

**From Riches to Raags: 3-Manifolds,
Right-Angled Artin Groups, and Cubical
Geometry**

Daniel T. Wise

DEPT. OF MATHEMATICS & STATISTICS, MCGILL UNIVERSITY, MON-
TREAL, QUEBEC, CANADA H3A 0B9
E-mail address: wise@math.mcgill.ca

2000 *Mathematics Subject Classification.* 53C23, 20F36, 20F55, 20F67,
20F65, 20E26

Key words and phrases. CAT(0) Cube Complex, Right-angled Artin
Group, Subgroup Separable, Small-Cancellation Theory, 3-Manifolds

Research supported by NSERC.

ABSTRACT. These notes aim to provide an introduction to the geometric group theory woven around nonpositively curved cube complexes. The main goal is to outline the proof that a hyperbolic group G with a quasiconvex hierarchy has a finite index subgroup that embeds in a right-angled Artin group. We sketch the supporting ingredients of the proof: The basics of nonpositively curved cube complexes, wallspaces and dual CAT(0) cube complexes, special cube complexes, the combination theorem for special cube complexes, the combination theorem for cubulated groups, cubical small-cancellation theory, and the malnormal special quotient theorem. We also discuss generalizations to relatively hyperbolic groups. Finally, we describe applications towards Baumslag's conjecture on the residual finiteness of one-relator groups with torsion, and to the virtual specialness and virtual fibering of certain hyperbolic 3-manifolds, including those with at least one cusp.

To Yael

Contents

Acknowledgement	1
Note to the reader	2
Chapter 1. Overview	3
1.1. Applications	6
1.2. A Scheme for Understanding Groups	7
Chapter 2. Nonpositively Curved Cube Complexes	8
2.1. Definitions	8
2.2. Some Favorite 2-Dimensional Examples	9
2.3. Right-Angled Artin Groups	13
2.4. Hyperplanes	14
Chapter 3. Cubical disk diagrams, hyperplanes, and convexity	16
3.1. Disk Diagrams	16
3.2. Properties of Hyperplanes	20
3.3. Local Isometries and Convexity	24
3.4. Background on Quasiconvexity	26
3.5. Cores, Hulls, and Superconvexity	28
Chapter 4. Special Cube Complexes	32
4.1. Hyperplane Definition of Special Cube Complex	32
4.2. Separability Criteria for Virtual Specialness	34
4.3. Canonical Completion and Retraction	36
4.4. Separability in the Hyperbolic Case	37
4.5. Wall-Injectivity and a Fundamental Commutative Diagram	40
4.6. Wall Projection Controls Retraction	41
Chapter 5. Virtual Specialness of Malnormal Amalgams	44
5.1. Specializing Malnormal Amalgams	44
5.2. Proof of the Isomorphic Elevation Lemma	49
Chapter 6. Wallspaces and their Dual Cube Complexes	54
6.1. Wallspaces	54
6.2. The Dual $CAT(0)$ Cube Complex	54
6.3. C is $CAT(0)$	56
6.4. Some Examples	57
6.5. Wallspaces from Codimension-1 Subgroups	59

Chapter 7. Finiteness properties of the dual cube complex	61
7.1. The Cubes of C :	61
7.2. The Bounded Packing Property and Finite Dimensionality:	62
7.3. Cocompactness in the Hyperbolic Case	63
7.4. Relative Cocompactness in the Relatively Hyperbolic Case	63
7.5. Properness of the G Action on $C(\tilde{X})$	65
7.6. The Cut-Wall Criterion for Properness	67
Chapter 8. Cubulating Malnormal Graphs of Cubulated Groups	68
8.1. A Wallspace for an Easy Non-Hyperbolic Group	68
8.2. Extending Walls	70
8.3. Constructing Turns	71
8.4. Cubulating Malnormal Amalgams	72
Chapter 9. Cubical Small Cancellation Theory	76
9.1. Cubical Presentations	77
9.2. The Fundamental Theorem of Small-Cancellation Theory	78
9.3. Combinatorial Gauss-Bonnet Theorem	80
9.4. Greendlinger's Lemma and the Ladder Theorem	81
9.5. Reduced Diagrams	83
9.6. Producing Examples	86
9.7. Rectified Diagrams	87
Chapter 10. Walls in Cubical Small-Cancellation Theory	92
10.1. Walls in Classical $C'(\frac{1}{6})$ Small-Cancellation Complexes	92
10.2. Wallspace Cones	92
10.3. Producing Wallspace Cones	93
10.4. Walls in \tilde{X}^*	94
10.5. Quasiconvexity of Walls in \tilde{X}^*	95
Chapter 11. Annular Diagrams	97
11.1. Classification of Flat Annuli	97
11.2. The Doubly Collared Annulus Theorem	99
11.3. Almost Malnormality	100
Chapter 12. Virtually Special Quotients	102
12.1. The Malnormal Special Quotient Theorem	102
12.2. Case Study: $F_2/\langle\langle W_1^{n_1}, \dots, W_r^{n_r} \rangle\rangle$	104
12.3. The Special Quotient Theorem	108
Chapter 13. Hyperbolicity and Quasiconvexity Detection	110
13.1. Cubical Version of Filling Theorem	110
13.2. Persistence of Quasiconvexity	112
13.3. No Missing Shells and Quasiconvexity	112
Chapter 14. Hyperbolic groups with a quasiconvex hierarchy	115
Chapter 15. The relatively hyperbolic setting	119

Chapter 16. Applications	123
16.1. Baumslag's Conjecture	123
16.2. 3-Manifolds	125
16.3. Limit Groups	126
Bibliography	127
Index of notation and defined terms	131

Acknowledgement

I am enormously grateful to Jason Behrstock for organizing the NSF-CBMS conference together with Abhijit Champanerkar. I am also grateful to the conference participants for their feedback, to the CUNY Graduate Center for hosting the conference, the CBMS for choosing this topic, and the NSF for funding it. My research has been supported by NSERC and undertaken at McGill University.

During my 2008-2009 sabbatical at the Hebrew University I benefited tremendously from the feedback and encouragement of Zlil Sela, with whom I went through the project on *The Structure of Groups with a Quasiconvex Hierarchy* that led to this conference. Valuable criticism and helpful corrections were subsequently provided by many other friends and colleagues, among them Ian Agol, Nicolas Bergeron, Hadi Bigdely, David Gabai, Mark Hagen Jason Manning, Eduardo, Martinez-Pedroza, Daniel Moskovitch, Piotr Przytycki, and Eric Swenson. I am grateful to Martin Bridson for initiating my interest in nonpositively curved square complexes. My daughter Talia helped with the figures and my parents Batya and Michael helped me with crayons and fractions respectively, and have been a source of continuing and enthusiastic support.

It is a great fortune to have wonderful collaborators who spent much time and energy working with me on the ideas presented here: Tim Hsu on linearity, separability in right-angled Artin groups, and cubulating malnormal amalgams; Chris Hruska on the finiteness properties of the dual cube complex, the bounded packing property, and the tower approach to the spelling theorem for one-relator groups with torsion. I've engaged with Michah Sageev's seminal work on CAT(0) cube complexes for the last 10 years and have learned many things working with him. My long-term collaboration with Frédéric Haglund on special cube complexes is at the heart of the matter here, and I am additionally building on his influential earlier ideas.

Finally, I am grateful to my adorable and loving wife Yael for giving me the time to write these notes, and for serving as a partner, a cheerleader, and a role model in all parts of my life.

Note to the reader

The target audience for these lecture notes consists of younger researchers at the beginning of their careers as well as seasoned researchers in neighboring areas interested in quickly acquiring the viewpoint. The requisite background for reading this text should be at the level of an introductory course on geometric group theory or even just hyperbolic groups, though some comfort with graphs of groups would be helpful. I have attempted to include all defined terms and notation in an index at the end of the document, and hope that the intrepid reader will be able to dive into a chapter of interest and work outwards.

CHAPTER 1

Overview

Our goal is to describe a stream of geometric group theory connecting many of the classically considered groups arising in combinatorial group theory with right-angled Artin groups. The nexus here are the “special cube complexes” whose fundamental groups embed naturally in right-angled Artin groups.

Nonpositively curved cube complexes, which Gromov introduced merely as a convenient source of examples [28], have come to take an increasingly central status within parts of geometric group theory – especially among groups with a comparatively small number of relations. Their ubiquity is explained by Sageev’s construction [65] which associates a dual cube complex to a group that has splittings or even ‘semi-splittings’ i.e. codimension-1 subgroups.

Right-angled Artin groups, which at first appear to be a synthetic class of particularly simple groups, have taken their place as a natural target – possibly even a “universal receiver” for groups that are well-behaved and that have good residual properties and many splittings or at least “semi-splittings”.

We begin by reviewing nonpositively curved cube complexes and a disk diagram approach to them – first entertained by Casson (above dimension 2). These disk diagrams are used to understand their hyperplanes and convex subcomplexes. While many of the essential properties of $\text{CAT}(0)$ cube complexes can be explained using the $\text{CAT}(0)$ triangle comparison metric, we have not adopted this viewpoint. It seems that the most important characteristic properties of $\text{CAT}(0)$ cube complexes arise from their hyperplanes, and these are exposed very well through disk diagrams – and the view will serve us further when we take up small-cancellation theory.

Special cube complexes are introduced as cube complexes whose immersed hyperplanes behave in an organized way and avoid various forms of self-intersections. $\text{CAT}(0)$ cube complexes are high-dimensional generalizations of trees, and likewise, from a certain viewpoint, special cube complexes play a role as high-dimensional “generalized graphs”. In particular they allow us to build (finite) covering spaces quite freely, and admit natural virtual retractions onto appropriate “generalized immersed subgraphs” just like ordinary 1-dimensional graphs. The fundamental groups of special cube complexes embed in right-angled Artin groups – because of a local isometry to the cube complex of a naturally associated raag. Since right-angled Artin

groups embed in (and are closely allied with) right-angled Coxeter groups, this means that one can obtain linearity and residual finiteness by verifying virtual specialness.

We describe some criteria for verifying that a nonpositively curved cube complex is virtually special – the most fundamental is the condition that double hyperplane cosets are separable. A deeper criterion [36] arises from a nonpositively curved cube complex that splits along an embedded 2-sided hyperplane into one or two smaller nonpositively curved cube complexes. Under good enough conditions the resulting cube complex is virtually special:

THEOREM 1.1 (Specializing Amalgams). *Let Q be a compact nonpositively curved cube complex with $\pi_1 Q$ hyperbolic. Let P be a 2-sided embedded hyperplane in Q such that $\pi_1 P \subset \pi_1 Q$ is malnormal and each component of $Q - N_o(P)$ is virtually special. Then Q is virtually special.*

A subgroup H of G is “codimension-1” if H splits G into two or more “deep components” – like an infinite cyclic subgroup of a surface group. In his PhD thesis, Sageev understood that when G acts minimally on a CAT(0) cube complex \tilde{X} , the stabilizers of hyperplanes are virtually codimension-1 subgroups of G . He contributed an important converse to this:

CONSTRUCTION 1.2 (dual CAT(0) cube complex). Given a group G and a collection H_1, \dots, H_r of codimension-1 subgroups, one obtains an action of G on a *dual CAT(0) cube complex* – whose hyperplane stabilizers are commensurable with conjugates of the H_i .

We review Construction 1.2 in the context of Haglund-Paulin wallspaces, and describe some results on the finiteness properties of the action of G on the CAT(0) cube complex \tilde{X} . The main point is that if we can produce sufficiently many quasiconvex codimension-1 subgroups in the hyperbolic group G , then we can apply Construction 1.2 to obtain a proper cocompact action of G on a CAT(0) cube complex. This is how we prove the following result [43]:

THEOREM 1.3 (Cubulating Amalgams). *Let G be a hyperbolic group that splits as $A *_C B$ or $A *_C B$ where C is malnormal and quasiconvex. Suppose A, B are each fundamental groups of compact cube complexes. And suppose that some technical conditions hold (and these hold when A, B are virtually special). Then G is the fundamental group of a compact nonpositively curved cube complex.*

A *hierarchy* for a group G is a way to repeatedly build it starting with trivial groups (but sometimes other basic pieces) by repeatedly taking amalgams $A *_C B$ and $A *_C B$ whose vertex groups have shorter length hierarchies. The hierarchy is *quasiconvex* if at each stage the amalgamated subgroup C is a finitely generated that embeds by a quasi-isometric embedding, and similarly, the hierarchy is *malnormal* if C is malnormal in $A *_C B$ or $A *_C B$.

Taken together, Theorem 1.1 and Theorem 1.3 inductively provide the following target for virtual specialness – a malnormal variant of our main result in Theorem 1.7.

THEOREM 1.4 (Malnormal Quasiconvex Hierarchy). *Suppose G has a malnormal quasiconvex hierarchy. Then G is virtually compact special.*

Cubical Small-cancellation Theory: A presentation $\langle a, b, \dots \mid W_1, \dots, W_r \rangle$ is $C'(\frac{1}{n})$ if for any “piece” P (i.e. a subword that occurs in two or more ways among the relators) in a relator W_i we have $|P| < \frac{1}{n}|W_i|$. For $n \geq 6$ the group G of the presentation is hyperbolic and disk diagram methods provide a very explicit understanding of many properties of G .

The presentation above can be reinterpreted as $\langle X \mid Y_1, \dots, Y_r \rangle$ where X is a bouquet of loops and each $Y_i \rightarrow X$ is an immersed circle corresponding to W_i , and the group G of the presentation is then $\pi_1 X / \langle\langle \pi_1 Y_1 \dots \pi_1 Y_r \rangle\rangle$. We generalize this to a setting where X is a nonpositively curved cube complex and each $Y_i \rightarrow X$ is a local isometry. We also offer a notion of $C'(\frac{1}{n})$ small-cancellation theory for such “cubical presentations”. The main results of classical small-cancellation theory – Greendlinger’s lemma and the ladder theorem (and other results involving annular diagrams) have quite explicit generalizations. In particular, we obtain the following result which generalizes the classification of finite trees: T is either a single vertex, is a subdivided arc, or has three or more leaves:

THEOREM 1.5. *If D is a reduced diagram in a cubical $C'(\frac{1}{24})$ presentation then either D is a single 0-cell or cone-cell, or D is a “ladder” consisting of sequence of cone-cells, or D has at least three spurs and/or cornsquares and/or shells.*

The undefined terms in Theorem 1.5 are described in Chapter 9, but the reader might wish to take a glimpse at Figure 9.5.

One motivation for introducing a cubical small-cancellation theory is that when the “relators” Y_i also have given wallspace structures, then there are natural walls – and hence usually codimension-1 subgroups – in the group G , generalizing the same phenomenon for $C'(\frac{1}{6})$ groups.

This cubical small-cancellation theory helps to coordinate the proof of the following result:

THEOREM 1.6 (Malnormal Special Quotient Theorem). *Let G be hyperbolic and virtually compact special. Let $\{H_1, \dots, H_r\}$ be an almost malnormal collection of subgroups. There exist finite index subgroups H'_1, \dots, H'_r such that $G / \langle\langle H'_1, \dots, H'_r \rangle\rangle$ is virtually compact special and hyperbolic.*

Most of our exposition circulates around the proof of Theorem 1.6. Assuming that $G = \pi_1 X$, we first choose a collection of local isometries $Y_i \rightarrow X$ with $\pi_1 Y_i = H_i$. We then choose appropriate finite covers \widehat{Y}_i (the H'_i will be $\pi_1 \widehat{Y}_i$) such that the group \bar{G} of $\langle X \mid \widehat{Y}_1, \dots, \widehat{Y}_r \rangle$ has a finite index subgroup \bar{G}' with a malnormal quasiconvex hierarchy (we have hidden a few steps

here) that can be obtained by cutting along hyperplanes in a finite cover \widehat{X} . Thus \bar{G}' is virtually special by Theorem 1.4.

THEOREM 1.7 (Quasiconvex Hierarchy). *Suppose G is hyperbolic and has a quasiconvex hierarchy. Then G is virtually compact special.*

Proving Theorem 1.7 depends on proving virtual specialness of the amalgamated free products and HNN's that arise at each stage of the hierarchy. Given a splitting, say $G = A *_C D$, the plan is to find a finite index subgroup G' with an almost malnormal quasiconvex hierarchy and conclude by applying Theorem 1.4. To do this, we verify separability of C by applying Theorem 1.6 to quotient subgroups of C using an argument inducting on $\text{Height}(G, C)$. This idea of repeatedly quotienting with an induction on height was independently discovered by Agol-Groves-Manning.

Generalizations of Theorem 1.7 hold in many (and conjecturally all) cases when G is hyperbolic relative to abelian subgroups. We describe how to deduce these generalizations from Theorem 1.7. The proof of separability essentially involves a generalization of Theorem 1.7 to provide virtually special peripheral quotienting. However cubulating requires some additional work.

1.1. Applications

We describe three notable classes of groups with quasiconvex hierarchies in Chapter 16:

Limit groups have hierarchies given by Kharlamovich-Miasnikov and by Sela, and are thus virtually special.

Every one-relator group has a Magnus-Moldavanskii hierarchy – and for one relator groups with torsion this hierarchy is a quasiconvex hierarchy. (Though technically one must pass to a torsion-free finite index subgroup). This resolves Baumslag's conjecture that every one-relator group with torsion is residually finite – indeed they are virtually special and thus linear and have separable quasiconvex subgroups.

For a hyperbolic 3-manifold M with an incompressible surface S the Haken hierarchy of M yields a quasiconvex hierarchy for $\pi_1 M$ provided $\pi_1 S$ is geometrically finite, and so $\pi_1 M$ is virtually special. When the hyperbolic 3-manifold has a geometrically finite incompressible surface, we thus find that $\pi_1 M$ is subgroup separable: Indeed, the geometrically finite subgroups are quasiconvex and hence separable using virtual specialness, and the virtual fiber subgroups are easily seen to be separable, and there are no other subgroups by the Tameness Theorem [1, 13]. A second corollary is that when the hyperbolic manifold M is Haken, in the sense that it has an incompressible surface S , then M is virtually fibered. Indeed, either S is a virtual fiber, or it is geometrically finite, and in the latter case $\pi_1 M$ is virtually in a raag and thus virtually RFRS and so Agol's fibering criterion applies [2].

1.2. A Scheme for Understanding Groups

The above discussions are instances of partial success in implementing the following “grand plan” for understanding many groups:

- Find codimension-1 subgroups in a group G .
- Produce the dual CAT(0) cube complex \tilde{C} upon which G acts.
- Verify that G acts properly and relatively cocompactly on \tilde{C} by examining the extrinsic nature of the codimension-1 subgroups.
- Consequently G is the fundamental group of a nonpositively curved cube complex $C = G \backslash \tilde{C}$. (Or C is an orbihedron if G has torsion.)
- Find a finite covering space \widehat{C} of C , such that \widehat{C} is special.
- The specialness reveals many structural secrets of G . For instance, G is linear since it embeds in $SL_n(\mathbb{Z})$, and the geometrically well-behaved subgroups of G are separable.

In conclusion, in many cases, especially when G has comparatively few relators we see that:

Though G might arise as the fundamental group of a small 2-complex or 3-manifold, in many cases one should sacrifice this small initial presentation in favor of a much larger and higher-dimensional object that is a nonpositively curved special cube complex, and has the advantage of being far more organized, thus revealing important structural aspects of G .

CHAPTER 2

Nonpositively Curved Cube Complexes

This chapter serves as a quick introduction to nonpositively curved cube complexes. We quickly review the basic definitions in Section 2.1. A variety of examples are provided in Section 2.2. The fundamental examples, which arise from right-angled Artin groups, are described in Section 2.3. While Gromov introduced nonpositively curved cube complexes as a source of examples with metric nonpositive curvature in the sense that the thin triangle comparison condition holds, we have adopted a combinatorial viewpoint here. Metric arguments are not usually critical here, and though a traditional geometric group theory attitude can sometimes coordinate or motivate a proof, it seems to make things messier to work simultaneously in two categories. The key feature that gives nonpositively curved cube complexes their characteristic properties are the hyperplanes which we describe in Section 2.4 and will continually revisit in Chapter 3 and subsequently.

2.1. Definitions

An n -cube is a copy of $[-1, 1]^n$. Its *faces* are restrictions of some coordinates to ± 1 . We regard the faces as cubes in their own right.



FIGURE 2.1. Two faces in a 3-cube

A *cube complex* X is a cell complex obtained by gluing cubes together along faces. The identification maps of faces are modeled by isometries – so this is entirely combinatorial.




FIGURE 2.2. A cube complex

The *link* of a 0-cube v of X is the simplex-complex whose n -simplices are corners of $(n + 1)$ -cubes adjacent with v . So $\text{link}(v)$ is the “ ϵ -sphere” about v in X .




FIGURE 2.3. Link within a cube complex

A *flag complex* is a simplicial complex with the property that $n + 1$ vertices span an n -simplex if and only if they are pairwise adjacent. Thus a flag complex is determined completely by a simplicial graph. Note that a graph Γ is flag if and only if $\text{girth}(\Gamma) \geq 4$.




FIGURE 2.4. A 2-dimensional and 1-dimensional flag complex

The cube complex X is *nonpositively curved* if $\text{link}(v)$ is a flag complex for each $v \in X^0$. We say \tilde{X} is a *CAT(0) cube complex* if it is nonpositively curved and simply-connected.

REMARK 2.1. A CAT(0) cube complex also has a genuine CAT(0) triangle comparison metric where each n -cube is isometric to $[-1, 1]^n$ however this is not usually the best viewpoint here.

We use the notation \tilde{X} since a CAT(0) cube complex \tilde{X} arises as the universal cover of a nonpositively curved cube complex X .

EXAMPLE 2.2 (Not nonpositively curved). The cube complex homeomorphic to an n -sphere obtained by identifying two n -cubes identified along their boundaries is not nonpositively curved for $n \geq 2$. Indeed, the link of each 0-cube is not simplicial.

The cube complex obtained by removing one of the eight 3-cubes around the origin in \mathbb{R}^3 is not nonpositively curved. The link of the central 0-cube is isomorphic to the simplicial complex obtained from an octahedron by removing a single open 2-simplex.

2.2. Some Favorite 2-Dimensional Examples

EXAMPLE 2.3 (Dehn Complex). A link projection is *alternating* if the curves travel alternately above and below at crossings. The projection is

prime if each embedded circle σ in the plane that intersects the projection P transversely in exactly two noncrossing points has the property that the part of P on either the inside or on the outside of σ consists of a single arc. A projection that is not prime is illustrated on the right of Figure 2.5. (Any link that is both prime and alternating in the usual sense has a projection that is both prime and alternating.)

Let L be a link in S^3 . The *Dehn complex* X of L is a square complex that embeds in $S^3 - L$. The 2-complex X has exactly two 0-cells v_b, v_t which are positioned at the “bottom” and at the “top” of the projection plane. We give the projection the checkerboard coloring. There is a 1-cell for each region of the projection – the 1-cells associated with black regions are oriented from v_t to v_b and the 1-cells associated with white regions are oriented from v_b to v_t . There is a 2-cell for each crossing of P , it is a square corresponding to a path that travels up and down around the crossing (following the boundary of a saddle).




FIGURE 2.5. An embedded Dehn complex




FIGURE 2.6. Squares correspond to crossings

Weinbaum discovered that the link projection is prime and alternating if and only if the Dehn complex X is nonpositively curved, but he formulated this in terms of $C(4)$ - $T(4)$ small-cancellation complexes [72, 52]. The explanation we give here is from [81].

We first observe that $\text{link}(v_t)$ embeds in the projection diagram (with the traditional omitted parts indicating the nature of the crossing points).

To see that X is nonpositively curved exactly when P is prime and alternating, we note that the checkerboard coloring shows that $\text{link}(v_t)$ is bipartite so it suffices to verify that there are no 2-cycles. But the two




FIGURE 2.7. $\text{link}(v_t)$ embeds in the projection diagram

different types of 2-cycles would show that P is either not alternating or not prime.




FIGURE 2.8. 2-cycles that indicate that P is not alternating or not prime.

EXAMPLE 2.4 (From graphs of graphs). Let X be a topological space that decomposes as a graph Γ of spaces where each vertex space X_v is a graph and each edge space $X_e \times [-1, 1]$ is the product of a graph and an interval. Suppose the attaching maps $\phi_{e^-} : X_e \times \{-1\} \rightarrow X_{i(e)}$ and $\phi_{e^+} : X_e \times \{+1\} \rightarrow X_{\tau(e)}$ are combinatorial immersions. Then X is a nonpositively curved square complex.




FIGURE 2.9. A graph of spaces with two vertex spaces and two edge spaces.

EXAMPLE 2.5 (Amalgam along \mathbb{Z}). A group of the form $F_2 *_{\mathbb{Z}} F_2$ always arises as a graph of groups which is the fundamental group of a graph of

spaces yielding a nonpositively curved square complex as above. To do this we can subdivide to make $|U| = |V|$ where U and V are words generating \mathbb{Z} in the two free groups. This subdivision rarely works for an HNN extension $F_2 *_{\mathbb{Z}^t = \mathbb{Z}}$. Understanding how to deal with such HNN examples was one of the motivations for this research.



FIGURE 2.10. Subdivide the left and right vertex spaces so that the cylinder edge space has bounding circles of the same length as on the right.

A \mathcal{VH} -complex is a square complex whose 1-cells are divided into two classes: *vertical* and *horizontal*, and where attaching maps of 2-cells are length 4 paths that alternate between vertical and horizontal edges.




FIGURE 2.11. A \mathcal{VH} -complex

When X arises as a graph of spaces as in Example 2.4, then $\text{link}(v)$ is a bipartite graph, and X is a nonpositively curved \mathcal{VH} -complex. The vertical 1-cells are in the vertex spaces, and the horizontal 1-cells project to edges of the underlying graph. Note that the nonpositive curvature holds because the attaching maps are immersions. We now draw attention to two classes of graphs of spaces that arise from restrictions on the nature of the attaching maps of edges spaces:

EXAMPLE 2.6 (Complete Square Complexes). When all attaching maps $\phi_{e_{\pm}}$ are covering maps of graphs, then X is a *complete square complex* which means that each $\text{link}(v)$ is a complete bipartite graph. In this case, the universal cover \tilde{X} is isomorphic to the cartesian product of two trees. These can be surprisingly complicated and we refer to [82], and to [47] for some simple such examples. Most notably, Burger-Mozes gave such examples where $\pi_1 X$ is infinite simple [12].

EXAMPLE 2.7 (Clean \mathcal{VH} -complexes). When all attaching maps $\phi_{e_{\pm}}$ are combinatorial embeddings, then X is *clean*. Sometimes X might not be clean but has a finite covering space that is clean.




FIGURE 2.12. A complete square complex with 6 squares.




FIGURE 2.13. A \mathcal{VH} -complex for a genus 2 surface, together with a degree 4 clean cover

REMARK 2.8. When X is a compact nonpositively curved \mathcal{VH} -complex, then X has a (single or) double cover \widehat{X} that splits as a graph of spaces as in Example 2.9, and so $\pi_1 \widehat{X}$ splits as a corresponding graph of groups. In this case X has a clean finite cover precisely when the edge groups of $\pi_1 \widehat{X}$ are separable [81].

2.3. Right-Angled Artin Groups

The *right-angled Artin group (raag)* or *graph group* $G(\Gamma)$ associated to the simplicial graph Γ has the following presentation:

$$(2.1) \quad \langle g_v : v \in \text{Vertices}(\Gamma) \mid [g_u, g_v] : (u, v) \in \text{Edges}(\Gamma) \rangle$$

EXAMPLE 2.9.

- (1) $G(\triangle) \cong \mathbb{Z}^3$
- (2) $G(\cdot \cdot \cdot) \cong F_3$
- (3) $G(\begin{smallmatrix} \times & \times \\ \times & \times \end{smallmatrix}) \cong F_2 \times F_3$
- (4) $G(\begin{smallmatrix} \rightarrow & \rightarrow \\ \rightarrow & \rightarrow \end{smallmatrix}) \cong \pi_1(S^3 - \infty)$

When $\text{girth}(\Gamma) \geq 4$, the standard 2-complex R of Presentation (2.1) is already a nonpositively curved square complex. In general, we must also add higher dimensional cubes. We define $R(\Gamma)$ to be the cube complex obtained from the standard 2-complex of Presentation (2.1) by adding an n -cube for

each collection of n pairwise commuting generators – i.e. for each n -clique in Γ . $R(\Gamma)$ is often called the *Salvetti complex* [16].

We record the above construction as follows:

THEOREM 2.10. *For each [finite] simplicial graph Γ , there is a [compact] nonpositively curved cube complex $R(\Gamma)$ such that the associated raag $G(\Gamma)$ can be identified with $\pi_1 R(\Gamma)$. In particular, the standard 2-complex of the defining presentation for $G(\Gamma)$ equals the 2-skeleton of $R(\Gamma)$.*

EXAMPLE 2.11. A 3-dimensional example that is worth thinking through to illustrate the construction proving Theorem 2.10 is indicated in Figure 2.14. Notice that $R(\Gamma)$ has only one 0-cell v , and that $\text{link}(v)$ contains two copies of Γ : one “ascending” and one “descending” and has additional simplices corresponding to corners of cubes that are mixed.




FIGURE 2.14. Γ is on the left, and $\text{link}(v)$ is on the right. Note that $\text{link}(v)$ contains an octahedron whose eight 2-simplices correspond to the eight corners of the added 3-cube.

PROPOSITION 2.12. *Raags have the following properties:*

- (1) *They are residually torsion-free nilpotent [22].*
- (2) *They are residually finite rational solvable [2].*
- (3) *They are linear [46].*
- (4) *They embed in right-angled Coxeter groups and hence in $SL_n(\mathbb{Z})$ [44, 20].*

We refer to Charney’s survey paper for more about right-angled and other Artin groups [15].

2.4. Hyperplanes

A *midcube* is a subspace of a cube obtained by restricting one coordinate to 0. A *hyperplane* is a connected subspace of a CAT(0) cube complex that intersects each cube in a single midcube or in \emptyset . The hyperplane H is said to be *dual* to each 1-cube that it intersects.

We record some fundamental properties of hyperplanes given in [65]. We will describe proofs in Chapter 3.

THEOREM 2.13. *Let \tilde{X} be a CAT(0) cube complex*

- (1) *Each midcube lies in a unique hyperplane of \tilde{X} .*




FIGURE 2.15. Midcubes in a 1-cube, 2-cube, and 3-cube




FIGURE 2.16. Some hyperplanes in a 3-dimensional CAT(0) cube complex, and in the plane, and in a tree.

- (2) A hyperplane H of \tilde{X} is itself a CAT(0) cube complex (regard midcubes as cubes).
- (3) The cubical neighborhood $N(H) \cong H \times [-1, 1]$ is a convex subcomplex of \tilde{X} called the carrier of H .
- (4) $\tilde{X} - H$ consists of two components.

REMARK 2.14. We usually use the L^1 metric on a CAT(0) cube complex \tilde{X} , where distance equals the length of the shortest path that is piecewise parallel to axes. The inclusion $\tilde{X}^1 \subset \tilde{X}$ is then an isometric embedding, where we use the graph metric on \tilde{X}^1 . We note that $d(p, q) = \#(p, q)$ for $p, q \in \tilde{X}^0$, where $\#(p, q)$ denotes the number of hyperplanes separating p and q .

A geodesic in \tilde{X} is a combinatorial edge path in the 1-skeleton of \tilde{X} that is a geodesic in \tilde{X}^1 with respect to the graph metric. A subcomplex \tilde{Y} of \tilde{X} is *convex* if for any geodesic γ in \tilde{X} , if the endpoints of γ are in \tilde{Y} then $\gamma \subset \tilde{Y}$.

CHAPTER 3

Cubical disk diagrams, hyperplanes, and convexity

The goal of this chapter is to acquaint the reader with disk diagrams in nonpositively curved cube complexes. Disk diagrams have long been understood to have widespread utility in combinatorial and geometric group theory, and especially within small-cancellation theory. Casson introduced the viewpoint of using disk diagrams for compact nonpositively curved cube complexes in lectures in the 1980's, where he used "bigon removal" to give a very simple elegant proof of a quadratic isoperimetric function. I am grateful to Yoav Moriah for sharing his notes on this topic with me, and to Michah Sageev for sending them to me and encouraging me to examine them in response to my paper [80]. The simple idea of "bigon removal" opened my eyes to the possibility of developing the cubical small-cancellation theory we discuss later. Sageev independently developed and utilized the disk diagram viewpoint to prove many of the fundamental properties of hyperplanes in [65]. Our own treatment of this is not meant to be efficient, but rather to provide the reader with a level of comfort and experience as we shall routinely adopt aspects and arguments arising for cubical disk diagrams later when we study diagrams in cubical presentations in Chapter 9.1.

We recall the definition of a disk diagram in Section 3.1. We use disk diagrams to prove the fundamental properties of hyperplanes in Section 3.2. We define local isometries between nonpositively curved cube complexes in Section 3.3. We review background and some useful lemmas concerning quasiconvexity in Section 3.4. Finally, in Section 3.5 we describe the H -cocompact convex subcomplex that arises for a quasiconvex subgroup H of a hyperbolic group G acting properly and cocompactly on a CAT(0) cube complex \tilde{X} .

3.1. Disk Diagrams

A *disk diagram* D is a compact contractible combinatorial 2-complex with a chosen planar embedding $D \subset \mathbb{R}^2$. Its *boundary path* or *boundary cycle* $\partial_p D$ is the attaching map of the 2-cell containing the point at ∞ (regarding $S^2 = \mathbb{R}^2 \cup \infty$).

