FROM RICHES TO RAAGS: 3-MANIFOLDS, RIGHT-ANGLED ARTIN GROUPS, AND CUBICAL GEOMETRY

DANIEL T. WISE

ABSTRACT. These notes aim to provide a bit more detail and background for some of the lectures given at the NSF-CBMS conference August 1-5, 2011. It is rough now, but will eventually be expanded upon to provide a humane account. I would be grateful for any corrections, suggestions, and additions.

CONTENTS

1. Overview 3
1.1. Applications: 5
1.2. A scheme for understanding groups 5
2. Nonpositively Curved Cube Complexes 7
2.1. Definitions 7
2.2. Some favorite 2-dimensional examples 7
2.3. Graph Groups 10
2.4. Hyperplanes 11
3. Disk diagrams over cube complexes, and applications to hyperplanes, and convexity 11
3.1. Disk Diagrams 11
3.2. Properties of Hyperplanes 13
3.3. Local isometries and convexity 16
3.4. Cores, Hulls, and Superconvexity 17
4. Special Cube Complexes 20
4.1. Hyperplane definition of special cube complex 20
4.2. Separability Criteria for Virtual Specialness 21
4.3. Canonical Completion and Retraction 22
4.4. Separability in the hyperbolic case 23
4.5. Wall-injectivity and a fundamental commutative diagram 24
4.6. Wall Projection Controls Retraction 26
5. Virtual Specialness of Malnormal Amalgams 27
5.1. Proof of Lemma 5.3 30
6. Wallspaces and their Dual Cube Complexes 33
6.1. Finiteness properties of the action on the dual cube complex 35
6.2. Properness of the G action on $C(\widetilde{X})$ 38
7. Cubulating Malnormal Graphs of Cubulated Groups 41
7.1. Examples of systems of walls in graphs of groups 41
7.2. Extending Walls 41
7.3. Constructing Turns 42
7.4. Cubulating malnormal amalgams 43
8. Cubical Small Cancellation Theory 46
8.1. Cubical presentations 46

Date: August 4, 2011.
2000 *Mathematics Subject Classification.* 53C23, 20F36, 20F55, 20F67, 20F65, 20E26.
Key words and phrases. CAT(0) Cube Complex, Right-angled Artin Group, Subgroup Separable.
Research supported by NSERC.

8.2. The Fundamental theorem of small-cancellation theory 47
8.3. Combinatorial Gauss-Bonnet Theorem 48
8.4. Greendlinger's Lemma and the ladder theorem for classical $C(6)$ small-cancellation theory 48
8.5. Reduced diagrams 50
8.6. Pieces and Reduced diagrams 52
8.7. Producing Examples 52
8.8. Rectified Diagrams 53
9. Walls in Cubical Small-Cancellation Theory 58
9.1. Walls in classical $C'(\frac{1}{6})$ small-cancellation 58
9.2. Wallspace Cones 58
9.3. Producing Wallspace Cones 58
9.4. Walls in $\widetilde{X}^*$ 59
9.5. Quasiconvexity of walls in $\widetilde{X}^*$ 59
10. Annular Diagrams 61
10.1. Classification of Flat Annuli 61
10.2. The Doubly Collared Annulus Theorem 62
11. The Malnormal Special Quotient Theorem 63
11.1. Case study: $F_2/\langle\langle W_1^{n_1}, \dots, W_r^{n_r}\rangle\rangle$ 64
11.2. Almost Malnormality 67
12. Cubical version of Filling Theorem 69
12.1. Persistence of quasiconvexity 70
12.2. No Missing Shells and Quasiconvexity 70
13. The structure of groups with a quasiconvex hierarchy 72
14. The relatively hyperbolic setting 74
15. Applications 76
15.1. Baumslag's conjecture 76
15.2. 3-manifolds 76
15.3. Limit groups 77
References 78

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 (apparently) merely as a convenient source of examples [Gro87], 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 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 disc diagrams are used to understand their hyperplanes and convex subcomplexes. While many of the essential properties of CAT(0) cube complexes can be explained using the CAT(0) triangle comparison metric, we have not adopted this viewpoint. It seems that the most important characteristic properties of 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. CAT(0) cube complexes are high-dimensional 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 [HW08] 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 $\widetilde{X}$, the stabilizers of hyperplanes are 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 $H_1, \ldots, H_r$ 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 $\widetilde{X}$. 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 [HWb]:

Theorem 1.3 (Cubulating Amalgams)**.** *Let G be a hyperbolic group that splits as $A *_C B$ or $A*_{C^t=D}$ where C is malnormal and quasiconvex. Suppose A, B are 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 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 $A*_{C^t=D}$ whose vertex groups have shorter length hierarchies. The hierarchy is *quasiconvex* if 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 $A*_{C^t=D}$.

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

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 $\langle a, b, \dots \mid W_1, W_2, \dots, W_r\rangle$ is $C'(\frac{1}{n})$ if for any "piece" P (i.e. a word 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 of the presentation is hyperbolic and disk diagram methods provide a very explicit understanding of many properties of the group G.

The presentation above can be reinterpreted as $\langle X \mid Y_1, Y_2, \dots, Y_r\rangle$ where X is a bouquet of loops and each $Y_i \to X$ is an immersed circle corresponding to W_i, and the group G of the presentation is $\pi_1 X/\langle\langle Y_i\rangle\rangle$. We generalize this to a setting where X is a nonpositively curved cube complex and each $Y_i \to X$ 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.*

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 (sufficiently small) 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 $Y_i \to X$ 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 AFP's and HNN's that arise at each stage of the hierarchy. Given a splitting, say $G = A *_{C^t = D}$, 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 $\mathrm{Height}(G, C)$ – this idea of repeatedly filling with an induction on height was independently discovered by Agol-Groves-Manning.

Theorem 1.7 holds in many (and conjecturally all) cases when G is hyperbolic relative to abelian subgroups. We describe how to deduce it from Theorem 1.7 – proof of separability essentially involves a generalization of Theorem 1.7 to provide virtually special peripheral fillings. However cubulating requires some additional work since the relatively hyperbolic version of Theorem 1.3 requires some aparabolicity of the edge group in the splitting in addition to the available relative quasiconvexity.

1.1. **Applications:** We describe three notable classes of groups with quasiconvex hierarchies

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

Every one-relator group has a Magnus hierarchy - and for one relator groups with torsion the Magnus hierarchy is a quasiconvex hierarchy. (Though one must use a torsion-free finite index subgroup). This resolves Baumslag's conjecture that every one-relator group with torsion is residaully 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, since 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 [Ago04, CG06]. A second corollary is that when the hyperbolic manifold M is Haken, in the sense that it has any incompressible surface S then it 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 [Ago08].

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 $\widetilde{C}$ upon which G acts.
- Verify that G acts properly and relatively cocompactly on $\widetilde{C}$ by examining the extrinsic nature of the codimension-1 subgroups.

- Consequently G is the fundamental group of a nonpositively curved cube complex. $C = G\backslash\widetilde{C}$ (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.

2. Nonpositively Curved Cube Complexes

2.1. **Definitions.** An *n-cube* is a copy of $[-1,1]^n$. Its *faces* are restrictions of some coordinates to ± 1. Note that faces are regarded as cubes.

FIGURE 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. 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 3. Two links within cube complexes

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 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 $v \in X^0$. $\widetilde{X}$ 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.

2.2. **Some favorite 2-dimensional examples.**

Example 2.2 (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. (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_b to v_t and the 1-cells associated with white regions are oriented from v_t to v_b. 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).

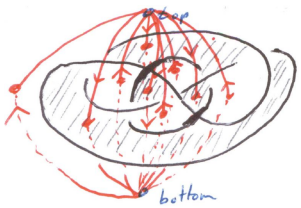

FIGURE 5. An embedded Dehn complex

FIGURE 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 [Wei71, LS77]. The explanation we give here is from [Wis06].

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

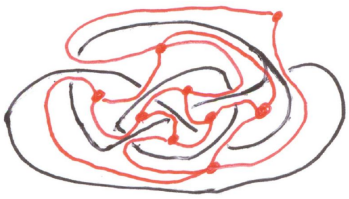

FIGURE 7. link(v_t) embeds in the projection diagram

To see that X is nonpositively curved exactly when P is prime and alternating, we note that the checkerboard coloring shows that link(v_t) is bipartite so it suffices to verify that there are no 2-cycles. But the two different types of 2-cycles would show that P is not alternating or not prime.

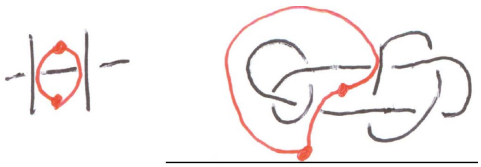

FIGURE 8. 2-cycles that indicate that P is not alternating or not prime

Example 2.3 (From graphs of graphs)**.** Let X decompose 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 $\phi_{e-} : X_e \times \{-1\} \to X_{\iota(e)}$ and $\phi_{e+} : X_e \times \{+1\} \to X_{\tau(e)}$ are combinatorial immersions. Then X is a nonpositively curved square complex.

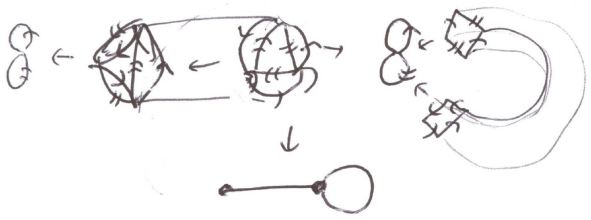

FIGURE 9. A graph of spaces with two vertex spaces and two edge spaces.

Example 2.4 (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 examples was one of the main motivations for this research.)

FIGURE 10. Subdivide the left and right graphs so that the cylinder edge space has bounding circles of the same length.

When X arises as a graph of spaces, then $\text{link}(v)$ is a bipartite graph, and X has the structure of a $\mathcal{VH}$*-complex*, which means that the 1-cells are divided into two classes: *vertical* and *horizontal* and attaching maps of 2-cells are length 4 paths that alternate between vertical and horizontal edges.

Let us 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 2.5 (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 $\widetilde{X}$ is isomorphic to the cartesian product of two trees. These can be surprisingly complicated and we refer to [Wis07], and to [JW] for some simple such examples. Most notably, Burger-Mozes gave such examples where $\pi_1 X$ is infinite simple [BM97].

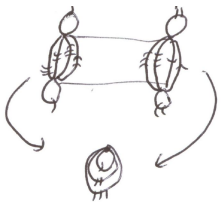

FIGURE 11. A complete square complex with 6 squares.

Example 2.6 (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.

Remark 2.7. When X is compact, having a clean finite cover is equivalent to the separability of the edge groups [Wis06].

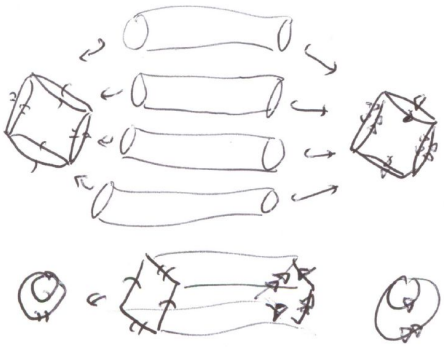

FIGURE 12. A $\mathcal{VH}$-complex for a genus 2 surface, together with a degree 4 clean cover

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

$$\langle\ g_v : v \in \text{Vertices}(\Gamma)\ \ \big|\ \ [g_u, g_v] : (u,v) \in \text{Edges}(\Gamma)\ \rangle \tag{1}$$

Example 2.8.

(1) $G($  $) \cong \mathbb{Z}^3$
(2) $G($  $) \cong F_3$
(3) $G($  $) \cong F_2 \times F_3$
(4) $G($  $) \cong \pi_1(S^3 -$ $)$

When girth$(\Gamma) \geq 4$, the standard 2-complex R of Presentation (1) is already a nonpositively curved square complex. In general, we must also add higher dimensional cubes, and define $R(\Gamma)$ to be the cube complex obtained by adding an n-cube for each collection of n pairwise commuting generators - i.e. for each n-clique in Γ. Then $R(\Gamma)$ is a nonpositively curved cube complex such that $\pi_1 R(\gamma) \cong G(\Gamma)$.

Example 2.9. A worthwhile 3-dimensional example to think about is indicated in Figure 13. Notice that $R(\Gamma)$ has only one 0-cell v, and that link(v) contains two copies of Γ: one "ascending" and one "descending" and has additional simplices corresponding to corners of cubes that are mixed.

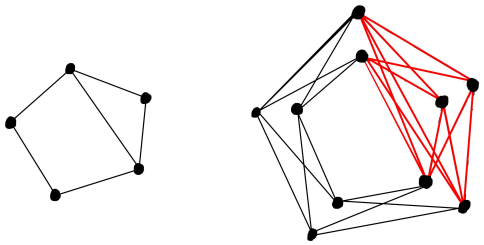

FIGURE 13. Γ is on the left, and link(v) is on the right. Note that link(v) contains an octahedron whose 8 2-simplices correspond to the 8 corners of the added 3-cube.

Proposition 2.10. *Graph groups have the following properties:*

(1) *They are residually torsion-free nilpotent* [Dro83].
(2) *They are residually finite rational solvable* [Ago08].
(3) *They are linear* [Hum94].

(4) *They embed in right-angled Coxeter groups and hence in* $SL_n(\mathbb{Z})$ [HW99, DJ00].

We refer to Charney's survey paper for more about right-angled and other Artin groups [Cha07].

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

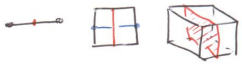

FIGURE 14. Midcubes in a 1-cube, 2-cube, and 3-cube

cube in a single midcube or in $\emptyset$.

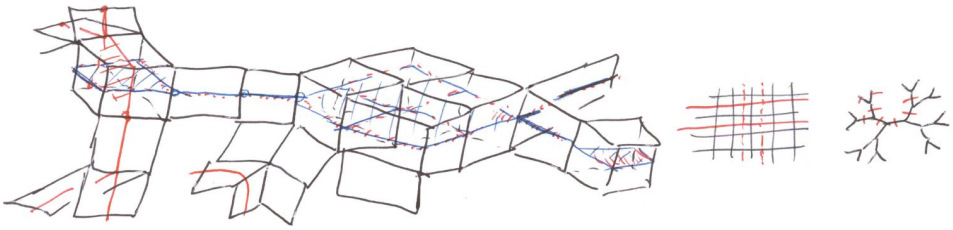

FIGURE 15. Some hyperplanes in a 3-dimensional CAT(0) cube complex, and in the plane, and in a tree.

We record some fundamental properties of hyperplanes. We will describe proofs in Section 3.

Theorem 2.11 (Sageev). (1) *Every midcube lies in a unique hyperplane of* $\widetilde{X}$.
(2) *A hyperplane* H *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* $\widetilde{X}$.
(4) $\widetilde{X} - H$ *consists of two components.*

Remark 2.12. We usually use the L^1 metric on a CAT(0) cube complex $\widetilde{X}$, where distance equals the length of the shortest path that is piecewise parallel to axes. $\widetilde{X}^1 \subset \widetilde{X}$ is then an isometric embedding, where we use the graph metric on $\widetilde{X}^1$. We note that $d(p,q) = \#(p,q)$ for $p, q \in \widetilde{X}^0$, where $\#(p,q)$ denotes the number of hyperplanes separating p and q.

3. Disk diagrams over cube complexes, and applications to hyperplanes, and convexity

3.1. **Disk Diagrams.** A *disc 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 $S^2 = \mathbb{R}^2 \cup \infty$).

FIGURE 16. Disk diagrams and their boundary paths

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

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

$$\begin{array}{ccc} \partial_p D & \to & D \\ \| & & \downarrow \\ P & \to & X \end{array}$$

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 17. Dual curves in a square disc diagram

When X is nonpositively curved and $D \to X$ has minimal area among all those disc 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 $D \to X$ be a disc 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 18.*

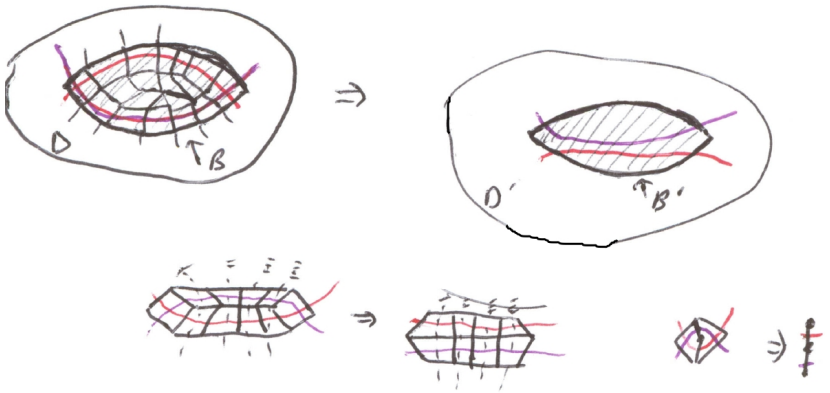

FIGURE 18. Bigonal subdiagrams can be replaced by subdiagrams with two fewer squares since $\mathsf{Area}(B') \leq \mathsf{Area}(B) - 2$.

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 19, it follows that:

(1) Each dual curve embeds - i.e. it cannot self-cross.
(2) A minimal area diagram D for $P \to X$ has no nonogons, oscugons, or bigons (see Figure 20).

FIGURE 19. A dual curve that passes through adjacent edges has another dual curve cutting through it to form a bigon.

FIGURE 20. A dual curve cannot self-cross as on the left in any diagram, moreover, minimal area diagrams cannot contain nonogons, oscugons, or bigons.

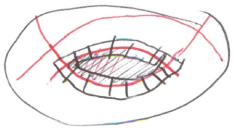

FIGURE 21.

Sketch of Theorem 3.2. We choose a smallest area bigon within D.

We then choose a smallest triangle of dual curves with one side on the bigon. There are two cases depending on whether it is not the case or it is the case that one of the vertices of this triangle is also a vertex of the bigon. These two possibilities are indicated in Figure 22. In each case we are able to homotope the diagram to a new diagram using a *hexagon move*

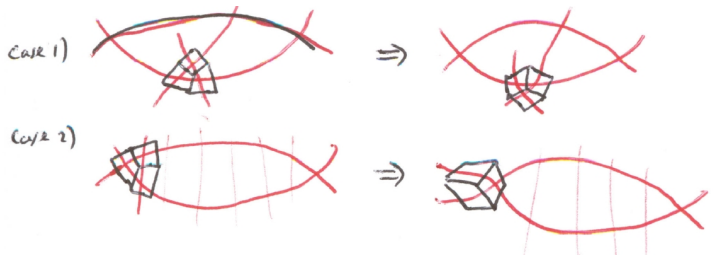

FIGURE 22.

illustrated in Figure 24, which pushes the front three squares in a cube to become the back three squares.

Repeating these hexagon moves we obtain a smaller and smaller area bigon until we arrive at the base case illustrated in Figure 24 when two squares are actually replaced by a length 2 path. This *cancelled pair* of squares plays a role similar to a "cancellable pair" in small-cancellation theory. □

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 $\widetilde{X}$. Then*

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

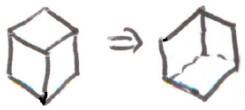

FIGURE 23. Hexagon move

FIGURE 24. Cancelled pair of squares

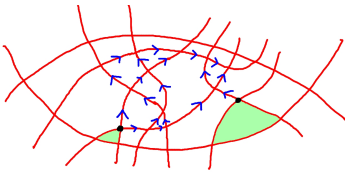

FIGURE 25. Consider the set of vertices in the graph of dual curves internal to a bigon. If there are no vertices then either there are no dual curves and we can cancel a pair of squares, or there are some dual curves and we can perform a leftward or rightward hexagon move at a corner. Otherwise, there 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.

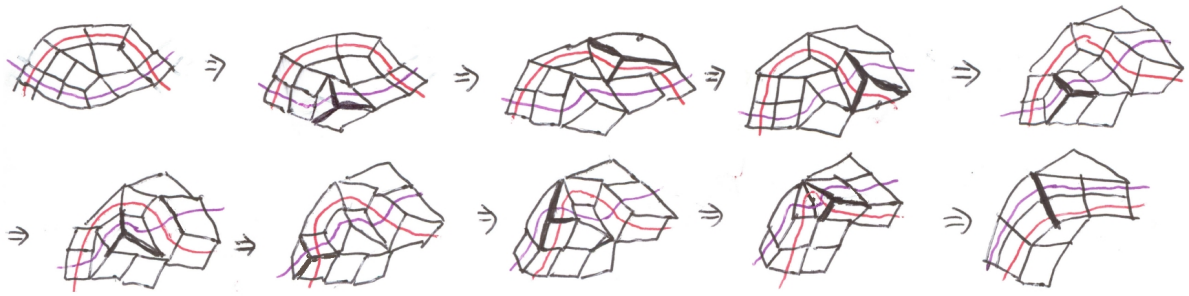

FIGURE 26. Complete example of bigon removal

Proof. 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. □

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

Proof. Suppose $\gamma \to \widetilde{X}$ is a path that starts and ends on opposite sides of H in $N(H)$. By passing to a proper subpath, we can assume γ never passes through H (but starts and ends on 1-cells *dual* to H in the sense that they H intersects them in a midcube).

Choose $\sigma \to N(H)$ not passing through H and e passing through H so that σe has the same endpoints as γ, and let D be a disc diagram for $\sigma e \gamma^{-1}$.

The dual curve starting on e must pass through γ since hyperplanes cannot self-cross (and hence this dual curve cannot end on σ).

This contradicts our assumption on γ. □

Lemma 3.5. *$\gamma \to \widetilde{X}$ is a geodesic $\Leftrightarrow$ the edges of γ are dual to distinct hyperplanes of $\widetilde{X}$.*

Proof. ($\Rightarrow$) holds since hyperplane carriers are convex. Indeed, if γ passed through two edges dual to the same hyperplane then we could shorten its length by 2.

FIGURE 27.

