathrm{CC}(O, \mathcal{P}(O))$ *is homotopy equivalent to a wedge of spheres of dimension* $\mathrm{rk}(O) - 1$*.*

Proof. By Proposition 5.22, there is a homotopy equivalence

$$\mathrm{CC}(O, \mathcal{P}(O)) \simeq *_{[v]_{\mathcal{G}} \in T_{\mathcal{G}}} \mathrm{CC}(O_v, \mathcal{P}(O_v)),$$

where for all $[v]_{\mathcal{G}} \in T_{\mathcal{G}}$, one has $O_v = \mathrm{Out}^0(A_{\succeq v}; \mathcal{G}_v, \mathcal{H}_v^t)$ such that:

1. $\mathcal{G}_v$ is saturated with respect to $(\mathcal{G}_v, \mathcal{H}_v^t)$,
2. $\mathcal{H}_v = \mathcal{H}_{\Gamma_{\succeq v}} \cup \{A_{\succeq w} \mid v \prec w\}$,
3. for $x \neq y \in \Gamma_{\succeq v}$, one has $x \leq_{\mathcal{G}_v} y$ if and only if $x \in [v]_{\mathcal{G}}$ and $x \preceq y$.

Let $[v]_{\mathcal{G}} \in T_{\mathcal{G}}$. Condition 3 implies that the $\mathcal{G}_v$-equivalence class of v in $\Gamma_{\succeq v}$ is equal to $[v]_{\mathcal{G}}$ and that all other $w \in \Gamma_{\succeq v}$ are greater than v with respect to $\leq_{\mathcal{G}_v}$. Now Condition 2 implies that for all w with $w >_{\mathcal{G}_v} v$, the group O_v acts trivially on $A_{\{w\}} \leq A_{\succeq v}$. Hence, the assumptions given on page 81 are fulfilled and Corollary 5.27 implies that $\mathrm{CC}(O_v, \mathcal{P}(O_v))$ is spherical of dimension $|[v]_{\mathcal{G}}| - 2$. It follows from Lemma 2.11 that the join of these complexes is spherical of dimension $\sum_{[v]_{\mathcal{G}} \in T_{\mathcal{G}}} (|[v]_{\mathcal{G}}| - 1) - 1 = \mathrm{rk}(O) - 1$. □

5.3 Summary of the inductive procedure and examples

5.3.1 Consequences for the induction of Day–Wade

The proof of Theorem 5.28 relies on the inductive procedure defined in [DW]. The authors there show that for every graph Γ, the group $\mathrm{Out}^0(A_\Gamma)$ has a

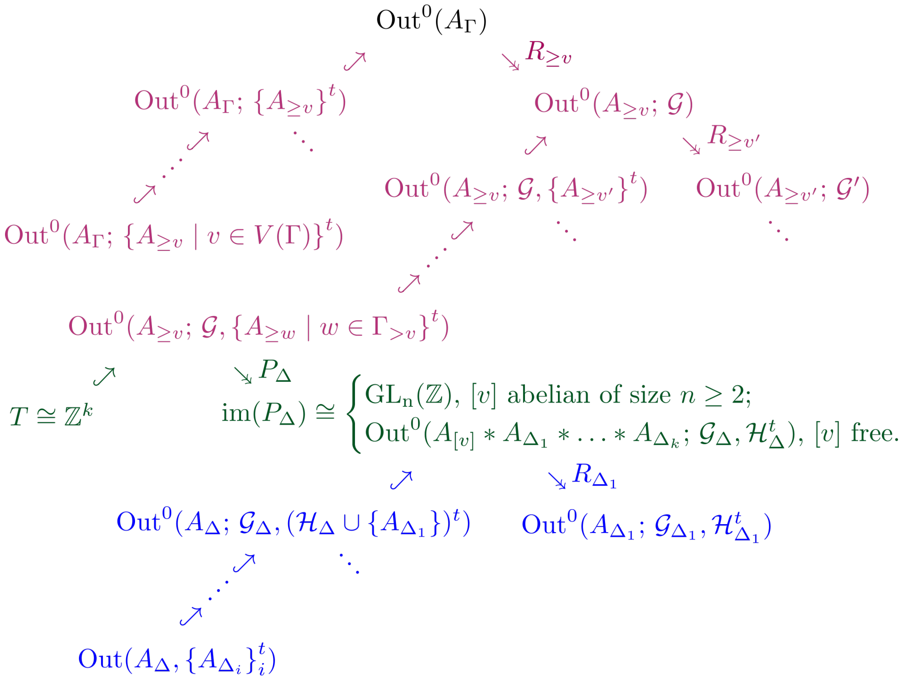

Figure 5.4: Decomposition of $\mathrm{Out}^0(A_\Gamma)$. Step 1 is coloured in magenta, Step 2 in green and Step 3 in blue.

subnormal series

$$1 = N_0 \leq N_1 \leq \ldots \leq N_k = \mathrm{Out}^0(A_\Gamma)$$

such that for all i, the quotient N_{i+1}/N_i is isomorphic to either a free abelian group, to $\mathrm{GL}_n(\mathbb{Z})$ or to a Fouxe-Rabinovitch group [DW, Theorem A]. The methods we use in Section 5.2.2 provide more detailed information about this inductive procedure which decomposes $\mathrm{Out}^0(A_\Gamma)$ in terms of short exact sequences related to restriction and projection homomorphisms: We are able to give an explicit description of the restriction and projection maps that one has to use during the induction and of the base cases one obtains this way. In what follows, we will give a summary of these results. See also Fig. 5.4.

To simplify notation, we will describe the decomposition of $O = \mathrm{Out}^0(A_\Gamma)$. However, all of this can also be stated in the more general case where O is any relative automorphism group of a RAAG.

Step 1 First one iteratively restricts to conical subgroups $A_{\geq v}$ until one is left with relative automorphism groups that act trivially on all of their proper conical subgroups—for this, one needs to apply exactly one restriction map for every (standard) equivalence class of $V(\Gamma)$ and the order in which one applies the corresponding restriction maps does not change the base cases of this first induction step. One of these base cases is given by the intersection of the kernels of all the conical restriction maps; this is the group $\mathrm{Out}^0(A_\Gamma; \{A_{\geq v} \mid v \in V(\Gamma)\}^t)$ which does not contain any inversions or transvections. The other base cases are all of the form $\mathrm{Out}^0(A_{\geq v}; \mathcal{G}, \{A_{\geq w} \mid w \in \Gamma_{>v}\}^t)$ for some $v \in V(\Gamma)$ and some family $\mathcal{G}$ of special subgroups of $A_{\geq v}$. There is exactly one such base case for every equivalence class $[v]$ of $V(\Gamma)$ and it is generated by all the restrictions to $A_{\geq v}$ of inversions, transvections and partial conjugations of $\mathrm{Out}^0(A_\Gamma)$ that act trivially on $A_{\geq w}$ for every $w > v$.

Step 2 Now for each of these groups, one applies the (possibly trivial) projection map P_Δ where $\Delta \coloneqq \Gamma_{\geq v} \setminus Z$ and Z is the full subgraph of $\Gamma_{\geq v}$ consisting of all those vertices of Γ which are adjacent to v and strictly greater than v with respect to the standard ordering on $V(\Gamma)$. The kernel of this projection map is given by the free abelian group T generated by all twists of elements in $[v]$ by elements in $\Gamma_{>v}$. We now have to distinguish two cases: If $[v]$ is an abelian equivalence class of size $n \geq 2$, then the image of P_Δ is given by $\mathrm{Out}(A_{[v]}) \cong \mathrm{GL}_n(\mathbb{Z})$. If this is not the case, we proceed with Step 3.

Step 3 If $[v]$ is a free equivalence class, the graph Δ decomposes as a disjoint union $\Delta = [v] \sqcup \bigsqcup \Delta_i$ where each Δ_i is a relative connected component of $\mathrm{im}(P_\Delta)$. One can show that the Δ_i are precisely the non-empty intersections $\Delta_i = (\Delta \setminus [v]) \cap \Gamma_i$ where Γ_i is a connected component of $\Gamma \setminus \mathrm{lk}(v)$. We now iteratively apply the restriction maps R_{Δ_i}. This yields two kinds of base cases: The first kind is given by the intersection of the kernels of all the R_{Δ_i} and can be described as the Fouxe-Rabinovitch group $\mathrm{Out}(A_\Delta, \{A_{\Delta_i}\}_i^t)$. The second kind is given by the images of the restriction maps. For each i, this is a relative automorphism group of A_{Δ_i}; as $\Delta_i \subseteq \Gamma_{>v}$, this group contains no inversions or transvections and is generated by partial conjugations.

