

RICE UNIVERSITY

**Automorphisms of nonpositively curved cube
complexes, right-angled Artin groups and
homology**

by

Corey Bregman

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Doctor of Philosophy

APPROVED, THESIS COMMITTEE:



Michael Wolf, Chair
Professor of Mathematics



Andrew Putman, Thesis Director
Professor of Mathematics, Notre Dame
University



Shelly Harvey
Professor of Mathematics



Ilinca Stanciulescu
Associate Professor of Civil Engineering

Houston, Texas

April, 2017

ABSTRACT

Automorphisms of nonpositively curved cube complexes, right-angled Artin groups
and homology

by

Corey Bregman

Recently, the geometry of CAT(0) cube complexes featured prominently in Agol's resolution of two longstanding conjectures of Thurston in low-dimensional topology: the virtually Haken and virtually fibered conjecture for hyperbolic 3-manifolds. A key step of the proof was to show that every hyperbolic 3-manifold group is virtually special, i.e. virtually the fundamental group of a special nonpositively curved (NPC) cube complex. In this thesis, we study algebraic properties of special groups as they relate to the geometry of special cube complexes.

In the first part of the thesis, we introduce a positive integer-valued invariant of special cube complexes called the genus, and show that having genus one is equivalent to having free abelian fundamental group. As a corollary, we obtain a new proof of the fact that every special group is either abelian or surjects onto a non-abelian free group. In the second part of the thesis, we turn our attention to automorphisms of NPC cube complexes. We give a criterion on a special cube complex which implies that any automorphism acts non-trivially on first homology, and show that a non-trivial action on homology can always be achieved by passing to covers. We then apply the criterion to provide a new geometric proof that the Torelli subgroup for a right-angled Artin group is torsion-free.

Acknowledgments

This thesis would of course not have been possible without the help of many people, and it is my great pleasure to thank them now. First, I would like to thank my advisor, Andrew Putman, for his persistent encouragement over the course of my PhD in various forms: to pursue what interests me, to collaborate with others, and strive for clarity in thought and exposition.

I am very grateful to my thesis committee: Mike Wolf, Shelly Harvey, and Ilinca Stanculescu, for many useful comments and their help with the final stages of my PhD. Special thanks is also due to Neil Fullarton for all his help and advice, for being a sounding board, a collaborator and a friend. For a time, we made up the geometric group theory contingent at Rice, and without him I would have had no one to talk to about cube complexes. I was very fortunate to be at Rice at the same time as him.

Junghwan Park and Jorge Acosta, my officemates and friends, were always around to talk about math and for this I am extremely grateful—even when I was mainly the one doing the talking, and all they wanted to do was play chess. My success at Rice would not have been possible without them. I would also like to thank David Cohen, alias Daco, my brooding academic brother, for always listening and offering advice, though I had little to offer in return.

Finally, I will always remember Rice for being far more than just the place where I obtained my PhD; it is also where I met my wife, Letao Zhang. Letao, you are my constant companion, my collaborator and my voice of reason. Thank you for always fighting for me. I will always be indebted to the Rice Math Department and Fulton and Harris' monograph for bringing us together.

Contents

Abstract	ii
Acknowledgments	iii
List of Illustrations	vi
1 Introduction	1
1.1 CAT(0) geometry and δ -hyperbolicity	2
1.2 NPC Cube complexes	3
1.3 Groups with a quasi-convex hierarchy	4
1.4 Right-angled Artin groups	7
2 Statement of Results	10
3 Background	14
3.1 Group theory, δ -hyperbolicity, CAT(0) geometry	14
3.1.1 Group theory preliminaries	14
3.1.2 Coarse geometry preliminaries	16
3.1.3 CAT(0) geometry	20
3.2 NPC cube complexes	23
3.3 Special cube complexes	27
4 The genus of a special cube complex	35
4.1 Collapsing hyperplanes	35
4.2 The cohomology group $H^1(X)$	39
4.3 The genus of a special group	40

5 Automorphisms of NPC cube complexes	54
5.1 Action of automorphisms on homology	54
5.2 Passing to covers	57
6 Applications to right-angled Artin groups	61
6.1 Statement of Results	61
6.2 Automorphisms of Raags	63
6.3 The Out_ℓ -spine K_Γ	64
6.4 The Torelli Subgroup for a Raag	65
6.5 Blow-ups of Salvetti complexes	66
6.6 Automorphisms of blow-ups	70
Bibliography	73

Illustrations

3.1	A schematic of the comparison map between triangle pqr in X and $p'q'r'$ in M_κ . The distance between x and y (the blue line) increases under f_κ	21
3.2	An example of a cube complex.	24
3.3	Hyperplanes in a cube complex, shown in red. Note that a hyperplane through an isolated edge is just its midpoint.	26
3.4	From left to right, the four hyperplane pathologies as listed in Definition 3.16. Image from [1].	29
3.5	An example of a Salvetti complex when Γ is a line segment of length 3. The base vertex is shown in red, and the four generators $a, b, c,$ and d label circles wedged at the base vertex. There are three tori, shown in yellow, blue, and purple, respectively, corresponding to the three edges of Γ and the three pairs of commuting generators of A_Γ	34
4.1	Possible configurations resulting in interosculation in the quotient.	49

To my wife, Letao, and my parents,
for their constant support.

Chapter 1

Introduction

A guiding principle of geometric group theory, essentially going back to Klein's Erlangen program, is that one can understand algebraic properties of finitely generated groups by studying the geometric spaces upon which they act. In the opposite direction, one can often use algebraic properties of groups to either construct or obstruct actions on interesting geometric spaces. The modern formulation of this idea has its roots in the work of Schwartz [2] and Milnor [3] (see Theorem 3.3) on growth of spaces and groups, and is formalized by the notion of a quasi-isometry (Definition 3.4) between metric spaces. Let X be a proper, geodesic metric space and G a discrete group. We will say G acts **geometrically** on X if G acts on X properly discontinuously and cocompactly by isometries.

Any group G with finite generating set S can be viewed as a metric space in a canonical way: the **Cayley graph** $\Gamma_{G,S}$ (Definition 3.3) is the graph with vertices labelled by elements of G , and with an edge between two vertices $g, h \in G$ if $g^{-1}h \in S$. The natural path metric $d_{G,S}$ on the Cayley graph induces a metric on the elements of G called the **word metric**. In this metric, the distance from any $g \in G$ to the identity is just the length of the shortest word in S which represents g . The Cayley graph, of course, depends on the generating set, but many large-scale aspects of the geometry are invariant under change of generating set. Indeed, if S, S' are two different generating sets, then $d_{G,S}$ and $d_{G,S'}$ are biLipschitz equivalent on G . In the sequel, we will be interested in understanding geometric properties of groups which

are independent of generating set. We call such properties **coarse** or **geometric** invariants (Definition 3.5).

1.1 CAT(0) geometry and δ -hyperbolicity

CAT(0) geometry was introduced by Alexandrov in the 1950s and popularized by Gromov [4] in the 1980s as an extension of nonpositive curvature in Riemannian geometry to the non-manifold setting. Given a geodesic metric space X , one can compare each geodesic triangle in X to a triangle with the same side lengths in the Euclidean plane \mathbb{E}^2 . If triangles in X are no “fatter” than their Euclidean counterparts, one says that X is CAT(0) (see §3.1.3 for a formal definition). More generally, for $\kappa \leq 0$, a geodesic metric space X is CAT(κ) if triangles in X are no larger than triangles in the complete Riemannian surface M_κ of constant sectional curvature κ . The notation “CAT” was coined by Gromov, and stands for the pioneering work of Cartan, Alexandrov and Topogonov in comparison geometry.

The CAT(0) condition often places strong restrictions on the structure of groups which can act on them geometrically. Like nonpositively curved Riemannian manifolds, CAT(0) spaces are contractible (a generalization of the Cartan-Hadamard theorem) and any finite isometric action on a CAT(0) space has a global fixed point (Theorem 3.6). It follows that if a group G acts freely on a finite-dimensional CAT(0) space then G has a finite-dimensional classifying space and is torsion-free. Moreover, it follows from the flat torus theorem ([5], Theorem II.7.1) that any solvable subgroup of a cocompact group of isometries on a CAT(0) space must be virtually abelian. The analogous result in the Riemannian setting is due independently to Gromoll–Wolf [6] and Lawson–Yau [7]. For a comprehensive introduction to CAT(0) geometry, see [5].

A notion related to the CAT(0) condition and also introduced by Gromov [4]

is that of δ -hyperbolicity (Definition 3.6). As in the definition of a CAT(0) space, hyperbolicity is also defined in terms of triangles. A group G is called **(word)-hyperbolic** if its Cayley graph is δ -hyperbolic for some generating set and some δ . Since δ -hyperbolicity is invariant under quasi-isometry and a fortiori, biLipschitz equivalence, the property of being hyperbolic is well-defined independent of generating set. For $\kappa < 0$, a CAT(κ) space is $\delta(\kappa)$ -hyperbolic, but the converse is far from being true. In fact, it is an open question whether every δ -hyperbolic group can act geometrically on a CAT(0) space. For more on δ -hyperbolic spaces and groups, see [4], [8], [9], and [10].

1.2 NPC Cube complexes

In general, it is difficult to verify whether a geodesic metric space is CAT(0), as one must check all geodesic triangles. Gromov [4] introduced nonpositively curved (NPC) cube complexes as a source of easily constructible examples of CAT(0) spaces. Cube complexes are constructed from Euclidean n -cubes $[-1, 1]^n$ by gluing them together along their faces by isometries. The metric is given by the path metric induced from the Euclidean metric on each cube. A cube complex X is said to be NPC if its universal cover is CAT(0). In this setting, Gromov showed that the CAT(0) condition reduces to a simple combinatorial condition at the links of vertices of X (see 3.2 for details).

Every CAT(0) cube complex comes equipped with many codimension one subspaces called **hyperplanes** (Definition 3.13). Hyperplanes are CAT(0) cube complexes in their own right and sit inside the bigger cube complex as a convex subspaces. Each hyperplane separates the cube complex into two disjoint halfspaces, and one can define a metric on the vertices by counting the number of hyperplanes crossed by any

combinatorial path. The rich structure of CAT(0) cube complexes comes from the existence of these two natural metrics associated to them: the CAT(0) path metric and the L^1 -metric (the combinatorial metric) on the vertices (Definition 3.14).

The most basic example of a CAT(0) cube complex is Euclidean n -space \mathbb{E}^n , tiled in the usual way by n -cubes. The hyperplanes are even-integer translates of the standard coordinate hyperplanes, and the L^1 -metric is the so-called taxicab metric on \mathbb{R}^n . One-dimensional CAT(0) cube complexes are just simplicial trees, and already in this case, work of Bass and Serre [11] shows that groups which act geometrically on simplicial trees have a definite structure: they can be decomposed as a **graph of groups**. More precisely, if G acts geometrically on a tree T , then G can be written as a free product for the vertex stabilizers of the action, amalgamated along the edge stabilizers. Conversely, if a group G decomposes as a graph of groups, one can construct an action of a group on a simplicial tree. For higher-dimensional cube complexes, there is a related decomposition, called a quasi-convex hierarchy, which we elucidate below.

1.3 Groups with a quasi-convex hierarchy

The relationship between groups and actions on CAT(0) cube complexes is intimately tied to the notions of quasi-convexity (Definition 3.20) and codimension one subgroups. In general, quasi-convexity depends on the generating set for G , but when G is hyperbolic, quasi-convexity is independent of generating set (Theorem 3.11). A subgroup $H \leq G$ is said to be of **codimension one** if all sufficiently large neighborhoods of H in the Cayley graph of G disconnect the Cayley graph into two or more unbounded components.

In his thesis, Sageev [12] gave a construction generalizing the one above for trees to

higher-dimensional CAT(0) cube complexes. Given a group G and a finite collection $\{H_1, \dots, H_k\}$ of codimension one subgroups, Sageev produces a dual CAT(0) cube complex which admits an action of G . When G is hyperbolic and the H_i are quasi-convex, Sageev showed that the dual cube complex will be finite dimensional and the action will be geometric. Conversely, if G acts geometrically on a CAT(0) cube complex, there are finitely many orbits of hyperplanes and the stabilizer of each hyperplane is a quasi-convex, codimension one subgroup of G .

Roller [13] formalized the Sageev construction by demonstrating the equivalence of actions of groups on pocsets (= poset with complementation) subject to certain finiteness conditions and actions on their dual CAT(0) cube complexes (see also work of Haglund and Paulin [14] on the action of a group on a space with walls). Sageev's construction provides an algebraic condition which allows one to build an action on a CAT(0) cube complex, but it does not give a method for classifying groups which can act geometrically on CAT(0) cube complexes. For this last step, we need to define special cube complexes.

Haglund and Wise [15] introduced **special cube complexes** as a class of NPC cube complexes whose hyperplanes behave in a controlled way (see §3.3 for a definition). **Special groups**, or groups which arise as fundamental groups of finite-dimensional special cube complexes, are known to enjoy many nice properties (see Theorem 3.10); in particular, when the cube complex is compact, they embed in $\mathrm{SL}(n, \mathbb{Z})$ [16] and are residually torsion-free nilpotent [17]. The latter implies moreover that such groups are **locally indicable**, *i.e.* every finitely generated nontrivial subgroup surjects onto \mathbb{Z} . All of the above stated properties are consequences of the fact that fundamental groups of compact special cube complexes embed into right-angled Artin groups (see §1.4 or Definition 3.19). Many families of groups commonly

occurring in low-dimensional topology are now known to be (virtually) compact special. Among them are:

1. Free groups and free abelian groups,
2. Surface groups,
3. Nonpositively curved 3-manifold groups ([18], [19], [20]),
4. Finitely generated Coxeter groups [1],
5. One-relator groups with torsion [21].

If a hyperbolic group G is the fundamental group of a special cube complex X , then any hyperplane Y in X corresponds to a quasi-convex, codimension one subgroup $H = \pi_1(Y)$ of G . Because of the properties guaranteed by specialness, we can cut open X along the hyperplane Y and write $X = X_1 \amalg_Y X_2$ or $X = X_1/(Y_1 \sim Y_2)$ according as H separates X into two components or one. We get a corresponding decomposition of G as an amalgamated product or HNN-extension over the quasi-convex subgroup H . After cutting once, we are left with a union of special cube complexes, each with hyperbolic fundamental groups, and we can continue cutting until eventually we are left with a collection of special cube complexes each with trivial fundamental group. The result is a way of decomposing G as a tower of iterated amalgamated products and HNN-extensions over quasi-convex subgroups, starting from the trivial group and ending in G . Such a decomposition is called a **quasi-convex hierarchy** for G (Definition 3.21).

Remarkably, Wise [21] showed (see also Agol–Groves–Manning [22]) that the converse is true. Namely, if G is hyperbolic and has a quasi-convex hierarchy terminating

in the trivial group, then G is virtually special, *i.e.* the group G has a finite index subgroup which is special (Theorem 3.12). Building on Wise’s result, Agol [18] proved the following theorem which allowed him to deduce the virtual Haken and virtual fibering conjecture for hyperbolic 3-manifolds:

Theorem 1.1 (Agol [18]). *If G is hyperbolic and acts geometrically on a $CAT(0)$ cube complex, then G is virtually special.*

Much still remains unknown about the structure of groups that act on $CAT(0)$ cube complexes once the assumption on hyperbolicity is removed, and part of the motivation for this thesis is to understand the properties of special groups in general.

1.4 Right-angled Artin groups

As noted above, Haglund–Wise [15] proved that the class of right-angled Artin groups (raag) are the universal receptors for special groups, in the sense that every compact special group embeds in a raag (Corollary 3.1). As such, right-angled Artin groups have a very rich subgroup structure, making them an interesting class of groups to study. Artin groups, more generally, are an infinite family of groups with similar presentations and properties, containing braid groups and other groups with strong connections to geometry and topology (see [23] for a survey). Raags represent one of the simplest classes of Artin groups which are not all of finite type, but for which the $K(\pi_1, 1)$ -conjecture of Arnol’d is known to hold [24].

Let $\Gamma = (V, E)$ be a finite simplicial graph with vertex set V and edge set E . The right-angled Artin group A_Γ associated to Γ is the group with presentation

$$A_\Gamma = \left\langle v \in V \mid [v, w], \text{ if } v, w \text{ share an edge in } \Gamma \right\rangle.$$

A **right-angled Artin group** is any group obtained in this way. If Γ has no edges, then there are no relations, hence A_Γ is free. On the other hand, if Γ is a complete graph, then all the generators commute and A_Γ is free abelian. Because of this, raags are said to interpolate between free groups F_n and free abelian groups \mathbb{Z}^n . In this thesis, we will chiefly be interested in what properties of the automorphism groups of raags are preserved as we pass from $\text{Aut}(F_n)$ to $\text{GL}_n(\mathbb{Z})$.

The study of $\text{GL}_n(\mathbb{Z})$ is classical, going back to the 1800s. This is primarily due to the close connection between $\text{GL}_n(\mathbb{Z})$ and the classification of lattices $\mathbb{Z}^n \hookrightarrow \mathbb{R}^n$, or equivalently, marked flat metrics on the n -torus \mathbb{T}^n . More precisely, one can identify the space of rank n positive definite quadratic forms \mathcal{Q}_n with the symmetric space $\text{GL}_n^+(\mathbb{R})/\text{SO}_n(\mathbb{R})$ via the mapping

$$A \in \text{GL}_n^+(\mathbb{R}) \mapsto A \cdot A^t = Q.$$

The representation $Q = A \cdot A^t$ gives the standard inner product on \mathbb{R}^n with respect to the basis given by the columns of A . Consequently, we can consider Q either as giving a positive definite quadratic form on \mathbb{R}^n or as representing an embedding of $\mathbb{Z}^n \hookrightarrow \mathbb{R}^n$ with a particular choice of basis, up to a rotation. After changing the basis for \mathbb{Z}^n by an element of $\text{GL}_n(\mathbb{Z})$, the image of the lattice does not change, but the quadratic form does change, *i.e.* there is an action of $\text{GL}_n(\mathbb{Z})$ on \mathcal{Q}_n by $M \cdot Q = MQM^t$ for $M \in \text{GL}_n(\mathbb{Z})$. This action is not cocompact, but it is proper. Since $\text{SO}_n(\mathbb{R})$ is a maximal compact subgroup of $\text{GL}_n^+(\mathbb{R})$, the quotient space \mathcal{Q}_n is contractible, hence for any torsion-free subgroup $\Gamma \leq \text{GL}_n(\mathbb{Z})$, the quotient \mathcal{Q}_n/Γ is an Eilenberg-MacLane space for Γ [25]. Thus, the space \mathcal{Q}_n provides a geometric model for understanding the group structure of $\text{GL}_n(\mathbb{Z})$.