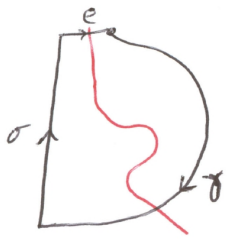

FIGURE 28. Hyperplanes are 2-sided

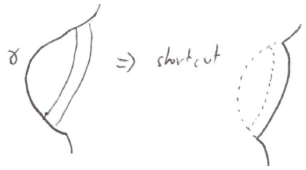

FIGURE 29. A geodesic γ cannot cross a hyperplane twice

($\Leftarrow$) 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 $\widetilde{X}$. □

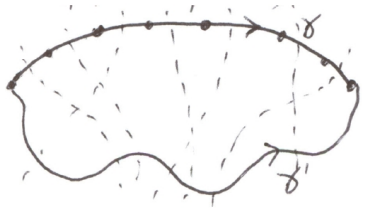

FIGURE 30. The hyperplanes dual to edges of a geodesic γ must be crossed by any other path γ' with the same endpoints.

Lemma 3.6 (No Inter-osculations: "crossing pair has a square")**.** *Let U, V be hyperplanes in $\widetilde{X}$. 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'.*

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 32, 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 32. We then use hexagon moves to push the square s leftwards to reveal the desired square s'. □

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

FIGURE 31. Crossing pair has a square. Hyperplanes cannot interosculate in $\widetilde{X}$

FIGURE 32. 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.

Remark 3.7 (Finding cornsquares)**.** Let D be a disc 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 (between a, b) such that a', b' cross in D.

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.

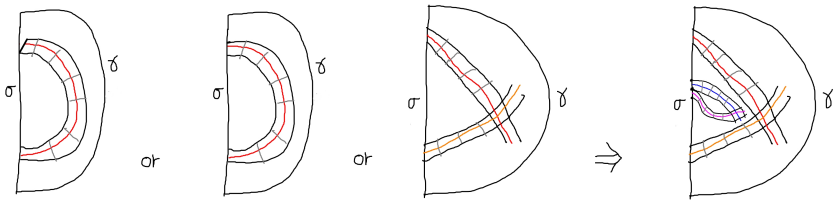

FIGURE 33. Finding a cornsquare

3.3. **Local isometries and convexity.**

Definition 3.8 (Local isometry)**.** A combinatorial map $\phi : Y \to X$ of nonpositively curved cube complexes is a *local isometry* if for each $y \in Y^0$, the map $\mathrm{link}(y) \to \mathrm{link}(\phi(y))$ is injective and moreover $\mathrm{link}(y) \subset \mathrm{link}(\phi(y))$ embeds as a *full subcomplex* in the sense that if u, v are vertices of $\mathrm{link}(y)$ that map to adjacent vertices of $\mathrm{link}(\phi(y))$ then u, v are already adjacent in $\mathrm{link}(y)$. (Note that the flag complex condition then implies that $(n+1)$ vertices of $\mathrm{link}(y)$ span an n-simplex iff their images in $\mathrm{link}(\phi(y))$ span an n-simplex.

A more concrete way to define a local isometry is that $Y \to X$ 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.)

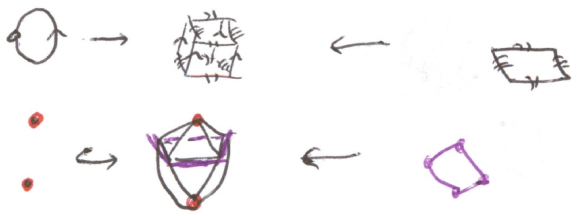

FIGURE 34. Local isometries of 1-torus and 2-torus into a 3-torus and the respective embeddings of links as full subcomplexes

Example 3.9. Any immersion of graphs is a local isometry. Any covering map is a local isometry. The fundamental example arises from the carrier $\bar{N} = N(\bar{H})$ of an immersed hyperplane $\bar{H}$ in a nonpositively curved cube complex X, where $\bar{N} \to X$ is a local isometry. In particular such immersed hyperplanes and their carriers are definable locally for a CAT(0) cube complex.

The following lemma was first noted in [Mos95] (in terms of the CAT(0) metric):

Lemma 3.10 (Local isometry ⇒ Convex embedding)**.** *If $\phi : Y \to X$ is a local isometry then $\widetilde{\phi} : \widetilde{Y} \to \widetilde{X}$ is an embedding as a convex subcomplex (hence in particular an isometric embedding).*

FIGURE 35. Considering a minimal diagram between a path $\tilde{\gamma}$ and a local isometry $\widetilde{Y} \to \widetilde{X}$ shows that $\tilde{\gamma}$ must lift to $\widetilde{Y}$ if $\tilde{\gamma}$ is a geodesic.

Proof. Consider a geodesic $\tilde{\gamma} \to \widetilde{X}$ that is homotopic in $\widetilde{X}$ to a path $\tilde{\sigma} \to \widetilde{Y}$. Let $D \to \widetilde{X}$ be a disc diagram for $\tilde{\gamma}\tilde{\sigma}^{-1}$, and suppose that D has minimal area among all possible such choices with $\tilde{\sigma} \to \widetilde{Y}$ 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 $\widetilde{Y}$ 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 35 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 35 is impossible by minimal area of D.

We are thus in the third situation of Figure 35, and so there is a cornsquare as in Remark 3.7. But then we can push this cornsquare towards $\widetilde{Y}$ and pass it through, leading to a new choice of $\tilde{\sigma}$ and having reduced the area of the disc diagram by 1. This is impossible. □

Corollary 3.11 (Carriers are convex)**.** *Let H be a hyperplane in $\widetilde{X}$ and let $N(H)$ be its carrier. Then $N(H)$ is a convex subcomplex of $\widetilde{X}$.*

Let H be a hyperplane in $\widetilde{X}$, let $\overleftarrow{H}$ and $\overrightarrow{H}$ be the closures of the two components of $\widetilde{X} - H$, and let $N(\overleftarrow{H})$ and $N(\overrightarrow{H})$ denote the smallest subcomplexes of $\widetilde{X}$ containing these subspaces. These are the two *halfspace carriers* associated with H. it follows from Corollary 3.11 that:

Corollary 3.12. *Halfspace carriers are convex*

Corollary 3.3 is useful in Section 3.4 since the intersection of convex subcomplexes of $\widetilde{X}$ is again convex.

3.4. **Cores, Hulls, and Superconvexity.**

Theorem 3.13. *Let $\widetilde{X}$ be a CAT(0) cube complex that is locally finite and δ-hyperbolic (thus finite dimensional). Let H act on $\widetilde{X}$ with a quasiconvex orbit $H\tilde{x}$. For each compact set $K \subset \widetilde{X}$ there is an H-cocompact subcomplex $\widetilde{Y}$ such that $HK \subset \widetilde{Y}$.*

Lemma 3.14. *There exists θ and D such that any metric-geodesic γ of length D crosses a hyperplane U of $\widetilde{X}$ with angle $\geq \theta$.*

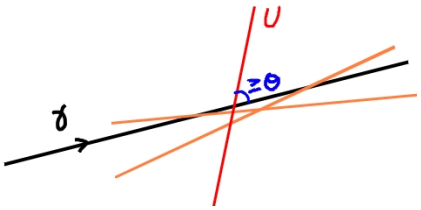

FIGURE 36. Some hyperplane U crosses γ with angle $\geq \theta$.

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

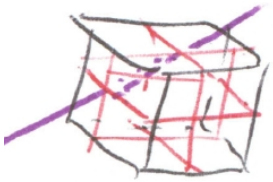

FIGURE 37. Some midcube makes a large angle with a geodesic in a d-cube.

Proof of Theorem 3.13. Consider the following convex subcomplex consisting of the intersection of halfspace carriers containing HK:

$$\widetilde{Y} = \mathrm{Hull}(HK) = \cap_{HK \subset \overrightarrow{U}} N(\overrightarrow{U})$$

Observe that $N(\overrightarrow{U}) \subset N_R(H\tilde{x})$ (as we can assume $\tilde{x} \subset K$) because if $d(p, HK) > R = D + \mathrm{diam}(K)$ then Lemma 3.14 provides a hyperplane U making a large angle with a geodesic γ from p to HK, and a quick computation using the quasiconvexity of HK shows that $U \cap HK = \emptyset$ because of δ-thin triangles. □

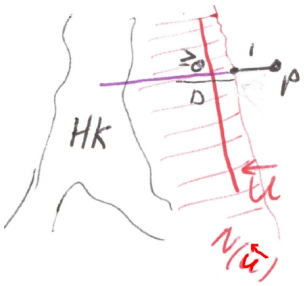

FIGURE 38. A point p that is far from HK is separated from HK by a hyperplane U.

Corollary 3.15. *Let X be a compact nonpositively curved cube complex with $G = \pi_1 X$ hyperbolic. Then for each quasiconvex subgroup H of G there exists a compact local isometry $Y \to X$ such that:*

$$\begin{array}{ccc} H & \subset & G \\ \| & & \| \\ \pi_1 Y & \to & \pi_1 X \end{array}$$

Proof. Choose $\widetilde{Y}$ using Theorem 3.13 and let $Y = H\backslash\widetilde{Y}$. □

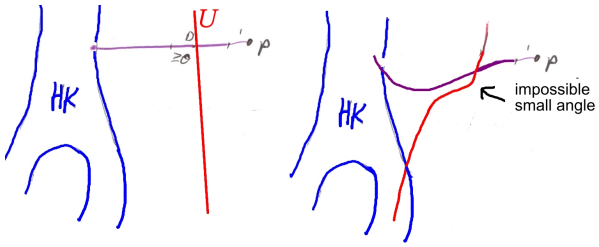

FIGURE 39. If U separates the quasiconvex subspace HK, then a triangle with two long sides consisting of part of γ and a geodesic in U would have fellow traveling at the corner near p which violates that $\sphericalangle \geq \theta$.

Remark 3.16. There is a relatively hyperbolic version of this as well which provides a "sparse" Y when G is hyperbolic relative to abelian subgroups and X is compact (or sparse).

Definition 3.17 (Superconvex)**.** $\widetilde{Y} \subset \widetilde{X}$ is *superconvex* if it is convex and for each bi-infinite geodesic $\widetilde{\gamma} \subset \widetilde{X}$, if $\widetilde{\gamma} \subset N_r(\widetilde{Y})$ for some $r \geq 0$, then $\widetilde{\gamma} \subset \widetilde{Y}$.

Lemma 3.18 (Superconvex core)**.** *Let $H \subset G$ be a quasiconvex subgroup with G acting properly and cocompactly on a CAT(0) cube complex $\widetilde{X}$. For each compact subcomplex K, the subspace HK lies in an H-cocompact superconvex subcomplex $\widetilde{Y} \subset \widetilde{X}$.*

Proof. δ-hyperbolicity and κ-quasiconvexity imply that any bi-infinite geodesic $\tilde{\gamma}$ that is r-close to $H\tilde{x}$ actually lies in $N_s(H\tilde{x})$ for some uniform s. Now let $\widetilde{Y} = \text{Hull}(N_s(HK))$. □

4. Special Cube Complexes

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 $\widetilde{X}$. Note that there is a natural map $\bar{H} \to X$ and that each midcube of a cube of X extends to a unique immersed hyperplane.

Figure 40. An immersed hyperplane in a nonpositively curved cube complex

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

The prohibited hyperplane pathologies are depicted in Figure 41

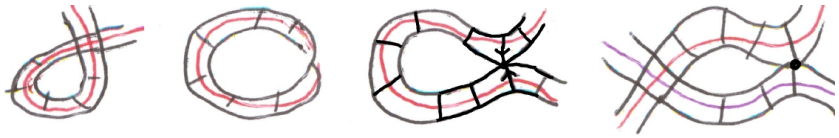

Figure 41. Self-crossing, 1-sidedness, self-osculation, and interosculation

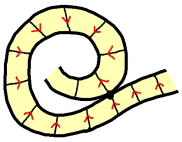

Figure 42. The above configuration is not prohibited in a special cube complex – note the orientation which makes it an "indirect self-osculation"

Example 4.3. (1) Any graph is special.
(2) Any CAT(0) cube complex is special.
(3) Any subcomplex of a product of two graphs $A \times B$ 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 [Wis06]). Even this turns out to be a surprisingly rich family of examples.

The fundamental example is:

The cube complex $R(\Gamma)$ of a graph group $G(\Gamma)$ is a special cube complex.

The following provides a useful alternate characterization:

Lemma 4.4. *Let X be a nonpositively curved cube complex. Then X is special iff there is a local-isometry $X \to R$ 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 graph group.*

Proof of Lemma 4.4. ($\Leftarrow$) If $A \to B$ 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 $X^1 \to R(\Gamma)$ which extends to a local isometry $X \to R$. □

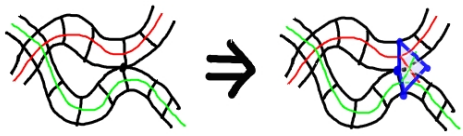

FIGURE 43. More careful depiction of the "no interosculation" condition

FIGURE 44. The crossing graph Γ of hyperplanes in X

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

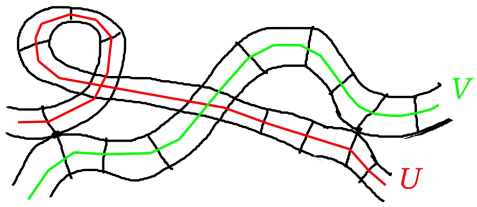

FIGURE 45. Three double hyperplane cosets $\pi_1 U \pi_1 V$ to consider here.

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.

In particular, a subgroup H is separable if and only if $H = \cap_{H \subset K, [G:K] < \infty} K$, and $\{1_G\}$ is separable if and only if the profinite topology on G is hausdorff.

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

Lemma 4.8. *Let $\widehat{X} \to X$ 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 $\widehat{X} \to X$ factors as $\widehat{X} \to \bar{X} \to X$ and such that K embeds in $\bar{X}$ as follows:*

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

A well-known consequence of Lemma 4.8 is:

Corollary 4.9. *Suppose Y is compact and $Y \to X$ is π_1-injective, and Y lifts to an embedding in the cover $\widehat{X}$ associated to $\pi_1 Y$, and $\pi_1 Y$ is separable. Then $Y \to X$ 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 $\widetilde{X}$ with finitely many hyperplane orbits and finitely many Stabilizer(U)-orbits of hyperplanes intersecting U and osculating U for each hyperplane U of $\widetilde{X}$. (We need Intersector(U, V) and 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 [NR03] of Coxeter groups in [HW10], and we used it to verify the virtual specialness of simple-type arithmetic hyperbolic lattices in [BHW11, BW].

4.3. **Canonical Completion and Retraction.**

Definition 4.11 ($A \otimes_R B$)**.** Let $\alpha : A \to R$ and $\beta : B \to R$ 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, r) where $r = \alpha(a) = r = \beta(b)$. 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 and whose morphisms are maps $C_1 \to C_2$ so that there is a further commutative diagram:

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

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

Construction 4.12. Let $f : Y \to X$ be a local isometry with X special and Y compact. There exists a finite cover $\rho : \mathsf{C}(Y \to X) \to X$ and an embedded lift $\widehat{f} : Y \to \mathsf{C}(Y \to X)$ of

$f : Y \to X$, and a retraction map $r : \mathsf{C}(Y \to X) \to Y$. The maps ρ and r are the *canonical completion* and *canonical retraction* associated to $f : Y \to X$.

Let $X \to R$ denote the local isometry to the nonpositively curved cube complex of the right-angled Artin group associated with X in Lemma 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 $\mathsf{C}(Y \to R)$. We then define $\mathsf{C}(Y \to X) = X \otimes_R \mathsf{C}(Y \to R)$. We refer the reader to the following diagram as well as to a corresponding concrete example in Figure 46:

$$\begin{array}{ccccc} \mathsf{C}(Y \to X) = & X \otimes_R \mathsf{C}(Y \to R) & \to & \mathsf{C}(Y \to R) \\ & \nwarrow \quad \nearrow & & \\ & \downarrow \qquad Y \qquad \downarrow & & \\ & \swarrow \quad \searrow & & \\ & X & \to & R \end{array} \tag{2}$$

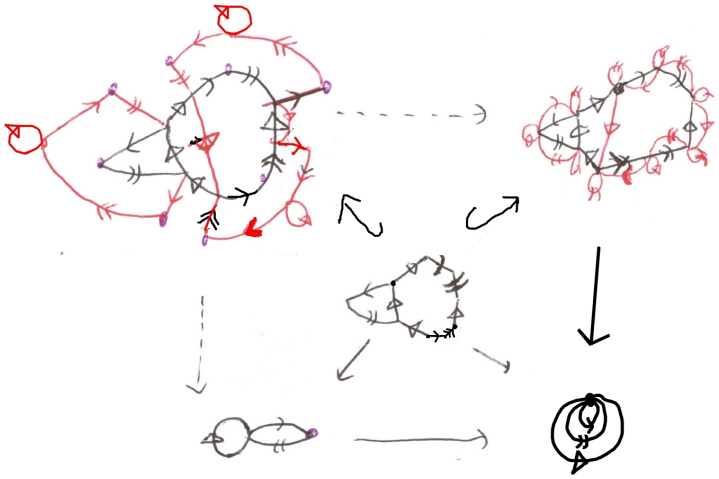

FIGURE 46. Canonical Completion in 1-dimensional case

We now turn to the general case. The key point is that for a local isometry $Y \to R$ the previously defined map $\mathsf{C}(Y^1 \to R^1) \to R^1$ extends to a covering map $\mathsf{C}(Y \to R) \to R$. We then again use the fiber-product to define $\mathsf{C}(Y \to R)$ as in the 1-dimensional case exactly as in Equation (2).

Finally we note that the retraction maps which we described explicitly when R is a bouquet of circles, are are defined using compositions in general.

4.4. **Separability in the hyperbolic case.**

Theorem 4.13. *If X is special and compact and $\pi_1 X$ is hyperbolic, then every quasiconvex subgroup H of $G = \pi_1 X$ is separable.*

Proof 1. Let $\sigma \notin H$. Choose $\widetilde{Y}$ that is H-cocompact and contains $\tilde{\sigma}$. Let $Y = H\backslash\widetilde{Y}$. Let $G' = \pi_1 \mathsf{C}(Y \to X)$. Then $\sigma \notin G' \supset H$. □

Proof 2, assumes residual finiteness, which is itself a bit easier. As before, let $Y \to X$ be an arbitrary compact local isometry representing H. Then Y is a retract of $\mathsf{C}(Y \to X)$. Thus H is a retract of $G' = \pi_1 \mathsf{C}(Y \to X)$. Thus H is closed in G' since a retract of a hausdorff topological space is closed. Hence H is separable in G'. □

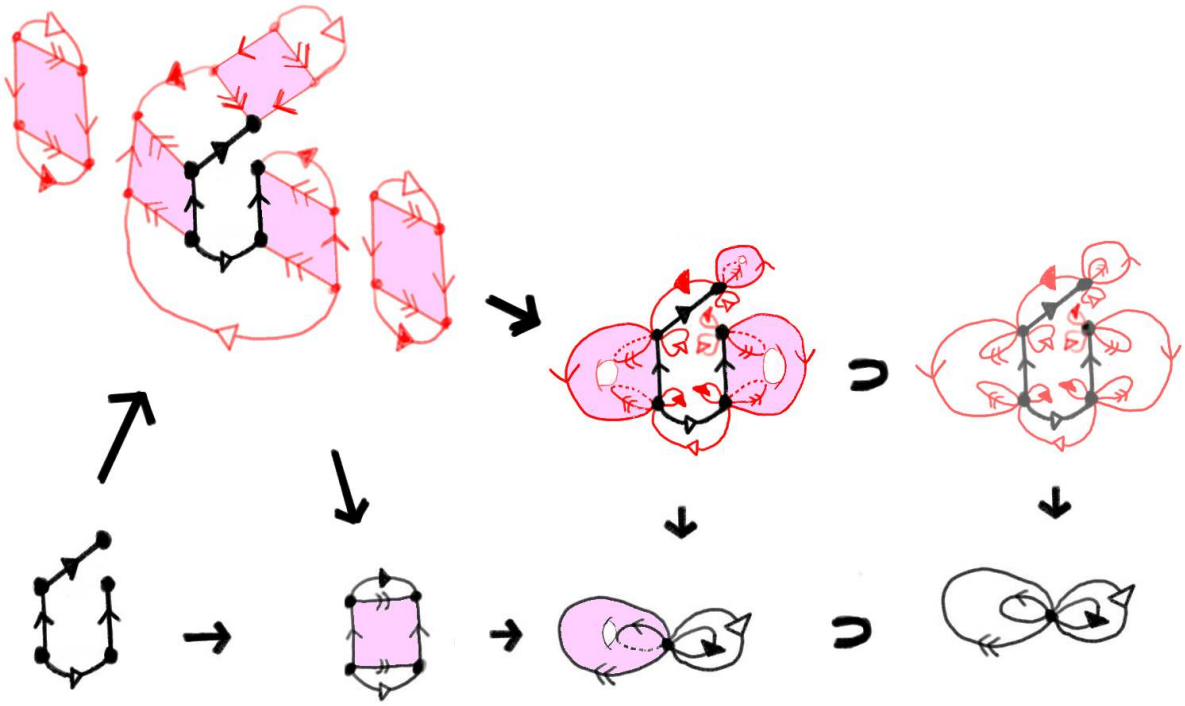

FIGURE 47. The figure corresponds to the following commutative diagram:

$$\begin{array}{ccccccc} & & \mathsf{C}(A \to B) & \to & \mathsf{C}(A \to R) & \supset & \mathsf{C}(A^1 \to R^1) \\ & \nearrow & \downarrow & & \downarrow & & \downarrow \\ A & \to & B & \to & R & \supset & R^1 \end{array}$$

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 3.11. If all quasiconvex subgroups are separable then all quasiconvex double cosets are separable by Theorem 4.15. The result thus follows from Theorem 4.6. □

The following result was obtained by Minasyan clarifying earlier work of Gitik [Min04, Git99].

