THE STRUCTURE OF GROUPS WITH A QUASICONVEX HIERARCHY

DANIEL T. WISE

Abstract. Let G be a word-hyperbolic group with a quasiconvex hierarchy. We show that G has a finite index subgroup G' that is a quasiconvex subgroup of a right-angled Artin group. It follows that every quasiconvex subgroup of G is a virtual retract, and is hence separable. The results are applied to certain 3-manifold and one-relator groups.

Contents

Date: February 29, 2012.
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. This document was completed while on sabbatical at the Hebrew University, and I am grateful to the Einstein Institute of Mathematics as well as the Lady Davis Foundation for its support, and grateful to Ian Agol, Mark Hagen, Frédéric Haglund, Jason Manning, Piotr Przytycki, and Eric Swenson for many corrections.

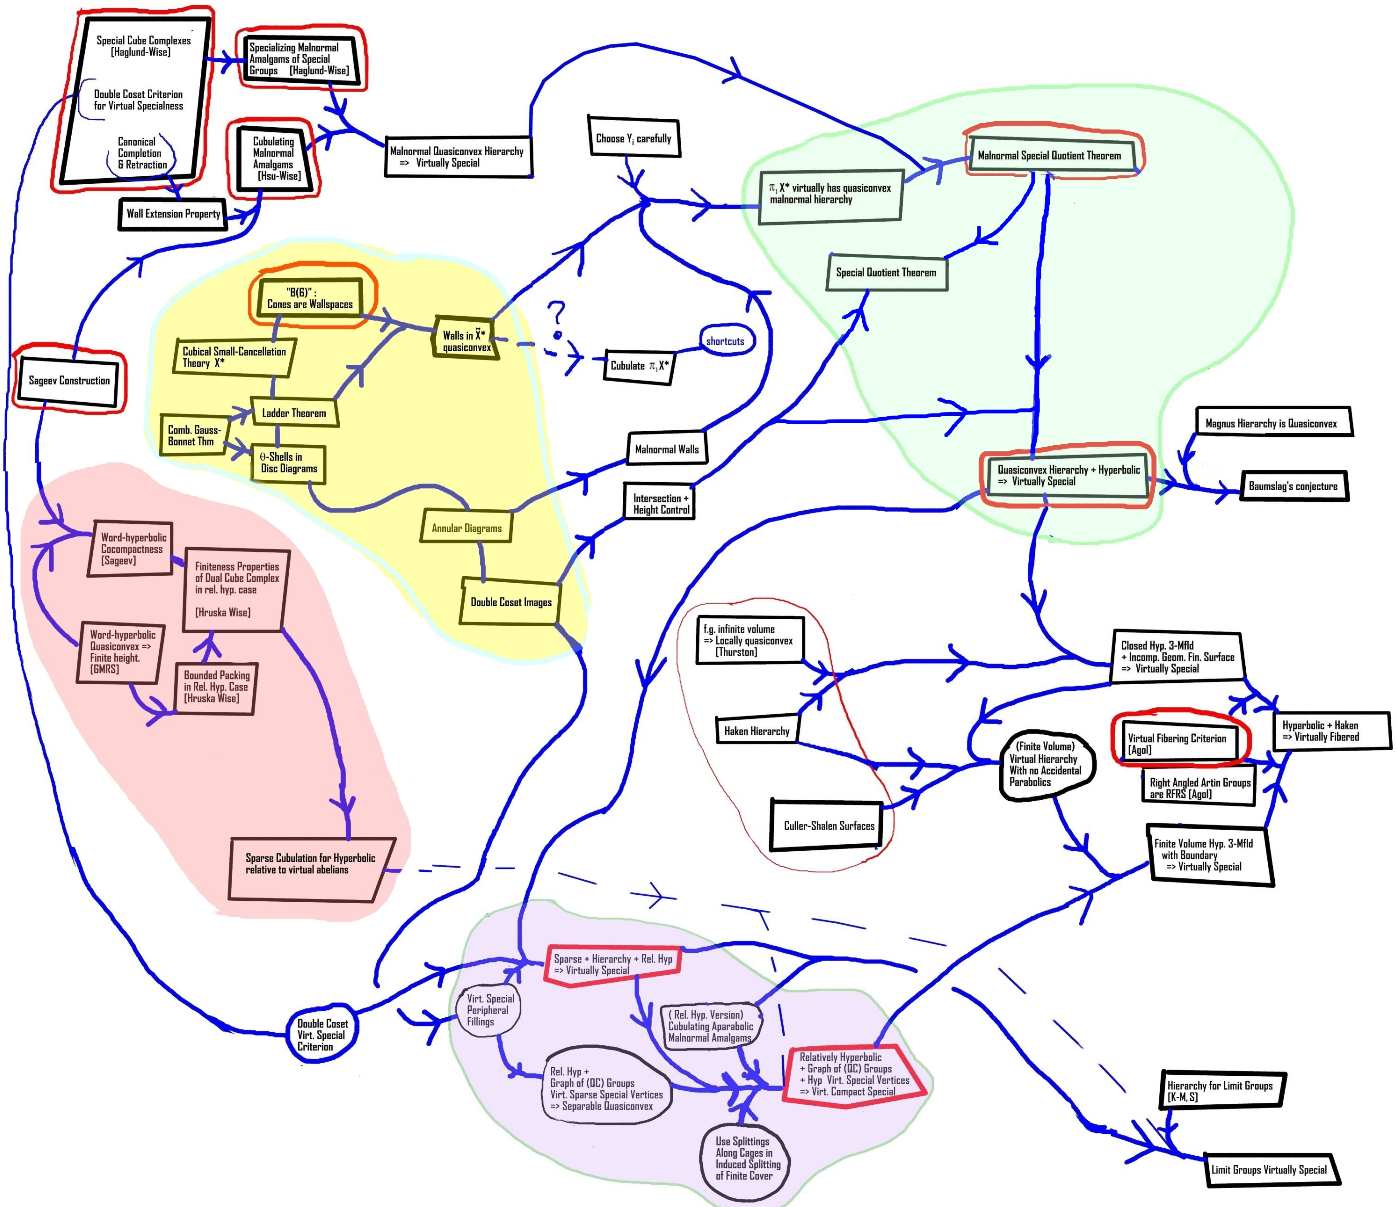

Figure 1. A flow chart indicating main points of the paper and some topical constellations.

1. Introduction

This paper has several parts:

In the first part of the paper we develop a small-cancellation theory over cube complexes. When the cube complex is 1-dimensional, we obtain the classical small-cancellation theory, as well as the closely related Gromov graphical small-cancellation theory.

It is hard to say what the main result is in the first part, since it seems the definitions are more important than the theorems. For this and the second part, the reader might wish to scan the table of contents to get a feel for what is going on. We give the following sample result to give an idea of the scope here. In ordinary small-cancellation theory, when $W_1, \ldots, W_r$ represent distinct conjugacy classes, the presentation $\langle a, b, \ldots \mid W_1^{n_1}, \ldots, W_r^{n_r} \rangle$ is "small-cancellation" for sufficiently large n_i. In analogy with this we have the following:

C6-Sample. *Let X be a nonpositively curved cube complex. Let $Y_i \to X$ be a compact local isometry for $1 \leq i \leq r$ such that each $\pi_1 Y_i$ is malnormal, and $\pi_1 Y_i, \pi_1 Y_j$ do not share any nontrivial conjugacy classes. Then $\langle X \mid \widehat{Y}_1, \ldots, \widehat{Y}_r \rangle$ is a "small-cancellation" cubical presentation for sufficiently large "girth" finite covers $\widehat{Y}_i \to Y_i$.*

Many other general small-cancellation theories have been propounded. For instance two such graded theories directed especially towards Burnside groups were produced by Olshanskii and McCammond. Stimulated by Gromov's ideas of small-cancellation over word-hyperbolic groups, there have been later important works of Olshanskii, followed by more recent theories "over relatively hyperbolic groups" by Osin [Osi06] and Groves-Manning [GM08]. The theory we propose is decidedly more geometric, and arguably favors explicitness over scope. However, although it may be more limited by presupposing a nonpositively curved cube complex as a starting point, it has the advantage of not presupposing (relative) hyperbolicity - yet some form of hyperbolicity must lurk inside for there to be any available small-cancellation.

In the second part of the paper we impose additional conditions that lead to the existence of a wallspace structure on the resulting small-cancellation complex. We can illustrate the nature of the results with the following sample:

B6-Sample. *Let G be an infinite word-hyperbolic group acting properly and cocompactly on a CAT(0) cube complex. Let $H_1, \ldots, H_k$ be quasiconvex subgroups that are not commensurable with G. And suppose that each H_i has separable hyperplane stabilizers. There exist finite index subgroups $H_1', \ldots, H_k'$ such that the quotient $G/\langle\langle H_1', \ldots, H_k' \rangle\rangle$ has a codimension-1 subgroup.*

In the third part of the paper, we probe further and seek a virtually special cubulation.

We then prove the following:

Theorem A (Special Quotient Theorem)**.** *Let G be a word-hyperbolic group that is virtually the fundamental group of a compact special cube complex. Let $H_1, \ldots, H_r$ be quasiconvex subgroups of G. Then there are finite index subgroups $H_i' \subset H_i$ such that: $G/\langle\langle H_1', H_2', \ldots, H_r' \rangle\rangle$ is virtually special.*

We then prove the following:

Theorem B (Quasiconvex Hierarchy ⇒ Virtually Special)**.** *Let G be a word-hyperbolic group with a quasiconvex hierarchy, in the sense that it can be decomposed into trivial groups by finitely many HNN extensions and amalgamated free products along quasiconvex subgroups. Then G is virtually special.*

There are two important applications of the virtual specialness of groups with a quasiconvex hierarchy: It is applied to hyperbolic 3-manifolds with a geometrically finite incompressible surface to reveal their virtually special structure. This resolves the subgroup separability problem for fundamental groups of such manifolds. It also completes a proof that Haken hyperbolic 3-manifolds are virtually fibered. It is also applied to resolve Baumslag's conjecture on the residual finiteness of one-relator groups with torsion.

The fourth part of the paper deals with groups that are hyperbolic relative to virtually abelian subgroups. The results are not yet complete, but are aimed at proving similar structural results for such groups that also have quasiconvex hierarchies.

The sections marked with ⊛ are not essential to the theory leading to the above applications.

2. CAT(0) cube complexes

2.1. **Basic definitions.** An *n-cube* is a copy of $[-1, 1]^n$, and a 0-*cube* is a single point. We regard the boundary of an n-cube as consisting of the union of lower dimensional cubes. A *cube complex* is a cell complex formed from cubes, such that the attaching map of each cube is combinatorial in the sense that it sends cubes homeomorphically to cubes by a map modeled on a combinatorial isometry of n-cubes. The *link* of a 0-cube v is the complex whose 0-simplices correspond to ends of 1-cubes adjacent to v, and these 0-simplices are joined up by n-simplices for each corner of an $(n + 1)$-cube adjacent to v.

A *flag complex* is a simplicial complex with the property that any finite pairwise-adjacent collection of vertices spans a simplex. A cube complex C is *nonpositively curved* if $\text{link}(v)$ is a flag complex for each 0-cube $v \in C^0$.

2.2. **Right-angled Artin groups.** Let Γ be a simplicial graph. The *right-angled Artin group* or *graph group* $G(\Gamma)$ associated to Γ is presented by:

$$\langle\ v : v \in \text{ vertices}(\Gamma) \ \ | \ \ [u, v] : (u, v) \in \text{ edges}(\Gamma)\ \rangle$$

For our purposes, the most important example of a nonpositively curved cube complex arises from a right-angled Artin group. This is the cube complex $C(\Gamma)$ containing a torus T^n for each copy of the complete graph $K(n)$ appearing in Γ [CD95, MV95]. Each added torus T^n is isomorphic to the usual product $(S^1)^n$ obtained by identifying opposite faces of an n-cube. Note that $\pi_1 C(\Gamma) \cong G(\Gamma)$ since the 2-skeleton of $C(\Gamma)$ is the standard 2-complex of the presentation above.

2.3. **Hyperplanes in CAT(0) cube complexes.** Simply-connected nonpositively curved cube complexes are called *CAT(0) cube complexes* because they admit a CAT(0) metric where each n-cube is isometric to $[-1, 1]^n \subset \mathbb{R}^n$; however we shall rarely use this metric.

The crucial characteristic properties of CAT(0) cube complexes are the separative qualities of their hyperplanes: A *midcube* is the codimension-1 subspace of the n-cube $[-1, 1]^n$ obtained by restricting exactly one coordinate to 0. A *hyperplane* is a connected nonempty subspace of C whose intersection with each cube is either empty or consists of one of its midcubes. The 1-cells intersected by a hyperplane are *dual* to it. We will discuss immersed hyperplanes within a nonpositively curved cube complex in Section 6.1.

Remark 2.1. Hyperplanes have several important properties [Sag95]:

(1) If D is a hyperplane of C then $C - D$ has exactly two components.
(2) Each midcube of a cube of C lies in a unique hyperplane.
(3) Regarding each midcube as a cube, a hyperplane is itself a CAT(0) cube complex.
(4) The union of all cubes that D passes through, is the *carrier* of D and is a convex subcomplex of C (see Section 2.4) that is isomorphic to $D \times I$.

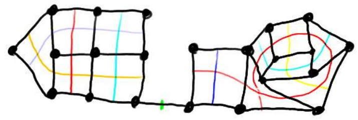

FIGURE 2. Dual curves in a square complex disk diagram

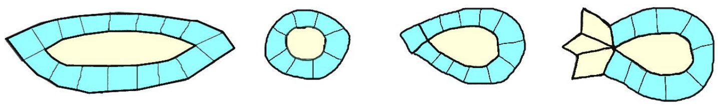

FIGURE 3. A bigon, nongon, monogon, and oscugon.

2.4. **Geodesics, local-isometries, and convexity.** While elsewhere in this paper, 1-cubes will be regarded as copies of $[-1, 1]$, to facilitate discussion of metric, we will regard each 1-cube as having length 1, and we let I_n denote the interval $[0, n]$ subdivided so that all integers are vertices. A *length n path from x to y* in a cube complex X is a combinatorial map $I_n \to X$ where $0, n \mapsto x, y \in X^0$. A path is a *geodesic* if there is no shorter length path with the same endpoints. We emphasize that geodesics are almost never unique when $\dim(X) \geq 2$, indeed there are $n!$ distinct geodesics connecting vertices at opposite corners of an n-cube. We define the *distance* between 0-cubes in a connected nonpositively curved cube complex to be the length of the geodesic between them. As usual, this provides a genuine metric on the 0-cells of the 1-skeleton. Moreover we are then able to declare the distance $d(A, B)$ between subcomplexes as the minimal distance $d(a, b)$ where $a, b \in A^0, B^0$. We also define the *diameter* diameter(Y) of a connected complex to be the supremum of the lengths of geodesics in Y.

An *immersion* is a local injection. A map $\phi : Y \to X$ between nonpositively curved cube complex is a *local-isometry* if it is an immersion and for each $y \in Y^0$, whenever u, v are ends of 1-cubes at y, if $\phi(u), \phi(v)$ form a corner of 2-cube in X at $\phi(y)$, then so do u, v. An embedding that is a local-isometry is *locally-convex*. A connected locally-convex subcomplex of a CAT(0) cube complex is called *convex*, and indeed, it can be deduced from the viewpoint in Section 2.5 that $Y \subset X$ is convex if and only if, each n-cube whose $(n-1)$-skeleton lies in Y lies in Y, and for any geodesic path $P \to X$ whose endpoints lie in Y^0, the path P lies in Y.

2.5. **Properties of minimal area cubical disk diagrams.** This section was motivated by lectures of Andrew Casson from Univ. of Texas at Austin in the 80's (apparently on generalized $C(4)$-$T(4)$ presentations related to Heegaard decompositions). I am grateful to Yoav Moriah who shared his notes with me and to Michah Sageev who encouraged me to take a look at this. [1] While the results are easy to obtain, I had not previously considered the relevance of disk diagrams to cubical complexes of dimension ≥ 3. The viewpoint here, and in particular Lemma 2.2, is due to Casson. We note that the properties listed in Remark 2.1 can be deduced from this viewpoint.

Let D be a disk diagram whose 2-cells are squares. The *dual curves* in D are paths which are concatenations of midcubes of squares of D. Note that when $D \to \widetilde{X}$ is a disk diagram in a CAT(0) cube complex, each dual curve maps to a hyperplane of $\widetilde{X}$.

The 1-cells crossed by a dual curve are *dual* to it. Note that each midcube lies in a unique maximal dual curve (or cycle). One simply extends outwards uniquely across dual 1-cells. A *bigon* is a pair of dual curves that cross at their first and last containing squares. A *monogon* is a single dual curve that crosses itself at its first and last containing squares. An *oscugon* is a single dual curve that starts and ends at distinct dual 1-cells that are adjacent but don't bound the corner of a square. A *nongon* is a single dual curve of length ≥ 1 that starts and ends on the same dual 1-cell, so it corresponds to an immersed cycle of midcubes. We refer the reader to Figure 3.

[1] I learned in August 2011 that much of this material was already explained in Sageev's thesis.[Sag95]

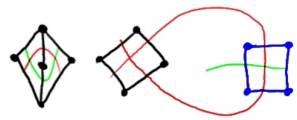

FIGURE 4. On the left is a smallest possible bigon. On the right is a monogon which must contain a smaller bigon.

Lemma 2.2. *Let $D \to X$ be a disk diagram in a nonpositively curved cube complex. If D contains a bigon or nongon among its dual curves, or if there is a pair of adjacent 1-cells that are dual to the same dual curve of length ≥ 1, then there is a new diagram D' such that:*

(1) *D' and D have the same boundary path, so $\partial_p D' \to X$ equals $\partial_p D \to X$,*
(2) $\mathsf{Area}(D') \leq \mathsf{Area}(D) - 2$ *and*
(3) *pairs of edges on $\partial_p D'$ that lie on the same dual curve of D' are precisely the same as pairs of edges on $\partial_p D$ that lie on the same dual curve of D.*

Corollary 2.3. *No disk diagram contains a monogon.*

If D has minimal area among all diagrams with boundary path $\partial_p D$, then D cannot contain a bigon, a nongon, or an oscugon.

Proof. The second statement follows immediately from Lemma 2.2. Consider a minimal area counterexample D to the first statement: So D is the union $S \cup D'$ where S is a rectangular strip $[-1, 1] \times I_n$ consisting of $n \geq 1$ squares, and carrying a dual curve σ at $\{0\} \times I_n$, and D' is a disk diagram, and $\partial_p D'$ is identified with the path $\{-1\} \times I_n$ along one side of this strip. Then a dual curve λ that is dual to a 1-cube on $\partial_p D'$ must cross σ in a second square. Let E denote the disk diagram bounded by $\lambda \cup \sigma$ and containing the cubes they pass through. Apply Lemma 2.2 to replace E by E', and obtain a smaller area counterexample. Note that Corollary 2.3 holds when $\mathsf{Area}(D) = 0$ since then each dual curve is the midcube of a 1-cube, and no dual curves cross, as there are no squares. □

Proof of Lemma 2.2. Consider a smallest area monogon, nongon, oscugon, or bigon in D. We will produce a new diagram D' with the same boundary path such that either $\mathsf{Area}(D') < \mathsf{Area}(D)$ or D' contains an even smaller such feature.

We first show that a nongon, monogon, or oscugon must contain a lower area bigon. Indeed, in each case, the dual curve α has length ≥ 1 (since squares locally embed, we see that even for a monogon, the dual curve must pass through at least one more square besides its self-crossing square). Thus, as illustrated on the right in Figure 4, a second dual curve β crosses α and travels through the diagram before crossing α a second time. The pair α, β provides a lower area bigon. (It is actually conceivable that β then itself forms a lower area oscugon (even without α), but of course we can repeat this procedure further to decrease the area further and obtain a minimal area bigon.)

Now we will show by induction on the number of squares inside a bigon that any bigonal disk diagram can be replaced by a disk diagram with smaller area.

Let α and β be the dual curves of the bigon, and let s_1 and s_2 denote their squares of intersection. We will show that either there is a smaller bigon inside, and hence by induction that bigon can be replaced by a disk diagram with at least 2 fewer squares. Or we will perform a slight modification to obtain a disk diagram with the same boundary but containing a smaller bigon, and hence this disk diagram itself can have its area reduced by 2.

The "base case", occurs when a smallest bigon arises from two squares meeting along a corner as on the left in Figure 4. Then these two squares map to the same square in X, and hence we can remove this cancellable pair to decrease the area, by replacing the pair of squares by a pair of edges glued together at a point.

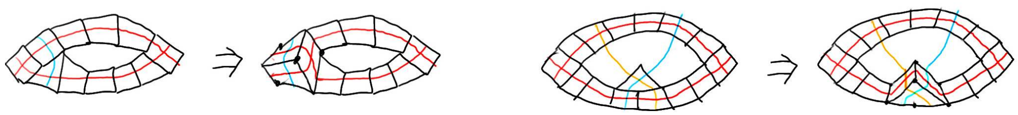

FIGURE 5. Some hexagon moves.

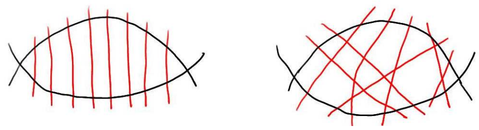

FIGURE 6.

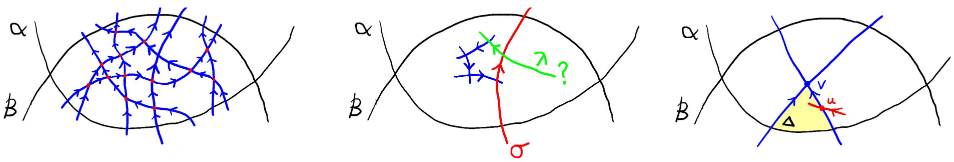

FIGURE 7.

Now suppose that α, β do not together bound a region containing a lower area bigon. Then cannot be a monogon, nongon, or oscugon, as above. Moreover, every dual curve passes through both α and β, since otherwise there would be a smaller bigon.

The plan is to find a (certain type of) minimal triangle in the complement of the dual curves, then we can perform a "hexagon move" to obtain a new disk diagram with a smaller bigon as in Figure 5. The first type of minimal triangle has one side on α and one side on β and no dual curves passing through it. The second type has its base on α, and neither of its two other sides are subsegments of β.

If the bigon contains no crossing pair of dual curves as on the left in Figure 6, then the first type of triangle occurs, and so we can perform a hexagon move of the first type. If there is at least one crossing pair of dual curves then we shall show below that the second type of triangle exists, and so we can perform a hexagon move of the second type. Hence by induction, the new diagram can have its area reduced by 2.

The collection of dual curves within our bigon forms a graph, and we make this into a directed graph by orienting all dual curves upwards from α to β, and thus orienting each edge of the graph (see the left of Figure 7). Observe that this directed graph has no directed cycle. Indeed, consider a simple directed cycle, (and we can even assume that it bounds a complementary region of the graph, for otherwise there would be a smaller area such directed cycle), and suppose that it travels counterclockwise - as an analogous argument works in the clockwise case. Among the dual curves contributing edges to the cycle, let σ denote the one having rightmost intersection with α. Let λ denote the dual curve corresponding to the next edge in the cycle. Then λ would intersect α even further to the right which is impossible (see the middle of Figure 7). Here we use that all dual curves intersect only once which follows from the minimality assumption on the bigon.

Each vertex of the graph (not on α, β) is the "top" of a triangle whose base is on α. Choose a vertex v that is minimal (excluding the leaf vertices on α) in the partial ordering induced by the cycle-free directed graph. Then the corresponding triangle Δ is our desired triangle of the second type. Indeed, if any other dual curve crosses either leg of Δ then there would be an even lower vertex u, which contradicts the minimality of v (as on the right of Figure 7). □

Remark 2.4 (Shuffling disk diagrams)**.** We shall later use the term *shuffle* to refer to an adjustment of a disk diagram obtained through a finite sequence of hexagon moves.

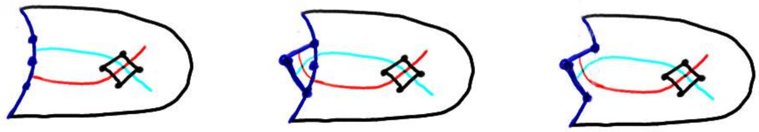

FIGURE 8.

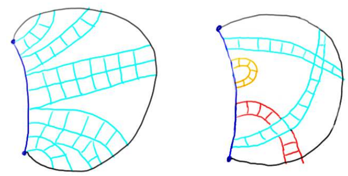

FIGURE 9.

2.6. **Splaying and Rectangles.** We now describe several related properties concerning the dual curves in minimal area cubical disk diagrams. We emphasize that our treatment focuses on subcomplexes and exclusively considers paths that are combinatorial, as discussed in Section 2.4. In particular, a subcomplex $Y \subset X$ of a CAT(0) cube complex is *convex* if: for each cube c of X, if an entire corner of c lies in Y then all of c lies in Y.

Lemma 2.5 (Splayed)**.** *Let $Y \subset X$ be a convex subcomplex of the CAT(0) cube complex. Let P be a path whose endpoints lie on Y, and let D be a disk diagram between P and Y, so there is a [geodesic] path $Q \to Y$ with the same endpoints as P and D is a diagram for PQ^{-1}. Suppose D has minimal area among all possible such choices fixing P and Y.*

Let a and b be consecutive 1-cells in Q. Then the dual curves in D starting at a and b do not intersect.

The statement of Lemma 2.5 holds with Q allowed to vary either among all such paths, or among all such geodesics. Indeed, the argument by contradiction given below provides a lower area diagram D without effecting the length of Q.

Proof. Suppose a, b are parallel in D to 1-cells a', b' that meet at the corner of a square c' in D. Since X is CAT(0), the 1-cells a, b must also meet at a square c. Since Y is convex, we see that $c \subset Y$.

We can thus adjust the diagram D to obtain a new diagram D' formed by attaching c to Q along a, b. Now $\mathsf{Area}(D') = \mathsf{Area}(D) + 1$. However, D' contains a bigon, and therefore by Lemma 2.2, its area can be reduced by two, to obtain a new diagram D'' with $\mathsf{Area}(D'') < \mathsf{Area}(D)$. This would contradict the minimality of D. See Figure 8. □

Corollary 2.6. *Let X, Y, D, Q be as in Lemma 2.5. Then there is no intersection in D between dual curves of distinct 1-cells of Q.*

The dual curves are splayed as on the left in Figure 9 (but not as on the right).

Proof. Consider an innermost pair of 1-cells whose dual curves are either equal or intersect. These 1-cells cannot be adjacent by Lemma 2.5. But any 1-cell on Q between them, would give another dual curve which either intersects one of these, or ends on another 1-cell of Q lying between them as in the right in Figure 9. This contradicts our innermost assumption. □

Corollary 2.7. *Let Y_1, Y_2 be convex subcomplexes of the CAT(0) cube complex X. Let P_1, P_2 be paths joining points on Y_1 and Y_2. Let D be a diagram with boundary path $P_1Q_1P_2Q_2$ where Q_i is a path in Y_i, and suppose D has minimal area among all such diagrams with Q_i allowed to vary.*

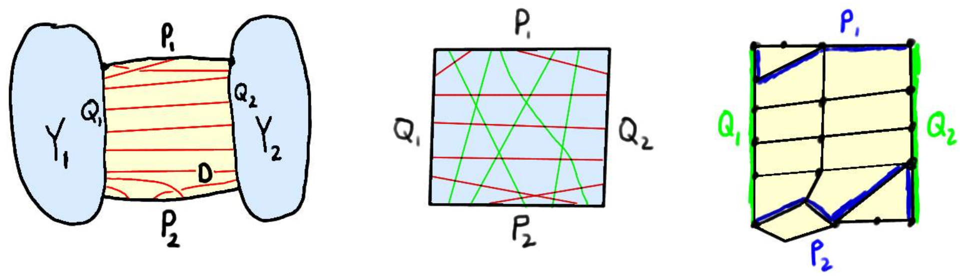

FIGURE 10.

Let S_1 and S_2 be a pair of dual curves joining edges of Q_1 to Q_2. Then the subdiagram F bounded by S_1, S_2 and the subtended portions Q_1', Q_2' of Q_1, Q_2 is a "flat rectangle" in the sense that pairs of dual curves joining Q_1', Q_2' don't cross.

Moreover, if we assume P_1, P_2 are geodesics, and that D has minimal area among disk diagrams with boundary path $P_1'Q_1P_2'Q_2$ where P_i' are geodesics with the same endpoints as P_i, then the dual curves from P_1 to P_2 have the same property: No two cross each other.

We refer the reader to Figure 10. We note that dual curves can start on Q_i and end on P_j, as illustrated on the first and second diagrams. The third diagram illustrates the second part of Corollary 2.7: Once we also choose the geodesic paths such that the area of the diagram is minimized, then all dual curves go from left to right, and from top to bottom. We obtain a genuine product rectangle in this case. The reader can view geodesic paths along the top and bottom, from which there is crossing of dual curves, and compare this to the paths P_1, P_2 that are pushed inwards somewhat.

Proof. We apply Corollary 2.6 twice: first from the point of view of Y_1, and then from the point of view of Y_2.

The proof of the second statement is similar, except now we minimize the area of a diagram with P_1 allowed to vary among paths P_1' in $Y = X$ (or even the combinatorial convex hull of the endpoints of P) but the path $Q_1P_2Q_2$ remains fixed. The same argument applies mutatis mutandis for P_2. □

Lemma 2.8 (Pushing beyond crossings)**.** *Let $D \to X$ be a minimal area disk diagram. Let S be a dual curve in D that starts and ends on 1-cells s_1, s_2 such that the boundary path of D is of the form s_1Ps_2Q. There exists a new diagram D' with the same boundary path and $\mathsf{Area}(D') = \mathsf{Area}(D)$ such that s_1, s_2 are still connected by a strip S' but the dual curves emanating from S' to P are splayed: No two cross each other on the P side of S'.*

We refer to the left pair of diagrams in Figure 12 indicating the total transformation from D to D'.

Proof. This follows by repeatedly using hexagonal replacement moves. Consider an innermost pair of a, b of edges along ∂S whose dual curves cross on the side bounded by P. If they are not adjacent, then there is an even more innermost pair. Note that the two dual curves cannot equal each other, or there would be a bigon with S, and thus the area can be reduced by Lemma 2.2.

Let c be the first square where the dual curves cross. Now add a cancellable pair of copies c'', c' of c along a, b. This increases the area by 2, and increases the area between S and P by 2. Perform a hexagonal replacement along S and the contiguous copy c' of c to obtain S', the area between P and S' is now one more than the area between P and S was. Finally the copy c'' of c has a bigon with c. We are able to reduce the area by 2. This area reduction is on the P side of S', and so the resulting diagram D' has the property that the area between P and S' has been reduced by one. Performing this repeatedly yields a new diagram whose wall has splayed strips on the P side, as claimed. See the sequence of pictures in Figure 11 for the single transformation. □

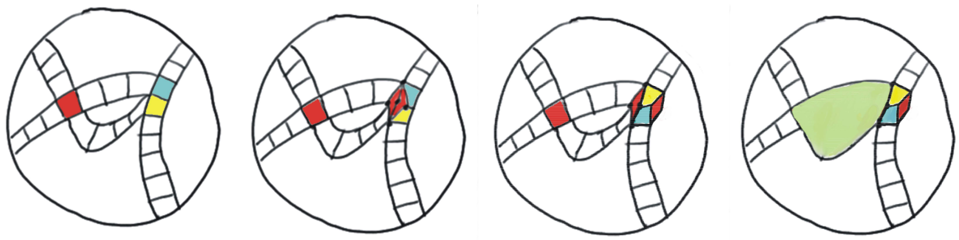

FIGURE 11.

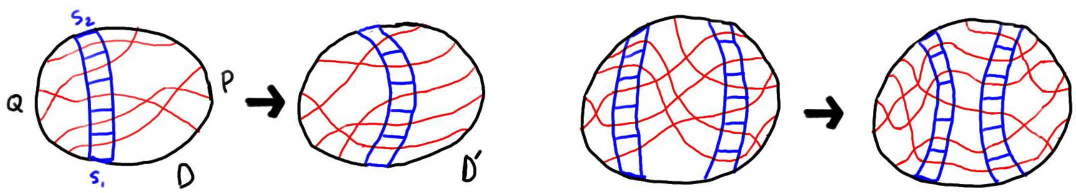

FIGURE 12.

Remark 2.9. We can apply Lemma 2.8 to understand the potential behavior between strips in disk diagrams. Let D be a diagram that has a pair of disjoint strips. Then we can replace it with a new diagram with the same boundary and at most as much area, such that the strips are moved inwards towards each other, but strips emanating from them are now splayed. See the transformation on the right in Figure 12.

This is particularly relevant when we consider a diagram between two convex subspaces Y_1, Y_2, and in particular, a diagram between a convex subspace and the carrier of a hyperplane. We are able to reach the conclusion of a "flat rectangle" between the rectangular strips.

2.7. **Annuli.** This section can be postponed until annuli arise in Section 15 and more importantly Section 5.15 and its sequels. "Generalized corners of squares" are introduced in Section 3.9.

Lemma 2.10 (Flat Annulus)**.** *Let $B \to X$ be an annular diagram. Suppose there are no generalized squares with outerpath on either the inner or outer boundary path of B. Then each dual curve starting on the inner boundary path ends on the outer boundary path (and vice-versa).*

The first annulus in Figure 13 illustrates a simple but typical example of the type of annulus examined in Lemma 2.10. This contrasts with the motivating case of a product, illustrated by the second annulus in Figure 13. The reader can imagine more elaborate examples.

Proof. An "innermost" pair of initial parts of dual curves that cross each other (and are simple curves up to that point) yields a generalized square. If two dual curves emanating from either the outer boundary circle, or the inner boundary circle, cross each other, then between these (inclusive) one can find an innermost pair of crossing dual curves (see the third annulus in Figure 13). The same holds for a simple dual curve which starts and ends on the same bounding circle.

Let d be a dual curve that starts and ends on the outer boundary path (the fourth annulus in Figure 13 contains three such scenarios). If d doesn't cross itself, then we consider the side of d not containing the inner boundary path, and obtain a generalized square as above. If d crosses itself, then since it is impossible for a dual curve to cross itself and bound a disk diagram inside, we see that some monogon in the image of d must bound the inner boundary path (the fifth annulus in Figure 13 indicates a more elaborate such scenario). The initial and terminal parts of d then provide a pair of dual curves starting on the outer boundary path that cross each other. As described at the beginning of the argument, an innermost such pair would yield a generalized square, which is impossible.

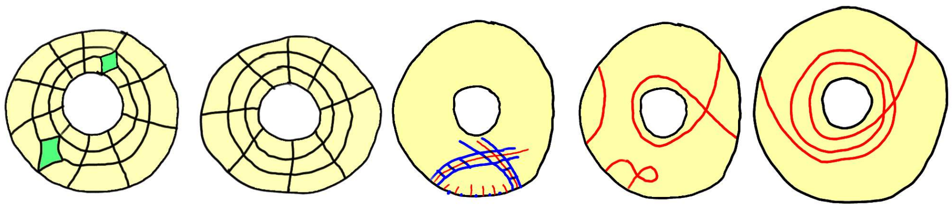

Figure 13. Square annuli.

We note that we can assume that there are no bigons on the interior of B because reducing by removing cancellable pairs corresponding to bigons doesn't effect the boundary pairing of dual curves. □

Remark 2.11. The other dual curves travel around B. When immersed hyperplanes of X do not cross themselves, we can avoid situations as in Figure 13. However, while minimal area of the diagram can help avoid some such self-crossing behavior, there is no way to avoid it in general, and we can only conclude that the "horizontal" dual curves travel "around" B, possibly multiple times.

2.8. **Rectangles and Superconvexity.** This subsection requires the definition of superconvexity from Section 8.1. We will use it later in Section 3.17 to produce examples.

Lemma 2.12. *Suppose Y is cocompact and superconvex and X is 2-dimensional, then there exists $n > 0$ with the following property: Let $F \subset X$ be a flat rectangle isomorphic to $I_m \times I_n$. If $\{0\} \times I_n$ lies in Y then F lies in Y.*

Proof. Consider a combinatorial rectangle $F = I_m \times I_n$ in X. Suppose the base $\{0\} \times I_n$ of F connects points that are quite far away, so $d_X(f(0,0), f(m,0)) > N$. Let σ be a geodesic with the same endpoints as the base, so $|\sigma| > N$. Then σ is parallel to the corresponding geodesic σ' at the top of F, and $|\sigma'| > N$.

By superconvexity we obtain a geodesic μ from the midpoint of σ to the midpoint of σ' that is contained in Y. The path μ is contained in a union of cubes stacked upon each other. Finally, parallelism and convexity shows that all of F lies in Y. □

Lemma 2.13. *Let Y be a convex and superconvex cocompact subcomplex of the CAT(0) cube complex X. There exists $D \geq 0$ such that the following holds: Let $I \times I_n \to X$ be a combinatorial strip. Suppose the base $\{0\} \times I_n$ of $I \times I_n$ lies in Y, and suppose that the distance between the endpoints of the base exceeds D, that is, $d((0,0),(0,n)) \geq D$. Then $I \times I_n$ lies in Y.*

Proof. By Lemma 8.6, the midpoint m at the top of the rectangle lies in Y. Now, each hyperplane passing through the rectangle either cuts through the base of the rectangle and hence crosses Y, or separates the top from the bottom, and hence separates m from a point in Y, and so crosses Y. It follows by convexity that the rectangle lies in Y. □

3. Small-cancellation theory over cube complexes

The goal of this section is to describe a "small-cancellation theory" for cubical presentations generalizing the usual small-cancellation theory for ordinary presentations.

3.1. **Introduction.** To orient the reader towards our eventual goals, we begin with parallel rough statements of the main theorem about classical small-cancellation diagrams and corresponding main theorem in cubical small-cancellation diagrams.

An *i-shell* in a disk diagram D is a 2-cell R with $\partial_p R = QS$, where the *outerpath* Q is a subpath of $\partial_p D$, and the *innerpath* S is internal to D except at its endpoints, and where S is the concatenation of

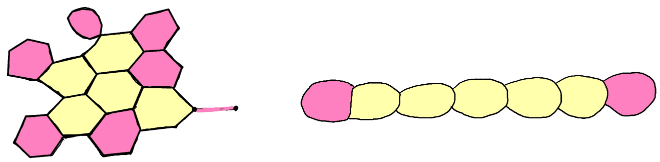

FIGURE 14. On the left in clockwise order we have a: 2-shell, 3-shell, spur, 3-shell, 1-shell, 1-shell, and 0-shell. On the right is a ladder.

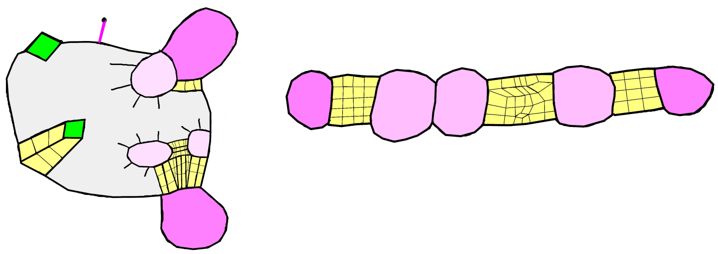

FIGURE 15. Two θ-shells, two generalized corners of squares, and a spur in a cubical small-cancellation diagram on the left. A cubical small-cancellation ladder on the right.

exactly i maximal pieces. See Figure 14 for a diagram containing i-shells. Any reduced disk diagram $D \to X$ satisfies the $C(6)$ condition provided that X is a $C(6)$ complex.

Following the language in [MW02], the main theorem of the classical small-cancellation theory in the $C(6)$ case is summarized by:

Classical small-cancellation diagrams: Let D be a $C(6)$ disk diagram. Then either D is a single 0-cell or a single 2-cell, or D contains a total of at least 2π worth of the following types of positively curved cells along its boundary:

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

Moreover, if there are exactly two such features of positive curvature, then the diagram is a ladder. See Figure 14.

In our generalization, a disk diagram D is built from *cone-cells* together with 2-cubes, 1-cubes, and 0-cubes. The cone-cells play the role that 2-cells did in the classical case, but roughly speaking, the squares are subsumed in a thickened 1-skeleton. The diagram D is "reduced" if it is locally minimal area in a certain sense (no square bigons, no generalized corners of squares on cone-cells), and it is "small-cancellation" if internal cone-cells are surrounded by many neighbors - generalizing the classical $C(6)$ condition that internal 2-cells have at least 6 sides.

The i-shells are replaced by positively curved cone-cells called *θ-shells*, where θ denotes the curvature of the θ-shell - reversing the piece-focused terminology in [MW02]. Spurs continue to play the same role, but there is now another source of positive curvature: "generalized corners of squares". Our main result is summarized by:

Cubical small-cancellation diagrams: Either D is a single 0-cell or cone-cell, or D contains at least 2π worth of positive curvature along its boundary of the form:

π for spurs
θ for θ-shells
$\frac{\pi}{2}$ for corners of generalized squares

Moreover, if there are exactly two features of positive curvature, then D is a ladder. See Figure 15.

3.2. **Cubical presentations.** A *cubical presentation* $\langle X \mid \{Y_i\}\rangle$ has "generators" consisting of a connected nonpositively curved cube complex X, and "relators" consisting of a collection $\{Y_i \to X\}$ of local-isometries of connected nonpositively curved cube complexes. The group of a cubical presentation is a quotient group $\pi_1 X/\langle\langle \pi_1 Y_i\rangle\rangle$ and is isomorphic to the fundamental group of a space X^* that we shall now discuss.

Associated to the cubical presentation $\langle X \mid \{\phi_i : Y_i \to X\}\rangle$ is a *coned-off space* X^* that is formed from the mapping cylinder $(X \cup \{Y_i \times [0,1]\})/\{(y_i, 1) \sim \phi_i(y_i)\}$ by identifying each $Y_i \times \{0\}$ to a single *cone-point* c_i. An application of the Seifert-Van Kampen lemma shows that the group of $\langle X \mid \{Y_i\}\rangle$ is isomorphic to $\pi_1 X^*$. The space X^* has a natural cell structure consisting of cubes in X together with "pyramids" consisting of cones of cubes. We refer to the image C_i of each $Y_i \times [0,1]$ as a *cone* of X^*.

An ordinary presentation $\langle a_1, \ldots, a_s \mid R_1, \ldots, R_t\rangle$ whose relators are reduced and cyclically reduced contains the data for a cubical presentation: Its generators correspond to the nonpositively curved cube complex consisting of a bouquet of s circles, and its relators correspond to a collection of t immersed cycles. The group of an ordinary presentation is the quotient of the free group on the generators of the presentation modulo the normal subgroup generated by the relators, and this group is isomorphic to the fundamental group of the standard 2-complex of the presentation. This standard 2-complex would be isomorphic to the coned-off space if we subdivide the i-th 2-cell into $|R_i|$ distinct 2-simplices meeting around the cone-point which is a new 0-cell added at the center of the 2-cell.

Remark 3.1 (Language Abuse)**.** We will use the pyramidal cells of the cones, or rather their 2-simplices, to initially discuss disk diagrams, but we will rarely use them later in the paper. Instead we will especially focus on the "data" of the cubical presentations. While we refer to the Y_i as "cones", they really function as "attaching maps" $Y_i \to X$ of the genuine cones $C(Y_i)$, and perhaps it would have been more suitable to term each Y_i as a "relator".

When we refer to $\widetilde{X}^*$, we sometimes have in mind only its underlying cube complex, which is the intermediate cover $\widetilde{X} \to \widetilde{X}^* \to X$ corresponding to the subgroup of $\pi_1 X$ equal to $\langle\langle \pi_1 Y_i\rangle\rangle$. However, sometimes we actually image $\widetilde{X}^*$ equipped with the various lifted cones gY_i, where g, i vary. Formally $\langle \widetilde{X}^* \mid gY_i :\ \text{where } g, i \text{ vary}\ \rangle$ would give the data of another cubical presentation that covers X^*. We emphasize that the underlying cube complex of $\widetilde{X}^*$ plays the role of the 1-skeleton of the universal cover of a 2-complex, whereas the various gY_i play the role of the various 2-cells.

3.3. **Pieces.** Two 1-cubes u, v are *parallel* in the CAT(0) cube complex $\widetilde{X}$ if they are dual to the same hyperplane. Let U, V be subcomplexes of the CAT(0) cube complex $\widetilde{X}$. The *wall projection* $\mathsf{WProj}_{\widetilde{X}}(U \to V)$ is a subcomplex of V defined as follows: The 1-skeleton of $\mathsf{WProj}_{\widetilde{X}}(U \to V)$ consists of the union of all closed 1-cubes of V that are parallel to 1-cubes of U. A cube of $\widetilde{X}$ lies in $\mathsf{WProj}_{\widetilde{X}}(U \to V)$ if and only if its 1-skeleton lies in $\mathsf{WProj}_{\widetilde{X}}(U \to V)$.

We declare cubes u, v (of the same dimension ≥ 2) to be *parallel* if there is a combinatorial map $c \times I_n \to \widetilde{X}$ where $c \times \{0\} \mapsto u$ and $c \times \{n\} \mapsto v$. One can show that $\mathsf{WProj}_{\widetilde{X}}(U \to V)$ is the union of all closed cubes of V that are parallel to cubes of U. The notion is examined more carefully in [HWa], where it is also defined for subcomplexes of a nonpositively curved cube complex that is not simply-connected. We refer the reader to Figure 16.

Let $\langle X \mid \{Y_i\}\rangle$ be a cubical presentation. Let $A = Y_i$ and $B = Y_j$. An *abstract cone-piece* of B in A is a nontrivial component $\mathcal{P}$ of $\mathsf{WProj}_{\widetilde{X}}(\widetilde{B} \to \widetilde{A})$, where either $i \neq j$ or $i = j$ but $\widetilde{B} \neq \widetilde{A}$. So either we are dealing with two distinct translates of the lift of the universal cover of A, or we are dealing with a base lift of $\widetilde{A}$ and an arbitrary lift of $\widetilde{B}$ with $i \neq j$. Each abstract cone-piece $\mathcal{P}$ comes equipped with a map $\mathcal{P} \to A$, and is "tagged" by data indicating the relative positions $\widetilde{A}, \widetilde{B}$. One way to indicate the data is to choose a representative *connecting strip*. This is a rectangular strip $S \cong I \times I_n$ for some $n \geq 0$ with initial and terminal edges $I \times \{0, n\}$ mapping to 1-cells in the lifts of $\widetilde{A}, \widetilde{B}$ to $\widetilde{X}$ that are associated with

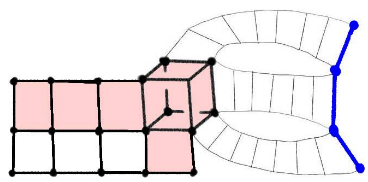

FIGURE 16. The above example illustrate a length 3 bold interval U whose wall projection onto a bold complex V is the shaded subcomplex of V. The ambient complex is not simply-connected.

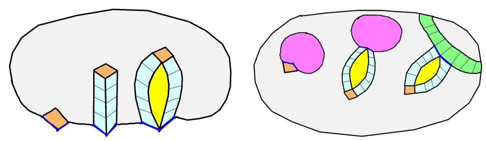

FIGURE 17. Generalized Corners of Squares: The boundary of the diagram on the left has three illustrated corners of generalized squares. The first is a genuine square with a corner on the boundary. The second is a remote square which could obviously be pushed to the boundary. The third is typical as the two rectangles from the square to the boundary bound a nontrivial square subdiagram. The diagram on the right illustrates corners of generalized squares on cone-cells and on a rectangle.

the piece. There is leeway in the choice of strip. Assuming the basepoints are chosen appropriately, S corresponds to a double coset representative $\pi_1 A g \pi_1 B$.

Continuing with the notation above, let $A = Y_i$, but now let $\widetilde{D}$ denote the carrier of some hyperplane H in $\widetilde{X}$ such that H is disjoint from the base copy of $\widetilde{A}$. In parallel to the earlier definitions, an *abstract wall-piece* of H in A is a nontrivial component $\mathcal{P}$ of $\mathsf{WProj}_{\widetilde{X}}(\widetilde{D} \to \widetilde{A})$. As above, there is some (possibly trivial) connecting strip S between $\widetilde{A}$ and $\widetilde{D}$. In parallel to the abstract cone-piece definitions, we note that we could have let D denote the immersed carrier of some immersed hyperplane in X, and then the "distinctness" condition between $\widetilde{A}$ and $\widetilde{B}$ above, becomes a requirement that the lift of $\widetilde{D}$ does not intersect $\widetilde{A}$ in a dual 1-cube of the hyperplane.

Remark 3.2. Let us look ahead to see how pieces are related to diagrams (see Figure 21 and Section 3.5 for the notion of "cone-cell".)

As suggested by Figure 18, we will bound sizes of wall-pieces by bounding sizes of contiguous wall-pieces. Given a wall-piece arising within a diagram as in the second diagram in Figure 18, we push corners of generalized squares (these are defined later, they arise from a square with emerging rectangles emerging from adjacent edges, and terminating on something - see Figure 51) past the three bounding rectangles, and we know there aren't any absorbable into the cone-cell by minimal area. We then arrive at a rectangular diagram (the third illustration) that is combinatorially a product, and one side of this rectangle lies on our cone-cell along the wall-piece of interest. Now we see that the wall-piece is also a wall-piece in a rectangle (in the fourth illustration) that is contiguous with the cone-cell.

This will enable us to produce examples by limiting the types of wall-pieces that must be examined to those arising from rectangles "based" on $\widetilde{A}$.

Moreover, by minimal area, we know that none of the squares in this contiguous rectangle can be absorbed into the cone-cell. This explains why we can ignore hyperplanes H that cross $\widetilde{A}$. Indeed, a fundamental property of CAT(0) cube complexes is that a hyperplane cannot "interosculate" with a convex subcomplex (see Lemma 3.10). Thus, if the hyperplane H of the dual curve of a contiguous rectangle actually crossed $\widetilde{A}$, then each of the squares in this contiguous rectangle would be contained in $\widetilde{A}$, and hence there would be squares absorbable into the cone-cell and a violation of the minimality of area.

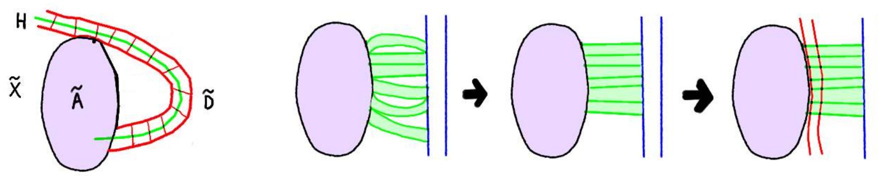

FIGURE 18.

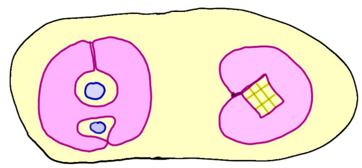

FIGURE 19.

A *cone-piece* of B in A is a path $P \to \mathcal{P}$ in some abstract piece, where $B = Y_j$ for some j. Likewise a *wall-piece* of H in A is a path $P \to \mathcal{P}$ in some abstract wall-piece. We will often refer to *cone-pieces in* A to mean a $P \to A$ without any reference to the various choices of $B, S, \mathcal{P}$, but there is always some possibly variable choice, and similarly we refer to *wall-pieces in* A. We will almost always assume that our pieces are *nontrivial* in the sense that they have length ≥ 1. We will sometimes use the word *piece* to mean either a cone-piece or a wall-piece, and likewise for *abstract piece*.

3.4. **Some small-cancellation conditions to bear in mind.** The main hypothesis leading to a useful theory is an appropriate bound on lengths of pieces in each Y_i. There are various formulations: Let $|Y_i|$ denote the infimum of the lengths of essential closed paths in Y_i. Let diameter($\mathcal{P}$) denote the diameter of an abstract piece $\mathcal{P}$. Let $\alpha > 0$ be a real number. The *absolute* $C'(\alpha)$ *condition* requires that diameter($\mathcal{P}$) $< \alpha|Y_i|$ for each abstract piece $\mathcal{P}$ in each Y_i.

A more general (contextual) $C'(\alpha)$ *condition* requires that $|P|_{Y_i} < \alpha|R_i|$ whenever R_i is an essential cycle in Y_i that contains a piece P as a subpath. (Here we use $|P|_{Y_i}$ to denote the distance in $\widetilde{Y_i}$ between the endpoints of $\widetilde{P}$; thus the above inequality is implied by diameter($\mathcal{P}$) $< \alpha|R_i|$ when $P \to Y_i$ factors through $\mathcal{P}$.) We will employ the absolute condition merely to verify the contextual condition. This contextual condition differs from the absolute condition both by focusing on pieces (which are paths) instead of abstract pieces, and by measuring pieces against cycles they occur in, instead of measuring them against arbitrary cycles.

The *combinatorial* $C(n)$ *condition* requires that no essential closed path in Y_i is the concatenation of fewer than n pieces.

Remark 3.3 (Graded Theory)**.** We will develop a *graded small-cancellation theory* and associated metric small-cancellation conditions later in Sections 3.19 and 3.20.

3.5. **Disk diagrams and cancellable pair removal and absorption.** Let $P \to X$ be a closed path that is nullhomotopic in X^*. Then we can consider various choices $\psi : (D, P) \to (X^*, X)$ of disk diagrams in X^* with boundary path P. Note that the 2-cells of D are either squares of X, or are triangles in some cone C_i of X^*. Moreover, since $P \to X$ avoids cone-points, each 0-cell a_i mapping to a cone-point c_i is internal to D, and hence the triangles adjacent to a_i (which must map to C_i) are grouped in cyclic collections meeting around a_i to form a subspace A_i that is a cone on its bounding cycle. We refer to each a_i as a *cone-point of* D. Note that $\partial A_i \to X$ factors through $Y_i \to X$. We refer to A_i as a *cone-cell* of D.

Remark 3.4. We caution the reader that ∂A_i might not embed, and moreover, A_i might not be an actual (sub)disk diagram. See Figure 19. However, by performing a very simple reduction, we can assume

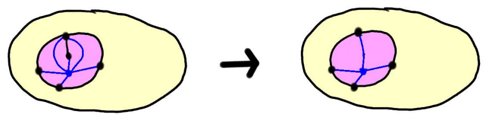

FIGURE 20.

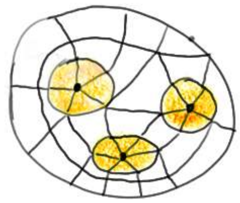

FIGURE 21. Cone-cells in $D \to X^*$

that the boundary path of A_i has no internal backtracks (see Figure 20). In the presence of adequate small-cancellation conditions, and minimal complexity properties of D, this will be the case, but needs some verification. We will pursue this later in Section 4.1. Without small-cancellation conditions on X^*, this holds under the assumption of minimal square area of D unless a non boundary component of $D - A_i$ contains a cone-cell.

We can choose D to satisfy the following successive minimal complexity requirements, so that the pair of numbers: (Conepoints(D), Squares(D)) is minimal in the lexicographical order.

(1) $|\psi^{-1}(\{c_i\})|$ is minimal among all possible disk diagrams with boundary path P,
(2) the number of squares in D is minimal among all possible choices with a minimal number of cone-points.

An immediate consequence of the first condition is that each $\partial A_i \to X$ is essential (and equivalently, each $\partial A_i \to Y_i$ is essential) for otherwise we could find a square diagram $A_i' \to X$ with the same boundary path, and substitute A_i' for A_i and thus decrease the number of cone-points in our disk diagram.

Remark 3.5 (Generalized cancellable pairs)**.** A "classical" cancellable pair in a disk diagram consists of a pair of 2-cells meeting along a 1-cell in a "mirror image" fashion. They are removed together with an open arc where they meet, and this is "zipped up" so only a "surgical scar" is left behind. In our more general situation, this is broken up more explicitly into two steps: Combining the two cone-cells to a single cone-cell, and then possibly replacing this cone-cell by a square disk diagram. For a cancellable pair, this disk diagram has zero area. In general, the new cone-cell might not be replaceable by a square diagram since its boundary path might be essential in the cone. Moreover, even if it is replaceable, it will usually require some squares. The boundary path of the new cone-cell arising in the classical case is a nullhomotopic path in a circle. However in our generalization, even in the case where cones are circles, we might combine two cone-cells wrapping around the same circle p and q times, so the boundary of the new cone-cell might be essential as it wraps around $(p - q)$ times.

We have illustrated these maneuvers in Figure 22. While a more general version of this is presented in Lemma 3.6, this "contiguous combination of cone-cells" is the case that we are usually concerned with.

Lemma 3.6. *Suppose that A_i and A_i' map to the same cone Y_i, and suppose there is an embedded path $P \to D$ whose endpoints are on ∂A_i and $\partial A_i'$, such that $P \to X$ is homotopic in X to a path $P' \to X$ that factors through Y_i. Then we can perform "surgery along P" to produce a lower complexity disk diagram with the same boundary path.*

Remark 3.7. While Lemma 3.6 still holds if P self-intersects around a square subdiagram as on the right of Figure 23, it may fail if P self-intersects around a subdiagram containing other cone-cells, as

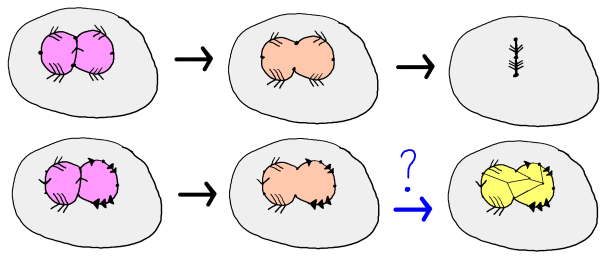

FIGURE 22. A pair of cone-cells can sometimes be combined and then possibly replaced by a square diagram.

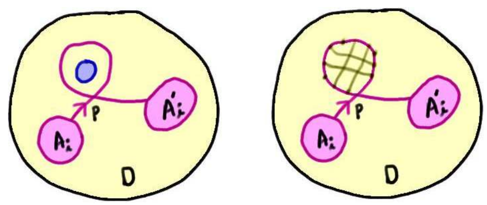

FIGURE 23. The path P wraps around a cone-cell in the diagram on the left, so we cannot combine A and A'. There is no problem on the right, as P wraps around a square diagram.

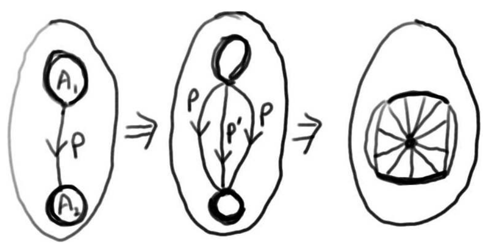

FIGURE 24. Combining Cones

in the diagram on the left. This second type of self-intersection cannot exist in a minimal complexity disk diagram under the small-cancellation condition we will examine below.

Proof. Let $K \to X$ be a square disk diagram for the path homotopy between P and P'. (Note that a path homotopy in X implies a path homotopy in Y_i by π_1-injectivity.) Now cut D along P, and insert two copies of K doubled along P'. Consider the subdiagram $B = A_i \cup P_i \cup A'_i$. Observe that we can cut along ∂B, and substitute a cone from a single cone-point associated with Y_i. We have thus reduced the complexity. □

Remark 3.8 (Not square homotopic)**.** When the path P of Lemma 3.6 is only homotopic into Y within X^*, then the replacement provides new cone-cells in the new diagram as in Figure 25. We have drawn a picture of a *drum* which is a thickened disc that has a square ladder around the outside, and disk diagram with cones on either membrane. The new diagram is obtained from the old by placing the drum along the connecting square diagram, and then pushing upwards through the drum. See Figure 25.

Assign a linear *ranking* to the cones $Y_1, Y_2, \dots$ and let X^*_r denote $\langle X \mid Y_1, \dots, Y_r\rangle$. Suppose that any (hyperplane) path $P \to \widetilde{X}^*$ that start and end on the same lift of some Y_i is actually path homotopic with a path $P' \to Y_i \subset \widetilde{X}^*$ where the path homotopy actually occurs in X^*_{i-1}.

Under this hypothesis, we can use a *ranked complexity* on disk diagram that: assigns a rank of 0 to each square; counts cone-cells according to the ranks of their cones; and uses the ordering where higher rank cells have the greatest value. In this case, the replacement we have described actually decreases complexity.

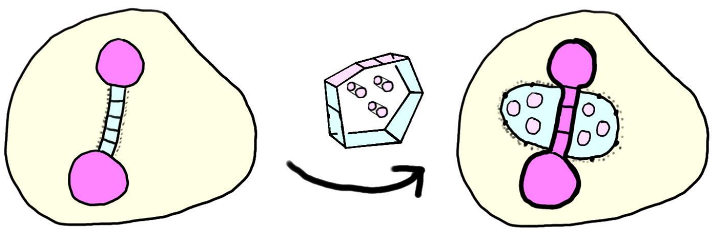

FIGURE 25. Pushing across a drum to obtain a graded replacement.

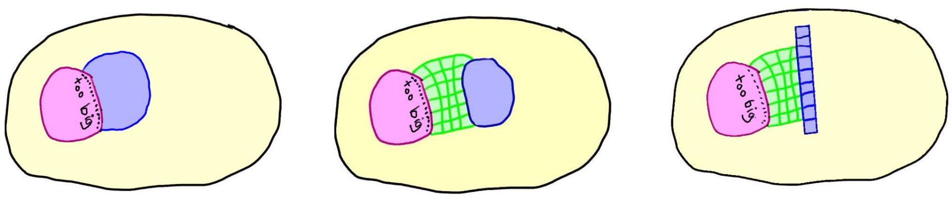

FIGURE 26. Too large to be real pieces. The left case is similar to the classical small-cancellation theory. It is a degenerate version of the middle case, which is itself covered by the rightmost case.

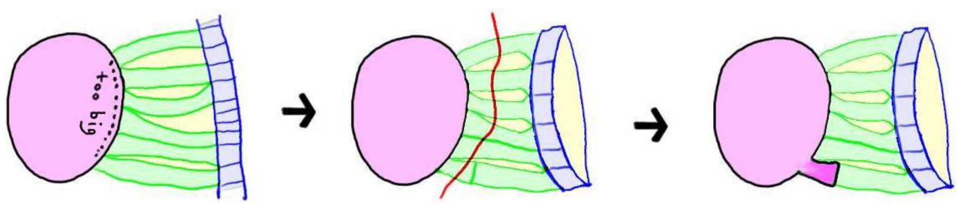

FIGURE 27. Big rectangle-piece yields an adjacent hyperplane with big piece, which must then cross $\widetilde{Y}_i$, and hence yields an absorbable square.

We can combine and possibly even remove "cancellable-pairs" of cone-cells that "meet" along an impossibly long piece as in Figure 26. Indeed, even when these cone-cells are not contiguous, when the piece between them is impossibly long, a rectangle joining them has a horizontal path homotopic into Y_i, and actually, the two cone-cells map to the same cone Y_i and lift to the same translate of $\widetilde{Y}_i$. We are thus able to apply Lemma 3.6, to reduce the number of cone-points which equals the number of cone-cells. We shall also deal with a cone-cell that has an impossibly long piece with a hyperplane, as on the right of Figure 26. In this case, some squares of this rectangle can be absorbed into the cone and so, again, the complexity of the disk diagram can be reduced. We note that this latter situation does not arise in the classical small-cancellation theory (since X is 1-dimensional).

While Lemma 3.6 can be used to reduce the complexity of a diagram whenever two cone-cells share an impossibly big piece, we must use a somewhat different argument to reduce the complexity when a cone-cell has an impossibly big wall-piece. In this case, we first rearrange the square part of the diagram so that our original rectangle is pushed pass any corners of generalized squares that lie between it and the cone-cell (this uses Lemma 2.8). We then find that the piece of our original rectangle actually lies in a piece of some other rectangle whose dual curve is dual to a 1-cell adjacent to our cone-cell. It is thus on such pieces that we place small-cancellation hypotheses. This discussion is described more formally in the following:

Lemma 3.9 (Contiguous wall-pieces dominate)**.** *Let D be a diagram, and let A be a cone-cell in D, and let S be a rectangular strip of D, let $Q = q_1q_2 \dots q_k$ be a path on ∂A, and suppose each rectangular strip R_i starting at the edge q_i of Q ends at a square of S, so that R_1 ends at the first square, and R_k ends*

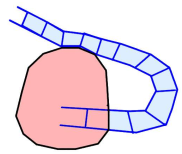

FIGURE 28. The subcomplex above interosculates with a hyperplane.

at the last square. Suppose that the squares of R_i, R_j are distinct for $i \neq j$. Let E be the subdiagram bounded by R_1, Q, R_k, and S, and suppose E is a square diagram.

There exists a new diagram E' with the same boundary path as E, such that E' contains a rectangular strip T' whose first square lies on q_1, and such that each rectangular strip Q_i' emerging from q_i passes through T'.

Proof. See Figure 27. We push the strip across any crossing dual curves of E that cross S and the dual curve of either Q_1 or Q_k. This is done in Lemma 2.8 where the path P (there) corresponds to a path U_1QU_k where U_1, U_k are external paths of R_1, R_k.

This gives us a new diagram E' and a new strip S'. The strip T' emerging from the square adjacent to q_1 has the desired property. □

We thus see that any hyperplane is behind a hyperplane touching the cone-cell. Consequently, if noncrossing hyperplanes have bounded projections, we get a bound on sizes of pieces unless there is a square absorption.

The following result rules out the behavior in Figure 28. It is proven along the lines of Lemma 2.5:

Lemma 3.10 (No Interosculating Hyperplanes)**.** *Let $\widetilde{Y} \subset \widetilde{X}$ be a convex subcomplex. Let U be a hyperplane that* osculates *with $\widetilde{Y}$ in the sense that it has a dual 1-cell with exactly one 0-cell in $\widetilde{Y}$. Then U cannot also be dual to a 1-cell that is contained in $\widetilde{Y}$.*

Consequently:

Lemma 3.11. *Let A be a cone-cell of D that maps to the cone Y. Suppose that D contains a square S with an edge e on ∂A, such that in a lift to $\widetilde{X}$ of S with $\widetilde{e}$ on $\partial\widetilde{A}$, an adjacent 1-cell $\widetilde{e}'$ of $\partial\widetilde{S}$ is dual to a hyperplane U that is also dual to a 1-cell that lies in $\widetilde{Y}$. Then $\widetilde{S}$ lies in $\widetilde{Y}$.*

Consequently, the square S can be absorbed into A to reduce the complexity of D.

Remark 3.12 (Alternate domination)**.** Let $\widetilde{A}, \widetilde{B}$ are convex subcomplexes of $\widetilde{X}$. Let γ is a geodesic between $\widetilde{A}, \widetilde{B}$ in the sense that it is a minimal length path whose endpoints are on $\widetilde{A}, \widetilde{B}$. For each 1-cell e traversed by γ, let U_e be the hyperplane dual to e, and let $N_e = N(U_e)$ denote the carrier of U_e. Then U_e separates $\widetilde{A}, \widetilde{B}$. Moreover, $\mathsf{WProj}_{\widetilde{X}}(\widetilde{B} \to \widetilde{A})$ is contained in $\mathsf{WProj}_{\widetilde{X}}(N_e \to \widetilde{A})$. The second claim from the first by noticing that each square ladder from $\widetilde{B}$ to $\widetilde{A}$ must contain a 1-cell dual to U_e.

By choosing e to be a 1-cell of γ that has a 0-cell on $\widetilde{A}$, we see again that contiguous wall-pieces bound the noncontiguous rectangle-pieces and wall-pieces.

3.6. **Rectified disk diagram to analyze the structure.** Consider the disk diagram $(D, P) \to (X^*, X)$. An *open cone-cell* A of D associated to some cone-point a is the union of open cells whose closures intersect a. Accordingly, A is the union of a together with a sequence of open 1-simplices and 2-simplices cyclically arranged about a. The *external boundary* of $I_n \times [-1, 1]$ is $I_n \times \{\pm 1\}$, and its *initial* and *terminal* 1-cells are $\{0\} \times [-1, 1]$ and $\{n\} \times [-1, 1]$. A *rectangle* of D is a combinatorial map $I_n \times [-1, 1] \to D$ that is injective except perhaps at the external boundary. We note that a rectangle contains (part of) a *dual curve* of D at $I_n \times \{0\}$, and that the above injectivity requirement implies that this dual curve embeds. Unless specifically indicated, we shall always assume that a rectangle is *nondegenerate* in the sense that $n \geq 1$.

We shall now assume that D is *nontrivial* in the sense that it doesn't consist of a single 0-cell. Thinking of D as embedded in S^2, we will regard the 2-cell at infinity A_∞ as the *cone-cell at infinity*.

We now assume there is a *linear ordering* on the cone-cells of D, and we shall assume that A_∞ is last, so for instance, we can label the cone-cells $A_1, \ldots, A_m, A_\infty$. Any choice will be adequate for our purposes, though distinct choices may lead to somewhat different rectified cell structures for D. We will also choose a linear ordering on the 1-cells in the attaching map of each cone-cell. Again, different choices will lead to slightly different results, but for instance, a first 1-cell together with a counterclockwise ordering is adequate.

Combining these choices and using the lexicographical ordering, we obtain a linear ordering on the set $\mathcal{S}$ of 1-cells in the attaching maps of all the cone-cells. (This is essentially an ordering on a subset of the 1-cells of D except where both sides of a 1-cell lie on a cone-cell, in which case it will not play an important role since it will only yield a degenerate rectangle below.) We use these orderings to determine which rectangles of D are *admitted*: Beginning with the first 1-cell in our ordering, we apply the following procedure to each 1-cell e as we proceed through the sequence. Traveling outwards from e away from its cone-cell, there is a maximal rectangle consisting of a sequence of distinct squares. This rectangle will either terminate at another 1-cell on the boundary of some cone-cell (possibly the same cone-cell that contained its initial 1-cell e), or it will terminate on the external boundary of some rectangle that was previously admitted, or it will terminate on its own external boundary, specifically, on the boundary of one of its own squares that had appeared earlier in the sequence. Following our ordering of 1-cells, we proceed in this way until each 1-cell in $\mathcal{S}$ lies in a (possibly degenerate) rectangle. We note that a degenerate rectangle may arise when either our 1-cell appears in two ways among the attaching maps (so it has a cone-cell on each side) or when it has a cone-cell on one side, and a square belonging to a previously admitted rectangle on the other side. In this case, we do not add any rectangle, but we will continue to refer to such a 1-cell as a degenerate rectangle. Each admitted rectangle has a linear orientation directed from its initial 1-cell to its terminal 1-cell.

We now describe a *rectified disk diagram* $\bar{D}$ that we will use to study D. Let $I_n \times (-1, 1)$ denote the *internal part* of the (possibly degenerate) rectangle $I_n \times [-1, 1]$. Let E denote the subspace of D consisting of the union of each open cone-cell and the internal part of each admitted rectangle. Note that the internal part of a degenerate rectangle is an open 1-cell.

Remark 3.13 (Back to the Future)**.** Under sufficiently strong small-cancellation conditions which we will later examine, we will be able to see that the components of $D - E$ are disk diagrams. Specifically, we must rule out the possibility that some component is not simply-connected. If there is a non-simply-connected component then there is an *innermost* one, in the sense that it doesn't separate the boundary of the diagram from some even smaller component. This innermost component would bound a rectified disk diagram D' (containing at least one cone-point) with E' as above, such that $D' - E'$ has simply-connected components. The boundary of D' has a particularly restrictive nature. Our assumption of minimal complexity of D (which implies minimal complexity of D') together with the small-cancellation conditions we will impose, will imply by Theorem 3.40 that D' does not exist.

The subdiagram D' has the property that it contains at least one cone-point, and since it is innermost, $D' - E'$ has simply-connected components, and so the construction (that we are in the midst of) can proceed. However, the boundary of D' consists entirely of subpaths of external boundaries of rectangles. See Figure 29 for some possibilities. The conclusion of Theorem 3.40 proscribe this, since there is neither a θ-shell, nor the outerpath of a generalized square.

For each trivial component (meaning a single 0-cell) of $D - E$ we have a 0-cell in our rectified disk diagram $\bar{D}$. For each nontrivial component F of $D - E$ we have a cycle in our rectified disk diagram that

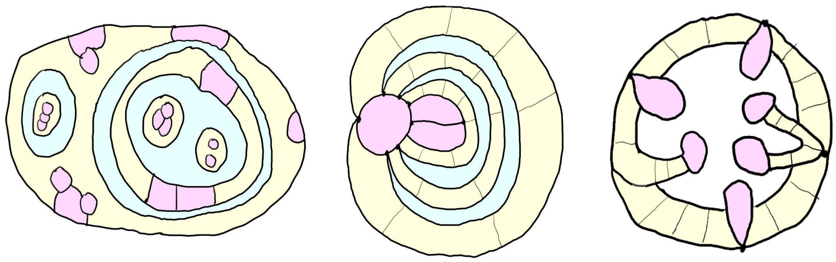

FIGURE 29. Two innermost non-simply-connected fragments on the left a fairly simple possibility for D' in the middle, and a partially illustrated more complicated possibility on the right.

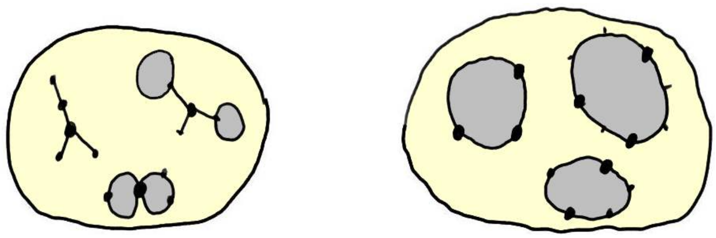

FIGURE 30. Nontrivial singular diagrams of $D - E$ "inflate" to the shard 2-cells of $\bar{D}$.

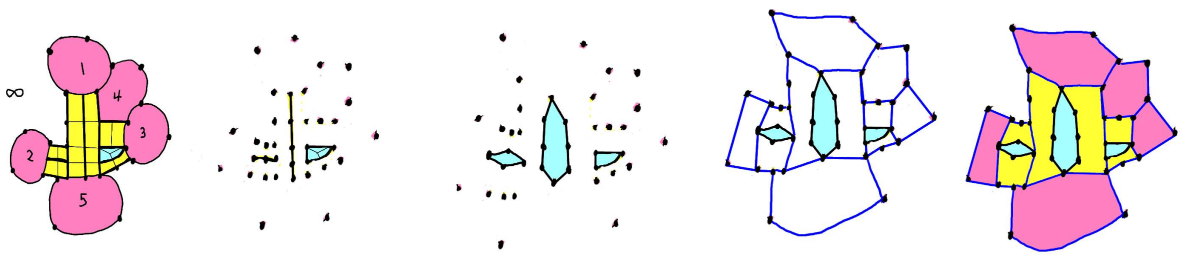

FIGURE 31.

is an embedded copy of $\partial_p F$. In a certain sense the "0-skeleton" of $\bar{D}$ is the disjoint union of boundary paths of components of $D - E$.

For each rectangle, its initial and terminal 1-cells are open 1-cells of $\bar{D}$. These are identical for a degenerate rectangle.

For each nontrivial component F of $D - E$ we add an open 2-cell $\bar{F}$ called a *shard*. The boundary path of F is a combinatorial path in D, and we first add a corresponding circle consisting of new 0-cells and 1-cells to $\bar{D}$. We then attach an open 2-cell $\bar{F}$ along this circle. Thus, an open 1-cell e of D (that is not an initial or terminal 1-cell of an admitted rectangle) will contribute zero, one, or two open 1-cells to $\bar{D}$ according to the number of ways it lies along the external boundary of an admitted rectangle. A 0-cell v of D will contribute d distinct 0-cells of $\bar{D}$ provided that v lies in d ways on the external boundaries of admitted rectangles. The process of inflating shards is illustrated in Figure 30.

We emphasize that the boundary path of each nontrivial shard is the concatenation of external subpaths of admitted rectangles.

We then add the admitted rectangles (in order of their admission) and then the cone-cells.

We note that there is a map $\bar{D} \to D$ that is combinatorial on the 1-skeleton, on the rectangles and on the cone-cells (after subdivision), but is not combinatorial on the shards (since the image of a shard can be a singular subdiagram of D).

We refer to Figure 31 for a diagram D, the complement $D - E$, the "0-skeleton" of $\bar{D}$, and then its "1-skeleton", and "2-skeleton".

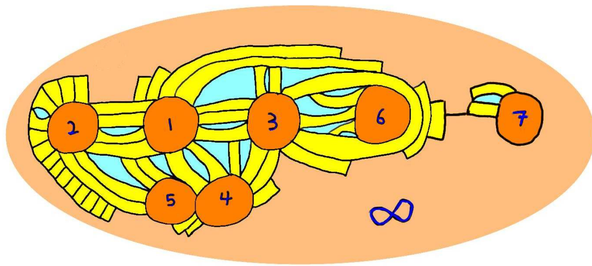

FIGURE 32. The rectified diagram: cone-cells in orange, rectangles in yellow, shards in blue. Shards corresponding to components containing no squares are not indicated.

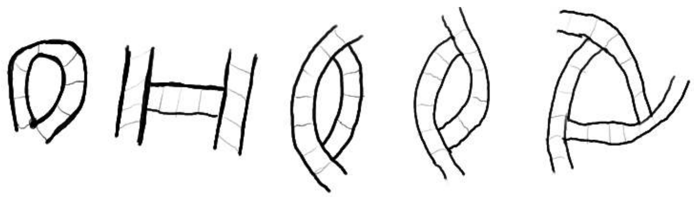

FIGURE 33. Some impossible configurations of admitted rectangles. The first diagram is impossible under the assumption that there are no cone-cells inside and this square diagram has minimal area. The other diagrams are impossible because of the logic of admitted rectangles.

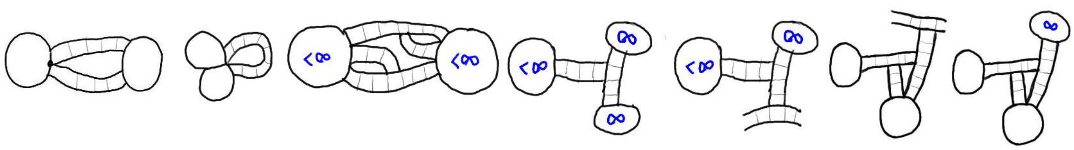

FIGURE 34. More impossible configurations of admitted rectangles and cone-cells. The first diagram is missing some admitted rectangles. The second diagram is impossible assuming that the bounded region is a shard and has minimal area.

We close this section by describing some configurations that cannot arise in a rectified diagram $\bar{D}$. Some impossible configurations are illustrated in Figure 33. The leftmost such configurations excludes many other cases from consideration. The rightmost configurations are impossible because of the recursive ordered way in which we admitted rectangles. Additional impossible configurations are illustrated in Figure 34.

Remark 3.14 (Genus 0 shards)**.** Our focus is on a situation where all shards are simply-connected, and according to Remark 3.13 this is guaranteed in the setting of a minimal complexity diagram under small-cancellation hypotheses. In general, it seems shards should be treated by adding a genus 0 surface - which may of course have multiple boundary components. It is unclear whether further examination might lead to utility.

3.7. **Gauss-Bonnet Theorem.** Let E be a combinatorial complex embedded in the sphere. Suppose an *angle* consisting of a real number $\sphericalangle(c)$ is assigned to each corner of each 2-cell of E.

The *curvature* of a 0-cell v of E is defined to be

$$\kappa(v) = 2\pi - \sum_{c \in Corners(v)} \sphericalangle(c) - \pi\chi(\text{link}(v)). \tag{1}$$

We note that when E is locally a surface without boundary at v then link(v) is a circle and so the final correction term vanishes yielding $\kappa(v) = 2\pi - \sum_{c \in Corners(v)} \sphericalangle(c)$. Similarly, when E looks like a surface with boundary at v, we have $\kappa(v) = \pi - \sum_{c \in Corners(v)} \sphericalangle(c)$. The *curvature* of a 2-cell f of E with $|f|$ sides is the sum of the angles of the corners of f minus the expected Euclidean angle sum for a Euclidean polygon with the same number of corners:

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

Alternately, letting $\mathrm{defect}(c) = \mathrm{defect}(\sphericalangle(c)) = \pi - \sphericalangle(c)$ we see that

$$\kappa(f) = 2\pi - \sum_{c \in \mathrm{Corners}(f)} \mathrm{defect}(c). \tag{2}$$

A simple computation verifies the following well-known fact lying at the heart of small-cancellation theory (see for instance [MW02]).

Theorem 3.15 (Combinatorial Gauss-Bonnet)**.** *Let E be a finite angled 2-complex embedded in the sphere then:*

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

3.8. **Assigning the angles.** We shall now assign angles to the corners of the 2-cells of $\bar{D}$ aiming to obtain nonpositive curvature at internal 0-cells, at rectangles, and at shards, and also at the cone-cells provided that certain small-cancellation conditions are met. There are actually two main angle assignments we will discuss here: The first is the *split-angling*, where the angle assigned to a corner of a cone-cell will depend upon the neighboring cells at that corner - its personality changes to suit its surroundings. The second angle assignment is the *grade-angling*, where cone-cells are treated a bit more like regular Euclidean polygons. In both cases, small-cancellation conditions we will examine later will provide nonpositive curvature of cone-cells. However, whereas the shards are automatically nonpositively curved in the split-angling, we will have to hypothesize this (later) for the grade-angling.

As we have just indicated, in the grade-angling the angle assignments depend upon a *grading* of the cone-cells, but the reader should focus especially on the case where each cone-cell has grade 6 except for the infinite cone-cell A_∞ whose grade is ∞. The grading is a map from the set of cone-cells to $\mathbb{N} \cup \{\infty\}$, and in practice, this will be induced by a map from the set of generalized relators $\{Y_i\}$ to $\mathbb{N}$.

The rectangles of $\bar{D}$ have the usual Euclidean angles - we assign an angle of $\frac{\pi}{2}$ to each of their constituent squares, and so the four corners at the initial and terminal 1-cells have angle $\frac{\pi}{2}$ and all other corners have angle π.

The reader should think of the ∞ cone-cell A_∞ as having an angle of π at each of its corners.

Split-angling assignment to cone-cell corners: We now describe the split-angling. All internal corners of cone-cells are assigned an angle of $\frac{\pi}{2}$ with several exceptions that we list and illustrate in detail below. Though there might appear to be a dizzying array of cases, they are actually simple degenerations and variations of several very natural choices modeled on familiar Euclidean scenarios, and we refer the reader to Figure 35:

Consider a pair of adjacent 1-cells on the boundary of a cone-cell. We say the pair of associated admitted rectangles *end in parallel on an admitted rectangle*, if the (possibly degenerate) subdiagram bounded by these rectangles is a shard: In particular, there is no cone-cell inside it. (It is possible that the terminal 1-cells of the rectangles are not adjacent but they will be in the following case.) The pair of admitted rectangles *end in parallel on a cone-cell* if there is a shard bounded by these two rectangles, and the edges they end on are adjacent to each other. This includes the degenerate case where one or both rectangles are degenerate. The qualifier *implicitly* indicates a related situation where one of the rectangles travels though a square that (at least locally) could have been the continuation of the target rectangle (see Figure 37).

An angle of $\frac{\pi}{2}$ is assigned to cone-cell corners except for the following cases:

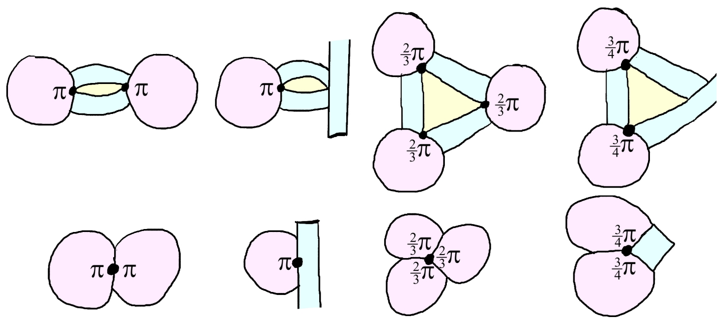

FIGURE 35. Internal cone-angles for the split-angling are modeled by the four cases above: All variants of the four cases above are provided in Figures 36, 37, 38, and 39.

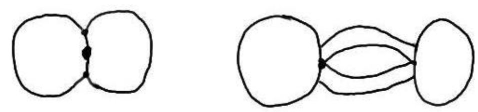

FIGURE 36. Parallel rectangles ending on the same cone-cell give π angles at cone-cell corner.

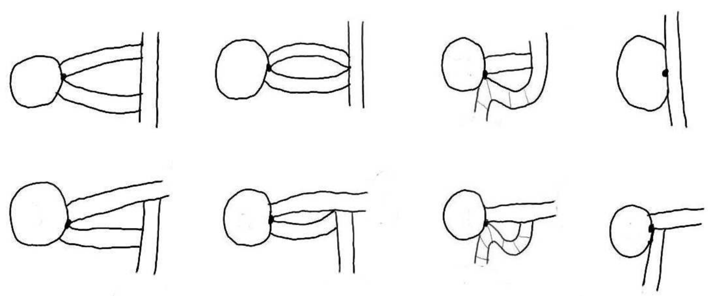

FIGURE 37. Rectangles (implicitly) parallel to the same admitted rectangle have π angle at the cone-cell corner.

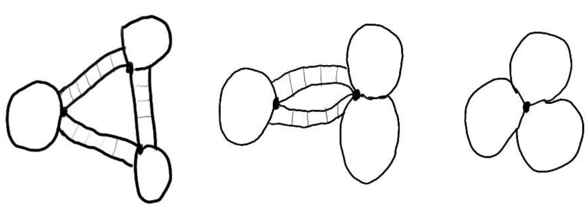

FIGURE 38. Three adjacent cone-cells.

(1) π when the associated admitted rectangles end in parallel on the same cone-cell. (See Figure 36.)
(2) π when the associated admitted rectangles (implicitly) end in parallel on the same admitted rectangle. (See Figure 37.)
(3) $\frac{2}{3}\pi$ when they end on a pair of finite cone-cells that are also adjacent along an admitted (possibly degenerate) rectangle, such that these three bound a (possibly degenerate) triangular shard. (See Figure 38.)
(4) $\frac{3}{4}\pi$ when they (implicitly) end on a rectangle-cone-cell combination (See Figure 39.)
(5) We emphasize that $\frac{\pi}{2}$ is assigned in the above two cases when either of the other cone-cells is infinite. (See Figure 40.)
(6) 0 is assigned at a corner when the rectangles terminate at a singular vertex on C_∞. (See Figure 41.)

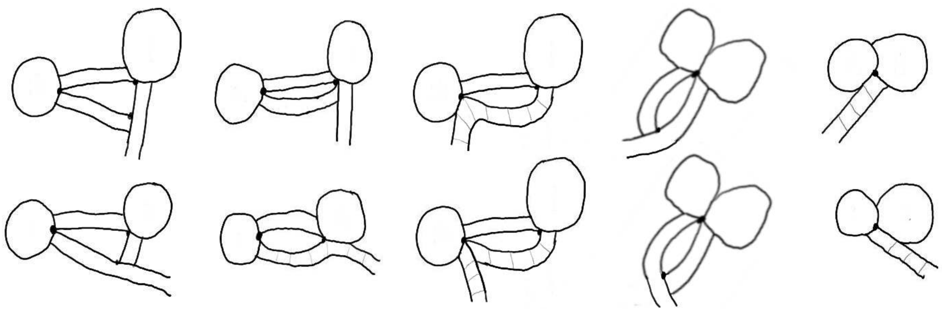

FIGURE 39. $\frac{3}{4}\pi$ is assigned to the internal angle between edges whose emerging rectangles end in parallel on a cone-cell/rectangle pair.

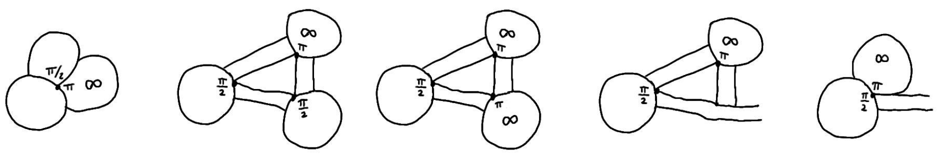

FIGURE 40.

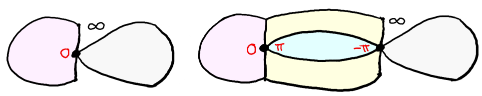

FIGURE 41. An angle of 0 is assigned to a corner facing a singular 0-cell.

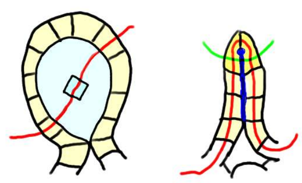

FIGURE 42.

We note that the choice of a 0-angle at a singular vertex is not very critical - we could have allowed $\frac{\pi}{2}$ here without much effect on the theory. In particular Theorem 3.20 would hold equally well with a $\frac{\pi}{2}$ choice here. But the 0-angle permits more general coverage in the hypothesis of Theorem 3.36. Curiously, the bigonal shard in this case takes angles of $-\pi, \pi$.

Remark 3.16 (The grade-angling variation)**.** In the *grade-angling* for a grade g cone-cell, we place internal angles of $\frac{g-2}{g}\pi$ in place of the angles $\frac{2}{3}\pi$ and $\frac{3}{4}\pi$ listed above. Note that this gives an angle of π when $g = \infty$.

In order to guarantee nonpositive curvature at vertices and at shards, we shall later be forced to make certain hypotheses on the possible collections of graded cone-cells and square surrounding a shard.

3.9. **Nonpositive curvature of shards.** In this subsection we will use minimality properties of the diagram D to conclude that shards have nonpositive curvature. In particular we will assume that there are no bigons in the square subdiagrams, and we will assume that there are no corners of generalized squares on cone-cells. By Corollary 2.3 and Lemma 2.5, this holds when D has minimal complexity.

The corners of a shard are of two types: *dull* corners which lie along a pair of edges on the external part of a single admitted rectangle, and otherwise *sharp* corners. The dull corners are not very interesting and we simply assign to them an angle of π.

FIGURE 43. Some simple bigonal shards The second diagram can occur with an infinite cone-cell. The third and fourth diagrams are impossible.

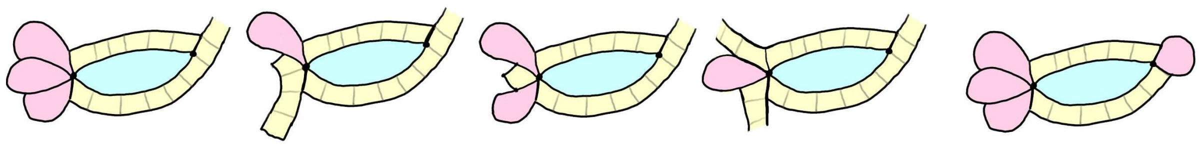

FIGURE 44. Three additional squares/cone-cells at one corner, and one at the other.

The sharp corners of a shard will have an angle of $\frac{\pi}{2}$ except in a few special cases that we discuss: Observe that there is no *monogonal shard* consisting of a nontrivial shard bounded by a single admitted rectangle, for then (see Figure 42) either the shard contains no squares and then there would be a backtrack along its boundary, or there would be some square inside the shard and hence a bigon of dual curves in the nonrectified D, and this would permit a square area reduction by Lemma 2.2.