The automorphism group $\text{Aut}(F_n)$ has been actively studied since the early part of the 20th century, beginning with Nielsen [26], [27]. The geometric picture analogous

to the one above for $\mathrm{GL}_n(\mathbb{Z})$ did not arise until Culler–Vogtmann’s [28] construction of **outer space** CV_n , defined as follows. Let R_n denote a wedge of n circles. A point in CV_n corresponds to a pair (G, ρ) , where G is a graph whose fundamental group is F_n and ρ is a **marking**, *i.e.* a homotopy equivalence $R_n \rightarrow G$. The graph G is not allowed to have vertices of valence one or two, and every edge of G has been assigned a positive length, inducing a path metric. The marking allows us to identify F_n with $\pi_1(G)$, but without specifying basepoints. Because of the ambiguity of basepoints, outer space CV_n is equipped with a proper action not by $\mathrm{Aut}(F_n)$ but by $\mathrm{Out}(F_n)$. Culler–Vogtmann [28] showed that CV_n is contractible and finite-dimensional, providing a geometric model analogous to \mathcal{Q}_n above.

One expects a similar picture generalizing the one above to arbitrary raags, but as yet only partial results hold. Many properties of $\mathrm{Out}(A_\Gamma)$, such as being virtually torsion free and having finite virtual cohomological dimension, have been proven by algebraic methods [29]. In the case of 2-dimensional raags without leaves, Charney–Margolis [30] proved that any minimal action on a rectangle complex is determined by its length function, but did not construct a moduli space of such actions. Recently, Charney–Stambaugh–Vogtmann [31] constructed a finite-dimensional, contractible simplicial complex on which a subgroup of $\mathrm{Out}(A_\Gamma)$ acts geometrically. One of our main motivations for studying automorphisms of NPC cube complexes was to extract useful information about the $\mathrm{Out}(A_\Gamma)$ from the above action.

Chapter 2

Statement of Results

Let G be a special group. In this thesis we will be interested in how the abelianization $H_1(G)$ determines the geometry of a special cube complex X with $\pi_1(X) \cong G$. It was shown by Wise [21] (see also Koberda–Suciuc [32]) that special groups which are not virtually abelian are **large**; they have finite index subgroups which surject onto the non-abelian free group F_2 . In particular, the rank of H_1 grows at least linearly after passing to finite index subgroups. Wise further asked ([21], pg. 143) whether any special group is either abelian or surjects onto a non-cyclic free group. A corollary of our main theorem answers this question in the affirmative

Theorem 2.1. *Let G be the fundamental group of a finite dimensional special cube complex. Then either G is abelian or surjects onto F_2 .*

Theorem 2.1 was originally proved by Antolín and Minasyan [33] who, using different methods, showed that any subgroup of a (finitely or infinitely generated) right-angled Artin group is either abelian or surjects onto F_2 . In our proof, we introduce an invariant of special cube complexes which we call the **genus**. This definition has a classical motivation, namely the original combinatorial genus of a surface due to Betti and Poincaré [34]: The **classical genus** of a closed surface Σ is the maximal number of disjoint non-separating simple closed curves whose union does not disconnect Σ . Analogously, if X is special then $g(X)$ is defined to be the number of pairwise disjoint, non-separating hyperplanes whose union does not disconnect X . We extend

this definition to special groups by defining $g(G)$ to be the maximum genus over all X with $\pi_1(X) = G$. If $g(G) = n$, then $G = \pi_1(X)$ surjects onto F_n . The geometric analogue of Theorem 2.1 characterizes low values of the genus explicitly:

Theorem 2.2. *Let X be special and finite dimensional. Then*

1. $g(X) = 0$ if and only if X is CAT(0).
2. $g(X) = 1$ if and only if $\pi_1(X)$ is abelian.
3. If Σ_γ denotes the closed orientable surface of classical genus γ , then

$$g(\pi_1(\Sigma_\gamma)) = \gamma.$$

In particular, the classical definition of genus agrees with ours. The geometric content of this theorem is that if G is special and not abelian, and X is any special cube complex with $\pi_1(X) = G$, there exists a map of cube complexes (Definition 3.9) $X \rightarrow S^1 \vee S^1$. We remark that for a general group G a notion related to the genus is the **corank**, *i.e.* the largest rank of a free group onto which G surjects. If $G = \pi_1(M)$ for some smooth manifold M , then the corank is the same as the **cut number**, the largest number of disjointly, properly embedded, 2-sided hypersurfaces in M whose union does not separate. This follows from the fact that the wedge of n circles is a $K(F_n, 1)$. Thus the genus of a special group G gives a lower bound for the corank. It would be interesting to know whether the genus is always equal to the corank.

In the second half of the thesis we investigate automorphisms of special groups and the action of the automorphisms of a cube complex on first homology. There are two parts to this problem: (1) which automorphisms of G can be realized as an automorphism of X , a compact cube complex with $\pi_1(X) = G$, and (2) when does an automorphism of X act non-trivially on $H_1(X) = H_1(G)$. Denote by $\text{Aut}(G)$ the

group of automorphisms of G , by $\text{Out}(G)$ the group of outer automorphisms of G , and by $\mathcal{I}(G) \leq \text{Out}(G)$ the subgroup of automorphisms acting trivially on $H_1(G)$.

The motivation for answering these questions comes from classical results on Riemann surfaces and free groups. Let $\Sigma = \Sigma_g$ be a surface of genus $g \geq 2$, and denote by $\text{Mod}(\Sigma)$ its mapping class group, *i.e.* the group of orientation-preserving diffeomorphisms of Σ up to homotopy. The Dehn–Nielsen–Baer theorem identifies $\text{Mod}(\Sigma)$ as an index 2 subgroup of $\text{Out}(\pi_1(\Sigma))$. If $\phi \in \text{Mod}(\Sigma)$ has finite order, it is a classical result that there exists a hyperbolic surface X diffeomorphic to Σ and an isometry $f : X \rightarrow X$ which *realizes* the homotopy class of ϕ (see for example, [35]). This fact can be used to show that the Torelli subgroup $\mathcal{I}(\pi_1(\Sigma))$ is torsion-free, by showing that any isometry of a hyperbolic surface acts non-trivially on first homology. The theorem reduces an algebraic question about subgroups of the mapping class groups to a geometric question about isometries of a compact surface.

Similarly, for free groups, Culler [36], Zimmermann [37], and Khramtsov [38] each independently showed that any finite order automorphism $\phi \in \text{Out}(F_n)$ can be realized as an automorphism of a simplicial graph Γ of rank n . An easy geometric argument then recovers the result of Baumslag–Taylor that $\mathcal{I}(F_n)$ is torsion-free for all n [39].

Recently, for each raag A_Γ , Charney, Stambaugh and Vogtmann [31] defined a contractible simplicial complex K_Γ on which a subgroup of $\text{Out}(A_\Gamma)$ acts properly discontinuously, cocompactly by simplicial automorphisms. Their space is defined in analogy with outer space for free groups, and if $A_\Gamma = F_n$, the complex K_Γ is just the spine of outer space. Using K_Γ , we show

Theorem 2.3. *Let $\phi \in \text{Out}(A_\Gamma)$ have finite order. Then ϕ acts non-trivially on $H_1(A_\Gamma)$. In particular, $\mathcal{I}(A_\Gamma)$ is torsion-free.*

This theorem is originally due to Wade [40], and independently Toinet [41], who proved the stronger result that the Torelli subgroup associated to $\text{Out}(A_\Gamma)$ is residually torsion-free nilpotent. However, both these proofs are almost entirely algebraic, while ours is geometric in the same spirit as those outlined for mapping class groups and free groups above. To prove Theorem 2.3, we first realize ϕ as a finite order automorphism of a compact special cube complex X whose fundamental group is A_Γ , then prove that any such automorphism acts non-trivially on $H_1(X) = H_1(A_\Gamma)$.

We also include a realization result for automorphisms of special groups which are δ -hyperbolic, and the following result about large groups which is elementary but which we nevertheless could not find in the literature.

Theorem 2.4. *Suppose G surjects onto F_2 and let $\phi \in \text{Out}(G)$ have finite order. Then there exists a finite index normal subgroup $N \trianglelefteq G$ and an outer automorphism $\psi \in \text{Out}(N)$ such that ψ acts non-trivially on $H_1(N)$ and $\psi_* = \phi_* \circ \iota_*$, where $\iota : N \rightarrow G$ is the inclusion.*

Rephrased in terms of spaces, if X is a $K(G, 1)$ and $\phi \in \text{Out}(G)$ we can represent ϕ as a homotopy equivalence $f : X \rightarrow X$. Then there exists a finite regular cover $p : \widehat{X} \rightarrow X$, a homotopy equivalence $\widehat{f} : \widehat{X} \rightarrow \widehat{X}$ such that $f \circ p = p \circ \widehat{f}$, and such that \widehat{f}_* acts non-trivially on $H_1(\widehat{X})$.

Outline of Results: After going over background material in Chapter 3, Theorems 2.1 and 2.2 are proven in Chapter 4. In Chapter 5, we prove Theorem 2.4 and other results related to automorphisms of special cube complexes and groups. Finally, in Chapter 6, we give our proof of Theorem 2.3.

Chapter 3

Background

In this chapter we describe relevant background in geometric group theory and $\text{CAT}(0)$ geometry which will be used for the rest of the thesis. In section 3.1, we provide some basic geometric group theory background, then give the definition of a δ -hyperbolic space and a $\text{CAT}(\kappa)$ metric space. In section 3.2, we give the definition of a nonpositively curved (NPC) cube complex, review Gromov's link condition, and describe the hyperplane structure of $\text{CAT}(0)$ cube complexes. Finally, in section 3.3, we introduce special cube complexes, right-angled Artin groups and quasi-convex hierarchies.

3.1 Group theory, δ -hyperbolicity, $\text{CAT}(0)$ geometry

Let G be a finitely generated group, with finite generating set S . In this thesis, we will often be concerned with various algebraic and geometric properties of groups, and the relationships between them.

3.1.1 Group theory preliminaries

Before we describe how one can consider G as a geometric object, we first define what we mean by algebraic properties of groups.

Definition 3.1. An **algebraic property** \mathcal{P} of the group G is a property which is invariant under isomorphism. The group G is said to **virtually have property** \mathcal{P} or

be **virtually** \mathcal{P} if there exists a finite index subgroup $G' \leq G$ which has \mathcal{P} . A group is said to **residually have property** \mathcal{P} or be **residually** \mathcal{P} if for every nontrivial element $g \in G$, there exists a group Q having property \mathcal{P} and a surjection $\phi : G \rightarrow Q$ such that $\phi(g) \neq 1$.

We now present some examples of algebraic properties to illustrate the definition just given, and which we will consider later on.

Example 3.1. A group G is **virtually torsion-free** if there exists a finite index subgroup $G' \leq G$ which is torsion-free.

The following theorem is sometimes referred as Selberg's Lemma:

Theorem 3.1 (Selberg [42]). *Let G be finitely generated. If G admits a faithful representation to $GL_n(k)$ for some field k of characteristic 0, then G is virtually torsion-free.*

We also present an important example of a residual property.

Example 3.2. A group G is **residually finite** if for every $g \neq 1 \in G$, there exists a finite group F and a surjection $\phi : G \rightarrow F$ such that $\phi(g) \neq 1 \in F$.

Roughly speaking, this says that one can distinguish non-trivial elements from the identity in G by looking at finite quotients. Mal'cev proved the following result:

Theorem 3.2 (Mal'cev [43]). *Let G be finitely generated. If G admits a faithful representation to $GL_n(k)$ for some field k , then G is residually finite.*

Taken together, Theorems 3.1 and 3.2 provide a method to certify that a given finitely generated group is both virtually torsion-free and residually finite: simply exhibit a faithful representation of G into $GL_n(k)$ for some field k of characteristic 0.

As we will see later on in this chapter, this is exactly how one shows that virtually special groups (Definition 3.18) have both of these properties.

There are two useful combinatorial group theory constructions for building new groups out of old ones, known as amalgamation and HNN-extension. Amalgamation is familiar to anyone who has invoked the Seifert-van Kampen theorem, but HNN-extension is probably less so. We define each now.

Definition 3.2. 1. Let A and B be two groups with isomorphic subgroups $C \leq A$, $C' \leq B$. If $\phi : C \rightarrow C'$ is an isomorphism, we define the **amalgamated product of A and B along C** to be the group with presentation

$$\langle A, B \mid c = \phi(c) \text{ for all } c \in C \rangle.$$

Denote the amalgamated product by $A *_C B$ or simply $A *_C B$.

2. If $C, C' \leq A$ are isomorphic subgroups with isomorphism $\phi : C \rightarrow C'$, we define the **HNN-extension of A along C** to be the group with presentation

$$\langle A, t \mid tct^{-1} = \phi(c) \text{ for all } c \in C \rangle.$$

Denote the HNN-extension by $A *_C A$ or $A *_C A$.

Amalgamation creates a new group by identifying the isomorphic subgroups C and C' , while HNN-extension adds a generator to A which conjugates C to C' via ϕ . The notation “HNN” comes from the three group theorists who first defined it, namely, Graham Higman, Bernhard Neumann, and Hanna Neumann.

3.1.2 Coarse geometry preliminaries

We now turn to the geometric side of things. We will often denote a metric space by a pair (X, d) where $d : X \times X \rightarrow \mathbb{R}$ is the metric. The metric space (X, d) is **proper**

if closed balls in X are compact. Any metric d on X induces a **length function** ℓ on paths $\gamma : [0, 1] \rightarrow X$ (see [5], Definition I.18). The length function in turn induces a **path metric** d_{path} where for two points $p, q \in X$:

$$d_{\text{path}}(p, q) = \inf\{\ell(\gamma) \mid \gamma \text{ a (rectifiable) path between } p \text{ and } q\}.$$

We call (X, d) a **path metric space** if $d = d_{\text{path}}$, and in addition, we say X is a **geodesic** metric space if for any two points $p, q \in X$, the distance $d(p, q)$ is achieved as the length of some path between p and q .

The way to study a finitely generated group geometrically is by studying its Cayley graph with respect to some generating set:

Definition 3.3. Given a finitely generated group with finite generating set S , the **Cayley graph** $\Gamma_{G,S}$ is the graph defined as follows. The vertices of $\Gamma_{G,S}$ are in one-to-one correspondence with the elements $g \in G$. Two vertices g, h are connected by an edge if $g^{-1}h \in S$.

The generating set S induces a norm on elements of G via

$$\|g\|_S = \min\{n \mid s_1 \cdots s_n = g, s_i \in S\}.$$

This norm is called the **word norm** with respect to S and further induces a **word metric** $d_{G,S}$ on G via the formula $d_{G,S}(g, h) = \|g^{-1}h\|_S$. By assigning the length one to each edge in $\Gamma_{G,S}$ it is easy to see that the minimum distance in the Cayley graph between two vertices labeled g and h is just $d_{G,S}(g, h)$. Moreover, the left action of G on itself induces an automorphism of $\Gamma_{G,S}$ and hence an isometry of the word metric. Indeed, we have

$$d_{G,S}(g \cdot g_1, g \cdot g_2) = \|(g_1^{-1}g^{-1}) \cdot (gg_2)\| = \|g_1^{-1} \cdot g_2\| = d_{G,S}(g_1, g_2).$$

The proper, geodesic metric space $(\Gamma_{G,S}, d_{G,S})$ is one of the basic objects of modern geometric group theory, and enables us to treat an algebraic object, namely G , as a geometric space. For a different generating set S' , the identity map on G is a biLipschitz equivalence between the two word metrics $d_{G,S}$ and $d_{G,S'}$. Unfortunately, the identity does not in general extend to a biLipschitz map $\Gamma_{G,S} \rightarrow \Gamma_{G,S'}$. In the large, however, the two metric spaces $\Gamma_{G,S}$ and $\Gamma_{G,S'}$ look very similar; this notion of similarity is made precise via the next definition.

Definition 3.4. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \rightarrow Y$ is called a **quasi-isometric embedding** if there exists a constant $C > 1$ such that for all points $x_1, x_2 \in X$:

$$\frac{1}{C}d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq Cd_X(x_1, x_2) + C.$$

If there exists a further constant $D > 0$ such that for all $y \in Y$, there exists an $x \in X$ satisfying

$$d_Y(y, f(x)) \leq D,$$

then we call f a **quasi-isometry** between X and Y .

One thinks of a quasi-isometric embedding as a Lipschitz map plus a small error constant; distances may be stretched and altered by a finite amount. It is not difficult to show that if $f : X \rightarrow Y$ is a quasi-isometry, then there exists a quasi-isometry $g : Y \rightarrow X$. Thus, the notion of quasi-isometry becomes an equivalence relation, which we denote by $X \simeq_{q.i.} Y$. If G is a finitely generated group, then any two Cayley graphs for G are quasi-isometric, hence it makes sense to ask whether two finitely generated groups are quasi-isometric to each other.

