Contents

CAT(0) Cube Complexes and Groups

Michah Sageev

IAS/Park City Mathematics Series
Volume XX, XXXX

CAT(0) Cube Complexes and Groups

Michah Sageev

Introduction

CAT(0) cube complexes are a particularly nice class of CAT(0) spaces. On the one hand, there is a rather broad class of groups which admit non-trivial actions on CAT(0) cube complexes. On the other hand, CAT(0) cube complexes have a combinatorial nature that give them the look and feel of trees. The added structure provided by hyperplanes has allowed for much more progress in understanding their geometry. For example, the Tits Alternative is known to hold for groups acting on CAT(0) cube complexes, but is still open in the setting of CAT(0) spaces generally. Finally, there are connections to other subjects, such as subgroup separability, 3-manifold theory, median algebras and spaces, and Kazhdan's Property (T).

The goal of these lectures is to give a brief introduction to world of CAT(0) cube complexes with the aim of giving the young geometric group theorist the tools to explore further directions of research. We assume no more than a cursory familiarity with CAT(0) spaces, with group actions on polyhedral complexes and covering space theory. For some of the applications, it may help to know something about 3-manifolds and hyperbolic groups, but these concepts are not necessary to understand the core material. Unfortunately, we did not have the time in these lectures to cover the connections to median spaces and algebras or Kazhdan's Property (T). One can learn about these topics in [**13**], [**35**], [**11**] and [**33**].

Dept. of Mathematics, Technion, Haifa 32000, Israel
Current address: Dept. of Mathematics, Technion, Haifa 32000, Israel
E-mail address: michah.sageev@gmail.com

LECTURE 1

CAT(0) cube complexes and pocsets

1. The basics of NPC and CAT(0) complexes

1.1. Cube complexes and links

We start by reviewing the basic notions of complexes and links that we will need. We will build complexes by gluing unit cubes along their faces by isometries. One is meant to imagine a disjoint union $\mathcal{C}$ of unit Euclidean cubes of various dimensions, together with a collection of isometries $\mathcal{F}$ between faces of cubes in $\mathcal{C}$. One then forms the quotient space $X = \mathcal{C}/\mathcal{F}$ obtained by identifying points in the domains of maps in $\mathcal{F}$ with their image. We will usually supress the $\mathcal{F}$ notation and label edges and faces we wish to identify. We thus have a quotient map $q : \mathcal{C} \to X$. The complex X is called a *cube complex*. The "cubes of X" are the images under q of the various faces of cubes in $\mathcal{C}$.

Recall that the *link* of a vertex in X is the simplicial complex which can be realized as a "small sphere" around the vertex. We describe what a link is a bit more precisely.

Note first that the 1-skeleton of X is a graph. It is possible that the 1-skeleton is not a simplicial graph, so that it may have loops and multiple edges. By a *local edge* of $\mathcal{C}$ we mean a subinterval of length $1/3$ of an edge of a cube in $\mathcal{C}$, one of whose endoints is a vertex of a cube in $\mathcal{C}$. A local edge in X is the image under q of a local edge in $\mathcal{C}$.

The link of a vertex v in X is a simplicial complex. The vertices of $lk(v)$ are the local edges of X containing v. A collection of vertices in $lk(v)$ span a simplex in $lk(v)$ if and only if the corresponding local edges of X are images under q of a collection of local edges all contained in the same cube in $\mathcal{C}$ and all of which share a vertex. See Figure 1.

1.2. Non-positively curved complexes

We are now going to restrict the classes of cube complexes that we will consider.

Definition 1.1. A *flag* complex is a simplicial complex with no "missing" simplices. This means that for each complete graph in the 1-skeleton of the complex, there is a simplex in the complex whose 1-skeleton is the given complete graph.