The base cases In summary, our induction yields the following base cases:

1. The “left-most” kernel $\mathrm{Out}^0(A_\Gamma; \{A_{\geq v} \mid v \in V(\Gamma)\}^t)$;
2. for every abelian equivalence class $[v]$ of size $n \geq 2$:
   (a) the free abelian group T generated by all twist of elements in $[v]$ by elements in $\Gamma_{>v}$, which has rank $n \cdot |\Gamma_{>v} \cap \mathrm{lk}(v)|$;
   (b) $\mathrm{Out}(A_{[v]}) \cong \mathrm{GL}_n(\mathbb{Z})$;
3. for every free equivalence class $[v]$:
   (a) the free abelian group T generated by all twist of elements in $[v]$ by elements in $\Gamma_{>v}$, which has rank $n \cdot |\Gamma_{>v} \cap \mathrm{lk}(v)|$;
   (b) for every connected component Γ_i of $\Gamma \setminus \mathrm{lk}(v)$ such that
   $$\Delta_i \coloneqq (\Gamma_{>v} \setminus \mathrm{lk}(v)) \cap \Gamma_i \neq \emptyset :$$
   a subgroup of $\mathrm{Out}^0(A_{\Delta_i})$ generated by partial conjugations;
   (c) a Fouxe-Rabinovitch group $\mathrm{Out}(A_\Delta, \{A_{\Delta_i}\}_i^t)$ where $\Delta = \Gamma_{\geq v} \setminus \mathrm{lk}(v)$ and the Δ_i are as in Item 3b.

The base cases having a non-empty set of parabolic subgroups are Item 2b and Item 3c if $|[v]| \geq 2$. Note that we allow $|v| = 1$ in Item 3c which results in $\mathrm{Out}(\langle v \rangle) \cong \mathbb{Z}/2\mathbb{Z}$ if v is maximal. One should mention that Item 1 and Item 3b

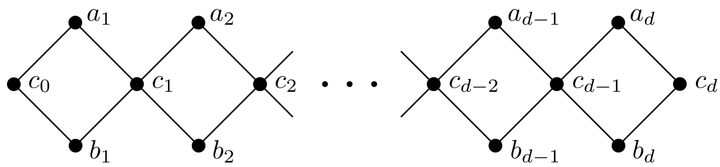

Figure 5.5: String of d diamonds.

are not necessarily base cases of the induction of Day–Wade: There might be further non-trivial restriction and projection maps and after applying them one can decompose these groups into Fouxe-Rabinovitch groups and free abelian groups generated by partial conjugations.

5.3.2 Examples

Strings of diamonds

Let Γ be the string of d diamonds (see Fig. 5.5), as considered in [CSV17, Section 5.3] and [DW, Section 6.3.1]. Assume $d \geq 2$. The standard equivalence classes of Γ are given by

$$[a_i] = [b_i] = \{a_i, b_i\}\,,\ 1 \leq i \leq d \qquad \text{and} \qquad [c_i] = \{c_i\}\,,\ 0 \leq i \leq d.$$

The conical subgroups here are

$$\begin{aligned} &\Gamma_{\geq a_i} = [a_i] = \{a_i, b_i\}\,,\ 1 \leq i \leq d, \\ &\Gamma_{\geq c_i} = [c_i],\ 1 \leq i \leq d-1, \\ &\Gamma_{\geq c_0} = [c_0] \cup [c_1] \cup [a_1] \text{ and } \Gamma_{\geq c_d} = [c_d] \cup [c_{d-1}] \cup [a_d]. \end{aligned}$$

If we order $[a_i]$ as (a_i, b_i), the family of maximal parabolic subgroups of the group $O := \mathrm{Out}^0(A_\Gamma)$ is given as

$$\mathcal{P}(O) = \{\mathrm{Stab}_O(\langle a_i \rangle) \mid 1 \leq i \leq d\}$$

and we have $\mathrm{rk}(O) = d$, i.e. $\mathrm{CC}(O, \mathcal{P}(O))$ is $(d-1)$-spherical.

After restricting to these conical subgroups (Step 1 of our induction), we are left with the following base cases:

1. The left-most kernel $\mathrm{Out}^0(A_\Gamma;\ \{A_{\geq v} \mid v \in V(\Gamma)\}^t)$; here, this is generated by all the partial conjugations of O;

2. for all $1 \leq i \leq d-1$: the group $\mathrm{Out}(\langle c_i \rangle) \cong \mathbb{Z}/2\mathbb{Z}$;

3. $\mathrm{Out}^0(\langle c_0, c_1, a_1, b_1 \rangle\,;\ \{\langle c_1 \rangle\,, \langle a_1, b_1 \rangle\}^t)$
$$= \mathrm{Out}^0(\langle c_0, c_1, a_1, b_1 \rangle\,;\ \{\langle c_1, a_1, b_1 \rangle\}^t);$$

4. $\mathrm{Out}^0(\langle c_d, c_{d-1}, a_d, b_d \rangle\,;\ \{\langle c_{d-1} \rangle\,, \langle a_d, b_d \rangle\}^t)$
$$= \mathrm{Out}^0(\langle c_d, c_{d-1}, a_d, b_d \rangle\,;\ \{\langle c_{d-1}, a_d, b_d \rangle\}^t);$$

5. for all $1 \leq i \leq d$: the group $\mathrm{Out}(\langle a_i, b_i \rangle) \cong \mathrm{Out}(F_2)$.

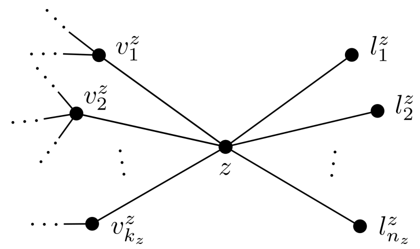

Figure 5.6: $\mathrm{st}(z) = \Gamma_{\geq l_i^z}$.

Only the groups of the last item have a non-empty set of parabolic subgroups (each given by the singleton $\{\mathrm{Stab}_{\mathrm{Out}(\langle a_i,b_i\rangle)}(\langle a_i\rangle)\}$). All items but the first one describe Fouxe-Rabinovitch groups, so the induction already ends here and we do not have to apply Step 2 and Step 3.

Trees

Let Γ be a tree, define $O := \mathrm{Out}^0(A_\Gamma)$ and, to simplify notation, assume that $|V(\Gamma)| \geq 3$. Let L denote the set of leaves of Γ. For each leaf l, let z_l denote the (unique) vertex adjacent to l and let $Z = \{z_l \mid l \in L\}$ be the set of vertices of Γ that are adjacent to some leaf. Then we have

$$[v] = \begin{cases} \{v\} & , v \in V(\Gamma) \setminus L; \\ \mathrm{lk}(z_v) \cap L & , v \in L. \end{cases}$$

The conical subgroups are given by

$$\Gamma_{\geq v} = \{v\}, \, v \in V(\Gamma) \setminus L \qquad \text{and} \qquad \Gamma_{\geq l} = \mathrm{st}(z_l), \, l \in L.$$

Now for each $z \in Z$, let $\{v_1^z, \ldots, v_{k_z}^z\} = \mathrm{lk}(z) \setminus L$ be the non-leaf vertices adjacent to z and $\{l_1^z, \ldots, l_{n_z}^z\} = \mathrm{lk}(z) \cap L$ the leaves adjacent to z (see Fig. 5.6). Then

$$\mathcal{P}(O) = \bigcup_{z \in Z} \left\{\mathrm{Stab}_O\left(\langle l_1^z, \ldots, l_i^z, v_1^z, \ldots, v_{k_z}^z, z\rangle\right) \,\middle|\, 1 \leq i \leq n_z - 1\right\}$$

and $\mathrm{rk}(O) = \sum_{z \in Z}(n_z - 1) = |L| - |Z|$, which implies that $\mathrm{CC}(O, \mathcal{P}(O))$ is $(|L| - |Z| - 1)$-spherical.