Definition 3.5. A **coarse** or **geometric** property of a finitely generated group G (or metric space X) is any property which is invariant under quasi-isometry.

The additive constant in the definition of quasi-isometry is exactly what is needed to make sense of two geometric spaces being isometric *in the large*, even though local structure may differ, many asymptotic invariants of both spaces are the same. The following foundational result due to Schwartz [2] and Milnor [3] is what allows one to connect the asymptotic geometry of a finitely generated group to a geometric space upon which it acts by isometries.

Theorem 3.3 (Schwartz [2], Milnor [3]). *Let X be a proper, geodesic metric space and suppose that G acts on X properly discontinuously, cocompactly by isometries. Then G is finitely generated and $G \simeq_{q.i.} X$.*

We call an action as in the above theorem **geometric**. The primary example of a geometric action is the left-action of a group on its own Cayley graph, as described above. Moreover, if $H \leq G$ is a finite index subgroup, it follows from the above theorem and the action of H on the Cayley graph of G that $H \simeq_{q.i.} G$. In special case of a free action, we obtain the following: if X is a compact geodesic metric space and $G \cong \pi_1(X)$ then G is quasi-isometric to the universal cover \tilde{X} . This observation tells us that many coarse invariants of G can be computed from a compact metric space whose fundamental group is G . One of the most useful coarse invariants is that of δ -hyperbolicity, originally defined by Gromov [4].

Definition 3.6. Let X be a geodesic metric space. Given any triple of points $p, q, r \in X$, a **geodesic triangle** T in X is any union of three geodesics connecting pairs of $\{p, q, r\}$. We denote the geodesic sides of this triangle by intervals $[p, q]$, $[q, r]$ and $[p, r]$. The metric space X is said to be **δ -hyperbolic** if there exists $\delta > 0$ such that for any geodesic triangle $T = [p, q] \cup [q, r] \cup [p, r]$, the δ -neighborhood of $[p, q] \cup [q, r]$ contains the third side $[p, r]$. We say X is **hyperbolic** if it is hyperbolic for some δ .

Theorem 3.4 (Gromov [4]). *If Y is hyperbolic and $f : X \rightarrow Y$ is a quasi-isometric embedding, then X is hyperbolic. In particular, δ -hyperbolicity is a coarse invariant.*

A finitely generated group G is **hyperbolic** if its Cayley graph is hyperbolic with respect to some (and hence any) generating set. Hyperbolicity is a strong condition and implies many nice properties for a group. For example, if G is hyperbolic then it is finitely presented and has solvable word and conjugacy problem. Many examples of hyperbolic groups come from fundamental groups of negatively curved compact Riemannian manifolds. Indeed, real hyperbolic n -space \mathbb{H}^n is δ -hyperbolic, hence cocompact lattices in $\text{Isom}(\mathbb{H}^n)$, such as fundamental groups of closed surfaces of classical genus $g \geq 2$ and fundamental groups of closed hyperbolic 3-manifolds are δ -hyperbolic, by Theorem 3.3.

3.1.3 CAT(0) geometry

A notion related to δ -hyperbolicity, also popularized by Gromov, is the idea of a CAT(κ) metric, where $\kappa \leq 0$. Let M_κ denote the complete, simply connected surface of constant sectional curvature $\kappa \leq 0$. If $\kappa = 0$, then $M_\kappa = \mathbb{E}^2$, the Euclidean plane, and if $\kappa = -1$, then $M_\kappa = \mathbb{H}^2$, the hyperbolic plane. Given a geodesic metric space X , and a triangle $T = [p, q] \cup [q, r] \cup [p, r]$ in X , one can always compare T to a triangle $T' = [p', q'] \cup [q', r'] \cup [p', r']$ in M_κ with equal side lengths, as a consequence of the triangle inequality. Moreover, there is a **comparison map** $f_\kappa : T \subset X \rightarrow T' \subset M_\kappa$ which is an isometry restricted to each side of T .

Definition 3.7. A geodesic metric space (X, d) is said to be **CAT(κ)** if the following condition holds for every triangle $T = [p, q] \cup [q, r] \cup [p, r]$: Given points $x \in [p, q]$ and

$y \in [p, r]$ which are equidistant from p , we have that

$$d_X(x, y) \leq d_{M_\kappa}(f_\kappa(x), f_\kappa(y)),$$

where f_κ is the comparison map.




Figure 3.1 : A schematic of the comparison map between triangle pqr in X and $p'q'r'$ in M_κ . The distance between x and y (the blue line) increases under f_κ .

Figure 3.1 shows a schematic of the comparison map between X and the model space M_κ . Heuristically, the $\text{CAT}(\kappa)$ condition means that triangles in X are no “fatter” than triangles in M_κ , where “fatter” means that distances under the comparison map are nondecreasing. The $\text{CAT}(\kappa)$ condition gives a way to extend notions of nonpositive curvature to non-manifold metric spaces and also provides a way to certify δ -hyperbolicity:

Theorem 3.5 ([5], Theorem II.1A.6, Theorem III H.1.2). *1. If M is a complete, Riemannian manifold of sectional curvature bounded above by $\kappa \leq 0$, then the universal cover \widetilde{M} is $\text{CAT}(\kappa)$.*

2. If X is $\text{CAT}(\kappa)$, for some $\kappa < 0$, then X is δ -hyperbolic for some δ .

We remark, however, that being quasi-isometric to a $\text{CAT}(\kappa)$ space is **NOT** a quasi-isometry invariant, nor is acting geometrically on a $\text{CAT}(\kappa)$ space, for a group.

By the above theorem, we can view any $\text{CAT}(\kappa)$ space as a $\text{CAT}(0)$ space and will often do so when a specific κ is not important. $\text{CAT}(0)$ spaces share much in common with simply connected, nonpositively curved Riemannian manifolds. In particular, the metric of a $\text{CAT}(0)$ space is convex, implying a version of the Cartan–Hadamard theorem:

Theorem 3.6 ([5], Theorem II.4.1, Corollary II.2.8). *Let X be $\text{CAT}(0)$. Then X is contractible, uniquely geodesic, and any finite group of isometries of X has a global fixed point.*

Here, **uniquely geodesic** means that there is a unique geodesic between any two points of X . In particular, the above result shows that if X is $\text{CAT}(0)$ then *a fortiori* X is simply connected. Like sectional curvature, one thinks of the $\text{CAT}(0)$ property is a local condition. A metric space (X, d) is said to be **locally $\text{CAT}(0)$** (resp. **$\text{CAT}(\kappa)$**) if every point $p \in X$ has connected neighborhood U such that the restriction of d to U is $\text{CAT}(0)$ (resp. $\text{CAT}(\kappa)$). We have following analogue of the local-to-global structure theorem from Riemannian geometry.

Theorem 3.7 ([5], Theorem II.4.1). *X is $\text{CAT}(0)$ if and only if it is connected, simply connected and locally $\text{CAT}(0)$.*

We will often exploit the above theorem to show spaces are $\text{CAT}(0)$. From the point of view of group theory, fundamental groups of locally $\text{CAT}(0)$ spaces have similar properties to fundamental groups of nonpositively curved Riemannian manifolds. In particular, it follows from Theorems 3.7 and 3.6 that if X is locally $\text{CAT}(0)$ and $G \cong \pi_1(X)$, then G is torsion-free and X is an Eilenberg-MacLane space $K(G, 1)$.

3.2 NPC cube complexes

In this section, we introduce a type of locally CAT(0) space called a nonpositively curved (NPC) cube complex. NPC cube complexes and their fundamental groups are a very rich class of spaces and groups, which will form the main focus of this thesis. They were first introduced by Gromov as a class of spaces for which it is easy to verify the CAT(0) condition. We first give the definition of NPC cube complexes and then explain some of their extra structure, which comes from the existence of hyperplanes.

A **Euclidean n -cube** is the space isometric to $[-1, 1]^n \subset \mathbb{E}^n$, for $n \geq 0$.

Definition 3.8. A **cube complex** X is any space obtained by identifying Euclidean cubes along their faces by isometries. The cube complex X comes with path metric coming from Euclidean metric on each cube. We say X is **locally finite** if each cube is contained in only finitely many cubes, and **finite dimensional** if the cubes in X are all of bounded dimension. We say X is **nonpositively curved (NPC)** if its universal cover \tilde{X} is CAT(0).

A cube complex does not have to be connected, but because we will often be interested in the fundamental group, they usually will be connected in what follows. Moreover, note that cubes appearing in X may be of different dimensions, and that identifying faces of the same cube is allowed. See Figure 3.2 for an example of a cube complex.

A cube complex structure on a metric space gives a CW-complex structure when one forgets about the metric. We can therefore consider continuous maps which preserve the combinatorial structure.

Definition 3.9. If X and Y are cube complexes and $f : X \rightarrow Y$ is a continuous map sending cubes to cubes and restricting to an isometry on each cube, then f is called




Figure 3.2 : An example of a cube complex.

a **map of cube complexes**. If f is bijective, we say that f is an **isomorphism** of the cube complexes and if further, $X = Y$, we say f is an **automorphism**. Denote the group of automorphisms of the cube complex structure by $\text{Aut}(X)$.

Note that an injective map of cube complexes need not be an isometric embedding, but each automorphism is an isometry of the underlying path metric.

In general, it is difficult to verify whether a geodesic space is $\text{CAT}(0)$, but for cube complexes, a theorem of Gromov reduces the $\text{CAT}(0)$ condition to a combinatorial condition at each vertex. Before we state Gromov's condition we need a definition.

Definition 3.10. Let X be a cube complex, and $v \in X^{(0)}$ a vertex. The ϵ -**link** of v for $\epsilon > 0$, denoted $\text{lk}_\epsilon(v)$, is the set of points at distance ϵ from v . For small ϵ , any two links are homeomorphic and we denote this space by $\text{lk}(v)$.

Given $v \in X^{(0)}$, the link $\text{lk}(v)$ naturally comes equipped with a triangulation: for small ϵ , the locus of points in a n -cube C at distance ϵ from a vertex is homeomorphic to an n -simplex. In particular, the intersection of the 2-skeleton of X with $\text{lk}(v)$ is

just the 1-skeleton of $\text{lk}(v)$, namely a graph.

Definition 3.11. The link of a vertex v is called **flag** if for each $k \geq 1$ and for each complete $(k+1)$ -subgraph D of the 1-skeleton $\text{lk}(v)^{(1)}$, there exists a k -simplex $\Delta^k \subset \text{lk}(v)$ such that the 1-skeleton of Δ^k is D .

We are now in a position to state Gromov's link condition:

Theorem 3.8 (Gromov [4]). *A cube complex X is NPC if and only if any vertex link of X is a flag, simplicial complex.*

Remark 3.1. Note that one must check two things at each vertex: first that the link is a simplicial complex, and second, that it is flag.

In the case of nonpositively curved Riemannian manifolds, a locally isometric immersion lifts to a convex embedding between universal covers. In the context of NPC cube complexes and maps of cube complexes, the analogous result can also be reduced to a combinatorial criterion, which we state here.

Definition 3.12. A map of cube complexes $f : X \rightarrow Y$ is called a **local isometry** if the following two conditions hold:

1. For every $x \in X^{(0)}$, the map $f : \text{lk}(x) \rightarrow \text{lk}(f(x))$ is injective.
2. If $u, v \in \text{lk}(x)^{(0)}$ and $f(u)$ and $f(v)$ are adjacent in $\text{lk}(f(x))$, then u and v are adjacent in $\text{lk}(x)$.

For NPC cube complexes, local isometries lift to convex embeddings of universal covers:

Lemma 3.1 ([44], Lemma 3.12). *If $f : X \rightarrow Y$ is a local isometry, then the induced map on universal covers $\tilde{f} : \tilde{X} \rightarrow \tilde{Y}$ is a convex embedding of CAT(0) spaces. In particular, $f_* : \pi_1(X) \hookrightarrow \pi_1(Y)$ is an injection.*

Much of the rich structure of NPC cube complexes comes from the existence of hyperplanes, which we define now.

Definition 3.13. A **midcube** of a cube $C = [-1, 1]^n$ is a subset of C obtained by restricting one of the coordinates to 0. Two midcubes $M_1 \subset C_1$ and $M_2 \subset C_2$ are **adjacent** if they intersect along a midcube of $C_1 \cap C_2$. A **hyperplane** of X is a maximal connected subset of X obtained by joining adjacent midcubes.




Figure 3.3 : Hyperplanes in a cube complex, shown in red. Note that a hyperplane through an isolated edge is just its midpoint.

Hyperplanes are important subsets of NPC cube complexes and will play a key role in what follows. See Figure 3.3 for an example of some hyperplanes from the cube complex in Figure 3.2. If $H \subset X$ and $r > 0$, we will use the notation $N_r(H)$ to denote the open r -neighborhood of H , and $\overline{N}_r(H)$ to denote the closed r -neighborhood. The inclusion map $\iota : H \hookrightarrow X$ induces a cube complex structure on H in a natural way, so that H is also NPC. One thinks of hyperplanes in X as codimension one subspaces of X , in the sense that they intersect each cube locally in a codimension one cube.

Moreover, hyperplanes are dual to the 1-skeleton of X in that a hyperplane meets any edge of X in at most one point. We say that an edge $e \subset X^{(1)}$ is **dual** to a hyperplane $H \subset X$ if $H \cap e \neq \emptyset$.

If X is NPC, then the universal cover \tilde{X} is CAT(0), and each hyperplane of \tilde{X} is a convex subspace in the CAT(0) metric. Hence, each hyperplane of the universal cover is a CAT(0) cube complex in its own right, and the inclusion map a convex embedding. Each hyperplane of \tilde{X} separates \tilde{X} into two components, which allows one to define a combinatorial metric defined on its 0-skeleton.

Definition 3.14. A hyperplane H **separates** two vertices $v_1, v_2 \in \tilde{X}$ if v_1 and v_2 lie in distinct components of $\tilde{X} \setminus H$. The **combinatorial or L^1 -distance** $d_{L^1}(v_1, v_2)$ between vertices $v_1, v_2 \in \tilde{X}$ is defined to be the number of hyperplanes separating v_1 and v_2 .

The CAT(0) condition implies that d_{L^1} is in fact a metric. A path γ in the 1-skeleton between two vertices is a **combinatorial geodesic** if it is a shortest length path in the 1-skeleton between v_1 and v_2 . Of course, any combinatorial geodesic in \tilde{X} must at least cross each of the hyperplanes which separate v_1 and v_2 at least once, but it turns out that γ is a combinatorial geodesic if and only if γ crosses each hyperplane at most once ([44], Lemma 3.9). A **combinatorial geodesic** in an NPC cube complex X is a path in the 1-skeleton of X which lifts to a combinatorial geodesic in \tilde{X} .

3.3 Special cube complexes

Hyperplanes in CAT(0) cube complexes exhibit very nice behavior; for example, they never intersect themselves, and in fact the closed unit neighborhood around any

hyperplane H is isometric to a product $H \times [-1, 1]$. In this section, we introduce special cube complexes, which are a class of NPC cube complexes whose hyperplanes mimic the behavior of those found in special cube complexes. They are defined as avoiding certain bad behaviors or pathologies, which we now define. First we have to develop some terminology concerning hyperplanes and 1-cubes.

Definition 3.15. Two distinct 1-cubes e_1 and e_2 are **adjacent** if they form the corner of a square $C \subset X$. Two hyperplanes H_1 and H_2 **intersect** if they are dual to adjacent 1-cubes. Two hyperplanes H_1 and H_2 are said to **osculate** at a vertex v if there exist a pair of distinct, non-adjacent 1-cubes e_1 and e_2 such that the hyperplane H_1 is dual to e_1 , the hyperplane H_2 is dual to e_2 , and the vertex v is an endpoint of both e_1 and e_2 .

Note that in the above definition we do not assume H_1 and H_2 are distinct. Given an NPC cube complex X and a hyperplane $H \subset X$, the inclusion $\iota : H \hookrightarrow X$ pulls back a locally trivial line bundle $\mathcal{L}(H)$ over H , defined locally by the fibers of the projection of a cube onto one of its midcubes. We call this line bundle $\mathcal{L}(H)$ the **normal bundle** to H . Note that if $\mathcal{L}(H)$ is a trivial bundle, then we can consistently orient all of the 1-cubes dual to H .

Definition 3.16 (Hyperplane pathologies). Let $H \subset X$ be a hyperplane with normal bundle $\mathcal{L}(H)$.

1. The hyperplane H **self-intersects** if H is dual to adjacent 1-cubes of X . Otherwise, the hyperplane H is **embedded**.
2. The hyperplane H is **one-sided** if its normal bundle $\mathcal{L}(H)$ is not a trivial line bundle. Otherwise, the hyperplane H is **two-sided**.

3. The hyperplane H **self-osculates** if it osculates at a vertex v . If H is two-sided and the edges e_1 and e_2 have opposite orientation, H is said to **directly** self-osculate at v . Otherwise, we say that H **indirectly** self-osculates.
4. Two hyperplanes H_1 and H_2 are said to **interoscuate** if they intersect in some cube and osculate at some vertex v .




Figure 3.4 : From left to right, the four hyperplane pathologies as listed in Definition 3.16. Image from [1].

See Figure 3.4 for a schematic of each of these four pathologies. Having defined each of the pathologies, we note that none of these occur in a CAT(0) cube complex.

Definition 3.17. An NPC cube complex X is called **special** if none of the four hyperplane pathologies of Definition 3.16 occur in X .

Although there are group theoretic reasons for wanting to avoid these pathologies in an NPC cube complex, we will content ourselves with just giving the definition. For more details, see [15] and [44].

