Cubical small-cancellation theory and large-dimensional hyperbolic groups

Macarena Covadonga Robles Arenas

Department of Pure Mathematics and Mathematical Statistics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Newnham College August 2023

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. It does not exceed the prescribed word limit for the relevant Degree Committee.

Macarena Covadonga Robles Arenas
August 2023

Abstract

Given a finitely presented group Q and a compact special cube complex X with non-elementary hyperbolic fundamental group, we produce a non-elementary, torsion-free, cocompactly cubulated hyperbolic group Γ that surjects onto Q, with kernel isomorphic to a quotient of $G = \pi_1 X$ and such that $max\{cd(G), 2\} \geq cd(\Gamma) \geq cd(G) - 1$.

Separately, we show that under suitable hypotheses, the second homotopy group of the coned-off space associated to a $C(9)$ cubical presentation is trivial, and use this to provide classifying spaces for proper actions for the fundamental groups of many quotients of square complexes admitting such cubical presentations. When the cubical presentations satisfy a condition analogous to requiring that the relators in a group presentation are not proper powers, we conclude that the corresponding coned-off space is aspherical.

A mi familia: Covadonga, Isabel, José Carlos, Rosario, y Juan Carlos,

a la memoria de mi abuelo Delfino,

y a la memoria de Miguel.

Acknowledgements

I am deeply grateful to my PhD supervisor, Henry Wilton, for being an inexhaustible source of mathematical knowledge, for his enthusiasm, empathy, and advice, for being a patient listener and a careful explainer, and – crucially – for striking the right balance between keeping me grounded in reality, and encouraging me to pursue and develop my own ideas. To put it simply, I could not have hoped for a better supervisor, nor for a kinder mentor.

I had never felt so valued and so stimulated at any other institution as I have felt at Cambridge. I thank all the people who have made this place what it is. In particular, I would like to thank the Pav E gang: Ana, Alexis, Julian, and Will; Andreas, Ismael, Samuel, Guillem, Thibault, Martí, Patrick, Dmitri, and Sam.

My work and my perspective of the field have benefited greatly from many mathematical conversations and interactions – I am especially thankful to Daniel Groves for many on-point remarks regarding the specifics of my projects, to Rémi Coulon, Mahan Mj, and Michah Sageev for being so generous with their time, and to Mladen Bestvina, Rachael Boyd, Francois Dahmani, Nir Lazarovich, Ian Leary, Ashot Minasyan, Nansen Petrosyan, Piotr Przytycki, Pierre Py, and Genevieve Walsh for all their insights. I also thank Daniel Wise for sharing his toolkit and his thoughts, and for all his support.

I thank Jack Button and Mark Hagen for their careful reading of this thesis, and for many helpful comments and suggestions.

I thank Montse Casals and Ilya Kazachkov, Claudio Llosa-Isenrich, Luis Jorge Sanchez Saldaña, Rob Kropholler, Dawid Kielak, and all the other mathematicians who have taken an interest in my work, and who have provided me with opportunities to talk about it.

I thank the rest of my friends in the geometric group theory community – Alice, Dario, Eduardo, Gil, JP, Jonah, Jone, Kevin, Kostas, Lawk, Marco, Markus, Martín, Monika, Nima, Rachel, Rylee, Sam, Suraj, and Thomas.

As ever, I am grateful to Max Neumann-Coto, Sami Douba, Abdul Zalloum, Julian Pavón, Josep Romagosa, Geneviève Young, and Isaac Ortigoza for their friendship and encouragement, and for lending a listening ear whenever one has been needed.

I owe everything I am to the love and support of my family. To my mother, Covadonga, my grandmother, Isabel, my brother, José Carlos, my aunt, Rosario, and my uncle, Juan

Carlos: you are the biggest blessing God has given me. This thesis is for you, and it is your achievement as much as it is mine.

I acknowledge the generous financial support of the Cambridge Trust and of Newnham College, Cambridge.

‘The most uncomplicated thing of all!’ he replied. ‘For someone well acquainted with the fifth dimension, it costs nothing to expand space to the desired proportions. I’ll say more, respected lady - to devil knows what proportions! I, however,’ Koroviev went on chattering, ‘have known people who had no idea, not only of the fifth dimension, but generally of anything at all, and who nevertheless performed absolute wonders in expanding their space. Thus, for instance, one city-dweller, as I’ve been told, having obtained a three-room apartment on Zemlyanoy Val, transformed it instantly, without any fifth dimension or other things that addle the brain, into a four-room apartment by dividing one room in half with a partition.’

The Master and Margarita
MIKHAIL BULGAKOV

And what there is to conquer
By strength and submission, has already been discovered
Once or twice, or several times, by men whom one cannot hope
To emulate—but there is no competition—
There is only the fight to recover what has been lost
And found and lost again and again: and now, under conditions
That seem unpropitious. But perhaps neither gain nor loss.
For us, there is only the trying. The rest is not our business.

Four Quartets, Part II: East Coker
T. S. ELIOT

Contents

List of Figures

Chapter 1

Introduction

The objective of this work is the study of quotients of cubulated groups, that is, quotients of groups that act nicely on CAT(0) cube complexes. Such quotients can often be viewed as *cubical presentations* $\langle X|\{Y_i \to X\}\rangle$ where X and the Y_i are non-positively curved cube complexes and $Y_i \to X$ are local isometries. Associated to a cubical presentation there is a space X^* – akin to a presentation complex for a group presentation – which is obtained by coning-off the images of the Y_i in X, and which allows us to speak of the fundamental group $\pi_1 X^* = \pi_1 X/\langle\langle\{\pi_1 Y_i\}\rangle\rangle$ of $\langle X|\{Y_i \to X\}\rangle$. These fundamental groups may be understood by combining the many available tools and results in cubical geometry with a theory analogous to the small-cancellation theory used to analyse group presentations where the relators have "small overlaps". This framework is known as *cubical small-cancellation theory*, and was originally developed by Wise [Wis21].

In "classical" small-cancellation theory, one defines for each $n \in \mathbb{N}$ a *metric* and a *non-metric* small cancellation condition, denoted by $C'(\frac{1}{n})$ and $C(n)$ respectively. The metric condition measures the length of overlaps (instances of repeated subwords) between relators, and the non-metric condition counts the minimum number of overlaps that are necessary to spell out each relator. Perhaps the main selling point of this theory is the theorem that finitely presented $C'(\frac{1}{6})$ and $C(7)$ small-cancellation groups are hyperbolic. Despite the fact that neither hyperbolic groups nor small-cancellation groups had formally been defined at the time, this result is essentially contained in Dehn's 1912 solution of the word problem for surface groups [Deh87]. Decades later, the study and development of small-cancellation theory was continued by, among others, Greendlinger [Gre60], Lyndon [Lyn66], Lyndon and Schupp [LS77], Rips [Rip82], Ol'shanskii [Ol′91], Gersten and Short [GS91], and later on by Wise [Wis04], who proved that $C'(\frac{1}{6})$ groups are cocompactly cubulated.

Just as classical small-cancellation theory is used to provide criteria for the hyperbolicity and cubulability of quotients of free groups, the cubical version can be used to provide criteria for the hyperbolicity and cubulability of quotients of cubulated groups [Wis21], suggesting a promising analogy between the classical and cubical versions of the theory.

Cubical small-cancellation theory has already been used to prove important results in geometric group theory, most saliently as a crucial stepping stone in Agol's celebrated proofs of the Virtual Haken and Virtual Fibering conjectures [Ago08, Ago13]. In many ways, however, the theory is largely underdeveloped, and numerous directions remain unexplored. We aim to make a contribution to remedy this situation, and focus on addressing two questions:

1. Can cubical small-cancellation theory be used to provide examples of **large-dimensional** (cubulated) hyperbolic groups with interesting properties?

2. Can cubical small-cancellation theory be used to provide good models for the classifying spaces, or classifying spaces for proper actions, of quotients of cubulated groups?

Let us provide some motivation for both of these endeavours.

Cubical small-cancellation as a means for producing examples

In his beautiful 1982 paper [Rip82], Rips proved that for every finitely presented group Q, there is a hyperbolic group Γ surjecting onto Q with finitely generated kernel. He then used this idea to produce examples of hyperbolic groups exhibiting a variety of exotic properties: hyperbolic groups with unsolvable membership problem, incoherent hyperbolic groups, hyperbolic groups having finitely generated subgroups whose intersection is not finitely generated, and hyperbolic groups containing infinite ascending chains of r-generated groups. Perhaps even more baffling than the fact that such a simple procedure could be used to produce such a wealth of examples, is the fact that in many of these instances, no other method is known for producing them. For example, as far as the author is aware, there are no known examples or constructions of hyperbolic groups with unsolvable membership problem that do not somehow utilise the Rips construction.

The idea of taking a finitely presented group Q having some pathological property, and building a nicer (usually hyperbolic) group Γ from Q in such a way that the same or some other closely related property persists in Γ, continues to be widely exploited: a list of many of the variations of Rips' original result, and their applications, is given in Section 3.1.

While hyperbolic groups exist in abundance [Gro93, Oll05], there are surprisingly few constructions in large dimensions [Gro87, Bes, JŚ03, FM10, Osa13]. One would hope to exploit Rips' theorem to produce large-dimensional hyperbolic groups with a variety of behaviours, but unfortunately this cannot be done: the presentation complexes associated to torsion-free $C(6)$ small-cancellation presentations are aspherical, and thus serve to show that these groups have geometric dimension at most equal to 2. Since all variations of the Rips construction utilise some form of small-cancellation theory, or otherwise construct directly a 2-dimensional classifying space, the prospects of using Rips-type techniques to construct large-dimensional hyperbolic groups seem rather poor.

Using cubical small-cancellation theory, we give in Chapter 3 a variation of the Rips construction that produces, given a finitely presented group Q and a compact special cube complex X with hyperbolic fundamental group, a cubulated hyperbolic group Γ that surjects onto Q with finitely generated kernel, and whose cohomological dimension can be effectively bounded in terms of the cohomological dimension of $\pi_1 X$:

Theorem A. *Let Q be a finitely presented group and G be the fundamental group of a compact special cube complex X. If G is hyperbolic and non-elementary, then there is a short exact sequence*

$$1 \to N \to \Gamma \to Q \to 1$$

where

1. *Γ is a hyperbolic, cocompactly cubulated group,*
2. *$N \cong G/K$ for some $K < G$,*
3. *$max\{cd(G), 2\} \geq cd(\Gamma) \geq cd(G) - 1$. In particular, Γ is torsion-free.*

This construction opens many routes of exploration: setting $Q = \mathbb{Z}$ and varying X, one obtains examples of cubulated hyperbolic groups which algebraically fibre. In a different direction, one can choose X with $cd(G) \geq 3$, and Q of type F_n but not F_{n+1}, and investigate under which conditions – if any – the subgroup $N < \Gamma$ is of type F_{n-1} but not F_n. While this is probably very hard, our theorem removes a key obstruction to this strategy: namely, when $cd(\Gamma) \leq 2$, results of Gersten [Ger96] imply that all finitely presented subgroups $N < \Gamma$ are in fact hyperbolic and hence of type F_n for all $n \in \mathbb{N}$.

Cubical small-cancellation as a means of controlling dimension

There are two classical, and closely related invariants that make precise the notion of dimension of a group: the cohomological dimension $cd(G)$, which is purely algebraic and is defined in terms of projective resolutions, and the geometric dimension $gd(G)$, which is the minimal dimension of a classifying space for G, and is therefore more geometric in flavour. It is always the case that $cd(G) \leq gd(G)$, and in general $cd(G) = gd(G)$ by the work of Eilenberg–Ganea [EG57] and Stallings and Swan [Sta68, Swa69]. The only possible exception, known as the Eilenberg–Ganea conjecture, is that there may exist groups with $cd(G) = 2$ and $gd(G) = 3$.

While these invariants can be extremely hard to compute, a general strategy usually has two parts: for an upper-bound, one tries to produce classifying spaces for G that arise in some natural way from the geometry or the combinatorics of G, and hopes that such classifying spaces are the smallest possible in terms of their dimension; for the lower-bound, one either computes directly some of the homology or cohomology groups of G, or at least tries to show that they are non-zero in large-enough dimension, or one finds a subgroup $G' < G$ whose cohomological dimension is either known, or easier to bound from below, and uses that $cd(G') \leq cd(G)$ to then bound from below the cohomological dimension of G.

Regardless of the group in question, an essential ingredient in this strategy is a good model for $K(G,1)$ – one whose dimension can be computed explicitly. For any group G, the first and most obvious candidate for a classifying space is the presentation complex associated to some presentation of G – but the presentation complex is at most 2-dimensional, and thus cannot be a classifying space for most groups. Even when $gd(G) = 2$, producing an explicit presentation for G that yields an aspherical presentation complex may be extremely difficult.

When a group has torsion, both $cd(G)$ and $gd(G)$ are infinite, and thus provide no additional information about G. This motivates the study of generalisations of $K(G,1)$, such as the classifying space $\underline{E}G$ for proper actions, and more generally, classifying spaces for families of subgroups of G. Every group admits a model for $\underline{E}G$, and all such models are G-homotopy equivalent. While not all groups admit a finite-dimensional model for $\underline{E}G$, many do, including for instance all groups having finite virtual cohomological dimension. One can define analogues of the cohomological and geometric dimension of a group using $\underline{E}G$, and relate these to other well-known dimensional invariants, such as the rational cohomological dimension $cd_{\mathbb{Q}}(G)$ or the virtual cohomological dimension $vcd(G)$. If a group has a finite dimensional $\underline{E}G$, these invariants often do provide meaningful information about G.

Classifying spaces for proper actions also appear in the statement of the Baum-Connes conjecture, which relates the equivariant K-homology of the classifying space of proper actions of a group to the K-theory of its corresponding reduced C*-algebra [Val02]. This is another motivation for trying to produce good models for $\underline{E}G$.

As mentioned a few paragraphs above, in the case of classical small-cancellation groups, it is a result of Lyndon [Lyn66] that the presentation complex associated to a $C(6)$ group presentation is aspherical if (and only if) none of the relators are proper powers. Continuing the analogy with cubical small-cancellation theory, this then begs the question of whether the "presentation complex" associated to a cubical small-cancellation presentation can, under suitable hypotheses, be a classifying space for its fundamental group.

In Chapter 4, we prove the first result in this direction:

Theorem B. *Let $X^* = \langle X | \{Y_i\}_{i\in I}\rangle$ be a minimal cubical presentation that satisfies the $C(9)$ condition. Let $\pi_1 X^* = G$. If $dim X \leq 2$ and each Y_i is homotopy equivalent to a graph, then X^* is a $K(G,1)$, so G is torsion-free and $gd(G) \leq 2$.*

The minimality hypothesis generalises prohibiting relators that are proper powers in the classical setting, and thus is a natural requirement if one expects to avoid torsion.

Relaxing minimality to allow torsion, but still retaining some control over how it can originate (this is the symmetry hypothesis in the statement below), we instead obtain classifying spaces for proper actions, which we can use to bound the rational cohomological dimension and the virtual cohomological dimension of some cubically presented groups:

Theorem C. *Let $X^* = \langle X | \{Y_i\}_{i=1}^{k}\rangle$ be a symmetric cubical presentation that satisfies the $C(9)$ condition. Let $\pi_1 X^* = G$. If $dim X \leq 2$ and each Y_i is homotopy equivalent to a graph, then there is a quotient $\bar{X}^*$ of the universal cover $\widetilde{X^*}$ of X^* that is an $\underline{E}G$, so $cd_{\mathbb{Q}}(G) \leq dim \bar{X}^* \leq 2$. If, in addition, X has a finite regular cover where each $Y_i \to X$ lifts to an embedding, then $vcd(G) \leq 2$.*

These two results hinge on the vanishing of the second homotopy group of a certain *reduced space* $\bar{X}^*$ that is a natural quotient of $\widetilde{X^*}$. This is Theorem 4.4.5 of Section 4.4, and holds for cubical presentations without any restrictions on dimension.

Theorems B and C, and more generally the technical result concerning the vanishing of π_2, together constitute the beginning of a programme to prove asphericity of $C(9)$ cubical presentations with no restriction on dimension. We pursue this programme in [Are].

Structure of the thesis

In Chapter 2 we give some general background on cube complexes and small-cancellation theory. In particular, we define all the notions related to cube complexes that are used later on in the text, we state some of the basic theorems related to these objects, and prove a few results that we like; Chapter 3 is devoted to Theorem A. Some background specific to this theorem is given in Section 3.2, the technicalities are handled in Section 3.3, most of Theorem A is proven in Section 3.4, and the last bit is proven in Section 3.5; Chapter 4 deals with Theorems B and C: in Section 4.2 we give a non-CAT(0)-geometric proof of the asphericity of non-positively curved cube complexes, in Section 4.3 we present some additional necessary background, in Section 4.4 we prove Theorems B and C, and in Section 4.5 we outline some examples to which these two results apply.

Chapter 2

General background

The content of this chapter can be found in a number of sources, including for instance in [CN05, Sag95, Wis12, BSV14].

2.1 Cube complexes

An *n-cube* is a copy of $[0,1]^n \subset \mathbb{R}^n$.

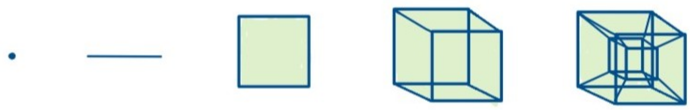

Figure 2.1: 0-cube, 1-cube, 2-cube, 3-cube, and 4-cube.

If c is an n-cube, the *faces* of c are the subcubes obtained by restricting one or several coordinates of c to $\{0\}$ or to $\{1\}$.

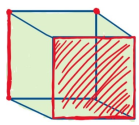

Figure 2.2: 0-,1-, and 2-faces in a 3-cube.

A *cube complex* is a CW-complex built from cubes of various dimensions by gluing them along faces by isometries. As topological spaces, cube complexes are equipped with the quotient topology. In other words, a cube complex is a quotient $\mathscr{C}/\mathscr{I}$ where $\mathscr{C}$ is a disjoint union of cubes and $\mathscr{I}$ is a collection of isometries between faces of cubes in $\mathscr{C}$. We assume also that faces are only glued to faces of the same dimension.

The *dimension* of a cube complex is the supremum of the dimensions of its constituent cubes. Note that a cube complex need not be finite dimensional. We will usually refer to a 1-dimensional cube complex as a *graph*, and to a 2-dimensional cube complex as a *square complex*. We say that a cube complex is *finite* if it is composed of finitely many cubes, and that it is *locally finite* if none of its vertices is contained in infinitely many cubes.

Definition 2.1.1. A *simplicial graph* is a graph with no *bigons* (cycles of length 2) or loops. A *flag complex* is a simplicial complex such that every set of $n+1$ vertices, with $n \geq 1$, span an n-simplex if and only if they are pairwise adjacent (i.e, any two vertices span a 1-simplex).

Intuitively, a simplicial complex is a flag complex if, whenever you see the 1-skeleton of a simplex, then the simplex is there. Note that, in this way, flag complexes are in one-to-one correspondence with simplicial graphs.

A *corner* in a cube c is an ε-sphere around a vertex of c, where $\varepsilon < \frac{1}{2}$. The *link* of a 0-cube v in a cube complex X is a simplex-complex with an n-simplex for each corner of an $(n+1)$-cube at v (glued accordingly).

Definition 2.1.2. A cube complex X is *non-positively curved* if $link(v)$ is a flag complex for each $v \in X^{(0)}$. A *CAT(0) cube complex* is a simply connected non-positively curved cube complex.

In this thesis, the notation $\widetilde{X}$ is always used to denote the universal cover of the corresponding cube complex (or more generally, CW complex) X. Accordingly, a CAT(0) cube complex shall be denoted $\widetilde{X}$, since CAT(0) cube complexes correspond precisely to the universal covers of non-positively curved cube complexes.

CAT(0) cube complexes are CAT(0) spaces in the metric sense of Cartan–Alexandrov–Toponogov [BH99], but we will not enter into any details regarding this aspect of the theory, and we will not use the CAT(0) metric for anything in this text. In the context of this thesis, the most important consequence of having a CAT(0) metric is that CAT(0) spaces are contractible, and thus, in particular, non-positively curved cube complexes are classifying spaces for their fundamental groups. We prove this without resorting to the CAT(0) metric in Section 4.2.

Example 2.1.3. Graphs are non-positively curved cube complexes, since $link(v)$ has no edges and therefore satisfies the flag condition vacuously.

Example 2.1.4. Euclidean space $\mathbb{E}^n$, tessellated in the usual way using n-cubes, is a CAT(0) cube complex – the link of each vertex is a $(n-1)$-sphere tessellated by 2^n simplices, and is thus flag. Similarly, the n-torus T^n is a non-positively curved cube complex when viewed as the quotient $\mathbb{E}^n/\mathbb{Z}^n$.

Example 2.1.5. Every closed surface with Euler characteristic $g \leq 0$ is homeomorphic to a non-positively curved cube complex by viewing it, for instance, as a $4g$-gon (in the orientable case) or a $2g$-gon (in the non-orientable case) with pairs of opposite sides identified and subdividing the $4g$-gon or the $2g$-gon into $4g$ or $2g$ squares. Such a subdivision for the non-orientable surface N_1 of Euler characteristic -1 is illustrated in Figure 2.4– since all closed surfaces with negative Euler characteristic cover N_1, one can also use this example to show that all such surfaces are homeomorphic to non-positively curved cube complexes. This observation goes back to the work of Scott [Sco78].

2.1.1 Hyperplanes in non-positively curved cube complexes

Let X be a non-positively curved cube complex.

Definition 2.1.6 (Midcube)**.** A *midcube* is the restriction of exactly one coordinate of an n-cube $c = [0,1]^n$ to $\{\frac{1}{2}\}$. In other words, a midcube is an $(n-1)$-cube that passes through the barycentre of c and is parallel to a pair of (necesarilly opposite) faces of c. Every n-cube c has n distinct midcubes of dimension $n-1$, all of which intersect in the barycentre of c.

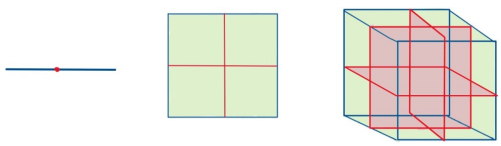

Figure 2.3: Midcubes in various cubes.

Definition 2.1.7 (Hyperplanes)**.** If e and e' are midcubes in distinct cubes of X, then their intersection is either empty or is a midcube in a cube of X. Let S be the disjoint union of all the midcubes in X. Let $\bar{S}$ be the quotient of S obtained by identifying midcubes along the faces in which they intersect. A *hyperplane* H is a connected component of $\bar{S}$. Note that the inclusion $S \hookrightarrow X$ induces a map $\bar{S} \to X$, and therefore also a map $H \to X$.

A hyperplane is *dual* to each 1-cube it intersects.

Examples 2.1.8. If X is a graph, then its hyperplanes are the barycentres of its edges. If $X = \mathbb{E}^n$, then its hyperplanes are copies of $\mathbb{E}^{n-1}$. If X is an n-torus obtained as a quotient of a single n-cube, then X has n hyperplanes, each of which is an $(n-1)$-torus.

Example 2.1.9. If X is a squaring of a closed surface S, for instance as in Example 2.1.5, then the hyperplanes of S are closed curves; in a more arbitrary squaring of S, the hyperplanes are graphs.

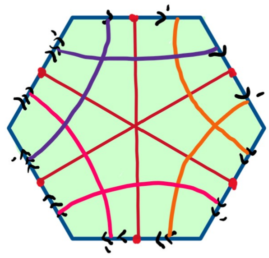

Figure 2.4: Hyperplanes (purple, orange, and pink) in a non-orientable genus 1 surface.

Definition 2.1.10 (Carrier)**.** The *carrier* $N(H)$ of a hyperplane H of X is the cubical neighbourhood of H. That is, it is the connected component containing H of the cube complex obtained from S by taking, for each cube c of S, the product $N(c) := c \times [0,1]$, considering the disjoint union $N(S)$ of all the $N(c)$, and taking the quotient $N(S) \to \overline{N(S)}$ induced by the quotient $S \to \bar{S}$.

As before, the inclusion $N(S) \hookrightarrow X$ induces a map $\overline{N(S)} \to X$, and therefore also a map $N(H) \to X$.

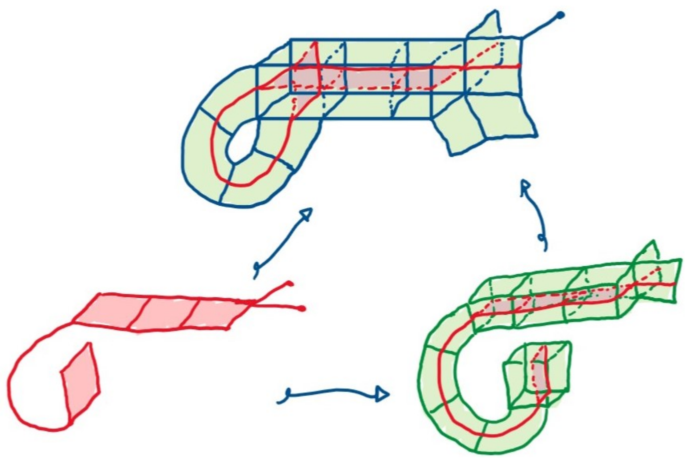

Figure 2.5: A hyperplane and its carrier.

For a graph B, a subgraph $A \subset B$ is *full* if whenever vertices $a_1, a_2 \in A$ are joined by an edge e of B, then $e \subset A$. In other words, A is the subgraph of B induced by A^0. A map $X \to Y$ between cell complexes is *combinatorial* if it maps cells to cells of the same dimension. An *immersion* is a local injection.