Step 1 of our induction leads to the following base cases:

1. The left-most kernel $\mathrm{Out}^0(A_\Gamma; \{A_{\geq v} \mid v \in V(\Gamma)\}^t)$ generated by all the partial conjugations of O acting trivially on $\mathrm{st}(z)$ for all $z \in Z$;
2. for all $v \in V(\Gamma) \setminus L$: the group $\mathrm{Out}(\langle v\rangle) \cong \mathbb{Z}/2\mathbb{Z}$;
3. for all $z \in Z$: the group $\mathrm{Out}^0(A_{\mathrm{st}(z)}; \{\langle v_1^z\rangle, \ldots, \langle v_{k_z}^z\rangle, \langle z\rangle\}^t)$.

In Step 2, we apply for each group of the last item the projection map P_Δ where $\Delta = \mathrm{lk}(z)$. Its kernel is a free abelian group of rank n_z, generated by the twists of the leaves adjacent to z. The image of this projection map is the Fouxe-Rabinovitch group

$$\mathrm{Out}^0(A_{\mathrm{lk}(z)}; \{\langle v_1^z\rangle, \ldots, \langle v_{k_z}^z\rangle\}^t).$$

It contains $n_z - 1$ maximal parabolic subgroups, given by the stabilisers of $\langle l_1^z, \dots, l_i^z, v_1^z, \dots, v_{k_z}^z \rangle$, $1 \leq i \leq n_z - 1$. Again, we do not have to apply Step 3.

Constructions

This subsection does not really contain examples but rather shows how to obtain new examples from known ones.

Direct product Let Γ_1 and Γ_2 be graphs and let $\Gamma := \Gamma_1 * \Gamma_2$ be their join. On the level of RAAGs, this means that $A_\Gamma = A_{\Gamma_1} \times A_{\Gamma_2}$. Let $\leq, \leq_1, \leq_2$ be the standard orderings and $[\cdot], [\cdot]_1, [\cdot]_2$ the corresponding equivalence classes of $\Gamma, \Gamma_1, \Gamma_2$, respectively. Let Z, Z_1, Z_2 be the (possibly trivial) subgraphs of $\Gamma, \Gamma_1, \Gamma_2$ consisting of all vertices that are adjacent to every vertex of the corresponding graph; clearly, $Z = Z_1 * Z_2$. It is easy to see that for $v_i \in \Gamma_i$, one has

$$[v_i] = \begin{cases} Z & , v_i \in Z_i; \\ [v_i]_i & , v_i \notin Z_i; \end{cases} \qquad \text{and} \qquad \Gamma_{\geq v_i} = (\Gamma_i)_{\geq_i v_i} \cup Z_j.$$

We do not spell out the consequences for all of the induction, but would like to point out the following implication for the ranks of the corresponding automorphism groups and hence the dimensions of the associated coset complexes:

$$\mathrm{rk}(\mathrm{Out}^0(A_\Gamma)) = \begin{cases} \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_1})) + \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_2})) & , Z_i = \emptyset \text{ for some } i; \\ \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_1})) + \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_2})) + 1 & , \text{otherwise.} \end{cases}$$

Note that $Z_i = \emptyset$ is equivalent to saying that the centre $Z(A_{\Gamma_i})$ is trivial.

Free product Let $\Gamma := \Gamma_1 \sqcup \Gamma_2$ be the disjoint union of the graphs Γ_1 and Γ_2, i.e. $A_\Gamma = A_{\Gamma_1} * A_{\Gamma_2}$, and keep the notation of the prior paragraph otherwise. Let D, D_1, D_2 be the (possibly trivial) subgraphs of $\Gamma, \Gamma_1, \Gamma_2$ consisting of all their isolated vertices; we have $D = D_1 \sqcup D_2$. For $v_i \in \Gamma_i$, one has:

$$[v_i] = \begin{cases} D & , v_i \in D_i; \\ [v_i]_i & , v_i \notin D_i; \end{cases} \qquad \text{and} \qquad \Gamma_{\geq v_i} = \begin{cases} \Gamma & , v_i \in D_i; \\ (\Gamma_i)_{\geq_i v_i} & , v_i \notin D_i. \end{cases}$$

Similar to the case of direct products, this implies:

$$\mathrm{rk}(\mathrm{Out}^0(A_\Gamma)) = \begin{cases} \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_1})) + \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_2})) & , D_i = \emptyset \text{ for some } i; \\ \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_1})) + \mathrm{rk}(\mathrm{Out}^0(A_{\Gamma_2})) + 1 & , \text{otherwise.} \end{cases}$$

This allows for example to generalise the example of tree-RAAGs given above to the setting of forests.

Complement graph For a graph Γ, let Γ^c denote its complement, i.e. the graph with vertex set $V(\Gamma)$ where v and w form an edge if and only if they do not form an edge in Γ. Let $\leq_c$ and $[\cdot]_c$ denote the standard ordering and its equivalence classes on Γ^c. Then it is easy to see that

$$[v] = [v]_c \qquad \text{and} \qquad v \leq_c w \Leftrightarrow w \leq v.$$

In particular, one has $\mathrm{rk}(\mathrm{Out}^0(A_{\Gamma^c})) = \mathrm{rk}(\mathrm{Out}^0(A_\Gamma))$. This also explains the analogy between the settings of direct and free products considered above.

5.4 Cohen–Macaulayness, higher generation and rank

In this section, we generalise the results of Section 5.2: We show that the coset complex of parabolic subgroups of a relative automorphism group O of a RAAG is not only spherical, but even Cohen–Macaulay. This is used to determine the degree of generation that families of (possibly non-maximal) parabolic subgroups provide. We also give an interpretation of the rank in terms of a "Weyl group" for O.

Notation and standing assumptions As before, let Γ be a graph, $\mathcal{G}$ and $\mathcal{H}$ families of special subgroups of A_Γ such that $\mathcal{G}$ is saturated with respect to $(\mathcal{G}, \mathcal{H})$, define $O := \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$ and set $\preceq := \leq_{\mathcal{G}}$ to be the $\mathcal{G}$-ordering on $V(\Gamma)$. Let $T_{\mathcal{G}}$ denote the set of $\sim_{\mathcal{G}}$-equivalence classes of vertices of Γ.

5.4.1 Cohen–Macaulayness

Recall from Theorem 3.20 that a coset complex $\mathrm{CC}(G, \mathcal{H})$ is homotopy Cohen–Macaulay if and only if every subset $\mathcal{H}' \subseteq \mathcal{H}$ is $(|\mathcal{H}'|-1)$-generating for G. This allows us to generalise Theorem 5.28 in the following way:

Theorem 5.29. *Let $O := \mathrm{Out}^0(A_\Gamma; \mathcal{G}, \mathcal{H}^t)$. The coset complex $\mathrm{CC}(O, \mathcal{P}(O))$ is homotopy Cohen–Macaulay.*

Proof. By Theorem 3.20, it suffices to show that for all $\mathcal{P}' \subseteq \mathcal{P}(O)$, the coset complex $\mathrm{CC}(O, \mathcal{P}')$ is $(|\mathcal{P}'|-1)$-spherical. This can be done following the induction of Section 5.2.2: We first iteratively apply restriction maps to conical subgroups as in Lemma 5.21. In each step, the parabolic subgroups in $\mathcal{P}'$ satisfy a dichotomy that allows us to apply Corollary 3.30. We get an analogue of Proposition 5.22: The coset complex $\mathrm{CC}(O, \mathcal{P}')$ is homotopy equivalent to the join $*_{[v]_{\mathcal{G}} \in T_{\mathcal{G}}} \mathrm{CC}(O_v, \mathcal{P}_v)$ where O_v is exactly as in Proposition 5.22 and $\mathcal{P}_v \subseteq \mathcal{P}(O_v)$. There is a one-to-one correspondence between the parabolic subgroups occurring in the join and the elements of $\mathcal{P}'$; in particular, $\sum_{[v]_{\mathcal{G}} \in T_{\mathcal{G}}} |\mathcal{P}_v| = |\mathcal{P}'|$. One now follows the arguments of Section 5.2.2 to show that if $[v]_{\mathcal{G}}$ is an abelian equivalence class of size n, we have $\mathrm{CC}(O_v, \mathcal{P}_v) \simeq \mathrm{CC}(\mathrm{GL}_n(\mathbb{Z}), \mathcal{Q})$ with $\mathcal{Q} \subseteq \mathcal{P}(\mathrm{GL}_n(\mathbb{Z}))$ and that if $[v]_{\mathcal{G}}$ is a free equivalence class, we have