The special condition has strong implications for the geometry, and as we will see in later chapters the homology, of special cube complexes. In particular, the fact that every hyperplane $H \subset X$ is embedded and two-sided implies that an ϵ -neighborhood $N_\epsilon(H)$ of H is isometric to a product $H \times (-\epsilon, \epsilon)$ for some $\epsilon > 0$ ($\epsilon = 1/2$ will do). Choosing an orientation on H we can consistently orient the 1-cubes dual to a

midcube of H . We say that two oriented 1-cubes e and e' are **parallel** if they are dual to the same hyperplane with the same orientation, denoted $e \parallel e'$. We use square brackets $[e]$ to designate the equivalence class of oriented edges parallel to e .

In the sequel, we will often cut open cube complexes along hyperplanes and consider the resulting cube complex. The **splitting of X along H** is the cube complex $X|H$ defined as follows. The set $X \setminus N_1(H)$ is a closed subcomplex of X , hence is compact special ([15], Corollary 3.9). There are natural inclusions $\iota^+, \iota^- : H \rightarrow X \setminus N_1(H)$. We define

$$X|H = H \times [0, 2] \coprod X \setminus N_1(H) \coprod H \times [3, 5] / (H \times \{2\} \sim \iota^-(H), H \times \{3\} \sim \iota^+(H)).$$

Note that it may be the case that $X = H \times \mathbb{S}^1$ in which case $\iota^+ = \iota^-$. We denote by H^- the image of $H \times \{0\}$ and by H^+ the image of $H \times \{5\}$ under this construction.

In this thesis, we will be studying groups which arise as fundamental groups of NPC cube complexes. Alternatively, these are exactly the groups which can act freely and geometrically on CAT(0) cube complexes.

Definition 3.18. A finitely generated group G is **(NPC) cubulated** if $G = \pi_1(X)$ for some compact NPC cube complex X . We say further that G is **(compact) special** if X is (compact) special.

The prototypical examples of compact special groups are right-angled Artin groups (raags), defined as follows.

Definition 3.19. Let $\Gamma = (V, E)$ be a finite simplicial graph. If $V = \{v_1, \dots, v_n\}$, the right-angled Artin group A_Γ associated to Γ is the group with presentation

$$A_\Gamma = \left\langle v_1, \dots, v_n \mid [v_i, v_j], \text{ if } v_i, v_j \text{ share an edge in } \Gamma \right\rangle.$$

A **right-angled Artin group** is any group obtained in this way.

To each raag A_Γ is associated a canonical NPC compact special cube complex called the **Salvetti complex** \mathbb{S}_Γ ([45], [46]). The Salvetti complex has the following cell structure:

- $\mathbb{S}_\Gamma^{(1)}$: Take a wedge of n circles, one for each vertex $v_1, \dots, v_n \in V$.
- $\mathbb{S}_\Gamma^{(2)}$: For each edge $(v_i, v_j) \in E$, attach a square $[-1, 1]^2$ along $v_i v_j v_i^{-1} v_j^{-1}$. Its image is a torus $\mathbb{T}^2 \subseteq \mathbb{S}_\Gamma^{(2)}$.
- $\mathbb{S}_\Gamma^{(k)}$: For each complete k -subgraph K of Γ , attach a k -cube $[-1, 1]^k$ by identifying its boundary with the k -many $(k-1)$ -tori in $\mathbb{S}_\Gamma^{(k-1)}$ corresponding to complete $(k-1)$ -subgraphs of K .

Figure 3.5 shows the Salvetti complex for A_Γ when Γ is a line segment of length 3. In this case, the 2-skeleton is the entire Salvetti complex.

In addition to being natural examples of special cube complexes, Salvetti complexes are also universal receptors for compact special cube complexes:

Theorem 3.9 (Haglund–Wise [15]). *Let X be compact special. Then there is a Salvetti complex \mathbb{S}_X and a local isometry $f_X : X \rightarrow \mathbb{S}_X$.*

Corollary 3.1. *Every compact special group is the subgroup of a raag.*

The corollary follows directly from Lemma 3.1 above. The Salvetti complex \mathbb{S}_X arises from the raag with defining graph $\Gamma(X)$ equal to the **crossing graph** of X : the vertices of $\Gamma(X)$ are in bijection with the hyperplanes of X , and there is an edge between two vertices if their corresponding hyperplanes cross.

Special and virtually special groups possess many nice properties, many of which follow from the fact that special groups embed in raags. We collect a few such properties of interest in a theorem.

Theorem 3.10. *Virtually special groups satisfy the following properties:*

1. *They admit faithful embeddings into $SL_n(\mathbb{Z})$ for some n [16].*
2. *They are residually torsion-free nilpotent [17].*
3. *They are virtually abelian or large [21].*

Here, a group is said to be **large** if it virtually surjects onto the free group F_2 . It follows from (1) that virtually special groups are linear over \mathbb{Q} ; hence, virtually torsion-free and residually finite by Theorems 3.1 and 3.2. Note that the first two properties above are proved by showing that raags also satisfy them, and invoking Corollary 3.1.

In the case of δ -hyperbolic group, work of Wise, Agol and many others has illuminated the connection between groups which act on CAT(0) cube complex and those which can be built up from the trivial group by iterated amalgamation and HNN-extension along quasi-convex subgroups. We conclude this chapter by briefly describing this connection here.

Let G be a finitely generated group with finite generating set S , and let $H \leq G$ be a subgroup. We can regard H as a subset of the Cayley graph $\Gamma_{G,S}$.

Definition 3.20. The subgroup H is said to be **quasi-convex** if there exists $R > 0$ such that for every pair of elements $h_1, h_2 \in H$, every geodesic $\gamma \subset \Gamma_{G,S}$ between h_1 and h_2 lies in the R -neighborhood of H .

Quasi-convexity usually depends on the generating set S . For example, consider \mathbb{Z}^2 with any basis $\{b_1, b_2\}$ as a generating set. Then with respect to these generators the coordinate axes $\langle b_i \rangle$ are quasi-convex, but the diagonal subgroup $\langle b_1 + b_2 \rangle$ is not. When G is δ -hyperbolic, however, quasi-convexity is independent of generating set:

Theorem 3.11 ([5], Corollary III Γ .3.6). *Let G be hyperbolic and $H \leq G$ a subgroup. Then H is quasi-convex if and only if H is finitely generated and the inclusion $\iota : H \hookrightarrow G$ is a quasi-isometric embedding.*

The independence of quasi-convexity on generating set for hyperbolic groups allows one to consider the class of hyperbolic groups which can be built up from the trivial group by iterated amalgamation and HNN-extension.

Definition 3.21. The class \mathcal{QH} of hyperbolic groups with a **quasi-convex hierarchy** is the smallest class of groups satisfying

- The trivial group $\{1\} \in \mathcal{QH}$.
- If $A, B \in \mathcal{QH}$, both $C \leq A$ and $C' \leq B$ are quasi-convex, and $G = A *_C=C' B$ is hyperbolic, then $G \in \mathcal{QH}$.
- If $A \in \mathcal{QH}$, both $C, C' \leq A$ are quasi-convex, and $G = A *_C=C'$ is hyperbolic, then $G \in \mathcal{QH}$.

The following theorem of Wise provides the link between hyperbolic groups with a quasi-convex hierarchy and special cube complexes:

Theorem 3.12 (Wise [44]). *If $G \in \mathcal{QH}$, then G is virtually special.*

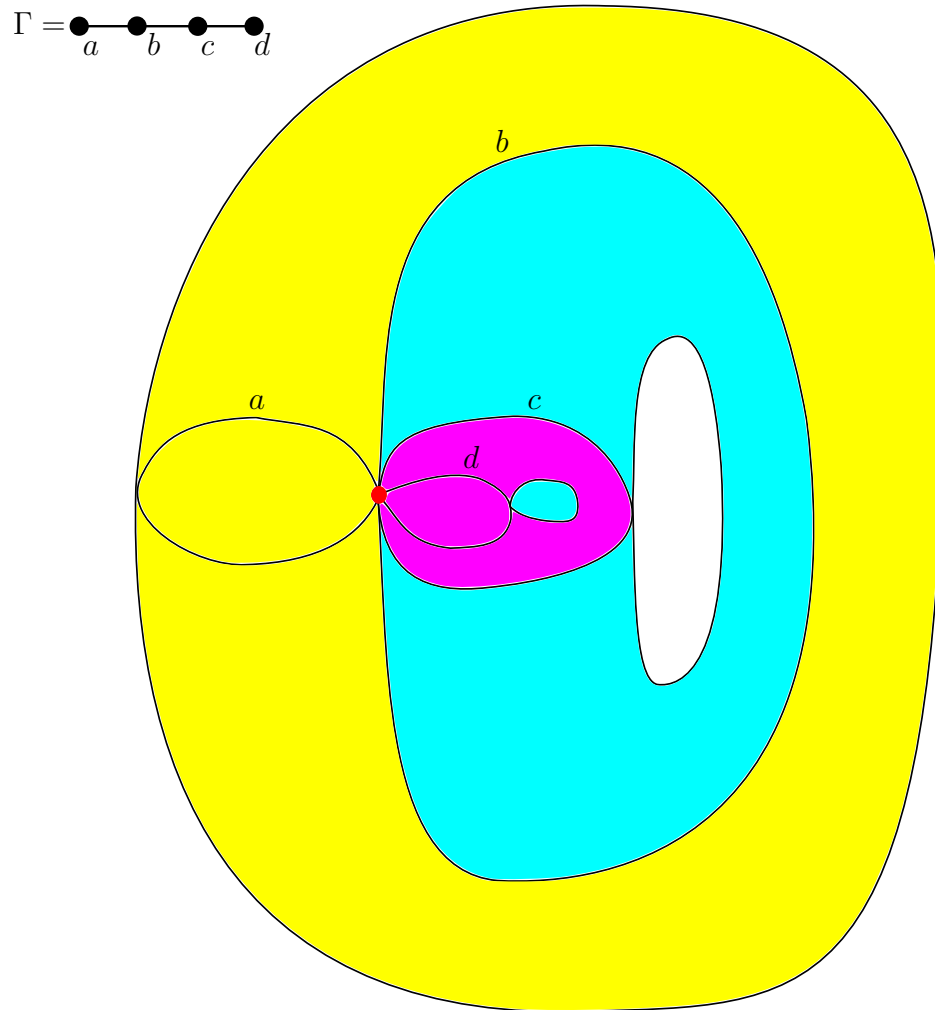


Figure 3.5 : An example of a Salvetti complex when Γ is a line segment of length 3. The base vertex is shown in red, and the four generators a , b , c , and d label circles wedged at the base vertex. There are three tori, shown in yellow, blue, and purple, respectively, corresponding to the three edges of Γ and the three pairs of commuting generators of A_Γ .

Chapter 4

The genus of a special cube complex

The goal of this chapter is to define a positive integer-valued invariant of special cube complex called the genus and prove Theorem 2.2. In section 4.1, we introduce a key technical tool which we call collapsing hyperplanes and use it to show that we can always collapse separating hyperplanes. In section 4.2, we show how nonseparating hyperplanes determine non-trivial 1-cocycles in special cube complexes. This observation illustrates an important relationship between the geometry of special cube complexes and homology, which we will exploit in this chapter in our study of the genus, and in later chapters in our study of automorphisms. In the last section, we define the genus of a special cube complex and a special group, and calculate the genus for several classes of groups. We then characterize genus zero (Proposition 4.3) and genus one (Theorem 4.1) special cube complexes in terms of their fundamental group, proving Theorem 2.2. We also deduce several corollaries, providing some examples of groups which are virtually special but not special (Corollary 4.6).

4.1 Collapsing hyperplanes

Definition 4.1. A hyperplane H is **separating** if $X \setminus H$ has more than one connected component. Otherwise, H is **non-separating**.

If H is separating and X is special, then $\overline{N_1(H)} \cong H \times [-1, 1]$. We will now describe a way of collapsing X along these product neighborhoods to obtain a NPC

cube complex with no separating hyperplanes. We first learned of the technique of collapsing hyperplanes in [31] and then adapted it to our setting.

Definition 4.2. Let $H \subset X$ be a separating hyperplane, and denote by $\overline{N_1}(H)$ its closed unit neighborhood. We define the **collapse X/H of X along H** to be the cube complex obtained by the identification $X/H = X/\{(x,t) \sim (x,s)\}$ where $(x,t), (x,s) \in N_1(H) \cong H \times [-1, 1]$. Let $\pi : X \rightarrow X/H$ denote the quotient map.

Proposition 4.1. *If X is special and H is separating, then the collapse X/H is special.*

Proof. The fact that $\overline{N_1}(H)$ is isomorphic to $H \times [-1, 1]$ implies that X/H has a cube complex structure. First we check that X/H is NPC. For this it suffices to check the Gromov link condition at each vertex. If a vertex does not meet $\overline{N_1}(H)$ then its link passes isometrically to the quotient, hence the link condition is still satisfied. If a vertex v_0 lies in $\overline{N_1}(H)$, then in the quotient v_0 is identified with exactly one other vertex v_1 which lies at the other end of an edge dual to H . Call e the edge joining v_0 and v_1 .

Denote the link of v_0 by $lk(v_0)$, and the full subcomplex of $lk(v_0)$ generated by cubes other than e which meet $\overline{N_1}(H)$ by $lk_H(v_0)$. Finally denote the full subcomplex generated by cubes in $lk(v_0)$ other than e by $lk_e(v_0)$. We similarly obtain complexes $lk(v_1)$, $lk_H(v_1)$ and $lk_e(v_1)$. Note that $lk_H(v_i)$ are exactly the edges in link of v_i which lie on the boundary of a cube containing e , for $i = 0, 1$. If m is the midpoint of e , then m is a vertex of H and $lk(m) \cong lk_H(v_0) \cong lk_H(v_1)$. Since H is NPC, the link $lk_H(v_i)$ is a flag simplicial complex, and hence a full subcomplex of $lk(v_i)$ and $lk_e(v_i)$. There are no monogons in the quotient because H does not self-intersect, and there are no bigons because X is NPC. Thus, in the quotient the link of the vertex corresponding

to the equivalence class of v_0, v_1 can be described as $lk_e(v_0) \amalg lk_e(v_1)$ identified along $lk_H(v_0) \cong lk_H(v_1)$. This is flag because it is made from two flag complexes glued along a full subcomplex.

No one-sided hyperplanes: Let $\pm[e_H]$ denote the equivalence class of parallel edges in X which are dual to H . Then in the quotient this class vanishes, and all other classes are preserved. Suppose the $e \parallel -e$ in X/H . Then $e, -e$ are dual to some hyperplane K and there is a path γ between the endpoints of e lying entirely within $N_1(K) \setminus K$. Let H_0 be the image of H under the collapse $\pi : X \rightarrow X/H$. Any path which meets H_0 has a lift to X , since the pre-image of any segment I lying in H_0 is a rectangle $I \times [-1, 1]$. Thus, $\pi^{-1}(K)$ is not 2-sided.

No self-intersection: Suppose K is a hyperplane in X which intersects itself in X/H . Then there are two squares in X with edges e, e' dual to K connected by an edge $e_0 \in [e_H]$. Then these two squares are opposite faces of a cube C containing e_0 as a dual edge, and e, e' extend to C to intersect in C . Thus K intersected itself in X .

No self-osculation: Suppose K is a hyperplane in X which self-osculates in X/H . Then there are two edges e, e' in X dual to K , and lying on opposite sides of an edge e_0 dual to H . If H and K do not intersect, then H is not separating. If they do intersect, then the fact that X is special implies that they meet in a square in X . Since it is not possible for K to self-intersect or self-osculate, there is a single square bounded on parallel sides by edges dual to K and on the other by H . It follows that after collapsing H , K does not self-osculate in the quotient.

No interosculation: Suppose K_1 and K_2 are hyperplanes of X which interosculate in the quotient. Note that as in the case of no-self-intersection, it is not possible for K_1 and K_2 to intersect in the quotient if they did not in X . Thus, K_1 and K_2 cross in X , and there are a pair of edges e_1 and e_2 , dual to K_1 and K_2 respectively, which lie at either ends of an edge e_0 dual to H . There are three cases, depending on whether or not K_1 and K_2 intersect or osculate H in X . If both K_1 and K_2 osculate H , then H does not separate. If exactly one of K_1 and K_2 intersects H , say K_1 , then K_1 and H cross in a square with boundary e_0 and e_1 . Then K_1 and K_2 interosculate in X , or they cross in a square at the other end of e_0 . It follows that under the collapse, no interosculation occurs. In the case where all three intersect, then in X there is a 3-cube containing e_0 , e_1 and e_2 and hence $\pi(K_1)$ and $\pi(K_2)$ cross in X/H .

Since none of the four hyperplane pathologies can occur in the quotient, X/H is NPC and special. □

Remark 4.1. Note that if K and H separate X then $\pi(K)$ still separates in X/H .

Definition 4.3. A special cube complex X is called **irreducible** if it has no separating hyperplanes. Otherwise X is **reducible**.

Proposition 4.2. *Every compact special cube complex X is homotopy equivalent to an irreducible compact special cube complex.*

Proof. An easy application of van-Kampen's theorem shows that collapsing separating hyperplanes in Proposition 4.1 induces an isomorphism on π_1 . Since both X and the quotient are NPC, they are each $K(\pi_1, 1)$'s, hence homotopy equivalent. By

compactness, there are only finitely many separating hyperplanes, and by Remark 4.1, we can collapse them in order. \square

4.2 The cohomology group $H^1(X)$

Let X be an NPC cube complex and suppose that every hyperplane is embedded and two-sided. If H is non-separating, then H defines a surjection $\phi_H: \pi_1(X) \rightarrow \mathbb{Z}$ as follows. First, choose an orientation on 1-cubes dual to H . This is possible since H is two-sided. For each 1-cube $e \in X$ define $\phi_H(e)$ to be the signed intersection of e with H and extend to 1-chains $C_1(X)$ by linearity. To see that ϕ_H is a cocycle, observe that the signed sum around any square, again by two-sidedness of H , is 0. Since H is non-separating, there exists a cycle in $X^{(1)}$ which meets H exactly once with positive orientation. For any $\alpha \in H_1(X)$, we define the intersection product $\alpha.H$ to be the integer $\phi_H(\alpha)$.