Consider a *bigonal shard* having exactly two sharp corners and hence exactly two admitted rectangles with external edges along its boundary. The 0-cells at the sharp corners where these rectangles meet already have at least two $\frac{\pi}{2}$ corners (from squares in these admitted rectangles, which have an edge along the sharp corners). If each of these 0-cells has at least two additional "admitted squares" (i.e. from within admitted rectangles) and/or cone-cells, then there is already a total angle sum of at least 2π around each and so we can assign an angle of 0 to each sharp corner of our shard.

It is impossible for one of these corners to have no further cone-cell or admitted square alongside it, for then the two admitted rectangles would be the same (they would continue around that corner), and we would have a monogonal shard - which was excluded earlier.

Suppose each of these sharp corners has exactly one additional square/cone-cell around it. There are essentially four possibilities illustrated in Figure 43, and the third and fourth are impossible. In the first case, we can assign an angle of 0 to each sharp corner, since the corresponding internal corners of the cone-cells are both π. The second case is the most interesting: For a finite cone, it is impossible by minimal complexity of D. Indeed, each $Y_i \to X$ is a local-isometry, and so in this case, we could absorb a square into the cone (without changing the boundary path). In the case where it is an infinite cone A_∞, we assign angles of $+\frac{\pi}{2}$ at the corner opposite the square, and assign an angle of $-\frac{\pi}{2}$ at the corner opposite the cone-cell at infinity. The positive curvature at the latter corner will be important later on, and the situation will be referred to as a *corner of a generalized square*.

Let us now consider the possibility where one sharp corner has exactly one further cone-cell/square, and the other sharp corner has three or more. In this case, we put an angle of $+\frac{\pi}{2}$ on the corner with one additional cone-cell/square, and we put an angle of $-\frac{\pi}{2}$ on the corner with three or more. A non-exhaustive selection of the possibilities is illustrated in Figure 44.

We now consider the possibility of one additional cone-cell/square at one sharp corner, and two additional cone-cells/squares at the other. First let us suppose there is a single additional square at one sharp corner. The various conceivable possibilities are illustrated in Figure 45, but some of these cannot arise. In the (possible) cases we assign an angle of $+\frac{\pi}{2}$ at the left corner of the bigon and $-\frac{\pi}{2}$ at the right corner of the bigon. The bottom four cases do not actually exist: The leftmost three do not exist since an admitted rectangle must have an initial 1-cell on a cone-cell. The rightmost bottom case does not exist because there is no consistent way of ordering the initial 1-cells of these three rectangles. In the top four cases, the internal angle of the corner of the cone-cell at our bigon is π since this is the case of

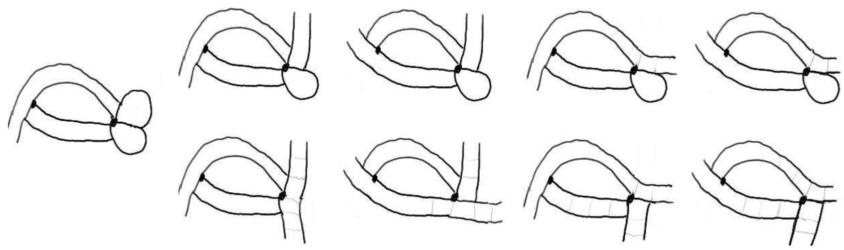

FIGURE 45. Square at one sharp corner, two cells at other.

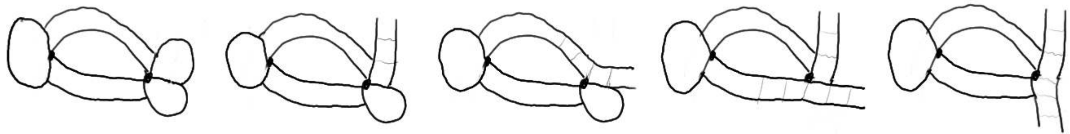

FIGURE 46. Cone-cell at one sharp corner, two additional cells at other.

two rectangles (implicitly) terminating on the same rectangle. We note however that the second case above (from the left) cannot exist because it has a rectangle without an initial 1-cell on a cone-cell.

We now suppose there is a single additional cone-cell at one sharp corner, and two additional cone-cells and/or squares at the other. The five possibilities that can arise are illustrated in Figure 46. In the fourth and fifth cases an angle of 0 is assigned to each sharp corner, as these two cases correspond to a π angle for the corner of the cone-cell since its rectangles (implicitly) terminate on another admitted rectangle. For the split-angling, we assign $\frac{2}{3}\pi$ to each cone-cell corner in the first case and hence use $\pm\frac{\pi}{3}$ for the sharp bigon corners, and in the second and third cases we assign $\frac{3}{4}\pi$ to each cone-cell corner and hence use $\pm\frac{\pi}{4}$ for the sharp bigon corners.

We emphasize that except for the cases discussed above, the corners of a shard receive a $\frac{\pi}{2}$ angle.

Remark 3.17 (Nonpositive shard for grade-angling)**.** For the grade-angling, we assign an angle of $\frac{g_i-2}{g_i}\pi$ to each grade g_i cone-cell corner, and then assign complementary angles to the sharp corners so that the 0-cells have zero-curvature, and hence there is a condition to check on the grades of the cells around this shard which would imply that its curvature is nonpositive. In the first case the condition is $\frac{2}{g_1}\pi + \frac{2}{g_2}\pi + \frac{2}{g_3}\pi \geq 2\pi$. In the second and third cases, the condition is $\frac{2}{g_1}\pi + \frac{2}{g_2}\pi + \frac{2}{4}\pi \geq 2\pi$ (as a square has been substituted for a cone-cell).

We now consider a *triangular shard* which is bounded by three admitted rectangles. As in the bigonal shard case, if one of its sharp corners has two additional cells, then we can assign an angle of 0 to that sharp corner, and an angle of $\frac{\pi}{2}$ to the remaining two corners. (Note that there must be at least one additional cell at each corner.) So, let us assume that each corner has only one additional cell. There are four cases according to whether there are zero, one, two, or three cone-cells.

Case zero, where each of these additional cells are squares is easily seen to be impossible: up to symmetry the locally possible situations are indicated in the first and second diagrams (counting from the left) in Figure 47. In the first there is no possible consistent way of ordering the initial 1-cells of the bounding admitted rectangles, and in the second there is a rectangle with no initial 1-cell on a cone-cell.

The possibilities for the case where there is exactly one cone-cell are illustrated in the third, fourth, and fifth diagrams (counting from the left) in Figure 47. The fifth diagram is impossible. The third and fourth diagrams are cases where the rectangles emerging from adjacent edges of a cone-cell (implicitly) end on the same rectangle. Thus the internal angle of this cone-cell is π, and we can assign an angle of zero to the sharp corner of the shard at the cone-cell, and angles of $\frac{\pi}{2}$ to each of the other two sharp corners.

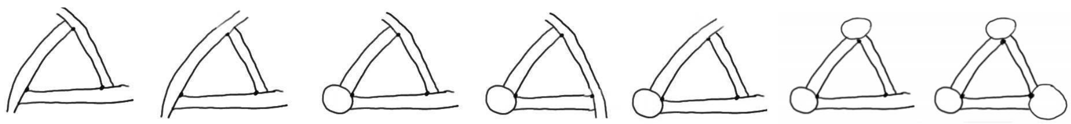

FIGURE 47. Locally possible configurations of triangular shards with a single additional cell at each corner.

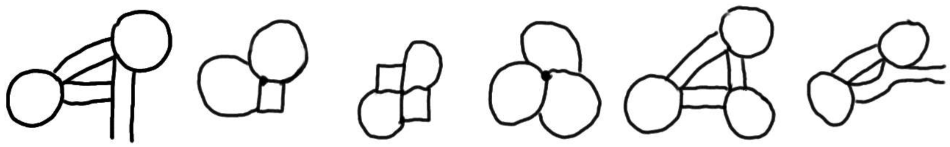

FIGURE 48. In most cases, when there is an infinite cone around a shard, then the shard (or one of its boundary vertices) will have negative curvature depending on whether we use the split-angling or the grade-angling.

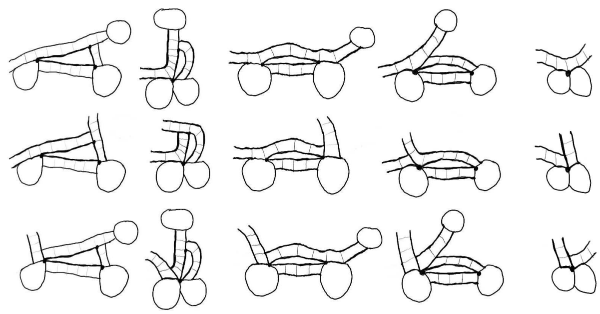

FIGURE 49. Some cases requiring a $-\frac{\pi}{2}$ angle . These figures should only be viewed by readers who are 18 years or older.

The possibilities for the cases where there are two or three cone-cells are illustrated in the sixth and seventh diagrams (from the left) in Figure 47. In the sixth case, in the split-angling, we use $\frac{3}{4}\pi$ for the internal cone-cell corners and $\frac{\pi}{4}, \frac{\pi}{4}, \frac{\pi}{2}$ for the sharp triangular shard corners. In the seventh case, we use $\frac{2}{3}\pi$ for each internal cone-cell corner, and $\frac{\pi}{3}$ for each shard corner.

Remark 3.18. In the grade-angling, we use $\frac{g_i-2}{g_i}\pi$ angles for the cone-cells, and so use $\frac{2}{g_i}\pi$ for the opposing sharp corner of the triangular shard. Accordingly, the condition on the grades resulting in nonpositive for the shard is similar to the one given before.

Remark 3.19. There are a variety of shards with infinite cone-cells at the corners such that the natural angle assignments lead to negative curvature of the shards and/or the vertices. Some of these are described in Figure 48.

3.10. **Tables of Small Shards.** Roughly twenty cases arise in the table in Figure 50. These are organized as they are obtained from the four figures on the left by contracting one or three parts of the rectangles around the triangle. Note that contracting two of these three sides results in rectangle around a monogonal shard, and this is not possible in our minimal square area situation.

One of the two ways that a 0-cell on the boundary of the cone-cell at infinity can have positive curvature is when it is the corner of a square. The other way is more interesting as it involves an opposing bigon with angles $\pm\frac{\pi}{2}$ (see Figure 51). In a certain sense, the first case is a degenerate version of the other case. Under the assumption of local convexity of the cone-cells (except for A_∞) and the assumption of minimal square area, neither of these configurations is possible along a cone-cell other than A_∞.

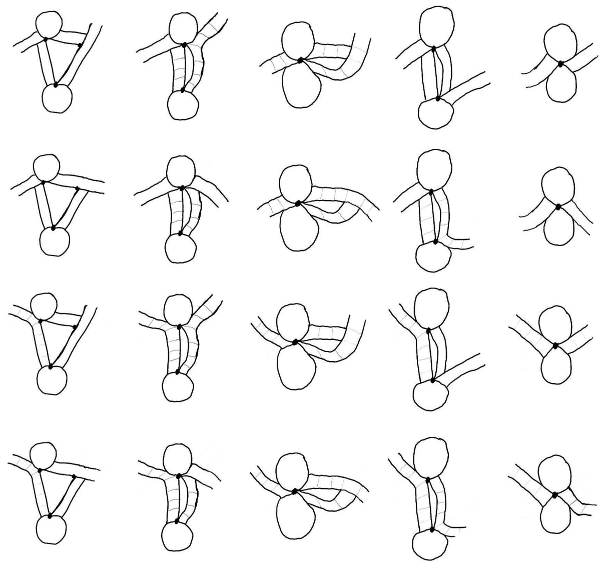

FIGURE 50. Some more cases requiring a 0 or $-\frac{\pi}{2}$ angle

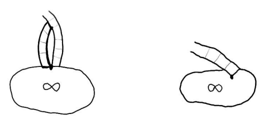

FIGURE 51. Corners of generalized squares: The nontrivial shard on the left requires angles of $\pm\frac{\pi}{2}$ in the bigon. Both cases have curvature $\frac{\pi}{2}$ at the vertex on the boundary of A_∞.

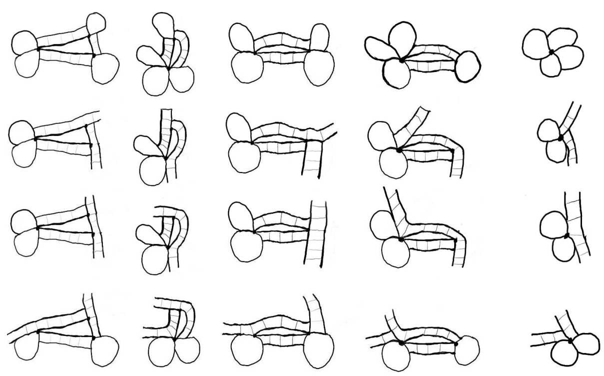

FIGURE 52.

3.11. **Nonpositive curvature of cone-cells via small cancellation.** We now discuss a metric small-cancellation condition on cone-cells in $\bar{D}$ that implies the nonpositive curvature of each of the cone-cells. There are similar but more complex combinatorial conditions that are a bit more general. The conditions are couched in terms of the "wall projections" (within a disk diagram) of cone-cells, rectangles and combinations of these onto a given cone-cell in $\bar{D}$. We refer to these as *cone-pieces* and *rectangle-pieces* in the boundary path of a cone-cell. Each such piece consists of a subpath of the boundary path of the cone-cell, that is a concatenation of edges, all of whose rectangles end in parallel on the same other cone-cell (respectively rectangle). That they end in *parallel* means that they bound a shard (possibly together with the rectangle where they end).

The curvature of a cone-cell with p sides is $\sum_{i=1}^{p} \sphericalangle_i - (p-2)\pi$ or alternatively, it is $2\pi - \sum(\pi - \sphericalangle_i) = 2\pi - \sum \text{defect}(\sphericalangle_i)$.

The internal angle $\sphericalangle$ between 1-cells of a cone-cell is related to the destination of the corresponding rectangles. If they both (implicitly) end on the same rectangle, or on the same cone-cell, then the internal angle is π so its defect is 0. We obtain positive deficiencies when there is a *transition* between the ends of these rectangles in the sense that they don't end in parallel on the same cone or implicitly on the same rectangle. Note that in the split-angling, the internal cone-cell angles that are not equal to π are $\frac{2}{3}\pi$, $\frac{3}{4}\pi$, $\frac{\pi}{2}$. In the grade-angling, they are $\frac{g-2}{g}\pi$, $\frac{\pi}{2}$.

Let us first establish a condition that leads to a quick conclusion:

Theorem 3.20. *Suppose that each rectangle-piece P_w in the cone-cell C satisfies $|P_w| \leq \frac{1}{12}|\partial C|$ and each cone-piece P_c satisfies $|P_c| \leq \frac{1}{12}|\partial C|$. Under the split-angling, each internal cone-cell has non-positive curvature. If the inequality is strict, then each internal cone-cell has negative curvature.*

A cone-cell is *internal* if its boundary path does not pass through a 1-cell in the diagram's boundary.

Remark 3.21. While the statement of Theorem 3.20 sets the correct tone, in practice we need the more flexible requirement that $\nabla_Y P < \frac{1}{12}\|\partial C\|_Y$ for each cone-piece or wall-piece P on ∂C. Here $\nabla_Y P$ denotes the distance between the endpoints of the lift $\widetilde{P}$ in $\widetilde{Y}$, or equivalently, the minimal length in the path-homotopy class of $P \to Y$, where Y is the cone supporting the cone-cell C, and $\|\partial C\|_Y$ denotes the length of the shortest closed path in the homotopy class of $\partial C \to Y$.

The metric conditions we treat later in Section 3.20 are used under analogous flexible restatements.

Proof. **Incorrect $\frac{1}{8}$:** The idea of the proof is that to pick up some defect for each transition between distinct pieces. This appears to lead to at least a $\frac{\pi}{4}$ defect for each transition, and hence nonpositive curvature when all pieces have length $\leq \frac{1}{8}|\partial C|$. However, a subtle problem with this argument is that there can be a 0 defect when the transition occurs between a cone-piece and rectangle-piece where rectangles implicitly end on the same rectangle. We refer the reader to Corollary 3.32.(2), Theorem 3.33.(2), and Problem 3.34. We remedy this argument below by grouping pieces together in a certain way to reach the slightly weaker conclusion.

Quick $\frac{1}{24}$: According to Scheme 3.24, it suffices to prove the statement under the assumption that internal cone-cells do not self-collide. In this case, there cannot be three consecutive transitions each of which has defect 0. This follows by observing that there cannot be a defect 0 transition between an incoming rectangle piece and another rectangle piece - for otherwise there would be a self-collision by Lemma 3.25.

If we assume that $|P| \leq \frac{1}{24}|\partial C|$, then there must be at least 24 pieces and hence transitions. Grouping these into triples, we see there are at least 8 such groups. The observation implies that each group has at least one nonzero defect, and hence has total defect at least $\frac{\pi}{4}$, which gives a total defect of $\geq 2\pi$.

Proof $\frac{1}{12}$: The $\frac{1}{24}$ argument employed a bound on lengths of cone-pieces and rectangle-pieces that explicitly arise in the rectified diagram $\bar{D}$. To obtain the $\frac{1}{12}$ result we use a slightly stronger interpretation of the hypothesis. Namely, there is a $\frac{1}{12}|C|$ bound on all cone-pieces and rectangle-pieces on C *occurring in D itself.* Such a piece would arise in some other rectification.

There is an orientation on each rectangle-piece determined by the orientation of the rectangle. By Scheme 3.24, it suffices to prove the Lemma under the assumption that internal cone-cells do not self-collide. Consequently we can assume the angle defect between successive rectangle-pieces is $\frac{\pi}{2}$, for otherwise Lemma 3.25 provides a self-collision.

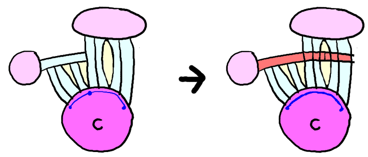

FIGURE 53. An explicit rectangle piece extends to include a cone piece.

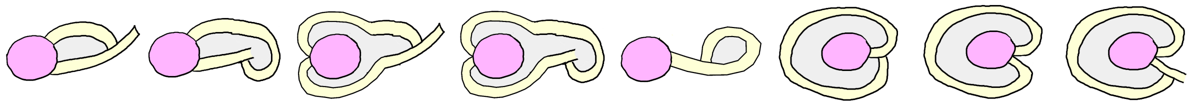

FIGURE 54. Diagram $\bar{D}_C$ obtained from the self-colliding cone-cell C.

The angle defect between an outgoing rectangle-piece and a cone-piece is $\geq \frac{\pi}{4}$. If the angle defect between an incoming rectangle-piece and a cone-piece is 0, then the cone-piece must actually be contained in the rectangle-piece, otherwise the angle defect is $\geq \frac{\pi}{4}$ there as well. The angle defect between consecutive cone-pieces is $\geq \frac{\pi}{3}$.

Accordingly, we associate $\frac{\pi}{6}$ to the outgoing vertex of each rectangle-piece. For a cone-piece, either it contains angle defects of one of $\frac{\pi}{4}, \frac{\pi}{2}, \frac{\pi}{3}$ on each side and we associate $\frac{\pi}{12}$ to it from each side (the most interesting case here is $\frac{\pi}{4} = \frac{\pi}{6} + \frac{\pi}{12}$), or it has an angle of 0 on one side, in which case it lies in an extension of the incoming rectangle-piece on that side and we include it as part of that (extended rectangle-piece). If this happens on each side of the cone-piece, then we include it on just one these sides. We illustrate this extension in Figure 53.

Thus ∂C is the concatenation of rectangle-pieces with associated angle defects of $\frac{\pi}{6}$, and cone-pieces with associated angle defects of $\geq \frac{\pi}{6}$.

If each such piece has length $\leq \frac{1}{12}|C|$ then there are at least 12 such pieces, and hence a total defect of $\geq 2\pi$. The argument for negative curvature is similar, as there are at least 13 such pieces, and hence a total defect of $\geq \frac{13}{6}\pi$. □

Let now now consider the grade-angling: Observe that if a grade $g \neq \infty$ cone-cell has at least g transitions, then the defect sum is $\geq 2\pi$ and it has nonpositive curvature.

Proposition 3.22. *Consider the grade-angling on $\bar{D}$. Suppose that for each g_1, g_2, g_3 grade cells meeting around a vertex, bigonal shard, or triangular shard, we have $\frac{1}{g_1} + \frac{1}{g_2} + \frac{1}{g_3} \leq 1$ (where $g_3 = 4$ for a square). Then each 0-cell has nonpositive curvature.*

Suppose that $|W| \leq \frac{1}{g}|C|$ for each rectangle-piece on the cone-cell C and $|C'| \leq \frac{1}{g}|C|$ for each cone-cell piece on the cone-cell C. Then each cone-cell has nonpositive curvature.

We note that the first hypothesis in Proposition 3.22 holds when each grade is ≥ 8.

3.12. **Internal cone-cells that do not self-collide.**

Definition 3.23 (Self-collision)**.** Let C be a cone-cell in a rectified diagram $\bar{D}$. We say C *self-collides* if there are admitted rectangles having initial or terminal 1-cells on ∂C such that one ends on the external boundary of the other. That C does not self-collision generalizes the idea of 2-cells in a disk diagram having embedded boundary paths.

We refer the reader to Figure 54, were we also include the case of a rectangle starting and ending on C, and a rectangle of C that collides with itself. All these self-collisions are excluded by the following:

Scheme 3.24. Consider a rectified disk diagram $\bar{D}$ arising from a diagram D with no bigonal square subdiagrams, and no corners of generalized squares on its cone-cells. Suppose angles have been assigned to the corners of $\bar{D}$ so that: all rectangle corners have angle $\frac{\pi}{2}$, nontrivial shards (and internal

0-cells) have nonpositive curvature, cone-cell angles are nonnegative, and singly-external corners have angle $\frac{\pi}{2}$.

Then all internal cone-cells of $\bar{D}$ have nonpositive curvature if and only if this holds for non-self-colliding internal cone-cells.

If the internal cone-cells of $\bar{D}$ that do not self-collide have nonpositive curvature, then no cone-cell of $\bar{D}$ self-collides.

Note that the closure of an internal cone-cell could contain 0-cells of ∂D.

Proof. An internal cone-cell C of $\bar{D}$ that self-collides determines a diagram $\bar{D}_C$ obtained from the immersed subdiagram subtended by C together with the two emerging colliding rectangles. The angle assignment on $\bar{D}_C$ is induced from the assignment on $\bar{D}$.

We will prove that there is no such C, by considering an example where $\bar{D}_C$ contains a minimal number of cone-cells. The minimality implies that, with the possible exception of C itself, no cone-cells of $\bar{D}_C$ self-collides. Indeed, if B is a cone-cell in $\bar{D}_C$ that is not strongly-embedded then this provides a smaller diagram $\bar{D}_B$, since C is not contained in $\bar{D}_B$. In each case the colliding rectangles of B are unable to encircle C. This is often obvious since C is an external cone-cell of $\bar{D}_C$, but an interesting case to consider here is the 4th diagram in Figure 55.

However, we note that there must be some other cone-cell within $\bar{D}_C$, for otherwise there would be a generalized corner of a square on C within $\bar{D}_C$ and hence within $\bar{D}$, or in the 5th diagram, there would be a self-crossing hyperplane in a square subdiagram.

In the 2nd diagram, $\kappa(C) \leq \pi$ and actually $< \pi$, as the cone-cell B guarantees at least one transition, and there is exactly one corner of generalized square at the opposite side, giving a $\frac{\pi}{2}$ curvature.

In the 3rd diagram, $\kappa(C) < \pi$ likewise, and there is actually also a vertex on the boundary with $\kappa = -\frac{\pi}{2}$.

In the 4th diagram, $\kappa(C) \leq 2\pi$, and moreover, since there is at least one transition around C, we have $\kappa(C) < 2\pi$. Moreover, there are no other positively curved cells, and actually one boundary vertex with curvature $-\frac{\pi}{2}$.

In the 5th diagram, $\kappa(C) = \pi$ and there is also a vertex with curvature $-\frac{\pi}{2}$.

In conclusion, even without having verified that there was an internal cone-cell which gives us the strictness of $\kappa(C)$ in the 2nd, 3rd, and 4th cases, we find that the total amount of identified curvature is $< 2\pi$. As all other cells of the diagram have nonpositive curvature, this yields a contradiction to Theorem 3.15. □

Lemma 3.25 (Cone-cell ordered rectified diagrams)**.** *Let $\bar{D}$ be a rectified diagram induced from an ordering on the bounding 1-cells of cone-cells of D that arises from an ordering of the cone-cells with C_∞ last.*

If the 1st subdiagram of Figure 55 is contained in $\bar{D}$, then it can only arise from the 2nd or 3rd subdiagrams of Figure 55, or their degenerate versions in the 4th and 5th subdiagrams.

Note that we do not require the cone-cell boundary ordering of Lemma 3.25 to also be consistent with a cyclic ordering of the 1-cells around each cone-cell.

Proof. This follows from the logic of the rectified diagram construction of Section 3.6. □

Corollary 3.26. *If a sequence of consecutively adjacent outgoing rectangles implicitly end on the same rectangle, then they explicitly end on the same rectangle. See Figure 56.*

We use the notation introduced in Definition 3.28 for the following:

Remark 3.27. $\overleftarrow{W}_0\overleftarrow{W}$ and $\overrightarrow{W}_0\overleftarrow{W}$, and $\overrightarrow{W}_0\overrightarrow{W}$, and $\overrightarrow{W}_0\hat{C}$, and $\hat{C}_0\overleftarrow{W}_0$ cannot arise because of Corollary 3.26.

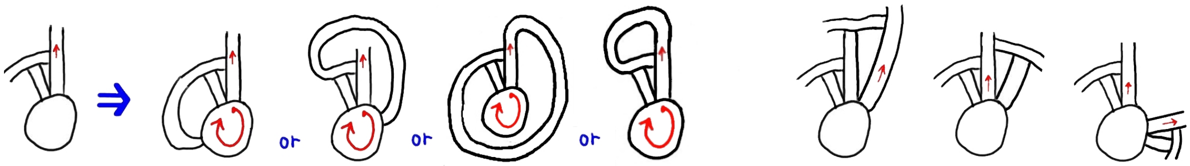

FIGURE 55. Limiting consecutive wall-pieces around internal cone-cell.

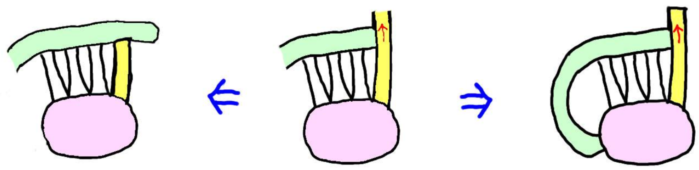

FIGURE 56. Using the argument of Lemma 3.29, if the middle diagram existed, then the hyperplane would be forced to also emerge from the cone as on the right. However, in applications, this is impossible by Scheme 3.24, so all the rectangles end explicitly on the same rectangle.

$\overleftarrow{W}_0\overrightarrow{W}$ cannot arise since rectangles initiate on a cone-cell, not on another rectangle.

$\overleftarrow{W}_{\frac{\pi}{4}}\hat{C}$ and $\hat{C}_{\frac{\pi}{4}}\overrightarrow{W}$ cannot arise following the proof of Corollary 3.26.

3.13. ⊛ **More general small-cancellation conditions and involved justification.** The goal of this section is to give a generalization of Theorem 3.20. This generalization, stated in Theorem 3.31 will not be used in the sequel and serves primarily as an experiment. The statement of the theorem involves notation which we will not use later.

Definition 3.28 (Destination Notation)**.** Consider the sequence of cone-cells and rectangles around a cone-cell. We combine together emerging rectangles with similar destination, and use the following notation $CWCCCWWWW$ for a sequence of cone-cell and rectangle destinations around our cone-cell. We will abuse the notation in the following way: Each C or W will also denote a path in the boundary of our cone-cell that is the concatenation of the initial or terminal edges of rectangles that start or end on the cone-cell C or rectangle W. We can refine this further by using notation $\hat{C}$ and $\check{C}$ to mean that the rectangles are oriented from our cone-cell upwards, or from the destination cone-cell towards our cone-cell. We refine this notation further by using $\bar{C}$ to mean that the rectangles between are degenerate trivial, and use $\tilde{C}$ to denote the cone-cell at infinity. We also use $C\vec{W}$ to mean that $\vec{W}$ emerges from C, and likewise $\overleftarrow{W}C$ to mean that it emerges from the left. Finally, we put numbers between terms: $C_{\frac{\pi}{3}}C$ and $C_{\frac{\pi}{2}}W$ and $W_{\frac{\pi}{2}}W$ and so forth, to indicate the defect of the angle in our cone-cell at the corresponding internal corner where there is a destination transition. We refer to Figure 57 for an example of this notation.

Lemma 3.29. *Suppose that the rectified diagram $\bar{D}$ was constructed using an ordering of the bounding edges of cone-cells that is induced by an ordering of the cone-cells followed by a cyclical order of the 1-cells on their boundary paths.*

Let e_1, e_3 be adjacent 1*-cells on the boundary of some cone-cell C, and their rectangles R_1, R_3 are both oriented outwards from C. Suppose that R_3 ends on a rectangle R_2 which implicitly crosses R_1. Then e_1 is the first 1-cell of ∂C and e_3 is the last.*

Consequently, in general there can be at most one such configuration for each cone-cell C. However, we shall reach stronger conclusions below under the assumption of small-cancellation conditions.

We refer the reader to Figure 55 which indicates the hypothesis situation on the left diagram, and the two (essentially the same) outcomes in the second and third diagrams - and their two subsequent degenerate cases. The last three diagrams indicate configurations that are consequently impossible.

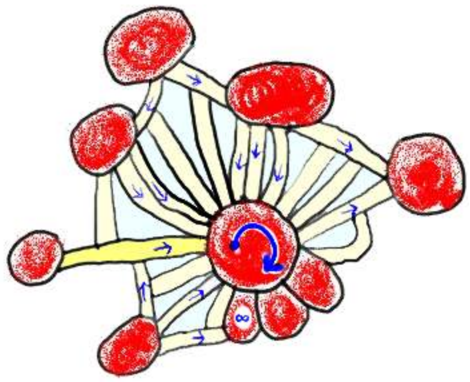

FIGURE 57. Clockwise from the yellow rectangle edge we have:
$\check{C}_{\frac{\pi}{4}}\check{C}_{\frac{\pi}{4}}\overleftarrow{W}_{\frac{\pi}{2}}\overrightarrow{W}_{\frac{\pi}{4}}\check{C}_{\frac{\pi}{4}}\overrightarrow{W}_{\frac{\pi}{4}}\hat{C}_{\frac{\pi}{4}}\bar{C}_{\frac{\pi}{3}}\bar{C}_{\frac{\pi}{2}}\tilde{C}_{\frac{\pi}{2}}\overleftarrow{W}_{\frac{\pi}{4}}\check{C}_{\frac{\pi}{4}}\overrightarrow{W}_0$

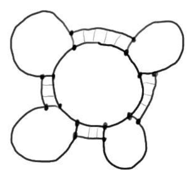

FIGURE 58. A diagram limiting generalizations of Theorem 3.31.

Proof. Let e_2 denote the initial 1-cell of R_2, and observe that $e_1 < e_2 < e_3$. Consequently, by our hypothesis on the structure of the ordering of 1-cells on boundaries of cone-cells, we see that e_2 lies on ∂C. Moreover, the cyclical orientation is then determined, and hence since e_1, e_3 are adjacent with $e_1 < e_3$ and the ordering increasing in the other direction from e_3, we see that e_1 is first as illustrated. □

Continuing with the constraints indicated in Remark 3.27, we now provide a table with minimal defects:

Table 3.30. (1) $\tilde{C}_{\frac{\pi}{2}}C$
(2) $\overleftarrow{W}_{\frac{\pi}{2}}\hat{C}$ and $\hat{C}_{\frac{\pi}{2}}\overrightarrow{W}$
(3) $W_{\frac{\pi}{2}}W$
(4) $C_{\frac{\pi}{3}}C$
(5) $C_{\frac{\pi}{4}}\overleftarrow{W}$ and $\overrightarrow{W}_{\frac{\pi}{4}}C$
(6) $\check{C}_0\overleftarrow{W}$ and $\overrightarrow{W}_0\check{C}$.

We note that the sums of adjacent subscripted angle defects is $\leq$ the actual angled defects. To obtain the nonpositive curvature of the cone-cells, it is therefore sufficient to show that there is a total of at least 2π after grouping. This is supplied by the following:

Theorem 3.31. *Assign the split-angling to $\bar{D}$ and suppose the following conditions hold for each cone-cell C. Then each internal finite cone-cell in $\bar{D}$ has nonpositive curvature:*

(1) $|\overrightarrow{W}| \leq \frac{1}{8}|C|$.
(2) $|\frac{\pi}{6}\bar{C}\frac{\pi}{6}| \leq \frac{1}{6}|C|$.
(3) $|C'_{\frac{\pi}{4}}\overrightarrow{W}| \leq \frac{3}{16}|C|$ *and* $|\overleftarrow{W}_{\frac{\pi}{4}}C'| \leq \frac{3}{16}|C|$.
(4) $|\overleftarrow{W}_{\frac{\pi}{4}}C'_{\frac{\pi}{4}}\overrightarrow{W}| \leq \frac{1}{4}|C|$.

Proof. Consider a cone-cell C. We decompose the destination rectangles using the following "grouping rules" taken in order:

(1) Any $C_{\frac{\pi}{3}}C_{\frac{\pi}{3}}C$ is decomposed into $C_{\frac{\pi}{6}})(_{\frac{\pi}{6}}C_{\frac{\pi}{6}})(_{\frac{\pi}{6}}C$.
(2) $\check{C}_0\overleftarrow{W}$ and $\overrightarrow{W}_0\check{C}$ are grouped together. In case of $\overrightarrow{W}_0\check{C}_0\overleftarrow{W}$ we make an arbitrary choice to group only $C_0\overleftarrow{W}$.
(3) We group $C_0\overleftarrow{W}_{\frac{\pi}{4}}C_{\frac{\pi}{4}}\overrightarrow{W}_0C$ as well as variations of this omitting one or more noncentral terms: $C_0\overleftarrow{W}_{\frac{\pi}{4}}C_{\frac{\pi}{4}}\overrightarrow{W}$; $\overleftarrow{W}_{\frac{\pi}{4}}C_{\frac{\pi}{4}}\overrightarrow{W}$; $\overleftarrow{W}_{\frac{\pi}{4}}C_{\frac{\pi}{4}}\overrightarrow{W}_0C$; $C_0\overleftarrow{W}_{\frac{\pi}{4}}C_{\frac{\pi}{4}}$; $C_{\frac{\pi}{4}}C_{\frac{\pi}{4}}\overrightarrow{W}_0C$; $C_0\overleftarrow{W}_{\frac{\pi}{4}}C_{\frac{\pi}{4}}C$; etc. Note that when a $\overleftarrow{W}$ or $\overrightarrow{W}$ term is omitted, there is still such a rectangle which is initiated from the central C term, and terminates on a rectangle with destination the left or right C term. However, this connecting rectangle is not visible from our cone-cell.
(4) Each singly grouped W is treated as $(_{\frac{\pi}{4}}\overrightarrow{W})$ or $(\overleftarrow{W}_{\frac{\pi}{4}})$.
(5) If C is not grouped with anything on one side - or doesn't have a broken $\frac{\pi}{6}$ on one side from $(_{\frac{\pi}{6}}C_{\frac{\pi}{6}})(_{\frac{\pi}{6}}C)$ - then it obtains a $\frac{\pi}{4}$ on that side. Each singly grouped C is treated as $(_{\frac{\pi}{8}}C_{\frac{\pi}{8}})$.

Observe that the minimal square area of D implies that: $|C'_{\frac{\pi}{4}}\overrightarrow{W}_0C|$ lies behind $C'_{\frac{\pi}{4}}\overrightarrow{W}$ (if we extend the rectangle $\overrightarrow{W}$) and likewise $|C_0\overleftarrow{W}_{\frac{\pi}{4}}C|$ lies behind $\overleftarrow{W}_{\frac{\pi}{4}}C$, and consequently, bounds on the length of the wall projection of the latter give bounds on the former.

Similarly, $C_0\overleftarrow{W}_0C'_{\frac{\pi}{4}}\overrightarrow{W}_0C$ lies behind $\overleftarrow{W}_0C'_{\frac{\pi}{4}}\overrightarrow{W}_0$.

We reach the conclusion, that the minimal defect provided by a grouped piece P exceeds $2\pi\frac{|P|}{|C|}$.

This proves nonpositive curvature for cone-cells that are internal. □

The following indicates consequences of the inequalities of Theorem 3.31. Corollary 3.32.(1) generalizes the classical $C'(\frac{1}{6})$ condition, and Corollary 3.32.(3) recovers Theorem 3.20.

Corollary 3.32. *If any of the following inequalities hold then internal 2-cells have nonpositive curvature (and negative curvature if the inequalities are strict).*

(1) $|C| \leq \frac{1}{6}$ *and* $|W| \leq \frac{1}{48}$.
(2) $|C| \leq \frac{1}{8}$ *and* $|W| \leq \frac{1}{16}$.
(3) $|C| \leq \frac{1}{12}$ *and* $|W| \leq \frac{1}{12}$.

The following result slightly strengthens Corollary 3.32.(1), and the same method of proof recovers Corollary 3.32.(2).

Theorem 3.33. *Assign the split-angling to $\bar{D}$ and suppose the following conditions hold for each cone-cell C. Then each internal finite cone-cell in $\bar{D}$ has nonpositive curvature, and if the inequalities are strict then each has negative curvature.*

(1) *Each (contiguous) cone-piece is* $\leq \frac{1}{6}$ *and each wall-piece is* $\leq \frac{1}{24}$.
(2) *Each (contiguous) cone-piece is* $\leq \frac{1}{8}$ *and each wall-piece is* $\leq \frac{1}{16}$.

Moreover, if $\dim(X) = 2$, *then one obtains nonpositive curvature (respectively negative) if either:*

(3) *Each contiguous cone-piece is* $\leq \frac{1}{6}$ *and each wall-piece is* $\leq \frac{1}{12}$.
(4) *Each contiguous cone-piece is* $\leq \frac{1}{8}$ *and each wall-piece is* $\leq \frac{1}{8}$.

Proof. We will focus on the case where $|C| \leq \frac{1}{6}$, and $\dim(X)$ is arbitrary. The case where $|C| \leq \frac{1}{8}$ is very similar and will use a $\frac{\pi}{8}$ angle instead of $\frac{\pi}{6}$ angle etc. The case where $\dim(X) = 2$ uses that a single wall-piece dominates the concatenation of pieces we focus on later in the proof. This is illustrated on the right in Figure 59.

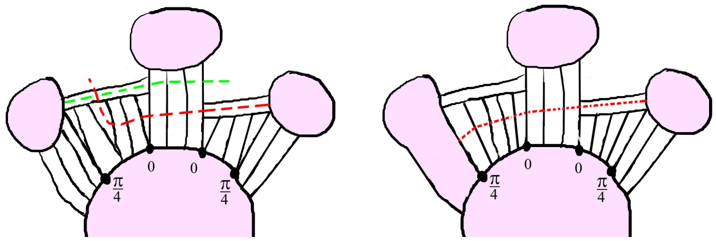

FIGURE 59. The concatenation of three pieces within $C_{\frac{\pi}{4}}W_0C_0W_{\frac{\pi}{4}}C$ is dominated by the concatenation of two wall-pieces. When $\dim(X) = 2$, it is dominated by a single wall-piece.

For each piece around an internal cone-cell, if it is a (contiguous) cone-piece there is an associated defect of $\geq \frac{2\pi}{6}$ and if it is a wall-piece then there is an associated defect of $\geq \frac{2\pi}{12}$. There are several cases where the defect at the transition between pieces is 0 that must be treated with special care. In particular, in the main exceptional case we will use that three pieces are dominated by two (crossing) wall-pieces, and our defect sharing system provides $\frac{\pi}{12}$ on each side. It is here that we require $|W| < \frac{1}{24}$ since a $\frac{2\pi}{12}$ defect is proportionate to $\frac{1}{24} + \frac{1}{24}$ of the circumference. It is also here that the bound improves when $\dim(X) = 2$, as since there is then only one dominating wall-piece it suffices to have $|W| < \frac{1}{12}$ here too.

In most of the cases outlined in Table 3.30, we distribute the defect evenly between the consecutive pieces: e.g. $\overline{C}_{\frac{\pi}{3}}\overline{C}_{\frac{\pi}{3}}\overline{C}$ decomposes into $\overline{C}_{\frac{\pi}{6}})(_{\frac{\pi}{6}}\overline{C}_{\frac{\pi}{6}})(_{\frac{\pi}{6}}\overline{C}$, and $\overrightarrow{W}_{\frac{\pi}{2}}\overrightarrow{W}_{\frac{\pi}{2}}\overleftarrow{W}$ decomposes into $\overrightarrow{W}_{\frac{\pi}{4}})(_{\frac{\pi}{4}}\overrightarrow{W}_{\frac{\pi}{4}})(_{\frac{\pi}{4}}\overleftarrow{W}$.

The exception to this rule involving a defect of 0 or $\frac{\pi}{4}$, when we will use an alternate distribution. In the case of a 0 defect, the critical observation is that the piece is dominated by one or two wall-pieces respectively. In the case of a $\frac{\pi}{4}$ defect, we will distribute this as $\frac{\pi}{6} + \frac{\pi}{12}$ towards the cone-piece and wall-piece respectively.

We use the following groupings:

(1) $_{\frac{\pi}{4}}\overrightarrow{W}_0\check{C}_0\overleftarrow{W}_{\frac{\pi}{4}}$ lies behind the following configuration: $_{\frac{\pi}{6}}(_{\frac{\pi}{12}}\overrightarrow{W}^{\times}\overleftarrow{W}_{\frac{\pi}{12}})_{\frac{\pi}{6}}$ as in Figure 59.
(2) $_{\frac{\pi}{4}}\check{C}_0\overleftarrow{W}$ and $\overrightarrow{W}_0\check{C}_{\frac{\pi}{4}}$ lie behind: $_{\frac{\pi}{6}})(_{\frac{\pi}{12}}\overleftarrow{W}_{\frac{\pi}{12}})(_?$ and $_?)(_{\frac{\pi}{12}}\overrightarrow{W}_{\frac{\pi}{12}})(_{\frac{\pi}{6}}$
(3) $C_{\frac{\pi}{4}}W$ and $W_{\frac{\pi}{4}}C$ are grouped as: $C_{\frac{\pi}{6}})_{\frac{\pi}{12}}W$ and $W_{\frac{\pi}{12}}(_{\frac{\pi}{6}}C$. □

Problem 3.34. Does nonpositive curvature (negative if <) hold in the following two cases:

(1) Each contiguous cone-piece is $\leq \frac{1}{8}$ and each wall-piece is $\leq \frac{1}{8}$.
(2) Each contiguous cone-piece is $\leq \frac{1}{6}$ and each wall-piece is $\leq \frac{1}{12}$.

The $4, 8, 8$ Euclidean tiling, and the tiling obtained by subdividing its square's edges, suggest that each statement in Problem 3.34 is sharp if true. The $4, 6, 6$ spherical tiling has $|C| < \frac{1}{6}$ and $|W| < \frac{1}{6}$.

3.14. **The ladder theorem.**

Definition 3.35 (Ladder)**.** A *ladder* is a disk diagram D with the property that there is a sequence of $n \geq 2$ closed cone-cells and/or vertices. $C_1, C_2, \ldots, C_n$ that are ordered so that C_j separates C_i from C_k when $i < j < k$. The diagram D is an alternating union of cone-cells (or vertices) and "pseudorectangles" R_i (possibly trivial or degenerate) in the following sense:

(1) $\partial_p D$ is a concatenation $P_1 P_2^{-1}$ where the initial and terminal points of P_1 lie on C_1 and C_n respectively.
(2) $P_1 = \alpha_1\rho_1\alpha_2\rho_2\ldots\alpha_n$ and $P_2 = \gamma_1\varrho_1\gamma_2\varrho_2\ldots\gamma_n$.
(3) $\partial_p C_i = \mu_i\alpha_i\nu_i^{-1}\beta_i^{-1}$ for each i where μ_1 and ν_n are trivial paths.
(4) $\partial_p R_i = \nu_i\rho_i\mu_{i+1}^{-1}\varrho_i^{-1}$ for each $i < n$.
(5) For each R_i, each dual curve starting on ν_i ends on μ_{i+1} and vice-versa.

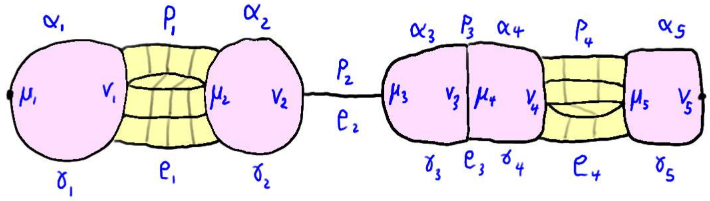

FIGURE 60. Ladder Definition Notation

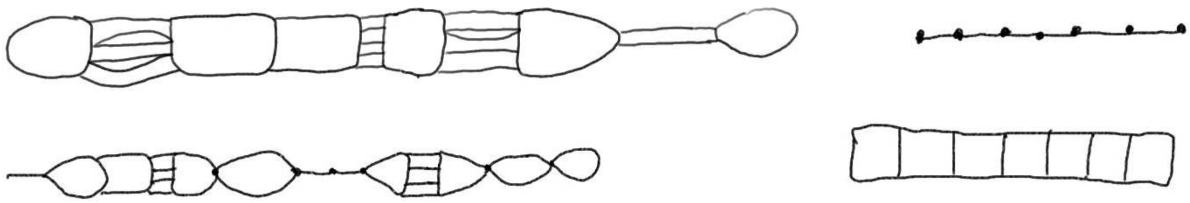

FIGURE 61. Some ladders.

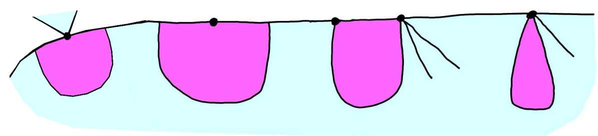

FIGURE 62. Singular doubly-external, doubly-external, singly-external, and nil-external corners.

(6) Moreover, a pair of dual curves in R_i cannot cross unless one (or both) start (and hence end) on ρ_i, ϱ_i.

There are two degenerate cases for R_i that should be noted: Firstly, it is possible that R_i is a vertical arc, and that $\nu_i = \mu_{i+1}$. Secondly, it is possible that R_i is a horizontal arc, in which case $\rho_i = \varrho_i$.

We refer to Figure 60 for help with the notation, and to Figure 61 for pictures of various ladders. A ladder is *nonsingular* if it has no cut-vertex, spur, or isolated 1-cell.

We regard a disk diagram consisting of a single 0-cell or cone-cell as a *trivial ladder*.

We emphasize that a cone-cell is *external* or *internal* according to whether or not its boundary contains a 1-cell in ∂D.

Theorem 3.36. *[Ladder Theorem] Suppose that the rectified diagram $\bar{D}$ was created using an ordering of 1-cells induced by an order of cone-cells (with C_∞ last), and suppose the angle assignment on $\bar{D}$ has the properties listed below. If $\bar{D}$ has exactly two positively curved cells then $\bar{D}$ is a ladder.*

1. *Each internal cone-cell, shard, and internal vertex has nonpositive curvature.*
2. *The rectangles have the usual angles: $\frac{\pi}{2}$ at the four corners, and π elsewhere.*
3. *$0 \leq \sphericalangle \leq \pi$ for each $\sphericalangle$ in a cone-cell.*
4. *When adjacent rectangles have distinct (even implicit) destinations the internal angle at the corner of the cone-cell at the corresponding vertex has nonzero defect.*
5. *The angle at a corner of a cone-cell at a vertex on ∂_p equals 0 when it is doubly-external but at a singular vertex, equals π when it is nonsingular doubly-external, equals $\frac{\pi}{2}$ when it is singly-external, and is > 0 when it is nil-external.*

The four types of external corners of a cone-cell are illustrated in Figure 62.

Proof. As there are exactly two positively curved cells, each has curvature exactly π. Indeed using Equation (1), the $0 \leq \sphericalangle$ hypothesis implies that the curvature of a boundary 0-cell v is $2\pi - \pi\chi(\text{link}(v)) - \sum \sphericalangle \leq \pi - \sum \sphericalangle \leq \pi$. Applying Equation (2), the hypotheses on external corner angles together with

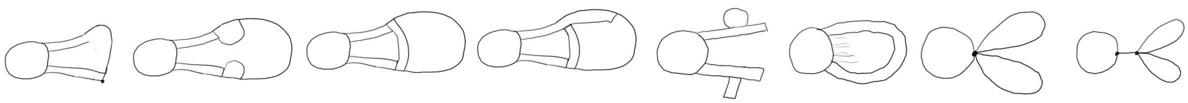

FIGURE 63. It must be a ladder.

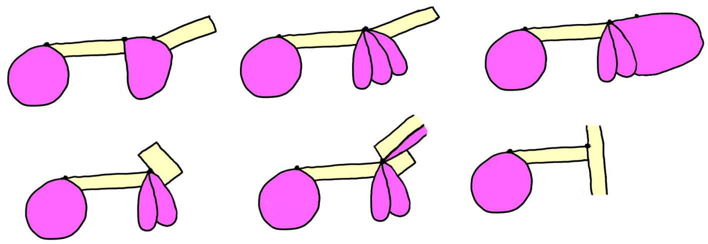

FIGURE 64. Woof Woof.

the $\sphericalangle \leq \pi$ hypothesis implies that the curvature of an external cone-cell is $\leq \pi$. Here each nontrivial boundary arc provides a defect of π from the two $\frac{\pi}{2}$ singly-external corners, and all other angles have nonpositive defect since $\sphericalangle \leq \pi$. To justify this count, the π-defect at the corner at a singular doubly-external vertex, is shared between the (locally) two external boundary paths ending there.

A case of interest is when the boundary path of a 2-cell consists entirely of two external boundary paths - in which case both corners are singular doubly-external with angle defect π.

The only π curvature 0-cell arises from a spur. The π-curvature cone-cells arise when the external cone-cell has a single external boundary path with a defect of $\frac{\pi}{2}$ on each end and a nontrivial innerpath whose internal corners have angle π, or when the external boundary path is the entire boundary path and the innerpath is trivial, in which case there is a defect of π at the singular doubly-external corner where the boundary path starts and ends.

Since each other cell has nonpositive curvature, Theorem 3.15 implies that each other cell has zero curvature.

Let us assume the diagram is nonsingular, otherwise an inductive argument would allow us to string together (possibly trivial) ladders to obtain a new ladder.

Consider a positively curved cell C_1. Either it is a 0-cell at the tip of a spur, or it is a cone-cell. Let us consider the latter case. Let e_1, e_2 denote the 1-cells on either side of ∂C_1 just outside of $\partial C_1 \cap \partial \bar{D}$. Let R_1, R_2 denote the rectangles at e_1, e_2. (It is possible that $e_1 = e_2$ and hence $R_1 = R_2$.) Traveling along $R_1 \cap \partial \bar{D}$ (see Figure 63 for various scenarios) we see that $\partial \bar{D}$ proceeds along the entire (possibly degenerate) external path on one side of R_1 until it ends on a rectangle or cone-cell. Indeed, another incoming (possibly degenerate and thus including another cone-cell) rectangle would give us a 0-cell on $\partial \bar{D}$ with negative curvature. The same reasoning holds for R_2.

By hypothesis, a transition in destination would give a positive defect along the interior path of C_1. Consequently, R_1 and R_2 both either have rectangle destinations or cone-cell destinations. Moreover, if they both have cone-cell destinations, then these cone-cells are the same, and likewise if they both have rectangle destinations, then they end implicitly on the same rectangle. We will now rule out the possibility that they have rectangle destinations (that are implicitly the same). Indeed, if R_1 ends on a rectangle, then its final boundary corner has exactly two squares along it as at the top of the third diagram in Figure 63 and no cone-cells. Indeed, further corners of squares would provide $\frac{\pi}{2}$ angles, and the last cone-cell or square as we go around this boundary vertex would provide another $\frac{\pi}{2}$ angle. All other angles are nonnegative, so the total angle sum would be $\geq \frac{3\pi}{2}$. The simplest possible such scenario is illustrated at the bottom right of Figure 64.

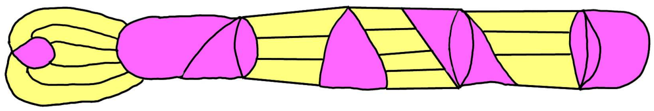

FIGURE 65. If 0-angles are assigned to nil-external corners, ladders could have cone-cells that intersect the boundary along trivial arcs.

We emphasize that the only way a shard could have a sharp corner at a boundary vertex, is if the rectangles on either side terminated at a pair of consecutive edges there. This is impossible in our case, for one of these boundary edges is an external edge of R_1. We therefore do not have to worry about negative angles coming from sharp corners of shards that occur at this vertex, and so the angle sum is $\geq \frac{3\pi}{2}$.

We thus find that the rectangle ends immediately on C_∞ at the square where R_1 meets it. Thus one end of this rectangle lies on C_∞. Now, if both ends of this rectangle lie on C_∞ then the ordering of initial 1-cells of rectangles is violated - since 1-cells on C_∞ come last. If one end lies on C_∞ and the other end lies on an external edge of R_2 then we get a similar contradiction.

We now find that we have reached some cone-cell C_2, and that all the rectangles leaving C_1 are parallel to each other and end on C_2. If $\kappa(C_2) = \pi$ then we are done, as the same reasoning applies in the reverse direction.

Otherwise $\kappa(C_2) = 0$. Note that $\partial C_2 \cap \partial_p(\bar{D})$ has exactly two components, for otherwise using Equation (2), our hypothesis that $\sphericalangle \leq \pi$ yields a total curvature $\leq -\pi$. Indeed, our hypothesis of $\frac{\pi}{2}$ at singly-external corners and ≥ 0 at degenerately-external corners, implies a total defect of $-\pi$ for each component of $\partial C_2 \cap \partial_p(\bar{D})$.

Observe that the 0-cells on $\partial\bar{D}$ where R_1 and R_2 intersect C_2 have curvature 0, as each has exactly two $\frac{\pi}{2}$ angles (one from R_i and one from the final cone-cell or square as we travel through the sequence of corners of cells until we get to the next singly-external corner of a cell). Various possibilities are illustrated in Figure 64, where the top diagrams are allowable and the bottom diagrams are impossible. The vertex has curvature 0, and after removing C_1 and the sequence of (possibly degenerate) rectangles between R_1 and R_2, we see that C_2 now intersects the new $\partial_p(\bar{D})$ in exactly one component, and we proceed by repeating the above argument on the subdiagram (where two angles of C_2 might be redefined if they went from being nil-external to singly-external).

Thus far we have obtained a ladder in the more general sense described in Remark 3.37, where cone-cells might intersect $\partial_p\bar{D}$ in trivial subpaths. This only used that nil-external corners have $\sphericalangle \geq 0$.

We now explain the consequence of our strengthened assumption that nil-external corners of cone-cells have $\sphericalangle > 0$. Any such corner provides a defect that is strictly $> \pi$. Consequently, if $\kappa(C_2) = 0$ then under this stricter $\sphericalangle > 0$ hypothesis, $\partial C_2 \cap \partial_p(\bar{D})$ contains exactly two *nontrivial* arcs.

Using similar reasoning, we note that in the strict case, the edges e_1, e_2 cannot intersect on the outside of ∂C_1 on $\partial_p(\bar{D})$. Indeed, then $\partial C_1 \cap \partial_p(\bar{D})$ consists of a single point. And if the corner at this point has $\sphericalangle > 0$ then $\kappa(C_1) = 2\pi - \pi - \sphericalangle < \pi$.

Without the $\sphericalangle > 0$ hypothesis, it is possible to have strings of cone-cells (looking like bigons and triangles in the middle, and possibly monogons at the end) having 0-angles along the boundary as in Figure 65. $\square$

Remark 3.37. The hypothesis of $\sphericalangle > 0$ at nil-external corners of cone-cells forces cone-cells in a ladder to intersect the boundary in nontrivial subpaths. See Figure 65.

Remark 3.38 (Effect of positive angles at singular doubly-external corners)**.** A version of Theorem 3.36 holds under the hypothesis that a singular doubly-external corner has angle > 0. Indeed, a sensible reinterpretation of the split-angling could assign a $\frac{\pi}{2}$ angle here, since the two (degenerate) rectangles do not end in parallel on the cone-cell at infinity. Another tempting interpretation of the split-angling would assign an angle of π, and a version of Theorem 3.36 would still hold.

When a positive angle is assigned to singular doubly-external corners, we must assume that the diagram is nonsingular to reach our initial enabling observation that each feature of positive curvature must have curvature exactly π, as a *dumbbell* consisting of two cone-cells connected by a (possibly trivial) arc of edges provides a counterexample to this. It appears that these dumbbells, possibly with spurs at their ends, are the only nonsingular ladders with exactly two positively curved cells. On the other hand, there are many singular ladders with alternating sequences of positively curved cone-cells and negatively curved 0-cells, and this suggests that the 0-angle assignment is more elegant.

Remark 3.39. There are similar results for annular diagrams obtained in Section 5.16. Note that an annular diagram is treated so that it has two cone-cells at infinity, and these both occur at the end in the ordering that is used to choose rectangles. The angle assignments are done in the same way.

The main difference is that in order to obtain an annulus that is as thin as a ladder, we must assume that internal cone-cells have negative curvature, and that (hard to obtain) the external cone-cells cannot have nonpositive curvature unless they touch both the inner and outer infinite-cones. Otherwise we obtain a thickness 2 situation that is hard to control.

3.15. **Positive curvature along boundary.** The possible positively curved cells are:

(1) A single isolated 0-cell has curvature 2π
(2) A single isolated cone-cell has curvature 2π
(3) The 0-cell at the end of a spur has curvature π
(4) The 0-cell at the center of the outerpath of a generalized square corner has curvature $\frac{\pi}{2}$
(5) A *shell* C is a positively curved external cone-cell. Its boundary path is a concatenation QS where the *outerpath* Q is a subpath of the boundary path of the diagram, and the *innerpath* S has all open 1-cells in the interior of the diagram. The curvature of C equals the sum of the defects of the angles along interior(S).

Theorem 3.40 (Positive curvature cells)**.** *Suppose $\bar{D}$ is a rectified disk diagram with an angle assignment satisfying the conditions enumerated in Theorem 3.36. Then one of the following hold:*

(1) *$\bar{D}$ consists of a single* 0*-cell or a single closed cone-cell.*
(2) *$\bar{D}$ is a ladder.*
(3) *$\bar{D}$ has at least three shells and/or spurs and/or generalized corners of squares along $\partial\bar{D}$.*

Proof. This follows from Theorem 3.15 since curvatures are $\leq \pi$ except for the two degenerate cases, and curvatures of internal cells and shards are nonpositive by hypothesis/construction. The case where there are only two positively curved cells was treated in Theorem 3.36. □

Definition 3.41 (Small-cancellation complex)**.** Consider a cubical presentation $\langle X \mid Y_i \rangle$ with an angle-assignment method such that in any reduced rectified diagram we have: internal 2-cells, internal 0-cells, and shards have nonpositive curvature. We will refer to $\langle X \mid Y_i \rangle$ as being a *cubical small-cancellation complex*.

3.16. **Examples.**

(1) Ordinary $C'(\frac{1}{6})$ small-cancellation theory where the cube complex is a graph and the relators are circles.
(2) Graphical Small-Cancellation theory. This was apparently first noticed in [RS87] but was rediscovered subsequently by Gromov [Gro03, Oll06].
(3) A $C'(\frac{1}{12})$ Cubical presentation.
(4) Large Dehn filling of Dehn complex of prime alternating link. See Subsection **??**.

(5) Small-cancellation theory over right-angled Artin groups with ordinary relators. (This gives a variety of simple examples that don't utilize relative hyperbolicity of $\pi_1 X$. But one must be careful about the cores.)

Specifically, a small-cancellation set of words alternating substantially around the generators. Pieces correspond to subwords that equal each other (after shuffling).

(6) Many Artin groups (almost all?). See Section 20. Every Artin group associated to a Coxeter group with no exponent of degree 2. More generally, this appears to work if there is no triangle of type: $(2, 2, n)$ or $(2, 3, 3)$ or $(2, 3, 4)$, or $(2, 3, 5)$.

It appears we can allow degree 2 if we exclude degree 3, but this requires complicated generalized relators that are centralizers, and perhaps a relative version of the theory.

We start with the right-angled Artin group. Then we add the relators. The small-cancellation conditions aren't satisfied if we just cone off the cycles. So instead, we cone off by the immersed graph associated to the normal closure of each relator in its two generator free subgroup. (This is just the Cayley graph of a 2-generator Artin group.) This gives us an immersed cube complex corresponding to the 1-skeleton of the universal cover of the standard 2-complex of the 2-generator Artin group.

Pieces correspond to $a_i^{n_j}$ paths, if they are pieces between two generalized relators of different types. The pieces correspond are arbitrary if they arrive from two distinct conjugates of the same generalized relator, where conjugation is by a path centralizing its two generators. In this case, we can cut and paste (or push around) to combine the two cone-cells in the disk diagram - and thus reduce area. We can therefore assume there are no such pieces.

Example 3.42. An interesting example arises from the 4-string braid group B_4. The kernel of its homomorphism to the underlying Coxeter group S_4 is called the 4-string pure braid group P_4. It is shown in [DLS91] that $P_4 \cong G \times \mathbb{Z}$ where G is the following quotient of a right angled Artin group:

$$\langle a_1, \ldots, a_5 \mid [a_i, a_{i+1}] \rangle \;/\; \langle\!\langle a_1 a_2^{-1} a_3 a_4^{-1} a_5 a_1^{-1} a_2 a_3^{-1} a_4 a_5^{-1} \rangle\!\rangle.$$

This appears to satisfy small-cancellation conditions, with a bit of care.

Higher degree examples (e.g. $n \geq 6$) work much better.

Example 3.43. Consider the presentation $\langle a, b, c \mid (ab)^2, (bc)^2, (ca)^2, (aaabbbccc)^2 \rangle$ which differs slightly from a lovely presentation for an index 3 subgroup of $PSL(2, \mathbb{Z}[\frac{1}{2}])$ discovered by Cameron Gordon. (His presentation has $a^3b^3c^3$ not raised to a power.)

There is an obvious homomorphism to $\mathbb{Z}_2^3$ in which the obvious torsion elements survive, and we can then collapse pairs of 2-cells corresponding to the relators. The result is a cubical presentation $\langle X \mid Y_1, Y_2, Y_3, Y_4 \rangle$ where X is an orientable genus 3 surface built from squares with 6 meeting around each 0-cell, and each Y_i is built from a lift of $(a^3b^3c^3)^2$ by adding six squares at the corners corresponding to the transitions ab, bc, ca, ab, bc, ca. It appears that all maximal cone-pieces are either of the form aa, bb, cc or are one of the added squares. Besides the wall-pieces consisting of single 1-cells, there are wall-pieces of the form cb^{-1}, ba^{-1}, and ac^{-1}.

It thus seems that the $C(6)$ condition is satisfied with the split-angling. There is an obvious (antipodal) wallspace structure on each Y_i, and it appears that together with it the $B(6)$ conditions (and probably $B(8)$ conditions) hold as well.

3.17. **Examples arising from special cube complexes.** The work here was motivated by the observation made in [Wis03] that for any finite immersed graph $\Lambda \to \Gamma$, there is a finite cover $\hat{\Lambda}$ such that $\langle \Gamma \mid \hat{\Lambda} \rangle$ satisfies Gromov's graphical $\frac{1}{6}$ small-cancellation theory (see [RS87]).

Theorem 3.44. *Let X be a compact nonpositively curved cube complex. Let $H_1, \ldots, H_k$ be residually finite subgroups of $\pi_1 X$, and for each i, let $Y_i \to X$ be a compact immersed complex with $\pi_1 Y_i \cong H_i$.*

Suppose each $\widetilde{Y}_i \subset \widetilde{X}$ is superconvex. Suppose that there is an upper bound on the diameters of intersections between distinct translates of $\widetilde{Y}_i, \widetilde{Y}_j$ in $\widetilde{X}$ (we allow $i = j$ here).

Then for each $\alpha > 0$ there are finite covers $\widehat{Y}_i$ such that $\langle X \mid \widehat{Y}_1, \dots, \widehat{Y}_k\rangle$ is $C'(\alpha)$.

Proof. By Lemma 2.13, let D be a uniform upper bound on the diameter of a flat strip that stands on Y_i. (This will also bound lengths of pieces between noncontiguous $\widetilde{Y}_i, \widetilde{Y}_j$).

Let E be a uniform upper bound on the diameters of intersections between translates of $\widetilde{Y}_i, \widetilde{Y}_j$ (of course we exclude the overlap between identical translates of $\widetilde{Y}_i$).

Note that by possibly passing to a finite index supergroup of each H_i, we can assume that each $H_i = \text{Stab}(\widetilde{Y}_i)$.

By residual finiteness, for each i, we can choose a finite regular cover $\widehat{Y}_i \to Y_i$ such that $|Y_i| > \frac{1}{\alpha}\max(D, E)$. □

Corollary 3.45. *Let X be a compact virtually special cube complex. Suppose $\pi_1 X$ is word-hyperbolic. Let $H_1, \dots, H_k$ be quasiconvex subgroups of $\pi_1 X$ that form a malnormal collection. Then for each $\alpha > 0$ there are finite index subgroups $H_i' \subset H_i$ and represented by compact local isometries $\widehat{Y}_i \to X$ such that $\langle X \mid \widehat{Y}_1, \dots, \widehat{Y}_k\rangle$ satisfies $C'(\alpha)$.*

Proof. By Lemma 8.5, there is an H_i-cocompact superconvex subcomplex $\widetilde{Y}_i \subset \widetilde{X}$ for each i. There is an upper bound on the overlap between translates of $\widetilde{Y}_i$ and $\widetilde{Y}_j$ by the malnormality assumption.

H_i is residually finite for each i since $\pi_1 X$ is virtually special.

We can thus apply Theorem 3.44. □

Lemma 3.46. *[Malnormal Controls Overlap] Let X be a compact nonpositively curved cube complex (with $\pi_1 X$ word-hyperbolic). For $1 \le i \le r$, let $Y_i \to X$ be a local-isometry with Y_i compact, and assume the collection $\{\pi_1 Y_1, \dots, \pi_1 Y_r\}$ is malnormal. Then there is a uniform upper bound D on the diameters of intersections $g\widetilde{Y}_i \cap h\widetilde{Y}_j$ between distinct $\pi_1 X$-translates of their universal covers in $\widetilde{X}$.*

Lemma 3.46 can be interpreted as saying that there is an upper bound on diameters of contiguous cone-pieces in $\langle X \mid Y_1, \dots, Y_r\rangle$. In practice, one applies Lemma 3.46 under the assumption that $\pi_1 X$ is word-hyperbolic for this enables the existence of compact cores Y_i for quasiconvex subgroups. A deeper investigation of Lemma 3.46 in the quasiconvex-malnormal and hyperbolic situation, shows that there is a uniform upper bound on the sizes of *all* cone-pieces - not just the contiguous ones. This relationship between malnormality and pieces is concealed by the noncontiguous cone-pieces which can sometimes be ignored in the small-cancellation theory since they are hidden behind and thus controlled by hyperplanes in a superconvex situation.

One way to prove Lemma 3.46 is to consider the nondiagonal components of $(\sqcup Y_i) \otimes_X (\sqcup Y_i)$. Each of these is contractible by malnormality, and of finite diameter since the finitely many Y_i are compact. We refer the reader to Section 8.2.

Proof. If $h\widetilde{Y}_i \cap g\widetilde{Y}_j$ has infinite diameter then applying the pigeon-hole principle, there is an infinite periodic path lying in both $h\widetilde{Y}_i$ and $g\widetilde{Y}_j$. This yields an infinite order element in both $\pi_1 Y_i^h$ and $\pi_1 Y_j^g$, thus violating malnormality unless $i = j$ and $gh^{-1} \in \pi_1 Y_i$. □

3.18. **Informal discussion of the limits of the theory.** The $C(6)$ condition is sufficient when $\dim(X) = 1$ but we cannot expect $C(6)$ to suffice in general. Indeed, consider the snub octahedron: it is a tiling of the sphere where each vertex is surrounded by a square and two hexagons. In any reasonable sense, each hexagon is not the concatenation of fewer than 6 pieces. The complement of a square does not seem to have any 3-shell.

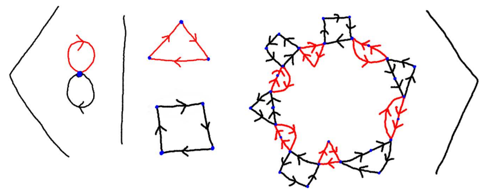

FIGURE 66. There are two grade-1 cones, and one grade-2 cone. All nonabsorbable pieces are between the grade-2 cone and itself. The example doesn't satisfy $C(p)$ for very large p, but suggests a plethora of further examples along similar lines. For a presentation of the form $\langle b, c \mid b^m, c^n, W\rangle$, one can attach copies of b^m and c^n to a copy of W along its syllables. When $m = 2$ and $n = 3$, one needs $|W|$ quite large before typical words give graded small-cancellation presentations.

Another example to bear in mind is the snub icosahedron, which has hexagons surrounding pentagons, and is otherwise similar.

The snub-$(4, 4)$ tiling of the plane, consisting of squares surrounding by octagons has nonpositive curvature, and seems like a reasonable limit here. An obvious grade-angling works fine for it, and (assuming the squares are already in X) the split-angling appears to work no matter how the admitted rectangles are oriented.

3.19. **Graded small-cancellation.** Thus far we have emphasized the version of cubical small-cancellation theory where cone-pieces between cones are small unless the associated cones in $\widetilde{X}$ are equal. We will now turn to a generalization which insists that cone-pieces are small unless one cone is contained in the other in $\widetilde{X}$.

This form of cubical small-cancellation theory provides the language for crucial to the (actual) proof of Theorem 13.1. Some simple examples are illustrated in Figures 66, 67, 68, and 69. Small-cancellation theory persists in this situation. Pieces that look large in the universal cover project to pieces that are small in the cone itself. Thus if there aren't enough pieces around a cone-cell, then the boundary path is already homotopic using the lower grade cone-cells. In a most general setting, we can replace a higher grade cone-cell by a diagram with lower grade cone-cells and squares. It is thus appropriate to use a minimal *graded complexity* which counts the number of cone-cells of each grade (followed by the number of squares). A lexicographical ordering is used here so $(3, 3, 2, 6, 5) > (3, 3, 2, 5, 17)$ etc.

Our main tools: Greendlinger's Lemma (in the weakened form of Theorem 3.40) and the Ladder Theorem, are about diagrams so they continue to hold in our context, as a reduced diagram will satisfy the small-cancellation conditions (with some angle assignment rule). The notion of *no missing θ-shells*, and maps $A^* \to X^*$, and a presentation A^* induced from X^* and a local isometry $A \to X$ proceed unchanged.

Gromov Polyhedra: Haglund's theory of "Gromov Polyhedra" fits nicely into this category. These are CAT(0) 2-complexes whose 2-cells are p-gons and whose vertex links are complete graphs K_r. Haglund shows how to produce many such examples with a proper cocompact group action. For instance, Haglund begins with a free product $A * B$ of finite groups, with generators consisting of the full set of nontrivial elements $\{A - 1\}, \{B - 1\}$. He then adds relators $R_1, \ldots, R_k$ with the property that each generator $a \in \{A - 1\}$ or $b \in \{B - 1\}$ appears exactly once, with the extra symmetry condition that if u and u^{-1} both appear in R_i then there is an automorphism of R_i sending one occurrence to the other. If $r = |A| = |B|$ then each link will be isomorphic to K_r, and if p is the syllabic length of each relator, then each 2-cell will have p sides. For instance, Haglund provides the following example where u, u^{-1} never recur: $\langle a, b \mid \ a^{17},\ b^{17},\ ab^4a^2b^5a^{12}b^6a^9b^7,\ a^3b^{16}a^4b^2a^6b^{14}a^{10}b^9\ \rangle$.

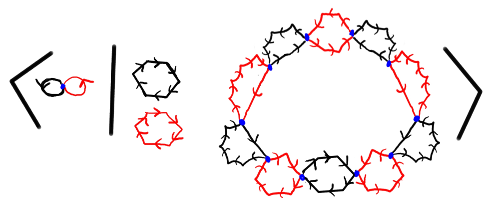

FIGURE 67. Haglund's group acting properly and cocompactly on a $(6, 10)$ Gromov polyhedron.

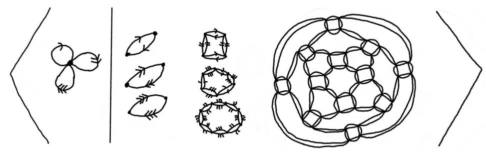

FIGURE 68. The 2-3-4 triangle group: A Coxeter group yields a graded presentation. Its generators are the usual generators. Its grade i relators are Cayley graphs of the i-generator Coxeter subgroups. For a relator Y of grade ≥ 3, the induced cubical presentation Y^* is already simply-connected so $|Y^*| = \infty$.

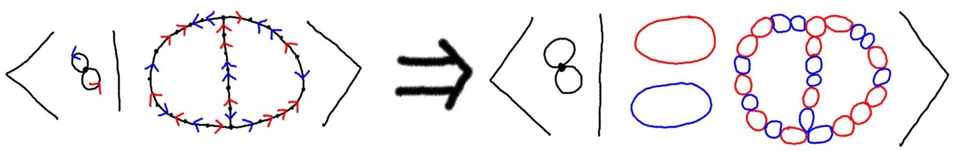

FIGURE 69. Substituting: Let $\langle B \mid C\rangle$ denote a graphical presentation where B is a bouquet of circles u, v, and $C \to B$ is an immersed theta-graph. Let X be a new bouquet of circles x, y, and let U, V denote immersed circles, and "substitute" copies of U, V for edges labelled by u, v in C and then fold to obtain a graph Y. Even preserving "orientation", there are many ways of doing this, and one could choose basepoints, or just do it randomly. When $\langle B \mid C\rangle$ and $\langle X \mid U, V\rangle$ are sufficiently small-cancellation, then so will the cubical presentation $\langle X \mid U, V, Y\rangle$ where U, V have grade 1 and Y has grade 2.

Haglund's examples naturally lead to cubical presentations that are graded, where X is a wedge on a bouquet of circles for the generators (optionally wedged along a new edge) and the grade 1 relators are the Cayley graphs $\Gamma(A), \Gamma(B)$ of A, B, and the grade 2 relators are copies of the circles corresponding to the R_i, but with copies of $\Gamma(A), \Gamma(B)$ attached along each generator. Haglund's symmetry condition implies that the maximal pieces only occur between these grade 1 relators, and are merely copies of $\Gamma(A), \Gamma(B)$.

Figure 67 illustrates one of the simplest of Haglund's examples: $\langle a, b \mid a^6, b^6, b^4a^5b^5a^2b^3a^4bab^2a^3\rangle$. We used a single generator for each $A, B \cong \mathbb{Z}_6$, but the reader can redraw with a wedge of two bouquets of 5 circles, grade 1 relators isomorphic to 6-simplices, and a grade 2 relator that looks like a string of 6-simplices glued around a decagon.

3.20. **Metric Small-Cancellation and Quasiconvexity.** This section can be skipped on a first reading. It is used prominently in Section 15. Corollary 3.50 is used very briefly in the proof of Theorem 12.1 and Theorem 13.1. The material could probably be used to give a simplified treatment of parts of Section 5.

For a path $S \to Y$ in a cone of X^*, we let $\nabla_Y(S)$ denote the distance in $\widetilde{Y}$ between the endpoints of a lift $\widetilde{S}$. For a graded cubical presentation $\langle X \mid Y_{ij}\rangle$ and a path $S \to Y_{ij}$, we define $\nabla_{Y_{ij}}(S)$ to equal the distance between the endpoints of a lift of S to $\widetilde{(Y_{ij})}^*$ where $(Y_{ij})^*$ is the cubical presentation induced from X^*_{i-1} by $Y_{ij} \to X$, where $X^*_{i-1} = \langle X \mid Y_{kj} : k < i\rangle$.

Definition 3.47 (Metric Small-Cancellation)**.** $\langle X \mid Y_i\rangle$ satisfies $C'(\alpha)$ if for each piece P between $g_i\widetilde{Y}_i$ and $g_j\widetilde{Y}_j$ either $\nabla_{Y_i}(P) < \alpha|Y_i^*|$ and $\nabla_{Y_j}(P) < \alpha|Y_j^*|$, or $\text{Grade}(Y_i) \leq \text{Grade}(Y_j)$ and $g_i\widetilde{Y}_i \subset g_j\widetilde{Y}_j$ or $\text{Grade}(Y_j) \leq \text{Grade}(Y_i)$ and $g_i\widetilde{Y}_i \supset g_j\widetilde{Y}_j$ (In fact, the above inclusions can be replaced by $g_i\widetilde{Y}_i \subset g_j\widetilde{Y}_j$ in $\widetilde{X}^*_m$ (and vice-versa). See Remark 3.8.)

It is a separate matter to use that $\nabla_{Y_i}(P) < \alpha|Y_i^*|$ for each piece in order to conclude small-cancellation. It seems we need adjacent outside rectangles to also have some bound: for instance a base path of any such rectangle in $\widetilde{X}$ along $\widetilde{Y}_i$ satisfies $\nabla_{Y_i}(P) < \alpha|Y_i^*|$.

The most important consequence of appropriate $C'(\alpha)$ (and respectively $C^\natural(\alpha)$) is that a small-cancellation complex has *short innerpaths* which means that for each essential θ-shell $SQ \to Y$ in a diagram $D \to X^*$ with $\Omega_Y(S) < \pi$, we have $\nabla_Y(S) < |Q|$.

There is a related condition $C^\natural(\alpha)$ requiring only that $\nabla_{Y_i}(P) \leq \alpha|Y_i^*|$. In this case *medium innerpaths* is used to mean that $\nabla_{Y_i}(S) \leq |Q|$.

Lemma 3.48. *There exists α such that if $\langle X \mid Y_i\rangle$ is $C'(\alpha)$ then it has short innerpaths under the split-angling.*

The proof of Theorem 3.20 should show that $C'(\frac{1}{12})$ suffices. Instead, we give a simplified proof of the same statement under the stronger condition of $C'(\frac{1}{24})$.

Proof in $C'(\frac{1}{24})$ case. In the split-angling, the largest concatenation of pieces without defect contributions from the transitions between pieces, consists of a (wall-piece)-(cone-piece)-(wall-piece). And the smallest defect contribution in the split-angling is $\frac{\pi}{4}$. Any path $S \to Y_i$ with $\Omega_{Y_i}(S) \leq \pi$ is thus the concatenation of at most 15 pieces. To see that $\Omega_{Y_i}(S) < \pi$ implies at most 12 pieces, observe that with n pieces, the minimal number of positive defect transitions is $\lfloor\frac{n-1}{3}\rfloor$, and this is at least 4 when $n > 12$.

Thus when $\nabla_{Y_i}(P) < \frac{1}{24}|Y_i|$ for each piece P then $\Omega_{Y_i}(S) < \pi$ implies that $\nabla_{Y_i}(S) < \frac{1}{2}|Y_i|$. □

Theorem 3.49. *Suppose the cubical presentation $\langle X \mid Y_i\rangle$ has an angling system making it into a small-cancellation complex with short innerpaths.*

Let $\langle A \mid B_j\rangle$ be another cubical presentation, and suppose the map $A^* \to X^*$ has no missing θ-shells.

Let $p, q \in \widetilde{A}^*$, (not cone points) and let $\bar{p}, \bar{q}$ be their images in $\widetilde{X}^*$. Let γ' be an arbitrary geodesic joining $\bar{p}, \bar{q}$. then there exists a geodesic γ homotopic to γ' in X (not $\widetilde{X}^*$). And there is a path $\sigma \to \widetilde{A}^*$ between p, q such that the image $\bar{\sigma}$ of σ in $\widetilde{X}^*$ has the following property: $\bar{\sigma}\gamma^{-1}$ is the boundary path of a ladder $L \to \widetilde{X}^*$. If the ladder doesn't consist of a single cone-cell or vertex, then p, q lie in the interior of the outerpath of the first and last cone-cells and/or spurs of L on each end.

Proof. Let D be a minimal complexity diagram between paths σ and γ, where σ varies among paths in $\widetilde{A}^*$ with endpoints p, q, and where γ varies among all geodesics in $\widetilde{X}$ path homotopic to γ' (so we aren't varying among geodesics in $\widetilde{X}^*$).

Observe that D has no corners of generalized squares along γ, or along σ.

Likewise D has no outerpaths of positively curved θ-shells along either. None along σ for such a θ-shell would map to $\widetilde{A}^*$ and thus σ could be passed through it. None along γ for such a θ-shell would violate that γ and hence γ' is a geodesic by the short innerpaths hypothesis.

There are thus only two positively curved cells: A spur or cone-cell at $\bar{p}$ and a spur or cone-cell at $\bar{q}$. □

Corollary 3.50. *Let $A^* \to X^*$ be as in Theorem 3.49 with A^*, X^* compact. Then $\widetilde{A}^* \to \widetilde{X}^*$ is a quasi-isometry.*

Suppose moreover that $\widetilde{X}$ is δ-hyperbolic. And the maximal diameter of a cone Y_i is κ. Then $\widetilde{A}^* \to \widetilde{X}^*$ is $(\delta + \kappa)$-quasiconvex.

Proof. Consider a minimal complexity ladder as in Theorem 3.49.

For the first statement, we show that the lengths of γ and σ are proportional since each cone-cell has at most μ edges on σ and at least one edge on γ, and the intermediate parts can be assumed to be rectangles, and thus have the same length. Let μ denote an upper bound on the total number of vertices in a cone. Thus if a subpath σ' of σ lying in some cone Y_i has length exceeding μ then we see that it has a subpath mapping to a closed path of Y_i. If this closed path is essential, then it lifts to a closed path in a cone of A^*, and if it is not essential then it bounds a square diagram in $\widetilde{A}$. Either way, we could reduce the complexity to get a smaller diagram between γ and σ.

For the second statement, note that cone-pieces have diameter $< \kappa$. The point is that a rectangle between successive cones in the ladder has a bounded height because if it is too tall, then an inner part of it can be replaced by a bounded part. Thus the geodesic γ' lies in the δ neighborhood of γ which lies on one side of a ladder L, with a path $\sigma \to \widetilde{A}^*$ on the other side. But then γ lies in the κ neighborhood of σ. □

Remark 3.51. The condition in Theorem 3.49 can be replaced by medium innerpaths: $\nabla(S) \leq |Q|$, but we must then allow γ to vary in $\widetilde{X}^*$ and not just in $\widetilde{X}$. So we conclude that γ exists, and still have a quasi-isometric embedding conclusion.

Lemma 3.52 (Convexity)**.** *Let us add the following hypotheses to the situation described in Theorem 3.49. Let $2\alpha + \beta = \frac{1}{2}$ where $\alpha, \beta > 0$ (we have in mind $\alpha = \frac{1}{8}$ and $\beta = \frac{1}{4}$).*

Suppose that cone-pieces in X^* are small in the sense that $\nabla_{Y_i}(P) < \alpha|Y_i^*|$ whenever P is a cone-piece of a translate of $\widetilde{Y}_j$ in a translate of $\widetilde{Y}_i$ (and one is not contained in the other).

Suppose that overlaps between cones Y_i and A are small in the sense that for any path P in the intersection of translates $\widetilde{Y}_i, \widetilde{A}$ in $\widetilde{X}$, either $\nabla_{Y_i}(P) < \beta|Y_i|$ or $\widetilde{Y}_i \subset \widetilde{A}$.

Then $A^* \to X^*$ is a local isometry, in the sense that $\widetilde{A}^* \to \widetilde{X}^*$ is a convex subcomplex.

Proof. Let γ be a geodesic in $\widetilde{X}^*$ whose endpoints lie on $\widetilde{A}^*$. Consider a minimal complexity diagram D between γ and a variable path σ in $\widetilde{A}^*$.

If D has no cone-cells then D must be trivial in the sense that $\gamma = \sigma$.

There is a geodesic path γ' in D with the same endpoints as γ such that γ', γ together bound a maximal square subdiagram of D. Let E denote the subdiagram of D bounded by σ, γ'. Observe that E cannot have any corners of generalized squares along σ or γ'. The former is impossible since σ is allowed to vary in $\widetilde{A}^*$, so a generalized square can be pushed through it. The latter is impossible since we could again push past the square to make E smaller.

There are no outerpaths of positively curved θ-shells along either γ' or σ. The latter is excluded since $A^* \to X^*$ has no missing θ-shells by hypothesis. The former is excluded by our hypothesis that X^* has short innerpaths, so if such a θ-shell existed it would contradict our hypothesis that γ' and hence γ is a geodesic. We conclude that there at most two features of positive curvature in E, so E is either a single cone-cell or 0-cell or ladder by Theorem 3.36.

In fact we seek to conclude that $\gamma' = \sigma$ so E is a subdivided interval. And we now apply our strengthened α, β hypothesis to reach this conclusion.

Consider a cone-cell C in E that maps to the cone Y. The path $\partial_p C$ is the concatenation $wxyz$ where w, y are either trivial (when C is an initial or terminal cone-cell in the ladder) or are cone-pieces, and w is a path between $\widetilde{C}$ and $\widetilde{A}$, and z is a subpath of γ'. By our hypotheses, $|wxy| < \alpha|Y| + \beta|Y| + \alpha|Y| < \frac{1}{2}|Y|$. The diagram D and hence E is reduced, so $wxyz$ is essential and hence $|wxyz| \geq |Y|$. Thus since $|wxy| < \frac{1}{2}|Y|$ we see that $|z| > \frac{1}{2}|Y|$ which contradicts that γ' is a geodesic. □

Here is a further variation and strengthening of the previous result:

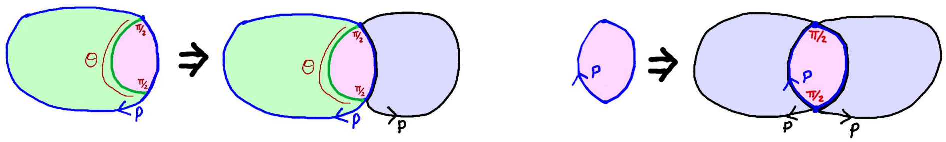

FIGURE 70. We glue a phony Y_i-cone-cell with boundary path P along the outerpath of a diagram to reach a contradiction.

Lemma 3.53 (Persistence of Superconvexity)**.** *Suppose X^* is small-cancellation with internal cone-cells having negative curvature.*

Suppose A_i^* satisfy a strong form of no missing θ-shells excluding those with $\theta \leq \pi$.

Then (conjugates of) intersections of conjugates of $\pi_1 A_i^*$ and $\pi_1 A_j^*$ in $\pi_1 X^*$ are represented by components of the fiber-product $A_i \otimes_X A_j$.

Suppose moreover that X^* has short innerpaths, then $\widetilde{A}_i^*$ is superconvex.

4. Torsion

4.1. **Cones Embed.**

Theorem 4.1. *Suppose the cubical presentation $\langle X \mid \{Y_i\}\rangle$ has an angling system so that the conditions of Theorem 3.36 are satisfied for each minimal complexity diagram $D \to X^*$, and such that Condition 3.36.(5) is strengthened so that nil-external corners have $\sphericalangle \geq \frac{\pi}{2}$. Then each Y_i embeds in $\widetilde{X}^*$.*

The requirement that nil-external corners have $\sphericalangle \geq \frac{\pi}{2}$ will ensure that an outerpath of a positively curved θ-shell is not a cone-piece. This property is at the heart of the proof. We note that such angles are exactly $\frac{\pi}{2}$ for the split-angling system.

Proof. Consider a minimal area disk diagram D whose boundary path P is an essential non-closed path in Y_i, and suppose the complexity of D is minimal among all diagrams whose boundary path is path-homotopic to P in Y_i.

By Theorem 3.40, $\bar{D}$ either consists of a single 0-cell or single cone-cell, or $\bar{D}$ contains a spur or generalized square with outerpath on $\partial_p D$, or $\bar{D}$ contains a positively curved shell.

The first two possibilities lead to an immediate contradiction: If $\bar{D}$ is a single 0-cell then $P = \partial_p D$ is a trivial path and hence closed in Y_i, and if $\bar{D}$ is a single cone-cell then this cone-cell must actually lie in Y_i, and hence again, P is closed in Y_i.

Since $Y_i \to X$ is an immersion, we can pass to a path $P' \to Y_i$ with backtracks removed (or we could have assumed P had no backtracks to begin with), so we can assume $\bar{D}$ has at most one spur - at the very basepoint of P.

We now consider the main case where $\bar{D}$ contains a shell or the corner of a generalized square. In the square case, since $Y_i \to X$ is a local isometry, we could push across this generalized square and obtain a new diagram with smaller area. In the shell case we find that it is either replaceable by a square diagram or it can be absorbed into Y_i (and hence in either case, D is not minimal). Indeed, the outerpath Q of a θ-shell R cannot be a cone-piece, for then, as in Figure 70, we can attach a cone-cell $E \to Y_i$ to D along Q there would be a disk diagram with an internal positively curved cone-cell, and likewise, when D is a single cone-cell, we can surround D by two such cone-cells E_1, E_2.

There is a technicality here: To argue in terms of disk diagrams we need a closed path $P' \to Y_i$ for $\partial_p E$ (and likewise P_i' for E_i). When Y_i is not CAT(0) then we can always extend any path in Y_i to an essential closed path. When Y_i is CAT(0) we would have to use an inessential path like PP^{-1} which is thus replaceable and unsatisfactory, or we could use phony cone-cells - and we refer to Lemma 5.27 for more on this approach. ◻

4.2. **Torsion.**

Theorem 4.2. *Let $\langle X \mid \{Y_i\}\rangle$ be a cubical presentation with the following properties: Then each torsion element in $\pi_1 X^*$ is conjugate to an element represented by a closed path γ such that γ^n is a closed path in some Y_i.*

1. *Each element γ of $\pi_1 X$ has a locally convex core $C = C(\gamma)$ with the property that any ball of radius r disconnects $\widetilde{C}$ for some $r = r(X)$.*
2. *For any essential positively curved θ-shell with outerpath Q and innerpath S, we have $|Q| - |S|$ exceeds twice the maximum length of a piece.*

Remark 4.3. We note that Hypothesis (1) holds when X is compact and $\widetilde{X}$ is δ-hyperbolic, and that $r = r(\dim(X), \delta(X))$ in this case.

We suspect that Theorem 4.2 holds for cubical presentations satisfying more general small-cancellation conditions. Perhaps one can proceed by proving an asphericity result, and then showing that a torsion element would fix a cone-point.

Proof. Let us first prove the Theorem when X is 1-dimensional, for then the constant $r = 0$, and things are simplified.

Let g be a torsion element in $\widehat{X}^*$. Let γ be a shortest combinatorial path in X representing g. Thus γ^n is nullhomotopic in X^* for some $n \geq 2$.

Consider a minimal area disk diagram D for γ^n. Note that D cannot be a diagram in X, for then γ would be a torsion element in $\pi_1 X$ in which case γ is trivial. In particular, we assume that D is nontrivial.

By Theorem 3.40, either D consists of a single cone-cell, and we have verified our conclusion, or D contains two or more positively curved shells.