Theorem 4.15. *If all quasiconvex subgroups of a hyperbolic group G are separable then all double quasiconvex cosets of G are separable.*

Remark 4.16. In the compact hyperbolic case X is special if and only if all hyperplane subgroups are separable. This holds because X has a finite cover $\widehat{X}$ with a malnormal hierarchy - something we will explore later. In the compact 2-dimensional $\mathcal{VH}$ case, X is virtually special if and only all hyperplane subgroups are separable, since separable hyperplanes implies virtual cleanliness, and compact clean complexes virtually embed in the product $A \times B$ of graphs [Wis06]. Very little is know about the general case.

4.5. **Wall-injectivity and a fundamental commutative diagram.**

Definition 4.17 (Wall-injective)**.** A combinatorial map of cube complexes $D \to C$ 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 be injective on the set of 1-cells. Figure 48 depicts an embedding that is not wall-injective.

FIGURE 48. The inclusion above is not wall-injective

Being wall-injective is not stable under taking covering spaces, as can be seen in Figure 49.

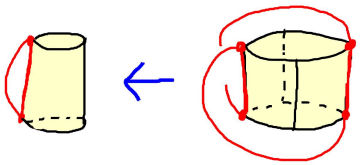

FIGURE 49. The preimage of a wall-injective subcomplex in a covering space might not be wall-injective, as is the case for the red circle above.

Remark 4.18. Consider the canonical completion $\mathsf{C}(Y \to X)$ of a local isometry $Y \to X$. Observe that $Y \hookrightarrow \mathsf{C}(Y \to X)$ is wall-injective. Indeed this follows by examining the behavior of the canonical retraction map.

Lemma 4.19 (Fundamental commutative diagram)**.** *Suppose $D \to C$ is a wall-injective local-isometric embedding with D finite and C special. Then there are commutative diagrams whose vertical maps are canonical completions, canonical retractions, and canonical inclusions.*

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

Sketch. In the 1-dimensional case, imagine first building $\mathsf{C}(D \to D)$ and then building $\mathsf{C}(D \to C)$ around it. The general case is similar. □

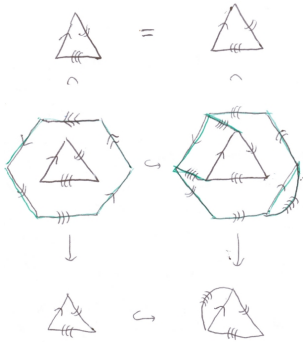

FIGURE 50. The fundamental commutative diagram

4.6. **Wall Projection Controls Retraction.**

Definition 4.20 (Wall Projection)**.** Let $A \subset X$ and $B \subset X$ be subcomplexes of a nonpositively curved cube complex X. The 1-skeleton of $\mathsf{WProj}_X(B \to A)$ 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 al cubes whose 1-skeleta are included. (We think of the cubes being likewise parallel into B.)

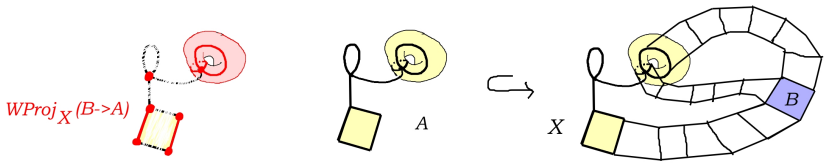

FIGURE 51. $\mathsf{WProj}_X(B \to A)$ is indicated on the left, where A, B are the indicated subcomplexes of the complex X on the right.

We note that $\mathsf{WProj}_X(B \to A)$ consists of a collection of locally-convex subcomplexes of A. In the 1-dimensional case, $\mathsf{WProj}_X(B \to A) = A^0 \cup (B \cap A)$.

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

Lemma 4.21. *Let A, B be subcomplexes of the special cube complex X with A compact. Consider an elevation $\breve{B}$ of B to $\mathsf{C}(A \to X)$. Then $r(\breve{B}) \subset \mathsf{WProj}_X(B \to A)$.*

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

Lemma 4.21 is powerful when $\mathsf{WProj}_X(B \to A)$ is *trivial* in the sense that each component of $\mathsf{WProj}_X(B \to A)$ is simply connected - i.e. CAT(0). For then $r(\breve{B})$ is null-homotopic for each $\breve{B}$.

Lemma 4.22 (Trivial Wall Projection)**.** *Let X be a compact (virtually) special nonpositively curved cube complex with $\pi_1 X$ hyperbolic. let $A \to X$ be a compact local isometry with $\pi_1 A \subset \pi_1 X$ malnormal. There exists a finite cover $A_o \to A$ such that any further cover $\bar{A} \to A_o$ can be completed to a finite special cover $\bar{X} \to X$ such that:*

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

Sketch. This is easy when X is 1-dimensional: Let $X_o \to X$ be a finite cover in which all elevations of A are embedded, and let A_o be the base elevation. Note that $\breve{A}_1 \cap \breve{A}_2$ is a forest because of malnormality. For each $\bar{A} \to A_o$ we let $\bar{X} = \mathsf{C}(\bar{A} \to X_o)$. □

The proof of Lemma 4.22 in higher dimensions is treacherous...

5. Virtual Specialness of Malnormal Amalgams

Theorem 5.1. *Let G be a hyperbolic group that acts properly and cocompactly on the CAT(0) cube complex $\widetilde{Q}$. Let $\widetilde{P}$ be a hyperplane such that* Stabilizer$(\widetilde{P})$ *is almost malnormal in G and such that $g\widetilde{P} \cap \widetilde{P} \neq \emptyset \Rightarrow g\widetilde{P} = \widetilde{P}$ and $g\overrightarrow{\widetilde{P}} = \overrightarrow{\widetilde{P}}$.*

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

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

We focus on illustrating the following geometric special case:

Theorem 5.2. *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.*

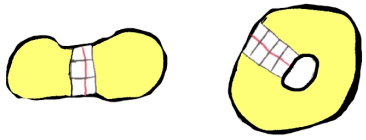

Figure 52. $P \subset N(P) \subset Q$ in the separating and nonseparating cases.

Note that $N_o(P)$ denotes the *open cubical neighborhood* of P consisting of all open cubes intersecting P. Theorem 5.2 is already interesting in the case that each component of $Q - N_o(P)$ is a graph, which is the main result in [Wis02]. Surprisingly, there is little formal difference between the proofs, though several easy steps in the 1-dimensional case turn out to be much deeper and technically challenging to verify in arbitrary dimensions.

A main ingredient in the proof of Theorem 5.2 is the following:

Lemma 5.3. *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 $\widehat{X} \to X$, there is a finite regular cover $\overset{\bowtie}{X}$ factoring through $\widehat{X}$ such that $\overset{\bowtie}{X}$ induces the same cover on each side A, B of P.

We prove Theorem 5.2 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 Q is separable.*

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 $A \to X$ and $B \to X$.

Let $\dot{Q} \to Q$ 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 $\dot{Q}$.

Let Y be a compact core containing $\dot{\sigma}$ (see Figure 53).

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 $\widehat{X}_i \to X$ (with $\dot{X}_i \to \widehat{X}_i \to X$) such that $Y_i \hookrightarrow \dot{X}_i$ projects to an

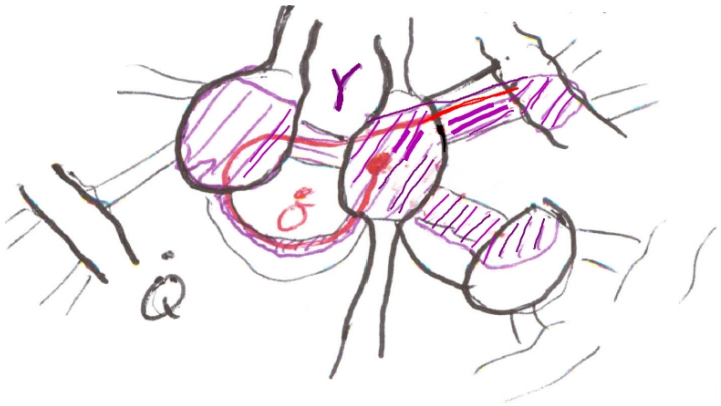

FIGURE 53. The compact core Y of $\dot{Q}$ is chosen to contain $\dot{\sigma}$

embedding $Y_i \hookrightarrow \widehat{X}_i$. And moreover, we use double quasiconvex coset separability 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 54.

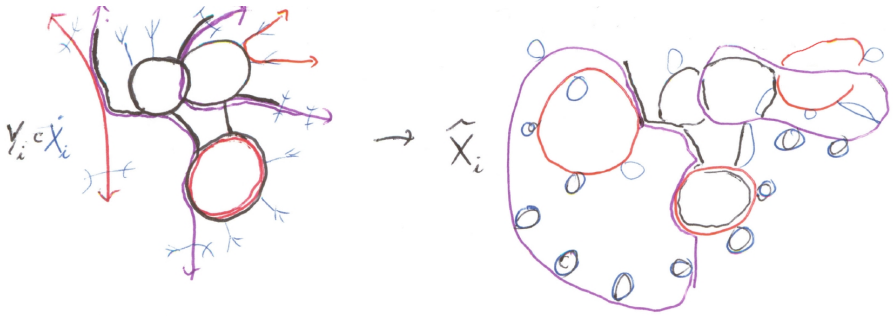

FIGURE 54. The intermediate cover $\widehat{X}_i$ maintains the embedding of Y_i and also maintains the distinctness of the attachment sites of those incoming and outgoing edgespaces at $\dot{X}_i$ that are represented in Y. We have sketched an example where X is a graph, and all edge spaces are cylinders, so the incoming (red) and outgoing (purple) edge spaces are lines and circles which project to circles in $\widehat{X}_i$

Let Z be obtained from Y by embedding each Y_i to $\widehat{X}_i$ (see Figure 55). Note that $Z \to X$

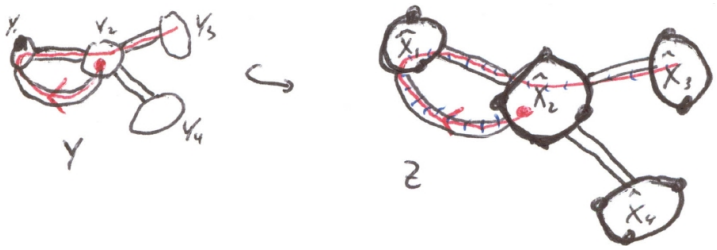

FIGURE 55. We then extend $Y \to X$ to $Z \to X$ by extending each vertex space Y_i to the chosen finite cover $\widehat{X}_i$.

is a finite cover on each vertex space, and contains some information on how to attach edge spaces. We now create a finite cover $\overset{\bowtie}{Z} \to Z$ such that all vertex spaces are isomorphic:

Let $\widehat{X}$ be a finite cover factoring through each $\widehat{X}_i$. Let $\overset{\bowtie}{X} \to \widehat{X}$ be as in Lemma 5.3. Let $\overset{\bowtie}{Q} \to Q$ be a finite cover whose unique vertex spaces is $\overset{\bowtie}{X}$ (see Figure 56).

FIGURE 56. The finite cover $\overset{\bowtie}{Q}$ has vertex space isomorphic to $\overset{\bowtie}{X}$ and an arbitrarily attached edge spaces associated to a one-to-one correspondence between elevations of A and B.

Let $\overset{\bowtie}{Z}=\overset{\bowtie}{Q}\otimes_Q Z$ so there is a commutative diagram:
$$\begin{array}{ccc} \overset{\bowtie}{Z} & \to & Z \\ \downarrow & & \downarrow \\ \overset{\bowtie}{Q} & \to & Q \end{array}$$

The partial one-to-one correspondence between elevations of A, B to vertex spaces of $\overset{\bowtie}{Z}$ 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 $\overset{\bowtie}{X}$.

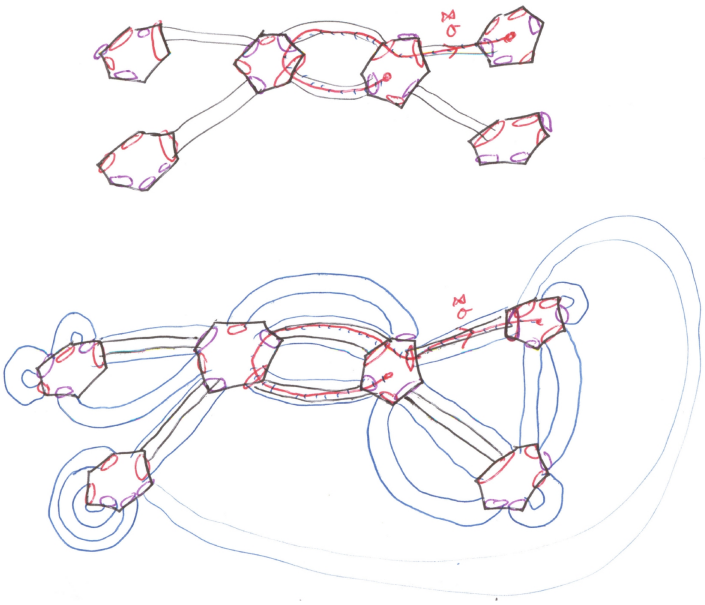

FIGURE 57. The edge spaces of $\overset{\bowtie}{Z}$ (above) are extended to finite covers, and the missing edge spaces are added to obtain $\overset{\bowtie}{X}$ (below).

Finally, σ is separated from $\pi_1 Q$ in the right coset representation in $\pi_1 \overset{\bowtie}{Q}$. Because each element of $\pi_1 Y$ sends the trivialcoset to a coset correspond to a preimage of the basepoint of Y (i.e. a lift of a basepoint of Z) in $\overset{\bowtie}{Z}$, but σ maps it to the preimage of the lift of an endpoint of σ. □

5.1. **Proof of Lemma 5.3.** We prove Lemma 5.3 in the separating case instead. So let $Q = N_o(P) = L \sqcup R$ and let $A \hookleftarrow L$ and $B \hookrightarrow R$ denote the attaching maps of $P \times \pm 1$ as in Figure 58

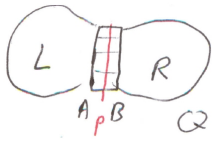

FIGURE 58. $Q = L \sqcup_{A \cong B} R$.

We use Lemma 4.22 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) $\mathsf{WProj}_{\widehat{L}}(A_i \to \widehat{A})$ and $\mathsf{WProj}_{\widehat{R}}(B_i \to \widehat{B})$ are trivial for all nonbase elevations A_i, B_i.

Form $\mathsf{C}(\widehat{A} \to \widehat{L})$ and $\mathsf{C}(\widehat{B} \to \widehat{R})$ and note that we have the following commutative diagram from Lemma 4.19:

$$\begin{array}{ccccccc} \mathsf{C}(\widehat{A} \to \widehat{L}) & \hookleftarrow & \mathsf{C}(\widehat{A} \to \widehat{A}) & \cong & \mathsf{C}(\widehat{B} \to \widehat{B}) & \hookrightarrow & \mathsf{C}(\widehat{B} \to \widehat{R}) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \widehat{L} & \hookleftarrow & \widehat{A} & \cong & \widehat{B} & \hookrightarrow & \widehat{R} \end{array}$$

Let $\bar{P} \to P$ be a finite cover of P factoring through all (corresponding) elevations of A to $\mathsf{C}(\widehat{A} \to \widehat{L})$ and elevations of B to $\mathsf{C}(\widehat{B} \to \widehat{R})$. Let $\bar{A}, \bar{B}$ be the corresponding covers of A, B. We use canonical retraction to induce the following covers:

$$\begin{array}{ccc} \overline{\mathsf{C}(\widehat{A} \to \widehat{L})} & \to & \bar{A} \\ \downarrow & & \downarrow \\ \mathsf{C}(\widehat{A} \to \widehat{L}) & \to & \widehat{A} \end{array} \quad \begin{array}{ccc} \overline{\mathsf{C}(\widehat{A} \to \widehat{A})} & \to & \bar{A} \\ \downarrow & & \downarrow \\ \mathsf{C}(\widehat{A} \to \widehat{A}) & \to & \widehat{A} \end{array} \quad \begin{array}{ccc} \overline{\mathsf{C}(\widehat{B} \to \widehat{B})} & \to & \bar{B} \\ \downarrow & & \downarrow \\ \mathsf{C}(\widehat{B} \to \widehat{B}) & \to & \widehat{B} \end{array} \quad \begin{array}{ccc} \overline{\mathsf{C}(\widehat{B} \to \widehat{R})} & \to & \bar{B} \\ \downarrow & & \downarrow \\ \mathsf{C}(\widehat{B} \to \widehat{R}) & \to & \widehat{R} \end{array}$$

We obtain the following commutative diagram:

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

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

of A to $\widehat{L}$.

$$\begin{array}{ccc} \bar{\widehat{A}}_i & \subset & \overline{\mathsf{C}(\widehat{A} \to \widehat{L})} \\ \| & & \downarrow \\ \widehat{A}_i & \subset & \mathsf{C}(\widehat{A} \to \widehat{L}) \\ \downarrow & & \downarrow \\ A_i & \to & \widehat{L} \\ \downarrow & & \downarrow \\ A & \to & L \end{array}$$

The key point is that $\bar{\widehat{A}}_i \cong \widehat{A}_i$ as, by Lemma 4.21, the triviality of $\mathsf{WProj}_{\widehat{L}}(A_i \to \widehat{A})$ implies that $\widehat{A}_i$ is nullhomotopic in the canonical retraction map $r : \mathsf{C}(\widehat{A} \to \widehat{L}) \to \widehat{A}$. Thus $\widehat{A}$ factors through $\bar{\widehat{A}}_i \cong \widehat{A}_i$ by choice of $\bar{A}$. Thus $\overset{\bowtie}{A}$ factors through $\bar{\widehat{A}}_i$ since $\overset{\bowtie}{A}$ factors through $\bar{A}$.

Similarly $\overset{\bowtie}{B}_L \cong \overset{\bowtie}{B}$.

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

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

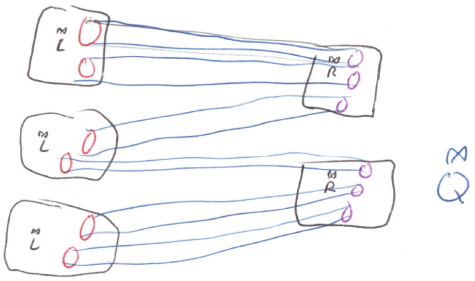

FIGURE 59. We form $\overset{\bowtie}{Q}$ by taking a balanced number of copies: $3 \cdot \overset{\bowtie}{L}$ and $2 \cdot \overset{\bowtie}{R}$ and gluing them together along copies of $\overset{\bowtie}{P} \times [-1, 1]$.

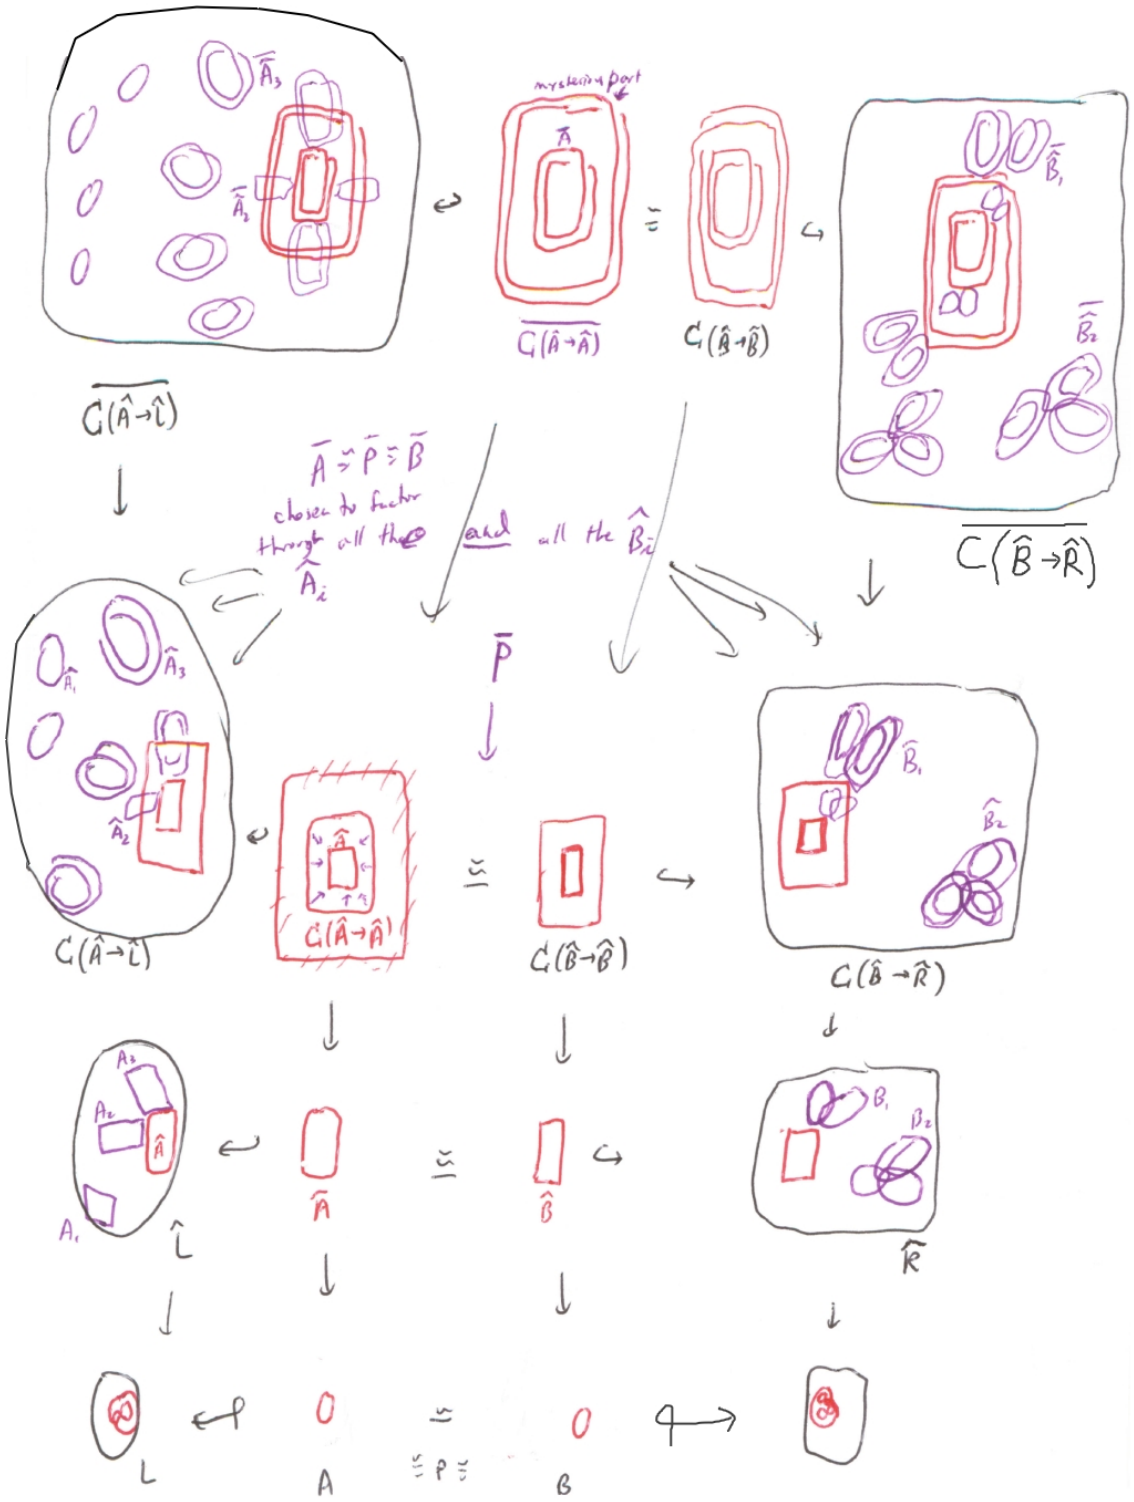

FIGURE 60. The proof of Lemma 5.3.