Combining the above observation with Proposition 4.1 we can characterize exactly when a special cube complex is CAT(0):

Corollary 4.1. *Suppose X is connected and special. Then the following are equivalent*

1. X is CAT(0).
2. $H_1(X) = 0$.
3. Every hyperplane is separating.

Proof. Recall that by Theorem 3.7, if X is NPC and simply connected, then X is CAT(0). If X is CAT(0) then $\pi_1(X)$ is trivial and hence $H_1(X)$ is as well. If X has a non-separating hyperplane then by the observation above $H^1(X)$ is non-trivial. Finally, suppose every hyperplane is separating. Any compact subset $K \subset$

X is contained in the closed unit neighborhoods of only finitely many hyperplanes hence Proposition 4.1 and Remark 4.1 imply that K can be collapsed to a point. In particular, $\pi_1(X)$ is trivial and hence X is CAT(0). \square

As a final corollary, we have the following curious observation about quasi-convex hierarchies for hyperbolic special groups.

Corollary 4.2. *If G is δ -hyperbolic and $G = \pi_1(X)$ for some compact special cube complex, then G has a quasi-convex hierarchy consisting only of HNN-extensions.*

Proof. Let H be a non-separating hyperplane. Then $\pi_1(H)$ is quasi-convex in G , and $G \cong \pi_1(X|H)*_{\pi_1(H)}$. Now collapse separating hyperplanes in $X|H$ and repeat. Note that H^+ and H^- are separating in X/H . Eventually we will end up with a complex which only has separating hyperplanes, since the total number of cubes decreases every time we split along hyperplanes and collapse. \square

4.3 The genus of a special group

As we saw in the previous section, each non-separating hyperplane of a special cube complex X contributes a free factor to $H^1(X)$, but in general these free factors may not be distinct. For example, if K_1 and K_2 are two disjoint non-separating hyperplanes such that $K_1 \cup K_2$ separates X , then every homology class which meets K_1 also meets K_2 and with the same algebraic intersection, hence $\phi_{K_1} = \phi_{K_2}$. Based on this observation we have the following

Definition 4.4. Let X be special. The **genus** $g(X)$ is the maximum number of disjoint hyperplanes in X whose union does not separate. If no maximum exists we say $g(X) = \infty$. If G is the fundamental group of a special cube complex, we define

the genus

$$g(G) = \sup\{g(X) : X \text{ is special and } \pi_1(X) = G\}.$$

The definition is motivated by the classical definition genus of compact surface: namely, the largest number of disjoint simple closed curves whose union does not disconnect the surface. The next proposition lists some properties of the genus.

Proposition 4.3. *Let X be a special cube complex (resp. special group). The genus enjoys the following properties:*

1. $g(X) \leq \text{rk}(H_1(X))$. In particular, $g(X)$ is finite whenever X is compact (resp. finitely generated).
2. $g(X) = 0$ if and only if X is $\text{CAT}(0)$ (resp. $X = \{1\}$).
3. If $g(X) = n$, then $\pi_1(X)$ surjects onto the free group F_n .

Proof. Let X be a special cube complex. If K_1 and K_2 are disjoint and do not separate, there are homology classes γ_1 and γ_2 in $H_1(X)$ such that $K_i \cdot \gamma_j = \delta_{ij}$ for $i, j = 1, 2$. Hence the K_i correspond to distinct free factors of $H_1(X)$. This proves (1). Property (2) follows directly from Corollary 4.1. For (3), let H_1, \dots, H_n be disjoint hyperplanes in X whose union does not separate. By specialness, we can choose closed neighborhoods $\overline{N}(H_i)$ around each hyperplane such that $\overline{N}(H_i) \cong H \times [-\epsilon, \epsilon]$ for some $\epsilon > 0$, and for $i \neq j$, the neighborhoods are disjoint: $\overline{N}(H_i) \cap \overline{N}(H_j) = \emptyset$. We now define a surjective map F from X to a wedge of n circles $R_n = \bigvee_{i=1}^n S^1$, as follows. Label the circles of R_n as S_i^1 for $1 \leq i \leq n$ and map $\overline{N}(H_i)$ to S_i^1 by first projecting onto its dual edge factor:

$$\overline{N}(H_i) \cong H \times [-\epsilon, \epsilon] \rightarrow [-\epsilon, \epsilon],$$

and then mapping $[-\epsilon, \epsilon]$ to S_i^1 by a degree one map. Finally, map the complement of $\cup_i \overline{N}(H_i)$ to the wedge point in R_n . By construction F is surjective, and since $X \setminus \cup_i \overline{N}(H_i)$ is connected, the homotopy fibers of F are all connected. It follows that F induces a surjection on fundamental groups, as desired. \square

We calculate the genus of some basic examples of special groups:

Example 4.1. $g(F_n) = n$.

Take the standard rose R_n as a cube complex with $\pi_1 = F_n$. Then $g(R_n) = n$ and property (1) implies this is best possible. In fact, any graph with $\pi_1 = F_n$ works.

Example 4.2. $g(\mathbb{Z}^n) = 1$.

This follows from the fact that if $g(G) = n$ then G surjects onto F_n by property (3).

Example 4.3. $g(\pi_1(\Sigma_\gamma)) = \gamma$, where Σ_γ is the closed surface of classical genus γ .

Note that this is not entirely obvious from the definition, since there may very well be high-dimensional cube complexes with the same fundamental group as Σ_γ . We observe, however, that if K_1, \dots, K_n are disjoint hyperplanes of X , then $\phi_{K_i} \cup \phi_{K_j} = 0 \in H^2(X)$ for all i, j . Hence, the ϕ_{K_i} span a Lagrangian subspace of $H^1(X) = H^1(\Sigma_\gamma)$. The maximal possible dimension of such a subspace is of course γ , and it is not hard to construct explicit 2-dimensional cube complex structures on Σ_γ which realize this maximum.

Definition 4.5. Hyperplanes K and L are called **parallel** if K and L do not meet.

Theorem 4.1. *Let X be compact special. Then $g(X) = 1$ if and only if $\pi_1(X) = \mathbb{Z}^n$.*

Proof. Without loss of generality, we can assume X is irreducible by Proposition 4.2, since collapsing separating hyperplanes does not change the genus. Let K be a

hyperplane of X . Since X has genus one, we know that if L is any hyperplane parallel to K , then $K \cup L$ separates. The idea is to imitate the proof of Proposition 4.1 by collapsing all hyperplanes parallel to K , while maintaining non-positive curvature and specialness. Define an equivalence relation on hyperplanes as follows: $K \sim L$ if K and L are parallel. As defined, this is just a symmetric relation, but we will consider the equivalence relation \sim^* that it generates and say that if $K \sim^* L$ then K and L are **ultra-parallel**. If K and L are ultra-parallel and X has genus one, then every minimal combinatorial loop which meets K once also meets L once. Moreover, we note that if K and L are ultra-parallel but not parallel, then K and L intersect.

Let K_1, \dots, K_m be the set of hyperplanes other than H which are ultra-parallel to H . We want to show that result of collapsing all of K_1, \dots, K_m is NPC and special. We remark that it may be the case that some subset of collapses fails to be special. The important point is that the full collapse is special and homotopy equivalent to X . First we need a lemma which implies that we can collapse hyperplanes at all.

Lemma 4.1. *If $K, L \subset X$ are distinct parallel hyperplanes, then $\overline{N_1}(K) \cong K \times [-1, 1]$ and $\overline{N_1}(L) \cong L \times [-1, 1]$.*

Proof. The lemma is symmetric in K and L . If $\overline{N_1}(K)$ is not embedded, then $K \cup L$ does not separate. □

From the lemma, if L is a hyperplane in an ultra-parallelism class $[H]$, $H \neq L$, we know that $\overline{N_1}(L)$ is embedded. A slight issue arises when we repeatedly collapse hyperplanes. Namely, if we collapse all the hyperplanes in $[H]$ which are actually parallel to L , then $\overline{N_1}(L)$ ceases to be embedded. However, by the next lemma, we can always collapse in such a way that every hyperplane neighborhood is embedded.

Lemma 4.2. *Given an equivalence class $[H]$ with $|[H]| \geq 2$, there always exists a sequence of collapses in which, at each stage, the remaining hyperplanes have embedded unit neighborhoods.*

Proof. Given a collection of hyperplanes \mathcal{H} in X , let $\Delta(\mathcal{H})$ be the graph obtained in the following way. The vertices of $\Delta(\mathcal{H})$ will be the elements of \mathcal{H} , and two vertices are connected if their corresponding hyperplanes are disjoint. If $\mathcal{H} = [H]$ is an ultra-parallelism class, then $\Delta([H])$ is connected. Suppose X' is obtained from X by collapsing a hyperplane L and let $[H]'$ be the image of $[H]$ in X' . Observe that $\Delta([H]')$ can be obtained from $\Delta([H])$ by deleting the vertex corresponding to L , and all of its incident edges. In particular, if $L \notin [H]$, then $\Delta([H]) = \Delta([H]')$.

The lemma, translated in terms of $\Delta([H])$, states that there is a sequence of vertex deletions such that, at each stage the complement is connected. The latter follows by induction and the well-known graph theoretic result that a connected graph always has at least two vertices which are not cut vertices, *i.e.* they do not disconnect the graph. \square

From the two previous lemmas, we know that if $[H]$ is an ultra-parallelism class, we can find $H \in [H]$ and an ordering of the hyperplanes $K_1, \dots, K_m \in [H] \setminus \{H\}$ such that when we collapse each K_i in order, the result at each stage will be an NPC cube complex homotopy equivalent to X . We remark that as before, two hyperplanes cannot cross in the quotient if they did not originally. The proof that in the quotient every hyperplane is two-sided and that no hyperplane self-intersects is exactly the same as above, and we do not need the hypothesis that X has genus one. The next lemma implies no self-osculation occurs in the quotient.

Lemma 4.3. *Fix an ultra-parallelism class $[H]$ which has cardinality at least 2. Then*

for all $L \in [H]$, and for every two vertices v_1 and v_2 in $\overline{N_1(L)}$ lying on the same side of L , there does not exist a combinatorial geodesic from v_1 to v_2 which crosses some edge dual to a hyperplane parallel to L .

Proof. We claim that either $g(X) \geq 2$ or no such path exists for any $L \in [H]$. Define \mathcal{S} to be the set of combinatorial geodesics γ such that there exists some $L \in [H]$ and satisfying

1. The geodesic γ does not cross L
2. The endpoints of γ lie on the same side of some hyperplane L
3. The geodesic γ crosses some hyperplane K parallel to L

We will show that if $g(X) = 1$ then \mathcal{S} is empty, by induction on a least length counterexample. For reasons of parity, we consider two base cases, when the combinatorial length $\ell(\gamma) = 1$ and $\ell(\gamma) = 2$. Let v_1 and v_2 be the endpoints of γ . If $\ell(\gamma) = 1$ then by condition (3), the single hyperplane K which γ crosses must be parallel to L . Hence, we can complete γ to a loop γ' by adding a path between v_1 and v_2 in $N_1(L)$ which meets K exactly once. Since L is non-separating, we conclude that $g(X) \geq 2$, a contradiction. If $\ell(\gamma) = 2$, then γ crosses two hyperplanes K_1 and K_2 . If $K_1 \neq K_2$, then as in the previous case, since one of K_1 and K_2 is parallel to L , we conclude that $g(X) \geq 2$. If $K_1 = K_2$, then K_1 is parallel to L and we must consider the orientations with which γ crosses K_1 . If γ crosses K_1 with the same orientation each time, then we again conclude that $g(X) \geq 2$. If, on the other hand, γ crosses K_1 first with one orientation, then the opposite we invoke the fact that K_1 does not directly self-osculate to conclude that γ backtracks, and hence is not a combinatorial geodesic.

Now suppose γ is a least length counterexample of length $n \geq 3$ occurring along a hyperplane $L \in [H]$. Let v_1 and v_2 be the endpoints of γ . Then γ crosses a sequence of

edges dual to hyperplanes $K_{i_1}^{\epsilon_1}, \dots, K_{i_r}^{\epsilon_r}$, where $\epsilon_j = \pm 1$ depending on the orientation with which γ crosses K_{i_j} .

Claim 1: L is parallel to K_{i_1} .

Otherwise, by no interosculation, there is a square with corner v_1 where K_{i_1} and L cross. Then both endpoints of the first edge e_1 of γ lie in $N_1(L)$. Writing $\gamma = e_1\gamma'$, we see that γ' is a shorter length counterexample.

Claim 2: γ crosses K_{i_1} algebraically (*i.e.* counted with sign) 0 times.

If not, then we can complete γ to a loop γ' as above, which crosses K_{i_1} algebraically non-zero times and crosses L geometrically 0 times. Since L is assumed non-separating, we conclude that $g(X) \geq 2$, a contradiction.

At this point, we can assume that γ must cross K_{i_1} algebraically 0 times. If we consider the sequence of crossings we can find an innermost pair with opposite sign, *i.e.* a subpath $e_1\alpha e_2 \subset \gamma$ such that e_1 and e_2 are both dual to K_{i_1} but with opposite orientation, and such that α does not cross K_{i_1} . Clearly α is not empty, otherwise γ would have backtracking.

We claim that either γ is not a combinatorial geodesic, or α is a shorter length counterexample. If α crosses some hyperplane parallel to K_{i_1} then $\alpha \in \mathcal{S}$, since it connects two vertices on the same side K_{i_1} , is a combinatorial geodesic since it is a subpath of γ and satisfies $\ell(\alpha) \leq \ell(\gamma) - 2$. It is therefore a shorter element of \mathcal{S} , contradicting our assumption on γ . Otherwise, K_{i_1} meets every hyperplane crossed by α . In this case however, we can replace γ by a combinatorially isotopic path with backtracking. To see this, consider the first edge f_1 of α . By no interosculation we

can find a square C bounded at a corner by e_1 and f_1 . We therefore replace $e_1\alpha$ by $\alpha' = f'_1e'_1\alpha''$, where e'_1 and f'_1 are the opposite edges of C , and α'' is the remainder of α after f_1 . Continuing in this way, we replace $e_1\alpha e_2$ by $\alpha_0e_0\bar{e}_0$, where e_0 is dual to K_{i_1} . Since γ was isotopic to a path with backtracking, we conclude that γ was not combinatorially geodesic, contradicting our assumption. Therefore \mathcal{S} is empty, as desired. \square

No self-osculation: Suppose that a hyperplane L directly self-osculates after collapsing some collection of the K_i . Then there is a path γ consisting of edges dual to some subcollection K_1, \dots, K_m , which connects two vertices lying on the same side of $\bar{N}_1(L)$. If L intersects each of K_1, \dots, K_m , then no self-osculation occurs in the quotient. Otherwise, L is parallel to some K_j . But then Lemma 4.3 implies that this is impossible.

At this point we have checked that after collapsing each of the K_i , the resulting space immerses in a Salvetti complex. For a local isometry, we need to further check that no interosculation occurs.


No interosculation: By the remark about intersecting hyperplanes above, we need only consider the case where hyperplanes L_1 and L_2 intersect in X and osculate in the quotient. In this case there is a path γ dual to hyperplanes K_1, \dots, K_m which are ultra-parallel to H and edges f_1 and f_2 dual to L_1 and L_2 , respectively at either end of γ . Moreover, L_1 and L_2 meet in some other square. If at least one of L_1 and L_2 intersects all of the K_i , then by no interosculation of X , after collapsing there is a square containing f_1 and f_2 . Finally, we have the case where both L_1 and L_2 are parallel to one of the K_i . There are three cases depending on which sides of L_1

and L_2 that γ connects. See Figure 4.1 for a schematic. Note that in this case all hyperplanes have embedded closed unit neighborhoods.


At most one of L_1 and L_2 is the chosen hyperplane H . In the case that neither L_1 nor L_2 is H then L_1 and L_2 are both eventually collapsed and no interosculation occurs in the quotient. Then assume that $L_1 = H$. In either case (1), (2), or (3) we find that no interosculation occurs in the quotient and either H directly self-oscultates, which we have already shown is impossible, or H indirectly osculates which does not contradict specialness of the quotient.

Let X' denote the result of collapsing all hyperplanes ultra-parallel to H , this is still NPC by Lemma 4.2. To see that X' has genus one, note that if \mathcal{H} is the set of hyperplanes of X , then $\mathcal{H} \setminus \{K_i\}$ is in one-to-one correspondence with the set of hyperplanes of X' and that if $L_1 \notin \{K_i\}$ and $L_2 \notin \{K_i\}$ are disjoint and separate X , then their images under the collapse are disjoint and separate X' . In X' , the image of H meets every other hyperplane.


To finish the proof, we proceed as follows. Choose some ultra-parallelism class $[H]$. If $[H]$ is a singleton, then H meets every hyperplane of X and we do not collapse anything. If $|[H]| \geq 2$, then by Lemma 4.2 we can find $H \in [H]$ such that if K_1, \dots, K_m are hyperplanes ultra-parallel to H , we can collapse the K_i in some order such that the resulting cube complex X' is NPC special, has genus one, and every hyperplane of X' meets the image of H . Now choose another ultra-parallelism class and repeat. Since X has only finitely many hyperplanes, in the end we obtain a special cube complex Y homotopy equivalent to X with the property that any two hyperplanes meet. It follows that the corresponding Salvetti \mathbb{S}_Y is a torus, and hence $\pi_1(X) = \pi_1(Y) \leq \mathbb{Z}^n$ for some n , by Lemma 3.9. In fact, since Y is compact, the



Case (1)



Case (2)



Case (3)

Figure 4.1 : Possible configurations resulting in interosculation in the quotient.

map $f_Y : Y \rightarrow \mathbb{S}_Y$ is a surjective, combinatorial local isometry, hence must be a finite covering. \square

Remark 4.2. If X is non-compact but has finitely many hyperplanes, the same proof as above works. However, it may be the case that the quotient complex Y is non-compact. Then the characteristic map $f_Y : Y \rightarrow \mathbb{S}_Y$ will be a surjective, combinatorial local isometry, but all we can conclude is that it is a covering map. Thus the image of $\pi_1(Y) \hookrightarrow \pi_1(\mathbb{S}_Y) \cong \mathbb{Z}^n$ may be a subgroup of infinite index.