If the outerpath of the shell lies in the concatenation of a cyclic permutation of γ with a path α whose endpoints are at distance (in $\widetilde{X}$) bounded by the length of a maximal piece, then a geodesic for α concatenated with the innerpath provides a shorter representative for g. So we can assume this doesn't happen.

Thus the outerpath is the concatenation of a cyclic permutation of γ followed by a long path β whose endpoints in $\widetilde{X}$ are at distance exceeding the length of a maximal piece. Now we can attach n copies of this shell to γ^n using the $\mathbb{Z}_n$ action. Successive copies overlap along β. But β cannot be a piece, and so successive copies of this shell map to the same cone Y_j. Consequently, the path γ^n maps to Y_j.

The idea for the case where X is d-dimensional is similar. We require that $|Q| - |S| - 2r$ exceeds twice the longest piece.

Now, either there is a way to shorten our representative for (a cyclic permutation of) γ by replacing some subpath of γ by a shorter path that runs along some some initial part of γ and then jumps across the annulus at the beginning and the end and then uses the innerpath S, or we find that there is an impossibly long piece at the continuation, and then we find that all the cone-cells map to the same cone as above. □

Corollary 4.4. *Let $\widehat{X} \to X$ be a regular cover, such that $Y_i \to X$ lifts to an embedding in $\widehat{X}$. Suppose the cubical presentation $\langle X \mid \{Y_i\}\rangle$ is small-cancellation. Then $\pi_1 X^*$ is virtually torsion-free.*

Proof. Consider the cover of X^* induced by $\widehat{X}$. A torsion element of $\pi_1 X^*$ is conjugate to a closed path γ such that γ^n is a closed path in some cone Y_i. Suppose $n \geq 2$ is minimal with this property, so that γ is not a closed path in Y_i (and is thus a nontrivial torsion element of $\pi_1 X^*$). Since $Y_i \to X$ lifts to an embedding $Y_i \to \widehat{X}$, we see that γ does not lift to a closed path in $\widehat{X}$ and so γ does not represent an element of $\pi_1 \hat{X}^*$. □

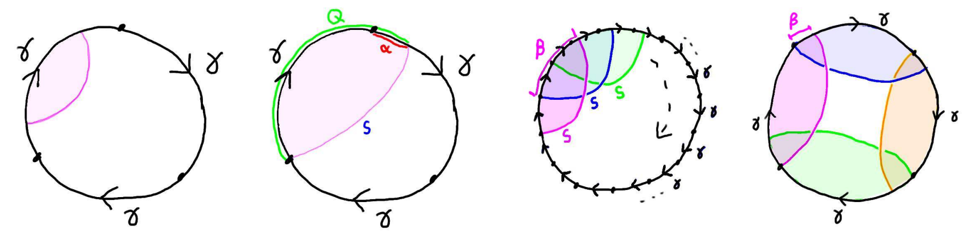

FIGURE 71.

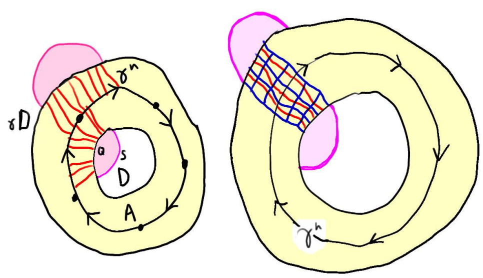

FIGURE 72. Torsion in cones in higher dimensions

Remark 4.5. Note that Corollary 4.4 holds under various more general hypotheses: For instance when each $\pi_1 Y_i \to \pi_1 \widehat{X}$ is malnormal.

4.3. **Relative Hyperbolicity of Quotient.**

Theorem 4.6. *Suppose the nonpositively curved cube complex X is compact, and that $\pi_1 X$ is word-hyperbolic. Let X^* be the complex associated to the cubical presentation $\langle X \mid Y_i \rangle$.*

Suppose that essential cone-cells in reduced diagrams $D \to X^*$ have negative curvature. Then $\pi_1 X^*$ is hyperbolic relative to the images of $\pi_1 Y_i \to \pi_1 X$.

Remark 4.7. When outerpaths of shells are $< \frac{1}{2}$ as in Condition 5.4, then one actually obtains a Dehn's algorithm, and so the conclusion is fairly immediate using Osin's criterion [Osi06] of a relative linear isoperimetric function, and is thus immediate when each Y_i is compact and there are finitely many Y_i.

Proof. This follows as in [Ger87] by applying Theorem 3.15. A variant of his proof utilizing an upper bound on the radius of a zero-curvature ball must be employed. □

5. NEW WALLS AND THE $B(6)$ CONDITION

5.1. **Introduction.**

In this section, we impose further hypotheses on the cones Y_i in a cubical presentation $X^* = \langle X \mid Y_i \rangle$. The main hypotheses is that each Y_i is a wallspace whose walls are collections of hyperplanes, and furthermore, the walls in Y_i have certain "convexity" properties - in the sense that any path in Y_i that starts and ends on the same wall, is either homotopic into that wall, or is "long" from a piece-count viewpoint. The wallspace cones and their properties allow us to define walls in $\widetilde{X}^*$, and these walls are the central focus of the section.

We define a notion of "length" of a path in a cone in Section 5.2 and then describe the $B(6)$ wallspace structure on cones in Section 5.3. The construction of walls in $\widetilde{X}^*$ and quasiconvexity properties of these walls in $\widetilde{X}^*$ are examined in Sections 5.5, 5.6 and 5.9. Conditions that imply that the set of walls is sufficiently rich to "fill" $\widetilde{X}^*$ are examined in Sections 5.11 and 5.12. Malnormality properties of the wall stabilizers are treated in Sections 5.14, 5.15, 5.16, and 5.17.

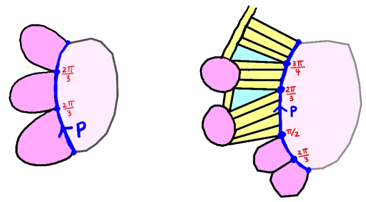

FIGURE 73. The total defect $\Omega_i(P)$ of the path $P \to Y_i$ on the right is $\leq \frac{\pi}{3} + \frac{\pi}{2} + \frac{\pi}{3} + \frac{\pi}{4}$.

5.2. **Total defects of paths in cones.** In classical $C(6)$ small-cancellation theory, a quick measure of the extent to which a path $P \to \partial R$ travels around R is the inifimal number n where $P = P_1 \dots P_n$ is an expression of P as the concatenation of pieces between other relators and our relator R. This can be thought of concretely in terms of a reduced diagram as on the left of Figure 73. If we assign $\frac{2\pi}{3}$ angles at the internal corners of R along P, then we obtain a total defect of $(n-1)\frac{\pi}{3}$. We shall now generalize the "piece length of P in R by focusing on the "total defect".

Consider the cubical presentation $\langle X \mid Y_i \rangle$, and suppose we have chosen a fixed method of assigning angles (for instance, the split-angling, or the grade-angling) on rectified disk diagrams in X^*.

Consider a path $P \to Y_i$. The *defect of P in Y_i* which we denote by $\Omega_i(P)$ is the infimum of the sum of defects of angles along the path P in a cone-cell C_i mapping to Y_i within angled rectified diagrams $\bar{D}$ that have P as an internal path on the boundary of C_i. Of course, we also assume that the original diagram D has minimal complexity (or as usual, it has no cancellable pairs of cone-cells, and no removable bigons or squares that are absorbable into cone-cells). See Figure 73.

It may be that P doesn't occur as an internal path along the boundary of a cone-cell mapping to Y_i. This could be remedied with a trick, by using a cone-cell with boundary path $P^{-1}P$, however $P^{-1}P$ is null-homotopic so the cone-cell can be removed (so we wouldn't have minimal complexity) without affecting the boundary path. A rigorous alternative (leading to the same result), is to add a new 1-cell to Y_i for each pair of vertices in Y_i, and to attach a copy of this new 1-cell to X along the images of these vertices. (We refer the reader to Lemma 5.27 where a similar approach is used to force 1-cells dual to the same wall to lie on the boundary path of a (newly added) essential cone-cell.) Doing so yields a new cubical presentation with the property that each such path P could now arise as the boundary of a (genuine) cone-cell that couldn't be compressed onto the boundary path of its containing disc-diagram and removed. We note that the new cubical presentation deformation retracts onto the original in a reasonable sense.

It may be that some edge e of $P \to Y_i$ cannot arise within the interior of a diagram D. Indeed, let $\bar{e}$ denote the image of this edge in X, then this happens precisely if $\bar{e}$ is not the image of some edge of Y_j or some other edge of Y_i, and that $Y_i \to X$ is a local surjection at e so each square along $\bar{e}$ lifts to Y_i (thus e does not arise within a wall-piece). In this case the infimum of defects along P as it arises in such diagrams, is the infimum of the empty set. Thus $\Omega_i(P) = +\infty$.

For 1-cells e, e' in Y_i, we define $\Omega_i(e, e')$ to be $\inf(\Omega_i(P))$ where $P \to Y_i$ is a path whose first and last 1-cells are e, e'. We define $\Omega_i(v, v')$ similarly for 0-cells v, v' in Y_i. For hyperplanes E, E' of Y_i, we define $\Omega_i(E, E')$ to equal $\inf(\Omega_i(e, e'))$ where e, e' vary among 1-cells dual to E, E'.

5.3. **Generalization of the $B(6)$ condition.** We will now add more elaborate structure to generalize other aspects of the classical $C'(\frac{1}{6})$-$T(3)$ and $C'(\frac{1}{4})$-$T(4)$ metric small-cancellation theories, to higher dimensions along the lines of the $B(6)$-$T(3)$ and $B(4)$-$T(4)$ theories we considered in [Wis04].

Definition 5.1. [Generalized $B(6)$]
We say $\langle X \mid \{Y_i\} \rangle$ satisfies the *generalized $B(6)$ condition* provided the following conditions hold:

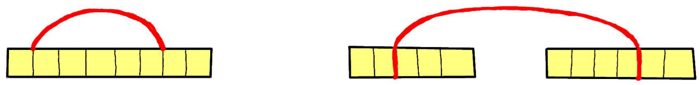

Figure 74. Hyperplane-convexity vs. wall-convexity in cones

(1) [Cubical Rel.Pres.] X is a nonpositively curved cube complex, and each $Y_i \to X$ is a local isometry of cube complexes.
(2) [NPC] X satisfies the nonpositive [negative] curvature condition for rectified diagrams: So $\Omega_i(P) \geq 2\pi$ [respectively $\Omega_i(P) > 2\pi$] for each essential closed path $P \to Y_i$. And each internal 0-cell and shard in each rectified reduced diagram has nonpositive curvature.
(3) [Wallspace Cones] Each Y_i is a wallspace (see Definition 7.1) where each *wall* in Y_i is the union of a collection of disjoint embedded hyperplanes, each of which maps to an immersed hyperplane in X. Each such collection separates Y_i. Each hyperplane in Y_i lies in a unique wall.
(4) [Hyperplane Convexity] If $P \to Y_i$ is a path that starts and ends on vertices on 1-cells dual to a hyperplane H of Y_i and $\Omega_i(P) < \pi$ then P is path-homotopic into its carrier $N(H)$ in Y_i.
(5) [Wall Convexity] Let S be a path in Y that starts and ends with 1-cells dual to the same wall of Y. If $\Omega_Y(S) < \pi$ then S is homotopic into that wall (and hence into the carrier of one of its hyperplanes).
(6) [Equivariance] For each Y the wallspace structure on Y is preserved by $\text{Aut}_X(Y)$.

We note that Condition (4) is not implied by Condition (5) since the latter requires that the path start and end with dual 1-cells, and not merely start and end at endpoints of dual 1-cells (see Figure 74).

Definition 5.2. We define $\text{Aut}_X(Y)$ to equal the group of automorphisms $\phi : Y \to Y$ such that we have a commutative diagram: $\begin{tikzcd} Y \arrow[r, "\phi"] \arrow[dr] & Y \arrow[d] \\ & X \end{tikzcd}$. Note that $\text{Aut}_X(\widetilde{Y})$ equals $\text{Stab}_{\pi_1 X}(\widetilde{Y})$, and that $\text{Aut}_X(Y) \cong \text{Normalizer}_{\text{Aut}_X(\widetilde{Y})}(\pi_1 Y) \,/\, \pi_1 Y$. In practice, this situation is simplified since we choose Y so that $\pi_1 Y$ is normal in $\text{Aut}_X(\widetilde{Y})$.

We emphasize that in the definition of "piece", we must treat two lifts of $\widetilde{Y}$ as identical if they differ by an element of $\text{Stab}_{\pi_1 X}(\widetilde{Y})$. This generalizes the way relators that are proper powers are treated in the classical case.

For instance, sticking a spur on to a relator Y can drastically destroy the small-cancellation properties, since it can artificially decrease the size of $\text{Aut}_X(Y)$. Indeed, two lifts $P \to Y$ that originally differ by an automorphism of Y will not form a piece, however after adding the spur to Y, the automorphism will no longer exist.

Remark 5.3. Definition 5.1.(6) is required to make the wall equivalence relation (defined below) on hyperplanes in $\widetilde{X}^*$ agree locally with wall structure on each Y. If some cone $Y \subset \widetilde{X}^*$ doesn't have this property, then the decomposition of the hyperplanes of $\widetilde{X}^*$ into equivalence classes of walls, would provide a coarser decomposition of the hyperplanes of Y than the actual wallspace of Y.

It is natural but apparently unnecessary to impose the global requirement that $gW \cap Y = W \cap Y$ for any wall W of $\widetilde{X}^*$ that crosses Y and any Y in $\widetilde{X}^*$. It is conceivable that gW hits Y accidentally within some hyperplane (so $g^{-1}Y$ plays the role of an inessential cone for W). However Condition (6) keeps things OK in the case that $g^{-1}Y$ is an essential cone of W, since the small-cancellation forces two essential cones meeting along distinct hyperplanes in the same wall of W to be equal to each other.

We impose another related condition later in Hypothesis (3) of Theorem 5.64.

Definition 5.4 (Short innerpaths)**.** An essential property of the classical $C'(\frac{1}{6})$ theory is generalized by the following "metric small-cancellation" condition, and we refer the reader to Definition 3.47 where this notion is treated in depth. The cubical presentation $\langle X \mid Y_i \rangle$ has *short innerpaths* if the following holds for each cone Y_i:

If $\Omega_i(S) < \pi$ then for any local geodesic $S' \to Y_i$ that is path-homotopic to S, and for any path $Q \to Y_i$ such that the concatenation SQ is an essential closed path, we have $|S'| < |Q|$.

5.4. **Cyclic Quotients and the $B(6)$ condition (to be developed).** Cyclic quotients are especially accessible to the $B(6)$ small-cancellation theory because a wallspace structure on the cones is then often readily achievable.

(Another case to consider are relators with large abelian automorphism groups. We've begun treating this in Section **??**.)

One subdivides the cube complex X, so that the local-isometry $Y \to X$ (representing a cone associated to the cyclic subgroup to quotient) has twice the number of separating hyperplanes. Now hyperplanes (that don't already separate) can be paired in a manner respecting almost all small-cancellation conditions provided a girth condition is satisfied.

Note:

1) While the $B(6)$ condition is easily satisfied here, hierarchical conditions might not be attainable without some sort of evenness assumptions at various levels.

2) An interesting example is a twisted product $Y = B \rtimes S$ where B is a CAT(0) ball and S is a subdivided circle (this is just a B-bundle over S.) Presumably $B \rtimes S$ is a product when X is special, but in general, it is possible to have some hyperplanes that wrap multiply around the circle basespace.

For Dehn fillings, it is natural to use the entire infinite cylinder as the relator.

Relative hyperbolicity gives an upper bound on diameters of pieces between planes. This gives an upper bound on diameters between (distinct) cylinders. A Dehn filling corresponding to an element that is long enough relative to these pieces and the walls gives us a small-cancellation quotient, and should be virtually special because of the subdivision and paired splicing.

3) These ideas can be generalized to abelian quotients.

5.5. **Embedding properties of the cones and hyperplane carriers.** The aim of this subsection is to show that certain very short circuits of cones and hyperplane carriers do not exist in X^*. We are using the split-angling, though similar statements hold in most cases for the grade-angling.

The following is a restatement of Theorem 4.1.

Lemma 5.5 (Cones embed again)**.** *Let $\langle X \mid \{Y_i\} \rangle$ satisfy Condition 5.1.(2) and the short innerpath condition of Definition 5.4. Then each Y_i embeds in $\widetilde{X}^*$.*

Proof. Consider a path $P' \to Y_i$ that lifts to a closed path in $\widetilde{X}^*$. Consider a disk diagram $D \to X^*$ that has minimal area among all diagrams whose boundary path P is path-homotopic to P' in Y_i. Specifically, the number of cone-cells in D is minimal, and the number of squares is minimal as well, for this fixed number of cone-cells.

By Theorem 3.40, $\bar{D}$ is either trivial in the sense that it is a single 0-cell or cone-cell, or else $\bar{D}$ has two or more generalized square corners, spurs, or shells on its boundary. Moreover, the outerpaths of these features of positive curvature have disjoint interiors.

If $\bar{D}$ is trivial then $\bar{D}$ cannot be a single 0-cell, or $P \to Y_i$ is closed. Since corners and spurs in $\bar{D}$ could be absorbed into Y_i making $\bar{D}$ even smaller, we can assume that $\bar{D}$ has shells. Now apply Condition 5.4 to the outerpath of a shell whose interior is disjoint from the endpoints of P. Its innerpath has angle sum $< \pi$ and its outerpath must have angle sum $< \pi$ which is impossible. We thus reach the conclusion that this shell must have been in Y_i itself, and we obtain a smaller complexity diagram - with fewer cone-cells, and this contradicts our choice of D.

Eventually we reach a contradiction. □

Lemma 5.6. *Suppose that $\langle X \mid Y_i \rangle$ is a small-cancellation complex, and that no essential path in a cone-cell is the concatenation of three pieces. Then the intersection of two cones in $\widetilde{X}^*$ is connected.*

Definition 5.7 (No acute corners)**.** An angling system has *no acute corners* if the defect of each cone-cell angle is $\leq \frac{\pi}{2}$.

We note that the split-angling has no acute corners, as does the grade-angling provided that each grade is ≥ 4.

Remark 5.8. The three piece hypothesis in Lemma 5.6 holds when the small-cancellation complex $\langle X \mid Y_i \rangle$ has no acute corners.

Indeed, suppose $P_1P_2P_3$ is an essential path in the cone Y that is the concatenation of three pieces. Let A be a 2-cell with $\partial_p A = P_1P_2P_3$ and form a disk diagram D by surrounding A by three rectangles and/or cone-cells according to whether P_i are wall-pieces, non-contiguous cone-pieces and/or contiguous cone-pieces. Since X^* is small-cancellation, the angles on the diagram D yield nonpositively curved internal cone-cells. However, as X^* has no acute corners, each internal angle of A has defect $\geq \pi$, so $\kappa(A) \geq \frac{\pi}{2}$. The analogous reasoning holds for fewer than 3 pieces.

Proof of Lemma 5.6. Suppose that $Y_1 \cap Y_2$ is not connected. Let D be a minimal area disk diagram between paths γ_1, γ_2 that start and end on points p, q in distinct components of $Y_1 \cap Y_2$. In particular, D is minimal complexity among all such possibilities where p, q are allowed to vary within two specific components of $Y_1 \cap Y_2$, and where the paths γ_1, γ_2 are allowed to vary among paths in Y_1, Y_2 starting and ending on p, q.

Since our aim is to show that D is a square diagram, and thus lies in both Y_1 and Y_2 by local convexity, we can assume without loss of generality that D is spurless. Likewise, we can assume that D has no outerpath of a generalized square in either γ_1 or γ_2, for then we could find a smaller square area diagram. Observe that D has no positively curved θ-shell with an outerpath in γ_i for then it could be absorbed into Y_i or replaced by a square subdiagram thus leading to a lower complexity diagram. Thus D has at most two locations for an outerpath of a positively curved cell. Consequently D is a ladder or a single cone-cell or a single vertex by Theorem 3.36.

A single vertex would imply that $p = q$, which is impossible. A single cone-cell is impossible because if its boundary path is the concatenation of two pieces, then its curvature is $< 2\pi$ since the defect at the two transitions is $< \pi$. Thus it can either be replaced by a square diagram, or it is absorbable into one of the two cones, and so either way, there is a smaller area complexity diagram.

If D is a ladder, then as above, there are no spurs, so there are 0-shells at each end. then a θ-shell at one end is bounded by the concatenation of three pieces. We now use the hypothesis that no essential path in a cone-cell is the concatenation of three pieces. □

Lemma 5.9 (Intersections of Cones)**.** *Let $\langle X \mid Y_i \rangle$ be as in Lemma 5.6. The intersection between distinct cones in $\widetilde{X}^*$ is CAT(0).*

Proof. Any path P in the intersection $Y_i \cap Y_j$ of cones has the property that $\Omega_i(P) = 0 = \Omega_j(P)$. Consequently, any such closed path P must bound a cubical diagram D in Y_i by Condition 5.1.(2). Then by local convexity, D is a diagram in Y_j as well so $D \subset Y_i \cap Y_j$ in $\widetilde{X}^*$. Now, $Y_i \cap Y_j$ is locally-convex and hence nonpositively curved since both Y_i and Y_j are. Thus, $Y_i \cap Y_j$ is CAT(0) since it is simply-connected and nonpositively curved. □

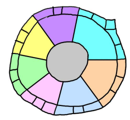

FIGURE 75. A self-crossing hyperplane

Definition 5.10 (Embedding Properties of Hyperplanes in Cones)**.** In many cases, the following conditions follow from Condition 5.1.(5) because single edges tend to be pieces, but simple examples show that in general these conditions must be hypothesized to achieve some of our desired conclusions.

(1) [2-sided] Each hyperplane of each Y_i is 2-sided, in the sense that its dual 1-cubes can be oriented so that dual 1-cubes that are opposite sides of a 2-cube are oriented the same way.
(2) [No self-intersection] No hyperplane H in Y_i is dual to all 1-cubes on the boundary of a 2-cube in Y_i.
(3) [No self-osculation] No hyperplane H in Y_i is dual to two distinct 1-cubes that share a 0-cube. (By nonpositive curvature, this implies Condition (2).)

Remark 5.11. Condition 5.10.(3) is a consequence of Definition 5.1.(4) by letting P denote the trivial path. We note that the same would hold if we restricted Definition 5.1.(4) to paths P starting and ending on the same side (in the 2-sided case). Accordingly, Condition 5.10.(2) holds as well.

Example 5.12 (A self-crossing hyperplane)**.** In Figure 75 we illustrate a 2-dimensional cubical presentation $\langle X \mid Y_1, \ldots, Y_8 \rangle$. X^* is C(6) and all cones are embedded and have well-behaved hyperplanes, and $\widetilde{X}^* = X^*$, but there is a hyperplane in $\widetilde{X}^*$ that self-crosses. By subdividing along the outside, one can arrange that $C'(\alpha)$ be satisfied for arbitrary $\alpha > 0$.

Reformulated appropriately, the following is a special case of Theorem 12.16.

Lemma 5.13 (Embedding Hyperplane Carriers)**.** *Let $\langle X \mid Y_i \rangle$ be a small-cancellation with no acute corners, and suppose Conditions 5.1.(4) and 5.1.(2) hold. Let H be an immersed hyperplane in the nonpositively curved cube complex of $\widetilde{X}^*$.*

(1) *If Condition 5.10.(1) is satisfied then H is 2-sided.*
(2) *If Condition 5.10.(2) is satisfied then H does not self-intersect.*
(3) *If Condition 5.10.(3) then H does not self-osculate.*

In conclusion, if all three conditions are satisfied, then $N(H) \to \widetilde{X}^*$ is an embedding and $N(H) \cong H \times [-1, 1]$.

Note that we define the *carrier* of an immersed hyperplane $Y \to Z$ to be the union of copies of cubes of Z whose midcubes are the cubes of Y. These copies of cubes are glued together along subcubes in $N(H)$ as they are in Z. The result is a (possibly twisted) I-bundle over Y. We note that $N(H)$ is nonpositively curved, and that $N(H) \to Z$ is a local-isometry.

Example 5.14. When the immersed hyperplanes of X are not 2-sided, in general, H might not be 2-sided so $N(H) \not\cong H \times I$. For instance, let X be a moebius strip obtained by identifying opposite sides of a square with a twist, and let Y_n denote a connected n-fold cover of X. Then $\langle X \mid Y_n \rangle$ is small-cancellation for sufficiently large n. But the cube complex of $\widetilde{X}^*$ is still a moebius strip for odd n.

Similar examples can be concocted to illustrate how $N(H) \to Y$ can fail to be an embedding.

Proof of Lemma 5.13. We will concentrate on establishing the conclusion where all three conditions are satisfied.

Consider two points $p, q \in N(H)$ that map to the same point of $\widetilde{X}^*$. Let $D \to X^*$ be a disk diagram of minimal complexity among all diagrams whose boundary path $P \to N(H)$ starts and ends on p, q. If D

has a corner of a generalized square, then since $N(H) \to X$ is a local-isometry, we see that this square could be absorbed into $N(H)$ and a lower complexity diagram could be produced.

Suppose that D has a θ-shell C with boundary path QS, with innerpath S and outerpath Q and such that C maps to the cone Y. First consider the case where $Q \to N(H)$ does not pass through any dual 1-cell of $H \to N(H)$. Then either Q is a wall-piece in the cone Y or the square ladder of Y containing Q along its external boundary maps into Y. In the former case, since there are no acute corners, we see that $\Omega(QS) \leq \Omega(Q) + \frac{\pi}{2} + \Omega(S) + \frac{\pi}{2} \leq 0 + \frac{\pi}{2} + \theta + \frac{\pi}{2} < 2\pi$, which implies by Condition 5.1.(2) that $\partial_p C = QS \to Y$ is null-homotopic, so C can be replaced by a square-diagram and we can reduce the complexity of D. In the latter case, Condition 5.1.(4) likewise implies that $QS \to Y$ is null-homotopic so C can be replaced by a square-diagram and we can reduce the complexity of D.

We next consider the case where Q passes through a dual 1-cell of H. Then there exists $L \to N(H)$ where $L \cong I_n \times I$ is a square-ladder that is dual to H, and $Q \to N(H)$ factors as $Q \to L \to N(H)$, and by the local convexity of the cone Y, we have $L \to \widetilde{X}^*$ actually factors as $L \to Y \to \widetilde{X}^*$. Therefore Condition 5.1.(4) again implies that C can be replaced by a square diagram, to reduce the complexity as above.

We emphasize that in the special case of the above situation where S is trivial and Q is a wall-piece in Y, and C is replaced by a square diagram, we see that the endpoints p', q' of Q in $N(H)$ are actually equal to each other, and moreover, following Lemma 2.2, the 1-cubes dual to H at $p' = q'$ must be equal to each other as well. In the case where Q is not a wall-piece (this also includes the situation where Q passes through a dual 1-cell) we see that Q lies on a square-ladder L dual to H such that $L \to C$. It is in this case that we must employ Condition 5.10.(3) to see that the 1-cubes dual to H at p', q' are actually equal in C (as L maps to a cylinder in C) and so these 1-cubes are equal in $\widetilde{X}^*$ and hence in H itself.

In particular, we are able to remove C, Q from D, P to obtain a lower complexity diagram D', P' such that P' has the same endpoints as P in $N(H)$.

We now focus on the 2-sidedness of H to see that $N(H) \cong H \times [-1, 1]$. Consider a minimal complexity diagram D whose boundary path P is a path on $N(H)$ that passes through an odd number of dual 1-cubes of H. The additional presence of Condition 5.10.(1), implies that for any hyperplane H_i of Y_i, any closed path in $N(H_i)$ passes through dual 1-cubes of H_i an even number of times. Therefore, when QS is the boundary path of a cone-cell C in Y_i that lies in $N(H_i)$, we see that Q and S pass through dual cubes of H the same number of times (where H_i is a component of $H \cap Y_i$). In particular, when S is the trivial path, then the smaller complexity diagram D', P' that we obtained above has the property that P and P' have the same parity.

We note that in the case where D is a square-diagram, it is clear that P passes through an even number of dual 1-cubes of H, since the dual curves in D provides a pairing of such dual 1-cubes. □

The following result will play a fundamental role in understanding the walls of $\widetilde{X}^*$ by revealing cancellable pairs of cone-cells in certain disk diagrams.

Lemma 5.15. *Each hyperplane H in $\widetilde{X}^*$ has connected intersection with each cone $Y \subset \widetilde{X}^*$.*
Moreover $N(H) \cap Y = M(H \cap Y)$, where $M(K)$ denotes the carrier in Y of a hyperplane K of Y.

Remark 5.16. When each Y_i is a pseudograph (see Definition 5.46) in the sense that its hyperplanes are CAT(0), one can reach the following stronger result:

Any path $P_h \to N(H)$ whose endpoints lie on Y, is path-homotopic through a square diagram in $\widetilde{X}$ to a path $P_y \to Y$. This can be used to simplify some of the proofs below.

However there are simple examples showing that such a square-diagram might not exists in general. For instance, the square ladder at the top of Figure 76 cannot be pushed towards the cone at the bottom without passing through three essential cone-cells. The reader can think of the light intermediate diagram as being made of squares.

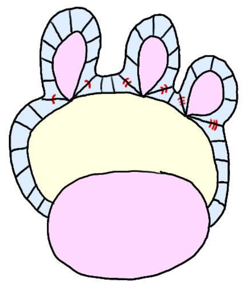

FIGURE 76.

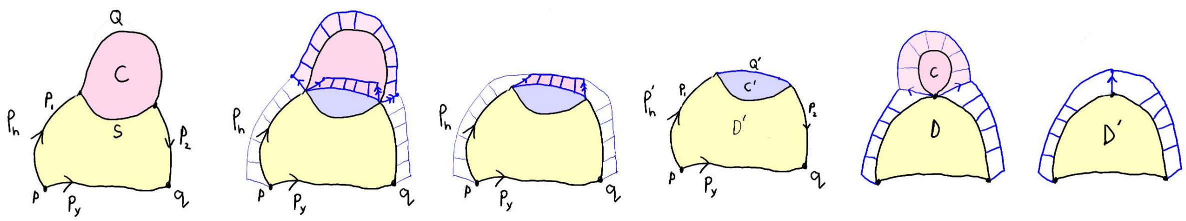

FIGURE 77. The four diagrams on the left illustrate how a hyperplane can be pushed across a cone-cell whose outerpath is not a real wall-piece. The right two figures illustrate the situation when the innerpath is trivial.

Proof. Let p, q be points in $N(H) \cap Y$ and let $D \to X^*$ be a diagram between paths $P_h \to N(H)$ and $P_y \to Y$ from p to q, and suppose that d has minimal complexity among all such diagrams and paths joining p, q.

Observe that D has no corner of a generalized square along P_h or P_y, for then, by local-convexity, we could absorb the square to produce a lower complexity diagram. Likewise, observe that D has no outerpath of a θ-shell along P_y, for then it could either be absorbed into Y or replaced by a square diagram by Condition 2, thus decreasing the complexity. Similarly, D cannot have an outerpath of a θ-shell C along P_h, where C maps to Y_i. Indeed, suppose $\partial_p C = QS$, where Q is the outerpath and S is the innerpath. Then $\Omega_i(S) \leq \theta < \pi$.

If Q is a wall-piece of H in Y_i, then $\Omega(QS) \leq \theta + \pi \leq 2\pi$ which contradicts Condition 2. Thus we can assume that the ladder in $N(H)$ containing Q lies in Y_i.

Thus Condition 5.1.(4) implies that S is square homotopic in X to path Q' on $N(H \cap Y)$. Let C' denote a square diagram whose boundary path is SQ'. Letting $P_h = P_1 Q P_2$, we observe that $P'_h = P_1 Q' P_2$ is still a path in $N(H)$. This uses Hypothesis 3 to see that Q and Q' are external boundary paths of square ladders whose initial and terminal 1-cubes are identical. Finally, by replacing the path Q with the path Q', and replacing the cone-cell C with the square diagram C', we obtain a new diagram $D' \to X^*$ between P_y and P'_h of lower complexity. We refer the reader to Figure 77. Note, that this procedure can be followed whenever an innerpath S of a cone-cell C, is homotopic into $N(H)$ (and we shall need to consider it again below). We also note that this procedure is meaningful even in the case where S is a trivial path. In that case, Condition 3 allows us to immediately short-cut P_h and replace Q by S itself to remove C and reduce the complexity.

We have shown that the outerpath of a positively curved cell must contain an endpoint p, q of P_y, P_h. Thus D is a ladder by Theorem 3.36, or it is a single cone-cell or 0-cell.

As we aim to show that P_h is actually also a path in Y, so p, q lie in the same component of $N(H) \cap Y$, we can assume that there is no spur on D (for otherwise, we can vary our choices of p, q and remove the spurs.

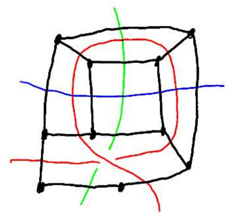

FIGURE 78. A nonembedded wall in $\widetilde{X}^*$.

Let us now consider the first cone-cell C at one side of D (the reasoning will also dispense with the case that D is a single cone-cell). The boundary path $\partial_p C$ is the concatenation $U_y U_c U_h$ where U_y is a cone-piece from Y, and U_h is a wall-piece from H, and U_c is a (possibly trivial) cone-piece from the next cone, or just a trivial path when D is a single 2-cell.

Since $\partial_p C$ cannot be the concatenation of fewer than four pieces (which holds as in Remark 5.8), we see that one of these pieces is fraudulent. It cannot be U_c, for otherwise their would be an absorbable pair of cone-cells in the diagram, and we would be in a simpler situation. If it is U_y, then we can absorb C into Y to obtain a smaller complexity diagram, and otherwise U_h is not a genuine wall-piece, and so using that $\Omega(U_y U_c) \leq \frac{\pi}{2} < \pi$, we see by Condition 5.1.(4) that $U_y U_c$ is square homotopic into a path J in $N(H \cap Y)$ and so we can perform the replacement we performed above to decrease the complexity.

This reasoning holds for the first cone-cell in D (so there was no reason to first remove spurs). We thus reach the conclusion that D has no cone-cells and is thus a (possibly trivial) arc of 1-cells, and hence $P_h = P_y$ is a path in Y, and p, q lie in the same component of $N(H) \cap Y$. $\square$

5.6. **Defining immersed walls in X^*.** We now work under the hypothesis that Condition 5.1.(3) holds.

We define an equivalence relation on hyperplanes in $\widetilde{X}^*$, which is generated by $A \sim B$ provided that for some translate of some cone Y_i in $\widetilde{X}^*$, we have $A \cap Y_i$ and $B \cap Y_i$ lie in the same wall of Y_i.

The *walls* of $\widetilde{X}^*$ are defined to be collections of hyperplanes of $\widetilde{X}^*$ corresponding to equivalence classes. Our main goal will be to show that these walls do indeed "embed" and separate.

Example 5.17. The simply-connected complex in Figure 78 whose five 2-cells are all wallspaces, indicates that some small-cancellation hypothesis will be necessary to ensure that walls embed.

Definition 5.18. The *structure graph* Γ_W of a wall W is a bipartite graph whose 0-cells Γ^0_W are in two classes Γ^h and Γ^c where Γ^h_W consists of *hyperplane vertices* corresponding precisely to the hyperplanes in W, and Γ^c_W consists of *cone vertices* corresponding precisely to the cones that are intersected by some hyperplane of W. Two vertices of Γ_W are connected by an edge precisely when the corresponding spaces have a nonempty intersection $H \cap Y$.

A combination of various results proven below will show that:

Theorem 5.19. (1) *Γ_W is a tree.*
(2) *If $u \neq v \in \Gamma^h_W$ then the corresponding spaces are disjoint.*
(3) *If $u \neq v \in \Gamma^c_W$ then the corresponding spaces U, V are disjoint unless u, v are both adjacent to some $h \in \Gamma^h_W$ corresponding to a hyperplane H and $U \cap V$ contains a 1-cell dual to H.*
(4) *$u \in \Gamma^c_W$ and $h \in \Gamma^h_W$ then the corresponding spaces U, H are disjoint unless u, h are adjacent, and U contains a 1-cell dual to H.*

The final three statements in Theorem 5.19 follow from the following:

Theorem 5.20. *Let W be a wall.*

At most one midcube of a square (and hence of a cube) can lie in a hyperplane of W.

If H_1, H_2 are hyperplanes in W, and Y is a cone, then $H_1 \cap Y$ and $H_2 \cap Y$ lie in the same wall of Y.

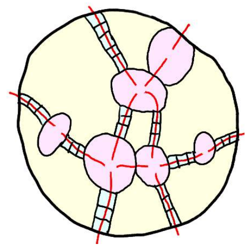

FIGURE 79.

Proof. Suppose that there is an alternating sequence $H_0, Y_1, H_1, Y_2, H_2 \dots, Y_r, H_r$ of hyperplanes and cones such that H_{i-1}, H_i belong to the same wall of Y_i for $1 \leq i \leq r$ but H_0, H_r are distinct hyperplanes that pass through the same (cube and hence) square or cone Z of $\widetilde{X}^*$. (We note that we permit back-tracking here, so it is possible that $H_i = H_{i+1}$ for many i. This facilitates the proof, which hinges upon minimal area instead of minimal length.)

Let P be a closed path that travels along the corresponding hyperplane carriers and cones, and starts and ends on this cube or cone Z. Let $D \to X^*$ be a diagram for P, and assume that D, P is minimal in the sense that D has minimal complexity among all such alternating sequences, and paths P, and diagrams for P.

Our argument has two stages: In the first stage we show that if D has minimal complexity, then D can be augmented to form a "collared diagram" E which is obtained by wrapping an annular ladder A around D such that A contains a self-intersecting path of W. In the second stage we show that removing and absorbing cancellable pairs preserves this collar structure, and so by passing to a minimal complexity collared diagram of this type, we see that it cannot exist.

Observe that P is a concatenation of (possibly trivial) paths $P^h_0 P^y_1 P^h_1 P^y_2 \dots P^y_r P^h_r P^z$ where each $P^y_i \to Y_i$ is path in a cone, and each $P^h_i \to N(H_i)$ is a path in the carrier of a hyperplane, and $P_z \to Z$ is a path in the square or cone Z that W crosses in two locally inequivalent ways. We note that when Z is a square dual to H_0, H_r, it actually lies in $N(H_0) \cap N(H_r)$. We can therefore assume that P^z is trivial in this case, by possibly absorbing it into P^h_0 (or into P^h_r).

We will show that P cannot contain a 1-cell e dual to some hyperplane in W in the sense that P^h_i does not pass through any 1-cell dual to H_i, and P^y_i does not pass through any 1-cell dual to the wall of W containing the hyperplanes $Y_i \cap H_{i-1}$ and $Y_i \cap H_i$. However, we do not impose any restrictions on P^z.

Indeed, if this were the case, then we would be able to produce a lower complexity counterexample. For the 1-cell e would be dual in D to a curve w of the wall W. Indeed, there is an immersed ladder L in D that is the concatenation of squares and cone-cells, such that L jumps across opposite 1-cells of a square, and L jumps across opposite 1-cells of a cone-cell in the sense that they are dual to the same wall of the ambient cone.

We note that since each cone-cell and each square contains an even number of 1-cells dual to W, there is really a dual graph, (as in Figure 79) whose vertices internal to D have even valence, and so there must be an even number of vertices on the boundary, and hence some dual curve starts and ends on $\partial_p D$.

There are several cases to consider, each leading to a lower complexity counterexample. We refer the reader to Figure 80.

Suppose w crosses itself within D by passing though the same square in two ways, or by passing through the same cone-cell in two 1-cells that don't map to 1-cells in the same wall of its ambient cone. Then a minimal such self-crossing dual curve is contained in a subdiagram of D, and hence provides a lower complexity counterexample: It bounds a path P' which bounds a subdiagram D' of D.

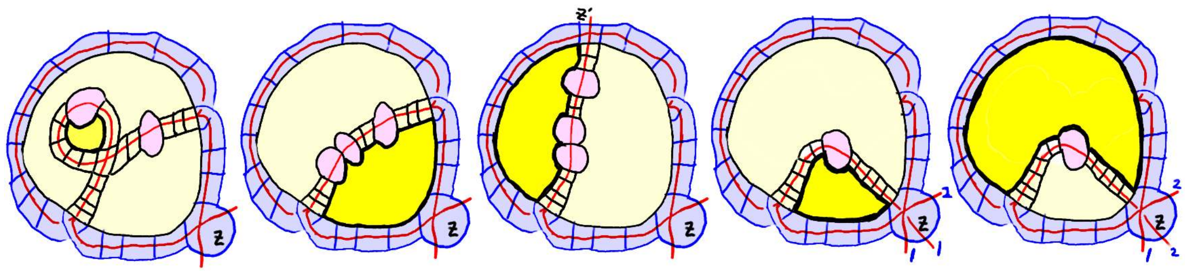

FIGURE 80.

Otherwise, we can choose w to simply start and end on boundary 1-cells of ∂D, and we obtain a lower complexity counterexample in one of the following ways:

If w ends on a 1-cell on P_i^y that is dual to the wall of Y_i containing $H_{i-1} \cap Y_i$ and $H_i \cap Y_i$ then $w \cup P$ bounds a lower complexity counterexample on the side of $D - w$ containing P_z. Similarly, if w ends on a 1-cell on P_i^h that is dual to the hyperplane H_i in $N(H_i)$, then $w \cup P$ bounds a lower complexity counterexample on the side of $D - w$ containing P_z. In the above two cases, we have the same self-intersection, and are merely taking a shortcut through the diagram.

If w ends on a 1-cell on P_i^y that is not dual to the wall of Y_i containing $H_{i-1} \cap Y_i$ and $H_i \cap Y_i$, then $w \cup P$ bounds a lower complexity counterexample on the side of $D - w$ not containing P_z. In this case the self-intersection is new and is at the cone Y_i.

Similarly, if w ends on a 1-cell on P_i^h that is not dual to the hyperplane H_i in $N(H_i)$, then $w \cup P$ bounds a lower complexity counterexample on the side of $D - w$ containing P_z. In this case the self-intersection is new and is at the square in $N(H_i)$ corresponding to the end of w and 1-cubes dual to H_i.

The final possibility is that w ends on a 1-cell f of P_z. (Note that because of our absorption convention above, this means that Z is a cone and not a square.) In this case, f is distinct from one of the two (or more) walls at Z corresponding to $H_0 \cap Z$ and $H_r \cap Z$. If it is distinct from $H_0 \cap Z$, then the initial part of P together with w bound a smaller complexity counterexample. if it is distinct from $H_r \cap Z$, then the terminal part of P together with w bounds a smaller complexity counterexample.

We have shown that by minimal complexity of D, the path P cannot pass through a dual 1-cell of W in the sense above. The next step of the proof is to "augment" D by adding an annular ladder A along P to obtain a new diagram $E = D \cup_P A$. We refer the reader to Figure 81.

Construction 5.21 (Collaring)**.**

The annulus A is a concatenation of square-ladders $L_i \to N(H_i)$ and cone-cells $C_i \to Y_i$ and a cone-cell or square $z \to Z$. Each L_i is the unique (possibly degenerate) ladder in $N(H_i)$ that contains P_i^h as a (possibly trivial) external arc. Each C_i is chosen so that $\partial_p C_i$ extends the path P_i^y so that it starts and ends on the terminal and initial 1-cells of L_{i-1} and L_i that are dual to $H_{i-1} \cap Y_i$ and $H_i \cap Y_i$ respectively.

We refer to E as a *collared diagram* in the sense that it has a single dual curve of a wall passing all the way around along its external 2-cells, except for one *corner* 2-cell where there is a transition, as the two hyperplanes do not belong to the same wall. More generally, we will later consider *collared diagrams with k-corners* in the sense that they have exactly k such corners. We note that when the dual curves braid with each other along the outside, there might be multiple ways of specifying the corners. The diagram E is a collared diagram with 1-corner. These ideas were treated in a simpler setup earlier in [OW11]. See Figure 82.

When Z is a square, we simply let z be a copy of Z, and when Z is a cone we let z be a cone-cell such that $\partial_p z$ contains the path $e_0 P^z e_r$ where e_0 is the initial 1-cell of L_0 and e_r is the terminal 1-cell of L_r (oriented appropriately).

The cone-cells satisfying the above prescriptions exist (after auxiliary extensions) by Lemma 5.27.

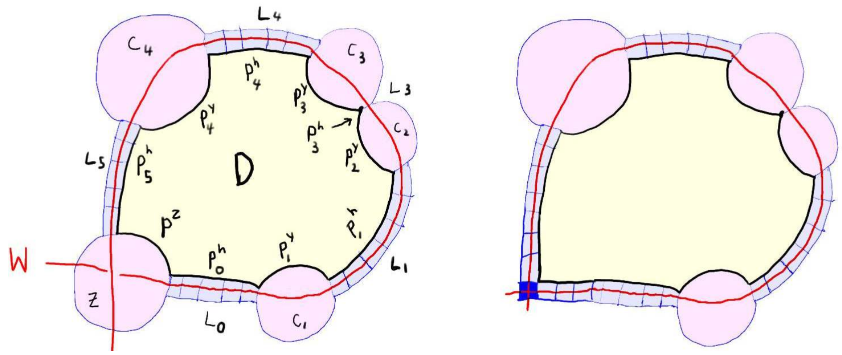

FIGURE 81.

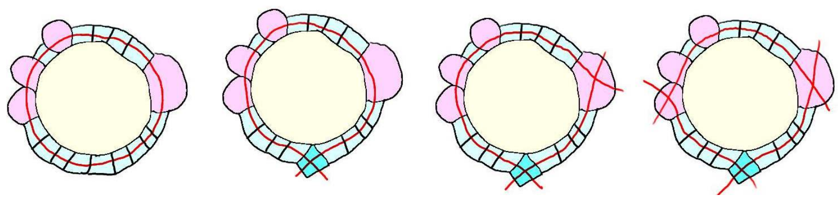

FIGURE 82.

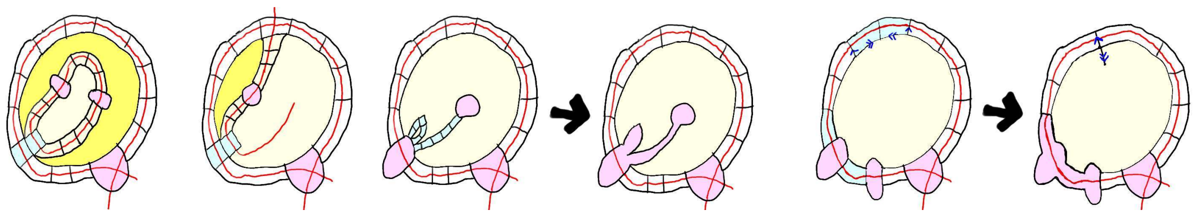

FIGURE 83.

We now verify that $P^z \to Z$ cannot pass through any 1-cell dual to W as above. Indeed, we follow the above argument, to produce a dual curve w leaving Z along some 1-cell of P^z. The details are similar.

Reductions Preserve Collar Structure: Having obtained the 1-corner collared diagram E, we reach the next stage of the argument which is to obtain a reduced 1-corner collared diagram. The idea is to show that if E is not reduced then, one can obtain a lower complexity 1-collared diagram E'. The essential thing to verify is that the collar structure is preserved. We refer the reader to Figure 83.

Reductions not involving an external 2-cell have no effect on the collared structure.

Absorption of an internal square into an external cone-cell.

Absorption of an external square into an adjacent external cone-cell.

Absorption of internal cone-cell with external cone-cell

Absorption of adjacent external cone-cells along a path on the boundary.

Absorption of internal squares from a fake wall-piece into an external cone-cell.

Absorption of external squares from a fake wall-piece into an internal cone-cell.

For the most part, absorptions into cone-cells have little effect, since they do not effect the diagram much. Even the combining of a cancellable pair of cone-cells has little effect: it just adds two small square subdiagrams, and then redraws the cone-cells boundaries.

Replacement of a cone-cell by a square diagram. In this case, if the cone-cell was external, then we might have to shave off one side of the square ladder containing the dual curve corresponding to the wall on the collar. See Figure 84.

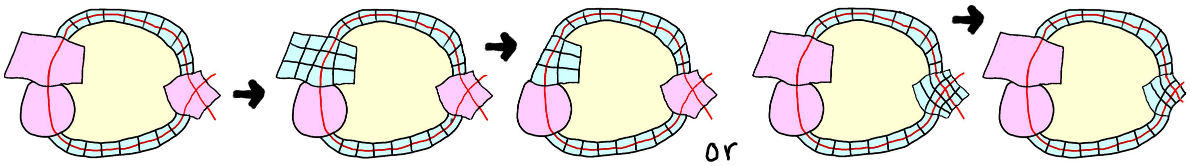

FIGURE 84.

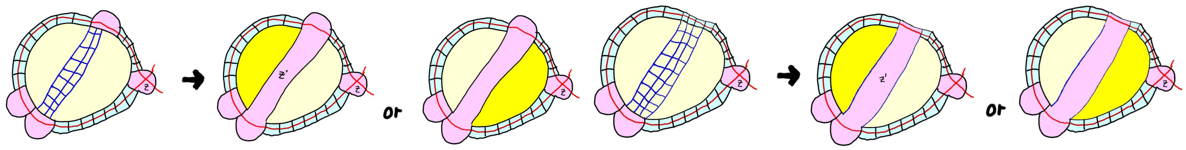

FIGURE 85. If there is a cancellable pair of cone-cells as on the first diagram, then there is a smaller complexity counterexample as shaded in the second or third. The second diagram is the case where the two walls in the cancellable pair are distinct, and the third where the two walls are the same. The fourth diagram indicates a fake wall-piece, which is absorbed analogously in the fifth and sixth diagrams.

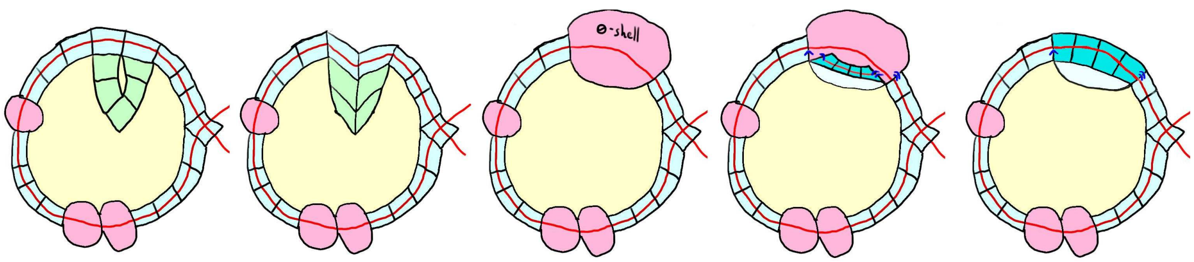

FIGURE 86.

An interesting case consists of the removal of a cancellable pair of squares where one is external and the other is not. In this case, one actually passes to substantially different diagram. There are several cases. If the ladder dual to the 1-cell of the internal square self-intersects (not illustrated) then there is an obvious lower complexity example. Otherwise, it either closes with itself, or ends on the boundary. Consideration of the various possibilities leads to a lower complexity diagram (after some gluing).

Absorption between external cells that shortcut the boundary is another interesting case that can lead to a substantially different diagram. We refer the reader to Figure 85.

Curvature along ∂E: We refer the reader to Figure 86. There are no spurs. Pushing past a corner of a generalized square provides lower complexity example.

A cone-cell C of positive curvature has innerpath S homotopic to the carrier of a hyperplane by Condition 5.1.(5). Therefore, letting $\partial_p C = QS$ we can replace C by a square diagram with boundary path $e_1 S e_2 Q'$ such that Q' lies along the boundary of the hyperplane carrier in the cone Y.

We thus conclude that only the corner 2-cells can support positive curvature. If the transition at a corner is a square, then there is a 0-cell with curvature $\frac{\pi}{2}$, and if the transition is a cone-cell, then there is a curvature of $\leq \frac{\pi}{2}$ unless the two collars are braided together there. If so, we can crop off some ladder, until we obtain a more genuine transition and a curvature of $\leq \frac{\pi}{2}$.

By Theorem 3.40, we see that there are at least three corners, unless the diagram is a ladder or a single cone-cell. Our 1-collared diagram cannot be any of these. □

We now prove the first statement listed in Theorem 5.19.

Theorem 5.22. *Let W be a wall, then Γ_W is a tree.*

Proof. This is similar to the proof of Theorem 5.20 but also uses its conclusion.

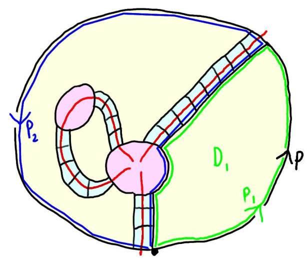

FIGURE 87.

Let P be the concatenation of an alternating sequence of paths $P_1^y P_1^h \dots P_r^h$, where each $P_i^y \to Y_i$ is a cone path and each $P_i^h \to N(H_i)$ is a path in a hyperplane carrier, and the hyperplane H_i intersects Y_{i-1} and Y_i in dual 1-cells, and each H_i belongs to W. By adjusting the starting point, we can assume that the first subpath is a cone path, and that the last subpath is a hyperplane path whose hyperplane H_r is dual to a 1-cell of the initial cone Y_1.

The path P determines a closed path $\bar{P} \to \Gamma_W$ whose vertices correspond to the subpaths and whose edges are transitions between these subpaths. Conversely, given a closed path $\bar{P} \to \Gamma_W$, we can choose a path P as above that projects to $\bar{P}$ in this sense.

We aim to show that any such path $\bar{P}$ is null-homotopic in Γ_W. We argue by contradiction to see that no such path exists. Consider a minimal complexity diagram D whose boundary path P induces an essential path $\bar{P} \to \Gamma_W$.

Observe that P cannot pass through a 1-cell e that is dual to W in the sense that e either lies in P_i^y and is dual to a hyperplane in a wall of $H_i \cap Y_i$ or $H_{i+1} \cap Y_i$, or e lies in P_i^h and e is dual to $H_i \subset N(H_i)$. Indeed if such a 1-cell existed, then as in the proof of Theorem 5.20, we consider its dual graph in D, and find that we can choose a dual curve w of a path in W that is carried by a ladder in D, such that w does not cross itself. Note that while it is possible for the dual graph to have high even valence where it might bifurcate at cone-cells, Theorem 5.20 shows that it cannot cross itself, and in particular, if it some path in the dual graph comes back to the same cone-cell then it returns along a hyperplane in the same wall of that cone.

We are therefore able to choose a simple curve in this dual graph that starts on e and ends on a 1-cell e' on P. Moreover, this curve is carried by a ladder within D.

If e' has the same property as e, then (after cyclically permuting) this decomposes the path P as homotopic to the concatenation of two paths P_1P_2, and consequently $\bar{P}$ is homotopic to the concatenation $\bar{P}_1\bar{P}_2$, and at least one of $\bar{P}_1, \bar{P}_2$ must be essential in Γ_W since $\bar{P}$ is essential. However, the diagrams D_1, D_2 for P_1, P_2 are lower complexity than D, so this is impossible. We refer the reader to Figure 87.

If e' does not share the same property as e, so it corresponds to a different wall in Y_i or corresponds to a different hyperplane crossing H_i in $N(H_i)$, then we obtain a self-crossing wall, which violates Theorem 5.20.

Since P does not pass through a 1-cell dual to W as above, we see that the diagram D can be augmented as in Construction 5.21 to form a collared diagram E no zero corners as in the first diagram in Figure 82.

As we did earlier, we can then choose a collared diagram E with no corners and of minimal complexity, so that it is reduced and has no θ-shells or corners of generalized squares along its boundary. However, such a collared diagram cannot exist by Theorem 3.40. $\square$

We define the *carrier* of the wall W as follows: First let B denote the disjoint union of the carriers $N(H_i)$ of its constituent hyperplanes together with the cones intersected by these. We then form a quotient of B by identifying pairs of subspaces $N(H)$ and Y along their intersection $N(H) \cap Y$ in $\widetilde{X}^*$,

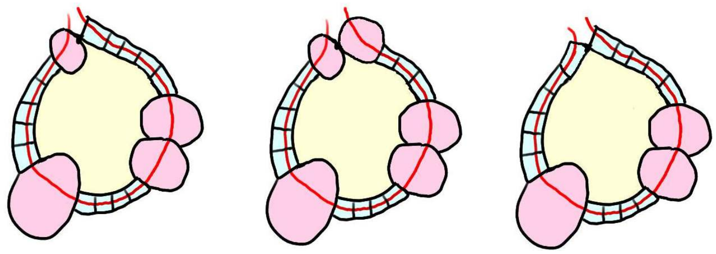

FIGURE 88.

and identifying Y_i and Y_j along their intersection in $\widetilde{X}^*$ provided that Y_i and Y_j share a 1-cell dual to some hyperplane H in W. Note that these intersections were proven to be connected in Lemma 5.6 and Lemma 5.15. We let $N(W)$ denote the resulting quotient. We note that there is an induced map $N(W) \to \widetilde{X}^*$.

Theorem 5.23 (Carriers Embed)**.** *For each wall W, the map $N(W) \to \widetilde{X}^*$ is an embedding.*

Proof. This is again a variant of the proof of Theorem 5.20.

Suppose there is a nonclosed path $P \to N(W)$ that maps to a closed path in $\widetilde{X}^*$. We choose such a path whose diagram D has minimal complexity in X^*.

We verify as above that P cannot pass through a dual 1-cell of one of the hyperplanes of W, for then there would either be a self-intersection which is precluded by Theorem 5.20, or we could cleave off part of D to obtain a new path P' with the same endpoints in $N(H)$ (by Theorem 5.22) such that P' had a smaller complexity diagram D'.

We then augment D by Construction 5.21 to obtain a collared diagram with one collar and two corners.

By passing to a minimal complexity such diagram, we can assume it is reduced and has no other features of positive curvature. The key point to verify here is that the collar structure is preserved when cancellable pairs are removed or absorbed. When there is a fake wall-piece or cone-piece from a square ladder or cone-cell in the collar on another side of the diagram (as in Figure 85) the absorbed wall from within the collar is the same as the wall in the absorbing cone, for otherwise we would contradict Theorem 5.20. Therefore we are merely able to cleave off a closed dual curve in the wall, and this merely provides a new path P' with the same endpoints as P, so the structure (and the original distinct pair of points in the carrier) are preserved. Other cases are treated similarly.

Note that the first diagram in Figure 83 provides a new path $P' \to N(H)$ with the same endpoints, but with lower complexity diagram. However, the second diagram in Figure 83 cannot occur, since it would yield a self-crossing which was ruled out in Theorem 5.20. Some variants of this are indicated in Figure 89. The first diagram explains how a cancellable pair of squares yields a new path P' with a lower complexity diagram, the second and third yield self-crossing contradictions, and the fourth diagram indicates that, no inner edge of the collar is dual to our wall, or there would have been a new path P' with a lower complexity diagram.

We then conclude that it must be a ladder by Theorem 3.36, and so the two corners must be adjacent cone-cells corresponding to cones that share a 1-cell dual to W. Indeed, they are already adjacent in the diagram. We emphasize that, as concluded in the Theorem 3.36, since a square corner only provides a curvature of $\frac{\pi}{2}$, both of these positively curved cells must be cone-cells, and so the first and third figures in Figure 88 are excluded.

□

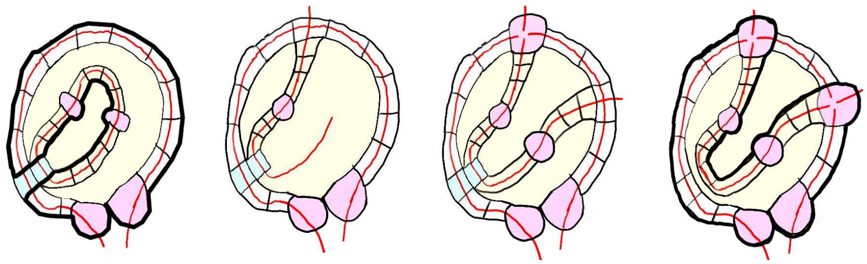

FIGURE 89.

Remark 5.24 (2-sidedness)**.** Since hyperplanes are 2-sided in their carriers by Lemma 5.13, and carriers of walls in cones are 2-sided by hypothesis, and each carrier is a tree union of cones and hyperplane carriers by Theorem 5.22, we see that each wall W is 2-sided in and hence separates $N(W)$.

Since $N(W)$ embeds in $\widetilde{X}^*$ by Theorem 5.23, we see that W is locally 2-sided in $\widetilde{X}^*$ and actually locally separates $\widetilde{X}^*$ since $N(W) \to \widetilde{X}^*$ is a local isomorphism along W by construction.

Now the 2-sidedness of the walls in $\widetilde{X}^*$ follows from their embeddedness, local 2-sidedness in $N(W)$ which embeds, and the simple-connectivity and hence 1-acyclicity of $\widetilde{X}^*$.

5.7. **No Inversions.** Even when each wall W of $\widetilde{X}^*$ is 2-sided, it might be that $\mathrm{Stab}(W)$ does not also stabilize each of the half spaces. We rectify this as follows:

Lemma 5.25 (No inversions)**.** *Suppose that $\langle X \mid Y_i \rangle$ satisfies the $B(6)$ conditions as well as the following condition. Then for each wall W, its stabilizer* $\mathrm{Stab}(W)$ *acts without inversions in the sense that it stabilizes each component of $N(W) - W$.*

(1) *[No inversions] The hyperplanes in X are 2-sided, and there is a choice of positive and negative side of each hyperplane, and a choice of positive and negative side of each wall in each Y, and these two notions are globally consistent, in the sense that the positive side of each hyperplane H equals the positive side of each wall represented by $H \cap Y$.*
   Moreover $\mathrm{Aut}_X(Y)$ *acts on the wallspace without inverting the sides of any wall.*

Proof. □

5.8. **Phony Cone Cells.**

Definition 5.26. $\langle X \mid Y_i \rangle$ has the *cycle property* if for each immersed path $P \to Y_i$ which starts and ends on 1-cubes dual to distinct hyperplanes in the same wall of Y_i, there exists a path Q with the same endpoints as P such that QP is essential.

Lemma 5.27. *Let $\langle X \mid Y_i \rangle$ be a cubical presentation. There exists a deformation retraction embedding $\langle X \mid Y_i \rangle \subset \langle X' \mid Y_i' \rangle$ such that $\langle X' \mid Y_i' \rangle$ has the cycle property, and such that all properties listed in Definition 5.1 are preserved.*

Proof. For each geodesic path $P \to Y_i$ with 1-cells not dual to the same hyperplane but dual to the same wall, and not connectable by a different path, we add an arc P'' of the same length as P' where $P = e_1 P' e_2$, to both Y_i and to X. The new hyperplanes are simply barycenters of new 1-cubes. We assign them to walls of Y_i according to the antipodal relation in the cycle $P'' e_1 P' e_2$. By concatenating cycles, and reducing or folding, one sees that being either dual to the same hyperplane or having a common local geodesic cycle is an equivalence relation. Hence no new classes are added when we add the arcs. □

Remark 5.28. We will implicitly assume that the construction of Lemma 5.27 has been implemented to support the collaring procedure that is especially used in the proof of Theorem 5.32. The pictures supporting the proof rely on the intuitive idea that pairs of 1-cells dual to distinct hyperplanes in the same wall of a cone actually lie opposite each other on some essential cone-cell.

5.9. **Carriers and Quasiconvexity.**

Definition 5.29 (Carrier of wall)**.** A wall W *crosses* a cone C if $W \cap C$ is nonempty. The *carrier* of a wall W of $\widetilde{X}^*$ is the union of all hyperplane carriers containing its hyperplanes, together with all cones that it crosses.

The most important cones in a carrier are those cones C crossed *essentially* by W in the sense that $W \cap C$ consists of two or more hyperplanes of C. Many of the properties of walls and carriers are explainable for the *essential carrier* which only includes essentially crossed cones, but we have decided to retain the point of view including all crossed cones instead of the essential ones. The essential carrier is a bit less "bumpy" then the carrier. In the "classical case" of $B(6)$ complexes, all crossed cones are essential.

Note that the essential cones correspond to non-leaves in the tree Γ_W. The *deeply essential cones* correspond to *deep vertices* in the tree which have the property that some bi-infinite geodesic passes through them.

It is a consequence of Theorem 5.20 that the carrier deformation retracts to the wall (if we add geometric walls in each cone in an appropriate fashion that we shall later explore) and that the carrier is simply-connected and is separated by the wall.

Definition 5.30 (Cubical ladders and consecutive cones)**.** Two cones in the carrier of W are *consecutive* if there is a hyperplane in W that has nonempty intersection with each of them. A *cubical ladder* in a cube complex X is a local isometry $I^n \times [0, m] \to X$ where $n \geq 1$. If $I^n \times \{0\}$ lies in the cone Y_i and $I^n \times \{1\}$ lies in the cone Y_j then we say it is a cubical ladder joining Y_i and Y_j.

Definition 5.31 (Thickened and Extended Carrier)**.** The *thickened carrier* $T(W)$ of W is the union of the carrier $N(W)$ together with each flat rectangle $R = I_n \times I_m$ whose left and right boundary paths $\{0, n\} \times I_m$ lie on cones Y_1, Y_2 of $N(W)$, and such that for some $0 \leq k < m$ the square ladder $I_n \times [k, k+1]$ is dual to a hyperplane in W.

The *extended carrier* $E(W)$ is the cubical local convex thickening of the thickened carrier $T(W)$. It is obtained by repeatedly adding any cube with an entire corner already present.

We refer to Figure 90 for heuristic illustrations of $N(W)$, $T(W)$, and $E(W)$ when X is 2-dimensional.

We will use the term *cladder* for the carrier of a wall within a disk diagram. We will later introduce the term W-ladder for the case when the corresponding dual curve is associated to a specific wall W in $\widetilde{X}^*$ (see Definition 5.57).

Theorem 5.32 (Thickened Carrier Isometric Embedding)**.** *Assume X^* satisfies the $B(6)$ condition. The thickened carrier of a wall isometrically embeds in $\widetilde{X}^*$.*

Proof. We can assume without loss of generality that γ is disjoint from T and in particular from N except at its endpoints. We refer the reader to Figure 92 for a guide to the notation used in the proof.

Step 1: Setting up the diagram: We first show that there exists a disk diagram E which contains (generalized) ladders $L_N \subset L_T$ at the top, and our geodesic γ at the bottom. This is something we have done routinely when γ starts and ends on N.

We begin by extending γ to $\gamma^+ = \gamma_0 \gamma \gamma_t$ where each γ_i is a minimal length path in T from an endpoint of γ whose outermost 1-cell is dual to W. We emphasize that this path can be chosen so that it travels through a sequence of 1-cells that are in T and are dual to hyperplanes joining (the same) pair of cones in $\widetilde{X}^*$. When the endpoint of γ lies on a cone of $N(W)$ then we choose γ_i to be a path in this cone to a dual 1-cell of W.

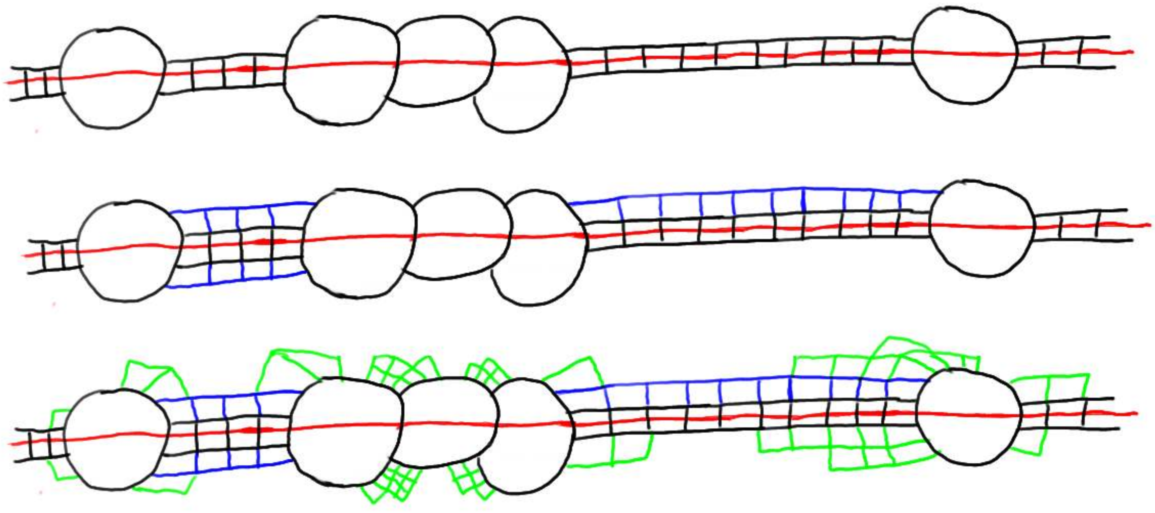

Figure 90. A carrier, a thickened carrier, and an extended carrier

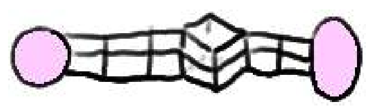

Figure 91. It seems harmless to include squares in T whose entire 1-skeleton lies in T.

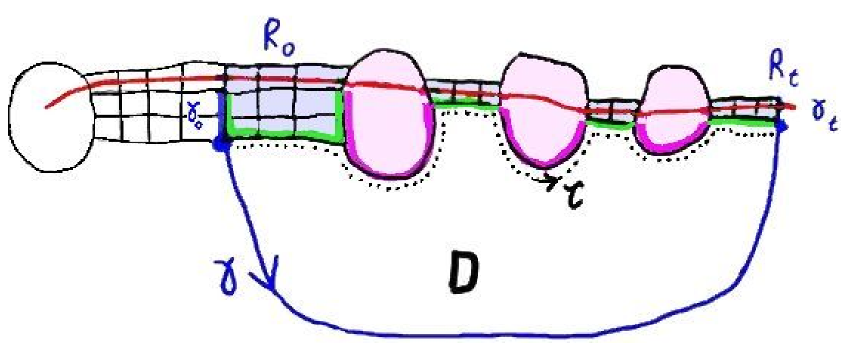

Figure 92.

For each of $i = 0$ and $i = t$, there is a rectangle R_i containing γ_i on one side and bounded by a square ladder mapping to a hyperplane of W on one side and ending on a path in a cone-cell of $N(W)$ on a third side. When γ_i lies on a cone (because the corresponding endpoint of γ lies on that cone), then R_i is degenerate, and is simply a copy of γ_i. Likewise, when γ_i is a single 1-cell (since γ ended on the carrier of a hyperplane of N) then we just let R_i be a copy of γ_i.

There is then a sequence of cones and square ladders in $N(W)$ that start and end at the initial and terminal 1-cells of γ^+. We choose the first and last such ladders to be those already lying in R_0 and R_t. Let V_i be the external boundary arc of our ladders (and initial and terminal rectangles). For each cone Y_i, we choose a path U_i that travels around Y_i from one dual 1-cell to another. (We choose the paths V_i, U_i to start and end on the same side of W as γ.) We consider the concatenation $V_0U_0V_1U_1 \cdots V_t$. We then consider a disk diagram D between γ and the above path. Adding on the above sequence of ladders, as well as cone-cells supporting the U_i paths, we can augment D to obtain the collared diagram E.

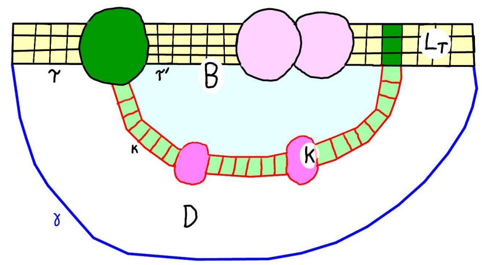

FIGURE 93.

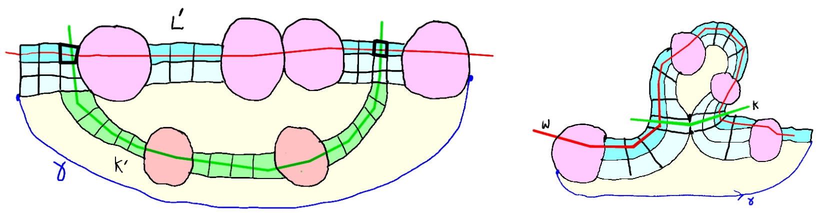

FIGURE 94.

Step 2: Estimating the distance: The ladder L_T in E contains the sequence of rectangles, square ladders, and cone-cells, and along the top of L_T is the ladder L_N mapping to $N(W)$. Let τ be the path on D that "opposes" γ, so that τ is also a path along the "bottom" of L_T.

Having shown above that one such diagram exists, we now assume that we have chosen D, E with the property that its complexity, Comp, is minimized where $\mathsf{Comp}(D, E) = (\#_c(D), \#_c(E), |\tau|)$ where $\#_c(X)$ denotes the number of cone-cells in X. We emphasize that $\#_c(D)$ equals the number of cone-cells in E that are *not* also in L_N (or equivalently in L_T). Note that we are not assuming that E is reduced or has minimal complexity in the usual sense. However, the minimality of Comp guarantees that no cone-cell in E can be replaced by a square diagram, and adjacent cone-cells cannot be combined. We are also not assuming that τ is the shortest possible path in some E, D, but rather, that it is shortest possible among all such paths arising from E, D with minimal numbers of cone-cells.

We claim that for each edge e in τ, the dual curve to e in D is a graph that ends on e in τ, and whose remaining endpoints all lie on γ. Since there are an odd number of such remaining endpoints, we see that $|\gamma| \geq |\tau|$, and since τ lies in T, our claim is proven.

Step 3: Choosing a minimal counterexample: Suppose that the dual graph ends at another edge e' in τ for some E, D with Comp minimal as above. Let K be the cladder in D carrying a dual curve that starts and ends on the edges e, e' of τ. Let τ' be the subpath of τ whose initial and terminal edges are e, e'. Let κ be the path on K with the same endpoints as τ' but doesn't pass through the dual curve of K, and in particular doesn't pass through e, e'. We extend the cladder K in D carrying a dual curve that starts and ends on τ, to a cladder K' that starts and ends on squares and/or cones of $L_N \subset L_T$. Let L' denote the subladder of L_N that begins and ends at the intersections with K'. See the diagram on the left of Figure 94. Let B be the subdiagram that is bicollared by K', L'. See Figure 93.

Among all minimal Comp examples for a given γ (or even a geodesic with the same endpoints as γ) exhibiting the contradictory feature represented by K, let us choose a contradiction with the additional property that B has ordinary minimal usual area complexity $(\#_c(B), \#_s(B)) = (\text{Conecells}(A), \text{Squares}(A))$.

We will now show by contradiction that such a minimal counterexample cannot exist.

Note that in the degenerate case where $e = e'$ illustrated on the right of Figure 94, there is automatically a smaller choice of B corresponding to an innermost backtracking cladder.

Step 4: Minimality of B implies square diagram:

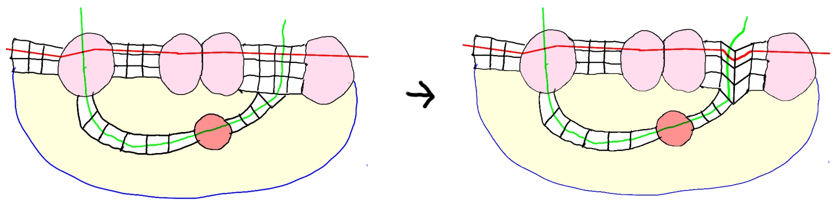

FIGURE 95.

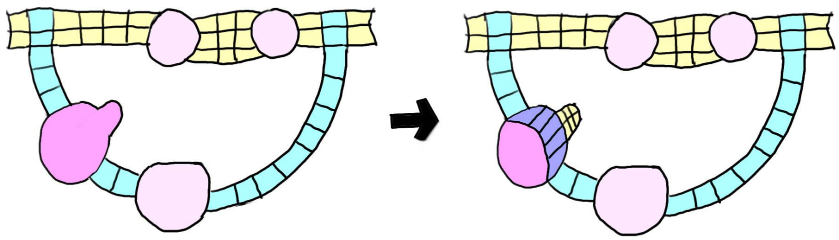

FIGURE 96.

We now focus on the subdiagram B that is bicollared by K', L' - we emphasize that $K', L' \subset B$. We shall show that any cone-cell in B actually lies in L'. Note that we cannot immediately apply Theorem 3.36 as we don't know that B is reduced along L' and has no θ-shells along L'.

The minimality of B implies that there are no corners of generalized squares on cone-cells within B, for they can be absorbed. There is also no corner of generalized square within B whose outerpath lies on ∂B except perhaps at the two corners of B where L', K' meet. Indeed, if a square in B has a generalized corner on a square ladder in K', then we can push K' upwards across this square while decreasing the area in B. Similarly, if there is a square in B with a generalized corner with outerpath along the top of L', then we can push the square upwards through L' (and hence L_T) until it is outside of E, and reduce the area of B. Neither of these moves affects $\mathsf{Comp}, |\tau|$. See Figure 95 for a trickier-than-usual pushing of a generalized corner of square.

Suppose that B contains a cone-cell C that does not lie in L'. Let us examine what happens to it as we perform a sequence of reductions to B. Note that we can assume from the outset that L' is reduced. It is impossible for any cone-cell C to be combined with another cone-cell by minimality of Comp and likewise impossible for C to eventually work its way into L'. Suppose after some sequence of reductions $B \mapsto \bar{B}$, a cone-cell C becomes a θ-shell $\bar{C}$ with outerpath Q along $\bar{\kappa} = \kappa$, and innerpath S. Let Y be the supporting cone, so $\Omega_Y(S) < \pi$. And note that $S = UfS'f'U'$ has a subpath $fS'f'$ where f, f' are 1-cells on $\partial\bar{C}$ that are dual to the dual curve of K, and also lie on the previous and next cells of K respectively. It is important to note that f, f' must be oriented in the same way with respect to this hyperplane, otherwise the path S' would contain another edge dual to this hyperplane, and following the associated cladder in D, we would be able to construct a backtracking cladder K_0 bounding a subdiagram B_0 of B. By Condition (5), the path $fS'f'$ of $\bar{C}$ is homotopic in Y to a path in the carrier of the hyperplane of Y dual to f, f', and hence S' is homotopic to a local geodesic J in this carrier. Let V denote the square diagram given the homotopy from S' to J. Let J' denote the corresponding local geodesic on the other side, so J, J' together form the side of a ladder M in the carrier with initial and final 1-cells f, f'. We refer the reader to Figure 96.

The path $U'QUJ'$ bounds a cone-cell C', and we replace $\bar{C}$ by the union $C' \cup_{J'} M \cup_J V$. This has the effect of substituting the square ladder M for $\bar{C}$ in $\bar{K}$, and hence decreases the number of cone-cells in the resulting bicollared diagram replacing $\bar{B}$. This violates our minimality assumption on B.

Consequently, since after reducing, there are no θ-shells along $\bar{\kappa}$, and likewise, none along the part of $\bar{B}$ collared by $\bar{L}'$ (since all those cone-cells are essential), we see that there are at most two positively

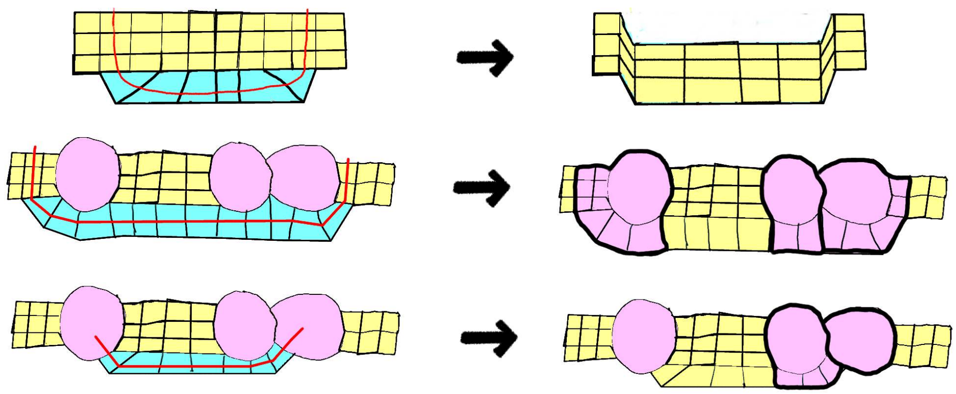

FIGURE 97.

curved cells at the corners of $\bar{B}$, and so $\bar{B}$ is a ladder by Theorem 3.36. In particular, $\bar{B}$ has no internal cone-cells either, so all cone-cells of B have been accounted for as arising in L'. We conclude that there are only squares in B outside of L', and in particular K is a square ladder.

Step 5: Minimality implies that B is very thin: Let A be the part of B between κ and τ', so $K \subset A$. We will show that $A = K$.

The minimality of B implies that it has no bigon of dual curves inside. Indeed, by Lemma 2.2, removing such a bigon decreases the area by two. In particular, a bigon with the dual curve of K itself can likewise be pushed through K, so that the important properties of the resulting diagram are preserved.

Each dual curve in A that starts on $\tau' - \{e, e'\}$ ends on a 1-cell on κ. Indeed, if such a dual curve ended on another edge of τ', then its corresponding cladder could be used instead of K, and would provide a lower complexity counterexample violating the minimality of B. Moreover, each dual curve in A starting on κ must end on τ' for otherwise there would be a bigon, and this was excluded earlier. In summary, each dual curve in A travels between τ' and κ, except for the dual curve within K itself. Hence $|\kappa| = |\tau'| - 2$.

Suppose that A contains a square s that doesn't lie in K. Then each dual curve of s has one end on κ and the other end on τ'. Thus some corner of s has a pair of dual curves on κ. This contradicts Lemma 2.5.

Step 6: Reducing $|\tau'|$ and hence $|\tau|$: If L' contains only squares then τ' can be replaced by the shorter path κ and K' can be pushed through L' (using several bigon removals). This reduces the length of τ' and hence violates the minimality of τ.

If L' contains cone-cells, then we can assume it is already reduced without affecting any of our minimized quantities. Let R' denote the part of L_T that is subtended by K, so R' is the union of L' and various additional rectangles in L_T. We then observe that the ladder K' wraps around L'.

We refer to Figure 97 on the left, as well as the resulting absorptions on the right.

Suppose the first square of K has one edge on a square of L_T and one edge on a cone-cell C of L_T. If the rectangle of K alongside C cannot be absorbed into C, then the innerpath S of C is the concatenation of two pieces: a rectangle-piece associated with an initial part of K' followed by a cone-piece. (If C is the unique cone-cell, then S is a single rectangle-piece). Letting Y denote the cone supporting C we have $\Omega_Y(S) < \pi$. This leads to a contradiction, because Condition (5) would then imply that S is homotopic into the carrier of a hyperplane associated to the wall of $L_T \cap Y$, which would imply that the cone-cell C could be replaced by a square diagram, violating the minimality of $\#_c(E)$.

Thus the contiguous rectangle is absorbable into C, and this includes the entire initial rectangle of K'.

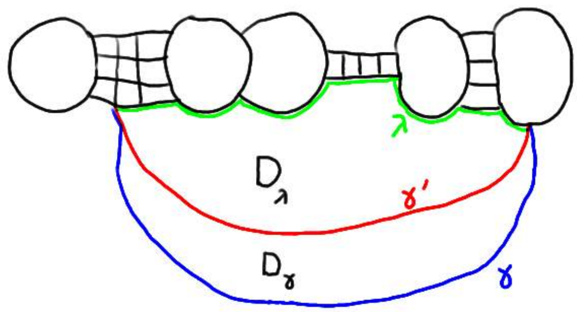

FIGURE 98.

We can thus assume that K starts and ends on cone-cells. But then the ladder theorem implies that B becomes a ladder after reducing, and hence each rectangle in K contiguous with a cone-cell of L_T is absorbed into that cone-cell upon reducing. The path τ' can thus be replaced by κ, yielding a shorter counterexample. Note that in the degenerate case where L_T consists of a single cone-cell C it is immediate that all of K absorbs into C since there are corners of generalized squares in K on C (one at each side). □

Theorem 5.33 (Extended Carrier Convexity)**.** *For any geodesic $\gamma \to \widetilde{X}^*$, if the endpoints of γ lie in $T(W)$ then $\gamma \subset T(W)$.*

I expect one can could prove that $E(W)$ is convex, but there was a gap in the original version of this proof, and only the weaker statement was covered.

Proof. We now use Condition 5.4 to prove the convexity of the extended carrier of W.

Let γ be a geodesic in $\widetilde{X}^*$ that starts and ends on the thickened carrier T of W. As provided in the proof of Theorem 5.32, let λ be a path on T with the same endpoints such that the disk diagram $D \to \widetilde{X}^*$ between them has minimal Comp among all such possibilities for D, λ. Note that λ can be chosen to be on a minimal length ladder in T containing a ladder in N.

Moreover, let γ' denote a geodesic in D with the same endpoints as γ such that γ' and γ together bound a square subdiagram of D. Thus D is the union of two diagrams D_λ and D_γ that meet along γ' (see Figure 98). We will assume that D_λ has minimal area among all possible choices with γ fixed.

Since γ' is a geodesic, Condition 5.4 implies that D_λ has no shell with outerpath on γ'. There is also no generalized square in D_λ with outerpath on γ', for then we could pass a square across γ' to increase $\mathsf{Area}(D_\gamma)$ and decrease $\mathsf{Area}(D_\lambda)$.

We now consider the diagram E bounded by γ' on one side, and bounded by and including the ladder L_T in $T(W)$ consisting of the sequence of cone-cells and flat rectangles between them. We will show that γ' lies on L_T by showing that D_λ is an arc. Since γ and γ' lie in the local (square) convex hull of each other, we see that γ lies in L_T as well, thus proving the theorem.

Without loss of generality, we can prove the statement for a subpath γ'_0 of γ' obtained by ignoring initial and terminal subpaths of γ' that already lies on L_T. We let E_0 denote the resulting diagram, and observe that cropping γ in this way, would necessitate a corresponding cropping of initial and terminal rectangles of L_T.

As above, minimality of complexity implies that there are no outerpaths of generalized squares at the "top" boundary path along L_T. Condition 5.1.(5) implies there is no θ-shell C within the cladder with outerpath at the top of L_T, since then the innerpath S and outerpath Q of C satisfy $\Omega_C(S) < \pi$ and and S is homotopic in the ambient cone to the carrier of a hyperplane and thus a square ladder U of length at most $|S| - 2$. So the C with $\partial_p C = QS$ can be replaced by a square diagram with boundary

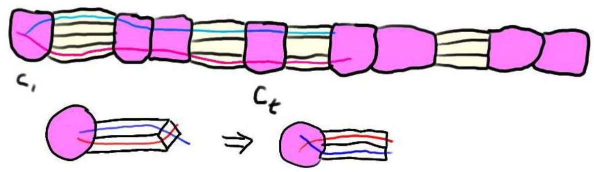

FIGURE 99.

path SQ' where Q' is along one side of the ladder U. This shortens the ladder L_N in N at the top of L_T since U replaces a cone-cell, and hence violates our earlier minimality assumption.

The lack of positively curved shells claimed above holds with the sole exception of a cone-cell and or square at either end of $L_N \subset L_T$. Note that cropping removed the possibility of a second corner of a generalized square. Moreover, the second dual curve of a square at the end of L_N must terminate on γ' (and not on a cone in L_T). Indeed, as explained in the proof of Theorem 5.32, all cladders emerging from λ terminate on γ (after passing through γ').

As there are only two positively curved cells, Theorem 3.36 implies that E_0 is a ladder. □

Corollary 5.34. *If X is 1-dimensional then carriers are convex.*

Proof. In this case the carrier is the same as the thickened carrier which is the same as the extended carrier, so the result follows from Theorem 5.33. □

Lemma 5.35 (Walls quasiisometrically embed)**.** *If pieces in cones (and hence cone-cells) are uniformly bounded in size, then each wall quasiisometrically embeds in $\widetilde{X}^*$.*

Proof. The hypothesis implies that the thickened carriers are in a finite neighborhood of the walls. A bit more thought shows the walls are quasi-isometric to their thickened carriers, which are isometrically embedded. □

5.10. ⊛ **Bigons.**

Definition 5.36. Let Y, H be a cone and a hyperplane intersecting it. The *wallray* based at Y in the direction of H consists of the part of W corresponding to the subtree of Γ_W that starts at the leaf consisting of the vertex y of Y, and contains the edge (y, h). The wallray is *carried* by the corresponding subspace of $N(W)$.

Theorem 5.37 (Intersection of Wallrays)**.** *Let Y, H_1, H_2 be a cone and a pair of hyperplanes in $\widetilde{X}^*$ with dual 1-cells in Y. Let $h_1 = Y \cap H_1$, and let $h_2 = Y \cap H_2$. Suppose that $\Omega_Y(h_1, h_2) > 0$ so that there are no 1-cells e_1, e_2 in h_1, h_2 that lie in the same cone-piece or wall-piece in Y. And suppose that h_1, h_2 do not cross in Y, so there is no 2-cube in Y with 1-cells dual to both h_1 and h_2.*

Then the wallrays (W_1, Y, H_1) and (W_2, Y, H_2) do not have crossing hyperplanes besides $H_1 \cap H_2$, and they do not both intersect any cone besides Y.

Remark 5.38. The statement needs to be reconciled with the proof, since we now allow some initial crossing in $H_1 \cap H_2$. Crossing at the last square implies crossing inside the last cone. See Figure 99.

The following Corollary is easier to prove than Theorem 5.37.

Corollary 5.39. *Let W_1, W_2 be walls intersecting a cone Y. Suppose that $\Omega_Y(W_1 \cap Y, W_2 \cap Y) > 0$. Then W_1 and W_2 do not have hyperplanes that intersect in any cone other than Y, or that intersect in any square except for those dual to codimension-2 hyperplanes emanating from Y.*

Proof of Theorem 5.37. Let Y^+ denote the union of Y and the carriers of codimension-2 hyperplanes consisting of components of $H_1 \cap H_2$ that intersect Y in a component of $h_1 \cap h_2$.

We say $p \leq_i Y^+$ if p lies on a cone or hyperplane carrier of $N(W_i)$ but does not lie on a cone or carrier of the wallray (W_i, Y, H_i), or $p \in Y^+$. If $p \not\leq_i Y^+$ then we write $p >_i Y^+$.

Consider paths P_1, P_2 in $N(W_1), N(W_2)$ with initial points p_1, p_2 and terminal points q_1, q_2 such that

(1) $p_i \leq_i Y^+$ for $i = 1$ or $i = 2$ or both,
(2) each $q_i >_i Y^+$.
(3) p_i, q_i lie on 1-cubes dual to W_i for each i.
(4) p_1, p_2 lie on a cone or square A containing the above 1-cubes, and q_1, q_2 lie on a cone or square B containing the above 1-cubes.
(5) there are paths $P_a \to A$ and $P_b \to B$ connecting p_1, p_2 and q_1, q_2.
(6) P_a is trivial if A is a square, and P_b is trivial if B is a square.

Note that by the conditions above, B cannot equal Y, and cannot be a square on Y^+ dual to a codimension-2 hyperplane $H_1 \cap H_2$ intersecting Y in some $h_1 \cap h_2$.

If there were an intersection contradicting the theorem, then there would be paths P_i in the carriers of (W_i, Y, H_i) that start on vertices along 1-cells in $A = Y$ dual to h_i, and a connecting path P_a joining their initial points, so that these paths end on vertices along 1-cells dual to W_1, W_2 on a cone $B \neq Y$ or on a square B not dual to a codimension-2 hyperplane emanating from Y, and we could let P_b connect their endpoints when B is not a square, and let their endpoints be identical when B is a square.