6. Wallspaces and their Dual Cube Complexes

6.0.1. *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 $p \in \overleftarrow{W} \cap \overrightarrow{W}$.

Remark 6.1. Haglund-Paulin's original wallspaces have a more elegant definition: A wall is simply a partition of X and the finite betwixting property holds automatically.

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

Figure 61. Two wallspaces.

6.0.2. *The dual CAT(0) cube complex.* The CAT(0) cube complex *dual* to a wallspace is defined as follows: It has a *0-cube* v for an *orientation* on all walls – that is, choosing one of the two halfspaces of each wall – such that:

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

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

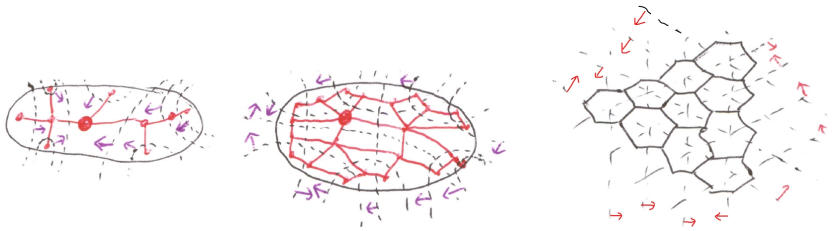

Figure 62. The dual cube complexes are: a tree, a square complex, and a 3-dimensional CAT(0) cube complex that is harder to depict.

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

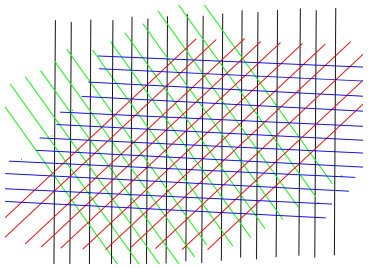

FIGURE 63. A system of n bi-infinite families of real lines in $\mathbb{R}^2$ has dual cube complex isomorphic to $\widetilde{T}^n$.

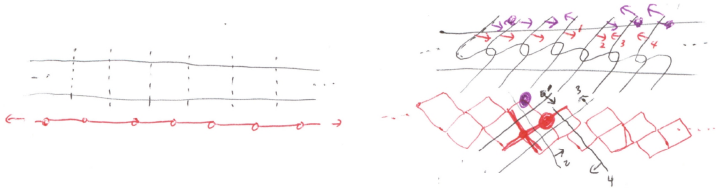

FIGURE 64. Two different wallspaces on an infinite strip and their dual cube complexes.

6.0.3. *Some examples:*

Example 6.3 (Wallspaces from tracks)**.**

Example 6.4 (Wallspaces from graphs of spaces)**.**

Example 6.5 (Wallspaces from small-cancellation theory)**.** Every simply connected $C'(\frac{1}{6})$ complex $\widetilde{X}$ is a wallspace: The walls are graphs that intersect 2-cells and 1-cells of $\widetilde{X}$ in *midcells*. We first subdivide the 1-skeleton so that all 2-cells have an even circumference. Some facts about these walls:

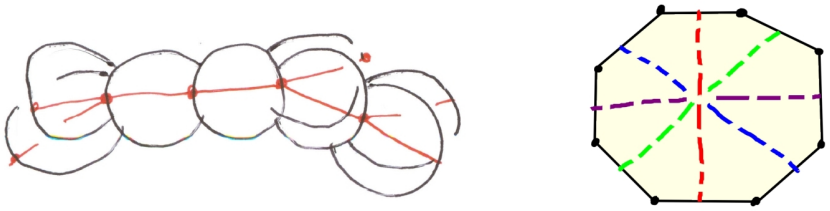

FIGURE 65. A wall in a $C'(\frac{1}{6})$ small-cancellation complex on the left, and the 4 midcells of an octagon on the right.

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

The above properties are proven using Greendlinger's lemma and the ladder theorem - See Section 8.2.1.

6.0.4. *Wallspaces from codimension-1 subgroups.* Let G be a finitely generated group. $H \subset G$ is a *codimension-1* subgroup if from some $r > 0$, there are two or more components K_i in $\Gamma - N_r(H)$ that are *deep* in the sense that $K_i \not\subset N_s(H)$ for any $s > 0$. This doesn't depend on the choice of Cayley graph $\Gamma = \Gamma(G, S)$.

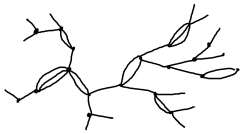

FIGURE 66. A wall in a $C'(\frac{1}{6})$ complex is a multi-tree.

Example 6.6. The following examples are illustrated in Figure 67

(1) $\mathbb{Z}^n \subset \mathbb{Z}^{n+1}$
(2) $\pi_1 S^1 \looparrowright \pi_1 M^2$
(3) $\pi_1 M^2 \looparrowright \pi_1 M^3$ with M^2, M^3 aspherical.
(4) $C \subset A *_C B$ with $C \subsetneq A, C \subsetneq B$.

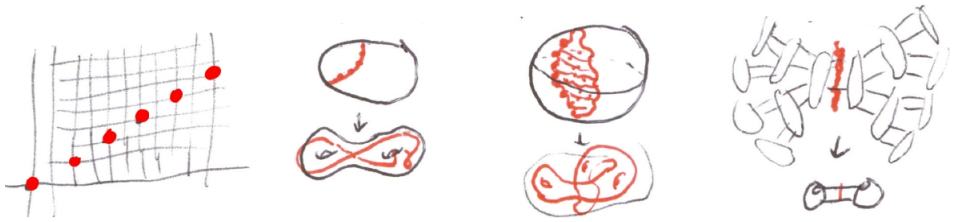

FIGURE 67. Some codimension-1 subgroups.

To obtain a wall from a codimension-1 subgroup, we partition the component index set I into $I = \overleftarrow{I} \sqcup \overrightarrow{I}$ and let $W = N_r(H)$ and $(\overleftarrow{W}) = W \cup_{i \in \overleftarrow{I}} K_i$ and $(\overrightarrow{W}) = \cup_{i \in \overrightarrow{I}} K_i$. This can be done for any subgroup, but we have in mind the case where $N_r(H)$ separates Γ, and where at least one deep component is assigned to each halfspace of W.

Remark 6.7 (Warning)**.** Stabilizer(W) and Stabilizer$(\overleftarrow{W})$ might only be commensurable with H instead of equal to it.

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.

Why is $\#(p, q) < \infty$? Why are there finitely many walls betwixted by any point p?

Example 6.8. 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 [Wisa].

Example 6.9. For closed hyperbolic M^3, Kahn-Markovic provide many immersed incompressible surfaces $S_i \looparrowright M$. Choosing finitely many gives 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 a its dual CAT(0) cube complex.

6.1. **Finiteness properties of the action on the dual cube complex.** Two walls $(\overleftarrow{W}_1, \overrightarrow{W}_1), (\overleftarrow{W}_2, \overrightarrow{W}_2)$ *cross* if all four quarterspaces are nonempty:

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

We emphasize that in the geometric wallspace case, the walls cross precisely when $W_1 \cap W_2 \neq \emptyset$.

6.1.1. *The cubes of C:* 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 $x \in X$.

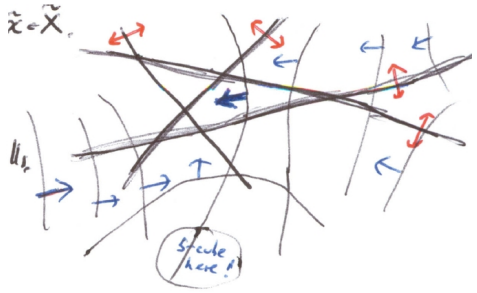

FIGURE 68. The four bold independently orientable walls together with orientations on all the other walls determine a 4-cube in the dual cube complex. That 4-cube is not maximal: there is a fifth wall that crosses each of these four.

Example 6.10. If walls in X arise from finite collections of embedded tracks in X, say $T_1, \dots, T_d$, then $\dim(C) \leq d$.

Example 6.11. Rubinstein-Wang found a graph manifold M and immersed incompressible surface $S \looparrowright M$ such that all translates of $\widetilde{S}$ intersect in the sense that $g_1\widetilde{S} \cap g_2\widetilde{S} \neq \emptyset$ in $\widetilde{M}$. Thus $\dim(C) = \infty$, and the dual cube complex C is actually an infinite cube.

6.1.2. *The bounded packing property and finite dimensionality:*

Definition 6.12. $H \subset G$ 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.

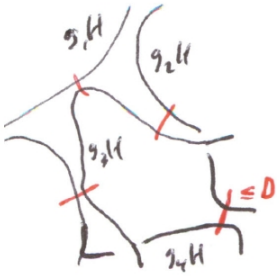

FIGURE 69. A collection of four pairwise D-close cosets of H.

Example 6.13. Finite subgroups, normal subgroups, f.g. separable subgroups [Yan] have bounded packing.

The Rubinstein-Wang example in Example 6.11 does not have bounded packing. Neither does a certain subgroup of Thompson's group T of the stabilizer of a hyperplane in the action of T on the infinite-dimensional CAT(0) cube complex discovered by Farley [Far03].

Lemma 6.14 (Sageev)**.** *If H is a quasiconvex subgroup of a hyperbolic group G then H has bounded packing.*

Proof. This follows from bounded height. □

In [HW09] we abstracted the notion of bounded packing, and generalized Sageev's result to:

Lemma 6.15. *If G is hyperbolic relative to $\{P_i\}$ and each P_i has bounded packing relative to its finitely generated subgroups then G has bounded packing relative to quasiconvex subgroups.*

6.1.3. *Cocompactness in Hyperbolic Case.*

Theorem 6.16 (Sageev 97)**.** *If X is δ-hyperbolic and G acts properly and cocompactly on $\widetilde{X}$, and there are finitely many G-orbits of walls, and each wall W is quasiconvex (i.e.* 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 Lemma 6.14 and hence we can apply the following lemma to see that there are finitely many G-orbits of maximal cubes. □

Lemma 6.17 (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.*

6.1.4. *Relative cocompactness in the relatively hyperbolic case.*

Theorem 6.18. *Let G be hyperbolic relative to $\{P_i\}$. Suppose G acts properly and cocompactly on $\widetilde{X}$ and each P_i cocompactly stabilizes some $\widetilde{F}_i \subset \widetilde{X}$ Suppose $\widetilde{X}$ is a wallspace each wall has quasiconvex stabilizer. There exists a compact subcomplex $K \subset C(\widetilde{X})$ such that:*

$$C(\widetilde{X}) = GK \cup GC(\widetilde{F}_i)$$

and

$$g_i C(\widetilde{F}_i) \cap g_j C(\widetilde{F}_j) \subset GK$$

unless $\widetilde{F}_i = \widetilde{F}_j$ and $g_j^{-1} g_i \in$ Stabilizer$(\widetilde{F}_i)$.

Above, we use the notation $C(\widetilde{F}_i)$ for the cube complex dual to the wallspace on $\widetilde{F}_i$ consisting of walls such that $\overleftarrow{W} \cap \widetilde{F}_i, \overrightarrow{W} \cap \widetilde{F}_i \neq \emptyset$.

Usually we consider the case where each $W \cap \widetilde{F}_i$ is connected or $\emptyset$ (and preferably $F_i \not\subset W$). There is a natural embedding $C(\widetilde{F}_i) \subset C(\widetilde{X})$ as a convex subcomplex – we orient the walls not crossing $\widetilde{F}_i$ towards $\widetilde{F}_i$.

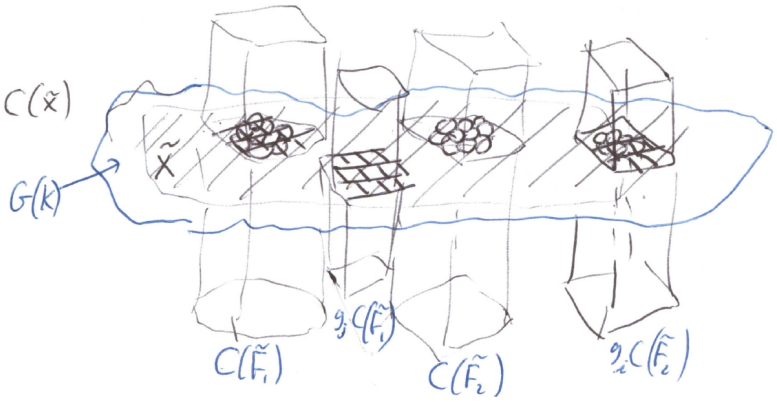

FIGURE 70. A relatively cocompact cubulation

Definition 6.19 (B6)**.** $\widetilde{X}$ satisfies the $B(6)$ small-cancellation condition if for each 2-cell R, the concatenation of three consecutive pieces $P_1P_2P_3$ in $\partial_p R$ satisfies $|P_1P_2P_3| \leq \frac{1}{2}|\partial_p R|$. Note that $\widetilde{X}$ satisfies $B(6)$ when each piece P satisfies $|P| \leq \frac{1}{6}|\partial_p R|$, and in particular $C'(\frac{1}{6})$ is $B(6)$.

Example 6.20. A *honeycomb* is a copy of the hexagonal tiling of the Euclidean plane, possibly with some of its 1-cells subdivided. See Figure 71.

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 $\widetilde{X}$ is $B(6)$ and contains a collection of honeycombs that are isolated in the sense that they have uniformly bounded overlap. The group $\pi_1 X$ hyperbolic relative to the stabilizers of these honeycombs which are virtually $\mathbb{Z}^2$. These honeycombs play the role of $\widetilde{F}_i$ in Theorem 6.18

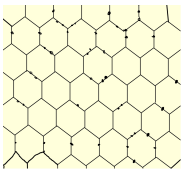

FIGURE 71. A honeycomb

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

6.2. **Properness of the G action on $C(\widetilde{X})$.** In a Haglund-Paulin wallspace $(\widetilde{X}, \mathcal{W})$, for each point $p \in \widetilde{X}$ 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 72. 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 $\widetilde{X} \to C(\widetilde{X})$ defined by $p \mapsto V_p$. Note that $\#(p, q) = d_C(V_p, V_q)$. It follows that properness of the action on C depends on relating $\#$ and $d_{\widetilde{X}}$.

Lemma 6.21. *G acts properly on C if and only if $\#(p, gp) \to \infty$ as $g \to \infty$.*

In practice we often verify the sufficient condition of *linear separation* which states that: there exists k_1, k_2 such that $\#(p, q) \geq k_1 d_{\widetilde{X}}(p, q) - k_2$.

This is normally done by finding $L > 0$ such that for any geodesic $\gamma \to \widetilde{X}$, 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.

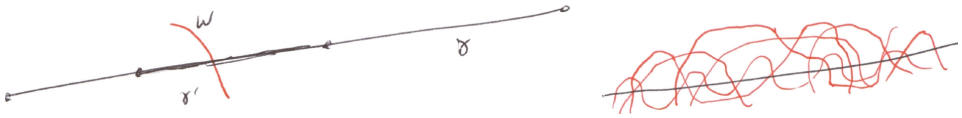

FIGURE 73. Verifying linear separation linear separation as on the left, can be a messy affair as on the right.

This can be tricky and messy to verify but in very extreme cases it is quite easy, for instance:

Lemma 6.22. *If $\widetilde{X}$ is a metric CAT(0) space, and each wall is a convex hyperplane (i.e. codimension-1 subspace), and components of $\widetilde{X} - \cup_{W\in\mathcal{W}} W$ have diameter uniformly bounded by D, then $\widetilde{X}$ 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. □

FIGURE 74. Linear separation is immediate when the walls are convex hyperplanes in a CAT(0) metric space, and the complementary regions have uniformly bounded diameter.

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

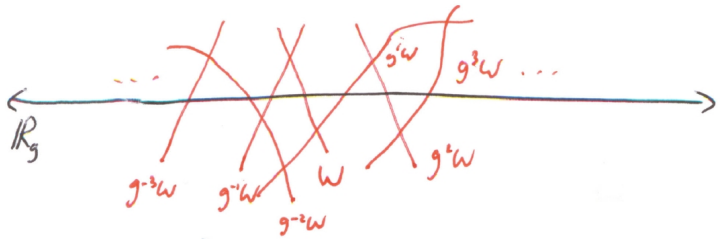

FIGURE 75. A cut-wall for g.

When G has finite torsion subgroups the following condition suffices to prove properness. It is a variation of some similar conditions examined in [HWa].

Lemma 6.23. *If each infinite order element $g \in G$ has a cut-wall then G acts with torsion stabilizers on C.*

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

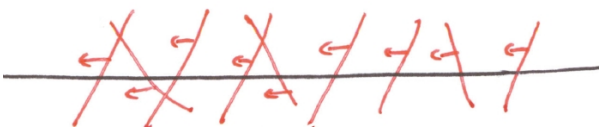

FIGURE 76. An infinite sequence of distinct walls that are directed away from a point on the line.

Example 6.24. The cut-wall criterion applies to the easy 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 6.25. The cut-wall criterion applies to wallspaces built from trees of wallspace with cut-wall criterion on vertex space and connected intersections of walls with vertex spaces. See Section 7.1.

7. Cubulating Malnormal Graphs of Cubulated Groups

7.1. Examples of systems of walls in graphs of groups.

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

As discussed in Section 6.0.4, the motivating case is where $\Gamma - N_r(H)$ has two (deep) components.

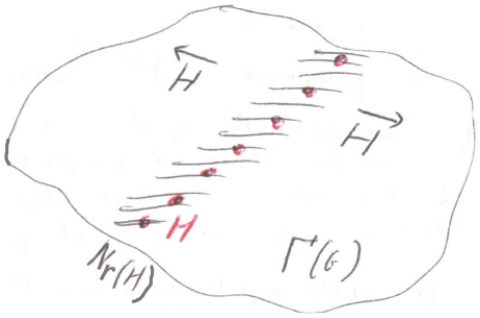

FIGURE 77. An H-wall.

Let $G \subset G'$. An H-wall in G *extends* to an H' wall in G' if $H = H' \cap G$ and $\overleftarrow{H} \subset N_r(\overleftarrow{H'})$ and $\overrightarrow{H} \subset N_r(\overrightarrow{H'})$ for some r, so $\overleftarrow{H}$ is coarsely contained in $\overleftarrow{H'}$ etc.

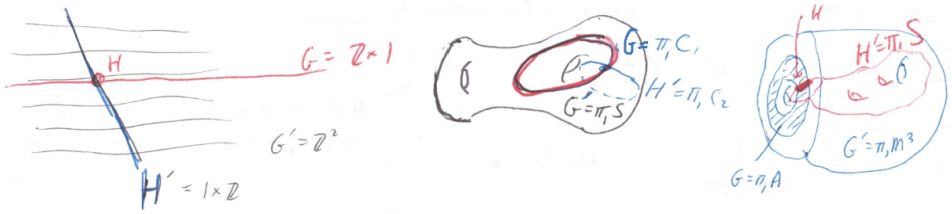

FIGURE 78. Some extensions of H-walls.

Theorem 7.1. *Let $H \subset G$ be a quasiconvex subgroup of G with $G = \pi_1 X$ where X is compact special and hyperbolic. Let A be a quasiconvex subgroup of H. Then any A-wall in H extends to a B-wall of G.*

Proof. Let $Y_H \to X$ be a compact core for H. Let $Y_A \to Y_H$ be a compact core for A, such that, moreover, the two sides of the A-wall coarsely lie in distinct halfspaces $\widetilde{Y}_H - \widetilde{Y}_A$ (partitioned accordingly).