As an immediate corollary we obtain:

Corollary 4.3. *If G is non-abelian then every compact special cube complex with $\pi_1(X) = G$ satisfies $g(X) \geq 2$.*

We can extend this result to the finite-dimensional case to give a proof of Theorem 2.1, stated in Chapter 2:

Corollary 4.4. *If X is special and finite-dimensional then either $\pi_1(X)$ is abelian or surjects onto F_2 .*

Proof. The compact case follows from Theorem 4.1. In the non-compact case, we observe that if there are infinitely many non-separating hyperplanes, then since X is finite-dimensional there are infinitely many disjoint non-separating hyperplanes. Take a loop γ which meets some non-separating hyperplane exactly once. Then γ only intersects finitely many non-separating hyperplanes. Thus we must have at least two disjoint non-separating hyperplanes whose union does not disconnect X and $g(X) \geq 2$. Otherwise, there are only finitely many non-separating hyperplanes and if $g(X) = 1$, we can apply the procedure of Theorem 4.1 to these hyperplanes. We obtain a homotopy equivalent special cube complex Y in which all non-separating hyperplanes meet.

Suppose γ_1 and γ_2 are two loops in the 1-skeleton $Y^{(1)}$ based a point $p \in Y$. Choose a compact subcomplex K containing $\gamma_1 \cup \gamma_2$. By collapsing all the separating hyperplanes which meet K , we get a complex Y' homotopy equivalent to Y , and in which the images γ'_1 and γ'_2 , of γ_1 and γ_2 respectively, only cross non-separating hyperplanes. Since the non-separating hyperplanes of Y' all cross we conclude that the homotopy classes $[\gamma'_1]$ and $[\gamma'_2]$ commute in $\pi_1(Y')$. We conclude that $\pi_1(Y) = \pi_1(X)$ is abelian. \square

The characterization of genus 1 special groups also has some immediate corollaries for groups with genus ≥ 2 .

Corollary 4.5. *Let G be a non-abelian special group. Then*

1. $rk(H_1(G)) \geq 2$.
2. G retracts onto F_2 .
3. The rank of H_1 grows at least linearly in finite index subgroups.
4. The growth of finite index subgroups in G is at least superexponential ($\succeq ne^{n \log(n)-n}$) in index.

Moreover, if G is not virtually abelian but is virtually special, (3) and (4) still hold.

Proof. Statements (1) and (2) follow from Corollary 4.4 and the fact that any surjection onto a free group splits. Statements (3) and (4) follow from the corresponding result for F_2 . The growth of finite index subgroups in a free group is due to Hall [47]. \square

Note that from (2) we get a quick proof of the Tits alternative for virtually special groups: either G contains a non-abelian free group, or it is virtually abelian. Indeed,

if G is virtual special, then there exists a finite index subgroup $G' \leq G$ which is special. By Corollary 4.4, either G' is abelian or it surjects onto F_2 . In the case of the latter, since F_2 is free, the surjection splits, hence G' contains a free subgroup F_2 , as does G .

The genus also restricts which groups can arise as fundamental groups of special cube complexes. Recall that if Σ_g is the closed surface of genus g , the group of orientation preserving diffeomorphisms of Σ_g is denoted $\text{Diff}^+(\Sigma_g)$. The mapping class group Mod_g is defined to be $\pi_0(\text{Diff}^+(\Sigma_g))$, the group of connected components of $\text{Diff}^+(\Sigma_g)$. Thus, two diffeomorphisms are identified if they are isotopic. The abelianization map $\pi_1(\Sigma_g) \rightarrow \mathbb{Z}^{2g}$ induces a surjective map $\text{Mod}_g \rightarrow \text{Sp}_{2g}(\mathbb{Z})$. We say $\phi \in \text{Mod}_g$ has **full rank** if the action of ϕ on $H_1(\Sigma_g)$ has finite quotient.

A theorem of Thurston and Nielsen states that every mapping class $\phi \in \text{Mod}_g$ falls into one of three categories: finite order, reducible, or pseudo-Anosov (see [35], Theorem 13.2). **Reducible** means that ϕ has a lift to Diff^+ which fixes some 1-submanifold. **Pseudo-Anosov** means that ϕ does not preserve any conjugacy class in $\pi_1(\Sigma_g)$. Finite order and reducible are not mutually exclusive, but both are disjoint from pseudo-Anosov. Thurston showed that if ϕ is pseudo-Anosov, then the mapping torus M_ϕ corresponding to any lift of ϕ to Diff^+ supports a constant curvature-(−1) Riemannian metric ([35], Theorem 13.4(3)). This construction provides many examples of hyperbolic 3-manifolds.

Corollary 4.6. *If G is any one of the following, then G is virtually compact special but not compact special:*

1. *virtually abelian but not abelian.*
2. *the fundamental group of a hyperbolic $\mathbb{Q}HS^3$.*

3. $\pi_1(M_\phi)$ where ϕ is pseudo-Anosov and has full rank.

Proof. 1. If G is virtually abelian and compact special then it does not contain F_2 , hence $g(G) = 1$. Theorem 4.1 then implies that $G \cong \mathbb{Z}^n$ for some n .

2. Results of Agol ([18], Theorems 1.1 and 9.3) imply that every closed hyperbolic 3-manifold group is virtually compact special. A hyperbolic rational homology 3-sphere has $\text{rk}(H_1) = 0$, hence cannot be compact special. We remark that it was already known that $\text{rk}(H_1) \geq 1$ for compact special cube complexes.

3. A straightforward application of Mayer-Vietoris shows that $\text{rk}(H_1(M_\phi)) = 1$. Thurston's theorem and Agol's theorem show that $\pi_1(M_\phi)$ is virtually compact special, but obviously $\pi_1(M_\phi)$ contains F_2 and hence is not virtually abelian.

□

Chapter 5

Automorphisms of NPC cube complexes

In this chapter, we investigate automorphisms of NPC cube complexes, and in particular, the action of automorphisms on first homology. In section 5.1, we provide a criterion on a special cube complex X which implies that any automorphism acts non-trivially on first homology. In the next chapter, we will verify this criterion for the class of blow-ups of Salvetti complexes in order to show that the Torelli subgroup for a raag is torsion-free. In section 5.2, we consider the action on first homology after lifting automorphisms to covers, and prove a general result about automorphisms of groups which surject onto the free group F_2 (cf. Theorem 2.4). We also show that finite-order outer automorphisms of a virtually special hyperbolic group G can always be realized as a cube complex automorphism of an NPC cube complex X with $\pi_1(X) = G$. To end the chapter, we provide an infinite family of examples of virtually special hyperbolic 3-manifold groups each of whose outer automorphism group acts trivially on first homology.

5.1 Action of automorphisms on homology

In this section we investigate the action of cube complex automorphisms on homology. Our starting point comes from two well-known examples:

1. Every non-identity torsion element of $\text{Out}(F_n)$ acts nontrivially on $H_1(F_n)$.
2. Every non-identity torsion element of Mod_g acts nontrivially on $H_1(\Sigma_g)$.

Observe that both F_n and $\pi_1(\Sigma_g)$ are compact special. This motivates the following

Question 5.1. *Suppose G is compact special. Does every non-identity finite order element of $\text{Out}(G)$ act non-trivially on $H_1(G)$?*

The first step is to ensure that $H_1(G)$ is non-trivial, but as we have seen, this is satisfied as soon as $G \neq 1$. We do not propose to answer this question fully here, but we will generalize the results above to cubulated hyperbolic and right-angled Artin groups. Our strategy will be two-fold. First, realize elements of $\text{Out}(G)$ as automorphisms of compact special cube complexes with fundamental group G . Second, use the geometry of the compact special cube complex to show certain automorphisms act non-trivially on homology. A question closely related to the one above is thus

Question 5.2. *When does an automorphism of a compact special cube complex X act non-trivially on $H_1(X)$?*

Since there are compact CAT(0) cube complexes with arbitrarily large (finite) automorphism groups, the answer to this question is not, in general, “always.” Nevertheless, we will provide circumstances under which every automorphism acts non-trivially on first homology.

The following proposition gives a useful criterion to guarantee that every automorphism acts non-trivially on $H_1(X)$.

Proposition 5.1. *Let X be compact special and $f : X \rightarrow X$ an automorphism. Suppose X satisfies the following three conditions for hyperplanes K_1 and K_2*

1. *If $K_1 \cap K_2 \neq \emptyset$, there exists $\alpha \in H_1(X)$ such that $\alpha.K_1 \neq \alpha.K_2$.*
2. *If $K_1 \cap K_2 = \emptyset$ and $K_1 \cup K_2$ separate X , every component of $X \setminus (K_1 \cup K_2)$ contains a non-separating hyperplane which does not meet K_1 or K_2 .*

3. If $f(K_1) = K_1$ for all K_1 then f is the identity.

Then if f acts trivially on $H_1(X)$, f is the identity.

Proof. Let $f : X \rightarrow X$ be an automorphism and suppose $f_* : H_1(X) \rightarrow H_1(X)$ is the identity. Assume for contradiction that f is not the identity. The order of f is finite, and by passing to a power we may assume it is a prime p . By condition (1) we know that for any hyperplane K_0 , the image $f(K_0)$ does not meet K_0 transversely, or else f_* would not be the identity. Then for every hyperplane K_0 , we have that $f(K_0) \cap K_0 = \emptyset$ or $f(K_0) = K_0$. Note that it is not possible for every hyperplane to be mapped to itself setwise without being the identity. Therefore, either $f = \text{id}_X$ and we are done or there exists some hyperplane K_0 such that $f(K_0) \cap K_0 = \emptyset$.

As we observed, the images $K_0 = f^0(K_0), K_1 = f(K_0), \dots, K_{p-1} = f^{p-1}(K_0)$ are all disjoint and f permutes the components of $X \setminus \cup_{i=0}^{p-1} N_1(K_i)$. We can assume that every cycle which meets K_0 algebraically non-trivially also meets each K_i with the same intersection. In particular, since K_0 is non-separating, there is a cycle which meets K_0 geometrically once. We conclude that any pair K_i, K_j with $0 \leq i \neq j \leq p-1$ separate X . We may also assume, after reordering and passing to a power, that for each i , $0 \leq i \leq p-1$, one of the components of $X \setminus (K_i \cup K_{i+1})$ does not contain any K_j , where $j \neq i, i+1$ and $K_p = K_0$. For if K_j lies in some component Y of $X \setminus (K_i \cup K_{i+1})$, then it does not meet K_i or K_{i+1} and it must separate K_i from K_{i+1} . Otherwise there is a path in X which meets all of the K_i except K_j geometrically once. Then we can assume that the K_i are permuted in order, and that one component of $X \setminus (K_i \cup K_{i+1})$ does not contain any other K_j .

By condition (2), we know that each component of $X \setminus (K_i \cup K_{i+1})$ contains a non-separating hyperplane which does not meet K_i or K_{i+1} . Choose one such hyperplane L_i in the component of $X \setminus (K_i \cup K_{i+1})$ which does not contain any other K_j . Then

since L_i does not meet the K_i or K_{i+1} , it is a hyperplane of X proper. Since it does not separate, there is a non-trivial cycle α_i contained in this component satisfying $\alpha_i.L_i = 1$. Then by naturality of the Kronecker pairing

$$1 = \alpha_i.L_i = \phi_{L_i}(\alpha_i) = f^*(\phi_{L_i})(f_*(\alpha_i)) = \phi_{f(L_i)}(\alpha_i) = \alpha_i.f(L_i) = 0.$$

This contradiction implies that f takes every hyperplane to itself, hence must be the identity by condition (3). \square

5.2 Passing to covers

Even if we cannot guarantee that every automorphism acts non-trivially on $H_1(X)$, in some cases it may be possible to pass to a cover and lift our automorphism so that it acts non-trivially on the homology of the cover. In fact, this is always the case. The result follows from the next proposition, which although not difficult to prove, does not seem to appear anywhere in the literature. We record it here for posterity.

Proposition 5.2. *Let G any finitely generated group which surjects onto F_2 , and let $\phi \in \text{Out}(G)$ have finite order. If ϕ_* is the induced map on the abelianization $H_1(G)$, then there exists a finite index normal subgroup $N \trianglelefteq G$ and an induced outer automorphism $\widehat{\phi}$ of N such that $\widehat{\phi}_*$ does not act trivially on $H_1(N)$.*

Remark 5.1. Informally, if G surjects onto F_2 , any finite order outer automorphism acts non-trivially on the abelianization of some finite index subgroup.

Proof. Let $\pi : G \rightarrow F_2$ be a surjection. Since G is finitely generated, $H_1(G)$ is a finitely generated abelian group say with first Betti number b_1 . Choose $d \gg b_1$ and find some finite-index normal subgroup $K' < F_2$ such that $\text{rk}(H_1(K')) = d$. Then $K = \pi^{-1}(K')$ is normal and of finite index in G . Finally let f be a lift of ϕ to $\text{Aut}(G)$

and define $N = K \cap \cdots \cap f^{n-1}(K)$ where n is the order of ϕ . It follows that $f(N) = N$, hence f induces an automorphism $\widehat{f} : N \rightarrow N$. We claim that after postcomposing f with conjugation by an element of G , the induced map \widehat{f} does not act trivially on $H_1(N)$.

This is just a straightforward application of the transfer homomorphism ([25] for a definition). If \widehat{f}_* acts non-trivially on $H_1(N)$, we are done. Otherwise, by the transfer, since $\text{rk}(H_1(N)) = d > b_1$, there exists $g \in G$ such that conjugation by g acts non-trivially on $H_1(N)$. Denote by c_g the automorphism of N induced by conjugation by g . Then $\widehat{f}' = c_g \circ \widehat{f}$ acts non-trivially on $H_1(N)$. Note that \widehat{f}' also has finite order. Setting $\widehat{\phi} = [\widehat{f}'] \in \text{Out}(N)$ finishes the proof. \square

Corollary 5.1. *Suppose G is finitely generated and virtually compact special. Then any finite order outer automorphism of G has a lift which acts non-trivially on the abelianization of some finite index subgroup.*

Proof. If G is not virtually abelian, then G virtually surjects onto F_2 and the result follows by Proposition 5.2. If G is virtually abelian but not abelian, then G contains \mathbb{Z}^n as a finite index normal subgroup, for some n . In this case it is not hard to show that $\text{rk}(H_1(G)) < n$. Thus the same proof as in Proposition 5.2 works here, too. If $G = \mathbb{Z}^n$, then $\text{Aut}(G) = \text{GL}_n(\mathbb{Z})$ and $\pi_1(G) = H_1(G)$, hence any automorphism clearly acts non-trivially on first homology. \square

In fact, when G is δ -hyperbolic we can say a little more. Let X be a compact NPC cube complex with $G = \pi_1(X)$ δ -hyperbolic and not virtually cyclic. By Theorem 1.1. of [18], X is virtually compact special. Note that the center $Z(G)$ is trivial. Let $\phi : G \rightarrow G$ be a finite-order outer automorphism of G . From the exact sequence

$$1 \rightarrow G \rightarrow \text{Aut}(G) \rightarrow \text{Out}(G) \rightarrow 1,$$

we can consider the extension given by pulling back the subgroup $\langle \phi \rangle$:

$$1 \rightarrow G \rightarrow E_\phi \rightarrow \langle \phi \rangle \rightarrow 1.$$

Since G is cubulated hyperbolic, E_ϕ is virtually cubulated hyperbolic. Hence by Lemma 7.15 of [44], we obtain a proper cocompact action of E_ϕ on a CAT(0) cube complex \tilde{Y} . Since G is torsion-free, the action of $G < E_\phi$ on \tilde{Y} is free. We therefore obtain a quotient Y with $\pi_1(Y) \cong G$, and a finite order automorphism $f : Y \rightarrow Y$ corresponding to ϕ . We have just proven

Proposition 5.3. *If G is cubulated and hyperbolic, every finite order element of $\text{Out}(G)$ can be realized as an automorphism of an NPC cube complex Y with $\pi_1(Y) = G$.*

We conclude this section with an application of Corollary 5.1 in the context of hyperbolic 3-manifolds. In order to describe the examples, however, we need to review a little knot theory in dimension 3.

Definition 5.1. An **oriented knot** K is a smooth embedding of the oriented circle S^1 in the 3-sphere S^3 . The knot K is said to be **hyperbolic** if its complement in S^3 supports a constant curvature-(-1) Riemannian metric. A oriented knot K is called **negative amphichiral** if it is isotopic to its reverse mirror-image *i.e.* change all crossings and reverse orientation.

Example 5.1. There exist infinitely many hyperbolic 3-manifold groups $\{G_i\}$ such that G_i is virtually special, $\text{Out}(G_i) \neq 1$, but $\text{Out}(G_i)$ acts trivially on $H_*(G_i)$.