Definition 2.1.11. A *local isometry* $\varphi : Y \to X$ between non-positively curved cube complexes is a combinatorial map such that for each $y \in Y^0$ and $x = \varphi(y)$, the induced map $\varphi : link(y) \to link(x)$ is an injection of a full subgraph of X when restricted to the 1-skeleton.

A more visual way to think about local isometries is the following: an immersion φ is a local isometry if whenever two edges $\varphi(e), \varphi(f)$ form the corner of a square in X, then e, f already formed the corner of a square in Y.

Figure 2.6: Left: non-local isometry; right: local isometry. The arrows here indicate where the 1-cubes of Y are being mapped to in X.

Examples 2.1.12. Any graph immersion is a local isometry; any covering map $\hat{X} \to X$ of non-positively curved cube complexes is also a local isometry.

Lemma 2.1.13. *For every hyperplane H of X, the map $N(H) \looparrowright X$ is a local isometry.*

Proof. Let $\bar{c}$ be a square in the image of the map $N(H) \looparrowright X$, with 1-cubes $\bar{e}, \bar{e}'$ and $\bar{f}, \bar{f}'$, so that $\bar{e}$ and $\bar{f}$ share a vertex $\bar{v}$. There are 2 cases: first, if H is dual to $\bar{e}$ or to $\bar{f}$ in $N(H)$, then H extends into $\bar{c}$, and is dual to the opposite face $\bar{e}'$ or $\bar{f}'$, so the corresponding square, c, lies in $N(H)$. If H is not dual to $\bar{e}$ nor to $\bar{f}$, then, as $\bar{c}$ is in the image of $N(H) \looparrowright X$, then H is dual to a 1-cube $\bar{g}$ that shares a vertex $\bar{v}$ with $\bar{e}$ and $\bar{f}$. Thus, e, g and f, g form corners of squares c' and c'' in $N(H)$. The corresponding squares in the image form, together with $\bar{c}$, the corner of a 3-cube, and since X is non-positively curved, this implies that there is a 3-cube $\bar{K}$ at the corresponding corner. Now, H extends to the interior of $\bar{K}$, so there is a 3-cube K in $N(H)$ that maps to $\bar{K}$, and e, f form the corner of a square – a face of K – in $N(H)$.

When X is a CAT(0) cube complex, we can restate the definition of a hyperplane in a more elegant way; we explain in Theorem 2.1.16 that both definitions are equivalent:

Definition 2.1.14. A *hyperplane* in a CAT(0) cube complex $\widetilde{X}$ is a connected subspace H of $\widetilde{X}$ satisfying that, for each cube c, the intersection $H \cap c$ is either empty or a unique midcube.

Carriers of hyperplanes are defined as one would expect:

Definition 2.1.15. The *carrier* $N(H)$ of a hyperplane H of $\widetilde{X}$ is the union of all the cubes that contain H.

Starting from Definition 2.1.14, the existence of (non-trivial) hyperplanes is not obvious: a priori, it might be the case that no subspace H of $\widetilde{X}$ satisfies that, for each cube in X having $H \cap c \neq \emptyset$, the intersection of c with H is a unique midcube. This is addressed in the theorem below, which also establishes important properties of hyperplanes that will play a central role in Subsection 2.1.2.

Theorem 2.1.16. *If $\widetilde{X}$ is a CAT(0) cube complex, then ([Sag95]):*

1. *Each midcube lies in a unique hyperplane of $\widetilde{X}$.*
2. *Each hyperplane H of $\widetilde{X}$ is a CAT(0) cube complex.*
3. $N(H) \cong H \times [0,1]$, *and is a convex subcomplex of $\widetilde{X}$ in the sense of Definition 4.2.3..*
4. $\widetilde{X} - H$ *consists of exactly 2 connected components.*

If X is a non-positively curved cube complex, then each midcube of $\widetilde{X}$ projects to a midcube of X under the covering map, thus, each hyperplane of $\widetilde{X}$, in the sense of Definition 2.1.14, projects to a hyperplane of X in the sense of Definition 2.1.7; likewise, by Theorem 2.1.16, the preimage in $\widetilde{X}$ of a hyperplane of X must be a collection of hyperplanes of $\widetilde{X}$ in the sense of Definition 2.1.14. In particular, both Definition 2.1.7 and Definition 2.1.14 describe the same objects when $\widetilde{X}$ is a CAT(0) cube complex.

The 1-skeleton $\widetilde{X}^{(1)}$ of a CAT(0) cube complex $\widetilde{X}$ is a metric space with the usual graph metric. A *geodesic* is a minimal-length path in this metric. We emphasise that in the course of this work, whenever we refer to a path or a geodesic in a complex X, we mean a combinatorial path or combinatorial geodesic between vertices of X.

The ensuing lemma is a standard result in the theory, which appears in various forms in different contexts. The version we use can be found in [Wis21, 2.16].

Lemma 2.1.17. *A combinatorial path $\sigma \to \widetilde{X}$ is a geodesic if and only if each edge of σ is dual to a distinct hyperplane of $\widetilde{X}$.*

Lemma 2.1.17 is proven by contradiction: one assumes that σ is a geodesic, and that some hyperplane H crosses σ more than once. Then one considers a minimal area disc diagram (see Definition 2.2.3) bounded by H and a subpath τ of σ with endpoints on $N(H)$, and uses it to show that the path σ' obtained by "going around" H rather than traversing τ is shorter than σ, as in Figure 2.7. A careful proof of this result uses the disc diagram techniques outlined in Section 4.3, and we will omit it as it would constitute too much of a detour.

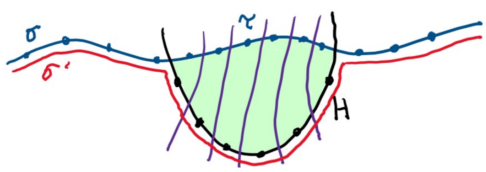

Figure 2.7: Idea of the proof of Lemma 2.1.17.

Definition 2.1.18. A full subcomplex $S \subset X$ is *convex* if for every geodesic $\sigma \to X$ whose endpoints lie in S, then $\sigma \subset S$.

A simple characterisation of convexity is presented in the next remark; it will be useful in Section 4.2.

Remark 2.1.19. Let Y be a subcomplex of a CAT(0) cube complex $\widetilde{X}$. Then Y is convex if and only if the following occurs: for each n-cube c with $n \geq 2$ in $\widetilde{X}$, whenever a corner of c lies in Y, then the cube c lies in Y.

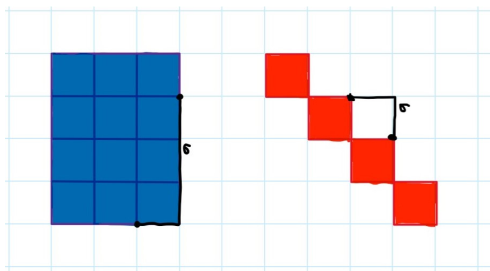

Figure 2.8: Convex (left) and non-convex (right) subcomplexes in the standard cubulation of $\mathbb{E}^2$.

The theorem below lies at the very foundations of much of what is discussed and proven in this text.

Theorem 2.1.20 ([Hag08]). *If $\varphi : Y \to X$ is a local isometry, then:*

1. *The induced map $\tilde{\varphi} : \widetilde{Y} \to \widetilde{X}$ is injective,*

2. *$\widetilde{Y} \hookrightarrow \widetilde{X}$ is a convex subcomplex, and*

3. *$\varphi_* : \pi_1 Y \to \pi_1 X$ is injective.*

Proof. Since the main point that concerns us is the π_1-injectivity, we explain only how it follows from item (1), and omit the rest of the proof.

If a path $\sigma \to Y$ is not nullhomotopic in Y, then it lifts to a non-closed path $\tilde{\sigma} \to \widetilde{Y}$. Since φ is a local isometry, item (1) of the theorem implies that $\tilde{\varphi}(\tilde{\sigma}) \to \widetilde{X}$ cannot be closed, and therefore $\varphi(\sigma) \to X$ is not nullhomotopic.

A handy consequence of the above is:

Lemma 2.1.21. *If X is a CAT(0) cube complex and $Y \subset X$ is a convex subcomplex, then Y is CAT(0).*

Proof. That Y is non-positively curved follows from the fact that links of vertices in Y are restrictions of the corresponding links in X, and since Y is a full subcomplex by definition, the restrictions are also full subcomplexes of links in X; that Y is simply-connected follows from item (3) of Theorem 2.1.20, since the inclusion $i : Y \to X$ is a local isometry.

2.1.2 Wallspaces and dual cube complexes

This section is mainly relevant to the contents of Chapter 3.

Definition 2.1.22. A *wallspace* is a set X together with a collection of *walls* $\{W_i\}_{i\in I} = \mathscr{W}$ where $W_i = \{\overleftarrow{W}_i, \overrightarrow{W}_i\}$ and such that:

1. $\overleftarrow{W}_i \cup \overrightarrow{W}_i = X$ and
2. $\overleftarrow{W}_i \cap \overrightarrow{W}_i = \emptyset$.

Moreover, $\mathscr{W}$ satisfies a *finiteness property*: For every $p,q \in X$ the number of walls separating p and q, denoted by $\#_{\mathscr{W}}(p,q)$, is finite. The $\overleftarrow{W}_i, \overrightarrow{W}_i$ above are the *half-spaces* of W_i.

Wallspaces, it turns out, have been staring us in the face since the very beginning of this journey:

Example 2.1.23. Let $\widetilde{X}$ be a connected CAT(0) cube complex and $\mathscr{W}$ a set of hyperplanes in $\widetilde{X}$ (or rather, a set of pairs of halfspaces determined by the set of hyperplanes). Then $(\widetilde{X}^{(0)}, \mathscr{W})$ is a wallspace. Indeed, hyperplanes separate $\widetilde{X}$ by Theorem 2.1.16, and the distance between pairs of vertices of $\widetilde{X}$ is exactly the number of hyperplanes separating them.

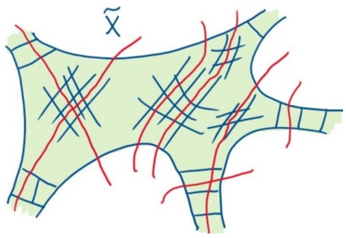

Figure 2.9: A wallspace associated to a CAT(0) cube complex (artistic rendition – prints available upon request).

Example 2.1.24. Let $\widetilde{S} = \mathbb{H}^2$ be be the universal cover of a closed surface S with negative Euler characteristic, and let Ω be a finite collection of curves in S. Let $\mathscr{W}$ be the set of half-spaces of $\mathbb{H}^2$ determined by the lifts of elements in Ω to $\mathbb{H}^2$. Then $(\mathbb{H}^2, \mathscr{W})$ is a wallspace.

Construction 2.1.25 (Sageev's Construction [Sag95])**.** We now describe a general technique that produces, starting from an abstract wallspace, a CAT(0) cube complex C called the *dual* of $(X,\mathscr{W})$. Each 0-cube $v \in C^0$ is a choice of one half-space for each wall W_i such that

1. Any two chosen half-spaces intersect,

2. for $x \in X$, at most finitely many chosen half-spaces do *not* contain x.

C has a 1-cube joining 0-cubes u,v if and only if u and v differ on a unique wall. Higher dimensional cubes are added inductively: for each $n \geq 2$, an n-cube is present whenever its $(n-1)$-skeleton is present.

For any set X, many choices of halfspaces satisfying conditions (1) and (2) above do in fact exist: these are the canonical 0-cubes described below.

We note also that walls in X are in bijective correspondence with hyperplanes in C.

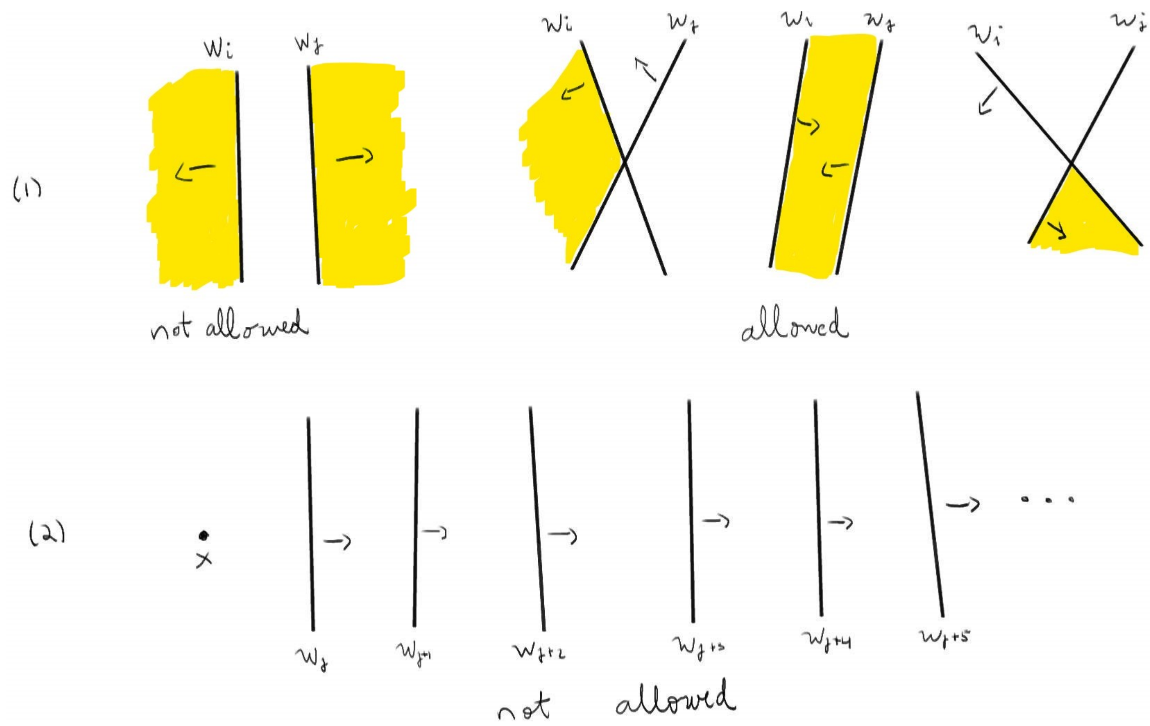

Figure 2.10: Conditions (1) and (2) in Construction 2.1.25.

Remark 2.1.26. When $\mathscr{W}$ is the set of all hyperplanes in a CAT(0) cube complex $\widetilde{X}$, then the dual to $(\widetilde{X},\mathscr{W})$ in Example 2.1.23 is isomorphic to $\widetilde{X}$.

A *canonical 0-cube* is a 0-cube v_x where $x \in X$ and such that for each W_i, the chosen half-space contains x. Note that v_x is indeed a 0-cube, as it satisfies the finiteness property in

Construction 2.1.25. A typical dual cube complex will have many canonical 0-cubes, but it will also have non-canonical ones. For instance, consider the wallspace $(\mathbb{R}^2, \mathscr{W})$ where $\mathscr{W}$ consists of 3 pairwise intersecting lines L_1, L_2, L_3 as in Figure 2.11: the dual cube complex is a 3-cube, seven of whose vertices are canonical 0-cubes, corresponding to the seven connected regions in $\mathbb{R}^2 - \cup_i L_i$. The eighth vertex is non-canonical, and does not correspond to a point in $\mathbb{R}^2$ in any natural way – intuitively, adding it in requires an extra dimension. This is consistent with the fact that a 3-cube cannot be realised as a subspace of $\mathbb{R}^2$.

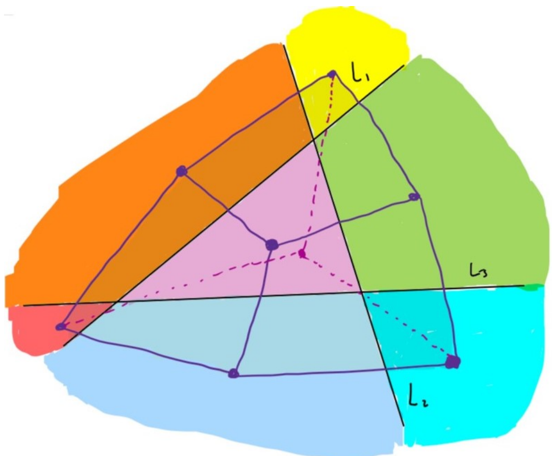

Figure 2.11: The wallspace $(\mathbb{R}^2, \mathscr{W})$ where $\mathscr{W}$ consists of 3 pairwise intersecting lines L_1, L_2, L_3, and its dual cube complex.

Two walls W_i, W_j *cross* if the four sets $\overleftarrow{W}_i \cap \overleftarrow{W}_j, \overleftarrow{W}_i \cap \overrightarrow{W}_j, \overrightarrow{W}_i \cap \overrightarrow{W}_j, \overrightarrow{W}_i \cap \overleftarrow{W}_j$ are non-empty. When no two walls in $(X, \mathscr{W})$ cross, the dual cube complex C is a tree.

Theorem 2.1.27 ([Sag95])**.** *C is a CAT(0) cube complex.*

The proof, though a little verbose, is straightforward. We need to show that C is non-positively curved and simply connected. To this end, we first check that $link(v)$ is simplicial and flag for each $v \in C^0$, and then check that C is connected, and that every closed path $P \to C$ is nullhomotopic.

Proof. First, note that Sageev's construction precludes the existence of two distinct squares glued along two consecutive edges – indeed, let u, v, u' be the consecutive common vertices of two such squares, and let w, z be the other two vertices. Since u, u' are adjacent to both w and z, both u, u' differ on a single choice of halfspace from both w and z, but there are only two choices of halfspaces with that property and one of them corresponds to v. Hence, $link(v)$ is simplicial.

We now show that $link(v)$ is flag. Let T be a 3-cycle in $link(v)$, whose vertices t_1, t_2, t_3 lie on edges e_1, e_2, e_3 of C. Let W_1, W_2, W_3 be the walls dual to e_1, e_2, e_3. Flipping exactly one of W_1,W_2,W_3 yields the terminal vertices of e_1,e_2,e_3. Flipping exactly two of W_1,W_2,W_3 yields three more vertices, and flipping all three W_1,W_2,W_3 yields an eighth vertex. Each of these 8 vertices is adjacent to 3 other vertices (which correspond to flipping one of the three walls) and so they define the 1-skeleton of a 3-cube. Hence, there is a 3-cube attached at v, and T bounds a 2-simplex. Proceeding by induction, we see that the $(n-1)$-skeleton of an n-simplex in $link(v)$ bounds an n-simplex.

Finally, we show that C is simply connected. Condition (2) in Sageev's construction implies that C is connected and path-connected, since a path between vertices p,q can be obtained by flipping a finite number of walls. Now, let $P \to C$ be a closed path and moreover choose P to be shortest within its homotopy class. Let e,e' be an innermost pair of edges in P that are dual to the same wall and let $P = \cdots eQe' \cdots$. If Q is trivial, then e and e' are adjacent and P has a backtrack. Eliminating the backtrack yields a path P' homotopic to P and having $|P'| = |P| - 2$, which contradicts the choice of P. Hence, Q is not trivial. Observe that each wall dual to an edge of Q crosses a wall W, so there is a rectangle in C with sides e,Q,e',Q', and P is homotopic to the path P' obtained by traversing Q' in place of eQe'. Once again this contradicts the initial choice of P. ◻

2.1.3 Cubulated groups and actions on wallspaces

We have been talking about cubulated groups since the introduction to this work. It's about time we formally introduced them.

Definition 2.1.28. A group G is *[cocompactly] cubulated* if it acts properly [and cocompactly] on a CAT(0) cube complex.

When cocompactness of the action is not assumed, the term "proper" should be taken to mean "metrically proper".

To cubulate a group, we often need to find or produce a wallspace on which it acts in the following sense:

Definition 2.1.29. A group G *acts on a wallspace* $(X,\mathscr{W})$ if it acts on X and permutes walls, so for each $W \in \mathscr{W}$ we have that $gW \in \mathscr{W}$ for all $g \in G$, and $\{g\overleftarrow{W}, g\overrightarrow{W}\} \in \mathscr{W}$ whenever $\{\overleftarrow{W}, \overrightarrow{W}\} \in \mathscr{W}$.

An action on a wallspace gives us for free an action on a CAT(0) cube complex:

Lemma 2.1.30. *If G acts on a wallspace $(X,\mathscr{W})$, then G acts on its dual C.*

Proof. The action of $g \in G$ on $(X,\mathscr{W})$, by definition, sends W to a wall $gW \in \mathscr{W}$, and the corresponding halfspaces $\{\overleftarrow{W}, \overrightarrow{W}\}$ to halfspaces $\{g\overleftarrow{W}, g\overrightarrow{W}\}$, so $g \in G$ acts on a vertex

$v \in C^{(0)}$, given by a choice of a halfpace $\overleftrightarrow{W}^v$ where $\leftrightarrow_v \in \{\leftarrow, \rightarrow\}$ for each wall W in $\mathscr{W}$, by sending it to the vertex $v' =: gv$ corresponding to the choice of the halfpace $g\overleftrightarrow{W}^v$ for each wall $W \in \mathscr{W}$. Since an edge e of C corresponds to a pair of choices of halfspaces that differ on a single wall, and the action of G on $(X, \mathscr{W})$ respects this property, the action of G on $C^{(0)}$ extends to $C^{(1)}$, and thus to $C^{(n)}$ for each $n \geq 2$, since the n-cubes of C are fully determined by its 1-skeleton.

Later on, in Proposition 3.2.18 and Theorem 3.2.27, we shall mention, and consequently utilise, criteria that guarantees properness and cocompactness for the action of a group on the dual cube complex associated to a collection of codimension-1 subgroups.

2.1.4 (Virtually) special cube complexes

We now discuss special cube complexes, which are a class of non-positively curved cube complexes that avoid certain hyperplane behaviours. Special cube complexes were introduced by Haglund–Wise in [HW08]; the contents of this section are relevant to Chapter 3, and specifically to the statement of the main result of the chapter, Theorem 3.4.1.

A *special cube complex X* is a non-positively curved cube complex whose immersed hyperplanes

1. are embedded,
2. are two-sided,
3. do not directly self-osculate, and
4. do not inter-osculate.

Condition (1) fails whenever there is "self-crossing", that is, whenever distinct cubes of an immersed hyperplane of X are midcubes of the same cube of X. This is illustrated in Figure 2.12 (far left).

Condition (2) holds if we can consistently direct all the 1-cubes dual to each hyperplane of X. Failure of this condition can be appreciated in Figure 2.12 (centre left).

A hyperplane *U directly self-osculates* if it is dual to distinct directed 1-cubes with the same initial or terminal 0-cube, such as in Figure 2.12 (centre right). We will often omit the adverb and refer to direct self-osculation simply as *self-osculation*.

Two hyperplanes U, V *inter-osculate* at $p \in X^0$ if they are dual to 1-cubes e, f at p that do not meet along the corner of a 2-cube (in other words, e, f are not adjacent in $link(p)$). Failure of Condition (4) is exemplified in Figure 2.12 (far right).

A cube complex X is *virtually special* if it has a finite degree covering that is a special cube complex.

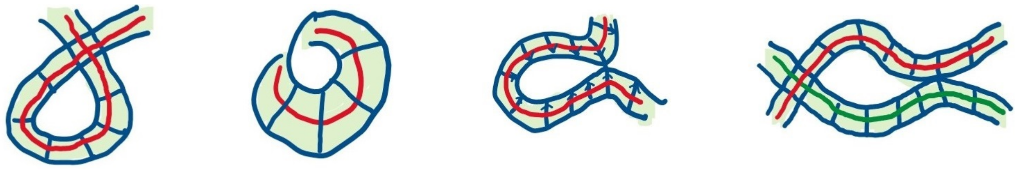

Figure 2.12: The 4 ~~Horsemen of~~ Hyperplane Pathologies.

Figure 2.13: Two non-pathologies: in the first picture there is no direct self-osculation, since the 1-cubes that are meeting along a 0-cube are directed in the same way (this is called *indirect self-osculation*); in the second one, there is no inter-osculation, since the hyperplanes extend and intersect along the square in the “back".

Example 2.1.31. Every graph is special: since all hyperplanes are midpoints of edges, all conditions are automatically satisfied.

We now introduce right angled Artin groups – their importance will presently be made clear.

Definition 2.1.32. Let Γ be a simplicial graph with vertex set V and edge set E, the *right angled Artin group* (“RAAG", for the sake of brevity) $G(\Gamma)$ is the group given by the presentation:

$$G(\Gamma) = \langle x_v : v \in V | [x_u, x_v] : \{u, v\} \in E \rangle \tag{2.1}$$

Salvetti complexes, which are the classifying spaces for RAAGs that we describe below, will be of utmost importance.

Theorem 2.1.33. *For any simplicial graph Γ, the RAAG $G(\Gamma)$ is isomorphic to $\pi_1 R$ for some non-positively curved cube complex R.*

Proof. To define R, we start with the presentation complex $R^{(2)}$ associated to the presentation (2.1) given above. If the girth of Γ is ≥ 4, then $R^{(2)} := R$ is already non-positively curved and is the desired cube complex. Otherwise, we complete $R^{(2)}$ to a non-positively curved cube complex R by gluing k-cubes inductively whenever their skeleton is present.