A *diagram in a complex* X is a combinatorial map $D \rightarrow X$.

LEMMA 3.1 (van Kampen). *A closed combinatorial path $P \rightarrow X$ is null-homotopic if and only if there exists a diagram $D \rightarrow X$ with $P \cong \partial_p D$ so that*



FIGURE 3.1. Disk diagrams and their boundary paths

there is a commutative diagram:

$$\begin{array}{ccc} \partial_p D & \rightarrow & D \\ \parallel & & \downarrow \\ P & \rightarrow & X \end{array}$$

When X is a cube complex, D is a square complex and we study the *dual curves* in D which are defined like hyperplanes. Each dual curve is an immersed circle or an immersed interval that starts and ends on distinct 1-cells of $\partial_p D$.




FIGURE 3.2. Dual curves in a square disk diagram

When X is nonpositively curved and $D \rightarrow X$ has minimal area among all those disk diagrams in X with the same boundary path, the behavior of dual curves is quite controlled. The main idea behind this control is summarized by the following result due to Casson whose viewpoint we have adopted here.

THEOREM 3.2 (Bigon removal). *Let $D \rightarrow X$ be a disk diagram with X a nonpositively curved cube complex. Any bigon B of dual curves can be replaced (through a homotopy) by a diagram B' with fewer squares, and the dual curve pairing of edges of $\partial_p B$ and $\partial_p B'$ remains the same. See Figure 3.3.*

Since whenever there is a dual curve that passes through two 1-cells that are adjacent to the same vertex there is then a bigon cutting through it as in Figure 3.4, it follows that:

- (1) Each dual curve embeds – i.e. it cannot self-cross.
- (2) A minimal area diagram D for $P \rightarrow X$ has no nonogons, oscugons, or bigons (see Figure 3.5).

SKETCH OF THEOREM 3.2. Choose a smallest area bigon within D . We then choose a lowest triangle of dual curves with at least one side on the bigon. Its existence is sketched in Figure 3.9. There are two cases depending




FIGURE 3.3. Bigonal subdiagrams can be replaced by subdiagrams with two fewer squares since $\text{Area}(B') \leq \text{Area}(B) - 2$.



FIGURE 3.4. A dual curve that passes through adjacent edges has another dual curve cutting through it to form a bigon.




FIGURE 3.5. A dual curve cannot self-cross as on the left in any diagram, moreover, minimal area diagrams cannot contain nonogons, oscugons, or bigons.

on whether or not one of the vertices of this triangle is also a vertex of the bigon. These two possibilities are indicated in Figure 3.6. In each case we are able to homotope the diagram to a new diagram using a *hexagon move* illustrated in Figure 3.8, which pushes the front three squares in a cube to become the back three squares.




FIGURE 3.6.

Repeating these hexagon moves we obtain a smaller and smaller area bigon until we arrive at the base case illustrated in Figure 3.8 when two squares are actually replaced by a length 2 path. This *anceled pair* of

squares plays a role similar to a “cancelable pair” in small-cancellation theory. □




FIGURE 3.7. Hexagon move




FIGURE 3.8. Canceled pair of squares




FIGURE 3.9. Consider the set of vertices in the graph of dual curves internal to a bigon. If there are no vertices then either there are no dual curves and we can cancel a pair of squares, or there are some dual curves and we can perform a leftward or rightward hexagon move at a corner. Otherwise, there is a partial ordering on the vertices where $u < v$ if the triangle with u on top contains v on one of its legs. A lowest vertex in this partial ordering gives a triangle corresponding to a hexagon move.




FIGURE 3.10. Complete example of bigon removal

3.2. Properties of Hyperplanes

The first part of the following statement shows that hyperplanes exist, as any midcube extends uniquely to an “immersed hyperplane” and it does not self-cross, and hence its only nonempty intersection with a cube is in a single midcube.

COROLLARY 3.3. *Let H be an immersed hyperplane in a $CAT(0)$ cube complex \tilde{X} . Then*

- (1) H does not self-cross.
- (2) H does not self-osculate.
- (3) H is simply-connected.

SKETCH. In each case a smallest area counterexample leads to an even smaller one by considering a bigon that must reside inside and applying Theorem 3.2. \square



FIGURE 3.11.

COROLLARY 3.4. *A hyperplane H of the $CAT(0)$ cube complex \tilde{X} is 2-sided, so $N(H) \cong H \times [-1, 1]$.*

Note that Corollary 3.4 follows from Corollary 3.3.(3), since an I -bundle over the simply-connected space H must be trivial. Another way of thinking about this is to observe that if H is 1-sided, then one could produce a connected double cover of X , contradicting that X is simply-connected.

To illustrate the principle, we give a diagrammatic proof

PROOF. If H is 1-sided, then there is a path $\gamma \rightarrow N(H)$ that starts and ends on opposite 0-cubes of a 1-cube e dual to H but does not traverse any 1-cube dual to H . Consider a disk diagram for the closed path γe . The dual curve λ in D that starts at the center of e must end at the center of some edge e' in γ which is impossible since λ maps to H , so e' is an edge of γ that is dual to H . \square

COROLLARY 3.5. *$\tilde{X} - H$ consists of exactly two components.*

Corollary 3.5 follows from a Meyer-Vietoris sequence using that \tilde{X} is simply-connected and \tilde{H} is connected. However, we again give a diagrammatic proof similar to the proof of Corollary 3.4.

PROOF. Note that $N(H) = H \times I$ so $\tilde{X} - H$ contains at most two components. If $\tilde{X} - H$ contained one component, then there would be a path

$\gamma \rightarrow (\tilde{X} - H)$ that starts and ends on opposite sides of a 1-cell e dual to H . Indeed, there is a path γ' that is disjoint from H and starts and ends on 0-cubes in opposite components of $N(H) - H$. And γ' can be extended to a path γ that starts and ends on opposite 0-cubes on an edge e dual to H . Here we use that H is path connected to extend γ' parallel to H in $N(H) \cong H \times I$, and since H does not self-cross we see that γ does not pass through any edge dual to H . Let D be a disk diagram for γe . Then the dual curve λ starting at e must end on a 1-cell in γ which is impossible. \square




FIGURE 3.12. The hyperplane H is 2-sided. And $\tilde{X} - H$ consists of two components

LEMMA 3.6 (No Inter-osculations: “crossing pair has a square”). *Let U, V be hyperplanes in \tilde{X} . If U, V cross in some square s – i.e. are dual to distinct 1-cells on ∂s , and U, V are dual to 1-cells e, f meeting at a vertex v , then e, f lie on $\partial s'$ for some square s' .*



FIGURE 3.13. Crossing pair has a square. Hyperplanes cannot interosculate in \tilde{X}

PROOF. Consider a diagram D surrounded with two dual curves on the top and bottom that cross on the right and osculate on the left. Note that such a diagram can always be built in the situation in the hypothesis of the Lemma, since we can take two crossing ladders that start at s and end at the vertex v , and then fill this in with some subdiagram.

Starting with the diagram on the left in Figure 3.14, we now proceed exactly as in Theorem 3.2 except that we will now only use hexagon moves pushing upwards downwards and leftwards until we get the diagram in the middle of Figure 3.14. We then use hexagon moves to push the square s leftwards to reveal the desired square s' . \square




FIGURE 3.14. To see that the crossing hyperplanes must cross again when they come close, we first push squares out of the bounded region in a diagram, and then slide the right-square over to the left.

DEFINITION 3.7 (Cornsquare). A *cornsquare* in a disk diagram D is a square s with a pair of dual curves emanating from s and ending on adjacent edges a, b in $\partial_p D$. We say that ab is the *outerpath* of the cornsquare s . See Figure 3.15.




FIGURE 3.15. Three cornsquares in a disk diagram D

Lemma 3.6 can be arranged to play a critical role because of the following:

REMARK 3.8 (Finding cornsquares). Let D be a disk diagram between a path σ and a path γ . Let a, b be dual curves that emanate from distinct edges of σ . If either $a = b$ or if a, b cross, then there exist dual curves a', b' emanating from distinct adjacent edges of σ (inclusively between a, b) such that a', b' cross in D . See Figure 3.16. The case where $a = b$ follows from the crossing case, since an innermost such situation is a crossing. The statement follows by considering an innermost pair of edges whose dual curves cross.




FIGURE 3.16. Finding a cornsquare

LEMMA 3.9. $\gamma \rightarrow \tilde{X}$ is a geodesic \Leftrightarrow the edges of γ are dual to distinct hyperplanes of \tilde{X} .

PROOF. (\Leftarrow) holds since any other path γ' with the same endpoints must travel through edges dual to these hyperplanes (and possibly some others) since hyperplanes separate \tilde{X} .




FIGURE 3.17. The hyperplanes dual to edges of a geodesic γ must be crossed by any other path γ' with the same endpoints.

(\Rightarrow) holds because if γ passed through two edges e_1, e_2 dual to the same hyperplane H then we could shorten its length by 2 contradicting that it is a geodesic. Indeed, first note that $e_1 \neq e_2$. Now consider a minimal area disk diagram D between the subpath of γ joining e_1, e_2 , and a path $\sigma \rightarrow N(H)$ with the same endpoints. By Remark 3.8, there is a cornsquare in D with outerpath on σ . Following Figure 3.14, we can use hexagon moves to push the square so that its corner is directly on σ , and then since σ is a path on $N(H)$ the square also lies in $N(H)$, and so we can push σ passed this square to obtain a shorter path $\sigma' \rightarrow N(H)$ with a smaller area diagram D' between σ' and the subpath of γ . \square




FIGURE 3.18. A geodesic γ cannot cross a hyperplane twice as there would be a shortcut

FIGURE 3.19. If γ passes through a hyperplane H twice, then we consider a minimal area diagram D between γ and a path in $N(H)$. A cornsquare in D shows that there was a smaller area choice.

We will revisit the method of this last proof as it bring us to discuss the idea that hyperplane carriers are convex.

3.3. Local Isometries and Convexity

DEFINITION 3.10 (Local isometry). A combinatorial map $\phi : Y \rightarrow X$ of nonpositively curved cube complexes is a *local isometry* if for each $y \in Y^0$, the map $\text{link}(y) \rightarrow \text{link}(\phi(y))$ is injective and moreover $\text{link}(y) \subset \text{link}(\phi(y))$ embeds as a *full subcomplex* in the sense that if u, v are vertices of $\text{link}(y)$ that map to adjacent vertices of $\text{link}(\phi(y))$ then u, v are already adjacent in $\text{link}(y)$. (The flag complex condition then implies that $(n + 1)$ vertices of $\text{link}(y)$ span an n -simplex iff their images in $\text{link}(\phi(y))$ span an n -simplex.

A more concrete way to define a local isometry is that $Y \rightarrow X$ is locally-injective and has no *missing corners of squares* in the sense that if two 1-cubes e, f at a 0-cube y map to edges $\phi(e), \phi(f)$ which bound the corner of a square at $\phi(y)$ then e, f already bound the corner of a square at y . (Careful: we really mean “ends of edges” here.)




FIGURE 3.20. The above map from an annulus to a surface is a local isometry that is not an embedding.




FIGURE 3.21. The above map is not a local isometry from an annulus into a cylinder. The failure is at the bold vertex.




FIGURE 3.22. Local isometries of 1-torus and 2-torus into a 3-torus and the respective embeddings of links as full sub-complexes

EXAMPLE 3.11. Any immersion of graphs is a local isometry. Any covering map is a local isometry. A fundamental example of a local isometry is a map $\bar{N} \rightarrow X$ where $\bar{N} = N(\bar{H})$ is the carrier of an immersed hyperplane $\bar{H} \rightarrow X$ in a nonpositively curved cube complex X . See Definition 4.1.

The following lemma was first noted in [58] in terms of the (noncombinatorial) CAT(0) metric popularized by Gromov:

LEMMA 3.12 (Local isometry \Rightarrow Convex embedding). *If $\phi : Y \rightarrow X$ is a local isometry then $\tilde{\phi} : \tilde{Y} \rightarrow \tilde{X}$ is an embedding as a convex subcomplex (hence in particular an isometric embedding).*




FIGURE 3.23. Considering a minimal diagram between a path $\tilde{\gamma}$ and a local isometry $\tilde{Y} \rightarrow \tilde{X}$ shows that $\tilde{\gamma}$ must lift to \tilde{Y} if $\tilde{\gamma}$ is a geodesic.

PROOF. Consider a geodesic $\tilde{\gamma} \rightarrow \tilde{X}$ that is homotopic in \tilde{X} to a path $\tilde{\sigma} \rightarrow \tilde{Y}$. Let $D \rightarrow \tilde{X}$ be a disk diagram for $\tilde{\gamma}\tilde{\sigma}^{-1}$, and suppose that D has minimal area among all possible such choices with $\tilde{\sigma} \rightarrow \tilde{Y}$ allowed to vary among paths with the same endpoints as $\tilde{\gamma}$.

Suppose D is not a subdivided interval (in which case $\tilde{\gamma} = \tilde{\sigma}$ lies in \tilde{Y} as desired).

Consider the dual curve emanating from the first internal edge of $\tilde{\gamma}$.

The dual curve cannot terminate at another edge of $\tilde{\gamma}$ as in the first diagram of Figure 3.23 since then γ would not be a geodesic. Indeed there are no bigons within D so the path γ' that tracks along this dual curve without crossing it is shorter than γ .

The second situation of Figure 3.23 is impossible by the minimal area of D .

We are thus in the third situation of Figure 3.23, and so there is a cornsquare as in Remark 3.8. But then as in Figure 3.19, we can push this cornsquare towards \tilde{Y} and pass it through, leading to a new choice of $\tilde{\sigma}$ and having reduced the area of the disk diagram by 1. This is impossible. \square

COROLLARY 3.13 (Carriers are convex). *Let H be a hyperplane in \tilde{X} and let $N(H)$ be its carrier. Then $N(H)$ is a convex subcomplex of \tilde{X} .*

Let H be a hyperplane in \tilde{X} , let \overleftarrow{H} and \overrightarrow{H} be the closures of the two components of $\tilde{X} - H$, and let $N(\overleftarrow{H})$ and $N(\overrightarrow{H})$ denote the smallest subcomplexes of \tilde{X} containing these subspaces. These are the two *halfspace carriers* associated with H . It follows from Corollary 3.13 that:

COROLLARY 3.14. *Halfspace carriers are convex.*

Corollary 3.14 is useful in Section 3.5 since the intersection of convex subcomplexes of \tilde{X} is again convex.

3.4. Background on Quasiconvexity

In this section we quickly recall some basic definitions and results about quasiconvex subgroups.

3.4.1. Quasiconvexity.

DEFINITION 3.15 (Quasiconvex). Let G be a group with a finite set of generators $S = \{s_1, \dots, s_r\}$, and let $\Gamma = \Gamma(G, S)$ be its Cayley graph. A subgroup H of G is *quasiconvex* if there exists $\mu > 0$ such that the following holds for any geodesic γ in Γ : If the endpoints of γ lie in H then $\gamma \subset \mathcal{N}_\mu(H)$. We use the notation $\mathcal{N}_\mu(H)$ for the closed μ -neighborhood of a subspace H .

In general, the quasiconvexity of H depends upon the choice of generators determining the Cayley graph as can be seen by considering cyclic subgroups of \mathbb{Z}^2 . However when G is hyperbolic, being quasiconvex is independent of the set of generators, although the constant κ may vary. Both finite and finite-index subgroups are always quasiconvex. A quasiconvex subgroup is finitely generated and the intersection of quasiconvex subgroups is quasiconvex. When G is hyperbolic, being quasiconvex is equivalent to being finitely generated and quasi-isometrically embedded. We refer the reader to [70].

When G is hyperbolic relative to subgroups $\{P_1, \dots, P_r\}$ there is also a notion of *relatively quasiconvex*. We refer the reader to [61] and to [39, 53] for a discussion of the definitions. The basic idea is that H is relatively quasiconvex if geodesics in a Cayley graph lie in a μ -neighborhood of the union of H with the parabolic subgroups that have nontrivial intersection with H . For the case of greatest interest, when the parabolic subgroups P_i are virtually f.g. abelian, being relatively quasiconvex is equivalent to being quasi-isometrically embedded [39]. We will often use the simple term “quasiconvex” instead of “relatively quasiconvex”.

3.4.2. Malnormality.

DEFINITION 3.16 (Malnormal Collection). A collection of subgroups $\{H_1, \dots, H_r\}$ of G is *malnormal* provided that $H_i^g \cap H_j = \{1_G\}$ unless $i = j$ and $g \in H_i$. Similarly, the collection is *almost malnormal* if intersections of nontrivial conjugates are finite.

Here we use the notation $H^g = gHg^{-1}$.

DEFINITION 3.17 (Height). Consider a collection $\{H_1, \dots, H_k\}$ of subgroups of G . We say $H_{m_i}^{g_i}$ and $H_{m_j}^{g_j}$ are *distinct conjugates* unless $m_i = m_j$ and $H_{m_i}g_i = H_{m_j}g_j$. We emphasize that each “conjugate” corresponds to a value of $(m_i, H_{m_i}g_i)$ and not just a subgroup $H_{m_i}^{g_i}$.

The collection has *height* $0 \leq h \leq \infty$ if h is the largest number so that there are h distinct conjugates of these subgroups whose intersection $H_{m_1}^{g_1} \cap H_{m_2}^{g_2} \cap \dots \cap H_{m_h}^{g_h}$ is infinite. If each H_i is finite, then the height of the collection is $h = 0$. If there is an infinite collection of distinct conjugates whose intersection is an infinite subgroup, then the height of the collection is $h = \infty$.

For instance, if G is infinite and $[G : H] < \infty$ then the height of H in G equals $[G : H]$. A collection of subgroups is almost malnormal as in Definition 3.16 precisely if its height is ≤ 1 .

The notion of height was introduced and studied in [27] where it was shown that quasiconvex subgroups of word-hyperbolic groups have finite height. It follows that a finite collection of quasiconvex subgroups also has finite height. We have explored the notion a bit further for relatively hyperbolic groups in [42]. Let us record this result of Gitik-Mitra-Rips-Sageev as follows:

LEMMA 3.18 (Quasiconvex \Rightarrow Finite Height). *Let $\{H_1, \dots, H_r\}$ be a collection of quasiconvex subgroups of the word-hyperbolic group G . Then $\{H_1, \dots, H_r\}$ has finite height in G .*

Finally, we record several observations about separability and intersecting conjugates. The reader can postpone this until it is used in Chapter 14. An element $g \in G$ is an *intersecting conjugator* of the subgroup C provided that $|C^g \cap C| = \infty$. Note that if g is an intersecting conjugator of C then C, gC must have infinite coarse intersection in the Cayley graph, and so gC has a bounded length representative (which is necessarily also an intersecting conjugator). Consequently:

LEMMA 3.19. *Let C be a quasiconvex subgroup of the hyperbolic group G . Then C has finitely many cosets $g_i C$ of intersecting conjugators.*

Separating C from the finitely many intersecting conjugators supplied by Lemma 3.19 we have the following:

LEMMA 3.20. *If C is separable and quasiconvex in the hyperbolic group G then there is a finite index subgroup G' such that C is malnormal in G' .*

Let G be hyperbolic relative to $\{P_1, \dots, P_r\}$. A subgroup H of G is *relatively malnormal* if $H^g \cap H$ is parabolic or finite for all $g \notin H$. Lemma 3.20 was generalized in [42] to:

LEMMA 3.21. *If C is separable and relatively quasiconvex in the hyperbolic group G then there is a finite index subgroup G' such that C is relatively malnormal in G' .*

3.4.3. Almost malnormal collections of commensurators. The reader can postpone this material until it is used in Section 12.3 and Chapter 14.

The groups H_1, H_2 are *abstractly commensurable* if they contain isomorphic finite index subgroups. We say the subgroups H_1, H_2 of G are *commensurable* if $H_1 \cap H_2$ has finite index in both H_1 and H_2 . This leads us to the notion of the *commensurator* $\mathbb{C}_G(H)$ of H in G which is defined by:

$$\mathbb{C}_G(H) = \{g \in G : [H : H^g \cap H] < \infty\}.$$

It is shown in [49] (see also [4]) that $[\mathbb{C}_G(H) : H] < \infty$ for any infinite quasiconvex subgroup H of the word-hyperbolic group G . Consequently, in this case $\mathbb{C}_G(H)$ is itself quasiconvex in G .

LEMMA 3.22. *Let $\{H_1, \dots, H_r\}$ be a collection of quasiconvex subgroups of a word-hyperbolic group G .*

A subgroup K of G is a “maximal-infinite intersection of conjugates” of $\{H_1, \dots, H_r\}$ if the intersection of K with any further conjugate is finite.

Let K_1, \dots, K_s be representatives of the finitely many distinct conjugacy classes of maximal-infinite intersection of conjugates of $\{H_1, \dots, H_r\}$. Then $\{\mathbb{C}_G(K_1), \dots, \mathbb{C}_G(K_s)\}$ is an almost malnormal collection of subgroups of G .

Note that this is indeed a finite collection by Lemma 3.18. To see that it is necessary to pass to commensurators, consider the example of a finite index subgroup such as: $n\mathbb{Z} \subset \mathbb{Z}$.

3.5. Cores, Hulls, and Superconvexity

Thinking of free groups as fundamental groups of graphs is helpful for understanding their subgroups. In particular, each finitely generated subgroup is naturally the fundamental group of an immersed finite graph [71]. We now sketch how this generalizes in higher dimensions. We refer to [33, 67] for more details.

THEOREM 3.23. *Let \tilde{X} be a CAT(0) cube complex that is locally finite and δ -hyperbolic (thus finite dimensional). Let H act on \tilde{X} with a quasiconvex orbit $H\tilde{x}$. For each compact set $K \subset \tilde{X}$ there is an H -cocompact convex subcomplex \tilde{Y} such that $HK \subset \tilde{Y}$.*

The statement of Theorem 3.23 covers the case where H is a quasiconvex subgroup of a hyperbolic group G that acts properly and cocompactly on \tilde{X} . The subcomplex \tilde{Y} is the intersection $\text{Hull}(HK)$ of the carriers of closed halfspaces containing HK and is depicted heuristically in Figure 3.24.

COROLLARY 3.24. *Let X be a compact nonpositively curved based cube complex with $G = \pi_1 X$ hyperbolic. Then for each quasiconvex subgroup H of G there exists a compact nonpositively curved cube complex Y and based local isometry $Y \rightarrow X$ such that:*

$$\begin{array}{ccc} H & \subset & G \\ \parallel & & \parallel \\ \pi_1 Y & \rightarrow & \pi_1 X \end{array}$$




FIGURE 3.24. The intersection $\tilde{Y} = \text{Hull}(HK)$ of closed half-spaces containing the orbit HK of the compact subcomplex K is an H -cocompact subcomplex of \tilde{X} .

PROOF. Choose \tilde{Y} using Theorem 3.23 and let $Y = H \backslash \tilde{Y}$. □

We refer to the local isometry $Y \rightarrow X$ as a *core* for H , and note that when \tilde{X} is the based cover of X with $H = \pi_1 \tilde{X}$, the space Y embeds as a compact subcomplex of \tilde{X} that is a deformation retract of \tilde{X} , just as in the simple scenario of the core of a cover of a graph.

REMARK 3.25. There is also a relatively hyperbolic version of Corollary 3.24 providing a “sparse” Y when G is hyperbolic relative to abelian subgroups and X is compact or sparse. (See the end of Section 7.4 for the notion of sparse.)

LEMMA 3.26. *There exists θ and D such that any metric-geodesic γ of length D crosses a hyperplane U of \tilde{X} with angle $\geq \theta$.*

We note that this proof uses the metric-geodesic and angles definable in the (noncombinatorial) CAT(0) metric. We refer to [11] for background.




FIGURE 3.25. Some hyperplane U crosses γ with angle $\geq \theta$.

SKETCH. This can be deduced by considering midcube intersections with geodesics in a d -cube with $d \leq \dim(\tilde{X})$. □




FIGURE 3.26. Some midcube makes a large angle with a geodesic in a d -cube.

PROOF OF THEOREM 3.23. Consider the following convex subcomplex consisting of the intersection of halfspace carriers containing HK :

$$\tilde{Y} = \text{Hull}(HK) = \bigcap_{HK \subset \vec{U}} N(\vec{U})$$

Observe that $N(\vec{U}) \subset \mathcal{N}_R(H\tilde{x})$ (as we can assume $\tilde{x} \subset K$) because if $d(p, HK) > R = D + \text{diam}(K)$ then Lemma 3.26 provides a hyperplane U making a large angle with a geodesic γ nearly from p to HK , and a computation using the quasiconvexity of HK shows that $U \cap HK = \emptyset$ because of δ -thin triangles. \square




FIGURE 3.27. Any point p that is sufficiently far from HK is separated from HK by a hyperplane U . If U intersects the quasiconvex subspace HK as on the right, then a triangle with two long sides consisting of part of γ and a geodesic in U would have fellow-traveling at the corner near p which violates that $\angle \geq \theta$.

DEFINITION 3.27 (Superconvex). $\tilde{Y} \subset \tilde{X}$ is *superconvex* if it is convex and for each bi-infinite geodesic $\tilde{\gamma} \subset \tilde{X}$, if $\tilde{\gamma} \subset \mathcal{N}_r(\tilde{Y})$ for some $r \geq 0$, then $\tilde{\gamma} \subset \tilde{Y}$.

LEMMA 3.28 (Superconvex core). Let $H \subset G$ be a quasiconvex subgroup with G acting properly and cocompactly on a $CAT(0)$ cube complex \tilde{X} . For each compact subcomplex K , the subspace HK lies in an H -cocompact superconvex subcomplex $\tilde{Y} \subset \tilde{X}$.

PROOF. δ -hyperbolicity and κ -quasiconvexity imply that any bi-infinite geodesic $\tilde{\gamma}$ that lies in $\mathcal{N}_r(H\tilde{x})$ for some r , actually lies in $\mathcal{N}_s(H\tilde{x})$ for some uniform s . Now let $\tilde{Y} = \text{Hull}(\mathcal{N}_s(HK))$. \square

There are variations on Lemma 3.28 that work when G is hyperbolic relative to subgroups $\{P_1, \dots, P_r\}$. A subgroup H is *full* in G if for each i and each $g \in G$, either $|H \cap P_i^g| < \infty$ or $[P_i : H \cap P_i^g] < \infty$. The following is proven in [67]:

LEMMA 3.29 (Full superconvex core). *Let G be hyperbolic relative to $\{P_1, \dots, P_r\}$. Suppose G acts properly and cocompactly on a $\text{CAT}(0)$ cube complex \tilde{X} . Let $H \subset G$ be a relatively quasiconvex subgroup that is full. For each compact subcomplex K , the subspace HK lies in an H -cocompact superconvex subcomplex $\tilde{Y} \subset \tilde{X}$.*

CHAPTER 4

Special Cube Complexes

A CAT(0) cube complex successfully generalizes the notion of a tree, and in many ways a nonpositively curved cube complex generalizes the notion of a graph. However, from the viewpoint of many of the properties arising from (locally) cutting and retracting, a more faithful generalization of a graph is a special cube complex. The difference is felt especially prominently on algebraic properties of the fundamental group. This chapter presents special cube complexes which were introduced in [36].

We first describe special cube complexes in terms of forbidden hyperplane pathologies in Section 4.1. We describe simple criteria for the virtual specialness of a nonpositively curved cube complex in Section 4.2. Given that the hyperplane pathologies are mostly concerned with some sort of failure of embeddedness, it is not surprising that the criteria for virtual specialness have to do with separability of hyperplane subgroups. Special cube complexes were devised precisely to permit us to construct the “canonical completion” and “retraction” which we describe in Section 4.3. Using the core described in Section 3.5, this enables us to see that quasiconvex subgroups are separable when X is compact and special and $\pi_1 X$ is hyperbolic. Thus in the presence of hyperbolicity, the virtual specialness criterion takes an even more definitive form as described in Section 4.4. Namely, X is virtually special iff all quasiconvex subgroups are separable. Finally, we discuss fundamental commutative diagrams that make the canonical completion and retraction behave predictably with respect to subcomplexes in Section 4.5. We then describe a lemma that will help us to aim the canonical retraction map in Section 4.6. These last two technical sections will support the proof sketched in Chapter 5.

4.1. Hyperplane Definition of Special Cube Complex

DEFINITION 4.1 (Immersed Hyperplane). An *immersed hyperplane* in a nonpositively curved cube complex X is $\bar{H} = \text{Stabilizer}(H) \backslash H$ where H is a hyperplane of \tilde{X} . Note that there is a natural map $\bar{H} \rightarrow X$ and that each midcube of a cube of X extends to a unique immersed hyperplane.

DEFINITION 4.2 (Special). A nonpositively curved cube complex X is *special* if:

- (1) Each immersed hyperplane embeds (and we will thus omit the term “immersed”).




FIGURE 4.1. An immersed hyperplane in a nonpositively curved cube complex

- (2) Each hyperplane is 2-sided.
- (3) No hyperplane self-osculates.
- (4) No two hyperplanes interosculate.

The prohibited hyperplane pathologies are depicted in Figure 4.2. The immersed hyperplane \bar{H} is *2-sided* if there is a way of consistently directing its dual 1-cubes so that 1-cubes on opposite sides of a square have the same direction. Using this direction on the 1-cubes, we say \bar{H} *self-osculates* if it is dual to two distinct 1-cubes with the same initial or terminal 0-cube. Finally, \bar{H}, \bar{H}' *interosculate* if they cross and they have dual 1-cubes that share a 0-cube but do not lie on a common 2-cube.




FIGURE 4.2. Self-crossing, 1-sidedness, self-osculation, and interosculation




FIGURE 4.3. The “no interosculation” condition means that if two hyperplanes cross and then come close together at some other vertex, then they also cross at that vertex. Thus if one sees a configuration as on the left, then there must have been an additional square allowing the hyperplanes to cross again as on the right.

EXAMPLE 4.3.

- (1) Any graph is special.
- (2) Any CAT(0) cube complex is special.




FIGURE 4.4. The above configuration is not prohibited in a special cube complex – note the orientation which makes it an “indirect self-osculation”, and it thus differs from the third pathology in Figure 4.2

- (3) Any subcomplex of a product of two graphs $A \times B$ is special.
- (4) **The cube complex $R(\Gamma)$ of a raag $G(\Gamma)$ is special.**

Any compact clean \mathcal{VH} -complex X has a finite cover \widehat{X} that is a subcomplex of the product $A \times B$ of graphs (see [81]). Even this turns out to be a surprisingly rich family of virtually special examples. However, it is the final class in Example (2.10) that is the most important: The fundamental examples of special cube complexes are those associated to raags as in Theorem 2.10.

The following result provides an alternate characterization of specialness. It provides a bridge from the organized cubical geometry of a special cube complex to the easy algebra of raags.

THEOREM 4.4. *Let X be a nonpositively curved cube complex. Then X is special iff there is a local isometry $X \rightarrow R$ for some $R = R(\Gamma)$.*

Since a local isometry is π_1 -injective we have:

COROLLARY 4.5. *If X is special then $\pi_1 X$ is a subgroup of a raag.*

PROOF OF THEOREM 4.4. (\Leftarrow) If $A \rightarrow B$ is a local isometry and B is special then A is special since the four pathologies project to pathologies under local isometries.

(\Rightarrow) Let Γ be the *crossing graph* of X . The vertices of Γ correspond to hyperplanes of X and two vertices are adjacent iff the corresponding hyperplanes cross.

The 2-sidedness of hyperplanes allows us to consistently direct 1-cells of X dual to the same hyperplane.

The labeling and directing of X^1 gives a map $X^1 \rightarrow R(\Gamma)$ which extends to a local isometry $X \rightarrow R$. It is an immersion since there is no self-osculation. It is a local isometry since there is no interosculation. \square

4.2. Separability Criteria for Virtual Specialness

THEOREM 4.6. *Let X be a compact nonpositively curved cube complex. Then X has a finite special cover if and only if*

- (1) $\pi_1 U$ is separable in $\pi_1 X$ for each hyperplane U




FIGURE 4.5. The crossing graph Γ of hyperplanes in X

- (2) $\pi_1 U \pi_1 V$ is separable in $\pi_1 X$ for each pair of crossing hyperplanes U, V .




FIGURE 4.6. Some double hyperplane cosets $\pi_1 U \pi_1 V$ to consider.

Note that we allow the basepoint to vary among all centers of all cubes in the above statement, and we also note that U, V might denote the same self-crossing hyperplane.

DEFINITION 4.7. $S \subset G$ is *separable* if it is closed in the *profinite topology* on G which is the topology whose basis consists of finite index cosets. Equivalently, S is *separable* if for each $g \notin S$, there is a finite quotient $G \rightarrow \bar{G}$ such that $\bar{g} \notin \bar{S}$. In particular, G is *residually finite* if each nontrivial element g has nontrivial image \bar{g} in some finite quotient \bar{G} of G . Thus G is residually finite iff $\{1_G\}$ is separable which holds when the profinite topology is Hausdorff. A subgroup H is separable iff H is the intersection of finite index subgroups of G .