Let K be a negative amphichiral knot such as the Figure 8 knot. If $M = S^3 \setminus K$ is the knot complement then there is an orientation-preserving involution $\sigma : M \rightarrow M$, induced by the amphichirality. If $T \subset M$ is a boundary parallel torus, then $\sigma|_{T^2}$ is just

the hyperelliptic involution on T^2 . In particular, σ sends every slope p/q to $-p/-q$. Let $M_{p/q}$ be the result of p/q -surgery on K . Since $M_{p/q} = M_{-p/-q}$, the action of σ on M extends to an involution $\widehat{\sigma} : M_{p/q} \rightarrow M_{p/q}$. If K is hyperbolic (*e.g.* the figure 8 knot), then a theorem of Thurston states that for all but finitely many slopes $M_{p/q}$ will also be hyperbolic 3-manifolds. Taking $p = 1$, a Mayer-Vietoris computation implies that $M_{1/q}$ will be an integral homology sphere, and will also be hyperbolic for infinitely many q . Then $\widehat{\sigma}$ is an automorphism of $\pi_1(M_{1/q})$ which is not inner since it has finite order, and non-trivial since it inverts the meridian and longitude of the knot. Moreover, as $\widehat{\sigma}$ is orientation-preserving, $\widehat{\sigma}_* : H_*(M_{1/q}) \rightarrow H_*(M_{1/q})$ is actually the identity.

To see that infinitely many pairs are non-isomorphic, observe that they are finite volume hyperbolic 3-manifolds, and therefore isometric if and only if their fundamental groups are isomorphic by Mostow rigidity. On the other hand, taking $q \rightarrow \infty$ implies $\text{vol}(M_{1/q}) \nearrow \text{vol}(S^3 \setminus K)$, by Thurston's theorem. As the volume of $S^3 \setminus K$ is strictly larger than that of any $M_{1/q}$, we can find infinitely many q with distinct volumes, hence which are not isometric. Thus, this construction gives infinitely many integral homology spheres $\mathbb{Z}HS^3$'s whose fundamental groups have non-trivial outer automorphism groups. By Agol's theorem, all of these virtually compact special, implying that Corollary 5.1 is best possible.

Chapter 6

Applications to right-angled Artin groups

In this chapter, we present applications of the results of the previous chapter to automorphisms of right-angled Artin groups. The main result is to give a new proof that the Torelli subgroup for a right-angled Artin group is torsion-free (cf. Theorem 2.3). The outline of the chapter is as follows. First, in section 6.1, we define the Torelli subgroup for a raag, state the main result, and discuss the strategy of its proof. In section 6.2, we review automorphisms of right-angled Artin groups, and in section 6.3, we recall Charney–Stambaugh–Vogtmann’s outer space for long-range automorphisms of raags. In section 6.4, we show that the Torelli subgroup is contained in the long-range automorphism group and prove that any prime order long-range automorphism is realized as a finite order automorphism of a blow-up of a Salvetti complex. In sections 6.5 and 6.6, we review the construction of blow-ups of Salvetti complexes and show that they satisfy the automorphism criterion from Chapter 5. We then use this to give a new proof that the Torelli subgroup for a raag is torsion-free.

6.1 Statement of Results

Let $\Gamma = (V, E)$ be a finite simplicial graph, with vertex set V and edge set E , and let A_Γ be the associated right-angled Artin group. If $V = \{v_1, \dots, v_n\}$ then $V^\pm = \{v_1^\pm, \dots, v_n^\pm\}$ is a generating set for A_Γ with the standard presentation, and the abelianization of A_Γ is $A_\Gamma^{ab} \cong \mathbb{Z}^n$. The abelianization map $\psi : A_\Gamma \rightarrow \mathbb{Z}^n$ induces

$\Psi : \text{Aut}(A_\Gamma) \rightarrow \text{GL}_n(\mathbb{Z})$ and we obtain short exact sequences

$$1 \rightarrow \text{IA}(A_\Gamma) \rightarrow \text{Aut}(A_\Gamma) \xrightarrow{\Psi} \text{GL}_n(\mathbb{Z})$$

$$1 \rightarrow \mathcal{I}(A_\Gamma) \rightarrow \text{Out}(A_\Gamma) \xrightarrow{\bar{\Psi}} \text{GL}_n(\mathbb{Z}).$$

The kernel $\mathcal{I}(A_\Gamma) = \ker \bar{\Psi}$ (resp. $\text{IA}(A_\Gamma) = \ker \Psi$) is called the **Torelli subgroup** of $\text{Out}(A_\Gamma)$ (resp. $\text{Aut}(A_\Gamma)$). The groups $\text{IA}(A_\Gamma)$ and $\mathcal{I}(A_\Gamma)$ are further related by the short exact sequence

$$1 \rightarrow \text{Inn}(A_\Gamma) \rightarrow \text{IA}(A_\Gamma) \rightarrow \mathcal{I}(A_\Gamma) \rightarrow 1$$

where $\text{Inn}(A_\Gamma) \cong A_\Gamma/Z(A_\Gamma)$ is the group of inner automorphisms of A_Γ .

The main goal of this chapter is to prove

Theorem 6.1 (Wade [40], Toinet [41]). *$\mathcal{I}(A_\Gamma)$ is torsion-free for all Γ .*

Here we present a geometric proof of this theorem using NPC cube complexes and the machinery developed in the previous section. We remark that Wade and Toinet's proofs almost exclusively use algebraic techniques. Before discussing the strategy of the proof, we list some immediate corollaries.

Corollary 6.1. *$\text{IA}(A_\Gamma)$ is torsion-free.*

Corollary 6.2 (Charney–Vogtmann [29]). *$\text{Out}(A_\Gamma)$ and $\text{Aut}(A_\Gamma)$ are both virtually torsion-free. In particular, each have finite virtual cohomological dimension.*

Both of these corollaries are straightforward consequences of the exact sequences above, Theorem 6.2 below, and Selberg's lemma (Theorem 3.1). We will prove the theorem in two steps. We assume for contradiction that $\phi \in \mathcal{I}(A_\Gamma)$ is torsion. Then

1. Realize ϕ as an automorphism $f : X \rightarrow X$ of some NPC cube complex X with $\pi_1(X) = A_\Gamma$. This means in particular that the induced map $f_* = \phi$ as an automorphism of $\pi_1(X)$.
2. Show that for any such X , every automorphism acts non-trivially on $H_1(X)$.

To carry out Step (1), we will make use of a contractible simplicial complex K_Γ on which $\mathcal{I}(A_\Gamma)$ acts. For Step (2), we will verify the criterion of Proposition 5.1 for certain NPC cube complexes.

6.2 Automorphisms of Raags

For each $v \in V$ we define two subsets of V :

$$lk(v) = \{w \in V \mid w \text{ is adjacent to } v\}$$

$$st(v) = \{v\} \cup lk(v).$$

Following [31], the relation $lk(v) \leq st(w)$ for $v, w \in V$ will be denoted $v \leq w$. In this case we say w **dominates** v . If $v \leq w$ and $w \leq v$ then we write $v \sim w$, in which case v and w are said to be **equivalent**.

Laurence [48] and Servatius [49] proved that the following four types of automorphisms generate $\text{Aut}(A_\Gamma)$:

1. **Inversions:** If $v \in V^\pm$, the automorphism i_v sends $v \mapsto v^{-1}$ and fixing all other generators.
2. **Graph Automorphism:** Any automorphism of Γ induces a permutation of V^\pm which extends to an automorphism of A_Γ .
3. **Transvections:** If $v \leq w$, the automorphism $\tau_{w,v}$ sends $v \mapsto vw$ and fixes all other generators.

4. **Partial Conjugations:** If C is a connected component of $\Gamma \setminus st(v)$ for some $v \in V$, the automorphism $\sigma_{v,C}$ maps $w \mapsto v w v^{-1}$ for every $w \in C$, and acts as the identity elsewhere.

If, in (3), v and w are adjacent, the automorphism $\tau_{w,v}$ is called an **adjacent** transvection. Otherwise, the automorphism $\tau_{w,v}$ is called a **non-adjacent** transvection. As in [50] and [31], we distinguish the subgroup of **long-range automorphisms** $\text{Out}_\ell(A_\Gamma) \subseteq \text{Out}(A_\Gamma)$ (resp. $\text{Aut}_\ell(A_\Gamma) \subseteq \text{Aut}(A_\Gamma)$) generated by automorphisms of type (1), (2), (4) and non-adjacent transvections.

6.3 The Out_ℓ -spine K_Γ

Recall the definition of the standard Salvetti complex $\mathbb{S} = \mathbb{S}_\Gamma$ associated to A_Γ from § 3.3. The cube complex \mathbb{S} is constructed as follows. Start with a single vertex x_0 . For every $v \in \Gamma$, attach both ends of a 1-cube e_v to x_0 . For every complete k -subgraph of Γ , we add in a k -cube C whose image is a k -torus with 1-skeleton the edges labelled by elements in the subgraph. Gromov's link criterion implies \mathbb{S} is an NPC cube complex whose fundamental group is A_Γ . In particular, we have that \mathbb{S} is a $K(A_\Gamma, 1)$.

In [31], Charney, Stambaugh and Vogtmann constructed a contractible simplicial complex K_Γ on which $\text{Out}_\ell(A_\Gamma)$ acts properly discontinuously cocompactly by simplicial automorphisms. Like outer space for free groups, one considers pairs (X, ρ) , where X is an NPC cube complex with fundamental group A_Γ , and ρ is a homotopy equivalence $\rho : X \rightarrow \mathbb{S}$. The pair (X, ρ) is called a **marked blow-up of a Salvetti complex**. Construction of such cube complexes will be described below. An automorphism $\phi \in \text{Out}_\ell(A_\Gamma)$ acts on $(X, \rho) \in K_\Gamma$ by changing the marking: Represent ϕ as a homotopy equivalence $h : \mathbb{S} \rightarrow \mathbb{S}$. Then $\phi.(X, \rho) = (X, h \circ \rho)$. We have

Theorem 6.2 ([31], Propositions 4.17, Theorem 5.24). *K_Γ is contractible and the action of $\text{Out}_\ell(A_\Gamma)$ on K_Γ is properly discontinuous.*

See [31] for details on the construction of K_Γ . For us, what will be important is that it is contractible, finite dimensional, and admits a properly discontinuous action of $\text{Out}_\ell(A_\Gamma)$.

6.4 The Torelli Subgroup for a Raag

Day has shown in [50] that $\text{IA}(A_\Gamma)$ is generated by automorphisms of the following two forms:

1. **(Partial Conjugation)** Let $v \in V^\pm$ be a generator, and $C \neq \emptyset$ a component of $\Gamma \setminus st(v)$. Then $\sigma_{v,C} : A_\Gamma \rightarrow A_\Gamma$ denotes the automorphism

$$\sigma_{v,C} : w \mapsto v w v^{-1}, \quad w \in C^\pm;$$

$$w \mapsto w, \quad \text{else.}$$

2. **(Commutator Transvection)** Let $v, w_1, w_2 \in V^\pm$ such that w_1, w_2 both dominate v , *i.e.* $lk(v) \subset lk(w_1), lk(w_2)$. Then there are non-adjacent transvections of v by w_1 and w_2 , and we can therefore transvect v by the commutator $[w_1, w_2]$.

Denote by $\tau_{w_1, w_2, v} : A_\Gamma \rightarrow A_\Gamma$ the automorphism

$$\tau_{w_1, w_2, v} : w \mapsto [w_1, w_2] w, \quad w = v;$$

$$w \mapsto w, \quad \text{else.}$$

We remark that in case (1), if $\Gamma \setminus st(v)$ is connected, then $\sigma_{v,C}$ is just conjugation by v . From this generating set, it is clear that $\text{IA}(A_\Gamma) \leq \text{Aut}_\ell(A_\Gamma)$, since partial conjugations lie in $\text{Aut}_\ell(A_\Gamma)$ by definition, and in order for the transvection in case

(2) to be non-trivial, we must have that v , w_1 and w_2 are pairwise non-adjacent. Passing to the outer automorphism group, we obtain $\mathcal{I}(A_\Gamma) \leq \text{Out}_\ell(A_\Gamma)$.

Since $\mathcal{I}(A_\Gamma) \leq \text{Out}_\ell(A_\Gamma)$, it follows that $\mathcal{I}(A_\Gamma)$ acts on K_Γ by simplicial automorphisms. Suppose $\phi \in \mathcal{I}(A_\Gamma)$ has finite order. Without loss we may assume the order is prime. K_Γ is finite-dimensional and by Theorem 6.2 it is contractible, hence the action of ϕ on K_Γ has a fixed point (otherwise, there would be a finite-dimensional Eilenberg-MacLane space for $\mathbb{Z}/n\mathbb{Z}$). Note that because K_Γ is the subdivision of a simplicial complex on which $\text{Out}_\ell(A_\Gamma)$ acts, the fixed point is actually vertex of K_Γ . The fixed point $(X, \rho) \in K_\Gamma$ corresponds to a marked blow-up of a Salvetti complex. By the definition of K_Γ , this means that ϕ is realized as an automorphism $f : X \rightarrow X$ which commutes with the marking ρ up to homotopy. Further, as $\phi \in \mathcal{I}(A_\Gamma)$, we know that the induced map $f_* : H_1(X) \rightarrow H_1(X)$ is the identity map. This will be the starting point for our investigation. We want to show that f itself must be the identity. We record the preceding discussion in

Proposition 6.1. *Let $\phi \in \mathcal{I}(A_\Gamma)$ have prime order. Then ϕ is realized as an automorphism $f : X \rightarrow X$ of some marked blow-up X of a Salvetti complex.*

6.5 Blow-ups of Salvetti complexes

At this point, we have realized torsion elements of $\mathcal{I}(A_\Gamma)$ as automorphisms of cube complexes which act trivially on first homology. In order to show that $\mathcal{I}(A_\Gamma)$ is torsion-free, it suffices to show that any automorphism of a blow-up X acts non-trivially on $H_1(X)$. To do this, we will show that every blow-up satisfies the hypotheses of Proposition 5.1.

Generalizing Whitehead partitions for free groups, Charney, Stambaugh, and Vogtmann define automorphisms of A_Γ which they call Γ -Whitehead partitions. The

reason for using a generating set consisting of Γ -Whitehead automorphisms instead of the standard generating set is that each Γ -Whitehead automorphism can be achieved by an expansion and collapse of a Salvetti complex for A_Γ .

Definition 6.1 ([31], Definition 2.1). Let $P \subset V^\pm$ have at least 2 elements, including some $m \in P$ with $m^{-1} \notin P$. Then (P, m) is a Γ -**Whitehead pair** if

1. No element of P is adjacent to m ,
2. If $v \in P$ and $v^{-1} \notin P$ then $v \leq m$,
3. If $v^\pm \in P$, then $w^\pm \in P$ for every w in the same component of $\Gamma \setminus st(m)$ as v .

A Γ -Whitehead pair (P, m) defines an automorphism $\phi = \phi_{(P, m)}$ defined by

$$\phi(v) = \begin{cases} m^{-1} & \text{if } v = m \\ vm^{-1} & \text{if } v \in P \text{ and } v^{-1} \notin P \\ mv & \text{if } v^{-1} \in P \text{ and } v \notin P \\ mvm^{-1} & \text{if } v^\pm \in P \\ v & \text{else} \end{cases}$$

The pair (P, m) also defines several important subsets of V^\pm

$$double(P) = \{v \in P \mid v^\pm \in P\}$$

$$single(P) = \{v \in P \mid v^{-1} \notin P\}$$

$$max(P) = \{v \in single(P) \mid v \sim m\}$$

$$lk(P) = lk(m)^\pm$$

The automorphism ϕ is clearly a product of an inversion of m , a transvection of elements of elements of $single(P)$ and a partial conjugation of elements of $double(P)$,

hence $\phi \in \text{Out}_\ell(A_\Gamma)$. Conversely, it is easy to see that $\text{Out}_\ell(A_\Gamma)$ is generated by all Γ -Whitehead automorphisms together with inversions and graph automorphisms. Note that each $\phi_{(P,m)}$ has order 2.

Condition (1) in the definition implies that $P \cap lk(P) = \emptyset$. The other side of P , denoted P^* , is the complement of $P \cup lk(P)$ in V^\pm . (P^*, m^{-1}) is also a Γ -Whitehead pair which defines the same *outer* automorphism of A_Γ . We therefore obtain a disjoint union

$$V^\pm = P \cup lk(P) \cup P^*.$$

Definition 6.2 ([31], Definition 2.4). The triple $\mathbf{P} = \{P, lk(P), P^*\}$ is called a Γ -**Whitehead partition** of V^\pm . P and P^* are the **sides** of \mathbf{P} .

Definition 6.3 ([31], Definition 3.3). Let $\mathbf{P} = \{P, lk(P), P^*\}$ and $\mathbf{Q} = \{Q, lk(Q), Q^*\}$ be two Γ -Whitehead partitions.

1. \mathbf{P}, \mathbf{Q} **commute** if $\max(P), \max(Q)$ are distinct and commute.
2. \mathbf{P}, \mathbf{Q} are **compatible** if either they commute or at least one of $P \cap Q, P^* \cap Q, P \cap Q^*$ or $P^* \cap Q^*$ is empty.

It is shown in [31] that if \mathbf{P}, \mathbf{Q} are compatible and do not commute, exactly one of the intersections is empty. A collection $\mathbf{\Pi} = \{\mathbf{P}_1, \dots, \mathbf{P}_k\}$ is called **compatible** if the \mathbf{P}_i are pairwise compatible. A **region** of $\mathbf{\Pi}$ is choice of side $P_i^\times \in \{P_i, P_i^*\}$ for each i , such that for any i, j , either \mathbf{P}_i and \mathbf{P}_j commute, or $P_i^\times \cap P_j^\times \neq \emptyset$.

We are now in a position to build the blow-up $\mathbb{S}_\mathbf{\Pi}$ associated to $\mathbf{\Pi}$. First we will construct a contractible complex $\mathbb{E}_\mathbf{\Pi}$ containing all the vertices of $\mathbb{S}_\mathbf{\Pi}$. To each region $R = P_1^\times \cap \dots \cap P_k^\times$ we associate a vertex $x_R = (a_1, \dots, a_k)$ of the k -cube $[0, 1]^k$ via

$$a_i = \begin{cases} 0 & \text{if } P_i^\times = P_i \\ 1 & \text{if } P_i^\times = P_i^* \end{cases}$$

Now we attach edges to $\mathbb{E}_{\mathbf{II}}^{(0)}$. If R and R' are two regions which differ exactly by switching sides along a single partition P_i , we attach an edge e_{P_i} from x_R to $x_{R'}$. The edge e_{P_i} is oriented from the region containing P_i to the region containing P_i^* . The rest of $\mathbb{E}_{\mathbf{II}}$ is formed by filling in cubes where their boundaries occur.