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

Consider $\mathsf{C}(\widehat{Y}_H \to X)$ and the canonical retraction $\mathsf{C}(\widehat{Y}_H \to X) \to \widehat{Y}_H$. The preimage of the locally convex subcomplex $Y_A \subset \widehat{Y}_H$ is a locally convex subcomplex $Y_B \subset \mathsf{C}(\widehat{Y}_H \to X)$, and $B = \pi_1 Y_B$. The B-halfspaces are the preimages of the A-halfspaces under the retraction $\pi_1\big(\mathsf{C}(\widehat{Y}_H \to X)\big) \to \pi_1 \widehat{Y}_H$. □

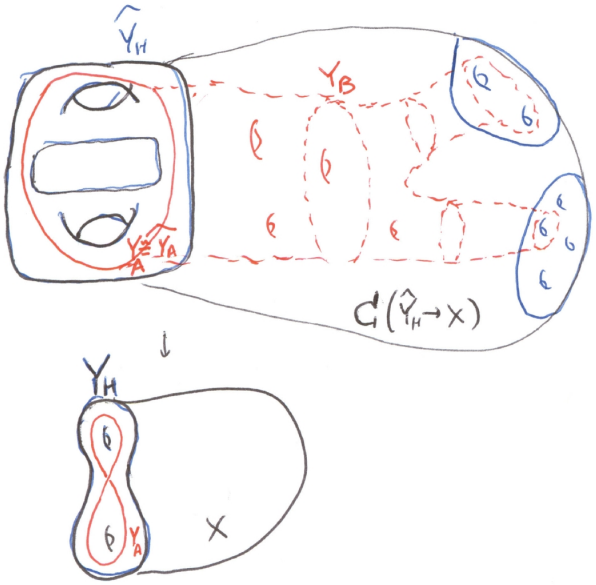

FIGURE 79. Using Canonical Completion and Retraction to Extend an H-wall.

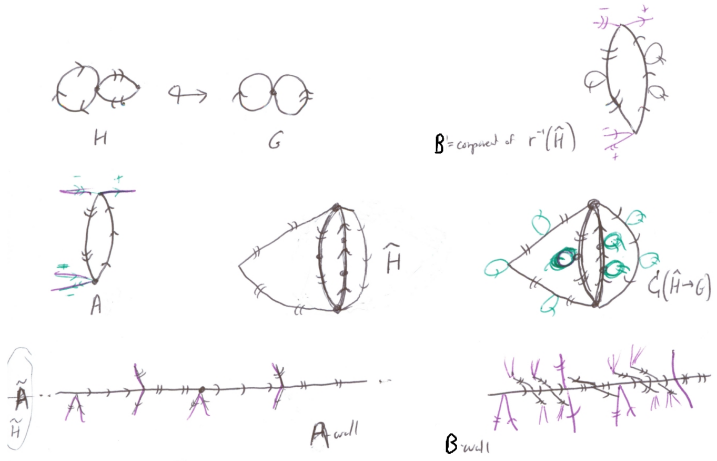

FIGURE 80. Wall extension using canonical completion and retraction.

7.3. **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 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 vertex space consisting of a surface and edge space consisting of a cylinder and an immersed wall consisting of a circle in the surface, which needs to extend into an immersed wall in the cylinder. This is illustrated in Figure 81.

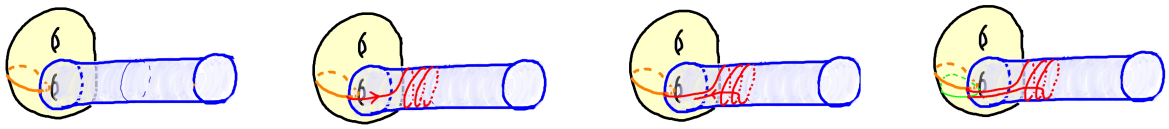

FIGURE 81. Turn in a cylindrical edge space

An H'-wall of A induces an H-wall of C where $H = C \cap H'$. Let us denote this by $X_H \to X_C$. (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 essentially ignore in our applications.) We use separability of H in C to choose a large girth cover (relative to H') of X_C then build the turn by cute-and-paste – combining the H-wall with the C-wall down the middle of $X_C \times [-1, 1]$.

Specifically, we consider the cover $\widehat{X}_C \times [-1, 1] \to X_C \times [-1, 1]$ and the wall $\widehat{X}_C \times \{0\}$ that resides inside. We now cut-and-paste $X_H \times [-1, 0]$ with $\widehat{X}_C \times \{0\}$ as illustrated in Figure 82.

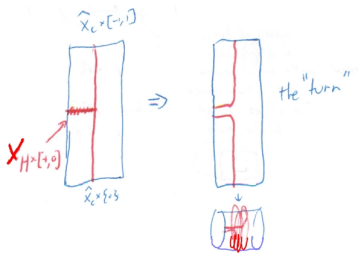

FIGURE 82. Cutting and pasting to form the turn.

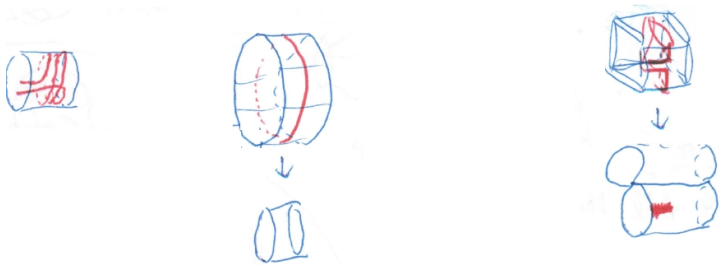

FIGURE 83. Eventually put a more legitimate high girth cover (cylinder and graph x interval) here.

7.4. **Cubulating malnormal amalgams.**

Theorem 7.2. *Let $G = A *_C B$ or $A*_{C^t=C'}$. 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.*
(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.*

Remark 7.3. There is a generalization of Theorem 7.2 which holds, for instance, when G is hyperbolic relative to virtually abelian subgroups, and C is relatively quasiconvex and almost malnormal.

Two instances where Theorem 7.2 takes a simplified form are as follows:

Example 7.4 ($\mathbb{Z}$ edge group)**.** Let $G = A *_{\mathbb{Z}} B$ or $A*_{\mathbb{Z}^t=\mathbb{Z}'}$ 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 7.5 (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 7.2. The plan for proving Theorem 7.2 is to produce sufficiently many quasiconvex codimension-1 subgroups, and then obtain a proper action on the dual cube complex. As discussed in Section 6.2, "sufficiently many" leads to a proper action and quasiconvexity leads to cocompactness.

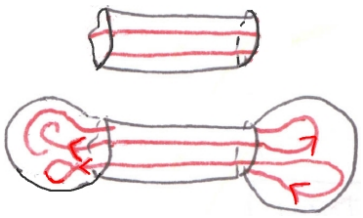

FIGURE 84. Start with a sufficient set of immersed walls in X_C, and extend these into immersed walls in X_A and X_B

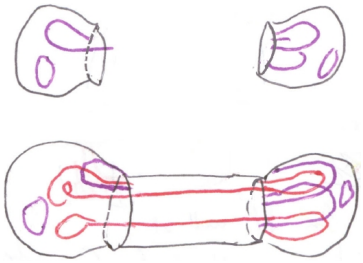

FIGURE 85. Consider the original sufficient immersed walls (arising from the cubulations) of X_A and X_B and add those as well.

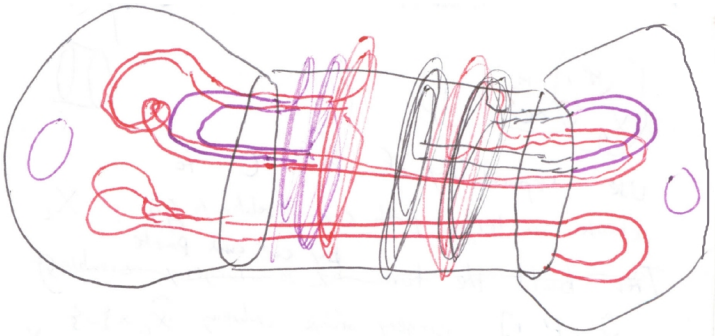

FIGURE 86. Double all these immersed walls, and turn those immersed walls that enter $X_C \times I$ on each side. (Some walls have not yet been turned.)

The malnormality of C and a choice of sufficiently "large" turns ensures that the new immersed walls correspond to quasi-isometrically subgroups and lift to genuine embedded walls in $\widetilde{X}$.

FIGURE 87. Add a vertical wall cutting through each $X_C \times I$

These new walls have the property that they intersect each collared vertex space $\widetilde{X}_A^+$ and $\widetilde{X}_B^+$ 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. □

8. Cubical Small Cancellation Theory

8.1. **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 $\phi_i : Y_i \to X$.

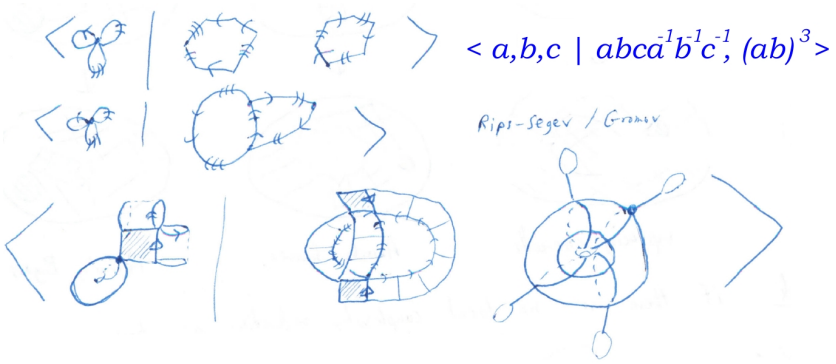

FIGURE 88. Classical, Rips-Segev (popularized by Gromov), and Cubical small-cancellation presentations.

The *group* G defined 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 $X^* = X \sqcup \mathrm{Cone}(Y_i) / \{(y_i, 0) \sim \phi_i(y_i) : \forall y_i \in Y_i\}$.

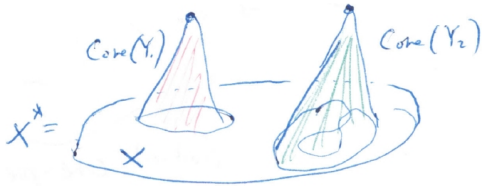

FIGURE 89. X^* is built from X by attaching cones over the relators

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 $\widetilde{X}^*$, and are mainly interested in its cubical part which is the cover of X associated to $\langle\langle \pi_1 Y_i \rangle\rangle$. Note that X plays the role of a bouquet of circles in a classical presentation, and the cubical part of $\widetilde{X}^*$ is a generalization of a Cayley graph.

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

The triangles belong in cyclic families each around a point mapping to the conepoint of some $\mathrm{Cone}(Y_i)$ and we combine each such family together to form a *cone-cell*.

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

FIGURE 90. A Disc diagram in X^* has squares mapping to X together with cone-cells consisting of a sequence of triangles around each cone point.

8.2. **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 known theorems are applications or variants of these statements (together with analogous statements for annular diagrams).

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

8.2.1. *Fundamental theorem of classical small-cancellation theory.*

Theorem 8.1. *Let X^* be a $C'(\frac{1}{6})$ small-cancellation presentation. A reduced disc diagram $D \to X^*$ is either:*

1. *a single 0-cell or 2-cell.*
2. *a "ladder".*
3. *or contains at least three shells and/or spurs.*

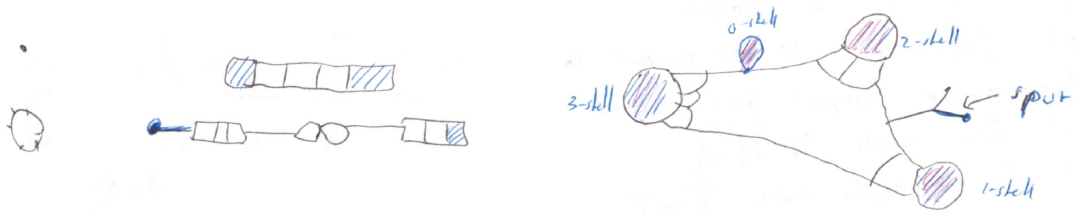

FIGURE 91. A single 0-cell or 2-cell, a possibly singular ladder, and a diagram with three or more spurs and or shells.

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.

8.2.2. *Fundamental theorem of cubical small-cancellation theory.* We now let X^* be a cubical presentation.

Theorem 8.2. *Let X^* be a $C'(\frac{1}{12})$ small-cancellation presentation. A reduced disc diagram $D \to X^*$ is either:*

1. *A single 0-cell or 2-cell.*
2. *A "ladder".*
3. *Contains at least three 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.

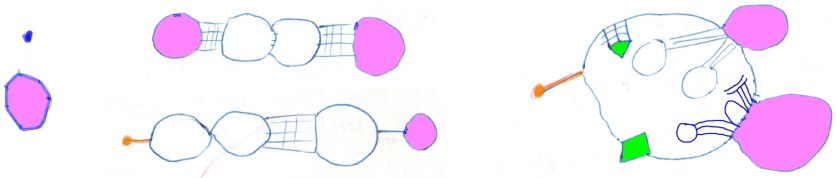

FIGURE 92. A single 0-cell or single cone-cell, a possibly singular ladder, and a diagram with three or more spurs, cornsquares, and/or shells.

8.3. **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).

The *curvature at a 0-cell* v of X is defined by:

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

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

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

where we use the notation $|f| = |\partial_p f|$ to denote the length of the boundary of f.

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

$$\kappa(v) = \pi(2 - \deg(v) - \sum c \in \mathrm{Corners}(v)\, \mathrm{def}(\sphericalangle(c))$$

$$\kappa(f) = 2\pi - \sum \mathrm{def}(\sphericalangle(c))_{c\in \mathrm{Corners}(f)} \sphericalangle(c)$$

Our cases of greatest interest are when X is a disc 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 93

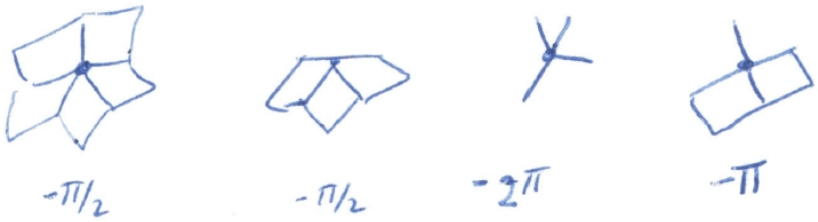

FIGURE 93. Some curvatures of 0-cells

Theorem 8.3. *Let X be a compact angled 2-complex then:*

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

Proof. This is just a way of distributing $\chi = V - E + F$ as "curvature" among the 0-cells and 2-cells. □

8.4. **Greendlinger's Lemma and the ladder theorem for classical $C(6)$ small-cancellation theory.** A *piece* in a disc diagram D is a path $P \to D$ that arises in two ways as a subpath of boundary path of a 2-cell. We normally ignore pieces that are *trivial* in the sense that they have length 0.

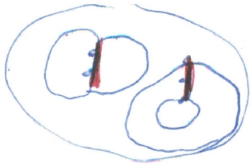

FIGURE 94. Pieces.

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

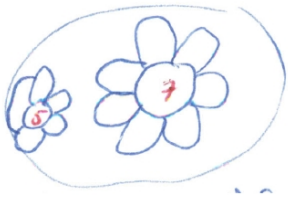

FIGURE 95. 5 and 7 pieces around a 2-cell.

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

FIGURE 96. A spur, 0-shell, 1-shell, and 2-shell.

Note that for an i-shell we have i *nontrivial* pieces in mind, or better said, the number i should indicate the minimal number of concatenated pieces necessary to express S, in which case a 0-shell is just a degenerate 1-shell (and this is consistent with the values they are assigned below).

Theorem 8.4 ("Greendlinger's Lemma")**.** *Let D be a $C(6)$ disc 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.

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

$\kappa(f) \leq 0$ 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 p-external $\kappa(f) \leq 2\pi - 2p\frac{\pi}{2}$.

FIGURE 97. Assigning angles to corners of 2-cells.

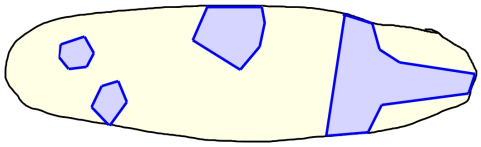

FIGURE 98. Internal, external, and multiply external 2-cells.

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).

Note that 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 8.3 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 8.5 (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 ≥ 3 to begin with. Thus the original D was a ladder. $\square$

FIGURE 99. 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 or a negatively curved feature as in the right two diagrams.

8.5. **Reduced diagrams.** For parallelism we will use the notation X^* for a 2-complex (but this is not quite in agreement with our usage for a cubical presentation, which subdivides

the 2-cells.) In classical small-cancellation theory, a diagram $D \to X^*$ 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:

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

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 would be some tails hanging off of the newly created diagram.)

We note that the $C(p)$ condition can be defined implicitly by requiring that no reduced diagram $D \to X^*$ can have a 2-cell whose boundary path is the concatenation of fewer than p pieces, and the $C'(\alpha)$ condition says that each piece P between 2-cells R_1, R_2 in a reduced diagram has $|P| < \alpha |R_i|$.

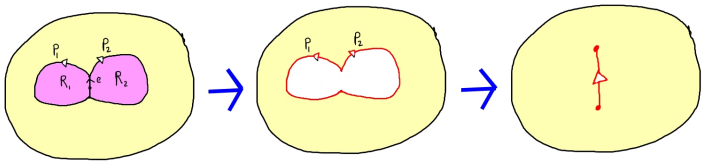

FIGURE 100. Removing a cancelable pair in classical case.

We now describe some generalizations of this reduction procedure, as well as some further reductions that arise for cubical presentations.

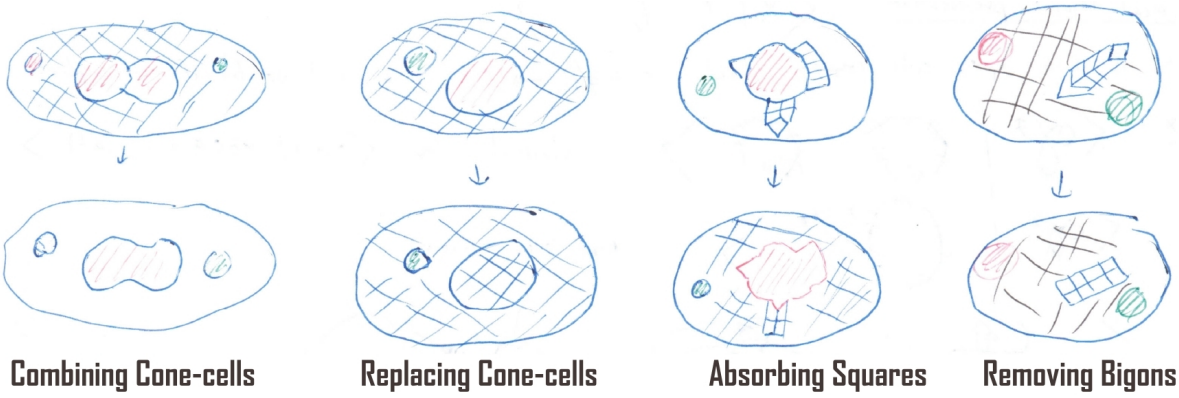

FIGURE 101. Reductions in diagrams $D \to X^*$.

Combining cone-cells: involves two cone-cells C_1, C_2 that map to the same Cone(Y) and that meet along a maximal open arc E such that $\partial_p C_i = eP_i$ and such that the concatenation $P_1 P_2^{-1}$ projects to a closed path in Cone(Y). We replace $C_1 \cup E \cup C_2$ by a new cone-cell C mapping to Cone(Y) and with $\partial_p(C) = P_1 P_2^{-1}$.

Replacing cone-cells: involves a cone-cell C mapping to some Cone(Y) such that $\partial_p C$ is already null-homotopic in Y, and so we can replace C by a square disc diagram $D_C \to Y$ 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 $\partial C \to Y$ extends to a map $\partial C \cup s \to Y$. 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 ad 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* $\mathsf{Comp}(D)$ and so reduced diagrams with a given boundary path always exist, since a minimal complexity diagram with a given boundary path must be reduced.

8.6. **Pieces and Reduced diagrams.** The *pieces* in a reduced diagram $D \to X^*$ are subpaths of the boundary path of a cone-piece $P \to \partial_p C$ whose outgoing dual curves (we focus on the *rectangles* carrying them) "fellow-travel" within the square part of D until they end together on the boundary of another cone-cell or on some rectangle.

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 such that C maps to $\mathrm{Cone}(Y_i)$. Here $|P|$ denotes the distance between the endpoints of $\widetilde{P}$ in $\widetilde{X}$, and $\|Y_i\|$ denotes the length of the shortest essential combinatorial path in Y_i.

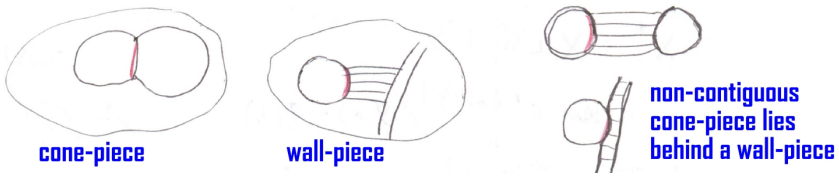

FIGURE 102. Cone-pieces and wall-pieces.

8.7. **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 8.6. *Given* $\langle a, b, \cdots \mid W_1, \ldots, W_r\rangle$ *with no* $W_i^p \nsim W_j^q$ *(for* $p, q \neq 0$*), there exists* M *such that for* $n_i \geq Mn$ *the following presentation is* $C'(\frac{1}{n})$.

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

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 $M > nD$. □

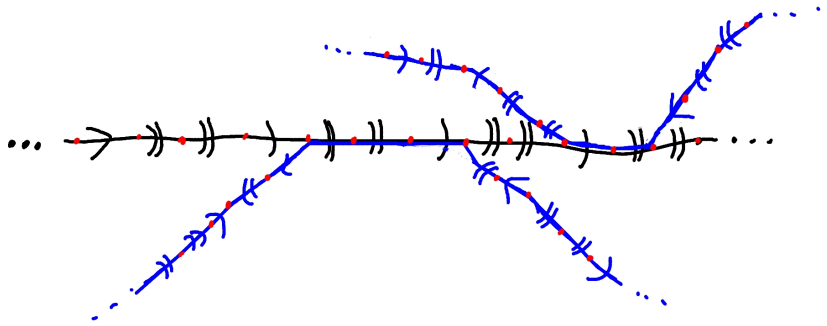

FIGURE 103. Some of the finitely many types of pieces illustrated in the universal cover of the bouquet of circles for $\langle a, b \mid (abb)^\infty, (aba^{-1}b^{-1}b^{-1})^\infty\rangle$.