Separability is related to lifting to embeddings because of the following lemma first made explicit by Peter Scott [68]:

LEMMA 4.8. *Let $\widehat{X} \rightarrow X$ be a covering map of complexes. Then $\pi_1 \widehat{X}$ is separable in $\pi_1 X$ (if and) only if the following holds: Then for each compact subcomplex $K \subset \widehat{X}$ there is an intermediate finite cover \bar{X} such that $\widehat{X} \rightarrow X$ factors as $\widehat{X} \rightarrow \bar{X} \rightarrow X$ and such that K embeds in \bar{X} as follows:*

$$\begin{array}{ccccc}
 & & K & & \\
 & \swarrow & \downarrow & \searrow & \\
 \widehat{X} & \rightarrow & \bar{X} & \rightarrow & X
 \end{array}$$

A well-known consequence of Lemma 4.8 is:

COROLLARY 4.9. *Suppose Y is compact and $Y \rightarrow X$ is π_1 -injective, and Y lifts to an embedding in the cover \tilde{X} associated to $\pi_1 Y$, and $\pi_1 Y$ is separable. Then $Y \rightarrow X$ lifts to an embedding in a finite cover.*

Corollary 4.9 explains how Conditions (1), (2), and (3) in Definition 4.2 are achieved in a finite cover using Condition (1) of Theorem 4.6. Condition (2) of Definition 4.6 allows us to obtain Condition (4) of Definition 4.2 in a finite cover.

REMARK 4.10. There is a version of Theorem 4.6 that works for a G action on \tilde{X} with finitely many hyperplane orbits and finitely many $\text{Stabilizer}(U)$ -orbits of hyperplanes intersecting U and osculating U for each hyperplane U of \tilde{X} . (We need $\text{Intersector}(U, V)$ and $\text{Osculator}(U, V)$ to be separable).

Theorem 4.6 is naturally achievable in many cases. For instance, we used it to verify the virtual specialness of the Niblo-Reeves cubulation [59] of Coxeter groups in [37], and we used it to verify the virtual specialness of simple-type arithmetic hyperbolic lattices in [6, 7].

4.3. Canonical Completion and Retraction

We now turn to the notions of canonical completions and retractions, and we refer to [36, 35]. These notions and some of their applications were originally explored in 1-dimension in the course of understanding separability properties of free groups in work published in [78], and subsequently [79]. Around 1994, when I was first examining this, it seemed to be an organized way of producing covering spaces of graphs, but turned out to have very convenient specific properties. An early attempt was made to generalize this to the cube complexes of raags in [45], and subsequently it was a nice surprise to find a clear-cut formulaic generalization with Haglund in 2002, when we created special cube complexes without initially realizing the connection to raags. It is substantially because of the canonical completion and retraction properties that special cube complexes deserve to be thought of as “generalized graphs”, at least from the viewpoint of the retractive properties of graphs that lead to many group theoretical properties of free groups.

DEFINITION 4.11 ($A \otimes_R B$). Let $\alpha : A \rightarrow R$ and $\beta : B \rightarrow R$ be local isometries of cube complexes. We define their *fiber-product* $A \otimes_R B$ to be the cube complex whose n -cubes are pairs of n -cubes in A and B respectively that map to the same cube in R . This is a frequently encountered definition when R is a bouquet of circles, and it is defined analogously in higher dimensions. We note that $A \otimes_R B$ may not be connected. A 0-cell of $A \otimes_R B$ corresponds to a pair $(a, b) \in A \times B$ with $\alpha(a) = \beta(b)$. When α, β are connected covering maps, the component of $A \otimes_R B$ containing (a, b) is

the based cover that is the smallest common cover of the based covers (A, a) and (B, b) of (R, p) where $\alpha(a) = p = \beta(b)$. We note that $A \otimes_R B$ is the universal receiver in the category whose objects are commutative diagrams of local isometries of cube complexes as displayed below. Morphisms in this category are maps $C_1 \rightarrow C_2$ so that there is the usual further commutative diagram.

$$\begin{array}{ccc} C & \rightarrow & B \\ \downarrow & & \downarrow \\ A & \rightarrow & R \end{array}$$

Definition 4.2 was crafted to enable the following generalization of M.Hall's theorem:

CONSTRUCTION 4.12. Let $f : Y \rightarrow X$ be a local isometry with X special and Y compact. There exists a finite cover $\rho : \mathbb{C}(Y \rightarrow X) \rightarrow X$ and an embedded lift $\tilde{f} : Y \rightarrow \mathbb{C}(Y \rightarrow X)$ of $f : Y \rightarrow X$, and a retraction map $r : \mathbb{C}(Y \rightarrow X) \rightarrow Y$. The maps ρ and r are the *canonical completion* and *canonical retraction* associated to $f : Y \rightarrow X$.

Let $X \rightarrow R$ denote the local isometry to the nonpositively curved cube complex of the right-angled Artin group associated with X in Theorem 4.4.

In the 1-dimensional case we complete each component of the preimage of each loop of R to a covering map. The new edges retract to the arcs of components they complete to circles. This provides $Y \hookrightarrow \mathbb{C}(Y \rightarrow R)$ and the retraction $\mathbb{C}(Y \rightarrow R) \rightarrow Y$. We then define $\mathbb{C}(Y \rightarrow X) = X \otimes_R \mathbb{C}(Y \rightarrow R)$, and we define the retraction map using the composition $\mathbb{C}(Y \rightarrow X) \rightarrow \mathbb{C}(Y \rightarrow R) \rightarrow Y$. We refer the reader to the commutative diagram and concrete example in Figure 4.7:

We now turn to the general case and refer to Figure 4.8. The key point is that for a local isometry $Y \rightarrow R$ when we consider the 1-skeleta, the previously defined map $\mathbb{C}(Y^1 \rightarrow R^1) \rightarrow R^1$ extends to a covering map $\mathbb{C}(Y \rightarrow R) \rightarrow R$. We again use the fiber-product to define $\mathbb{C}(Y \rightarrow X)$ as in the 1-dimensional case exactly as in the commutative diagram in Figure 4.7. Finally the retraction maps defined on the 1-skeleton extend naturally to the higher cubes allowing us to define the canonical retraction $\mathbb{C}(Y \rightarrow R) \rightarrow Y$, and we then define the canonical retraction $\mathbb{C}(Y \rightarrow X) \rightarrow Y$ using the composition $\mathbb{C}(Y \rightarrow X) \rightarrow \mathbb{C}(Y \rightarrow R) \rightarrow Y$.

4.4. Separability in the Hyperbolic Case

We can now obtain some interesting applications towards separability by combining the core result described in Section 3.5 with the canonical completion tool of Section 4.3.

THEOREM 4.13. *Suppose X is special and compact and that $\pi_1 X$ is hyperbolic. Then every quasiconvex subgroup H of $G = \pi_1 X$ is separable.*




FIGURE 4.7. Canonical Completion in 1-dimensional case.

$$\begin{array}{ccccc}
 \mathbb{C}(Y \rightarrow X) = X \otimes_R \mathbb{C}(Y \rightarrow R) & \rightarrow & \mathbb{C}(Y \rightarrow R) & & \\
 & \searrow & \nearrow & & \\
 \downarrow & & Y & & \downarrow \\
 & \swarrow & \searrow & & \\
 X & \rightarrow & R & &
 \end{array}$$

PROOF 1. Let $\sigma \notin H$. By Theorem 3.23, there is an H -cocompact convex subcomplex \tilde{Y} that contains $\tilde{\sigma}$. Let $Y = H \backslash \tilde{Y}$. Let $G' = \pi_1 \mathbb{C}(Y \rightarrow X)$. Then $H \subset G'$ but $\sigma \notin G'$. \square

PROOF 2. (assumes residual finiteness of raags) By Corollary 3.24, let $Y \rightarrow X$ be a compact local isometry with $H = \pi_1 Y$. Then Y is a retract of $\mathbb{C}(Y \rightarrow X)$. Thus H is a retract of $G' = \pi_1 \mathbb{C}(Y \rightarrow X)$. Thus H is closed in the profinite topology on G' since a retract of a Hausdorff topological space is closed. Hence H is separable in G' , and thus separable in G . \square

As a quick consequence of Theorem 4.13 we see that: If G is virtually compact special hyperbolic and H is quasiconvex in G , then H is separable. Indeed, let G' be the finite index special subgroup, and let $H' = G' \cap H$. Note that H' is separable in G' , and thus the intersection of finite index subgroup and hence also separable in G . Thus H separable in G as it is the union of finitely many closed sets consisting of the cosets of H' in H .

THEOREM 4.14. *Let X be a compact nonpositively curved cube complex with $\pi_1 X$ hyperbolic. Then X is virtually special if and only if each quasiconvex subgroup is separable.*

PROOF. (\Rightarrow) Holds by Theorem 4.13. (\Leftarrow) Note that hyperplanes have quasiconvex subgroups by Corollary 3.13. If all quasiconvex subgroups are




FIGURE 4.8. The above figure corresponds to the following commutative diagram:

$$\begin{array}{ccccc}
 \mathcal{C}(Y \rightarrow X) & \rightarrow & \mathcal{C}(Y \rightarrow R) \supset \mathcal{C}(Y^1 \rightarrow R^1) & & \\
 & \nwarrow & \nearrow & & \\
 \downarrow & & Y & & \downarrow \\
 & \swarrow & \searrow & & \downarrow \\
 X & \rightarrow & R \supset R^1 & &
 \end{array}$$

separable then all quasiconvex double cosets are separable by Theorem 4.16. The result thus follows from Theorem 4.6. \square

We obtain the following useful consequence of Theorem 4.14. Note that while generalizations of this are known in relatively hyperbolic situations, it is an open problem whether virtual specialness is always an invariant of the fundamental group.

COROLLARY 4.15. *Suppose X and X' are compact nonpositively curved cube complexes with isomorphic fundamental groups that are hyperbolic. If X is virtually special then X' is virtually special.*

The following result was obtained by Minasyan clarifying earlier work of Gitik [56, 26].

THEOREM 4.16. *Let G be an arbitrary hyperbolic group. If all quasiconvex subgroups of G are separable then all double quasiconvex cosets of G are separable.*

We close this section with a conjecture about virtual specialness:

CONJECTURE 4.17. *Let X be a compact nonpositively curved cube complex. Suppose $\pi_1 P$ is separable for each hyperplane P of X . Then X is virtually special.*

REMARK 4.18. We note that Conjecture 4.17 asks whether the criteria of Theorems 4.6 and 4.14 can be simplified. Conjecture 4.17 has an affirmative answer in the very simple scenario where X is a nonpositively curved \mathcal{VH} -complex [81]. For in that case, separable hyperplanes implies virtual cleanliness, and compact clean complexes virtually embed in the product $A \times B$ of graphs.

When $\pi_1 X$ is hyperbolic, separability allows us to pass to a finite cover where all hyperplanes are malnormal and embedded. Then Theorem 5.1 proves Conjecture 4.17 by induction. Similarly, in the relatively hyperbolic case, Theorem 15.3 proves Conjecture 4.17 when $\pi_1 X$ is hyperbolic relative to f.g. free abelian subgroups. Little is known about the general case, and a result not depending on some form of hyperbolicity will push the envelope here.

4.5. Wall-Injectivity and a Fundamental Commutative Diagram

This section is only necessary to support the details in an argument in Chapter 5.

DEFINITION 4.19 (Wall-injective). A combinatorial map of cube complexes $D \rightarrow C$ is *wall-injective* if distinct hyperplanes of D map to distinct hyperplanes of C .

When D and C are 1-dimensional, wall-injectivity is equivalent to injectivity on the set of 1-cells. Figure 4.9 depicts an embedding that is not wall-injective.



FIGURE 4.9. The inclusion of cube complexes above is not wall-injective

Being wall-injective is not stable under taking covering spaces, as can be seen in Figure 4.10.

REMARK 4.20. Consider the canonical completion $\mathcal{C}(Y \rightarrow X)$ of a local isometry $Y \rightarrow X$. Observe that $Y \hookrightarrow \mathcal{C}(Y \rightarrow X)$ is wall-injective. Indeed this follows by examining the behavior of the canonical retraction map.

LEMMA 4.21 (Fundamental commutative diagram). *Suppose $D \rightarrow C$ is a wall-injective local-isometric embedding with D finite and C special. Then*




FIGURE 4.10. The preimage of a wall-injective subcomplex in a covering space might not be wall-injective, as is the case for the red circle above.

there are commutative diagrams whose vertical maps are canonical completions, canonical retractions, and canonical inclusions.

$$\begin{array}{ccc}
 D & = & D \\
 \cap & & \cap \\
 \mathbb{C}(D \rightarrow D) & \subset & \mathbb{C}(D \rightarrow C) \\
 \downarrow & & \downarrow \\
 D & \hookrightarrow & C
 \end{array}
 \qquad
 \begin{array}{ccc}
 D & = & D \\
 \uparrow & & \uparrow \\
 \mathbb{C}(D \rightarrow D) & \subset & \mathbb{C}(D \rightarrow C)
 \end{array}$$

SKETCH. In the 1-dimensional case, imagine first building $\mathbb{C}(D \rightarrow D)$ and then building $\mathbb{C}(D \rightarrow C)$ around it. The general case is similar. \square




FIGURE 4.11. The fundamental commutative diagram.

4.6. Wall Projection Controls Retraction

This section is only necessary to support the details in an argument in Chapter 5.

DEFINITION 4.22 (Wall Projection). Let $A \subset X$ and $B \subset X$ be subcomplexes of a nonpositively curved cube complex X . Define the 1-skeleton of $WProj_X(B \rightarrow A)$ to be the union of A^0 and all 1-cells a of A that are *parallel* to 1-cells of B in the sense that there is 1-cell b of B such that a, b are dual to the same hyperplane of X . To this 1-skeleton we add all cubes whose 1-skeleta were included. (We think of the cubes being likewise parallel into B .)




FIGURE 4.12. $WProj_X(B \rightarrow A)$ is indicated on the left, where A, B are the indicated subcomplexes of the complex X on the right.

$WProj_X(B \rightarrow A)$ consists of a collection of locally-convex subcomplexes of A . In the 1-dimensional case, $WProj_X(B \rightarrow A) = A^0 \cup (B \cap A)$.

Let $\widehat{B} \rightarrow B$ be a covering map, and let $A \rightarrow B$ be a map with A connected. An *elevation* of $A \rightarrow B$ is a minimal cover $\widehat{A} \rightarrow A$ and a lift of $\widehat{A} \rightarrow B$ to a map $\widehat{A} \rightarrow \widehat{B}$. There are usually multiple choices of elevation - and one can specify by choosing basepoints. The choice is usually clear from the context and often allowed to vary. We abuse the terminology as is done for the more standard term “lift” which is an elevation with $\widehat{A} = A$, and when the elevation $\widehat{A} \rightarrow \widehat{B}$ is an embedding, we often refer to its image as the elevation of A to \widehat{B} .

Our reason for being interested in wall projections is the following:

LEMMA 4.23. *Let A, B be subcomplexes of the special cube complex X with A compact. Consider an elevation \check{B} of B to $\mathbb{C}(A \rightarrow X)$. Then $r(\check{B}) \subset WProj_X(B \rightarrow A)$.*

SKETCH. Since the canonical retraction map is “label preserving” (though it is not orientation preserving and can collapse dimension) a 1-cell \check{b} of \check{B} either collapses to a 0-cell or maps to a 1-cells in A dual to the same hyperplane as b in X . \square

Lemma 4.23 is powerful when $WProj_X(B \rightarrow A)$ is *trivial* in the sense that each component of $WProj_X(B \rightarrow A)$ is simply connected - i.e. $CAT(0)$. For then $r(\check{B})$ is null-homotopic for each \check{B} .

The following is proven in [35]:

LEMMA 4.24 (Trivial Wall Projection). *Let X be a compact (virtually) special nonpositively curved cube complex with $\pi_1 X$ hyperbolic. let $A \rightarrow X$ be a compact local isometry with $\pi_1 A \subset \pi_1 X$ malnormal. There exists a finite cover $A_o \rightarrow A$ such that any further cover $\bar{A} \rightarrow A_o$ can be completed to a finite special cover $\bar{X} \rightarrow X$ such that:*

- (1) *All elevations of $A \rightarrow X$ to \bar{X} are embeddings.*
- (2) *The base elevation \bar{A} is wall-injective in \bar{X} .*
- (3) *Every elevation \check{A} of A that is distinct from $\bar{A} \subset \bar{X}$ has $WProj_{\bar{X}}(\check{A} \rightarrow \bar{A})$ trivial.*

SKETCH. This is easy when X is 1-dimensional: Let $X_o \rightarrow X$ be a finite cover in which all elevations of A are embedded, and let A_o be the base elevation. Note that each $\check{A}_1 \cap \check{A}_2$ is a forest because of malnormality. Now for any $\bar{A} \rightarrow A_o$ we let $\bar{X} = \mathbb{C}(\bar{A} \rightarrow X_o)$. \square

The proof of Lemma 4.24 in higher dimensions is treacherous...

CHAPTER 5

Virtual Specialness of Malnormal Amalgams

In this chapter we employ results described in Chapter 4 about special cube complexes and particularly about canonical completion and retraction, to give a sketch of the main theorem in [35]. Theorem 5.2 states this theorem in the case where G acts properly and cocompactly on \tilde{X} , but we focus on the simpler version in Theorem 5.1 which gives a criterion for a nonpositively curved cube complex X to be virtually special if it can be cut along a malnormal hyperplane into virtually special pieces. Section 5.1 proves that quasiconvex subgroups of $\pi_1 X$ are separable, and hence X is virtually special by the criterion in Theorem 4.14. The key ingredient enabling this proof of separability is the covering space property stated in Lemma 5.3, which we focus on in Section 5.2. As the proof of Lemma 5.3 is hard to motivate, we first present a simpler but incorrect sketch of proof relying on an unrealizable fantasy. The reader is encouraged to consider the fantasy and real versions in Figures 5.10 and 5.11 when attempting to navigate through the details.

Though the proof of Theorem 5.1 should generalize to various relatively hyperbolic settings, and indeed, the statement of Theorem 5.1 is covered by a mild generalization of Theorem 15.3, the statement of Theorem 5.1 is difficult to generalize without some measure of hyperbolicity and/or malnormality. We refer to Conjecture 4.17.

Although we have attempted to give sketches, the proofs are still quite involved, and this chapter is recommended only for the reader who wishes to see the an account of the details of subgroup separability proofs that are executed by building finite covering spaces. The general reader should absorb the statement of Theorem 5.1 and move on to subsequent material. It is only the statement of Theorem 5.1 (or actually the generalization allowing torsion in Theorem 5.2) will be used later, and we will not use the methodology employed in these proofs.

5.1. Specializing Malnormal Amalgams

We will sketch a proof of the following torsion-free case of Theorem 5.2:

THEOREM 5.1. *Let Q be a compact nonpositively curved cube complex with $\pi_1 Q$ hyperbolic. Let P be a 2-sided embedded hyperplane with $\pi_1 P$ malnormal in $\pi_1 Q$. Suppose each component of $Q - N_o(P)$ is virtually special. Then Q is virtually special.*




FIGURE 5.1. $P \subset N_o(P) \subset Q$ in the separating and nonseparating cases.

Note that $N_o(P)$ denotes the *open cubical neighborhood* of P consisting of all open cubes intersecting P . Theorem 5.1 is already interesting in the case that each component of $Q - N_o(P)$ is a graph, which is the main result in [79]. Surprisingly, there is little formal difference between the proofs, though several easy steps in the 1-dimensional case turn out to be deeper and technically challenging to verify in arbitrary dimensions. This is especially the case for Lemma 4.24.

The following is one of the two ingredients providing Theorem 12.2. An “orbihedron style” reformulation would make its statement parallel to Theorem 5.1. Its proof requires a bit of bookkeeping to see that the torsion is irrelevant right from the onset, and we refer to [35] for the details. We use the notation \vec{P} for a halfspace associated to a hyperplane \tilde{P} .

THEOREM 5.2. *Let G be a hyperbolic group that acts properly and cocompactly on the $CAT(0)$ cube complex \tilde{Q} . Let \tilde{P} be a hyperplane such that $\text{Stabilizer}(\tilde{P})$ is almost malnormal in G and such that $g\tilde{P} \cap \tilde{P} \neq \emptyset \Rightarrow g\tilde{P} = \tilde{P}$ and moreover $g\vec{P} = \vec{P}$.*

Suppose that for each component \tilde{X} of $\tilde{Q} - N_o(\tilde{P})$, the group $\text{Stabilizer}(\tilde{X})$ has a finite index torsion-free subgroup H such that $H \backslash \tilde{X}$ is compact special.

Then G has a finite index torsion-free subgroup J such that $J \backslash \tilde{Q}$ is compact special.

A main ingredient in the proof of Theorem 5.1 is the following:

LEMMA 5.3 (Isomorphic Elevations). *Let Q be a compact nonpositively curved cube complex and let P be a hyperplane in Q such that:*

- (1) $\pi_1 Q$ is hyperbolic
- (2) P is an embedded, nonseparating, 2-sided hyperplane in Q .
- (3) $\pi_1 P$ is malnormal in $\pi_1 Q$.
- (4) $X = Q - N_o(P)$ is virtually special.

Then for any finite cover $\widehat{X} \rightarrow X$, there is a finite regular cover $\widehat{X} \xrightarrow{\boxtimes}$ factoring through \widehat{X} such that $\widehat{X} \xrightarrow{\boxtimes}$ induces the same cover on each side A, B of P .

We prove Theorem 5.1 by showing that $\pi_1 Q$ has separable quasiconvex subgroups and applying Theorem 4.14.

THEOREM 5.4. *Let Q be as in Lemma 5.3. Then every quasiconvex subgroup of $\pi_1 Q$ is separable.*




FIGURE 5.2. $\overset{\boxtimes}{X} \rightarrow \widehat{X} \rightarrow X$. The isomorphism between A, B arising from the product $N(P)$ lifts to isomorphisms between the elevations $\overset{\boxtimes}{A}, \overset{\boxtimes}{B}$ to the finite cover $\overset{\boxtimes}{X}$ that factors through the cover \widehat{X} .

PROOF. We think of Q as a graph of spaces, with vertex space $X = Q - N_o(P)$ and with open edge space $N_o(P)$ having attaching maps $A \rightarrow X$ and $B \rightarrow X$.

Let $\dot{Q} \rightarrow Q$ be a based cover associated to some quasiconvex subgroup H , and let $\sigma \in \pi_1 Q - \pi_1 \dot{Q}$, and let $\dot{\sigma}$ be the based lift of σ to \dot{Q} .




FIGURE 5.3. The compact core Y of \dot{Q} is chosen to contain $\dot{\sigma}$.

By Corollary 3.24, let Y be a compact core containing $\dot{\sigma}$ (see Figure 5.3). For each vertex space Y_i of Y contained in \dot{X}_i , we use separability in $\pi_1 X$ to choose an intermediate finite cover $\widehat{X}_i \rightarrow X$ (with $\dot{X}_i \rightarrow \widehat{X}_i \rightarrow X$) such that $Y_i \hookrightarrow \dot{X}_i$ projects to an embedding $Y_i \hookrightarrow \widehat{X}_i$. And moreover, we use the double quasiconvex coset separability of Theorem 4.16 to ensure that, furthermore, the various incoming and outgoing edge space of Y_i project to distinct elevations of A, B to \widehat{X}_i . See Figure 5.4.

Let Z be obtained from Y by extending each Y_i to \widehat{X}_i (see Figure 5.5). Note that $Z \rightarrow X$ is a finite cover on each vertex space, and contains some information on how to attach edge spaces. We now create a finite cover $\overset{\boxtimes}{Z} \rightarrow Z$ such that all vertex spaces are isomorphic:




FIGURE 5.4. The intermediate cover \widehat{X}_i maintains the embedding of Y_i and also maintains the distinctness of the attachment sites of those incoming and outgoing edge spaces at \widehat{X}_i that are represented in Y . We have sketched an example where X is a graph, and all edge spaces are cylinders, so the incoming (red) and outgoing (purple) edge spaces are lines and circles which project to circles in \widehat{X}_i .




FIGURE 5.5. We then extend $Y \rightarrow X$ to $Z \rightarrow X$ by extending each vertex space Y_i to the chosen finite cover \widehat{X}_i .

Let \widehat{X} be a finite cover factoring through each \widehat{X}_i . Let $\overset{\boxtimes}{X} \rightarrow \widehat{X}$ be as in Lemma 5.3. Let $\overset{\boxtimes}{Q} \rightarrow Q$ be a finite cover whose unique vertex space is $\overset{\boxtimes}{X}$ (see Figure 5.6).

Let $\overset{\boxtimes}{Z} = \overset{\boxtimes}{Q} \otimes_Q Z$ so there is a commutative diagram:

$$\begin{array}{ccc}
 \overset{\boxtimes}{Z} & \rightarrow & Z \\
 \downarrow & & \downarrow \\
 \overset{\boxtimes}{Q} & \rightarrow & Q
 \end{array}$$

The partial one-to-one correspondence between the various elevations of A, B to vertex spaces of $\overset{\boxtimes}{Z}$ extends to a complete one-to-one correspondence, and we use this to attach all missing edge spaces. Note that there are as many missing incoming as missing outgoing because A, B have the same number of elevations to $\overset{\boxtimes}{X}$. Let \widehat{Q} denote the finite cover of Q constructed in this manner, and note that $\overset{\boxtimes}{Z} \subset \widehat{Q}$.




FIGURE 5.6. The finite cover \hat{Q} has vertex space isomorphic to \hat{X} and attached edge spaces associated to a one-to-one correspondence between elevations of A and B .




FIGURE 5.7. The edge spaces of \hat{Z} (above) are extended to finite covers, and the missing edge spaces are added to obtain \hat{Q} (below).

Finally, σ is separated from $\pi_1 Q$ in the right coset representation in $\pi_1 \hat{Q}$. Indeed, each element of $\pi_1 Y$ sends the trivial coset to a coset corresponding to a preimage of the basepoint of Y (i.e. a lift of a basepoint of Z) in \hat{Z} , but σ translates the trivial coset to the coset represented by the endpoint of the lift $\hat{\sigma}$ of σ to \hat{Q} . \square

5.2. Proof of the Isomorphic Elevation Lemma

We now describe the proof of Lemma 5.3 in the separating case. Let $Q = N_o(P) = L \sqcup R$ and let $A \leftarrow P$ and $B \rightarrow P$ denote the attaching maps of $P \times \{\pm 1\}$ as in Figure 5.8.




FIGURE 5.8. $Q = L \sqcup_{A \cong B} R$.

We refer to Figure 5.9 for a schematic summary of the entire proof under rigorous assumptions.

REMARK 5.5 (Fantasy Proof). We refer the reader to Figure 5.10 for the initial motivation under simplified assumptions that cannot be obtained: Namely, that there is a cover \widehat{L} and an elevation \widehat{A} and a retraction $\widehat{L} \rightarrow \widehat{A}$ such that all other elevations of A to L are nullhomotopic under the retraction. (And likewise, a cover \widehat{R} and elevation \widehat{B} to \widehat{R} with $\widehat{B} \cong \widehat{A}$, and a retraction $\widehat{R} \rightarrow \widehat{B}$.)

To see that this fantasy cannot be realized, consider the case where L is a bouquet of two circles labelled by a, b and A is an immersed circle corresponding to the commutator $[a, b]$. Then any finite \widehat{L} corresponds to a branched cover of a surface whose 2-cells have boundary paths that are elevations of $[a, b]$. Now for any cover, the elevation \widehat{A} is homologically equal to the sum of all the other elevations, and so these others cannot all be simultaneously null-homotopic under a retraction map to A .

We use Lemma 4.24 and some canonical completions and retractions to obtain finite covers \widehat{L}, \widehat{R} with based elevations \widehat{A}, \widehat{B} such that:

- (1) Each elevation of A to \widehat{L} and of B to \widehat{R} is an embedding
- (2) The isomorphism $A \cong B$ lifts to an isomorphism $\widehat{A} \cong \widehat{B}$
- (3) The base elevations \widehat{A}, \widehat{B} are wall-injective in \widehat{L}, \widehat{R}
- (4) $\text{WProj}_{\widehat{L}}(A_i \rightarrow \widehat{A})$ and $\text{WProj}_{\widehat{R}}(B_i \rightarrow \widehat{B})$ are trivial for all nonbase elevations A_i, B_i .

Form $\mathbb{C}(\widehat{A} \rightarrow \widehat{L})$ and $\mathbb{C}(\widehat{B} \rightarrow \widehat{R})$ and note that we have the following commutative diagram from Lemma 4.21:

$$\begin{array}{ccccccc}
 \mathbb{C}(\widehat{A} \rightarrow \widehat{L}) & \leftarrow & \mathbb{C}(\widehat{A} \rightarrow \widehat{A}) & \cong & \mathbb{C}(\widehat{B} \rightarrow \widehat{B}) & \rightarrow & \mathbb{C}(\widehat{B} \rightarrow \widehat{R}) \\
 \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 \widehat{L} & \leftarrow & \widehat{A} & \cong & \widehat{B} & \rightarrow & \widehat{R}
 \end{array}$$

Let $\bar{P} \rightarrow P$ be a finite cover of P factoring through all (corresponding) elevations of A to $\mathbb{C}(\widehat{A} \rightarrow \widehat{L})$ and elevations of B to $\mathbb{C}(\widehat{B} \rightarrow \widehat{R})$. Let \bar{A}, \bar{B} be

the corresponding covers of A, B . We use canonical retraction to induce the following covers:

$$\begin{array}{ccccccc} \overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})} & \rightarrow & \overline{A} & \overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{A})} & \rightarrow & \overline{A} & \overline{\mathbb{C}(\widehat{B} \rightarrow \widehat{B})} & \rightarrow & \overline{B} & \overline{\mathbb{C}(\widehat{B} \rightarrow \widehat{R})} & \rightarrow & \overline{B} \\ \downarrow & & \downarrow & \downarrow & & \downarrow & \downarrow & & \downarrow & \downarrow & & \downarrow \\ \mathbb{C}(\widehat{A} \rightarrow \widehat{L}) & \rightarrow & \widehat{L} & \mathbb{C}(\widehat{A} \rightarrow \widehat{A}) & \rightarrow & \widehat{A} & \mathbb{C}(\widehat{B} \rightarrow \widehat{B}) & \rightarrow & \widehat{B} & \mathbb{C}(\widehat{B} \rightarrow \widehat{R}) & \rightarrow & \widehat{R} \end{array}$$

We obtain the following commutative diagram:

$$\begin{array}{ccccccc} \overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})} & \hookrightarrow & \overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{A})} & \cong & \overline{\mathbb{C}(\widehat{B} \rightarrow \widehat{B})} & \hookrightarrow & \overline{\mathbb{C}(\widehat{B} \rightarrow \widehat{R})} \\ \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \widehat{L} & \hookrightarrow & \widehat{A} & \cong & \widehat{B} & \hookrightarrow & \widehat{R} \end{array}$$

Let $\overline{\mathbb{A}}_L$ and $\overline{\mathbb{A}}$ denote the smallest regular covers factoring through all elevations of A to $\overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})}$ and $\overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{A})}$ respectively. We claim that $\overline{\mathbb{A}}_L$ and $\overline{\mathbb{A}}$ are isomorphic covers of A . It is immediate that $\overline{\mathbb{A}}_L$ factors through $\overline{\mathbb{A}}$ since $\overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{A})} \hookrightarrow \overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})}$ so each elevation of A to the former is also an elevation to the latter. To see that $\overline{\mathbb{A}}$ factors through $\overline{\mathbb{A}}_L$ we note that elevations of A to $\overline{\mathbb{A}}_L$ are either contained in $\overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{A})}$ or factor through non-base elevations of A to \widehat{L} .

Consider an elevation \widehat{A}_i of A to $\overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})}$, and note that it factors through elevations to intervening spaces as in the following diagram:

$$\begin{array}{ccc} \widehat{A}_i & \subset & \overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})} \\ \parallel & & \downarrow \\ \widehat{A}_i & \subset & \mathbb{C}(\widehat{A} \rightarrow \widehat{L}) \\ \downarrow & & \downarrow \\ A_i & \rightarrow & \widehat{L} \\ \downarrow & & \downarrow \\ A & \rightarrow & L \end{array}$$

The key point is that $\widehat{A}_i \cong \widehat{A}_i$ as, by Lemma 4.23, the triviality of $\text{WProj}_{\widehat{L}}(A_i \rightarrow \widehat{A})$ implies that \widehat{A}_i is nullhomotopic in the canonical retraction map $r : \mathbb{C}(\widehat{A} \rightarrow \widehat{L}) \rightarrow \widehat{A}$. Thus \widehat{A} factors through $\widehat{A}_i \cong \widehat{A}_i$ by choice of $\overline{\mathbb{A}}$. Thus $\overline{\mathbb{A}}$ factors through \widehat{A}_i since $\overline{\mathbb{A}}$ factors through $\overline{\mathbb{A}}$.

Similarly $\overline{\mathbb{B}}_R \cong \overline{\mathbb{B}}$.

Let $\overline{\mathbb{L}}$ and $\overline{\mathbb{R}}$ denote the smallest regular covers factoring through $\overline{\mathbb{C}(\widehat{A} \rightarrow \widehat{L})}$ and $\overline{\mathbb{C}(\widehat{B} \rightarrow \widehat{R})}$. Note that elevations of A to $\overline{\mathbb{L}}$ are isomorphic to $\overline{\mathbb{A}}_L \cong \overline{\mathbb{A}}$ and similarly for elevations of B .