That is, for each complete subgraph K_n in Γ, there is an n-torus in R whose 1-skeleton is given by the vertices of Γ in K_n, and whose 2-skeleton is given by the edges of Γ in K_n. The result of this procedure is a non-positively curved cube complex because if the boundary of an n-simplex is present in $link(w)$, then that means that there are n generators in (2.1) that commute, and therefore a copy of K_n corresponding to these generators in Γ, so R has an n-cube corresponding to each K_n in $link(w)$. ◻

In fact, Salvetti complexes avoid the four hyperplane pathologies introduced at the start of this section:

Theorem 2.1.34. *The Salvetti complex of a RAAG is special.*

Sketch. First observe that, since each 1-cube in R corresponds to a generator of the RAAG, then each hyperplane must be dual to a unique 1-cube. This implies in particular that no hyperplane can self-osculate, and that no pair of hyperplanes can inter-osculate either. Moreover, hyperplanes are 2-sided because each 2-cube in R corresponds to a pair of commuting generators, and hyperplanes are embedded because, since the defining graph of the RAAG is simplicial, then each commutator comes from a distinct pair of generators (that is, we do not label any square with the same generator on all 4 sides, and no pair of generators arises as the labels of more than 1 square). ◻

A pleasing and fundamental observation is that local isometries interact well with the notion of specialness [HW08]:

Lemma 2.1.35. *If $A \looparrowright B$ is a local isometry and B is special, then A is special. In particular, any covering space $\widehat{B} \to B$ of a special cube complex B is also special.*

Theorem 2.1.34 and Lemma 2.1.35 allow us to deduce:

Theorem 2.1.36. *[HW08, 4.2] A non-positively curved cube complex X is special if and only if there exists a local isometry $X \to R$ where R is a Salvetti complex. In fact, $R = R(\Gamma)$ where Γ is a simplicial graph with a vertex for each hyperplane and such that two vertices are joined by an edge if and only if their corresponding hyperplanes cross.*

Proof. As stated in the theorem, we define a simplicial graph Γ whose vertices are in bijective correspondence with the hyperplanes of X, and whose edges indicate pairs of crossing hyperplanes (note that there is a single edge for each pair of crossing hyperplanes, regardless of the number of connected components of their intersection).

We then construct $R(\Gamma)$ as in Theorem 2.1.33. We label each 1-cube of X with the hyperplane it is dual to, and we orient all 1-cubes so that the attaching map of each 2-cube is a commutator. Indeed, this is always possible because hyperplanes of X are 2 sided. In the same manner, and using that Salvetti complexes are special, we label and orient the edges of $R(\Gamma)$. Thus there is a combinatorial map $X^{(2)} \to R(\Gamma)$ that preserves labels and orientations. To extend this map to X, note that each family of n-pairwise intersecting hyperplanes in X corresponds to an n-clique in Γ, and therefore corresponds to an n-cube of $R(\Gamma)$, in particular every n-cube c of X can be mapped to the n-cube of $R(\Gamma)$ associated to the n hyperplanes crossing at c. Since X has no self-osculating hyperplanes, no two outgoing or ingoing 1-cubes have the same label, so the map is an immersion.

Finally, since there is no inter-osculation in X, if two 1-cubes e, f at a 0-cube of X map to 1-cubes of $R(\Gamma)$ that bound the corner of a square in $R(\Gamma)$, then there is an edge in Γ connecting v_e and v_f, so the hyperplanes associated to e and f have to cross in X, therefore e and f form the corner of a square in X, so the map is a local isometry.

The converse is a consequence of Lemma 2.1.35 and Theorem 2.1.34. ◻

We thus obtain the main algebraic consequence of specialness:

Corollary 2.1.37. *X is special if and only if $\pi_1 X$ is a subgroup of a RAAG.*

Proof. One direction follows from Theorem 2.1.36, together with the π_1-injectivity of local isometries stated in Theorem 2.1.20. The other direction follows from Theorem 2.1.34 and Lemma 2.1.35. ◻

One must pause to appreciate the importance of Corollary 2.1.37: it implies, firstly, that starting with a class of cube complex defined in what appears to be a completely arbitrary manner, we arrive at an algebraic characterisation in terms of a class of groups that is relatively well-understood and that satisfies many desirable properties. This also implies that "being special" is a property that groups satisfy **abstractly (!!!)**.

The previous discussion motivates the following definition:

Definition 2.1.38. A group G is *[virtually] special* if it is isomorphic to the fundamental group of a [virtually] special cube complex.

In the presence of hyperbolicity, virtual specialness is independent of the cubulation:

Theorem 2.1.39 ([HW08])**.** *If X and Y are compact non-positively curved cube complexes with isomorphic, hyperbolic fundamental groups, then X is virtually special if and only if Y is virtually special.*

When discussing specialness, one cannot neglect to mention Agol's fantastic result [Ago13]:

Theorem 2.1.40. *If G is hyperbolic and cocompactly cubulated, then G is virtually special.*

Outside of the hyperbolic context, this fails fabulously: there are examples of Wise [Wis07] of non-positively curved cube complexes that are not special and have no finite-degree coverings at all, and examples with similar[1] properties due to Burger and Mozes [BM00]. These results and more are collected in [Cap19].

2.2 Classical small-cancellation theory

As a warm-up to the cubical small-cancellation theory that will be used in the later chapters of this work, we outline in what follows some of the basic definitions and results in the "classical" version of the theory.

Let $P = \langle S | R \rangle$ be a presentation for a group G, and let $\mathscr{X}(P)$ denote its presentation complex. This is a 2-complex that has a single vertex, an edge for each $s \in S$, and a 2-cell for each $r \in R$, so that $\pi_1 \mathscr{X}(P) = G$. The Cayley graph $Cay(G,S)$ is the 1-skeleton of the universal cover $\widetilde{\mathscr{X}(P)} =: \widetilde{\mathscr{X}}(P)$. The definitions below are stated in terms of arbitrary 2-complexes, but the reader may take $X = \widetilde{\mathscr{X}}(P)$ for the remainder of this section.

Definition 2.2.1. A map $f : X \longrightarrow Y$ between 2-complexes is *combinatorial* if it maps open cells homeomorphically to open cells. A complex is *combinatorial* if all attaching maps are combinatorial (possibly after subdividing).

Definition 2.2.2 (Pieces)**.** Let X be a combinatorial 2-complex. A non-trivial combinatorial path $p \to X$ is a *piece* if there are 2-cells C_1, C_2 in X such that $p \to X$ factors as $p \to \partial C_1 \to X$ and $p \to \partial C_2 \to X$ but there does not exist a homeomorphism $\partial C_1 \to \partial C_2$ such that the following diagram commutes

$$\begin{tikzcd} p \arrow[r] \arrow[d] & \partial C_1 \arrow[d] \arrow[dl] \\ \partial C_2 \arrow[r] & X \end{tikzcd}$$

Definition 2.2.3 (Disc diagram)**.** A *disc diagram D* is a compact contractible combinatorial 2-complex, together with an embedding $D \hookrightarrow S^2$. A *disc diagram in a complex X* is a combinatorial map $D \to X$. The *area* of a disc diagram D is the number of 2-cells in D.

A disc diagram D might map to X in an 'ineffective' way: it might, for instance, be quite far from an immersion. Sometimes it is possible to replace a given diagram with a simpler diagram D' having the same boundary path as D, as we now explain.

[1] Frightful!

Definition 2.2.4. A *cancellable pair* in D is a pair of 2-cells C_1, C_2 in D meeting along a path e such that their boundary paths $\partial C_1, \partial C_2$ map to the same closed path in X, thus making the following diagram commute:

$$\begin{tikzcd} & e \arrow[dl] \arrow[dr] & \\ \partial C_1 \arrow[rr] \arrow[dr] & & \partial C_2 \arrow[dl] \\ & X & \end{tikzcd}$$

A cancellable pair leads to a smaller area disc diagram via the following procedure: cut out $e \cup Int(C_1) \cup Int(C_2)$ and then glue together the paths $\partial C_1 - e$ and $\partial C_2 - e$ to obtain a diagram D' with $\mathsf{Area}(D') = \mathsf{Area}(D) - 2$ and $\partial D' = \partial D$. A diagram is *reduced* if it has no cancellable pairs.

Definition 2.2.5 (Small cancellation conditions)**.** An *arc* in a diagram is a path whose internal vertices have valence 2 and whose initial and terminal vertices have valence ≥ 3. A *boundary arc* is an arc that lies entirely in ∂D. A complex X satisfies the $C(n)$ condition if for every reduced disc diagram $D \to X$, the boundary path of each 2-cell in D either contains a non-trivial boundary arc, or is the concatenation of at least n pieces. The complex X satisfies the $C'(\frac{1}{n})$ condition if for each 2-cell $R \to X$, and each piece $p \to X$ which factors as $p \to R \to X$, then $|p| < \frac{1}{n}|\partial R|$.

Definition 2.2.6 (Ladder)**.** A disc diagram L is a *ladder* if it is the union of a sequence of closed 1-cells and 2-cells $c_1, \dots, c_n$, such that for $1 < j < n$, there are exactly two components in $L - c_j$, and exactly one component in $L - c_1$ and $L - c_n$. Moreover, if c_i is a 1-cell then it is not contained in any other closed c_j.

Definition 2.2.7 (Shells and spurs)**.** A *shell* of D is a 2-cell $C \to D$ whose boundary path $\partial C \to D$ is a concatenation $qp_1 \cdots p_k$ for some $k \leq 3$ where q is a boundary arc in D and $p_1, \dots, p_k$ are non-trivial pieces in the interior of D. The arc q is the *outerpath* of C and the concatenation $p_1 \cdots p_k$ is the *innerpath* of C. A *spur* is a vertex of degree 1 on ∂D.

We now state the two fundamental technical results of small-cancellation theory; proofs of which can be found in [Jan16], for instance.

Theorem 2.2.8 (Greendlinger's Lemma)**.** *Let X be a $C(6)$ complex and $D \to X$ be a minimal area disc diagram, then either*

1. *D is a single cell,*

2. *D is a ladder,*

3. *D has at least three shells and/or spurs.*

Theorem 2.2.9 (The Ladder Theorem)**.** *Let X be a $C(6)$ complex and $D \to X$ be a minimal area disc diagram. If D has exactly 2 shells or spurs, then D is a ladder.*

As mentioned in the introduction, the main motivation for studying small-cancellation groups is that they provide a wealth of examples of hyperbolic groups:

While not stated in this form, the theorem below goes back to the work of Dehn [Deh87]. The reader can also consult [Gro87, Str90].

Theorem 2.2.10. *Let G be a group with a finite $C'(\frac{1}{6})$ or $C(7)$ presentation. Then G is hyperbolic.*

In Chapter 4 we investigate generalisations of the next classical result [Lyn66]:

Theorem 2.2.11. *Let G be a group with a finite $C(6)$ presentation $P = \langle S|R \rangle$. Then the presentation complex $\mathscr{X}(P)$ is aspherical provided that no $r \in R$ is a proper power.*

We also record, for completeness, the following result, which gives sufficient criteria for a classical small-cancellation group to be cocompactly cubulated. A generalisation of this result is explained and used in Chapter 3.

Theorem 2.2.12 ([Wis04])**.** *Let G be a group with a finite $C'(\frac{1}{6})$ or $C'(\frac{1}{4}) - T(4)$ presentation. Then G acts properly and cocompactly on a CAT(0) cube complex C.*

The question of whether one can extend the above result to cubulate general $C(6)$ or $C(7)$ groups is still unresolved.

2.2.1 Rips' exact sequence

This subsection provides some additional context to the discussion and results of Chapter 3.

Given a finitely presented group Q, it is possible to produce a short exact sequence

$$1 \to N \to \Gamma \to Q \to 1$$

where Γ is a small-cancellation group and N is finitely generated. As simple as it may seem, this construction is very useful for producing groups with unexpected combinations of properties – for instance, we will see shortly that it implies the existence of incoherent hyperbolic groups.

Theorem 2.2.13 ([Rip82]). *Let Q be a finitely presented group. For all $n > 0$, there is a $C'(\frac{1}{n})$ group Γ and a short exact sequence*

$$1 \to N \to \Gamma \to Q \to 1 \tag{2.2}$$

where N is generated by 2 elements.

Sketch. Let $Q = \langle a_1, \dots, a_m | r_1, \dots, r_s \rangle$. We will obtain G from Q by adding noise to the presentation as follows. Let

$$\Gamma = \langle a_1, \dots, a_m, x, y | r_1\theta_1, \dots, r_s\theta_s, \{a_i x a_i^{-1}\theta_{i,x}\}, \{a_i^{-1} x a_i \theta_{-i,x}\}, \{a_i y a_i^{-1}\theta_{i,y}\}, \{a_i^{-1} y a_i \theta_{-i,y}\} \rangle,$$

where each θ_j, $\theta_{\pm i,x}$ and $\theta_{\pm i,y}$ is a distinct word on x and y that alternates sufficiently between the letters. Concretely, let

$$\begin{aligned} \theta_1 &= xyx^2yx^3y \dots x^\lambda y \\ \theta_2 &= x^{\lambda+1}yx^{\lambda+2}yx \dots x^{2\lambda}y \\ &\vdots \\ \theta_s &= x^{s\lambda+1}yx^{s\lambda+2}yx \dots x^{s\lambda}y \\ \theta_{1,x} &= x^{(s+1)\lambda+1}yx^{(s+1)\lambda+2}yx \dots x^{(s+1)\lambda}y \\ &\vdots \\ \theta_{-1,x} &= x^{(s+m)\lambda+1}yx^{(s+m)\lambda+2}yx \dots x^{(s+m)\lambda}y \\ &\vdots \end{aligned} \tag{2.3}$$

and so on, where λ is chosen in function of $\max\{|r_i|\}$ and n to force the $C'(\frac{1}{n})$ condition on Γ. There is a map $\Gamma \to Q$ that sends $x, y \mapsto 1$. It is well-defined because the relations $r_1\theta_1, \dots, r_s\theta_s$ map to $r_1, \dots, r_s$. The rest of the relations in the presentation force the subgroup $\langle x, y \rangle$ to be normal in Γ, so $N = ker(G \to Q) = \langle x, y \rangle$ is finitely generated by x and y. $\square$

Remark 2.2.14. For $n \geq 6$, Theorem 2.2.11 implies that the presentation complex associated to the presentation of Γ given above is aspherical. Thus, $cd(G) = gd(G) \leq 2$. Since finitely presented normal subgroups of groups of cohomological dimension 2 must be free or of finite index by [Bie81, 8.6], it follows that if Q is infinite, then N is not finitely presented. Indeed, N is 2-generated so it is straightforward to check that the construction in the proof of Theorem 2.2.11 prevents it from being free (this also follows from [Bie81, 8.7]).

A group is *coherent* if all its finitely generated subgroups are finitely presented; it is *incoherent* if it is not coherent. While coherence might seem like a rather restrictive property, it is satisfied by many naturally occurring classes of groups, including free groups, surface groups, and 3-manifold groups. Since many of these groups are hyperbolic, one might wonder if all hyperbolic groups are coherent.

Rips used his construction to answer this question in [Rip82].

Corollary 2.2.15. *There exist incoherent hyperbolic groups.*

Proof. Take $n \geq 6$, apply the Rips construction to any infinite finitely presented group Q, and use Remark 2.2.14. ∎

The *membership problem for finitely generated subgroups*, also known as the *generalised word problem for finitely generated subgroups* is the questions of whether, for a given group Γ, there exists an algorithm that decides for all finitely generated subgroups $H < \Gamma$ and for all elements $g \in \Gamma$, if $g \in H$. The membership problem reduces to the word problem by setting $H = \{1\}$.

The *generalised rank problem* is the questions of whether, given a class of groups $\mathscr{G}$, there exists an algorithm that decides for all groups $\Gamma \in \mathscr{G}$ and for each $n \in \mathbb{N}$, if $rk(\Gamma) \leq 2$ or $rk(\Gamma) \geq n$. As suggested by the name, this is a variation of the *rank problem*, which asks whether there exists an algorithm that decides, for each $n \in \mathbb{N}$, if $rk(\Gamma) = n$.

In the case of hyperbolic groups, we can address both of these problems using the Rips construction [Rip82, BMS94]:

Corollary 2.2.16. *There exist hyperbolic groups with unsolvable membership problem for finitely generated subgroups, and hyperbolic groups for which the generalised rank problem is undecidable.*

Proof. For the membership problem, let Q be a finitely presented group with unsolvable word problem (such groups exist by [Nov55] and [Boo58]), choose $n \geq 6$ and apply the Rips construction to Q. Note that $g \in N \triangleleft \Gamma$ if and only if $g =_Q 1$. Since the word problem is undecidable in Q, it is impossible to establish whether the second equality holds. Thus, it is impossible to establish whether $g \in N$, so the membership problem is undecidable for $N \triangleleft \Gamma$.

For the generalised rank problem, let $Q_n = Q * Q * \cdots * Q$ where Q is given by an arbitrary finite presentation, and again choose $n \geq 6$. Apply the Rips construction to Q_n. By Grushko's Theorem, either $rk(\Gamma) = rk(N) = 2$ or $rk(\Gamma) \geq n$, depending on whether Q presents the trivial group or not. But this is undecidable for arbitrary finitely presented groups, thus, we cannot determine whether $rk(\Gamma) = 2$ or $rk(\Gamma) \geq n$. ∎

In contrast to the above, the solvability of the membership problem for finitely **presented** subgroups of hyperbolic groups is still very much an open question.

Chapter 3

A cubical Rips construction

3.1 Introduction

The Rips exact sequence, first introduced by Rips in [Rip82], is a useful tool for producing examples of groups satisfying combinations of properties that are not obviously compatible. It works by taking as an input an arbitrary finitely presented group Q, and producing as an output a hyperbolic group Γ that maps onto Q with finitely generated kernel. The "output group" Γ is crafted by adding generators and relations to a presentation of Q, in such a way that these relations create enough "noise" in the presentation to ensure hyperbolicity. One can then lift pathological properties of Q to (some subgroup of) Γ. For instance, Rips used his construction to produce the first examples of incoherent hyperbolic groups, hyperbolic groups with unsolvable generalised word problem, hyperbolic groups having finitely generated subgroups whose intersection is not finitely generated, and hyperbolic groups containing infinite ascending chains of r-generated groups.

The purpose of this chapter is to present a new variation of the Rips exact sequence. Our main result is:

Theorem 3.1.1 (Theorem 3.4.1)**.** *Let Q be a finitely presented group and G be the fundamental group of a compact special cube complex X. If G is hyperbolic and non-elementary, then there is a short exact sequence*

$$1 \to N \to \Gamma \to Q \to 1$$

where

1. *Γ is a hyperbolic, cocompactly cubulated group,*

2. *$N \cong G/K$ for some $K < G$,*

3. *max*$\{cd(G), 2\} \geq cd(\Gamma) \geq cd(G) - 1$. *In particular,* Γ *is torsion-free.*

Remark 3.1.2. By Agol's Theorem [Ago13], the group Γ obtained in Theorem 3.4.1 is in fact virtually special.

Many variations of Rips' original construction have been produced over the years by a number of authors, including Arzhantseva-Steenbock [AS14], Barnard-Brady-Dani [BBD07], Baumslag-Bridson-Miller-Short [BBMS00], Belegradek-Osin [BO08], Bridson-Haefliger [BH99], Bumagin-Wise [BW05], Haglund-Wise [HW08], Ollivier-Wise [OW07], and Wise [Wis03, Wis98]. Below is a sample of their corollaries:

- There exist non-Hopfian groups with Kazhdan's Property (T) [OW07].
- Every countable group embeds in the outer automorphism group of a group with Kazhdan's Property (T) [OW07, BO08].
- Every finitely presented group embeds in the outer automorphism group of a finitely generated, residually finite group [Wis03].
- There exists an incoherent group that is the fundamental group of a compact negatively curved 2-complex [Wis98].
- There exist hyperbolic special groups that contain (non-quasiconvex) non-separable subgroups [HW08].
- Property (T) and property FA are not recursively recognizable among hyperbolic groups [BO08].

The groups in Rips' original constructions are cubulable by [Wis04], as are the groups in [HW08]; on the other extreme, the groups produced in [OW07] and in [BO08] can have Property (T), so will not be cubulable in general (see [NR97]).

A notable limitation of all available Rips-type techniques is that the hyperbolic group Γ surjecting onto Q will have cohomological dimension at most equal to 2. This is unsurprising, since, in a precise sense, "most" hyperbolic groups are 2-dimensional (see [Gro93] and [Oll05]). Moreover, examples of hyperbolic groups having large cohomological dimension are scarce: Gromov conjectured in [Gro87] that all constructions of high-dimensional hyperbolic groups must utilise number-theoretic techniques, and later on, Bestvina [Bes] made this precise by asking whether for every $K > 0$ there is an $N > 0$ such that all hyperbolic groups of (virtual) cohomological dimension $\geq N$ contain an arithmetic lattice of dimension $\geq K$. Both of these questions have been answered in the negative by work of a number of people, including Mosher-Sageev [MS97], Januszkiewicz-Swiatkowski [JŚ03], and later on

Fujiwara-Manning [FM10], and Osajda [Osa13], but flexible constructions are still difficult to come by.

Theorem 3.4.1 produces cocompactly cubulated hyperbolic groups containing quotients of arbitrary special hyperbolic groups. While our construction particularises to Rips' original result, it can also produce groups with large cohomological dimension. Thus, it serves to exhibit a collection of examples of hyperbolic groups that is new and largely disjoint from that produced by all other Rips'-type theorems.

Most versions of Rips' construction, including the original, rely on some form of small cancellation. This is what imposes a bound on the dimension of the groups thus obtained: group presentations are inherently 2-dimensional objects, and one can prove that the presentation complexes associated to (classical and graphical) small cancellation presentations are aspherical.

We rely instead on *cubical* presentations and *cubical* small cancellation theory, which is intrinsically higher-dimensional. Roughly speaking, cubical small cancellation theory considers higher dimensional analogues of group presentations: pieces in this setting are overlaps between higher dimensional subcomplexes, and cubical small cancellation conditions measure these overlaps. This viewpoint allows for the use of non-positively curved cube complexes and their machinery, and has proved fruitful in many contexts, most notably in Agol's proofs of the Virtual Haken and Virtual Fibred Conjectures [Ago13, Ago08], which build on work of Wise [Wis21] and his collaborators [BW12, HW12, HW15].

While many groups have convenient cubical presentations, producing these, or proving that they do satisfy useful cubical small cancellation conditions, is difficult in general. Some examples are discussed in [Wis21], [AH22], and [JW22]. Other than these, we are not aware of instances where explicit examples of non-trivial cubical small cancellation presentations are given, nor of many results producing families of examples with some given list of properties. This note can be viewed as one such construction, and can be used to produce explicit examples that are of a fundamentally different nature to those already available.

3.1.1 Structure of the chapter

In Section 3.2 we present the necessary background on hyperbolicity, quasiconvexity and cubical small cancellation theory. In Section 3.3 we state and prove Theorem 3.3.2, which is the main technical result, and also state and prove some auxiliary lemmas. In Section 3.4 we give the proof of Theorem 3.4.1. Finally, in Section 3.5 we review some standard material on the cohomological dimension of groups, and analyse the cohomological dimension of Γ.

3.2 Background

We utilise the following theorem of Arzhantseva [Arz01]:

Theorem 3.2.1. *Let G be a non-elementary torsion-free hyperbolic group and H a quasi-convex subgroup of G of infinite index. Then there exist infinitely many $g \in G$ for which the subgroup $\langle H, g\rangle$ is isomorphic to $H * \langle g\rangle$ and is quasiconvex in G.*

Since cyclic subgroups of a hyperbolic group G are necessarily quasiconvex, one can repeatedly apply Theorem 3.2.1 to produce quasiconvex free subgroups of any finite rank:

Corollary 3.2.2. *If G is a non-elementary torsion-free hyperbolic group, then for every $n \in \mathbb{N}$ there exists a quasiconvex subgroup $F_n < G$.*

Recall that for a graph B, a subgraph $A \subset B$ is *full* if whenever vertices $a_1, a_2 \in A$ are joined by an edge e of B, then $e \subset A$. In other words, A is the subgraph of B induced by A^0. A map $X \to Y$ between cell complexes is *combinatorial* if it maps cells to cells of the same dimension. An *immersion* is a local injection.

Definition 3.2.3. A *local isometry* $\varphi : Y \to X$ between non-positively curved cube complexes is a combinatorial map such that for each $y \in Y^0$ and $x = \varphi(y)$, the induced map $\varphi : link(y) \to link(x)$ is an injection of a full subgraph.

A more visual way to think about local isometries is the following: an immersion φ is a local isometry if whenever two edges $\varphi(e), \varphi(f)$ form the corner of a square in X, then e, f already formed the corner of a square in Y.

A key property of local isometries is that they are π_1-injective. It is then natural to ask *which* subgroups of the fundamental group of a non-positively curved cube complex can be realised by local isometries with compact domain. In the setting of non-positively curved cube complexes with hyperbolic fundamental group, one large class of subgroups having this property is the class of quasiconvex subgroups. This is proved in [Hag08], and collected in [Wis21, 2.31 and 2.38] as presented below. A generalisation to the setting of relatively hyperbolic groups is treated in [SW15].

Definition 3.2.4. A subspace $\widetilde{Y} \subset \widetilde{X}$ is *superconvex* if it is convex and for every bi-infinite geodesic line L, if $L \subset N_r(\widetilde{Y})$ for some $r > 0$, then $L \subset \widetilde{Y}$. A map $Y \to X$ is *superconvex* if the induced map between universal covers $\widetilde{Y} \to \widetilde{X}$ is an embedding onto a superconvex space.

Proposition 3.2.5. *Let X be a compact non-positively curved cube complex with $\pi_1 X$ hyperbolic. Let $H < \pi_1 X$ be a quasiconvex subgroup and let $C \subset \widetilde{X}$ be a compact subspace. Then there exists a superconvex H-cocompact subspace $\widetilde{Y} \subset \widetilde{X}$ with $C \subset \widetilde{Y}$.*

Proposition 3.2.6. *Let X be a compact non-positively curved cube complex with $\pi_1 X$ hyperbolic. Let $H < \pi_1 X$ be a quasiconvex subgroup. Then there exists a local-isometry $Y \to X$ with $\pi_1 Y = H$.*

3.2.1 Cubical small cancellation theory

Cubical presentations, cubical small cancellation theory, and many related notions were introduced in [Wis21]. We recall them below.