We aim to produce a contradiction ruling out a minimal counterexample by applying Theorem 3.36 to a suitably produced bicollared diagram capturing the above data in a minimal case. Let us first assume that P_1, P_2, P_a, P_b are chosen with the above properties so that their concatenation is a quadrilateral that bounds a minimal complexity disk diagram $D \to \widetilde{X}^*$ among all possible such counterexamples to the statement of the theorem.

Our first aim is to show that the minimal complexity of D implies that $P_a P_1 P_b^{-1}$ does not pass through any dual 1-cell of W_1, and likewise, $P_a^{-1} P_2 P_b$ does not pass through any dual 1-cell of W_2.

As the argument is symmetric, let us consider the first scenario. Indeed, suppose P_1 did pass through such a 1-cell m dual to W_1, then we could produce a lower complexity example, following Figure 103. Consider an embedded ladder L in D carrying a dual curve to this 1-cell m and terminating at another dual 1-cell m' on $\partial_p D$. If m' also lies on $P_a P_1 P_b^{-1}$, then we can produce a smaller complexity counterexample, by traveling along $P_a P_1 P_b^{-1}$ until we hit m, then traveling along the bottom of L (that is the side facing P_2), and then continuing along the subpath of $P_y P_1$ after m'.

If m lies on P_y then P_y is replaced by the initial subpath preceding m, and the initial point of P_1 is adjusted. Likewise if m' lies on P_z then P_z is replaced by the terminal subpath following m', and the terminal point of P_1 is adjusted.

Also note that in the degenerate case where m is an isolated 1-cell of D along a backtrack of $P_y P_1 P_z$, then $m = m'$ and the procedure described above simply crops off the part of this arc in D from m outwards, corresponding to the subpath of $\partial_p D$ between m and m' inclusive. (So among our minimal counterexamples, we first choose one with a shortest boundary path.)

A more interesting case arises when the ladder L terminates on a 1-cell on P_2, for in this case we will not only produce a lower complexity diagram, but we will make a more substantial shift of our choice of initial or terminal positions while retaining our hypotheses. Let v be the endpoint of the path along the side of L facing P_a. There are two cases according to whether $v \leq_1 Y^+$ or $v >_1 Y^+$.

Let v' denote the other endpoint of m', and observe that $v \leq_2 Y^+$ if and only if $v' \leq_2 Y^+$. Indeed, if v lies on Y or on a square s dual to a codimension-2 hyperplane of $H_1 \cap H_2$ emanating from Y then so does v' (and vice-versa). This is because wall carriers cannot interosculate with cones or squares by Theorem 5.23. See Figure 100.

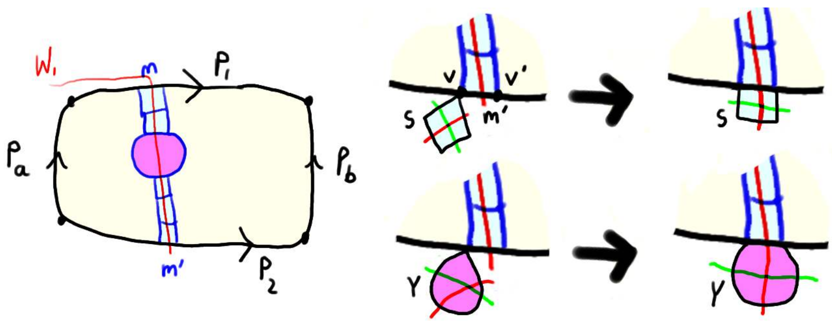

FIGURE 100.

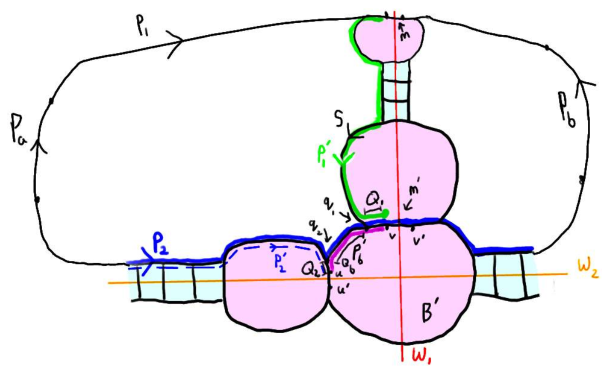

FIGURE 101.

Consider the second case where $v >_1 Y^+$. The path P_2 is supported by a ladder in $N(W_2)$. Let B' denote the cone or square on this ladder which contains the 1-cell m', and note that as above, $B' \neq Y$ in the cone case, and B' is not a square dual to $H_1 \cap H_2$ emanating from Y.

Let q_2 be the first point on P_2 that lies on B'. Let Q_2 be a shortest path in B from q to a 1-cell dual to W_2, and let u be the endpoint of Q_2. Let S be the path on the ladder L from m to m'. Let q_1 be the first point of S that lies on B', and let Q_2 be the terminal subpath of S from q_1 to v. Let Q_b be the subpath of P_2 from q_1 to q_2.

We now define P'_2 to be the path obtained from P_2 by substituting Q_2 for the part after q_2. We define P'_b to be the concatenation $Q_2^{-1}Q_bQ_1$. The path P'_1 will be the concatenation of the part of P_1 preceding m, followed by the initial part of S until q_1, followed by Q_1.

We note that when B' is a square instead of a cone, the path Q_b is trivial. The construction is similar if m and/or m' lie on P_a and/or P_b.

The new quadrilateral $P'_aP_1P'^{-1}_bP_2^{-1}$ bounds a proper subdiagram D' of D.

In the second case where $v > Y^+$, we will replace A instead of B. The details are similar. We refer to Figure 102.

We note, that as in the fourth figure, since we have shown in Theorem 5.20 that there is no self-crossing of walls, a dual curve (in a dual graph) that passes through the same cone-cell twice, is actually passing through the same wall of the corresponding cone. The third and fourth diagrams represent a dichotomy between the possibility that an endpoint of the ladder carrying the dual curve w does or does not lie on Y^+. In the third diagram it does not, so we produce a smaller complexity example with a different pair of endpoints. In the fourth diagram it does lie on Y, and so we choose a different initial point but the same endpoint. We emphasize that in this case, since a hyperplane carrier cannot interosculate with a cone, if one endpoint of the dual 1-cell of the ladder lies in Y, then both do. We could have thus chosen an even smaller diagram (on the opposite side of this ladder) with the same

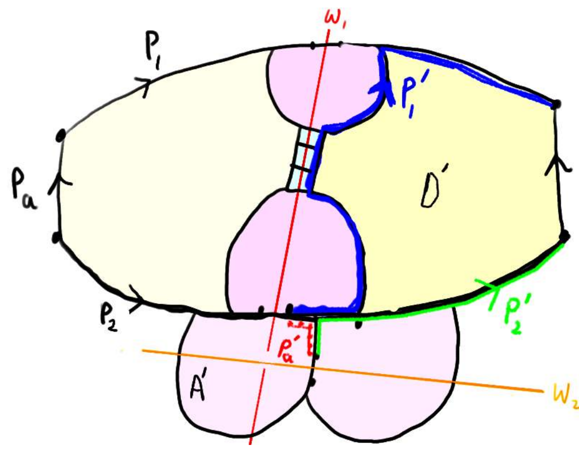

FIGURE 102.

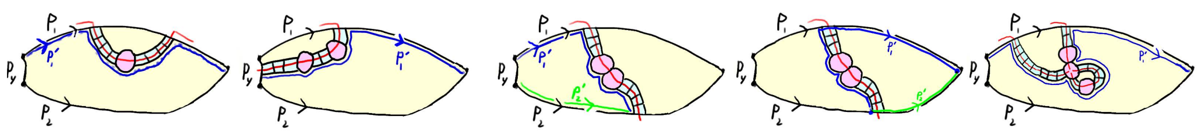

FIGURE 103.

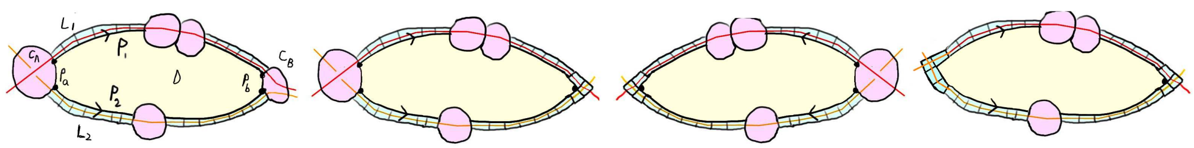

FIGURE 104.

properties. A degenerate version of this situation occurs when the initial 1-cell of P_1 is dual to W_1 and is a spur of D.

Having shown that D does not contain dual 1-cells of W_1 along $P_aP_1P_b^{-1}$ or dual 1-cells of W_2 along $P_a^{-1}P_2P_b$, we now follow Construction 5.21 to obtain a bicollared diagram with two corners.

We first attach a square or cone-cell C_A along P_a where C_a maps to A and has boundary path $\partial_p C$ extending the path $e_1P_ae_2$ where each e_i is a 1-cell dual to W_i at the initial point of P_i. Similarly, we attach a square or cone-cell C_B mapping to B along the path $f_1P_bf_2$ where f_1, f_2 are 1-cells dual to W_2 at the initial and terminal vertices of P_2. We then attach a ladder L_i starting with C_A and ending with C_B along each P_i carrying a dual curve w_i in W_i. (Note that these are only ladders in a general sense allowing initial and terminal cells to be squares...) In this way we form the diagram E as in one of the diagrams in Figure 104.

The next stage of the argument is to remove cancellable pairs in E to obtain a reduced diagram that is a counterexample with the same structure, in the sense that it is bicollared by ladders associated to W_1, W_2, but that the corner at A maps to a cone or square with a vertex satisfying $v \leq_i Y^+$ for some i, but the corner at B maps to a cone or square with a vertex satisfying $v >_1 Y^+$ and $v >_2 Y^+$

We refer the reader to Figure 105 for various possibilities that arise when removing cancellable pairs. There are several situations where we must pass to a new bicollared (almost sub) diagram that starts at a cone-cell, and does not end on Y or on a square that is dual to a component of a codimension-2 hyperplane $H_1 \cap H_2$ that contains a component of $h_1 \cap h_2$.

Note that we are also forced to consider "singular bicollared diagrams" as in Remark 5.40 and Figure 107.

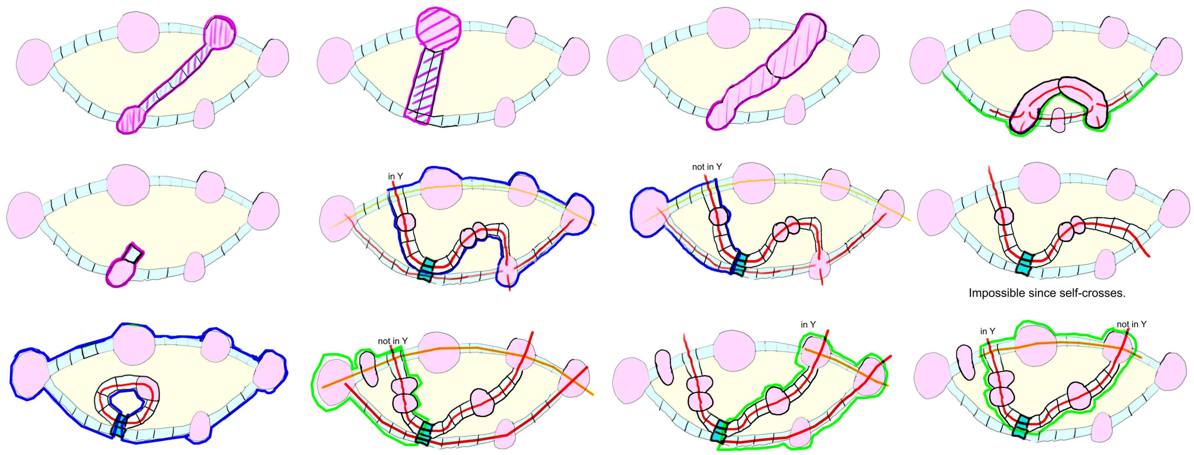

FIGURE 105.

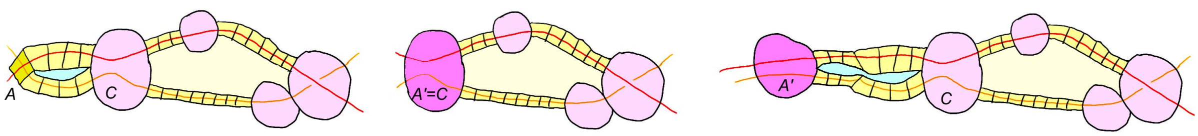

FIGURE 106.

The main point here is to consider a cancellable pair from a fake wall-piece or cone-piece which was part of the collar on the other side of the diagram E. If so, there is a dichotomy between this cone-cell occurring at Y or not. If it doesn't then we obtain a smaller complexity diagram by adjusting the endpoints, and if it does occur on Y then we obtain a smaller complexity diagram by adjusting the initial point.

There is one possible type of cancellable pair that deserves special mention: This is when A is a square, and forms a corner of a generalized square whose outerpath lies on some cone-cell C in E. Note that C is necessarily in both collars. A similar such situation can occur with B a square. We illustrate the A case on the left in Figure 106.

In the A case, if C is a cone-cell mapping to Y, the as is illustrated in the second diagram in Figure 106, we simply crop off the initial part of the diagram, and allow $A' = C$ to play the role of A.

If C does not map to Y, but A maps to a square that is dual to a codimension-2 hyperplane of $H_1 \cap H_2$ emanating from Y, then we remove the interior and boundary path of A from the diagram, and add a cone-cell A' mapping to Y, and attach it to the internal path of A along a pair of square ladders mapping to H_1, H_2. The modified result is on the right in Figure 106. Note that from the viewpoint of "reducing the diagram" we can regard the square A as having been first absorbed into C, before we add the new cone-cell that has a length-2 piece with C.

Let $p = p_1$ (which equals p_2 in this case). By hypothesis, for at least one value $i \in \{1, 2\}$ we have $p \leq_i Y^+$.

If C maps to Y, then we still have $p' \leq_i Y^+$, where p' is at the "inside" of the initial corner $A' = C$ of our new collared diagram (the second picture)

If C doesn't map to Y, and $p \in Y^+$ then $p' \in Y$ since we augmented the diagram by adding a cone-cell A' mapping to Y together with two rectangles.

If C doesn't map to Y, and $p \notin Y^+$ then we use the second picture possibility. The crucial point that must be verified is that C does not contain a 1-cell dual to the wallray (W_i, Y, H_i), or in other words, that C is outside the carrier of this wallray. Observe that to enter this wallray from a point not on Y and

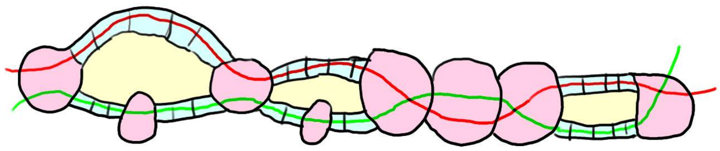

FIGURE 107.

arrive at a point within the carrier of this wallray that is not on Y, one must first travel to Y and then travel around Y to jump across the wall in $W_i \cap Y$ represented by $H_i \cap Y$, by traveling from a piece on Y along one hyperplane of $W_i \cap Y$ (not equal to $H_i \cap Y$) to a piece on Y along $H_i \cap Y$.

Note that there is a path from p to C along a single hyperplane, and this certainly cannot jump across a wall of Y, and so the cone of C cannot lie within $N(W_1, Y, H_1)$ Similar arguments show that a path along a rectangle followed by a path around C (starting and ending on the same wall) cannot travel around Y but I'm not sure if this is necessary.

In the B case, we simply remove B together with the width 2 square ladder from C to B, and allow $B' = C$ to play the role of B. Note that C cannot be a cone-cell mapping to Y because then B would be dual to a codimension-2 hyperplane of $H_1 \cap H_2$ emanating from Y, so the endpoint $q = q_1 = q_2$ of D lying on B would contradict our hypothesis that $q >_i Y^+$ for each i.

Let F denote the new bicollared diagram.

First note that F must be a ladder by Theorem 3.36.

Then note that by construction, F has the property that its initial cone-cell does not lie in $N(W_k, Y, H_k)$ for at least one of $k = 1, 2$, however its terminal cone-cell lies in $N(W_i, Y, H_i)$ for each i.

The bicollar of F along W_k consists of a (sub)ladder which must have some cone-cell mapping to Y. Regarding A, B as the leftmost and rightmost cone-cells of the ladder, we consider the rightmost cone-cell C_r mapping to Y, and note that it cannot be the final cone-cell B since $B > Y^+$.

Finally, observe that the dual curves of the square ladders carrying the part of W_1 and W_2 emanating rightwards from C_r to C_{r+1} must map to H_1, H_2. Indeed, otherwise we would backtrack into a cone-cell also not in $N(W_k, Y, H_k)$ and there would be a Y cone-cell appearing further to the right.

However, this means that these dual curves lie in a piece between C_r and C_{r+1} which contradicts the hypothesis of the theorem that $\Omega_Y(H_1, H_2) > 0$. □

Remark 5.40. We note that in the course of the proof it is natural to consider *singular collared diagrams* which have the property that the collars can occupy the same 2-cell as in Figure 107. These arise when a cone-cell or square in one collar is absorbed into a cone-cell on the other collar because of a cancellable pair.

Remark 5.41. In general, it is possible for two hyperplanes H_1, H_2 to cross two different cones Y_1, Y_2 such that Y_1, Y_2 are not simply joined together by a generalized rectangle, but could have some sequence of cones between them. This can be avoided by extra hypotheses on the hyperplanes within a cone. We refer the reader to Figure 108.

I believe the following special case of Theorem 5.37 could have been obtained more directly along the lines of Lemma 5.6 and Lemma 5.15, and might have facilitated the proof of Theorem 5.37.

Corollary 5.42. *Let Y be a cone and let H_1, H_2 be hyperplanes in $\widetilde{X}^*$. Suppose that $\Omega_Y(H_1 \cap Y, H_2 \cap Y) > 0$. Then $H_1 \cap H_2$ is the union of codimension-2 hyperplanes emanating from Y.*

5.11. ⊛ **1-dimensional linear separation.** The wallspace $\widetilde{X}^*$ satisfies the *linear separation property* if $\#(p, q) \geq Ld(p, q) - M$ for some constants $L, M \geq 1$. Here $\#(p, q)$ denotes the number of walls separating the 0-cells p, q, and $d(p, q)$ is the distance between p, q in the 1-skeleton of $\widetilde{X}$.

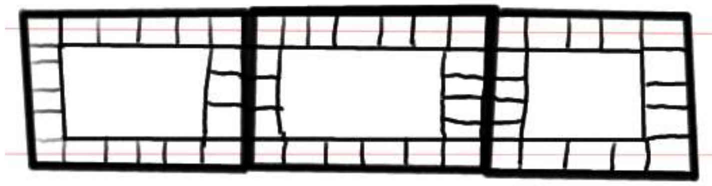

FIGURE 108. There is a ladder containing three cone-cells in the complex above such that the first and last cone-cells are connected by a pair of hyperplanes without a square diagram between them.

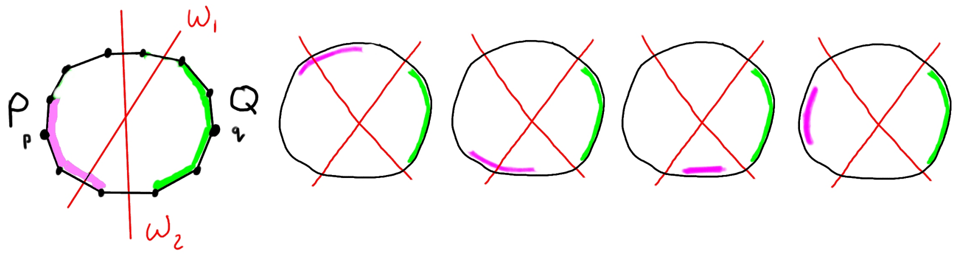

FIGURE 109.

Definition 5.43. Let X be 1-dimensional. We say $p, q \in Y^0$ are *strongly separated* by the wall w of Y if no 1-cell e dual to a hyperplane of w lies in the same piece as p or q.

We say the cone Y in $\widetilde{X}^*$ has the $\frac{\pi}{2}$*-strong separation property* if p, q are strongly separated by some wall whenever $\Omega_Y(p, q) \geq \frac{\pi}{2}$.

Lemma 5.44. *Let $\langle x_1, \ldots, \mid R_1, \ldots \rangle$ be an ordinary presentation, satisfying the ordinary $B(6)$ small-cancellation condition. So more than half a relator cannot be the concatenation of fewer than four pieces. Use the split-angling (which is the same as the grade-angling here with each 2-cell having grade* 6*).*

Then each relator has the $\frac{\pi}{2}$-strong separation property.

Proof. Suppose p, q are points on the boundary of a 2-cell (i.e. a cone) that do not lie in the concatenation of two pieces. Let P and Q be "piece neighborhoods" of p and q in the sense that P contains every edge that lies in a piece with p, and similarly for Q. By hypothesis, $P \cap Q = \emptyset$. By possibly exchanging the notation, assume that $|Q| \geq |P|$. Let w_1, w_2 be the walls dual to the edges immediately before and after Q. If w_1 and w_2 are both dual to edges of P then $|P| > |Q|$ which is impossible. If neither w_1 nor w_2 is dual to an edge of P, then P is separated from Q by one or both of w_1, w_2.

We refer the reader to Figure 109. The diagram on the left indicates the notation, and it is then obvious that if Q passes through both hyperplanes then its length exceeds $|P| + 2$. The diagrams on the left indicate some of the possible combinatorial fashions that one of the walls w_1, w_2 might separate. □

Theorem 5.45. *Suppose X is 1-dimensional, and $\langle X \mid Y_i \rangle$ is a cubical presentation with an angle assignment satisfying the $B(6)$ condition and short innerpaths, and the $\frac{\pi}{2}$-strong separation property for each cone.*

Then $\widetilde{X}^*$ satisfies the linear separation property.

Proof. Let γ be a geodesic in $\widetilde{X}^*$. Let e_1 be a 1-cell on γ. We will show that either the wall W_1 in $\widetilde{X}^*$ that is dual to e_1, crosses γ in no other dual 1-cell, or else there is a 1-cell e_2 within a uniformly bounded distance of e_1, so that the wall W_2 dual to it has the same property. Either way, this shows that each edge of γ is within a uniform distance of a wall that γ crosses at a single 1-cell. This is easily seen to imply the linear separation property.

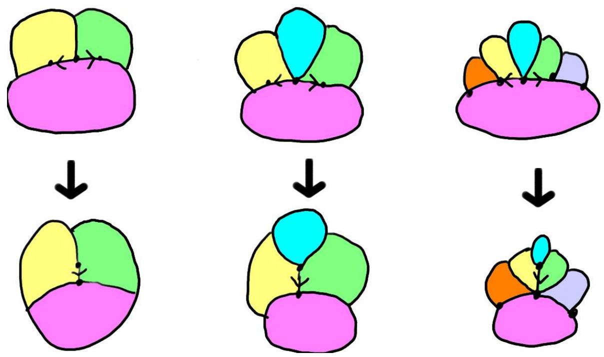

FIGURE 110.

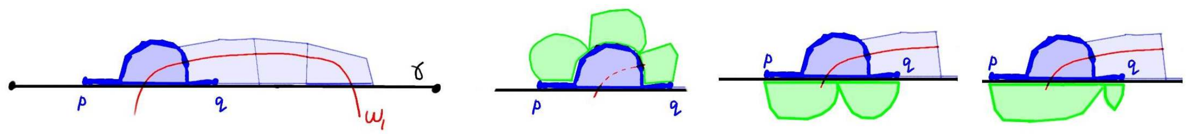

FIGURE 111.

By Corollary 5.34, the geodesic γ lies in $T(W_1)$. Let C_1 be an initial cone-cell in a ladder in $T(W_1)$ that carries γ and starts on the dual 1-cell e_1 and ends on a second dual 1-cell e_1'. Let p, q denote outermost points in $\gamma \cap Y_1$ where Y_1 is the cone supporting C_1.

We now show that $\Omega_Y(p, q) \geq \frac{\pi}{2}$. Let γ' denote the subpath of γ joining p, q. Suppose there is a path σ from p to q with $\Omega_Y(\sigma) < \pi$. Since γ is a geodesic in $\widetilde{X}^*$, we must have $|\gamma'| \leq |\sigma|$. Since $\Omega_Y(\sigma) < \pi$ and it is impossible that $|\sigma| < |\gamma'|$ we see that σ and γ' are path-homotopic in Y by Condition 5.4. Consequently, if $\Omega_Y(\sigma) < \frac{\pi}{2}$ then $\Omega_Y(\gamma') \leq \Omega_Y(\sigma) < \frac{\pi}{2}$ by the following observation:

Note that $\Omega(\alpha) \geq \Omega(\beta)$ whenever Y is 1-dimensional and $\alpha \to Y$ and $\beta \to Y$ are path homotopic, and β is an immersion. This is suggested by Figure 110. Note that the number of internal corners along α decreases as backtracks are folded outwards. So our claim is clear for the grade-angling. For the split-angling, one finds that two defects of $\frac{\pi}{3}$ and/or $\frac{\pi}{2}$ can become a single defect of $\frac{\pi}{3}$ or $\frac{\pi}{2}$. Therefore the total defect does not increase.

However, we refer the reader to the two rightmost diagrams in Figure 111 for an explanation of why $\Omega_Y(\gamma') > \frac{\pi}{2}$. Indeed, we can concatenate part of γ' with part of the piece between C_1 and the next cone-cell in the ladder for W_1, to obtain a path δ with $\Omega_Y(\delta) < \pi$ (the part of γ' contributed a defect of $< \frac{\pi}{2}$ and at most an extra $\frac{\pi}{2}$ of defect occurs at the transition between these subpaths because there are no acute corners). Since δ crosses the wall of $W_2 \cap Y$ in two distinct hyperplanes, we obtain a contradiction of Condition 5.1.(5).

By the $\frac{\pi}{2}$-strong separation hypothesis, there exists a wall W_2 separating p, q. We refer the reader to Figure 112 for an explanation of why W_2 cannot double cross γ. All possible situations lead to either p, q lying in the same piece as one of the hyperplanes of $W_2 \cap Y$, or lead to a path μ with $\Omega_Y(\mu) < \pi$ that passes through distinct hyperplanes in the same wall of a cone-cell Y thus contradicting Condition 5.1.(5). □

5.12. ⊛ **Linear Separation when X is a pseudograph.**

Definition 5.46 (Pseudograph)**.** A nonpositively curved cube complex X is a *pseudograph* if each hyperplane of X is a compact CAT(0) space.

Lemma 5.47. *Suppose X is a compact pseudograph. Then $\pi_1 X$ is free.*

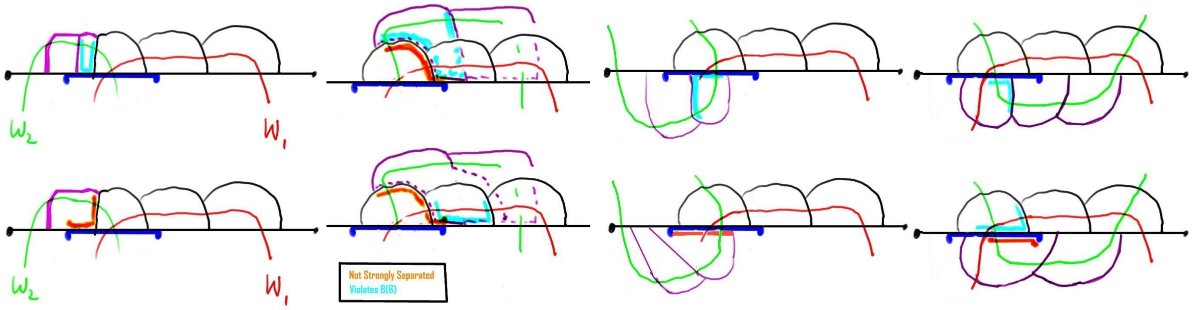

FIGURE 112. I believe this covers most of the variety of cases.

I don't know if Lemma 5.47 holds when X is a nonpositively curved cube complex whose immersed hyperplanes are simply-connected (without a compactness hypotheses).

The following proof can probably be replaced by a direct elementary proof that $\widetilde{X}$ is quasiisometric to a tree.

Proof. Let H be a hyperplane in $\widetilde{X}$. If each component of $X - H$ is deep, then we see that $\widetilde{X}$ and hence X has more than one end, and hence $\pi_1 X$ is infinite cyclic or splits as a free product. The factors are quasiconvex subgroups.

Thus each factor is again a pseudograph with the same property. Hence each factor is free (by induction on rank).

If exactly one component of $X - H$ is not deep, then $\pi_1 X$ leaves invariant the intersection of the deep halfspaces, and so the theorem is true by induction on the number of cells in X.

If neither component is deep, then $\pi_1 X$ stabilizes H and so $\pi_1 X$ is trivial. □

Remark 5.48. The main difficulty in generalizing the proof of Theorem 5.45 is that it is unclear what plays the role of p, q (e.g. min and max points in $C \cap S$), and it is difficult to ensure that geodesic paths in $S \cap L_i$ overlap with p, q or some substitute.

We say a wall w in Y is *contiguous* with a hyperplane u if one of the following hold for some hyperplane v of w:

(1) $v = u$.
(2) v crosses u.
(3) v and u have dual 1-cells sharing an endpoint (this covers the previous two cases).
(4) v has a dual 1-cell with an endpoint lying on a piece in Y that contains a dual 1-cell of u.

Theorem 5.49. *If $\langle X \mid Y_i \rangle$ satisfies the following properties then $\widetilde{X}^*$ has linear separation.*

(1) *X is a compact pseudograph.*
(2) *X^* satisfies the $B(6)$ condition.*
(3) *All cones are finite.*
(4) *For distinct hyperplanes w_1, w_2 in the same wall of Y, there does not exist a path $\varepsilon \to Y$ with $\Omega_Y(\varepsilon) < \frac{\pi}{2}$ whose first and last edges e_p, e_q are dual to hyperplanes h_p, h_q, such that h_p equals w_1 and h_q has a dual 1-cell with an endpoint on a dual 1-cell of w_2 or a piece containing a dual 1-cell of w_2.*
(5) *For each Y, and each CAT(0) interval $J \subset Y$, and hyperplanes $\dot{h}, \dot{k} \subset J$ consisting of components of intersections of hyperplanes h, k with J the following holds: If $\Omega_Y(\dot{h}, \dot{k}) \geq \frac{\pi}{2}$ then there is a wall w in Y that separates h, k but is not contiguous with h or k.*

Remark 5.50. The hypothesis that all cones are finite can be replaced by a strengthening of Condition (5). Specifically, we assume that the wall w_2 of Y separating h, k should have a hyperplane lying lying within a uniformly bounded distance of w_1.

Remark 5.51. The statement of the theorem suggests that we should have declared $\Omega_Y(e, f)$ to be the infimum of defects of paths that start and end at 1-cells dual to the same hyperplanes as e and f.

Proof. Let $S \to \widetilde{X}$ be the local convex hull of the geodesic $\gamma \to \widetilde{X}^*$ from p to q.

We will show that each edge of S lies within a uniform distance of a hyperplane whose wall in $\widetilde{X}^*$ intersects S in a single hyperplane. Since there are uniformly many hyperplanes in S within a uniform distance of any point, we obtain the desired linear separation.

Let e_1 be an edge of S that is dual to a wall W_1 of $\widetilde{X}^*$ intersecting S in more than one hyperplane. Let h_e be the hyperplane of S dual to e_1. By Theorem 5.33, considering the subpath of γ starting and ending on edges dual to distinct hyperplanes in W_1, we find that this subpath is homotopic to a subpath λ that lies on a ladder $L_1 \to T(W_1)$. Accordingly, the ladder L_1 begins with a square ladder in S that starts at e_1 and terminates at a 1-cell f_1, and is followed by a cone-cell C mapping to a cone Y, such that f_1 maps to Y and is dual to a hyperplane w_1 of Y, and w_1 lies in the same wall as a hyperplane w_1' of Y with $w_1' \neq w_1$. And the cone-cell C either meets S in a 1-cell e_1' dual to w_1' or else the cone-cell C has a piece (with a subsequent cone-cell in L_1), so that the piece contains the 1-cell e_1' dual to w_1', and an endpoint of the piece contains a point of λ which is necessarily in S.

Let J be the component of the preimage of Y in S that intersects h_e.

Let h_p be a hyperplane in S passing through J that is extreme in the following sense: for any geodesic σ_p from a 1-cell dual to $h_p \cap J$ to the point p, all 1-cells of the initial subpath $\sigma_p \cap J$ are dual to hyperplanes crossing h_p (except for the initial 1-cell dual to h_p). We observe that for any geodesic in S from J to p whose intersection with J consists precisely of its initial 1-cell, the hyperplane dual to this 1-cell has the property above, and so we can choose h_p in this manner. We choose the hyperplane h_q analogously.

Note that we can assume that we chose h_p so that either $h_p = h_e$ or so that h_p is disjoint from h_e and separates p, h_e. Indeed, suppose that h_e is not a final hyperplane in J to p. Then there is a geodesic α from a 1-cell of J dual to h_e to the point p, so that the final 1-cell of $\alpha \cap J$ is dual to a hyperplane that doesn't cross h_e. We then let h_p be the hyperplane dual to this 1-cell.

Let $\dot{h}_p = h_p \cap J$ and $\dot{h}_q = h_q \cap J$. Our hypothesis that Condition (4) holds implies that $\Omega_Y(\dot{h}_p, \dot{h}_q) \geq \frac{\pi}{2}$. Indeed, let κ be a path starting and ending at 1-cells dual to $\dot{h}_p$ and $\dot{h}_q$. Since either $h_p = w_1$ or w_1 separates p, h_p, we see that κ passes through a 1-cell dual to w_1. Our choice of h_q implies that it is dual to a 1-cell with an endpoint in J that is either also on a 1-cell dual to w_2 or also on a piece in Y containing a 1-cell dual to w_2. Thus Condition (4) implies that $\Omega_Y(\kappa) \geq \frac{\pi}{2}$.

By the strong-separation hypothesis in Condition (5), there exists a wall w_2 of Y that strongly separates h_p, h_q in the sense that w_2 is not contiguous with either h_p or h_q.

Let W_2 be the wall of $\widetilde{X}^*$ corresponding to w_2. We will show that W_2 cannot intersect S in a second hyperplane. Indeed, suppose that W_2 intersects S in a second hyperplane. Let σ be a geodesic from p to q that passes through a 1-cell e_2 in J that is dual to w_2.

There are two cases to consider according to whether σ intersects W_2 in another hyperplane in S on its way to p or on its way to q. We consider the former case and note that there is an analogous argument for the latter case holding for q and σ_q and h_q.

Let σ_p denote the subpath of σ that starts on w_2 and ends on a second 1-cell dual to W_2 (in the direction of p).

We refer the reader to Figure 113 for diagrams indicating the notations.

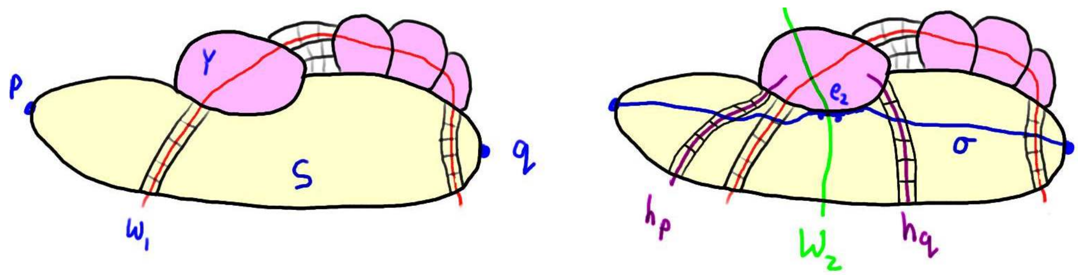

FIGURE 113.

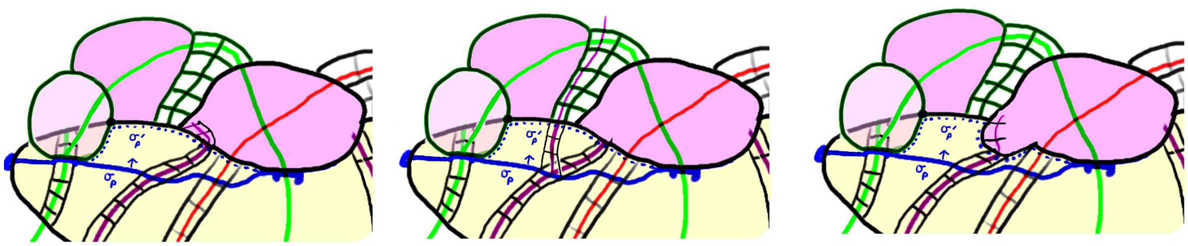

FIGURE 114.

By Theorem 5.33, the path σ_p is square-homotopic to a path σ'_p that lies along a ladder $L_2 \to T(W_2)$. The path σ'_p must pass through a 1-cell f_2 of S that is dual to the hyperplane h_p, and this 1-cell f_2 actually lies in $J \subset S \cap Y$. Indeed, by choice of w_2 we know that $\Omega_Y(h_p, w_2) \neq 0$ and h_p, w_2 do not cross in Y, and so by Theorem 5.37, even the wallray from Y emanating in the direction of h_p cannot cross the wall W_2 outside of Y. The second diagram in Figure 114, illustrates the impossible scenario that h_p veers into L_2 in a way that makes it cross W_2 outside of Y. The third diagram indicates that the cone-cells might lie substantially inside S - more than was illustrated for the wall W_1.

Let δ denote the subpath of σ'_p that starts on e_2 and ends on f_2. By the property characterizing h_p, all the vertices of δ are endpoints of 1-cells dual to h_p. Thus an endpoint of a 1-cell dual to h_p lies on the same piece as a 1-cell dual to w_2 which is impossible. □

5.13. ⊛ **Codimension-1 subgroup preserved.** Show that (under certain hypotheses) each new wall is cut by an infinite new wall that proceeds infinitely deeply on each side.

This certainly fails if we consider $\langle X \mid Y \rangle$ where Y is a very high girth cover of X satisfying all the properties. The main hypothesis will ensure that new walls are infinitely extendible.

5.14. **Elliptic Annuli.**

Definition 5.52 (Elliptic Annulus)**.** An element $g \in \pi_1 X^*$ is *elliptic* if $g\widetilde{Y}_i = \widetilde{Y}_i$ for some lift $\widetilde{Y}_i$ of some cone Y_i of X^*.

An annular diagram $A \to X^*$ is *elliptic* if the lift of its universal cover $\widetilde{A} \to \widetilde{X}^*$ actually lifts to a cone $\widetilde{Y}_i$.

Note that when A is elliptic, both boundary paths P_+, P_- of A represent elliptic elements (or rather, conjugacy classes) in $\pi_1 X^*$.

A typical elliptic annulus $A \to X^*$ contains a single cone-cell that overlaps with itself at an internal path that is not a piece. As opposed to the typical situation in a disk diagram, it is impossible to "reduce" A by combining the cone-cell with itself. When a cone Y_i of X^* contains an interesting square annulus, then one obtains an elliptic annulus which is disguised by being built out of squares - though it could be replaced by a single self-overlapping cone-cell as above. There are two special cases worth mentioning: Any elliptic element yields an elliptic annulus that is isomorphic to a subdivided circle. If

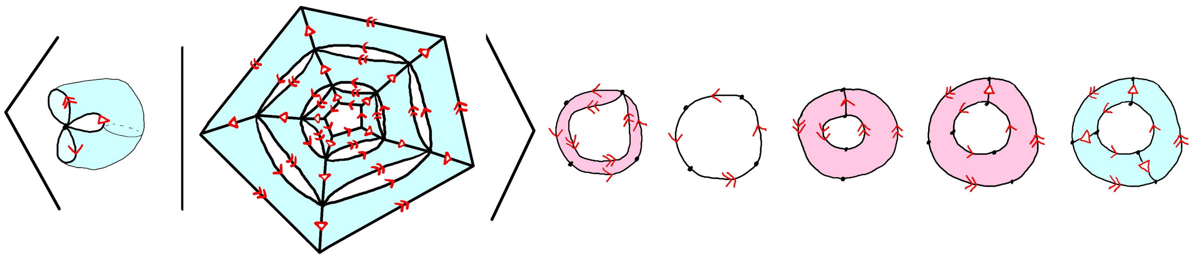

FIGURE 115. Elliptic Annuli.

Y_i contains a hyperplane that is stabilized (without inversion) by some element $g \in \text{Aut}(Y_i)$, then this data yields an interesting elliptic annulus that looks like a closed square annuladder. More generally, in the $B(6)$ case, we regard an elliptic annuladder carrying the dual curve of some wall as a *length* 0 annuladder. See Figure 115.

Lemma 5.53 (Elliptic subannuli merge)**.** *Suppose X^* is a small-cancellation complex, and internal cone-cells have negative curvature, and $\widetilde{Y}$ is superconvex for each cone Y, and contiguous cone-pieces are finite (or more generally, have trivial stabilizers in $\pi_1 X$).*

If an annular diagram A contains two or more essential elliptic annular subdiagrams. Then they are associated to the same cone $\widetilde{Y}$.

An annulus is *essential* if its boundary paths are essential in X^*.

Proof. Reduce as much as possible, and consider the annular subdiagram bounded by two nontrivial elliptic annuli. We can assume there is no essential elliptic annular subdiagram between them (otherwise consider consecutive pairs).

Let B denote the annular subdiagram between (but not including) them, and consider a generalized corner of a square or a nonnegatively curved θ-shell along ∂B. The former would absorb into one of the two elliptic annuli, and likewise, the latter must be absorbable or replaceable since the defect sum around it would be exactly 2π. (In both cases, the interesting situation is where the corner or outerpath is at the transition in the elliptic annulus.)

Thus B is not just a flat annulus as in Lemma 2.10 but an actual product. If B contains squares, then the superconvexity of cones absorbs the square ladders along the outside. If B contains no squares, then $\widetilde{B}$ would be an infinite contiguous cone-piece between the two distinct cones associated to the elliptic annuli on either side of ∂B, thus these cones are the identical, and the elliptic annuli merge. Note that the assumption that the elliptic annuli are nontrivial implies that B represents a nontrivial conjugacy class in X^* and hence in X, so $\widetilde{B}$ is infinite. □

We now briefly focus on elliptic W-annuladders which are defined in Definition 5.67. The following is similar to Lemma 5.53:

Lemma 5.54. *Suppose X^* is small-cancellation and $B(6)$ and superconvex cones. Suppose $A \to X^*$ is a reduced annular diagram, containing an essential elliptic annulus and an essential W-annullader. Then they are the same.*

Sketch. Let B denote the diagram between the W-annuladder and the elliptic annulus such that B includes the former but not the latter. Then B has nonpositive curvature on the W-annuladder side, and negative curvature on the elliptic annulus side.

One first concludes that B equals the W-annuladder, and one then concludes that the W-annuladder absorbs into the elliptic annulus. □

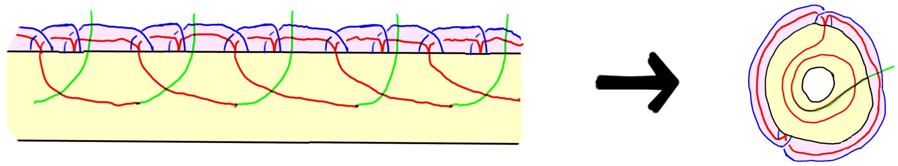

FIGURE 116.

5.15. **Annular Diagrams and the $B(8)$ condition.**

Definition 5.55. [Generalized $B(8)$ condition] The following conditions will restrict the structure of certain annuli:

(1) [Strict metric small-cancellation] For each closed path SQ in a cone Y, if $\Omega_Y(S) \leq \pi$ then either $|S| < |Q|$ or SQ is nullhomotopic in Y.
(2) [Negative curvature relative to walls] Let S be a path in Y whose first and last 1-cell are dual to hyperplanes in the same wall. If $\Omega_Y(S) \leq \pi$ then these 1-cells are dual to the same hyperplane, and S is homotopic into this hyperplane in Y.
(3) [Negative Curvature] For each closed cycle $P \to Y$, if $\Omega_Y(P) \leq 2\pi$ then $P \to Y$ is nullhomotopic.

Remark 5.56 (Comparison with Definition 5.1)**.**
Condition 5.55.(3) strengthens Condition 5.1.(2).
Condition 5.55.(1) strengthens Condition 5.4.
Condition 5.55.(2) strengthens Condition 5.1.(5).

Definition 5.57. A *W-ladder in a diagram* is a ladder containing a dual curve mapping to a wall W of $\widetilde{X}^*$, it provides a generalization of a square ladder containing a dual curve. Since a cone-cell can offer multiple continuations of a W-ladder, we will often choose continuations that yield a simple dual curve, and more specifically, a *simple* W-ladder is an embedding on its interior. On the other hand, we will sometimes consider an entire dual graph instead of a dual curve, and then obtain a *branched* W-ladder.

We will often use the term *cladder* when suppressing the associated wall W.

Lemma 5.58 (W-ladders cannot self-cross)**.** *Let $A \to X^*$ be an annular diagram whose outside boundary λ is collared by a wall W, in the sense that $\widetilde{\lambda}$ lies on $N(W)$.*

Let $L \to A$ be a W-ladder that is dual to a 1-cell of λ dual to W. Then L cannot self-cross, in the sense that it passes through the same square in non-parallel 1-cells, or the same cone-cell in 1-cells that map to distinct walls in the corresponding cone.

Moreover, L cannot end on another 1-cell of the outside boundary of λ unless it too is dual to W.

Proof. This is true for a disk diagram because of Theorem 5.20.

Let $\widetilde{A} \to A$ be the universal cover of A, and observe that $\widetilde{A}$ is collared by the universal cover of the collar of A. If L self-crosses in A then a pair of distinct lifts of $\widetilde{L}$ in $\widetilde{A}$ cross each other. Connecting them along the collar of $\widetilde{A}$, this gives a self-crossing W-ladder within a disk diagram, which is impossible by Theorem 5.20.

Similarly, if L enters A on a 1-cell on the outside boundary that is dual to W but exits at a 1-cell on the outside boundary that is not dual to W within the collar, then the lift of $\widetilde{L}$ in $\widetilde{A}$ does the same thing, which is impossible. We refer the reader to Figure 116. □

Theorem 5.59 (⊛ Conjugate into wall)**.** *Suppose that $\langle X \mid \{Y_i\}\rangle$ satisfies the conditions of Definition 5.1 and additionally, the strengthened conditions of Definition 5.55, and in particular Definition 5.55.(1).*

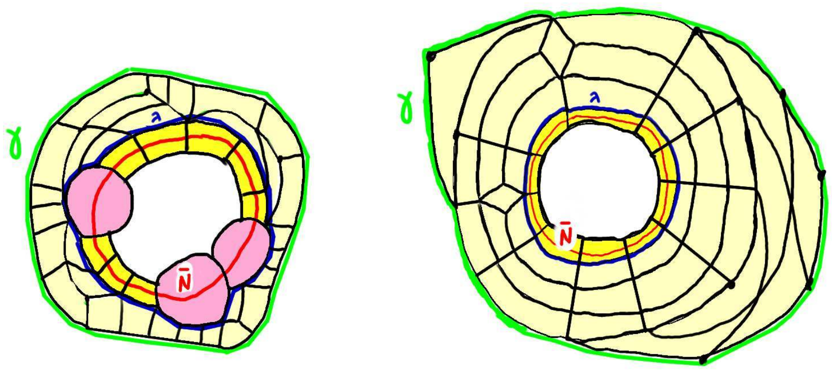

FIGURE 117.

Let $\gamma \to X^*$ be an immersed essential circle that has minimal length within its homotopy class in X^*. Suppose γ is homotopic to a closed path λ that lies in the carrier of a wall of $\widetilde{X}^*$. So $\widetilde{\lambda} \subset N = N(W)$, or equivalently, $\lambda \to X$ factors as $\lambda \to \bar{N} \to X^*$, where $\bar{N} = \langle\gamma\rangle\backslash N$.

Let A be an annular diagram between γ and λ, and suppose that A is reduced and of minimal complexity subject to $\lambda \to \bar{N}$ varying within its homotopy class in $\bar{N}$, and subject to $\widetilde{A}$ having one side equal to the fixed copy $\widetilde{\gamma} \subset \widetilde{X}^*$ of the universal cover of γ, and the other side lying on a fixed copy of $N \subset \widetilde{X}^*$.

Then A is a square diagram in X, and γ lies on one side of a flat subannulus, whose other side γ' lies in the local convex hull of λ in A.

We refer the reader to Figure 117 for two possible scenarios. In the second case, γ is homotopic to a path in a square annulus in $\bar{N}$. In the first case γ is homotopic to a more general annulus in $\bar{N}$ containing cone-cells.

Remark 5.60. We note that $\widetilde{\gamma}$ lies in $T(W)$ if $\widetilde{\gamma}$ intersects $N(W)$ in a point. Indeed this follows by combining the above result with Theorem 5.33. However, in general it is possible that $\widetilde{\gamma}$ is disjoint from $T(W)$ and even $E(W)$. Indeed, the simplest example is the subdivided cylinder $[0, n] \times S^1$. We let γ denote $\{0\} \times S^1$, and we let W denote the universal cover of $\{n - \frac{1}{2}\} \times S^1$. Then $\widetilde{\gamma}$ does not lie in $N(W) = T(W) = E(W)$ which equals $[n-1, n] \times \widetilde{S}^1$.

The proof shows that this is essentially the only possible failure, in the sense that if λ does not lie along a square annulus of $\bar{N}$ then $\widetilde{\gamma}$ is forced to intersect $N(W)$.

Proof. Let A be an annular diagram as in the hypothesis of the theorem. So the inside boundary path of A is γ which is of minimal length (among closed paths in X) in its homotopy class in X^*, and the outside boundary path of A is a path $\lambda \to N(W)$, and A is reduced in the sense that it doesn't have removable square bigons, or absorbable squares or replaceable cone-cells or combinable cone-cells, and furthermore there does not exist a local complexity reduction achieved by pushing λ across some square or cone-cell that maps to $N(W)$. We note that all such reductions preserve the "class" of A in the sense that the lift of $\widetilde{A}$ at $\widetilde{\gamma}$ has the property that the outside boundary path lifts to the same fixed copy of $N(W)$ no matter what local reduction is performed above.

If γ' has a common point with λ then the result follows from Theorem 5.33 which states that a geodesic starting and ending on the carrier is square homotopic into the carrier.

As suggested by the first and second diagrams in Figure 118, our choice of A implies that λ cannot pass through any dual 1-cell of W. For then, following a simple W-ladder emanating from this dual curve, we are able to find a lower complexity annulus of the same class, using a different choice of λ but the same γ. Note that this W-ladder cannot self-cross within A_λ, as each of the various self-intersection

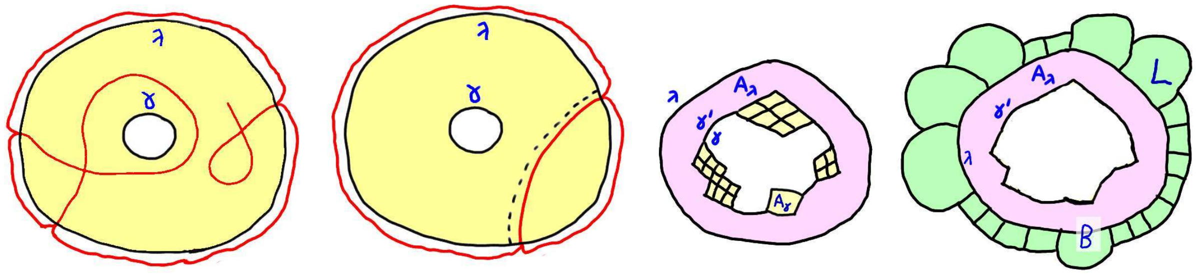

FIGURE 118.

possibilities are ruled out by Theorem 5.20 as established in Lemma 5.58. We note that pushing λ inwards to absorb a single cell of A that maps to $\bar{N}$ reduces the complexity without effecting the class. However, the more global explanation we offer that follows an entire W-ladder within A also covers the case of an isolated 1-cell on A which is traversed twice by λ.

We shall now choose annular subdiagrams A_γ and A_λ with the following properties: A_γ will be a square annular diagram within A between γ and another path γ' of the same length, and A_λ will be the remaining annular diagram, so that A equals the union of A_γ and A_λ along γ'. Moreover, we choose A_λ, A_γ such that the (#(cone-cells), #(squares)) complexity $\mathsf{Area}(A_\lambda)$ is minimal among all possible decomposition choices with γ fixed but other features varying as above.

There is a W-ladder consisting of a sequence of cones and squares in $\bar{N}$ which support the path λ. We choose this sequence so that it is minimal in the sense that it uses as few cones as possible, and then as few squares as possible. Choosing a basepoint the path λ is expressible as a concatenation of subpaths that either travel along connecting square ladders or paths in cones that start and end on 0-cells dual to distinct hyperplanes in the same wall (of the intersection of W and a cone). In the latter case, it is possible for the path to start and/or end on a piece containing a dual 1-cell, but the starting and ending dual 1-cells are dual to distinct hyperplanes in the same wall. This can be remedied by first adding some arcs or "backtracks" to λ.

We now augment A by attaching along λ an "annular" W-ladder L corresponding to the above concatenation of paths. As illustrated in the fourth diagram of Figure 118, let $B = A_\lambda \cup_{\gamma'} L$. We refer the reader to Construction 5.21.

We now form the rectified annular diagram $\bar{B}$ from B as in Section 3.6, and then examine the possible positively curved cells in $\bar{B}$. There is no outerpath of a generalized square in $\bar{B}$ along γ', for then we could push the square across γ' to reduce the area of A_λ (at the expense of increasing $\mathsf{Area}(A_\gamma)$. We refer to Figure 119 for illustrations of this and the next few excluded possibilities. In each case, the possibility is illustrated in a diagram above, and the "reducing action" is illustrated directly below it.

For $\theta \le \pi$, there is no θ-shell with outerpath Q along γ', for otherwise, denoting its boundary by QS, with Q outer and S inner, we would have $\Omega(S) \le \pi$ and hence by Condition 5.55.(1), either $|S| < |Q|$ so γ' and hence $\gamma \to X^*$ is not of minimal length in its homotopy class, or the θ-shell bounded by QS can be replaced by a square diagram, thus reducing the complexity.

There is no outerpath of a generalized square along the part of the boundary of B in $\partial L - \lambda$. Indeed, any such outerpath would necessarily lie along one of the connecting square ladders of L, but then λ can be pushed across this generalized square to reduce $\mathsf{Area}(A_\lambda)$ while remaining on $\bar{N}$.

Generalizing our earlier reduction when γ intersects λ at some point, we observe that: if some cone-cell C of L intersects γ', then we can cut B along C leaving copies C_1, C_2 of C at each end, to form a diagram D whose boundary path is a concatenation $\gamma' c_1 \ell c_2$ where ℓ is a path along $\partial L - \lambda$, and c_1, c_2 are paths on C. This is illustrated in the final pair of diagrams in Figure 119. The arguments above show that D has no positively curved cells except for C_1, C_2, and hence Theorem 3.36 shows that D is a

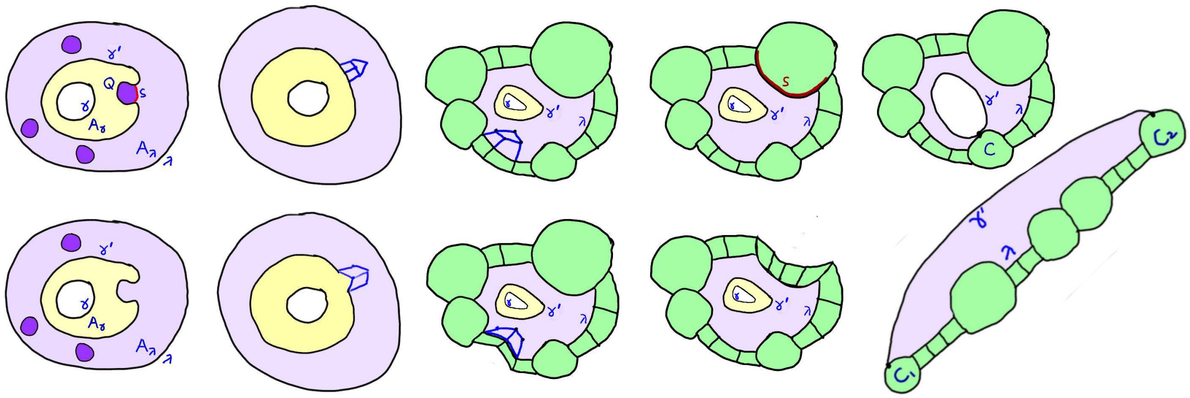

FIGURE 119.

ladder. This completes the proof in this case, as it follows that $\gamma' = \lambda$ lies on $\bar{T} = \langle\gamma\rangle\backslash T(W)$, and hence γ lies in the local convex hull of λ within A.

Now assume that no cone-cell C of L intersects γ', so in particular, each cone-cell in $\bar{B}$ is either internal, or is a θ-shell consisting of a cone-cell not in L that has outerpath along the inside of B on γ', or is a θ-shell consisting of a cone-cell in L whose outerpath lies on the opposite side of L from λ.

By Condition 5.55.(3), each internal cone-cell has negative curvature. Since we have already excluded θ-shells with $\theta \leq \pi$ having outerpath along γ' above, we see that any such θ-shell along γ' would have negative curvature. For each cone-cell in L, its innerpath S passes through 1-cells dual to distinct hyperplanes in the same wall. Thus, by Condition 5.55.(2), if $\Omega(S) \leq \pi$, then the 1-cells are dual to the same hyperplane, and S is homotopic into this hyperplane by a square diagram. We could therefore have used a connecting square ladder in place of this cone-cell, in this case. This is impossible as it violates the minimality of our choice of L. We conclude that $\Omega(S) > \pi$ for each innerpath of cone-cell in L, and so any such cone-cell yields negative curvature.

In summary, there are no outerpaths of generalized squares, each θ-shell has $\theta > \pi$ and is hence negatively curved, and each internal cone-cell is negatively curved. Since $\chi(B) = 0$, we see from Theorem 3.15 that there are no cone-cells in B.

We conclude that either L contains a cone-cell, in which case γ' is forced to lie on (or rather factor through) $\bar{N}$ as above (since each such cone-cell must intersect γ'), and $A = A_\gamma$ is a square diagram with γ lying in the local convex hull of λ. Or, L is a width-1 square annulus consisting of the product of a 1-cube and a subdivided circle. In this case, by Lemma 2.10, B is then a "flat annulus" in the sense of Lemma 2.10. Thus A is a square annulus. □

5.16. **Doubly Collared Annular Diagrams.**

Theorem 5.61 (Doubly Collared Annulus)**.** *Let $\langle X \mid \{Y_i\}\rangle$ satisfy the $B(6)$ conditions of Definitions 5.1, as well as the no inversion condition given in Lemma 5.25. Let $A \to X^*$ be an annular diagram with boundary paths α_1, α_2 that are essential in X^*. Suppose that the lift of its universal cover $\widetilde{A} \to \widetilde{X}^*$ has the property that the induced lifts of $\widetilde{\alpha}_i \to \widetilde{X}^*$ lie on carriers of walls $N_1 = N(W_1)$ and $N_2 = N(W_2)$.*

There exists a new annular diagram B such that B is reduced *in the sense that it has no:*

1. *cancellable square bigons,*
2. *absorbable squares along cone-cells,*
3. *absorbable (cancellable) pairs of cone-cells,*
4. *outerpaths of generalized corners of squares on its boundary,*
5. *cone-cells that are replaceable by square diagrams,*

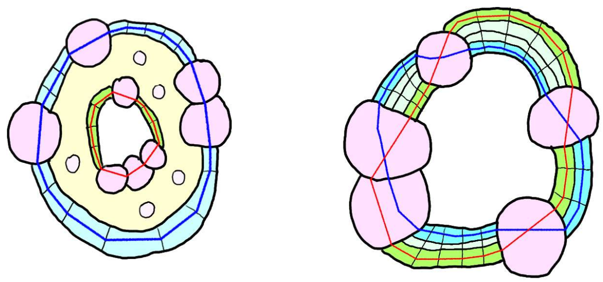

FIGURE 120. B is thick on the left, and B is thin on the right.

(6) *cone-cells on its boundary that are replaceable by square ladders (with no adjustment of internal boundary path).*

B contains a pair of annuladders L_1, L_2 mapping to $\bar{N}_1, \bar{N}_2$ where L_i is a W_i-annuladder. Each $L_i \to B$ is an embedding and B deformation retracts to L_i. For each i, the path β_i is one of the boundary paths of L_i. Then B has one of the following two structures:

(1) *Either B is* thick *in which case L_1, L_2 have disjoint interiors and lie along the boundary of B in the sense that β_1, β_2 are the boundary paths of B.*
(2) *Or B is* thin *and B is itself an annuladder. We emphasize that in this case β_1, β_2 might not equal the boundary paths of B.*

Finally, B lies in the same class *as A in the sense that the lifts of $\widetilde{A}$ and $\widetilde{B}$ have: $\tilde{\alpha}_i, \tilde{\beta}_i$ lie in the same translate of N_i for each i, and α_i, β_i represent the same conjugacy class in each* $\mathrm{Stab}(N_i)$ *(and hence in $\pi_1 X^*$ as well).*

We refer the reader to Figure 120 for illustrations of the thick and thin cases of an annular diagram B together with the annuladders L_1, L_2 inside it.

Proof of Theorem 5.61. **Preliminary Note:** The thick case arises when the translates of W_1 and W_2 under consideration do not cross the same cone or square. Then the diagram B can be chosen to be the union of a minimal (#(Cone-cells), #(Squares)) complexity annular diagram $A' \to X^*$ in the same class as A, together with minimal complexity annuladders $L_i \to \bar{N}_i$ each having a common boundary path with A'.

The thin case arises when W_1 and W_2 cross, or are equal to each other. Note that there is a degenerate case where $L_1 = L_2$ that arises in this situation.

A minimal annular diagram in the class: Let E be an annular diagram in the same class as A with boundary paths ε_i that represent elements conjugate to α_i in $\mathrm{Stab}(N_i)$. Suppose moreover, that the complexity (#(Cone-cells), #(Squares)) of E is minimal among all possible such diagrams.

Properties conferred by minimality: The minimality ensures that E is reduced in the usual sense, of having no bigons, cancellable pairs of cone-cells, absorbable squares, generalized corners of squares on cone-cells, or replaceable cone-cells. However, the minimality ensures there are other properties related to "compressions into the boundary".

If E has a square or cone-cell with a 1-cell along ε_i and this square or cone-cell maps to N_i under the map $\widetilde{E} \to \widetilde{X}^*$, then we could push ε_i through this square or cone-cell and obtain a lower complexity diagram in the same class. Thus no such configuration exists.

If E has a cutpoint, then this cutpoint subtends a subpath of either ε_1 or ε_2 which bounds a disk diagram that is a subdiagram of E. Chopping off this disk diagram, and replacing this subpath by a point reduces the complexity but does not effect the class. We can therefore assume that no such configuration exists.

Accordingly, if e_i is a 1-cell on ε_i whose lift is dual to a 1-cell of W_i then e_i must be an isolated 1-cell of E and the other boundary path ε_j must also pass through e_i (which cannot be a cut-cell as above).

Finally, if ε_i contains the outerpath cd of a corner of a generalized square of E, and this length-2 path lies along the external boundary of a length-2 square W_i-annuladder (more concretely, there is a 1-cell e_i dual to W_i that forms a square with corners ce_i and $e_i^{-1}d$ in $\widetilde{X}^*$) then this generalized square can be pushed out of the diagram, to decrease the complexity while maintaining the class.

The thick case: In case W_1 and W_2 don't cross the same square or cone and aren't equal, then ε_i cannot pass through a 1-cell e_i lifting to a dual 1-cell of W_i. Indeed, as above, such a 1-cell is forced to be an isolated 1-cell of A that also lies on ε_j (here $i \neq j$). A cone or square dual to W_j in N_j that contains e_i on its boundary is crossed by both W_i and W_j.

We shall now assume that no such 1-cells e_1, e_2 are traversed by $\varepsilon_1, \varepsilon_2$.

Minimal W_i-annuladders: For each i, (following a procedure similar to Construction 5.21) we let $L_i \to \bar{N}_i$ be a W_i-annuladder having ε_i as one of its boundary paths, and moreover, assume that L_i is chosen to have minimal (#(Cone-cells), #(Squares)) complexity among all such choices with ε_i fixed.

Forming B in the thick case: We now form the thick annuladder $B = L_1 \cup_{\varepsilon_1} E \cup_{\varepsilon_2} L_2$ by gluing L_1, L_2 along $\varepsilon_1, \varepsilon_2$ to E.

We are assuming here that the lifts $\widetilde{L}_1$ and $\widetilde{L}_2$ do not "cross" the same cone or the same square (possibly even along the same dual 1-cell).

We now verify that B is reduced and has the desired properties. Since E is reduced, we need only consider the interaction between cells in L_1, L_2, and the interaction between cells in E with L_1, L_2. A cancellable pair of cells between L_1, L_2 would lift to the same cell (i.e. square or cone) in $\widetilde{X}$ and hence $\widetilde{X}^*$, and would imply that either W_1, W_2 cross in some square or both cross some cone or are equal. A cancellable pair of cells between L_i, L_i would violate minimality if they are already adjacent in L_i. Minimality would also be violated if we could replace a cone-cell by a square diagram without effecting ε_i. A cancellable pair of cells between L_i, L_i that are not already adjacent would mean that these cells are adjacent in E, and hence ε_i passes through a cutpoint which would violate the minimality of E as above. A cancellable pair between a cell in L_i and a cell in E would imply that there is a cell of N_i within E along the boundary of ε_i, contradicting the minimality of E, as above. Finally, a corner of a generalized square along B would have to come from a square within E, since for any square in L_i, one of the dual curves of this square lies entirely in L_i and doesn't meet ∂B. But, such a corner of a generalized square within E was ruled out by the minimality of E above.

Forming B in the thin case: The case where $\widetilde{L}_1, \widetilde{L}_2$ cross the same cone or square is very closely related to, and could probably be treated by applying Theorem 5.37 but we will give an independent treatment here.

Bouncing 1-cells: The *bouncing 1-cells* of W_i are those 1-cells $b_1^i, b_2^i, b_3^i \ldots$ in ε_i that are dual to W_i in the W_i-annuladder L_i, and actually lie in its interior (note that it is possible for such dual 1-cells to occur on the boundary of L_i along cone-cells). The proof hinges upon reducing the number of bouncing 1-cells, so that the 2-complex $B = L_1 \cup_{\varepsilon_1} E \cup_{\varepsilon_2} L_2$ is in fact an annular diagram since ε_i is a boundary path of L_i. We can nevertheless form the complex B above when there are bouncing 1-cells, and a typical such situation is illustrated in Figure 121.

The *eventful 1-cells* of E are those 1-cells with the property that they are isolated, and either in ε_1 and dual to W_1 or in ε_2 and dual to W_2, or are the location of a cancellable or absorbable pair of 2-cells between L_i, L_j, if we were to form B. We emphasize that a 1-cell may be eventful for more than one reason. As explained above, and as is clear from the cancellable pair possibility, each eventful 1-cell is an isolated 1-cell of E that is traversed by ε_i on one side and ε_j on the other side. There is thus a cyclic ordering of eventful 1-cells in E.

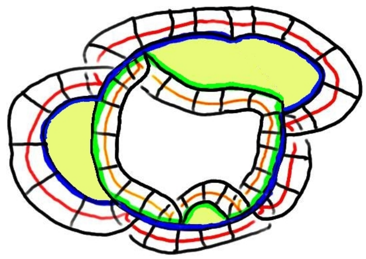

FIGURE 121. The nonplanar 2-complex $B = L_1 \cup_{\varepsilon_1} E \cup_{\varepsilon_2} L_2$ will be adjusted to produce a thin annular diagram.

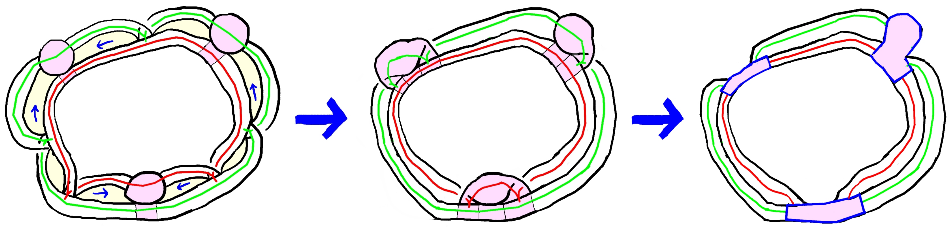

FIGURE 122. Slide bouncing events to the cone-cells and then rechoose cone-cells: We slide the events in the diagram on the left to obtain the diagram on the right. Note that the shaded parts of E are forced to have no squares. The cone-cells in the adjusted diagram are a bit larger and absorb more squares from the other ladder than they did before. We then rechoose the cone-cells to absorb squares where bouncing events occur, and thus obtain a planar B.

Outline of construction of thin annular diagram: As the construction of the thin annulus has complex supporting details, we outline the plan before proceeding. See Figure 122.

◎1 **Configurations that violate minimality of L_1, L_2:** We will show below that our minimality conditions on E, L_1, L_2 restricts the nature of the sequence of eventful 1-cells as follows:

(1) There does not exist a cancellable pair of squares at a non-bouncing eventful 1-cell unless $W_1 = W_2$.
(2) There does not exist a cancellable pair of squares at a bouncing eventful 1-cell.
(3) If a, b, c are consecutive bouncing eventful 1-cells and b is dual to W_i and b lies on a square in L_j then either a or c has a cone-cell on the L_i side, and the event is absorption of a square of L_j into this L_i cone-cell.

◎2 **Sliding events to cone-cells:** Each bouncing eventful 1-cell of L_i that lies on a square is either preceded or succeeded (or both) by a cone-event. We now use this to "slide" the bouncing eventful 1-cells so that they all occur on a cone-cell instead of a square. This involves rechoosing E, L_1, L_2.

◎3 **Rechoosing cone-cells:** After absorbing squares and cone-cells at the absorption-events, we obtain a complex B' that has all its events on cone-cells, and in particular, has all its bouncing events on cone-cells. The parts of B' between consecutive cone-cells are collections of rectangles forming a contiguous or non-contiguous cone-piece. We remedy the nonplanarity of B by rechoosing the cone-cells to build a new complex B that is the desired (planar) thin annular diagram.

◎4 **"Zipping" when $W_1 = W_2$:** A simple argument shows that L_1, L_2 can be "zipped together" and that E is a subdivided circle when there is an absorption of squares or cone-cells that implies that $W_1 = W_2$.

We now describe the argument. An important first step is to restrict the types of events.

There is no cancellable pair of squares (unless $W_1 = W_2$): The minimality of E, and our hypothesis that we are considering consecutive bouncing 1-cells will imply there cannot be a cancellable pair of squares in B between L_1, L_2. Indeed, let s be such a square in L_i intersecting L_j in a bouncing 1-cell,

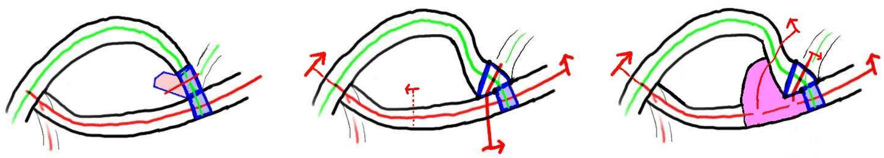

FIGURE 123. The left diagram shows that a cancellable pair of squares between L_1, L_2 cannot occur except at a bounce as illustrated in the middle. In this case, the orientations shows that the square-bounce is consecutive with a cone-event as on the right.

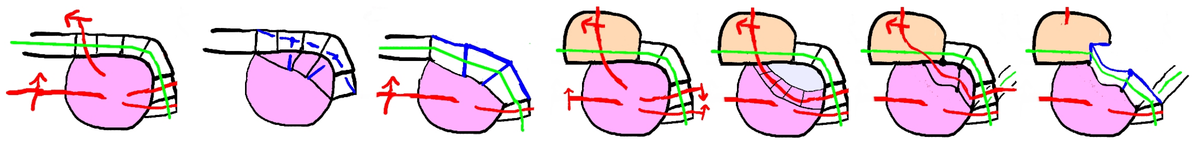

FIGURE 124. Replacing the ladder L along C with a shorter square ladder L'.

and suppose that the 2-cell in L_j across from this 1-cell is s' with s and s' canceling across it. As in Figure 123, we could either absorb a cell within E into L_j as illustrated on the left (thus contradicting minimality of E) or s is incident with L_j at another 1-cell as illustrated on the right. In this case, we see that L_j has an earlier bouncing 1-cell, and moreover (keeping track of the illustrated orientations) using 2-sidedness of $\bar{W}_j$ in $\bar{N}_j$ there must have been an even earlier one besides this. Note that the cell prior to s' in L_j must be a cone-cell for this to occur.

With a bit more analysis, we now show that this bouncing square-square cancellable pair preceded by the cone-cell C is impossible as it violates our minimality assumptions. As in the third diagram in Figure 123, the wall W_i enters L_j at the top of the C. Consider the ladder L in L_j between its square s_j that cancels with a square s_i in L_i and the first cell in L_j that either W_i enters or is a square that absorbs into C or a cone-cell that combines with C. See the fourth diagram in Figure 123.

In the case where W_i enters L_j or a square absorbs into C (and both of these are essentially cancellable pair situations), then we note that L is entirely formed from squares, since a cone-cell could be compressed out of L_i thus reducing its complexity. Consequently, we find that s_j can be absorbed into C.

The other case to consider is where there is no cancellable pair, and W_i simply enters L_j.

Firstly consider the case where it enters at a square. Now, L is again a square ladder, since as above, otherwise a cone-cell can be compressed out of L_i. As the bottom external boundary path of L forms a single wall-piece with C, and this piece starts and ends on 1-cells dual to the same wall, applying Conditions 5.1.(4) and 5.1.(5), we find that the path in C compresses onto the W_2 wall in C. We are thus able to adjust L_i by replacing L with a ladder L' having two fewer squares. See the first three diagrams in Figure 124.

Secondly consider the case where it enters at a cone-cell. Now we apply Conditions 5.1.(5) and find that there is a corner of a generalized square in L_i on C, and we use this to decrease the length of L_i as in the last four diagrams in Figure 124. This is very similar to the arguments used in the 2nd and 3rd cases of Figure 128.

Consecutive events: We now show that for (maximal) consecutive events at 1-cells a, b, after absorbing cells one obtains a diagram that is either: a ladder having exactly two cone-cells at the two events; or there is a cone-cell at one of these events and a bouncing square at the other and the reduced subtended diagram looks like a generalized corner of square on a cone-cell; or there are two bouncing squares and the subtended diagram looks like a bigonal square diagram; or a single cone-cell - arising from a ladder that bounces twice on a cone-cell and absorbs into it - see Figure 127. We will later restrict these options further using our minimality assumptions.

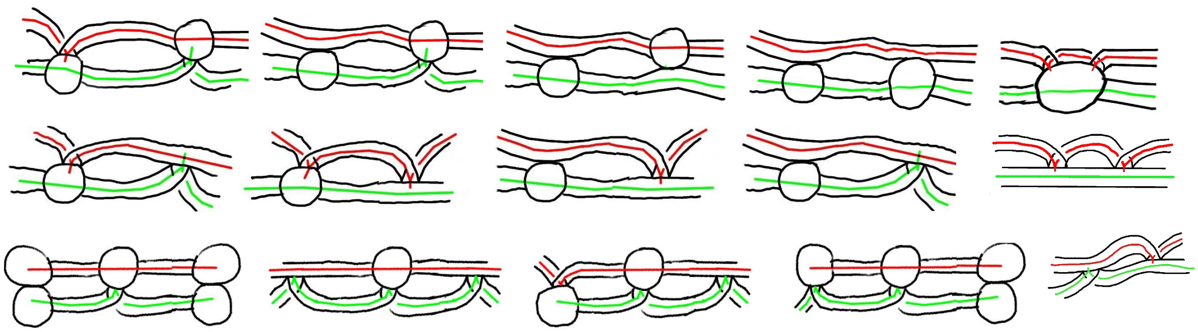

FIGURE 125. Some of these consecutive events cannot arise in a minimal situation.

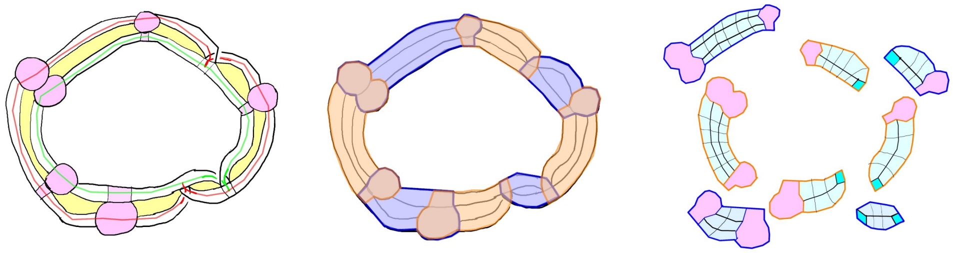

FIGURE 126. Each subdiagram of B starting and ending at an event (after absorptions) is either a single cone-cell, a ladder with exactly two cone-cells, a corner of generalized square on a cone-cell, or a bigonal square diagram.

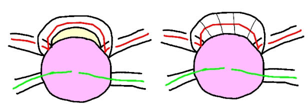

FIGURE 127. It is possible for L_i to bounce twice in a row on a cone-cell C, but the minimality of E implies that there are no squares between, and the minimality of L_j implies that this part of L_i absorbs into C.

The key point is that the subdiagram subtended by consecutive events is already nearly reduced, and has very little possibility for positive curvature. The parts within each of L_1, L_2, E are reduced, and the parts between L_i and E are reduced, so the only possible reduction is between L_1, L_2 - at the two ends of the subdiagram. When two cone-cells are combined there is no further reduction, (except that such an absorption between cone-cells could provide a new cone-cell that is replaceable by a square diagram - but we can make this replacement at the end of the entire process) and when a cone-cell in L_i absorbs a square in L_j, it is a priori possible for some sequence of further absorptions of a sequence of squares of L_j to absorb into this "prolonged" cone-cell. A subsequent absorption of a cell in E with this cone-cell remains impossible, since it could have been absorbed to begin with, thus decreasing the complexity of E. We note that the degenerate case where everything (a square ladder) on one side is absorbed into the cone-cell is one of our possible conclusions.

The only global features that could prevent the diagram from being reduced are: a generalized corner of a square - necessarily from a square at one corner to a cone-cell at the other; or a bigonal square diagram - between squares at the two corners.

Having obtained a reduced diagram, the collars from L_1, L_2 ensure that there are only two possible places sources of positive curvature, and so Theorem 3.36 ensures that we obtain a ladder between two cone-cells.

Restricting further: Bouncing 1-cells on the same side or an alternating triple: A consecutive pair of square events contradicts the minimality if they both bounce from L_i to L_j. We refer the reader

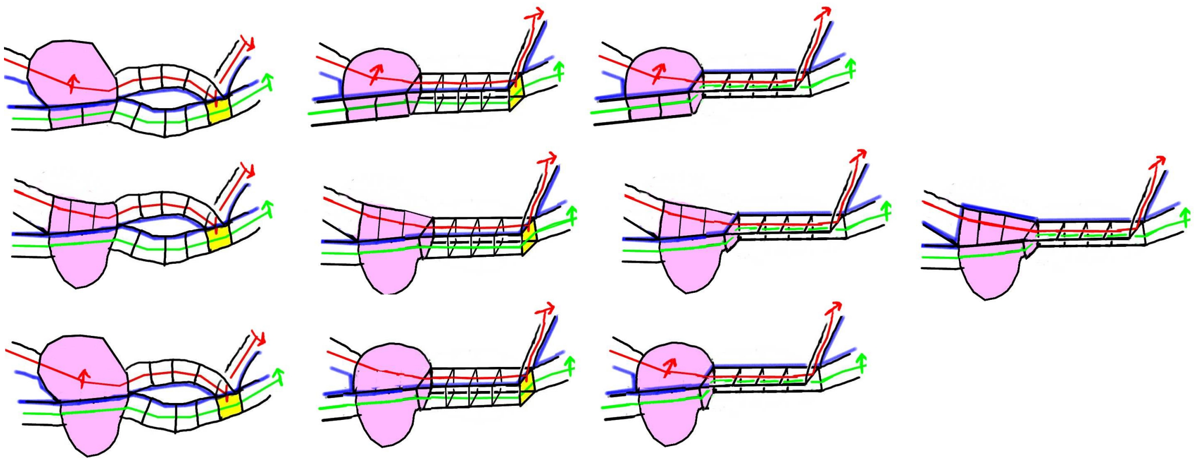

FIGURE 128. If a square bounce is consecutive with a cone absorption event, then we slide it over to bounce on the cone. The 2nd and 3rd cases lead to a decrease in the L_1, L_2 complexity and thus do not arise.

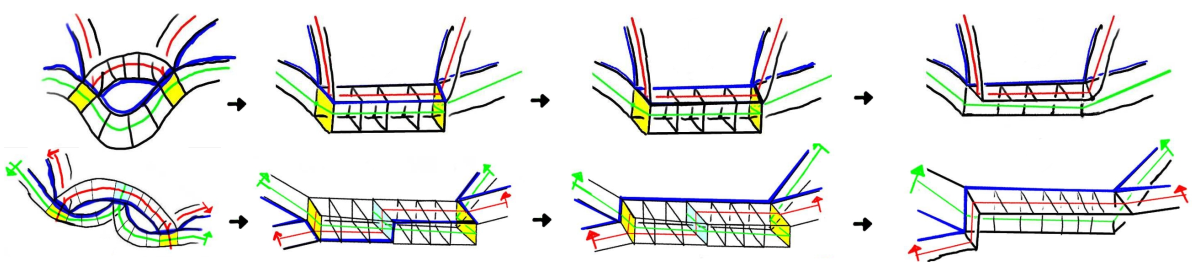

FIGURE 129. Two consecutive square events on the same side, or three alternating square events allow the reduction of the L_1, L_2 complexity.

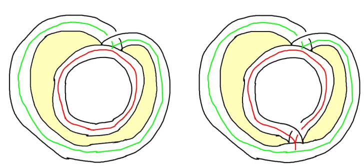

FIGURE 130. One or two (alternate) square bounces cannot be all, for otherwise a carrier would be 1-sided.

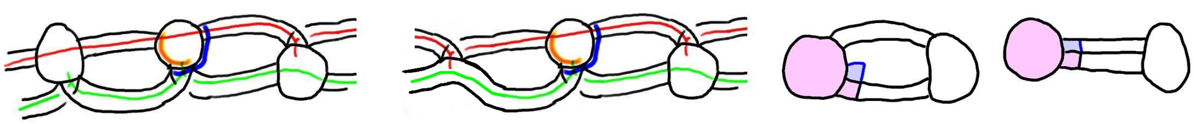

FIGURE 131. On the left, the innerpath σ of the middle cone-cell, is the concatenation of two pieces so $\Omega(\sigma) < \pi$, and hence the cone-cell can be replaced by a square ladder. A similar argument holds for the second diagram. The pair of diagrams on the right, illustrate that a cone-cell cannot be prolonged by absorbing squares from the other ladder.

to the top of Figure 129. Moreover, if there is a bounce from L_j to L_i, followed by L_i to L_j, followed by L_j to L_i again, one also yields a reduction as on the bottom of Figure 129. As illustrated in Figure 130, we use that $\bar{W}_i$ is 2-sided in $\bar{N}_i$ to see that a single square bounce or two alternate square bounces are not everything.

Sliding bouncing 1-cells to cones: For each i, we are able to slide the bouncing 1-cells of L_i so they occur on a cone-cell of L_i. We refer the reader to Figure 128. We note that by assuming first that we are in a situation with a minimal total length of L_i, L_j, we can assume that the bottom sequence can never

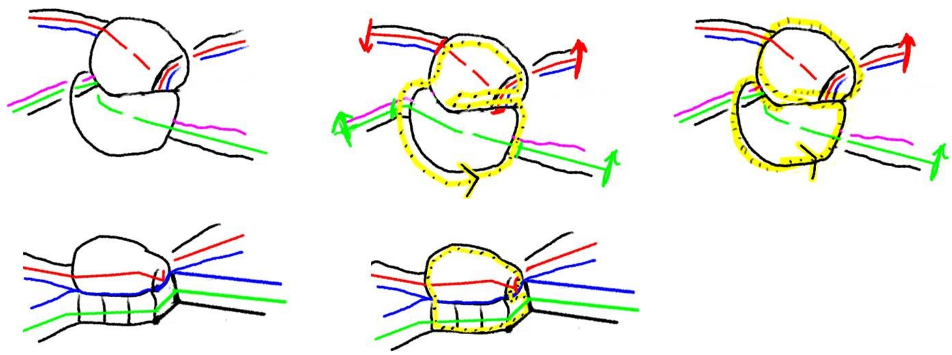

FIGURE 132. We rechoose the cone-cells so that square ladders end on disjoint arcs. There can be many ways of doing this, as illustrated on the top two.

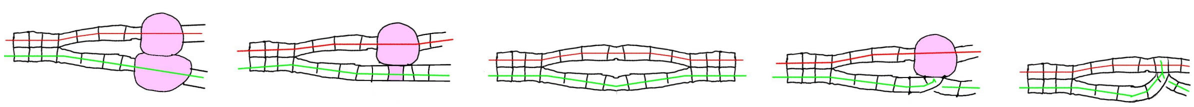

FIGURE 133. If $W_1 = W_2$ are identified because of a cancellable pair, then we consider the next event to prove that L_1, L_2 do not diverge. Not that the 4th and 5th diagrams are impossible.

occur, since it reduces one of the lengths, since ultimately a square can be absorbed into a cone. The top sequence is typical. The middle sequence can be assumed not to occur as we did for the bottom sequence. Note that the last two diagrams in the middle sequence are interchangeable (but differ on ε_i), indicating that bouncing 1-cells can be passed across a rectangle in L_i that is absorbable into a cone-cell in L_j.

Choosing new cone-cells: Having performed the above moves, we can now assume that each event is either an absorbable pair of cone-cells between L_i, L_j, or a square (or sequence of squares) in L_i that is absorbable into L_j (or vice-versa).

Moreover, the diagram between successive events is a (grid) rectangle.

We now show how to choose a cone-cell for each such event, such that the rectangular diagram on its left and right can be attached to it in a way that allows us to create an annular diagram. We refer the reader to Figure 132.

When W_1, W_2 fold into each other: We now consider the case where there is a cancellable pair of squares between L_1, L_2 along a 1-cell that is not dual to either W_1, W_2. A similar situation arises when there is a cancellable pair of squares or cone-cells along a 1-cell dual to both W_1, W_2. These situations are very special as they immediately imply that $W_1 = W_2$, and have global consequences as we construct the annulus.

Consider the subdiagram starting at such a cancellable pair and ending at the next event. See Figure 133. A priori, we do not know that (as illustrated) the W_i-ladders from this cancellable pair to the next event consist entirely of squares, however we shall conclude that this is the case. We note that the final case, indicating that W_i, W_j cross in a square is impossible by Theorem 5.20. The second to last case will also be impossible since we will show that the two 1-cells on the terminal cone-cell are adjacent, and they are also adjacent or in the same piece on the other side. It follows that the cone-cell compresses which violates our minimal choice of L_i, L_j.

There is some initial sequence of cancellable pairs of cells, and we consider the first two that do not cancel. For instance, in Figure 133 there might be three such consecutive cells. We consider the diagram obtained by chopping off the sequence of cancellable pairs, and identifying the boundary along the last pair of dual 1-cells of W_i, W_j, as in Figure 134. If the next event is a similar cancellable pair of

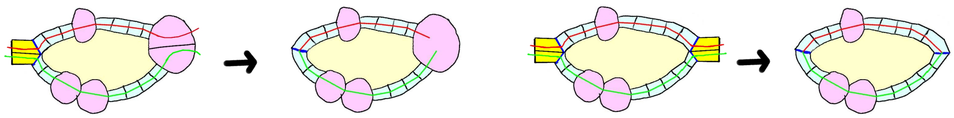

FIGURE 134. The parts where L_1, L_2 diverge would yield impossible collared diagrams after identifying 1-cells or absorbing on each end.

squares, then we perform the same operation on that side. If it is a cone event, then we simply absorb into one cone, and let it be the final cell of our diagram.

The resulting diagram is reduced but has no positive curvature. It is thus impossible by Theorem 3.15, or in fact Theorem 3.40.

We reach the conclusion that the two hyperplanes in the cone-cell of the next event, are not only in the same wall, but in the same hyperplane, as in fact, the dual 1-cells are adjacent and fold into each other. See Remark 5.11.

This conclusion passes from cone-cell to cone-cell as we go around the diagram. Indeed, at each stage, successive pairs of dual 1-cells are dual to the same wall in the cone (from the previous stage), but lie in the same piece (from the next stage) and thus lie in the same hyperplane. We can thus apply the argument above to choose appropriate cone-cells. Moreover, if desired, we can simplify each cone-cell so that the W_i-ladders occupy the same position all the way around, and the result is a single annuladder that is both a W_i-annuladder and a W_j-annuladder. □

Corollary 5.62. *Suppose $B(8)$-Conditions 5.55.(2) and 5.55.(3) are added to the hypotheses of Theorem 5.61. Then either B is thin, or B has no cone-cells and is thus a square annular diagram $B \to X$.*

Proof. Suppose B is a thick ladder containing W-annuladders L_1 and L_2 on its inside and outside, such that L_1, L_2 have disjoint interiors. Condition 5.55.(3) implies that any internal cone-cell is negatively curved and Condition 5.55.(2) implies that any cone-cell with a single boundary arc is a negatively curved θ-shell (since it lies on L_1 or L_2). These would be impossible by Theorem 3.15, and so these cone-cells are replaceable. Thus in the thick case, B must be a square diagram. □

Remark 5.63 (Elliptical degeneracy)**.** There is an elliptical annuladder case to consider in the proof (or at least the conclusion) of Theorem 5.61. This arises when the L_1 or L_2 complexity is either $(1, 0)$ or $(0, m)$ and moreover the initial and terminal dual 1-cells of the annuladder do not form a piece where they are joined up in the $(1, 0)$ case (and something similar in the $(0, m)$ case). In this case, one or two of L_1, L_2 is an elliptical annuladder. Then L_1, L_2 are forced to merge within the annuladder B into a single elliptical annuladder following Lemma 5.53.

When only one is elliptical, say L_2, then we need a bit more work - and we must assume that internal cone-cells have negative curvature. We let $B_1 = L_1 \cup_{\varepsilon_1} E$ be the subannulus of B that meets L_2 along ε_2. Observe that B_1 has nonpositive curvature along L_1. Any generalized corner of square along ε_2 could be absorbed into the elliptical annulus L_2, and likewise any θ-shell with $\theta \leq 0$ would be replaceable by a square diagram or absorbable into L_2 since otherwise it would yield an internal cone-cell. Since there is no positive curvature along L_1, and no nonnegatively curved cone-cell along ε_2 nor any corner of generalized square on ε_2, we see that E is a product square annulus. If we assume cones are superconvex, then if L_2 is nontrivial, then E would absorb into L_2, so we can assume E is a subdivided circle. If L_1 has no cone-cells, then it too would likewise be absorbed into L_2. So let us assume it has at least one cone-cell C. In the split-angle case Condition 5.1.(5) suffices to show that C is replaceable or absorbable into L_2. In general, Condition 5.55.(2) is sufficient.