FIGURE 5.9. We form $\overset{\boxtimes}{Q}$ by taking a balanced number of copies: $3 \cdot \overset{\boxtimes}{L}$ and $2 \cdot \overset{\boxtimes}{R}$ and gluing them together along copies of $\overset{\boxtimes}{P} \times [-1, 1]$.

Now build Q by taking a balanced disjoint union of copies: $\deg(\overset{\boxtimes}{R}) \cdot \overset{\boxtimes}{L} \sqcup \deg(\overset{\boxtimes}{L}) \cdot \overset{\boxtimes}{R}$, and choose a one-to-one correspondence between elevations of A and B to attach the various edge spaces $\overset{\boxtimes}{P} \cong [-1, 1]$.




FIGURE 5.10. The fantasy proof of Lemma 5.3. (Read upwards.)




FIGURE 5.11. The real proof of Lemma 5.3. (Read upwards.)

CHAPTER 6

Wallspaces and their Dual Cube Complexes

We begin this chapter by describing the wallspaces which were introduced by Haglund-Paulin in [34]. We define wallspaces in Section 6.1 where we give a variant of their notion. In Section 6.2 we describe Sageev's construction of the cube complex dual to a wallspace, and in Section 6.3 we sketch a proof that it is CAT(0). In Section 6.4 we describe a variety of sources of wallspaces. While Sageev originally applied his construction to the wallspaces built using codimension-1 subgroups described in Section 6.5, the dual cube complex of an arbitrary cube complex is implicit in his construction and was described explicitly in [60, 17], and was already being employed in [59, 80].

6.1. Wallspaces

A *wall partition* of a space X is a decomposition $X = \overleftarrow{W} \cup \overrightarrow{W}$ into *halfspaces*. We let $W = \overleftarrow{W} \cap \overrightarrow{W}$. In our viewpoint, W is usually nonempty, and is referred to as a *wall*, and we refer to $\overleftarrow{W} - W$ and $\overrightarrow{W} - W$ as its *open halfspaces*. Usually X is a geodesic metric space, W is connected, $X - W$ has exactly two components, and no two wall partitions are associated with the same wall.

A *wallspace* (X, \mathcal{W}) is a space X together with a collection of wall partitions such that: Firstly, $\#(p, q) < \infty$ for all $p, q \in X$, where $\#(p, q)$ denotes the number of walls *separating* p, q in the sense that p, q lie in distinct open halfspaces. Secondly, each $p \in X$ *betwixts* finitely many walls, in the sense that $p \in \overleftarrow{W} \cap \overrightarrow{W}$. In other words, p lies on finitely many walls.

REMARK 6.1. Haglund-Paulin's original notion of wallspace in [34] has a more elegant definition: They define a wall to be a genuine partition of X and the finite betwixting property holds automatically. We have found it convenient to offer the more flexible variation on their idea described above.

EXAMPLE 6.2. Every CAT(0) cube complex is a wallspace – the walls correspond to hyperplanes.

6.2. The Dual CAT(0) Cube Complex

The CAT(0) cube complex C *dual* to a wallspace is defined as follows: A θ -cube v is an *orientation* on each wall – that is, $v(W)$ is a choice of one




FIGURE 6.1. Two wallspaces.

of the two halfspaces $\overleftarrow{W}, \overrightarrow{W}$ for each wall $W \in \mathcal{W}$. And moreover this choice must satisfy the following two properties:

- (1) No two walls are oriented away from each other (i.e. all chosen halfspaces intersect, so $v(W) \cap v(W') \neq \emptyset$ for $W, W' \in \mathcal{W}$.)
- (2) All but finitely many walls are oriented towards some (and hence any) point $x \in X$ (i.e. $x \in v(W)$ for almost all $W \in \mathcal{W}$).

A 1-cube joins two 0-cubes in C precisely when they differ on exactly one wall. An n -cube is attached exactly when its $(n - 1)$ -skeleton is present.




FIGURE 6.2. A 0-cube of C cannot have walls oriented away from each other, nor can it have infinitely many walls oriented away from a single point $x \in X$.




FIGURE 6.3. The dual cube complexes are: a tree, a square complex, and a 3-dimensional CAT(0) cube complex that is harder to depict.

G acts on the wallspace (X, \mathcal{W}) if it acts on X and permutes the walls (or technically, permutes the wall partitions). There is then a natural action of G on the dual cube complex C .




FIGURE 6.4. A system of n bi-infinite families of real lines in \mathbb{R}^2 has dual cube complex isomorphic to \tilde{T}^n .




FIGURE 6.5. Two different wallspaces on an infinite strip and their dual cube complexes.

6.3. C is CAT(0)

To see that C is nonpositively curved we verify that $\text{link}(v)$ is a flag complex for each $v \in C^0$. That $\text{link}(v)$ has no loops is because we attached a 2-cube when its 1-skeleton is there – and implicitly assumed that its skeleton is an embedded subcomplex. That $\text{link}(v)$ has no bigons is because two consecutive edges ab at the corner of a 2-cube uniquely determine the other two edges $a'b'$ where a', b' are associated to the same walls as a, b respectively and the 2-cube has boundary $aba'b'$. To see that $\text{link}(v)$ is a flag complex, we consider a collection of $(n + 1)$ pairwise adjacent vertices in $\text{link}(v)$ and observe that they bound an n -simplex. This n -simplex is a corner of the n -cube which corresponds to independently varying the orientations of the associated n distinct walls, while leaving the orientations of all other walls fixed.

To see that C is simply-connected, we consider a closed edge path $e_1 e_2 \cdots e_t$ and show that we can either remove a backtrack or else we can push a pair of edges $e_i e_j$ across a 2-cube to replace them by $e'_j e'_i$ and obtain a “lower complexity” homotopic edge path. Each edge corresponds to a wall, and since the path is closed, each wall must occur an even number of times. Our complexity will be the least number of edges between two edges associated to the same wall, as might occur for a subpath of our cyclic path of the form $e_p e_{p+1} e_{p+2} \cdots e_{p+q-1} e_{p+q}$ where e_p, e_{p+q} are associated to the same wall,

but no e_{p+i} is associated to the e_p wall for $1 \leq i < q$. When this complexity is zero there is a backtrack. Otherwise one argues that each intervening edge e_{p+i} is associated to a wall that crosses the wall of e_p – indeed all four possible orientations arise among the four vertices on e_p, e_{p+r} . Consequently, we can push across a square to replace $e_p e_{p+1}$ by $e'_{p+1} e'_p$ and hence have a lower complexity subpath of $e'_{p+1} e'_p e_{p+2} \cdots e_{p+q-1} e_{p+q}$ by omitting the first edge e'_{p+1} .

6.4. Some Examples

DEFINITION 6.3 (Wallspace from tracks). A *track* T in a 2-complex A is an embedded graph which is disjoint from the 0-skeleton, intersects each closed 1-cell in finitely many points, and intersects each closed 2-cell in finitely many properly embedded intervals: their boundary points map to the boundary of the 2-cell and their interior maps to its interior. The definition we give is slightly more general than the standard “Dunwoody track”. The key property of a track T is that it has a neighborhood $N(T)$ in A such that $N(T)$ is homeomorphic to a $(-1, 1)$ -bundle over T , with T corresponding to the section at 0.

Each component $\tilde{T} \subset \tilde{A}$ in the preimage of T provides a wall in \tilde{A} . Indeed, $N(\tilde{T}) \cong (-1, 1) \times T$, so T has a 2-sided neighborhood, and thus $\tilde{A} - \tilde{T}$ consists of two open halfspaces meeting along \tilde{T} .

Thus a locally finite collection of tracks in A naturally makes \tilde{A} into a wallspace. A great many of the examples that we describe arise in this form, or some higher dimensional variation.

EXAMPLE 6.4 (A tricky wallspace from tracks). Let A be the standard 2-complex of the following presentation:

$$\langle a, b, u, c, d, v \mid ab^{-1}u^{-1}, b^{-1}au, [b, c], cd^{-1}v^{-1}, d^{-1}cv^{-1} \rangle.$$

Let B be the standard 2-complex of:

$$\langle a, b, u, c, d, v, t \mid ab^{-1}u^{-1}, b^{-1}au, [b, c], cd^{-1}v, d^{-1}cv, a^t = a, d^t = d, (uv)^t = vu \rangle$$

We regard A as a subcomplex of B and note that the 2-cells of A occur on the left of Figure 6.6 which (also) illustrates the 2-cells of B . The group $\pi_1 B \times \mathbb{Z}$ is isomorphic to the pure mapping class group of the five punctured sphere [23].




FIGURE 6.6. Tracks in the 2-complex B .

There are five tracks corresponding to the embedded graphs that are transverse to B^1 illustrated in Figure 6.6. Our wallspace is the 2-complex

\tilde{B} together with the graphs that are components of the preimages of these graphs in \tilde{B} .

I had incorrectly hoped to show that $\pi_1 B$ acts freely on this wallspace, but Piotr Przytycki found an error in my proof – there is actually a noncyclic free group acting trivially on the dual cube complex.

EXAMPLE 6.5 (Examples from Rhombi). Any CAT(0) metric space that is a 2-complex built from rhombi has a wallspace structure investigated in [47]. The walls arise from trees that cut through the rhombi along midcells. This structure is consistent with de Bruijn’s beautiful treatment of the Penrose tilings in [21], which remarkably anticipates Sageev’s dual cube complex construction. The carriers of walls in a Penrose tiling by rhombi are called *worms*. See Figure 6.7.




FIGURE 6.7. The cube complex dual to the wallspace on a Penrose tiling is a copy of \mathbb{R}^5 since there are five bi-infinite families of parallel walls. This explains how de Bruijn interpreted the Penrose tiling as a 2-flat traveling in a copy of \mathbb{R}^5 .

The following example gives a taste of the material described in greater generality in Chapter 10. We refer the reader to Definition 9.1 for the notion of $C'(\frac{1}{6})$.

EXAMPLE 6.6 (Wallspaces from small-cancellation theory). Every simply connected $C'(\frac{1}{6})$ complex \tilde{X} is a wallspace: The walls are graphs that intersect 2-cells and 1-cells of \tilde{X} in *midcells*. We first subdivide the 1-skeleton so that all 2-cells have an even circumference. Some facts about these walls:




FIGURE 6.8. A wall in a $C'(\frac{1}{6})$ small-cancellation complex on the left, and the 4 midcells of an octagon on the right.

- (1) Each immersed wall in \tilde{X} embeds,
- (2) is 2-sided (as it is locally 2-sided and $H^1(\tilde{X}) = 0$),
- (3) is a tree – or possibly a multi-tree when \tilde{X} has duplicate 2-cells as often arises from presentations when relators are proper powers;
- (4) and has a convex *carrier* – i.e. the smallest subcomplex containing the wall is a convex subcomplex.

The above properties are proven using Greendlinger’s lemma and the ladder theorem – see Section 9.2.1.




FIGURE 6.9. A wall in a $C'(\frac{1}{6})$ complex is a multi-tree.

6.5. Wallspaces from Codimension-1 Subgroups

Let G be a finitely generated group. $H \subset G$ is a *codimension-1* subgroup if for some $r > 0$, there are two or more H -orbits of components K_i in $\Gamma - \mathcal{N}_r(H)$ that are *deep* in the sense that $K_i \not\subset \mathcal{N}_s(H)$ for any $s > 0$. This doesn’t depend on the choice of Cayley graph $\Gamma = \Gamma(G, S)$.

EXAMPLE 6.7. The following examples are illustrated in Figure 6.10:

- (1) $\mathbb{Z}^n \subset \mathbb{Z}^{n+1}$
- (2) $\pi_1 S^1 \curvearrowright \pi_1 M^2$
- (3) $\pi_1 M^2 \curvearrowright \pi_1 M^3$ with M^2, M^3 aspherical.
- (4) $C \subset A *_C B$ with $C \not\subset A, C \not\subset B$.




FIGURE 6.10. Some codimension-1 subgroups.

To obtain a wall from a codimension-1 subgroup, first let $\{K_i\}_{i \in I}$ denote a complete set of representatives of the H -orbits in $\Gamma - \mathcal{N}_r(H)$. Partition the index set into $I = \overleftarrow{I} \sqcup \overrightarrow{I}$. Let $W = \mathcal{N}_r(H)$ and $\overleftarrow{W} = (W \cup \bigcup_{i \in \overleftarrow{I}} HK_i)$ and $\overrightarrow{W} = (W \cup \bigcup_{i \in \overrightarrow{I}} HK_i)$. This can be done for any subgroup, but we have in mind the case where $\mathcal{N}_r(H)$ is connected and separates Γ , and where the translates of at least one deep component is assigned to each halfspace of W .

REMARK 6.8. $\text{Stabilizer}(W)$ and $\text{Stabilizer}(\overleftarrow{W})$ might only be commensurable with H .

We obtain a wallspace with a G -action by letting $X = \Gamma$ and letting $\mathcal{W} = \{(g\overleftarrow{W}, g\overrightarrow{W})\}_{g \in G}$. This construction can also be implemented with finitely many subgroups H_j . The reader should check that $\#(p, q) < \infty$ and that any point p betwixts finitely many walls.

EXAMPLE 6.9. Any finitely generated infinite index subgroup of F_2 is codimension-1. However there are infinitely many choices of walls for a given subgroup. In fact, there are so many choices that for any finitely generated infinite index subgroup, one can choose a sufficiently complicated wall so that the action of F_2 on the associated dual cube complex is free [75].

EXAMPLE 6.10. For a closed hyperbolic M^3 , Kahn-Markovic provide many immersed incompressible surfaces $S_i \looparrowright M$ [48]. Choosing finitely many and applying the above construction gives a wallspace $(\widetilde{M}^3, \{g\widetilde{S}_i\})$. We then obtain an action of $\pi_1 M^3$ on its dual CAT(0) cube complex. We describe a “soft approach” using $\partial\widetilde{M}^3$ to choose sufficiently many Kahn-Markovic surfaces in order to obtain a free action in [7]. This approach sidesteps the linear separation verifications which can be arduous.

CHAPTER 7

Finiteness properties of the dual cube complex

Finding codimension-1 subgroups of G and choosing associated walls allows us to produce an action on the dual cube complex but it does not necessarily guarantee that this action will be illuminating and that the complex won't be pathological. In this chapter we investigate conditions on G and the wall space ensuring that the dual will have valuable finiteness properties. The broad theme is that “sufficiently many” walls leads to properness of the group action and negative curvature coupled with metric tameness of the walls leads to cocompactness of the group action. We begin by discussing the cubes of the dual in Section 7.1. We then describe the “bounded packing property” for a subgroup of G in Section 7.2, and show that when the wall stabilizers have the bounded packing property then the dual cube complex is finite dimensional. Following Sageev, in Section 7.3 we explain the key point that the action on the dual cube complex is cocompact when G is hyperbolic and the wall stabilizers are quasiconvex. In Section 7.4 we describe a generalization of this to relative cocompactness of the action on the dual when the wall stabilizers are relatively quasiconvex and G is relatively hyperbolic. The reader focused on the hyperbolic setting should skip Section 7.4. We turn to properness of the action in Section 7.5, and show that G acts properly on the dual if the number of walls separating g_1x, g_2x grows as $d(g_1, g_2)$ grows. Finally, in Section 7.6 we describe the “cut-wall” criterion for showing that G acts with finite stabilizers, that is sometimes easy to verify.

7.1. The Cubes of C :

Two walls $(\overleftarrow{W}_1, \overrightarrow{W}_1), (\overleftarrow{W}_2, \overrightarrow{W}_2)$ *cross* if all four quarterspaces are nonempty:

$$\overleftarrow{W}_1 \cap \overleftarrow{W}_2, \overleftarrow{W}_1 \cap \overrightarrow{W}_2, \overrightarrow{W}_1 \cap \overleftarrow{W}_2, \overrightarrow{W}_1 \cap \overrightarrow{W}_2 \neq \emptyset$$

In many cases, walls cross precisely when $W_1 \cap W_2 \neq \emptyset$. Indeed, this holds for “geometric wallspaces” where the space X is a topological space, each wall W is closed and connected, and $X - W$ consists of exactly two components which are precisely the open halfspaces of W .

Each n -cube of the dual cube complex C corresponds to a cardinality n collection of pairwise crossing walls together with an orientation on all other walls towards these (i.e. to intersect both their halfspaces), and such that all but finitely many walls are oriented towards a basepoint $x \in X$.




FIGURE 7.1. The above wallspace contains four bold walls, one dotted wall, and eight other walls. The four bold independently orientable walls together with orientations on all the other walls determine a 4-cube in the dual cube complex. That 4-cube is not maximal, since the dotted wall crosses each of these four.

EXAMPLE 7.1. Following Definition 6.3, if walls in X arise from finite collections of embedded tracks in X , say T_1, \dots, T_d , then $\dim(C) \leq d$.

7.2. The Bounded Packing Property and Finite Dimensionality:

DEFINITION 7.2. $H \subset G$ has *bounded packing* if for each D there exists $N = N(D)$ bounding the cardinality of collections $\{g_1H, \dots, g_rH\}$ of pairwise D -close left cosets.




FIGURE 7.2. A collection of four pairwise D -close cosets of H .

EXAMPLE 7.3. Finite subgroups, normal subgroups, and separable subgroups [83] have bounded packing.

It was pointed out in [66], that the following result is implicit in [27]. An explicit argument using a bounded height argument is given in [42].

THEOREM 7.4 (Sageev). *If H is a quasiconvex subgroup of a hyperbolic group G then H has bounded packing.*

We abstracted the above notion of bounded packing in [42], and generalized Theorem 7.4 to:

THEOREM 7.5. *If G is hyperbolic relative to $\{P_i\}$ and each P_i has bounded packing relative to its subgroups then G has bounded packing relative to quasiconvex subgroups.*

Since cubes of the dual cube complex $C(X)$ correspond to collections of pairwise crossing walls in X , the bounded packing property leads to finite dimensionality as follows [40]:

THEOREM 7.6. *Suppose the finitely generated group G has bounded packing with respect to finitely generated subgroups H_1, \dots, H_r . Let $C(X)$ be the dual CAT(0) cube complex with respect to a system of H_i -walls in $X = \Gamma(G)$. Then $C(X)$ is finite dimensional.*

EXAMPLE 7.7. Rubinstein-Wang found a graph manifold M and immersed incompressible surface $S \looparrowright M$ such that all translates of \tilde{S} intersect in the sense that $g_1\tilde{S} \cap g_2\tilde{S} \neq \emptyset$ in \tilde{M} . Thus $\dim(C) = \infty$, and the dual cube complex C is actually an infinite cube.

The Rubinstein-Wang example $\pi_1 M$ does not have bounded packing with respect to the surface subgroup $\pi_1 S$. Thompson's group T has a subgroup without bounded packing: Namely, the stabilizer of a hyperplane in the action of T on the infinite-dimensional CAT(0) cube complex discovered by Farley [25].

7.3. Cocompactness in the Hyperbolic Case

THEOREM 7.8 (Sageev 97). *If \tilde{X} is δ -hyperbolic and G acts properly and cocompactly on \tilde{X} , and there are finitely many G -orbits of walls, and each wall W is quasiconvex (i.e. $\text{Stabilizer}(W)$ is quasiconvex). Then G acts cocompactly on the dual cube complex $C = C(X)$.*

PROOF. There is an upper bound on the dimension of C because of Theorem 7.4 and Theorem 7.6. Consequently we can apply the following key lemma to see that there are finitely many G -orbits of maximal cubes. \square

The following was proven in [59]:

LEMMA 7.9 (bounded in-center). *For each κ, D, δ there exists R such that if W_1, \dots, W_t is a collection of pairwise D -close κ -quasiconvex subspaces of the δ -hyperbolic space X , then there is a point $p \in X$ such that: $d(W_i, p) \leq R$ for all i .*

7.4. Relative Cocompactness in the Relatively Hyperbolic Case

The following theorem is the main result in [40]:

THEOREM 7.10. *Let G be hyperbolic relative to $\{P_i\}$. Suppose G acts properly and cocompactly on \tilde{X} and each P_i cocompactly stabilizes some $\tilde{F}_i \subset \tilde{X}$. Suppose \tilde{X} is a wallspace and each wall has relatively quasiconvex stabilizer. There exists a compact subcomplex $K \subset C(\tilde{X})$ and convex P_i -invariant subcomplexes $C(\tilde{F}_i) \subset C(\tilde{X})$ such that:*

- (1) $C(\tilde{X}) = GK \cup \bigcup_i GC(\tilde{F}_i)$
- (2) $g_i C(\tilde{F}_i) \cap g_j C(\tilde{F}_j) \subset GK$ unless $\tilde{F}_i = \tilde{F}_j$ and $g_j^{-1}g_i \in \text{Stabilizer}(\tilde{F}_i)$.
- (3) Moreover, each $C(\tilde{F}_i) \cap GK$ lies in a finite neighborhood of $P_i K$.

It is sometimes the case that each $C(\tilde{F}_i)$ is the dual cube complex of an induced wallspace structure on \tilde{F}_i whose walls are intersections of \tilde{F}_i with walls of \tilde{X} . Indeed, this is the case when the following property holds: for pairs of walls W, W' of \tilde{X} that intersect \tilde{F}_i , if W, W' cross then $(W \cap \tilde{F}_i), (W' \cap \tilde{F}_i)$ cross. The general case involves the dual cube complex $C(\tilde{F}_i)$ of an induced “hemiwallspace” and this is studied in detail in [40].




FIGURE 7.3. A relatively cocompact cubulation

We give the following definition to permit us to describe Example 7.12

DEFINITION 7.11 (B(6)). \tilde{X} satisfies the $B(6)$ *small-cancellation condition* if for each 2-cell R , the concatenation of three consecutive pieces $P_1 P_2 P_3$ in $\partial_p R$ satisfies $|P_1 P_2 P_3| \leq \frac{1}{2} |\partial_p R|$. Note that \tilde{X} satisfies $B(6)$ when each piece P satisfies $|P| \leq \frac{1}{6} |\partial_p R|$, and in particular $C'(\frac{1}{6})$ is $B(6)$. We discuss small-cancellation in Section 9.2.1.

EXAMPLE 7.12. A *honeycomb* is a copy of the hexagonal tiling of the Euclidean plane, possibly with some of its 1-cells subdivided. See Figure 7.4.




FIGURE 7.4. A honeycomb

Let X be the standard 2-complex of:

$$\langle a, b, c, d, e \mid abca^{-1}b^{-1}c^{-1}, cdec^{-1}d^{-1}e^{-1}, adbea^{-1}d^{-1}b^{-1}e^{-1} \rangle$$

Then \tilde{X} is $B(6)$ and contains a collection of honeycombs that are isolated in the sense that they have uniformly bounded overlap. The group $\pi_1 X$ is

hyperbolic relative to the stabilizers of these honeycombs which are virtually \mathbb{Z}^2 . These honeycombs play the role of \tilde{F}_i in Theorem 7.10.

Another $B(6)$ example that arises naturally is the following presentation for the fundamental group of the figure 8 knot complement: $\langle a, b \mid baa = abb \rangle$.

When G is hyperbolic relative to f.g. virtually abelian subgroups $\{P_i\}$, the dual cube complexes $C(\tilde{F}_i)$ are *quasiflats* that are quasi-isometric to \mathbb{E}^{n_i} . G is said to act *cosparcely* on $C(\tilde{X})$ in this case and $G \backslash C$ is *sparse*.

In conclusion, the most important case covered by Theorem 7.10 is where G is hyperbolic relative to virtually abelian subgroups $\{P_i\}$, in which case $C(\tilde{X})$ looks like a finite cubular neighborhood of G , with some cubulated copies of \mathbb{E}^{n_i} hanging off. These are the $C(\tilde{F}_i)$ translates. They are coarsely disjoint and attached along coarse copies of P_i cosets in G .

7.5. Properness of the G Action on $C(\tilde{X})$

In a Haglund-Paulin wallspace (\tilde{X}, \mathcal{W}) , for each point $p \in \tilde{X}$ there is a *canonical 0-cube* V_p defined by orienting all walls towards p . In general V_p is a *canonical cube* whose independent walls are those betwixted by v .



FIGURE 7.5. The canonical cube V_p is a 0-cube on the left and a 2-cube on the right.

There is a G -equivariant map $\tilde{X} \rightarrow \text{Cubes}(C(\tilde{X}))$ defined by $p \mapsto V_p$. Note that $\#(p, q) = d_C(V_p, V_q)$ for $p, q \in \tilde{X}$. Consequently, to understand properness of the action on C , we are led to examine the relationship between $\#$ and $d_{\tilde{X}}$.

LEMMA 7.13. G acts metrically properly on C if and only if $\#(p, gp) \rightarrow \infty$ as $g \rightarrow \infty$.

An action of a group G on a metric space C is *metrically proper* if for each $r > 0$ and each $c \in C$ the set $\{g : gB_r(c) \cap B_r(c) \neq \emptyset\}$ is finite. When C has compact closed balls, this is the same as acting *properly* in the sense $\{g : gS \cap S \neq \emptyset\}$ is finite whenever $S \subset C$ is compact.

In practice we often verify that $\#(p, gp) \rightarrow \infty$ by verifying the sufficient condition of *linear separation* which states that: there exists $k_1, k_2 > 0$ such that $\#(p, q) \geq k_1 d_{\tilde{X}}(p, q) - k_2$ for all $p, q \in \tilde{X}$.

This is normally done by finding $L > 0$ such that for any geodesic $\gamma \rightarrow \tilde{X}$, each length L subsegment $\gamma' \subset \gamma$ is cut at a single point by a wall W such that W doesn't cut γ anywhere else.



FIGURE 7.6. Verifying linear separation as on the left, can be a messy affair as on the right. Moving along the geodesic γ one passes through many walls locally, but one might pass through only a few walls globally.

While this can be tricky to verify in general, there are situations where it is routine. For instance:

LEMMA 7.14. *If \tilde{X} is a metric $CAT(0)$ space, and each wall is a convex hyperplane (i.e. codimension-1 subspace), and components of $\tilde{X} - \cup_{W \in \mathcal{W}} W$ have diameter uniformly bounded by D , then \tilde{X} satisfies the linear separation property.*

PROOF. $\#(p, q)$ equals the number of points of intersection between γ and the collection of walls $\{W\}$, and successive points are within D of each other. \square




FIGURE 7.7. Linear separation is immediate when the walls are convex hyperplanes in a $CAT(0)$ metric space, and the complementary regions have uniformly bounded diameter.

We can now record the commensurability invariance of being cubulated:

LEMMA 7.15. *Let G be a finitely generated group that contains a finite index subgroup G' that acts metrically properly on a $CAT(0)$ cube complex. Then G acts metrically properly on a $CAT(0)$ cube complex.*

Moreover, if G is hyperbolic and G' acts properly and cocompactly on a $CAT(0)$ cube complex then so does G .

PROOF. Observe that the finite index subgroup G' of G gives a sufficient system of quasiconvex walls in the Cayley graph of G . By Theorem 7.8 and Lemma 7.13 the group G again acts properly and cocompactly on the dual $CAT(0)$ cube complex \tilde{X} . \square

Without hyperbolicity, there are examples where the finite index subgroup G' is the fundamental group of a compact nonpositively curved cube complex, but the torsion-free group G is not. We refer to [32] for a discussion of this failure for crystallographic groups.

7.6. The Cut-Wall Criterion for Properness

An *axis* \mathbb{R}_g for an element g acting on \tilde{X} is a g -invariant copy of \mathbb{R} in \tilde{X} . A *cut-wall* for g is a wall W such that $g^n W \cap \mathbb{R}_g = \{n\}$ for all $n \in \mathbb{Z}$.




FIGURE 7.8. A cut-wall for g .

When G has no infinite torsion-subgroup the following condition suffices to prove properness. It is a variation of similar conditions examined in [40].

LEMMA 7.16. *If each infinite order element $g \in G$ has a cut-wall then the action of G on C has the property that the stabilizer of each point is a torsion group.*

PROOF. Suppose $hc = c$ for some point $c \in C$, then $h^{d^l}v = v$ where v is a 0-cube on the smallest d -cube containing c . Let $g = h^{d^l}$ and consider an axis \mathbb{R}_g and cut-wall W for g . That $gv = v$ means that the orientations of walls $\{g^n W\}$ is preserved by the action of g . Moreover $\{g^{2^n} W\}$ are all oriented towards $+\infty$ or all towards $-\infty$ as in Figure 7.9. This contradicts the 2nd axiom defining 0-cubes of C , since infinitely many walls are oriented away from $0 \in \mathbb{R}_g$. \square




FIGURE 7.9. An infinite sequence of distinct walls that are directed away from a point on the line.

EXAMPLE 7.17. The cut-wall criterion applies to metric CAT(0) wallspaces with convex hyperplane walls. But now we need not assume that complementary regions are uniformly bounded, just that each is bounded, and that each element has an axis. (We note that in general it is possible to have parabolic elements with no axis.)

EXAMPLE 7.18. The cut-wall criterion applies to wallspaces built from trees of wallspace with a cut-wall criterion satisfied by each vertex space and with connected intersections of walls with vertex spaces. See Chapter 8.1.

CHAPTER 8

Cubulating Malnormal Graphs of Cubulated Groups

The goal of this chapter is to describe how to cubulate a hyperbolic group G that splits as a graph of cubulated hyperbolic groups. Our exposition sketches results presented in [43], and we refer there for the details. The edge groups of G immediately provide a substantial amount of the codimension-1 subgroups needed to effectively apply the dual cube complex construction. Indeed, any element of G that is “hyperbolic” with respect to the action on the Bass-Serre tree is cut by an edge group. It remains to find codimension-1 subgroups that are sufficient from the viewpoint of each vertex group G_v of G . We do this by “extending” sufficiently many codimension-1 subgroups of each G_v into codimension-1 subgroups of G . This is illustrated for a simplistic non-hyperbolic example in Section 8.1 which aims to give the reader some intuition about how the immersed walls in a base space develop to walls in the tree of spaces that is its universal cover. Two techniques are employed to deal with more complicated examples: “extensions” and “turns”. In Section 8.2 we describe how a codimension-1 subgroup in an edge group extends into a codimension-1 subgroup in a neighboring vertex group. In Section 8.3 we describe how a codimension-1 subgroup in a vertex group can be “turned around” within an adjacent edge group so that it returns and is absorbed in the original vertex group. In Section 8.4 we combine these techniques to prove the main result by constructing sufficiently many quasiconvex codimension-1 subgroups of the group G .

8.1. A Wallspace for an Easy Non-Hyperbolic Group

To motivate some aspects of the constructions described later in this section, we will examine a system of walls in a group that is not hyperbolic but splits as an amalgamated product $\mathbb{Z}^2 *_\mathbb{Z} \mathbb{Z}^2$. In [74], we have provided a detailed case-study of the cubability of any group G splitting as a graph of groups with \mathbb{Z}^2 vertex groups with \mathbb{Z} edge groups. Many such groups that admit free actions on CAT(0) cube complexes. However, though G acts freely, the dual cube complex is generally not locally finite nor even finite dimensional.

Let $\langle a, b \rangle$ and $\langle c, d \rangle$ denote copies of \mathbb{Z}^2 and consider a group G that splits as the following amalgamated free product.

$$\langle a, b \rangle *_{a=c^2} \langle c, d \rangle$$

Observe that G is the fundamental group of a compact nonpositively curved cube complex obtained by starting with appropriate cubical subdivisions of standard tori, and then gluing a cubical cylinder using attaching maps that are local isometries. However, to motivate our later constructions, we provide an alternate cubulation dual to a wallspace provided by a system of codimension-1 subgroups.

Let X denote a graph of spaces corresponding to the splitting of G : so X has two vertex spaces, each of which is homeomorphic to a torus, and X has a single edge space homeomorphic to a cylinder. And this cylinder is attached on each side along a closed curve representing a and c^2 respectively.

There are three orbits of walls in \tilde{X} . Let S denote a circle cutting the cylinder $C \times I$ of X in half. The first type of wall in \tilde{X} is a line consisting of any lift of the universal cover of \tilde{S} . Its two halfspaces correspond to the parts of \tilde{X} lying on opposite sides of the strip $\tilde{C} \times I$. The second and third types of walls are trees corresponding to lift of the two graphs graphs indicated in Figure 8.1. It is convenient that in this simple case, these underlying “immersed walls” are embedded subspaces, but one can imagine more elaborate examples where the underlying immersed walls do not embed in the base space.




FIGURE 8.1. The green, red, and blue graphs in the space X provide walls in \tilde{X} corresponding to components of their preimages.

Let us consider the resulting walls in \tilde{X} of the second and third type. Observe that \tilde{X} is a tree of spaces whose vertex spaces are planes and whose edge spaces are strips. One of our walls intersects a vertex space in a line, and then proceeds through adjacent strips in segments, after which it becomes a line in the next vertex space, and so forth. In particular, we note that a wall intersects a vertex space in either a line or in \emptyset .

One sees that these are indeed walls since \tilde{X} is simply connected, and an open neighborhood of each wall W is of the form $W \times (-1, 1)$ with W identified with $W \times \{0\}$.

To see that G acts freely on the dual CAT(0) cube complex, we note that the axis criterion can be applied. Any element $g \in G$ is either elliptic or hyperbolic with respect to the Bass-Serre tree. If it is hyperbolic, then a wall of the first type serves as a cut-wall for g . If it is elliptic, then either a blue or red line within the plane stabilized by g will serve as a cut-wall for an axis of g within that plane. And this cut-wall functions as a cut-wall within all of \tilde{X} since walls intersect a vertex space in at most one line.




FIGURE 8.2. Three walls in \tilde{X} .