$$\mathrm{CC}(O_v, \mathcal{P}_v) \simeq \mathrm{CC}(\mathrm{Out}(A; \mathcal{A}^t), \mathcal{Q})$$

where $A = F_n * A_1 * \ldots * A_k$, $\mathcal{A} = \{A_1, \ldots, A_k\}$ and $\mathcal{Q} \subseteq \mathcal{P}(\mathrm{Out}(A; \mathcal{A}^t))$. Both $\mathrm{CC}(\mathrm{GL}_n(\mathbb{Z}), \mathcal{P}(\mathrm{GL}_n(\mathbb{Z})))$ and $\mathrm{CC}(\mathrm{Out}(A; \mathcal{A}^t), \mathcal{P}(\mathrm{Out}(A; \mathcal{A}^t)))$ are homotopy Cohen–Macaulay: In the first case, this holds because the coset complex is isomorphic to the building associated to $\mathrm{GL}_n(\mathbb{Q})$ (see Proposition 3.15), in the second case this follows from Proposition 4.7 and Theorem 4.41. Hence, Theorem 3.20 implies that $\mathrm{CC}(O_v, \mathcal{P}_v)$ is spherical of dimension $|\mathcal{P}_v| - 1$. It now follows from Lemma 2.11 that the complex $\mathrm{CC}(O, \mathcal{P}')$ is spherical of dimension $(\sum_{[v]_{\mathcal{G}} \in T_{\mathcal{G}}} |\mathcal{P}_v|) - 1 = |\mathcal{P}'| - 1$. □

An immediate consequence of this is that $\mathrm{CC}(O, \mathcal{P}(O))$ is a chamber complex (see Remark 2.14).

5.4.2 Parabolic subgroups of lower rank

Definition 5.30. Let $r \coloneqq \mathrm{rk}(O)$ and $1 \leq m \leq r-1$. We define the family of *rank-m standard parabolic subgroups of O* as the set of all intersections of $(r-m)$ distinct maximal standard parabolic subgroups,

$$\mathcal{P}_m(O) \coloneqq \{P_1 \cap \ldots \cap P_{r-m} \mid P_1, \ldots, P_{r-m} \text{ distinct elements of } \mathcal{P}(O)\}.$$

In particular, we have $\mathcal{P}(O) = \mathcal{P}_{r-1}(O)$.

Every parabolic subgroup of O is itself a relative automorphism group of A_Γ. The term "rank-m" parabolic subgroup is justified by the following:

Proposition 5.31. *For all $P \in \mathcal{P}_m(O)$, we have $\mathrm{rk}(P) = m$.*

Proof. For every $P \in \mathcal{P}_m(O)$, there is a $V \subset V(\Gamma)$ and for every $v \in V$ a subset $J_v \subset \{1, \ldots, |[v]_\mathcal{G}|\}$ such that

$$P = \mathrm{Out}^0(A_\Gamma; \mathcal{G}', \mathcal{H}^t), \text{ where } \mathcal{G}' = \mathcal{G} \cup \left\{ A_{\Delta_v^j} \,\middle|\, v \in V,\, j \in J_v \right\},$$
$$\Delta_v^j = \{v_1, \ldots, v_j\} \cup \Gamma_{\succ v} \text{ and } \sum_{v \in V} |J_v| = r - m.$$

As $\mathcal{G}$ contains $P(\mathcal{H})$, so does $\mathcal{G}'$. It is easy to check that if $v \in V$, the $\mathcal{G}$-equivalence class $[v]_\mathcal{G}$ can be written as the disjoint union of $(|J_v|+1)$-many $\mathcal{G}'$-equivalence classes and that otherwise, one has $[v]_\mathcal{G} = [v]_{\mathcal{G}'}$. From this, the claim follows immediately. □

The following is an easy corollary of Theorem 5.29 and the results of Chapter 3:

Corollary 5.32. *The family $\mathcal{P}_m(O)$ of rank-m parabolic subgroups of O is m-generating, the corresponding coset complex $\mathrm{CC}(O, \mathcal{P}_m(O))$ is m-spherical.*

Proof. As the coset complex is homotopy Cohen–Macaulay by Theorem 5.29, this is an immediate consequence of Theorem 3.11. □

In the case where $O = \mathrm{GL}_n(\mathbb{Z})$, this is a special case of Theorem 3.13. Observe that although $\mathrm{CC}(O, \mathcal{P}_m(O))$ is m-spherical, it is a priori a complex of dimension

$$|\mathcal{P}_m(O)| - 1 = \binom{r}{m} - 1.$$

Presentations for O A consequence of higher generation is that one can obtain presentations of O from presentations of the parabolic subgroups as follows: Write $\mathcal{P}_m(O) = \left\{P_1, \ldots, P_{\binom{r}{m}}\right\}$. For each i, let L_i be the set of all inversions, transvections and partial conjugations of O that are contained in P_i. By Theorem 5.2, the set L_i generates P_i. Let $P_i = \langle L_i \mid R_i \rangle$ be a presentation for P_i. Then we have:

Corollary 5.33. *Let $1 \leq m \leq r-1$ and $k \coloneqq \binom{r}{m}$. Then:*

1. *The union $\bigcup_{i=1}^k L_i$ is a generating set for O.*
2. *If $m \geq 2$, a presentation for O is given by $O = \left\langle \bigcup_{i=1}^k L_i \,\middle|\, \bigcup_{i=1}^k R_i \right\rangle$.*

Proof. This follows from Corollary 5.32 and Theorem 3.8. □

5.4.3 Interpretation of rank in terms of Coxeter groups

The rank of a group with a BN-pair is given by the rank of the associated Weyl group W, which is a Coxeter group (see Section 3.2.1). This is also true in the setting of relative automorphism groups of RAAGs as we will see in what follows.

Definition 5.34. Let $\mathrm{Aut}(\Gamma)$ denote the group of graph automorphisms of Γ. This group embeds in $\mathrm{Out}(A_\Gamma)$ and we define $\mathrm{Aut}^0(\Gamma)$ as the intersection $\mathrm{Aut}(\Gamma) \cap O$.

If O is equal to $\mathrm{Out}(F_n)$ or $\mathrm{GL}_n(\mathbb{Z})$, we have $\mathrm{Aut}^0(\Gamma) = \mathrm{Aut}(\Gamma) = \mathrm{Sym}(n)$, the Weyl group associated to $\mathrm{GL}_n(\mathbb{Q})$, which has rank $n-1$. In general, $\mathrm{Aut}^0(\Gamma)$ can be seen as the group of "algebraic" graph automorphisms of Γ. It appears as "$\mathrm{Sym}^0(\Gamma)$" in [CCV07, Section 3.2] where it is studied under the additional assumption that Γ is connected and triangle-free.

Lemma 5.35. *The group* $\mathrm{Aut}^0(\Gamma)$ *is naturally isomorphic to the direct product*

$$\mathrm{Aut}^0(\Gamma) \cong \bigoplus_{[v]_\mathcal{G} \in T_\mathcal{G}} \mathrm{Sym}([v]_\mathcal{G}).$$

Proof. If $|[v]_\mathcal{G}| > 1$, the group O contains for all $x, y \in [v]_\mathcal{G}$ the transvection ρ_x^y and the inversion ι_y. It follows that the full group $\mathrm{Sym}([v]_\mathcal{G})$ of permutations of $[v]_\mathcal{G}$ is contained in $\mathrm{Aut}^0(\Gamma)$, so the direct product $\bigoplus_{[v]_\mathcal{G} \in T_\mathcal{G}} \mathrm{Sym}([v]_\mathcal{G})$ is a subgroup of $\mathrm{Aut}^0(\Gamma)$.