The above shows that the thick case is ruled out when one of L_1, L_2 is elliptic. The thin case is already dealt with in Lemma 5.53.

This is dealt with in Lemma 5.54 as well.

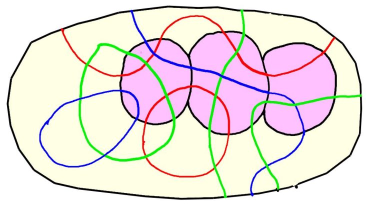

FIGURE 135. To verify Condition (3). Each wall in cones of $\langle X \mid Y_i \rangle$ maps to a single equivalence class of hyperplanes in X, but a single equivalence class has connected intersection with each piece in each cone.

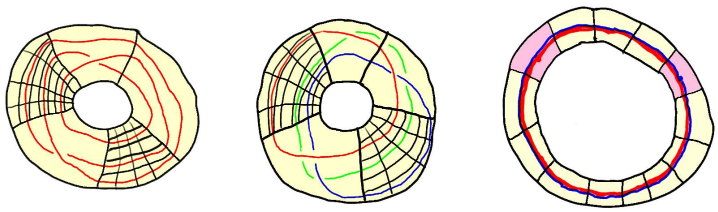

FIGURE 136.

5.17. Malnormality of wall stabilizers.

Theorem 5.64. *Let $\langle X \mid Y_i \rangle$ be a cubical presentation satisfying the conditions below. Then* $\text{Stab}(W)$ *is almost malnormal in $\pi_1 X^*$ for each wall W.*

1. *$B(6)$ and $B(8)$ conditions of Definitions 5.1 and 5.55, and the no inversion condition of Lemma 5.25.*
2. *$\{\pi_1 H_1, \ldots, \pi_1 H_k\}$ is a malnormal collection of subgroups of $\pi_1 X$ when $\{H_1, \ldots, H_k\}$ are distinct immersed hyperplanes of X that are images of hyperplanes in the same wall W of $\widetilde{X}^*$. (See Definition 12.2.)*
3. *For each $Y \in \{Y_i\}$ and hyperplanes $h_1 \neq h_2 \subset Y$, and lifts $\widetilde{h}_1, \widetilde{h}_2$ contained in a lift $\widetilde{Y}$ of Y to $\widetilde{X}^*$: if $\widetilde{h}_1, \widetilde{h}_2$ lie in hyperplanes H_1, H_2 of $\widetilde{X}^*$ that belong to the same wall of W, then h_1, h_2 cannot lie in the same cone-piece of Y.*

Remark 5.65. Condition (3) is difficult to check in practice, but is instead verified through the following scenario which represents a stronger condition. The hyperplanes of X are partitioned into equivalence classes. For each i, the hyperplanes of each wall of Y_i map to hyperplanes in X that are in the same equivalence class. The union of all hyperplanes of X in an equivalence class have connected intersection with (or more generally, preimage in) each piece of each Y_i. This is illustrated heuristically in Figure 135.

For instance, we will later use this to verify Condition (3) in a situation where all hyperplanes of X embed, each wall of Y maps to the same hyperplane in X, and the injectivity radius of hyperplanes of X exceeds the diameter of the largest cone-piece or wall-piece in any Y.

Remark 5.66. Condition (3) is equivalent to the following: For any cone Y in $\widetilde{X}^*$, if W_1, W_2 are walls in $\widetilde{X}^*$ intersecting Y in hyperplanes h_1, h_2: If $W_1 = gW_2$ for some $g \in \pi_1 X^*$, then h_1, h_2 cannot lie in the same cone-piece of Y.

Condition (3) implies that the dual curves illustrated in the first diagram of Figure 136 actually lie within the same hyperplane in each of the cones (receiving the cone-cells). If this is not the case for a pair of dual curves in the second diagram of Figure 136, then that pair of dual curves does not have lifts contained in the same wall of $\widetilde{X}^*$.

Definition 5.67 (Annuladder)**.** We use the term *annuladder* for an annulus with the structure of a ladder as in Definition 3.35 but we now include the possibility of a square ladder. We use the term

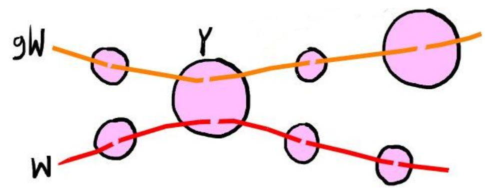

FIGURE 137. If $gW \cap Y$ and $W \cap Y$ diverge (by not passing through the same piece) then $N(W) \cap N(gW) = Y$.

W-annuladder to mean that its universal cover maps to $N(W)$ and there is a W dual curve within it that generates π_1.

Proof of Theorem 5.64. Let $N = N(W)$ be the carrier of the wall. Let B be the diagram provided by Theorem 5.61. By Corollary 5.62, B is either thin or it has no cone-cells and is a square diagram.

In the latter case, B represents a homotopy in X between closed curves in distinct immersed hyperplanes $\bar{H}_1, \bar{H}_2$ of X where H_1, H_2 are components of their preimages in $\widetilde{X}^*$ that lie in the same wall W. Condition (2) on malnormality implies that $\bar{H}_1 = \bar{H}_2$ and that B can be homotoped into the carrier of $\bar{H}_1$ (relative to β_1, β_2). In particular, B is homotopic into $\bar{N} = \text{Stab}(W)\backslash N$ relative to β_1, β_2 and we find that β_1, β_2 are conjugate in Stab(W) and hence the same holds for α_1, α_2. Note that the above argument deals with the case that B contains no cone-cell (whether it is thick or thin).

We now examine the case where B is thin. The case where B contains no cone-cell is dealt with above. Now B contains a pair of annular W-ladders L_1, L_2 passing through each cone-cell of B and traveling around B so that each generates $\pi_1 B$. We now use Condition 5.64.(3) to see that the dual curves of L_1, L_2 lie in the same hyperplane in each cone (at entrance and exit). We refer the reader to Figure 138. This gives us a homotopy between β_1 and β_2 within $\bar{N}$.

Alternate argument in the hyperbolic case (our intended application): When H is a quasiconvex subgroup of a word-hyperbolic group, almost malnormality of H is equivalent to: $H^g \cap H$ contains an infinite order element if and only if $g \in H$. Thus, in our case of greatest interest, when X^* has compact cones, and $\widetilde{X}^*$ is δ-hyperbolic and N is quasi-isometrically embedded, it suffices to verify the above condition. When the cones of X^* are finite, an infinite order element cannot be represented by an elliptic annulus. There are then pieces between consecutive cone-cells in B, and it is this which allows us to see that $B \to X^*$ factors as $B \to \bar{N} \to X^*$. (Or alternately, that $\widetilde{B}$ lifts to N.)

When one or both of the annuladders L_1, L_2 in B are elliptical annuli, then Lemma 5.54 and Lemma 5.53 imply that $L_1 = B = L_2$ is a single elliptical annulus. (Moreover, we can assume that it is built from a cone-cell and not a square ladder, as the latter case was dispensed with above.) Now, however, there is no cone-piece between consecutive cones - and so we are no longer able to deduce from Condition (3) that $W_1 = W_2$.

Instead we argue as follows: The elliptic annulus B has $\widetilde{B}$ lifting to a cone Y in $\widetilde{X}^*$ where $Y \subset (N \cap gN)$. The subgroup $H \cap gHg^{-1}$ equals Stab$(N) \cap$ Stab(gN), which is a subgroup of Stab$(N \cap gN)$, and this is a useful statement since $N \cap gN \neq \emptyset$.

There are then two cases: Either $N \cap gN = N$, in which case $g \in$ Stab$(N) =$ Stab(W) so $g \in$ Stab(W). Or $N \cap gN = Y$, in which case Stab$(N) \cap$ Stab$(gN) \subset$ Stab(Y). The former case arises when 1-cells dual to $(Y \cap W)$ lie in the same piece as 1-cells dual to $(Y \cap ghW)$. The latter arises when they are in different pieces, for then $N \cap gN = Y$ by Corollary 5.39. We use here that $(Y \cap gW)$ doesn't cross or osculate with $(Y \cap W)$ in Y, and we use that a piece of Y does not contain 1-cells e, f such that e is dual to $(Y \cap W)$ and f is dual to $(Y \cap gW)$ unless $(Y \cap W) = (Y \cap gW)$. □

Lemma 5.68. *If W is 2-sided in N and its hyperplanes map to a collection of embedded disjoint hyperplanes in X then* Stab$(N) =$ Stab(W).

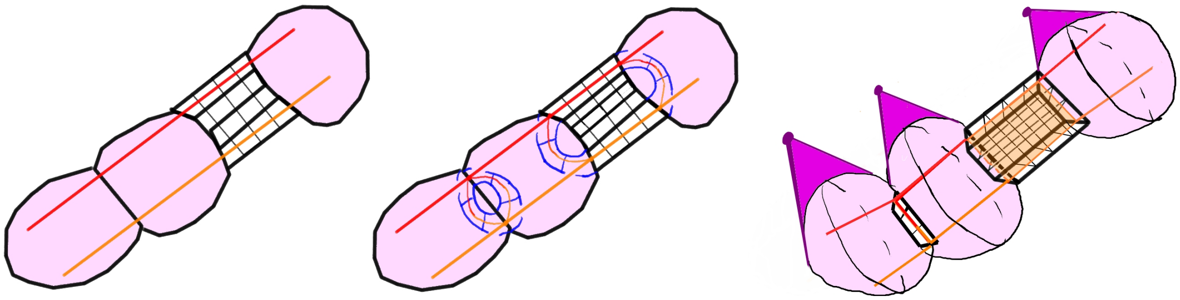

FIGURE 138. Since the annulus B is thin, Condition 5.64 implies that the two annuladders travel through the same immersed wall $\bar{W} \to X^*$, and hence $B \to X^*$ factors through the immersed carrier $\bar{N} \to X^*$ consisting of cones and carriers of hyperplanes.

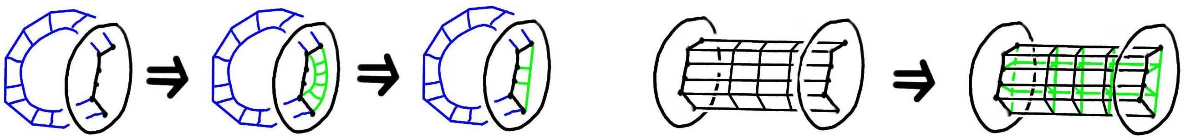

FIGURE 139.

Proof. It is obvious that $\text{Stab}(W) \subset \text{Stab}(N)$. To see that $\text{Stab}(N) \subset \text{Stab}(W)$ we suppose that $gN = N$ and need to show that $gW = W$. If there are no cones in N then we use that $\bar{W}$ is a single embedded hyperplane of $\bar{N}$ to see that $\pi_1 X$ automorphisms of W are in one-to-one correspondence with $\pi_1 X$ automorphisms of N. If there is a cone Y in N, then we can translate by an element $h \in \text{Stab}(W)$, so that $hgY = Y$. It then suffices to verify that $hg \in \text{Stab}(W)$. But $hg(W \cap Y)$ and $W \cap Y$ are translates of the same wall in Y that lie in the same piece, and are thus the same wall by Condition (3), so $hgW = W$. □

Remark 5.69 (Concrete Version)**.** Consider a (possibly degenerate) rectangular square diagram between two cone-cells corresponding to a piece between the corresponding cones. Suppose the hyperplanes dual to a side S of the rectangle lying along the cone Y are all distinct, except for the first and last 1-cells that are dual to the same hyperplane H. Thus, the intermediate path joining these 1-cells is a local geodesic (and hence a geodesic in the CAT(0) piece).

Then S lies along $N(h)$ where $h = H \cap Y$. We refer the reader to Figure 139.

More generally, we should be able to slice off parts of the rectangle whenever there are duplicated hyperplanes in S, until we reach the situation above. Alternately, we might be able to choose our cone-cells so that S has the desired properties a priori.

6. Special Cube Complexes

In [HW08] we defined "special cube complexes" and examined some of their properties. In Section 6.2 we review the definition of special cube complexes in terms of illegal hyperplane pathologies, and we state the characterization in terms of local isometries to the cube complex of a right-angled Artin group. In Section 6.4 we review the definition of canonical completion and retraction. The hyperplane pathology definition of special cube complex arose from our desire to define canonical completion and retraction above dimension one. Subsequently, we realized that this was equivalent to a local isometry to the cube complex of a right-angled Artin group. Many other aspects of special cube complexes are explored in [HW08], [HWa] and [HW10] including various conditions which imply that a cube complex is virtually special. The material in Section 6.5 is new.

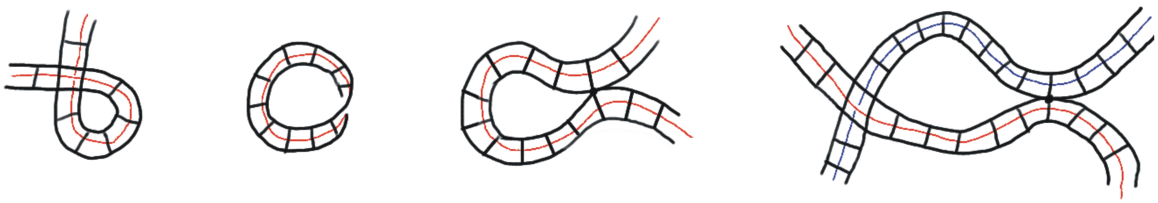

FIGURE 140. Immersed Hyperplane Pathologies

6.1. **Hyperplanes.** A *midcube* of the n-cube $[-1,1]^n$ is the subspace obtained by restricting exactly one of the coordinates to 0. A *hyperplane* Y in the CAT(0) cube complex C, is a connected subspace whose intersection with each cube is either a midcube or is empty. The 1-cubes intersected by Y are *dual* to Y. For a CAT(0) cube complex, there exists a hyperplane dual to each 1-cube, and moreover, hyperplanes are themselves CAT(0) cube complexes with respect to the cell structure induced by intersection, and are convex subspaces in the CAT(0) metric [Sag95].

We now define an immersed hyperplane in an arbitrary cube complex C. Let M denote the disjoint union of the collection of midcubes of cubes of C. Let D denote the quotient space of M induced by identifying faces of midcubes under the inclusion map. The connected components of D are the *immersed hyperplanes* of C.

6.2. **Hyperplane Definition of Special Cube Complex.** We shall define a special cube complex as a cube complex which does not have certain pathologies related to its immersed hyperplanes.

An immersed hyperplane D *crosses itself* if it contains two different midcubes from the same cube of C.

An immersed hyperplane D is *2-sided* if the map $D \to C$ extends to a map $D \times I \to C$ which is a combinatorial map of cube complexes.

A 1-cube of C is *dual* to D if its midcube is a 0-cube of D. When D is 2-sided, it is possible to consistently orient its dual 1-cubes so that any two dual 1-cubes lying (opposite each other) in the same 2-cube are oriented in the same direction.

An immersed 2-sided hyperplane D *self-osculates* if for one of the two choices of induced orientations on its dual 1-cells, some 0-cube v of C is the initial 0-cube of two distinct dual 1-cells of D.

A pair of distinct immersed hyperplanes D, E *cross* if they contain midcubes lying in the same cube of C. We say D, E osculate, if they have dual 1-cubes which contain a common 0-cube, but do not lie in a common 2-cube. Finally, a pair of distinct immersed hyperplanes D, E *inter-osculate* if they both cross and *osculate*, meaning that they have dual 1-cubes which share a 0-cube but do not lie in a common 2-cube.

A cube complex is *special* if all the following hold:

(1) No immersed hyperplane crosses itself
(2) Each immersed hyperplane is 2-sided
(3) No immersed hyperplane self-osculates
(4) No two immersed hyperplanes inter-osculate

Example 6.1. Any graph is a 1-dimensional cube complex that is special.

Any CAT(0) cube complex is special.

The cube complex associated to a right-angled Artin group is special.

6.3. **Right-angled Artin group characterization.** A special cube complex is automatically nonpositively curved. In fact, we give the following characterization of special cube complexes in [HW08]:

Proposition 6.2. *A cube complex is special if and only if it admits a combinatorial local isometry to the cube complex of a right angled Artin group.*

A quick explanation of Proposition 6.2 is that for a local isometry $B \to C$, the prohibited hyperplane pathologies on C induce the same prohibited pathologies in B. On the other hand, if C is special, then we define a graph Γ whose vertices are the immersed hyperplanes of C, and whose edges correspond to intersecting hyperplanes. Then there is a natural map $C \to C(\Gamma)$ which is a local isometry.

6.4. **Canonical completion and retraction.** We refer to [HW08, HWa] for more details about the following fundamental property of special cube complexes:

Proposition 6.3 (Canonical completion and retraction)**.** *Let $f : Y \to X$ be a local isometry from a compact nonpositively curved cube complex to a special cube complex. There exists a finite degree covering space $\mathsf{C}(Y \to X) \to X$ called the* canonical completion *of $f : Y \to X$ such that $f : Y \to X$ lifts to an embedding $\widehat{f} : Y \to \mathsf{C}(Y \to X) \to X$, and there is a retraction map $\mathsf{C}(Y \to X) \to Y$ called the* canonical retraction*.*

It follows that $\pi_1 Y$ is separable in $\pi_1 X$ since virtual retracts of residually finite groups are separable. As explained in [HW08], when X is virtually special compact and word-hyperbolic, any quasiconvex subgroup $H \subset \pi_1 X$ is separable. Indeed, H can be represented by a compact based local isometry $Y \to X$ by Proposition 8.2. Thus H is a virtual retract. The analogous separability results hold in the sparse case as discussed in Theorem 16.23.

6.5. **Extensions of quasiconvex codimension-1 subgroups.**

Definition 6.4 (K-partitions and K-walls in G)**.** Let G be a finitely generated group with Cayley graph $\Gamma(G, S)$ and let K be a subgroup of G. A *coarse K-partition* of G is a collection of subsets $\{G_1, \ldots, G_m\}$ with $G = G_1 \cup \cdots \cup G_m$, that is *$K$-stable* so for each $k \in K$ the translate kG_i coarsely equals some G_j in the sense that the symmetric difference $kG_i \Delta G_j$ lies in KC for some compact C. We shall assume that each pair $G_i \neq G_j$ is *K-coarsely r-separated* in the sense that for each $r \geq 0$ we have $N_r(G_i) \cap N_r(G_j)$ lies in KC where C is some compact subset of $\gamma(G, S)$. With these last properties in mind, we will regard two partitions as *K-equivalent* if their differences lie in KC for some C. We shall also assume that there is some $r \geq 0$ so that each G_i is *coarsely r-connected* in the sense that G_i lies in a single component of $N_r(G_i)$ in $\Gamma(G, S)$.

While our arguments work in general, we have in mind the case of a *K-wall* which is a K-partition consisting of two subsets, and especially the situation where K is a codimension-1 subgroup, and each of these sets is *K-deep* in the sense that it doesn't lie in $N_r(K)$ for any r. For instance one set might be a K-deep component of $\Gamma(G, S) - N_r(K)$, and the other set equals its complement.

Parts of $\{G_1, \ldots, G_m\}$ that are not already K-deep, could be absorbed into other parts without affecting the intention of the situation, so one normally assumes that each part is K-deep.

We note that if G' is a finite index subgroup of G, then we obtain a K-partition $\{G'_1, \ldots, G'_m\}$ of G' by setting $G'_i = G_i \cap G$ for each i, and if each part of G is K-deep, then so is each part of G'. Conversely, if $\{G'_1, \ldots, G'_m\}$ is a K-partition of G' with K-deep parts, then (up to K-equivalence) it is associated to a unique K-partition of G. Indeed, we let $G_i = N_r(G'_i) \cap G$ where r is chosen so that $G \subset N_r(G')$. Note that each G_i is K-deep since G'_i is K-deep, and note that K-stability holds since if kG'_i is coarsely the same as G'_j then kG_i is coarsely the same as kG_j.

Definition 6.5 (Extension of K-partitions)**.** Let H be a subgroup of G, and let K be a subgroup of H, and let $H_1, \ldots, H_m$ be a coarse K-partition of H. We say that $H_1, \ldots, H_m$ *extends* to a coarse K'-partition $\{G_1, \ldots, G_m\}$ of G if K' is a subgroup of G with $K' \cap H = K$, and $\{H \cap G_1, \ldots, H \cap G_m\}$ is K-equivalent to $\{H_1, \ldots, H_m\}$.

We say that the subgroup $H \subset G$ has the *extension property for K-partitions* if each K-partition of H extends to a K'-partition of G.

The goal of this section is to prove the following:

Theorem 6.6 (Quasiconvex Extension Property)**.** *Let G be the fundamental group of a compact special cube complex X. Let H be a subgroup represented by a compact based local isometry $Y \to X$. Let K be a subgroup represented by a compact based local isometry $Z \to Y$. Then any K-partition of H extends to a K'-partition of G such that K' is represented by a compact local isometry $Z' \to X$.*

Remark 6.7. It is conceivable that Theorem 6.6 holds under the weaker assumption that H is a virtual retract of G, but some work would be needed to produce and ensure that K' is quasiconvex.

We will prove Theorem 6.6 focusing on the case where the K-partition of H has no shallow parts.

Proof. Let $\{H_1, \ldots, H_m\}$ be a K-partition of H. Let $\widetilde{z}$ be a lift of the basepoint z of Z to $\widetilde{Y}$. Choose r so that distinct $H_i\widetilde{z}$ do not intersect the same component of $\widetilde{Y} - N_r(\widetilde{Z})$. (The intuition here is that each $H_i\widetilde{z}$ lies in the connected subspace $\cup_{\{S\ :\ S\cap H_i\widetilde{z}\neq\emptyset\}} S \cup N_r(\widetilde{Z})$ where S varies over the components of $\widetilde{Y}-N_r(\widetilde{Z})$. We would insist that a shallow $H_i\widetilde{z}$ lie within $N_r(\widetilde{Z})$.)

By Proposition 8.3, $N_r(\widetilde{Z})$ lies in a K-cocompact convex subcomplex $\widetilde{Z}_1 \subset \widetilde{Y}$ and let $Z_1 = K\backslash\widetilde{Z}_1$.

Pass to a finite based cover $\widehat{Y} \to Y$ such that $Z_1 \to Y$ lifts to an embedding in $\widehat{Y}$. For instance, we could let $\widehat{Y} = \mathsf{C}(Z_1 \to Y)$. Let $\widehat{H} = \pi_1\widehat{Y}$, and assign the K-partition $\{\widehat{H}_1, \ldots, \widehat{H}_m\}$ where $\widehat{H}_i = \widehat{H} \cap H_i$ for each i.

Consider the canonical completion $\mathsf{C}(\widehat{Y} \to X)$ and the canonical retraction $\mathsf{C}(\widehat{Y} \to X) \to \widehat{Y}$. It induces an equivariant retraction map $\widetilde{X} \to \widetilde{Y}$ that fixes $\widetilde{Y}$.

Let Z_1^+ denote the base-component of the preimage of Z_1 in $\mathsf{C}(\widehat{Y} \to X)$, and let $K' = \pi_1 Z_1^+$. Note that Z_1^+ is locally convex by Lemma 6.9.

Observe that $K' \cap \widehat{H} = K$ by Remark 6.10 (with $\widehat{Y}$ playing the role of Y).

Let $\widetilde{Z}_1^+$ denote the base component of the preimage of Z_1^+ in $\widetilde{X}$. We consider the subsets $\widetilde{X}_1, \ldots, \widetilde{X}_m$ of $\widetilde{X}-\widetilde{Z}_1^+$ that retract to the distinct coarsely connected subsets of $\widetilde{Y}-\widetilde{Z}$ corresponding to $\{\widehat{H}_1\widetilde{z}, \ldots, \widehat{H}_m\widetilde{z}\}$.

Let $J \subset G$ be the subset with $J\widetilde{z} \subset \widetilde{Z}_1^+$. Partition J into $J_1 \cup \cdots \cup J_m$ such that $H_i \cap J \subset J_i$ (there are many ways to do this, but they are all coarsely the same). Now define G_i so that it contains J_i and all elements g_i with $g_i\widetilde{z} \subset \widetilde{X}_i$.

This yields a K'-partition of G that extends the K-partition of $\widehat{H}$. We have thus reached our goal of extending our original K-partition of H, as the K partition of $\widehat{H}$ extends to a unique K partition of H.

The construction extends deep sets of the K-partition to deep sets of the K'-partition. This is because the combinatorial retraction map is distance non-increasing, but fixes the subcomplex $\widetilde{Y}$. Therefore, if some part G_i lies in $N_s(K')$ for some s, then its image H_i lies in $N_s(K)$ and is thus not deep. □

The wall case of the following corollary is an important ingredient in the hypotheses of Proposition 7.8.

Corollary 6.8. *Let G be a word-hyperbolic group with a finite index subgroup $\bar{G}$ that is the fundamental group of a compact special cube complex. Let H be a quasiconvex subgroup of G and let K be a quasiconvex subgroup of H that also lies in $\bar{G}$. Then any K-partition of H extends to a K'-partition of G such that K' is quasiconvex.*

Proof. Let $\bar{H} = \bar{G} \cap H$, and note that the K-partition $H_1 \sqcup \cdots \sqcup H_m$ of H induces a K-partition $\bar{H}_1 \sqcup \cdots \sqcup \bar{H}_m$ of $\bar{H}$ where $\bar{H}_i = \bar{H} \cap H_i$.

Theorem 6.6 provides a K'-partition $\bar{G}_1 \sqcup \cdots \sqcup \bar{G}_m$ of $\bar{G}$ extending the K-partition of $\bar{H}$. Finite neighborhoods $G_i = N_s(\bar{G}_i)$ provide a K'-partition of G itself that extends the K-partition of H. □

Lemma 6.9 (Locally Convex Canonical Preimage)**.** *Let X be special. Let $Y \to X$ be a local isometry with Y compact. Let $Z \subset Y$ be a locally convex subspace.*

Consider the canonical completion and its canonical retraction map $\phi : \mathsf{C}(Y \to X) \to Y$. Let $Z^+ = \phi^{-1}(Z)$. Then Z^+ is a locally convex subspace of $\mathsf{C}(Y \to X)$.

Before proceeding with the proof, we note that, after subdividing some of the cubes in $\mathsf{C}(Y \to X)$ that are not in Y, the map ϕ is a *cubical map* - it maps cubes to cubes (of possibly lower dimension) by maps modeled on $I^{n+m} \to I^n$ that collapse dimensions. This subdivision is denoted $\mathsf{C}_{\boxminus}(Y, X)$ and discussed in [HWa].

Proof. Suppose two 1-cubes e_1, e_2 are adjacent along a 0-cube v with $e_1, e_2, v \in Z^+$. Suppose e_1, e_2 form the corner of a 2-cube c at v in $\mathsf{C}(Y \to X)$. If $\phi(e_1), \phi(e_2)$ form the corner of a 2-cube at $\phi(v)$ then it must be $\phi(c)$, and moreover $\phi(c)$ lies in Z by local convexity of Z in Y, and so c is in Z^+. The alternative is that one of $\phi(e_1)$ or $\phi(e_2)$ collapses to $\phi(v)$, or perhaps that both collapse to $\phi(v)$. In the former case $\phi(c)$ must collapse to $\phi(e_2)$ or $\phi(e_1)$, and in the latter case $\phi(c)$ collapses to $\phi(v)$. But either way, we again see that $\phi(c) \subset Z$ so c is in Z^+. □

Remark 6.10. Choose a basepoint in $Z \subset Z^+ \subset \mathsf{C}(Y \to X)$, then: $\pi_1 Z^+ \cap \pi_1 Y = \pi_1 Z$.
Indeed $(\pi_1 Z^+ \cap \pi_1 Y) \supset \pi_1(Z^+ \cap Y) = \pi_1 Z$.
And $(\pi_1 Z^+ \cap \pi_1 Y) = \phi(\pi_1 Z^+ \cap \pi_1 Y) \subset \phi(\pi_1 Z^+) = \pi_1 Z$.

The following result presumes familiarity with sparse cube complexes as treated in Definition 16.4.

Theorem 6.11. *Let X be a sparse special nonpositively curved cube complex. The quasiconvex extension property holds for quasi-isometrically embedded subgroups of $\pi_1 X$ whose intersections with parabolic subgroups have a compact core. In particular, it holds when the subgroup is quasi-isometrically embedded and has either trivial or finite index intersection with each parabolic subgroup.*

Proof. By Proposition 8.2 there is a compact geometric representative $Y \to X$. Now the proof follows as before. □

Remark 6.12. When G is locally-quasiconvex and virtually sparse special, every K-wall in H extends to a quasiconvex K'-wall in G. The difficulties in the proof of Theorem 6.11 are merely in ensuring the quasiconvexity of the subgroup K' induced by retracting.

Remark 6.13 (No new parabolic slopes)**.** Let P be a parabolic subgroup of $\pi_1 X$, and observe that the hyperplanes of $\widetilde{X}$ determine a wall structure on P.

Let K' be a new wall of $\pi_1 X$ arising from a K-wall in a quasiconvex subgroup $H \subset \pi_1 X$ as produced by Theorem 6.6.

Then K' does not contribute essentially new walls on P in the following sense: Let $\widetilde{F}$ be the cubical convex hull of some orbit $P\widetilde{x}$. Then, each new wall in $\widetilde{X}$ either cuts $\widetilde{F}$ parallel to some hyperplane of $\widetilde{X}$, or $\widetilde{F}$ lies in one (or both) halfspaces associated to the new wall.

Algebraically, we have $P \cap (K')^g$ is either finite or is equal to P, or is commensurable to $P \cap \mathrm{Stab}(D)$ where D is some hyperplane in $\widetilde{X}$.

The reason for this is as follows: By the construction given in Theorem 6.6, the wall is represented by a compact local isometry $Z' \to X$, and hence by a convex combinatorial subcomplex $\widetilde{Z'} \subset \widetilde{X}$. It must therefore intersect $\widetilde{F}$ in a wall that is a convex combinatorial subcomplex, and is thus limited by the original combinatorially available possibilities.

Masters gave a pretty argument showing that every finitely generated free subgroup of a 3-dimensional closed right-angled hyperbolic Coxeter group lies in a quasifuchsian surface subgroup [Mas08]. We can use the material developed here to give a variation on his result:

Theorem 6.14. *Let X be a compact special cube complex with $\pi_1 X$ word-hyperbolic. Let H be a quasiconvex subgroup of $\pi_1 X$ that is not of finite index. Then H lies in a codimension-1 quasiconvex subgroup of $\pi_1 X$.*

Proof. By Proposition 8.2, let $Y \to X$ be a compact local-isometry with $\pi_1 Y = H$. Since H is quasiconvex and has infinite index in $\pi_1 X$ we can choose an element a such that $\langle H, a\rangle \cong H * \mathbb{Z}$. (We could likewise choose a, b to obtain $H * F_2$ and produce an amalgamated product over H below.) Let $K = \langle H, a\rangle$. Let $Z \to X$ be a compact core for K so $\pi_1 Z = K$, and note that we can assume that $Y \subset Z$ (indeed, we can rechoose the core for H in K at this stage). Now extend the codimension-1 (splitting) subgroup $H \subset K$ to a codimension-1 (splitting) subgroup $H' \subset G'$ of a finite index subgroup G' of G and hence to a codimension-1 subgroup of G itself. Note that H' is just π_1 of the based preimage of Y under the canonical retraction of $\mathsf{C}(Z \to X) \to Z$. □

Problem 6.15. Does Masters' theorem hold for quasiconvex free subgroups of higher dimensional right-angled Coxeter groups?

Is every quasiconvex free subgroup of a word-hyperbolic compact special group contained in a surface subgroup?

Can one recover his theorem for special hyperbolic 3-manifold groups using our method?

Problem 6.16. Let G be virtually sparse special. Does G have the extension property with respect to any quasiconvex H-wall in an arbitrary quasiconvex K subgroup of G? In particular, does this hold when K is a parabolic subgroup?

6.6. **The Malnormal Combination Theorem.**

In [HWa] we prove the following:

Proposition 6.17. *Let Q be a compact nonpositively curved cube complex with an embedded 2-sided hyperplane H. Suppose that $\pi_1 Q$ is word-hyperbolic. Suppose that $\pi_1 H \subset \pi_1 Q$ is malnormal. Let $N_o(H)$ denote the open cubical neighborhood of H. Suppose that each component of $Q - N_o(H)$ is virtually special. Then Q is virtually special.*

Here is a more general formulation that permits torsion: Let G act properly and cocompact on a CAT(0) cube complex $\widetilde{Q}$. Suppose there is a 2-sided hyperplane $\widetilde{H}$ such that:

(1) $\mathrm{Stab}(H)$ *is almost malnormal in G,*
(2) *$g\widetilde{H} \cap \widetilde{H} = \emptyset$ for each $g \in G - \mathrm{Stab}(\widetilde{H})$.*
(3) $\mathrm{Stab}(H)$ *preserves each component of $N_o(\widetilde{H}) - \widetilde{H}$*
(4) *For each component $\widetilde{X}$ of $\widetilde{Q} - GN_o(\widetilde{H})$, the group* $\mathrm{Stab}(\widetilde{X})$ *has a finite index torsion-free subgroup J such that $J\backslash\widetilde{X}$ is special.*

We note that Condition (3) holds in a finite index subgroup so long as Condition (2) holds.

7. Cubulations

7.1. **Wallspaces.**

Definition 7.1. Haglund and Paulin introduced the notion of a wallspace to abstract a property that arises in many natural scenarios, and especially for Coxeter groups [HP98]. A *wallspace* is a set X together with a collection of *walls* each of which is a partition $X = \overleftarrow{N} \sqcup \overrightarrow{N}$ into *halfspaces*, and such that moreover, $\#(p, q) < \infty$ for each $p, q \in X$ where $\#(p, q)$ equals the number of walls separating p, q.

The fundamental example of a wallspace is the 0-skeleton of a CAT(0) cube complex, together with a system of walls associated to the hyperplanes.

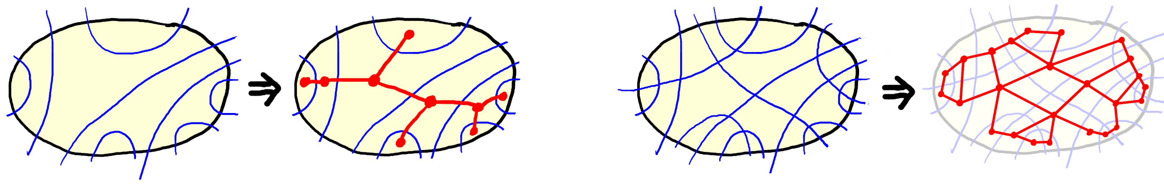

FIGURE 141. The dual cube complex on the left is a tree, the one on the right is 2-dimensional. The n-cubes correspond to certain n-fold collections of pairwise crossing walls.

7.2. **Sageev's construction.**

Definition 7.2. Let G be a finitely generated group with Cayley graph Γ. A subgroup $H \subset G$ is *codimension-1* if it has a finite neighborhood $N_r(H)$ such that $\Gamma - N_r(H)$ contains at least two H-orbits of components that are *deep* in the sense that they do not lie in any $N_s(H)$.

For instance any $\mathbb{Z}^n$ subgroup of $\mathbb{Z}^{n+1}$ is codimension-1, and any infinite cyclic subgroup of a closed surface subgroup is as well. We note that an edge group of a nontrivial splitting is codimension-1. We also note that if the coset diagram $H\backslash\Gamma$ has more than one topological end, then H is codimension-1. There is a closely related notion: H is *divisive* if $\Gamma - N_r(H)$ has two or more deep components. Every codimension-1 subgroup is divisive, however there are divisive subgroups that are not codimension-1 The difficulty is that the action of H on Γ might permute the deep components of $\Gamma - N_r(H)$. When $\Gamma - N_r(H)$ has finitely many deep components, there is a finite index subgroup $H' \subset H$ whose action stabilizes each of these components, and one obtains a multi-ended coset diagram $H'\backslash\Gamma$, which is equivalent to H' being codimension-1 in G.

Given a finite collection of codimension-1 subgroups $H_1, \dots, H_k$ of G, Michah Sageev introduced a simple but powerful construction that yields an action of G on a CAT(0) cube complex C that is *dual* to a system of walls associated to these subgroups [Sag95].

For each i, let $N_i = N_{r_i}(H_i)$ be a neighborhood of H_i that separates Γ into at least two deep components. The *wall* associated to N_i is a fixed partition $\{\overleftarrow{N}_i, \overrightarrow{N}_i\}$ consisting of one of these deep components $\overleftarrow{N}_i$ together with its complement $\overrightarrow{N}_i = \Gamma - \overleftarrow{N}_i$, and more generally, the translated *wall* associated to gN_i is the partition $\{g\overleftarrow{N}_i, g\overrightarrow{N}_i\}$. The two parts of the wall are *halfspaces*.

We presume a certain degree of familiarity with the details of Sageev's construction here, but hope that any interested reader will mostly be able to follow the arguments. We shall not describe the structure of the dual cube complex C here but will describe its 1-skeleton. A 0-cube of C is a choice of one halfspace from each wall such that each element of G lies in all but finitely many of these chosen halfspaces. A wall is thought of as *facing* the points in its chosen halfspace. Two 0-cubes are joined by a 1-cube precisely when their choices differ on exactly one wall. See Figure 141 for two particularly simple dual cube complexes.

The walls in Γ are in one-to-one correspondence with the hyperplanes of the CAT(0) cube complex C given by Sageev's construction, and the stabilizer of each such hyperplane equals the codimension-1 subgroup that stabilizes the associated translated wall: The stabilizer of the hyperplane corresponding to a translated wall associated to gN_i is commensurable with gH_ig^{-1}.

Sageev's construction naturally decomposes into two separate ideas. The first is that a collection of codimension-1 subgroups yields a wallspace. The second is that a wallspace yields a dual cube complex (see [CN05, Nic04] for more details on the latter).

7.3. **Finiteness properties of the dual cube complex.**

Cocompactness properties of the action of G on the CAT(0) cube complex dual to a wallspace associated to a collection of codimension-1 subgroups was analyzed in [Sag97] where Sageev proved that:

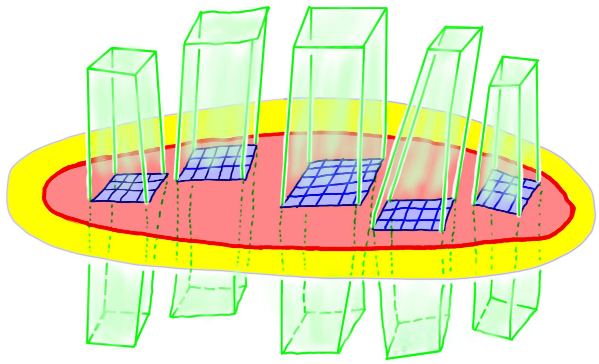

FIGURE 142. Heuristic picture of a cosparse cubulation.

Proposition 7.3. *Let G be a word-hyperbolic group, and $H_1, \ldots, H_k$ be a collection of quasiconvex codimension-1 subgroups. Then the action of G on the dual cube complex is cocompact.*

We refer to [HWb] for a more elaborate discussion of the finiteness properties of the action obtained from Sageev's construction, as well as for background and an account of the literature.

In parallel to Proposition 7.3, is a properness criterion which we state as follows (see for instance [HWb]). We use the notation $\#(p, q)$ for the number of walls separating p, q.

Proposition 7.4. *If $\#(1, g) \to \infty$ as $d_\Gamma(1, g) \to \infty$ then G acts properly on C.*

The main theorem in [HWb] generalizes Proposition 7.3 to a relatively hyperbolic context as follows:

Proposition 7.5. *Let G be a f.g. group that is hyperbolic relative to a collection of parabolic subgroups $P_1, \ldots, P_s$. Let $H_1, \ldots, H_k$ be a collection of quasi-isometrically embedded codimension-1 subgroups of G. Let C denote the CAT(0) cube complex dual to the G-translates of $W_1, \ldots, W_k$. For each i, let C_i denote the CAT(0) cube complex dual to the walls in P_i corresponding to the nontrivial walls obtained from intersections with translates of the W_i, and note that C_i embeds in C as a convex subcomplex. Then:*

(1) *there exists a compact subcomplex K such that $C = GK \cup_{i=1}^{s} GC_i$, and*
(2) *$g_iC_i \cap g_jC_j \subset GK$ unless $i = j$ and $g_j^{-1}g_i \in \mathrm{Stab}(C_i)$.*

Remark 7.6. We may further assume that GK is connected. Indeed, since G is finitely generated, and C is connected, one can add a collection of paths S_i to K such that each S_i starts at the basepoint in K and ends at the translate of this basepoint by the i-th generator of G.

We are most interested in the following corollary which is the source of cosparse CAT(0) cube complexes (see Definition 16.4). The first explicit appearance of such cosparse CAT(0) cube complexes was in the CAT(0) cube complexes associated to $B(6)$ small-cancellation groups [Wis04] (they were called "cofinite" there). However the notion is certainly a general phenomenon associated to relative hyperbolicity.

Corollary 7.7. *Let G be hyperbolic relative to virtually free-abelian subgroups. Let W_i be a finite collection of quasi-isometrically embedded codimension-1 subgroups (or rather walls in $\Gamma(G)$). Then G acts cosparsely on the CAT(0) cube complex associated to this system of walls.*

7.4. **Cubulating Amalgams.** We now summarize the results from [HWc]. As the main statement is rather technical, we first give an imprecise simplification that disregards certain facilitating conditions.

Quasi-Theorem. *Let G split as $A *_C B$ or $A*_{C^t=C'}$. And suppose that G is hyperbolic relative to virtually abelian subgroups, and that C is malnormal and quasiconvex in G. If A, B act properly and* *cocompactly on CAT(0) cube complexes then G acts properly on a CAT(0) cube complex dual to a system of walls associated to quasi-isometrically embedded subgroups.*

A subgroup $H \subset G$ of a relatively hyperbolic group is *aparabolic* if it has finite intersection with each (noncyclic) parabolic subgroup of G.

A slightly simplified version of the main result from [HWc] is:

Proposition 7.8. *If G has all the following properties then G acts properly on a CAT(0) cube complex, and the stabilizers of the hyperplanes are quasi-isometrically embedded.*

(1) *G is hyperbolic relative to f.g. virtually abelian subgroups.*
(2) *G splits as an amalgamated product $G \cong A *_C B$.*
(3) *A and B are fundamental groups of compact (or more generally, sparse) nonpositively curved cube complexes (or more generally, they act properly on CAT(0) cube complexes with corresponding quotient).*
(4) *The two embeddings C_+, C_- of C are relatively quasiconvex in their vertex groups.*
(5) *the embeddings $C_- \subset A$ and $C_+ \subset B$ are almost malnormal.*
(6) *the embeddings $C_- \subset A$ and $C_+ \subset B$ are aparabolic.*
(7) *C has separable quasiconvex subgroups.*
(8) *There are quasiconvex subgroups $H_1, \dots, H_r$ of C and H_i-walls in C such that C acts properly on the resulting cube complex.*
(9) *Each H_i-wall of C, extends to an H_i^A-wall of A and an H_i^B-wall of B, where H_i^A, H_i^B are quasiconvex subgroups of A, B.*

Alternately, we can assume the following slightly more flexible possibility:

($5'$) and ($6'$) *C_+ is almost malnormal in B and C_+ is aparabolic in B.*
($8'$) and ($9'$) *There is a system of walls for A so that A acts properly and cosparsely on the resulting cube complex such that there are induced H_i-walls for C, and each such H_i-wall extends to an H_i^B-wall in B.*

In the HNN case we have the following adjusted statements:

($\bar{2}$) *G splits as $A*_C$.*
($\bar{4}$) *C is quasi-isometrically embedded in G.*
($\bar{5}$) *$\{C_+, C_-\}$ are an almost malnormal pair of subgroups of A.*
($\bar{6}$) *$\{C_+, C_-\}$ are aparabolic in A.*
($\bar{8}$) *There are quasiconvex subgroups $H_1, \dots, H_r$ and H_i-walls of C, so that C acts properly on the resulting cube complex.*
($\bar{9}$) *Each H_i-wall of C_+ and each H_i-wall of C_- extends to an H_i^A-wall.*

A more general version requires that only C_+ be malnormal and aparabolic at the expense of assuming there is a cubulation of A with only unpaired excess walls at C_+. Specifically, the walls induced by $C_- \subset A$ are precisely the walls of C, and the walls induced by $C_+ \subset A$ are the walls of C together with some additional excess. We will illustrate this in a specific case in the proof of Theorem 17.6 (and not use it elsewhere).

8. Local-convexity and cores

Definition 8.1. A combinatorial map $\phi : Y \to X$ between nonpositively curved cube complexes is a *local-isometry* provided that for each pair of 1-cubes e_1, e_2 of Y with initial vertex y, mapping to a pair $\phi(e_1), \phi(e_2)$ of 1-cubes with initial vertex $\phi(y)$, if $\phi(e_1), \phi(e_2)$ bound a corner of a square at $\phi(y)$ then e_1, e_2 bound the corner of a square at y.

If the inclusion map $\phi : Y \hookrightarrow X$ of a subspace is also a local-isometry then we say that Y is a *locally-convex* subcomplex.

It is not hard to verify that a local-isometry $\phi : Y \to X$ lifts to an isometric embedding $\widetilde{Y} \to \widetilde{X}$ between universal covers, and is hence π_1-injective.

We proved in [SW] (see also [Hag08]) that:

Proposition 8.2. *Let X be a compact nonpositively curved cube complex with $\pi_1 X$ word-hyperbolic. Let $H \subset \pi_1 X$ be a quasiconvex subgroup, and let $C \subset \widetilde{X}$ be a compact subspace. Then there exists an H-cocompact subspace $\widetilde{Y} \subset \widetilde{X}$ containing C. We refer to $\widetilde{Y}$ as a* locally convex core *containing C. Likewise, when C is a compact subspace of $H\backslash\widetilde{X}$, we term $H\backslash\widetilde{Y} \subset H\backslash\widetilde{X}$ a* locally convex core *containing C.*

The analogous result holds for G acting cosparsely on a sparse CAT(0) cube complex $\widetilde{X}$ with G hyperbolic relative to abelian subgroups: If H is a quasi-isometrically embedded subgroup then there exists $\widetilde{Y} \subset \widetilde{X}$ such that Y is H-cosparse.

Moreover, if H intersects each free-abelian subgroup of G in either a finite index or trivial subgroup, then we can choose Y to be H-cocompact.

One can also prove the following (see [HW08] for a strong form of this).

Proposition 8.3. *Let $\widetilde{Y} \subset \widetilde{X}$ be a convex subcomplex of a CAT(0) cube complex. For each $a \geq 0$ there exists $b \geq 0$ and a convex subcomplex $\widetilde{Y}^+$ such that $N_a(\widetilde{Y}) \subset \widetilde{Y}^+ \subset N_b(\widetilde{Y})$. Moreover, assuming $\widetilde{X}$ has a proper group action, we can assume $\widetilde{Y}^+$ is stabilized by the stabilizer of $\widetilde{Y}$.*

8.1. **Superconvexity.**

Definition 8.4. Let X be a metric space. A subset $Y \subset X$ is *superconvex* if for any bi-infinite geodesic γ, if $\gamma \subset N_r(Y)$ for some $r > 0$, then $\gamma \subset Y$. A map $Y \to X$ is *superconvex* if the map $\widetilde{Y} \to \widetilde{X}$ is an embedding onto a superconvex subspace.

Lemma 8.5. *Let H be a quasiconvex subgroup of a word-hyperbolic group G. And suppose that G acts properly and cocompactly on a CAT(0) cube complex X. For each compact subcomplex $D \subset X$ there exists a superconvex H-cocompact subcomplex $K \subset X$ such that $D \subset K$.*

Proof. It follows from δ-hyperbolicity that any infinite geodesic lying in $N_r(H\widetilde{x})$ actually lies in $N_{2\delta}(H\widetilde{x})$. Now apply Proposition 8.2 to obtain a convex cocompact core Y containing $N_r(H\widetilde{x})$. (We remind the reader that we use the combinatorial metric.)

To see that Y is superconvex, observe that any geodesic in a finite neighborhood of Y, is actually contained in a 2δ neighborhood of $H\widetilde{x}$, and hence in Y. □

Lemma 8.6. *If $Y \subset X$ is superconvex and cocompact, then there is a non-negative function $f : \mathbb{R} \to \mathbb{R}$ such that for any length $f(r)$ geodesic segment σ that lies in $N_r(Y)$, the midpoint of σ lies in Y.*

Consequently, for any geodesic segment σ lying in $N_r(Y)$, if σ' is obtained by removing the initial and terminal subsegments of length $\frac{f(r)}{2}$ then $\sigma' \subset Y$.

8.2. **Fiber Products.**

Definition 8.7 (fiber product)**.** Given a pair of combinatorial maps $A \to X$ and $B \to X$ between cube complexes, we define their *fiber product* $A \otimes_X B$ to be a cube complex, whose i-cubes are pairs of i-cubes in A, B that map to the same i-cube in X. There is a commutative diagram:

$$\begin{array}{ccc} A \otimes_X B & \to & B \\ \downarrow & & \downarrow \\ A & \to & X \end{array}$$

Note that $A \otimes_X B$ is the complex in $A \times B$ that is the preimage of the diagonal $D \subset X \times X$ under the map $A \times B \to X \times X$. Recall that D has a natural structure as a cube complex since for any cube Q, the diagonal of Q^2 is isomorphic to Q by either of the projections.

Our description of $A \otimes_X B$ as a subspace of a cartesian product endows the fiber product with the property of being a universal receiver in the following sense: Consider a commutative diagram as below. Then there is an induced map $C \to A \otimes_X B$ such that all the obvious compositions commute.

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

When the target space X is understood, we will simply write $A \otimes B$. When A, B, and X have basepoints and the maps are basepoint preserving, then $A \otimes B$ has a basepoint, and it is often natural to consider only the base component.

When $A \to X$ and $B \to X$ are covering maps, then so is $A \otimes B \to X$. Moreover, let $(a, b) \in A \times B$ reflect choices of the preimage of the basepoint $x \in X$, then the component of $A \otimes B$ containing (a, b) is the covering space of X corresesponding to $\pi_1(A, a) \cap \pi_1(B, b)$.

We will now generalize this:

Lemma 8.8. *Suppose that X is a nonpositively curved cube complex, and that $A \to X$ and $B \to X$ are connected and locally convex, and that A is super-convex. Then the [noncontractible] components of $A \otimes B$ correspond precisely to the [nontrivial] intersections of conjugates of $\pi_1(A, a)$ and $\pi_1(B, b)$ in $\pi_1 X$.*

Proof. Let X_A and X_B denote the based covers of X corresponding to $\pi_1(A, a)$ and $\pi_1(B, b)$, and note that there are locally convex embeddings $A \subset X_A$ and $B \subset X_B$.

Suppose that for some $\alpha, \beta \in \pi_1 X$ the conjugates $\pi_1 A^\alpha \cap \pi_1 B^\beta = \pi_1 X_A^\alpha \cap \pi_1 X_B^\beta$ have nontrivial intersection. □

Remark 8.9. When $A \to X$ and $B \to X$ are superconvex, then bi-infinite local geodesic pieces between A and B are precisely the same as bi-infinite local geodesics in $A \otimes B$.

Lemma 8.10. *If A and B are superconvex then so is each component of $A \otimes B$.*

When A is connected and superconvex, then $\pi_1 A$ is malnormal if and only if $A \otimes A$ consists of a diagonal component (that is an isomorphic copy of A) together with various contractible components. More generally, when each component A_i of A is superconvex, the malnormality of the collection of conjugacy class representatives $\pi_1 A_i \subset \pi_1 X$ corresponds to contractibility of all nondiagonal components of $A \otimes A$.

Remark 8.11 (Small-cancellation and Superconvexity)**.** When $\langle X \mid Y_i \rangle$ satisfies a small-cancellation condition asserting that wall-pieces have bounded length, then $\widetilde{Y}_i \to \widetilde{X}$ is superconvex.

Indeed, consider a geodesic γ in $\widetilde{X}$. Suppose that γ lies in $N_r(\widetilde{Y})$ for some $r > 0$, but that $\gamma \not\subset \widetilde{Y}$.

We can find an arbitrarily large subpath γ' that is disjoint from $\widetilde{Y}$. We will show that γ' is parallel to an impossibly long piece of $\widetilde{Y}$.

Let D be a minimal area diagram between γ' and $\widetilde{Y}$ in the sense that the boundary path of D equals $\sigma_1 \gamma' \sigma_2 \lambda$ where σ_i are paths from γ' to the endpoints of a path λ in $\widetilde{Y}$, and these three remaining paths can vary among all such paths and D with these properties.

We can ignore arcs of bounded length consisting of spurs at the corners of D (just redefine γ, λ etc.)

Note that dual curves emanating from a pair of edges on one of the four sides cannot cross each other. This is because γ' is a geodesic in $\widetilde{X}$, because $\widetilde{Y}$ is convex, so we could reduce area by pushing λ inwards, and σ_i can likewise be pushed inwards to reduce the area.

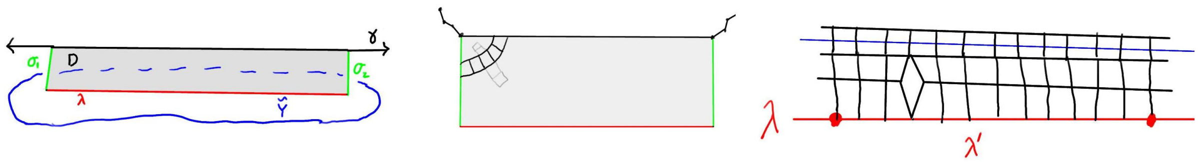

FIGURE 143.

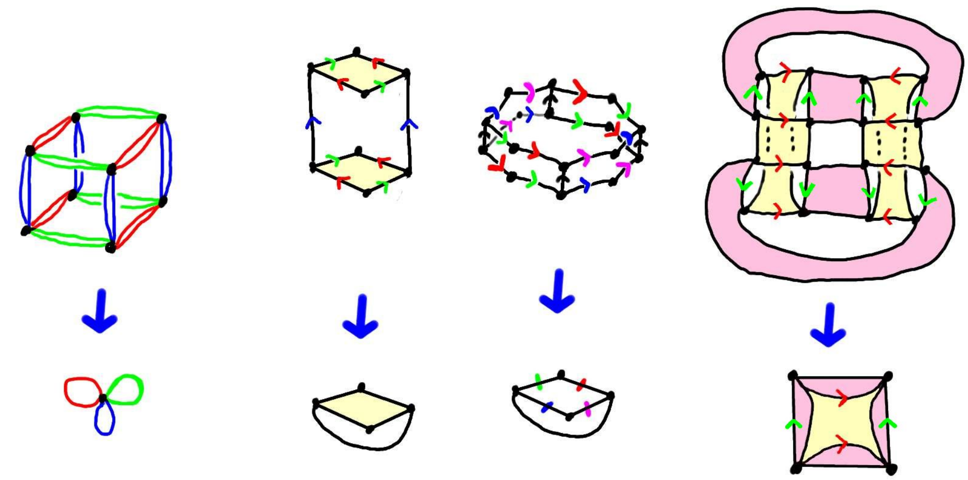

FIGURE 144. The preimages in $\ddot{Y} \to Y$ of hyperplanes in Y are walls in $\ddot{Y}$. We have indicated some walls with corresponding labelled edges. The most interesting case is on the right, where $\#_W \neq \mathsf{H}_1(-;\mathbb{Z}_2)$.

Consequently, D is a flat rectangle. In particular, each dual curve travels between σ_1, σ_2 or between γ', λ. Indeed, a minimal (inward) counterexample going across a corner, would yield a generalized outerpath of a square on one side or the other, by considering dual curves that it bounds. See Figure 143.

By the pigeon-hole principle, since Y is compact, there are arbitrarily long subpaths of λ whose endpoints project to the same point of Y. Let λ' be such a subpath, and note that it projects to a closed essential path in Y.

Using the fact that D is not a line segment, observe that λ' has a hyperplane piece with a hyperplane represented by a dual curve from σ_1 to σ_2. This contradicts Condition 5.4.

9. SPLICING WALLS

9.1. **A finite cover that is wallspace.**

Construction 9.1 (Splicing)**.** Let Y be a connected nonpositively curved cube complex whose hyperplanes are 2-sided and embedded. Let $\Lambda(Y)$ denote the set of hyperplanes of Y. Let $q : \Lambda(Y) \to S$ be a map with the property that for each $s \in S$, no two hyperplanes in $q^{-1}(s)$ cross each other.

Consider the homomorphism $\#_q : \pi_1 Y \to \mathbb{Z}_2^S$ induced by $\#_q(e) = v_{q(\Lambda_e)}$ where Λ_e is the hyperplane dual to the 1-cube e and where $\mathbb{Z}_2^S$ has basis $\{v_s : s \in S\}$.

Let $\ddot{Y}$ denote the cover of Y corresponding to the kernel of $\#_q$. Let $s \in S$ lie in the image of q. Let W_s denote the collection of hyperplanes of $\ddot{Y}$ that map to hyperplanes of Y which map to s.

Remark 9.2. We have in mind the following situation: $Y \hookrightarrow X$ is an embedded locally convex subcomplex of a nonpositively curved cube complex X whose hyperplanes are embedded and 2-sided.

The set S equals the collection $\Lambda(X)$ of hyperplanes of X. The map $q : \Lambda(Y) \to \Lambda(X)$ sends a hyperplane of Y to the hyperplane of X containing it. The map $\#_q : \pi_1 X \to \mathbb{Z}_2^{\Lambda(X)}$ sends each path

σ to the $\mathbb{Z}_2$-vector whose v_Λ coordinate is the number of times (modulo 2) that σ passes through the hyperplane Λ. The map $\#_q : \pi_1 Y \to \mathbb{Z}_2^{\Lambda(X)}$ is induced by composition with the natural map $\pi_1 Y \to \pi_1 X$.

Lemma 9.3. *For each $s \in q(\Lambda(Y))$, the collection of hyperplanes W_s separates $\ddot{Y}$.*

The motivation is to produce a collection of walls in $\ddot{Y}$ each of which corresponds to some W_s.

Proof. Consider a closed edge path $\ddot{\sigma}$ in $\ddot{Y}$. We must show that $\ddot{\sigma}$ cannot pass through hyperplanes of W_s an odd number of times. If this were the case, then the image σ of $\ddot{\sigma}$ in Y would pass through hyperplanes of $q^{-1}(s)$ an odd number of times. But then σ would not be in the kernel of $\pi_1 Y \to \mathbb{Z}_2^S$, for $\#_q(\sigma)$ would take the value 1 on the v_s coordinate. □

We will also employ the following simple method of inducing a wallspace on a cover.

Construction 9.4 (Cover induced wallspace)**.** , Let Y be a nonpositively curved cube complex that is a wallspace. Let $Y' \to Y$ be a covering space. There is an *induced wallspace structure* on Y' where each wall of Y' is the preimage of a wall of Y, and each halfspace of that wall in Y' is the preimage of a halfspace in Y.

9.2. **Preservation of small-cancellation and obtaining wall convexity.**

Lemma 9.5 (Obtaining π-wall separation)**.** *Let $\langle X \mid Y_1, \ldots, Y_k \rangle$ be a cubical presentation. Suppose that X has finitely many immersed hyperplanes D, and that for each $\widetilde{D}$ and $r > 0$, there are finitely many distinct translates $g\widetilde{D}$ with $d_{\widetilde{X}}(\widetilde{D}, g\widetilde{D}) \leq r$. Suppose that $\pi_1 D$ is separable in $\pi_1 X$ for each hyperplane D, and suppose that there is a uniform upper bound on the diameters of wall-pieces and cone-pieces. Then for each θ there is a finite regular cover $\widehat{X} \to X$ whose induced covers $\widehat{Y}_1, \ldots, \widehat{Y}_k$ have the following property: Consider $\langle \widehat{X} \mid g\widehat{Y}_i : 1 \leq i \leq k, g \in \text{Aut}(\widehat{X}) \rangle$. For any path $S \to \widehat{Y}_i$ that starts and ends on 1-cells of a hyperplane $\widehat{D}$ of $\widehat{X}$ but is not homotopic into $\widehat{D}$, we have $\Omega_{Y_i}(S) \geq \theta$.*

The condition on the hyperplanes is naturally formulated in terms of double cosets with short representatives and it obviously holds when X is compact.

Proof. By the upper bound on diameters of pieces, for each hyperplane D and each Y_i, there are finitely many nontrivial cosets Dg_i represented by elements S with $\Omega_{Y_i}(S) < \theta$ for some Y_i. By separability of $\pi_1 D$, we can choose a cover $\hat{X}_D$ separating $\pi_1 D$ from each of these cosets, so that there is no path $S \to Y_i$ that lifts to a path $\widehat{S}$ in $\widehat{X}_D$ that starts and ends on D, and has $\Omega_{Y_i}(S) < \theta$, unless S is homotopic into D.

Let $\widehat{X}$ be a regular cover factoring through each $\widehat{X}_D$. Then $\widehat{X}$ has the desired property. □

Lemma 9.6 (Obtaining and preserving wall convexity)**.** *Let $\ddot{Y}$ be obtained as in Construction 9.1 from $Y \to X$.*

Suppose that for each path $P \to Y$ whose endpoints are on 1-cells dual to the same hyperplane of X, either $\Omega_i(P) \geq \pi$ or P is homotopic in X (and hence in Y by local convexity) into the carrier of this hyperplane. Then the same holds for paths $\ddot{P} \to \ddot{Y}$.

Similarly, suppose $\Omega_Y(P) \geq \theta$ whenever $P \to Y$ whose first and last edges are dual to hyperplanes in the same wall of Y. Then the same holds for paths $\ddot{P} \to \ddot{Y}$ whose first and last edges are dual to the same wall of $\ddot{Y}$.

Proof. In each case, a path $\ddot{P} \to \ddot{Y}$ with the above start-end property, projects to a path $P \to Y$ with the same property. The defect $\Omega_{\ddot{Y}}(\ddot{P}) = \Omega_Y(P)$. □

Lemma 9.7. *Let $\{Y_i \to X\}$ be local isometries of cube complexes. Let $\widehat{Y_i} \to Y_i$ be regular covers with the following symmetry property: Each automorphism ϕ of $Y_i \to X$ lifts to an automorphism $\widehat{\phi}$ of $\widehat{Y_i} \to X$.*

For each condition listed in Definition 5.1, if $\langle X \mid Y_i\rangle$ satisfies this condition then so does $\langle X \mid \widehat{Y_i}\rangle$.

Proof. For Condition 5.1.(3), we utilize Construction 9.4 to produce the wallspace structure on each $\widehat{Y_i}$, and so if walls of Y_i agree with walls of X then so do walls of $\widehat{Y_i}$.

The other conditions are almost immediate. The main point is that $\widehat{Y_i}$ has fewer essential curves than Y_i, their lengths are the same, and for a path P, we have $\Omega_{Y_i}(P) = \Omega_{\widehat{Y_i}}(P)$. □

9.3. ⊛ **Obtaining the separation properties for pseudographs.**

Lemma 9.8. *Suppose that Y has a complete disjoint system of CAT(0) hyperplanes.*

Assume something about girth. and maybe something about pieces and hyperplanes in a wall.

Conclude that $\widehat{Y}_{\mathcal{W}}$ has the correct separation properties:

If $\widehat{\Lambda}, \widehat{\Lambda}'$ are distinct hyperplanes in the same wall, then they are at least 4 pieces away from each other.

For pairs of hyperplanes $\widehat{\Lambda}_1$, $\widehat{\Lambda}_2$, that project to hyperplanes Λ_1, Λ_2 that are far from each other (more than two pieces away), they are separated by a wall $[\widehat{\Lambda}_3]$ such that Λ_3 is far from both Λ_1, Λ_2.

Let us prove Lemma 9.8 for the case where Y is a graph and the pieces are trees.

Graph Case. We assume that the "pieces" are collections of finite trees $\{T_j\}$.

If σ is a path that starts and ends on the same hyperplane Λ of Y (note that Λ is an edge in this case), then either $\#_{\mathbb{W}}(\sigma) = 0$, or $\#_{\mathbb{W}}(\sigma)$ has a nonzero coordinate on some hyperplane Λ_f dual to an edge f outside any tree neighborhood T of Λ. In the former case σ lifts to a path $\widehat{\sigma}$ in $\widehat{Y}_{\mathcal{W}}$ which starts and ends on the same hyperplane.

In the latter case, $\widehat{\sigma}$ connects distinct hyperplanes $\widehat{\Lambda}, \widehat{\Lambda}'$, but there is a wall $[\widehat{\Lambda}_f]$ that separates $\widehat{\Lambda}, \widehat{\Lambda}'$ since $\widehat{\sigma}$ passes through $[\widehat{\Lambda}_f]$ an odd number of times. Furthermore, $[\widehat{\Lambda}_f]$ is disjoint from the tree neighborhoods $\widehat{T}, \widehat{T}'$ of Λ_e, Λ'_e.

A similar statement holds for a pair of 1-cells e, d, that lie far apart. Any path σ between them must have an odd value on some edge f in the complement of the disjoint neighborhoods of e, d. Consequently, any path $\widehat{\sigma}$ projecting to σ cuts through $[\widehat{\Lambda}_f]$ an odd number of times. □

General Case. Suppose Y is a special cube complex, and some subset V_i of the hyperplanes of Y form a disjoint complete set of hyperplanes. Then there is a finite cover of Y which is splicable and satisfies the $B(6)$ separation conditions.

After passing to a finite cover $\widehat{Y}$ there is a combinatorial map $\widehat{Y} \to \Gamma$ such that Γ is the graph dual to the complete disjoint set of hyperplanes. The map is just the map that forgets about (the orientations of) the other hyperplanes. It is definable in the universal cover, and the map is equivariant, so it is definable in the quotient space as well.

The remaining hyperplanes map to finitely generated subgroups of $\pi_1\Gamma$, so without loss of generality we may assume that they map to free factors, and that moreover, they homotopically map to subgraphs of Γ. Indeed, this would be the case after passing to a finite cover. ("Homotopically" is practically replaced by the idea that they map π_1-isomorphically to, and their images are contained in a slight π_1-injective uniform thickening of various subgraphs.)

The separation properties between hyperplanes dual to our edges is identical to the proof in the special case - if they are far enough apart then there is a separating hyperplane - chosen from the disjoint complete family.

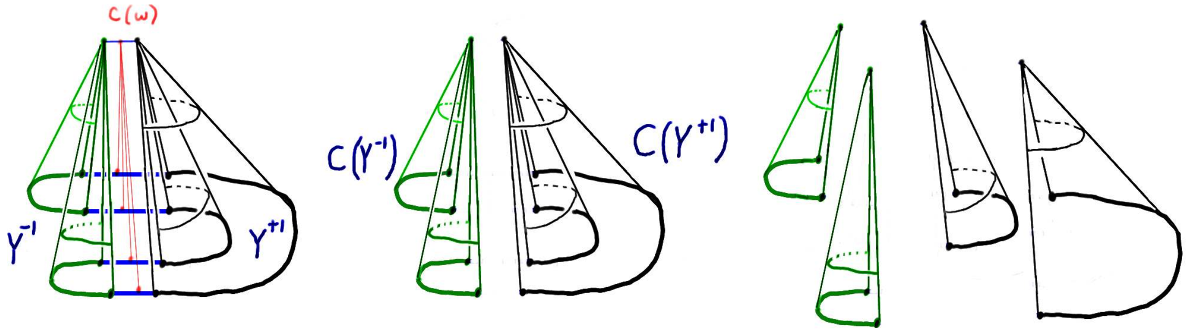

FIGURE 145.

Applying separability results, we should be able to arrange a cover of Γ such that there is very large girth relative to Γ_v subgraph. This means, any path that starts and ends on a Γ_v subgraph is either homotopic into it, or travels quite far away. Thus any path that is nonzero in the wall-count homomophism would have an odd value on some edge far from Γ_v.

Similarly, when Γ_v and Γ_u are sufficiently far away the same holds. □

Remark 9.9. It is important to keep track of a few issues related to torsion.

Firstly there is a trick to get rid of the first layer of torsion- just choose the large girth covers of the objects we are quotienting, so that they are induced by a particular cover of the special cube complex. That allows us to get rid of that torsion!

Now, the $\mathbb{Z}_2^{\mathcal{W}}$ torsion can either be left as is in the case of searching for an actual virtual special quotient. (It leaves walls invariant, so it won't effect the malnormal hierarchy business).

When we are only interested in a special quotient after applying the forgetful functor to the cube complex, we might be able to use the $\mathbb{Z}_2^{\mathcal{W}}$ cover induced on each particular quotiented subcomplex. This would group walls together according to their grouping in the larger complex. We would therefore need to assume that walls in the subcomplex that map to the same wall in the main complex are VERY far apart. This plays the role of $B(6)$ small-cancellation, I reckon, so that our walls are quasiconvex - and perhaps gives malnormality.

We only need geometric separation.

We need hyperplane separation to get linear separators!

The following example has $\pi_1 X$ free, yet there is no complete set of cuts in any finite cover $\widehat{X} \to X$.

Example 9.10. Let X be the square complex obtained by identifying two squares along their vertices. Then $\#_W \ncong \mathsf{H}_1(-;\mathbb{Z}_2)$ for any finite cover $\widehat{X}$. See the rightmost diagram in Figure 144.

10. CUTTING X^*

Construction 10.1 (Inflating a single cone with respect to wall)**.** Let Y be a nonpositively curved cube complex. Let w be a wall consisting of the union of a separating collection of disjoint 2-sided embedded hyperplanes that do not self-osculate.

Let $N^o(w)$ be the open cubical neighborhood of w, and note that $N^o(w) \cong w\times(-1,1)$ by assumption. There is a natural map $w\times[-1,1] \to Y$. Let Y^{-1} and Y^{+1} denote the parts of $Y - N^o(w)$ on opposite sides of the wall w so that $Y^{\pm 1}$ contains the image of $w\times\{\pm 1\}$.

We define $C_w(Y)$ to be the union of the ordinary cones $C(Y^{+1})$ and $C(Y^{-1})$ glued together with $C(w)\times[-1,1]$ by identifying $C(w)\times\{-1\}$ with its image in $C(Y^{-1})$ and identifying $C(w)\times\{+1\}$ with its image in $C(Y^{+1})$. We identify $C(w)$ with $C(w)\times\{0\} \subset C(w)\times(-1,1)$. We refer the reader to Figure 145.

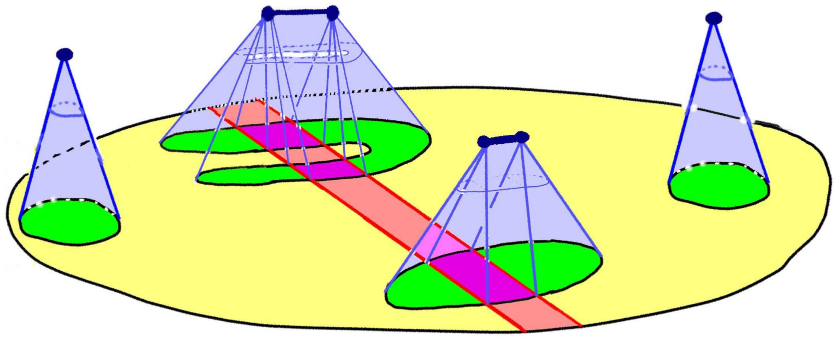

FIGURE 146. $C_w(\langle X \mid Y_1, Y_2, Y_3, Y_4\rangle)$ has $Y_i \cap w$ with $0, 2, 1, 0$ components respectively.

There is a natural combinatorial map $C_w(Y) \to C(Y)$ induced by quotienting the cone-edge to the cone-point.

Construction 10.2 (Inflated coned space with respect to wall)**.** We now consider $\langle X \mid Y_i\rangle$, and suppose that each Y_i is a wallspace. Let w be a *base-wall* of X^* consisting of a collection of disjoint hyperplanes with the property that $w_i = w \cap Y_i$ is a wall in the wallspace structure of Y_i for each i. Generalizing the situation where $Y_i \to X$ is an embedding, we use the notation $w \cap Y_i$ for the preimage of w under the map $Y_i \to X$. We shall *inflate* the coned off space X^* which equals the union $X \cup \bigcup_i C(Y_i)$ as follows: Let $C_{w_i}(Y_i)$ denote the inflated cone space of Y_i with respect to w_i. When $w_i = \emptyset$ we simply let $C_{w_i}(Y_i) = C(Y_i)$. The *base* of $C_{w_i}(Y_i)$ is Y_i as usual. Now define $C_w(\langle X \mid Y_i\rangle)$ to be $X \cup \bigcup_i C_{w_i}(Y_i)$ where for each i we attach $C_{w_i}(Y_i)$ to X along its base using the map $Y_i \to X$.

Note that the map $C_w(\langle X \mid Y_i\rangle) \to X^*$ that sends each cone-edge to a cone-point and each inflated cone space to a cone space is a homotopy equivalence.

Observe that the base-wall w has a natural geometric extension w' in $C_w(\langle X \mid Y_i\rangle)$ consisting of the union of hyperplanes in w, together with the cone $C(w_i)$ on each w_i in Y_i, which lies in $C(w_i) \times \{0\} \subset C(w_i) \times (-1, 1))$. In particular, w' has an open cubical neighborhood $N^o(w')$ in $C_w(\langle X \mid Y_i\rangle)$ consisting of the union of open cubical neighborhoods of hyperplanes in w together with the open cubical neighborhoods $C(w_i) \times (-1, 1)$ of $C(w_i)$ in each $C_{w_i}(Y_i)$ with $w_i \neq \emptyset$.

We will examine the image of $\pi_1 w'$ below by interpreting it, or rather $N(w') = w' \times I$ as a certain π_1-injective subcomplex (it is a subcomplex provided certain non self-osculation assumptions hold), and use this to understand the splitting of $\pi_1 X^*$.

The splitting of $\pi_1 X^*$ is represented geometrically by cutting $C_w(\langle X \mid Y_i\rangle)$ along w'. In particular, $C_w(\langle X \mid Y_i\rangle) - w'$ deformation retracts to $C_w(\langle X \mid Y_i\rangle) - N^o(w')$ which equals the space obtained from X by coning off the various Y_i^{+1} and Y_i^{-1}. More precisely $C_w(\langle X \mid Y_i\rangle) - N^\circ(w') = C(\langle X - N^\circ(W) \mid Y_j : w_j = \emptyset, Y_i^{+1}, Y_i^{-1} : w_i \neq \emptyset\rangle)$. It is often the case that Y_i^{+1} or Y_i^{-1} is not connected, but is treated as one unit and receives a single cone-point in $C_w(\langle X \mid Y_i\rangle)$. See Figure 145.

By adding additional splittings (over trivial groups) we can pass to a space which is the coned off space of the complex (which may have two components) consisting of $\langle X - N^o(w) \mid Y_i^{\pm 1}\rangle$. The wallspace structures on Y_i^{+1} and Y_i^{-1} are induced by the original wallspace Y_i but we note that a wall in Y_i^{+1} might have more hyperplanes than in Y_i.

Lemma 10.3 (Persistence after cutting)**.** *Small-cancellation conditions on $\langle X \mid Y_i\rangle$ are preserved after cutting along a wall, and then cutting along cone-points.*

In particular, the $B(6)$ and $B(8)$ conditions persist.

The malnormality of a collection of hyperplanes in a base-wall persists.

The injectivity radius of hyperplanes exceeding half the diameters of cone-pieces persists.

Condition 5.64.(3) ensuring that distinct hyperplanes in the same base-wall cannot be dual to 1-cells in the same cone-piece persists.

Proof. Each component of $Y_i - N^o(w)$ is a nonpositively curved cube complex. And continues to embed in $X - N^o(w)$ as a locally convex subspace.

If Y_i was a wallspace, then the walls of a component Z_i of $Y_i - N^o(w)$ are defined to be the intersection of Z_i with walls of Y_i.

Lemma 10.7 implies the persistence of $B(6)$ and $B(8)$.

Lemma 10.6 verifies that the malnormality of a collection of hyperplanes persists.

We verify in Lemma 10.10 the persistence of Condition 5.64.(3). □

Definition 10.4. [Geometric malnormal collection] We say that two annular diagrams $f_1 : A_1 \to X$ and $f_2 : A_2 \to X$ are *equivalent* if they have the same boundary cycles P, P' (in an orientation preserving manner) so that the diagram below commutes, and moreover, such that there are lifts $\widetilde{f_i} : \widetilde{A_i} \to \widetilde{X}$ that restrict to the same lifts of $\widetilde{P}, \widetilde{P}'$.

$$\begin{array}{ccccc} & & P \sqcup P' & & \\ & \swarrow & \downarrow & \searrow & \\ A_1 & \to & X & \leftarrow & A_2 \end{array}$$

A map $Z \to X$ of cell complexes is *malnormal* if there is no *essential* annular diagram $(A, \partial A) \to (X, Z)$ in the sense that any such map is equivalent to a map $(A', \partial A') \to (Z, Z)$.

The above notation is meaningful when $Z \subset X$ is a subspace. More generally, we require that for each commutative diagram below on the left, there exists $A' \to Z$ such that the middle diagram commutes and such that the two annular diagrams on the right are equivalent. (Note that $\partial_p A$ consists of two boundary cycles, and we identify $\partial_p A' = \partial_p A$.)

$$\begin{array}{ccc} \partial_p A & \to & Z \\ \downarrow & & \downarrow \\ A & \to & X \end{array} \qquad \begin{array}{ccc} \partial_p A & \to & Z \\ \downarrow & \nearrow & \\ A' & & \end{array} \qquad \begin{array}{c} A \to X \\ A' \to Z \to X \end{array}$$

A prestidigitative use of Definition 10.4 shows that $Z \to X$ is π_1-injective on each component if it is malnormal. Indeed, a disk diagram in X for a path P in Z can be reconsidered as an annular diagram between P and the trivial path (which could be adjusted to a nontrivial path in Z to accommodate the annular diagram definition). Homotoping the diagram into Z shows that P is already null-homotopic in Z.

We now relate the geometric notion of malnormality with the notion from Definition 12.2 of a malnormal collection (of conjugacy classes) of subgroups:

Lemma 10.5 (Algebraic and geometric forms of malnormality)**.** *Suppose $Z \to X$ is a map where X is connected, and Z is the disjoint union of its components: $Z = \sqcup_{i \in I} Z_i$. Then $\{\pi_1 Z_i\}$ is a malnormal collection in $\pi_1 X$ if and only if the map is malnormal.*

Lemma 10.6. *Let X be a nonpositively curved cube complex. Let $\{h_1, \ldots, h_k\}$ be a malnormal collection of immersed 2-sided hyperplanes in X. Let $\{w_1, \ldots, w_\ell\}$ be a collection of disjoint embedded 2-sided hyperplanes. Let $X' = X - N^o(\cup w_i)$. Let $\{h_{ij}\}$ be hyperplanes in X' mapping to $\{h_1, \ldots, h_k\}$.*

Then $\{h_{ij}\}$ is a malnormal collection of immersed hyperplanes in X'.

We note that when X' is not connected, the algebraic version of malnormality must be suitably reinterpreted.

Downstairs proof. An essential annular diagram $A \to X'$ with boundary cycles mapping to the $\cup N(h_{ij})$ determines an annulus in X with boundary cycles in $\cup N(h_i)$. By malnormality of $\cup N(h_i) \to X$,

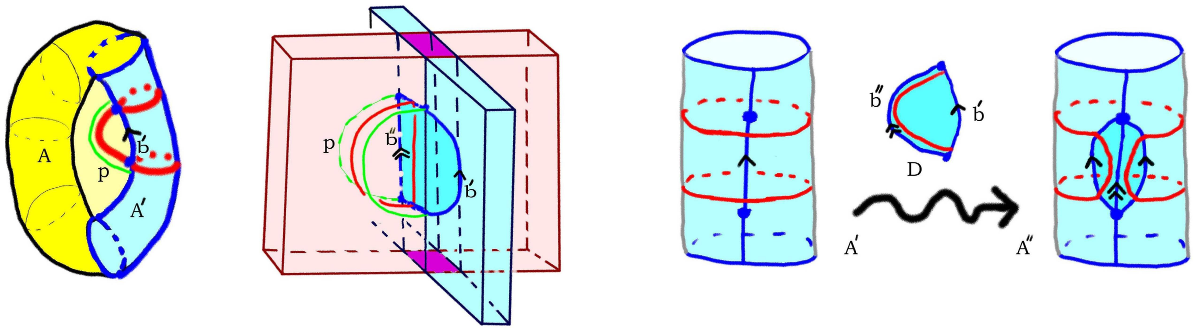

FIGURE 147.

Lemma 10.5 implies that A is homotopic to an annulus $A' \to \cup N(h_i)$, and in particular, into a single $N(h_i)$. This is illustrated on the left in Figure 147.

Suppose that A' is chosen so that it intersects $\{w_1, \ldots, w_\ell\}$ in a minimal number of components. Note that all w-dual curves in A' are closed since $\partial A' \cap w = \emptyset$, and therefore each of these components is a dual w-circle that is either essential or null-homotopic.

By considering innermost null-homotopic w-circles first, each w-circle bounding a disc in A' can obviously be removed through a homotopy without introducing further w-circles. The essential w-circles come in "facing pairs" as in the annulus A' on the left of Figure 147 that are joined by a dual curve in a disk diagram C between the conjugators c, c' of A, A'. This diagram C is illustrated between the annuli A, A' in the left diagram.

The combinatorial path p along the outside of this dual curve is homotopic to the subpath b' of c' whose initial and terminal edges are dual edges. Let E denote the disk diagram between p and b'. It is illustrated as a subspace of the diagram C. Considering lifts $\widetilde{E}$ and $\widetilde{A'} \subset N(\tilde{h}_i)$, we see that $\tilde{p}$ has the same endpoints as $\tilde{b}'$. We can choose a path b'' in $N(\tilde{w}_j) \cap N(\tilde{h}_i)$ that is homotopic to p and hence b'. Let D be the disk diagram in $N(\tilde{h}_i)$ between b' and b''. We illustrate D within the configuration of two crossing hyperplane carriers in $\widetilde{X}$ in the middle of Figure 147. Note that since p doesn't cross $\tilde{w}_i$ we can assume that b'' has the same property. This simplifies our next step, for otherwise extra null-homotopic w-circles would be introduced. Moreover, we can choose D to avoid w_j circles by removing them as above.

Finally we can drag A' along b' into D to obtain a new annular diagram $A'' \to N(h_i)$ as on the right of Figure 147. This turns the pair of essential circles into the single nonessential circle. We then remove this as we did earlier. □

Upstairs proof. If there is an annulus $(A, \partial A) \to (X', \cup N(h_{ij}))$, then the malnormality hypothesis gives an equivalent annulus $(A, \partial A) \to (\cup N(h_i), \cup N(h_{ij}))$, so these maps restrict to the same map on ∂A and have equal conjugators, or equivalently, they have lifts to $\widetilde{X}$ that restrict to the same lifts of $\partial\widetilde{A}$. Our goal is to find another equivalent annulus of the form $(A, \partial A) \to (\cup N(h_{ij}, \cup N(h_{ij}))$.

We work in the universal cover $\widetilde{N}$ of the component $N = N(h_i)$ containing the image of A. Observe that $\widetilde{M} = \widetilde{N} - N^o(\widetilde{w})$ is a convex subcomplex of N. Identifying $\mathbb{Z}$ with $\pi_1 A$, there is a $\mathbb{Z}$-equivariant retraction map $\widetilde{N} \to \widetilde{M}$, which takes a point $x \in \widetilde{N}$ to the unique point in $\widetilde{M}$ that is nearest to it.

The composition $\widetilde{A} \to \widetilde{N} \to \widetilde{M}$ yields the desired map $\widetilde{A} \to \widetilde{N}$. □

We note that if we used the CAT(0) metric (and not the cubical metric) in the above proof, we would produce a $\mathbb{Z}$-equivariant homotopy, and this projects to a homotopy $A \to X$ to $A \to X'$ fixing ∂A.

Lemma 10.7. *The B(6) and B(8) conditions persist.*

Proof. We verify the persistence of Condition 5.1.(5) as the other properties are proven similarly. Let X_1^* be obtained from X^* by cutting along w and then along cone-points, and let Y_1 be a cone of X_1^* obtained in this way from a cone Y of X^*. Let S be a path in Y_1 that starts and ends on a wall v_1 of Y_1 and suppose $\Omega_{Y_1}(S) < \pi$. Then S is a path in Y starting and ending on a wall v inducing v_1 such that $\Omega_Y(S) < \pi$. Consequently, by Condition 5.1.(5) we see that S is homotopic into the carrier of a single hyperplane u of v.

Consider a minimal area square diagram D between S and a path P on $N(u)$. Obviously, no 1-cell of S is dual to w. Consequently, no 1-cell of P is dual to w for otherwise, we could reduce the complexity of D. Indeed, a dual curve in D that is dual to w must start and end on 1-cells of P, and hence we could push P past it to obtain a smaller diagram D' between S and a new path P' that is still on $N(u)$. (This is essentially an application of Lemma 2.2) . Since D is disjoint from w, it is a diagram in Y_1 and we are done as P' lies on the carrier of a single hyperplane $u_1 \subset u$. □

We refer the reader to Definition 3.47 for the notions here. I was unable to resolve the following problem, but found that the stronger condition in Lemma 10.9 is a workable substitute.

Problem 10.8. Does the metric small-cancellation property of short innerpaths persist?

Lemma 10.9. *The following properties persist and they together imply the short innerpath property.*

(1) *There is a bound B_Y on diameters of pieces in each cone Y.*
(2) *An innerpaths S with $\Omega_Y(S) < \pi$ is the concatenation of a uniformly bounded number of pieces (For instance, is 5 enough in the split-angling?)*
(3) *Each essential path in a cone Y has length exceeding twice the sum of the maximal diameter B_Y of such pieces, and hence exceeding $2\nabla_Y(S)$ for any candidate innerpath S.*

Proof. Lower bounds on systoles persist since each induced cone maps to its ancestral cone by a local isometry. Similarly the upper bound on $2\nabla_Y(S)$ persists since pieces in the derived complex are pieces in the ancestor. □

Lemma 10.10 (Wall intersection persists)**.** *Suppose that X^* satisfies Condition 5.64.(3). Let w be a base-wall in X^*, and let X_1^* (or X_1^*, X_2^* in the separating case) be the cubical presentation obtained by cutting along w. Then X_1^* (and also X_2^* in the separating case) also satisfies Condition 5.64.(3).*

Proof. Let Y_1 be a cone of X_1^*, and let h_1, k_1 be hyperplanes in the same wall of Y_1, and suppose that h_1 and k_1 have dual 1-cells in the same cone-piece of Y_1, say associated to a cone Y_1' along some rectangular diagram between Y_1, Y_1'.

By construction, Y_1 is a component of a cone Y obtained by cutting along $N(w)$, so Y_1 is either Y^+ or Y^-. And likewise Y_1' arises from some Y', and the piece between them Y_1, Y_1' arises from a cone-piece between Y and Y' accordingly. In particular, h_1, k_1 arise from hyperplanes h, k in Y that are dual to 1-cells in this cone-piece. Since h_1, k_1 lie hyperplanes in the same base-wall v_1 of X_1^*, and v_1 is induced from a base-wall v of X^*, we see that h, k lie in hyperplanes in the same base-wall v.

Condition 5.64.(3) applied to h, k, v in X^* implies that h, k are in the same hyperplane of Y. We refer the reader to Figure 148. On the left we illustrate that h, k actually lie in the same wall of Y, and consequently, h', k' lie in the same (induced) wall of Y'.

On the right of Figure 148 we illustrate that since pieces are CAT(0) and twice the injectivity radius of hyperplanes exceeds the diameters of pieces, hyperplanes have connected intersection with pieces. □

Remark 10.11 (Algebraic Splitting)**.** Under appropriate small-cancellation conditions, the geometric splitting of $C_w(\langle X \mid Y_i \rangle)$ along w' induces a splitting of $\pi_1 X^*$ along the image of $\pi_1 w'$.

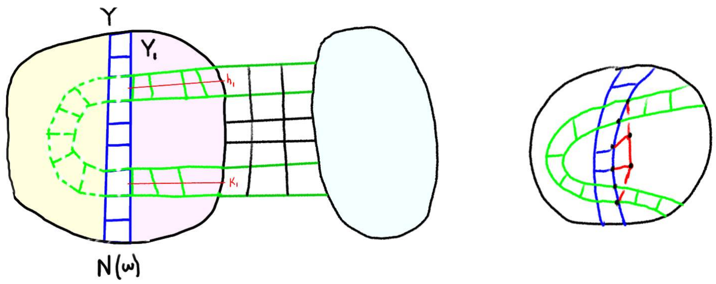

FIGURE 148. Persistence of Condition 5.64.(3)

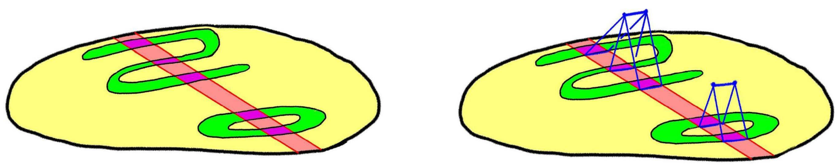

FIGURE 149. Each $T_i \times I$ added on the right connects the wall in a cone.

Adding Dummy Squares to form $\bar{X}^*$: It is inconvenient that $w \cap Y_i$ might be disconnected for some values of i. We remedy this by adding *dummy squares* as follows: Let $w_{i1}, \ldots, w_{ip_i}$ denote the hyperplanes that are the components of $w_i = w \cap Y_i$. Choose 1-cells $e_{i1}, \ldots, e_{ip_i}$ in Y_i that are dual to these. Let T_i be a p_i-*pod* which is a tree consisting of a base-vertex v_i of valence p_i, whose edges end at leaves $v_{i1}, \ldots, v_{ip_i}$. Consider the square complex $T_i \times I$, and now attach a copy of $T_i \times I$ to Y_i with the edge $v_{ik} \times I$ attached to e_{ik} (so that the orientations are consistent), and likewise attach a copy of $T_i \times I$ to X with the edge $v_{ik} \times I$ attached to the image of e_{ik} under the map $Y_i \to X$. We do not perform this procedure when $p_i = 0$ in which case $w \cap Y_i = \emptyset$.

Following the above procedure we obtain $\bar{Y}_i$ for each i, and $\bar{X}$ which contains a $T_i \times I$ contribution for each Y_i. We let $\bar{Y}_i = Y_i$ when $w_i = \emptyset$. We let $\bar{w}$ denote the hyperplane of $\bar{X}$ containing w, so the carrier $N(\bar{w})$ contains $N(w)$ and each $T_i \times I$. The advantage of $\bar{w}$ over w is that the wall $\bar{W}$ in $\widetilde{\bar{X}}^*$ corresponding to $\bar{w}$ has a unique hyperplane. Note that $\bar{W} \cap \widetilde{X}^* = W$, a wall mapping to the base-wall w.

Observe that $\langle \bar{X} \mid \bar{Y}_i \rangle$ deformation retracts to $\langle X \mid Y_i \rangle$ by collapsing along free faces.