Definition 3.2.7. A *cubical presentation* $\langle X | \{Y_i\}\rangle$ consists of a connected non-positively curved cube complex X together with a collection of local isometries of connected non-positively curved cube complexes $Y_i \xrightarrow{\varphi_i} X$. In this setting, we shall think of X as a "generator" and of the $Y_i \to X$ as "relators". The fundamental group of a cubical presentation is defined as $\pi_1 X / \langle\langle \{\pi_1 Y_i\} \rangle\rangle$.

Associated to a cubical presentation $\langle X | \{\varphi_i : Y_i \to X\}\rangle$ there is a *coned-off space* X^* obtained from $(X \cup \{Y_i \times [0,1]\})/\{(y_i, 1) \sim \varphi_i(y_i)\}$ by collapsing each $Y_i \times \{0\}$ to a point. This space has a natural cell-structure consisting of the cubes of X and pyramids over the images of the Y_i with cone-vertices corresponding to the apexes. By the Seifert-Van Kampen Theorem, the group $\pi_1 X / \langle\langle \{\pi_1 Y_i\} \rangle\rangle$ is isomorphic to $\pi_1 X^*$. Thus, the coned-off space is a presentation complex of sorts for $\langle X | \{Y_i\}\rangle$. In practice, when discussing cubical presentations, we often have in mind the coned-off space X^*, rather than the abstract cubical presentation.

Remark 3.2.8. A group presentation $\langle a_1, \dots, a_s | r_1, \dots, r_m\rangle$ can be interpreted cubically by letting X be a bouquet of s circles and letting each Y_i map to the path determined by r_i. On the other extreme, for every non-positively curved cube complex X there is a "free" cubical presentation $X^* = \langle X | \ \rangle$ with fundamental group $\pi_1 X = \pi_1 X^*$.

In the cubical setting, there are 2 types of pieces: *wall-pieces* and *cone-pieces*. Cone-pieces are very much like pieces in the classical sense – they measure overlaps between relators in the presentation. On the other hand, wall-pieces measure the overlaps between cone-cells and *rectangles* (hyperplane carriers) – wall-pieces are always trivial in the classical case, since the square part of X^* coincides with the 1-skeleton of the presentation complex.

Precise definitions follow.

Definition 3.2.9 (Elevations)**.** Let $Y \to X$ be a map and $\hat{X} \to X$ a covering map. An *elevation* $\hat{Y} \to \hat{X}$ is a map satisfying

1. $\hat{Y}$ is connected,

2. $\hat{Y} \to Y$ is a covering map, the composition $\hat{Y} \to Y \to X$ equals $\hat{Y} \to \hat{X} \to X$, and

3. assuming all maps involved are basepoint preserving, $\pi_1\hat{Y}$ equals the preimage of $\pi_1\hat{X}$ in $\pi_1 Y$.

Notation 3.2.10. In the entirety of this text, a path $\sigma \to X$ is assumed to be a combinatorial path mapping to the 1-skeleton of X.

Definition 3.2.11 (Pieces)**.** Let $\langle X|\{Y_i\}\rangle$ be a cubical presentation. An *abstract contiguous cone-piece* of Y_j in Y_i is an intersection $\widetilde{Y}_j \cap \widetilde{Y}_i$ where $\widetilde{Y}_j, \widetilde{Y}_i$ are elevations of Y_j and Y_i to $\widetilde{X}$, and either $i \neq j$ or $i = j$ but $\widetilde{Y}_j \neq \widetilde{Y}_i$. A *cone-piece* of Y_j in Y_i is a path $p \to P$ in an abstract contiguous cone-piece of Y_j in Y_i. An *abstract contiguous wall-piece* of Y_i is an intersection $N(H) \cap \widetilde{Y}_i$ where $N(H)$ is the carrier of a hyperplane H that is disjoint from $\widetilde{Y}_i$. To avoid having to deal with empty pieces, we shall assume that H is dual to an edge with an endpoint on $\widetilde{Y}_i$. A *wall-piece* of Y_i is a path $p \to P$ in an abstract contiguous wall-piece of Y_i[1].

A *piece* is either a cone-piece or a wall-piece.

Remark 3.2.12. In Definition 3.2.11, two lifts of a cone Y are considered identical if they differ by an element of $Stab_{\pi_1 X}(\widetilde{Y})$. This is in keeping with the conventions of classical small cancellation theory, where overlaps between a relator and any of its cyclic permutations are not regarded as pieces. This hypothesis facilitates replacing relators by their proper powers to achieve good small cancellation conditions in some cases.

The $C(p)$ and $C'(\frac{1}{p})$ conditions are now defined as in the classical case (making no distinction between the two types of pieces when counting them). Namely:

Definition 3.2.13. A cubical presentation X^* satisfies the $C(p)$ *small cancellation condition* if no essential closed path $\sigma \to Y_i$ is the concatenation of fewer than p pieces, and the $C'(\frac{1}{p})$ *condition*[2] if whenever $\mu \to X^*$ is a piece in an essential closed path $\sigma \to Y_i$, then $|\mu| < \frac{1}{p}|\sigma|$, where $|\mu|$ is the distance between endpoints of $\widetilde{\mu} \subset \widetilde{X}$.

As in the classical case, if the fundamental group of X in a cubical presentation $X^* = \langle X|\{Y_i\}\rangle$ is hyperbolic, sufficiently good small cancellation conditions lead to hyperbolicity. The following form of [Wis21, 4.7] follows immediately from the fact that a cubical presentation that is $C'(\frac{1}{\alpha})$ for $\alpha \geq 12$ can be endowed with a non-positively curved angling rule that satisfies the short innerpaths condition when $\alpha \geq 14$ ([Wis21, 3.32 and 3.70]):

[1]The "abstract contiguous cone-piece" and "abstract contiguous wall-piece" terminology comes from the fact that it is also a priori necessary to consider "non-contiguous cone-pieces" and "non-contiguous wall-pieces", however [Wis21, 3.7] shows that one can limit oneself to the analysis of contiguous pieces.

[2]These are sometimes called the *contextual* $C(p)$ and $C'(\frac{1}{p})$ conditions – see [Wis21, 3.d].

Theorem 3.2.14. *Let X^* be a cubical presentation satisfying the $C'(\frac{1}{p})$ small cancellation condition for $p \geq \frac{1}{14}$. Suppose $\pi_1 X$ is hyperbolic and X^* is compact. Then $\pi_1 X^*$ is hyperbolic.*

Definition 3.2.15. A collection $\{H_1, \ldots, H_r\}$ of subgroups of a group G is *malnormal* provided that $H_i^g \cap H_j = 1$ unless $i = j$ and $g \in H_i$.

Compactness, malnormality and superconvexity will together guarantee the existence of a uniform bound on the size of both cone-pieces and wall-pieces. This is the content of [Wis21, 2.40 and 3.52], which we recall below:

Lemma 3.2.16. *Let X be a non-positively curved cube complex with $\pi_1 X$ hyperbolic. For $1 \leq i \leq r$, let $Y_i \to X$ be a local-isometry with Y_i compact, and assume the collection $\{\pi_1 Y_1, \ldots, \pi_1 Y_r\}$ is malnormal. Then there is a uniform upper bound D on the diameters of intersections $g\widetilde{Y}_i \cap h\widetilde{Y}_j$ between distinct $\pi_1 X$-translates of their universal covers in $\widetilde{X}$.*

Lemma 3.2.17. *Let Y be a superconvex cocompact subcomplex of a CAT(0) cube complex X. There exists $D \geq 0$ such that the following holds: For each $n \geq 0$, if $I_1 \times I_n \to X$ is a combinatorial strip whose base $0 \times I_n$ lies in Y, and such that $\mathrm{d}((0,0),(0,n)) \geq D$, then $I_1 \times I_n$ lies in Y.*

Recall that a *wallspace* is a set X together with a collection of *walls* $\{W_i\}_{i \in I} = \mathscr{W}$ where $W_i = \{\overleftarrow{W}_i, \overrightarrow{W}_i\}$ and $\overleftarrow{W}_i, \overrightarrow{W}_i \subset X$ for each $i \in I$, and such that:

1. $\overleftarrow{W}_i \cup \overrightarrow{W}_i = X$ and
2. $\overleftarrow{W}_i \cap \overrightarrow{W}_i = \emptyset$.

Moreover, $\mathscr{W}$ satisfies a *finiteness property*: For every $p, q \in X$ the number of walls separating p and q, denoted by $\#_{\mathscr{W}}(p,q)$, is finite. The $\overleftarrow{W}_i, \overrightarrow{W}_i$ above are the *half-spaces* of W_i.

Once we have specified a cubical presentation, we will cubulate its fundamental group $\pi_1 X^*$ via Sageev's construction, which produces a CAT(0) cube complex C that is dual to a wallspace. We will assume the reader is familiar with this procedure. Good references include [Sag95, CN05, BSV14]. Cocompactness of the action on the dual cube complex will readily follow from Proposition 3.2.18, which is a well-known result of Sageev ([Sag97]). Properness is more delicate, and will follow from Theorem 3.2.27 once we know that $\pi_1 X^*$ is hyperbolic, since in that case $\pi_1 X^*$ has no infinite torsion subgroups.

Proposition 3.2.18. *Let G be hyperbolic and $\{H_1, \ldots, H_n\}$ be a collection of quasiconvex subgroups. Then the action of G on the dual CAT(0) cube complex C is cocompact.*

Before stating Theorem 3.2.27, we need some definitions.

Definition 3.2.19. Let $Y \to X$ be a local isometry. $Aut_X(Y)$ is the group of automorphisms $\psi : Y \to Y$ such that the diagram below is commutative:

$$\begin{tikzcd} Y \arrow[r, "\psi"] \arrow[dr] & Y \arrow[d] \\ & X \end{tikzcd}$$

If Y is simply connected, then $Aut_X(Y)$ is equal to $Stab_{\pi_1 X}(Y)$.

Definition 3.2.20. A cubical presentation $\langle X|Y_i\rangle$ satisfies the B(6) condition if it satisfies the following conditions:

1. (Small Cancellation) $\langle X|Y_i\rangle$ satisfies the $C'(\frac{1}{\alpha})$ condition for $\alpha \geq 14$.
2. (Wallspace Cones) Each Y_i is a wallspace where each wall in Y_i is the union $\sqcup U_j$ of a collection of disjoint embedded 2-sided hyperplanes in Y_i, and there is an embedding $\sqcup N(U_j) \to Y_i$ of the disjoint union of their carriers into Y_i. Each such collection separates Y_i. Each hyperplane in Y_i lies in a unique wall.
3. (Hyperplane Convexity) If $P \to Y_i$ is a path that starts and ends on vertices on 1-cells dual to a hyperplane U of Y_i and P is the concatenation of at most 7 pieces, then P is path homotopic in Y_i to a path $P \to N(U) \to Y_i$.
4. (Wall Convexity) Let S be a path in Y_i that starts and ends with 1-cells dual to the same wall of Y_i. If S is the concatenation of at most 7 pieces, then S is path-homotopic into the carrier of a hyperplane of that wall.
5. (Equivariance) The wallspace structure on each cone Y is preserved by $Aut_X(Y)$.

Historical Remark 3.2.21. In the setting of classical small cancellation theory, the $B(2n)$ condition was defined in [Wis04]. Specifically, the “classical" $B(2n)$ condition states that for each 2-cell R in a 2-complex X, and for each path $S \to \partial R$ which is the concatenation of at most n pieces in X, we have $|S| \leq \frac{1}{2}|\partial R|$. The classical $B(2n)$ condition is intermediate to the $C'(\frac{1}{2n})$ and $C(2n)$ conditions in the sense that $C'(\frac{1}{2n}) \Rightarrow B(2n) \Rightarrow C(2n)$. While not a perfect parallel with the notion considered here, the notation is meant to suggest the fact that, in the classical setting, the $B(6)$ condition is sufficient to guarantee the existence of a wallspace structure on X that leads to cocompact cubulability.

The B(6) condition is extremely useful because it facilitates producing a wallspace structure on the coned-off space X^* by starting *only* with a wallspace structure (satisfying some extra conditions) on each of the cones. This is done by defining an equivalence relation $\sim$ on the hyperplanes of $\widetilde{X^*}$ as explained below.

Definition 3.2.22. Let U and U' be hyperplanes in the underlying cube complex of $\widetilde{X^*}$. Then $U \sim U'$ provided that for some translate of some cone Y_i in $\widetilde{X^*}$, the intersections $U \cap Y_i$ and $U' \cap Y_i$ lie in the same wall of Y_i. A *wall* of $\widetilde{X^*}$ is a collection of hyperplanes in the underlying cube complex of $\widetilde{X^*}$ corresponding to an equivalence class.

That this equivalence relation does in fact define a wallspace structure on X^* when the B(6) condition is satisfied is the content of [Wis21, 5.f].

Definition 3.2.23. A hyperplane U is *m-proximate* to a 0-cube v if there is a path $P = P_1,\dots,P_m$ such that each P_i is either a single edge or a piece, v is the initial vertex of P_1 and U is dual to an edge in P_m. A wall is *m-proximate* to v if it has a hyperplane that is m-proximate to v. A hyperplane is *m-far* from a 0-cube if it is not m'-proximate to it for any $m' \leq m$.

Definition 3.2.24. A hyperplane U of a cone Y is *piecefully convex* if the following holds: For any path $\tau\rho \to Y$ with endpoints on $N(U)$, if τ is a geodesic and ρ is trivial or lies in a piece of Y containing an edge dual to U, then $\tau\rho$ is path-homotopic in Y to a path $\mu \to N(U)$.

The following is remarked upon in [Wis21, 5.43]. We write $\widetilde{N}(U) := \widetilde{N(U)}$.

Proposition 3.2.25. *Let K be the maximal diameter of any piece of Y_i in X^*. Then a hyperplane U of Y_i is piecefully convex provided that its carrier $N(U)$ satisfies: $\mathsf{d}_{\widetilde{Y}_i}(g\widetilde{N}(U), \widetilde{N}(U)) > K$ for any translate $g\widetilde{N}(U) \neq \widetilde{N}(U) \subset \widetilde{Y}_i$.*

Definition 3.2.26 (Cut by a wall)**.** Let $g \in G$ be an element acting on $\widetilde{X}$. An *axis* $\mathbb{R}_g$ for g is a g-invariant copy of $\mathbb{R}$ in $\widetilde{X}$. An element g is *cut by a wall* W if $g^n W \cap \mathbb{R}_g = \{n\}$ for all $n \in \mathbb{Z}$.

The theorem below is a restatement of [Wis21, 5.44], together with [Wis21, 5.45], and the fact that the short innerpaths condition is satisfied when $C'(\frac{1}{\alpha})$ holds for $\alpha \geq 14$.

Theorem 3.2.27. *Suppose $X^* = \langle X | \{Y_i\} \rangle$ satisfies the following hypotheses:*

1. *X^* satisfies the B(6) condition.*
2. *Each hyperplane U of each cone Y_i is piecefully convex.*
3. *Let $k \to Y \in \{Y_i\}$ be a geodesic with endpoints p,q. Let U_1 and U_1' be distinct hyperplanes in the same wall w_1 of Y. Suppose k traverses a 1-cell dual to U_1, and either U_1' is 1-proximate to q or k traverses a 1-cell dual to U_1'. Then there is a wall w_2 in Y that separates p,q but is not 2-proximate to p or q.*
4. *Each infinite order element of $Aut(Y_i)$ is cut by a wall.*

Then the action of $\pi_1 X^*$ on its dual cube complex C has torsion stabilizers.

3.3 Cubical noise

In the classical setting, there are two essentially distinct strategies for producing group presentations satisfying good small cancellation conditions: taking large enough powers of the relators, and adding “noise" to the presentation by multiplying each relator by a sufficiently long, suitably chosen word. In the cubical setting, taking powers of relators translates to taking finite-degree covers of the cycles that represent the relators, and this method generalises to taking finite-degree covers of cube complexes with more complicated fundamental groups. This is the line of inquiry that has been most explored, and for which there exist useful theorems producing cubical small cancellation. This will, however, not be suitable for our applications, because once a cubical presentation $\langle X | \{\hat{Y}_i \to X\}\rangle$ has been obtained by taking covers, any modifications of the cones (other than taking further covers) will dramatically affect the size of pieces, and possibly invalidate whatever small cancellation conclusions had been attained. Thus, we instead prove a cubical small cancellation theorem that builds on the idea of adding noise to a presentation. The procedure we describe will be more stable, in the sense that slightly perturbing the choice of cones will not affect the small cancellation conclusions.

Remark 3.3.1. We state and prove Theorem 3.3.2 in more generality than that which is needed for later applications. In practice the reader may take Y to be a bouquet of finitely many circles, as this is all that is required for the proof of Theorem 3.4.1. It is also worth noting that while the statement of Theorem 3.3.2 (3) requires that X and Y be special, the proof only uses that hyperplanes are embedded and 2-sided.

Theorem 3.3.2. *Let X and Y be compact non-positively curved cube complexes with hyperbolic fundamental groups and let $\mathscr{H}$ be the set of hyperplanes of X. Let $\{H_1,\dots,H_r\}$ be a malnormal collection of free, non-abelian quasiconvex subgroups of $\pi_1 X$, and suppose that $H_i \cap Stab(\widetilde{U})$ is trivial or equal to H_i for all $U \in \mathscr{H}$. Let $y_1,\dots,y_r$ be words in $\pi_1 X * \pi_1 Y$. Then for each $\alpha \geq 1$ there are cyclic subgroups $\langle w_i\rangle \subset H_i * \langle y_i\rangle$ such that $w_i = w_i' y_i$ where $w_i' \in H_i$ for each $i \in \{1,\dots r\}$ and:*

1. *The group $\pi_1 X * \pi_1 Y / \langle\langle w_1,\dots,w_r\rangle\rangle$ has a cubical presentation satisfying the $C'(\frac{1}{\alpha})$ condition.*
2. *If $\alpha \geq 14$, then the group $\pi_1 X * \pi_1 Y / \langle\langle w_1,\dots,w_r\rangle\rangle$ is hyperbolic.*
3. *If X and Y are special, there is an $\alpha_0 \geq 14$ such that if $\alpha \geq \alpha_0$, then the group $\pi_1 X * \pi_1 Y / \langle\langle w_1,\dots,w_r\rangle\rangle$ acts properly and cocompactly on a CAT(0) cube complex.*

Remark 3.3.3. The reader might wish to compare Theorem 3.3.2 with [Wis21, 5.48], which is the analogous result for finite-degree coverings, and whose proof informs the proof below.

Proof. **Obtaining small cancellation:** Since $H_1,\dots,H_r$ are quasiconvex subgroups of $\pi_1 X * \pi_1 Y$ and $\pi_1 X * \pi_1 Y$ is hyperbolic, then Corollary 3.2.6 implies that there are based local isometries $C_1 \to X \vee Y,\dots,C_r \to X \vee Y$ of superconvex subcomplexes with $\pi_1 C_i \cong H_i$ for each $i \in \{1,\dots,r\}$. So there is a cubical presentation $(X \vee Y)^* = \langle X \vee Y | \{C_i \to X \vee Y\}_{i=1}^r \rangle$ with fundamental group $\pi_1 X * \pi_1 Y / \langle\langle \pi_1 C_1,\dots,\pi_1 C_r \rangle\rangle$. By Lemma 3.2.16 and Lemma 3.2.17, malnormality of $\{H_1,\dots,H_r\}$ and superconvexity of the $\{C_1,\dots,C_r\}$ ensures that there is a uniform upper bound K on the diameter of pieces.

By hyperbolicity, any cyclic subgroup of $\pi_1 X * \pi_1 Y$ is quasiconvex. So for any choice of cyclic subgroups $\langle w_i \rangle < \pi_1 X * \pi_1 Y$ with $i \in \{1,\dots,r\}$ there are based local isometries $W_i \to X \vee Y$ with $\pi_1 W_i \cong \langle w_i \rangle$.

Let $\alpha \geq 1$, and choose each $\sigma_i \to X \vee Y$ so that $\sigma_i = \sigma_i' \gamma_i$ and satisfying:

1. σ_i' is a based closed path in $C_i \subset X$,
2. σ_i' is not a proper power, and does not contain subpaths of length $\geq K\alpha$ that are proper powers,
3. $\sigma_i' \gamma_i$ does not have any backtracks,
4. the W_i corresponding to $\langle w_i \rangle := \langle \sigma_i \rangle$ has diameter $||W_i|| \geq K\alpha^2$.

For instance, one can choose σ_i' to be of the form $\sigma_i' = \lambda_1\lambda_2\lambda_1\lambda_2^2 \dots \lambda_1\lambda_2^{K\alpha}$ where λ_1,λ_2 are paths representing distinct generators of the corresponding $H_i < \pi_1 X$. Without loss of generality, we can assume that the λ_1,λ_2 and the $\gamma_1,\dots,\gamma_r$ have minimal length in their homotopy classes, and therefore that none of the λ_i or γ_i have any backtracks, so any backtracks in $\sigma_i'\gamma_i$ arise from cancellation between σ_i' and γ_i. If any cancellation happens, we can rechoose λ_1 and λ_2 to eliminate it (for instance, by shortening λ_1 and λ_2).

Pieces in each C_i have size bounded by K, and each $W_i \to X \vee Y$ factors through the corresponding C_i, so the size of pieces between different cone-cells or between cone-cells and rectangles is bounded by K, the size of pieces between a cone-cell and itself is bounded by $K\alpha - 1$ and $||W_i|| \geq (K\alpha)! + K\alpha \geq K\alpha^2$, so $\langle X \vee Y | \{W_i \to X \vee Y\}_{i=1}^r \rangle$ satisfies the $C'(\frac{1}{\alpha})$ condition.

Obtaining hyperbolicity: As explained above, the $\langle w_i \rangle$ can be chosen so that $(X \vee Y)^* = \langle X \vee Y | \{W_i \to X \vee Y\}_{i=1}^r \rangle$ satisfies the $C'(\frac{1}{\alpha})$ condition for $\alpha \geq 14$. Since $\pi_1 X * \pi_1 Y$ is hyperbolic and $(X \vee Y)^*$ is compact, Theorem 3.2.14 then implies that $(X \vee Y)^*$ is hyperbolic.

Obtaining cocompact cubulability: Define a wallspace structure on $(X \vee Y)^*$ as follows. Firstly, by subdividing $X \vee Y$, we may assume that each W_i has an even number of hyperplanes cutting the generator of $\langle w_i \rangle$. The specialness hypothesis ensures that all hyperplanes of $X \vee Y$ and of each W_i are embedded and 2-sided, moreover, since each W_i has cyclic fundamental group and $H_i \cap Stab(\widetilde{U})$ is trivial or equal to $Stab(\widetilde{U})$ for each $\widetilde{U}$, all the hyperplanes of each W_i are contractible or have the homotopy type of a circle representing the generator of $\pi_1 W_i$. Hence, we can define a wallspace structure on each of the cones by defining a wall to be either a single hyperplane if the hyperplane does not cut the generator of the corresponding $\langle w_i \rangle$, or by defining a wall to be an equivalence class consisting of two antipodal hyperplanes cutting the generator. Concretely, if the generator of $\langle w_i \rangle$ is a cycle $\sigma \to X \vee Y$ of length $2n$, then letting $\sigma = e_1 \dots e_{2n}$, hyperplanes U and U' are in the same equivalence class if and only if U is dual to e_j and U' is dual to e_{j+n} (mod n) for some $j \in \{1, \dots, 2n\}$. These choices are exemplified in Figure 3.1.

We now check condition (1) of Theorem 3.2.27, which will allow us to extend the wallspace structure on the cones to a wallspace structure for $(X \vee Y)^*$ using the equivalence relation in Definition 3.2.22.

Choose the $\langle w_i \rangle$ so that $(X \vee Y)^* = \langle X \vee Y | \{W_i \to X \vee Y\}_{i=1}^r \rangle$ satisfies the $C'(\frac{1}{\alpha})$ condition for $\alpha \geq 14$. With the choice of walls described above, each cone is a wallspace satisfying condition (2) of Definition 3.2.20. The $C'(\frac{1}{14})$ condition is also sufficient to ensure that condition (3) is met. Indeed, the only way for a path with endpoints on the carrier of a hyperplane U to not be homotopic into the carrier of the hyperplane is if the path is homotopic into a power of the generator of W_i, and such a path would have to traverse at least 14 pieces. Moreover, condition (4) is met by rechoosing the cyclic subgroups so that the cubical presentation satisfies $C'(\frac{1}{16})$. To wit, since pairs U, U' of hyperplanes lying on the same wall W are antipodal and $(X \vee Y)^*$ satisfies $C'(\frac{1}{\alpha})$, the number of pieces in a path $\sigma \to C_i$ with endpoints on distinct hyperplanes of W is at least $\frac{\alpha}{2}$, so choosing $\alpha \geq 16$ ensures that such a path traverses at least 8 pieces. The choice of wallspace on each cone also ensures that condition (5) is met: any automorphism of $X \vee Y$ sends a wall not cutting a generator to a wall not cutting a generator, and a wall cutting a generator to a wall cutting a generator.

Thus, condition (1) of Theorem 3.2.27 is satisfied. Since each wall arises from a quasiconvex subgroup, Proposition 3.2.18 ensures cocompactness of the action on the dual cube complex. To ensure properness of the action, we check the rest of the conditions of Theorem 3.2.27.