Theorem 8.7. *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 quasiconvex subgroups. There exist compact based local isometries $Y_i \to X$ with $H_i = \pi_1 Y_i$ and finite subsets $S_i \subset H_i - \{1\}$ such that for any regular covers $\widehat{Y}_i \to Y_i$ 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 3.18, let $\widetilde{Y}_i$ be an H_i-cocompact superconvex core. Wall-pieces in $\widetilde{Y}_i$ have length $\leq R$ – this uses cocompactness of $\widetilde{Y}_i$ together with superconvexity. Note that non-contiguous cone-pieces and non-contiguous wall-pieces are dominated by contiguous wall-pieces. (Contiguous) cone-pieces between translates of $\widetilde{Y}_i$ and $\widetilde{Y}_j$ are bounded by L – this uses malnormality of $\{H_1, \dots, H_r\}$ together with H_i-cocompactness of each $\widetilde{Y}_i$.

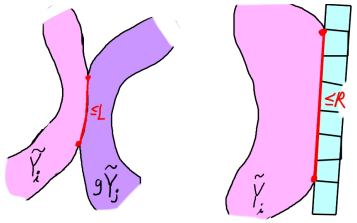

FIGURE 104. Contiguous cone-pieces and wall-pieces are bounded in $\widetilde{X}$ by respectively using cocompactness and malnormality, and using cocompactness and superconvexity.

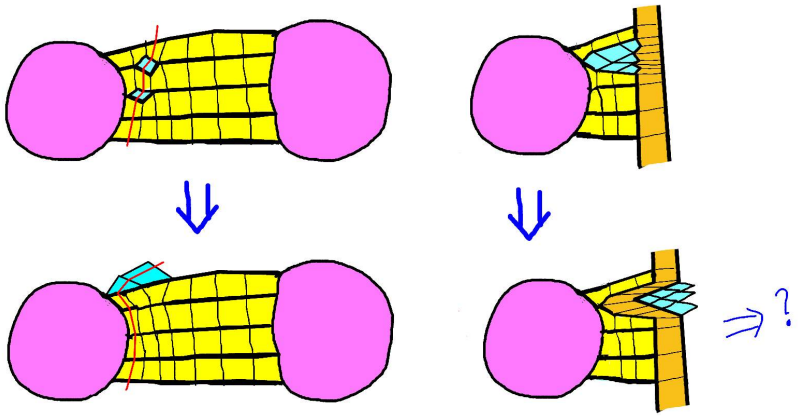

FIGURE 105. Noncontiguous pieces are bounded by contiguous wall-pieces, since squares within the shards can pushed out of the way.

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

8.8. **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.

8.8.1. *Admitting rectangles.* We begin by declaring a linear ordering of the cone-cells $C_1, C_2, \ldots$ 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.

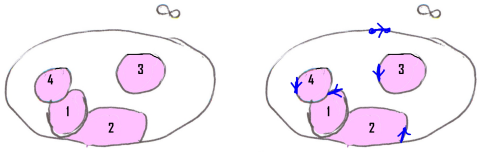

FIGURE 106. Ordering the 1-cells in boundary paths of cone-cells.

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

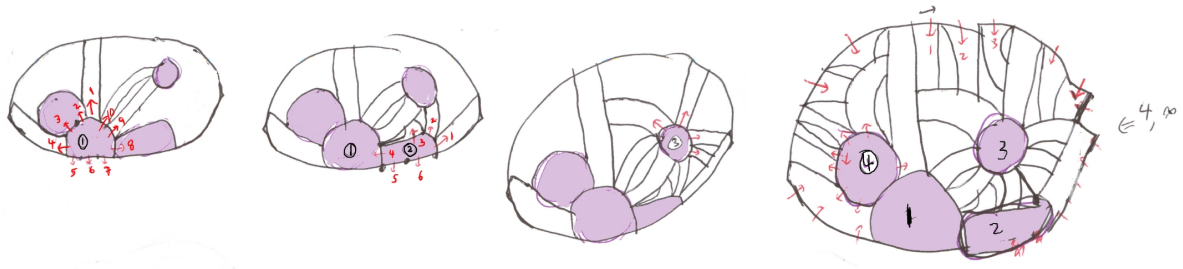

FIGURE 107. Admitting rectangles.

Admitted rectangles are of the form $[-1,1] \times [0,n]$ with degenerate case of $n = 0$, a have *internal part* of form $(-1,1) \times [0,n]$.

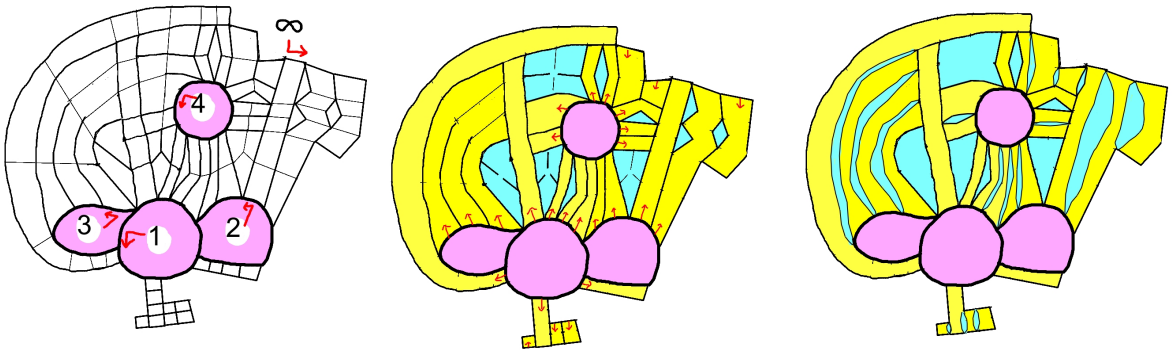

FIGURE 108. The diagram D is on the left. We have indicated the ordering of cone-cells and cyclic ordering of attaching maps. The admitted rectangles within D are indicated in the middle, and the rectified diagram $\bar{D}$ is on the right. (We have not highlighted 0-cell shards).

8.8.2. *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 suitable conditions (in fact, under the small-cancellation conditions we pursue) each shard is a simply-connected square disc diagram.

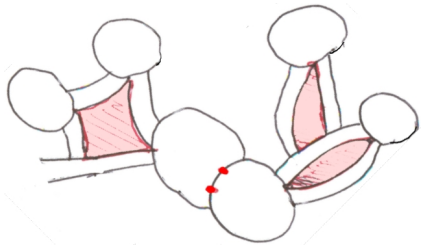

FIGURE 109. Five shards. Note that that two of these shards are just 0-cells.

8.8.3. *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*.

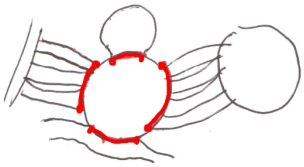

FIGURE 110. Four pieces.

8.8.4. *Assigning Angles.* When no cone-cell C has $\partial_p C$ the concatenation of < 24 pieces[1], then there is a nice way to assign angles at corners of cone-cells and shards.

All admitted rectangles are assigned ordinary 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 111:

- 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.

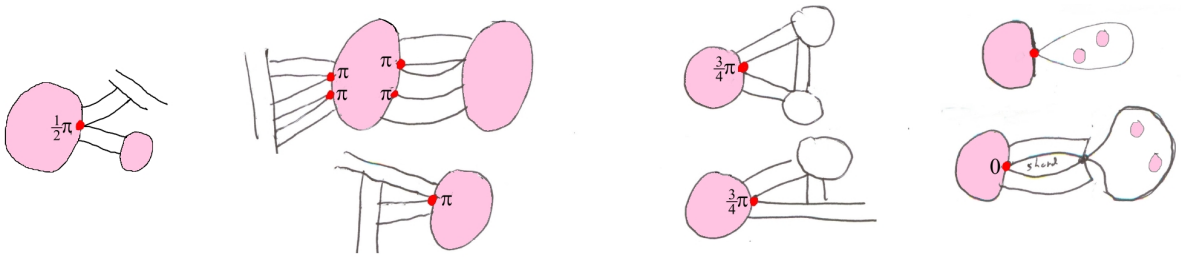

FIGURE 111. Four pieces.

The corners of shards in D are assigned the "obvious" angles so that the vertices v at their corners have $\kappa(v) = 0$. The shards are (almost always) automatically nonpositively curved – i.e. $\sum \sphericalangle - (|f| - 2)\pi \leq 0$. 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 112.

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

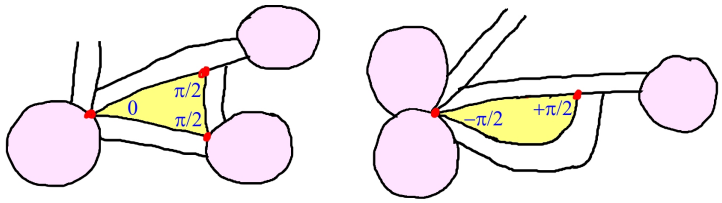

FIGURE 112. Choose angles so that $\kappa(v) = 0$ for each vertex v, and the shard will be automatically nonpositively curved.

8.8.5. *The positively curved shards.* The exceptional situations of positively curved shards can arise from cornsquares on cone-cells, or "monkey-tails" within the diagram consisting of a rectangle that terminates on itself bounding a shard. However these cannot arise for a reduced diagram. The former cannot arise because there would be a square absorption, the latter cannot arise because there would be a square bigon to remove.

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.

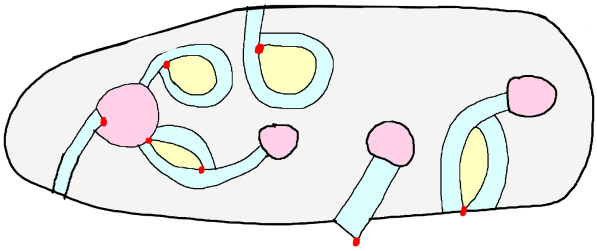

FIGURE 113. The exceptional situations of positively curved shards arise from cornsquares on cone-cells and also "monkey-tails" within the diagram. Neither of these arise for a reduced diagram. Cornsquares terminating on ∂C_∞ correspond to a $+\frac{\pi}{2}$ curvature.

8.8.6. *Curvatures of cone-cells with* $< \frac{1}{24}$ *small-cancellation.* We now examine the curvature of a cone-cell in $\bar{D}$. We first observe that 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 8.8. *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 on the rectified diagram $\bar{D}$ has $\kappa(C) \leq 0$. (Moreover 25 or more pieces implies $\kappa(C) < 0$.)*

Proof.

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

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 $\sum \mathrm{def}(\triangleleft) \geq 8\frac{\pi}{4} = 2\pi$ as needed. We note that it is possible for their to be three consecutive pieces with no nonzero defects between them as on the left in Figure 114.

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 114. □

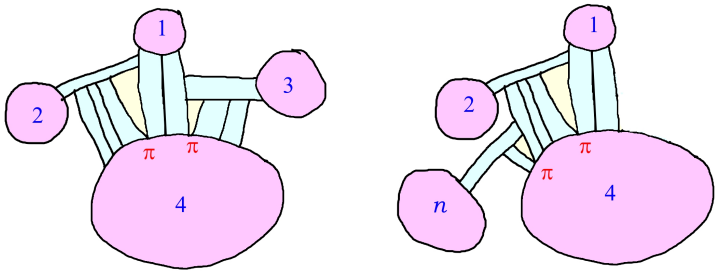

FIGURE 114. While it is possible to have three consecutive pieces with zero transitional defects as on the left, it is impossible to have more than this because, as on the right, the rules for building rectified diagrams would be violated when one considers the ordering of cone-cells. Both $n > 4$ and $n < 4$ lead to contradictions.

9. Walls in Cubical Small-Cancellation Theory

9.1. **Walls in classical $C'(\frac{1}{6})$ small-cancellation.** 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 $\widetilde{X^*}$ as follows: Firstly, by possibly subdividing all 1-cells, we assume tat all attaching maps of 2-cells have even length. The walls of $\widetilde{X}^*$ are graphs (see Figure 115) that intersect 1-cells and 2-cells in midcells (see Figure 65) as described in Example 6.5.

Figure 115. The carrier of a wall in a $C(\frac{1}{6})$-complex $\widetilde{X}^*$.

Walls are 2-sided, embedded, multi-trees, with convex carriers. Being a multi-tree and not self-crossing are consequences of Greendlinger's lemma, and the ladder theorem implies the convexity of the carriers. A minimal area diagram D for a path corresponding to a self-crossing could not have enough shells. And 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 (since there is no room for shells on γ or σ, and then a contradiction arises by consider the relationship between a 1-shell on the top or bottom of L with the ladder L' that resides in $N(W)$ and contains σ.

9.2. **Wallspace Cones.** The key ingredients that allow us to define walls in the classical case are:

(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 $Y \to X$ is a *convex wallspace* as follows:

(1) Each hyperplane in Y is 2-sided and embedded.
(2) The collection of $\{H_i\}$ of hyperplanes are partitioned into subcollections called *walls*
(3) Two hyperplanes in the same wall are disjoint from each other.
(4) If $P \to Y$ 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 disc diagram $D \to X$ between P and a path $P' \to N(W)$.

9.3. **Producing Wallspace Cones.** In the classical case, we could immediately turn all relators into wallspace by subdividing the 1-skeleton. It is much harder to do this naturally in a higher dimensional setting. Another way to turn relators into cones is to take 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 $Y \to X$ where hyperplanes of X embed. We now

Add pictures explaining these contradictions here...

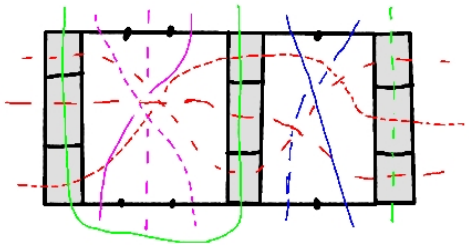

FIGURE 116. The hyperplanes of the cube complex Y are partitioned into walls.

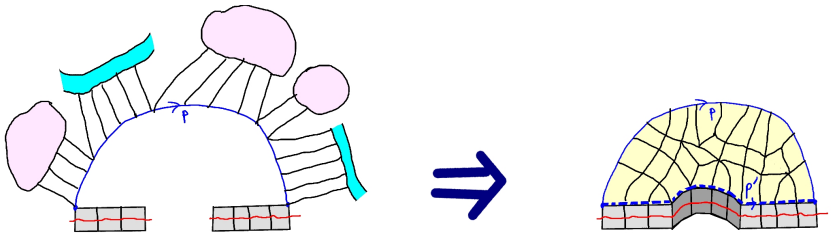

FIGURE 117. If $P \to Y$ is the concatenation of too few pieces, then P is square-homotopic to a path $P' \to N(W)$.

consider the covering space $\ddot{Y} \to Y$ induced by the homomorphism $\#_{\mathbb{W}} : \pi_1 Y \to \mathbb{Z}_2^{\mathbb{W}}$ which counts the number of times a path travels through an edge dual each of the respective classes.

The space $\ddot{Y}$ has a natural wallspace structure: Each wall is the preimage of all hyperplanes in $\mathbb{W}$. For instance, when Y is a circle and $\mathbb{W}$ is the discrete partition, $\ddot{Y}$ is just the usual $\mathbb{Z}_2$ cover. We refer the reader to Figure 118.

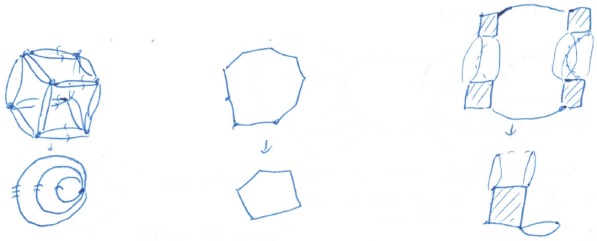

FIGURE 118. Given a cone $Y \to X$, we can obtain a wallspace cone $\ddot{Y}$ that is the finite cover associated to a homomorphism $\pi_1 Y \to \mathbb{Z}_2^{\mathbb{W}}$.

9.4. **Walls in $\widetilde{X}^*$.** Suppose each Y_i in X^* has a wallspace structure. A *wall* W in $\widetilde{X}^*$ is a collection of disjoint hyperplanes such that the intersection with each relator $Y_i \subset X^*$ is either $\emptyset$ or a wall V of Y_i. If we think of $\widetilde{X}^*$ as the actual universal cover, then we imagine the wall so that $W \cap \mathrm{Cone}(Y_i) = \mathrm{Cone}(V)$.

Theorem 9.1. *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 $\widetilde{X}^*$ lies in a unique wall of $\widetilde{X}^*$.*

Proof. This follows from the fundamental theorem of cubical small-cancellation theory. □

9.5. **Quasiconvexity of walls in $\widetilde{X}^*$.** For a wall W of $\widetilde{X}^*$ its *carrier* $N(W)$ is the union of all carriers of the hyperplanes of W and all cones intersecting W. Its *thickened carrier* $T(W)$

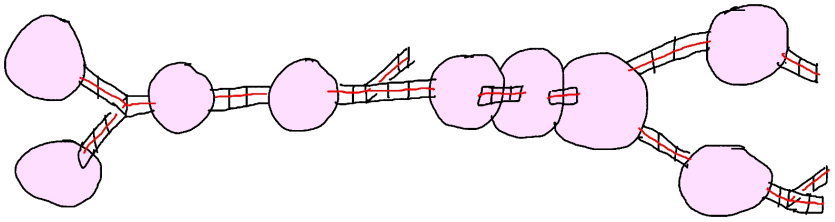

FIGURE 119. A wall in $\widetilde{X}^*$ is a collection of hyperplanes that intersects each Y_i in a wall of Y_i (which generalizes a midcell or midcube).

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 crosses both.

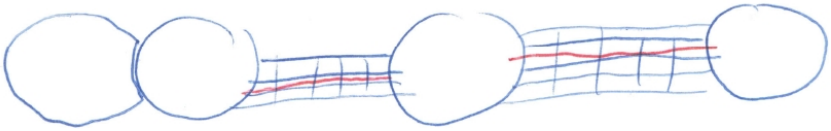

FIGURE 120. The thickened carrier $T(W)$ is the union of $N(W)$ together with minimal length ladders that start and end on consecutive cones.

The thickened carrier $T(W)$ interpolates between $N(W)$ and $\widetilde{X}^*$. 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 $\widetilde{X}^*$. We are thus able to obtain the following:

Corollary 9.2. *When X and the cones Y_i are compact the inclusion* Stabilizer$(W) \subset \pi_1 X^*$ *is a quasi-isometric embedding.*

Theorem 9.3 (Thickened Carrier Isometrically Embeds)**.** *Let $T(W)$ be the thickened carrier of a wall W in $\widetilde{X}^*$. Then $T(W) \subset \widetilde{X}^*$ 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 disc diagram that has γ on one side and has a path $\sigma \to T(W)$ 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 the the same wall of the cone. Note that the cladder can bifurcate at cone-cells. We refer to Figure 121.

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 8.2 and the minimality of D, since if a cladder travels back to cross σ again (see Figure 122), we would either find that there was a way to thicken $T(W)$ within D, or we would obtain a reduced diagram with only positively curved features that is not a ladder. □

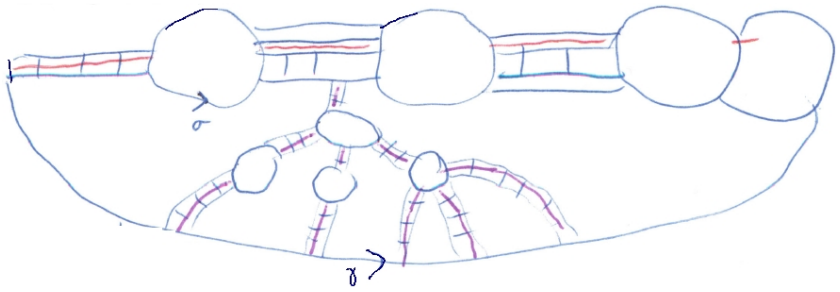

FIGURE 121. Let D be a minimal complexity diagram between $T(W)$ and a geodesic γ. Each cladder in D starting on an edge of σ ends entirely on γ. Thus $|\gamma| \geq |\sigma|$ and $T(W)$ isometrically embeds.

FIGURE 122. The cladder cannot end at another edge of σ or it would lie in $T(W)$ or violate the ladder theorem.

10. Annular Diagrams

10.1. **Classification of Flat Annuli.** An *annular diagram* is a compact 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 disc 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 $C(\alpha)$ conditions are defined as for disc diagrams D.

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

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

Theorem 10.1. *: Let A be a $C(7)$ annular diagram. Then either A has a spur or an i-shell with $i \leq 3$ or A looks like one of the diagrams in Figure 123:*

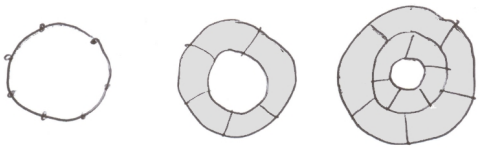

FIGURE 123. A $C(7)$ annulus with no feature of positive curvature must have width 0,1, or 2 as above.

Remark 10.2. We note that the $C(7)$ condition implies that if there is an internal 2-cell then there are 2-cells with negative curvature and hence there would have to be a positively curved feature. 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 124.

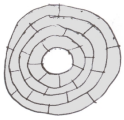

FIGURE 124. A flat $C(6)$ annulus can have internal 2-cells and thus arbitrarily many rings.