Moreover, there is a deformation retraction of $C_{\bar{w}}(\langle \bar{X} \mid \bar{Y}_i \rangle)$ onto $C_w(\langle X \mid Y_i \rangle)$ that pulls each base-edge of T_i upwards to the cone-edge of $C_{\bar{w}_i}(Y_i)$.

Dummy Squares Preserve Small-Cancellation: If $\langle X \mid Y_i \rangle$ satisfies a small-cancellation condition predicated on an upper bound on the diameters of all wall-pieces and cone-pieces, then the same holds for $\langle \bar{X} \mid \bar{Y}_i \rangle$.

First note that the 1-cells parallel to each T_i are certainly not involved in any pieces.

As each $T_i \times I$ is contained in $\bar{Y}_i$, the only way a 1-cell e from another cone with e parallel to the I factor of $T_i \times I$ could be involved with a new piece is as a wall-piece from one of the new hyperplanes dual to an edge of T_i. But such a piece would also be a cone-piece.

(It appears that it would also be a wall-piece (in some other way) unless that cone contained all of w. This would probably violate the $B(6)$ condition 5.1.(5), since a path in w from one component of $w \cap Y_i$ to another such component would lie in a single cone-piece.)

Nevertheless, any small-cancellation condition depending upon sizes of pieces is preserved.

Examining the Splitting: Let $\bar{A} = N(\bar{w}) \cong \bar{w} \times [-1, 1]$, and consider the map $\bar{A} \to \bar{X}$. As in Construction 12.17, we let $\langle \bar{A} | \bar{A} \otimes_{\bar{X}} \bar{Y}_i \rangle$ denote the induced cubical presentation. Our $T_i \times I$ additions have now made the cones connected, whereas $A \otimes_X Y_i$ has components in one-to-one correspondence with components of $w \otimes_X Y_i$.

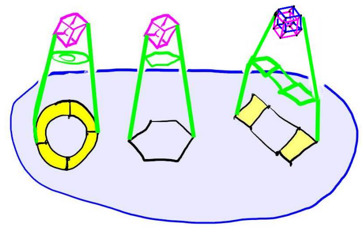

FIGURE 150. Cubical cone points:

(1) Let $S \to Y$ be a path that starts and ends on vertices on the carrier of the same wall W of Y. If $\Omega_Y(S) < \pi$ then S is path-homotopic into $N(W)$.

Condition 5.1.(1) implies that the map $\bar{A}^* \to \bar{X}^*$ has no missing θ-shells, and is thus π_1-injective by Theorem 12.16.

Observe that $\bar{A}^*$ is isomorphic to a product $D \times I$ and that the deformation retraction map $C_{\bar{w}}(\langle \bar{X} \mid \bar{Y}_i \rangle) \to C_w(\langle X \mid Y_i \rangle)$ sends D to a space whose fundamental group is the same as $\mathrm{Stab}(W)$ where W is the wall in $\widetilde{X}^*$ corresponding to w.

We thus see that $\pi_1 \bar{A}^* = \mathrm{Stab}(W)$, and that $\pi_1 X^* = \pi_1 \bar{X}^*$ splits as an HNN extension or AFP over the subgroup $\mathrm{Stab}(W) = \mathrm{Stab}(\bar{W}) = \pi_1 \bar{A}^*$.

The reader might prefer to consider the two maps $D \times \{\pm 1\}$ to $\bar{X} - N^o(\bar{w})$, and verify that there are no missing θ-shells, and that the induced presentations of $D \times \{+1\}$ and $D \times \{-1\}$ are the same.

Remark 10.12 (Cubes instead of conepoints)**.** A promising viewpoint that we have not pursued is to extend the idea of Construction 10.2 to a more complete setting. I describe this in the hope that some variation on this idea will be useful.

As each cone Y_i is a wallspace, there is an embedding $Y_i \to C(Y_i)$ where $C(Y_i)$ is the dual CAT(0) cube complex. (Note that it might not be a cube.) For each i, let M_i be the associated mapping cylinder, and then instead of coning off the cones, we attach a copy of M_i to X along its base $X \leftarrow Y_i \to M_i$. See Figure 150.

11. HIERARCHIES

Definition 11.1. The class $\mathcal{MQH}$ of groups with an *almost malnormal quasiconvex hierarchy* is the smallest class of groups closed under the following conditions:

(1) If $|G| < \infty$ then $G \in \mathcal{MQH}$.
(2) If $G \cong A *_B C$ and $A, C \in \mathcal{MQH}$ and B is quasiconvex and almost malnormal in G, then $G \in \mathcal{MQH}$.
(3) If $G \cong A*_B$ and $A \in \mathcal{MQH}$ and B is quasiconvex and almost malnormal in G then $G \in \mathcal{MQH}$.

It follows from the Bestvina-Feighn Combination Theorem that every group in $\mathcal{MQH}$ is word-hyperbolic [BF92].

In conjunction with Definition 11.1, a *hierarchy* for G is a specific sequence of splittings leading to terminal groups, and the *length* of this hierarchy is the longest maximal (nontrivial) subsequence. Depending on the context we might allow the splittings to be general graphs of groups, or we might insist that each is either an HNN extension or AFP. The hierarchy is *malnormal*, *quasiconvex*, etc. if each of its constituent splittings has this property. The *terminal* groups could be trivial as in $\mathcal{QH}$, or they could be finite as for $\mathcal{MQH}$. Other possibilities are considered in Section 16.

Theorem 11.2 ($\mathcal{MQH}$ is virtually special)**.** *Each group in $\mathcal{MQH}$ is virtually special. More generally, if G has an almost malnormal quasiconvex hierarchy terminating at virtually compact special hyperbolic groups then G is virtually compact special.*

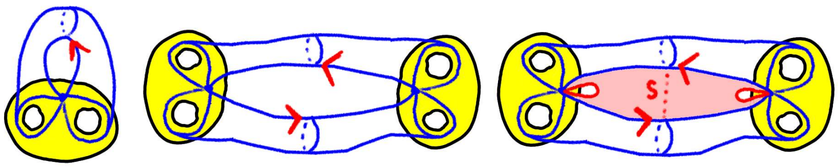

FIGURE 151. $X \leftarrow \widehat{X} \hookrightarrow \widehat{X}'$

Proof. This holds by induction on the "length" of a hierarchy using Theorem 11.3.

(We note that every compact special cube complex with word-hyperbolic fundamental group has a finite cover with a malnormal quasiconvex hierarchy [HWa]. Thus such a hierarchy can be extended further to a virtual malnormal quasiconvex hierarchy.) □

Theorem 11.3. *Let $G = A *_C B$ or $G = A*_C$ where C is almost malnormal and quasiconvex and A, B (respectively A) is virtually compact special. Then G is virtually compact special.*

Proof. Proposition 6.17, together with Proposition 7.8 and the wall case of Corollary 6.8. For the amalgamated product $A *_C B$ one lets A', B' be special finite index subgroups of A, B, and then let $C' = A' \cap C \cap B'$. Then one chooses a collection $\{H_1, \ldots, H_r\}$ of quasiconvex codimension-1 subgroups cubulating C' (and consequently cubulate C). By Corollary 6.8, this collection is extendible to quasiconvex codimension-1 subgroups $\{H_1^A, \ldots, H_r^A\}$ in A and $\{H_1^B, \ldots, H_r^B\}$ in B.

The HNN case is similar but also follows from the AFP case as described in Remark 11.4. □

Remark 11.4. Revising the notation, let X be a graph of spaces with a vertex space A and edge space $C \times I$ corresponding to the HNN extension, so $G = \pi_1 X$ splits as $\pi_1 A *_{\pi_1 C}$. Let $\widehat{X}$ denote the double cover of X corresponding to the map sending the stable letter t of the HNN extension to the generator of $\mathbb{Z}_2$ and sending $\pi_1 A$ to the identity element. Let S be a "dummy-square" such that $\partial_p S$ has label $a_1 t_1 a_2 t_2$. We add S to $\widehat{X}$ to form $\widehat{X}'$ by identifying t_1 and t_2 with the two copies of the t-edge in $\widehat{X}$. We do this so that their orientations agree. See Figure 151. Finally, there is a splitting of $\widehat{X}'$ as a graph of spaces: each vertex space equals $\widehat{A_i} \cup a_i$ where $\widehat{A_1}, \widehat{A_2}$ are the two components of the preimage of A in $\widehat{X}$, and a_1, a_2 are the other two 1-cells on ∂S. And the sole edge space consists of the two copies $(C \times I)_1, (C \times I)_2$ of the preimage of $C \times I$, joined together with the square S. Note that $\pi_1 \widehat{X}' \cong \pi_1 \widehat{X} * \mathbb{Z}$ is word-hyperbolic if $\pi_1 X$ is. The reader can check that our splitting as a graph of spaces induces a splitting of $\pi_1 \widehat{X}'$ as an amalgamated free product, whose vertex groups are each isomorphic to $\pi_1 A * \mathbb{Z}$ and hence virtually special, and whose edge group is isomorphic to $\pi_1 C * \pi_1 C$ and is easily seen to be quasiconvex.

The height of the pair of inclusions of $\pi_1 C$ into $\pi_1 A$ inside the HNN extension equals the height of each edge group $\pi_1 C * \pi_1 C$ in each vertex group of the amalgamated free product. In particular, almost malnormality of the splitting is preserved.

If the amalgamated free product is virtually compact special, then the original group is as well, since it has a finite index subgroup with this property. (Here one needs word-hyperbolicity to guarantee the compactness. See Remark 12.6).

The amalgamated free product case of Theorem 13.1 thus implies that $\pi_1 \widehat{X}'$ is virtually compact special, and hence the same holds for the quasiconvex subgroup $\pi_1 \widehat{X}$, and hence $\pi_1 X$.

Theorem 11.2 generalizes to the class $\mathcal{QMVH}$ of groups with a *quasiconvex malnormal virtual hierarchy*. It is closed under the following additional condition:

(4) If $[G : H] < \infty$ and $H \in \mathcal{QMVH}$ then $G \in \mathcal{QMVH}$.

We aim to analyze the class of groups that are hyperbolic relative to tori, and that have quasiconvex hierarchies.

Each compact special cube complex X has $\pi_1 X \in \mathcal{QH}$. It is therefore interesting to examine possible converses, but it is unclear exactly what extra ingredients are needed to ensure that a group with a hierarchy is virtually special.

Definition 11.5. Let $\mathcal{QH}$ denote the smallest class $\mathcal{G}$ of groups that is closed under the following three operations, and let $\mathcal{QVH}$ denote the related class obtained by adding the fourth operation.

(1) $\{1\} \in \mathcal{G}$.
(2) If $G = A *_B C$ and $A, C \in \mathcal{G}$ and B is f.g. and embeds by a quasi-isometry, then G is in $\mathcal{G}$.
(3) If $G = A*_B$ and $A \in \mathcal{G}$ and B is f.g. and embeds by a quasi-isometry, then G is in $\mathcal{G}$.
(4) Let $H \subset G$ with $[G : H] < \infty$. If $H \in \mathcal{G}$ then $G \in \mathcal{G}$.

Remark 11.6. The class of groups with a hierarchy such that each edge group is f.g. (and free) is nastier than one might expect, and certainly contains groups that are not virtually special. For instance it is known that such groups can have undecidable word problem. Even among 2-dimensional cubical complexes whose fundamental groups exhibit such a hierarchy, there are examples that are not virtually special [BM97, Wis07]. The simplest examples with a hierarchy that are not virtually special are certain Baumslag Solitar groups (see Problem 11.8).

One of the main results in this paper is that every virtually torsion-free word-hyperbolic group in $\mathcal{QVH}$ is virtually special. It is unclear what hypothesis (besides hyperbolic relative to abelian subgroups) could replace hyperbolicity in a variant of the above statement.

Conjecture 11.7 ($\mathcal{QSVH}$ is virtually special)**.** *Let us restrict the third operation as follows:*

(3') *If $G = A*_B$ and $A \in \mathcal{G}$ and B is f.g. and embeds by a quasi-isometry and B is separable in G, then G is in $\mathcal{G}$.*

Let $\mathcal{QSVH}$ denote the class $\mathcal{G}$ of groups that is closed under the four operations. These are the groups with a Separable Quasi-isometric Virtual Hierarchy. *Then for every group $G \in \mathcal{QSVH}$ there is a finite index subgroup $H \subset G$ such that H is the fundamental group of a special cube complex.*

Problem 11.8. A good test case is the class of one-relator groups. The strongest possible conjecture one could hope for here would be that a one-relator group is virtually special if and only if it contains no subgroup $BS(n, m)$ with $n \neq \pm m$ nonzero, and where $BS(n, m) = \langle a, t \mid (a^n)^t = a^m \rangle$.

Another important test case are free-by-cyclic groups and more generally ascending HNN extensions of free groups that are hyperbolic relative to virtually abelian subgroups.

A special case of Conjecture 11.7 is:

Conjecture 11.9. *Let X be a (compact) nonpositively curved cube complex. If $\pi_1 D$ is separable in $\pi_1 X$ for each hyperplane D of X, then X is virtually special.*

12. Virtually Special Quotient Theorem

The goal of this section is to prove the following theorem:

Theorem 12.1 (Virtually Special Quotient)**.** *Let G be a word-hyperbolic group with a finite index subgroup that is the fundamental group of a compact special cube complex. Let $H_1, \ldots, H_r$ be quasiconvex subgroups of G. Then there are finite index subgroups $H'_1, \ldots, H'_r$ such that the quotient: $G' = G/\langle\langle H'_1, \ldots, H'_r \rangle\rangle$ is a word-hyperbolic group with a finite index subgroup that is the fundamental group of a compact special cube complex.*

12.1. **The malnormal special quotient theorem.**

Definition 12.2 (Malnormal Collection)**.** A collection of subgroups $\{H_1, \ldots, H_r\}$ of G is *malnormal* provided that $H_i^g \cap H_j = \{1_G\}$ unless $i = j$ and $g \in H_i$. Similarly, the collection is *almost malnormal* if intersections of nontrivial conjugates are finite (instead of trivial). Note that this condition implies that $H_i \neq H_j$ (unless they are finite in the almost malnormal case).

Theorem 12.1 will be a consequence of the following special case which is our main focus in this section, and probably the crux of this paper:

Theorem 12.3 (Malnormal Virtually Special Quotient)**.** *Let G be a word-hyperbolic group with a finite index subgroup J that is the fundamental group of a compact special cube complex. Let $\{H_1, \ldots, H_r\}$ be an almost malnormal collection of quasiconvex subgroups of G. Then there are finite index subgroups $\ddot{H}_1, \ldots, \ddot{H}_r$ such that: For any finite index subgroups $H_i', \ldots, H_r'$ contained in the $\ddot{H}_1, \ldots, \ddot{H}_r$ the quotient: $G' = G/\langle\langle H_1', \ldots, H_r' \rangle\rangle$ is a word-hyperbolic group with a finite index subgroup J' that is the fundamental group of a compact special cube complex.*

The statement will be proven using subgroups H_i' and $\ddot{H}_i$ that are normal in H_i, but consequently holds (as stated) without assuming these subgroups are normal. We also note that the statement allows us to choose the H_i' to lie in any pre-chosen finite index subgroups H_i° of H_i.

Remark 12.4. There is a tricky point in the algebraic statement. When passing to a nonpositively curved cube complex X for a finite index subgroup J of G, each subgroup H_i appears in several ways as a subgroup $H_{ij} = H_i^{g_j} \cap J$. We could keep the subgroups H_{ij} abstractly isomorphic by assuming that J is normal in G. However, they might not have the same representative geometry when we pass to a nonpositively curved cube complex X with $\pi_1 X \cong J$. This can be remedied by recubulating with a cube complex containing an action of G/J. One way to do this is to use an action of G on $X^{[G:J]}$. Another way uses the following lemma:

Lemma 12.5. *Let G be word-hyperbolic and suppose that G has a finite index subgroup that acts properly and cocompactly on a CAT(0) cube complex. Then G acts properly and cocompactly on a CAT(0) cube complex.*

Remark 12.6. The situation for a relatively hyperbolic group is more delicate. A case in point, is the Coxeter group $\langle a_1, a_2, a_3 \mid a_i^2, (a_i a_j)^3 \rangle$ which doesn't act properly and cocompactly on a CAT(0) cube complex but has an index 6 subgroup isomorphic to $\mathbb{Z}^2$.

Proof of Lemma 12.5. Let G' be a finite index subgroup of G that acts properly and cocompactly on a CAT(0) cube complex $\widetilde{Y}$. To simplify matters, we can assume that $\widetilde{Y}'$ does not contain a G'-invariant CAT(0) cubical subcomplex. (For otherwise we could pass to such a subcomplex $\widetilde{Y}'$, which leads to a slightly smaller result for G.) Let $U_1, \ldots, U_r$ denote representatives of the G'-orbits of hyperplanes in $\widetilde{Y}$. Then $\text{Stab}(U_i)$ is a codimension-1 subgroup of G' for each i. Accordingly $\text{Stab}(U_i)$ is a codimension-1 subgroup of G as well, and a finite neighborhood $N_i = N_{r_i}(\text{Stab}(U_i))$ separates the Cayley graph $\Gamma(G, S)$, so we obtain a wall consisting of a deep component of $\Gamma(G, S) - N_i$ together with its complement. We assume that we have chosen the component so that the wall is in agreement with the way U_i separates an orbit of $G'\widetilde{y}$.

We consider the full collection of G-translates of such walls, and observe that G acts properly on the resulting cube complex since G' does, and acts cocompactly because each subgroup is quasiconvex [Sag97]. □

Proof of Theorem 12.3. By Lemma 12.5, there is a proper cocompact action of G on a CAT(0) cube complex $\widetilde{X}_0$.

Since G has separable quasiconvex subgroups (as it is virtually special), it has a finite index subgroup J that acts specially on $\widetilde{X}_0$ (see [HW10]), so $X_0 = J\backslash\widetilde{X}_0$ is a compact special cube complex. Note that specialness is preserved by arbitrary further covers.

By Proposition 8.2, for each i, let $\widetilde{Z}_i \subset \widetilde{X}_0$ denote an H_i-cocompact superconvex subcomplex. Let $H_{i0} = H_i \cap J$ for each i. Let $Z_i = H_{i0}\backslash\widetilde{Z}_i$ so $Z_i \to X_0$ is a compact superconvex local isometry representing $H_{i0} \subset J$. (We can ignore the basepoint here as only the conjugacy class in J interests us).

A regular cover $X \to X_0$ determines a cubical presentation in the following way: Let $\widehat{Z}_i$ denote an elevation of Z_i to X, so $\widehat{Z}_i$ is the cover of Z_i corresponding to $\pi_1 Z_i \cap \pi_1 X$. We shall choose X so that $\pi_1 X$ is normal in G. For $1 \leq i \leq r$ we let $Y_i = \widehat{Z}_i$, and employing the G-action on X to obtain other elevations, we extend this correspondence to let $\{Y_1, \ldots, Y_j\} = \{g\widehat{Z}_i : 1 \leq i \leq k, g \in G\}$ denote the collection all possible such elevations, and we obtain the cubical presentation:

$$\langle X \mid Y_1, \ldots, Y_j \rangle$$

Our initial goal below, is to show how to choose a finite regular cover $X \to X_0$ so that using the above notation, the cubical presentation $\langle X \mid Y_1, \ldots, Y_j \rangle$ satisfies the following properties:

(1) Each hyperplane M of X has $\pi_1 M$ malnormal in $\pi_1 X$.
(2) Each cone $Y_i \to X$ is an embedding.
(3) There exists D such that each cone-piece and wall-piece is a CAT(0) subcomplex of X with diameter $\leq D$.
(4) The length of the smallest essential cycle in Y_i is $> 12D$ (so $C(6)$ (metric) small-cancellation with split-angling holds).
(5) Thus internal cone-cells in reduced disk diagrams have negative curvature.
(6) Twice the injectivity radius of each hyperplane of X is $> D$ and thus exceeds the diameter of any cone-piece.
(7) For each hyperplane $M \subset X$ and path $S \to Y_i$ starting and ending on 0-cells at endpoints of 1-cells dual to M if $\Omega_{Y_i}(S) \leq \pi$ then S is homotopic into $N(M) \cap Y_i$.

We note that the main conditions above are *stable* in the sense that they are preserved by passage to a further finite cover (or finite regular cover), this greatly facilitates achieving the conditions, as we can resolve each specific issue at some finite cover, and then conclude by taking a finite cover factoring through all of these whose fundamental group is normal in G.

Malnormality: For each hyperplane M_ℓ of X_0, the subgroup $\pi_1 M_\ell$ is separable in $\pi_1 X_0$. And there are finitely many cosets $\pi_1 M_\ell g_i$ such that $\pi_1 M_\ell^{g_i} \cap \pi_1 M_\ell$ is infinite. We can therefore separate $\pi_1 M_\ell$ from each g_i in a finite index subgroup to make $\pi_1 M_\ell$ malnormal there. We obtain a finite cover $X_{M_\ell} \to X_0$ for each hyperplane, and any finite regular cover of X_0 factoring through all of these has the desired property. The above sketch is explained in [HW09] in a relatively hyperbolic context.

Cone Injectivity: By separability of quasiconvex subgroups, we can pass to a finite cover $X_i \to X_0$ such that Z_i embeds. (Indeed, $\mathsf{C}(Z_i \to X_0)$ does the trick.) Any finite regular cover of X_0 factoring through each X_i has the desired property.

Piece Size and Embedding: The almost malnormality of $\{H_1, \ldots, H_r\}$ and superconvexity of $\{\widetilde{Y}_1, \ldots, \widetilde{Y}_j\}$ bounds the cone-pieces and wall-pieces in each $g\widetilde{Y}_i$, in that each such cone-piece or wall-piece lies in a diameter D CAT(0) subcomplex of $\widetilde{X} = \widetilde{X}_0$. (See Lemma 3.46, and note that each $\widetilde{Y}_p$ is a translate of a $\widetilde{Z}_q$.) By residual finiteness of $\pi_1 X_0$, and local finiteness of $\widetilde{X}$, we can pass to a finite regular cover of X_0 such that all diameter $\leq D$ CAT(0) subcomplexes of $\widetilde{X}$ embed.

Small-Cancellation and Negative Curvature: For each i, by residual finiteness, we can choose a finite cover of $\bar{Z}_i \to Z_i$ such that $|Z_i| > 12D$. We then choose a regular cover $\bar{X}_i$ such that each elevation of Z_i to $\bar{X}_i$ factors through $\bar{Z}_i$. Any further cover of $\bar{X}_i$ has the desired property for Z_i, and thus any cover X_0 factoring through each $\bar{X}_i$ has the desired property.

Theorem 3.20 applies to show that the split-angling system has nonpositive curvature for (essential) internal cone-cells in reduced disk diagrams, and moreover, essential internal cone-cells have negative curvature since the inequality is strict.

Injectivity radius > $\frac{1}{2}$ Cone-piece diameter: This is a simple consequence of the separability of hyperplane subgroups. The *injectivity radius* of a hyperplane $M \subset X$ is half the length of the shortest path σ starting and ending on 1-cells of M such that σ is not homotopic into M. This equals half the minimal distance between lifts $g_1\widetilde{M}, g_2\widetilde{M}$ in $\widetilde{X}$.

π-Wall Separation: Like the malnormality claim and the injectivity radius claim, this is a consequence of separability of hyperplanes.

Any path $S \to Y_i$ with $\Omega_{Y_i}(S) \leq \pi$ is the concatenation of at most five pieces. Indeed, using the split-angling, the minimal defect for any transition between pieces on S is $\frac{\pi}{4}$. Thus, it suffices to separate from M any path $S \to Y_i$ that is the concatenation of at most five pieces, and that starts and ends on 1-cells dual to the hyperplane M of X_0 but is not homotopic into M. There are finitely many such situations to dispose of as above.

Splicing the cones to obtain $B(8)$:

We now pass to a finite cover $\ddot{Y}_i \to Y_i$ that is induced by a cover $\ddot{X} \to X$ such that $\langle X \mid \ddot{Y}_i\rangle$ satisfies the conditions above (except (2)) as well as the following:

(8) The $B(6)$ and $B(8)$ small-cancellation conditions in Definitions 5.1 and 5.55.

Let $\mathbb{W} = \mathbb{W}(X)$ denote the family of hyperplanes of X. Consider the map $\#_{\mathbb{W}} : \pi_1 X \to \mathbb{Z}_2^{\mathbb{W}}$ induced by counting (modulo 2) the number of times a path passes through a hyperplane. As in Subsection 9.1, let $\ddot{X}$ denote the cover of X associated to $\#_{\mathbb{W}}$. Let $\ddot{Y}_i \to Y_i$ denote the cover induced by $\ddot{X} \to X$ for each i. According to Construction 9.1, each $\ddot{Y}_i$ is a wallspace where each wall of $\ddot{Y}_i$ is the preimage of a single hyperplane of X.

We define $\ddot{H}_i$ to equal $\pi_1\ddot{Y}_i$. Note that $\pi_1\ddot{X}$ is a normal subgroup of G, because the intrinsic definition of $\ddot{X} \to X$ implies that it preserves the symmetries of the G-action on X. Consequently each $\ddot{H}_i$ is a normal subgroup of H_i.

For any finite index normal subgroup $H'_i \subset H_i$ that is contained in $\ddot{H}_i$, we let Y'_i denote the corresponding cover of Y_i. Following Construction 9.4, we let each Y'_i have the wallspace structure induced from Y_i.

We use the notation $X^* = \langle X \mid gY'_i : g \in \text{Aut}(X), 1 \leq i \leq r\rangle$.

The walls of $\widetilde{X^*}$ embed in $\widetilde{X^*}$ by Theorem 5.20, and are quasi-isometrically embedded by Lemma 5.35.

By construction, the new walls embed as base-walls in X^*. Indeed, the hyperplanes of a wall of Y'_i is the entire preimage of a hyperplane of X under the map $Y'_i \to X$.

$\langle X \mid gY'_i\rangle$ inherits the small cancellation properties of $\langle X \mid Y_i\rangle$ by Lemma 9.7 and in addition, B(8)-Condition 5.55.(2) is now satisfied by Lemma 9.6.

The Hierarchy: We now apply the cutting instructions of Section 10 where at each stage we alternately cut along a wall and then along various cone-points that are cutpoints. In the former case, the $B(8)$ conditions are satisfied and in the latter case, we are merely splitting along trivial groups. The result after each such pair of cuts are new complexes that are smaller in the sense that they have fewer 1-cells. After finitely many cuts, we arrive at 0-cells.

Since the $B(8)$ conditions are preserved by the cuts, at each stage the walls we are cutting along are quasiconvex by Theorem 5.35 and are almost malnormal by Theorem 5.64.

We thus have an almost malnormal quasiconvex hierarchy, and so the group $\pi_1 X^*$ is virtually special by Theorem 11.2. Finally, observe that $G/\langle\langle H_1', \ldots, H_r' \rangle\rangle$ contains $G' = \pi_1 X^* = \pi_1 X/\langle\langle \pi_1 g Y_i' : g \in \mathrm{Aut}(X), 1 \leq i \leq r \rangle\rangle$ as a finite index subgroup. $\square$

12.2. **Height and Virtual Almost Malnormality.**

Definition 12.7 (Height)**.** Consider a collection $\{H_1, \ldots, H_k\}$ of subgroups of G. We use the notation $H^g = g^{-1}Hg$. We say $H_{m_i}^{g_i}$ and $H_{m_j}^{g_j}$ are *distinct conjugates* unless $m_i = m_j$ and $H_{m_i} g_i = H_{m_j} g_j$. We emphasize that each "conjugate" corresponds to a value of $(m_i, H_{m_i} g_i)$ and not just a subgroup $H_{m_i}^{g_i}$.

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

For instance, if G is infinite and $[G : H]$ is finite then the height of H in G equals $[G : H]$.

The notion of height was introduced and studied in [GMRS98] where it was shown that quasiconvex subgroups of word-hyperbolic groups have finite height. It follows that a finite collection of quasiconvex subgroups also has finite height. We have explored the notion a bit further for relatively hyperbolic groups in [HW09].

A collection of subgroups is almost malnormal as in Definition 12.2 precisely if its height is ≤ 1.

We will later make use of the following closely related result which also hinges on the key point implying the height finiteness. It holds because for conjugates whose intersection contains an infinite order element, the corresponding cosets must lie within a uniformly bounded distance of an axis, and because there are finitely many A-translates of gB cosets within a finite distance of A.

The following was first proven in [GMRS98]:

Proposition 12.8. *Let G be a hyperbolic group and A, B be quasiconvex subgroups. There are finitely many double cosets AgB such that $A^g \cap B$ is infinite.*

Definition 12.9 (Commensurator)**.** The *commensurator* $\mathbb{C}_G(H)$ of H in G is defined by: $\mathbb{C}_G(H) = \{g \in G : [H : H^g \cap H] < \infty\}$. It is shown in [KS96] (see also [Arz01]) that $[\mathbb{C}_G(H) : H] < \infty$ for any infinite quasiconvex subgroup H of the word-hyperbolic group G. Consequently, in this case $\mathbb{C}_G(H)$ is itself quasiconvex in G.

Definition 12.10 (Virtually Almost Malnormal)**.** The subgroup H is *virtually almost malnormal* in G if for each $g \in G$ either $H^g \cap H$ is finite, or $g \in \mathbb{C}_G(H)$.

The notions of height (and *width*) might have been better crafted in terms of intersections up to commensurability. We remedy this with the following simple observation:

Lemma 12.11. *If $H \subset G$ is virtually almost malnormal then $\mathbb{C}_G(H)$ is almost malnormal in G.*

Proof. Suppose that $\mathbb{C}_G(H)^g \cap \mathbb{C}_G(H)$ is infinite. Since $H^g \cap H$ has finite index in $\mathbb{C}_G(H)$ we see that $[(\mathbb{C}_G(H)^g \cap \mathbb{C}_G(H)) : (\mathbb{C}_G(H)^g \cap \mathbb{C}_G(H)) \cap (H^g \cap H)] = [(\mathbb{C}_G(H)^g \cap \mathbb{C}_G(H)) : (H^g \cap H)] < \infty$. Thus $H^g \cap H$ is infinite, and so $[H^g : H] < \infty$ by virtual almost malnormality. But then $g \in \mathbb{C}_G(H)$, so $\mathbb{C}_G(H)$ is almost malnormal as claimed. $\square$

Lemma 12.12. *Let $H_1, \ldots, H_r$ be a collection of quasiconvex subgroups of a word-hyperbolic group G. Let $K_1, \ldots, K_s$ be representatives of the finitely many distinct conjugacy classes of subgroups consisting of intersections of collections of distinct conjugates of $H_1, \ldots, H_k$ in G that are maximal with respect to having infinite intersection. Then $\{\mathbb{C}_G(K_1), \ldots, \mathbb{C}_G(K_s)\}$ is an almost malnormal collection of subgroups of G.*

Note that this is indeed a finite collection by [GMRS98, HW09].

Proof. First observe that by maximality, for each K_i, H_j, and g_ℓ we have either:

(1) $K_i \cap H_j^{g_\ell}$ is finite.
(2) $K_i \subset H_j^{g_\ell}$.

Suppose $\mathbb{C}_G(K_s) \cap \mathbb{C}_G(K_t)^g$ is infinite for some $g \in G$. Since each $[\mathbb{C}_G(K) : K] < \infty$ we see that $K_s \cap K_t^g$ is infinite as well. Now $K_t = \bigcap_{i=1}^m H_{i_t}^{g_{i_t}}$ so:

$$K_s \cap K_t^g = K_s \cap \bigcap_{i=1}^{m} H_{i_t}^{g_{i_t} g} = \bigcap_{i=1}^{m} (K_s \cap H_{i_t}^{g_{i_t} g}).$$

Since the intersection is infinite, we must have alternative (2) that $K_s \subset H_{i_t}^{g_{i_t} g}$ for each i. But then $K_s \subset K_t^g$. Likewise $K_t^g \subset K_s$. Since we have chosen representatives of *distinct* conjugacy classes of maximal intersections (that is, a maximal number of factors, but a smallest intersection) we see that $s = t$, and that $K_s^g = K_s$, and so $g \in \mathbb{C}_G(K_s)$. □

Lemma 12.13 (Malnormal Intersection)**.** *Let $\{B_1, \dots, B_k\}$ be an almost malnormal collection of subgroups of G. For each i, let $\{B_i g_{ij} A : j \in J_i\}$ denote a collection of distinct double cosets such that $B_i^g \cap A$ is infinite if and only if $g \in B_i g_{ij} A$ for some ij. For each ij, let $M_{ij} = B_i^{g_{ij}} \cap A$.*

Then $\{M_{ij} : 1 \le i \le k, j \in J_i\}$ is an almost malnormal collection of subgroups of A.

Proof. Suppose ${M_{ij}}^a \cap M_{pq}$ is infinite for some $a \in A$. Then $B_i^{g_{ij} a} \cap B_p$ is infinite and so by almost malnormality of $\{B_1, \dots, B_k\}$ in G, we must have $i = p$. Thus $M_{ij} = B_i^{g_{ij}} \cap A$ for some ij, and $M_{iq} = B_i^{g_{iq}} \cap A$ for some iq. Since $(B_i^{g_{ij}})^a \cap B_i^{g_{iq}} \supset (B_i^{g_{ij}} \cap A)^a \cap (B_i^{g_{iq}} \cap A) = M_{ij}^a \cap M_{iq}$, we see that $(B_i^{g_{ij}})^a \cap B_i^{g_{iq}}$ is infinite, and so $g_{ij} a g_{iq}^{-1} \in B_i$ by almost malnormality. Thus $B_i \subset [B_i g_{ij} A][B_i g_{iq} A]^{-1}$ so $B_i g_{ij} A = B_i g_{iq} A$ and hence $j = q$. Finally, $B_i^a \cap B_i \supset M_{ij}^a \cap M_{ij}$ is infinite and so $a \in B_i$ by almost malnormality. □

12.3. **The proof of the Special Quotient Theorem.**

Proof of Theorem 12.1. Let $\{K_1, \dots, K_s\}$ denote the collection of infinite maximal intersections of conjugates given in the statement of Lemma 12.12. For each i, let $\mathcal{K}_i$ denote $\mathbb{C}_G(K_i)$.

According to Lemma 12.5, G acts properly and cocompactly on a CAT(0) cube complex $\widetilde{X}$. Let $J \subset G$ be a finite index torsion-free normal subgroup and let $X = J \backslash \widetilde{X}$. Let R be the diameter of a finite ball $U \subset \widetilde{X}$, such that $JU = \widetilde{X}$, and such that U contains the basepoint $\tilde{x}$ of $\widetilde{X}$.

We apply Lemma 8.5 to obtain H_i-cocompact superconvex subcomplexes $\widetilde{Y}_i \subset \widetilde{X}$ and $\mathcal{K}_i$-cocompact superconvex subcomplexes $\widetilde{Z}_i \subset \widetilde{X}$, and such that each contains $\widetilde{U}$. We do this so that $g^{-1}\widetilde{Z}_j \subset \widetilde{Y}_i$ whenever $K_j^g \subset H_i$ (using the notation $x^g = g^{-1} x g$). For instance, we can first choose $\mathcal{K}_i$-cocompact superconvex subcomplexes (which are then K_i-cocompact as well), and then apply Lemma 8.5 again to ensure that the H_i-cocompact superconvex subcomplexes contain the various translates required above.

There is an upper bound $D = D(\widetilde{X})$ on diameters of wall-pieces in translates $g_j \widetilde{Z}_j$ and of cone-pieces between G-translates of any $g_i \widetilde{Y}_i$ and $g_j \widetilde{Z}_j$ with the exception of the case where $g_j^{-1}\widetilde{Z}_j \subset g_i^{-1}\widetilde{Y}_i$ whence $K_i \subset H_j^{g_i g_j^{-1}}$. (Note that this will correspond to a situation below where $g_j Z_j \subset g_i Y_i$ under the G action on X.) The diameters of wall-pieces are bounded because of the superconvexity - and this bounds the noncontiguous cone-pieces as well. The diameters of the above contiguous cone-pieces are bounded using the reasoning of Lemma 5.35, since if there is an infinite contiguous cone-piece then since K_j is a maximal intersection it would be contained in the H_i. (One sees this more explicitly in the fiber-product interpretation; since an infinite contiguous cone-piece corresponds to a noncontractible component in a $Y_i \otimes Z_j$ fiber-product, which would then necessarily be diagonal by maximality.)

The *basepoint translation length* of $g \in G$ is the combinatorial distance $d_{\widetilde{X}}(\tilde{x}, g\tilde{x})$. Note that a torsion-element g fixing the basepoint has length 0. (As $g \notin J$, the appearance of such a torsion element g will be reflected later by nontrivial conjugates of the original subgroups, and hence translates gZ_j etc. in the cubical presentation.)

A conjugate $(J \cap H_i)^g = J \cap H_i^g$ corresponds to $J \cap H_i^{jf}$ which is the stabilizer in J of $jf\widetilde{Y}_i$. Since $JU = \widetilde{X}$, the translate $f\widetilde{Y}_i$ contains the basepoint $\widetilde{x}$, and we thus choose the basepoint of $fY_i \to X$ so that it maps to the basepoint of X. Therefore $(J \cap H_i)^g = (\pi_1(fY_i))^j$.

By Proposition 12.8, for each i_a, i_b, f_a, f_b there are finitely many double cosets $\pi_1(f_aY_{i_a})j\pi_1(f_bY_{i_b})$ in J such that $\pi_1(f_aY_{i_a})^j \cap \pi_1(f_bY_{i_b})$ is infinite. Let B denote an upper bound on the basepoint translation length of a minimal length representative for each.

By residual finiteness, for each j let $\mathcal{K}_j^\circ$ denote a finite index subgroup of $(K_j \cap J) \subset \mathcal{K}_j$ such that $|Z_j^\circ| > \max(8D, 4B, \frac{D}{\alpha})$ where $Z_j^\circ = \mathcal{K}_j^\circ \backslash \widetilde{Z}_j$. Looking ahead, we note that:

(1) $> \frac{D}{\alpha}$ with $\alpha = \frac{1}{24}$ is to ensure the short innerpath property using Theorem 3.20.
(2) In particular this implies $> \frac{1}{12}D$ implying $C'(\frac{1}{12})$ and thus small-cancellation with the split-angling by Theorem 3.20.
(3) $> D'$ is to verify that $Y_i \to X$ have no missing θ-shells.
(4) $> \max(8D, 4B)$ is to satisfy the criteria of Lemma 12.22.

We now apply the construction of Theorem 12.3 to $(G, \mathcal{K}_1^\circ \subset \mathcal{K}_1, \ldots, \mathcal{K}_k^\circ \subset \mathcal{K}_k)$ to obtain finite index subgroups $\mathcal{K}_1', \ldots, \mathcal{K}_k'$ contained in $\mathcal{K}_1^\circ, \ldots, \mathcal{K}_k^\circ$ (and hence contained in J) such that $\bar{G} = G/\langle\langle \mathcal{K}_1', \ldots, \mathcal{K}_k' \rangle\rangle$ is virtually special and word-hyperbolic.

Passing to the induced quotient $\bar{J}$ of $J \subset G$ we obtain the end result of this construction: a cubical presentation of the form $\langle X \mid g\ddot{Z}_i \rangle$ where $\pi_1 X = J$ as above and each $\pi_1 \ddot{Z}_i = K_i'$, but the cubical presentation contains g translates of them where g varies over the various double cosets JgK_i'.

In particular, the group $\bar{G}$ has $\bar{J} \cong \pi_1 X^*$ as a finite index subgroup, and we shall use the associated cubical presentation to verify that each $\bar{H}_i$ is quasiconvex and that $\bar{h} = \text{Height}_{\bar{G}}\{\bar{H}_1, \ldots, \bar{H}_k\} < h = \text{Height}_G\{H_1, \ldots, H_k\}$.

For each i, let $Y_i = (H_i \cap J)\backslash \widetilde{Y}_i$ so the quasiconvexity of $\bar{H}_i$ follows by verifying the quasiconvexity of the image of the finite index subgroup $\pi_1 Y_i \to \pi_1 X$ in $\bar{J}$. However the induced cubical presentation $Y_i^* \to X^*$ has no missing θ-shells and is thus π_1-injective by Corollary 12.18. Since X^* has short innerpaths by Theorem 3.20, we are then able to obtain the quasiconvexity of each $\pi_1 \widehat{Y_i^*} \to \pi_1 X^*$ using Corollary 3.50.

We now verify the height decrease by computing the intersection of conjugates in J and $\bar{J}$:

$$J \cap \bigcap_{i=1}^{p} H_{n_i}^{g_i} = \bigcap_{i=1}^{p} (J \cap H_{n_i})^{g_i} = \bigcap_{i=1}^{p} (J \cap H_{n_i})^{j_i f_i} = \bigcap_{i=1}^{p} (\pi_1(f_i Y_{n_i}))^{j_i}.$$

Since each $|j_i| \leq B$ and $|\ddot{Z}_i| \geq |\ddot{Z}_i^o| > \max(4B, 8D)$, the criteria of Lemma 12.22 are met, and we see that $\overline{\cap_{i=1}^p \pi_1(f_i Y_{n_i})^{j_i}} = \cap_{i=1}^p \overline{\pi_1(f_i Y_{n_i})}^{\bar{j}_i}$. Since intersections of conjugates in $\bar{J}$ are images of intersections of conjugates in J itself, and since the maximally infinite such intersections project to finite subgroups of $\bar{J}$ by construction, we see that $\bar{h} < h$ as claimed.

Consequently, the theorem follows by induction on the height of $\{H_1, \ldots, H_r\}$ in G. The quotient group has trivial kernel in the base case where $h = 0$. □

12.4. **Missing θ-shells and injectivity.**

Definition 12.14 (No missing θ-shells)**.** Let $\langle A \mid B_j \rangle$ and $\langle X \mid Y_i \rangle$ be cubical presentations. A *map between cubical presentations* $f : A^* \to X^*$ is a local isometry $f : A \to X$ such that for each j there is an induced map $f : B_j \to Y_{f(j)}$ so that there is a commutative diagram:

$$\begin{array}{ccc} B_j & \to & Y_{f(j)} \\ \downarrow & & \downarrow \\ A & \to & X \end{array}$$

This data corresponds to a combinatorial map $A^* \to X^*$ sending the base to the base by a local isometry, sending cone-points to cone-points, and sending vertical 1-cells to vertical 1-cells etc.

Suppose that X^* is a small-cancellation complex for some angling system. We say $f : A^* \to X^*$ has *no missing θ-shells*, provided that the following two conditions hold:

Firstly, each essential closed curve in a component of $A \otimes_X Y_i$ lifts to some B_j.

Secondly, for any positively curved θ-shell R (so $\partial_p R = QS$ is necessarily essential) in a reduced diagram $D \to X^*$, if the outerpath Q of R lifts to A, then S is path-homotopic in Y_i to a path S' so that the lift of Q extends to a lift of $QS' = \partial_p R'$ to $\partial_p R' \to B_j$ for some j with $i = f(j)$. So that there is a commutative diagram:

$$\begin{array}{ccccc} Q & \to & A & \leftarrow & B_j \\ & \searrow & & \nearrow & \downarrow \\ & & \partial_p R' & \to & Y_i \end{array}$$

We note that the first condition is a special case of the second if we include the degenerate situation of a diagram D consisting of a single cone-cell with $\partial_p D = QS$ where Q is the outerpath and S is the trivial path. This is a degenerate example of a θ-shell whose innerpath is trivial.

Remark 12.15 (Graded Generalization)**.** In the graded case, we can generalize the conditions of Definition 12.14 as follows: The inner path $S \to Y$ is path-homotopic to $S' \to X$ through a disk diagram $D_S \to X^*_{\text{grade}(Y_i)-1}$. The outerpath Q is path-homotopic to a path $Q' \to A$ through a disk diagram $D_Q \to A^*$ where A^* is induced from $A \to X^*_{\text{grade}(A)-1}$. (For instance, it suffices that Q be path-homotopic to Q' in Y_i^*.)

As before we then require that $Q'S' = \partial_p R'$ lifts to a cone of A.

Theorem 12.16. *Let $f : A^* \to X^*$ be a map of cubical presentations, and suppose that X^* satisfies the $C(6)$ small-cancellation condition. If f has no missing θ-shells then f is π_1-injective. Moreover the map $\widetilde{A}^* \to \widetilde{X}^*$ is injective on the base.*

Proof. A path is *essential* if its lift to the universal cover is not closed. Suppose some essential path in A^* projects to a path in X^* that is not essential. Let (D, P) be a minimal complexity such example, where $P \to A^*$ is essential, but projects to a path $P \to X^*$ bounding a disk diagram $D \to X^*$, that has minimal complexity among all such (D, P).

By Theorem 3.40, D must have a cell with positive curvature. We can obviously exclude the case where D is trivial. If D has a corner of a generalized square on P, then this square lifts to A since $A \to X$ is a local isometry. This allows us to push P across the square to produce a smaller complexity counterexample (D', P'). The same holds if D has a spur, and more generally we can assume that there are no backtracks in $\partial_p D$ along the outerpath of some θ-shell, for then we could fold to expose and remove a spur. Note that one of these cases (a spur or a corner of a generalized square) must hold if D is a square diagram, when by the local-isometry hypothesis, the entire diagram D would lift to A.

Now suppose that D has some cone-cell. If D consists entirely of this cone-cell, say associated with $Y_i \to X$, then since there are no missing θ-shells, P is path-homotopic to P' in A^* and P' lifts to a closed path in some B_j, and hence $P \to A^*$ is not essential.

If D has some positively curved θ-shell R associated with $Y_i \to X$ and with outerpath Q and innerpath S, then again, since $f : A^* \to X^*$ has no missing θ-shells, we see that Q is path-homotopic to Q' in A^*, and S is path-homotopic to S' in X^* through a diagram D_S with complexity(D_S) < Grade(R), and the path $Q' \to A$ factors through a map $Q' \to B_j$ that extends to a map $Q'S' = \partial_p R' \to B_j$ such that $f(j) = i$. This contradicts the minimality of the complexity of D.

Since the cone-cell R can be replaced by the union of R' and the diagrams D_S between S and S' and D_Q between Q and Q', and since R' can be absorbed into A, we see that there is a contradiction of the minimality of the complexity of D. □

Construction 12.17 (Induced Presentation)**.** Let $\langle X \mid Y_i\rangle$ denote a cubical presentation. Let $A \to X$ be a local isometry. The *induced cubical presentation* is $\langle A \mid A \otimes_X Y_i\rangle$. We let A^* denote its associated complex. Note that there is a natural map $A^* \to X^*$.

Corollary 12.18. *Let $\langle X \mid Y_i\rangle$ be a cubical presentation satisfying small-cancellation with short innerpaths. Let $A \to X$ be a local isometry. Suppose that for each i, each component of $A \otimes_X Y_i$ is either a copy of Y_i or is a contractible complex K with diameter(K) $\leq \frac{1}{2}|Y_i|$. Let $A^* = \langle A \mid A \otimes Y_i\rangle$. Then the natural map $A^* \to X^*$ has no missing θ-shells and is thus π_1-injective.*

More generally: each component of $A \otimes Y_i$ either maps isomorphically to Y_i or is a complex K such that diameter(K) < $|Y_i^*|$ and $\pi_1 K_m^* = 1$, where $m =$ Grade$(Y_i) - 1$ and the cubical presentation K_m^* is induced from $K \to X$ and X_m^* which denotes the subpresentation containing only cones of grade < m.

Note that the generalization is arranged so that K lifts to $\widetilde{Y}_i^*$. A natural scenario of the generalization is the particular case when each component of $A \otimes Y_i$ is either a copy of Y_i or has diameter(K) < $|Y_i|$ and is either a contractible cube complex K or a copy of a cone Y_j with Grade(Y_j) < Grade(Y_i).

Proof. Consider a positively curved θ-shell R with outerpath Q and innerpath S. If the two maps $A \leftarrow Q \to Y$ determine a map $Q \to A \otimes Y$ whose image lies in a component K that is a copy of Y, then the θ-shell is not missing. Otherwise, let $Q' \to K$ denote a geodesic that is path-homotopic to $Q \to K$ in K_m^*. So $|Q'| \leq$ diameter(K) $\leq \frac{1}{2}|Y|$ where the first inequality holds since $\pi_1 K_m^* = 1$ and the second by hypothesis. Let $S' \to \widetilde{Y}^*$ be a geodesic with the same endpoints as $S \to \widetilde{Y}^*$, so S' is path-homotopic to S in Y^*, and note that $|S'| = \nabla_Y(S) < |Q'|$ by short innerpaths. But then $Q'S' \to Y^*$ is null-homotopic since $|Q'S'| = |Q'| + |S'| < |Q'| + |Q'| \leq \frac{1}{2}|Y^*| + \frac{1}{2}|Y^*|$. □

Problem 12.19. Is no missing θ-shells preserved by induced presentations? More specifically, suppose that $A \to X, B \to X, C \to X$ are local-isometries, and C is the base component of $A \otimes_X B$. Suppose that $A^* \to X^*$ and $B^* \to X^*$ have no missing θ-shells. Does $C^* \to X^*$ have no missing θ-shells?

This holds under a strong form of no missing θ-shells (for instance when X is 1-dimensional) where S lifts to A whenever $QS \to X$ is a θ-shell R in a diagram $D \to X$ with S an innerpath and Q outerpath lifting to A (and no changes on S).

The more natural weak form of no missing θ-shells allows R to be replaced by a diagram $R' \cup_S E$ where E has complexity lower than R (so all its cells have lower grade), and where $\partial_p R'$ does lift to A. I do not know if this weaker form is also preserved.

We note that if $D \to A$ is a local-isometry. and $A^* \to X^*$ has no missing θ-shells, it is possible for the composition $D^* \to X^*$ to have missing θ-shells. For instance, let $D \to A$ be the outerpath of a θ-shell.

12.4.1. ⊛ *Adding Higher Grade Relators.* We refer to Figure 152 for an example suggesting the construction in the following:

Lemma 12.20 (Adding Higher Grade Small-Cancellation Relators And Preserving Quasiconvexity)**.** *Consider $X^* = \langle X \mid \{Z_i\}\rangle$, and finite collections of compact local isometries $\{Y_j \to X\}$ and $\{W_k \to X\}$ such that for some $\alpha \leq \frac{1}{24}$:*

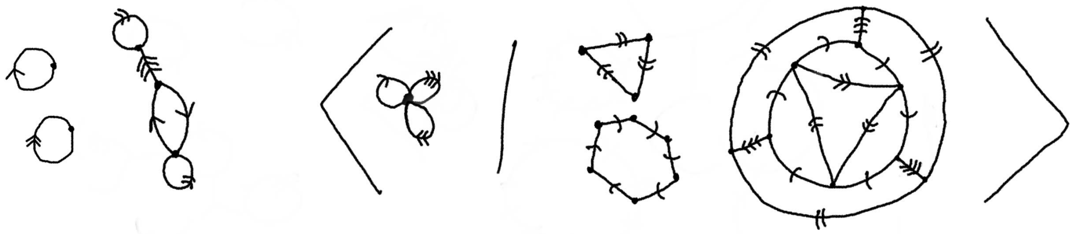

FIGURE 152. Let $\widetilde{Y}_{11}$ and $\widetilde{Y}_{12}$ and $\widetilde{Y}_{21}$ denote the universal covers of the three graphs on the left, and let $\widetilde{X}$ denote the universal cover of the bouquet of three circles. The graded cubical presentation on the right could be obtained from this initial data following Lemma 12.20.

(1) *Each $\widetilde{Y}_j$ and $\widetilde{W}_k$ is superconvex.*
(2) *If $g_j\widetilde{Y}_j \not\subset g_k\widetilde{W}_k$ then either diameter$(g_j\widetilde{Y}_j \cap g_k\widetilde{W}_k) < M$ or $g_j\widetilde{Y}_j \cap g_k\widetilde{W}_k = g_i\widetilde{Z}_i$ for some $g_i\widetilde{Z}_i$.*
(3) *If $g_i\widetilde{Z}_i \not\subset g_j\widetilde{Y}_j$ then diameter$(g_i\widetilde{Z}_i \cap g_j\widetilde{Y}_j) < \alpha|Z_i^*|$.*
(4) *If $g_i\widetilde{Z}_i \not\subset g_k\widetilde{W}_k$ then diameter$(g_i\widetilde{Z}_i \cap g_k\widetilde{W}_k) < \alpha|Z_i^*|$.*
(5) *$\{\pi_1 Y_j^*\}$ is an almost malnormal collection in $\pi_1 X^*$.*
(6) *Each $\pi_1 Y_j^*$ is residually finite.*

There exist covers $\check{Y}_j^* \to Y_j^*$ such that for any regular covers $\widehat{Y}_j^* \to Y_j^*$ that factor through $\check{Y}_j^*$ we have:

(1) *$X^{\bar{*}} = \langle X \mid Z_i, \widehat{Y}_j\rangle$ is $C'(\alpha)$.*
(2) *Each $Y_j^{\bar{*}} \to X^{\bar{*}}$ and $W_k^{\bar{*}} \to X^{\bar{*}}$ has no missing θ-shells.*
(3) *If $g_j\widetilde{Y}_j \not\subset g_k\widetilde{W}_k$ then diameter$(g_j\widetilde{Y}_j \cap g_k\widetilde{W}_k) < \alpha|\widehat{Y}_j^*|$.*

Here $\widehat{Y}_j$ denotes the cover of Y_j that is the cubical part of the cover $\widehat{Y}_j^* \to Y_j^*$. In our application, Condition (5) holds because $g_j\widetilde{Y}_j \cap g_{j'}\widetilde{Y}_{j'}$ either equals some $g_i\widetilde{Z}_i$ or has diameter uniformly bounded by L.

Proof. By Lemma 2.13, the superconvexity and cocompactness yields a uniform upperbound R on diameters of rectangles with base on each $\widetilde{Y}_j$.

Hypothesis (3) ensures that for any $\langle X \mid Z_i, \widehat{Y}_j\rangle$ a contiguous cone-piece P between a $\widehat{Y}_j$ and a Z_i satisfies $\nabla_{Z_i^{\bar{*}}}(P) < \alpha|Z_i^*| \leq \alpha|Z_i^{\bar{*}}|$ as the length of the shortest essential path does not decrease when cones are added. (In any case $Z_i^{\bar{*}} = Z_i^*$ since the newly added cones of X^* have higher grade than Z_i.)

By almost malnormality and quasiconvexity, there exists a uniform L such that diameter$(\bar{g}_j\widetilde{Y}_j^* \cap \bar{g}_{j'}\widetilde{Y}_{j'}^*) < L$ unless $\bar{g}_j\widetilde{Y}_j^* = \bar{g}_{j'}\widetilde{Y}_{j'}^*$ and $j = j'$.

By residual finiteness, we choose regular covers $\check{Y}_j^* \to Y_j^*$ such that $\max(L, M, R, \text{diameter}(Z_i)) < \alpha|\check{Y}_j^*|$.

A contiguous cone-piece P between Z_i and $\widehat{Y}_j$ satisfies: $\nabla_{\widehat{Y}_j^*}(P) \leq \text{diameter}(Z_i) < \alpha|\check{Y}_j^*| \leq \alpha|\widehat{Y}_j^*|$.

A contiguous cone-piece P between $\widehat{Y}_j$ and $\widehat{Y}_{j'}$ satisfies: $\nabla_{\widehat{Y}_j^*}(P) \leq \text{diameter}(g_j\widetilde{Y}_j^* \cap g_{j'}\widetilde{Y}_{j'}^*) \leq L < \alpha|\check{Y}_j^*| \leq \alpha|\widehat{Y}_j^*|$.

Conclusion (3) follows from Hypothesis (2) and our choice of $\check{Y}_j^*$.

$X^{\bar{*}}$ has short innerpaths by Lemma 3.48, and so Corollary 12.18 guarantees that each $W_k^{\bar{*}} \to X^{\bar{*}}$ has no missing θ-shells. Indeed, for each $\widehat{Y}_j$ the components of $W_k \otimes_X \widehat{Y}_j$ that are not copies of $\widehat{Y}_j$ (or copies of some lower grade Z_i) arise from contractible intersections having diameter $< M$ and diameter$(K) \leq$

$M < \alpha|\check{Y}_j^*| \leq \alpha|\widehat{Y}_j^*| < \frac{1}{2}|\widehat{Y}_j^*|$. Similarly, a component K of $W_k \otimes_X Z_i$ satisfies diameter$(K) \leq \alpha|Z_i^{\bar{*}}| < \frac{1}{2}|Z_i^{\bar{*}}|$. The analogous explanation shows that Corollary 12.18 applies to $Y_j^{\bar{*}} \to X^{\bar{*}}$. □

12.5. **Controlling Intersections in Quotient.**

Example 12.21. It is possible for $A \cap B = 1$ but $\bar{A} \cap \bar{B}$ to be infinite under some quotient $G \to \bar{G}$. For instance, let $G = \langle a, b \mid a = Wb\rangle$ for some small-cancellation word W. And let $\bar{G} = G/\langle\langle W\rangle\rangle$. Then usually, $\langle a\rangle \cap \langle b\rangle = 1$ but $\langle \bar{a}\rangle \cap \langle \bar{b}\rangle \cong \mathbb{Z}$. However, the intersection of the images remains trivial if we quotient by $\langle\langle W^n\rangle\rangle$ for large n. This motivates a final step in the proof of Theorem 12.1, where it is important to know that $\overline{A \cap B} = \overline{A} \cap \overline{B}$.

It is also possible for $\bar{A} \cap \bar{B}^{\bar{g}}$ to be larger than expected because of identification of conjugators under the quotient. Here is an example illustrating this:

Let $G = \langle a_1, a_2, c, d \mid c = Wd\rangle$ where W is some small-cancellation word in the generators. Let $A = \langle a_1, a_2\rangle$ and let $B = \langle a_1^c, a_2^d\rangle$. Then for most choices of W the intersection $A \cap B = 1$. Let $\bar{G} = G/\langle\langle W^n\rangle\rangle$. For most values of W, and for $n = 1$, we have $\bar{A} \cap \bar{B} = \langle \bar{a}_1, \bar{a}_2\rangle$. But for large values of n, we have $\bar{A}^{\bar{c}} \cap \bar{B} = \langle \bar{a}_1\rangle$.

Lemma 12.22. *Let $\langle X \mid Y_1, \ldots, Y_k\rangle$ be a cubical presentation satisfying small-cancellation hypotheses of Theorem 3.36 as well as the short innerpath condition of Definition 5.4. (I don't think one uses negatively curved internal cone-cells here.) Let $A_1 \to X$ and $A_2 \to X$ be based local isometries.*

Suppose the cubical presentation has "small pieces" in the sense that each cone-piece P in Y_i satisfies: $\nabla_{Y_i}(P) < \frac{1}{4}|Y_i|$, and each wall-piece satisfies $\nabla_{Y_i}(P) < \frac{1}{8}|Y_i|$. (In the graded case we use the notation $\nabla_{Y_i}(P) < \frac{1}{4}|Y_i^*|$ and $\nabla_{Y_i}(P) < \frac{1}{8}|Y_i^*|$.)

Suppose the cubical presentation has small pieces relative to A_1, A_2 in the following sense for each $1 \leq j \leq 2$ and $1 \leq i \leq k$: For each pair of lifts $\widetilde{A}_j, \widetilde{Y}_i$ to $\widetilde{X}$, the piece P between $\widetilde{A}_j, \widetilde{Y}_i$ satisfies: Either $\nabla_{Y_i}(P) < \frac{1}{8}|Y_i|$ or $\widetilde{Y}_i \subset \widetilde{A}_j$.

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

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

1. *[Double Coset Separation] $\overline{\pi_1 A_1}\bar{g}_i\overline{\pi_1 A_2}$ and $\overline{\pi_1 A_1}\bar{g}_j\overline{\pi_1 A_2}$ are distinct for $i \neq j$.*

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

2. *[Square Annular Diagram Replacement] If $\overline{\pi_1 A_1}^{\bar{g}} \cap \overline{\pi_1 A_2}$ is infinite for some $\bar{g} \in \bar{G}$, then $\overline{\pi_1 A_1}\bar{g}\overline{\pi_1 A_2} = \overline{\pi_1 A_1}\bar{g}_i\overline{\pi_1 A_2}$ for some i (which is unique by (1)).*
3. *[Intersections of Images] For each g_i we have: $\overline{\pi_1 A_1^{g_i} \cap \pi_1 A_2} = \overline{\pi_1 A_1}^{\bar{g}_i} \cap \overline{\pi_1 A_2}$.*

We will assume in the proof that each g_i is a minimal length representative for its double coset.

Proof of Claim (1). Choose a minimal complexity disk diagram $D \to X^*$ with boundary path $P = a_1 g_i a_2 g_j^{-1}$, where a_1, a_2 are closed based local geodesics in A_1, A_2 and g_i, g_j are representatives as above. If D has no cone-cells then D actually factors as a cubical diagram $D \to X \to X^*$ and hence g_i, g_j represent the same double coset in $\pi_1 X$ so $i = j$. Let us assume that this is not the case, so D contains at least one cone-cell.

Intuitively, we will now produce a maximal annular diagram B containing the boundary path of D such that $B \to X$ is a cubical diagram, and B is in the local convex hull of the boundary path within D, in the sense that we obtain B by continually adding squares of D whose corners are already present, and adding closed edges corresponding to spurs.

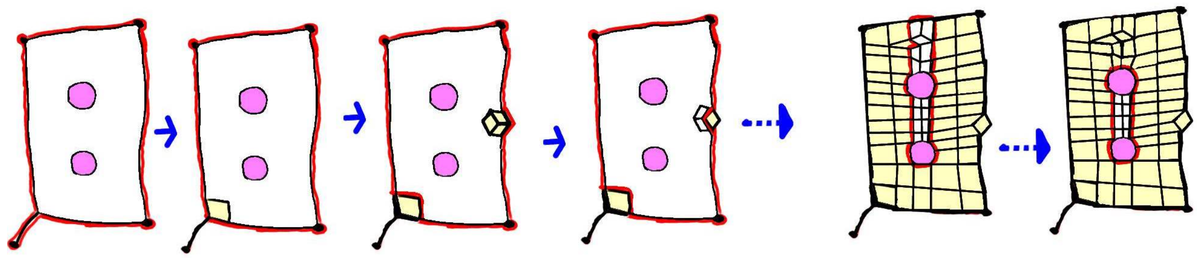

FIGURE 153.

More precisely let us begin with $E_0 = D$ and $Q_0 = \partial_p D$. We will construct a sequence of paths Q_i bounding disk diagrams E_i that are contained in slight modifications D_i of D. See Figure 153.

For each $i \geq 0$, either:

(1) Q_i contains the outerpath of a spur of E_i.
(2) Q_i contains the outerpath of a generalized square of E_i.
(3) Q_i contains neither.

In the first case, we remove the terminal 0-cell and 1-cell of the spur from E_i to form E_{i+1} and we remove this backtrack to form Q_{i+1}. We continue to repeat the first case until there are no such backtracks remaining. We then examine the second possibility.

In the second case, we first adjust the interior of E_i to obtain E_i' (and thus adjust the ambient diagram D_i to obtain D_i') so that the square actually lies along Q_i, we then remove from E_i' the open square and the corner consisting of the two open 1-cells and the open 0-cell. The path Q_{i+1} is obtained from Q_i by pushing across that square so that we replace the two edges from the outerpath corner by the opposite two edges. Note that if the outerpath of the square meeting Q_i is longer, then Q_{i+1} will obtain backtracks which we will omit when we return to the first case. After performing the second case once, we return to the first.

Since the number of cells in E_i is decreasing, after t steps we terminate at E_t and $Q_t = \partial_p E_t$. Observe that E_t contains all cone-cells originally in D, and in particular E_t is nontrivial provided D contains at least one cone-cell.

At each stage, let B_i denote the annular subdiagram in D_i' bounded by P and Q_i. We think of the diagram B_i as the "local convex hull" of ∂P in D, but this is misleading as some cubical replacement moves were necessary to create it, and D is evolving.

It is important to note that at each stage, and in particular, in the final result, each dual curve in B_i emerging from an edge of Q_i terminates on P. Indeed, this holds initially, and is preserved by each of the two replacement moves indicated above.

By construction E_t cannot have the corner of a generalized square on its boundary path, nor can it have a spur. As was explained for B_i above, each dual curve emerging from a 1-cube on the boundary path of E_t must terminate on a 1-cell of the boundary of D.

Now apply Theorem 3.40 to see that either E_t is a single cone-cell, or E_t has at least two positive curvature θ-shells.

A dual curve emanating from an edge in Q_t cannot cross itself in B_t before terminating on P. Indeed, after crossing itself inwards it must eventually cross itself back outwards since it terminates on P. So we see that there is either a monogon or bigon dual curve subdiagram in B_t and hence a way to reduce the area, which is impossible. See the first diagram in Figure 154.

Now consider the family of dual curves emanating from the outerpath on Q_t of a single θ-shell R of E_t. As explained above, each such dual curve terminates on P. We claim that they cannot cross each other, and are thus splayed as in the second diagram in Figure 154.

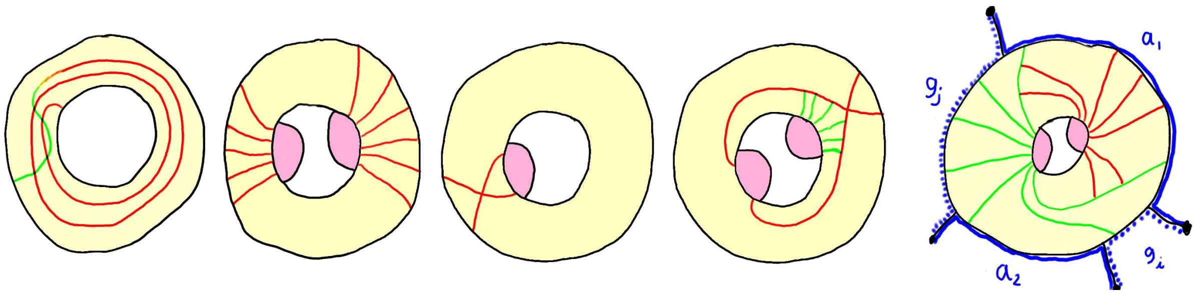

FIGURE 154.

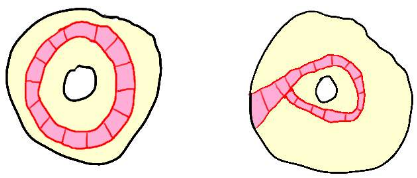

FIGURE 155. Rectangles can self-intersect in the rectified annuli

Indeed, if two such curves crossed each other on the same side of the annulus as in the third diagram in Figure 154 then there would be an outerpath of a generalized square in B_t on R itself. This yields a complexity reduction by absorbing the square into the cone of R. This is impossible by the minimal area of D. (We say that a dual curve emanating from the θ-shell R crosses itself on the "same side" as R if it crosses itself before crossing a dual curve of the other θ-shell R'.)

We note that this argument shows that if E_t is a single cone-cell then its emanating dual curves are splayed.

Let us now assume there are at least two shells, say R and R' (of course, there are at least three if E_t is not a ladder). Two such dual curves emanating from R cannot cross by going around to the other side of the annulus B_t as in the fourth diagram in Figure 154. For then, all the dual curves emanating from the outerpath of R' would cross one or the other of these dual curves. We thus see that the outerpath of R' is in the union of at most two wall-pieces. Each hyperplane piece H in Y satisfies $\nabla_Y H < \frac{1}{4}|Y|$, so the outerpath would not be at least half the length of ∂R. This contradicts Condition 5.4 (for instance), unless ∂R is not essential in Y, in which case D was not of minimal area.

In a similar manner, we see that two dual curves leaving the outerpath of the shell R cannot travel around the annulus and end on the same subpath a_1, a_2, g_i, g_j of P. See the fifth diagram in Figure 154 where the curves both end on a_1. For then consider another shell R', and note that the dual curves from its outerpath end entirely on two dual curves together with a_1, and hence the outerpath has length $< (\frac{1}{8} + \frac{1}{8} + \frac{1}{8})|\partial R'|$ which contradicts Condition 5.4.

We conclude that the dual curves from the outerpath of each shell lie in a sequence of at most four "pieces" in a_1, a_2, g_i, and g_j. Since by hypothesis each such "piece" is $< \frac{1}{8}$ of ∂R, we see that the outerpath of R is $< \frac{1}{2}$ of $|\partial R|$ which contradicts Condition 5.4.

A similar argument works to reach a contradiction when E_t is a single cone-cell. Indeed in this case it suffices to see that each "piece" with one of a_1, a_2, g_i, g_j has length $< \frac{1}{4}|Y|$. □

Proof of Claim (2). Let $D \to X^*$ be an annular diagram for $P_1^Q = P_2$ where $P_i \to A_i$ are based paths and P_1, P_2 are boundary paths of D and Q is a path in D joining the basepoints, and Q is homotopic to a based path g in X that projects to $\bar{g}$, and P_i represent infinite order elements in $\pi_1 X^*$.

We can compress replaceable cone-cells and absorb cancellable pairs while preserving the element representing the conjugating path in $\bar{G}$. Similarly, can also push P_1, P_2 past any corners of generalized

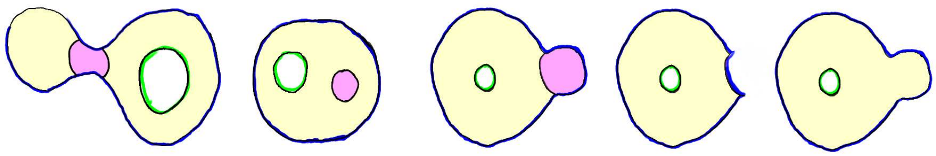

FIGURE 156.

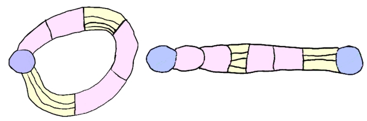

FIGURE 157. An annular diagram with a doubly external cone-cell must be thin, since cutting along that cone-cell must yield a ladder.

squares, which could be absorbed into A_1, A_2 since $A_i \to X$ are local isometries, and so no P_i contains the outerpath of a corner of a generalized square in D. We can also assume that D is *minimal* in the sense that none of its cone-cells can be absorbed into A_i through P_i on either side of ∂D. (It is here that our hypothesis on $A_i \to X$ are applied.)

We pass to a rectified annular diagram as in Section 3.6. So we can assume that D has no cancellable pairs, and by local convexity, there are no corners of generalized squares on either boundary path. We also assume that D is *minimal*, in the sense that its cone-cells cannot be replaced by square diagrams, and cannot be absorbed into either boundary path.

A cone-cell of D whose boundary path intersects ∂D in more than two components is negatively curved. Consider a cone-cell R whose boundary path intersects P_i in exactly two subpaths as in the first diagram in Figure 156. This cone-cell subtends a subdiagram of D containing R that is a disk diagram containing at least two positively curved cells, and we would contradict minimality (as any such cell could be absorbed into A_i). Finally, consider a cone-cell whose boundary path intersects ∂D in a single subpath of P_i. If it is a (nonnegatively curved) θ-shell then by hypothesis, since it cannot be replaced by a square diagram, it can be absorbed into A_i thus contradicting minimality. We conclude that each cone-cell in D must intersect each of P_1 and P_2 in a single subpath of its boundary, and so D has the form illustrated on the right in Figure 158. Thus D has the property that all of its cone-cells are external, and hence it has the form illustrated on the right in Figure 158, though it might not be singular.

We conclude that if D has a cone-cell, then it must be a width 1 annuladder. (Prove this by cutting out one cone-cell as in Figure 157.) However, this is ruled out by a count on the lengths of boundary paths of cone-cells, and the lengths of their pieces with each other and with A_1, A_2. Indeed, each overlap with the boundary projects to a length of $< \frac{1}{4}$ of the girth (of its associated cone), and cone-pieces or wall-pieces likewise project to length $< \frac{1}{4}$ of the girth.

The only remaining possibility is that D is a square diagram. But then we have shown that Q is path homotopic in X^* to the element Q' such that $P_1^{Q'} = P_2$ in $\pi_1 X$ where P_i have infinite order. Consequently, Q' is in the same double coset class as one of our finitely many representatives g_i.

Finally, this double coset class representative is uniquely determined by Claim (1). □

Proof of Claim (3). It is immediate that $\overline{\pi_1 A_1^{g_i} \cap \pi_1 A_2} \subset (\overline{\pi_1 A_1}^{\bar{g}_i} \cap \overline{\pi_1 A_2})$ so we must only verify that $\overline{\pi_1 A_1^{g_i} \cap \pi_1 A_2} \supset (\overline{\pi_1 A_1}^{\bar{g}_i} \cap \overline{\pi_1 A_2})$.

For each element represented by a closed path $P_2 \to A_2$ in the intersection $\overline{\pi_1 A_1}^{\bar{g}} \cap \overline{\pi_1 A_2}$, there is a closed based path $P_1 \to A_1$ and a closed based path $Q \to X$, such that there is an annular diagram $D \to X^*$ for $P_1^Q = P_2$. As in the proof of Claim (2), there is a sequence of removals and absorptions of cancellable pairs, and absorptions of θ-shells to the boundary of the annulus (or replacements of these by square diagrams). We can pass to a reduced and minimal (in the sense of Claim (2)) annular diagram

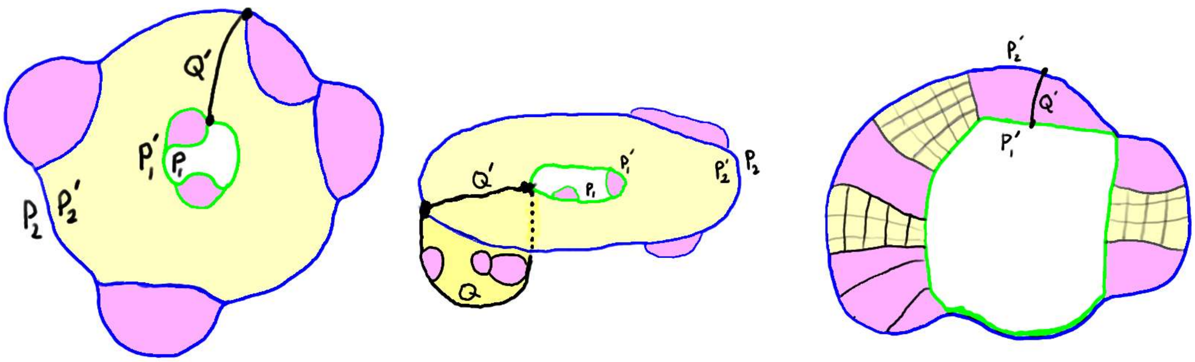

FIGURE 158. Before and after: The middle figure illustrates a typical situation where an original annulus is turned into a square annulus by pushing out cone-cells through the two bounding circles and pushing the conjugator out (so it is connected to the new conjugator by a "flap" consisting of a disk diagram in X^*). No extra flap is illustrated on the left diagram - were $Q = Q'$. The diagram on the right represents the final impossible annuladder we consider.

$D' \to X$ for ${P'_1}^{Q'} = P'_2$ where each P'_i is a closed path in A_i that is path homotopic to P_i in X^* (and actually in A_i^*), and likewise $Q' \to X$ is path-homotopic to Q in X^*.

As concluded in the proof of Claim (2), $D' \to X$ is a square annular diagram, and so we have found an element g' represented by Q' such that $\overline{g'} = \overline{g}$ and $\overline{P}_2 \in \overline{\pi_1 A_1^{g'} \cap \pi_1 A_2}$, as P_2 is path-homotopic to P'_2 in X^* (and actually in A_2^*), and D' demonstrates that $P'_2 \in \pi_1 A_1^{g'} \cap \pi_1 A_2$.

If P_2 represents a nontrivial element in $\pi_1 X^*$, then it represents a nontrivial element in $\pi_1 X$, and hence so does P'_2. But then the conjugacy ${P'_1}^{Q'} = P'_2$ shows that Q' lies in one of the double cosets $\pi_1 A_1 g_i \pi_1 A_2$ in $\pi_1 X$. The images in $\pi_1 X^*$ of these double cosets are distinct by Claim (1). However, since each such Q' is path-homotopic in X^* to our conjugator Q (representing $\bar{g}$), we see that Q' maps to the same element $\bar{g}$ in $\pi_1 X^*$. Thus, for any P_2 chosen above representing a nontrivial element in $\pi_1 X^*$ in the intersection, since the associated conjugator Q' represents an element mapping to $\bar{g}$, Claim (1) ensures that each Q' lies in the same double coset $\pi_1 A_1 g_f \pi_1 A_2$ for some fixed f.

In conclusion, we have thus found that each nontrivial element in the intersection of the images in $\pi_1 X^*$ (initially represented by $P_1^Q = P_2$), is the image of an element (represented by $(P'_1)^{Q'} = P'_2$) in the intersection within $\pi_1 X$. And moreover, each element Q' lies in the same double coset $\pi_1 A_1 g_f \pi_1 A_2$ in $\pi_1 X$. Since elements in the same double coset correspond to the same intersection of subgroups, we reach our desired conclusion:

$$\text{Nontrivial Elements}\{\overline{\pi_1 A_1}^{\bar{g}} \cap \overline{\pi_1 A_2}\} \subset \overline{\pi_1 A_1^{g_f} \cap \pi_1 A_2}. \qquad \square$$

13. AMALGAMS OF VIRTUALLY SPECIAL GROUPS

13.1. **Virtually Special Amalgams.**

Theorem 13.1 (Special Hyperbolic Amalgam)**.** *Let $G = A *_C B$ be an amalgamated product where G is word-hyperbolic and C is quasiconvex. If A, B are virtually compact special then so is G. A similar result holds for graphs of groups, and in particular: Let $G = A*_C$ be an HNN extension such that G is word-hyperbolic and C is quasiconvex in G. If A is virtually compact special then so is G.*

Remark 13.2. Following the lead that Theorem 13.1 holds when C is already malnormal by Theorem 11.2, it is attractive to attempt to deduce it from Theorem 11.2 by passing to a finite index subgroup of G whose induced graph of groups has the property that each of its edge groups is malnormal. In fact, this holds when C is separable.

Theorem 13.3 (Word-hyperbolic $\mathcal{QVH}$ special)**.** *Every word-hyperbolic group in $\mathcal{QVH}$ is virtually special.*

Proof. This holds by induction on the length of a hierarchy by applying Theorem 13.1. □

We use the notation $x^y = yxy^{-1}$ here but warn the reader that $x^y = y^{-1}xy$ is used in other places in the paper.

Proof of Theorem 13.1. **Reduction to HNN case:** There is an embedding $A *_C B \to D*_C$ where $D = A * B$ and the HNN extension $D*_C$ has C embedded in the first factor of $A * B$, and has C^t embedded in the second factor. This embedding is quasiconvex, and so C is itself quasiconvex in D. We claim that if A, B are word-hyperbolic and virtually compact special then so is $D = A * B$. We recommend that the reader prove this using a simple covering space argument. Moreover, if $D*_C$ is virtually compact special, then so is the quasiconvex subgroup $A *_C B$. Thus the AFP case follows from the HNN case.

Note that if G' is a torsion-free finite index subgroup, then its induced graph of groups decomposition has the required properties, and the virtual specialness of G' implies the virtual specialness of G.

Strategy of proof: Consider $G = A*_{C^t=D}$ where A is virtually special, and G is word-hyperbolic and C is quasiconvex in G. Let $h = \text{Height}_G(C)$. We will form a quotient $G = A*_{C^t=D} \to \bar{A}*_{\bar{C}^{\bar{t}}=\bar{D}} = \bar{G}$ where $\bar{A}$ is virtually special, $\bar{G}$ is word-hyperbolic and $\bar{C}$ is quasiconvex, and $\bar{h} = \text{Height}_{\bar{G}}(\bar{C}) < h$. The group $\bar{G}$ will therefore be virtually special by induction on height. The virtual specialness of the base case where $h = 0$ so C is finite was treated above. (Note that the case $h = 1$ where the edge group is almost malnormal is in Theorem 11.2.)

The quotient $G \to \bar{G}$ will also be chosen to have the property that each intersecting conjugator k_i of C has $\bar{k}_i \notin \bar{C}$, where k is an *intersecting conjugator* if $C^k \cap C$ is infinite but $k \notin C$. Note that there are finitely many C cosets of intersecting conjugators by Proposition 12.8.

Since $\bar{G}$ is virtually compact special and word-hyperbolic, and $\bar{C}$ is quasiconvex, there is a finite quotient $\bar{G} \to Q$ such that $\bar{C}$ is separated from each $\bar{k}_i$ in Q, and hence $\phi(C)$ is separated from each $\phi(k_i)$ in the resulting composition quotient $G \overset{\phi}{\to} Q$. Let G' be the kernel of ϕ and observe that each edge group of the induced splitting of G' is almost malnormal. Consequently, G' is virtually special by Theorem 11.2.

The collection of tree stabilizers: For each i, let $\{T_{ik} : k \in J_i\}$ denote representatives of G-orbits of the finitely many i-edge subtrees with infinite pointwise stabilizer $\text{Fix}(T_{ik})$. Let $B_{ik} = \text{Stab}(T_{ik})$. Each B_{ik} acts on T_{ik} with a fixed point, and we choose one fixed vertex and declare it to be the *root* of T_{ik}. Without loss of generality we assume each T_{ik} is positioned so that its root lies at the base vertex v of T. Note that T_{ik} already has the structure as a directed graph, since it is a subgraph of the directed graph T. However, T_{ik} has an additional structure where we regard the root as the *highest* vertex and where each edge *ascends* from its lower vertex to its higher vertex (towards the root). Accordingly, for each edge e of T_{ik} ascending from the vertex a to the vertex b we have $\text{Stab}_{B_{ik}}(a) = \text{Stab}_{B_{ik}}(e)$ and $[\text{Stab}_{B_{ik}}(b) : \text{Stab}_{B_{ik}}(e)] < \infty$. For each k we let $S_{ik} \subset T_{ik}$ denote a (fundamental domain) subtree consisting of unique representatives of the B_{ik}-orbits of vertices and edges of T_{ik}. One constructs S_{ik} by beginning with the root, and at each stage, including unique representatives of the next highest closed edges that meet vertices already included. For each $k \in J_i$ let $L_{ik} \subset G$ index the vertices of S_{ik}, so each vertex of S_{ik} equals xv for a unique $x \in L_{ik}$. For notational comfort we index the root vertex by $1 = 1_G$. We let H_{ikx} denote the subgroup of A whose conjugate H_{ikx}^x equals $A^x \cap B_{ik}$ which equals $\text{Stab}_{B_{ik}}(xv)$.

As explained above, if xv ascends to yv in S_{ik} then H_{ikx}^x is a finite index subgroup of H_{iky}^y. Hence, conjugation by the element xy^{-1} induces a finite-index inclusion $H_{ikx}^{xy^{-1}} \hookrightarrow H_{iky}$.

The almost malnormal collection in A and the quotient $A \to \bar{A}$:

Note that $B_{hk} = \mathbb{C}_G(\mathrm{Fix}(T_{hk}))$. The collection $\{B_{hk} : k \in J_h\}$ is almost malnormal in G since it is a subcollection of a collection that is almost malnormal by Lemma 12.12. Consequently, $\{H_{hkx} : k \in J_h, x \in L_{hk}\}$ is an almost malnormal collection of subgroups of A by Lemma 12.13.

We will apply Theorem 12.3 to A where the role of $\{H_1, \dots, H_r\}$ is played by $\{H_{hkx}\}$. For each k, the constraint subgroups $\{\ddot{H}_{hkx}\}$ supplied by Theorem 12.3, embed as finite index subgroups of the root H_{hk1}. Indeed there is a sequence of monomorphisms following the ascending path in S_{hk} from xv to $1v$. Likewise, the preference subgroups H°_{hkx} ascend to finite index subgroups of H_{hk1}. We can thus choose a finite index subgroup $H^\bullet_{hk1}$ of H_{hk1} that is contained in the intersection of all these constraint and preference subgroups. We choose H'_{hk1} to be a subgroup of $H^\bullet_{hk1}$ that is normal in H_{hk1}, and this descends to a subgroup H'_{hkx} for each x. For each directed edge (xv, yv) in S_{hk}, we thus have:

$$(H'_{hkx})^{xy^{-1}} = H'_{hky}.$$

The torsion-free finite index subgroup $A^{♀}$: The group A acts properly and cocompactly on a CAT(0) cube complex $\widetilde{X}$. Let $A^{♀}$ denote a torsion-free finite index normal subgroup of A, and we let $X = A^{♀}\backslash\widetilde{X}$. The subgroup $A^{♀}$ will facilitate our cubical small-cancellation theory computations below. We let $U = (C \cap A^{♀})\backslash\widetilde{U}$ and $V = (D \cap A^{♀})\backslash\widetilde{V}$. Note that it is possible that $(\pi_1 U)^t \neq \pi_1 V$.

The preference subgroups: We will choose the subgroups $H^\circ_{hkx} \subset H_{hkx}$ to be small enough that:

(1) each H°_{hkx} lies in $A^{♀}$.
(2) cosets corresponding to the transition elements (i.e. between stable letters) in the normal forms of intersecting conjugators in G of the edge group C are separated by $A \to \bar{A}$.
(3) double cosets $H_{ikx}aH_{\ell jy}$ corresponding to intersecting conjugators in A of the various H_{ikx} are separated by $A \to \bar{A}$ where $1 \leq i \leq h,\ k \in J_i,\ x \in L_{hk}$ and $1 \leq \ell \leq h,\ j \in J_\ell, y \in L_{\ell j}$.
(4) $\bar{C}$ and $\bar{D}$ are quasiconvex in $\bar{A}$.

The conjugation isomorphism $C^t = D$ projects to $\bar{C}^{\bar{t}} = \bar{D}$: We now show how to verify that the kernel of $C \to \bar{C}$ maps isomorphically to the kernel of $D \to \bar{D}$ under the map $C \to D$, and hence the conjugation isomorphism $C^t = D$ projects to $\bar{C}^{\bar{t}} = \bar{D}$. We begin with some notation: Let (xv, yv) denote an edge of S_{hk} that is directed from xv to yv. The element $x^{-1}y$ is of the form $\alpha t\beta$ for some $\alpha, \beta \in A$ where $\alpha = \alpha(x, y)$ and $\beta = \beta(x, y)$. There are two points here:

Firstly, as we shall later confirm, we have the following presentations for $\bar{C}$ and $\bar{D}$.

$$\bar{C} = C/\langle\langle (H'_{hkx})^{\alpha} \ :\ k \in J_h,\ (xv, yv) \in \text{Di-Edges}(S_{hk}),\ \alpha = \alpha(x, y)\rangle\rangle$$
$$\bar{D} = D/\langle\langle (H'_{hky})^{\beta^{-1}} \ :\ k \in J_h,\ (xv, yv) \in \text{Di-Edges}(S_{hk}),\ \beta = \beta(x, y)\rangle\rangle$$

Secondly, the map $C \to D$ induced by conjugation by t, sends $(H'_{hkx})^{\alpha}$ to $(H'_{hky})^{\beta^{-1}}$. Indeed, $(H_{hkx})^{xy^{-1}} \subset H_{hky}$ holds by construction, and our "descending choice" of the H'_{hkx} ensures that: $(H'_{hkx})^{xy^{-1}} = H'_{hky}$. Since $xy^{-1} = \alpha t\beta$ we have: $((H'_{hkx})^{\alpha})^t = ((H'_{hkx})^{xy^{-1}})^{\beta^{-1}} = (H'_{hky})^{\beta^{-1}}$. Thus the conjugation isomorphism $C^t = D$ maps the generators of kernel$(C \to \bar{C})$ isomorphically to the generators of kernel$(D \to \bar{D})$. Thus $C^t = D$ projects to an isomorphism $\bar{C}^{\bar{t}} = \bar{D}$ as claimed.

The associated superconvex subcomplexes: Let $\{H_{ikx} : x \in L_{ik}\}$ denote representatives of the conjugacy classes in A of intersections with conjugates of B_{ik}. Specifically, each $H_{ikx} = B_{ik}^{x^{-1}} \cap A$ for some $x \in G$. Each H_{ikx} acts properly and cocompactly on a superconvex basepoint containing subcomplex $\widetilde{Y}_{ikx}$. Choosing $H_{111} = C$ and $H_{11t} = D$ we let $\widetilde{U} = \widetilde{Y}_{111}$ and $\widetilde{V} = \widetilde{Y}_{11t}$. For $i > i'$ if H_{ikx} is conjugate into $H_{i'k'x'}$ then the corresponding translate of $\widetilde{Y}_{ikx}$ lies in $\widetilde{Y}_{i'k'x'}$. This property can be arranged by making the choices with $i = h$ first and then proceeding towards $i = 1$ to make the choices inclusively at each stage. There is a constant M such if diameter$\left(a\widetilde{Y}_{ijk} \cap a'\widetilde{Y}_{i'j'k'}\right) > M$ then this intersection of A-translates equals some $a''\widetilde{Y}_{i''j''k''}$ with $i'' \geq i, i'$ and moreover if $i'' = i$ then

$a''\widetilde{Y}_{i''j''k''} = a\widetilde{Y}_{ijk}$ and similarly for $i'' = i'$. There is a constant R bounding the diameters of base rectangles on these superconvex subcomplexes of $\widetilde{X}$.

Computing presentations for $\bar{C}$ and $\bar{D}$: Assume each $H^{\circ}_{hjk} \subset H^{♀}_{hjk}$ and assume each H°_{hjk} has $|H^{\circ}_{hjk} \backslash \widetilde{Y}_{hjk}| > 24\max(R, M)$. Under these assumptions, we verify that our claimed presentations for $\bar{C}$ and $\bar{D}$ are correct. The normal subgroup $N = \ker(A \to \bar{A})$ we obtain is $N = \langle\langle \pi_1 \widehat{Y}_{hjk} \rangle\rangle$ and is contained in the normal subgroup $A^{♀}$ of A since each $\pi_1 \widehat{Y}_{hjk} = H'_{hjk} \subset H^{\circ}_{hjk} \subset A^{♀}$. The presentation $X^* = \langle X \mid a\widehat{Y}_{hjk} : a \in A\rangle$ accurately determines the presentation for $\bar{A}^{♀}$. Indeed, $\langle\langle \pi_1 \widehat{Y}_{hjk}\rangle\rangle_A$ is a subgroup of $A^{♀}$ since each $\pi_1\widehat{Y}_{hjk}$ is, and moreover it is the smallest normal subgroup of $A^{♀}$ containing each $\pi_1\widehat{Y}_{hjk}$ that is also invariant under the A-action by conjugation. Accordingly, $\langle\langle \pi_1\widehat{Y}_{hjk}\rangle\rangle_A = \langle\langle (\pi_1\widehat{Y}_{hjk})^a : a \in A\rangle\rangle_{A^{♀}}$.

The key to computing $C \cap N$ is to first observe that $C \cap N = C \cap (A^{♀} \cap N) = (C \cap A^{♀}) \cap N = \pi_1 U \cap N$. Applying Corollary 12.18, we use no missing θ-shells to compute an induced presentation U^* from $U \to X$ and X^*. It is of the form: $U^* = \langle U \mid c\widehat{Y}_{hjk} : c \in C\rangle$. And consequently $C \cap N = \langle\langle \pi_1\widehat{Y}_{hjk}\rangle\rangle_C$. Likewise $D \cap N = (D \cap A^{♀}) \cap N = \pi_1 V \cap N$, and there is an induced presentation $V^* = \langle V \mid d\widehat{Y}_{hjk} : d \in D\rangle$, and we see that $D \cap N = \langle\langle \pi_1\widehat{Y}_{hjk}\rangle\rangle_D$.