It remains to show that this group does not contain any other elements, i.e. that every element of $\mathrm{Aut}^0(\Gamma)$ preserves all the $\mathcal{G}$-equivalence classes of $V(\Gamma)$. To see this, assume that $\phi \in \mathrm{Aut}(A_\Gamma)$ represents an element of O such that $\phi(v) = v'$ for some $v, v' \in V(\Gamma)$. We will show that $v \sim_\mathcal{G} v'$:

For $x \in V(\Gamma)$ and a word w in the alphabet $V(\Gamma)^{\pm 1}$, let $\#_x(w) \in \mathbb{Z}$ denote the number of occurrences of x in w, counted with sign. For $g \in A_\Gamma$, let $\#_x(g) \in \mathbb{Z}/2\mathbb{Z}$ be the image of $\#_x(w)$ in $\mathbb{Z}/2\mathbb{Z}$, where w is a word representing g—this number only depends on g and not on the chosen representative w. Now assume that $v \neq v'$. Then clearly, $\#_{v'}(v) = 0$ and $\#_{v'}(\phi(v)) = \#_{v'}(v') = 1$. Writing ϕ as a product of inversions, transvections and partial conjugations, it follows that there must be such a Laurence generator $[\psi] \in O$ with $\#_{v'}(\psi(v)) = 1$. This is only possible if ψ is given by the transvection $\rho_v^{v'}$. However, if this is contained in O, we know that $v \preceq v'$. As ϕ^{-1} sends v' to v, we also have $v' \preceq v$, hence $v \sim_\mathcal{G} v'$. □

Recall that the *rank of a Coxeter system* (W, S) is given by $\mathrm{rk}(W, S) = |S|$.

Corollary 5.36. *There is a subset* $S \subset \mathrm{Aut}^0(\Gamma) \leq O$ *such that* $(\mathrm{Aut}^0(\Gamma), S)$ *is a Coxeter system of rank equal to* $\mathrm{rk}(O)$.

Proof. The symmetric group on a set of n elements is the Coxeter group of type $\mathtt{A}_{n-1}$, so the claim follows from Lemma 5.35. □

Additional comments on this can be found in the "BN-pairs" paragraph of Section 5.5.

5.5 Closing comments and open questions

We conclude with comments on the limitations of our constructions and on open questions related to the complex $\mathcal{CC} = \mathrm{CC}(O, \mathcal{P}(O))$.

Description as a subgroup poset in A_Γ Both in the setting of $\mathrm{GL}_n(\mathbb{Z})$ and of Fouxe-Rabinovitch groups $\mathrm{Out}(A; \mathcal{A}^t)$, we studied the coset complex of parabolic subgroups by finding an isomorphic poset of subgroups of A_Γ and then determined its homotopy type. These were the poset of direct summands of $\mathbb{Z}^n$ and the relative free factor complex $\mathcal{F}(A, \mathcal{A})$, respectively. In general, however, the author is not aware of a natural description of $\mathcal{CC}$ which looks similar.

Limitations of our construction It seems that our definition of parabolic subgroups and the corresponding coset complex capture well the aspects of $\mathrm{Out}(A_\Gamma)$ that come from similarities of this group with $\mathrm{GL}_n(\mathbb{Z})$ and $\mathrm{Out}(F_n)$: Firstly, our definitions recover the Tits building as $\mathrm{CC}(\mathrm{GL}_n(\mathbb{Z}), \mathcal{P}(\mathrm{GL}_n(\mathbb{Z})))$ and the free factor complex as $\mathrm{CC}(\mathrm{Out}(F_n), \mathcal{P}(\mathrm{Out}(F_n)))$. Secondly, the results we obtain show strong similarities in behaviour between the general situation of $\mathrm{Out}(A_\Gamma)$ and these special cases: The associated coset complex is spherical, even Cohen–Macaulay (Theorem 5.29) and families of parabolic subgroups are highly generating with the degree of generation depending on the rank of these subgroups (Corollary 5.32). Another strong indication which suggests a certain optimality of our definitions is the description of $\mathrm{rk}(\mathrm{Out}^0(A_\Gamma))$ in terms of a Coxeter subgroup (Corollary 5.36). Furthermore, our induction leads to well-suited families of parabolic subgroups in all those "components" of $\mathrm{Out}^0(A_\Gamma)$ that closely resemble general linear groups and automorphism groups of free groups; i.e. the base cases that are given by $\mathrm{GL}_n(\mathbb{Z})$, $n \geq 2$, and Fouxe-Rabinovitch groups containing transvections (Item 2b and Item 3c in Section 5.3.1).

However, our construction is rather transvection-based in the sense that the standard ordering of $V(\Gamma)$—which is used to define the parabolic subgroups—is entirely determined by the transvections that $\mathrm{Out}(A_\Gamma)$ contains. This makes our definition of parabolic subgroups quite *local*: Whether or not $v \leq w$ can be read off from the one-balls around these vertices. This is also reflected by the fact that the conical subgraphs $\Gamma_{\geq v}$, which play a central role in our induction, are contained in the two-ball around v if v is not an isolated vertex. In contrast, certain aspects of $\mathrm{Out}(A_\Gamma)$ seem not to be mere generalisations of phenomena in arithmetic groups and automorphism groups of free groups. For example, $\mathrm{Out}(A_\Gamma)$ contains partial conjugations which cannot be written as a product of transvections. The existence of these partial conjugations is a *global* phenomenon in the sense that the shape of the connected components of $\Gamma \setminus \mathrm{st}(v)$ is not determined by local conditions on Γ. These aspects are not very well represented in $\mathcal{CC}$: The base cases of our induction that correspond to them do not contain any parabolic subgroups. In the extremal case where there is no equivalence class of $V(\Gamma)$ that has size greater than one, $\mathcal{P}(\mathrm{Out}^0(A_\Gamma))$ is even empty. For specific applications, one might try to overcome this by introducing further parabolic subgroups that capture these global aspects. However, the author currently does not see a canonical way to do this.

BN-pairs The existence of a "Weyl group" $\mathrm{Aut}^0(\Gamma)$ as described in Section 5.4.3 suggests that one might be able to transfer additional notions from the theory of groups with BN-pair to automorphism groups of RAAGs. It does for instance seem reasonable to define a "Borel-subgroup" by taking the intersection of all standard parabolic subgroups or to use the Weyl group to define apartments in $\mathcal{CC}$. For this, it might be helpful to use the *standard representation* $\mathrm{Out}(A_\Gamma) \to \mathrm{GL}_{|V(\Gamma)|}(\mathbb{Z})$ induced by the abelianisation. The question that has yet to be clarified is to what extent this point of view might be fruitful for studying automorphism groups of RAAGs; one should keep in mind that all this can also be done for $O = \mathrm{Out}(F_n)$ which is far away from having a BN-pair.

Curve complex The free factor complex was introduced as an analogue of the more classical Tits building in the setting of $\mathrm{Out}(F_n)$. The same is true for the curve complex $\mathcal{C}(S)$ associated to a surface S—it is an analogue of a Tits building in the setting of mapping class groups (see the remarks in Section 4.7, p. 64). In this sense, $\mathcal{CC}$ can also be seen as an $\mathrm{Out}(A_\Gamma)$-analogue of $\mathcal{C}(S)$.

Boundary structures As explained in Section 4.7, both buildings and free factor complexes can be seen as boundary structures of classifying spaces—in the first case, this classifying space is the associated symmetric space; in the second case, it is the Outer space defined by Culler–Vogtmann. In the RAAG-setting, one may ask whether a similar statement holds and $\mathcal{CC}$ can be seen as a boundary structure of the RAAG Outer space defined in [CSV17] or a similar space.

Geometric aspects This text focuses on the topology of $\mathcal{CC}$. It also seems very reasonable, however, to ask what can be said about the geometry of this complex. Motivated by the work of Masur and Minsky [MM99], who showed that the curve complex $\mathcal{C}(S)$ is hyperbolic, Bestvina and Feighn [BF14] proved that the free factor complex is hyperbolic as well. This is only one of many results in the study of $\mathrm{Out}(F_n)$ from a geometric point of view, which has become popular in recent years. On the other hand, there is also a rich theory concerning metric aspects of buildings (for an overview, see [AB08, Section 12]). Combining these two theories should be an interesting topic for further investigations.

Appendix A

Appendix: Graph posets

The following example illustrates Lemma 4.44 in the case where $n = 2$.