Theorem 10.3. *Let A be a annular diagram satisfying the $C'(\frac{1}{24})$ cubical small-cancellation condition. Then either A has a spur or shell or else: Either A is a square annulus, or A is an annuladder, or A is a square annulus bicollared by a pair of annuladders.*

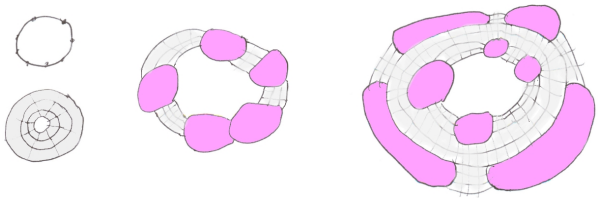

FIGURE 125. A flat annulus in a cubical $C'(\frac{1}{24})$ complex either has a positively curved feature, or else is either a square annulus, or a thin annuladder, or is thick and bicollared by annuladders.

10.2. **The Doubly Collared Annulus Theorem.** We now describe a result which is especially catered to understand conjugate elements that stabilize two walls in $\widetilde{X}^*$. It will allow us to verify almost malnormality of wall stabilizers under certain conditions, which plays a role in the proof Theorem 11.2.

Theorem 10.4 (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 $A \to X^*$ be an annular diagram with boundary cycles α_1, α_2 that are essential. Suppose $\widetilde{A} \to \widetilde{X}^*$ has $\tilde{\alpha}_1, \tilde{\alpha}_2$ lifting to wall carriers $N_i = N(W_i)$.

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

And 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 $\widetilde{L}_i \subset N_i$.

And B is in the same class *as A in the sense that $\partial\widetilde{B}$ lies on N_1, N_2 and β_i is conjugate to α_i in* 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).*

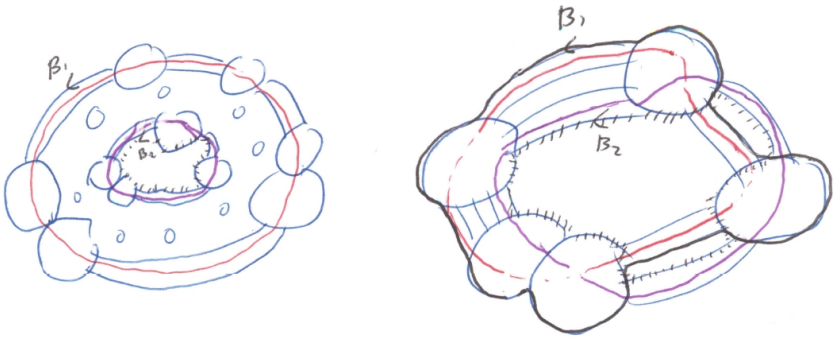

FIGURE 126. B is either thick and bicollared by two annuladders L_1, L_2, or is a thin and is a single annuladder with L_1, L_2 wandering within B.

11. The Malnormal Special Quotient Theorem

G has an *almost malnormal quasiconvex hierarchy [terminating in virtually compact special hyperbolic groups]* if it can be built from finite groups [virtually compact special hyperbolic groups] by HNN extensions and AFP along almost malnormal quasiconvex subgroups.

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

Theorem 11.1 (Almost malnormal quasiconvex hierarchy)**.** *If G has an almost malnormal quasiconvex 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 7.2 and Theorem 5.1. □

The following central result is the main tool used to prove our main theorem:

Theorem 11.2 (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 $\bar{G} = G/\langle\langle H_i' \rangle\rangle$ is virtually compact special and hyperbolic.*

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 11.1.

Let $\widetilde{X}$ be a CAT(0) cube complex upon which G acts properly and cocompactly. For each H_i let $\widetilde{Y}_i$ be an H_i-cocompact superconvex subcomplex. Let R, L be upper bounds on the diameters of contiguous rectangle-pieces and cone-pieces between the translates of the $\widetilde{Y}_i$ in $\widetilde{X}$.

We choose J to be a normal finite-index subgroup of G so that:

(1) $J = \pi_1 X$ where $X = J\backslash\widetilde{X}$ is a compact special cube complex.
(2) Moreover, using separability of the hyperplanes, we can assume that each such hyperplane subgroup is malnormal,
(3) and moreover, the hyperplanes have very high injectivity radius relative to R, L – say $\geq 15 \max R, L$ – so "short paths" that start and end on a hyperplane carrier are square-homotopic into that carrier.
(4) Moreover, we assume that $Y_i \to X$ embeds for each i, where $Y_i = (H_i \cap J)\backslash\widetilde{Y}_i$.

We now consider the homomorphism $\#_{\mathbb{W}} : \pi_1 X \to \mathbb{Z}_2^{\mathbb{W}}$ that counts the number of times a path passes through the various hyperplanes modulo 2. It induces a covering space $\ddot{X} \to X$

and we let $\ddot{Y}_i$ be the induced covers. Finally we let $H_i' = \pi_1 \ddot{Y}_i$, and claim that $\bar{G} = G/\langle\langle H_i' \rangle\rangle$ is virtually special.

To see this first check that $\bar{J}$ has a finite index subgroup K with the following cubical presentation:

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

Each hyperplane of X induces a splitting of K along an almost malnormal quasiconvex subgroup (almost malnormality by analysis of annuli and small-cancellation etc. and quasiconvexity holds by Theorem 9.3.) Taken in some order these splittings induce an almost malnormal quasiconvex hierarchy for K. □

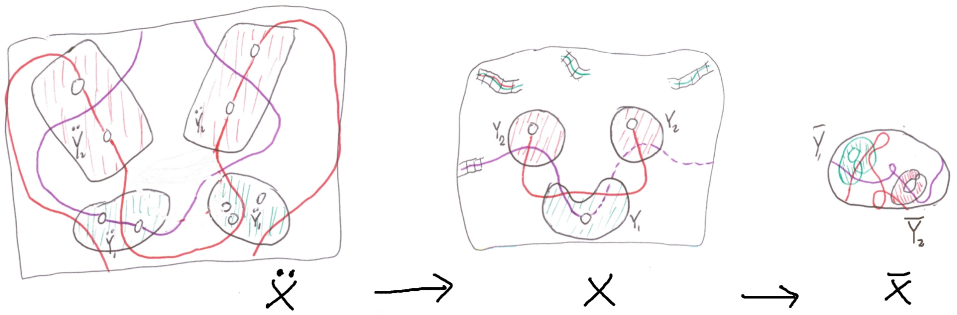

FIGURE 127. We pretend that G is torsion-free and let $\bar{X} = G\backslash\widetilde{X}$ and $\bar{Y}_i = H_i\backslash\widetilde{Y}_i$. We have partially illustrated $\ddot{X}_i \to \ddot{X} \to \bar{X}$.

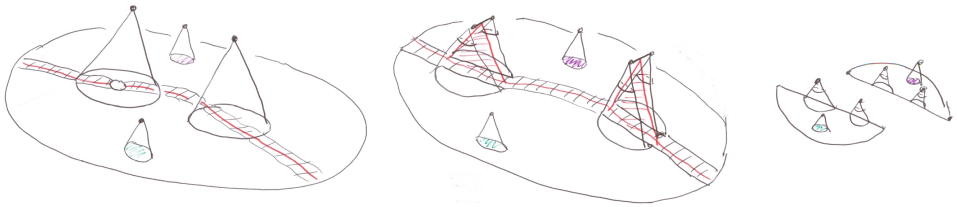

FIGURE 128. A hyperplane of X provides a splitting of $\ddot{X}^*$. We inflate conepoints of relevant $\ddot{Y}_i$ (add dummy squares) and then cut to obtain one or two smaller cubical relative presentations with the same properties that $\ddot{X}$ had.

11.1. **Case study: $F_2/\langle\langle W_1^{n_1}, \ldots, W_r^{n_r} \rangle\rangle$.** We now describe the special case of Theorem 11.2 where G is free and each H_i is cyclic. This argument is extracted from [Wisc] where it arose naturally. We warn the reader that we have been a bit cavalier about (identifying and ignoring) multiple 2-cells in covers – these arise from proper powers.

Theorem 11.3. *Consider the presentation $\langle a, b \mid W_1, \ldots, W_r \rangle$ and assume that $W_i^p \nsim W_j^q$ for $i \neq j$ unless $p, q = 0$. There exists n_i which can be chosen arbitrarily large such that $\langle a, b \mid W_1^{n_1}, \ldots, W_r^{n_r} \rangle$ is virtually special.*

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

Proof. We shift immediately to a geometric viewpoint, of $\langle \bar{X} \mid \bar{Y}_1, \ldots, \bar{Y}_r \rangle$ 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 11.2.) 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 $\breve{X} \to X$ such that each elevation $\breve{Y}_i \to Y_i$ embeds in $\breve{X}$. Note that there might be several distinct elevations of each such immersed circle.

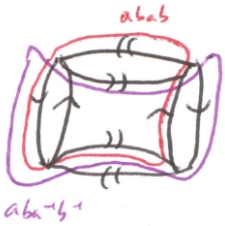

FIGURE 129. We first pass to a finite cover $\breve{X}$ where all elevated circles embed.

Let $\widetilde{X}$ denote the universal cover $\bar{X}$, and likewise $\widetilde{Y}_i$ are universal covers of $\bar{Y}_i$. Let M denote an upper bound on the lengths of overlaps between the various translates of the $\widetilde{Y}_i$ in $\widetilde{X}$. (M is just 2 in our case, but in general M can be a bit bigger than the maximal length of a piece in the presentation – consider e.g. $\langle a, b | (ab), ababababb \rangle$).

Step 2: We now pass to a further regular cover of $\bar{X}$ (factoring through $\breve{X}$) so that the resulting complex $\langle X \mid gY_k : g \in \mathrm{Aut}(X), 1 \leq i \leq r \rangle$ has the property that $|P| \leq \frac{1}{8}|\partial R|$ whenever P is a piece on R. This is simply a matter of using residual finiteness of $\pi_1 \bar{X}$ to choose a finite index normal subgroup contained in $\pi_1 \breve{X}$ that contains no elements of length $\leq 8M$. In our case a very simple cover illustrated in Figure 130 suffices.

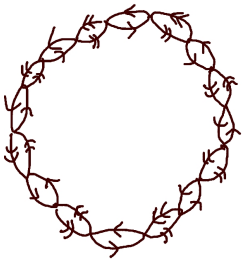

FIGURE 130. The degree 16 cover $X \to \bar{X}$ pictured above has the property that one obtains a ($\leq \frac{1}{8}$)-complex when 2-cells corresponding to the elevations $(ab)^8$ and $(aba^{-1}b^{-1})^4$ are added.

Step 3: At this point the group of $\langle \bar{X} \mid Y_1, \ldots, Y_r \rangle$ acts properly and cocompactly on a CAT(0) cube complex, and likewise, so does the finite index subgroup associated to $\langle X \mid gY_i : g \in \mathrm{Aut}(X), 1 \leq i \leq r \rangle$. Indeed, we can use the midcell wall-structure discussed in Example 6.5, and there are sufficiently many walls to get a proper action as described in [Wis04]. However we seek virtual specialness, and will obtain it by finding an almost malnormal quasiconvex hierarchy – but this comes at the expense of a further cover $\ddot{Y}_i \to Y_i$ 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 $\mathbb{W}$ is essentially the set of edges, and consider the homomorphism $\#_{\mathbb{W}} : \pi_1 X \to \mathbb{Z}_2^{\mathbb{W}}$ which counts the number of times a path traverses each edge modulo 2. Let $\ddot{X} \to X$ denote the associated finite cover, and let $\ddot{Y}_i \to Y_i$ denote the induced double covers (since $Y_i \hookrightarrow X$ is an embedded cycle, $\ddot{Y}_i$ is indeed a double cover.)

Our desired quotient group corresponds to the following cubical presentation:

$$\langle \bar{X} \mid \ddot{Y}_1, \ldots, \ddot{Y}_r \rangle \tag{3}$$

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

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

The cover $\ddot{X} \to \bar{X}$ induces a covering space of the standard 2-complex of Presentation (3) associated to the following cubical presentation:

$$\langle \ddot{X} \mid g\ddot{Y}_i : g \in \mathrm{Aut}(\ddot{X}), 1 \leq i \leq r \rangle. \tag{4}$$

Let $(\ddot{X})^*$ denote the 2-complex associated to Presentation (4).

Each edge e of X determines a track T_e in the 2-complex associated to the cubical presentation:

$$\langle X \mid g\ddot{Y}_i : g \in \mathrm{Aut}(X), 1 \leq i \leq r \rangle \tag{5}$$

The vertices of the track are the points at the centers of the preimages e. The edges of the track are midcells of the 2-cells corresponding to the $\ddot{Y}_i$, which we emphasize are attached by double covers of the simple cycles Y_i. See Figure 131.

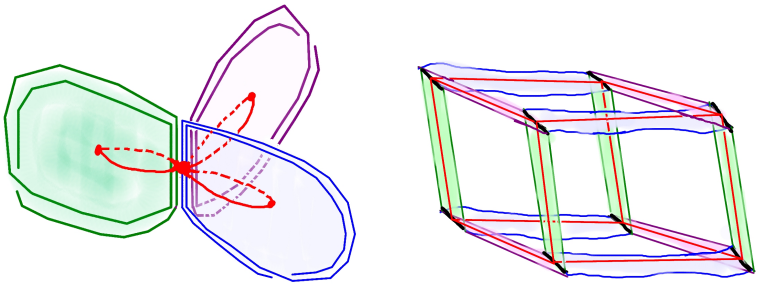

FIGURE 131. We give a heuristic picture of T_e in our example on the left, and the track $\ddot{T}_e$ on the right. (The 2-cells would actually be folded with each other a bit more in some places since there are some length 2 pieces.) Note that T_e has exactly three edges since an edge e lies on the 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 $\ddot{T}_e$ of the track T_e is a 3-valent graph that is 2-sided (without involutions) cutting through $(\ddot{X})^*$ and providing a splitting of $\pi_1(\ddot{X})^*$.)

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

Remark 11.4. Note that 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 5.2 and it is not necessary to invoke Theorem 5.1. I was expecting similar reasoning to hold in the higher dimensional case, but had difficulties verifying that the action on the cubulation was proper, and thus resorted to the alternate (and less direct) cubulation afforded by Theorem 7.2 in Theorem 11.1.

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. The track is readily seen to have a 2-sided preimage in the cover $\ddot{N}_e \to N_e$ that is induced by $\ddot{X}$. A confusing point is that $\pi_1 X^*$ acts with inversions on the walls wereas $\pi_1(\ddot{X})^*$ acts without inversions.

That one obtains a hierarchy follows 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 [Wis04] since these antipodal walls are of the type considered there. Consequently, our group theoretical

hierarchy is not needed, and the hierarchy of the cube complex itself is a malnormal hierarchy. We are thus able to obtain virtual specialness by repeatedly applying Theorem 5.2. It is not necessary to use Theorem 11.1 and in particular, we do not need Theorem 7.2 here.

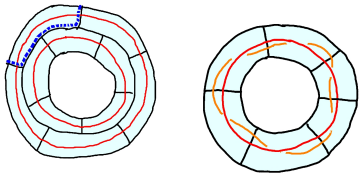

FIGURE 132. Annular diagrams. Width 2 is impossible because of the four pieces. Width 1 implies that $g\widetilde{N}_e = \widetilde{N}_e$. Width 0 implies that $g\widetilde{N}_e \cap \widetilde{N}_e \neq \emptyset$.

Let us focus on the malnormality. Consider an annular diagram $A \to X^*$ that has both boundary cycles on the same track carrier (so $\widetilde{A}$ 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 10.4, we replace A with an annuladder B in the same class such that B is either thin or thick and contains annuladders L_1, L_2 etc. However, the thick (width 2) case is excluded since there is no essential path $P \to N(T_e)$ such that P is the concatenation of ≤ 4 pieces and such that the initial and terminal edge of P is e. 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. Consequently, as $\ddot{T}_e$ is a deformation retract of $\ddot{N}_e$, the two elements are actually conjugates of an element of $\pi_1\ddot{N}_e = \pi_1\ddot{T}_e$ as required by malnormality.

We note that the malnormality could have been deduced from Theorem 10.1. Indeed, as above we can assume by minimality of A and the fact that carriers have no missing shells that A has no features of positive curvature, and hence it is a width 0, 1, or 2 annulus by Theorem 10.1. Now as above the $\leq \frac{1}{8}$ condition implies that A cannot be width 2, and a similar argument considering pieces excludes the width 1 case since $\widetilde{A}$ cannot travel alongside the wall carrier $\widetilde{N}_e$ along $\tilde{\alpha}_i$. The width 0 case shows that $\widetilde{N}_e \cap g\widetilde{N}_e \neq \emptyset$, but then Stabilizer$(\widetilde{N}_e) \cap$Stabilizer$(g\widetilde{N}_e) \subset$ Stabilizer$(\widetilde{N}_e \cap g\widetilde{N}_e)$ and the latter is trivial $\widetilde{N}_e \cap g\widetilde{N}_e$ is compact - just a finite tree in this case. □

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

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

1. $C'(\frac{1}{24})$.
2. strongly convex *wallspace cones (in the sense that if $P \to Y$ 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 square compress to a path $P' \to N(V)$.)*
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 connected intersection with each piece.*

(6) *Suppose each Y_i is compact.*

Sketch. Let $A \to X^*$ be an annular diagram whose boundary cycles α_1, α_2 both lift to the carrier of the same wall $N(W)$ under the lift $\widetilde{A} \to \widetilde{X}^*$. Applying Theorem 10.4, we can pass to $B \to \widetilde{X}^*$ 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 that B is thin (it cannot have any reductions by minimality) and there is a map $\widetilde{B} \to N(W)$.

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 $\widetilde{L}_1, \widetilde{L}_2$ map to $N(W_1) = g_1 N(W)$ and $N(W_2) = g_2 N(W)$, 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 $N(W_1) = N(W_2)$. $\square$

12. Cubical version of Filling Theorem

Theorem 12.1. *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 $G/\langle\langle P_i' \rangle\rangle$ is hyperbolic whenever $P_i/\langle\langle P_i' \rangle\rangle$ is hyperbolic and $S_i \cap P_i' = \emptyset$.*

Proof. Let $\widetilde{F}_i$ be a P_i-cocompact superconvex subcomplex of $\widetilde{X}$. Let R be an upper bound on the length of (contiguous) wall-pieces on the $\widetilde{F}_i$ (i.e. rectangular flaps hanging off are bounded by superconvexity and cocompactness). Let M be an upper bound on the lengths of contiguous cone-pieces (comes from cocompactness and almost malnormality).

Choose S_i to be all "short" elements in P_i – i.e. those of length $< 24 \max R, M$.

Then $G/\langle\langle P_i' \rangle\rangle = \pi_1 X^*$ where $X^* = \langle X \mid F_1, \dots, F_r \rangle$ and $F_i = P_i' \backslash \widetilde{F}_i$.

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. Consider a pair of geodesics γ_1, γ_2 in $\widetilde{X}^*$ 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.

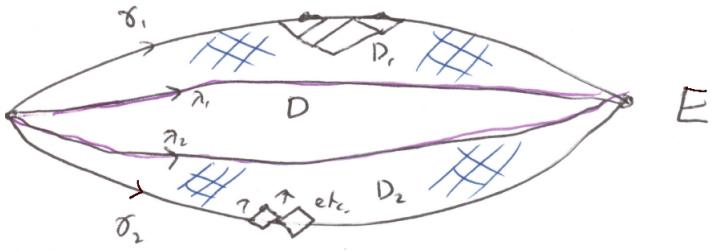

Figure 133. $E = D_{1\lambda_1} \cup_{\lambda_2} D_2$.

D is a ladder (or single cone-cell or single point) since it has no cornsquares (by maximality of D_1, D_2 or shells along λ_1, λ_2 (by short innerpaths and geodesicity of γ_1, γ_2).

γ_i, λ_i must σ-fellowtravel relative to parabolics because of relative hyperbolicity. So we have a situation as illustrated in Figure 134.

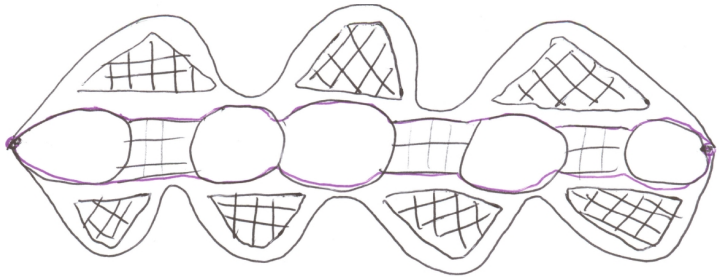

Figure 134. E consists of a ladder sandwiched between two parabolic fellow-travelling bigons.

The cone-pieces between successive cones in the ladder are bounded with net length $\leq M$ so there is L-fellow traveling at the rectangles in the ladder D, and "around" cone-cells since each F_i is δ_i-hyperbolic so geodesics that start and end $\leq L$ from each other $\kappa_i L$-fellow travel in F_i. See the two diagrams at the left of Figure 135.

FIGURE 135.

And $\lambda_i\gamma_i$ uniformly fellow travel since they do so around the projection of each $\widetilde{F}_i$ The diagram on the right of Figure 135 lies in $N_\sigma(\widetilde{F}_i)$ and projects to a pair of geodesics with boundedly close endpoints in the hyperbolic space $P_i'\backslash N_\sigma(\widetilde{F}_i)$ and thus fellow travel. □

12.1. **Persistence of quasiconvexity.** We follow the terminology of Theorem 12.1.

Theorem 12.2. *Let $H \subset G$ be a* full *relatively quasiconvex subgroup (so $|H \cap P_i^g| = \infty$ implies that $[P_i^g : P_i^g \cap H] < \infty$ for each g).*

Then there exists slightly larger sets S_i^+ (without 1_G) such that $\bar{H}$ is quasiconvex in $\bar{G}$ whenever $P_i' \cap S_i^+ = \emptyset$.