8.2. Extending Walls

An H -wall in a group G is a finite neighborhood $\mathcal{N}_r(H)$ of H in $\Gamma = \Gamma(G)$ and a decomposition $\Gamma = \overleftarrow{H} \cup \overrightarrow{H}$ with $\overleftarrow{H} \cap \overrightarrow{H} = \mathcal{N}_r(H)$ – so each component of $\Gamma - \mathcal{N}_r(H)$ lies in \overleftarrow{H} or \overrightarrow{H} .

As discussed in Section 6.5, the motivating case is where $\Gamma - \mathcal{N}_r(H)$ has two components, each of which is deep.




FIGURE 8.3. An H -wall.

Let $G \subset G'$. An H -wall in G extends to an H' wall in G' if $H = H' \cap G$ and $\overleftarrow{H} \subset \mathcal{N}_r(\overleftarrow{H}')$ and $\overrightarrow{H} \subset \mathcal{N}_r(\overrightarrow{H}')$ for some r .




FIGURE 8.4. Some extensions of H -walls. The second and third examples are indicated using fundamental groups of subspaces.

THEOREM 8.1. *Let X be a compact special cube complex, and let $G = \pi_1 X$ be hyperbolic. Let $H \subset G$ be a quasiconvex subgroup of G . Let A be a quasiconvex subgroup of H . Then any A -wall in H extends to a B -wall of G .*

PROOF. By Theorem 3.23, let $Y_H \rightarrow X$ be a compact core for H . Let $Y_A \rightarrow Y_H$ be a compact core for A , such that, moreover, the two sides of the A -wall coarsely lie in distinct halfspaces $\tilde{Y}_H - \tilde{Y}_A$ (partitioned accordingly).

Let $\hat{Y}_H \rightarrow Y_H$ be a finite cover such that Y_A embeds in \hat{Y}_H .

Consider $\mathcal{C}(\hat{Y}_H \rightarrow X)$ and the canonical retraction $\mathcal{C}(\hat{Y}_H \rightarrow X) \rightarrow \hat{Y}_H$. The preimage of the locally convex subcomplex $Y_A \subset \hat{Y}_H$ is a locally convex subcomplex $Y_B \subset \mathcal{C}(\hat{Y}_H \rightarrow X)$, and let $B = \pi_1 Y_B$. The B -halfspaces are the preimages of the A -halfspaces under the retraction $\tilde{X} \rightarrow \tilde{Y}_H$ induced by $\mathcal{C}(\hat{Y}_H \rightarrow X) \rightarrow \hat{Y}_H$. \square




FIGURE 8.5. Using Canonical Completion and Retraction to Extend an H -wall.

8.3. Constructing Turns

We now describe how to construct “turns” which are immersed walls in an edge space which permit an immersed wall in a vertex space to enter the edge space, then wander around for a while, and then return to the vertex space parallel to the way that they entered. Before describing how to do this in general, we illustrate the situation with a simple example of a surface vertex space X_A and cylinder edge space $X_C \times I$ where the immersed wall consists of a circle $X'_H \rightarrow X_A$ in the surface, that needs to extend into an immersed wall in the cylinder. This is illustrated in Figure 8.7.

We will now work with *immersed walls* which are locally 2-sided, π_1 -injective immersions whose universal covers lift to actual walls. We will




FIGURE 8.6. Wall extension using canonical completion and retraction.




FIGURE 8.7. Turn in a cylindrical edge space

use notation like $X_C \looparrowright X_A$ to denote such an immersed wall, with the understanding that $A = \pi_1 X_A$ and $C = \pi_1 X_C$ etc.

Consider the general scenario of a vertex space X_A with edge space $X_C \times I$ attached to it. We will describe a recipe to mimic the construction of a turn that we gave above in the intuitive setting of a surface and cylinder.

An H' -wall of A induces an H -wall of C where $H = C \cap H'$. Let us denote this by $X_H \rightarrow X_C$. (We are interested in the case where C is not coarsely contained on one side of the H' -wall, in which case the induced wall would be a trivial wall – which we essentially ignore in our applications.) We use separability of H in C to choose a cover $\widehat{X}_C \rightarrow X_C$ that has large girth (relative to H'). We now cut-and-paste to build the turn: Consider the cover $\widehat{X}_C \times [-1, 1] \rightarrow X_C \times [-1, 1]$ and the wall $\widehat{X}_C \times \{0\}$ that resides inside. We cut-and-paste $X_H \times [-1, 0]$ with $\widehat{X}_C \times \{0\}$ as illustrated in Figure 8.8.

8.4. Cubulating Malnormal Amalgams

THEOREM 8.2. *Let $G = A *_C B$ or $A *_C B'$. Then G acts properly and cocompactly on a $CAT(0)$ cube complex provided the following hold:*

- (1) G is hyperbolic.
- (2) C is quasiconvex and almost malnormal in G .




FIGURE 8.8. Cutting and pasting to form the turn.




FIGURE 8.9. Building turns.

- (3a) A, B are virtually π_1 of a nonpositively curved cube complex.
- (3b) C has separable quasiconvex subgroups. [and likewise for C']
- (3c) There is a sufficient system of quasiconvex H_i -walls in C that extend to H'_i -walls in A, B . [and likewise for C']

We note that Conditions (3a), (3b) and (3c) follow from the following:

- (3) A, B are each virtually π_1 of a compact special cube complex.

We describe a relatively hyperbolic generalization of Theorem 8.2 in Theorem 8.5.

Two instances where Theorem 8.2 takes a simplified form are as follows:

EXAMPLE 8.3 (\mathbb{Z} edge group). Let $G = A *_{\mathbb{Z}} B$ or $A *_{\mathbb{Z}^t = \mathbb{Z}'} B$ where A, B are fundamental groups of nonpositively curved cube complexes and \mathbb{Z} is a malnormal subgroup of G . Then G is the fundamental group of a nonpositively curved cube complex. This is because Conditions (3b) and (3c) automatically hold when $C \cong \mathbb{Z}$.

EXAMPLE 8.4 (Free vertex groups). $F_2 *_{\langle U, V \rangle = \langle U', V' \rangle} F'_2$ is cubulated whenever $\langle U, V \rangle$ is malnormal in F_2 and $\langle U', V' \rangle$ is malnormal in F'_2 . (In fact, malnormality on one side suffices.)

SKETCH OF PROOF OF THEOREM 8.2. The plan for proving Theorem 8.2 is to produce sufficiently many quasiconvex codimension-1 subgroups, and

then obtain a proper action on the dual cube complex. As discussed in Chapter 7.5, “sufficiently many” leads to a proper action on the dual cube complex and quasiconvexity leads to its cocompactness.




FIGURE 8.10. Start with a sufficient set of immersed walls in X_C , and extend these into immersed walls in X_A and X_B




FIGURE 8.11. Consider the original sufficient immersed walls (arising from the cubulations) of X_A and X_B and add those as well.




FIGURE 8.12. Double all these immersed walls, and turn those immersed walls that enter $X_C \times I$ on each side.

The malnormality of C and a choice of sufficiently “large” turns ensures that the new immersed walls correspond to quasi-isometrically embedded subgroups and lift to genuine embedded walls in \tilde{X} . The quasiconvexity within the vertex groups is one of the technical points verified in [43], and then the quasiconvexity within all of $\pi_1 X$ is ensured by a criterion in [9].




FIGURE 8.13. Add a vertical wall cutting through each $X_C \times I$

These new walls have the property that they intersect each collared vertex space \tilde{X}_A^+ and \tilde{X}_B^+ in a single component or in \emptyset . The linear separation property is verified to make sure that elliptic elements in the vertex groups act freely, and the cut-criterion using the vertical walls ensures that the hyperbolic elements act freely. \square

We employ the following relatively hyperbolic version of Theorem 8.2. Its proof is similar but requires more care.

THEOREM 8.5. *Let $G = L *_E R$ be hyperbolic relative to f.g. virtually abelian subgroups. Suppose L, R are quasiconvex and are fundamental groups of compact [sparse] nonpositively curved cube complexes that are virtually special. Suppose $E \subset L$ is a parabolic and almost malnormal. Then G acts properly and cocompactly [cosparsely] on a $CAT(0)$ cube complex.*

We actually do not require above that the cube complex for R is virtually special – just the one for L . The following variant of a special case of Theorem 8.5 can be proven easily and independently:

COROLLARY 8.6. *Let G be hyperbolic relative to f.g. free abelian subgroups. Suppose G splits as $L *_E R$ with E quasiconvex in G . Suppose $R = E \times \mathbb{Z}^n$ and that E is a full quasiconvex subgroup of L . Suppose L is cocompactly [cosparsely] cubulated. Then G acts properly and cocompactly [cosparsely] on a $CAT(0)$ cube complex.*

CHAPTER 9

Cubical Small Cancellation Theory

Following Gromov, who introduced his “rotating-family” method in [29], a number of approaches have been given for obtaining a hyperbolic quotient $\bar{G} = G/\langle\langle H_1, \dots, H_r \rangle\rangle$ where H_i are appropriate quasiconvex subgroups of the hyperbolic group G . Most notably, this idea has been developed in [62] and [30]. We describe an approach within the cubical category: so $G = \pi_1 X$ where X is a nonpositively curved cube complex, and the H_i are represented by local isometries $Y_i \rightarrow X$ having large systoles compared to their overlaps. While the approach is limited in its applicability, it is more combinatorial than geometric, and the conclusions one obtains are concrete and explicit. The approach also represents a faithful generalizations of the classical small-cancellation theory. When X is 1-dimensional, this approach was already developed in early work of Rips. See, for instance [64].

We begin by defining cubical presentations in Section 9.1, and roughly indicating how to think of disk diagrams in this setting. In Section 9.2 we recall the fundamental theorem of small-cancellation theory in the classical case, and then describe it in the cubical case. We note that statements requiring less stringent small-cancellation hypotheses are proven in [76]. We review the Combinatorial Gauss-Bonnet Theorem in Section 9.3. We sketch a proof of the fundamental theorem of small-cancellation theory in the classical case in Section 9.4. In Section 9.5, we generalize the notion of reduced diagram from the classical case, and then define the $C'(\frac{1}{n})$ small-cancellation condition in the cubical case. We show how to construct examples of $C'(\frac{1}{24})$ cubical presentations in Section 9.6. In Section 9.7 we describe how to decompose a disk diagram $D \rightarrow X^*$ into conecells, rectangles, and shards. This allows us to assign angles and apply the Combinatorial Gauss-Bonnet Theorem to prove the fundamental theorem of small cancellation theory in the cubical case following the same method as in the classical case.

Since the intention of this chapter is to generalize the classical small-cancellation $C'(\frac{1}{n})$ condition, let us conclude the introduction to this chapter by recalling it:

DEFINITION 9.1 (Classical $C'(\frac{1}{n})$). Let X be a 2-complex. A nontrivial combinatorial path $P \rightarrow X$ is a *piece* if the lift $\tilde{P} \rightarrow \tilde{X}$ is a subpath of two distinct cycles in \tilde{X}^1 that are the boundary cycles of 2-cells of \tilde{X} . A 2-complex X satisfies the $C'(\frac{1}{n})$ *metric small-cancellation* condition if $|P| <$

$\frac{1}{n}|\partial_p R|$ whenever P is a piece that is a subpath of the boundary path ∂_p of the 2-cell R .

We often obtain the 2-complex X as the standard 2-complex of a presentation $\langle a, b, \dots, | W_1, W_2, \dots \rangle$. In this case, P is a piece in X provided that the corresponding word in the generators appears in two distinct ways as a subword among the cyclic words $W_i^{\pm 1}$. The meaning of “distinct” is subtle when some $W_i = V_i^{n_i}$ is a proper power or when some relator appears multiple times in the presentation. Two occurrences of P are not regarded as distinct if they differ by an automorphism of a relator or by an isomorphism between the two relators. For instance in the following presentation, a , cc , and d^3 are pieces, but ccc , ab , and aba are not.

$$\langle a, b, c \mid (ab)^5, accc, accc, ddc^{-1}d^{-3}c \rangle$$

This agrees with the definition above that used \tilde{X} .

9.1. Cubical Presentations

A cubical presentation $\langle X \mid Y_1, Y_2, \dots \rangle$ consists of a nonpositively curved cube complex X together with local isometries of cube complexes $\phi_i : Y_i \rightarrow X$.




FIGURE 9.1. Classical, Rips-Segev (popularized by Gromov), and Cubical presentations.

The group G determined by the cubical presentation is $G = \pi_1 X / \langle\langle \pi_1 Y_1, \pi_1 Y_2, \dots \rangle\rangle \cong \pi_1 X^*$ where $X^* = X \sqcup \text{Cone}(Y_i) / \{(y_i, 0) \sim \phi_i(y_i) : \forall y_i \in Y_i\}$. The notation $\langle\langle H_1, H_2, \dots \rangle\rangle$ denotes the smallest normal subgroup containing $\cup H_i$.

We will often use X^* to denote the cubical presentation $\langle X \mid Y_1, Y_2, \dots \rangle$.

Note that X^* is built from cubes and pyramids. We denote its universal cover by \tilde{X}^* , and are mainly interested in the cubical part of \tilde{X}^* which is the cover of X associated to $\langle\langle \pi_1 Y_1, \pi_1 Y_2, \dots \rangle\rangle$. Note that X plays the role of a bouquet of circles in a classical presentation, and the cubical part of \tilde{X}^* generalizes the Cayley graph.




FIGURE 9.2. X^* is built from X by attaching cones over the relators

9.1.1. Disk Diagrams in X^* . A *disk diagram* D in X^* is a combinatorial map of a disk diagram $D \rightarrow X^*$ such that $\partial_p D \rightarrow X$. We decompose D into 0-cubes, 1-cubes, and 2-cubes mapping to X and triangles mapping to the various $\text{Cone}(Y_i)$.

The triangles are partitioned into cyclic families each around a point mapping to the conevertex of some $\text{Cone}(Y_i)$ and we combine each such family together to form a *cone-cell*.

We define $\text{Comp}(D) = (\#\text{cone-cells}, \#\text{squares})$.




FIGURE 9.3. A disk diagram in X^* has squares mapping to X together with cone-cells formed from sequences of triangles around each conevertex.

9.2. The Fundamental Theorem of Small-Cancellation Theory

We now state and illustrate the main theorem of small-cancellation theory in the classical and cubical cases. Most theorems in small-cancellation theory are applications of or variants of these statements (together with analogous statements for annular diagrams).

For instance, by considering minimal complexity diagrams whose boundary path is a (nonclosed) path in Y_i , one can show that each $Y_i \rightarrow X$ lifts to an embedding in \tilde{X}^* .

9.2.1. Fundamental theorem of classical small-cancellation theory.

THEOREM 9.2. *Let X be a $C'(\frac{1}{6})$ classical presentation. A reduced disk diagram $D \rightarrow X$ is either:*

- (1) a single 0-cell or 2-cell;
- (2) a “ladder”;
- (3) or contains three or more shells and/or spurs.




FIGURE 9.4. A single 0-cell or 2-cell, a possibly singular ladder, and a diagram with three or more spurs and or shells.

Here a *shell* R in D is a 2-cell with $\partial_p R = QS$ where Q is a subpath of $\partial_p D$ and $|S| < |Q|$. A *spur* is a valence 1 vertex that is the endpoint of a 1-cell in ∂D .

9.2.2. Fundamental theorem of cubical small-cancellation theory. We will define the $C'(\frac{1}{n})$ condition for cubical presentations as well as the notion of reduced diagram in that context in Section 9.5. The rough idea is the same as in the classical case: $C'(\frac{1}{n})$ says that pieces are $< \frac{1}{n}$ of the boundary cycles. However, we now have “cone-cells” instead of 2-cells, and there are pieces with other cone-cells and also pieces with “rectangles”. Roughly speaking, a diagram is reduced if it is (locally) of minimal area.

THEOREM 9.3. *Let X^* be a $C'(\frac{1}{12})$ cubical presentation. A reduced disk diagram $D \rightarrow X^*$ is either:*

- (1) a single 0-cell or cone-cell;
- (2) a “ladder”;
- (3) or contains three or more shells, cornsquares, and/or spurs.

Note that we have strengthened the classical hypothesis that pieces have length $< \frac{1}{6}$ to a hypothesis that they have length $< \frac{1}{12}$. We have also added “cornsquares” to the list of positive curvature features.




FIGURE 9.5. A single 0-cell or single cone-cell, a possibly singular ladder, and a diagram with three or more spurs, cornsquares, and/or shells.

9.3. Combinatorial Gauss-Bonnet Theorem

An *angled 2-complex* X is a combinatorial 2-complex that has an *angle* $\sphericalangle(c) \in \mathbb{R}$ associated to each corner c of each 2-cell (or equivalently, associated to each edge in each link).

$$\kappa(v) = 2\pi - \pi\chi(\text{link}(v)) - \sum_{c \in \text{Corners}(v)} \sphericalangle(c)$$

The *curvature* at a 2-cell f of X is defined by:

$$\kappa(f) = \sum_{c \in \text{Corners}(f)} \sphericalangle(c) - (|\partial_p f| - 2)\pi$$

where $|\partial_p f|$ denotes the length of the boundary path of f .

Letting $\text{def}(\sphericalangle) = \pi - \sphericalangle$ we have the following equivalent formulas that are sometimes useful:

$$(9.1) \quad \kappa(v) = \pi(2 - \text{deg}(v)) - \sum_{c \in \text{Corners}(v)} \text{def}(\sphericalangle(c))$$

$$(9.2) \quad \kappa(f) = 2\pi - \sum_{c \in \text{Corners}(f)} \text{def}(\sphericalangle(c))$$

Our cases of greatest interest are when X is a disk diagram and the formula for $\kappa(v)$ often simplifies. In particular when v is internal $\kappa(v) = 2\pi - \sum \sphericalangle$ and when v is at a nonsingular boundary point $\kappa(v) = \pi - \sum \sphericalangle$. We indicate some common scenarios in Figure 9.6.




FIGURE 9.6. Some curvatures of 0-cells

THEOREM 9.4. *Let X be a compact angled 2-complex then:*

$$2\pi\chi(X) = \sum_{v \in 0\text{-cells}} \kappa(v) + \sum_{f \in 2\text{-cells}} \kappa(f)$$

PROOF. This is just a way of distributing $\chi = V - E + F$ as “curvature” among the 0-cells and 2-cells. We refer to [55] for a proof consistent with our terminology. □

9.4. Greendlinger's Lemma and the Ladder Theorem

We now sketch proofs of Greendlinger's Lemma and the Ladder Theorem for the classical $C(6)$ small-cancellation setting. Combined, these results provides Theorem 9.2.

A *piece* in a disk diagram D is a path $P \rightarrow D$ that arises in two ways as a subpath of a boundary path of a 2-cell. We normally ignore *trivial* pieces having length 0.




FIGURE 9.7. Pieces.

D satisfies the $C(p)$ *small-cancellation condition* if no 2-cell R in D has $\partial_p R$ the concatenation of $< p$ pieces.




FIGURE 9.8. Five and seven pieces around 2-cells.

An *i-shell* in D is a 2-cell R such that $\partial_p R = QS$ where the *outerpath* Q is a subpath of $\partial_p D$ and the *innerpath* S is the concatenation of i nontrivial pieces. A *spur* is a 1-cell ending at a 0-cell of valence 1.




FIGURE 9.9. A spur, 0-shell, 1-shell, 2-shell, and 3-shell.

We emphasize that the innerpath of an i -shell is the concatenation of i *nontrivial* pieces. In some sense, a 0-shell is a degenerate 1-shell, as its innerpath consists of one trivial piece in D . This viewpoint is consistent with the values assigned below.

THEOREM 9.5 (“Greendlinger’s Lemma”). *Let D be a $C(6)$ disk diagram. Then either D is a single 0-cell or D is a single closed 2-cell or D has at least 2π worth of i -shells and spurs whose values are:*

- π for spurs
- π for 0-shells and 1-shells
- $\frac{2\pi}{3}$ for 2-shells
- $\frac{\pi}{3}$ for 3-shells

SKETCH. We assign angles to corners of 2-cells using the following rules:

- $\frac{2\pi}{3}$ at internal corners at vertices of valence ≥ 3 .
- π at valence 2 vertices (i.e. where two 2-cells meet, or a single 2-cell meets the boundary)
- $\frac{\pi}{2}$ at singly external corners
- 0 at doubly external corners of valence > 2 .




FIGURE 9.10. Assigning angles to corners of 2-cells.

A 2-cell R is *internal* if it has no bounding 1-cell in ∂D , and it is *external* if ∂R contains a 1-cell in ∂D . We say R is *multiply-external* if $\partial_p R$ has more than one maximal subpath that is a subpath of $\partial_p D$. The terms *singly-external*, *doubly-external*, *triplly-external* etc. have the obvious meanings.




FIGURE 9.11. Internal, external, and multiply external 2-cells.

$\kappa(f) \leq 0$ when f is internal or multiply-external, and $\kappa(f) < 0$ when f is triply-external (or more). Indeed, when f is internal we have $\kappa(f) = 2\pi - |\partial_p f|(\pi - \frac{2\pi}{3}) = 2\pi - |\partial_p f|\frac{\pi}{3}$, and when f is m -external $\kappa(f) \leq 2\pi - 2m\frac{\pi}{2}$.

When f is an i -shell, it is singly-external and treating $i = 0$ like $i = 1$ we have $\kappa(f) = 2\pi - \frac{\pi}{2} - \frac{\pi}{2} - (i - 1)\frac{2\pi}{3} = \pi - (i - 1)\frac{2\pi}{3}$ which agrees with the listed values.

An internal 0-cell v has $\kappa(v) \leq 0$ (there are two cases to consider: valence(v) > 2 and valence(v) = 2). A boundary 0-cell v has $\kappa(v) = \pi$ if v is at the tip of a spur. Otherwise it has $\kappa(v) \leq 0$ (again there are several cases to consider).

In the exceptional cases where D is a single 0-cell v or single closed 2-cell f we have $\kappa(v) = 2\pi$ and $\kappa(f) = 2\pi$ respectively. Let us now assume that D is not exceptional.

Applying Theorem 9.4 we have:

$$2\pi = 2\pi\chi(D) = \sum_v \kappa(v) + \sum_f \kappa(f) \\ \leq \pi(\#\text{(spurs)}) + \pi(\#\text{(0-shells)}) + \pi(\#\text{(1-shells)}) + \frac{2\pi}{3}(\#\text{(2-shells)}) + \frac{\pi}{3}(\#\text{(3-shells)}) \quad \square$$

THEOREM 9.6 (The Ladder Theorem). *If the $C(6)$ diagram D has exactly two features of positive curvature (i.e. spurs, and i -shells with $0 \leq i \leq 3$) then D is a ladder.*

SKETCH. Both features must have curvature equal to exactly π , and all other 0-cells and 2-cells must have curvature exactly 0. Removing a spur or 0-shell or 1-shell from one side, the remainder is a single 0-cells or 2-cell, or is a ladder by induction. Indeed, we uncovered exactly one positive curvature feature, and the removed feature was attached on the interior of the outerpath of the newly exposed positively curved feature – otherwise there would have been more than two positively curved features to begin with. Thus the original D was a ladder. \square




FIGURE 9.12. Removing one of the positively curved features exposes at most one positively curved feature, and thus leaves behind a smaller ladder (or a single 0-cell or 2-cell). The removed positively curved feature is attached on the interior of the outerpath of the exposed positively curved feature in the smaller ladder – otherwise there would have been an additional feature of positive curvature as in the right two diagrams.

9.5. Reduced Diagrams

For parallelism we will use the notation X^* for a 2-complex (though this is not in exact agreement with our usage for a cubical presentation, where the cones are built from pyramids.) In classical small-cancellation theory, a diagram $D \rightarrow X^*$ is *reduced* if it has no *cancelable pair*, which is a pair of 2-cells R_1, R_2 that meet along an edge e (and possibly other edges) such that starting at e , their boundary paths eP_1, eP_2 project to the same closed path in X , so there is a commutative diagram:

$$\begin{array}{ccc} e & \rightarrow & \partial_p R_2 \\ \downarrow & \nearrow & \downarrow \\ \partial_p R_1 & \rightarrow & X \end{array}$$

The cancelable pair is *removed* by deleting the open cells R_1, R_2 and the maximal open piece that e extends to, and identifying P_1, P_2 . (If we only remove the open 1-cell e , there could be some tails hanging off of the newly created diagram.)

The $C(p)$ condition on X^* is defined implicitly by requiring that no reduced diagram $D \rightarrow X^*$ can have a 2-cell whose boundary path is the concatenation of fewer than p pieces. Equivalently, for each reduced diagram $D \rightarrow X^*$, the diagram D is $C(p)$. The $C'(\frac{1}{n})$ condition on X^* says that each piece P between 2-cells R_1, R_2 in a reduced diagram has $|P| < \frac{1}{n}|\partial_p R_i|$.




FIGURE 9.13. Removing a cancelable pair in classical case.

We now describe some generalizations of this cancelable pair removal procedure, as well as some further available reductions that arise for cubical presentations.




FIGURE 9.14. Reductions in diagrams $D \rightarrow X^*$.

Combining cone-cells: involves two cone-cells C_1, C_2 that map to the same $\text{Cone}(Y)$ and that meet along a maximal open arc E such that $\partial_p C_i = EP_i$ and such that the concatenation $P_1P_2^{-1}$ projects to a closed path in $\text{Cone}(Y)$. We replace $C_1 \cup E \cup C_2$ by a new cone-cell C mapping to $\text{Cone}(Y)$ and with $\partial_p C = P_1P_2^{-1}$.

Replacing cone-cells: involves a cone-cell C mapping to some $\text{Cone}(Y)$ such that $\partial_p C$ is already null-homotopic in Y , and so we can replace C by a square disk diagram $D_c \rightarrow Y$ with $\partial_p D_c = \partial_p C$.

Absorbing squares: There are really two essentially different possibilities here: In the first case, a cone-cell C absorbs a “contiguous square” s that lies along an edge e on ∂C , and such that the map $\partial C \rightarrow Y$ extends to a map $\partial C \cup s \rightarrow Y$. We replace C by a new cone-cell C' such that $\partial_p C'$ has e replaced by $\partial s - e$.

In the second case, s is a cornsquare whose outerpath lies on C . As in Lemma 3.6, we can homotope the square part of D without changing the number of squares so that there is now a square s' with a corner along $\partial_p C$ in the sense that $\partial s' = abcd$ and ab is a subpath of $\partial_p C$. We replace C by the cone-cell C' where $\partial_p C'$ is obtained from $\partial_p C$ by replacing ab by cd .

Removing bigons: involves a bigonal square diagram in D which necessarily can be replaced by a smaller square diagram as in Theorem 3.2.

We emphasize that: *All these reductions reduce $\text{Comp}(D)$.* Consequently, reduced diagrams with a given boundary path always exist, since a minimal complexity diagram with a given boundary path must be reduced.

DEFINITION 9.7 (Reduced Diagram). For a cubical presentation X^* , a diagram $D \rightarrow X^*$ is *reduced* if:

- (1) there is no bigon in a square subdiagrams
- (2) there is no cornsquares whose outerpath lies on a cone-cell
- (3) no square sharing an edge with the boundary of a cone-cell can be absorbed into it.
- (4) no adjacent cone-cells sharing an edge can be combined to form a single cone-cell.
- (5) $\partial_p C \rightarrow X$ is essential for each cone-cell C .

DEFINITION 9.8 (Pieces and the $C'(\frac{1}{n})$ condition). The *pieces* in a reduced diagram $D \rightarrow X^*$ are subpaths $P \rightarrow \partial_p C$ of the boundary path of a cone-cell C such that the dual curves initiating at edges of P “fellow-travel” within the square part of D until they adjacently terminate on the boundary of another cone-cell or on a rectangle. We will focus on the *rectangles* carrying these outgoing dual curves.

X^* is $C(\frac{1}{n})$ if $|P| < \frac{1}{n}|Y_i|$ whenever P is a piece in a cone-cell C of D where C maps to $\text{Cone}(Y_i)$. The *size* $|P|$ denotes the diameter of the lift \tilde{P} to $\tilde{Y}_i \subset \tilde{X}$, and $|Y_i|$ denotes the length of the shortest essential combinatorial path in Y_i .




FIGURE 9.15. Cone-pieces and wall-pieces.

9.6. Producing Examples

A quick source of examples of classical $C'(\frac{1}{6})$ groups arise by raising relators to sufficiently high powers as follows:

THEOREM 9.9. *Given $\langle a, b, \dots \mid W_1, \dots, W_r \rangle$ with $W_i^p \not\sim W_j^q$ (for $p, q \neq 0$), there exists M such that for $n_i \geq M$ the following presentation is $C'(\frac{1}{n})$:*

$$\langle a, b, \dots \mid W_1^{n_1}, \dots, W_r^{n_r} \rangle$$

SKETCH. Let D be an upper bound on the size of pieces between the various W_i^∞ and W_j^∞ – we note that there are finitely many possible pieces. Now choose $M > nD$. □




FIGURE 9.16. Some of the finitely many types of pieces illustrated in the universal cover of the bouquet of circles for $\langle a, b \mid (abb)^\infty, (aba^{-1}b^{-1}b^{-1})^\infty \rangle$.

THEOREM 9.10. *Let X be a compact nonpositively curved cube complex with $\pi_1 X$ hyperbolic. Let $\{H_1, \dots, H_r\}$ be a malnormal collection of quasi-convex subgroups. There exist compact based local isometries $Y_i \rightarrow X$ with $H_i = \pi_1 Y_i$ and finite subsets $S_i \subset H_i - \{1\}$ such that for any regular covers $\widehat{Y}_i \rightarrow Y_i$ with $S_i \cap \pi_1 \widehat{Y}_i = \emptyset$ the following cubical presentation is $C'(\frac{1}{n})$:*

$$\langle X \mid \widehat{Y}_1, \dots, \widehat{Y}_r \rangle$$

SKETCH. By Lemma 3.28, let \widetilde{Y}_i be an H_i -cocompact superconvex core. For any hyperplane U disjoint from \widetilde{Y}_i , the intersection $\widetilde{Y}_i \cap N(U)$ has diameter uniformly bounded by some constant R . This uses cocompactness of \widetilde{Y}_i together with superconvexity. Thus R is an upper bound on the size of any contiguous wall-piece. A key point is that non-contiguous cone-pieces and non-contiguous wall-pieces are dominated by contiguous wall-pieces, and so these also have size $\leq R$ as well. Intersections between nonequal translates of \widetilde{Y}_i and \widetilde{Y}_j are bounded by L . This uses malnormality of $\{H_1, \dots, H_r\}$ together with the H_i -cocompactness of each \widetilde{Y}_i . Thus contiguous cone-pieces have size $\leq L$.

Choose S_i to be representatives of conjugacy classes of closed paths $A \rightarrow Y_i$ where $Y_i = H_i \backslash \widetilde{Y}_i$ such that $|A| \leq n \max(R, L)$. □




FIGURE 9.17. Contiguous cone-pieces and wall-pieces are bounded in \tilde{X} by respectively using cocompactness and mal-normality, and using cocompactness and superconvexity.




FIGURE 9.18. Noncontiguous pieces are also contiguous wall-pieces, since squares within the shards can be pushed out of the way.

9.7. Rectified Diagrams

To prove the fundamental theorem of cubical small-cancellation theory, we will assign angles to corners of 2-cells in D . However, as a consequence of the hexagon-move flexibility of the square part of D , there are multiple ways of declaring the pieces within D , and this creates some technical problems. We remedy this by fixing one specific piece structure – which concomitantly decomposes D into “cone-cells”, “rectangles”, and “shards”. We sketch this “rectification” \bar{D} of D in this subsection.

9.7.1. Admitting rectangles. We begin by declaring a linear ordering of the cone-cells C_1, C_2, \dots of D with the cone-cell at infinity C_∞ appearing last. We then cyclically order the 1-cells in each $\partial_p C_i$, by choosing a first 1-cell and then proceeding counterclockwise. Together these two orderings provide a linear ordering of all 1-cells appearing in all the cone-cell boundary cycles.

We now *admit* the rectangle emanating from the ij -th 1-cell e_{ij} terminating on either a cone-cell or on the side of a previously admitted rectangle.




FIGURE 9.19. Ordering the 1-cells in boundary paths of cone-cells.

If a rectangle had already terminated at e_{ij} then we ignore the rectangle emanating from e_{ij} (as it was already admitted) and we proceed to the next 1-cell in the ordering.




FIGURE 9.20. Admitting rectangles.

Admitted rectangles are of the form $[-1, 1] \times [0, n]$, where the case $n = 0$ is *degenerate*, and admitted rectangles have *internal part* of form $(-1, 1) \times [0, n]$. (We regard $[-1, 1]$ as a single 1-cube but $[0, n]$ as a string of n 1-cubes.)




FIGURE 9.21. The diagram D is on the left. We have indicated the ordering of cone-cells and cyclic ordering of attaching maps. The admitted rectangles within D are indicated in the middle, and the rectified diagram \bar{D} is on the right. (We have not highlighted 0-cell shards.)

9.7.2. Shards. Let E denote the union of the open cone-cells and the internal parts of admitted rectangles. The components of $D - E$ are the *shards*. Under sufficient small-cancellation conditions, each shard is a simply-connected square disk diagram.




FIGURE 9.22. Five shards. Note that that two of these shards are just 0-cells.