Example A.1. For $n = 2$ and $G \in L$, the preimage $p_1^{-1}(G)$ is either contractible or a wedge of 0-spheres. Suppose G is a theta graph. Then a 0-sphere s_i in $p_1^{-1}(G)$ is isomorphic to $\{G\} \times X(G_i)$, where G_i is a rose obtained from G by collapsing a maximal forest, for $i = 1, 2, 3$. We claim that each such 0-sphere is contractible in $\mathcal{Z}$. Indeed, for each rose G_i consider the dumbbell graph G_i' obtained by blowing up G_i to have a separating edge. Then $p_1^{-1}(G_i')$ is contractible. Now in $\mathcal{Z}$, the sphere s_i can be homotoped into $p_1^{-1}(G_i')$. Thus each s_i is contractible in $\mathcal{Z}$. See Fig. A.1.

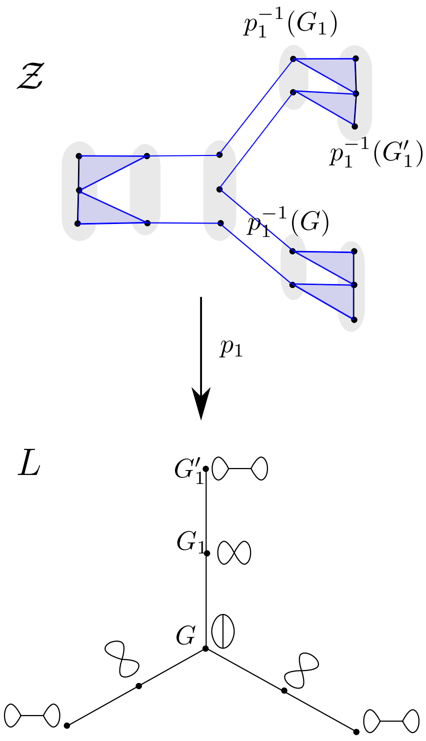

Figure A.1: Projection map $p_1 \colon \mathcal{Z} \to L$ for Example A.1.

Proof of Proposition 4.14 for $n = 3$. Fig. A.2 shows all possible combinatorial types of graphs in $\mathcal{CV}_3$. We want to show that for each such graph G, the poset $\mathrm{X}(G)$ is homotopy equivalent to a (non-trivial) wedge of circles if G does not contain a separating edge and is contractible otherwise. Using Lemma 4.13, it suffices to show the same statement for the poset $\mathrm{C}(G)$ of all core subgraphs.

If G is a rose, the realisation of $\mathrm{X}(G) = \mathrm{Sub}(G)$ is the boundary of a triangulated 1-sphere. For the graphs b) – e) in Figure A.2, the complex $\mathrm{X}(G)$ is depicted in Figure A.3.

As the graphs f), g) and h) do not contain any disconnected core subgraphs, the claim here follows from Proposition 4.19. The only disconnected core subgraph of i) consists of the edges $1, 4$ and 5. Hence, $\mathrm{C}(G)$ is derived from $\mathrm{cC}(G)$ by attaching the star of the vertex $\{1, 4, 5\}$ along its link. It is an easy exercise to check that the result is homotopy equivalent to a circle. The same is true for j) whose only non-connected core subgraph is $\{1, 2, 4, 6\}$.

For the remaining graphs k) – p), the following tables define Morse functions $\phi : \mathrm{C}(G) \to \mathbb{R}$ with contractible descending links:

k)

vertex v	$\phi(v)$
$\mathrm{st}(\{1, 2, 3, 4\})$	0
$\mathrm{cC}(G) \setminus \mathrm{st}(\{1, 2, 3, 4\}) = \{\{1, 2, 3, 5\}, \{1, 2, 4, 5\}, \{1, 3, 4, 5\}\}$	1

l)

vertex v	$\phi(v)$
$\mathrm{st}(\{1, 2, 3, 4\})$	0
$\mathrm{cC}(G) \setminus \mathrm{st}(\{1, 2, 3, 4\})$	1

m)

vertex v	$\phi(v)$
$\mathrm{st}(\{1, 2, 3, 4, 5\})$	0
$\mathrm{cC}(G) \setminus \mathrm{st}(\{1, 2, 3, 4, 5\})$	1

n)

vertex v	$\phi(v)$
$\mathrm{st}(\{1, 2, 3, 4\})$	0
$\{1, 2, 3, 5\}$	1
$\{2, 3, 5\}, \{1, 3, 4, 5\}$	2

o) & p)

vertex v	$\phi(v)$
$\mathrm{st}(\{1, 2, 3, 4, 5\})$	0
$\{1, 2, 3, 4, 6\}$	1
all other core subgraphs containing 6	2

As an illustration, we explain why all the descending links for n) are contractible: $\mathrm{st}(\{1, 2, 3, 4\})$ is obviously contractible as this is true for any star in a simplicial complex. The vertices of $\mathrm{C}(G)$ not contained in this star are precisely the proper core subgraphs of G containing the (separating) edge 5. The descending link of $\{1, 2, 3, 5\}$ contains a unique maximal element and hence is contractible; this cone point is given by $\{1, 2, 3\}$ which is the unique maximal core subgraph of $\{1, 2, 3, 5\}$ not containing 5. As $\{2, 3, 5\}$ does not contain 4, it is contained in $\{1, 2, 3, 5\}$ which hence forms a cone point of its descending link. Lastly, the link of $\{1, 3, 4, 5\}$ is coned off by $\{1, 3\}$.

The interested reader may complete this argument to an alternative proof of Proposition 4.14 for arbitrary $n \geq 2$ in the case where G contains at least one separating edge. □

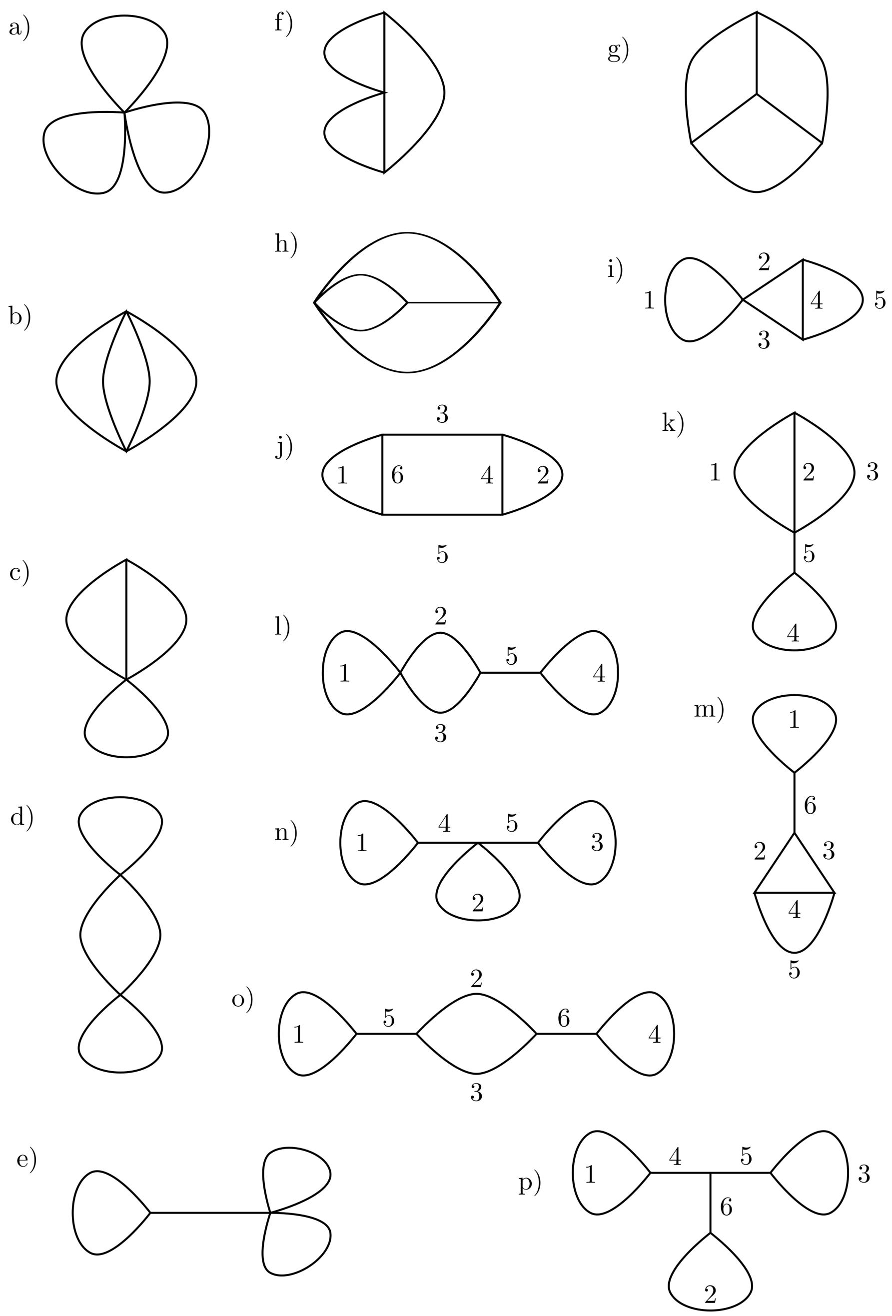

Figure A.2: All combinatorial types of graphs in $\mathcal{CV}_3$.