Proof. Let $\widetilde{A}$ be a superconvex H-cocompact core, and $A = H\backslash\widetilde{A}$. There exists D such that $\operatorname{diam}(g_i\widetilde{F}_i \cap \widetilde{A}) < D$ unless $g_i\widetilde{F}_j \subset \widetilde{A}$ (by fullness, cocompactness etc.)

Choose S_i^+ so that the induced presentation $A^* \to X^*$ has no missing shells and short innerpaths. Now $\widetilde{A}^* \subset \widetilde{X}^*$ is an isometric embedding. □

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

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

The interested reader can recognize a more general formulation when X is $C(6)$ that functions when S is the concatenation of three or fewer pieces.

A map $Y \to X$ 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 [MW05].

Theorem 12.3. *Let X be a $C'(\frac{1}{6})$-complex. If $Y \to X$ has no missing shells then $\widetilde{Y} \to \widetilde{X}$ is an injection. In particular $\pi_1 Y \to \pi_1 X$ is injective.*

Proof. Suppose there are two distinct 0-cells u, v of $\widetilde{Y}$ that map to the same 0-cell of $\widetilde{X}$. Let $D \to \widetilde{X}$ be a minimal area diagram among all those with a boundary path $P \to \widetilde{Y}$ from u to v. By Theorem 8.1, 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 P. Since $Y \to X$ has no missing shells, the path $Q \to \widetilde{Y}$ extends to $R \to \widetilde{Y}$ and this allows us to produce a small area diagram D' obtained by removing R and replacing Q by S. □

In this document we have emphasized the scenario where shells have *short inner paths* in the sense $|S| < |Q|$ - we note that this is not always the $C(6)$ case where the innerpath S is the concatenation of at most 3 pieces.

Theorem 12.4. *Let X be a compact $C'(\frac{1}{6})$-complex. If $Y \to X$ has no missing shells then $\widetilde{Y} \to \widetilde{X}$ is a quasiconvex subcomplex.*

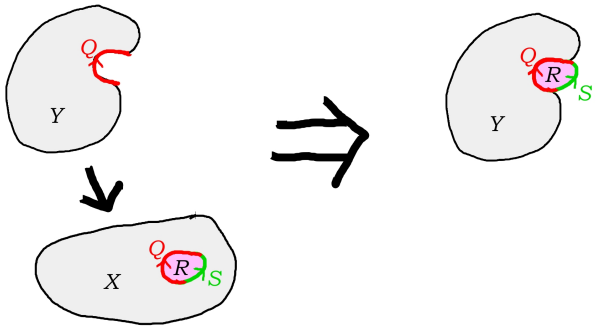

FIGURE 136. $Y \to X$ has no missing shell.

Proof. Let D be a minimal area diagram (and minimal number of 1-cells) between a geodesic $\gamma \to \widetilde{X}$ and a variable path $\sigma \to \widetilde{Y}$ with the same endpoints as γ. Theorem 8.1 implies that D is a ladder. Indeed, there is no shell or spur of D on σ because of 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 $N_r(\widetilde{Y})$ where r is half the maximal circumference of a 2-cell of X. □

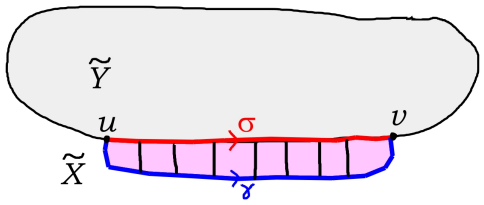

FIGURE 137. $\widetilde{Y} \to \widetilde{X}$ is quasiconvex.

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 $X^* = \langle Y_1, \dots \rangle$ be cubical presentations. A *map of cubical presentations* $f : A^* \to X^*$ has the property that $f : A \to X$ is a local isometry of nonpositively curved cube complexes, and for each B_i there is $Y_{f(i)}$ such that B_i maps to X through $Y_{i'}$ as follows:

$$\begin{array}{ccc} B_i & \to & Y_{i'} \\ \downarrow & & \downarrow \\ A & \to & Y \end{array}$$

The map $f : A^* \to X^*$ has *no missing shells* if the following holds: For each closed essential path $QS \to X$ mapping to some Y_j with $|Q| > |S|$ and $S \to \widetilde{X}$ a geodesic, if there is a lift $Q \to A$, 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:

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

The proof of the following is analogous to that of Theorems 12.3 and 12.4.

Theorem 12.5. *Let X^* be a cubical presentation that satisfies the $C'(\frac{1}{24})$ small-cancellation condition. Let $A^* \to X^*$ have no missing shells. Then $\widetilde{A}^* \to \widetilde{X}^*$ is an embedding. Moreover, if X is compact then $\widetilde{A}^* \to \widetilde{X}^*$ is quasiconvex.*

13. The structure of groups with a quasiconvex hierarchy

The main result of these notes is that:

Theorem 13.1. *If G has a quasiconvex hierarchy then G is virtually special.*

This follows from the following by induction on the length of the hierarchy:

Theorem 13.2. *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.*

Remark 13.3. We will give a proof that is structurally correct and probably supportable but has a gap - at a particular computational point.

Plan: We seek to separate C from its intersecting conjugators in a finite quotient $G \to F$. 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 11.1 and thus G is as well.

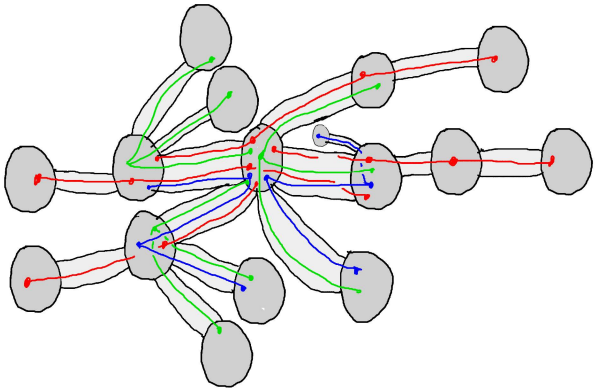

Figure 138. Partially illustrated finite cover of the graph of spaces corresponding to the HNN extension in Figure 139. The immersed subspaces corresponding to failure of malnormality are locally separated by edge spaces, and so each individual edge space is malnormal.

Lemma 13.4. *A quasiconvex subgroup $C \subset G$ has finitely many cosets Cg_i of* intersecting conjugators *having the property that $|g^{-1}Cg \cap C| = \infty$.*

Lemma 13.5. *If C is separable and quasiconvex in the hyperbolic group G then there exists a subgroup G' such that C is malnormal in G'.*

Inductively achieving plan: It suffices to consider the HNN extension case $G = A*_{C^t=D}$. We list the finitely many intersecting conjugators $g_1, \ldots, g_r \in G$, each of which has a normal form $g_i = (a_{i1}t^{\pm\epsilon_{i1}}a_{i2}t^{\pm\epsilon_{i2}}\ldots)$. Let $\{H_1, \ldots, H_r\}$ denote the representatives of maximal infinite intersections of conjugates of C. The idea is to use Theorem 11.2 to choose finite index subgroups $\{H_1', \ldots, H_r'\}$ such that there is a quotient:

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

where $\bar{A} = A/\langle\langle H_i' \rangle\rangle$ and $\bar{C}, \bar{D}$ are the images of C, D in $\bar{A}$; and the conjugation homomorphism $t : C \to D$ projects to $t : \bar{C} \to \bar{D}$; and $\bar{C} \subset \bar{G}$ is quasiconvex; and $\text{Height}_{\bar{G}}(\bar{C}) < \text{Height}_G(C)$; and the intersecting conjugators do not map to $\bar{C}$.

Most of the above are fairly natural for large girth choices – requiring: a filling theorem, intersection computation lemmas, NFT and small-cancellation theory.

Main Difficulty: How do we choose $A \to \bar{A}$ so that it induces:

$$\begin{array}{ccc} C & \overset{t}{\to} & D \\ \downarrow & & \downarrow \\ \bar{C} & \overset{t}{\to} & \bar{D} \end{array}$$

To understand the problem, even if C were malnormal in $A_C B$, Theorem 11.2 doesn't (immediately) allow us to control $C \to \bar{C} = C/\langle\langle C' \rangle\rangle$ so that there is an induced quotient $A *_C B \to \bar{A} *_{\bar{C}} \bar{B}$ using the same quotient of C on both sides. To address this we have:

Auxiliary HNN extension: $K = A*_{C_{ij}^{t_{ij}}} = D_{ij}$ which is a multiple HNN extension corresponding to collections of maps of maximal trees T_i with ∞-stabilizer in Bass-Serre tree of $A*_{C^t=D}$. K is virtually special (using a malnormal quasiconvex hierarchy) and the H_i subgroups correspond to the T_i. Theorem 11.2 applies to $\big(K, \{H_1, \ldots, H_r\}\big)$ gives us the strengthened form we need. In the torsion-free case, small-cancellation theory allows us to conclude that $\bar{C} = C/\langle\langle (C \cap H_i')^C \rangle\rangle$ and $\bar{D} = D/\langle\langle (D \cap (H_i')^t)^D \rangle\rangle$.

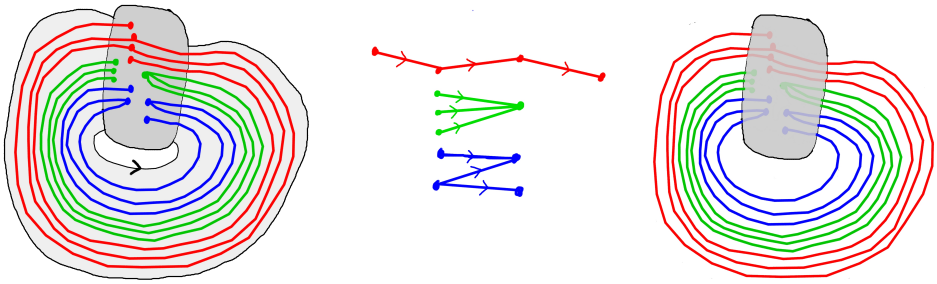

FIGURE 139. Immersed spaces on the left corresponding to the three orbits of maximal ∞-stabilizer trees T_1, T_2, T_3 in the middle. On the right is a graph of spaces for the corresponding auxiliary HNN extension.

14. The relatively hyperbolic setting

Lemma 14.1 (Virtually Special Fillings)**.** *Let G be hyperbolic relative to $\{P_i\}$ and virtually sparse special. There exists P_i^o such that for $P_i' \subset_f P_i^o$ the quotient $G/\langle\langle P_i' \rangle\rangle$ is virtually compact special hyperbolic.*

Sketch. This is nearly the same as the proof of Theorem 11.2. □

Corollary 14.2 (Controlling Cusps)**.** *Let M be a hyperbolic 3-manifold with boundary components $T_1, \dots, T_r$. There exist finite covers $\hat{T}_i^o \to T_i$ such that for any further finite covers $\widehat{T}_i \to \hat{T}_i^o$, there is a finite regular cover $\widehat{M} \to M$ such that the induced covers of the boundary components are $\widehat{T}_i$.*

Theorem 14.3. *Let X be a nonpositively curved cube complex then X is special provided:*

1. *X is sparse*
2. *$\pi_1 X$ is hyperbolic relative to virtually abelian subgroups*
3. *X has a cubical hierarchy (e.g. X is compact and has finitely many hyperplanes and all are 2-sided and embedded).*

Sketch. We apply Lemma 14.1 to verify double hyperplane coset separability, and thus obtain virtual specialness by a version of Theorem 4.6. The key point is to fill the parabolics so that the quotient still has a hierarchy, and is thus virtually special by Theorem 13.1. □

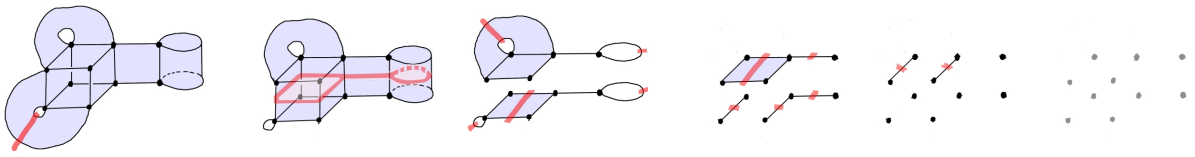

Figure 140. A cubical hierarchy of length 14.

Theorem 14.4. *Let G be hyperbolic relative to abelian subgroups and suppose G splits as a graph of groups with special vertex groups and quasiconvex edge groups. Then quasiconvex subgroups of G are separable.*

Proof. Use Lemma 14.1 to quotient the vertex groups to hyperbolic virtually special groups – note that we also quotient cyclic groups in the vertex groups which are 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 14.1 can accommodate this, and we obtain a quotient $G \to \bar{G}$. 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 13.1 (or actually a version of it terminating in compact virtually special hyperbolic groups) we see that $\bar{H}$ can be separated from $\bar{g}$ and we are done. □

We hope to be able to prove:

Conjecture 14.5. *Let G be hyperbolic relative to virtually abelian groups. Suppose G has a quasiconvex hierarchy. Then G is virtually (compact) sparse special.*

Theorem 14.6. *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.*

Add picture showing the cages that are separated and then used to give alternate splitting of G'.

Add picture to explain Theorem 14.7

Sketch. Apply Theorem 14.4 to obtain G' that splits as a graph of groups with relatively malnormal quasiconvex edge groups. The group G' can be cubulated (cocompactly) using Theorem 14.7. Now apply a variant of Theorem 11.1 to see that double hyperplane cosets are separable. □

The following is a relatively hyperbolic generalization of Theorem 7.2.

Theorem 14.7. *Let G split as $A *_C B$ with G hyperbolic relative to abelian subgroups, and $C \subset B$ malnormal and aparabolic, and A isomorphic to a multiple central HNN extension of C. If B is cocompactly cubulated then so is A.*

15. Applications

15.1. **Baumslag's conjecture.** A *one-relator group* is a group having a presentation $\langle a, b, \cdots \mid W^n \rangle$ 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 [LS77] 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 15.1 ([Bau67])**.** *Every one-relator group with torsion is residually finite.*

The main tool for studying one-relator groups is the *Magnus 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 G. The hierarchy terminates at a virtually free group of the form $\mathbb{Z}_n * F$.

For one-relator groups with torsion, the subgroups M, M' are quasiconvex at each level of the hierarchy in [Wisb]. This result depends upon a variant of the Newman spelling theorem [HW01, Lau07]. Remarkably, computer studies done by several researchers, most notably Dunfield-Thurston [DT06], suggest that most (e.g. 94% !?) one-relator groups are of the form $F_n \ltimes \mathbb{Z}$ and in particular, the Magnus hierarchy terminates quickly but is not a quasiconvex hierarchy.

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 Theorem 13.1.

Theorem 15.2. *Every one-relator group with torsion is virtually special.*

Because of Proposition 2.10, a virtually special hyperbolic has very strong properties, and in particular it is residually finite, so Conjecture 15.1 follows from Theorem 15.2.

15.2. **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 surfaces* is a 2-sided π_1-injective surface along which $\pi_1 M$ splits 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 [CS84]. In general, an incompressible surface in a hyperbolic 3-manifold is either geometrically finite or virtually corresponds to a fiber (see [Bon86]). 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 [Can94]).[2] Finally, the geometrical finiteness of an incompressible surface where the 3-manifold splits corresponds precisely to the quasi-isometric embedding of the corresponding

[2]Most 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

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 15.3. *If M is a hyperbolic 3-manifold with an incompressible geometrically finite surface then $\pi_1 M$ is virtually special.*

One hopes that all hyperbolic fibered 3-manifolds have finite covers with incompressible geometrically finite surfaces, so that all Haken hyperbolic 3-manifolds are virtually special.

Corollary 15.4. *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 [Ago08]:

Proposition 15.5 (Agol's fibering criterion)**.** *Let M be a compact 3-manifold, and suppose that $\pi_1 M$ is residually finite $\mathbb{Q}$-solvable. 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 graph group which is residually finite rational solvable and so:

Corollary 15.6. *Every hyperbolic Haken 3-manifold is virtually fibered.*

15.3. **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 $G \to \bar{G}$ 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.

(1) $A *_Z B$ where Z is cyclic and malnormal in A, and A, B have such hierarchies.
(2) $A*_{Z^t=Z'}$ 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 [KM98], and is also implicit in Sela's retractive tower description of limit groups [Sel03]. This hierarchy allows one to prove that limit groups are hyperbolic relative to free abelian subgroups [Dah03, Ali05]. Using this cyclic hierarchy and the relative hyperbolicity one concludes that:

Corollary 15.7. *Every limit group is virtually special.*

References

[Ago04] Ian Agol. Tameness of hyperbolic 3-manifolds, 2004. Preprint.

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

[Ali05] Emina Alibegović. A combination theorem for relatively hyperbolic groups. *Bull. London Math. Soc.*, 37(3):459–466, 2005.

[Bau67] Gilbert Baumslag. Residually finite one-relator groups. *Bull. Amer. Math. Soc.*, 73:618–620, 1967.

[BHW11] Nicolas Bergeron, Frédéric Haglund, and Daniel T. Wise. Hyperplane sections in arithmetic hyperbolic manifolds. *J. Lond. Math. Soc. (2)*, 83(2):431–448, 2011.

[BM97] Marc Burger and Shahar Mozes. Finitely presented simple groups and products of trees. *C. R. Acad. Sci. Paris Sér. I Math.*, 324(7):747–752, 1997.

[Bon86] Francis Bonahon. Bouts des variétés hyperboliques de dimension 3. *Ann. of Math. (2)*, 124(1):71–158, 1986.

[BW] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *American Journal of Mathematics*. To Appear.

[Can94] Richard D. Canary. Covering theorems for hyperbolic 3-manifolds. In *Low-dimensional topology (Knoxville, TN, 1992)*, Conf. Proc. Lecture Notes Geom. Topology, III, pages 21–30. Internat. Press, Cambridge, MA, 1994.

[CG06] Danny Calegari and David Gabai. Shrinkwrapping and the taming of hyperbolic 3-manifolds. *J. Amer. Math. Soc.*, 19(2):385–446 (electronic), 2006.

[Cha07] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata*, 125:141–158, 2007.

[CS84] M. Culler and P. B. Shalen. Bounded, separating, incompressible surfaces in knot manifolds. *Invent. Math.*, 75(3):537–545, 1984.

[Dah03] François Dahmani. Combination of convergence groups. *Geom. Topol.*, 7:933–963 (electronic), 2003.

[DJ00] Michael W. Davis and Tadeusz Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *J. Pure Appl. Algebra*, 153(3):229–235, 2000.

[Dro83] Carl Droms. *Graph Groups*. PhD thesis, Syracuse University, 1983.

[DT06] Nathan M. Dunfield and Dylan P. Thurston. A random tunnel number one 3-manifold does not fiber over the circle. *Geom. Topol.*, 10:2431–2499, 2006.

[Far03] Daniel S. Farley. Finiteness and CAT(0) properties of diagram groups. *Topology*, 42(5):1065–1082, 2003.

[Git99] Rita Gitik. On the profinite topology on negatively curved groups. *J. Algebra*, 219(1):80–86, 1999.

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

[Hum94] S. P. Humphries. On representations of Artin groups and the Tits conjecture. *J. Algebra*, 169:847–862, 1994.

[HWa] G. Christopher Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. pages 1–42. Preprint 2010.

[HWb] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. Preprint.

[HW99] Tim Hsu and Daniel T. Wise. On linear and residual properties of graph products. *Michigan Math. J.*, 46(2):251–259, 1999.

[HW01] G. Christopher Hruska and Daniel T. Wise. Towers, ladders and the B. B. Newman spelling theorem. *J. Aust. Math. Soc.*, 71(1):53–69, 2001.

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

[HW09] G. Christopher Hruska and Daniel T. Wise. Packing subgroups in relatively hyperbolic groups. *Geom. Topol.*, 13(4):1945–1988, 2009.

[HW10] Frédéric Haglund and Daniel T. Wise. Coxeter groups are virtually special. *Adv. Math.*, 224(5):1890–1903, 2010.

[JW] David Janzen and Daniel T. Wise. Cubulating rhombus groups. pages 1–19. Preprint, 2011.

[KM98] O. Kharlampovich and A. Myasnikov. Hyperbolic groups and free constructions. *Trans. Amer. Math. Soc.*, 350(2):571–613, 1998.

[Lau07] Joseph Lauer. Cubulating one relator groups with torsion. Master's thesis, McGill University, 2007.

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

[Min04] Ashot Minasyan. Separable subsets of gferf negatively curved groups. Preprint, 2004.

[Mos95] Lee Mosher. Geometry of cubulated 3-manifolds. *Topology*, 34(4):789–814, 1995.
[MW05] J. P. McCammond and D. T. Wise. Coherence, local quasiconvexity, and the perimeter of 2-complexes. *Geom. Funct. Anal.*, 15(4):859–927, 2005.
[NR03] G. A. Niblo and L. D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.
[Sel03] Z. Sela. Diophantine geometry over groups. II. Completions, closures and formal solutions. *Israel J. Math.*, 134:173–254, 2003.
[Wei71] C. M. Weinbaum. The word and conjugacy problems for the knot group of any tame, prime, alternating knot. *Proc. Amer. Math. Soc.*, 30:22–26, 1971.
[Wisa] Daniel T. Wise. Recubulating free groups. *Israel J. Math.*, pages 1–7. To appear.
[Wisb] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. pages 1–189.
[Wisc] Daniel T. Wise. Virtual cleanliness. pages 1–21.
[Wis02] Daniel T. Wise. The residual finiteness of negatively curved polygons of finite groups. *Invent. Math.*, 149(3):579–617, 2002.
[Wis04] Daniel T. Wise. Cubulating small cancellation groups. *GAFA, Geom. Funct. Anal.*, 14(1):150–214, 2004.
[Wis06] Daniel T. Wise. Subgroup separability of the figure 8 knot group. *Topology*, 45(3):421–463, 2006.
[Wis07] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.
[Yan] Wen-Yuan Yang. *Proc. Amer. Math. Soc.* To appear.

Dept. of Mathematics & Statistics, McGill University, Montreal, Quebec, Canada H3A 2K6
E-mail address: wise@math.mcgill.ca