We complete the construction of $\mathbb{S}_{\mathbf{II}}$ by attaching cubes to $\mathbb{E}_{\mathbf{II}}$, starting with the 1-cubes. Set $\overline{P}_i^\times = P_i^\times \cup lk(P_i)$. For each region R , define a subset V^\pm

$$I(R) = \overline{P}_1^\times \cap \cdots \cap \overline{P}_k^\times.$$

Compatibility implies each $I(R)$ is non-empty, and Lemma 3.10 (1) of [31] states that ever $v \in V^\pm$ occurs in some $I(R)$. If $v^\pm \in I(R)$, attach both vertices of an edge e_v at x_R . Suppose $v \in I(R)$ and $v^{-1} \notin I(R)$, and v is a single in $P_{i_1}^\times, \dots, P_{i_r}^\times$. By Lemma 3.10 (2) of [31], there is a region R_v obtained from R by switching sides along the $P_{i_j}^\times$, and $v^{-1} \in I(R_v)$. In this case we therefore attach an edge e_v from x_R to x_{R_v} . Note that $e_{v^{-1}} = \bar{e}_v$.

Every edge of $(\mathbb{S}_{\mathbf{II}})^{(1)}$ carries a **label** which is either some generator $v \in V^\pm$ or some partition \mathbf{P}_i . Two edges e_{l_1}, e_{l_2} have **commuting labels** if one of the following holds

1. $l_1 = v \in V^\pm, l_2 = w \in V^\pm$ and v, w are distinct and commute in A_Γ ,
2. $l_1 = v \in V^\pm, l_2 = P_i$ and $v \in lk(P_i)$,
3. $l_1 = P_i, l_2 = P_j$ and $\mathbf{P}_i, \mathbf{P}_j$ are distinct and commute.

With commuting labels defined as above, any collection of k edges with commuting labels at a vertex x_R forms the corner of the 1-skeleton of a k -cube in $(\mathbb{S}_{\mathbf{II}})^{(1)}$, with parallel edges carrying the same label ([31], Corollary 3.12). To finish the construction of $\mathbb{S}_{\mathbf{II}}$, we fill in all such k -cubes where they occur. $\mathbb{S}_{\mathbf{II}}$ is called the **blow-up** of \mathbb{S}_Γ along \mathbf{II} . We have

Theorem 6.3 ([31], Theorem 3.14). *The blow-up \mathbb{S}_Π is connected, locally CAT(0) and $\pi_1(\mathbb{S}_\Pi) \cong A_\Gamma$.*

After crossing an edge e_v labelled by a generator $v \in V$, there is a path in $E_\Pi^{(1)}$ connecting the two endpoints of e_v . This path crosses edges labelled by every partition containing v as a singleton. We call such a path a **characteristic loop** γ_v .

6.6 Automorphisms of blow-ups

Proposition 6.2. *Every blow-up X satisfies the hypotheses of Proposition 5.1.*

Proof. Without loss of generality, we assume $X = \mathbb{S}_\Pi$ is a blow-up of the standard Salvetti with the identity marking, hence comes equipped with some labeling of the 1-skeleton by generators $v_1, \dots, v_n \in V$ or partitions $\mathbf{P}_1, \dots, \mathbf{P}_k \in \Pi$. The hyperplanes of X are in one-to-one correspondence with these labels, so we check them one by one. To this end, let l_1 and l_2 be labels with corresponding hyperplanes K_{l_1}, K_{l_2} .

$K_{l_1} \cap K_{l_2} \neq \emptyset$: Observe that K_{l_1}, K_{l_2} intersect if and only if their corresponding labels commute. Consider a square bounded on by edges e_1 and e_2 dual to K_{l_1} and K_{l_2} , respectively. Each edge e_i can be completed to a characteristic loop γ_i as follows. If $l_i = v$ is a generator, then take a characteristic loop γ_v . If $l_i = P$ is a partition, choose some $m \in \max(P)$, and complete this to a characteristic loop γ_m . Next observe that since $lk(l_1) \subset st(l_2)$ for every label l occurring on γ_1 , l_2 commutes with every such l , and similarly for γ_2 and l_1 . It follows that $\gamma_1 \subseteq K_{l_2}$ and $\gamma_2 \subseteq K_{l_1}$. Thus, $\gamma_i \cdot K_{l_j} = \delta_{ij}$ and this case is satisfied.

$K_{l_1} \cap K_{l_2} = \emptyset$ and $K_{l_1} \cup K_{l_2}$ separates: Since E_Π contains all of the vertices of \mathbb{S}_Γ , it is easy to see that if the l_i both correspond to generators, then $K_{l_1} \cup K_{l_2}$ cannot separate. Thus the only possibilities for pairs of separating hyperplanes are

one generator, one partition or two partitions.

First suppose $l_1 = v$ and $l_2 = P$. We know that K_P disconnects E_{Π} into two components, corresponding to vertices which contain P and those which contain P^* . If $K_v \cup K_P$ separate, then we must have $v \in \text{single}(P)$, and in fact $\{v\} = \text{single}(P) = \text{max}(P)$. Then (P, v) is one of the Γ -Whitehead partitions in Π . By assumption, this partition is non-degenerate; hence there must be $w_1^{\pm} \in \text{double}(P)$ and $w_2^{\pm} \in \text{double}(P^*)$. The hyperplanes corresponding to w_1 and w_2 do not separate there respective components.

Now assume $l_1 = P$ and $l_2 = Q$. Then P and Q are compatible and do not commute, hence without loss of generality we have $P \subset Q$ and $Q^* \subset P^*$ by Lemma 3.4 of [31]. Then $P^* \cap Q \neq \emptyset$. In E_{Π} , deleting K_P and K_Q leaves three components E_1 , E_2 and E_3 whose vertices correspond to regions containing $P \cap Q$, $P^* \cap Q$, and $P^* \cap Q^*$, respectively. Elements of $\text{single}(P) \setminus \text{single}(Q)$ connect E_1 and E_2 , elements of $\text{single}(Q) \setminus \text{single}(P)$ connect E_2 and E_3 , while elements of $\text{single}(P) \cap \text{single}(Q)$ connect E_1 and E_3 . If $\text{max}(Q) \neq \text{max}(P)$ then $K_P \cup K_Q$ does not separate. Then if $\text{max}(P) = \text{max}(Q)$, the only way $K_P \cup K_Q$ separates is if actually $\text{single}(P) = \text{single}(Q)$. As $P^* \cap Q \neq \emptyset$ we must have that $\text{double}(Q) \neq \text{double}(P)$, or else $P = Q$, which is impossible. Then the component containing E_2 has a non-separating hyperplane labeled by some $w^{\pm} \in \text{double}(Q) \setminus \text{double}(P)$. If $\text{single}(P)$ is not a single element, then the hyperplane corresponding to any element of $\text{single}(P)$ does not disconnect the component containing $E_1 \cup E_3$. Otherwise, $\{m\} = \text{max}(P) = \text{single}(P) = \text{single}(Q)$ is a single generator. In this case, since (P, m) is non-trivial, there exists $v^{\pm} \in \text{double}(P)$, and the hyperplane K_v does not separate the component containing $E_1 \cup E_3$.

Finally, to see that condition (3) of Proposition 5.1 is satisfied, observe that for

each maximal collection of pairwise commuting hyperplanes, there is a unique cube in which they all meet. If $f : X \rightarrow X$ is an automorphism which preserves every hyperplane, the automorphism f must fix each of these cubes pointwise. Since the union of these cubes covers X , we deduce that f is the identity. This completes the proof. \square

Corollary 6.3. *Every automorphism of a blow-up a Salvetti acts nontrivially on H_1 .*

We are now in a position to finish the proof of Theorem 6.1:

Proof. Suppose for contradiction there exists $\phi \neq 1 \in \mathcal{I}(A_\Gamma)$ such that $\phi^n = 1$. Passing to a power, we may assume n is prime. By Proposition 6.1, there exists a blow-up of a Salvetti X and an automorphism $f : X \rightarrow X$ such that $f_* = \phi \in \text{Out}(A_\Gamma)$. Corollary 6.3 now implies that if f acts trivially on $H_1(X)$, f is the identity. Hence, $\phi = 1$, a contradiction. \square

Bibliography

- [1] F. Haglund and D. T. Wise, “Coxeter groups are virtually special,” *Adv. Math.*, vol. 224, no. 5, pp. 1890–1903, 2010.
- [2] A. S. Švarc, “A volume invariant of coverings,” *Dokl. Akad. Nauk SSSR (N.S.)*, vol. 105, pp. 32–34, 1955.
- [3] J. Milnor, “Growth of finitely generated solvable groups,” *J. Differential Geometry*, vol. 2, pp. 447–449, 1968.
- [4] M. Gromov, “Hyperbolic groups,” in *Essays in group theory*, vol. 8 of *Math. Sci. Res. Inst. Publ.*, pp. 75–263, Springer, New York, 1987.
- [5] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, vol. 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [6] D. Gromoll and J. A. Wolf, “Some relations between the metric structure and the algebraic structure of the fundamental group in manifolds of nonpositive curvature,” *Bull. Amer. Math. Soc.*, vol. 77, pp. 545–552, 1971.
- [7] H. B. Lawson, Jr. and S. T. Yau, “Compact manifolds of nonpositive curvature,” *J. Differential Geometry*, vol. 7, pp. 211–228, 1972.
- [8] E. Ghys and P. de la Harpe, “Panorama,” in *Sur les groupes hyperboliques d’après Mikhael Gromov (Bern, 1988)*, vol. 83 of *Progr. Math.*, pp. 1–25, Birkhäuser

Boston, Boston, MA, 1990.

- [9] M. Coornaert, T. Delzant, and A. Papadopoulos, *Géométrie et théorie des groupes*, vol. 1441 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1990. Les groupes hyperboliques de Gromov. [Gromov hyperbolic groups], With an English summary.
- [10] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, “Notes on word hyperbolic groups,” in *Group theory from a geometrical viewpoint (Trieste, 1990)*, pp. 3–63, World Sci. Publ., River Edge, NJ, 1991. Edited by Short.
- [11] J.-P. Serre, *Trees*. Springer-Verlag, Berlin-New York, 1980. Translated from the French by John Stillwell.
- [12] M. Sageev, “Ends of group pairs and non-positively curved cube complexes,” *Proc. London Math. Soc. (3)*, vol. 71, no. 3, pp. 585–617, 1995.
- [13] M. A. Roller, “Poc sets, median algebras and group actions.” Preprint: <https://arxiv.org/abs/1607.07747>, 2016.
- [14] F. Haglund and F. Paulin, “Simplicité de groupes d’automorphismes d’espaces à courbure négative,” in *The Epstein birthday schrift*, vol. 1 of *Geom. Topol. Monogr.*, pp. 181–248, Geom. Topol. Publ., Coventry, 1998.
- [15] F. Haglund and D. T. Wise, “Special cube complexes,” *Geom. Funct. Anal.*, vol. 17, no. 5, pp. 1551–1620, 2008.
- [16] T. Hsu and D. T. Wise, “On linear and residual properties of graph products,” *Michigan Math. J.*, vol. 46, no. 2, pp. 251–259, 1999.

- [17] C. G. A. Droms, *GRAPH GROUPS (ALGEBRA, KIM, ROUSH, MAGNUS)*. ProQuest LLC, Ann Arbor, MI, 1983. Thesis (Ph.D.)–Syracuse University.
- [18] I. Agol, “The virtual Haken conjecture,” *Doc. Math.*, vol. 18, pp. 1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.
- [19] Y. Liu, “Virtual cubulation of nonpositively curved graph manifolds,” *J. Topol.*, vol. 6, no. 4, pp. 793–822, 2013.
- [20] P. Przytycki and D. T. Wise, “Mixed 3-manifolds are virtually special.” Preprint: <https://arxiv.org/abs/1205.6742>, 2013.
- [21] D. T. Wise, “The structure of groups with a quasiconvex hierarchy.” Preprint: <http://math.hunter.cuny.edu/olgak/groupThNotes/wiseHierarchy2013.pdf>, October 2012.
- [22] I. Agol, D. Groves, and J. F. Manning, “An alternate proof of Wise’s malnormal special quotient theorem,” *Forum Math. Pi*, vol. 4, pp. e1, 54, 2016.
- [23] R. Charney, “An introduction to right-angled Artin groups,” *Geom. Dedicata*, vol. 125, pp. 141–158, 2007.
- [24] R. Charney and M. W. Davis, “The $K(\pi, 1)$ -problem for hyperplane complements associated to infinite reflection groups,” *J. Amer. Math. Soc.*, vol. 8, no. 3, pp. 597–627, 1995.
- [25] K. S. Brown, *Cohomology of groups*, vol. 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.
- [26] J. Nielsen, “Die Isomorphismen der allgemeinen, unendlichen Gruppe mit zwei Erzeugenden,” *Math. Ann.*, vol. 78, no. 1, pp. 385–397, 1917.

- [27] J. Nielsen, “Die Isomorphismengruppe der freien Gruppen,” *Math. Ann.*, vol. 91, no. 3-4, pp. 169–209, 1924.
- [28] M. Culler and K. Vogtmann, “Moduli of graphs and automorphisms of free groups,” *Invent. Math.*, vol. 84, no. 1, pp. 91–119, 1986.
- [29] R. Charney and K. Vogtmann, “Finiteness properties of automorphism groups of right-angled Artin groups,” *Bull. Lond. Math. Soc.*, vol. 41, no. 1, pp. 94–102, 2009.
- [30] R. Charney and M. Margolis, “Length functions of 2-dimensional right-angled Artin groups,” *Geom. Dedicata*, vol. 166, pp. 31–45, 2013.
- [31] R. Charney, N. Stambaugh, and K. Vogtmann, “Outer space for untwisted automorphisms of right-angled Artin groups.” Preprint: <http://arxiv.org/abs/1212.4791>, 2012.
- [32] T. Koberda and A. I. Suciú, “Residually finite rationally p groups.” Preprint: <http://arxiv.org/abs/1604.02010>, 2016.
- [33] Y. Antolín and A. Minasyan, “Tits alternatives for graph products,” *J. Reine Angew. Math.*, vol. 704, pp. 55–83, 2015.
- [34] H. Poincaré, *Œuvres. Tome VI. Les Grands Classiques Gauthier-Villars*. [Gauthier-Villars Great Classics], Éditions Jacques Gabay, Sceaux, 1996. Géométrie. Analysis situs (topologie). [Geometry. Analysis situs (topology)], Reprint of the 1953 edition.
- [35] B. Farb and D. Margalit, *A primer on mapping class groups*, vol. 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

- [36] M. Culler, “Finite groups of outer automorphisms of a free group,” in *Contributions to group theory*, vol. 33 of *Contemp. Math.*, pp. 197–207, Amer. Math. Soc., Providence, RI, 1984.
- [37] B. Zimmermann, “Über Homöomorphismen n -dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen,” *Comment. Math. Helv.*, vol. 56, no. 3, pp. 474–486, 1981.
- [38] D. G. Khramtsov, “Finite groups of automorphisms of free groups,” *Mat. Zametki*, vol. 38, no. 3, pp. 386–392, 476, 1985.
- [39] G. Baumslag and T. Taylor, “The centre of groups with one defining relator,” *Math. Ann.*, vol. 175, pp. 315–319, 1968.
- [40] R. D. Wade, “Johnson homomorphisms and actions of higher-rank lattices on right-angled Artin groups,” *J. Lond. Math. Soc. (2)*, vol. 88, no. 3, pp. 860–882, 2013.
- [41] E. Toinet, “Conjugacy p -separability of right-angled Artin groups and applications,” *Groups Geom. Dyn.*, vol. 7, no. 3, pp. 751–790, 2013.
- [42] A. Selberg, “On discontinuous groups in higher-dimensional symmetric spaces,” in *Contributions to function theory (internat. Colloq. Function Theory, Bombay, 1960)*, pp. 147–164, Tata Institute of Fundamental Research, Bombay, 1960.
- [43] A. Malcev, “On isomorphic matrix representations of infinite groups,” *Rec. Math. [Mat. Sbornik] N.S.*, vol. 8 (50), pp. 405–422, 1940.
- [44] D. T. Wise, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, vol. 117 of *CBMS Regional Conference Series in Mathematics*.

Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.

- [45] M. Salvetti, “The homotopy type of Artin groups,” *Math. Res. Lett.*, vol. 1, no. 5, pp. 565–577, 1994.
- [46] R. Charney and M. W. Davis, “Finite $K(\pi, 1)$ s for Artin groups,” in *Prospects in topology (Princeton, NJ, 1994)*, vol. 138 of *Ann. of Math. Stud.*, pp. 110–124, Princeton Univ. Press, Princeton, NJ, 1995.
- [47] M. Hall, Jr., “Subgroups of finite index in free groups,” *Canadian J. Math.*, vol. 1, pp. 187–190, 1949.
- [48] M. R. Laurence, “A generating set for the automorphism group of a graph group,” *J. London Math. Soc. (2)*, vol. 52, no. 2, pp. 318–334, 1995.
- [49] H. Servatius, “Automorphisms of graph groups,” *J. Algebra*, vol. 126, no. 1, pp. 34–60, 1989.
- [50] M. B. Day, “Peak reduction and finite presentations for automorphism groups of right-angled Artin groups,” *Geom. Topol.*, vol. 13, no. 2, pp. 817–855, 2009.