9.7.3. Pieces. The *pieces* in D are paths on cone-cells (not C_∞) consisting of a sequences of edges whose admitted rectangles end in parallel on the same cone-cell (not C_∞) or rectangle – hence the terms *cone-piece* and *wall-piece*. To say the rectangles end *in parallel* means that only a shard consisting of a square diagram lies between them.




FIGURE 9.23. Four pieces.

9.7.4. Assigning Angles. When no cone-cell C has $\partial_p C$ the concatenation of < 24 pieces¹, then there is a nice way to assign angles at corners of cone-cells and shards.

All admitted (non-degenerate) rectangles are assigned the “usual” angles: π along the sides and $\frac{\pi}{2}$ at the corners. The corners of cone-cells are assigned angles of $\frac{\pi}{2}$ except for the situations indicated in Figure 9.24:

- π is assigned to a corner when the emerging rectangles end in parallel on a cone-cell or rectangle, or if they “implicitly” end in parallel on a rectangle as illustrated.
- $\frac{3}{4}\pi$ is assigned to a corner when the emerging rectangles bound a (possibly degenerate) triangular shard with two or three cone-cells at its corners.
- 0 is assigned in the unusual case of two emerging rectangles ending in parallel (a shard between them) at a singular vertex on ∂D .

The corners of shards in D are assigned the “obvious” angles so that the vertices v at their corners have $\kappa(v) = 0$. A shard f is (almost always) automatically nonpositively curved – i.e. $\sum \angle - (|\partial_p f| - 2)\pi \leq 0$. This assignment does occasionally require negative angles and there are many cases to consider, but in each case, the obvious choices work, and the “shards take care of themselves”. Two cases are indicated in Figure 9.25.

¹actually < 12 suffices but requires a more complex angle assignment




FIGURE 9.24. Four pieces.




FIGURE 9.25. Choose angles so that $\kappa(v) = 0$ for each vertex v , and the shard will be nonpositively curved.

9.7.5. The positively curved shards. The exceptional situations of positively curved shards arise from cornsquares on cone-cells and from *monkey-tails* within the diagram consisting of a rectangle that terminates on itself bounding a shard. However neither of these arise for a reduced diagram. Such cornsquares are explicitly prohibited in the definition of reduced, and monkey-tails cannot arise because there would be a square bigon arising from a rectangle cutting through the monkey-tail.

Cornsquares terminating on ∂C_∞ are (together with spurs) an important feature of positive curvature that can arise in the diagram. In the case where there is an actual corner of a square on ∂C_∞ the 0-cell at the corner is the entire shard and has $\kappa = \frac{\pi}{2}$. In the general case, one could insist that the shard be flat by assigning angles of $\pm \frac{\pi}{2}$ in which case the positive curvature is again concentrated at v , or one could make $\kappa(v) = 0$ and have the shard have $\kappa = \frac{\pi}{2}$ since its angles are both $\frac{\pi}{2}$. I am not sure which is a more natural viewpoint, but they lead to the same conclusion. Technically, the 0-cell at the endpoint of a spur is a positively curved shard, and so is the isolated 0-cell arising in the case of a trivial diagram.

9.7.6. Curvatures of cone-cells with $< \frac{1}{24}$ small-cancellation. We now examine the curvature of a cone-cell in \bar{D} . As in the classical case, if C is doubly-external, we have $\kappa(C) \leq 0$, and if C is triply-external or more then $\kappa(C) < 0$.

For internal cone-cells we have:

LEMMA 9.11. *Let C be an internal cone-cell in a reduced diagram D . If $\partial_p C$ is not the concatenation of fewer than 24 pieces, the angle assignment*




FIGURE 9.26. The exceptional situations of positively curved shards arise from cornsquares on cone-cells and monkey-tails within the diagram. Neither of these arise for a reduced diagram. Cornsquares terminating on ∂C_∞ correspond to a $+\frac{\pi}{2}$ curvature.

on the rectified diagram \bar{D} has $\kappa(C) \leq 0$. (Moreover 25 or more pieces implies $\kappa(C) < 0$.)

PROOF. Consider Equation (9.2):

$$\kappa(C) = 2\pi - \sum_{\triangleleft} \text{def}(\triangleleft)$$

We show that there is a defect of at least $\frac{\pi}{4}$ for every three pieces, and thus with at least eight such defects we obtain $\sum \text{def}(\triangleleft) \geq 8\frac{\pi}{4} = 2\pi$ as needed. Note that it is possible to have three consecutive pieces with no nonzero defects between them as on the left in Figure 9.27. However the ordering on the cone-cells which underlies our rectification makes it impossible to have a sequence of four pieces with zero angle defect for the corresponding transitional corners. See the right of Figure 9.27. \square




FIGURE 9.27. While it is possible to have three consecutive pieces with zero transitional defects as on the left, it is impossible to have more than this because, as on the right, the rules for building rectified diagrams would be violated when one considers the ordering of cone-cells. Both $n > 4$ and $n < 4$ lead to contradictions.

Walls in Cubical Small-Cancellation Theory

In this chapter we describe and investigate a situation where \tilde{X}^* has a natural wallspace structure. We first explain these wallspaces in the classical $C'(\frac{1}{6})$ setting in Section 10.1. The enabling feature on $X^* = \langle X \mid Y_1, Y_2, \dots \rangle$ that will allow us to build a wallspace structure for \tilde{X}^* , is to hypothesize that each Y_i is itself a wallspace, and we describe this in Section 10.2. We explain how to build Y_i with a wallspace structure by passing to a finite cover in Section 10.3. We build the walls in \tilde{X}^* in Section 10.4. Finally, we show that these walls are quasiconvex in Section 10.5.

10.1. Walls in Classical $C'(\frac{1}{6})$ Small-Cancellation Complexes

Let X^* denote the standard 2-complex of a classical $C'(\frac{1}{6})$ small-cancellation complex. There is a natural system of walls in \tilde{X}^* as follows: Firstly, by possibly subdividing all 1-cells, we assume that all attaching maps of 2-cells have even length. The walls of \tilde{X}^* are graphs (see Figure 10.1) that intersect 1-cells and 2-cells in midcells (see Figure 6.8) as described in Example 6.6.




FIGURE 10.1. The carrier of a wall in a $C(\frac{1}{6})$ -complex \tilde{X}^* .

Walls are 2-sided, embedded multi-trees with convex carriers. Being a multi-tree and not self-crossing are consequences of Greendlinger's lemma, as a minimal area diagram D for a path corresponding to a self-crossing could not have enough shells (see Figure 10.2). The ladder theorem implies the convexity of the carriers, as a minimal area diagram between a geodesic γ and a path σ on the carrier $N(W)$ of the wall W would be forced to be a ladder L . Indeed, no shell lies on γ or σ , and then a contradiction arises by considering how the initial or terminal 1-shell of L relates to the ladder $L' \subset N(W)$ containing σ .

10.2. Wallspace Cones

The key ingredients that allow us to define walls in the classical case are:




FIGURE 10.2. Walls in the $C'(\frac{1}{6})$ complex \tilde{X} cannot self-cross, and cannot contain cycles since a minimal area disk diagram for such a situation would have at most one shell.

- (1) An even length circle is a wallspace, whose walls correspond to pairs of centers of antipodal edges
- (2) Under the $C'(\frac{1}{6})$ hypothesis, these walls are “convex” – one must pass through many pieces to travel essentially from a wall back to itself.

We generalize this to require that each relator $Y \rightarrow X$ is a *convex wallspace* as follows:

- (1) Each hyperplane in Y is 2-sided and embedded.
- (2) The collection $\{H_i\}$ of hyperplanes are partitioned into subcollections called *walls*.
- (3) Two hyperplanes in the same wall are disjoint from each other.
- (4) If $P \rightarrow Y$ is a path that starts and ends on the carrier $N(W)$ of a wall W , and P is the concatenation of fewer than 15 pieces then there is a disk diagram $D \rightarrow X$ between P and a path $P' \rightarrow N(W)$.




FIGURE 10.3. The hyperplanes of the cube complex Y are partitioned into walls.

10.3. Producing Wallspace Cones

In the classical case, we could immediately turn all relators into wallspaces by subdividing the 1-skeleton. It is harder to do this naturally in a higher dimensional setting. Another way to turn relators into wallspaces is to take




FIGURE 10.4. If $P \rightarrow Y$ is the concatenation of few pieces, then P is square-homotopic to a path $P' \rightarrow N(W)$.

their double covers. This has the disadvantage that it substantially changes the group of the presentation, but it is sufficient for our purposes (as we are interested in producing quotients with walls).

Let \mathbb{W} denote a partition of the hyperplanes of Y – we have in mind the case where hyperplanes in the same equivalence class do not cross each other. Such a partition arises by taking the preimages of hyperplanes under a map $Y \rightarrow X$ where hyperplanes of X embed. Consider the covering space $\check{Y} \rightarrow Y$ induced by the homomorphism $\#_{\mathbb{W}} : \pi_1 Y \rightarrow \mathbb{Z}_2^{\mathbb{W}}$ that counts the number of times a path travels through an edge dual to respective classes.

The space \check{Y} has a natural wallspace structure: Each wall is the preimage of all hyperplanes in a \mathbb{W} class. For instance, when Y is a circle and \mathbb{W} is the discrete partition, \check{Y} is just the usual \mathbb{Z}_2 cover. We refer the reader to Figure 10.5.




FIGURE 10.5. Given a cone $Y \rightarrow X$, we can obtain a wallspace cone \check{Y} that is the finite cover associated to a homomorphism $\pi_1 Y \rightarrow \mathbb{Z}_2^{\mathbb{W}}$.

10.4. Walls in \tilde{X}^*

Suppose each Y_i in X^* has a wallspace structure. A *wall* W in \tilde{X}^* is a collection of disjoint hyperplanes such that the intersection with each relator

$Y_i \subset X^*$ is either \emptyset or is a wall V_i of Y_i . If we think of the universal cover \tilde{X}^* together with its cones, and want the walls to be connected and separating, then we can imagine extending each wall W slightly within each cone so that $W \cap \text{Cone}(Y_i) = \text{Cone}(V_i)$ for each Y_i .




FIGURE 10.6. A wall in \tilde{X}^* is a collection of hyperplanes that intersects each Y_i in a wall of Y_i (which generalizes a midcell or midcube).

THEOREM 10.1. *Suppose $X^* = \langle X \mid Y_1, Y_2, \dots \rangle$ is $C'(\frac{1}{24})$ and each Y_i has the structure of a wallspace which is convex with respect to the pieces of X^* . Then each hyperplane of \tilde{X}^* lies in a unique wall of \tilde{X}^* .*

PROOF. This follows from the fundamental theorem of cubical small-cancellation theory. \square

10.5. Quasiconvexity of Walls in \tilde{X}^*

For a wall W of \tilde{X}^* its *carrier* $N(W)$ is the union of all carriers of the hyperplanes of W and all cones intersecting W . Its *thickened carrier* $T(W)$ is the union of $N(W)$ together with all minimal square ladders that start and end on cones of $N(W)$ that are *consecutive* in the sense that some hyperplane of W crosses both.




FIGURE 10.7. The thickened carrier $T(W)$ is the union of $N(W)$ together with minimal length square ladders that start and end on consecutive cones.

The thickened carrier $T(W)$ interpolates between $N(W)$ and \tilde{X}^* . It is easy to see that $N(W) \subset T(W)$ is a quasi-isometry when cones are finite, and we show below that $T(W)$ isometrically embeds in \tilde{X}^* . We are thus able to obtain the following:

COROLLARY 10.2. *When X and the cones Y_i are compact the inclusion $\text{Stabilizer}(W) \subset \pi_1 X^*$ is a quasi-isometric embedding.*

THEOREM 10.3 (Thickened Carrier Isometrically Embeds). *Let $T(W)$ be the thickened carrier of a wall W in \tilde{X}^* . Then $T(W) \subset \tilde{X}^*$ is an isometric embedding.*

PROOF. We are actually proving that there is an isometric embedding between their 1-skeleta. Let γ be a geodesic that start and ends on $T(W)$ and let D be a minimal complexity disk diagram that has γ on one side and has a path $\sigma \rightarrow T(W)$ on the other side.

For each edge e of σ there is a *cladder* (i.e. “cone-ladder”) within D that starts on e and travels through square ladders (like a dual curve) and enters and leaves cone-cells along edges dual to the same wall of the cone. Note that the cladder can bifurcate at cone-cells. We refer to Figure 10.8.

We claim that all terminal edges of the cladder end on γ . And so since this holds for each e we have $|\sigma| \leq |\gamma|$. Our claim follows from the ladder characterization in Theorem 9.3 and the minimality of D , since if a cladder travels back to cross σ again (see Figure 10.9), we would either find that there was a way to thicken $T(W)$ within D , or we would obtain a reduced diagram with only two positively curved features that is not a ladder. \square




FIGURE 10.8. Let D be a minimal complexity diagram between $T(W)$ and a geodesic γ . Each cladder in D starting on an edge of σ ends entirely on γ . Thus $|\gamma| \geq |\sigma|$ and $T(W)$ isometrically embeds.




FIGURE 10.9. The cladder cannot end at another edge of σ or it would lie in $T(W)$ as in the situation, or violate the ladder theorem as in the first and second situations.

Annular Diagrams

In this Chapter we begin by examining “annular diagrams”. These were originally introduced to study the conjugacy problem in parallel with disk diagrams which were introduced to study the word-problem. The fundamental theorem of small-cancellation theory asserts that a disk diagram with few features of positive curvature has a specific structure. In parallel with this, we explain in Section 11.1 that an annular diagram with no features of positive curvature has a restricted nature. In Section 11.2, we examine annular diagrams reflecting a conjugacy between elements stabilizing walls in \tilde{X}^* . In this case the annular diagrams exist with a particularly specific structure. We apply this in Section 11.3 to conclude that under certain conditions, the stabilizer of each wall of \tilde{X}^* is almost malnormal in $\pi_1 X^*$.

11.1. Classification of Flat Annuli

An *annular diagram* is a compact connected combinatorial 2-complex A with a fixed embedding in S^2 such that $\pi_1 A \cong \mathbb{Z}$. Note that A has two boundary paths corresponding to the attaching maps of the additional 2-cells that would be added to obtain S^2 . As disk diagrams are used to study the triviality of elements of $\pi_1 X^*$, annular diagrams are used to study conjugacy between nontrivial elements.

The $C(p)$ and $C'(\frac{1}{n})$ conditions are defined as for disk diagrams D .

An *annular diagram in X^** is defined as for a disk diagram, and there is an existence theorem for annular diagrams representing conjugacy between elements.

An annular diagram $A \rightarrow X^*$ is *reduced* if there are no reductions as for disk diagrams, however one must be a bit more careful about combining cone-cells - since as in Figure 11.1, it is possible for a cone-cell to reach around the annulus and touch itself along an edge in a manner that \tilde{A} would have a reduction. Such configurations lead to “elliptic annuli” which are studied more carefully in [76].




FIGURE 11.1. The cone-cell cannot be combined with itself.

We now discuss the restrictive form of an annular diagram that does not have features of positive curvature on its boundary. We begin with the classical case and then move towards increasingly complicated statements in the cubical case.

THEOREM 11.1 (Classical Thin Annuli). *Let A be a $C(7)$ annular diagram. Then either A has a spur or an i -shell with $i \leq 3$ or A is of one of the three forms indicated in Figure 11.2:*




FIGURE 11.2. A $C(7)$ annulus with no feature of positive curvature must have width 0,1, or 2 as above.

SKETCH. Using the angle assignment in the proof of Theorem 9.5, the $C(7)$ condition implies an internal 2-cell would be negatively curved. Hence there would also be a positively curved feature by Theorem 9.4. If A contains no 2-cell then A is a subdivided circle. If one A contains a separating 2-cell C , then cutting along C one obtains a ladder, and hence regluing one obtains a width 1 annulus. In the last case, the external 2-cells on each side form a sequence of 3-shells which must then be staggered to yield a width 2 annulus. \square

REMARK 11.2. A variant of Theorem 11.1 also holds in the $C(6)$ case: Either there are spurs or i -shells or else A is an arbitrary width flat annulus. We have illustrated a width 4 flat annulus in Figure 11.3.




FIGURE 11.3. A flat $C(6)$ annulus can have internal 2-cells and thus arbitrarily many rings.

We next describe a generalization of Theorem 11.1 to an annulus arising from a cubical small-cancellation presentation. The proof is analogous to the proof in the classical situation.

THEOREM 11.3 (Cubical Thin Annuli). *Let A be an annular diagram satisfying the $C'(\frac{1}{24})$ cubical small-cancellation condition. Then either A has a spur, cornsquare or shell, or else: Either A is a square annulus, or A is an annuladder, or A is a square annulus bicolored by a pair of annuladders.*




FIGURE 11.4. A flat annulus in a cubical $C'(\frac{1}{24})$ complex either has a positively curved feature, or else is either a square annulus, or is “thin” and consists of a single annuladder, or is “thick” and is bicollared by annuladders with a square annular diagram in between.

We use the term *annuladder* for an annular diagram that looks “locally” like a ladder. It is formed from a sequence of cone-cells and rectangular grids – as one would get by taking ladder and identifying its initial and terminal cone-cells. To say that A is *bicollared* by annuladders just means that there are annuladder subdiagrams L_1, L_2 in A having disjoint interiors and running along the two sides of A in the sense that the two boundary paths α_1, α_2 of A are boundary paths of L_1, L_2 .

11.2. The Doubly Collared Annulus Theorem

We now describe a result which is especially catered to understand conjugate elements that stabilize two walls in \tilde{X}^* . It will allow us to verify almost malnormality of wall stabilizers under certain conditions, which plays a role in the proof Theorem 12.3.

THEOREM 11.4 (Doubly Collared Annuli). *Let $\langle X \mid Y_1, \dots, Y_r \rangle$ satisfy $C'(\frac{1}{24})$ and have cones Y_i that are convex wallspaces.*

Let $A \rightarrow X^$ be an annular diagram with boundary cycles α_1, α_2 that are essential in \tilde{X}^* . Suppose $\tilde{A} \rightarrow \tilde{X}^*$ has $\tilde{\alpha}_1, \tilde{\alpha}_2$ lifting to wall carriers N_1, N_2 where each $N_i = N(W_i)$ for some wall W_i in \tilde{X}^* .*

There exists a new annular diagram $B \rightarrow X^$ that is reduced.*

B contains two annuladders L_1, L_2 that are deformation retracts of B , with each L_i having (outside) boundary path β_i and each $\tilde{L}_i \subset N_i$.

B is in the same class as A in the sense that $\partial \tilde{B}$ lies on N_1, N_2 and β_i is conjugate to α_i in $\text{Stabilizer}(N_i)$ for each i .

Finally, B is either thick and is a bicollared annular diagram with L_1, L_2 on either side and $\partial B = \beta_1 \sqcup \beta_2$, or B is thin in which case B is itself an annuladder (but the β_i might not be the boundary paths of B in this case).




FIGURE 11.5. B is either “thick” and bicollared by two annuladders L_1, L_2 , or is “thin” and is a single annuladder with L_1, L_2 wandering within B .

11.3. Almost Malnormality

We now examine a technical condition ensuring that X^* has the property that $\text{Stabilizer}(W)$ is almost malnormal for each wall W in \tilde{X}^* .

LEMMA 11.5. *Let $X^* = \langle X \mid Y_1, \dots, Y_r \rangle$ be a cubical presentation satisfying the following conditions. Then $\text{Stabilizer}(W)$ is almost malnormal for each wall W of \tilde{X}^* .*

- (1) $C'(\frac{1}{24})$.
- (2) strongly convex wallspace cones (in the sense that if $P \rightarrow Y$ is a path that starts and ends on some wall carrier $N(V)$, and P is the concatenation of fewer than 16 pieces, then P is path homotopic in Y to a path $P' \rightarrow N(V)$.)
- (3) Suppose the hyperplanes of X are 2-sided and are partitioned into families of noncrossing hyperplanes so that each family $\{H_1, \dots, H_k\}$ gives a malnormal collection $\{\pi_1 H_1, \dots, \pi_1 H_k\}$ in $\pi_1 X$.
- (4) Suppose the walls in each Y_j are precisely intersections of Y_j with these hyperplane families.
- (5) Suppose each hyperplane family carrier $\sqcup_i N(H_i)$ has empty or connected intersection with each piece in each Y_j .
- (6) Suppose each Y_i is compact.

SKETCH. Let $A \rightarrow X^*$ be an annular diagram whose boundary cycles α_1, α_2 lift to translates of the same wall carrier $N(W)$ under the lift $\tilde{A} \rightarrow \tilde{X}^*$. Applying Theorem 11.4, we can pass to $B \rightarrow \tilde{X}^*$ which contains the same information. Let us assume B is chosen to have minimal complexity within its class.

If B is a square annular diagram, then since the hyperplane families of X are already malnormal, we see (by minimality) that B is thin – it is an annulus that is at most one square thick – and there is a map $\tilde{B} \rightarrow N(W)$.

B cannot be thick and doubly collared since a cone-cell along the boundary in L_1 or L_2 would provide a reduction because of strong convexity.

Finally, if B is thin, then since \tilde{L}_1, \tilde{L}_2 map to $N(W_1) = g_1N(W)$ and $N(W_2) = g_2N(W)$, we see that two hyperplanes in the same family pass through some piece between consecutive cones in B , and so they must be the same hyperplane by hypothesis, so $N(W_1) = N(W_2)$. \square

Virtually Special Quotients

In this chapter we show that being virtually compact special hyperbolic is preserved by appropriate quotienting. In Section 12.1 we sketch a proof of the “Malnormal Special Quotient Theorem”. The main idea is to arrange things so that the quotient virtually has a malnormal quasiconvex hierarchy and so Theorem 12.2 applies. In Section 12.2, we review this proof for an easy motivating case: For elements W_1, W_2 of the free group F_2 , the quotient $F_2/\langle\langle W_1^{n_1}, W_2^{n_2} \rangle\rangle$ is virtually special for sufficient n_i . In Section 12.3, we describe the “Special Quotient Theorem” which shows that virtual specialness is preserved when we quotient by sufficient finite index subgroups of arbitrary quasiconvex subgroups.

12.1. The Malnormal Special Quotient Theorem

DEFINITION 12.1 (Almost malnormal quasiconvex hierarchy). The class of groups \mathcal{MQH} with a *malnormal quasiconvex hierarchy* is the smallest class such that:

- (1) $1 \in \mathcal{MQH}$
- (2) $A *_C B \in \mathcal{MQH}$ whenever $A, B \in \mathcal{MQH}$ and C is malnormal and quasiconvex in $A *_C B$
- (3) $A *_C t=D \in \mathcal{MQH}$ whenever $A \in \mathcal{MQH}$ and C is malnormal and quasiconvex in $A *_C t=D$

Note that \mathcal{MQH} is the class of groups that can be built from trivial groups using finitely many HNN extension and free products with amalgamation along subgroups that are malnormal and quasiconvex.

More generally, we obtain two other increasingly general classes of groups by varying the base (1) and allowing *almost* malnormality in (2) and (3):

G has an *almost malnormal quasiconvex hierarchy* if G can be built from finite groups using finitely many HNN extension and free products with amalgamation along subgroups that are almost malnormal and quasiconvex. Finally, in Chapter 14 we will need the slightly larger class of groups with an *almost malnormal quasiconvex hierarchy terminating in virtually compact special hyperbolic groups*. These can be built from virtually compact special hyperbolic groups by HNN extensions and amalgamated free products along almost malnormal quasiconvex subgroups.

The following theorem combines two key results to provide a natural virtually special target:

THEOREM 12.2 (AMQH \Rightarrow VS). *If G has an almost malnormal quasi-convex hierarchy [terminating in virtually compact special hyperbolic groups] then G is virtually compact special.*

PROOF. This follows by induction on the length of the hierarchy by combining Theorem 8.2 and Theorem 5.2. We note that the edge group walls employed in the proof of Theorem 8.2 yield the embedded 2-sided malnormal hyperplanes needed for Theorem 5.2 to be applicable. \square

The following central result is the main tool used to prove the main result in Chapter 14:

THEOREM 12.3 (Malnormal Special Quotient Theorem). *Let G be hyperbolic and virtually compact special. Let $\{H_1, \dots, H_r\}$ be an almost malnormal collection of quasiconvex subgroups. There exist H'_i with $[H_i : H'_i] < \infty$ such that $\bar{G} = G/\langle\langle H'_1, \dots, H'_r \rangle\rangle$ is virtually compact special and hyperbolic.*

REMARK 12.4 (Stronger form). A stronger form of Theorem 12.3 asserts that there exist finite index normal subgroups $\ddot{H}_i \subset H_i$ such that $\bar{G} = G/\langle\langle H'_1, \dots, H'_r \rangle\rangle$ is virtually compact special and hyperbolic whenever $H'_i \subset H_i$ are finite index normal subgroups with each $H'_i \subset \ddot{H}_i$. The proof is similar.

SKETCH. The plan is to choose H'_i such that \bar{G} has a finite index subgroup \bar{J} with an almost malnormal quasiconvex hierarchy, so \bar{J} is virtually compact special by Theorem 12.2.

Since G has a finite index subgroup that is the fundamental group of a compact nonpositively curved cube complex, Lemma 7.15 implies that G acts properly and cocompactly on a CAT(0) cube complex \tilde{X} .

Applying Lemma 3.28, let \tilde{Y}_i be an H_i -cocompact superconvex subcomplex for each i . Let R, L be upper bounds on the diameters of contiguous wall-pieces and cone-pieces between the translates of the \tilde{Y}_i in \tilde{X} .

Choose J to be a torsion-free finite-index normal subgroup of G such that:

- (1) $X = J \backslash \tilde{X}$ is a compact special cube complex.
- (2) $\pi_1 U \subset \pi_1 X$ is malnormal for each hyperplane U of X .
- (3) Each hyperplane U of X has “high injectivity radius” in the sense that: For any path $P \rightarrow X$ whose endpoints are on $N(U)$, if $|P| \leq 15 \max(R, L)$ then P is path homotopic into $N(U)$.
- (4) $Y_i \rightarrow X$ embeds for each i , where $Y_i = (H_i \cap J) \backslash \tilde{Y}_i$.

To obtain Property (1), note that G has a torsion-free finite index subgroup G_o by hypothesis, and the compact nonpositively curved cube complex $G_o \backslash \tilde{X}$ has a finite special cover by Corollary 4.15. Of course, any further cover is also special. Properties (2) (3) and (4) follow using separability and quasiconvexity.

Let \mathbb{W} denote the set of hyperplanes of X and consider the homomorphism $\#_{\mathbb{W}} : \pi_1 X \rightarrow \mathbb{Z}_2^{\mathbb{W}}$ that counts the number of times a path passes

through the various hyperplanes. It induces a covering space $\ddot{X} \rightarrow X$ and we let \ddot{Y}_i be the induced covers. Finally we let $H'_i = \pi_1 \ddot{Y}_i$, and claim that $\bar{G} = G/\langle\langle H'_1, \dots, H'_r \rangle\rangle$ is virtually special.

Let \bar{J} be the image of J . Observe that \bar{J} has a finite index subgroup K with the following cubical presentation:

$$\langle \bar{X} \mid g\ddot{Y}_i : g \in G, 1 \leq i \leq r \rangle$$

Each hyperplane of X induces a splitting of K along an almost malnormal quasiconvex subgroup. The almost malnormality follows from Lemma 11.5 and the quasiconvexity holds by Theorem 10.3. Taken in some order these splittings induce an almost malnormal quasiconvex hierarchy for K . \square




FIGURE 12.1. We pretend that G is torsion-free and let $\bar{X} = G \backslash \ddot{X}$ and $\bar{Y}_i = H_i \backslash \ddot{Y}_i$. We have partially illustrated $\ddot{X} \rightarrow X \rightarrow \bar{X}$ as well as $\ddot{Y}_i \rightarrow Y_i \rightarrow \bar{Y}_i$.




FIGURE 12.2. A hyperplane of X provides a splitting of $K = \pi_1 \ddot{X}^*$. We inflate cone points of relevant \ddot{Y}_i and then cut to obtain one or two smaller cubical presentations with the same properties that \ddot{X}^* had.

12.2. Case Study: $F_2/\langle\langle W_1^{n_1}, \dots, W_r^{n_r} \rangle\rangle$

We now describe the special case of Theorem 12.3 where G is free and each H_i is cyclic. This argument is extracted from [77] where it arose naturally. We warn the reader that we have been a bit cavalier about (identifying and ignoring) repeated 2-cells in covers – these arise from proper powers.

THEOREM 12.5. *Consider the presentation $\langle a, b \mid W_1, \dots, W_r \rangle$ and assume that $W_i^p \neq W_j^q$ for $i \neq j$ unless $p, q = 0$. There exists n_i which can be chosen arbitrarily large such that $\langle a, b \mid W_1^{n_1}, \dots, W_r^{n_r} \rangle$ is virtually special for all $n_i \geq 1$.*

Echoing Remark 12.4, with a bit more work we can show that there exists K such that $\langle a, b \mid W_1^{n_1 K}, \dots, W_r^{n_r K} \rangle$ is virtually special.

PROOF. We shift to the geometric viewpoint of $\langle \bar{X} \mid \bar{Y}_1, \dots, \bar{Y}_r \rangle$ where \bar{X} is a bouquet of circles and each \bar{Y}_i is an immersed circle. (We mostly maintain parallel notation with the proof of Theorem 12.3.) We will demonstrate the proof with the following very explicit example in mind: $\langle a, b \mid aba^{-1}b^{-1}, ab \rangle$.

Step 1: We first pass to a finite regular cover $\check{X} \rightarrow \bar{X}$ such that each elevation $\check{Y}_i \rightarrow \bar{Y}_i$ embeds in \check{X} . Note that there might be several distinct elevations of each such immersed circle.




FIGURE 12.3. We first pass to a finite cover \check{X} where all elevated circles embed.

Let \tilde{X} denote the universal cover of \bar{X} , and likewise \tilde{Y}_i are universal covers of \bar{Y}_i . Let M denote an upper bound on the lengths of overlaps between the various translates of the \tilde{Y}_i in \tilde{X} . While M is just 2 in our case, in general M is bounded by twice the maximal length of a relator in the presentation [31].

Step 2: We now pass to a further regular cover of $X \rightarrow \bar{X}$ factoring through \check{X} so that $\langle X \mid gY_i : g \in \text{Aut}(X), 1 \leq i \leq r \rangle$ has the property that $|P| \leq \frac{1}{8}|\partial Y_i|$ whenever P is a piece in Y_i . Here SY_i denotes the elevation of \bar{Y}_i to X . In general, one could construct X by using the residual finiteness of $\pi_1 \bar{X}$ to choose a finite index normal subgroup contained in $\pi_1 \check{X}$ that contains no elements of length $\leq 8M$. In our case the easy cover illustrated in Figure 12.4 suffices.

Step 3: The group of $\langle \bar{X} \mid Y_1, \dots, Y_r \rangle$ acts properly and cocompactly on a CAT(0) cube complex, and likewise, so does the finite index subgroup associated to $\langle X \mid gY_i : g \in \text{Aut}(X), 1 \leq i \leq r \rangle$. Indeed, we can use the mid-cell wall-structure discussed in Example 6.6, and there are sufficiently many walls to get a proper action as described in [80]. However we seek virtual specialness, and will obtain it by finding an almost malnormal quasiconvex



FIGURE 12.4. The degree 16 cover $X \rightarrow \bar{X}$ above has the property that all pieces are $\leq \frac{1}{8}$ when the 2-cells corresponding to the elevations $(ab)^8$ and $(aba^{-1}b^{-1})^4$ are added.

hierarchy – but this comes at the expense of further covers $\ddot{Y}_i \rightarrow Y_i$ corresponding to multiplying the exponent m_i of $W_i^{m_i}$ by an additional factor of 2.

Let \mathbb{W} denote the set of hyperplanes in X – so \mathbb{W} corresponds to the set of edges, and consider the homomorphism $\#_{\mathbb{W}} : \pi_1 X \rightarrow \mathbb{Z}_2^{\mathbb{W}}$ that counts the number of times a path traverses the distinct edges modulo 2. Let $\check{X} \rightarrow X$ denote the associated finite cover, and let $\check{Y}_i \rightarrow Y_i$ denote the induced double covers (since $Y_i \hookrightarrow X$ is an embedded cycle, \check{Y}_i is indeed a double cover.)

Our desired quotient group corresponds to the presentation

$$(12.1) \quad \langle \bar{X} \mid \check{Y}_1, \dots, \check{Y}_r \rangle$$

which corresponds to the following presentation where $n_i = \deg(\check{Y}_i \rightarrow \bar{Y}_i)$.

$$\langle a, b \mid W_1^{n_1}, \dots, W_r^{n_r} \rangle$$

The cover $\check{X} \rightarrow \bar{X}$ induces a covering space of the standard 2-complex of Presentation (12.1), and this cover corresponds to the following presentation:

$$(12.2) \quad \langle \check{X} \mid g\check{Y}_i : g \in \text{Aut}(\check{X}), 1 \leq i \leq r \rangle$$

Let \check{X}^* denote the 2-complex associated to Presentation (12.2).

Following Definition 6.3, each edge e of X determines a (1-sided) track T_e in the 2-complex associated to the cubical presentation:

$$(12.3) \quad \langle X \mid g\check{Y}_i : g \in \text{Aut}(X), 1 \leq i \leq r \rangle$$

The vertices of the track are the points at the centers of the preimages of e . The edges of the track are midcells of the 2-cells corresponding to the $g\check{Y}_i$, which we emphasize are attached by double covers of the simple cycles gY_i . See Figure 12.5.