Examples:

(1) A simplicial graph is a flag complex if and only if it has no cycle of length three.
(2) Any simplicial graph is the 1-skeleton of a unique flag complex, obtained by attaching a simplex to each complete subgraph.
(3) The first barycentric subdivision of any simplicial complex is a flag complex. (exercise)

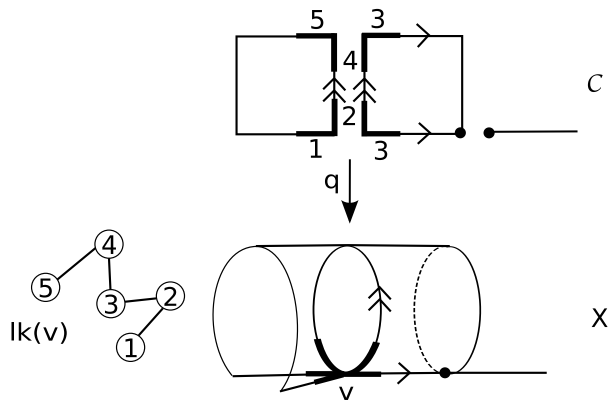

FIGURE 1. A square complex and a link.

Definition 1.2. A *non-positively curved (NPC) cube complex* is a cube complex whose vertex links are simplicial flag complexes. A 1-connected NPC complex is called a CAT(0) *cube complex*.

If we focus on NPC square complexes, we see that the above definition rules out the following identifications:

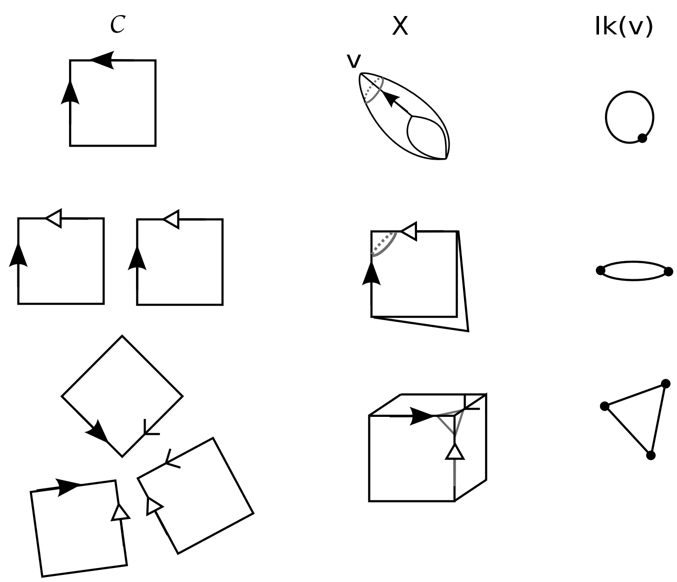

FIGURE 2. What is not allowed in NPC square complexes.

For NPC complexes of dimensions bigger than 2, the first two identifications are still not allowed, and the third identification is allowed only if there is a 3-dimensional cube containing the 3-squares in its boundary.

In these notes, we will work with this definition. However, since there is already a well established notion of CAT(0) which has to do with thin triangles in a geodesic metric space, some remarks are in order.

(1) One can define a path metric on a cube complex in the usual way as follows. We define a *rectifiable* path in X as one that can be broken into finitely many subpaths each of which is contained in some cube of X. If these paths are themselves rectifiable (in the classical sense) we can now define the length of the original path as the sum of the length of the subpaths. The distance between p and q is then defined as the infimum of the lengths of the rectifiable paths joining p and q.
(2) In the event that the complex is finite dimensional, a result of Bridson [**6**] tells us that the above indeed defines a metric and that with this metric, the complex is a complete, geodesic metric space. (The case of locally finite complexes was treated by Moussong [**30**].)
(3) A result of Gromov [**20**] then tells us that with this metric a finite dimensional NPC complex is locally CAT(0).
(4) The Cartan-Hadamard theorem then tells us that if the space is 1-connected it is CAT(0) in the usual sense.
(5) More recently, Leary [**28**], showed that all this makes sense in the infinite dimensional case as well. In particular, he showed that with the above metric, an NPC cube complex (possibly infinite dimensional) is a geodesic, locally CAT(0) metric space. Moreover, he showed that it is complete if and only if every ascending sequence of cubes terminates.