Separating intersecting conjugators: By Proposition 12.8, there are finitely many double cosets Cg_sC in G such that $C^{g_s} \cap C$ is infinite. Each representative g_s can be chosen to be of the form: $g_{s0}t^{\epsilon_{s1}}g_{s1}t^{\epsilon_{s2}} \dots t^{\epsilon_{sr_s}}g_{sr_s}$ where $g_{sk} \notin C$ if $\epsilon_{sk} = -1, \epsilon_{s(k+1)} = 1$, and $g_{sk} \notin D$ if $\epsilon_{sk} = 1, \epsilon_{s(k+1)} = -1$, and $g_{s0}, g_{sr_s} \notin C$. Since C and D are separable subgroups of A, there is a finite index normal subgroup $A^{\diamond} \subset A$ such that $g_sC \not\subset A^{\diamond}C$ and $g_sD \not\subset A^{\diamond}D$ as appropriate. We add the intersections $H^{\diamond}_{hjx} = H_{hjx} \cap A^{\diamond}$ to our list of preferences. Consequently, $\bar{g}_s \notin \bar{C}$ and likewise $\bar{g}_s \notin \bar{D}$ etc. and thus in the quotient $A*_{C^t=D} \to \bar{A}*_{\bar{C}^{\bar{t}}=\bar{D}}$, each normal form maps to a normal form of the same length. Thus length ≥ 1 intersecting conjugators will have length ≥ 1 images, and length 0 intersecting conjugators are also separated from $\bar{C}$.

Reduction in height: There is an equivariant map $T \to \bar{T}$ from the G-action on the Bass-Serre tree T of $A*_C$ to the $\bar{G}$ action on the Bass-Serre tree $\bar{T}$ of $\bar{A}*_{\bar{C}}$. Consider the finitely many subgroups $\{H_{ikx} : 1 \leq i \leq h, k \in J_i, x \in L_{hk}\}$. By Proposition 12.8, there are finitely many double cosets $H_{ikx}aH_{\ell jy}$ in A of intersecting conjugators with the property that $H_{ikx} \cap H^a_{\ell jy}$ is infinite. We will show below that for suitable $A \to \bar{A}$, if an intersection $\bar{H}_{ikx} \cap \bar{H}^{\bar{a}}_{\ell jy}$ is infinite, then $\bar{H}_{ikx}\bar{a}\bar{H}_{\ell jy}$ is the image of one of the finitely many $H_{ikx}aH_{\ell jy}$ above. Consequently any finite subtree $\bar{F} \subset \bar{T}$ with infinite pointwise stabilizer can be lifted to a finite subtree $F \subset T$ with infinite pointwise stabilizer and $\text{Fix}_{\bar{G}}(\bar{F}) = \overline{\text{Fix}_G(F)}$. Consequently any such subtree $\bar{F}$ has at most h edges, since F does as $h = \text{Height}_G(C)$. Moreover, if F has h edges and infinite pointwise stabilizer then $\text{Stab}_G(F)$ is conjugate to B_{hk} for some k, and so $\text{Stab}_{\bar{G}}(\bar{F})$ is finite.

Towards height control: We will choose $A \to \bar{A}$ so that $\overline{H_{ikx}} \cap \overline{H_{\ell jy}}^{\bar{a}}$ is commensurable with $\overline{H_{ikx} \cap H^a_{\ell jy}}$ for each such double coset. Indeed, the commensurabilities below are immediate. The equality holds by Lemma 12.22.(3) applied to $A^{♀} \to \bar{A}^{♀} = \pi_1\langle X \mid \widehat{Y}_{hkx}\rangle$ with respect to $\{H^{♀}_{ikx} : i < h\}$. (The extra preference is that each $|\widehat{Y}_{hkx}| > 8$ times the length of the shortest representatives of the various intersecting conjugators.)

$$\overline{H_{ikx}} \cap \overline{H_{\ell jy}}^{\bar{a}} \sim \overline{H^{♀}_{ikx}} \cap \overline{H^{♀}_{\ell jy}}^{\bar{a}} = \overline{H^{♀}_{ikx} \cap (H^{♀}_{\ell jy})^a} \sim \overline{H_{ikx} \cap H^a_{\ell jy}}.$$

Moreover each intersecting conjugator in $\bar{A}$ is the image of an intersecting conjugator in A in the sense that if $\overline{H_{ikx}} \cap \overline{H_{\ell jy}}^{\bar{b}}$ is infinite then $\bar{b} = \overline{a'}$ for some a' in a double coset $H_{ikx}aH_{\ell jy}$ of intersecting conjugators. For this we consider the larger finite set of subgroups $\{(H^{♀}_{ikx})^g\}$ where g varies over representatives of the left cosets $\{gA^{\female}\}$ in A. We apply Lemma 12.22.(2) to a full set of double cosets $\{(H^{\female}_{ikx})^g q (H^{\female}_{\ell jy})^f\}$ representing intersecting conjugators: Observe that for some chosen coset representative g and some $\beta \in A^{\female}$ we have:

$$\overline{(H^{\female}_{ikx})} \cap \overline{(H^{\female}_{\ell jy})}^{\bar{b}} = \overline{(H^{\female}_{ikx})} \cap \overline{(H^{\female}_{\ell jy})}^{\overline{g\beta}} = \overline{(H^{\female}_{ikx})} \cap \overline{(H^{\female}_{\ell jy})^g}^{\bar{\beta}}$$

Thus by Lemma 12.22.(2), $\bar{\beta} = \overline{c'}$ for some $c' \in H^{\female}_{ikx} c (H^{\female}_{\ell jy})^g$ where c and hence c' is an intersecting conjugator so $H^{\female}_{ikx} \cap ((H^{\female}_{\ell jy})^g)^{c'}$ is also infinite. Consequently $H_{ikx} \cap ((H_{\ell jy})^{gc'})$ is infinite so $a' = gc' \in H_{ikx} a (H_{\ell jy})$ for one of the known intersecting conjugators a. Thus $\bar{b} = \overline{a'}$ as claimed.

Hyperbolicity of $\bar{G}$ and Quasiconvexity of $\bar{C}$: Observe that $\bar{A}$ is hyperbolic as indicated in Theorem 12.3, and $\bar{C}, \bar{D}$ are quasiconvex in $\bar{A}$ since $\bar{C}^{\female}$ and $\bar{D}^{\female}$ are quasiconvex in $\bar{A}^{\female}$ by Corollary 3.50. Consequently the finite height of $\bar{C}$ in $\bar{G}$ implies that $\bar{G}$ is hyperbolic by the "no long annulus" criterion of [BF92], and then $\bar{C}$ is quasiconvex in $\bar{G}$ by [Mit04]. □

Remark 13.4. An alternative way to structure the proof is to note that the procedure can also be iterated until, at the last step, C and D are quotiented to a finite group, and we have separated C from the intersecting conjugators in a group that is an HNN extension of a virtually-compact special hyperbolic group along a finite subgroup. This accomplishes the objective which enables the strategy - but without induction on height to know that $\bar{G}$ is virtually special. Instead the onus is on a strong form of the special quotient theorem.

Problem 13.5. Let G act properly on a CAT(0) cube complex. Is $G \cong \pi_1 X^*$ where X^* is a small-cancellation cubical presentation? If G is word-hyperbolic and acts properly and cocompactly, can X^* be chosen so that $\widetilde{X}$ is δ-hyperbolic? Perhaps the canonical associated cube of spaces or some variation of it works. The grading should be in reverse order of the dimension.

14. Hyperbolic 3-manifolds with a geometrically finite incompressible surface

Let M be a finite volume hyperbolic 3-manifold with an incompressible geometrically finite 2-sided surface S. There exists a (topological) hierarchy for M which begins by cutting the 3-manifold along S and proceeds by cutting along further incompressible surfaces until only balls remain. It is known that geometrical finiteness is equivalent to quasiconvexity (even when there are cusps) [Hru10]. All further cuts in the hierarchy correspond algebraically to splittings along quasiconvex subgroups, as it is a theorem of Thurston's that finitely generated subgroups of a geometrically finite infinite volume group are themselves geometrically finite [Can94]. In the hyperbolic case, the 3-manifold hierarchy induces a quasiconvex hierarchy of $\pi_1 M$, and we thus have:

Theorem 14.1. *Let M be a closed hyperbolic 3-manifold with an incompressible surface S such that $\pi_1 S \subset \pi_1 M$ is geometrically finite. Then $\pi_1 M$ has a finite index subgroup that is the fundamental group of a compact special cube complex.*

A group G is *residually finite rational solvable* (RFRS) if there is a decreasing sequence of finite index normal subgroups $G = G_0 > G_1 > G_2 \cdots$ such that $\cap G_i = \{1_G\}$ and such that for each i we have $G_{i+1} > K_i$ where G_i / K_i is torsion-free abelian. To absorb the definition, imagine repeatedly passing to a finite index subgroup by pulling back a large index subgroup of the free abelianization at each stage. The following is proven in [Ago08]:

Lemma 14.2. *Every right-angled Artin group is (virtually) RFRS.*

Corollary 14.3. *Every closed hyperbolic 3-manifold M with a geometrically finite incompressible surface satisfies:*

(1) *$\pi_1 M$ is a subgroup of a right-angled Artin group (right-angled Coxeter group).*
(2) *$\pi_1 M$ is virtually fibered.*
(3) *$\pi_1 M$ is subgroup separable.*

Proof. Since $\pi_1 M$ has a quasiconvex hierarchy, we see that $\pi_1 M$ is virtually special by Theorem 13.3.

Since $\pi_1 M$ has a finite index subgroup that lies in a right-angled Artin group it is RFRS by Lemma 14.2. Agol's virtual fibering criterion [Ago08] applies to the finite cover $\widehat{M}$ with $\pi_1\widehat{M}$ special.

A consequence of the tameness theorem [Ago04, CG06], is that subgroups of $\pi_1 M$ are either geometrically finite or virtual fibers. The latter are easily seen to be separable using an index 2 normal subgroup with $\mathbb{Z}$ quotient. We proved in [HW08] that quasiconvex subgroups of a compact word-hyperbolic special cube complex are separable, as they are virtual retracts. But geometrically finite subgroups are precisely the same as quasiconvex subgroups [Hru10]. □

We will handle the case of a cusped hyperbolic 3-manifold in Section 14.4.

Problem 14.4. Let M be hyperbolic 3-manifold with an incompressible geometrically finite subgroup. Does M (virtually) have a hierarchy where the tori don't get cut until the last steps?

Problem 14.5. Show that every fibered hyperbolic 3-manifold has a finite cover with an incompressible geometrically finite surface.

14.1. **Virtual separation and largeness.**

Lemma 14.6. *Let $D \subset X$ be a separating hyperplane in a connected special cube complex, and suppose that D is* essential *in the sense that neither $D \to X_+$ nor $D \to X_-$ is π_1-surjective so $\pi_1 X$ splits nontrivially along $\pi_1 D$. There is a finite cover $\widehat{X} \to X$ such that no component of $\widehat{D}$ separates $\widehat{X}$.*

Proof. Let $\alpha_+ \to X_+$ and $\alpha_- \to X_-$ be closed based paths representing elements outside of $\pi_1 D$. Use the separability of $\pi_1 D$ in $\pi_1 X$ to pass to a finite quotient of $\pi_1 X$ such that $\bar{\alpha}_+ \notin \overline{\pi_1 D}$ and likewise such that $\bar{\alpha}_- \notin \overline{\pi_1 D}$. Let $\widehat{X}$ denote the corresponding finite cover. Let $D_1, \ldots, D_r$ denote the components of the preimage of D, and consider the dual graph Γ whose vertices correspond to components of $\widehat{X} - \cup_i D_i$ and whose edges correspond to components D_i. Note that Γ is connected since $\widehat{X}$ is connected. By construction, each vertex has valence ≥ 2. And by regularity, $\pi_1 X$ acts transitively on the edges of Γ. If some edge separated, then each edge would separate, and so the finite regular tree would be an n-pod for some n. Consequently, no D_i separates. □

Remark 14.7. The properties discussed here hold in the framework where D is a "codimension-1 subspace" in the sense that D separates any open neighborhood $N(D)$ of D in X. And D is compact and π_1-injective and $\pi_1 D$ is separable in $\pi_1 X$. In our situation separability automatically holds because hyperplanes are always separable.

Lemma 14.8. *Let $D \subset X$ be a hyperplane in a special cube complex. Let $Y = X - N^o(D)$ where $N^o(D)$ denotes the open cubical neighborhood. Let D^+ and D^- denote the sides of $N(D) \cong D \times [-1, 1]$, so there are maps $D^+ \to Y$ and $D^- \to Y$.*

Suppose that D is nonseparating (Y is connected), and at least one of $D^+ \to Y$ or $D^- \to Y$ is not π_1-surjective. Or, suppose that D is separating (let $Y = Y_+ \sqcup Y_-$), and $[\pi_1 Y_+ : \pi_1 D^+] \geq 3$ and $[\pi_1 Y_- : \pi_1 D^-] \geq 2$.

Then there is a finite cover $\widehat{X} \to X$ such that the dual graph of hyperplanes in the preimage of D has negative euler characteristic. Consequently, $\pi_1\widehat{X}$ surjects onto the free group F_2.

Proof. This is similar to the proof of Lemma 14.6. When D is nonseparating and say $\alpha \in \pi_1 Y_+ - \pi_1 D_+$, then we choose a finite regular cover $\widehat{X}$ such that $\bar{\alpha} \notin \overline{\pi_1 D}$. Then the dual graph for $D_1, \ldots, D_r$ in $\widehat{X}$ has each vertex of valence ≥ 3, since each component of $\widehat{X} - \cup_i D_i$ has at least one edge for a D_- type side, and at least two edges corresponding to D_+ type sides that are connected by a lift of α.

In the case where D is separating, we let $\alpha_0, \alpha_1, \alpha_2$ be representatives of distinct nontrivial cosets of $\pi_1 D_+$ in $\pi_1 Y_+$, and we let β_0, β_1 be representatives of distinct nontrivial cosets of $\pi_1 D_-$ in $\pi_1 Y_-$. We choose a finite quotient of $\pi_1 X$ maintaining the distinctness of the three cosets and the two cosets, and we let $\widehat{X}$ be the corresponding finite cover. Then each vertex in the dual graph corresponding to a component of the preimage of Y_+ has valence at least three, and each vertex in the dual graph corresponding to a component of the preimage of Y_- has valence at least two.

Thus the dual graph has negative euler characteristic in both cases, and the quotient of $\widehat{X}$ to the dual graph yields a noncyclic free quotient. □

Example 14.9. We now describe a finite dimensional example X where $F_2 \cong \pi_1 X$ but X does not admit a compact local isometry $Y \to X$ where $\pi_1 Y$ is nontrivial.

Let $\bar{X}$ be the standard 2-complex of $\langle a_1, a_2, b_1, b_2 : [a_i, b_j] \rangle$, so $\bar{X}$ is isomorphic to the cartesian product of two bouquets of two circles and $\pi_1 \bar{X} \cong F_2 \times F_2$.

Let $X \to \bar{X}$ be the based covering space with $\pi_1 \bar{X} = \langle a_1 b_1, a_2 b_2 \rangle$. Note that $\pi_1 \bar{X}$ has trivial intersection with $\langle a_1, a_2 \rangle$ and with $\langle b_1, b_2 \rangle$. Consequently, for each nontrivial element $g \in \pi_1 \bar{X}$, there is a unique flat plane F_g in $\widetilde{X}$ such that is stabilized by g, and in fact, g acts by translations along its diagonal. Moreover, F_g does not contain a nonempty proper convex subcomplex stabilized by g.

It follows that for any compact local isometry $Y \to X$, we have $\pi_1 Y = 1$.

The punchline of this section is the following strengthening of the Tits alternative for groups acting properly on finite dimensional CAT(0) cube complexes [SW05].

Theorem 14.10. *Let X be a finite dimensional virtually special cube complex. Then $\pi_1 X$ is either large or virtually abelian.*

Proof when X is compact. By possibly passing to a finite cover, we can assume without loss of generality that X is special. Consider a hyperplane $D \subset X$. By induction on dimension, either $\pi_1 D$ is large or virtually abelian. In the former case, we see that $\pi_1 X$ is large since there is a virtual retraction $\mathsf{C}(D \to X) \to D$. In the latter case, either $\pi_1 X$ is virtually $\pi_1 D \rtimes \mathbb{Z}$ and hence virtually abelian, or the splitting along D has some large index on at least one side, and so there is a finite index subgroup with a free quotient by Lemma 14.8. (This follows the plan used in [SW05] but avoids the algebraic torus theorem.) □

Proof in the general case. We can pass to a finite cover where all hyperplanes are 2-sided and embed. We can also assume that X is *minimal* in the sense that it does not contain a locally convex π_1-isomorphic subcomplex. Hence for each hyperplane D, either D is nonseparating or $\pi_1 D$ has index ≤ 2 on each side of $X - D$.

Suppose X contains a hyperplane D where $\pi_1 D$ is virtually abelian. Consideration of the $\pi_1 D$ on its two sides in X, we can either apply Lemma 14.8 to obtain largeness, or we obtain one of the following: Either $\pi_1 D$ has index 2 on both sides in the separating case, or $\pi_1 D$ has index 1 on each side in the nonseparating case. These two situations lead to short exact sequences of the form: $1 \to \pi_1 D \to \pi_1 X \to \mathbb{Z}$ or $1 \to \pi_1 D \to \pi_1 X \to \mathbb{Z}_2 * \mathbb{Z}_2 \to 1$. And in each case, we obtain that $\pi_1 X$ is virtually polycyclic and hence virtually abelian.

Now suppose some hyperplane D is not virtually abelian. Then D itself must contain a hyperplane E where $\pi_1 E$ has nontrivial index on both of its sides within D. We then let D' be the hyperplane of X with $D \cap D' = E$, and we see that $\pi_1 D'$ has nontrivial index on both its sides within X. And so again, largeness follows from Lemma 14.8. □

Perhaps Theorem 14.10 can be strengthened to: If X is special and finite dimensional then: Either $\pi_1 X$ is abelian or $\pi_1 X$ has a noncyclic free quotient. In particular, I suspect that $G \cong \mathbb{Z}^n$ whenever G is both special and virtually abelian. Note that the Klein bottle group is virtually special and virtually abelian, but not special as it is not residually torsion-free nilpotent.

Consideration of the proof of Theorem 14.10 suggests the following might hold.

Conjecture 14.11. *Let X be a finite-dimensional nonpositively curved cube complex. Suppose that $\pi_1 D$ is virtually abelian for each hyperplane D of X. Then $\pi_1 X$ is either A or A-by-S where A is virtually abelian and S is a non-elementary quasifuchsian group.*

14.2. **Cutting all tori with first surface.**

Lemma 14.12 (Together in free-abelianization)**.** *Let G be virtually compact special and word-hyperbolic. Let $\sigma_1, \ldots, \sigma_k$ be elements of infinite order in G. There exists an index n characteristic subgroup G' such that $\sigma_1^n, \ldots, \sigma_k^n$ (and all their conjugates by G) map to nontrivial elements in the free-abelianization $G' \to H^1(G'; \mathbb{Z})$.*

Proof. Let G act properly and cocompactly on the CAT(0) cube $\widetilde{X}$, and let J be a finite index characteristic subgroup such that $X = J \backslash \widetilde{X}$ is special and compact. See Lemma 12.5.

We can assume without loss of generality that each σ_i generates a maximal cyclic subgroup (pass to a smaller and better generator), and we can assume that no σ_i, σ_j have conjugate powers (remove one from the list). After this assumption, we know by hyperbolicity that no σ_i is conjugate to itself by an infinite order element.

By Lemma 8.5, for each i, let $\widetilde{Y}_i$ be a $\langle \sigma_i \rangle$-cocompact superconvex subcomplex of $\widetilde{X}$, and let Y_i be the quotient of $\widetilde{Y}_i$ by $J \cap \langle \sigma_i \rangle$ so there is a map $Y_i \to X$. For each i, let $\widehat{X}_i = \mathsf{C}(Y_i \to X)$, so there is a retraction $\widehat{X}_i \to Y_i$. Let $\widehat{X}$ denote a finite cover with $\pi_1 \widehat{X}$ characteristic in G factoring through each $\widehat{X}_i$ and let $n = [G : \pi_1 \widehat{X}]$. Then for each i, the composition $\widehat{X} \to \widehat{X}_i \to Y_i$ shows that σ_i^n is nontrivial in $\mathsf{H}^1(\widehat{X})$. □

Remark 14.13. It appears that Lemma 14.12 holds for a collection of tori in a group that is hyperbolic relative to tori, that is virtually the fundamental group of a sparse special cube complex.

Proposition 14.14. *Let M be a finite volume hyperbolic 3-manifold. There exists a filling $\bar{M}$ of M such that $\pi_1 \bar{M}$ is word-hyperbolic virtually compact special.*

Remark 14.15. It is adequate for the later applications to prove a weaker form of Proposition 14.14 that allows one to pass to a finite cover of M with a large Dehn filling that is virtually special. This is considerably easier.

Proof. One way to prove this is to use the fact that there exists a hyperbolic filling $\bar{M}$ that contains a geometrically finite immersed incompressible surface which lifts to an embedding in a finite cover [**?**]. Thus $\pi_1 \bar{M}$ is word-hyperbolic and compact virtually special by Theorem 14.1.

Another approach is to choose large fillings of all but one cusp and thus obtain a hyperbolic one-cusp manifold. Then following ideas that we will later develop, observe that one-cusp manifolds are virtually special, and that their large fillings are as well. □

Corollary 14.16. *Let M be a finite volume hyperbolic 3-manifold with cusps $T_1, \ldots, T_k$ (where $k \geq 1$). There exists a finite cover $\widehat{M} \to M$ and an incompressible surface $S \subset \widehat{M}$ such that each cusp $\widehat{T}_i$ of $\widehat{M}$ contains a boundary curve τ_i of S.*

Proof. By Proposition 14.14, there exists large fillings of $T_1, \ldots, T_k$ to obtain a closed hyperbolic 3-manifold $\bar{M}$ with $\pi_1 \bar{M}$ virtually compact special and word-hyperbolic, and such that $\pi_1 T_1, \ldots, \pi_1 T_k$ map to nontrivial cyclic subgroups generated by $\sigma_1, \ldots, \sigma_k$.

By Lemma 14.12, there is a finite index subgroup G' of $G = \pi_1 \bar{M}$ such that $\sigma_1^n, \dots, \sigma_k^n$ are nontrivial in $\mathsf{H}^1(G'; \mathbb{Z})$.

Let $\widehat{M}$ denote the cover of M induced by the finite index subgroup G'. Let $\mathsf{H}^1(G'; \mathbb{Z}) \to \mathbb{Z}$ be a quotient such that each σ_i^n has nontrivial image.

Consider the resulting map $\widehat{M} \to S^1$, make this transversal to a point of S^1 after doing the same to each $\widehat{T}_i$. The preimage is a surface. Now compress to an incompressible surface with the same boundary. (It might not be connected.) □

Remark 14.17. The statement and proof of Corollary 14.16 is unnecessarily complicated. In particular it is unnecessary to pass to a finite cover $\widehat{M}$. Given a 3-manifold M with boundary consisting of a collection of tori $T_1, \dots, T_k$, for each i, the map $\mathsf{H}_1(T_i) \to \mathsf{H}_1(M)$ has infinite image. Indeed, this follows immediately by considering the Mayer-Vietoris long exact sequence associated to a Dehn filling $M_i = D_i \cup_{T_i} M$ where D_i is a solid torus. Consequently, each $\mathsf{H}_1(T_i)$ has nontrivial image in the free abelianization of $\mathsf{H}_1(M)$. But for any finite collection of nontrivial elements in a free-abelian group, there is an infinite cyclic quotient in which they all survive. This gives a map $\phi : M \to S^1$ such that each $T_i \to S^1$ is essential. We homotope ϕ so that it is transversal to a point of S^1, and let S be the preimage surface, and then compress to obtain an incompressible surface whose boundary cuts each T_i as claimed.

It will be useful to know that (each component of) the incompressible surface S produced in Corollary 14.16 is geometrically finite. Indeed, this will support the proof of virtual specialness in the finite volume case. The above surface S is sufficient if the goal is merely to show that M virtually fibers since we stop here if S is a virtual fiber, and otherwise proceed with the proof of Theorem 14.29 if it is geometrically finite. We shall now be a bit more careful in a more elaborate construction of S, I believe all of this might be done more directly with constructions involving surfaces in 3-manifolds but it is interesting that we can essentially proceed with a few fundamental 3-manifold facts: Firstly: Haken 3-manifolds have a hierarchy. Secondly: This is a quasiconvex hierarchy in the hyperbolic case provided that the first cut is geometrically finite. Thirdly: The one cusp case contains an incompressible geometrically finite surface (Proposition 14.25.(1)).

We now provide a geometrically finite strengthening of Corollary 14.16.

Lemma 14.18. *Let M be a finite volume hyperbolic 3-manifold with cusps $T_1, \dots, T_k$ (where $k \geq 1$). There exists a finite cover $\widehat{M} \to M$ and an incompressible geometrically finite surface $S \subset \widehat{M}$ such that each cusp $\widehat{T}_i$ of $\widehat{M}$ contains a boundary curve τ_i of S.*

A key point in the proof below is the production of a codimension-1 quasiconvex subgroup W such that no power of any σ_i is conjugate into W. See Problem 14.19.

Proof. **Adding extra wall:** There exists an incompressible geometrically finite surface S with no accidental parabolics (see Definition 14.21). Indeed, a geometrically finite incompressible surface in the 1-cusp case exists by Proposition 14.25.(1), and consequently, such a surface exists in the general cusped case, because we can choose hyperbolic large Dehn fillings for all but one cusp. Note that accidental parabolics of S that are different from a boundary slope of S in some cusp T cannot exist. Indeed, if σ is an accidental parabolic and τ is a boundary slope, then the intersection number $\#_T(\sigma, \tau) = 0$ since $\#_M(S, \tau) = 0$ as σ can be homotoped away from S in $N(S) \cong S \times I$. We can avoid arbitrary accidentals (including those that are powers of the boundary slope) after we avoid boundary slope accidentals by applying Lemma 14.22.

We can pass to a finite cover M' so that all the cusps are very large. Now quotient by the boundary slopes of S' together with arbitrary slopes of other bounding tori. Small-cancellation shows that S'

maps to a closed g.f. surface of the filled manifold $\bar{M}$, and that the element σ generating the π_1-image of each $\pi_1 T_i'$ is either split by the surface (in the former case), or not homotopic into it (in the latter case).

By Lemma 14.14, let $\bar{M}$ be a Dehn filling of M such that $\pi_1\bar{M}$ is hyperbolic. Let $\sigma_1, \ldots, \sigma_k$ be generators for the images of $\pi_1 T_1, \ldots, \pi_1 T_k$.

A quotient group: Let J be a finite index subgroup of G such that $X = J\backslash\widetilde{X}$ is compact and special. Let W be a hyperplane of X (say corresponding to the codimension-1 subgroup $\pi_1\bar{S}'$ with the property that no $\hat{\sigma}_i^n$ is homotopic into W. Let $H = \pi_1 W$. (If we wish we can assume that H is malnormal in J by passing to a further finite index subgroup.) Let H^o be a finite cover of H such that each σ_i maps to an infinite order element in $G/\langle\langle H'\rangle\rangle$, and hence in $J/\langle\langle H'\rangle\rangle$ (normal closure in different groups). Such H^o exists by small-cancellation theory.

Apply Theorem 12.1 to J, H' relative to H^o to obtain a finite cover Y of a superconvex hull of $N = N(W)$, and let $\hat{N}$ denote the corresponding cover of N. Note that $H' = \pi_1 Y = \pi_1\widehat{N}$. The resulting quotient $J/\langle\langle H'\rangle\rangle$ is represented by the cubical small-cancellation presentation $\langle X \mid Y\rangle$.

Preparing the vertex group: Let $\alpha_1, \ldots, \alpha_m$ denote all elements $(\sigma_i^{g_j})^n$ where $n = [G : J]$ and $G = \cup g_i J$, which constitute a full collection of generators representing commensurability classes of conjugates of cyclic subgroups $J \cap \langle\sigma_i\rangle^g$. By Lemma 14.12 we can pass to an index r normal subgroup K of $\bar{J}$ such that each $(\alpha_i^r)^j$ is nontrivial in the free-abelianization $K \to \mathsf{H}^1(K;\mathbb{Z})$.

Preparing the edge group: Observe that $\bar{J}$ splits as a graph of groups $U*_F$ where $U = \langle N \mid \widehat{N}\rangle$ and F is represented by the image of $\pi_1(X - W)$. Consequently K has an induced splitting as a graph of groups with trivial edge groups. We can pass to a finite subgroup L of K with the property that the elevations of α_i pass through edges spaces at most once. (Either count intersection number with hyperplane of the edge space, or readjust K so that it is the fundamental group of a graph of spaces where all the edge spaces are isolated edges. Then use residual finiteness in similarly to the analogous argument for a collection of immersed cycles in a graph.)

Using a single edge space in the induced splitting of L we see that $L \cong V*_E$ where E is trivial. The elevations of σ_i are now of two types: Those that are conjugate into V, and those that pass through the stable letter of E exactly once. The homomorphism $K \to \mathsf{H}^1(K;\mathbb{Z})$ that projects each $(\alpha_i^r)^j$ to a nontrivial element, shows that $V \to \mathsf{H}^1(V;\mathbb{Z})$ has the same property. There is thus a quotient $V \to \mathbb{Z}$ with the property that for each element β in V with a power β^p conjugate to some σ_i^q in G, the element β maps to a nontrivial element of $\mathbb{Z}$.

The large action: This yields a homomorphism $V*_E \to \mathbb{Z}*_E \cong F_2$ with the property that each β above survives.

Pulling back the splitting to the corresponding finite cover $\widehat{M}$ of M yields the desired surface. (Make $\widehat{M} \to B_2$ transversal to the midpoints of 1-cells, and the preimage is a surface. Then compress to an incompressible surface with the same boundary.) The resulting incompressible surface cannot be a virtual fiber because of the dual action on the large tree. □

Problem 14.19 (Missing Wall)**.** Let $\sigma_1, \ldots, \sigma_k$ be infinite order elements of a virtually compact special word-hyperbolic group G. Does there exist a quasiconvex codimension-1 subgroup $H \subset G$ such that no σ_i has a power conjugate into H?

Let $F \subset G$ be an infinite index quasiconvex subgroup (e.g. $\langle\sigma_1^n, \ldots, \sigma_k^n\rangle$ for some large n). Does there exist a quasiconvex codimension-1 subgroup H such that no nontrivial element of F is conjugate to an element of H?

The following variant of Problem 14.19 has an affirmative answer when there is only one cusp [MZ08]. Presumably the multiple cusp case can be deduced from the single cusp case, or can be reproven using the same methodology.

Problem 14.20. Let M be a cusped hyperbolic 3-manifold. Does M contain the fundamental group of a closed surface containing no nontrivial parabolic elements?

Definition 14.21. Let S be an incompressible surface in a hyperbolic 3-manifold M. An *accidental parabolic* σ in S is a closed essential path in S that is homotopic to a path in a torus of ∂M but not homotopic to a path in ∂S.

Lemma 14.22 (No accidental parabolics)**.** *Let $S \subset M$ be an incompressible surface. Then there is an incompressible surface S' with $\partial S \subset \partial S'$ such that S' has no accidental parabolics.*

Remark 14.23. We note that if each component of the original surface has geometrically finite fundamental group then the same holds for each component of the result. Indeed, the fundamental group of the surface obtained by compressing the annulus maps to a subgroup of the fundamental group of the original. Consequently, if the original is geometrically finite then so is the result.

Proof. Consider the manifold $\bar{M}$ obtained by cutting M along S. An essential annulus in $\bar{M}$ with one bounding circle on a copy of S and another bounding circle on a cusp T, would imply an embedded essential annulus A in $\bar{M}$ with the same property. The result of compressing along $N(A)$ yields a new surface S' with the same euler characteristic but with two additional boundary components. Thus the complexity consisting of genus of components decreases. Therefore this procedure can only be implemented finitely many times. Note that after performing this procedure we can pass back to an incompressible surface by compressing along discs without affecting the boundary. □

Remark 14.24. The goal of this subsection was to produce a splitting of $\pi_1 \widehat{M}$ as a graph of groups where the edge groups are quasiconvex and have no accidental parabolics, and the vertex groups are word-hyperbolic.

In my initial planning, this was an intermediate goal towards a shortcut version of cubulation which sidestepped the malnormality of edge groups (by substituting aparabolicity on one side). In retrospect, it appears that it was unnecessarily strong, and that there is an adequate quasiconvex hierarchy available that jives with Theorem 16.28.

Indeed, M has a geometrically finite incompressible surface because we can choose large Dehn fillings of all but one cusp, to obtain a new hyperbolic manifold $\bar{M}$, and then note that $\bar{M}$ has an incompressible geometrically finite surface $\bar{S}$ by Proposition 14.25.(1). A preimage S of this surface in M is thus likewise geometrically finite and incompressible. (If it were a virtual fiber it would have virtually cyclic index in M and hence in $\bar{M}$ as well.) We can cut S along annuli and increase its boundary to ensure that it has no accidental parabolics as in Lemma 14.22 and the result is again geometrically finite by Remark 14.23. This exact procedure is followed for subsequent splittings. We thus see that M has a hierarchy whose edge groups are geometrically finite and have no accidental parabolics.

A source of geometrical finiteness for the first cut is the following result of Culler and Shalen [CS84].

Proposition 14.25. *Let M be a hyperbolic 3-manifold with one cusp. Then:*

(1) M contains a (separating) geometrically finite incompressible surface with boundary.

(2) M contains a second separating geometrically finite incompressible surface with boundary having a different slope.

14.3. **Omnipotence.** An ordered set of elements $\{g_1, \ldots, g_r\}$ is *independent* in G if each g_i has infinite order, and the subgroups $\langle g_i \rangle$, $\langle g_j \rangle$ do not have conjugates with nontrivial intersection for $i \neq j$.

A group G is *omnipotent* if for each independent set of elements $\{g_1, \ldots, g_r\}$ there is a number $K \geq 1$ such that for each set of positive natural numbers $\{n_1, \ldots, n_r\}$ there is a quotient $G \to \bar{G}$ such that $\bar{g}_i$ has order $n_i K$ for each i.

We note that the conclusion of omnipotence is similar to the special quotient theorem applied to a collection of cyclic subgroups of G. In [Wis00] we proved that free groups are omnipotent and we now generalize this to:

Theorem 14.26. *Let G be a virtually special word-hyperbolic group. Then G is omnipotent.*

Proof. We can assume without loss of generality (it just affects the constant K) that each $H_i = \langle g_i \rangle$ is a maximal cyclic subgroup, so that $\{H_1, \dots, H_r\}$ is a malnormal collection of subgroups. For each i we apply Theorem 12.3 to $(G, \{H_1, \dots, \widehat{H_i}, \dots, H_r\})$, where the notation $\widehat{H_i}$ indicates that we have omitted the i-th subgroup, to obtain a virtually special quotient $G \to G_i$ where each H_j has finite image. Moreover, by choosing the $(H_j)'$ large enough, we can guarantee that H_i injects in G_i.

We now apply Lemma 14.12 to G_i to obtain a characteristic finite index subgroup $J_i \subset G_i$ such that conjugates of $H_i^g \cap J_i$ survive in the free-abelianization $J_i \to Z_i$.

There is a constant K_i, such that for each m_i, there is a quotient $Z_i \to Q_i = Z_i/m_iZ_i$ such that $g_i^{p_i}$ has order m_iK_i, where $p_i = [G_i : J_i]$.

We now use the quotients $G \to \prod_i G_i \to \prod_i Z_i \to \prod_i Q_i$ that are parametrized by m_i.

With a bit of fiddling with the constants, (and at the expense of making it quite large) we can determine a uniform constant K claimed above, and use the images of the homomorphisms $G \to \prod_i Q_i$ to obtain the desired $G \to \bar{G}$. □

Remark 14.27. What is needed here is that each G is the fundamental group of a virtually special cube complex, such that each cyclic subgroup has a compact core associated with it.

In the relatively hyperbolic case, we could use maximal abelian subgroups to obtain a corresponding result.

We can avoid using the malnormal special quotient theorem, if we can obtain trivial wall projections of elevations of these cores, and then apply Lemma 14.12. This gives an elementary proof, that is a more direct generalization of the proof in [Wis00].

Conjecture 14.28. *Let G be a virtually special word-hyperbolic group. Let $g_1, \dots, g_r$ be an independent set of elements. There exists $N_1, \dots, N_r$ such that for any $n_i \geq N_i$ the group $\bar{G} = G/\langle\langle g_1^{n_1}, \dots, g_r^{n_r} \rangle\rangle$ satisfies:*

1. *$\bar{G}$ is virtually torsion-free.*
2. *$\bar{G}$ is residually finite.*
3. *$\bar{G}$ virtually splits as a graph of groups.*
4. *$\bar{G}$ is virtually special.*

We note that $\bar{G}$ acts properly on a CAT(0) cube complex for large N_i (see Section 5.4). This already follows from [Wis04] when G is free, which is the most important test case for Conjecture 14.28. The euler characteristic calculation for a finite index torsion-free subgroups shows that (for large n_i, the virtual torsion-freeness yields virtually large first betti number and hence some splittings. One hopes that this might be enough to get a hierarchy, which would be a quasiconvex hierarchy by [MW08].

14.4. **The cusped case.**

Theorem 14.29. *Let M be a finite volume cusped hyperbolic 3-manifold. Then $\pi_1 M$ is virtually the fundamental group of a compact special cube complex.*

Proof. By Lemma 14.18, there is a finite cover $\hat{M}$ of M such that $G = \pi_1\hat{M}$ splits as a graph of groups where each vertex group is word-hyperbolic and each edge group is quasi-isometrically embedded in G. Moreover, by Lemma 14.22, we can assume that the splitting has no accidental parabolics in the edge groups. Consequently, for an edge group E with images E_+, E_- in the vertex group V of its associated

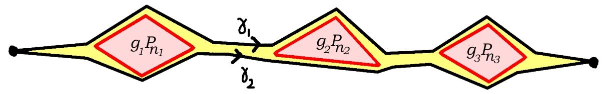

FIGURE 159. γ_1 and γ_2 fellow-travel relative to parabolic subgroups

splitting, E_+ and E_- have trivial intersection with rank 2 subgroups of V. (Note however, that they can have infinite intersection with rank 2 subgroups of G, and in particular, E_+, E_- have conjugates in V that intersect in an infinite cyclic subgroup of a rank 2 subgroup of G.)

The vertex groups have quasiconvex hierarchies since they are fundamental groups of irreducible hyperbolic 3-manifolds with boundary, and so each vertex group is virtually compact special. We thus find that G is virtually compact special by Theorem 16.28. Consequently $\pi_1 M$ is as well. □

One can deduce that $\pi_1 M$ is consequently the fundamental group of a sparse cube complex that is virtually special.

Problem 14.30. Let M be a cusped hyperbolic 3-manifold. Is $\pi_1 M$ the fundamental group of a compact nonpositively curved cube complex.

Problem 14.31. Let M be a hyperbolic 3-manifold. Does $\pi_1 M$ have a finite index subgroup that is the fundamental group of a 3-dimensional nonpositively curved cube complex?

Let M be a hyperbolic 3-manifold with boundary. Does $\pi_1 M$ have a finite index subgroup that is π_1 of a 2-dimensional nonpositively curved cube complex?

In general, the 3-manifold itself may not be homeomorphic to a nonpositively curved cube complex. Indeed, Li showed that if M is orientable and irreducible with ∂M a torus, and M contains no closed nonperipheral embedded incompressible surfaces, then only finitely many Dehn fillings of M yield 3-manifolds homeomorphic to nonpositively curved cube complexes [Li02]. However, for some valid 3-dimensional cases see the work of [AR90]. A motivating 2-dimensional case is Weinbaum's observation that the Dehn complex of a prime alternating link is nonpositively curved (see [Wei71] and [Wis06]).

15. LARGE FILLINGS ARE HYPERBOLIC

We now give a self-contained explanation of the word-hyperbolicity of hyperbolic fillings of cubulated groups that are relatively hyperbolic. We also show prove the persistence of quasiconvexity of certain subgroups under suitably large fillings. This is a limited special case of results obtained by Osin and Groves-Manning [Osi07, GM08] but there is a fairly natural proof within our framework.

We will use the following result of Papasoglu proven in [Pap95]:

Proposition 15.1. *Let Γ be a graph. Suppose each geodesic bigon in Γ is* ϵ-thin*, in the sense that each side lies in the ϵ-neighborhood of the other side. Then each geodesic triangle is δ-thin, and so Γ is δ-hyperbolic.*

We also need the following result suggested by Figure 159. It can be deduced from [DS05, Thm 1.12].

Proposition 15.2. *Let G be hyperbolic relative to subgroups $P_1, \ldots, P_r$. Suppose G acts properly and cocompactly on the geodesic metric space $\widetilde{X}$, and each P_i cocompactly stabilizes a nonempty connected subspace $\widetilde{Z}_i$. There exists a constant κ with the following property:*

Let γ_1, γ_2 be a pair of geodesics in $\widetilde{X}$ with the same endpoints. Then there is a sequence of cosets $\{g_i P_{n_i}\}$ such that γ_1, γ_2 asynchronously κ-fellow travel *relative to $\{g_i \widetilde{Z}_{n_i}\}$ in the sense that for $0 \leq t_1 \leq |\gamma_1|$ and reparametrization θ, either $d(\gamma_1(t), \gamma_2(\theta(t))) \leq \kappa$, or $\gamma_1(t), \gamma_2(\theta(t))$ both lie in $N_\kappa(g_i \widetilde{Z}_{n_i})$, where $g_i \widetilde{Z}_{n_i}$ varies with t.*

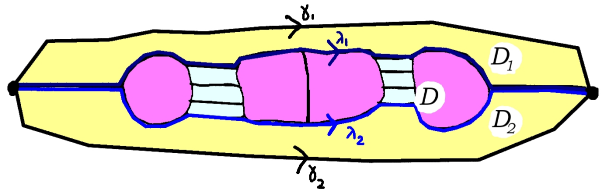

FIGURE 160. The diagram E between γ_1, γ_2 must decompose as the union of a "stemmed ladder" D surrounded by two square diagrams D_1, D_2.

With a bit of care, the proofs generalize to a sparse cube complexes, and we have separately sketched the details in that setting. In addition to the Propositions stated here, we will need the "isometric core property" and Lemma 16.6 to support the proof in the cosparse case.

Theorem 15.3. *Suppose that X is a nonpositively curved cube complex such that:*

(1) *$\widetilde{X}$ is sparse with flats $F_1, \ldots, F_k$,*
(2) *$G = \pi_1 X$ is hyperbolic relative to subgroups $P_1, \ldots, P_r$.*

For each i, there is a finite subset $S_i \subset P_i - \{1_G\}$ such that if P'_i is a hyperbolic-index normal subgroup of P_i that is disjoint from S_i, then the quotient $G/\langle\langle P'_1, \ldots, P'_r\rangle\rangle$ is word-hyperbolic.

Remark 15.4. Theorem 15.3 applies in the relatively hyperbolic case provided an isometric cocompact subcomplex $\widetilde{X}_o \subset \widetilde{X}$ exists. It is only to guarantee its existence that we use require the sparse hypothesis here.

Proof in compact case. By Lemma 8.5, for each i, let $\widetilde{Y}_i$ be a superconvex subcomplex that P_i acts on cosparsely.

The superconvexity and malnormality of the collection $\{P_i\}$ ensures that the wall-pieces and cone-pieces between the $\widetilde{Y}_i$ are of uniformly bounded diameter $\leq B$. Let S_i consist of the nontrivial elements of P_i with translation distance $< \frac{1}{24}B$ in $\widetilde{Y}_i$.

Let P'_i be a hyperbolic-index normal subgroup disjoint from S_i for each i. Let $\widehat{Y}_i = P'_i \backslash \widetilde{Y}_i$, so the cubical presentation $\langle X \mid \widehat{Y}_1, \ldots, \widehat{Y}_r\rangle$ satisfies the $C'(\frac{1}{24})$ small cancellation condition.

Consider geodesics γ_1, γ_2 in $\widetilde{X}^*$. Let $E \to X^*$ be a minimal complexity disk diagram between γ_1 and γ_2, with the following properties: There are geodesics $\lambda_i \to E$ with the same endpoints as γ_i such that each pair λ_i, γ_i bounds a square subdiagram D_i of E, and $E = D_1 \cup_{\lambda_1} D \cup_{\lambda_2} D_2$, and where D has minimal complexity among all possible such choices.

We note that γ_i, λ_i are geodesics in $\widetilde{X}$, for otherwise γ_i would not be a geodesic in $\widetilde{X}^*$.

We will show that D is α-thin and show that each D_i is β-thin, and consequently E is $(\alpha + 2\beta)$-thin. Since α, β will depend only on X and the choices of the P'_i, we see that $\widetilde{X}^*$ has $(\alpha + 2\beta)$-thin bigons. We thus conclude that the 1-skeleton of $\widetilde{X}^*$ has $(\alpha + 2\beta)$-thin geodesic bigons, and is thus δ-hyperbolic by Proposition 15.1.

Observe that there is no θ-shell in D whose outerpath is a subpath of either λ_i for this would contradict that λ_i is a geodesic, indeed X^* is $C'(\frac{1}{24})$ so has short innerpaths by Lemma 3.48. Observe that there is no generalized corner of a square in D whose outerpath is on λ_i, for we could then push the square out to D_i, and reduce the complexity of D.

We conclude from Theorem 3.36 that D is either an arc, or is a ladder or single cone-cell with a (possibly trivial) arc attached at each end. Moreover, we note that the square parts of D are actually grids, since a square within a shard of D could be passed into D_1 or D_2, thus violating our minimal choice of D. We refer the reader to Figure 160 for a picture of a typical D within E.

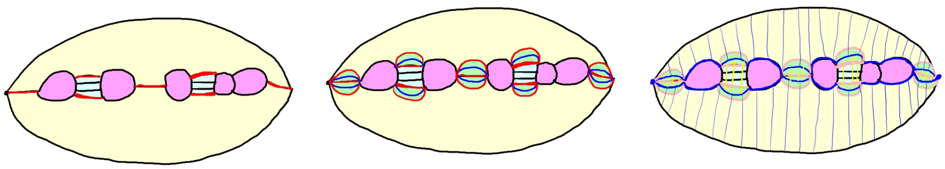

FIGURE 161. $\widetilde{X}_o^*$ has thin bigons, for large girth hyperbolic fillings.

Let α be a constant so that for any pair of geodesics in one of the finitely many $\widehat{Y}_j$, if the endpoints of these geodesics are within B of each other then the geodesics are within α of each other, and note that $\alpha \geq B$.

As in Definition 3.35.(2), each λ_i is a concatenation $\lambda_{i1}\lambda_{i2}\ldots\lambda_{i\ell}$, where λ_{1j} and λ_{2j} are on opposite sides of successive cone-cells and grids within the ladder. Observe that the distance between endpoints in λ_{1j} and λ_{2j} is bounded by the maximal diameter B of a piece. Consequently, the geodesics $\lambda_{1j}, \lambda_{2j}$ accordingly either B-fellow travelers within a grid, or α-fellow travelers within the cone that is a translate of $\widehat{Y}_{n_j}$. Thus λ_1, λ_2 are in α-neighborhoods of each other.

By Proposition 15.2, there exists μ such that γ_i, λ_i must μ-fellow travel relative to translates $g_j\widetilde{Y}_{n_j}$. Consider subpaths γ_i', λ_i' that bound a geodesic rectangle in $N_\mu(\widetilde{Y})$ whose left and right sides have length $\leq \mu$. There exists δ such that the finitely many spaces $P_i'\backslash N_\mu(\widetilde{Y}_i)$ are δ-hyperbolic (they are just thickenings of the $\widehat{Y}_i$). Let β be such that for any pair of geodesics in one of these $\widehat{Y}_i$ thickenings, if the endpoints are μ-close then the geodesics β-fellow travel, and note that $\beta \geq \mu$. Then γ_i, λ_i must β-fellow travel since they piecewise β-fellow travel. □

Proof in cosparse case. By Lemma 16.6, there is a connected π_1-surjective subcomplex $X_o \subset X$ such that $\widetilde{X}_o \subset \widetilde{X}$ is an isometric embedding, and moreover, the complementary components $\widetilde{X} - \widetilde{X}_o$ are contained in unique quasiflats $g_j\widetilde{F}_j$ of $\widetilde{X}$. Let $\widetilde{Y}_i^o$ denote the intersection $\widetilde{X}_o \cap \widetilde{Y}_i$, and we will later let $\widehat{Y}_i^o$ denote the quotient of $\widetilde{Y}_i^o$ corresponding to $\widehat{Y}_i$.

As in the compact case, we consider a pair of geodesics γ_1, γ_2 with the same endpoints, but now assume that they lie in $\widetilde{X}_o^*$, which denotes the preimage of X_o in $\widetilde{X}^*$. We then consider a minimal complexity $E = D_1 \cup_{\lambda_1} D \cup_{\lambda_2} D_2 \to \widetilde{X}^*$. As before D is a ladder. By the minimality of E, the boundary path of each cone-cell C of D lies in $\widetilde{X}_o^*$, and so can be regarded as a path in $\widetilde{X}_o$. Indeed, for an edge e in $\partial_p C$, either e is on ∂E which implies that e is in $\widetilde{X}_o$ since each γ_i is in $\widetilde{X}_o$, or there is a subsequent cell s or C' meeting C along e. If it is a neighboring cone-cell C' then e lies in $g\widetilde{Y} \cap g'\widetilde{Y}'$ which lies in $\widetilde{X}_o^*$. If it is a square s, then either $s \subset \widetilde{X}_o^*$ or s lies in $\widetilde{Y} - \widetilde{X}_o$ and so could be absorbed into C.

We now consider the (possibly degenerate) rectangular parts of D. Their top and bottom boundaries are geodesics in $\widetilde{X}$ that are path-homotopic to paths in $\widetilde{X}_o$. We replace E by a diagram E' that contains these $\widetilde{X}_o$ paths as illustrated in Figure 161. This gives us geodesics λ_1', λ_2' with the same endpoints as λ_1, λ_2 that bound diagrams D_1' and D_2' within E', and also a diagram D' that is a thickening of the ladder D.

D_1', D_2' are thin by Proposition 15.2 applied to $\widetilde{X}_o$ as before. And D' is thin as before, by breaking it up into cone-cells and thickened rectangles, each of which has a geodesic top and bottom in $\widehat{Y}_{n_j}^o$ or in $\widetilde{X}_o$ respectively. As in the compact case, we regard each cone-cell as mapping to a word-hyperbolic $\widehat{Y}_{n_j}^o$. However, the thickened rectangles are treated again as we treated D_1, D_2. □

Remark 15.5. It is likely that there is a more general version of Theorem 15.3 that takes as input a small-cancellation cubical presentation X^* with $\pi_1 X^*$ relatively hyperbolic.

Theorem 15.6. *Continuing with the notation of Theorem 15.3. Let H be a relatively quasiconvex subgroup of $\pi_1 X$ that is full. There exist (slightly larger) finite subsets $S_i^+ \subset P_i - \{1\}$ such that H projects to a quasiconvex subgroup of $\pi_1 X^*$.*

Proof in compact case. By Proposition 8.2, let $\widetilde{A}$ be an H-cocompact superconvex core, and let $A = H\backslash\widetilde{A}$. There exists M such that for each g, i either $g\widetilde{Y}_i \subset \widetilde{A}$, or diameter$(\widetilde{A} \cap g\widetilde{Y}_i) \leq M$.

We can thus choose S_i^+ so that the map of the induced presentation $A^* \to X^*$ has no missing θ-shells, and is hence π_1-injective by Corollary 12.18.

Suppose that cone-pieces in X^* are small in the sense that $\nabla_{Y_i}(P) < \frac{1}{8}|Y_i^*|$ whenever P is a cone-piece between distinct $\widetilde{Y}_i, g\widetilde{Y}_j$.

Suppose that overlaps between cones Y_i and A are small in the sense that for any path P in the intersection of translates $\widetilde{Y}_i, \widetilde{A}$ in $\widetilde{X}$, either $\nabla_{Y_i}(P) < \frac{1}{4}|Y_i|$ or $\widetilde{Y}_i \subset \widetilde{A}$.

Then $\widetilde{A}^* \to \widetilde{X}^*$ is a convex subcomplex by Lemma 3.52. □

Proof in sparse case. Let $\widetilde{A}$ denote an H-invariant superconvex subcomplex that is cosparse.

The complex $\widetilde{X}$ equals $GK \cup_i g\widetilde{F}_i$ where each $\widetilde{F}_i$ is a quasiflat, and distinct $g_i\widetilde{F}_i, g_j\widetilde{F}_j$ intersect in GK. By Lemma 16.6, there is a connected π_1-surjective subcomplex $X_o \subset X$ such that $\widetilde{X}_o \subset \widetilde{X}$ is an isometric embedding, and moreover, $GK \subset X_o$ so the complementary components $\widetilde{X} - \widetilde{X}_o$ are contained in unique quasiflats of $\widetilde{X}$.

We make our choices as in the cocompact case. As before $\widetilde{A}^*$ is a convex subcomplex of $\widetilde{X}^*$. Let $\widetilde{X}_o^*$ denote the preimage of X_o in $\widetilde{X}^*$. Let $\widetilde{A}_o^*$ denote the intersection between $\widetilde{A}^*$ and $\widetilde{X}_o^*$. We will show that $\widetilde{A}_o^*$ is connected and convex in $\widetilde{X}_o^*$. Let γ be a geodesic in $\widetilde{X}_o^*$ whose endpoints lie in $\widetilde{A}_o^*$. By the convexity of $\widetilde{A}^*$ in $\widetilde{X}$, we see that γ lies in $\widetilde{A}^*$. But then $\gamma \subset (\widetilde{A}^* \cap \widetilde{X}_o^*) = \widetilde{A}_o^*$.

As $\widetilde{A}_o^*$ is an $\bar{H}$-invariant subcomplex of the $\bar{G}$-cocompact complex $\widetilde{X}_o^*$ (where $\bar{G}$ denotes $\pi_1 X^*$), this proves the quasiconvexity of $\bar{H}$ in $\bar{G}$. □

16. Relatively Hyperbolic Case

16.1. **Introduction.** The main goal of this section is to prove Theorem 16.28 whose most important case is the following:

Theorem 16.1. *Let G be torsion-free and hyperbolic relative to virtually abelian subgroups. Suppose that G splits as a graph of groups where each edge group is quasiconvex and each vertex group is word-hyperbolic and virtually compact special. Then G is virtually compact special.*

Theorem 16.28 generalizes Theorem 13.3 to a setting that includes many groups that are hyperbolic relative to abelian subgroups. Our motivation is to the 3-manifold application since by Corollary 14.16, for each cusped hyperbolic 3-manifold M there is a finite cover $\widehat{M}$ such that $\pi_1\widehat{M}$ has the type of hierarchy specified in Theorem 16.1.

We expect that a complete generalization of Theorem 13.3 would be as follows:

Conjecture 16.2. *Let G be (torsion-free) and hyperbolic relative to virtually abelian subgroups. Suppose that G has a quasiconvex hierarchy terminating at finite groups. (So all edge groups are quasi-isometrically embedded.) Then G is virtually cosparse/compact special.*

A potential direct approach towards Conjecture 16.2 would require reworking the results in [HWa] and [HWc] so that they apply in a relatively hyperbolic context. The main obstacle towards a relatively hyperbolic generalization of [HWa] is to provide a relatively hyperbolic version of the "trivial wall projection lemma". What is missing from [HWc] is a cubulation criterion that allows more flexibility

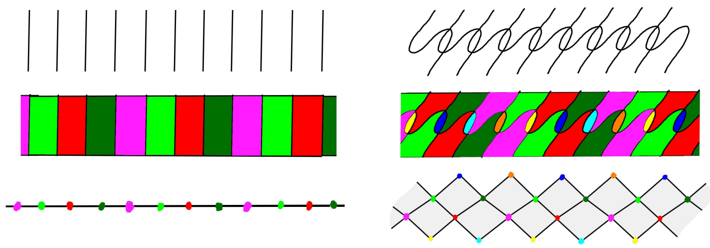

FIGURE 162. Jiggled Walls: The wallspace on an infinite strip on the left has dual cube complex homeomorphic to $\mathbb{R}$. If we "jiggle" the walls so that they cross their neighbors, then one obtains a thicker cube complex as on the right. Similar "thickenings" occur if we jiggle typical wallspaces on $\mathbb{R}^n$.

in the interaction between the edge groups and the parabolic subgroups. In particular, we would need to drop the aparabolicity requirement on the edge groups.

We will will follow a less direct approach that circumvents these obstacles at the expense of developing some additional tools to prove (a limited form of) the relatively hyperbolic case as a consequence of the hyperbolic case. We emphasize that for both (2) and (3) below we prove separability in a relatively hyperbolic group by carefully quotienting to a virtually compact special hyperbolic group. This separability is used in (4) to pass to a finite cover with an easier hierarchy.

(1): A virtually special parabolic filling theorem that is an analog of the malnormal special quotient theorem except that only parabolic subgroups are filled.

(2): We show that sparse cube complexes with hierarchies (i.e. all their hyperplanes embed) are virtually special by combining the parabolic special quotient theorem with the double coset criterion for separability. There are some complications here because the small-cancellation method of Lemma 12.22.(1) doesn't quite apply without a patch.

(3): Subgroup separability of graphs of relatively hyperbolic groups whose vertex groups are virtually special. Here we apply (1) to the vertex groups to obtain a quotient that is a graph of small-cancellation virtually special groups. This quotient group has a suitable quasiconvex hierarchy, and is hence virtually special, so separating quasiconvex subgroups from elements in such quotients is sufficient.

(4): We then prove the theorem by using the separability to pass to a finite index subgroup that has a revised splitting along a "cage" with more controlled interactions with the parabolic subgroups. This better splitting is aparabolic on at least one side, and so we are able to cubulate. Virtual specialness then holds by (2).

16.2. ⊛ **Sparse complexes.**

Definition 16.3 (Quasiflat)**.** A *quasiflat* is a CAT(0) cube complex dual to a proper action of a virtually abelian group P on a wallspace, with finitely many orbits of walls.

The canonical example of a quasiflat $\widetilde{F}_i$ is isomorphic to the universal cover of the standard n_i-torus T^{n_i} for some n_i, perhaps corresponding to a collection of n_i parallelism classes of hyperplanes in $\mathbb{R}^m$ with a $P_i = \mathbb{Z}^m$ action. The general case is not much more complicated: The typical situation is obtained by "jiggling" these hyperplanes somewhat so that they aren't convex (see Figure 162). Many such situations arise naturally in [Wis04].

Definition 16.4 (Sparse)**.** A nonpositively curved cube complex X is *sparse* if it is a finite union $K \cup \bigcup_i F_i$ where K is compact and $F_i \cap F_j \subset K$ for $i \neq j$, and each F_i equals $P_i \backslash \widetilde{F}_i$ where P_i is a f.g. virtually abelian group acting freely on a quasiflat.

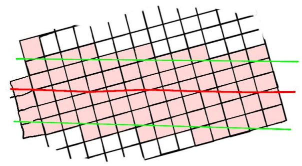

Figure 163. The smallest subcomplex containing a convex subspace is isometrically embedded.

Likewise we say that a group G acts *cosparsely* on a CAT(0) cube complex $\widetilde{X}$ if there is a compact subcomplex K, and finitely many quasiflats $\widetilde{F}_i$ such that $\widetilde{X} = GK \cup \bigcup_i G\widetilde{F}_i$ and $g_i\widetilde{F}_i \cap g_j\widetilde{F}_j \subset GK$ unless $i = j$ and $g_i^{-1}g_j \in \mathrm{Stab}(\widetilde{F}_i)$. See Figure 142 for a "picture" of a cosparse cube complex, and see Example 16.11 for potential overcubulated groups.

Note that when G is torsion-free then $X = G\backslash\widetilde{X}$ is sparse as above. However, when there is torsion the quotient space X can be a "sparse orbihedron". Nevertheless, X will be quasi-isometric to the wedge of finitely many flats and half-flats.

Remark 16.5. In practice we can assume that each $\widetilde{F}_i$ has finitely many P_i-orbits of hyperplanes.

$\widetilde{X}$ is *minimal* if it has no G-invariant convex proper nonempty subcomplex. By assuming that $\widetilde{X}$ is minimal with respect to the action of $G = \pi_1 X$ one can assume that each hyperplane H of $\widetilde{X}$ crosses GK. For otherwise we can replace $\widetilde{X}$ by the intersection of all cubical halfspaces $g\overrightarrow{H}$ corresponding to translates of the halfspace of H containing GK.

The following plays a role in the sparse versions of the results in Section 15. The CAT(0) cube complex $\widetilde{F}$ has the *isometric core property* with respect to the action of a group H, if each H-cocompact subcomplex $\widetilde{K} \subset \widetilde{F}$ lies in an H-cocompact isometrically embedded subcomplex $\widetilde{K}'$. The property is immediate when $\widetilde{F}$ is H-cocompact. This property always holds whenever H is virtually abelian and $\widetilde{F}$ is a quasiflat. To see this, note that there exists an H-cocompact CAT(0) convex subspace $\widetilde{U} \subset \widetilde{F}$, and moreover, $N_r(U)$ is likewise H-cocompact and CAT(0) convex and any $\widetilde{K}$ lies in some $N_r(U)$. The smallest subcomplex containing $N_r(U)$ is our desired isometric core. See Figure 163. (It appears that this property holds whenever there is an H-cocompact CAT(0) convex subspace (not subcomplex!) of $\widetilde{F}$.) Let G act properly and cocompactly on a CAT(0) cube complex $\widetilde{X}$, and let $H \subset G$ be a finitely generated subgroup that is not quasi-isometrically embedded. Then $\widetilde{X}$ does not have the isometric core property with respect to H. Conversely, it appears likely that the isometric core property holds exactly when H is quasi-isometrically embedded in this case.

Lemma 16.6. *Let G act on the CAT(0) cube complex $\widetilde{X}$ which is a union $GK_o \cup_{i=1}^r G\widetilde{F}_i$ with K_o compact, and where $g_i\widetilde{F}_i \cap g_j\widetilde{F}_j \subset GK_o$ unless $g_i\widetilde{F}_i = g_j\widetilde{F}_j$. If each $\widetilde{F}_i$ has the isometric core property with respect to its stabilizer, then $\widetilde{X}$ has the isometric core property.*

In particular, if $\widetilde{X}$ is sparse, then it has the isometric core property.

Proof. Let $\widetilde{B}$ be a G-cocompact subcomplex of $\widetilde{X}$, and assume without loss of generality that $GK \subset B$. Let $\widetilde{B}_i = \widetilde{X} \cap \widetilde{F}_i$. By the isometric core property of $\widetilde{F}_i$, the subcomplex $\widetilde{B}_i$ extends to a $\mathrm{Stab}(\widetilde{F}_i)$-cocompact subcomplex $\widetilde{B}'_i \subset \widetilde{F}_i$.

Let $\widetilde{B}' = \widetilde{B} \cup_i G\widetilde{B}'_i$. We now show that $\widetilde{B}' \subset \widetilde{X}$ is isometrically embedded (though perhaps not convex). Indeed, for any geodesic γ in $\widetilde{X}$ joining points in $\widetilde{B}'$, consider a subpath γ_i where γ departs $\widetilde{B}'$ and travels within $\widetilde{F}_i$. Then γ_i has the same length as a path γ'_i in $\widetilde{B}_i \cap \widetilde{B}'$ with the same endpoints. Successively replacing γ_i by γ'_i, we pull γ' to path with the same endpoints in $\widetilde{B}'$ such that $|\gamma'| = |\gamma|$. $\square$

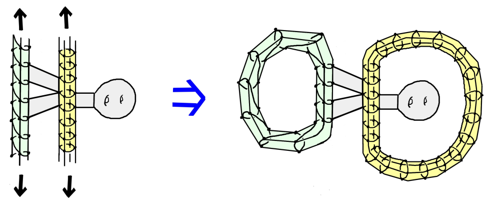

FIGURE 164. Closing up two infinite cylinders in base space.

16.3. ⊛ **Closing up infinite quasiflats.** The following result is sometimes helpful for providing a shortcut via a cocompact interpolation. In particular, it can be used to simplify versions of Theorem 16.20 and Theorem 15.3.

Lemma 16.7. *Let X be a sparse cube complex with fundamental group G. There exists a group G' such that:*

1. *G' splits as a tree T_r of groups where T_r is an r-pod.*
2. *The central vertex group of T_r is G.*
3. *Each edge group of T_r is a peripheral subgroup P_i of G.*
4. *G' acts properly and cocompactly on a CAT(0) cube complex $\widetilde{X}'$.*
5. *Moreover, if each P_i is free-abelian then we choose each $P_i' \cong P_i \times \mathbb{Z}^s$ for some s.*
6. *Moreover, if X has a hierarchy then X' has a hierarchy.*

The motivating case here is where the cube complexes F_i^n associated to the peripheral subgroups, have $\widetilde{F}_i^n$ isomorphic to standard cubulations F_i^n of $\mathbb{E}^n$, though the peripheral subgroup might only be virtually $\mathbb{Z}^m$ with $m \ll n$. In this case, it is easy to imagine quotienting each such F_i^n by a $\mathbb{Z}^{n-m}$ subgroup, to obtain cocompactness, and with a large enough fundamental domain so that the new space looks locally like the old - so nonpositive curvature is maintained. We give a very simple example of this in Figure 164. Since each $\widetilde{F}_i^n$ looks only coarsely like $\mathbb{E}^n$, we will instead adopt an approach using the method of Section 7.2.

Proof. The proper action of G on $\widetilde{X}$ provides a collection of subgroups H together with H-walls. Indeed these H-walls arise from the hyperplanes of $\widetilde{X}$. (We could assume that $\widetilde{X}$ is minimal and hence each such H is actually codimension-1 but the proof does not require this).

The plan of the proof is to construct G' as indicated by extending each peripheral subgroup P to a peripheral subgroup P', so that each H-wall in G extends to an H'-wall in G', and moreover, each P' will act properly and cocompactly on the cube complex $C(P')$ that is dual to the induced wallspace on P'. It follows from Proposition 7.5 that G' acts properly and cocompactly on the CAT(0) cube complex $\widetilde{X}' = C(G')$ that is dual to our wallspace on G'.

Besides actually being able to extend the H-walls of G into the relevant P' subgroups of G', to ensure that the action of P' on $C(P')$ is proper and cocompact, we must ensure that the wallspace of P' is "sufficient" so that each element of P' is transverse to some wall, and also "efficient" in the sense that the number of distinct commensurability classes of walls in P' equals the dimension of P'.

We identify P with $P\widetilde{x}$ where $\widetilde{x} \in \widetilde{F}$ is a 0-cell in a minimal P-invariant subcomplex $\widetilde{E}$. An H-wall can interact with P in several ways: Either it partitions P into two deep parts, and hence cuts P along a codimension-1 subgroup, or P lies entirely on one side of the H-wall, or P stabilizes the H-wall and index 2 cosets of P lie on opposite sides. Indeed, we can pass to a finite index subgroup of P_o so that the translates of the H-hyperplane in $\widetilde{E}$ do not cross each other. The P_o-translates of the H-hyperplanes then yield an action on a tree. We then either find that $H \cap P_o$ is codimension-1, or the action has a fixed vertex, or a stabilized edge. These yield the three cases above.

It is the first type above - the *essential* walls in P that will guide us in constructing P'.

We now let $H_1, \ldots, H_n$ be representatives of the P-orbits of essential walls in P coming from various H-walls of G. We emphasize that a single H-wall in G can contribute several H_i-walls in P, since there might be several G-translates of the H-wall that are not also P-translates. However, when the H-wall in G does not cross any of its G translates, its resulting H_i, H_j-walls for P will not cross each other in P (be they essential walls or not).

We will now give an explicit construction under the assumption that each peripheral subgroup P is free-abelian, the general case will be treated later.

Again, $H_1, \ldots, H_n$ are the distinct stabilizers of walls in $P\mathbb{Z}^m$, and let $P'_i = P_i \times \mathbb{Z}^{n-m}$. For each k, we extend H_k to a codimension-1 subgroup H'_k of P'_k, moreover, we do this so that $\{H'_1, \ldots, H'_n\}$ are sufficient to cubulate P'_i.

Indeed, let $\{v_1, \ldots, v_n\}$ be vectors in $\mathbb{Z}^m$ that span $\mathbb{R}^m$ and such that each v_i is orthogonal to the hyperplane of H_i, then we can choose $\{v'_1, \ldots, v'_n\}$ to be vectors in $\mathbb{Z}^n$, such that the first coordinates of v'_k yield v_k, and such that they span $\mathbb{R}^n$, and let H'_k be the intersection of $H_k \times \mathbb{Z}^{n-m}$ with the stabilizer of the orthogonal complement of v'_k. (Remark: The vectors v_i, v_j are parallel whenever H_i, H_j are commensurable. We could alternately then let n be the number of commensurability classes. We would then choose our vectors $\vec{v}'_i$ so that it is a "basis with multiplicity" so that $\vec{v}'_i = \vec{v}'_j$ when $\vec{v}_i = \vec{v}_j$, and commensurable H_i, H_j would extend (in parallel) to commensurable H'_i, H'_j in P'.)

Assuming H_k is deep, the H_k-wall of P_i has a unique coarse extension to an H'_k-wall of P'_i.

For an H-wall that overlaps P_i at a finite index subgroup H_i of P, we use the separability of $H_i \subset P'$ to let H'_i be a finite index subgroup of P' such that $H'_i \cap P = H_i$.

We build G' from G by amalgamating each P_i with the corresponding subgroup of P'_i.

For each H-wall of G, we first extend the subgroup $H \subset G$ to a subgroup $H' \subset G$ by forming an r-pod of groups whose central vertex is H and whose type P_i leaves are subgroups H'_k amalgamated with H along edge groups H_k supplied above at P by H. It is easy to see that each H' is quasi-isometrically embedded in G'.

Let us now examine how we extend H_i-walls of P to H'_i-walls of P'. For an essential H_i-wall of P, there is coarsely a unique way to extend this to an H'_i-wall of P'. When $H \cap P$ is commensurable with P, then $H'_i \cap P' = H'_i \cap P$ and we extend the H_i-wall of P to an H'_i-wall of P' that contains P' on both sides.

A limited special case of Proposition 7.8′ provides the wallspace structure on G', with a proper action of G' on the dual cube complex. By construction, the induced cubulations of each P'_i is cocompact, and so the dual cube complex is cocompact by Proposition 7.5.

We now consider the case where P_i is virtually abelian but not necessarily free-abelian. By Remark 16.9 the virtually $\mathbb{Z}^m$ group P_i embeds in the virtually $\mathbb{Z}^n$ group P'_i in such a way, that by construction, each codimension-1 subgroup in $H_k : 1 \leq k \leq m$ of P_i extends to a codimension-1 subgroup H'_k of P'_i, and furthermore, the action of P'_i on the associated dual cube complex $C(P'_i)$ is free and cocompact. The rest of the explanation is similar. □

Generalizing the statement of Lemma 16.7 would require a generalized version of Lemma 16.8. See for instance Problem 16.10 and Example 16.11.

Lemma 16.8. *Every torsion-free virtually $\mathbb{Z}^m$ group H is a subgroup of a virtually $\mathbb{Z}^n$ group G that acts freely and cocompactly on the standard n-dimensional cube complex.*

Proof. Suppose H has a finite index subgroup K with $K \cong \mathbb{Z}^m$. Then H acts properly and cocompactly on $\mathbb{E}^m$ as proven by Zassenhaus [Rat94, Thm 7.4.5]. Choose m distinct walls in $\mathbb{E}^m$ that are isometric copies of $\mathbb{E}^{m-1}$ and dual to a basis of K. The action of H on this collection of (cocompact) walls gives a (possibly larger but still locally finite) family of n parallelism classes of walls. As nonparallel walls cross, applying Sageev's construction of Section 7.2, the dual cube complex is the standard cubulation $\widetilde{X}$ of $\mathbb{E}^n$, and moreover, H acts properly on $\widetilde{X}$.

Let $G_\circ = \text{Aut}(\widetilde{X})$, and note that $G_\circ$ acts properly and cocompactly on $\widetilde{X}$. We have thus shown that an arbitrary virtually $\mathbb{Z}^m$ group H has a finite kernel quotient which is a subgroup of a group $G_\circ$ acting properly and cocompactly on the standard cubulation of $\mathbb{E}^n$ for some n. In particular, as H is torsion-free it embeds in $G_\circ$.

It was proven in [BL77] that each maximal torsion-free subgroup of a polycyclic-by-finite group is of finite index. We therefore conclude the proof by letting G be a maximal torsion-free subgroup of $G_\circ$ that contains H. □

Remark 16.9 (Virtually $\mathbb{Z}^m$ extension property)**.** The proof of Lemma 16.8 shows the following: Let P be virtually $\mathbb{Z}^m$. Let $\{H_1, \dots, H_r\}$ be a collection of codimension-1 subgroups of P, with chosen H_k-walls, such that P acts properly on the associated dual cube complex. Then there exists a virtually $\mathbb{Z}^n$ group P', such that each codimension-1 subgroup H_k extends to a codimension-1 subgroup H'_k of P', and the H_k-walls in P extends to H'_k-walls in P', so that P' acts properly and cocompactly on the associated dual cube complex.

Problem 16.10 (Stabilization)**.** Continuing with the notation of Lemma 16.8: Can we moreover choose G so that $G \cong H \times \mathbb{Z}^d$ for some d?

Example 16.11 (Overcubulated Euclidean groups)**.** Allowing torsion, the $(3, 3, 3)$ and $(2, 3, 6)$ triangle groups are virtually $\mathbb{Z}^2$ but do not act on the standard cubulation of $\mathbb{E}^2$. On the other hand, the $(2, 4, 4)$ triangle group does act properly on the standard cubulation of $\mathbb{E}^2$.

I am grateful to Bill Dunbar for explaining the following torsion-free example to me: Let $G = \mathbb{Z}^2 \rtimes_\phi \mathbb{Z}$ be the group where $\phi : \mathbb{Z}^2 \to \mathbb{Z}^2$ is an order-6 automorphism. So $G \cong \langle a, b, t \mid [a, b], a^t = b, b^t = a^{-1}b\rangle$. Then G is virtually $\mathbb{Z}^3$, but G does not act freely on the standard cubulation of $\mathbb{E}^3$. Indeed, in an action of G on $\mathbb{E}^3$, the element t acts by a screw-motion rotating by $\frac{\pi}{3}$ about its axis. Now, G acts by automorphisms on the sphere at ∞, and $\langle t\rangle/\langle\langle t^6\rangle\rangle$ would act faithfully there, but this action also gives rise to an action on a 3-cube which does not admit an order 6 orientation preserving automorphism.

Consequently G cannot act freely and cocompactly on any CAT(0) cube complex $\widetilde{X}$, for then by Lemma 16.12, G would act freely and cocompactly on the standard cubulation of $\mathbb{E}^3$ which is impossible.

Lemma 16.12. *Suppose that G is virtually $\mathbb{Z}^n$ and suppose that G acts properly and cocompactly on a CAT(0) cube complex $\widetilde{X}$. Then G acts properly and cocompactly on the standard cubulation of $\mathbb{E}^n$.*

We note that $\widetilde{X}$ might not contain a convex subcomplex isomorphic to $\widetilde{T}^n$. For instance, let $\widetilde{X}$ denote the cartesian product of tree copies of the cube complex in Figure 162 with the associated $\mathbb{Z}^3$ action.

Proof. The flat torus theorem of [BH99] shows that $\widetilde{X}$ contains a G-invariant isometric copy F of $\mathbb{E}^n$. There is a wallspace structure on F arising from the intersections of hyperplanes $\widetilde{H} \subset \widetilde{X}$ with F. If X were minimal then it would be impossible for $\widetilde{H} \cap F = F$ or $\widetilde{H} \cap F = \emptyset$, but without assuming minimality, we simply ignore these intersections. Each intersection $\widetilde{H} \cap F$ is an isometric copy of $\mathbb{E}^{n-1}$ and if several hyperplanes of $\widetilde{X}$ intersect F in the same wall then we simply identify these from the viewpoint of the wallspace, so a single wall of F could correspond to several hyperplanes of $\widetilde{X}$. We emphasize that crossing hyperplanes of $\widetilde{X}$ might intersect F in parallel hyperplanes.

Suppose there are exactly m parallelism classes of walls in F, and observe that $m \leq n$ for otherwise $\widetilde{X}$ would coarsely contain a copy of $\mathbb{E}^m$, and this is impossible since $\widetilde{X}$ is quasi-isometric to $\mathbb{Z}^n$. Applying Sageev's construction of Section 7.2, we see that G acts on the cube complex $\widetilde{T}^m$ dual to the wallspace of F. Note that $\widetilde{T}^m$ is the standard cubulation of $\mathbb{E}^m$ since there are m-parallelism classes of walls. The action is proper by Proposition 7.4. The dimension $m = n$ after all, since $\mathbb{Z}^n \subset G$ cannot act freely on $\widetilde{T}^m$ with $m < n$. And finally it is cocompact since $\mathbb{Z}^n$ must act cocompactly on $\widetilde{T}^n$ since it is of finite index in $\mathrm{Aut}(\widetilde{T}^n)$. (I expect the subdivision of F is isomorphic to $\widetilde{T}^n$.) □

16.4. **Parabolic fillings that are virtually special.** We now describe a variant of Theorem 12.3 for relatively hyperbolic groups that is restricted to quotienting only by finite index subgroups of the parabolic groups. Our focus, which is Lemma 16.13, provides a relatively hyperbolic variant of Theorem 12.3 and a special case of Conjecture 19.1. (The full case of Conjecture 19.1 can probably be proven by combining Lemma 16.13 with Theorem 12.3 as in the proof of Theorem 12.1.) The reader not yet familiar with sparseness can assume the cube complexes are compact.

Lemma 16.13. *Let X be a sparse nonpositively curved cube complex. Suppose that $\pi_1 X$ is hyperbolic relative to subgroups $P_1, \ldots, P_r$ stabilizing quasiflats $\widetilde{F}_1, \ldots, \widetilde{F}_r$, where X is sparse relative to $F_1, \ldots, F_r$. Suppose that X is virtually special. There exist finite index subgroups $P_1^o, \ldots, P_r^o$ such that for any further finite index subgroups $P_i^c \subset P_i^o$ the quotient group $G/\langle\langle P_1^c, \ldots, P_r^c\rangle\rangle$ is a word-hyperbolic group virtually having a quasiconvex hierarchy within the cubical small-cancellation category.*

Specifically, let $\widetilde{E}_i$ denote the superconvex hull of $\widetilde{F}_i$, and let $E_i^c = P_i^c \backslash \widetilde{E}_i$ for each i. Then the cubical presentation $\langle X \mid E_1^c, \ldots, E_r^c\rangle$ is $B(6)$ small-cancellation, and has a finite cover satisfying the $B(6)$ hierarchical conditions.

Remark 16.14. Lemma 16.13 should hold more generally for any smaller *virtually cyclic index* subgroups $P_i' \subset P_i^o$ in the sense that P_i/P_i' is virtually cyclic.

Proof. Note that the P_i-action on $\widetilde{F}_i$ extends to an action on $\widetilde{E}_i$. The definition of sparse implies that the $\widetilde{F}_i$ are isolated from each other, and after passing to $\widetilde{E}_i$, they are isolated from external hyperplanes as well, since any large overlap would be absorbed in $\widetilde{E}_i$. There is thus an upper bound D on the diameter of wall-pieces in the $\widetilde{E}_i$, and of cone-pieces between distinct translates of the $\widetilde{E}_i, \widetilde{E}_j$.

Let $\widehat{X}$ be a finite special cover, so that the induced covers $\widehat{E}_i$ have systole exceeding $12D$. Let $\ddot{X}$ be the finite cover corresponding to the wall homomorphism $\#_{\mathbb{W}}$, and likewise let $\ddot{E}_i$ be the induced covers. There is a wallspace on each $\ddot{E}_i$ induced by the hyperplanes of $\widehat{X}$. This (or rather $\langle \ddot{X} \mid \ddot{E}_i\rangle$) satisfies the $B(6)$ hierarchical condition. For any further cover $E_i^c \to \ddot{E}_i$, we give E_i^c the induced wallspace structure of Construction 9.4. We are proving a limited version of the Malnormal Special Quotient Theorem in a relatively hyperbolic framework. Hyperbolicity holds for $\pi_1 X^*$ by Theorem 15.3. The hierarchical $B(6)$ condition persists - but for $\langle \ddot{X} \mid E_i^c\rangle$. Letting $P_i^o = \pi_1 \ddot{E}_i$ gives the desired finite index subgroups. □

Implicit in the statement of Lemma 16.13 is the following variation on omnipotence (see Section 14.3) that is related to the work of Behrstock-Neumann on quasi-isometric classification of 3-manifolds.

Corollary 16.15. *Let M be a finite volume hyperbolic manifold and let $\partial M = \partial_1 \sqcup \partial_2 \sqcup \cdots \sqcup \partial_r$. There exist finite covers $\partial_i^\circ \to \partial_i$ such that for any further finite covers $\partial_i^c \to \partial_i^\circ$, there is a finite cover $\widehat{M} \to M$ such that for each i, the cover of ∂_i induced by $\widehat{M}$ is isomorphic to ∂_i^c.*

Proof. We use that $\pi_1 M = \pi_1 X$ where X is a virtually special sparse cube complex, as proven in Theorem 14.29.

We can assume without loss of generality that M is orientable and hence each component of ∂M is a torus. Indeed, if $\widehat{M} \to M$ is a finite cover, and ∂_{ij} are the various preimages in $\widehat{M}$ of some ∂_i, then we can choose ∂_i° to be a finite cover factoring through all covers ∂_{ij}° as j varies. Thus proving the statement for $\widehat{M}$ implies the statement for M itself.

For each i, let $P_i = \pi_1\partial_i$. By Lemma 16.13, there are finite index subgroups P_i° such that for each $\pi_1 M/\langle\langle P_1^c, \ldots, P_r^c\rangle\rangle$ is virtually special for any finite index subgroups $P_i^c \subset P_i^\circ$. Moreover, we can assume that $\langle X \mid E_1^\circ, \ldots, E_r^\circ\rangle$ satisfies the $C'(\frac{1}{12})$ small-cancellation condition (and in fact, some such assumption enables the proof of Lemma 16.13). It follows that $\langle X \mid E_1^c, \ldots, E_r^c\rangle$ is also $C'(\frac{1}{12})$ for any finite index normal subgroups $P_i^c \subset P_i$ with $P_i^c \subset P_i^\circ$. It is here were we use that each P_i is free-abelian, and hence $E_i^c \to E_i$ is automatically regular, so pieces between E_i^c and itself, already arise between E_i° and itself.

Let $X^* = \langle X \mid E_1^c, \ldots, E_r^c\rangle$. Observe that E_i^c lifts to an embedding in $\widetilde{X}^*$ by Theorem 4.1. It follows that the kernel of $P_i \to \pi_1 X^*$ equals P_i^c. Alternately, this follows from Theorem 12.16 since the map $\langle E_i \mid E_i^c\rangle \to \langle X \mid E_1^c, \ldots, E_r^c\rangle$ has no missing θ-shells.

Since $\pi_1 X^*$ is virtually special, it is in particular, virtually torsion-free. Let $\pi_1 X^* \to Q$ be a finite quotient whose kernel is torsion-free. The kernel of the composition $\pi_1 M \to \pi_1 X^* \to Q$ corresponds to a cover $\widehat{M} \to M$ whose boundary tori are ∂_i^c as desired. □

16.5. Separability for relatively hyperbolic hierarchies.

Theorem 16.16. *Let G be hyperbolic relative to abelian subgroups and suppose that G splits as a graph of groups whose edge groups are quasiconvex and whose vertex groups are fundamental groups of virtually special sparse cube complexes. Then each quasiconvex subgroup of G is separable.*

Remark 16.17. The proof of separability for the edge groups and probably for the cages (used in the proof of Theorem 16.28) is simpler.

Proof. Let G split as a graph Γ of groups with vertex groups G_v and edge groups G_e, and let H be a quasiconvex subgroup and $g \notin H$. By suitably filling the parabolic subgroups, we will produce a word-hyperbolic quotient $G \to \bar{G}$ with $\bar{g} \notin \bar{H}$, such that $\bar{H}$ is still quasiconvex and is also separable since $\bar{G}$ has a quasiconvex hierarchy and is thus virtually compact special. The existence of a consistent choice of parabolic fillings where all vertex quotient groups $\bar{G}_v$ are virtually compact special holds by Lemma 16.13. The hyperbolicity of $\bar{G}$ and the quasiconvexity of $\bar{H}$ hold by [Osi07, Thm 1.1]. (See also Theorem 15.3 and Theorem 15.6.) The virtual specialness of $\bar{G}$ follows from Theorem 13.3.

Choosing P^o: We now describe how to choose quotients $G_v \to \bar{G}_v$ that are virtually compact special hyperbolic, and that consistently induce quotients $G_e \to \bar{G}_e$ of the edge groups, so there is an induced quotient $G \to \bar{G}$ that splits as a graph of groups each of whose vertex groups $\bar{G}_v$ is virtually compact special hyperbolic. Each parabolic subgroup P of G intersects the vertex groups G_v in finitely many conjugacy classes of subgroups (some of which might be cyclic) denoted by P_{vi}. By Lemma 16.13, there are finite index subgroups P_{vi}^o of these P_{vi}, so that any further normal finite index subgroups P_{vi}^c yield virtually compact special hyperbolic quotients $\bar{G}_v/\langle\langle P_{vi}^c\rangle\rangle$. Note that P can have multiple intersections (i.e. with more than one vertex groups and/or conjugacy class within a vertex group). When P is elliptic relative to Γ we let P^o be a finite index subgroup such that $P^o \cap P_{vi} \subset P_{vi}^o$ for each P_{vi}. Otherwise P is split nontrivially by Γ and all the P_{vi} are commensurable with subgroups contained in a single codimension-1 normal subgroup P^* of P, and we let P^o be a characteristic finite index of P^* such that $(P^o \cap P_{vi}) \subset P_{vi}^o$ for each P_{vi}.

We now choose the the subgroups P^c so that $\bar{G}$ is word-hyperbolic, so that the quasiconvexity of each $\bar{G}_v$ in $\bar{G}$ is maintained, so that $\bar{H}$ is quasiconvex, and so that $\bar{g} \notin \bar{H}$. Consequently $\bar{G}$ is word-hyperbolic and virtually compact special by Theorem 13.1, and thus the quasiconvex subgroup $\bar{H}$ can be separated from $\bar{g}$ in a finite quotient of $\bar{G}$ and hence of G.

□

16.6. Residually verifying the double coset criterion.

Proposition 16.18 (Double Coset Criterion)**.** *Let X be a nonpositively curved cube complex with finitely many immersed hyperplanes. Then X is virtually special if (and only if) for each pair of immersed hyperplanes A, B and choice of basepoint $x \in A \cap B$, the double coset $\pi_1 A \pi_1 B$ is separable in $\pi_1 X$.*

We refer the reader to [HW08] and especially [HW10] where a version of this criterion is given that works in the presence of torsion.

Definition 16.19. A *hierarchy* for a nonpositively curved cube complex X is a sequence of cube complexes $X_0, X_1, \ldots, X_r$ where $X_0 = X$ and $X_r = X^0$ is a set of 0-cells, and for $0 \le i < r$ there is a 2-sided embedded hyperplane D_i in the cube complex X_i such that $X_{i+1} = X_i - N^o(D_i)$, where $N^o(D)$ denotes the open cubical neighborhood of D.

Every compact nonpositively curved cube complex whose hyperplanes are 2-sided and embedded has a hierarchy, so in particular this holds for any compact special cube complex. To accommodate infinite situations (and with an eye towards the corresponding hierarchy for $\pi_1 X$) we could allow the possibility that D_i is a collection of disjoint hyperplanes. In the compact case, this could permit a lower length hierarchy. In the infinite case, this might be the only way for the hierarchy to be of finite length. With this level of generality, any X with finitely many hyperplanes each of which is 2-sided and embedded has a hierarchy.

There are also relevant generalizations where each component of X_r has some property, and we would say the hierarchy terminates with this property - e.g., a hierarchy terminating at tori.

Theorem 16.20. *If X has all of the following properties then X is virtually special.*

1. *X is sparse.*
2. *$\pi_1 X$ is hyperbolic relative to the collection of virtually abelian subgroups P_i stabilizing the quasiflats in the action of $\pi_1 X$ on $\widetilde{X}$.*
3. *X has finitely many distinct immersed hyperplanes (see Remark 16.5).*
4. *X has a hierarchy.*

When X is compact and $\pi_1 X$ is hyperbolic, Theorem 16.20 is a geometric special case of Theorem 13.3.

Proof. Let us first describe our strategy for showing that $G = \pi_1 X$ is virtually special. For each nontrivial double hyperplane coset $\pi_1 A g \pi_1 B \neq \pi_1 A \pi_1 B$, let $U = \pi_1 A$ and $V = \pi_1 B$. We will produce a quotient $G \to \bar{G}$ where $\bar{U}\bar{g}\bar{V} \neq \bar{U}\bar{V}$, and $\bar{G}$ is hyperbolic and virtually compact special, and $\bar{U}, \bar{V}$ are quasiconvex subgroups. Since quasiconvex double cosets are separable for word-hyperbolic groups with separable quasiconvex subgroups [Git99, Min04], we see that there is a finite quotient $\bar{G} \to Q$ where these double cosets are separated, and this separates our original double cosets of G. (In fact, it appears that $\bar{U}, \bar{V}$ correspond directly to a double hyperplane coset in a word-hyperbolic virtually compact special cube complex, in which case the separability would follow from the "only if" part of Proposition 16.18.) Consequently, X is virtually special by Proposition 16.18.

With this strategy in mind, let us examine how to choose $\bar{G}$. For each P_i, we let $\widetilde{E}_i$ denote the superconvex hull of $\widetilde{F}_i$ in $\widetilde{X}$ and note that $\widetilde{E}_i \subset \widetilde{F}_i \cup GK$ for each i. We will let $\bar{E}_i$ denote a high-systole quotient of $\widetilde{E}_i$ by a normal subgroup P_i^c of P_i that is either of finite index or virtually $\mathbb{Z}$ index. Our quotient $\bar{G}$ is $\pi_1 X^*$ where $X^* = \langle X \mid \bar{E}_i \rangle$.

The double cosets remain distinct by a form of Lemma 12.22.(1) as indicated in Remark 16.21.

By induction on the length of the hierarchy, the vertex group(s) corresponding to the cube complex obtained by cutting along the first hyperplane are virtually special. As described in the proof of Theorem 16.16, we are thus able to apply Lemma 16.13 to choose normal subgroups P_i^o that are finite index or virtually-cyclic index in P_i depending on whether or not P_i is elliptic with respect to the splitting of the first hyperplane. These choices ensure that for any finite index characteristic subgroups $P_i^c \subset P_i^o$, the result will yield cubical small-cancellation word-hyperbolic quotients for all the vertex groups, that is consistent on the edge groups so that we obtain a quotient to a graph of word-hyperbolic groups. We emphasize that we are not quotienting by a finite index subgroup of entire parabolic subgroups that are not elliptic. This avoids "damaging" the stable letters of our graph of groups - as we aim to provide a graph of quotiented groups.