Similar modifications to X^* will ensure that conditions (2) and (3) of Theorem 3.2.27 are met. For condition (2), by Proposition 3.2.25, it suffices to ensure that $\mathrm{d}_{\widetilde{W}_i}(g\widetilde{N}(U), \widetilde{N}(U)) > K$ for any translate $g\widetilde{N}(U) \neq \widetilde{N}(U) \subset \widetilde{W}_i$. Since each piece of W_i contains at least 1 edge, this can be guaranteed by rechoosing the w_i's so that X^* satisfies the $C'(\frac{1}{K'})$ condition, where $K' = \max\{K, 16\}$. Condition (3) also follows, because any two hyperplanes in the same

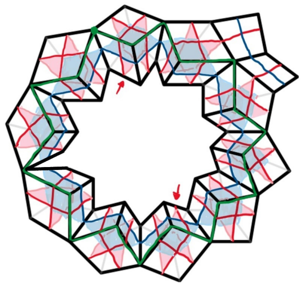

Figure 3.1: A potential cone-cell and its hyperplanes. The generator of its fundamental group is drawn in green, the hyperplanes that cross it are drawn in red, and the hyperplanes that do not cross it are drawn in blue. A pair of antipodal hyperplanes is indicated.

wall are at least 8-far, so there is a hyperplane V crossing the generator of $\langle w_i \rangle$ that is 2-far from both U and from U', and one can ensure that the antipodal hyperplane V' is also 2-far from both U and from U' by rechoosing X^* so that it satisfies the $C'(\frac{1}{K''})$ condition, where $K'' = \max\{2K, 16\}$.

Finally, the choice of walls implies that condition (4) also holds: since $\pi_1(W_i)$ is cyclic for each $i \in I$, every element $g \in Aut(W_i)$ has an axis, which is cut by a wall of X^* crossing the generator of $\pi_1(W_i)$.

Definition 3.3.4 (Height)**.** The *height* of $H \leq G$ is the maximal $n \in \mathbb{N}$ such that there exist distinct cosets $g_1 H, \ldots, g_n H$ for which $H^{g_1} \cap \ldots \cap H^{g_n}$ is infinite.

In [GMRS98] it was proven that:

Theorem 3.3.5. *Quasiconvex subgroups of hyperbolic groups have finite height.*

Definition 3.3.6. The *commensurator* $C_G(H)$ of a subgroup H of G is the subgroup $C_G(H) = \{g \in G : [H : H^g \cap H] < \infty \text{ and } [H^g : H^g \cap H] < \infty\}$.

Remark 3.3.7. If G is hyperbolic and the subgroup H is infinite and quasiconvex, then $[C_G(H) : H] < \infty$ by [KS96]. In particular, $C_G(H)$ is also a quasiconvex subgroup of G, and if G is torsion-free and H is free and non-abelian then so is $C_G(H)$.

The following result will be used in the proof of Lemma 3.3.9:

Lemma 3.3.8. *[Wis21, 8.6] Let G be hyperbolic and torsion-free and $H_1, \ldots, H_r$ be a collection of quasiconvex subgroups of G. Let $K_1, \ldots, K_s$ be representatives of the finitely many distinct conjugacy classes of subgroups consisting of intersections of collections of distinct conjugates of $H_1, \ldots, H_r$ in G that are maximal with respect to having infinite intersection. Then $\{C_G(K_1), \ldots, C_G(K_s)\}$ is a malnormal collection of subgroups of G.*

A version of the ensuing lemma can be found in [Kap99], but we include a proof for completeness:

Lemma 3.3.9. *Let G be a non-elementary, torsion-free hyperbolic group. For every $k \in \mathbb{N}$, G contains a malnormal collection $\{H_1, \ldots, H_k\}$ of infinite-index quasiconvex non-abelian free subgroups.*

Before proving Lemma 3.3.9, we make a few observations about malnormal subgroups of free groups. Recall that a subgroup $H < G$ is *isolated* if whenever $g^n \in H$ for $g \in G$, then $g \in H$; a subgroup $H < G$ is *malnormal on generators* if for any generating set $\{a_1, \ldots, a_n\}$ for H and $g \in G$ and $g \notin H$, then $a_i^g \notin H$ for any $i \in \{1, \ldots, n\}$.

Lemma 3.3.10. *[FMR02, Lemma 1] Let F be a free group and $H \subset F$ a 2-generator subgroup. Then H is malnormal if and only if H is isolated and malnormal on generators.*

Claim 3.3.11. *Let F be a finite rank non-abelian free group and let $\{h_1, \ldots, h_k\}$ be a finite collection of non-trivial elements of F. Then there is a subgroup $J < F$ for which $\{J, \langle h_i \rangle\}$ is a malnormal collection for each $i \in \{1, \ldots, k\}$.*

Proof. Assume that a basis is given for F, and, abusing notation, write also $h_1, \ldots, h_k$ to denote reduced words on the basis representing the finite set of elements $h_1, \ldots, h_k$. We may assume also that $h_i \neq h_j^m$ whenever $m \in \mathbb{Z} - \{0\}$ and $i \neq j$. Let $J = \langle a, b \rangle$ where $a = h_1\beta_1 \ldots h_k\beta_k, b = \beta_1' h_1 \ldots \beta_k' h_k$, and where each β_i and β_i' is a reduced word on the basis satisfying that, for each $i \in \{1, \ldots, k\}$,

1. $\beta_i \neq \beta_j^m$, $\beta_i' \neq (\beta_j')^m$, and $\beta_i \neq (\beta_j')^m$ whenever $i \neq j$ and $m \in \mathbb{Z} - \{0\}$,
2. no β_i, β_i' is a product of h_i's or their inverses,
3. for each β_i, its first letter is not equal to the last letter of h_i, and its last letter is not equal to the first letter of h_{i+1} (modulo k). Similarly, for each β_i', its last letter is not equal to the first letter of h_i, and its first letter is not equal to the last letter of h_{i+1} (modulo k). Moreover, the last letter of β_k is not equal to the inverse of the last letter of h_k, the first letter of β_1' is not equal to the inverse of the first letter of h_1, and the last letter of β_k is not equal to the first letter of β_1'.

As there is no cancellation, J is a rank-2 free group. By Lemma 3.3.10, to prove that J is malnormal it suffices to show that J is isolated and malnormal on generators. The choice of the β_i, β_i' implies in particular that a and b are not proper powers, and this implies in turn that J is isolated, since F is free. We now show that J is malnormal on generators. Consider a conjugate $a^g = (h_1\beta_1 \dots h_k\beta_k)^g$ where $g \notin J$ (the case of b^g is analogous). Since $g \notin J$, g cannot be written as a non-trivial product of powers of a and b and their inverses. If g cannot be written as a subword of a product of a's and b's and their inverses, then a^g cannot be an element of J. Indeed, this is because the choices made above preclude any cancellation, and because a^g is not a product of powers of a and b and their inverses; the choice of a and b implies that no cyclic permutation of a, b, their product or their proper powers lies in J, so no conjugate of a by a subword of a product of a's and b's and their inverses lies in J.

Finally, consider a conjugate h_i^g of h_i. If $\langle h_i^g \rangle \cap J$ is non-trivial, then since F is free, it follows that g must be a subword of some $j \in J$, and even in this case h_i^g can only be a (non-trivial) cyclic permutation of of a, b, their product or their proper powers, but no such cyclic permutation is an element of J, so J intersects all conjugates of h_i's trivially. ◻

Proof of Lemma 3.3.9. It suffices to show that G contains a malnormal quasiconvex free subgroup J of arbitrarily high rank, for then if $J = \langle a_1, \dots, a_k, b_1, \dots, b_k \rangle$, the collection $\{H_1, \dots, H_k\}$ where $H_i = \langle a_i, b_i \rangle$ is also malnormal and quasiconvex. For this, it suffices to show that G contains a malnormal quasiconvex free group J of *some* rank ≥ 2. Indeed, for any $n \in \mathbb{Z}$, J contains infinitely many subgroups of rank n, all of which are malnormal and quasiconvex.

By Theorem 3.2.1, G contains a free non-abelian quasiconvex subgroup J_0; we may assume further that J_0 has infinite-index in G. Let $\mathscr{J}_0$ be the lattice of infinite intersections of conjugates of J_0: this lattice is finite by Theorem 3.3.5. If $\mathscr{J}_0$ contains a non-abelian free group J_1 then we replace J_0 with J_1, and we can repeat this process a finite number of times until we either reach a maximal intersection of conjugates of some J_i that is itself free non-abelian or until all subgroups in the lattice $\mathscr{J}_i$ are cyclic. In the former case, the commensurator $C_G(J_i)$ is malnormal and quasiconvex by Lemma 3.3.8. In the latter case, by Claim 3.3.11, J_i contains a free non-abelian subgroup J that forms a malnormal collection with each of these cyclic subgroups, hence J is malnormal in G. ◻

Lemma 3.3.9 can be improved to control intersections with quasiconvex subgroups:

Corollary 3.3.12. *Let G be a non-elementary, torsion-free hyperbolic group and let $\{S_1, \dots, S_\ell\}$ be a collection of quasiconvex subgroups of G. Then the collection $\{H_1, \dots, H_k\}$ from the conclusion of Lemma 3.3.9 can be chosen so that $H_i \cap S_j$ is either trivial or equal to H_i for each $i \in \{1, \dots, k\}$ and each $j \in \{1, \dots, \ell\}$.*

Remark 3.3.13. In particular, if G is the fundamental group of a compact non-positively curved cube complex X, and $\mathscr{H}$ is the set of hyperplanes of X, then $\{H_1, \dots, H_k\}$ can be chosen so that $H_i \cap Stab(\widetilde{U})$ is either trivial or equal to H_i for each $U \in \mathscr{H}$ and each

$i \in \{1, \ldots, k\}$. Indeed, since X is compact, it has finitely many hyperplanes, and hence finitely many hyperplane stabilisers. Each hyperplane stabiliser is quasi-isometrically embedded and $\pi_1 X$ is hyperbolic, so each hyperplane stabiliser is quasiconvex.

Proof of Corollary 3.3.12. It suffices to prove the result for a single malnormal, quasiconvex non-abelian subgroup $J \leq G$. Indeed, as explained in the first paragraph of the proof of Lemma 3.3.9, the subgroups in the malnormal collection $\{H_1, \ldots, H_k\}$ are produced as subgroups of a single non-abelian free subgroup J, so if we ensure that $J \cap S_j$ is either trivial or equal to J for each $j \in \{1, \ldots, \ell\}$, then this will also be the case for $\{H_1, \ldots, H_k\}$.

We now proceed by induction on ℓ. Assume that $\ell = 1$ and consider the intersection $J \cap S_1$, where J is as provided in Lemma 3.3.9. If this intersection is trivial, then there is nothing to show, so suppose that $K_1 := J \cap S_1$ is non-trivial, so K_1 is either cyclic or free of rank ≥ 2. If K_1 is cyclic, say generated by k_1, then as J is free, by Claim 3.3.11 there exists a $J' \leq J$ such that $\{J', \langle k_1 \rangle\}$ is malnormal, so in particular $J' \cap S_1$ is trivial; if K_1 is free of rank ≥ 2, then since K_1 is quasiconvex, Lemma 3.3.9 implies that there exists a quasiconvex, non-abelian free subgroup $J'' \leq K_1$ that is malnormal in G, so that $J'' \cap S_1 = J''$.

Now assume that the result holds for $m = \ell - 1$ and let $\{S_1, \ldots, S_\ell\}$ be a collection of quasiconvex free non-abelian subgroups. Then by the induction hypothesis and Lemma 3.3.9, there is a quasiconvex non-abelian subgroup $J < G$ such that $J \cap S_i$ is trivial or equal to J for each $i \in \{1, \ldots, m\}$. As before, if $K_\ell := J \cap S_\ell = \{1\}$ then there is nothing to show, if K_ℓ is cyclic then by Claim 3.3.11 there exists a $J' \leq J$ such that $J' \cap S_\ell$ is trivial, and since $J' \leq J$ then it is still the case that $J' \cap S_i$ is trivial or equal to J' for each $i \in \{1, \ldots, m\}$. Finally, if K_ℓ is non-abelian then since it is quasiconvex, Lemma 3.3.9 produces a new J'' inside K_ℓ for which $J'' \cap S_\ell = J''$. Since $J'' < J$, then $J'' \cap S_i$ is trivial or equal to J'' for each $i \in \{1, \ldots, m\}$, and the result follows. ◻

3.4 Main theorem

Theorem 3.4.1. *Let Q be a finitely presented group and G be the fundamental group of a compact special cube complex X. If G is hyperbolic and non-elementary, then there is a short exact sequence*

$$1 \to N \to \Gamma \to Q \to 1 \tag{3.1}$$

where

1. *Γ is a hyperbolic, cocompactly cubulated group,*

2. *$N \cong G/K$ for some $K < G$,*

3. *$max\{cd(G), 2\} \geq cd(\Gamma) \geq cd(G) - 1$. In particular, Γ is torsion-free.*

In what follows, we prove all parts of Theorem 3.4.1 except for (3), which we explain in the next section.

Remark 3.4.2. Cocompact cubulability of Γ hinges on the specialness of X as this is the hypothesis that is used in Theorem 3.3.2. However, as noted in Remark 3.3.1, in reality all that is needed is for the hyperplanes of X to be embedded and 2-sided.

Proof. Choose a finite presentation $\langle a_1, \ldots, a_s | r_1, \ldots, r_k \rangle$ for Q, let B be a bouquet of s circles $a_1, \ldots, a_s$, and let X be a compact non-positively curved cube complex with $\pi_1 X = G = \langle x_1, \ldots, x_m \rangle$.

By Lemma 3.3.9 and Remark 3.3.13, there is a malnormal collection $\{H_\ell\} \cup \{H'_{ij}\} \cup \{H''_{ij}\}$ of quasiconvex free subgroups of rank ≥ 2 of $\pi_1(X \vee B)$, so we can apply Theorem 3.3.2 to X and B, where the y_i's are given by:

1. $y_\ell = r_\ell$ for each $1 \leq \ell \leq k$,

2. $y'_{ij} = a_i x_j a_i^{-1}$ for each $1 \leq i \leq s$, $1 \leq j \leq m$,

3. $y''_{ij} = a_i^{-1} x_j a_i$ for each $1 \leq i \leq s$, $1 \leq j \leq m$,

Hence, there are words $w_1, \ldots, w_\ell, w'_{1,1}, \ldots, w'_{ij}, w''_{1,1}, \ldots, w''_{ij} \in \pi_1 X$ for which the group

$$\Gamma = G * \pi_1 B / \langle\langle \{r_\ell w_\ell\}_{\ell=1}^k, \{a_i x_j a_i^{-1} w'_{ij}\}_{i=1,j=1}^{s,m}, \{a_i^{-1} x_j a_i w''_{ij}\}_{i=1,j=1}^{s,m} \rangle\rangle$$

is hyperbolic and acts properly and cocompactly on a CAT(0) cube complex.

There is a map $\Gamma \xrightarrow{\phi} Q$ that sends every generator of $\pi_1 X$ to 1. It is a well-defined homomorphism because the set of relations $\{r_\ell w_\ell = 1\}_\ell$ map exactly to the relations $\{r_\ell = 1\}_\ell$ in $\pi_1 B$.

The relations $\{a_i x_j a_i^{-1} w'_{ij} = 1\}_{i \in S, j \in M}$, and $\{a_i^{-1} x_j a_i w''_{ij} = 1\}_{i \in S, j \in M}$ ensure that $\langle x_1, \ldots, x_m \rangle$ is a normal subgroup of Γ so $N = ker\phi = \langle x_1, \ldots, x_m \rangle$ and $N \cong \pi_1 X / K$ for some subgroup $K < G$. ◻

3.5 Cohomological dimension bounds for Γ

We briefly recall some standard facts about the cohomological dimension of groups. We refer the reader to [Bro94] for more details and proofs.

A *resolution* for a module M over a ring R is a long exact sequence of $R-$modules

$$\cdots \to M_n \to M_{n-1} \to \cdots \to M_1 \to M_0 \to M \to 0$$

A resolution is *finite* if only finitely many of the M_i are non-zero. The *length* of a finite resolution is the maximum integer n such that M_n is non-zero. A resolution is *projective* if each M_i is a projective module.

Definition 3.5.1. The *cohomological dimension* $cd(G)$ of a group G is the length of the shortest projective resolution of $\mathbb{Z}$ as a trivial $\mathbb{Z}G$-module.

There is a natural topological analogue of cohomological dimension:

Definition 3.5.2. The *geometric dimension* $gd(G)$ of a group G is the least dimension of a classifying space for G.

Remark 3.5.3. The cellular chain complex of a classifying space for a group G yields a free (in particular, projective) resolution of $\mathbb{Z}$ over $\mathbb{Z}G$, the length of which is equal to the dimension of the classifying space. This implies immediately that $cd(G) \leq gd(G)$ for any group G. In particular, if G is free, then $cd(G) = 1$.

Remark 3.5.4. The universal covers of non-positively curved cube complexes are CAT(0) spaces, and hence are contractible, so every non-positively curved cube complex X is a classifying space for its fundamental group. Therefore, if $G = \pi_1 X$ for a compact non-positively curved cube complex X, then the cohomological dimension of G is bounded above by the dimension of X.

The next proposition is classical, and the reader may consult [Bro94, VIII.2.4] for a proof:

Proposition 3.5.5. *The following hold for any group G:*

1. *If $G' < G$ then $cd(G') \leq cd(G)$ and equality holds provided that $cd(G) < \infty$ and $[G : G'] < \infty$.*

2. *If $1 \longrightarrow G' \longrightarrow G \longrightarrow G'' \longrightarrow 1$ is exact, then $cd(G) \leq cd(G') + cd(G'')$.*

3. *If $G = G_1 * G_2$, then $cd(G) = max\{cd(G_1), cd(G_2)\}$.*

The following result is a consequence of Corollary 3.5.11, stated below.

Proposition 3.5.6. *Let G and Q be as in Theorem 3.4.1, B and Γ be as in its proof, and let $q : G * \pi_1 B \to \Gamma$ be the natural quotient. Then $Ker(q)$ is free.*

Proposition 3.5.7. Γ *can be chosen so that* $cd(\Gamma) \geq cd(G) - 1$.

Proof. There is a short exact sequence

$$1 \longrightarrow Ker(q) \longrightarrow G * \pi_1 B \longrightarrow \Gamma \longrightarrow 1.$$

Since $Ker(q)$ is a free group and $cd(G * \pi_1 B) = \max\{cd(G), 1\} = cd(G)$, then $cd(G) \leq cd(\Gamma) + cd(Ker(q)) = cd(\Gamma) + 1$. ∎

The torsion-freeness of Γ will follow from having a finite upper bound on its cohomological dimension:

Proposition 3.5.8. Γ *can be chosen so that* $max\{cd(G), 2\} \geq cd(\Gamma)$

Before proving Proposition 3.5.8, we state some auxiliary results.

Definition 3.5.9 (The Cohen-Lyndon property)**.** Let G be a group, $\{H_i\}_{i\in I}$ a family of subgroups and $N_i \lhd H_i$ for each $i \in I$. The triple $(G, \{H_i\}, \{N_i\})$ has the *Cohen-Lyndon property* if for each $i \in I$ there exists a left transversal T_i of $H_i\langle\langle \cup_{j\in I} N_j\rangle\rangle$ in G such that $\langle\langle \cup_{j\in I} N_j\rangle\rangle$ is the free product of the subgroups N_j^t for $t \in T_i$, so

$$\langle\langle \cup_{j\in I} N_j\rangle\rangle = *_{j\in I, t\in T_i} N_j^t.$$

The Cohen-Lyndon property was first defined and studied in [CL63], where it was proven to hold for triples $(F, C, \langle c\rangle)$ where F is free, C is a maximal cyclic subgroup of F, and $c \in C - \{1\}$. This was later generalised in [EH87] to the setting of free products of locally indicable groups. Most remarkably, it was recently proven in [Sun20] that triples $(G, \{H_i\}, \{N_i\})$ have the Cohen-Lyndon property when the H_i are "hyperbolically embedded" subgroups of G and the N_i avoids a finite set of "bad" elements depending only on the H_i. We will not define hyperbolically embedded subgroups here, and instead state only the particular case of the theorem that is required for our applications:

Theorem 3.5.10. *[Sun20] Let G be hyperbolic,* $\{H_i\}$ *be malnormal and quasiconvex subgroups of G, and* $N_i \lhd H_i$ *for each i. Then there exists a finite set of elements* $\{g_1, \dots, g_n\} \in \cup_i H_i - \{1\}$ *such that the triple* $(G, \{H_i\}, \{N_i\})$ *has the Cohen-Lyndon property provided that* $N_i \cap \{g_1, \dots, g_n\} = \emptyset$ *for all i.*

To simplify notation, let $\{H_\ell\} \cup \{H'_{ij}\} \cup \{H''_{ij}\} := \mathbf{H}$ and write $H_\iota \in \mathbf{H}$. Let $S \subset \mathbf{H}$ be a finite set. It is clear from the constructions in Theorems 3.4.1 and 3.3.2 (say, by applying Claim 3.3.11 to S before producing the cyclic subgroups in the proof of Theorem 3.3.2) that the elements $\{r_\ell w_\ell\}_{\ell=1}^{k}, \{a_i x_j a_i^{-1} w'_{ij}\}_{i=1,j=1}^{s,m}, \{a_i^{-1} x_j a_i w''_{ij}\}_{i=1,j=1}^{s,m}$ can be chosen so that for each $c_\iota \in \{r_\ell w_\ell\}_{\ell=1}^k \cup \{a_i x_j a_i^{-1} w'_{ij}\}_{i=1,j=1}^{s,m} \cup \{a_i^{-1} x_j a_i w''_{ij}\}_{i=1,j=1}^{s,m}$, where $c_\iota \in H_\iota$, the intersection $\langle\langle c_\iota \rangle\rangle_{H_\iota} \cap S$ is empty, and hence:

Corollary 3.5.11. *Γ can be chosen so that the triple $(G * \pi_1 B, \{H_\iota\}, \{\langle\langle c_\iota\rangle\rangle_{H_\iota}\})$ has the Cohen-Lyndon property.*

The following is proven in [PS21]:

Proposition 3.5.12. *If $(G, \{H_i\}, \{N_i\})$ has the Cohen-Lyndon property, then*

$$cd(G/\langle\langle \cup_{i\in I} N_i \rangle\rangle) \leq \max\{cd(G), \sup\{cd(H_i)+1\}, \sup\{cd(H_i/N_i\}\}.$$

Definition 3.5.13. A *graphical* presentation is a 1-dimensional cubical presentation. Namely, a cubical presentation $X^* = \langle X | \{Y_i\}\rangle$ where X is a graph and $Y_i \to X$ are graph immersions.

In the particular setting of graphical presentations, it is well-known that the coned-off space X^* is aspherical. Concretely, the following result is a special case of Theorem 4.4.7 from Chapter 4 of this thesis. See also Example 4.5.2 in the same chapter.

Theorem 3.5.14. *Let $X^* = \langle X | \{Y_i\}\rangle$ be a $C(9)$-graphical small cancellation presentation. Then X^* is aspherical.*

Theorem 3.5.14 is due to Gromov and Ollivier [Gro93, Oll06] in the graphical $C'(\frac{1}{6})$ setting, and to Gruber [Gru15] in the graphical $C(6)$ setting. A proof of the latter (more general) case is given in Gruber's aforementioned paper, though we caution the reader that the language utilised there differs in many ways from that in this text.

Remark 3.5.15. As in the statement of Corollary 3.5.11, we simplify the notation so that $\{H_\ell\} \cup \{H'_{ij}\} \cup \{H''_{ij}\} := \mathbf{H}$, $c_\iota \in H_\iota$ equals the corresponding $r_\ell w_\ell, a_i x_j a_i^{-1} w'_{ij}$, or $a_i^{-1} x_j a_i w''_{ij}$, and $C_\iota \to X \vee B$, $W_\iota \to X \vee B$ are the corresponding local isometries defined in the proof of Theorem 3.3.2 having $\pi_1(C_\iota) = H_\iota, \pi_1(W_\iota) = \langle c_\iota \rangle$ for each ι.

Proof of Proposition 3.5.8. By the Cohen-Lyndon property for $(G * \pi_1 B, \{H_\iota\}, \{\langle\langle c_\iota\rangle\rangle_{H_\iota}\})$, then $cd(G * \pi_1 B / \langle\langle \cup_\iota c_\iota \rangle\rangle) \leq \max\{cd(G) * \pi_1 B, \sup\{cd(H_\iota)+1\}, \sup\{cd(H_\iota/\langle\langle c_\iota\rangle\rangle_{H_\iota}\})\}$, and since each H_ι is free, then $cd(H_\iota) = 1$ for all $H_\iota \in \mathbf{H}$. We claim that each of the quotients $H_\iota/\langle\langle c_\iota\rangle\rangle_{H_\iota}$ has a $C(6)$-graphical small cancellation presentation, and so $cd(H_\iota/\langle\langle c_\iota\rangle\rangle_{H_\iota}) \leq 2$ for all $H_\iota \in \mathbf{H}$.