For a treatment of general CAT(0) spaces and how cubical complexes fit into that general theory, see Bridson and Heafliger's book [**7**].

We will sometime use the term *cubed group* for a group that admits a proper, cocompact action on a CAT(0) cube complex. Typically one constructs cubed groups by building compact NPC complexes. Their fundamental groups are then cubed groups.

Sometimes we will be interested in a CAT(0) cube complex that admits a cocompact group action, but we don't care too much about the group in question. Thus we use the term *cocompact* CAT(0) *cube complex* to mean a CAT(0) cube complex whose automorphism group acts cocompactly on it.

Let us now look at some examples.

Examples:

(1) **Graphs.** The link of a vertex in a graph has no edges, so every graph (simplicial or not) is an NPC complex. The universal cover of a graph is a tree, which is the model CAT(0) cube complex.
(2) **Tori.** A torus is obtained from a square, by identifying opposite edges and is thus naturally a cube complex. It is easy to check that the link of the sole vertex in this complex is a cycle of length 4. Thus a torus is an NPC complex. The reader should check that a torus of every dimension is naturally an NPC complex. What is the link of a vertex?
(3) **Surfaces.** Consider now an orientable surface of genus $g > 1$. Recall that it is obtained by taking a $4g$-gon and identifying faces in pairs in a suitable way (see Figure 3.) We can now subdivide the $4g$-gon into squares by adding the barycenters of the edges, a vertex in the center of the $4g$-gon, and an edge between each new edge-barycenter and the center of the $4g$-gon. The reader should check the vertex links to see that the complex obtained is indeed an NPC complex.

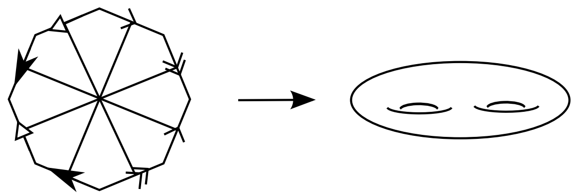

FIGURE 3. Squaring a surface of genus 2.

(4) **Products.** We first review the notion of a join of two simplicial complexes. Recall that an abstract simplicial complex K is simply a collection of finite subsets of some underlying set S, so that this collection is closed under taking subsets. The join of two complexes K_1 and K_2, with underlying sets S_1 and S_2, is the complex obtained by taking as the underlying set the disjoint union $S_1 \cup S_2$ and taking as the collection of subsets all pairwise unions of elements in K_1 and K_2. For example, the join of two simplicial complexes of dimension 0 is a complete bipartite graph. The reader should now check that the join of two flag complexes is a flag complex.

Now consider two cube complexes X and Y. Their product $X \times Y$ is naturally also a cube complex. The reader should check that if (v, w) is a vertex in $X \times Y$, then the link of (v, w) is the join of $lk(v)$ and $lk(w)$. So for example, the link of a product of two trees is a complete bipartite graph. Now since the join of two flag complexes is flag, the product of two NPC complexes is NPC. See Figure 4.

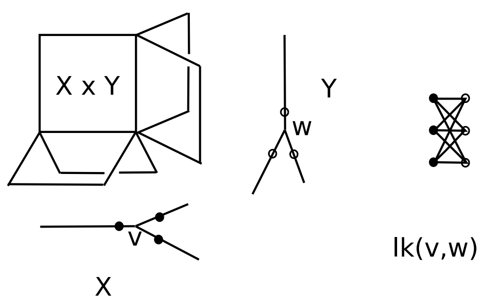

FIGURE 4. The local structure of a product of two trees.

Exercise 1.1. Show that a product of two trees does not embed in $\mathbf{R}^3$, even locally.

Exercise 1.2. Prove that a simply connected square complex whose link is a complete bipartite graph is a product of two trees.