By Theorem 15.3, the group $\bar{G}$ is word-hyperbolic since we have performed a large filling on each peripheral subgroup. By Theorem 15.6, the subgroups $\bar{U}, \bar{V}$ are quasiconvex in $\bar{G}$.

Finally, the group $\bar{G}$ has a quasiconvex hierarchy because we chose each $\bar{E}_i$ so that $\langle X \mid \bar{E}_i \rangle$ satisfies the hierarchical $B(6)$ conditions. Some care is needed here to make sure that the original hierarchy for X projects to a hierarchy for X^*, and that the images of hyperplane groups have the expected presentations and are quasiconvex. (Use no missing θ-shells for the former property, and additionally short innerpaths for the latter property.) □

Remark 16.21. The statement and proof of Lemma 12.22.(1) does not quite apply in the proof of Theorem 16.20. To fit into our framework, we would need to add a graded set of cones that also includes the various hyperplanes of various codimensions in a filled torus. The grade of a cone is $(d-i)$ when it is associated to i-hyperplanes.

One then follows the proof of Lemma 12.22.(1) using a minimal complexity diagram. Even so, there is an extra issue to explain that has not arisen before. Recall that the diagram D for $a_1 g_1 a_2 g_2^{-1}$ that corresponds to a double coset intersection in $\pi_1 X^*$, decomposes into an annulus B_t surrounding a subdiagram D_t (we avoid the notation E_t used there) and that the strategy of the argument is to show that if D_t has a positively curved cone-cell then there is a contradiction.

We assume that each $\bar{E}_i$ is chosen so that the injectivity radius of carriers of hyperplanes in $\bar{E}_i$ exceeds five times the maximal diameter of wall-pieces and cone-pieces arising in the various $\widetilde{E}_i$, and moreover, the same should hold for intersections of $\bar{E}_i$ with the boundaries of such carriers. This assumption ensures that the following statement holds:

A typical positively curved θ-shell in D_t, as illustrated on the left in Figure 165, would be absorbable into A_1 and allow us to produce a lower complexity example. And such a θ-shell exists unless D_t is a single cone-cell or ladder with one end on g_1 and the other end on g_2, as illustrated in the second diagram of Figure 165.

However, in this ladder case, there is no guarantee that the two π-shells of the ladder absorb (or the single cone-cell absorbs in the degenerate case), as it is impossible to ensure that the overlap of each A_i with a flat F_j is bounded - indeed some hyperplanes do pass through flats. In particular, the "no missing θ-shell" style hypothesis bounding the overlap between $\widetilde{A}_i$ and $\widetilde{E}_i$ (called $\widetilde{Y}_i$ in the statement of Lemma 12.22.(1)) cannot immediately be achieved.

Instead we shall further assume that the cones $\bar{E}_i$ are product tori (in the cosparse case, either first apply Lemma 16.7 or require a product "up to the extra dimensions"). Now each cone-cell in the ladder is an essential path in $\bar{E}_i$ that is the concatenation of paths $\alpha_1 \to A_1$ and $\alpha_2 \to A_2$ with bounded subpaths concatenated between α_1, α_2. It therefore corresponds to a meridian-longitude path as illustrated in the fourth diagram of Figure 165, or perhaps a path that is just a single meridian or longitude. Consequently, that cone-cell can be decomposed into the concatenation of two "lower grade" cone-cells (and perhaps some squares). Subsequently, these two lower grade cone-cells can be absorbed into A_1, A_2 as on the right in Figure 165.

The reason why the fourth diagram in Figure 165 is applicable in higher dimensions is because we collapse the dimension of the torus corresponding to the codimension-2 hyperplane $A_1 \cap A_2$. So $\bar{E}_i$ fibers over the illustrated torus with fiber corresponding to the $A_1 \cap A_2$ parallelism class. The path in our drawing can be interpreted as being closed "up to a path in $N(A_1) \cap N(A_2)$".

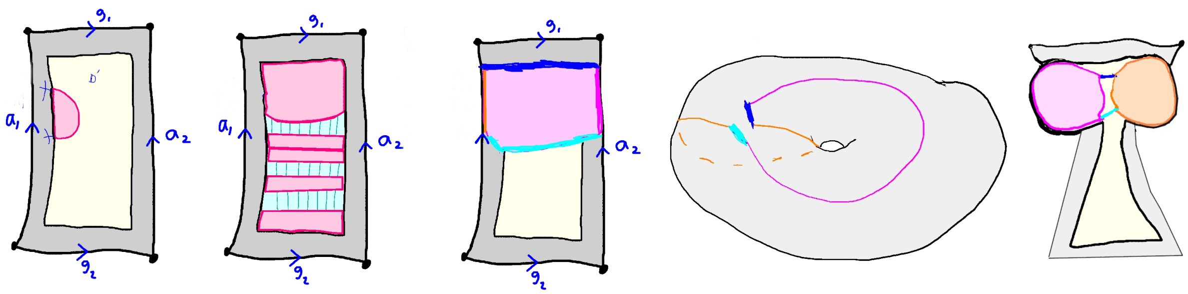

FIGURE 165. Peripherally Separating Hyperplanes:

Remark 16.22. Theorem 16.20 is a bit neater to prove under the additional direct assumption that hyperplanes are separable. In that case one can use a co-finite filling of the peripheral subgroups instead of a co-cyclic filling.

We note that the vertex groups are virtually special by induction on the dimension, and so we could apply the separability argument in Theorem 16.16 to see that the edge groups are separable, and use this to pass to a large cover relative to the initial splitting. This would enable a total filling the vertex peripheral subgroups.

Theorem 16.23 (Separability in sparse case)**.** *Let X be sparse and special, and let $H \subset \pi_1 X$ be a quasi-isometrically embedded subgroup. Then H is separable in $\pi_1 X$.*

Let H_1, H_2 be quasi-isometrically embedded finitely generated subgroups of $\pi_1 X$. Then each double coset $H_1 H_2$ is separable.

Proof. It is interesting to collect various available modes of proof.

Let us first sketch the proof given in [SW]: First apply Lemma 16.7 to assume without loss of generality that X is compact. Then extends H to a full subgroup by amalgamating with finite index subgroups of the parabolic subgroups where necessary, to obtain a subgroup H^+ (that retracts to H). The subgroup H^+ is represented by a compact based local isometry $Y \to X$ by Proposition 8.2. We then apply Proposition 6.3 to obtain a finite cover $\mathsf{C}(Y \to X)$ that retracts to Y, to see that $\pi_1 Y$ is a virtual retract, and is hence separable. This mode of proof was first introduced in a very primitive form in [Wis00].

Another proof is available through Lemma 16.13, since the image of $\bar{H}$ is quasiconvex for a large enough filling by Theorem 15.6, and $\bar{g} \notin \bar{H}$ when $g \notin H$ is a small element compared to the size of the filling.

Of course, this is a specific reenactment of the proof of Theorem 16.16, which proves Theorem 16.23 without knowing cubulation or even specialness.

Finally, Theorem 16.20 can be applied to the cube complex $X_\ominus$ that is the mapping cylinder of based sparse local isometry $Y \to X$. One then sees that $X_\ominus$ is virtually special, and now a finite index subgroup of H is now associated with a hyperplane.

This last viewpoint allows us to prove the double coset separability as well: Indeed, having formed $X_\ominus$ as above with $H = H_1$, we now let $Y_2 \to X$ be a sparse based local isometry that was thickened up to contain at least one 1-cell dual to the newly added hyperplane corresponding to Y_1. We then form $X_\oplus$ to be the mapping cylinder of $Y_2 \to X_\ominus$. The sparse cube complex $X_\oplus$ has a hierarchy (since X does) and is thus virtually special. The double hyperplane coset provides the double coset separability. □

16.7. **Relative malnormality and separability.** We now revisit the notions of height from Section 12.2 in the relatively hyperbolic case. This will be used in the proof of Theorem 16.28.

Definition 16.24 (Relative Height and Malnormality)**.** Let G be relatively hyperbolic.

The *relative height* of $H \subset G$ is the smallest number h such that for any $h+1$ distinct cosets $\{Hg_1, \ldots, Hg_{h+1}\}$ the intersection $\cap_i H^{g_i}$ is either parabolic or elliptic.

The subgroup H of G is *relatively malnormal* if $H \cap gHg^{-1}$ is either elliptic or parabolic for each $g \in G - H$.

Thus H is relatively malnormal exactly when its relative height is 0 or 1. In particular, any malnormal subgroup and any parabolic or elliptic subgroup is relatively malnormal.

The following was proven in [HW09]:

Proposition 16.25. *Let $H \subset G$ be a relatively quasiconvex subgroup of a relatively hyperbolic group. Then H has finite relative height.*

The following was proven in [HW09, Thm 9.3]:

Proposition 16.26. *Let H be a separable, relatively quasiconvex subgroup of the relatively hyperbolic group G. Then there is a finite index subgroup K of G containing H such that H is relatively malnormal in K.*

A subgroup $H \subset G$ is *isolated from parabolics* if for each parabolic subgroup P and each $g \in G$, the intersection $H \cap P^g$ is either finite or equal to P^g. When H is isolated from parabolics, the following Lemma strengthens the conclusion of Proposition 16.26 to obtain almost malnormality.

Lemma 16.27. *Let G be relatively hyperbolic, and let H be a quasiconvex subgroup that is isolated from parabolics. If H is separable then there exists a finite index subgroup $G' \subset G$ such that H is an almost malnormal subgroup of G'.*

Sketch. If Hg_iH is a double coset with the property that $H^{g_i} \cap H$ is infinite, then H and g_iH have infinite coarse intersection in the sense that $N_r(H) \cap g_iH$ has infinite diameter for some r. Consider a geodesic rectangle of width r and length $R \gg r$ whose four vertices are in H and g_iH. We cut this rectangle into two relatively δ-thin triangles and find that either there are large overlapping peripherals in the middle that also have large overlaps with H and g_iH, or there is a bounded short cut from one side to the other down the middle. Thus, either H and g_iH have large fellow travel with some parabolic P and hence H, H^{g_i} contain P, or H, g_iH are within a bounded distance of each other.

Consider the finitely many double cosets Hg_iH corresponding to the bounded distance case. The separability of H allows us to choose a finite index normal subgroup $N \subset G$ such that the images of these double cosets are separated in G/N. We let $G' = NH$.

Let $g \in G'$ be such that $H^g \cap H$ is infinite, and note that our choice of G' ensures that gH and H have large overlap with the same parabolic P. Note that by isolation, $P \subset H$.

Considering an infinite rectangle whose sides are stabilized by h and h^g we see that both h and h^g are in P, and thus since $|h| = \infty$ and since P is almost malnormal, we see that $g \in P$.

But then $g \in H$ which was our desired conclusion. □

16.8. **Hierarchy with all peripherals busted at first stage.** We now describe a generalization of Theorem 13.1 that holds in particular for a group that is hyperbolic relative to abelian subgroups, and that splits over a collection of quasiconvex subgroups with word-hyperbolic virtually compact special vertex groups. Our motivation is to the 3-manifold application since by Corollary 14.16, for each cusped hyperbolic 3-manifold M there is a finite cover $\widehat{M}$ such that $\pi_1\widehat{M}$ has such a hierarchy. As usual, the case of a graph of groups can be deduced from the case of a loop or edge of groups.

Theorem 16.28. *Let G be torsion-free and hyperbolic relative to virtually abelian subgroups. Suppose that G splits as a graph of groups where each edge group is quasiconvex and each vertex group is word-hyperbolic and virtually compact special. Then G is virtually compact special.*

Definition 16.29 (Accidental parabolic)**.** Let G be hyperbolic relative to peripheral subgroups $\{P_i\}$. Suppose G splits as a graph of groups. Let e be an edge in the Bass-Serre tree T of the splitting. An element g in $G_e = \text{Stab}(e)$ is an *accidental parabolic* if g is conjugate to an infinite order element in some parabolic subgroup P, but e lies outside a (minimal) nonempty P-invariant subtree of the Bass-Serre T of G.

From the viewpoint of the graph of groups, an accidental parabolic corresponds to an infinite order parabolic element of an edge group such that the parabolic subgroup isn't essentially "split" by that edge. This notion agrees with the notion of accidental parabolic in an incompressible surface S by considering the splitting of $\pi_1 M$ along $\pi_1 S$.

Remark 16.30. As stated, Theorem 16.28 is limited to the case where all abelian subgroups have rank ≤ 2. However, we actually prove a slightly more general statement handling the case where: G is hyperbolic relative to virtually abelian groups, and G splits as a graph of groups where each edge group is quasiconvex, and there are no accidental parabolics in edge groups, and moreover each inclusion of each edge group into π_1 of the subgraph(s) of groups omitting it does not intersect any noncyclic parabolic subgroups.

Moreover, if the vertex groups are virtually compact/sparse special then so is G.

Strategy of proof: If each edge group were aparabolic, then combining Theorem 16.16 with Proposition 16.26, we could pass to a finite index subgroup G' whose edge groups are malnormal and aparabolic. We could then use the hierarchy associated to a sequence of splittings associated with the graph of groups Γ'. At each stage Proposition 7.8 can be applied to cubulate, and then Theorem 16.20 is applied to obtain virtual specialness.

The trickier situation where edge groups are not aparabolic is instead handled using a "variant hierarchy" associated to the sequence of edges of Γ' yielding a sequence of subgraphs Γ'_r. When an edge group E of Γ'_r is malnormal (and aparabolic) then we use the usual splitting as above. Otherwise we use a different splitting $E^+ *_{K_E} V_E$ along a "cage" K_E that is malnormal and aparabolic in $V_E = \pi_1 \Gamma'_{r+1}$. The group E^+ is readily seen to be virtually special. This allows us to employ the general version of Proposition 7.8. The resulting cube complex at each stage is sparse (or compact in favorable circumstances) and we can thus apply Theorem 16.20 to obtain virtual specialness at each stage.

Proof. For an edge group E of a group G splitting as a graph of groups Γ, the *expanded edge group E^+* consists of the multiple HNN extension with base vertex group E and with edge groups corresponding to the various infinite intersections of maximal parabolic subgroups with E. We think of E^+ as a graph of groups with "initial" and "terminal" vertex groups isomorphic to an edge group E, and then a collection of arcs starting and ending on these initial and terminal vertex groups. The internal vertex and edge groups of each such arc are all isomorphic to the same subgroup of a parabolic group which is associated to as many arcs as the number of its conjugates intersecting E. Finally, there is a homomorphism $E^+ \to G$ which is induced by a map of underlying graphs of groups. See Figure 166.

The *cage* K_E associated to the edge E is the graph of groups obtained from the graph of groups for E^+ by removing the open edge associated to E itself. It thus has two distinguished vertex groups - the initial and terminal image groups of E, and otherwise has arcs with parabolic stabilizers that start and end on these respectively. We will not consider the cage K_E in the disconnected case.

Step 1: Cage π_1-injectivity, quasiconvexity, and malnormality: By first passing to a finite index subgroup G' of G we can assume that the cages have the following properties:

1. Each cage maps injectively into G'.
2. Each cage maps quasi-isometrically into G'.
3. Each cage maps to a malnormal subgroup of the group $V_E \subset G'$ that is the graph of groups obtained by removing the open edge E.

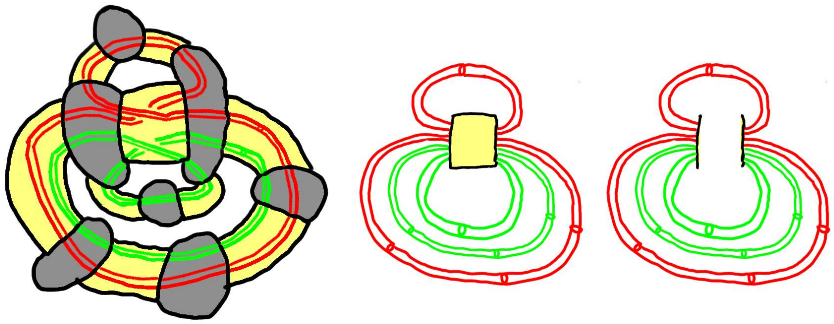

FIGURE 166. A graph of spaces on the left (the peripherals are highlighted on inside), an expanded edge space in the middle, and its associated cage on the right. In this example, the two parabolic subgroups are represented by immersed tori. The bold vertical lines in the cage correspond to the initial and terminal images of the edge space, the small circles in the cage correspond to places where the parabolic torus passes through vertex spaces, and the cylinders in the cage correspond to places where the parabolic torus passes through edge spaces.

Explanation in our setting: Each edge group in the graph of groups is quasi-isometrically embedded and thus relatively quasiconvex. It therefore has finite relative height by Proposition 16.25. Let n be the number of edge groups, and let m be the maximal height of these edge groups. Then for any arc A in the Bass-Serre tree T, if $|A| > mn$ then Stab(A) is parabolic or elliptic. Indeed, one notes that stabilizers of distinct edges of A correspond to distinct cosets of edge groups, and then applies the pigeon-hole principle.

We choose a finite index subgroup of G with the property that in its associated graph of spaces, if two tori enter a vertex space through the same edge space, and leave through different edge spaces in the universal cover, then they leave through different edge spaces in this finite cover. In particular, tori cannot backtrack from an edge space by entering a vertex space, and returning essentially to the same edge space. This involves two steps: Firstly there is a finite index subgroup of each vertex group which separates the finitely many double cosets of edge groups corresponding to the overlaps that must be distinguished. Secondly, we use separability of the finite index subgroups of the vertex groups to obtain the separation in a finite index subgroup of the whole group. Finally, we pass to a finite cover induced by covers of the underlying graph to obtain girth $> nm$.

Now expanded edge groups correspond to nearly embedded subspaces in the induced cover of graph of spaces. It is then easy to recognize that expanded edge groups inject using the normal form theorem for graphs of groups. They can also be proven to quasi-isometrically embed along these lines.

More general explanation: By residual finiteness, maximal abelian subgroups are separable (both are consequences of Theorem 16.16), and so we can pass to a finite index subgroup such that all peripheral subgroups are free-abelian. So let us assume that G has this property to begin with.

For each expanded edge group E^+ and each n, there is an *immersed expanded edge group* E^+_n and a map $E^+_n \to E^+$. The group E^+_n is a multiple HNN extension of E where each stable letter centralizes the intersection between E and a peripheral subgroup. Moreover, each peripheral subgroup of E^+_n wraps n times around the corresponding peripheral subgroup of E^+. An exact definition is a bit tedious, but we refer the reader to Figure 167.

It follows from basic results about quasigeodesics in relatively hyperbolic spaces that $E^+_n \to G$ is a quasi-isometric embedding for all sufficiently large n (see e.g. [HWc]). Moreover, each E^+_n is *isolated* from peripheral subgroups, in the sense that $\widetilde{E^+_n} \cap N_L(P)$ has finite diameter unless P is commensurable with one of the peripheral subgroups appearing in E^+_n.

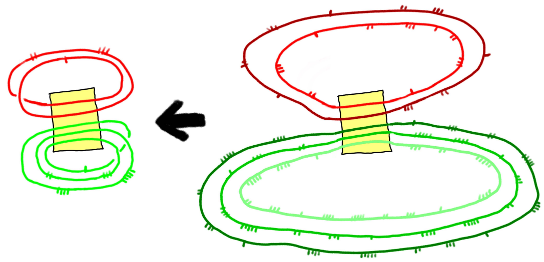

FIGURE 167. The high girth expanded edge space E_3^+.

Since (immersed) expanded edge groups are separable by Theorem 16.16, by Lemma 16.27, for each E there exists a finite index subgroup $J_E \subset G$ such that E_n^+ is a malnormal subgroup of J_E. We let G' denote a finite index normal subgroup of G that is contained in each J_E.

Since each expanded edge group E^+ is malnormal in G', it follows that each cage K_E is malnormal in the group V_E (the cage splitting would have underlying graph a bouquet of two edges). Moreover, this statement is inherited by further cages and subgraphs as we proceed down the hierarchy. Assuming there are no accidental parabolics, if $E^+ = E$ then E is already malnormal and aparabolic in G, and we will use the usual splitting of G. When $E^+ \neq E$ then since each torus passing through E is actually split by E, we see that E does not disconnect the underlying graph of V_E and we will use the splitting $E^+ *_{K_E} V_E$ along the associated cage.

Step 2: Splitting along cages: Each cage corresponds to a graph of spaces whose two vertex spaces are the original incoming and outgoing edge spaces (the two sides of the associated edge space), and whose edge spaces look like split peripherals (codimension-1 subgroups of peripherals intersecting edge groups).

For each edge E of the underlying graph of the splitting of G', we now describe an alternate splitting of G' as an amalgamated free product along the cage associated to E. One vertex group of this splitting is the group V_E corresponding to the graph of groups obtained by omitting E. The other vertex group is the expanded edge group E^+. The edge group is the cage K_E. So we claim that $G' \cong E^+ *_{K_E} V_E$ where the inclusions of K_E are the natural ones.

To verify this we will rethink it geometrically as follows. Let $\phi : Y^+ \to X$ be the inclusion of the graph of spaces corresponding to E^+ into the graph of spaces X corresponding to the splitting of G'. The mapping cylinder M_ϕ deformation retracts to X by pushing the cylinder forwards, so $\pi_1 M_\phi \cong \pi_1 X$. There is another deformation retraction which pushes the edge space corresponding to E upwards to the copy of this edge space in Y^+. The resulting space is isomorphic to the graph of spaces naturally corresponding to an amalgamated product $E^+ *_{K_E} V_E$.

Step 3: The new vertex groups are cubulated: We now show how to cubulate each expanded edge group E^+. We contract the edge associated to E so that the edge group isomorphic to its initial and terminal vertex groups is now a single base vertex group. Each arc in the underlying graph is associated to a particular parabolic subgroup, and collecting them together according to this equivalence relation, we find that they form a collection of "parabolic cycles" in the underlying graph of the splitting, so that each edge lies in a unique cycle. We reform the underlying graph (see Figure 168) so that it consists of a disjoint union of circles corresponding to the parabolic cycles above, and there is an edge joining the base vertex to a vertex on the circle when the corresponding point on the parabolic cycle passes through the base vertex.

Easy case: There is an easier way to proceed in the special case where each circle has exactly one connecting edge to the base vertex, and where each parabolic is a trivial torus bundle over the circle (but without assuming that parabolics are virtually $\mathbb{Z}^2$). We choose a fixed cubulation B of the base vertex group. Each torus T_i, is attached to the base along a subgroup represented by $B_i \to Y$. We let

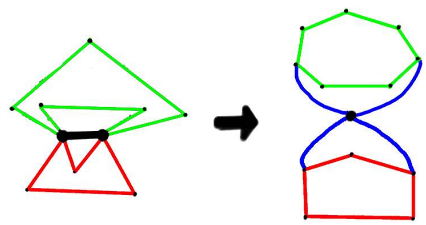

FIGURE 168.

$T_i \cong B_i \times (S^1)^{r_i}$ and then attach the two using an edge space $B_i \times [-1, 1]$. The motivating case is where $r_i = 1$ but the argument works in general. Note that the resulting cube complex is compact if and only if the base cube complex is compact as this implies that B_i is compact.

In general: The general case is handled similarly.

Assume that the finite index subgroup $G' \subset G$ is normal and assume that all peripheral subgroups are actually free-abelian and not just virtually abelian.

If each edge group $\bar{E}$ of G has a finite index subgroup E_o with a compact cubulation, then by Theorem 16.16, we can assume that each edge group E of G' factors through the corresponding E_o.

Recall that we assumed the covering space corresponding to G' factors through covers that contained expanded edge groups E^+_n each of whose tori intersects E in a single subtorus. This property is not stable under further covers, but it does imply that in G', each E^+ has the property that the automorphism group of the covering space $E \to \bar{E}$ between edge spaces of the graph of spaces, extends to an action on E^+. (We use $E, \bar{E}$ for corresponding edge spaces associated to G', G.)

Thus, each torus intersects E in a collection of "isomorphic places". Choosing a cubulation of $\bar{E}$ there is an induced cubulation of E, and this extends to a cubulation of E^+. The easiest way of doing this is to choose local isometries representing the intersections with the peripherals, and then gluing these together to a common torus. In our case of interest this amounts to attaching a collection of $A \times [-1, 1]$ to the cubulation B of E using the various isomorphic cores to attach $A \times \{\pm 1\}$. A more general form of this which works without the restriction on accidental parabolics must handle tori attached along higher codimension summands. In this case we form a graph of spaces akin to that illustrated on the right in Figure 168. The central vertex space is the cubulation B of E. Each peripheral subgroup has a cubulation $A \times (S^1)^{n_A}$ for some n_A, where A is isomorphic to the way this peripheral subgroup (possibly accidentally) overlaps with the cubulation B of E, and then we attach various copies $A \times [-1, 1]$ to corresponding locations in $A \times (S^1)^{n_A}$ and B.

We note that when E has a compact symmetric cubulation (as is the case when $\bar{E}$ does, then the cubulation of E^+ is compact.

Step 4: The cubulation: Since each cage K_E is malnormal in the old vertex group V_E and quasiconvex in the whole group, we obtain a cubulation as follows:

Choose cubulations of the new vertex groups (the expanded edge groups) as in Step 3. By passing to a finite index subgroup $(E^+)'$ of E^+, and using the walls from *its* cubulation corresponding to the intersection of the original codimension-1 subgroups with $(E^+)'$, we can assume that the stabilizers of walls of the cubulation of the cages from the new vertex groups lie in special subgroups of the old vertex groups.

Observe that by Proposition 8.2, the embeddings of the cages are represented by compact local isometries in the cubulations of the old vertex groups, since the cages are malnormal and quasiconvex. More precisely, when the old vertex group V_E has a compact cubulation then we obtain compactness of these representations, and otherwise we obtain sparseness. In the stated version of the theorem, cages are aparabolic and thus we obtain a compact representation in the cube complex of V_E.

We can thus apply the extension property of Theorem 6.11 to extend walls into the old vertex groups. We note that V_E is virtually special [compact/sparse] and cubulated by induction.

We will thus be able to apply the variant of Proposition 7.8 indicated by Restatements (5′) and (6′). (We refer to this as Proposition 7.8′.)

To preserve the cocompactness, we note that as explained in Remark 6.13, the new walls do not provide new commensurability classes of codimension-1 subgroups in the parabolic subgroups, and so we will not "overcubulate" them. So the relative cocompactness provided by Theorem 7.5 is genuine cocompactness, since the cubulations C_i of P_i remain cocompact after adding the new walls.

Conclusion: The result now follows by induction using the hierarchy obtained by successively splitting along the edge groups of the induced splitting of G'. While cages might be a bit smaller within a subgraph of groups, their malnormality and injectivity properties persist. To clarify, the splitting of G' yields a hierarchy (just split along the various edges). Again, as we progress through this hierarchy, a separating edge yields a splitting whose edge group is malnormal on each side, and while the vertex groups are relatively hyperbolic, the edge group contains no noncyclic parabolic subgroups and so the result can be cubulated directly using Proposition 7.8. For a nonseparating edge, we use the splitting along the cage. Either way, we obtain the desired cubulation. After each level in the hierarchy, after cubulating, we apply Theorem 16.20 to verify the virtual specialness of the cubulation that was produced. □

Repeating the theme of Problem 14.4 we have:

Problem 16.31. Let G be hyperbolic relative to virtually abelian subgroups, and suppose that G is virtually special. Does G (virtually) have a hierarchy terminating at the parabolic subgroups?

17. ⊛ Limit Groups and Abelian Hierarchies

The goal of this section is to examine some extremely simple quasiconvex hierarchies which are of specific interest. We apply the results of Section 16 to prove that limit groups are virtually compact special in Section 17.1. We prove that relatively hyperbolic groups with abelian hierarchies are virtually sparse special in Section 17.2.

17.1. **Limit groups.** A *fully residually free group* or *limit group* is a group J with the property that for any finite set of nontrivial elements $\{j_1, \dots, j_k\}$ there is a free quotient $J \to \bar{J}$ such that $\bar{j}_i$ is nontrivial for each $i \in \{1, \dots, k\}$. It was shown in [KM98] that every limit group is a subgroup of a group G_r constructed through the following hierarchical structure:

(1) G_0 is a trivial group.
(2) For each $i \geq 0$, we have $G_{i+1} \cong G_i *_C A$ where C is a malnormal abelian subgroup of G_i and $A \cong C \times B$ is a finite rank free-abelian group.

Remark 17.1. Considering the corresponding topological space (a bouquet of circles with a sequence of tori attached to it along maximal pre-existing tori), we see that each maximal torus in the resulting space can be attached at the initial stage where its first rank ≥ 2 torus is attached along some cyclic root. We thus find that we can assume that each subgroup C above is cyclic.

Lemma 17.2. *Each group G_r above is the fundamental group of a compact cube complex that is virtually special.*

Proof. The cubulation of G_i induces a cubulation of C whose walls extend to walls of $A \cong C \times B$ using the product structure, so we can assume they don't cross in the flat corresponding to A. The other walls of A are induced from B using the product structure. The cubulation now follows from Proposition 7.8.

We note that it is unnecessary to utilize turns, so a compact cubulation of G_i provides a compact cubulation of G_{i+1}. □

Since finitely generated subgroups of G_r are quasi-isometrically embedded [Kap02], we can apply Proposition 8.2 to obtain the following consequence:

Corollary 17.3. *Every f.g. limit group is the fundamental group of a virtually special sparse cube complex.*

Using Sela's retractive tower description of limit groups [Sel03], one obtains the stronger result that:

Theorem 17.4. *Every limit group is the fundamental group of a compact nonpositively curved cube complex that is virtually special.*

Proof. We must verify the preservation of the virtually compact special property under two constructions:

In the first case $G = A *_C (C \times \mathbb{Z}^n)$ where A is compact and virtually special, and C is a malnormal subgroup of A. In this case let X be a compact nonpositively curved cube complex with $A = \pi_1 X$, and let $Y \to X$ be a compact representation of C by Proposition 8.2. Let $Z = (X \sqcup (Y \times (S^1)^n))/(y, 0, \ldots, 0) \sim \phi(y) : y \in Y$. Then Z is a nonpositively curved cube complex with $\pi_1 Z \cong G$. Let $\hat{X}$ be a special cover of X, and let $\hat{Z}$ be the cover of Z induced by the retraction $Z \to X$. Then $\hat{Z}$ is a nonpositively curved cube complex with a quasiconvex hierarchy and $\pi_1 \hat{Z}$ is hyperbolic relative to virtually free-abelian subgroups, so $\hat{Z}$ is virtually special by Theorem 16.20.

In the second case G splits as a graph of groups with two vertex groups A, B where A is compact virtually special, and $B \cong \pi_1 S$ with S an orientable surface. The edges of the graph correspond to the boundary circles of S. Each edge group is infinite cyclic and maps to A on one side, and embeds in B as the fundamental group of the corresponding boundary circle. Finally, there is a retraction $G \to A$.

If S is a cylinder, then the retraction map shows that the two boundary circles represent elements that are conjugate in A, and hence this was covered by the first case (with $n = 1$). Otherwise, each edge group is malnormal and aparabolic in S. There is a version of Proposition 7.8′ that can handle this graph of groups situation, but instead, we add dummy squares to obtain a new splitting as an amalgamated product $A *_{F_r} (B * F_{r-1})$ over a free group (where ∂S has r components).

Of course, as the group has not changed, there is still a retraction to A. Topologically, this corresponds to the dummy squares having two opposite edges on edge spaces, and one edge on a topological space for A, and the free face in the vertex space for $B * F_{r-1}$. An easy fiber product argument shows that F_r is malnormal in $B * F_{r-1}$, and of course F_r is aparabolic in $B * F_{r-1}$ and quasiconvex in G. We apply Proposition 7.8 to cubulate, and then pass to a finite cover so that the A part of the cube complex is special.

We must also make sure that the $B * F_{r-1}$ part of the cube complex has a hierarchy. That can be done using the details of the proof of Proposition 7.8′. It can also be achieved by using Wilton's result of separability which follows from Corollary 17.3. The last possibility is to avoid this by proving a form of Theorem 16.20 that allows some virtualness in the hierarchy. □

I don't know if passing to a finite index subgroup is necessary:

Problem 17.5. Is every limit group the fundamental group of a special cube complex?

In [KM98] it was shown that every limit group can be produced from free-abelian groups and free groups using a *separated cyclic hierarchy* whose splittings are all free products, or HNN extensions or amalgamated free products over an infinite cyclic subgroup where the cyclic edge group is malnormal on at least one side, and its two inclusions do not have intersecting conjugates in the vertex group (in the HNN case). This is also apparent from the retractive tower viewpoint.

Theorem 17.6. *Every group with a separated cyclic hierarchy is the fundamental group of a sparse nonpositively curved cube complex that is virtually special.*

Proof. We described how to cubulate a free-abelian extension of an infinite cyclic subgroup above. Of course, if the original group has a sparse cubulation, then the extension is sparse as well, since the geometric result is of the form $X *_A (A \times (S^1)^n)$ for some n where $A \to X$ is a local isometry representing the maximal abelian subgroup (that is cyclic in this case).

Now consider an HNN extension $G \cong H*_{C^t=D}$ where H is the fundamental group of a sparse nonpositively curved cube complex that is virtually special, and C is infinite cyclic, and C is malnormal in H, and C, D do not have conjugates in H with nontrivial intersection.

The reader will require an understanding of the HNN version of Proposition 7.8′ to understand this proof.

By Lemma 16.13 and Remark 16.14, we pass to a quotient $H \to \bar{H}$ with the property that $\{\bar{C}, \bar{D}\}$ form a malnormal collection of subgroups. (The key is for $\bar{D}$ to equal the image of the centralizer of D.) Following the style of proof in Section 14.3, we choose a finite index normal subgroup $\bar{N} \subset \bar{H}$ and a quotient $\bar{N} \to \mathbb{Z}$ such that $\bar{D} \cap \bar{N}$ has trivial image in $\mathbb{Z}$ but $\bar{C} \cap \bar{N}$ has infinite image. This allows us to produce walls in H that hit C but miss D. These can be represented by finitely generated codimension-1 subgroups (which must then be quasi-isometrically embedded because of the structure of H: indeed all its finitely generated subgroups are quasi-isometrically embedded (see [Kap02] and the references therein). We can thus recubulate H so that C passes through more hyperplanes than D. The cubulation result now follows from the HNN version of Proposition 7.8.

The details of Proposition 7.8 are comparatively easy here: Cubulate G by extending all walls of D into C, and then turn the excess walls from H into C back to H. (*Turns* are discussed in [HWc].)

The cubulation is sparse by Corollary 7.7, and thus virtually special by Theorem 16.20. □

Problem 17.7. Is every group with a separated cyclic hierarchy virtually a limit group?

17.2. **Abelian hierarchies (preliminary).**

Definition 17.8. G has an *abelian hierarchy* if it has a hierarchy where at each step, the edge group C in $A *_C B$ or $A*_{C^t=C'}$ is a finitely generated free-abelian group.

Remark 17.9. Let G be hyperbolic relative to free-abelian subgroups. We note that G then has the property that any nontrivial abelian subgroup $C \subset G$ is contained in a unique maximal abelian subgroup of G. Consequently, if G also has an abelian hierarchy then it has a hierarchy terminating at free-abelian subgroups where at each stage either:

(1) C is a free-abelian subgroup of $A *_C B$ and C is malnormal in A and lies in a maximal malnormal abelian subgroup of B
(2) C is a free-abelian subgroup of $A*_{C^t=C'}$ and C' lies in a maximal abelian subgroup $D \subset A$, and $\{C, D\}$ is malnormal in A.

The following is proven in [Dah03, BW]:

Lemma 17.10. *Suppose that G is hyperbolic relative to free-abelian subgroups, and that G has an abelian hierarchy. Then every finitely generated subgroup of G is quasi-isometrically embedded.*

Lemma 17.11 (Absorbing walls). *Let X be a virtually special sparse cube complex, where $G = \pi_1 X$ is hyperbolic relative to $\{P = P_0, P_1, \ldots, P_r\}$, and each P_i is free-abelian. Let H be a codimension-1 subgroup of P with associated H-wall W_H of P. There exists a finite cover $\widehat{X}$ and an embedded "base" K-wall W_K in $\widehat{X}$ that extends W_H, so that W_K crosses copies of P only in one or more immersed submanifolds parallel to H, and crosses no P_i with $i \geq 1$.*

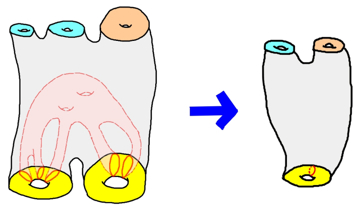

FIGURE 169. The wall W_H is indicated within the torus P_0 at the bottom on the right. Its "absorbing wall" extension is indicated in the covering space on the left.

Proof. Let P'_0 be a sufficient finite index subgroup of H, and let P'_i be a sufficient finite index subgroup of P_i for each $i \geq 1$, so that the group $\bar{G} = \pi_1 X / \langle\langle P'_0, P'_1, \ldots, P'_r \rangle\rangle$ is virtually compact special hyperbolic by Lemma 16.13 (as varied in Remark 16.14.

Let J be a special finite index subgroup of $\bar{G}$, and let Z_0 denote $J \cap \bar{P}_0$. Proposition 6.3 provides a finite index subgroup $J' \subset J$ and a retraction $J' \to Z_0$.

For each i, let $\phi_i : F_i \to X$ be a local isometry with $P_i = \pi_1 F_i$, after adjusting basepoints. Let $X_{\circledcirc} = (X \sqcup \sqcup_i F_i \times [-1,1])/(\phi_i(x) \sim (x,-1) \colon x \in F_i)$ be obtained from X by attaching a copy of $F_i \times [-1,1]$ along each F_i. We let $T_i = F_i \times \{1\}$ be the corresponding subspaces of the cube complex $X_{\circledcirc}$, and we note that the T_i are disjoint from each other. Let $\widehat{X_{\circledcirc}}$ be the based cover of $X_{\circledcirc}$ corresponding to J'.

Identify Z_0 with $\pi_1 S^1$, where S^1 is a circle consisting of a single 0-cell and 1-cell with barycenter b. The retraction $J' \to Z_0$ allows us to choose a map $\widehat{X_{\circledcirc}} \to S^1$ such that covers $\widehat{T}_{0j}$ of T_0 map to S^1 transversely at b. In particular, the preimage of b in each $\widehat{T}_{0j}$ consists of subtori "parallel" to H. For $i \geq 1$, the various $\widehat{T}_{ij}$ can be chosen to map to the 0-cell of S^1. The map $\sqcup \widehat{T}_{ij} \to S^1$ extends to a map $\psi : \widehat{X_{\circledcirc}} \to S^1$ that is transverse to b.

The base component W of $\psi^{-1}(b)$ yields our desired K-wall in $\widehat{X_{\circledcirc}}$, where $K = \pi_1 W$ is the fundamental group of the base component. The image of K in $X_{\circledcirc}$ gives an immersed K-wall extending several immersed copies of covers of our H-wall. □

When G is a finite volume hyperbolic 3-manifold group with a cusp P, we can interpolate a step in the proof, so that there are multiple cusps, and then the construction shows that K cannot be a virtual fiber (since it misses some cusps!). Consequently K is geometrically finite. We are then able to obtain the following:

Corollary 17.12. *Let M be a compact 3–manifold whose interior admits a complete hyperbolic metric of finite volume. Let C be a circle in ∂M. Then there is a finite cover $\widehat{M}$ with an oriented properly embedded incompressible geometrically finite surface S with ∂S covering C.*

Theorem 17.13. *Suppose that G is hyperbolic relative to free-abelian subgroups, and that G has an abelian hierarchy terminating in virtually special groups. Then G is virtually sparse special.*

In particular, we obtain the following:

Corollary 17.14. *Suppose that G is hyperbolic relative to free-abelian subgroups, and that G has an abelian hierarchy. Then G is virtually sparse special.*

Proof of Theorem 17.13. We will prove that G has a finite index subgroup G' such that $G' \cong \pi_1 S$ where S is a sparse cube complex with a hierarchy. It follows that S is virtually special by Theorem 16.28.

By Remark 17.9, G has a hierarchy terminating in free-abelian subgroups (these base cases are obviously cubulated), with certain malnormality properties. It suffices to consider the case of an HNN extension: $A*_{C'=C'}$. The case where C is trivial is easy, so let us assume that C is infinite. Then C' lies in a maximal abelian subgroup D, and $\{C, D\}$ is malnormal in A.

Let $A = \pi_1 X$ where X is sparse and virtually special - by induction on the length of the hierarchy. By Lemma 17.11, we can extend the system of immersed walls in X, or rather $X_{\circledcirc}$ as in Lemma 17.11, to a new system such that considering the isomorphism between C and C' represented by $T \to T'$, each commensurability class of wall in C is represented by an immersed wall in C' that extends to an "absorption" wall in $X_{\odot}$ that does not interact at all with C, and similarly, each commensurability class of wall in C' is represented by an immersed wall in C that extends to an "absorption" wall in $X_{\odot}$ that does not interact at all with C'. Since the new walls have quasi-isometrically embedded stabilizers, the cube complex dual to the associated wallspace in $X_{\circledcirc}$ is cosparse, and we obtain a new cube complex Y with $\pi_1 Y \cong \pi_1 X$, and with the property that C, C' are now represented by local isometries $E \to Y$ and $E' \to Y$.

A finite special cover of X has a hierarchy, and the corresponding cover of Y has this same cubical hierarchy, so Y has a finite special cover Y_1 by Theorem 16.28.

As in Corollary 16.15, we now apply Lemma 16.13 to Y_1 in order to obtain a cover Y_2 so that $\pi_1 Y_2$ induces the same finite index subgroups of C, C'. Finally, we subdivide the absorption hyperplanes of C and of C' so that there are the same number of hyperplanes on these tori in Y_2. (It is easier to think of subdividing within Y, and letting that induce subdivisions in Y_2.)

In conclusion, we can now choose a one-to-one correspondence between the E and E' tori elevations in Y_2, and the walls within them of various types, and cubulate a finite index subgroup of G as planned. Again the cubulation is sparse since the wall stabilizers are quasi-isometrically embedded by Lemma 17.10 and so Proposition 7.5 applies. □

Problem 17.15. Suppose G is relatively hyperbolic and has an abelian hierarchy. Does G virtually have a compact cubulation?

We now attempt to provide a more cutting version of Lemma 17.11 that is more along the lines of Section 6.5. This is preliminary.

Lemma 17.16. *Let G be virtually special and hyperbolic relative to abelian subgroups. Let P be a full quasiconvex malnormal subgroup with a codimension-1 quasiconvex subgroup H and H-wall W_H within P. Suppose $P_1, \ldots, P_k$ are other subgroups that contain no infinite order element conjugate into P. There exists a quasiconvex codimension-1 subgroup K and an extension of W_H to a K-wall W_K, so no infinite order element of P_i is conjugate into K for $i \geq 1$, and W_K cuts P in a single component (within the universal cover).*

Proof. We work under the assumption that G is torsion-free. We can apply Lemma 16.7 to assume that G lies in a subgroup G' that is virtually compact special. The subgroup P equals $G \cap P'$ where P' is full and quasiconvex in G'. Moreover, a retraction $P' \to P$ shows that the quasiconvex subgroup H of P' extends to a quasiconvex subgroup H' of P', and the H-wall extends to an H'-wall in P'.

The group G' has the extension property with respect to P', and so the H'-wall extends to a K'-wall in G'. We now recubulate G' using this additional wall to obtain a dual cube complex $X' = G' \backslash \widetilde{X}'$. This additional wall is then associated to a hyperplane stabilized by K', which then intersects the P'-cocompact convex subcomplex $\widetilde{Y}'$ in a hyperplane $\widetilde{V}'$ stabilized by H'.

Let $Y' = P' \backslash \widetilde{Y}'$, and form $\mathsf{C}(Y' \to X')$. As explained in [HWa], the canonical retraction ensures that the canonical copy of Y' that is a retract of $\mathsf{C}(Y' \to X')$ is *wall-injective*, in the sense that each hyperplane of $\mathsf{C}(Y' \to X')$ intersects Y' in at most one hyperplane of Y'. Furthermore, the preimage of a hyperplane U' of Y' is equal to a collection U'_i of hyperplanes of $\mathsf{C}(Y' \to X')$ that map to the hyperplane of X' containing U'. (More correctly, they are parallel within a cubical subdivision of $\mathsf{C}(Y' \to X')$). In particular, the preimage of V' in $\mathsf{C}(Y' \to X')$ is just the cover of $\widehat{W}'$ containing V', where W' is the hyperplane of X' containing U'.

If Y' were already wall-injective in X', then W' would be our desired wall extension. However, this might not be the case. We therefore work with the wall-injective copy of Y' within $\mathsf{C}(Y' \to X')$, which yields W' with the desired properties.

We now apply the malnormal special quotient theorem to crush the other elevations of Y'. The details of the malnormal special quotient theorem proof ensure that in the new cubulation, the hyperplane W' (or rather its lift in a cubulation of a torsion-free finite index subgroup) will only cross the living elevations of Y' in copies of covers of V' - which was our goal.

The codimension-1 subgroup we obtained induces one for P by taking intersections. □

Remark 17.17. Next thing needed to complete the relatively hyperbolic case is to clarify that we can proportionately readjust to make E, E' isomorphic in a cover after cubulating. One option is to first fill the parabolics (proportionately) and then apply the special quotient theorem etc.

18. Application towards one-relator groups

18.1. **Overview.** We now provide an application resolving the conjecture of Gilbert Baumslag on the residual finiteness of one-relator groups with torsion.

In analogy with Haken 3-manifolds, every one-relator group has a hierarchy terminating in a group isomorphic to $F_r * \mathbb{Z}_n$ for some r, n.

We describe the hierarchy in Section 18.2, and prove in Lemma 18.8 that it is a quasiconvex hierarchy when $n \geq 2$. Without the torsion assumption, there are simple examples where the hierarchy is not quasiconvex, as described in Example 18.3.

Theorem 18.1. *The Magnus-Moldavanskii hierarchy is quasiconvex for any one-relator group with torsion.*

Combining Theorem 18.1 with Theorem 13.3, we obtain the following result strongly affirming Baumslag's conjecture:

Corollary 18.2. *Every one-relator group with torsion has a finite index subgroup that is the fundamental group of a compact special cube complex.*

We now describe a word-hyperbolic one-relator group whose Magnus hierarchy is not a quasiconvex hierarchy. I wonder if there is a natural characterization of such examples (probably not).

Example 18.3. Consider the presentation $\langle a, b, c, t \mid abc^{-1}, a^t = b, b^t = c\rangle$. Its group is an HNN extension of the free group $\langle a, b, c \mid abc^{-1}\rangle \cong \langle a, b \mid -\rangle \cong F_2$, and in fact, a semi-direct product $F_2 \rtimes_\phi \mathbb{Z}$ using the automorphism induced by: $\phi(a) = a^t = b$ and $\phi(b) = b^t = c = ab$.

By choosing words with no repeated letters (or even primitive words), that are complicated relative to the index shift isomorphism between Magnus subgroups, we can obtain a small-cancellation group, with similar behavior.

For instance, letting x^y denote yxy^{-1} the HNN extension: $\langle a_1, \ldots, a_{r+1}, t \mid a_1 \ldots a_{r+1}, a_i^t = a_{i+1} : 1 \leq i \leq r\rangle$ is isomorphic to the one-relator group: $\langle a_1, t \mid a_1 a_1^t a_1^{t^2} \ldots a_1^{t^r}\rangle$ which equals $\langle a_1, t \mid a_1 t a_1 t^2 a_1 t^3 \ldots a_1 t^r\rangle$ which is a $C'(\frac{1}{6})$ small-cancellation group for $r > 21$. Indeed, pieces have length at most $2r$ but the word has length $r + \frac{r(r+1)}{2} = \frac{r^2+3r}{2} > 12r$. (In fact the group is $C'(\frac{1}{4}) - T(4)$ when $r > 13$.)

Since the HNN extension is a semi-direct product, the vertex group and hence the edge groups are normal of infinite index and thus cannot be quasiconvex. It may be that some alternate one-relator presentation provides a quasiconvex malnormal hierarchy, but I reckon there are examples that have no quasiconvex hierarchy. Nevertheless, I expect the following holds:

Conjecture 18.4. *Every word-hyperbolic one-relator group has a finite index subgroup with a quasiconvex hierarchy.*

18.2. **The Magnus-Moldavanskii Hierarchy.**

Construction 18.5 (Magnus-Moldavanskii Construction)**.** Following [LS77], we shall now describe a variation on Moldavanskii's variant of Magnus's inductive construction of one-relator groups. I am grateful to Hadi Bigdely and Eduardo Martinez-Pedroza for editing this section.

We start with a one-relator group $\langle S \mid W^n \rangle$ where the generating set is S, and W is cyclically reduced word in the generators that is not equal to a proper power. We define the *repetition complexity* of W to be $|W|$ minus the number of distinct letters that occur in W. Equivalently, this is the sum of the number of times that letters occurring in W are repeated, so the complexity of $aba^{-1}bbc$ would be $1+2+0 = 3$. Having a given presentation in mind, we refer to the *complexity* of a one-relator group as the repetition complexity of the relator in this presentation.

Note that when the complexity is zero, and more generally, when some generator appears exactly once in W, then the group presented by $\langle S \mid W^n \rangle$ is isomorphic to a virtually free group $F * \mathbb{Z}_n$ where the rank of F equals $|S| - 1$.

Assuming that no generator appears exactly once, we show that $G' = G * \mathbb{Z}$ splits as an HNN extension $K*_M$ of a one-relator group K over a Magnus subgroup. The complexity of K is less than the complexity of G, and so this process stops after finitely many steps. It is clear that this induces an actual hierarchy (that terminates at finite cyclic groups), since a splitting of $G*\mathbb{Z}$ induces a splitting of G, and if it is a quasiconvex splitting, then it induces a quasiconvex splitting. The free factor with $\mathbb{Z}$ is an artifice that facilitates the combinatorial group theory description. The stable letter t of the HNN extension conjugates one Magnus subgroup to another by an "index shift". A *Magnus subgroup* of a one-relator group $\langle S \mid W^n \rangle$ is a subgroup generated by a subset of the generators omitting some generator appearing in the relator W^n. Magnus's "Freiheitsatz" theorem states that these specific generators freely-generate the Magnus subgroup of a one-relator group. We state a general formulation of the Freiheitsatz in Proposition 18.10.

Let us now describe this process a bit more carefully. We add a new further generator t to the presentation, so that the resulting finitely presented group G' is isomorphic to $G * \langle t \rangle$. We let $\bar{S}$ denote a new set of generators in one-to-one correspondence with the generators in S, so $s \leftrightarrow \bar{s}$. We will choose an integer p_s for each $s \in S$, and perform a substitution $s \mapsto \bar{s}t^{p_s}$, that rewrites the relator W as a word $\bar{W}'$ in terms of $\bar{S} \cup \{t\}$, so we have a new presentation $\langle t, \bar{s} \in \bar{S} \mid \bar{W}' \rangle$ for G'. There are Tietze transformations justifying this: we first add generators $\bar{s}$ with relators $s^{-1}\bar{s}t^{p_s}$, and then rewrite the relator by substituting $\bar{s}t^{p_s}$ for each s in W, and finally, we remove the old generators s and relators $s^{-1}\bar{s}t^{p_s}$.

For each generator x of a free group F, there is a homomorphism $\#_x : F \to \mathbb{Z}$ induced by sending x to 1 and all other generators to 0. For a word V, the value $\#_x(V)$ is the exponent sum of the letter x in the word V,

We shall now assume that the integers p_s are chosen so that the resulting word satisfies $\#_t(\bar{W}) = 0$. Indeed, assuming that W contains at least two letters, there is always a way of doing this: If $\#_a(W) = 0$ for some generator $a \in S$, then we let $p_a = 1$ and let $p_s = 0$ for all $s \neq a$. Otherwise, we can choose $a, b \in S$ with $\#_a(W) \neq 0$ and $\#_b(W) \neq 0$, and we then define $p_a = \#_b(W)$ and define $p_b = -\#_a(W)$, and define $p_s = 0$ for all $s \neq a, b$.

In this way, our word W in S becomes a word $\bar{W}'$ in $\bar{S} \cup \{t\}$. For each $\bar{s} \in \bar{S}$ let $L_{\bar{s}} \leq R_{\bar{s}}$ be the smallest and greatest values of $\#_t(U)$ where U is a prefix of $\bar{W}'$ preceding an occurrence of $\bar{s}^{\pm 1}$ so $U\bar{s}^{\pm 1}V = \bar{W}'$. In the degenerate case where $\bar{W}'$ contains no occurrence of $\bar{s}^{\pm 1}$ then we assign $L_{\bar{s}}, R_{\bar{s}} = 0$.

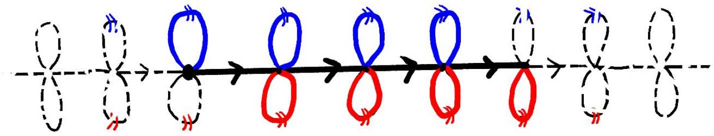

FIGURE 170. The subcomplex $Y^1 \subset \widehat{X}$ for the word $\bar{W}' = at^3bt^{-2}at^3bt^{-4}$ is illustrated in bold. The reader should check that the base lift of $\bar{W}'$ at the basepoint only passes through two b edges and two a edges, but we include their intermediate $\mathbb{Z}$-translates.

We introduce *new generators* $\bar{s}_i : L_{\bar{s}} \leq i \leq R_{\bar{s}}$, and *new relators* $\bar{s}_i = t^i \bar{s} t^{-i} \ : \ L_{\bar{s}} \leq i < R_{\bar{s}}$. Adding these generators and relators, we thus obtain a new presentation for G'. Finally, since $\#_t(\bar{W}') = 0$, we see that $\bar{W}'$ is freely equivalent to a word $\bar{W}''$ in $(t^i \bar{s} t^{-i}) : \bar{s} \in \bar{S}$, and we use the introduced relators to rewrite $\bar{W}''$ as a word $\bar{W}$ in our new generators $\bar{s}_i$.

Exchanging the relations $\{\bar{s}_i = t^i \bar{s} t^{-i} : L_{\bar{s}} \leq i < R_{\bar{s}}\}$ for relations of the form $\bar{s}_i^t = \bar{s}_{i+1}$ (here we use the notation $x^y = y^{-1}xy$), the resulting presentation for G' is the following:

$$\langle t, \bar{s}_i \ : \ L_{\bar{s}} \leq i \leq R_{\bar{s}},\ \bar{s} \in \bar{S} \ \mid \ \bar{W}^n, \bar{s}_{i+1}^t = \bar{s}_i \ : \ L_{\bar{s}} \leq i \leq R_{\bar{s}},\ \bar{s} \in \bar{S}\,\rangle$$

Thus G' is an HNN extension of the one-relator group K presented by $\langle \bar{S} \mid \bar{W}^n \rangle$ with stable letter t, and Magnus subgroups $M_+^t = M_-$ where: $M_- = \langle \bar{s}_i : L_{\bar{s}} \leq i < R_{\bar{s}} \rangle$, and $M_+ = \langle \bar{s}_i : L_{\bar{s}} < i \leq R_{\bar{s}} \rangle$. We note that when G and hence G' is f.g. then so is K.

Observe that $|W| = |\bar{W}|$ and the number of generators occurring in $\bar{W}$ exceeds the number occurring in W by $\sum_{\bar{s} \in \bar{S}} R_{\bar{s}} - L_{\bar{s}}$. Recall that the complexity is the difference between the relator length and the number of occurring generators, and hence the complexity decreases by $\sum_{\bar{s} \in \bar{S}} R_{\bar{s}} - L_{\bar{s}}$. If t appears at least once in $\bar{W}'$ we can hope that $R_{\bar{s}} > L_{\bar{s}}$ for some $\bar{s}$, and hence the complexity will decrease. To this end we must choose our p_s integers a bit more carefully, and will explain how to do so below.

Recapitulation and Geometric Interpretation: We began with a one relator group G with presentation $\langle S \mid W \rangle$. We define $G' = G * \langle t \rangle$, which obviously has the presentation $\langle S, t \mid W \rangle$. Using different generators $\bar{s} = st^{-p_s}$, and rewriting W in terms of these new generators, we obtain a new presentation for G' of the form $\langle \bar{S}, t \mid \bar{W}' \rangle$. The word $\bar{W}'$ has occurrences of $\bar{s}$ in one-to-one correspondence with occurrences of s in W, but there are additional occurrences of t. We then express G' as an HNN extension whose base group K is a new one-relator group. Let X denote the standard 2-complex of the presentation $\langle \bar{S}, t \mid \bar{W}' \rangle$. Let $\widehat{X} \to X$ be the $\mathbb{Z}$-covering space induced by $t \mapsto 1, \bar{s} \mapsto 0$, and note that this extends the homomorphism $\phi : G \to \mathbb{Z}$ induced by $s \mapsto p_s$.

Consider the based lift of the closed path $\bar{W}'$ to $\widehat{X}$. Let Y denote the subcomplex of $\widehat{X}$ that contains the closure of the based lift of the 2-cell of X, and also contains all $\mathbb{Z}$-translates of a 1-cell $\bar{s}$ that lie between the first and last such $\bar{s}$ 1-cell that $\bar{W}'$ passes through. We refer the reader to Figure 170.

The group $K = \pi_1 Y \subset \pi_1 \widehat{X}$. To obtain a presentation for K we contract the t-tree within Y, and our $\bar{s}_i$ generators correspond to the loops attached to this t-tree. The new relator $\bar{W}$ is closely related to our original word W, but with some of its generators subscripted - the temporary t-letters have now disappeared. More precisely, regarding a t-path as a maximal tree in Y^1, there is a natural generator system for $\pi_1 Y^1$ consisting of conjugates of the original generators of $\pi_1 X^1$. Namely: $\{\bar{s}_i = t^{-i} \bar{s} t^i : L_{\bar{s}} \leq i \leq R_{\bar{s}}\}$. We also include a single loop $\bar{s}_0$ for each generator of $\bar{S}$ that doesn't appear in $\bar{W}'$.

Finally, the complexity of the presentation for K decreases precisely if $\bar{W}$ contains more distinct generators than occurred in W. This happens precisely if the closed path $\bar{W}'$ passes through two non-t one-cells of $\widehat{X}$ that map to the same 1-cell of X. This occurs precisely if $\phi : G \to \mathbb{Z}$ "separates two occurrences of a generator" which we will define and examine in the explanation below of how to make the choices so that the complexity decreases.

Complexity Reduction: We will now show that for suitable choices of the integers p_s above, the new one relator group K has lower complexity than the original one-relator group G.

A homomorphism $\phi : F \to \mathbb{Z}$ with $\phi(W) = 0$ corresponds to a choice of integers $p_s = \phi(s)$. Consider two occurrences of a generator s in W corresponding to a subword of the form $s^{\pm 1}Us^{\pm 1}$ in W. We say that ϕ *separates* these two occurrences provided that either: $\phi(sU) \neq 0$ in case of sUs; or $\phi(s^{-1}U) \neq 0$ in case of $s^{-1}Us^{-1}$; or $\phi(U) \neq 0$ in case of sUs^{-1} or $s^{-1}Us$. We refer the reader to the later geometric interpretation below. When ϕ separates occurrences of s, we have $R_{\bar{s}} > L_{\bar{s}}$ and hence this choice yields a complexity reduction.

Suppose there is a generator a with $\#_a(W) = 0$. If some pair of successive appearances of $a^{\pm 1}$ in W have the same sign, then they are separated by $\#_a$. So let us consider the alternating case where $W = aA_1a^{-1}A_2aA_3\cdots$ and each A_i is nonempty and contains no $a^{\pm 1}$. If some generator b occurs in both A_i, A_j with $i \not\equiv_2 j$ then the $\#_a$ homomorphism separates these two occurrences of b. So, let us assume that no generator appears in both an odd and an even syllable.

We call b a *zero letter* of W if $\#_b(W) = 0$ and call b a *nonzero letter* of W if $\#_b(W) \neq 0$. Considering $aA_1a^{-1}A_2aA_3\cdots$, we see that either each A_i contains at least one nonzero letter b, or some A_i consists entirely of zero letters.

Suppose some A_i consists entirely of zero letters, and consider an innermost pair $b^{\pm 1}Cb^{\pm 1} \subset A_i$ with the property that $b \notin C$. We can assume that C is nonempty for in the case bb or $b^{-1}b^{-1}$ the $\#_b$ homomorphism separates the occurrences of b. Now for any $c \in C$, we have $\#_c(C) \neq 0$ so $\#_c$ separates the surrounding occurrences of $b^{\pm 1}$.

Let us therefore consider the case where each A_i contains a nonzero letter. There is thus an "even" nonzero letter b in A_2 and an "odd" nonzero letter c in A_1, and we use the $\phi = w_c\#_b - w_b\#_c$ homomorphism, where we use the notation $w_x = \#_x(W)$. Let A_i be a syllable with $\#_b(A_i) \neq 0$, and note that $c \notin A_i$ since i is even, so ϕ separates the two occurrences of a around A_i.

We now assume that $\#_a(W) \neq 0$ for each letter a occurring in W. Suppose some letter occurs with both signs. Then we can choose $a^\epsilon Ba^{-\epsilon}$ such that $a^{\pm 1}$ does not occur in B, and such that $|B|$ is minimal among all such choices with a, B allowed to vary. Observe that $\#_b(B) \neq 0$ for any letter b occurring in B, for then b^{+1}, b^{-1} both occur in B and this would violate the minimality of $|B|$. We then use the homomorphism $w_a\#_b - w_b\#_a$ which separates the two occurrences of a around B.

Without loss of generality, we may assume that each occurring letter occurs only positively. Express W as $aA_1aA_2aA_n$ where each A_i is nonempty and has no occurrence of a. If $\#_b(A_i) = 0$ but b occurs in W, then $w_a\#_b - w_b\#_a$ separates the a's around aA_ia. If $\#_b(A_i) \geq 2$, but b occurs in W, then $w_a\#_b - w_b\#_a$ separates the occurrences of b in A_i. We thus have $\#_b(A_i) = 1$ for each b occurring in W.

We may thus assume that each A_i contains each letter occurring in W, and moreover it occurs exactly once. It follows that $|A_i|$ is constant, so each occurrence of a within W occurs in the same position modulo the total number of letters occurring in W. The same statement applies for each letter occurring in W. Consequently, $W = U^n$ where each letter occurs exactly once in U.

18.3. **Quasiconvexity using the strengthened spelling theorem.** Let X be a staggered 2-complex. Such 2-complexes are discussed in Section 18.4, but in particular, we have in mind the motivating case where Let X be the standard 2-complex of $\langle a, b, \dots \mid W^n \rangle$. In this case, a 2-cell in a disk diagram $D \to X$ is *extreme* if more than $\frac{n-1}{n}$ of its boundary path is a subpath of the boundary path $\partial_p D$ of D. Such a subpath is the *outerpath* of the extreme 2-cell. More generally, for a staggered 2-complex X we say r is an *extreme 2-cell* if this outerpath has length exceeding $\frac{n-1}{n}|\partial_p r|$ where n is the *exponent* of r.

The Newman spelling theorem (see [LS77] and the references therein and [HW01] for the staggered case) states that a reduced diagram $D \to X$ that is nontrivial, spurless, and not a single 2-cell, must contain at least two extreme 2-cells. The following amplification of this is proven in [LW]:

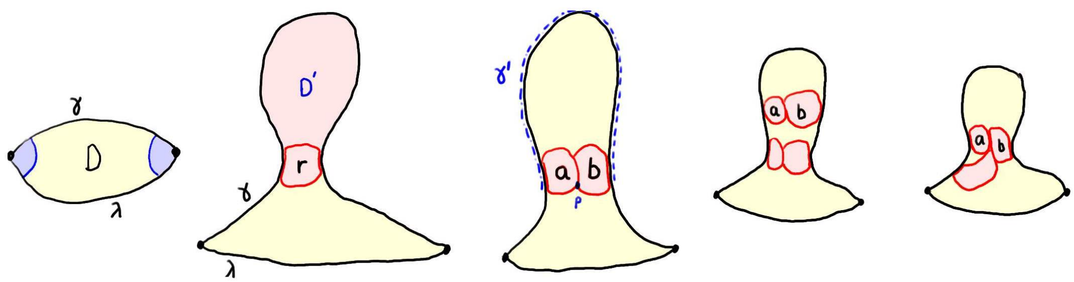

FIGURE 171. Pac-Man: The first figure shows the only two places where an extreme cell could lie, the second figure shows the subdiagram D' subtended by r, the third figure illustrates a bridge, and the fourth and fifth figures illustrate bridges that are not outermost.

Proposition 18.6. *Let X be a staggered 2-complex. Let $D \to X$ be a reduced spurless disk diagram with an exponent n internal 2-cell r (meaning that $\partial r \cap \partial D \subset D^0$). Then D contains at least $2n$ extreme 2-cells.*

Proposition 18.6 is most interesting when each 2-cell of X has exponent ≥ 2, as arises in the motivating case of a one-relator group with torsion.

Lemma 18.7 (Fellow Traveling)**.** *Let X be a staggered 2-complex with torsion. whose 2-cell attaching maps have the form $W_i^{n_i}$ with $n_i \geq 2$. Let $M = \max\{n_i|W_i|\}$ and let $\kappa = \frac{3}{2}M$.*

Let $D \to X$ be a reduced diagram between λ and γ. Suppose that γ lifts to a geodesic in the 1-skeleton of $\widetilde{X}$. Suppose λ and γ are immersed combinatorial paths, and that neither contains a subpath that is the outerpath of an extreme 2-cell.

Then $\gamma \subset N_\kappa(\lambda)$.

Proof. Without loss of generality, we can assume that D is spurless. Indeed, the only possible spurs occur at an initial or terminal agreement between λ and γ, and removal of these does not effect the claim. By Proposition 18.6, D cannot contain an exponent n internal 2-cell, because then D would have $2n$ extreme 2-cells, but D can accommodate at most two such 2-cells - one at each end, by hypothesis on λ, γ.

Claim: For each open 2-cell r in D, the intersection $\partial r \cap \gamma$ is connected.

Indeed, otherwise r would subtend a subdiagram D' with $D' \neq \bar{r}$, so by the spelling theorem, D' has at least two extreme 2-cells which is impossible by hypothesis on γ.

A *bridge* in D is a pair of 2-cells $a \neq b$ such that $\bar{a} \cap \bar{b} \neq \emptyset$ and $\partial a, \partial b$ both intersect γ. See Figure 171. The *base vertex* p of the bridge is the endpoint of $\bar{a} \cap \bar{b}$ on the λ-side of the bridge. It is conceivable that $\bar{a} \cap \bar{b}$ is disconnected, in which case we choose the lowest such point.

Observe that the subpath $\gamma' \subset \gamma$ subtended by the bridge satisfies $|\gamma'| \leq M$ and so $\gamma' \subset N_M(p)$. Indeed, γ is a geodesic, and thus so is γ'. The endpoints of γ' are joined by a pair of subpaths of $\partial a, \partial b$, which are of length $\leq \frac{1}{2}M$. Thus $|\gamma'|$ is bounded by the sum of these lengths.

A bridge is *outermost* if its subtended path γ' is not properly contained in the path subtended by another bridge.

Let e be an edge on γ and suppose that e is not on λ. Then e lies on some 2-cell r of D. Either:

(1) $\partial r \cap \lambda$ contains a 1-cell,
(2) some 2-cell r' is incident with both r and λ [This becomes case (1) when $r' = r$]
(3) no 2-cell is incident with both r and λ.

We consider case (3). Since there are no internal 2-cells in D, each 2-cell around r must have a 1-cell on γ. Traveling around r, we denote the 2-cells $r_1, r_2, \ldots, r_t$, where $\partial r_1 \cap \gamma$ precedes $\partial r_2 \cap \gamma$ in the orientation on γ as in the first diagram in Figure 172.

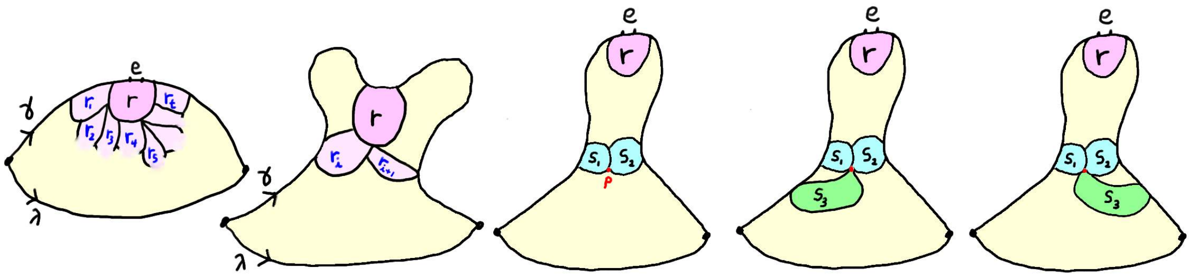

FIGURE 172.

Since r_1 intersects γ before r, and r_t intersects γ after r, there must be some i such that r_i intersects γ before r and r_{i+1} intersects γ after r. We thus find that r and hence e lies in a subdiagram subtended by a bridge as in the second diagram in Figure 172.

As in the third diagram in Figure 172, let s_1, s_2 be 2-cells of a outermost bridge subtending a subdiagram containing r, and let p be the basepoint of this bridge. Either:

(1) $p \in \lambda$,
(2) p lies on a 2-cell containing a 1-cell of λ,
(3) p lies on a 2-cell s_3 distinct from s_1, s_2, such that s_3 contains a 1-cell of γ.

Note that Case (3) is impossible because then, as in the fourth and fifth diagrams of Figure 172, either s_3, s_2 or s_1, s_3 would form an even larger bridge contradicting that s_1, s_2 is maximally outermost.

We are thus left with cases (1) and (2).

We conclude that e lies in $N_M(p)$ and p lies in $N_{\frac{1}{2}M}(\lambda)$, so e lies in $N_\kappa(\lambda)$. Thus γ lies in $N_\kappa(\lambda)$ as claimed. □

Lemma 18.8. *K is quasiconvex in G [or $G * \mathbb{Z}$].*

Proof. Suppose that $G = \langle t, a_1, a_2, \cdots \mid W^n \rangle$ such that $\#_t(W) = 0$ where $\#_x(Y)$ denotes the exponent sum of the letter x in the word W. Let X be the standard 2-complex of the presentation for G, and let $\widehat{X} \to X$ be the $\mathbb{Z}$-cover corresponding to the quotient induced by $a_i \mapsto 0$ and $t \mapsto 1$. Let Y denote the subcomplex of $\widehat{X}$ that equals the closure of the based lift of the 2-cell of X. By the Moldavanskii-Magnus construction, $\pi_1 Y \cong K$.

Note that $Y \subset \widehat{X}$ is π_1-injective. When $n \geq 2$ this follows from the Newman spelling theorem [LS77, HW01]. Indeed, any other lift of W is shifted over from the based lift. Consequently $Y \to X$ has no missing extreme cell bounded by a path W^n, since W^n would travel $2n$-times through a lift of a translate of a t-edge that is not in Y.

Let $\widetilde{Y} \subset \widetilde{X}$ be a component of the preimage of Y in $\widetilde{X}$. Note that $\widetilde{Y}$ also has no missing W^{n-1} paths. Let γ be a geodesic whose endpoints lie on $\widetilde{Y}$. Let D be a minimal area diagram between γ and a path $\lambda \to \widetilde{Y}$ with the same endpoints, such that $\mathsf{Area}(D)$ is minimal among all possible choices of λ. Observe that D has no extreme cells whose outerpaths lie in λ or γ, for then we could find a smaller area diagram in the first case, and contradict that γ is a geodesic in the second. Consequently, Lemma 18.7 shows that $\gamma \subset N_\kappa(\lambda) \subset N_\kappa(\widetilde{Y})$.

We note that the Magnus subgroup corresponding to the edge group of the HNN extension is the intersection of two conjugates of K and is hence itself quasiconvex. (But the proof we gave can be applied to it directly.) □

18.4. **Staggered 2-complex with torsion.** A 2-complex X is *staggered* if there is a linear order on a subset of the 1-cells, and a linear ordering on the 2-cells, such that each 2-cell has an ordered 1-cell in its attaching map and $\max(a) < \max(b)$ and $\min(a) < \min(b)$ whenever $a < b$ are 2-cells.

Here we let max(a) denote the greatest 1-cell in ∂a and let min(a) denote the least. The spelling theorem [HP84, HW01] shows that a finite staggered 2-complex has word-hyperbolic fundamental group provided that the attaching map of each 2-cell is a proper power w^n for some path $w \to X^1$ and some nontrivial closed immersed path w, and some $n \geq 2$. We refer to such a 2-complex as a *staggered 2-complex with torsion*.

Finally, for each finite staggered 2-complex X, either X has a single 2-cell, in which case $\pi_1 X$ is virtually special by Corollary 18.2, or else $\pi_1 X$ splits as an amalgamated product of staggered 2-complexes with fewer 2-cells, where, as we shall explain, the amalgamated subgroup is quasiconvex and malnormal. It follows that $\pi_1 X$ is virtually compact special by Theorem 11.2. We thus obtain the following result:

Theorem 18.9. *Every staggered 2-complex with torsion has virtually special π_1.*

By initially adding some additional unordered 1-cells, we can assume that all 0-cells are connected by a path of unordered 1-cells. This has the effect of adding a free factor to the fundamental group and doesn't effect the existence of a quasiconvex hierarchy, but facilitates its description by maintaining connectedness. The splitting now arises geometrically as follows: One factor is π_1 of the 2-complex consisting of the closure of the top 2-cell together with the unordered 1-skeleton, and the second factor is π_1 of the 2-complex consisting of the closure of all other 2-cells together with the unordered 1-skeleton.

Let X be a staggered 2-complex. A *Magnus subcomplex* $Z \subset X$ is a connected subcomplex with the following properties:

(1) if C is a 2-cell of X, and all the ordered boundary 1-cells of C lie in Z, then $C \subset Z$.
(2) the ordered 1-cells of X contained in Z form an interval.

We note that the intersection of Magnus subcomplexes is a Magnus subcomplex. The following generalization of the Freiheitsatz was shown in [HW01, Thm 6.1] (though the result is incorrectly stated there) for Magnus subcomplexes as defined above:

Proposition 18.10. *If Z is a Magnus subcomplex of a staggered 2-complex, then $Z \to X$ is π_1-injective.*

We augment Proposition 18.10 with the following two statements, the second of which generalizes Newman's result that Magnus subgroups of one-relator groups with torsion are malnormal.

Lemma 18.11. *Let $Z \subset X$ be a Magnus subcomplex of a staggered 2-complex with torsion. Then*

(1) *$\widetilde{Z} \to \widetilde{X}$ is convex.*
(2) *$\pi_1 Z$ is a malnormal subgroup of $\pi_1 X$.*

The following proof presumes familiarity with the mode of proof in [HW01] rather than its main theorem.

Proof of convexity. Let γ be a geodesic in $\widetilde{X}$ such that $\gamma \cap \widetilde{Z}$ consists precisely of the endpoints of γ Let $D \to \widetilde{X}$ be a disk diagram between γ and a path $\lambda \to \widetilde{Z}$, and assume that D is of minimal area among all such (D, λ) with γ fixed, so in particular D is reduced. We can assume that D does not map entirely into $\widetilde{Z}$, so can assume that some 2-cell of D maps to a 2-cell above or below the Magnus complex Z - and shall assume that it maps above, without loss of generality.

We choose a maximal cyclic tower lift $\phi : D \to T$, and then consider the greatest 2-cell r in T and note that by Howie's Lemma [How87], $\partial_p r$ is of the form $(Qe)^n$ where Q does not traverse e, and n is the exponent of r, and no other 2-cell is incident with e. The 1-cells in $\phi^{-1}(e)$ are all in ∂D, for otherwise there would be a cancellable pair. Moreover, since r is above Z, we see that $\phi^{-1}(e)$ lies entirely in γ. However, if $n \geq 2$, we see that consideration of a single 2-cell in $\phi^{-1}(r)$ shows that γ is not a geodesic, which is a contradiction. □

Proof of malnormality. Consider a reduced annular diagram $A \to X$ where the two boundary paths of A map to Z. We will show that any such A maps to Z, and since Z is a connected subcomplex of X, this shows that $\pi_1 Z$ is malnormal.

Choose a shortest path $\gamma \to A$ whose endpoints lie on the disjoint circles in ∂A. Choose $d \in \mathbb{Z} = \mathrm{Aut}(\widetilde{A} \to A)$ sufficiently large that the translates $0\widetilde{\gamma}$ and $d\widetilde{\gamma}$ in $\widetilde{A}$ do not both intersect a common 2-cell.

Cutting $\widetilde{A}$ along $0\widetilde{\gamma}$ and $d\widetilde{\gamma}$, we obtain a disk diagram $D \to X$ with $\partial_p D = \lambda_1 \gamma \lambda_2^{-1} \gamma^{-1}$ where λ_1, λ_2 each travel d times around the distinct boundary cycles of A.

As in the proof of convexity, A and hence D must contain a 2-cell that does not map to Z. Consideration of a maximal tower lift, shows that the preimage of a highest 2-cell in a maximal tower lift, has all of its highest 1-cells on ∂D (since D is reduced), and hence on γ, γ^{-1} since λ_1, λ_2 map to Z. Moreover, our choice of d ensures that the 2-cell has all its highest 1-cells on either γ or on γ^{-1}. When the exponent $n \geq 2$, we can thus travel along the inner path of this 2-cell to find a shorter path in $\widetilde{A}$, and hence in A between the two boundary paths.

The base case, where A is singular, and $|\gamma| = 0$ yields a contradiction, since there is no room at all for a highest 2-cell in D, so we see that there was a cancellable pair of 2-cells. □

19. Problems

Generalize this to relatively hyperbolic CAT(0) cube complexes. Are they always virtually special? Do they have pseudograph CAT(0) quotients? Perhaps there is an argument by induction on dimension and/or depth. An important test case are the negatively curved square complexes.

Generalize the virtual special quotient theorem so that there exists finite index subgroups H_i' such that $G/\langle\langle H_i'' \rangle\rangle$ is virtually special for any finite index subgroups $H_i'' \subset H_i'$. This appears to work in the cyclic case.

Let H be a codimension-1 subgroup of G. Does there exist a finite index subgroup $H' \subset H$ such that $H/\langle\langle H' \rangle\rangle$ is also codimension-1 in $G/\langle\langle H' \rangle\rangle$?

The following should be proven following the scheme of the proof of Theorem 12.1, by repeatedly using a relatively malnormal special quotient theorem in the sparse case. Alternatively, the graded small-cancellation approach should work (assuming G initially has this structure).

Conjecture 19.1 (Relatively Hyperbolic Virtually Special Quotient Theorem)**.** *Let G be hyperbolic relative to virtually abelian subgroups. And suppose G is virtually sparse special. Let $H_1, \ldots, H_k$ be quasiconvex subgroups. And let $H_i^o \subset H_i$ be finite index subgroups. There exist finite index subgroups $H_i' \subset H_i$ such that $G/\langle\langle H_1', \ldots, H_k' \rangle\rangle$ is virtually sparse special.*

Problem 19.2. Suppose G is hyperbolic relative to virtually abelian subgroups. Suppose G is virtually sparse special. Is G virtually compact special?

Finding general conditions under which the following problem has an affirmative answer will yield many applications.

Problem 19.3 (Virtual Haken problem for cube complexes)**.** Let X be a compact nonpositively curved cube complex. Does there exist a finite cover $\widehat{X}$ with the property that the immersed hyperplanes of $\widehat{X}$ are actually embedded?

The following example shows that there exists a nonpositively curved cube complex with no finite cover whose hyperplanes embed.

Example 19.4. Let Y denote a compact nonpositively curved cube complex with no finite cover but with $\pi_1 Y$ nontrivial. We can choose Y so that link(y) is a complete bipartite graph for each $y \in Y^0$ (see [BM97, Wis07]). We can choose a closed path $\sigma \to Y$ such that the lift $\widetilde{\sigma} \to \widetilde{Y}$ is an isometry

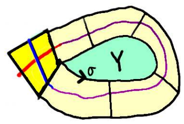

FIGURE 173.

(the path σ travels through ends of 1-cells in the same class of the bipartite structure in each link). Let $A = [-1,1] \times [0,n]$ be a flat strip of length $n = |\sigma|$, and glue A to Y along σ by identifying σ with $\{-1\} \times [0,n]$. Finally, add an extra square S at the basepoint along the two 1-cells $[-1,1] \times \{0\}$ and $[-1,1] \times \{n\}$ corresponding to the first and last 1-cells of A. Let X be the resulting complex which is heuristically illustrated in Figure 173. Since X deformation retracts to Y, it has no finite covers. Since there are no corners of squares along the interior of σ we see that X is a nonpositively curved square complex. However, the hyperplane of X containing $\{0\} \times [0,n]$ self-crosses within the square S, and hence within X.

Conjecture 19.5. *Let G be a word-hyperbolic group acting properly and cocompactly on a CAT(0) cube complex C. Then G has a finite index subgroup F acting specially on C.*

Problem 19.6 (Codimension-1 subgroups in $C(6)$ groups)**.** Is there an example of a $C(6)$ presentation with infinite fundamental group but with no codimension-1 subgroup? Is there such an example with Property-(T)? Can such examples be identified using a generalized or appropriately aimed version of Zuk's spectral gap criterion [Żuk96].

20. ⊛ ARTIN GROUPS

Let $(x,y)^m$ denote the initial half of the word $(xy)^m$. A f.g. *Artin* group is presented by $\langle a_1, a_2, \ldots \mid (a_i, a_j)^{m_{ij}} = (a_j, a_i)^{m_{ji}} : i \neq j \rangle$ where M is a symmetric square matrix whose entries are natural numbers that are ≥ 2. We allow $m_{ij} = \infty = m_{ji}$, in which case there is no relation between a_i, a_j among the relators

For each $i < j$, let Y_{ij} denote the 1-skeleton of the universal cover of the standard 2-complex of $\langle a_i, a_j \mid (a_i, a_j)^{m_{ij}} = (a_i, a_j)^{m_{ij}} \rangle$. When $m_{ij} = \infty$, we omit Y_{ij}.

Let $A = A(M)$ be an Artin group as above. Then A has the following cubical presentation satisfying the $C(6)$ condition:

$$\langle a_1, a_2, \ldots \mid Y_{ij} : i < j \text{ and } m_{ij} < \infty \rangle \tag{3}$$

Theorem 20.1. *Suppose that $3 \leq m_{ij} \leq \infty$ for each i, j. Then the cubical presentation in Equation (3) satisfies the $C(6)$ property.*

More generally, suppose that there is no triangle of form $(2,3,3)$, $(2,3,4)$, $(2,3,5)$, or $(2,2,n)$, in the graph associated to the Artin presentation. Then the cubical presentation satisfies the $C(6)$ conditions.

Proof. Observe that pieces between distinct translates of Y_{ij} and $Y_{k\ell}$ in $\widetilde{X}^*$ correspond to lines a_p^∞ where one of i, j equals p, and one of k, ℓ equals p.

In the first case, at least 6 pieces are needed to form an essential cycle in Y_{ij}.

In the more general case, use the angle grading. □

Remark 20.2. Jon McCammond reported on a similar idea at a talk he gave in Albany around 2000. He grouped the relations of the same type in an Artin group within a disk diagram to obtain small-cancellation behavior under "large enough type" conditions. Perhaps the earliest occurrences of this idea is in the work of Appell and Schupp [AS83].

Theorem 20.3. *Suppose that no exponent satisfies $3 \leq m_{ij} \leq 5$ then we use the above construction starting with X denoting the natural cube complex of the underlying right-angled Artin group, and we should obtain a $C(6)$ cubical generalized presentation.*

Remark 20.4 (Not $B(6)$)**.** There are several natural wallspace structures on Y_{ij}.

The first arises from the induced map to the associated Coxeter group. Namely, two edges are dual to the same wall if and only if they lie in the same equivalence class generated by: two edges are equivalent if they are opposite edges in the same $2m_{ij}$-gon. However, this wallspace is not Hausdorff, and in fact, corresponds to an action of the 2-generator Artin group on a cubulated copy of $\mathbb{E}^3$.

A second structure arises from the map to $\mathbb{Z}$ induced by sending each generator $a_k \mapsto 1$. In this case, the walls don't even cross each other, and the cubulation gives an action on $\mathbb{E}^1$.

In particular, we note that the above cubical presentations are not $B(6)$.

To attempt to achieve the $B(6)$ condition it will be necessary to replace the bouquet of circles corresponding to the free group, by a cube complex X (it will probably be better if X is a compact pseudograph or at least has a sufficient family of finite hyperplanes). For instance, we can take a finite CAT(0) cube complex, like a high-dimensional cube, and attach edges joining opposite vertices, for the generators. Another attractive option is an r-dimensional handlebody corresponding to an r-cube for the 0-cell, and an extra r-cube connecting a pair of opposite faces for each generator. There are a host of different possibilities here. However, I don't think this can work for the walls corresponding to the natural $\mathbb{Z}$ quotient.

We then form a local-isometry $Y_{ij} \to X$. There will now be extra walls in the wallspace of Y_{ij} which make it Hausdorff, and have appropriate separation conditions. The $C(6)$ small-cancellation conditions also persist.

To handle general Artin groups, we will probably have to add one such object for each finite-type Artin group. Possibly this should start with the underlying right-angled Artin group, and possibly they should be added one at a time together.

Remark 20.5. Can this theory be applied to every Artin group? Note that cones make pieces from 2-cells related to themselves not so relevant. A key point will be the superconvexity: For right angled Artin groups, each relator: $ababa = babab$ will be superconvex only after adding infinite a and b "strips" corresponding to their centralizers. A combinatorial theory will be necessary.

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.

[AR90] I. R. Aitchison and J. H. Rubinstein. An introduction to polyhedral metrics of nonpositive curvature on 3-manifolds. In *Geometry of low-dimensional manifolds, 2 (Durham, 1989)*, pages 127–161. Cambridge Univ. Press, Cambridge, 1990.

[Arz98] G. N. Arzhantseva. Generic properties of finitely presented groups and Howson's theorem. *Comm. Algebra*, 26(11):3783–3792, 1998.

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

[AS83] K. I. Appel and P. E. Schupp. Artin groups and infinite Coxeter groups. *Invent. Math.*, 72(2):201–220, 1983.

[BF92] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

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

[BL77] D. C. Brewster and J. C. Lennox. Maximal torsion free subgroups of polycyclic by finite groups. *Arch. Math. (Basel)*, 29(1):39–40, 1977.

[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.

[BW] Hadi Bigdely and Daniel T. Wise. Quasiconvexity and relatively hyperbolic groups that split. 16 pp. Preprint 2011.

[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.

[CD95] Ruth Charney and Michael W. Davis. Finite $K(\pi, 1)$s for Artin groups. In *Prospects in topology (Princeton, NJ, 1994)*, volume 138 of *Ann. of Math. Stud.*, pages 110–124. Princeton Univ. Press, Princeton, NJ, 1995.

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

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

[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.

[DLS91] Carl Droms, Jacques Lewin, and Herman Servatius. Tree groups and the 4-string pure braid group. *J. Pure Appl. Algebra*, 70(3):251–261, 1991.

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

[DS05] Cornelia Druţu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis Osin and Sapir.

[Ger87] S. M. Gersten. Reducible diagrams and equations over groups. In *Essays in group theory*, pages 15–73. Springer, New York-Berlin, 1987.

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

[GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.

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

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

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

[How87] James Howie. How to generalize one-relator group theory. In S. M. Gersten and John R. Stallings, editors, *Combinatorial group theory and topology*, pages 53–78, Princeton, N.J., 1987. Princeton Univ. Press.

[HP84] J. Howie and S. J. Pride. A spelling theorem for staggered generalized 2-complexes, with applications. *Invent. Math.*, 76(1):55–74, 1984.

[HP98] Frédéric Haglund and Frédéric Paulin. Simplicité de groupes d'automorphismes d'espaces à courbure négative. In *The Epstein birthday schrift*, pages 181–248 (electronic). Geom. Topol., Coventry, 1998.

[Hru10] G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[HWa] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. Submitted.

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

[HWc] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. pages 1–20. Preprint.

[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.

[Kap02] Ilya Kapovich. Subgroup properties of fully residually free groups. *Trans. Amer. Math. Soc.*, 354(1):335–362 (electronic), 2002.

[KM98] O. Kharlampovich and A. Myasnikov. Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups. *J. Algebra*, 200(2):517–570, 1998.

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

[Li02] Tao Li. Boundary curves of surfaces with the 4-plane property. *Geom. Topol.*, 6:609–647 (electronic), 2002.

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

[LW] Joseph Lauer and Daniel T. Wise. Cubulating one-relator groups with torsion. pages 1–19. Preprint, 2011.

[Mas08] Joseph D. Masters. Kleinian groups with ubiquitous surface subgroups. *Groups Geom. Dyn.*, 2(2):263–269, 2008.

[Min04] Ashot Minasyan. Separable subsets of gferf negatively curved groups. Preprint, 2004.
[Mit04] Mahan Mitra. Height in splittings of hyperbolic groups. *Proc. Indian Acad. Sci. Math. Sci.*, 114(1):39–54, 2004.
[MV95] John Meier and Leonard VanWyk. The Bieri-Neumann-Strebel invariants for graph groups. *Proc. London Math. Soc. (3)*, 71(2):263–280, 1995.
[MW02] Jonathan P. McCammond and Daniel T. Wise. Fans and ladders in small cancellation theory. *Proc. London Math. Soc. (3)*, 84(3):599–644, 2002.
[MW08] Jonathan P. McCammond and Daniel T. Wise. Locally quasiconvex small-cancellation groups. *Trans. Amer. Math. Soc.*, 360(1):237–271 (electronic), 2008.
[MZ08] Joseph D. Masters and Xingru Zhang. Closed quasi-Fuchsian surfaces in hyperbolic knot complements. *Geom. Topol.*, 12(4):2095–2171, 2008.
[Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309 (electronic), 2004.
[Oll06] Yann Ollivier. On a small cancellation theorem of Gromov. *Bull. Belg. Math. Soc. Simon Stevin*, 13(1):75–89, 2006.
[Osi06] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.
[Osi07] Denis V. Osin. Peripheral fillings of relatively hyperbolic groups. *Invent. Math.*, 167(2):295–326, 2007.
[OW11] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than $1/6$. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.
[Pap95] P. Papasoglu. Strongly geodesically automatic groups are hyperbolic. *Invent. Math.*, 121(2):323–334, 1995.
[Rat94] John G. Ratcliffe. *Foundations of hyperbolic manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.
[RS87] Eliyahu Rips and Yoav Segev. Torsion-free group without unique product property. *J. Algebra*, 108(1):116–126, 1987.
[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
[Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.
[Sel03] Z. Sela. Diophantine geometry over groups. II. Completions, closures and formal solutions. *Israel J. Math.*, 134:173–254, 2003.
[SW] Michah Sageev and Daniel T. Wise. Cores for quasiconvex actions.
[SW05] Michah Sageev and Daniel T. Wise. The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.*, 37(5):706–710, 2005.
[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.
[Wis] Daniel T. Wise. Virtual cleanliness. pages 1–21.
[Wis00] Daniel T. Wise. Subgroup separability of graphs of free groups with cyclic edge groups. *Q. J. Math.*, 51(1):107–129, 2000.
[Wis03] Daniel T. Wise. Sixtolic complexes and their fundamental groups. Preprint, 2003.
[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.
[Żuk96] Andrzej Żuk. La propriété (T) de Kazhdan pour les groupes agissant sur les polyèdres. *C. R. Acad. Sci. Paris Sér. I Math.*, 323(5):453–458, 1996.

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