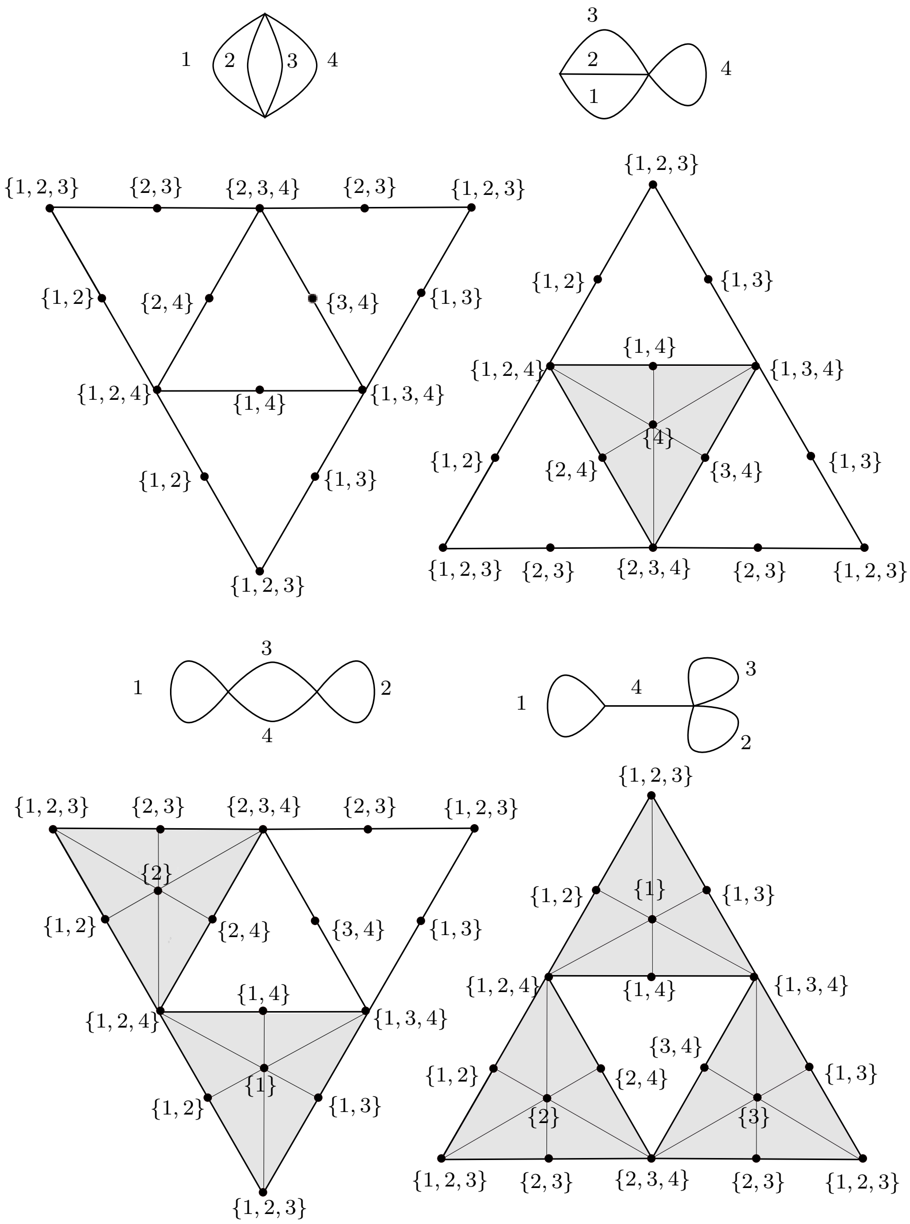

Figure A.3: The realisation of X(G) for rank three graphs with four edges (the tetrahedra have been unfolded for better visibility). The first three are homotopy equivalent to a wedge of circles while the last is contractible.

Bibliography

[AB08] Peter Abramenko and Kenneth S. Brown. *Buildings*, volume 248 of *Graduate Texts in Mathematics*. Springer, New York, 2008.

[Ago13] Ian Agol. The virtual Haken conjecture. *Documenta Mathematica*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[AH93] Herbert Abels and Stephan Holz. Higher generation by subgroups. *Journal of Algebra*, 160(2):310–341, 1993.

[BF00] Mladen Bestvina and Mark Feighn. The topology at infinity of $\mathrm{Out}(F_n)$. *Inventiones Mathematicae*, 140(3):651–692, 2000.

[BF14] Mladen Bestvina and Mark Feighn. Hyperbolicity of the complex of free factors. *Advances in Mathematics*, 256:104–155, 2014.

[BFM+16] Kai-Uwe Bux, Martin G. Fluch, Marco Marschler, Stefan Witzel, and Matthew C. B. Zaremsky. The braided Thompson's groups are of type F_∞. *Journal für die Reine und Angewandte Mathematik. [Crelle's Journal]*, 718:59–101, 2016. With an appendix by Zaremsky.

[BG] Benjamin Brück and Radhika Gupta. Homotopy type of the complex of free factors of a free group. http://arxiv.org/abs/1810.09380v2.

[BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Annals of Mathematics. Second Series*, 135(1):1–51, 1992.

[Bjö84] Anders Björner. Some combinatorial and algebraic properties of Coxeter complexes and Tits buildings. *Advances in Mathematics*, 52(3):173–212, 1984.

[Bjö95] Anders Björner. Topological methods. In *Handbook of combinatorics, Vol. 1, 2*, pages 1819–1872. Elsevier Sci. B. V., Amsterdam, 1995.

[Bog08] Oleg Bogopolski. *Introduction to group theory.* EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2008. Translated, revised and expanded from the 2002 Russian original.

[Bre72] Glen E. Bredon. *Introduction to compact transformation groups*. Academic Press, New York-London, 1972. Pure and Applied Mathematics, Vol. 46.

[Bro00] Kenneth S. Brown. The coset poset and probabilistic zeta function of a finite group. *Journal of Algebra*, 225(2):989–1012, 2000.

[Brüa] Benjamin Brück. Between buildings and free factor complexes: A Cohen–Macaulay complex for Out(RAAGs). http://arxiv.org/abs/1906.05606v1.

[Brüb] Benjamin Brück. Higher generating subgroups and Cohen–Macaulay complexes. https://arxiv.org/abs/1901.04204v3. To appear in the Proceedings of the Edinburgh Mathematical Society.

[BS73] A. Borel and J.-P. Serre. Corners and arithmetic groups. *Commentarii Mathematici Helvetici*, 48:436–491, 1973. Avec un appendice: Arrondissement des variétés à coins, par A. Douady et L. Hérault.

[BSV18] Kai-Uwe Bux, Peter Smillie, and Karen Vogtmann. On the bordification of outer space. *Journal of the London Mathematical Society. Second Series*, 98(1):12–34, 2018.

[BWW09] Anders Björner, Michelle Wachs, and Volkmar Welker. On sequentially Cohen-Macaulay complexes and posets. *Israel Journal of Mathematics*, 169:295–316, 2009.

[CCV07] Ruth Charney, John Crisp, and Karen Vogtmann. Automorphisms of 2-dimensional right-angled Artin groups. *Geometry & Topology*, 11:2227–2264, 2007.

[CD95] Ruth Charney and Michael W. Davis. The $K(\pi, 1)$-problem for hyperplane complements associated to infinite reflection groups. *Journal of the American Mathematical Society*, 8(3):597–627, 1995.