To see this, consider the cubical presentation $\langle X \vee B | \{W_\iota \to X \vee B\}_\iota\rangle$ constructed in the proof of Theorem 3.4.1. As explained in the proof of Theorem 3.3.2, each H_ι is carried by a local isometry $C_\iota \to X \vee B$. Since each C_ι is itself a non-positively curved cube complex and $\pi_1 C_\iota \cong H_\iota$ is free, then each C_ι is homotopy equivalent to a graph $\bar{C}_\iota$ in $C_\iota^{(1)}$. Similarly, each W_ι is homotopy equivalent to a cycle $\bar{W}_\iota$ in W_ι and we can further assume that each $\bar{W}_\iota$ lies in $\bar{C}_\iota$. The intersections between pieces of each $\bar{W}_\iota$ are contained in the corresponding intersections between pieces of W_ι, and the length of each $\bar{W}_\iota$ is bounded below by the diameter of the corresponding W_ι, so each $\bar{W}_\iota$ has at least as many pieces as the corresponding W_ι. Since the W_ι are chosen to satisfy at least $C'(\frac{1}{16})$, and in particular $C'(\frac{1}{16}) \Rightarrow C(9)$, then each $\langle \bar{C}_\iota | \bar{W}_\iota \to \bar{C}_\iota \rangle$ satisfies the $C(6)$ condition, and the proof is complete. $\square$

This finishes the proof of Theorem 3.4.1.

Chapter 4

Asphericity of cubical presentations: the 2-dimensional case

4.1 Introduction

The aim of this chapter is to explore the asphericity of certain cubical presentations $X^* = \langle X | \{Y_i \to X\} \rangle$ where X is a compact, non-positively curved cube complex and $Y_i \to X$ are local isometries of compact non-positively curved cube complexes. A finite cubical presentation is a natural generalisation of a finite group presentation $\mathscr{P} = \langle s_1, \dots s_k | r_i, \dots, r_\ell \rangle$.

It is a well-known result of Lyndon [Lyn66] that classical $C(6)$ presentations are aspherical if no relators are proper powers, and thus that the groups admitting such presentations have cohomological dimension at most equal to 2. This was generalised to the setting of graphical $C'(\frac{1}{6})$ small-cancellation presentations by Gromov [Gro03] and Ollivier [Oll06], and graphical $C(6)$ presentations by Gruber [Gru15], and to the setting of small-cancellation for rotation families of groups by Coulon [Cou11]. In what follows, we take the first steps towards a generalisation of these results to the setting of cubical small-cancellation theory, and show:

Theorem 4.1.1. *Let $X^* = \langle X | \{Y_i\}_{i \in I} \rangle$ be a minimal cubical presentation that satisfies the $C(9)$ condition. Let $\pi_1 X^* = G$. If $dim(X) \leq 2$ and each Y_i is homotopy equivalent to a graph, then X^* is a $K(G,1)$, so G is torsion-free and $gd(G) \leq 2$.*

The above hypothesised minimality is the cubical analogue of requiring, in the classical setting, that none of the relators are proper powers, and is necessary to avoid torsion. See Definition 4.3.10 below.

Replacing minimality with the weaker hypothesis that $X^* = \langle X | \{Y_i\}_{i \in I} \rangle$ be symmetric in the sense of Definition 4.3.8, we obtain instead classifying spaces for proper actions for

$\pi_1 X^*$. These spaces arise in connection to the Baum-Connes conjecture, and thus this is part of the motivation for finding good models for $\underline{E}G$.

Theorem 4.1.2. *Let $X^* = \langle X | \{Y_i\}_{i=1}^k \rangle$ be a symmetric cubical presentation that satisfies the $C(9)$ condition. Let $\pi_1 X^* = G$. If $dim(X) \leq 2$ and each Y_i is homotopy equivalent to a graph, then there is a quotient $\bar{X}^*$ of $\widetilde{X^*}$ that is an $\underline{E}G$, so $cd_{\mathbb{Q}}(G) \leq dim\bar{X}^* \leq 2$. If, in addition, X has a finite regular cover where each $Y_i \to X$ lifts to an embedding, then $vcd(G) \leq 2$.*

Since some of the cubically presented groups satisfying the hypotheses above have torsion, Theorem 4.1.2 is the best statement possible in this generality.

Theorems 4.1.1 and 4.1.2 will follow from Theorem 4.4.5, which applies to cubical presentations of arbitrary dimension and is of independent interest.

Assuming minimality or symmetry, and under the hypotheses on the dimension of X and the Y_i's stated above, Theorems 4.1.1 and 4.1.2 answer a question posed by Wise in [Wis21, 4.5]. The general case of the theorem, where the dimensions of X and Y_i are arbitrary, and which we use to derive near-sharp bounds on the (virtual) cohomological dimension of X^*, will be treated in forthcoming work [Are].

Classical and graphical small-cancellation groups are hyperbolic as soon as they satisfy either the $C'(\frac{1}{6})$ or the $C(7)$ conditions, and the groups produced in [Cou11] are also hyperbolic. Thus, our main theorem is particularly interesting in that hyperbolicity is neither assumed nor deduced: while one can conclude hyperbolicity for a sufficiently good cubical small cancellation presentation if one assumes hyperbolicity of $\pi_1 X$, it is nevertheless the case that non-hyperbolic, and even non-relatively-hyperbolic groups can satisfy strong cubical small-cancellation conditions. Thus, Theorems 4.1.1 and 4.1.2 apply to a large and varied family of quotients of cubulated groups.

Despite their indisputable value as a source of examples of interesting group-theoretic behaviour, the classical and graphical small-cancellation theories are limited in their applicability in that they can only serve to understand "nice" quotients of free groups; cubical small-cancellation theory transfers some of the complexity arising from the relators in a group presentation to the "cubical generator" in the cubical presentation – that is, the cube complex X that is being quotiented – making it possible to exploit the tools of cubical geometry to prove more general theorems.

This approach has already proven fruitful in many instances: particularly, it plays a key role in Agol's celebrated proofs of the Virtual Haken and Virtual Fibred Conjectures [Ago13, Ago08], which build on work of Wise [Wis21] and his collaborators [BW12, HW12, HW15]; it is used by Arzhantseva and Hagen in [AH22] to show that many groups that arise as quotients of cubulated groups are acylindrically hyperbolic; it is used by Jankiewicz and Wise in [JW22] to construct fundamental groups of compact nonpositively curved cube complexes that do not virtually split, and it is used by the author in [Are23] – Chapter 3 of this thesis – to produce a version of the Rips construction that provides cocompactly cubulated hyperbolic groups of arbitrarily large finite cohomological dimension, many of which algebraically fibre.

4.1.1 Structure of the chapter

In Section 4.2, as a sort of base-case for our main theorem, we reprove without resorting to CAT(0) geometry, the well-known result that CAT(0) cube complexes are contractible. In Section 4.3 we present the necessary background on cubical small cancellation theory, diagrams in cell complexes, and asphericity. In Section 4.4 we give the proof of our main technical result, Theorem 4.4.5, and use it to deduce Theorems 4.1.1 and 4.1.2. Finally, in Section 4.5 we outline some examples to which our theorem applies.

4.2 On the asphericity of non-positively curved cube complexes

A *CAT(0) cube complex* is a simply connected, non-positively curved cube complex. CAT(0) cube complexes are homotopy equivalent to CAT(0) spaces (in the metric sense of Cartan–Alexandrov–Toponogov) by [BH99, II.4] and [Lea13], so they are therefore contractible, which in turn implies that non-positively curved cube complexes are aspherical spaces. While the CAT(0) metric is sometimes useful, in practice one can prove most results about non-positively curved cube complexes without resorting to the CAT(0) metric. Because of this, and since non-positively curved cube complexes correspond to cubical presentations having no relators, and as such, satisfy the $C(9)$ condition in the sense of Definition 4.3.11, it is therefore worthwhile to investigate whether the asphericity of non-positively curved cube complexes (or equivalently, the contractibility of their universal covers) can be proved via purely cubical/combinatorial means. This is the purpose of the present section. We stress that alternative proofs of the contractibility of CAT(0) cube complexes, while not present in the literature as far as the author is aware, are well-known to experts.

Theorem 4.2.1. *Non-positively curved cube complexes are aspherical.*

Definition 4.2.2. Let X be a CW complex. Suppose that c, c' are two cubes of X such that $c' \subset c$, c is a maximal face of X, and no other maximal face of X contains c'. Then c' is a *free face* of X. A *collapse* is the removal of all open cells c'' such that $c' \subseteq c \subseteq c''$, where c' is a free face. A cell complex is *collapsible* if there is a sequence of collapses leading to a point.

All collapsible cell-complexes are contractible [Whi39], but the converse is not necessarily true, even in the setting of cube complexes: Bing's *house with two rooms* [Hat02, pg. 4] is an example of a contractible cube complex that isn't collapsible (and not non-positively curved).

In what follows we will show that finite CAT(0) cube complexes are collapsible. From this, we will be able to deduce that all CAT(0) cube complexes are contractible via a nested subcomplex argument. Note that collapsibility of finite CAT(0) cube complexes follows from more general work of Adiprasito and Benedetti [AB20]. However, their proof heavily relies on CAT(0) geometry.

Lurking behind our proof is a form of convexity associated to the ℓ^1 metric on the 1-skeleton of a CAT(0) cube complex X. We define it below in a manner that elucidates its combinatorial nature.

Definition 4.2.3 (Combinatorial convexity)**.** A connected subcomplex Y of a CAT(0) cube complex X is *convex* if, for each n-cube c with $n \geq 2$ in X, whenever a corner of c lies in Y, then the cube c lies in Y. The *convex hull* $Hull(Z)$ of a subcomplex $Z \subset X$ is the smallest convex subcomplex of X containing Z.

We remind the reader of the following observation, stated in Chapter 2 as Lemma 2.1.21.

Remark 4.2.4. A convex subcomplex Y of a CAT(0) cube complex X is itself a CAT(0) cube complex.

The result below can be readily derived from [Hag08, 2.28].

Theorem 4.2.5. *The convex hull of a finite subcomplex of a CAT(0) cube complex is finite.*

We also make use of the following proposition, which can be found in [Sag95].

Proposition 4.2.6. *A hyperplane in a CAT(0) cube complex is a CAT(0) cube complex, and separates it into exactly two connected components.*

Remark 4.2.7. As we are aiming to reprove a basic fact about CAT(0) cube complexes, some care must be taken to ensure that our arguments are not circular. That is, that none of the result we utilise hinge on the contractibility of CAT(0) cube complexes to begin with. In the case of Theorem 4.2.5 and Proposition 4.2.6 above, the proofs only utilise combinatorial convexity, disc diagrams, and the fact that CAT(0) cube complexes are simply connected.

Proof of Theorem 4.2.1 via collapsibility of finite CAT(0) cube complexes.

Finite CAT(0) cube complexes: We first show, by induction on dimension, that a finite CAT(0) cube complex is either a 0-dimensional cube or has a free face. A connected, 0-dimensional cube complex is necessarily a single vertex; a finite, 1-dimensional CAT(0) cube complex is a finite tree, and therefore has a free face (a leaf). Assume that the result holds for all CAT(0) cube complexes of dimension $< k$. Let Y be a k-dimensional CAT(0) cube complex, and let H be a hyperplane in Y. Then H has dimension less than k and is also a CAT(0) cube complex, so H has a free face f_H. Note that f_H lies in a face f_Y of Y, and f_Y cannot be glued to any other cubes of Y, as this would imply that f_H extends to a midcube in any such cube. Hence, f_Y is a free face of Y.

Now we show, by induction on dimension, that for any CAT(0) cube complex Y and for any free face f of Y, there is a sequence of collapses starting with f that results in a subcomplex Y' that is again a CAT(0) cube complex. Note that this claim implies that Y is collapsible, since Y' must have a free face and then the process can be repeated, and terminates in finitely many steps because Y was finite to begin with. The base case is $dim(Y) = 1$, which is immediate as Y is then a tree, so every connected subcomplex of Y is a tree and thus CAT(0). Now assume that the claim holds for all CAT(0) cube complexes of dimension $< k$ and let $dim(Y) = k$.

Let f be a free face of Y and F be the maximal face containing f, let H be a hyperplane of Y having a free face that is a midcube m of f and note that m is a free face of H. Since $dim(H) < dim(Y)$ and H is a CAT(0) cube complex, the induction hypothesis implies that there is a sequence of collapses starting with m that results in a CAT(0) cube complex H', moreover, after repeating this process finitely many times, we may assume that the sequence of collapses terminates at a point p that is a vertex of H. We claim that this sequence of collapses extends to a sequence of collapses for the carrier $N(H)$ of H in Y which terminates with an edge e of Y.

Indeed, each collapse of H results in a subcomplex that has a free face m', so the corresponding collapse in Y results in a subcomplex that has a free face (the face having m' as a midcube). Thus, the sequence of collapses terminating on p extends to a sequence of collapses terminating in a subcomplex of Y in which $N(H)$ collapses to $N(p)$, which is an edge e. Now, since Y is CAT(0), then the interior $N(H)^o$ of $N(H)$ separates Y by Proposition 4.2.6, and each connected component Z_1, Z_2 of $Y - N(H)^o$ is CAT(0): each Z_i is non-positively curved since for each vertex v of $Y - N(H)^o$, the link of v in $Y - N(H)^o$ is the restriction $link(v)|_{Y-N(H)^o}$, and is either identical to $link(v)$, or is obtained from $link(v)$ by removing the flag subcomplex corresponding to $link(v)|_{N(H)^o}$ from it; each Z_i is simply connected because $Y = Z_1 \cup N(H) \cup Z_2$ is simply connected, and is homotopy equivalent to the wedge $Z_1 \vee Z_2$.

Finally, note that Y' is obtained from $Y - N(H)^o$ by reattaching e to it along $(Y - N(H)^o) \cap e$, so it now suffices to check that $link(v)$ is flag for each vertex v of $(Y - N(H)^o) \cap e$ in Y', that is, for the endpoints of e. But the links of these vertices are flag in $Y - N(H)^o$, and attaching

e only adds a disconnected vertex to each of these links, so the flag condition continues to hold, and Y' is simply connected because both connected components of $Y - N(H)^o$ are, thus Y' is homotopy equivalent to the wedge of 2 simply connected cube complexes.

Infinite CAT(0) cube complexes: Since S^n is compact, any map $S^n \longrightarrow Y$ for $n \geq 1$ has compact image $Im_Y(S^n)$, so $Im_Y(S^n)$ is contained in a finite subcomplex $Z \subset Y$. Thus, $Im_Y(S^n) \subset Hull(Z)$, which is finite by Theorem 4.2.5 and Remark 4.2.4, and contractible by the argument above. Hence, $S^n \longrightarrow Y$ is nullhomotopic, and $\pi_n(Y) = 0$ for all $n \geq 1$. ◻

4.3 Background

4.3.1 Cubical small-cancellation theory

Unless otherwise noted, all definitions and results concerning cubical small-cancellation theory recounted in Subsections 4.3.1 and 4.3.2 originate in [Wis21].

We begin this section by recalling our main objects of study:

Definition 4.3.1 (Cubical presentation)**.** A *cubical presentation* $\langle X | \{Y_i\} \rangle$ consists of a connected non-positively curved cube complex X together with a collection of local isometries of connected non-positively curved cube complexes $Y_i \xrightarrow{\varphi_i} X$. Local isometries of non-positively curved cube complexes are π_1-injective, so it makes sense to define the *fundamental group of a cubical presentation* as $\pi_1 X / \langle\langle \{\pi_1 Y_i\} \rangle\rangle$. By the Seifert-Van Kampen Theorem, this group is isomorphic to the fundamental group of the space X^* obtained by coning off each Y_i in X. Hereinafter, all auxiliary definitions and results about cubical presentations are in practice statements about their associated coned-off spaces.

Remark 4.3.2. A group presentation $\langle a_1, \dots, a_s | r_1, \dots, r_m \rangle$ can be interpreted cubically by letting X be a bouquet of s circles and letting each Y_i map to the path determined by r_i. On the other extreme, for every non-positively curved cube complex X there is a “free" cubical presentation $X^* = \langle X | \ \rangle$ with fundamental group $\pi_1 X = \pi_1 X^*$.

Definition 4.3.3 (Elevations)**.** Let $Y \to X$ be a map where Y is connected, and let $\hat{X} \to X$ be a covering map. An *elevation* $\hat{Y} \to \hat{X}$ is a map satisfying

1. $\hat{Y}$ is connected,
2. $\hat{Y} \to Y$ is a covering map, the composition $\hat{Y} \to Y \to X$ equals $\hat{Y} \to \hat{X} \to X$, and
3. assuming all maps involved are basepoint preserving, $\pi_1 \hat{Y}$ equals the preimage of $\pi_1 \hat{X}$ in $\pi_1 Y$.

Definition 4.3.4 (Pieces)**.** Let $\langle X|\{Y_i\}\rangle$ be a cubical presentation. An *abstract contiguous cone-piece* of Y_j in Y_i is an intersection $\widetilde{Y}_j \cap \widetilde{Y}_i$ where $\widetilde{Y}_j, \widetilde{Y}_i$ are fixed elevations to $\widetilde{X}$ and either $i \neq j$ or where $i = j$ but $\widetilde{Y}_j \neq \widetilde{Y}_i$. A *cone-piece* of Y_j in Y_i is a combinatorial path $p \to P$ in an abstract contiguous cone-piece of Y_j in Y_i. An *abstract contiguous wall-piece* of Y_i is an intersection $N(H) \cap \widetilde{Y}_i$ where $\widetilde{Y}_i$ is a fixed elevation and $N(H)$ is the carrier of a hyperplane H that is disjoint from $\widetilde{Y}_i$. To avoid having to deal with empty pieces, we shall assume that H is dual to an edge with an endpoint on $\widetilde{Y}_i$. A *wall-piece* of Y_i is a combinatorial path $p \to P$ in an abstract contiguous wall-piece of Y_i.

A *piece* is either a cone-piece or a wall-piece.

Remark 4.3.5. In Definition 3.2.11, two elevations of a cone Y are considered identical if they differ by an element of $Stab_{\pi_1 X}(\widetilde{Y})$. This is in keeping with the conventions of classical small cancellation theory, where overlaps between a relator and any of its cyclic permutations are not regarded as pieces.

Definition 4.3.6. Let $Y \to X$ be a local isometry. $Aut_X(Y)$ is the group of combinatorial automorphisms $\psi : Y \to Y$ such that the diagram below is commutative:

$$\begin{tikzcd} Y \arrow[r, "\psi"] \arrow[dr] & Y \arrow[d] \\ & X \end{tikzcd}$$

If Y is simply connected, then $Aut_X(Y)$ is equal to $Stab_{\pi_1 X}(Y)$. In general, $Aut_X(Y) \cong (N_{Aut_X(\widetilde{Y})}\pi_1 Y)/\pi_1 Y$, where $N_G(H)$ is the normaliser of H in G.

See [AH22] for a detailed discussion.

Convention 4.3.7. Let $X^* = \langle X|\{Y_i\}_I\rangle$. For convenience we assume, as is standard to assume in this framework [Wis21, 3.3], that $\pi_1 Y_i$ is normal in $Stab_{\pi_1 X}(\widetilde{Y}_i)$ for each $i \in I$. Throughout, we will assume also that X is finite dimensional and locally finite, and that each Y_i is compact and connected.

For the purposes of this paper, we will need to impose further restrictions on a cubical presentation; these are described below.

Definition 4.3.8. A local-isometry $Y \to X$ with Y connected and superconvex is symmetric if for each component K of the fibre product $Y \otimes_X Y$, either K maps isomorphically to each copy of Y, or $[\pi_1 Y : \pi_1 K] = \infty$. A cubical presentation $\langle X|\{Y_i\}\rangle$ is *symmetric* if each $Y_i \to X$ is symmetric.

In [Wis21, 8.12], the following is observed:

Lemma 4.3.9. *Let $X^* = \langle X | \{Y_i\} \rangle$ be a cubical presentation. Then X^* is symmetric if and only if $C_{\pi_1 X}(\pi_1 Y_i) = Stab_{\pi_1 X}(\widetilde{Y}_i)$, $[Stab_{\pi_1 X}(\widetilde{Y}_i) : \pi_1 Y_i] < \infty$, and $\pi_1 Y_i \lhd Stab_{\pi_1 X}(\widetilde{Y}_i)$ for each $Y_i \to X$.*

See Definition 3.2.4 for the notion of superconvexity, and Definition 3.3.6 for that of commensurator subgroup. Neither of these concepts is used directly in what remains of this thesis.

To the best of our knowledge, the following notion was introduced in [AH22], where it is used in the course of proving acylindrical hyperbolicity for certain cubical presentations satisfying strong cubical small-cancellation conditions.

Definition 4.3.10. A cubical presentation $\langle X | \{Y_i\}_I \rangle$ is *minimal* if the following holds for each $i \in I$: let $\widetilde{Y}_i \to Y_i$ be the universal cover, and let $\widetilde{Y}_i \to Y_i \to X$ be an elevation of $Y_i \to X$. Fix a basepoint x_0 in X. Then $Stab_{\pi_1(X,x_0)}(\widetilde{Y}_i) = \pi_1(Y_i, y_{i_0})$.

Minimality generalises prohibiting relators that are proper powers in the classical small-cancellation case. In our setting, minimality is used to avoid the "obvious" torsion that could be created in a quotient $\pi_1 X / \langle\langle \{\pi_1 Y_i'\} \rangle\rangle$ if the $Y_i' \to X$ are themselves non-trivial finite covers $Y_i' \to Y_i \to X$. For instance, if Y is a finite degree covering of X, the cubical presentation $X^* = \langle X | Y \rangle$ will have finite fundamental group, even if it satisfies the $C(9)$ condition defined below (or any other small cancellation condition).

Definition 4.3.11. A cubical presentation $X^* = \langle X | \{Y_i\} \rangle$ satisfies the $C(n)$ *small-cancellation condition* if no essential closed path $\sigma \to Y_i$ is the concatenation of fewer than n pieces.

An analogue of the $C'(\frac{1}{n})$ condition can also be defined in this setting, but we won't need it in the present work. We remark only that, as in the classical case, the cubical $C'(\frac{1}{n})$ condition implies the cubical $C(n+1)$ condition, but the implication is very much not reversible: regardless of the choices of n and n', the cubical $C(n)$ condition does not necessarily imply the cubical $C'(\frac{1}{n'})$ condition.

4.3.2 Diagrams in cube complexes and cubical presentations

In this section we introduce disc and spherical diagrams, and the general related jargon – the analysis of the possible diagrammatic behaviours in $C(9)$ cubical presentations, and the use of the tools already available in this context, will be the main ingredients utilised in Section 4.4.

Definition 4.3.12. A map $f : X \longrightarrow Y$ between 2-complexes is *combinatorial* if it maps cells to cells of the same dimension. A complex is *combinatorial* if all attaching maps are combinatorial, possibly after subdividing the cells.

Definition 4.3.13 (Disc diagram)**.** A *disc diagram D* is a compact contractible combinatorial 2-complex, together with an embedding $D \hookrightarrow S^2$. The *boundary path* ∂D is the attaching map of the 2-cell at infinity. A *disc diagram in a complex X* is a combinatorial map $D \to X$. A *square disc diagram* is a disc diagram that is also a cube complex (though **not necessarily** non-positively curved!). Note that any disc diagram in a cube complex is a square disc diagram.

We introduce some of the phenomena that may arise in square-disc diagrams; unlike in the general setting of 2-complexes, problematic behaviour of diagrams can be described and classified precisely in the cubical setting.

Definition 4.3.14 (Square-disc behaviours)**.** A *dual curve* in a square disc diagram is a path that is a concatenation of midcubes. The 1-cells crossed by a dual curve are *dual* to it. A *bigon* is a pair of dual curves that cross at their first and last midcubes. A *monogon* is a single dual curve that crosses itself at its first and last midcubes. A *nonogon* is a single dual curve of length ≥ 1 that starts and ends on the same dual 1-cell, thus it corresponds to an immersed cycle of midcubes. A *spur* is a vertex of degree 1 on ∂D.

A *corner* in a diagram D is a vertex v that is an endpoint of consecutive edges a, b on ∂D lying in a square s. A *cornsquare* – short for "generalised corner of a square"– consists of a square c and dual curves p, q emanating from consecutive edges a, b of c that terminate on consecutive edges a', b' of ∂D. The *outerpath* of the cornsquare is the path $a'b'$ on ∂D. A *cancellable pair* in D is a pair of 2-cells R_1, R_2 meeting along a path e such that the following diagram commutes:

$$\begin{tikzcd} & e \arrow[dl] \arrow[dr] & \\ \partial R_1 \arrow[rr] \arrow[dr] & & \partial R_2 \arrow[dl] \\ & X & \end{tikzcd}$$

A cancellable pair leads to a smaller area disc diagram via the following procedure: cut out $e \cup Int(R_1) \cup Int(R_2)$ and then glue together the paths $\partial R_1 - e$ and $\partial R_2 - e$ to obtain a diagram D' with $\mathsf{Area}(D') = \mathsf{Area}(D) - 2$ and $\partial D' = \partial D$.

By performing the procedure just described to cancellable pairs, diagrams in non-positively curved cube complexes can often be simplified to avoid certain pathologies.

Lemma 4.3.15. *[Wis21, 2.3+2.4] Let $D \to X$ be a disc diagram in a non-positively curved cube complex. If D contains a bigon or a nonogon, then there is a new diagram D' having* *the same boundary path as D, so $\partial D' \to X$ equals $\partial D \to X$, and such that $Area(D') \leq Area(D) - 2$. Moreover, no disc diagram in X contains a monogon, and if D has minimal area among all diagrams with boundary path ∂D, then D cannot contain a bigon nor a nonogon.*

Some additional terminology is necessary when describing diagrams in cubical presentations:

Definition 4.3.16. Recall that the coned-off space X^* introduced in Definition 4.3.1 consists of X with a cone on Y_i attached to X for each i. The vertices of the cones on Y_i's are the *cone-vertices* of X^*. The cellular structure of X^* consists of all the original cubes of X, and the pyramids over cubes in Y_i with a cone-vertex for the apex. Let $D \to X^*$ be a disc diagram in a cubical presentation. The vertices in D which are mapped to the cone-vertices of X^* are the *cone-vertices* of D. Triangles in D are naturally grouped into cyclic families meeting around a cone-vertex. Each such family forms a subspace of D that is a cone on its bounding cycle. A *cone-cell* of D is a cone that arises in this way. To simplify the theory, when analysing diagrams in a cubical presentation we "forget" the subdivided cell-structure of a cone-cell C and regard it simply as a single 2-cell.