Step 4: We claim that each track \check{T}_e in \check{X}^* is π_1 -injective, 2-sided, and malnormal in $\pi_1 \check{X}^*$. Moreover, the collection of splittings along the $\pi_1 \check{T}_e$ tracks provides a hierarchy. We therefore get a malnormal quasiconvex hierarchy for $\pi_1 \check{X}^*$.

REMARK 12.6. We thus get a malnormal hierarchy for the cubulation dual to the walls associated with these tracks, and so the proof of virtual specialness relies completely on repeatedly applying Theorem 5.1 and it is




FIGURE 12.5. We give a heuristic picture of T_e in our example on the left, and the track \tilde{T}_e on the right. (The 2-cells would actually be folded with each other a bit more in some places since there are some length 2 pieces.) Note that T_e has exactly three edges since an edge e lies on the boundary of a 2-cell in three ways. Each 2-cell is attached along a double cover of an embedded cycle Y_i in X , but we have drawn each 2-cell as the cone on the boundary of a very skinny moebius strip. An edge of the track travels from the center of e to the conepoint of the 2-cell on the “top” layer and then travels back to the center of e on the bottom layer. (If we collapse multi 2-cells, then a preimage \tilde{T}_e of the track T_e is a 3-valent graph that is 2-sided (without involutions) cutting through \check{X}^* and providing a splitting of $\pi_1\check{X}^*$.)

not necessary to invoke Theorem 12.2. I had expected similar reasoning to hold in the higher dimensional case, but there are difficulties verifying that the action on the dual cube complex is proper. This difficulty is side-stepped by employing the alternate cubulation afforded by using Theorem 8.2 within Theorem 12.2.

The π_1 -injectivity holds by observing that the 2-complex $N_e = N(T_e)$ carrying the track T_e immerses by a map that has no missing shells, so we can apply Theorem 13.3, and likewise apply Theorem 13.4 to verify quasiconvexity. The track is readily seen to have a 2-sided preimage in the cover $\check{N}_e \rightarrow N_e$ that is induced by \check{X} . A subtle point is that $\pi_1\check{X}^*$ acts with inversions on the walls whereas $\pi_1\check{X}^*$ acts without inversions.

We obtain a hierarchy since each edge e will be cut by some track. However, in this case, our tracks, or rather our walls correspond precisely to those studied in [80] since these antipodal walls are of the type considered there. Thus $\pi_1\check{X}^*$ acts freely and cocompactly on the cube complex $\tilde{C} = \tilde{C}(\check{X})$ dual to the wallspace obtained from these tracks. Let $C = \pi_1\check{X}^*\backslash\tilde{C}$. Then C has a hierarchy obtained by repeatedly cutting along 2-sided malnormal

hyperplanes. We thus obtain virtual specialness by repeatedly applying Theorem 5.1. (In this case we have avoided Theorem 12.2 and in particular, we do not need Theorem 8.2.)




FIGURE 12.6. Thick and thin classical annular diagrams. Width 2 is impossible because of the four pieces. Width 1 implies that $g\tilde{N}_e = \tilde{N}_e$. Width 0 implies that $g\tilde{N}_e \cap \tilde{N}_e \neq \emptyset$.

We now examine the malnormality. Consider an annular diagram $A \rightarrow X^*$ that has both boundary cycles on the same track carrier (so \tilde{A} lifts to have both boundary lines on translates of the same wall carrier). Since the carrier is immersed and has no missing shells, we can choose A to have no spurs or shells. By Theorem 11.4, we can replace A with an annular diagram B in the same class such that B is either thin or is thick and contains annuladders L_1, L_2 etc. However, the thick (width 2) case is excluded since there is no essential path $P \rightarrow N(T_e)$ such that P is the concatenation of ≤ 4 pieces and such that the initial and terminal edge of P is e . Indeed, each piece is $\leq \frac{1}{8}$. In the thin case, successive 2-cells of B meet along pieces, and so the annuladders L_1, L_2 containing the tracks within B must “line up” since no track T_e travels twice through the same piece (as the Y_i embed in X). Consequently, as \ddot{T}_e is a deformation retract of \ddot{N}_e , the two elements are actually conjugated by an element of $\pi_1 \ddot{N}_e = \pi_1 \ddot{T}_e$ as required by malnormality. \square

12.3. The Special Quotient Theorem

We now discuss the following consequence of Theorem 12.3:

THEOREM 12.7 (Special Quotient Theorem). *Let G be hyperbolic and virtually compact special. Let $\{H_1, \dots, H_r\}$ be a collection of quasiconvex subgroups. There exist H'_i with $[H_i : H'_i] < \infty$ such that $\bar{G} = G/\langle\langle H'_1, \dots, H'_r \rangle\rangle$ is virtually compact special and hyperbolic.*

SKETCH. Let $\{K_1, \dots, K_s\}$ denote the collection of infinite maximal intersections of conjugates given in the statement of Lemma 3.22. For each i , let \mathcal{K}_i denote $\mathbb{C}_G(K_i)$, so we have the almost malnormal collection $\{\mathcal{K}_1, \dots, \mathcal{K}_s\}$.

We apply Theorem 12.3 to obtain finite index subgroups $\mathcal{K}'_i \subset \mathcal{K}_i$ and a virtually compact special hyperbolic quotient $\bar{G} = G/\langle\langle \mathcal{K}'_1, \dots, \mathcal{K}'_s \rangle\rangle$. We choose this quotient so that $\text{Height}_{\bar{G}}\{\bar{H}_1, \dots, \bar{H}_r\} < \text{Height}_G\{H_1, \dots, H_r\}$.

Consequently, the theorem follows by induction on the height. The quotient group has trivial kernel in the base case when the height is 0.

There is a computation to ensure that the height does actually decrease. For this we must choose the \mathcal{K}'_i to avoid certain elements of \mathcal{K}_i – they then have sufficiently large systole so that small-cancellation methods are applicable. The tool supporting this computation is stated in Lemma 12.8. It is applied after passing to a finite index torsion-free subgroup of G that is the fundamental group of a special cube complex X . The Y_i correspond to our subgroups \mathcal{K}'_i , and the A_j correspond to the H_j (and their intersections). \square

LEMMA 12.8 (Controlling Intersections and Double Cosets). *Let $\langle X \mid Y_1, \dots, Y_k \rangle$ be a $C'(\frac{1}{24})$ cubical presentation. Let $A_1 \rightarrow X$ and $A_2 \rightarrow X$ be based local isometries. Suppose that for each pair of lifts \tilde{A}_j, \tilde{Y}_i to \tilde{X} , we have either $\tilde{Y}_i \subset \tilde{A}_j$ or $\text{diam}(\tilde{Y}_i \cap \tilde{A}_j) < \frac{1}{8}\|Y_i\|$.*

Let $\{\pi_1 A_1 g_i \pi_1 A_2\}$ be a collection of distinct double cosets in $\pi_1 X$. Suppose that for each chosen representative g_i and each Y_j we have $|g_i| < \frac{1}{8}\|Y_j\|$.

Let $G \rightarrow \bar{G}$ denote the quotient $\pi_1 X \rightarrow \pi_1 X^$. Then:*

- (1) $\overline{\pi_1 A_1 \bar{g}_i \pi_1 A_2}$ and $\overline{\pi_1 A_1 \bar{g}_j \pi_1 A_2}$ are distinct for $i \neq j$.

Suppose now that the cosets $\pi_1 A_1 g_i \pi_1 A_2$ form a complete set of double cosets with the property that $\pi_1 A_1^{g_i} \cap \pi_1 A_2$ is infinite. Then:

- (2) *If $\overline{\pi_1 A_1^{\bar{g}} \cap \pi_1 A_2}$ is infinite for some $\bar{g} \in \bar{G}$, then $\overline{\pi_1 A_1 \bar{g} \pi_1 A_2} = \overline{\pi_1 A_1 \bar{g}_i \pi_1 A_2}$ for some i .*
- (3) *For each g_i we have: $\overline{\pi_1 A_1^{g_i} \cap \pi_1 A_2} = \overline{\pi_1 A_1^{\bar{g}_i} \cap \pi_1 A_2}$.*

Hyperbolicity and Quasiconvexity Detection

In this Chapter we describe how to apply the cubical small-cancellation theory developed in Chapter 9, to see that hyperbolicity is preserved under sufficiently small-cancellation quotienting. As for classical small-cancellation groups, we expect one could directly verify with a disk diagram argument that the quotient has δ -thin triangles. However, the Ladder Theorem fits neatly with Papasoglu’s “thin bigon criterion”, and so we prove hyperbolicity following this approach in Section 13.1. Such results are proven in [62, 30] for a general relatively hyperbolic group G , but it is interesting to find that there is a simple approach in the cubical context.

In Section 13.2 we show that quasiconvexity of a subgroup persists in a sufficiently small-cancellation quotient. The method employed to verify quasiconvexity is the “no missing shell criterion” described in Section 13.3. This is a generalization of the notion of local isometry discussed in Section 3.12.

13.1. Cubical Version of Filling Theorem

THEOREM 13.1. *Let X be a compact nonpositively curved cube complex with $G = \pi_1 X$ hyperbolic relative to $\{P_1, \dots, P_r\}$. There exist finite subsets $S_i \subset P_i - \{1\}$ such that $G/\langle\langle P'_1, \dots, P'_r \rangle\rangle$ is hyperbolic whenever P'_i are normal subgroups with P_i/P'_i hyperbolic and $S_i \cap P'_i = \emptyset$.*

PROOF. By Lemma 3.28, let \tilde{F}_i be a P_i -cocompact superconvex subcomplex of \tilde{X} for each i . Let R be an upper bound on the sizes of contiguous wall-pieces on the \tilde{F}_i (as rectangular flaps hanging off \tilde{F}_i are bounded using superconvexity and cocompactness). Let M be an upper bound on the sizes of contiguous cone-pieces between the \tilde{F}_i and their translates (which exists by cocompactness and almost malnormality).

Choose S_i to represent a conjugate of each nontrivial “short” element in P_i – i.e. those whose translation length in \tilde{F}_i is $\leq 24 \max(R, M)$.

Consider normal subgroups P'_i such that $S_i \cap P'_i = \emptyset$ and P_i/P'_i is hyperbolic for each i . Consider the quotient $G/\langle\langle P'_1, \dots, P'_r \rangle\rangle = \pi_1 X^*$ where $X^* = \langle X \mid F_1, \dots, F_r \rangle$ and $F_i = P'_i \setminus \tilde{F}_i$. Note that X^* is $C'(\frac{1}{24})$.

We show that $\pi_1 X^*$ is hyperbolic by verifying Papasoglu’s thin bigon criterion for hyperbolicity of graphs [63]. It states that a graph is δ -hyperbolic if there exists ρ such that any two combinatorial geodesics with the same endpoints must ρ -fellow-travel.

Consider a pair of geodesics γ_1, γ_2 in \tilde{X}^* and a minimal complexity disk diagram E between them. The square convex hulls of γ_1, γ_2 in E determine two square subdiagrams D_1, D_2 with a diagram D between them such that D has no cornsquares at the top or bottom λ_1, λ_2 . Note that $|\lambda_i| = |\gamma_i|$ for each i , and in particular λ_i are geodesics.




FIGURE 13.1. $E = D_1 \cup_{\lambda_1} D \cup_{\lambda_2} D_2$.

By Theorem 9.3, D is a ladder (or single cone-cell or 0-cell). Indeed, D has no cornsquares (by maximality of D_1, D_2) or shells along λ_1, λ_2 (by short innerpaths and geodesicity of λ_1, λ_2).

γ_i, λ_i must σ -fellow-travel relative to parabolics for some $\sigma > 0$ by the relative hyperbolicity of \tilde{X} . We thus have a situation as illustrated in Figure 13.2.




FIGURE 13.2. E consists of a ladder D sandwiched between two parabolic fellow-traveling bigons D_1, D_2 .

The cone-pieces in the ladder D have size $\leq M$ and the wall-pieces have size $\leq R$. Thus λ_1, λ_2 R -fellow-travel along each rectangular grid and $\kappa_i M$ -fellow-travel along cones. Here we use that P/P'_i is hyperbolic to see that $P'_i \setminus \tilde{F}_i$ is δ_i -hyperbolic so geodesics that start and end within M must κ_i -fellow-travel in $P'_i \setminus \tilde{F}_i$. See the two diagrams at the left of Figure 13.3.

Finally, λ_i, γ_i uniformly fellow-travel since they σ -fellow-travel relative to parabolics, and so we can break them up into a sequence of corresponding segments that are either: images of σ -close pairs in \tilde{X} , or pairs of geodesics in $P'_i \setminus \mathcal{N}_\sigma(\tilde{F}_i)$ that start and end within σ of each other. They thus uniformly fellow-travel. We refer to the diagram on the right of Figure 13.3. \square




FIGURE 13.3. The paths γ_1, γ_2 fellow-travel since their subpaths fellow-travel around each rectangular grid and cone. In the latter case, the subpaths are geodesics in $P'_i \setminus \tilde{F}_i$ whose endpoints are $\max(M, R)$ -close. The paths γ_i, λ_i fellow-travel in \tilde{X}^* since for each parabolic \tilde{F}_i that they travel around, their subpaths γ'_i, λ'_i where on either side, have σ -close endpoints, and thus fellow-travel within $P'_i \setminus \mathcal{N}_\sigma(\tilde{F}_i)$.

13.2. Persistence of Quasiconvexity

Recall that $H \subset G$ is *full* if H intersects each parabolic subgroup P_i^g in a finite or finite index subgroup of P_i^g . We follow the terminology of Theorem 13.1.

THEOREM 13.2. *Let G be hyperbolic relative to $\{P_1, \dots, P_r\}$. Suppose G acts freely and cocompactly on a $CAT(0)$ cube complex \tilde{X} . Let $H \subset G$ be a full relatively quasiconvex subgroup. There exist (slightly larger) finite sets S_i^+ of nontrivial elements such that \tilde{H} is quasiconvex in \tilde{G} whenever $P_i' \cap S_i^+ = \emptyset$.*

PROOF. Apply Lemma 3.29 to obtain superconvex P_i -cocompact subcomplexes \tilde{F}_i with $\tilde{F}_i = P_i K_i$ where K_i is compact. As H is relatively quasiconvex, there are finitely many P_i^g with $|H \cap P_i^g| = \infty$. Applying Lemma 3.29 again, we choose a superconvex H -cocompact core \tilde{A} that contains gK_i and hence $g\tilde{F}_i$ whenever $|H \cap P_i^g| = \infty$. Let $A = H \setminus \tilde{A}$. It follows that there exists D such that $\text{diam}(g_j \tilde{F}_j \cap \tilde{A}) < D$ unless $g_j \tilde{F}_j \subset \tilde{A}$ (by fullness and cocompactness).

Choose S_i^+ so that for any allowable choice of P_i' , the induced presentation $A^* \rightarrow X^*$ has no missing shells and has short innerpaths. Then $\tilde{A}^* \subset \tilde{X}^*$ is an isometric embedding. \square

13.3. No Missing Shells and Quasiconvexity

Let X be a 2-complex satisfying the $C'(\frac{1}{6})$ condition. We say $Y \rightarrow X$ has *no missing shells* if for any 2-cell $R \rightarrow X$ with $\partial_p R = QS$ and $|Q| > |S|$, any lift of $Q \rightarrow R \rightarrow X$ to $Q \rightarrow Y$ extends to a lift $R \rightarrow Y$:

$$\begin{array}{ccc} Q & \rightarrow & Y \\ \downarrow & \nearrow & \downarrow \\ R & \rightarrow & X \end{array}$$

A map $Y \rightarrow X$ with no missing shells behaves very much like a local isometry (e.g. of cube complexes). In particular, we have the following property that was intensely exploited in [54].

THEOREM 13.3. *Let X be a $C'(\frac{1}{6})$ -complex. If $Y \rightarrow X$ has no missing shells then $\tilde{Y} \rightarrow \tilde{X}$ is an injection. In particular $\pi_1 Y \rightarrow \pi_1 X$ is injective.*

PROOF. Suppose there are two distinct 0-cells u, v of \tilde{Y} that map to the same 0-cell of \tilde{X} . Let $D \rightarrow \tilde{X}$ be a minimal area diagram among all those whose boundary path lifts to a path in \tilde{Y} from u to v . By Theorem 9.2, D is either trivial, or is a single 2-cell, or contains at least two spurs and/or shells. D cannot be trivial since $u \neq v$. In the remaining cases, there is a shell R whose outerpath Q is a subpath of $\partial_p D$. Since $Y \rightarrow X$ has no missing shells, the path $Q \rightarrow \tilde{Y}$ extends to $R \rightarrow \tilde{Y}$ and this allows us to produce a smaller area diagram D' obtained by removing R and replacing Q by S . \square




FIGURE 13.4. $Y \rightarrow X$ has no missing shell.

THEOREM 13.4. *Let X be a compact $C'(\frac{1}{6})$ -complex. If $Y \rightarrow X$ has no missing shells then $\tilde{Y} \rightarrow \tilde{X}$ is a quasiconvex subcomplex.*

PROOF. Let D be a minimal area diagram (and minimal number of 1-cells) between a geodesic $\gamma \rightarrow \tilde{X}$ and a variable path $\sigma \rightarrow \tilde{Y}$ with the same endpoints as γ . Theorem 9.2 implies that D is a ladder. Indeed, there is no shell or spur of D on σ by minimality and the no missing shell hypothesis. And there is no shell or spur on γ because of the geodesic hypothesis. Finally we see that the geodesic γ lies in $\mathcal{N}_r(\tilde{Y})$ where r is half the maximal circumference of a 2-cell of X . \square

The no missing shell definition as well as the accompanying consequences have natural cubical small-cancellation generalizations.

Let $A^* = \langle A \mid B_1, \dots \rangle$ and $X^* = \langle X \mid Y_1, \dots \rangle$ be cubical presentations. A map of cubical presentations $f : A^* \rightarrow X^*$ has the property that $f : A \rightarrow X$ is a local isometry of nonpositively curved cube complexes, and for each B_i




FIGURE 13.5. $\tilde{Y} \rightarrow \tilde{X}$ is quasiconvex.

there is $Y_{i'}$ such that B_i maps to X through $Y_{i'}$ as follows:

$$\begin{array}{ccc} B_i & \rightarrow & A \\ \downarrow & & \downarrow \\ Y_{i'} & \rightarrow & X \end{array}$$

The map $f : A^* \rightarrow X^*$ has *no missing shells* if the following holds: For each closed essential path $QS \rightarrow X$ mapping to some Y_j with $|Q| > |S|$ and $S \rightarrow \tilde{X}$ a geodesic, if there is a lift $Q \rightarrow A$, then there exists B_i with $Y_{i'} = Y_j$ and a lift of QS to B_i so that there is a commutative diagram:

$$\begin{array}{ccccc} & & B_i & \rightarrow & A \\ & \nearrow & \downarrow & & \downarrow \\ QS & \rightarrow & Y_j & \rightarrow & X \end{array}$$

The proof of the following is analogous to that of Theorems 13.3 and 13.4.

THEOREM 13.5. *Let X^* be a cubical presentation that satisfies the $C'(\frac{1}{24})$ small-cancellation condition. Let $A^* \rightarrow X^*$ have no missing shells. Then $\tilde{A}^* \rightarrow \tilde{X}^*$ is an embedding. Moreover, if X^* is compact then $\tilde{A}^* \rightarrow \tilde{X}^*$ is quasiconvex.*

Hyperbolic groups with a quasiconvex hierarchy

DEFINITION 14.1 (Quasiconvex Hierarchy). A *quasiconvex hierarchy* for a group G is a specific way of building G by starting with trivial groups, and then using a finite sequence of HNN extensions and/or Amalgamated Free Products over finitely generated quasi-isometrically embedded subgroups. The *length* of the hierarchy is the number of splittings used to build G .

Alternatively, we can think of the hierarchy as a way of decomposing the group G by splitting along f.g. quasi-isometrically embedded subgroups, and then repeating the process on the result, until one arrives at trivial subgroups.

For many purposes, it is natural to allow these terminal subgroups to be finite instead of trivial.

EXAMPLE 14.2. The fundamental group of a closed genus 2 surface has a quasiconvex hierarchy: Infinite cyclic groups are built as HNN extension over a trivial groups, and then one builds free groups by taking free products of cyclics, and finally one builds the group G using an amalgamated product over a cyclic subgroup.

Recall that within a hyperbolic group, being a quasiconvex subgroup is equivalent to being f.g. and quasi-isometrically embedded. The trivial subgroup is obviously quasiconvex, and so is any infinite cyclic subgroup of a hyperbolic group.

The main result described in these notes is:

THEOREM 14.3. *If G is a hyperbolic group with a quasiconvex hierarchy. Then G is virtually special.*

While generalizations of Theorem 14.3 hold under various ways of relaxing the hypotheses, we cannot go too far. For instance, for $n \neq \pm 1$, the Baumslag-Solitar group $BS(1, n)$ presented by $\langle a, t \mid a^t = a^n \rangle$ is not virtually special as it cannot be cubulated. However $BS(1, n)$ does have a hierarchy, but not a quasiconvex hierarchy. There are other examples which can be cubulated but are still not virtually special – as there is not sufficient hyperbolicity to enable critical points in the proof. An interesting class of non-virtually special examples arise from complete square complexes that are not virtual products. See Example 2.6.

Among the various ways that one could hope to generalize Theorem 14.3 is to prove the following variant omitting the quasiconvexity requirement:

CONJECTURE 14.4. *Let G be a hyperbolic group with a hierarchy terminating in finite subgroups. Then G is virtually [compact] special.*

Theorem 14.3 holds by induction on the length of the hierarchy as a consequence of the following:

THEOREM 14.5. *Let $G = A *_C B$ or $A *_C t=D$ be a hyperbolic group such that C is quasiconvex and A, B are virtually compact special. Then G is virtually compact special.*

PROOF. **Plan:** We will separate C from its intersecting conjugators in a finite quotient $G \rightarrow F$. The kernel G' splits as a graph of groups where all edge groups are malnormal and quasiconvex. Thus G' has a malnormal quasiconvex hierarchy terminating at virtually compact special groups. So G' is virtually compact special by Theorem 12.2 and thus G is as well.

Recall that an element $g \in G$ is an *intersecting conjugator* of C provided that $|C^g \cap C| = \infty$. We refer to Lemmas 3.20 for a motivational observation about intersecting conjugates and separability.

Inductively achieving plan: It suffices to consider the HNN extension case $G = A *_C t=D$. By Lemma 3.19, let $g_1, \dots, g_r \in G$ denote the finitely many intersecting conjugators of C in G . Each g_i has a normal form $g_i = a_{i1} t^{\pm \epsilon_{i1}} a_{i2} t^{\pm \epsilon_{i2}} \dots$ for some $a_{ij} \in A$.

Let $\{H_1, \dots, H_r\}$ denote the representatives of maximal infinite intersections of conjugates of C . The idea is to use Theorem 12.3 to choose finite index subgroups $\{H'_1, \dots, H'_r\}$ such that there is a quotient:

$$\begin{array}{ccc} G & = & A *_C t=D \\ \downarrow & & \downarrow \\ \bar{G} & = & \bar{A} *_C t=\bar{D} \end{array}$$

where $\bar{A} = A / \langle\langle H'_1, \dots, H'_r \rangle\rangle$ and \bar{C}, \bar{D} are the images of C, D in \bar{A} ; and the conjugation homomorphism $t : C \rightarrow D$ projects to $t : \bar{C} \rightarrow \bar{D}$; and $\bar{C} \subset \bar{G}$ is quasiconvex; and $\text{Height}_{\bar{G}}(\bar{C}) < \text{Height}_G(C)$; and the intersecting conjugator images \bar{g}_i are outside of \bar{C} . As \bar{G} is virtually special by induction, we can then apply Theorem 4.13 to separate \bar{C} from all the \bar{g}_i in a finite quotient $\bar{G} \rightarrow F$. The composition $G \rightarrow \bar{G} \rightarrow F$ yields the desired finite quotient enabling our plan.

Most of the above are fairly natural for large systole choices of the H'_i . The justification requires: the normal form theorem, small-cancellation theory, criteria for quasiconvexity and hyperbolicity, and a lemma to compute the intersection of images of subgroups.

A key point is to choose $A \rightarrow \bar{A}$ so that it induces:

$$(14.1) \quad \begin{array}{ccc} C & \xrightarrow{t} & D \\ \downarrow & & \downarrow \\ \bar{C} & \xrightarrow{t} & \bar{D} \end{array}$$

Let T denote the Bass-Serre tree of $G = A *_C t=D$. Let $\{S_i\}$ denote a collection of representatives of G -orbits of finite maximal subtrees of T with the

property that $\text{Stabilizer}(S_i)$ is infinite. For each i , let $\{H_{i1}, \dots, H_{iq_i}\}$ denote a collection of subgroups of A conjugate to the subgroups of $\text{Stabilizer}(S_i)$ stabilizing distinct $\text{Stabilizer}(S_i)$ orbits of vertices of S_i – we choose one such subgroup for each orbit.

One checks that their commensurators $\{\mathbb{C}_A(H_{ij}) : 1 \leq i \leq r, 1 \leq j \leq q_i\}$ form an almost malnormal collection in A . We then apply Theorem 12.3 to A relative to $\{\mathbb{C}_A(H_{ij}) : 1 \leq i \leq r, 1 \leq j \leq q_i\}$ to obtain the virtually compact special hyperbolic quotient $A \rightarrow \bar{A}$. With some care, this can be done so that the commutative diagram in Equation (14.1) is satisfied. A useful point here is that for each tree S_i , its vertex stabilizers form a poset under inclusion, where the the largest element is the stabilizer of a vertex that is fixed by all of $\text{Stabilizer}(S_i)$. For each i , these inclusions provide corresponding embeddings between the subgroups $\{H_{i1}, \dots, H_{iq_i}\}$ of A . We refer to Example 14.6. We use these isomorphisms to “transfer” the \check{H}_{ij} constraints given by Theorem 12.3 and Remark 12.4. In addition, we transfer the desired systole properties ensuring small-cancellation and the separation properties to facilitate the plan.

Theorem 13.5 applies to compute the presentations for the images \bar{C}, \bar{D} , of C, D in \bar{A} . Specifically, we see that \bar{C} is the quotient of C by the normal closures in C of the intersections of C with the \check{H}_{ij} . The analogous statement holds for \bar{D} . The isomorphism between the various H'_{ij} subgroups that we are quotienting by ensures that the isomorphism $C^t = D$ projects to an isomorphism $\bar{C}^t = \bar{D}$.

We see that $\text{Height}_{\bar{G}}(\bar{C}) < \text{Height}_G(C)$ by arguing that the subtrees of \bar{T} with infinite stabilizer are images of subtrees of T with infinite stabilizer under the equivariant map $T \rightarrow \bar{T}$ between the Bass-Serre trees of HNN splittings of G and \bar{G} . As \bar{C} has finite height in \bar{G} , the quasiconvexity of \bar{C} in $\bar{G} = \bar{A} *_{\bar{C}^t = \bar{D}}$ holds by a criterion in [57], and the hyperbolicity of \bar{G} holds by a criterion in [8]. \square




FIGURE 14.1. The two spaces on the right immerse in the graph of spaces corresponding to $A *_{C^t = D}$ on the left. Each corresponds to the maximal ∞ -stabilizer tree S_1, S_2 illustrated above it.

EXAMPLE 14.6 (Transferring the constraints). Consider the HNN extension $G \cong A \star_{C^t=D}$ depicted on the left in Figure 14.1. The group A is free on $\langle a, b, c, x, y, z \rangle$ where a, b, c are the simple arrows, and x, y, z are the triangled arrows. There are two G orbits of maximal finite trees with infinite stabilizer. In this simple example, these stabilizers are infinite cyclic. The first is associated to the subgroups $\{\langle ab \rangle, \langle a \rangle, \langle b \rangle\}$ of A . The second is associated to the subgroups $\{\langle cyy \rangle, \langle y \rangle, \langle z \rangle, \langle x \rangle, \langle xcb \rangle\}$. Together, these form a malnormal collection of subgroups of A , to which we will apply Theorem 7.4.

For the sake of exposition, let us assume that virtual compact specialness is guaranteed when we quotient by any powers that are multiples of $\{3, 5, 7\}$ and $\{4, 5, 6, 7, 8\}$ respectively, and moreover, assume that the other supporting properties are achieved when we quotient by powers that are multiples of $\{10, 11, 13\}$ and $\{10, 11, 12, 13, 14\}$. We then transfer this information using the trees to obtain consistent quotients. For the first tree, any powers of the form

$$\{n \cdot (4) \cdot 5 \cdot (2 \cdot 7) \cdot 10 \cdot 11 \cdot 13, n \cdot (3 \cdot 4) \cdot 5 \cdot (2 \cdot 7) \cdot 10 \cdot 11 \cdot 13, n \cdot (3 \cdot 4) \cdot 5 \cdot (7) \cdot 10 \cdot 11 \cdot 13\}$$

will be appropriate. The second tree is handled similarly. More generally, in a noncyclic situation, one must transfer arbitrary finite index subgroups, and these are not indexed by integers, but the same reasoning applies.

The relatively hyperbolic setting

The goal of this chapter is to map out a relatively hyperbolic generalization of Theorem 14.3. The explanations are much more sketchy than in the chapters on other topics. We hope the reader will be able to parse the statements and see how the pieces fit together, and we refer to the exposition in [76] for the details. We refer to sparse cube complexes which were discussed in Section 7.4, but the reader should read with the simpler compact case in mind.

The main theme that arises in the proofs is to obtain separability information about G using quotients $G \rightarrow \bar{G}$ where \bar{G} is virtually special since it is a hyperbolic group with a quasiconvex hierarchy. One uses separability to pass to a finite index subgroup G' that can be cubulated in the sense that $G' = \pi_1 X$ where X is a nonpositively curved cube complex. One again uses separability to pass to a finite cover \hat{X} with embedded 2-sided hyperplanes, and thus a cubical hierarchy. One again uses the virtually compact special quotients $\pi_1 \hat{X} \rightarrow \pi_1 \hat{X}^*$ to verify the double hyperplane coset criterion, and obtain virtual specialness of \hat{X} .

Here is a more detailed description of the results and how they are related. When G_v is virtually special and hyperbolic relative to abelian subgroups, Lemma 15.1 permits us to quotient the parabolic subgroups to obtain a hyperbolic virtually special quotient \bar{G}_v . Thus if G is hyperbolic relative to abelians and has a quasiconvex hierarchy, we inductively assume its vertex groups G_v are virtually special, and we apply Lemma 15.1 simultaneously to all vertex groups of G to obtain quotients $G \rightarrow \bar{G}$. The quotient \bar{G} splits as a graph of virtually special hyperbolic groups with quasiconvex edge groups and so \bar{G} is virtually special by Theorem 14.3. We use this strategy in Theorem 15.4 to verify the separability of quasiconvex subgroups $H \subset G$: for $g \notin H$, we find a special hyperbolic quotient \bar{G} with \bar{g} outside \bar{H} which is separable since it is quasiconvex. Separability allows us to pass to a finite index subgroup G' of G such that the induced splitting of G' has relatively malnormal edge groups. Under additional simplifying assumptions, Theorem 8.5 cubulates G' , and we find that $G' = \pi_1 X$ where X is a nonpositively curved cube complex with a hierarchy. This cube complex is virtually special since the parabolic fillings allow us to verify the double coset criterion for virtual specialness of Theorem 4.6.

LEMMA 15.1 (Virtually Special Fillings). *Let G be hyperbolic relative to the virtually abelian subgroups $\{P_1, \dots, P_r\}$, and suppose G is virtually*

sparse special. There exist finite index subgroups P_i^o such that for any finite index subgroups $P_i' \subset P_i^o$, the quotient $\bar{G} = G/\langle\langle P_1', \dots, P_r' \rangle\rangle$ is virtually compact special hyperbolic.

SKETCH. This follows the same strategy as the proof of Theorem 12.3. The subgroups P_i^o are chosen so that \bar{G} has a finite index subgroup with a quasiconvex hierarchy, and so \bar{G} is virtually special by Theorem 14.3. We then follow the proof of Theorem 12.3.

The proof utilizes a cubical presentation $X^* = \langle X \mid F_1, \dots, F_r \rangle$ where each $F_i = P_i' \backslash \tilde{F}_i$, and \tilde{F}_i is a superconvex P_i -cocompact complex afforded by a variant of Lemma 3.29. Sufficient systole and wallspace properties are chosen for $P_i^o \backslash \tilde{F}_i$, and these properties are preserved by the further cover $F_i = P_i' \backslash \tilde{F}_i$. The quotient $\bar{G} = \pi_1 X^*$ is hyperbolic by Theorem 13.1. \square

The following consequence of Lemma 15.1 follows since \bar{G} is virtually torsion-free, and each $\bar{P}_i = P_i \backslash P_i'$.

COROLLARY 15.2 (Controlling Cusps). *Let M be a hyperbolic 3-manifold with boundary components T_1, \dots, T_r . There exist finite covers $\widehat{T}_i^o \rightarrow T_i$ such that for any further finite covers $\widehat{T}_i \rightarrow \widehat{T}_i^o$, there is a finite regular cover $\widehat{M} \rightarrow M$ such that the induced covers of the boundary components are \widehat{T}_i .*

Consider a nonpositively curved cube complex X with a 2-sided embedded hyperplane P , and note that $X - N^o(P)$ consists of nonpositively curved cube complexes. These are said to be obtained from X by *cutting along a hyperplane*. We say X has a *cubical hierarchy* if we arrive at 0-cubes after cutting along hyperplanes finitely many times.