[CSV17] Ruth Charney, Nathaniel Stambaugh, and Karen Vogtmann. Outer space for untwisted automorphisms of right-angled Artin groups. *Geometry & Topology*, 21(2):1131–1178, 2017.

[CV86] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones Mathematicae*, 84(1):91–119, 1986.

[CV09] Ruth Charney and Karen Vogtmann. Finiteness properties of automorphism groups of right-angled Artin groups. *Bulletin of the London Mathematical Society*, 41(1):94–102, 2009.

[DR12] Andrew J. Duncan and Vladimir N. Remeslennikov. Automorphisms of partially commutative groups II: Combinatorial subgroups. *International Journal of Algebra and Computation*, 22(7):1250074, 44, 2012.

[DW] Matthew B. Day and Richard D. Wade. Relative automorphism groups of right-angled Artin groups. http://arxiv.org/abs/1712.01583v3. To appear in the Journal of Topology.

[FR40] D. I. Fouxe-Rabinovitch. Über die Automorphismengruppen der freien Produkte. I. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):265–276, 1940.

[GH] Vincent Guirardel and Camille Horbez. Algebraic laminations for free products and arational trees. http://arxiv.org/abs/1709.05664v2. To appear in Algebraic & Geometric Topology.

[GL07] Vincent Guirardel and Gilbert Levitt. The outer space of a free product. *Proceedings of the London Mathematical Society. Third Series*, 94(3):695–714, 2007.

[GL17] Vincent Guirardel and Gilbert Levitt. JSJ decompositions of groups. *Astérisque*, (395):vii+165, 2017.

[Gri11] James T. Griffin. *Automorphisms of free products of groups.* PhD thesis, University of Cambridge, 2011.

[GS18] Vincent Guirardel and Andrew Sale. Vastness properties of automorphism groups of RAAGs. *Journal of Topology*, 11(1):30–64, 2018.

[GV09] Andrei Gabrielov and Nicolai Vorobjov. Approximation of definable sets by compact families, and upper bounds on homotopy and homology. *Journal of the London Mathematical Society. Second Series*, 80(1):35–54, 2009.

[Har81] W. J. Harvey. Boundary structure of the modular group. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume 97 of *Ann. of Math. Stud.*, pages 245–251. Princeton Univ. Press, Princeton, N.J., 1981.

[Har86] John L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Inventiones Mathematicae*, 84(1):157–176, 1986.

[Hat95] Allen Hatcher. Homological stability for automorphism groups of free groups. *Commentarii Mathematici Helvetici*, 70(1):39–62, 1995.

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[HH17] Arnaud Hilion and Camille Horbez. The hyperbolicity of the sphere complex via surgery paths. *Journal für die Reine und Angewandte Mathematik. [Crelle's Journal]*, 730:135–161, 2017.

[HK18] Sebastian Hensel and Dawid Kielak. Nielsen realisation for untwisted automorphisms of right-angled Artin groups. *Proceedings of the London Mathematical Society. Third Series*, 117(5):901–950, 2018.

[HM] Michael Handel and Lee Mosher. Relative free splitting and free factor complexes I: Hyperbolicity. http://arxiv.org/abs/1407.3508v2.

[Hol85] Stephan Holz. *Endliche Identifizierbarkeit von Gruppen*. PhD thesis, Universität Bielefeld, 1985.

[HV98a] Allen Hatcher and Karen Vogtmann. Cerf theory for graphs. *Journal of the London Mathematical Society. Second Series*, 58(3):633–655, 1998.

[HV98b] Allen Hatcher and Karen Vogtmann. The complex of free factors of a free group. *The Quarterly Journal of Mathematics. Oxford. Second Series*, 49(196):459–468, 1998.

[HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, 2008.

[HW12] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. *Annals of Mathematics. Second Series*, 176(3):1427–1482, 2012.

[Iva87] N. V. Ivanov. Complexes of curves and Teichmüller modular groups. *Akademiya Nauk SSSR i Moskovskoe Matematicheskoe Obshchestvo. Uspekhi Matematicheskikh Nauk*, 42(3(255)):49–91, 255, 1987.

[Lau95] Michael R. Laurence. A generating set for the automorphism group of a graph group. *Journal of the London Mathematical Society. Second Series*, 52(2):318–334, 1995.

[Lut99] Frank Hagen Lutz. *Triangulated Manifolds with Few Vertices and Vertex-Transitive Group Actions*. PhD thesis, Technische Universität Berlin, 1999.

[MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Inventiones Mathematicae*, 138(1):103–149, 1999.

[MMV98] John Meier, Holger Meinert, and Leonard VanWyk. Higher generation subgroup sets and the Σ-invariants of graph groups. *Commentarii Mathematici Helvetici*, 73(1):22–44, 1998.

[Nie17] J. Nielsen. Die Isomorphismen der allgemeinen, unendlichen Gruppe mit zwei Erzeugenden. *Mathematische Annalen*, 78(1):385–397, 1917.

[Nie24] Jakob Nielsen. Die Isomorphismengruppe der freien Gruppen. *Mathematische Annalen*, 91(3-4):169–209, 1924.

[Qui73] Daniel Quillen. Higher algebraic K-theory. I. pages 85–147. Lecture Notes in Math., Vol. 341, 1973.

[Qui78] Daniel Quillen. Homotopy properties of the poset of nontrivial p-subgroups of a group. *Advances in Mathematics*, 28(2):101–128, 1978.

[Ram05] Daniel A. Ramras. Connectivity of the coset poset and the subgroup poset of a group. *Journal of Group Theory*, 8(6):719–746, 2005.

[Seg93] Yoav Segev. Group actions on finite acyclic simplicial complexes. *Israel Journal of Mathematics*, 82(1-3):381–394, 1993.

[Ser89] Herman Servatius. Automorphisms of graph groups. *Journal of Algebra*, 126(1):34–60, 1989.

[Ser03] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[Sol69] Louis Solomon. The Steinberg character of a finite group with BN-pair. In *Theory of Finite Groups (Symposium, Harvard Univ., Cambridge, Mass., 1968)*, pages 213–221. Benjamin, New York, 1969.

[Spa66] Edwin H. Spanier. *Algebraic topology*. McGraw-Hill Book Co., New York-Toronto, Ont.-London, 1966.

[SR] Yuri Santos Rego. On the finiteness length of some soluble linear groups. http://arxiv.org/abs/1901.06704v1.

[Sta82] Richard P. Stanley. Some aspects of groups acting on finite posets. *Journal of Combinatorial Theory. Series A*, 32(2):132–161, 1982.

[Sta96] Richard P. Stanley. *Combinatorics and Commutative Algebra*, volume 41 of *Progress in Mathematics*. Springer, 1996.

[SW16] John Shareshian and Russ Woodroofe. Order complexes of coset posets of finite groups are not contractible. *Advances in Mathematics*, 291:758–773, 2016.

[Tit74] Jacques Tits. *Buildings of spherical type and finite BN-pairs*. Lecture Notes in Mathematics, Vol. 386. Springer-Verlag, Berlin-New York, 1974.

[vH03] Anja von Heydebreck. Homotopy properties of certain complexes associated to spherical buildings. *Israel Journal of Mathematics*, 133:369–379, 2003.

[Vog90] Karen Vogtmann. Local structure of some $\mathrm{Out}(F_n)$-complexes. *Proceedings of the Edinburgh Mathematical Society. Series II*, 33(3):367–379, 1990.

[Vog08] Karen Vogtmann. What is... outer space? *Notices of the American Mathematical Society*, 55(7):784–786, 2008.

[Wac07] Michelle L. Wachs. Poset topology: tools and applications. In *Geometric combinatorics*, volume 13 of *IAS/Park City Math. Ser.*, pages 497–615. Amer. Math. Soc., Providence, RI, 2007.

[Wad14] Richard D. Wade. Folding free-group automorphisms. *The Quarterly Journal of Mathematics*, 65(1):291–304, 2014.

[Wel18] Cora Welsch. *On coset posets, nerve complexes and subgroup graphs of finitely generated groups.* PhD thesis, WWU Münster, 2018.

[Whi36] J. H. C. Whitehead. On Certain Sets of Elements in a Free Group. *Proceedings of the London Mathematical Society. Second Series*, 41(1):48–56, 1936.