A situation that may occur with diagrams in coned-off spaces is that two cone-cells might come from the same coned-off relation Y. When this happens, it is often possible to fuse this two cone-cells together into a single cone-cell, as explained below. This simplification does not arise in the "purely cubical" setting, but will be useful for analysing diagrams in cubical presentations.

Definition 4.3.17. A pair of cone cells C, C' in D is *combinable* if they map to the same cone Y of X^* and ∂C and $\partial C'$ both pass through a vertex v of D, and map to closed paths at the same point of Y when regarding v as their basepoint.

As the name suggests, such a pair can be combined to simplify the diagram by replacing the pair with a single cone-cell mapping to Y and whose boundary is the concatenation $\partial C \partial C'$.

Definition 4.3.18. A *disc diagram $D \to X^*$* is *reduced* if the following conditions hold:

1. There is no bigon in a square subdiagram of D.

2. There is no cornsquare whose outerpath lies on a cone-cell of D.

3. There does not exist a cancellable pair of squares.

4. There is no square s in D with an edge on a cone-cell C mapping to the cone Y, such that $(C \cup s) \to X$ factors as $(C \cup s) \to Y \to X$.

5. For each internal cone-cell C of D mapping to a cone Y, the path ∂C is essential in Y.

6. There does not exist a pair of combinable cone-cells in D.

A closely-related notion to that of disc diagram reducibility is that of complexity:

Definition 4.3.19. The *complexity* $Comp(D)$ of a disc diagram $D \to X^*$ is the ordered pair (#Cone-cells, #Squares). We order the pairs lexicographically: namely $(\#C, \#S) < (\#C', \#S')$ whenever $\#C < \#C'$ or $\#C = \#C'$ and $\#S < \#S'$. A disc diagram $D \to X^*$ has *minimal complexity* if no disc diagram $D' \to X^*$ having $\partial D = \partial D'$ has $Comp(D') < Comp(D)$.

Since the objective of this work is to understand the second homotopy group of a cubical presentation, it will come as no surprise that in addition to disc diagrams, their spherical analogues will also play a key role in this endeavour.

Definition 4.3.20 (Spherical diagram)**.** A *spherical diagram* Σ is a compact simply-connected combinatorial 2-complex, together with a homeomorphism $\Sigma \hookrightarrow S^2$. A *spherical diagram in a complex X* is a combinatorial map $\Sigma \to X$. We define *reduced* spherical diagrams and *complexity* in the same way as in Definitions 4.3.18 and 4.3.19, replacing all instances of "disc diagram" by "spherical diagram". A spherical diagram has *minimal complexity* if no spherical diagram $\Sigma' \to X^*$ homotopic to Σ has $Comp(\Sigma') < Comp(\Sigma)$.

All diagrams in this paper are either square disc or square spherical diagrams, or are disc or spherical diagrams in the coned-off space associated to a cubical presentation.

Remark 4.3.21. If a disc diagram or spherical diagram has minimal complexity, then it is also reduced, as any of the pathologies in Definition 4.3.18 would indicate a possible complexity reduction. For instance, a cornsquare whose outerpath lies on a cone-cell can be absorbed, after a sequence of hexagon moves – each of which corresponds to pushing a hexagon on one side of a 3-cube to obtain the hexagon on the other side – into said cone-cell; a cancellable pair of squares can be cancelled as described in Definition 4.3.14; a square bigon leads to a cancellable pair of squares after a sequence of hexagon moves; and a pair of combinable cone-cells in D can be combined to reduce the number of cone-cells by 1. A detailed case-by-case analysis of all possible reductions can be found in [Wis21, 3.e].

Definition 4.3.22 (Shell)**.** A *shell* of D is a 2-cell $C \to D$ whose boundary path $\partial C \to D$ is a concatenation $QP_1 \cdots P_k$ for some $k \leq 4$ where Q is a boundary arc in D and $P_1, \ldots, P_k$ are non-trivial pieces in the interior of D. The arc Q is the *outerpath* of C and the concatenation $S := P_1 \cdots P_k$ is the *innerpath* of C.

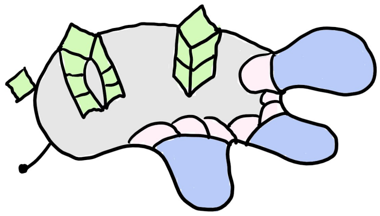

Figure 4.1: A potential disc diagram with a range of features: a spur, shells, corners, cornsquares and a cut-vertex.

Remark 4.3.23. Note that if a cubical presentation X^* satisfies the $C(n)$ condition and D is a minimal complexity diagram, then the outerpath of a shell in D is the concatenation of $\geq n-4$ pieces.

In almost all forms of small-cancellation theory, the main technical result that facilitates proving theorems is a form of diagram classification. In the classical setting, this is known as Greendlinger's Lemma [Gre60], in the cubical setting this is known as diagram trichotomy, or as the "Fundamental Lemma". We state it in a simplified form that is sufficient for our applications, and which we shall call *diagram dichotomy*:

Theorem 4.3.24 (Diagram Dichotomy)**.** *[Jan16] Let $X^* = \langle X | \{Y_i\} \rangle$ be a cubical presentation satisfying the $C(9)$ condition, and let $D \to X^*$ be a disc diagram. Then one of the following holds:*

- *D consists of a single cell,*
- *D has at least two shells and/or corners and/or spurs.*

Theorem 4.3.25. *[Jan16] Let $X^* = \langle X | \{Y_i\} \rangle$ be a cubical presentation satisfying the $C(9)$ condition. Then each Y_i embeds in the universal cover $\widetilde{X^*}$ of the coned-off space X^*.*

4.3.3 Asphericity

We will require the following fundamental fact about combinatorial paths in 2-complexes.

Theorem 4.3.26 (The Van-Kampen Lemma)**.** *Let X be a combinatorial 2-complex. Let $P \to X^1$ be a closed combinatorial path. Then P is nullhomotopic if and only if there exists a disc diagram D in X with $\partial D \cong P$ so that there is a commutative diagram:*

$$\begin{array}{ccc} \partial D & \longrightarrow & D \\ \downarrow & & \downarrow \\ P & \longrightarrow & X \end{array}$$

For our purposes, the spherical version of Van-Kampen's Lemma will also be extremely useful. That is, we need to be able to assert that all elements of π_2 have diagram representatives:

Theorem 4.3.27 (The Spherical Van-Kampen Lemma)**.** *Let X be a combinatorial 2-complex. Then every homotopy class of maps $S^2 \to X$ is represented by a spherical diagram $\Sigma \to X$.*

Theorem 4.3.27 is presented (with or without a proof) in slightly different ways in various sources. The version we use can be extracted from [Fen83, Section 2] via the viewpoint of "pictures", which are planar graph representations of homotopy elements and are dual to spherical diagrams.

4.4 Main theorem

We commence this section with a definition.

Definition 4.4.1 (Classifying space for proper actions)**.** Let G be a group. A *classifying space for proper actions* for G is a G-CW-complex $\underline{E}G$ satisfying:

1. every cell stabiliser is finite,
2. for each finite subgroup $H < G$, the fixed point space $\underline{E}G^H$ is contractible.

Every group G admits an $\underline{E}G$, and all models for $\underline{E}G$ are G-homotopy equivalent[1]; we also have the inequality $cd_{\mathbb{Q}}(G) \leq dim(\underline{E}G)$, for any group G admitting a finite dimensional $\underline{E}G$. More details, and a general construction of classifying spaces for proper actions, can be found in [LÖ5].

To prove our asphericity result, we must slightly simplify the coned-off space X^* associated to a cubical presentation. Intuitively, we do this in $\widetilde{X^*}$ by "collapsing" or "squashing together" various cones into a single one whenever their base spaces correspond to elevations of a $Y_i \to X$ having the same preimage in $\widetilde{X}$. The precise construction is as follows:

Construction 4.4.2. Let $X^* = \langle X | \{Y_i\}_{i \in I} \rangle$ be a cubical presentation and consider the universal cover $\widetilde{X^*}$ of the coned-off space. Note that X is a subspace of X^*, so the preimage $\hat{X}$ of X in $\widetilde{X^*}$ is a covering space of X, namely the regular cover corresponding to $ker(\pi_1 X \to \pi_1 X^*)$. Consider the universal cover $\widetilde{X}$ of X. For each $i \in I$, and for any fixed elevation $\widetilde{Y}_i \to \widetilde{X}$ of $Y_i \to X$, we have that $\pi_1(Y_i, y_{i_0}) < Stab_{\pi_1(X,x_0)}(\widetilde{Y}_i)$. If $X^* = \langle X | \{Y_i\}_{i \in I} \rangle$ is minimal, then by definition $\pi_1(Y_i, y_{i_0}) = Stab_{\pi_1(X,x_0)}(\widetilde{Y}_i)$, but in general this is not the case.

[1] Two G-spaces X and Y are G-homotopy equivalent if there are G-equivariant maps $f : X \to Y, g : Y \to X$ such that the the compositions $f \circ g, g \circ f$ are G-homotopic to the corresponding identity maps id_Y, id_X.

Let $\{g_\ell \pi_1 Y_i\}$ be coset representatives of $\pi_1(Y_i, y_{i_0})$ in $Stab_{\pi_1(X,x_0)}(\widetilde{Y}_i)$. The elevations $\{g_\ell Y_i\}$ have the same image in $\hat{X} \subset \widetilde{X^*}$, so their cones are all isomorphic in $\widetilde{X^*}$. Thus, there is a quotient $\cup_\ell g_\ell C(Im_{\hat{X}}(\widetilde{Y}_i)) \to C(Im_{\hat{X}}(\widetilde{Y}_i))$ where all cones over $Im_{\hat{X}}(\widetilde{Y}_i)$ are identified to a single cone. This extends to a quotient $\widetilde{X^*} \to \bar{X}^*$, which we call the *reduced space* of $\widetilde{X^*}$.

Remark 4.4.3 (Cubically presenting the trivial group). Another way to think about $\bar{X}^*$ is as the cubical presentation $\langle \hat{X} | \{Im_{\hat{X}}(\widetilde{Y}_i)\}\rangle$, where the $\hat{Y}_i$ range over all $i \in I$ and over all elevations of Y_i with distict image. Of course, $\hat{X}$ is not compact, and there are infinitely many $Im_{\hat{X}}(\widetilde{Y}_i)$'s, so this is an "infinitely generated" cubical presentation with infinitely many "relators". It presents, in fact, the trivial group, as we show in Lemma 4.4.4. It is immediate from Construction 4.4.2 that if $\widetilde{X^*}$ satisfies the $C(n)$ condition for $n > 0$, then so does $\bar{X}^*$.

Thus, there are no hidden technicalities in applying the theory from Section 4.3 to $\bar{X}^*$ rather than to $\widetilde{X^*}$. In particular, it makes sense to talk about minimal complexity diagrams and pieces, and Diagram Dichotomy and Theorem 4.3.25 all hold for $\bar{X}^*$.

As mentioned above, the following is a quick but important observation:

Lemma 4.4.4. *$\bar{X}^*$ is simply-connected.*

Proof. Choosing a representative $g_0 C(Im_{\hat{X}}(\widetilde{Y}_i))$ for each collection $\{g_\ell C(Im_{\hat{X}}(\widetilde{Y}_i))\}_{g_\ell}$ described above, we may view the reduced space $\bar{X}^*$ as a subspace of $\widetilde{X^*}$. Thus, $\bar{X}^*$ is a retract of $\widetilde{X^*}$, and $\pi_1 \bar{X}^*$ injects into $\pi_1 \widetilde{X^*}$, and is therefore trivial. ◻

We can now state and prove our main technical theorem.

Theorem 4.4.5. *Let $X^* = \langle X | \{Y_i\}\rangle$ be a cubical presentation that satisfies the $C(9)$ condition. Let $\bar{X}^*$ be the reduced space of $\widetilde{X^*}$. Then $\pi_2 \bar{X}^* = 0$.*

Proof. Let $\Sigma \to \bar{X}^*$ be a minimal complexity spherical diagram, where the minimum is taken over all spherical diagrams representing a fixed non-zero homotopy class in $\pi_2 \bar{X}^*$. If Σ has no cone-cells, then it lies on the cubical part of $\bar{X}^*$ and it is therefore contractible by Theorem 4.2.1. Thus, we may assume that Σ has a cone-cell.

Let C_∞ be a cone-cell in Σ and let D_0 be obtained from Σ by removing from it the interior of C_∞. Note that D_0 is a disc diagram. Indeed, since Σ has minimal complexity, then C_∞ embeds in Σ, as by Theorem 4.3.25 the cone of $\bar{X}^*$ containing C_∞ embeds in $\bar{X}^*$, so a non-embedded cone-cell in Σ would have to be nullhomotopic, contradicting Condition 5 of Definition 4.3.18. Moreover, D_0 has minimal complexity as any complexity reduction in D_0 would lead to a complexity reduction in Σ.

By Theorem 4.3.24, either

1. D_0 is a single vertex or a single cone-cell,

2. D_0 has at least 2 shells and/or corners and/or spurs.

The first case would imply that Σ consists exactly of two cone-cells C_1 and C_2 with $\partial C_1 = \partial C_2 = \sigma$. Assume $\partial C_1 \to Y_1$ and $\partial C_2 \to Y_2$. If $Y_1 = Y_2$, then $\Sigma \to \bar{X}^*$ factors as $\Sigma \to C(Y_1) \to \bar{X}^*$. But $C(Y_1)$ is contractible, because it is a cone, so Σ is nullhomotopic in $C(Y_1)$, and therefore also in $\bar{X}^*$, which contradicts the choice of Σ.

Otherwise, $Y_1 \neq Y_2$. If σ is essential in either Y_1 or Y_2, then either σ is a piece, and thus an essential path that is the concatenation of < 9 pieces, contradicting the $C(9)$ condition, or σ is not a piece, which implies that Y_1 and Y_2 are elevations of the same Y_i that differ by an element of $Stab_{\pi_1 X}(\widetilde{Y})$. In this case, Y_1 and Y_2 have the same image, so $C(Y_1)$ and $C(Y_2)$ are identified in $\bar{X}^*$, and thus $C_1 = C_2$, again contradicting the minimal complexity of Σ. If σ is not essential in neither Y_1 nor Y_2, then there are disc diagrams $D_1 \to Y_1, D_2 \to Y_2$ with $\partial D_1 = \partial C_1$ and $\partial D_2 = \partial C_2$, which together bound a spherical diagram $\Sigma' \to \hat{X}$. As $C(Y_1)$ and $C(Y_2)$ are contractible, Σ' and Σ are homotopic. Since $\hat{X}$ is aspherical – it is a covering space of a non-positively curved cube complex, and thus also non-positively curved – then Σ' is nullhomotopic in $\hat{X}$, and therefore also in $\bar{X}^*$. This in turn implies that Σ is nullhomotopic in $\bar{X}^*$.

In the third case, the presence of spurs immediately implies that D_0 is not a minimal complexity diagram. Likewise, the presence of a corner on ∂D_0 implies the presence of a corner on ∂C_∞, contradicting that the spherical diagram Σ has minimal complexity, as such a corner could be absorbed into C_∞, reducing the number of squares in Σ by Condition 2 of Definition 4.3.18. In particular, this implies that D_0 cannot be a square disc diagram. We can therefore conclude that D_0 is a disc diagram with at least two shells – in fact, all we need to use now is that D_0 has at least one shell.

Let C be a shell of D_0. Then the innerpath of C is at most 4 pieces, and the outerpath of ∂C coincides with a subpath of ∂C_∞ in Σ. Thus, the outerpath of C is a single piece, contradicting the $C(9)$ condition. □

We can now deduce asphericity for the reduced space $\bar{X}^*$ associated to a low-dimensional $C(9)$ cubical presentation, and in particular for X^* when the cubical presentation is minimal:

Corollary 4.4.6. *Let $X^* = \langle X | \{Y_i\} \rangle$ be a cubical presentation that satisfies the $C(9)$ condition. If $dim(X) \leq 2$ and $\max\{cd(Y_i)\} = 1$, then $\bar{X}^*$ is contractible.*

Proof. As explained in Construction 4.5.1, the preimage $\hat{X}$ of X in $\widetilde{X^*}$ is a covering space of X, corresponding to $ker(\pi_1 X \to \pi_1 X^*)$. By Theorem 4.3.25, each Y_i embeds in $\widetilde{X^*}$, and thus also in the quotient $\bar{X}^*$. Viewing $\bar{X}^*$ as a cubical presentation, $\bar{X}^*$ decomposes as the union of its cubical part $\hat{X}$ and the cones over all elevations of Y_i's with distinct images in $\widetilde{X^*}$. To simplify notation, let

$$\bigsqcup_i \bigsqcup_{gStab_{\pi_1(X,x_0)}(\widetilde{Y}_i)\in\pi_1X/\pi_1Y} gY_i := \mathbf{Y} \text{ and } \bigsqcup_i \bigsqcup_{gStab_{\pi_1(X,x_0)}(\widetilde{Y}_i)\in\pi_1X/\pi_1Y} gC(Y_i) := \mathbf{C}(\mathbf{Y}),$$

And note that

$$\hat{X}\cap\mathbf{C}(\mathbf{Y}) = \mathbf{Y}.$$

We have the Mayer-Vietoris sequence:

$$\cdots \longrightarrow H_n(\mathbf{Y}) \longrightarrow H_n(\hat{X})\oplus H_n(\mathbf{C}(\mathbf{Y})) \longrightarrow H_n(\bar{X}^*) \longrightarrow$$

$$\longrightarrow H_{n-1}(\mathbf{Y}) \longrightarrow \cdots \qquad \cdots \longrightarrow H_0(\hat{X}^*) \longrightarrow 0$$

Since $\mathbf{Y}$ is homotopy equivalent to a graph, then $H_n(\mathbf{Y}) = 0$ whenever $n\geq 2$, so we get isomorphisms $H_n(\hat{X})\cong H_n(\bar{X}^*)$ for each $n\geq 2$, since $C(gY_i))$ is contractible for each $g\in\pi_1X^*$ and $i\in I$. Now, $H_2(\bar{X}^*)\cong\pi_2\bar{X}^* = 0$ by Theorem 4.4.5 and Hurewicz's Theorem, and $H_3(\bar{X}^*) = 0$ since the sequence is exact and the terms on the right and left of $H_3(\bar{X}^*)$ are equal to zero. Since $\bar{X}^*$ is 2-connected, we may apply Hurewicz again to conclude that $\pi_3\bar{X}^* = 0$.

The proof is now finished: since by Theorem 4.4.5, $\pi_1\bar{X}^* = 0$, and $\bar{X}^*$ has no cells of dimension ≥ 3, then $H_n(\bar{X}^*)\cong\pi_n\bar{X}^* = 0$ for all $n\in\mathbb{N}$. ◻

The first of our two theorems is now established:

Theorem 4.4.7. *Let $X^* = \langle X|\{Y_i\}_{i\in I}\rangle$ be a minimal cubical presentation that satisfies the $C(9)$ condition. Let $\pi_1X^* = G$. If $dim(X)\leq 2$ and each Y_i is homotopy equivalent to a graph, then X^* is a $K(G,1)$, so G is torsion-free and $gd(G)\leq 2$.*

Proof. This is a direct consequence of Corollary 4.4.6, since then, by minimality, $\bar{X}^* = \widetilde{X^*}$ is contractible. ◻

Lemma 4.4.8. *Let $X^* = \langle X|\{Y_i\}_{i\in I}\rangle$ be a symmetric cubical presentation. Then $\bar{X}^*$ is a π_1X^*-CW-complex, and if v is a cone-vertex of $\bar{X}^*$ corresponding to a cone over some Y_i, then $Stab_{\pi_1X^*}(v) = Stab_{\pi_1X}(\widetilde{Y}_i)/\pi_1Y_i$.*

Proof. The group π_1X^* acts on $\widetilde{X^*}$ by permuting the cones. The action is a covering space action, and in particular is free, and induces an action on $\bar{X}^*$. For each cone-vertex v in $\bar{X}^*$, $Stab_{\pi_1X^*}(v) = Stab_{\pi_1X^*}(C(Im_{\hat{X}}(\widetilde{Y}_i)) = \{g_\ell\pi_1Y_i\}$ where $C(Im_{\hat{X}}(\widetilde{Y}_i))$ is the cone with v as its cone-vertex, and $\{g_\ell\pi_1Y_i\}$ are left coset representatives of $\pi_1(Y_i,y_{i_0})$ in $Stab_{\pi_1(X,x_0)}(\widetilde{Y}_i)$. This is exactly the quotient $Stab_{\pi_1X}(\widetilde{Y}_i)/\pi_1Y_i$. ◻

Lemma 4.4.9. *If $X^* = \langle X | \{Y_i\} \rangle$ is a symmetric cubical presentation then the action of $\pi_1 X^*$ on $\bar{X}^*$ has finite cell-stabilisers.*

Proof. All stabilisers of cells in $\hat{X} \subset \bar{X}^*$ are trivial; since the cubical presentation is symmetric, then $[Stab_{\pi_1 X}(\widetilde{Y}) : \pi_1 Y_i] < \infty$, so Lemma 4.4.8 implies that the stabilisers of cone-vertices are finite. ∎

So far we have shown that if $X^* = \langle X | \{Y_i\} \rangle$ satisfies the $C(9)$ condition, then $\bar{X}^*$ is contractible, and that if, in addition, X^* is symmetric, then the action of $\pi_1 X^*$ on $\bar{X}^*$ has finite cell-stabilisers. To show that $\bar{X}^*$ is a classifying space for proper actions for $\pi_1 X^*$, we still must prove that every finite subgroup $H < \pi_1 X^*$ fixes a point in $\bar{X}^*$, and that this point is unique. Note that since $\bar{X}^*$ is aspherical and finite-dimensional, then the action of a finite subgroup H on $\bar{X}^*$ cannot be free, as otherwise the quotient $\bar{X}^*/H$ would be a classifying space for H. However this does not necessarily imply that the action has a global fixed point.

Lemma 4.4.10. *If a non-trivial subgroup $H < \pi_1 X^*$ fixes a point in $\bar{X}^*$, then H is conjugate into $Stab_{\pi_1 X^*}(Y_i)$ for some elevation $Y_i \to \hat{X}$ of some $Y_i \to X$.*

Proof. Let H be a non-trivial subgroup of $\pi_1 X^*$. Since the action of H on $\widetilde{X^*}$ is free and $\widetilde{X^*}$ coincides with $\bar{X}^*$ outside of the cones, then a fixed point ξ under the action of H on $\bar{X}^*$ must be a cone-vertex corresponding to some left coset $gStab_{\pi_1 X^*}(Y_i)$ for some elevation $Y_i \to \hat{X}$ of some $Y_i \to X$. Thus, $HgStab_{\pi_1 X^*}(Y_i) = gStab_{\pi_1 X}(Y_i)$, so $g^{-1}HgStab_{\pi_1 X^*}(Y_i) = Stab_{\pi_1 X^*}(Y_i)$ and H is conjugate into $Stab_{\pi_1 X^*}(Y_i)$. ∎

Lemma 4.4.11. *If a non-trivial subgroup $H < \pi_1 X^*$ fixes a point in $\bar{X}^*$, then that point is unique, and in particular the fixed-point space $(\bar{X}^*)^H$ is contractible.*

Before proving the lemma, we need an auxiliary result.

Proposition 4.4.12. *Let $X^* = \langle X | \{Y_i\}_{i=1}^k \rangle$ be a cubical presentation that satisfies the $C(n)$ condition for some $n \geq 2$, and where each Y_i is compact, then there is a uniform upper-bound $\mathscr{L} \geq 0$ on the size of pieces in $\widetilde{X}^*$.*

Proof. Since each Y_i is compact, $\pi_1 Y_i$ is finitely generated for each $i \in \{1, \ldots, k\}$, and since the cubical presentation has finitely many relations $Y_i \to X$, then finitely many elements of $\pi_1 X$ suffice to generate all the $\pi_1 Y_i$'s. The compactness of the Y_i's implies additionally that there is an uniform upper-bound ℓ_i for the length of minimal-length paths representing the generators of $\pi_1 Y_i$, thus, there is a bound $\mathscr{L} = \max\{\ell_i\}$ for the size of pieces arising in generators of $\pi_1 Y_1, \ldots, \pi_1 Y_k$. Since X^* satisfies the $C(n)$ condition for $n \geq 2$, then each path representing a generator is a concatenation of at least 2 pieces, and in particular any such piece p satisfies $|p| < \ell_i \leq \mathscr{L}$ whenever it arises in a generator of a $\pi_1 Y_i$.

Assume that there is an essential closed path $\sigma \to Y_i$ that is a concatenation of pieces having a piece p_r with $|p_r| > \mathscr{L}$. By the discussion above, we can assume in particular that σ does not represents a generator of a $\pi_1 Y_{i'}$ for any $i' \in \{1,\dots,k\}$. Note that any expression for σ as a concatenation of pieces can be further expressed as a concatenation of pieces in generators of $\pi_1 Y_i$. Write $\sigma = \alpha_1 \cdots \alpha_m = p_1 \cdots p_r \cdots p_n$, where each α_j is a (not necessarily distinct) generator of $\pi_1 Y_i$. Since $|p_r| > \mathscr{L}$, then there is an α_j with $j \in \{1,\dots,m\}$ for which $p_r \cap \alpha_j = \alpha_j$. Thus, α_j can be expressed as a single piece $q_r \subset p_r$, contradicting the discussion in the previous paragraph, and hence the $C(n)$ condition. ◻

The proof of Lemma 4.4.11 now combines Proposition 4.4.12 with the hypothesised symmetry of the cubical presentation under consideration.