FIGURE 15.1. A cubical hierarchy of length 14.

THEOREM 15.3. *Let X be a nonpositively curved cube complex then X is virtually special provided:*

- (1) X is sparse
- (2) $\pi_1 X$ is hyperbolic relative to virtually abelian subgroups
- (3) X has a cubical hierarchy

SKETCH. A collection of disjoint hyperplanes of X induces a splitting of $G = \pi_1 X$ as a graph of groups, where the vertex groups are virtually special by induction on the length of the cubical hierarchy. We apply Lemma 15.1 simultaneously to each of these vertex groups to get a quotient $G = \bar{G}$ that splits as a graph of virtually special hyperbolic groups with quasiconvex edge groups. The hyperbolicity of \bar{G} follows from Theorem 13.1 and the

quasiconvexity of the edge groups follows from a variant of Lemma 13.4. Although a parabolic subgroup can intersect multiple vertex groups and intersect each of these in several distinct conjugacy classes, the flexibility of Lemma 15.1 allows us to choose the parabolic quotients of the vertex groups so that they induce the same quotient on each parabolic subgroup of G .

A variant of Lemma 12.8 applies to ensure that double hyperplane cosets can be separated in each such \bar{G} quotient. As the images are double quasiconvex cosets, they can be separated in a finite quotient $\bar{G} \rightarrow F$ by Theorem 4.16. We thus obtain virtual specialness of the cube complex by Theorem 4.6. \square

THEOREM 15.4. *Let G be hyperbolic relative to abelian subgroups and suppose G splits as a graph of groups with virtually special vertex groups and quasiconvex edge groups. Then quasiconvex subgroups of G are separable.*

PROOF. Use Lemma 15.1 to quotient the vertex groups to hyperbolic virtually special groups. Note that we also quotient cyclic groups in the vertex groups that arise as intersections of vertex groups and non-cyclic parabolic subgroups of G . A key point is to choose the quotienting so that there is a consistently induced manner to quotient parabolics in edge groups, but Lemma 15.1 can accommodate this, and we obtain a quotient $G \rightarrow \bar{G}$. Large enough fillings ensure that the edge groups of G have quasiconvex images in \bar{G} , and that our quasiconvex subgroup H has quasiconvex image \bar{H} and an element $g \notin H$ maps to $\bar{g} \notin \bar{H}$. Since \bar{G} is virtually special by Theorem 14.3 (or rather, a variant of it terminating in virtually compact special hyperbolic groups) we see that \bar{H} can be separated from \bar{g} by Theorem 4.13. \square

THEOREM 15.5. *Let G be hyperbolic relative to abelian groups. Suppose G splits as a graph of groups with quasiconvex edge groups. Suppose each edge group is hyperbolic and has trivial intersection with \mathbb{Z}^2 subgroups of the vertex groups. If each vertex group is virtually compact [sparse] special, then so is G .*

We record the following specific case of Theorem 15.5:

COROLLARY 15.6. *Suppose G splits as a graph of groups with quasiconvex edge groups and hyperbolic virtually compact special vertex groups. Then G is virtually compact special.*

SKETCH OF PROOF OF THEOREM 15.5. We apply Theorem 15.4 and Lemma 3.21 to each edge group of G , to obtain a torsion-free finite index subgroup G' that splits as a graph of groups with relatively malnormal quasiconvex edge groups. The main difficulty is to show that G' can be co-compactly cubulated so $G' = \pi_1 X$ where X is a nonpositively curved cube complex, and we will describe how to do that below. Separability allows us to pass to a finite cover $\widehat{X} \rightarrow X$ where the cube complex \widehat{X} has a hierarchy. Theorem 15.3 then implies that \widehat{X} is virtually special.

An *expanded edge group* E^+ in a graph of groups is the group obtained by combining an edge group E together with the parabolic subgroups that have




FIGURE 15.2. A graph of spaces on the left (the parabolics are highlighted immersed tori), an expanded edge space in the middle corresponds to E^+ , and its associated cage on the right corresponds to K_E .

infinite intersection with it. Using separability, we can assume we passed to a finite index subgroup so that each expanded edge group E^+ is π_1 -injective, quasiconvex, and malnormal, and so we assume that G' already has this property. The cage K_E is the graph of groups obtained by “removing” the original edge group E from the expanded edge group E^+ . We refer to Figure 15.2. The group K_E naturally splits as a graph of groups with two vertex groups that are copies of E , and edge groups corresponding to the various ways a parabolic subgroup intersects E .

We will obtain the cubulation by an induction on the number of edge groups in the induced splitting of G' . When E is a separating edge, we cocompactly cubulate G' by applying Theorem 8.5 to the splitting as an amalgamated product along E . Note that the resulting vertex groups had splittings of the same form as G' but with fewer edges. For a nonseparating edge E , consideration of the mapping cylinder associated to $E^+ \hookrightarrow G'$ shows that G' has another splitting as $L *_{K_E} E^+$ where L is the group obtained by deleting the edge group E from G' . We apply Theorem 8.5 to $L *_{K_E} E^+$ to obtain the cubulation of G' . To support this, note that L has fewer edges, and that E^+ is more easily cubulated – and also covered by a variant of Theorem 8.5. \square

We expect that with additional effort one can prove the following conjecture. The main hurdle will be to generalize the cubulation results described in [43]. Although known generalizations of Theorem 8.2 can handle a relatively hyperbolic situation, it currently requires a parabolicity of the edge groups in the vertex groups as in Theorem 8.5. One expects that this can be generalized to the less stringent requirement that the edge groups are full subgroups of the vertex groups.

CONJECTURE 15.7. *Let G be hyperbolic relative to virtually abelian groups. Suppose G has a quasiconvex hierarchy. Then G is virtually (compact) sparse special.*

CHAPTER 16

Applications

In this chapter we describe applications to three classes of relatively hyperbolic groups with quasiconvex hierarchies: One-relator groups with torsion, limit groups, and fundamental groups of hyperbolic 3-manifolds with a geometrically finite incompressible surface.

16.1. Baumslag’s Conjecture

A *one-relator group* is a group having a presentation $\langle a, b, \dots \mid W^n \rangle$ with a single defining relation. Assuming that W is cyclically reduced and is not a proper power, the one-relator group has torsion if and only if $n \geq 2$. In this case, all torsion is conjugate into $\langle W \rangle \cong \mathbb{Z}_n$ and the group is virtually torsion-free (since free groups are potent). We refer to [52] for more information on one-relator groups. A significant feature of one-relator groups with torsion is that they are hyperbolic, since the Newman Spelling Theorem provides very strong small-cancellation behavior. It became clear in the 60’s that one-relator groups with torsion are better behaved than general one-relator groups, and to test this Gilbert Baumslag posed the following:

CONJECTURE 16.1 ([5]). *Every one-relator group with torsion is residually finite.*

The main tool for studying one-relator groups is the *Magnus-Moldavanskii hierarchy* (M-M hierarchy). Roughly speaking, every one-relator group G is an HNN extension $H *_M t = M'$ of a simpler one-relator group H where M and M' are free subgroups generated by subsets of the generators of the presentation of H . The M-M hierarchy terminates at a virtually free group of the form $\mathbb{Z}_n * F$.

A *Magnus subgroup* M of the one-relator group $\langle a, b, \dots \mid W^n \rangle$ is a subgroup generated by a subset of the generators omitting at least one generator that occurs in W^n . The crucial point in the M-M hierarchy is the following result proven by Magnus (see [52, 38]).

THEOREM 16.2 (Freiheitsatz). *Every Magnus subgroup of a one-relator group is free. Moreover, its generators form a basis.*

Instead of explicitly describing the M-M hierarchy in general, we give the following example. The reader can also refer to [52].

EXAMPLE 16.3. Consider the one-relator group G_1 presented by:

$$\langle a, s \mid as^{-1}a^{-3}sas^{-1}a^2sa^{-1}s^{-2}asa^{-1}sas^{-1}a^2s^{-1}a^{-1}s^2 \rangle$$

Using the notation $x^y = y^{-1}xy$, we prepare the next step by rewriting the relator to suggest that s will be the stable letter in an HNN splitting, and have the following equivalent presentation:

$$\langle a, s \mid a(a^{-3})^s a(a^2)^s a^{-1} a^{s^2} (a^{-1})^s a(a^2)^s (a^{-1})^{s^2} \rangle$$

Treating s as the “stable letter” conjugating the Magnus subgroup $M_2 = \langle a, t \rangle$ to the Magnus subgroup $M'_2 = \langle t, u \rangle$ with $a^s = t, t^s = u$, the group G_1 splits as $G_2 *_{M_2^s = M'_2}$ where G_2 has the following presentation:

$$\langle a, t, u \mid at^{-3}at^2a^{-1}ut^{-1}at^2u^{-1} \rangle$$

Rewriting the relator to suggest that t will be a stable letter for the next splitting we have:

$$\langle a, t, u \mid aa^{t^3} (a^{-1})^t u^t a^{t^2} u^{-1} \rangle$$

G_2 splits as an HNN extension $G_1 *_{M_1^t = M'_1}$ where $M_1 = \langle a, b, c, u \rangle$ and $M'_1 = \langle b, c, d, v \rangle$ with $a^t = b, b^t = c, c^t = d, u^t = v$ and where G_1 is the following one-relator group, which is obviously free since no generator occurs more than once in the relator.

$$\langle a, b, c, d, u, v \mid adb^{-1}vcu^{-1} \rangle$$

A similar example with torsion arises by raising the relator to the n -th power in the above example. In general, following Moldavanskii, it is convenient to sometimes replace G by $G * \mathbb{Z}$ during the above construction as this enables the combinatorial group theory to easily deal with the case where no generator has exponent sum zero. We refer to [52] and [76].

Remarkably, computer studies done by several researchers, most notably Dunfield-Thurston [24], suggest that most (e.g. 94% !?) one-relator groups are of the form $F_r \rtimes \mathbb{Z}$ and in particular, the M-M hierarchy terminates quickly but is not a quasiconvex hierarchy.

For one-relator groups with torsion, the subgroups M, M' are quasiconvex at each level of the hierarchy [76]. This result is proven using a diagrammatic argument that depends upon a variant of the Newman spelling theorem [41, 51]. Instead of reproducing the argument, we give a short proof of the special case where $n \geq 6$ using small-cancellation theory. We note that $\langle a, b, \dots \mid W^n \rangle$ satisfies the $C'(\frac{1}{n})$ condition, and is hence a $C'(\frac{1}{6})$ small-cancellation group for $n \geq 6$. We can thus use small-cancellation theory to prove the quasiconvexity of the Magnus subgroups as well as Theorem 16.2 for high torsion:

THEOREM 16.4. *Let $G = \langle a_1, a_2, \dots \mid W^n \rangle$ be a one-relator group presentation with $n \geq 6$. Let $M = \langle a_i : i \in I \rangle$ be a Magnus subgroup. Then M is a quasiconvex subgroup of G that is freely generated by its generators.*

PROOF. Let X be the standard 2-complex of G . Let Y be a bouquet of circles corresponding to the generators of M . Let $Y \subset X$ be the based inclusion corresponding to $M \rightarrow G$.

Observe that $Y \rightarrow X$ has no missing shells. Consequently $\pi_1 Y \rightarrow \pi_1 X$ is injective and quasiconvex by Theorem 13.3 and Theorem 13.4. \square

When G is a one-relator group with torsion, and G' is a torsion-free finite index subgroup, the induced hierarchy for G' is a quasiconvex hierarchy that terminates at trivial groups (instead of finite groups) and is thus covered by the torsion-free case of Theorem 14.3.

THEOREM 16.5. *Every one-relator group with torsion is virtually special.*

Because of Proposition 2.12, a virtually special hyperbolic group has very strong properties, and in particular it is residually finite, so Conjecture 16.1 follows from Theorem 16.5.

16.2. 3-Manifolds

Prior to Thurston's work, the main tool used to study 3-manifolds is a *hierarchy* which is a sequence of splittings along incompressible surfaces until only 3-balls remain. An *incompressible surface* is a 2-sided π_1 -injective surface along which $\pi_1 M$ splits essentially as either an HNN extension or amalgamated free product.

It is well-known that every irreducible 3-manifold with an incompressible surface has a hierarchy and every irreducible 3-manifold with boundary has an incompressible surface. It is a deeper result that for a finite volume 3-manifold with cusps, there is always an incompressible geometrically finite surface [18]. In general, an incompressible surface in a hyperbolic 3-manifold is either geometrically finite or virtually corresponds to a fiber (see [10]). A fundamental result of Thurston's about subgroups of fundamental groups of infinite volume hyperbolic manifolds ensures that if the initial incompressible surface is geometrically finite, then the further incompressible surfaces in (any) hierarchy are geometrically finite (see the survey in [14]).¹ Finally, the geometrical finiteness of an incompressible surface where the 3-manifold splits corresponds precisely to the quasi-isometric embedding of the corresponding subgroup along which the fundamental group splits. Thus, if M has an incompressible surface then $\pi_1 M$ has a quasiconvex hierarchy and hence:

THEOREM 16.6. *If M is a hyperbolic 3-manifold with an incompressible geometrically finite surface then $\pi_1 M$ is virtually special.*

PROOF. In view of the above discussion, this follows from Theorem 14.3 when $\pi_1 M$ is hyperbolic. However, when M has cusps and is thus only relatively hyperbolic, we require Lemma 16.7 and then apply Corollary 15.6. \square

¹Most of this discussion is now subsumed into the dichotomy between virtual fiberness and geometrical finiteness for finitely generated subgroups – a consequence of the Tameness Theorem of Agol, Calegari-Gabai

LEMMA 16.7. *Let M be a compact irreducible atoroidal 3-manifold. There exists a finite cover $\widehat{M} \rightarrow M$ and a surface S that is a disjoint union $S = S_1 \sqcup \cdots \sqcup S_k$ of incompressible surfaces such that the fundamental group of each component $\widehat{M} - S$ is hyperbolic.*

COROLLARY 16.8. *If M is a hyperbolic 3-manifold with an incompressible geometrically finite surface then $\pi_1 M$ is subgroup separable.*

In the 80's Thurston suggested that perhaps every hyperbolic 3-manifold is virtually fibered. The key to proving the virtual fibering is the following beautiful result which weaves together several important ideas from 3-manifold topology [2]:

PROPOSITION 16.9 (Agol's fibering criterion). *Let M be a compact 3-manifold, and suppose that $\pi_1 M$ is RFRS. Then M has a finite cover that fibers.*

For a Haken hyperbolic 3-manifold M , either it virtually fibers, or the first incompressible surface is geometrically finite. Virtual specialness then implies that M has a finite cover with $\pi_1 \widehat{M}$ contained in a raag which is RFRS and so:

COROLLARY 16.10. *Every hyperbolic Haken 3-manifold is virtually fibered.*

16.3. Limit Groups

Fully residually free groups or *limit groups* have been a recent focal point of geometric group theory. These are groups G with the property that for every finite set g_1, \dots, g_k of nontrivial elements, there is a free quotient $G \rightarrow \bar{G}$ such that $\bar{g}_1, \dots, \bar{g}_k$ are nontrivial. Among the many remarkable properties proved for these groups is that they have a rather simple *cyclic hierarchy* terminating at free groups and free abelian groups. Each splitting in the cyclic hierarchy has one of the following forms:

- (1) $A *_Z B$ where Z is cyclic and malnormal in A ;
- (2) $A *_Z B$ where Z is cyclic and malnormal in A and Z, Z' do not have nontrivially intersecting conjugates.
- (3) $A *_Z B$ where Z is cyclic and malnormal in A and $B \cong Z \times \mathbb{Z}^n$ for some n .

This hierarchy was obtained in [50], and is also implicit in Sela's retractive tower description of limit groups [69]. This hierarchy allows one to prove that limit groups are hyperbolic relative to free abelian subgroups [19, 3]. Using the relative hyperbolicity, we can follow the cyclic hierarchy and repeatedly apply Theorem 15.5 and Corollary 8.6 to obtain:

COROLLARY 16.11. *Every limit group is virtually compact special.*

Combined with subgroup separability results for virtually special groups that are hyperbolic relative to abelian subgroups one recovers Wilton's result that limit groups are subgroup separable [73].

Bibliography

1. Ian Agol, *Tameness of hyperbolic 3-manifolds*, 2004, Preprint.
2. ———, *Criteria for virtual fibering*, J. Topol. **1** (2008), no. 2, 269–284. MR MR2399130 (2009b:57033)
3. Emina Alibegović, *A combination theorem for relatively hyperbolic groups*, Bull. London Math. Soc. **37** (2005), no. 3, 459–466. MR MR2131400
4. G. N. Arzhantseva, *On quasiconvex subgroups of word hyperbolic groups*, Geom. Dedicata **87** (2001), no. 1-3, 191–208. MR MR1866849 (2003h:20076)
5. Gilbert Baumslag, *Residually finite one-relator groups*, Bull. Amer. Math. Soc. **73** (1967), 618–620.
6. Nicolas Bergeron, Frédéric Haglund, and Daniel T. Wise, *Hyperplane sections in arithmetic hyperbolic manifolds*, J. Lond. Math. Soc. (2) **83** (2011), no. 2, 431–448. MR 2776645
7. Nicolas Bergeron and Daniel T. Wise, *A boundary criterion for cubulation*, American Journal of Mathematics, To Appear.
8. M. Bestvina and M. Feighn, *A combination theorem for negatively curved groups*, J. Differential Geom. **35** (1992), no. 1, 85–101.
9. Hadi Bigdely and Daniel T. Wise, *Quasiconvexity and relatively hyperbolic groups that split*, 16 pp. Preprint 2011.
10. Francis Bonahon, *Bouts des variétés hyperboliques de dimension 3*, Ann. of Math. (2) **124** (1986), no. 1, 71–158.
11. Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Springer-Verlag, Berlin, 1999. MR 1 744 486
12. Marc Burger and Shahar Mozes, *Finitely presented simple groups and products of trees*, C. R. Acad. Sci. Paris Sér. I Math. **324** (1997), no. 7, 747–752. MR 98g:20041
13. Danny Calegari and David Gabai, *Shrinkwrapping and the taming of hyperbolic 3-manifolds*, J. Amer. Math. Soc. **19** (2006), no. 2, 385–446 (electronic). MR MR2188131 (2006g:57030)
14. Richard D. Canary, *Covering theorems for hyperbolic 3-manifolds*, Low-dimensional topology (Knoxville, TN, 1992), Conf. Proc. Lecture Notes Geom. Topology, III, Internat. Press, Cambridge, MA, 1994, pp. 21–30.
15. Ruth Charney, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158. MR MR2322545 (2008f:20076)
16. Ruth Charney and Michael W. Davis, *Finite $K(\pi, 1)$ s for Artin groups*, Prospects in topology (Princeton, NJ, 1994), Ann. of Math. Stud., vol. 138, Princeton Univ. Press, Princeton, NJ, 1995, pp. 110–124. MR MR1368655 (97a:57001)
17. Indira Chatterji and Graham Niblo, *From wall spaces to CAT(0) cube complexes*, Internat. J. Algebra Comput. **15** (2005), no. 5-6, 875–885. MR MR2197811 (2006m:20064)
18. M. Culler and P. B. Shalen, *Bounded, separating, incompressible surfaces in knot manifolds*, Invent. Math. **75** (1984), no. 3, 537–545.
19. François Dahmani, *Combination of convergence groups*, Geom. Topol. **7** (2003), 933–963 (electronic). MR MR2026551 (2005g:20063)

20. Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235.
21. N. G. de Bruijn, *Algebraic theory of Penrose’s nonperiodic tilings of the plane. I, II*, Nederl. Akad. Wetensch. Indag. Math. **43** (1981), no. 1, 39–52, 53–66. MR 609465 (82e:05055)
22. Carl Droms, *Graph groups*, Ph.D. thesis, Syracuse University, 1983.
23. Carl Droms, Jacques Lewin, and Herman Servatius, *Tree groups and the 4-string pure braid group*, J. Pure Appl. Algebra **70** (1991), no. 3, 251–261. MR MR1104843 (92b:20040)
24. Nathan M. Dunfield and Dylan P. Thurston, *A random tunnel number one 3-manifold does not fiber over the circle*, Geom. Topol. **10** (2006), 2431–2499. MR 2284062 (2007k:57033)
25. Daniel S. Farley, *Finiteness and CAT(0) properties of diagram groups*, Topology **42** (2003), no. 5, 1065–1082. MR 2004b:20057
26. Rita Gitik, *On the profinite topology on negatively curved groups*, J. Algebra **219** (1999), no. 1, 80–86. MR MR1707664 (2000g:20072)
27. Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev, *Widths of subgroups*, Trans. Amer. Math. Soc. **350** (1998), no. 1, 321–329.
28. M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263.
29. ———, *CAT(κ)-spaces: construction and concentration*, Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI) **280** (2001), no. Geom. i Topol. 7, 100–140, 299–300. MR 1879258 (2002j:53045)
30. Daniel Groves and Jason Fox Manning, *Dehn filling in relatively hyperbolic groups*, Israel J. Math. **168** (2008), 317–429. MR MR2448064
31. V. S. Guba, *The word problem for the relatively free semigroup satisfying $T^m = T^{m+n}$ with $m \geq 3$* , Internat. J. Algebra Comput. **3** (1993), no. 3, 335–347. MR 94m:20116
32. M.F. Hagen, *Cocompactly cubulated crystallographic groups*, (2012).
33. Frédéric Haglund, *Finite index subgroups of graph products*, Geom. Dedicata **135** (2008), 167–209. MR MR2413337
34. Frédéric Haglund and Frédéric Paulin, *Simplicité de groupes d’automorphismes d’espaces à courbure négative*, The Epstein birthday schrift, Geom. Topol., Coventry, 1998, pp. 181–248 (electronic). MR 2000b:20034
35. Frédéric Haglund and Daniel T. Wise, *A combination theorem for special cube complexes*, Ann. of Math., To appear.
36. ———, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1 551–1620. MR MR2377497
37. ———, *Coxeter groups are virtually special*, Adv. Math. **224** (2010), no. 5, 1890–1903. MR 2646113
38. James Howie, *How to generalize one-relator group theory*, Combinatorial group theory and topology (Princeton, N.J.) (S. M. Gersten and John R. Stallings, eds.), Princeton Univ. Press, 1987, pp. 53–78.
39. G. Christopher Hruska, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. **10** (2010), no. 3, 1807–1856. MR 2684983
40. G. Christopher Hruska and Daniel T. Wise, *Finiteness properties of cubulated groups*, 1–52, Preprint 2010.
41. ———, *Towers, ladders and the B. B. Newman spelling theorem*, J. Aust. Math. Soc. **71** (2001), no. 1, 53–69. MR 1 840 493
42. ———, *Packing subgroups in relatively hyperbolic groups*, Geom. Topol. **13** (2009), no. 4, 1945–1988. MR MR2497315 (2010c:20053)
43. Tim Hsu and Daniel T. Wise, *Cubulating malnormal amalgams*, 1–20, Preprint.

44. ———, *On linear and residual properties of graph products*, Michigan Math. J. **46** (1999), no. 2, 251–259. MR 1 704 150
45. ———, *Separating quasiconvex subgroups of right-angled Artin groups*, Math. Z. **240** (2002), no. 3, 521–548. MR 1 924 020
46. S. P. Humphries, *On representations of Artin groups and the Tits conjecture*, J. Algebra **169** (1994), 847–862.
47. David Janzen and Daniel T. Wise, *Cubulating rhombus groups*, Groups Geom. Dyn., 1–19, To appear.
48. Jeremy Kahn and Vladimir Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Preprint, 2009.
49. Ilya Kapovich and Hamish Short, *Greenberg’s theorem for quasiconvex subgroups of word hyperbolic groups*, Canad. J. Math. **48** (1996), no. 6, 1224–1244. MR MR1426902 (98g:20060)
50. O. Kharlampovich and A. Myasnikov, *Hyperbolic groups and free constructions*, Trans. Amer. Math. Soc. **350** (1998), no. 2, 571–613.
51. Joseph Lauer, *Cubulating one relator groups with torsion*, Master’s thesis, McGill University, 2007.
52. Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, Springer-Verlag, Berlin, 1977, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.
53. Eduardo Martínez-Pedroza and Daniel T. Wise, *Relative quasiconvexity using fine hyperbolic graphs*, Algebr. Geom. Topol. **11** (2011), no. 1, 477–501. MR 2783235
54. J. P. McCammond and D. T. Wise, *Coherence, local quasiconvexity, and the perimeter of 2-complexes*, Geom. Funct. Anal. **15** (2005), no. 4, 859–927. MR MR2221153
55. Jonathan P. McCammond and Daniel T. Wise, *Fans and ladders in small cancellation theory*, Proc. London Math. Soc. (3) **84** (2002), no. 3, 599–644. MR 1 888 425
56. Ashot Minasyan, *Separable subsets of gferf negatively curved groups*, Preprint, 2004.
57. Mahan Mitra, *Height in splittings of hyperbolic groups*, Proc. Indian Acad. Sci. Math. Sci. **114** (2004), no. 1, 39–54. MR 2040599 (2005d:20073)
58. Lee Mosher, *Geometry of cubulated 3-manifolds*, Topology **34** (1995), no. 4, 789–814.
59. G. A. Niblo and L. D. Reeves, *Coxeter groups act on CAT(0) cube complexes*, J. Group Theory **6** (2003), no. 3, 399–413. MR 2004e:20072
60. Bogdan Nica, *Cubulating spaces with walls*, Algebr. Geom. Topol. **4** (2004), 297–309 (electronic). MR 2 059 193
61. Denis V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Mem. Amer. Math. Soc. **179** (2006), no. 843, vi+100. MR MR2182268 (2006i:20047)
62. ———, *Peripheral fillings of relatively hyperbolic groups*, Invent. Math. **167** (2007), no. 2, 295–326. MR MR2270456 (2008d:20080)
63. P. Papasoglu, *Strongly geodesically automatic groups are hyperbolic*, Invent. Math. **121** (1995), no. 2, 323–334. MR 1346209 (96h:20073)
64. Eliyahu Rips and Yoav Segev, *Torsion-free group without unique product property*, J. Algebra **108** (1987), no. 1, 116–126. MR MR887195 (88g:20071)
65. Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) **71** (1995), no. 3, 585–617. MR 97a:20062
66. ———, *Codimension-1 subgroups and splittings of groups*, J. Algebra **189** (1997), no. 2, 377–389. MR 98c:20071
67. Michah Sageev and Daniel T. Wise, *Cores for quasiconvex actions*.
68. Peter Scott, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2) **17** (1978), no. 3, 555–565.
69. Z. Sela, *Diophantine geometry over groups. II. Completions, closures and formal solutions*, Israel J. Math. **134** (2003), 173–254. MR MR1972179 (2004g:20061)

70. Hamish Short, *Quasiconvexity and a theorem of Howson's*, Group theory from a geometrical viewpoint (Trieste, 1990) (É. Ghys, A. Haefliger, and A. Verjovsky, eds.), World Sci. Publishing, River Edge, NJ, 1991, pp. 168–176.
71. John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565.
72. C. M. Weinbaum, *The word and conjugacy problems for the knot group of any tame, prime, alternating knot*, Proc. Amer. Math. Soc. **30** (1971), 22–26.
73. Henry Wilton, *Hall's theorem for limit groups*, Geom. Funct. Anal. **18** (2008), no. 1, 271–303. MR MR2399104 (2009d:20101)
74. Daniel T. Wise, *Cubular tubular groups*, 1–14, Preprint 2011.
75. ———, *Recubulating free groups*, Israel J. Math., 1–9, To appear.
76. ———, *The structure of groups with a quasiconvex hierarchy*, 1–189.
77. ———, *Virtual cleanliness*, 1–21.
78. ———, *Subgroup separability of graphs of free groups with cyclic edge groups*, Q. J. Math. **51** (2000), no. 1, 107–129. MR 1 760 573
79. ———, *The residual finiteness of negatively curved polygons of finite groups*, Invent. Math. **149** (2002), no. 3, 579–617.
80. ———, *Cubulating small cancellation groups*, GAFA, Geom. Funct. Anal. **14** (2004), no. 1, 150–214.
81. ———, *Subgroup separability of the figure 8 knot group*, Topology **45** (2006), no. 3, 421–463. MR MR2218750
82. ———, *Complete square complexes*, Comment. Math. Helv. **82** (2007), no. 4, 683–724. MR MR2341837 (2009c:20078)
83. Wen-Yuan Yang, Proc. Amer. Math. Soc., To appear.

Index of notation and defined terms

- 2-sided, 32
- $A \otimes_R B$, 35
- $B(6)$, 63
- $C'(\frac{1}{n})$, 75, 83
- $C(p)$, 83
- H -wall, 69
- $N(H)$, 14
- $WProj_X(B \rightarrow A)$, 41
- $\prec(c)$, 79
- $\partial_p D$, 15
- $\mathbb{C}_G(H)$, 27
- $Hull(HK)$, 29
- \mathcal{MQH} , 101
- \mathcal{VH} -complex, 11
- $\langle\langle H_1, H_2, \dots \rangle\rangle$, 76
- $\mathcal{N}_\mu(H)$, 25
- $|Y_i|$, 84
- i -shell, 80
- n -cube, 7
- $\#(p,q)$, 14
- $\mathbb{C}(Y \rightarrow X)$, 36
- $C(p)$ small-cancellation condition, 80
- Absorbing squares, 83
- acts, 54
- admit, 86
- Admitted rectangles, 87
- almost malnormal, 25
- almost malnormal quasiconvex
 - hierarchy, 101
- alternating, 8
- angle, 79
- angled 2-complex, 79
- annuladder, 98
- annular diagram, 96
- annular diagram in X^* , 96
- axis, 66
- $B(6)$, 63
- betwixts, 53
- bicollared, 98
- boundary cycle, 15
- boundary path, 15
- bounded packing, 61
- cage, 121
- cancelable pair, 82
- canceled pair, 17
- canonical 0-cube, 64
- canonical completion, 36
- canonical cube, 64
- canonical retraction, 36
- carrier, 14, 58, 94
- CAT(0) cube complex, 8
- cladder, 95
- class, 98
- clean, 11
- codimension-1, 58
- Combining cone-cells, 83
- commensurable, 27
- commensurator, 27
- complete square complex, 11
- cone-cell, 77
- cone-piece, 88
- consecutive, 94
- convex, 14
- convex wallspace, 92
- core, 28
- cornsquare, 21
- cosparsely, 64
- cross, 60
- crossing graph, 33
- cube complex, 7
- cubical hierarchy, 119
- cubical presentation, 76
- curvature at a 2-cell, 79
- cut-wall, 66
- cutting along a hyperplane, 119
- cyclic hierarchy, 125

- deep, 58
- degenerate, 87
- Dehn complex, 9
- diagram in a complex X , 15
- disk diagram, 15
- disk diagram D in X^* , 77
- distinct conjugates, 25
- doubly-external, 81
- dual, 13, 53
- dual CAT(0) cube complex, 3
- dual curves, 16

- elevation, 41
- expanded edge group, 120
- extends, 69
- external, 81

- faces, 7
- fiber-product, 35
- flag complex, 8
- full, 30, 111
- full subcomplex, 23
- Fully residually free groups, 125

- geodesic, 14
- graph group, 12

- halfspace carriers, 24
- halfspaces, 53
- height, 26
- hexagon move, 17
- hierarchy, 3, 124
- honeycomb, 63
- horizontal, 11
- hyperplane, 13

- immersed hyperplane, 31
- immersed walls, 70
- in parallel, 88
- incompressible surfaces, 124
- innerpath, 80
- internal, 81
- internal part, 87
- interosculate, 32
- intersecting conjugator, 26, 115

- length, 114
- limit groups, 125
- linear separation, 64
- link, 8
- local isometry, 23

- Magnus subgroup, 122
- Magnus-Moldavanskii hierarchy, 122
- malnormal, 25
- malnormal quasiconvex hierarchy, 101
- map of cubical presentations, 112
- metric small-cancellation, 75
- midcells, 57
- midcube, 13
- missing corners of squares, 23
- monkey-tails, 89
- multiply-external, 81

- no missing shells, 111, 113
- nonpositively curved, 8
- nontrivial, 80

- one-relator group, 122
- open cubical neighborhood, 44
- open halfspaces, 53
- orientation, 53
- outerpath, 21, 80

- parallel, 40
- piece, 75, 80
- pieces, 84, 88
- prime, 9
- profinite topology, 34
- properly, 64

- quasiconvex, 25
- quasiconvex hierarchy, 114
- quasiflats, 64

- raag, 12
- rectangles, 84
- reduced, 82, 84, 96
- relatively malnormal, 26
- removed, 83
- Removing bigons, 84
- Replacing cone-cells, 83
- residually finite, 34
- right-angled Artin group, 12

- self-oscultates, 32
- separable, 34
- separating, 53
- shards, 87
- shell, 78
- singly-external, 81
- size, 84
- sparse, 64
- special, 31
- spur, 78, 80
- strongly convex, 99
- superconvex, 29

- thick, 98
- thickened carrier, 94

thin, 98
track, 56
triply-external, 81
trivial, 41, 80

vertical, 11

wall, 53, 93
wall partition, 53
wall-injective, 39
wall-piece, 88
walls, 92
wallspace, 53
worms, 57