Proof of Lemma 4.4.11. Let $H < \pi_1 X^*$. Assume $\xi \neq \xi'$ are cone-vertices of $\bar{X}^*$ fixed by H, so $H\xi = \xi$ and $H\xi' = \xi'$. Then Lemma 4.4.10 implies that $H \subset g^{-1} Stab_{\pi_1 X^*}(Y_i) g \cap g'^{-1} Stab_{\pi_1 X^*}(Y_{i'}) g'$ for some $i,i' \in \{1,\dots,k\}$ and $g,g' \in \pi_1 X^*$. We claim that the collection $\{Stab_{\pi_1 X^*}(Y_i)\}$ is malnormal, and thus H is the trivial subgroup. To this end, we first show that the collection $\{Stab_{\pi_1 X}(\widetilde{Y}_i)\}$ is malnormal in $\pi_1 X$.

Suppose that the intersection $Stab_{\pi_1 X}(\widetilde{Y}_i)^{\tilde{g}} \cap Stab_{\pi_1 X}(\widetilde{Y}_j)^{\tilde{g}'}$ is infinite for some $i,i' \in \{1,\dots,k\}$ and $\tilde{g},\tilde{g}' \in \pi_1 X$. Since X^* is symmetric, then $[Stab_{\pi_1 X}(\widetilde{Y}_i) : \pi_1 Y_i] < \infty$ and $[Stab_{\pi_1 X}(\widetilde{Y}_j) : \pi_1 Y_j] < \infty$, so $\pi_1 Y_i^{\tilde{g}} \cap \pi_1 Y_j^{\tilde{g}'}$ has finite index in $Stab_{\pi_1 X}(\widetilde{Y}_i)^{\tilde{g}} \cap Stab_{\pi_1 X}(\widetilde{Y}_j)^{\tilde{g}'}$. Thus, $\pi_1 Y_i^{\tilde{g}} \cap \pi_1 Y_j^{\tilde{g}'}$ is infinite, and in particular contains an infinite order element $\tilde{h}$. Therefore, the axis of $\tilde{h}$ in $\widetilde{X}$ is an unbounded piece between elevations $\widetilde{Y}_i$ and $\widetilde{Y}_j$ of Y_i and Y_j, contradicting Proposition 4.4.12, and thus the $C(9)$ condition. We conclude that $Stab_{\pi_1 X}(\widetilde{Y}_i)^{\tilde{g}} \cap Stab_{\pi_1 X}(\widetilde{Y}_j)^{\tilde{g}'}$ is finite; since $\pi_1 X$ is torsion-free, then in fact $Stab_{\pi_1 X}(\widetilde{Y}_i)^{\tilde{g}} \cap Stab_{\pi_1 X}(\widetilde{Y}_j)^{\tilde{g}'}$ must be trivial, so the collection $\{Stab_{\pi_1 X}(\widetilde{Y}_i)\}$ is malnormal in $\pi_1 X$.

To promote this to malnormality of $\{Stab_{\pi_1 X^*}(Y_i)\}$ in $\pi_1 X^*$, let $h \in Stab_{\pi_1 X^*}(Y_i)^g \cap Stab_{\pi_1 X^*}(Y_j)^{g'}$, and let $\tilde{h}_i \in Stab_{\pi_1 X}(\widetilde{Y}_i)^{\tilde{g}}$ and $\tilde{h}_j \in Stab_{\pi_1 X}(\widetilde{Y}_j)^{\tilde{g}'}$ be preimages of h, so $\tilde{h}_i \tilde{h}_j^{-1}$ is an element of $\pi_1 \hat{X}$. Since $\tilde{h}_i \in Stab_{\pi_1 X}(\widetilde{Y}_i)^{\tilde{g}}$ and $\tilde{h}_j \in Stab_{\pi_1 X}(\widetilde{Y}_j)^{\tilde{g}'}$ and as mentioned before, by symmetry, $[Stab_{\pi_1 X}(\widetilde{Y}_i) : \pi_1 Y_i] < \infty$ and $[Stab_{\pi_1 X}(\widetilde{Y}_j) : \pi_1 Y_j] < \infty$, then there are paths $\sigma_i \to Y_i, \sigma_j \to Y_j$ such that for some $k,k' \in \mathbb{N}$, the paths $\sigma_i^k, \sigma_j^{k'}$ bound disc diagrams $D_i \to Y_i, D_j \to Y_j$ in $\widetilde{X^*}$, each of which can be assumed to have minimal complexity, and therefore consisting of a single cone-cell. Moreover, by the discussion above, σ_i and σ_j represent lifts of conjugates of $\tilde{h}_i$ and $\tilde{h}_j$, so there are paths σ_i', σ_j' such that $\sigma_i' \sigma_j'^{-1}$ is a closed path that represents $\tilde{h}_i \tilde{h}_j^{-1}$. Thus, $\sigma_i' \sigma_j'^{-1}$ bounds a disc diagram D in $\widetilde{X^*}$, and there are paths $\ell_i, \ell_i', \ell_j, \ell_j'$ that are lifts of the corresponding conjugating elements, so that $\sigma_i \ell_i \sigma_i' \ell_i'$ and $\sigma_j \ell_j \sigma_j' \ell_j'$ bound square-disc diagrams R_i, R_j in $\widetilde{X^*}$. We can further assume that R_i, R_j and D are chosen to have minimal area and minimal complexity amongst all possible diagrams with the same corresponding boundary. For R_i and R_j, this means in particular that there are no cornsquares on any of $\sigma_i, \sigma_i', \sigma_j, \sigma_j'$, so each of these paths determines a hyperplane carrier in the corresponding diagram R_i or R_j, and more generally that R_i and R_j are *grids*, in the sense that they are cubically isomorphic to a product of combinatorial intervals $I_n \times I_m$ for suitable $m, n \in \mathbb{N}$. We note that, of course, R_i or R_j could be degenerate square diagrams, in which case $\sigma_i = \sigma_i'$ or $\sigma_j = \sigma_j'$, and the argument in the next paragraph becomes a bit simpler.

If D is a square disc diagram, then σ_1' and σ_j' are homotopic rel their endpoints in $\widetilde{X^*}$, so $\tilde{h}_i = \tilde{h}_j$, contradicting malnormality in $\pi_1 X$. We claim that this is the only possibility for D. Assume towards a contradiction that D contains at least one cone-cell, and observe that since any corners or spurs on ∂D could be removed, improving the choices made in the previous paragraph, that cone-cell must be a shell S. Consider the disc diagram $F = D_i \cup_{\sigma_i} R_i \cup_{\sigma_i'} D \cup_{\sigma_j'} R_j \cup_{\sigma_j} D_j$. Then, since S is a shell of D, its innerpath is a concatenation of at most 4 pieces, and either the outerpath of S is a subpath of $\sigma_i' \subset R_i$, or a subpath of $\sigma_j' \subset R_j$, or a subpath of their concatenation. In the first case, since $\partial S \cap R_i$ in F is the outerpath of S in D, then $\partial S \cap R_i$ is a single piece (a wall piece if R_i is not a degenerate diagram, and a cone piece otherwise), so ∂S is the concatenation of at most 5 pieces, contradicting the $C(9)$ condition. Similarly if the outerpath of S is a subpath of σ_j'. In the last case, ∂S is the concatenation of at most 6 pieces: the 4 pieces coming from its innerpath and a piece coming from each of σ_i', σ_j', contradicting again the $C(9)$ condition.

So the collection $\{Stab_{\pi_1 X^*}(Y_i)\}$ is malnormal in $\pi_1 X^*$, and $H = \{1\}$ or $\xi = \xi'$. ◻

Remark 4.4.13. In the proof above, we use that $\pi_1 X$ is torsion-free, which follows trivially from the assumption that X is finite-dimensional, since X is aspherical, but is true also for infinite dimensional non-positively curved cube complexes, and is a consequence of CAT(0) geometry [BH99, Lea13].

We can finally conclude that, in fact, the finite subgroups of $\pi_1 X^*$ have non-empty fixed-point sets under their action on $\bar{X}^*$:

Corollary 4.4.14. *Every finite subgroup $H < \pi_1 X^*$ fixes a point in $\bar{X}^*$.*

Proof. Let H be a finite subgroup of $\pi_1 X^*$. If H is trivial, then it fixes all of $\bar{X}^*$. Otherwise, since the action of H on $\bar{X}^*$ cannot be free, for any non-trivial $h \in H$, there exists $k \in \mathbb{N}$ for which h^k fixes a point in $\bar{X}^*$. This implies that each cyclic subgroup $\langle h \rangle$ fixes a point, and in particular that every $h \in H$ fixes a point in $\bar{X}^*$. Indeed, if $h^k v = v$ for some $k \in \mathbb{N}$ and $v \in \bar{X}^*$, then $hv = h(h^k v) = h^k(hv)$, so h^k fixes hv and by Lemma 4.4.11 $hv = v$.

To see now that every element of H fixes the same point, we use Burnside's Lemma: since H is a finite group acting on a finite set S (the union of the orbits of the fixed points of non-trivial elements), and each non-trivial element fixes a unique point, if m is the number of orbits of elements, then $m = \frac{|S|+|H|-1}{|H|}$, so $|S| = (m-1)|H|+1$ and in particular, as the size of each orbit must divide $|H|$, at most one orbit of points can have size $< |H|$, and therefore the remaining $m-1$ orbits must have size 1, and H must fix a point in $\bar{X}^*$. ☐

Putting together the previous results, we obtain:

Theorem 4.4.15. *Let $X^* = \langle X | \{Y_i\}_{i=1}^k \rangle$ be a symmetric cubical presentation that satisfies the $C(9)$ condition. Let $\pi_1 X^* = G$. If $dim(X) \leq 2$ and each Y_i is homotopy equivalent to a graph, then there is a quotient $\bar{X}^*$ of $\widetilde{X^*}$ that is an $\underline{E}G$, so $cd_{\mathbb{Q}}(G) \leq dim(\bar{X}^*) \leq 2$.*

Proof. This follows from Corollary 4.4.6, Lemma 4.4.11, and Corollary 4.4.14. ☐

We can go a step further if we assume that X has a finite regular cover where every $Y_i \to X$ lifts to an embedding, as is observed in [Wis21, 4.4] for $C'(\frac{1}{20})$ cubical presentations.

Lemma 4.4.16. *Let $X^* = \langle X | \{Y_i\}_{i \in I} \rangle$ be a symmetric $C(9)$ cubical presentation where $dim(X) \leq 2$ and each Y_i is homotopy equivalent to a graph. If there exists a finite regular cover $\widehat{X} \to X$ where $Y_i \to X$ lifts to an embedding for each $i \in I$, then $\pi_1 X^*$ is virtually torsion-free.*

Proof. Consider the covering $\widehat{X^*} \to X^*$ induced by $\widehat{X} \to X$, which is a finite-degree covering because $\widehat{X} \to X$ is. If $g \in \pi_1 X^* - \{1\}$ satisfies $g^n = 1$ for some $n \in \mathbb{N}$, then Corollary 4.4.14 implies that if $\sigma \to X^*$ is a closed path representing g, then σ^n is conjugate to a closed path in some Y_i. Now since g is non-trivial, σ does not lift to a closed path in Y_i. Since $Y_i \to X$ lifts to an embedding $Y_i \to \widehat{X}$, then σ cannot lift to a closed path in $\widehat{X}$ either, so σ is not closed in $\widehat{X^*}$, and $g \notin \pi_1 \widehat{X^*}$. ☐

If a group G has finite virtual cohomological dimension, then $vcd(G) \leq dim(\underline{E}G)$, since a torsion-free subgroup $G' < G$ of finite index has $cd(G') = vcd(G)$, and $dim(\underline{E}G)$ gives an upper bound for the cohomological dimension of any torsion-free subgroup of G. In particular:

Corollary 4.4.17. *Let $X^* = \langle X | \{Y_i\}_{i \in I} \rangle$ be a symmetric $C(9)$ cubical presentation with $G = \pi_1 X^*$. If $dim(X) \leq 2$ and each Y_i is homotopy equivalent to a graph, and there exists a finite regular cover $\widehat{X} \to X$ where $Y_i \to X$ lifts to an embedding for each $i \in I$, then $vcd(G) \leq 2$.* ☐

4.5 Examples

Using Theorems 4.4.7 and 4.4.15, we can provide classifying spaces, or classifying spaces for proper actions, in the situations outlined below:

Example 4.5.1. (Forcing small-cancellation by taking covers) If X is a compact non-positively curved cube complex with hyperbolic fundamental group, and $H_1,...,H_k$ are quasiconvex subgroups of $\pi_1 X$ that form a malnormal collection, then for each $n > 0$ there are finite index subgroups $H_i' \subset H_i$ and local isometries $Y_i \to X$ with Y_i compact and $\pi_1 Y_i = H_i'$ such that $X^* = \langle X|\{Y_i\}\rangle$ satisfies $C(n)$. How one achieves this condition is explained in [Wis21, 3.51]. To guarantee that X^* is symmetric, it suffices to require that $H_i = C_{\pi_1 X}(H_i)$. When $n \geq 9$, $dim(X) \leq 2$ and $H_1,...,H_k$ are free, we conclude that $\bar{X}^*$ is a classifying space for proper actions for $\pi_1 X^*$, and in particular that $\pi_1 X^*$ has rational cohomological dimension $cd_{\mathbb{Q}}\pi_1 X^* \leq 2$.

Example 4.5.2 (Forcing small-cancellation by creating noise). Let X be a compact non-positively curved cube complex with non-elementary hyperbolic fundamental group, as above. Then by Theorem 3.3.2 of Chapter 3, for all $k \geq 1$ there exist (infinitely many choices of) free non-abelian subgroups $\{H_1,\dots,H_k\}$, and cyclic subgroups $z_i < H_i$ such that there exist local isometries $Y_i \to X$ with Y_i compact and $\pi_1 Y_i = z_i$, so that the cubical presentation $X^* = \langle X|\{Y_i\}\rangle$ is minimal, and satisfies the $C(9)$ condition. Thus, if $\dim(X) \leq 2$, then X^* is a classifying space for all such examples.

When X is a square complex whose fundamental group is not hyperbolic, general constructions of $C(9)$ cubical presentations are harder to come by, but can be produced by hand in specific situations. For instance, if X_Γ is the Salvetti complex of a RAAG whose defining graph Γ has no triangles, then $dim(X) \leq 2$, and one can produce a wealth of cubical presentations $X^* = \langle X_\Gamma|\{Y_i\}\rangle$ that satisfy the $C(9)$ condition and are minimal. This is outlined in [Wis21, 3.s (4)].

Thus, Theorems 4.4.7 and 4.4.15 are widely applicable, even when one start with a non-positively curved cube complex X whose fundamental group is far from being hyperbolic.

The End.

Bibliography

[AB20] Karim Adiprasito and Bruno Benedetti. Collapsibility of CAT(0) spaces. *Geom. Dedicata*, 206:181–199, 2020.

[Ago08] Ian Agol. Criteria for virtual fibering. *J. Topol.*, 1(2):269–284, 2008.

[Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[AH22] Goulnara N. Arzhantseva and Mark F. Hagen. Acylindrical hyperbolicity of cubical small cancellation groups. *Algebr. Geom. Topol.*, 22(5):2007–2078, 2022.

[Are] Macarena Arenas. Asphericity of cubical presentations: the general case. *in preparation*.

[Are23] Macarena Arenas. A cubical rips construction. 2023. To appear in Algebraic & Geometric Topology. Available at https://arxiv.org/pdf/2202.01048.pdf.

[Arz01] G. N. Arzhantseva. On quasiconvex subgroups of word hyperbolic groups. *Geom. Dedicata*, 87(1-3):191–208, 2001.

[AS14] Goulnara Arzhantseva and Markus Steenbock. Rips construction without unique product, 2014.

[BBD07] Josh Barnard, Noel Brady, and Pallavi Dani. Super-exponential distortion of subgroups of CAT(-1) groups. *Algebr. Geom. Topol.*, 7:301–308, 2007.

[BBMS00] Gilbert Baumslag, Martin R. Bridson, Charles F. Miller, III, and Hamish Short. Fibre products, non-positive curvature, and decision problems. *Comment. Math. Helv.*, 75(3):457–477, 2000.

[Bes] Mladen Bestvina. Questions in geometric group theory. http://www.math.utah.edu/ bestvina/eprints/questions-updated.pdf.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

[Bie81] Robert Bieri. *Homological dimension of discrete groups*. Queen Mary College Department of Pure Mathematics, London, second edition, 1981.

[BM00] Marc Burger and Shahar Mozes. Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.*, (92):151–194 (2001), 2000.

[BMS94] G. Baumslag, Charles F. Miller, III, and H. Short. Unsolvable problems about small cancellation and word hyperbolic groups. *Bull. London Math. Soc.*, 26(1):97–101, 1994.

[BO08] Igor Belegradek and Denis Osin. Rips construction and Kazhdan property (T). *Groups Geom. Dyn.*, 2(1):1–12, 2008.

[Boo58] William W. Boone. The word problem. *Proc. Natl. Acad. Sci. USA*, 44:1061–1065, 1958.

[Bro94] Kenneth S. Brown. *Cohomology of groups*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[BSV14] Mladen Bestvina, Michah Sageev, and Karen Vogtmann, editors. *Geometric group theory*, volume 21 of *IAS/Park City Mathematics Series*. American Mathematical Society, Providence, RI; Institute for Advanced Study (IAS), Princeton, NJ, 2014. Including lecture notes from the Graduate Summer School held at the Park City Mathematics Institute (PCMI), Park City, UT, July 1–21, 2012.

[BW05] Inna Bumagin and Daniel T. Wise. Every group is an outer automorphism group of a finitely generated group. *J. Pure Appl. Algebra*, 200(1-2):137–147, 2005.

[BW12] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[Cap19] Pierre-Emmanuel Caprace. Finite and infinite quotients of discrete and indiscrete groups. In *Groups St Andrews 2017 in Birmingham*, volume 455 of *London Math. Soc. Lecture Note Ser.*, pages 16–69. Cambridge Univ. Press, Cambridge, 2019.

[CL63] D. E. Cohen and R. C. Lyndon. Free bases for normal subgroups of free groups. *Trans. Amer. Math. Soc.*, 108:526–537, 1963.

[CN05] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005.

[Cou11] Rémi Coulon. Asphericity and small cancellation theory for rotation families of groups. *Groups Geom. Dyn.*, 5(4):729–765, 2011.

[Deh87] Max Dehn. *Papers on group theory and topology*. Springer-Verlag, New York, 1987. Translated from the German and with introductions and an appendix by John Stillwell, With an appendix by Otto Schreier.

[EG57] Samuel Eilenberg and Tudor Ganea. On the Lusternik-Schnirelmann category of abstract groups. *Ann. of Math. (2)*, 65:517–518, 1957.

[EH87] Martin Edjvet and James Howie. A Cohen-Lyndon theorem for free products of locally indicable groups. *J. Pure Appl. Algebra*, 45(1):41–44, 1987.

[Fen83] Roger A. Fenn. *Techniques of geometric topology*, volume 57 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1983.

[FM10] Koji Fujiwara and Jason Fox Manning. CAT(0) and CAT(-1) fillings of hyperbolic manifolds. *J. Differential Geom.*, 85(2):229–269, 2010.

[FMR02] Benjamin Fine, Alexei Myasnikov, and Gerhard Rosenberger. Malnormal subgroups of free groups. *Comm. Algebra*, 30(9):4155–4164, 2002.

[Ger96] S. M. Gersten. Subgroups of word hyperbolic groups in dimension 2. *J. London Math. Soc. (2)*, 54(2):261–283, 1996.

[GMRS98] Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.

[Gre60] Martin Greendlinger. On Dehn's algorithms for the conjugacy and word problems, with applications. *Comm. Pure Appl. Math.*, 13:641–677, 1960.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gro93] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[Gro03] M. Gromov. Random walk in random groups. *Geom. Funct. Anal.*, 13(1):73–146, 2003.

[Gru15] Dominik Gruber. Groups with graphical $C(6)$ and $C(7)$ small cancellation presentations. *Trans. Amer. Math. Soc.*, 367(3):2051–2078, 2015.

[GS91] S. M. Gersten and H. Short. Small cancellation theory and automatic groups. II. *Invent. Math.*, 105(3):641–662, 1991.

[Hag08] Frédéric Haglund. Finite index subgroups of graph products. *Geom. Dedicata*, 135:167–209, 2008.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1 551–1620, 2008.

[HW12] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. *Ann. of Math. (2)*, 176(3):1427–1482, 2012.

[HW15] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. *Invent. Math.*, 199:293–331, 2015.

[Jan16] Kasia Jankiewicz. The fundamental theorem of cubical small cancellation theory. *Trans. Amer. Math. Soc.*, 2016.

[JŚ03] Tadeusz Januszkiewicz and Jacek Świątkowski. Hyperbolic Coxeter groups of large dimension. *Comment. Math. Helv.*, 78(3):555–583, 2003.

[JW22] Kasia Jankiewicz and Daniel T. Wise. Cubulating small cancellation free products. *Indiana Univ. Math. J.*, 71(4):1397–1409, 2022.

[Kap99] Ilya Kapovich. A non-quasiconvexity embedding theorem for hyperbolic groups. *Math. Proc. Cambridge Philos. Soc.*, 127(3):461–486, 1999.

[KS96] Ilya Kapovich and Hamish Short. Greenberg's theorem for quasiconvex subgroups of word hyperbolic groups. *Canad. J. Math.*, 48(6):1224–1244, 1996.

[LÖ5] Wolfgang Lück. Survey on classifying spaces for families of subgroups. In *Infinite groups: geometric, combinatorial and dynamical aspects*, volume 248 of *Progr. Math.*, pages 269–322. Birkhäuser, Basel, 2005.

[Lea13] Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013.

[LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.

[Lyn66] Roger C. Lyndon. On Dehn's algorithm. *Math. Ann.*, 166:208–228, 1966.

[MS97] Lee Mosher and Michah Sageev. Nonmanifold hyperbolic groups of high cohomological dimension, 1997.

[Nov55] P. S. Novikov. Über die algorithmische Unentscheidbarkeit des Wortproblems in der Gruppentheorie. Tr. Mat. Inst. Steklova 44, 140 S. (1955)., 1955.

[NR97] Graham Niblo and Lawrence Reeves. Groups acting on CAT(0) cube complexes. *Geom. Topol.*, 1:approx. 7 pp. (electronic), 1997.

[Ol′91] A. Yu. Ol′shanskiĭ. *Geometry of defining relations in groups*, volume 70 of *Mathematics and its Applications (Soviet Series)*. Kluwer Academic Publishers Group, Dordrecht, 1991. Translated from the 1989 Russian original by Yu. A. Bakhturin.

[Oll05] Yann Ollivier. *A January 2005 invitation to random groups*, volume 10 of *Ensaios Matemáticos [Mathematical Surveys]*. Sociedade Brasileira de Matemática, Rio de Janeiro, 2005.

[Oll06] Yann Ollivier. On a small cancellation theorem of Gromov. *Bull. Belg. Math. Soc. Simon Stevin*, 13(1):75–89, 2006.

[Osa13] Damian Osajda. A construction of hyperbolic Coxeter groups. *Comment. Math. Helv.*, 88(2):353–367, 2013.

[OW07] Yann Ollivier and Daniel T. Wise. Kazhdan groups with infinite outer automorphism group. *Trans. Amer. Math. Soc.*, 359(5):1959–1976, 2007.

[PS21] Nansen Petrosyan and Bin Sun. Cohomology of group theoretic dehn fillings ii, 2021.

[Rip82] E. Rips. Subgroups of small cancellation groups. *Bull. Lond. Math. Soc.*, 14(1):45–47, 1982.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.

[Sco78] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.

[Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.

[Str90] Ralph Strebel. Appendix. Small cancellation groups. In *Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988)*, volume 83 of *Progr. Math.*, pages 227–273. Birkhäuser Boston, Boston, MA, 1990.

[Sun20] Bin Sun. Cohomology of group theoretic Dehn fillings I: Cohen-Lyndon type theorems. *J. Algebra*, 542:277–307, 2020.

[SW15] Michah Sageev and Daniel T. Wise. Cores for quasiconvex actions. *Proc. Amer. Math. Soc.*, 143(7):2731–2741, 2015.

[Swa69] Richard G. Swan. Groups of cohomological dimension one. *J. Algebra*, 12:585–610, 1969.

[Val02] Alain Valette. *Introduction to the Baum-Connes conjecture*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2002. From notes taken by Indira Chatterji, With an appendix by Guido Mislin.

[Whi39] J. H. C. Whitehead. Simplicial Spaces, Nuclei and m-Groups. *Proc. London Math. Soc. (2)*, 45(4):243–327, 1939.

[Wis98] Daniel T. Wise. Incoherent negatively curved groups. *Proc. Amer. Math. Soc.*, 126(4):957–964, 1998.

[Wis03] Daniel T. Wise. A residually finite version of Rips's construction. *Bull. London Math. Soc.*, 35(1):23–29, 2003.

[Wis04] Daniel T. Wise. Cubulating small cancellation groups. *GAFA, Geom. Funct. Anal.*, 14(1):150–214, 2004.

[Wis07] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.

[Wis12] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012.

[Wis21] Daniel T. Wise. *The structure of groups with a quasiconvex hierarchy*, volume 209 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2021.