As simple as they appear, groups acting on products of trees can be quite interesting. For example, Wise produced examples of groups acting on products of trees which had no finite index subgroups [**40**]. Burger and Mozes [**8**] produced the first example of finitely presented infinite simple groups as groups acting properly on a product of trees.

(5) **RAAGs** A *right angled Artin group* (RAAG) is a group with the following simple presentation. Start with a finite simplicial graph Γ and define

$$A(\Gamma) = \langle \Gamma^{(0)} | [v, w] \Leftrightarrow (v, w) \text{ is an edge of } \Gamma \rangle$$

So all abelian groups and free groups are RAAGs as are all products of free groups.

There is a natural complex associated to a RAAG, called the *Salvetti* complex $R(\Gamma)$, which we can build in the following way. Start with a single vertex and add a loop for each vertex of Γ. This is the 1-skeleton of $R(\Gamma)$. Now for every maximal n-clique in Γ we want to attach an n-torus to the complex. The n-torus can be seen as the quotient of a cube by identifying opposite sides. See Figure 5 for a simple example. The edges

FIGURE 5. The Salvetti complex of a RAAG.

of the cube descend to a collection of n loops in the torus intersecting at a simple loop. Identify these n loops in the torus with the n-loops in the 1-skeleton of $R(\Gamma)$ associated to our n-clique.

2. Hyperplanes

Hyperplanes are natural subspaces which "cut up" a CAT(0) cube complex into halfspaces. In this section we define them and describe some of their basic properties.

Definition 1.3. Given an n-dimensional cube σ in X, a *midcube* of a sigma is an $n-1$-dimensional unit cube running through the barycenter of σ and parallel to one of the faces of σ. Thus each n-cube has n midcubes, all intersecting at the barycenter.

Let X be an NPC cube complex. Let $\Box$ denote the equivalence relation on the edges of X generated by $e \Box f$ if and only if e and f are opposite edges of some square in X.

Definition 1.4. Given an equivalence class of edges $[e]$, the *hyperplane* dual to $[e]$ is the collection of mid cubes which intersect edges in $[e]$.

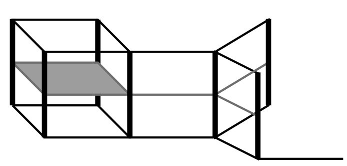

FIGURE 6. A hyperplane in a cube complex.

It is useful to go back to the previous examples and think about what the hyperplanes there look like.

(1) When X is a graph, the hyperplanes are simply midpoints of edges.
(2) When X is a torus, each hyperplane is a simple closed curve. The universal cover of X is simply the Euclidean plane tiled by squares, and the hyperplanes there are horizontal or vertical lines.
(3) When X is a higher genus surface, squared as above, it is easy to check that the hyperplanes are simple closed curves lifting to lines in the universal cover of X.
(4) When X is a product, the hyperplanes are preimages (under the natural projection maps) of the hyperplanes in each of the factors.
(5) Regarding Salvetti complexes, you should do the following exercise.

Exercise 1.3. Show that when X is a Salvetti complex, each hyperplane is itself a Salvetti complex. Describe the RAAG that it is the Salvetti complex of.

Here are some more exercises.

Exercise 1.4. Find an NPC square complex with a single hyperplane.

Exercise 1.5. Find an infinite NPC complex with two hyperplanes. Then find one with a single hyperplane.

In an NPC complex, hyperplanes can be immersed in complicated ways, in a CAT(0) cube complex, they are much better behaved. This is captured in the following basic theorem, which we will make much use of.

Theorem 1.1. *Let X be a* CAT(0) *cube complex and $\hat{\mathfrak{h}}$ a hyperplane in X. Then the following statements hold.*

(1) Every hyperplane of X is embedded.
(2) Every hyperplane of X separates X into precisely two components.
(3) Every hyperplane is a CAT(0) *cube complex.*
(4) Every collection of pairwise intersecting hyperplanes intersects.

This theorem can be proved using CAT(0) geometry (see [**36**]) or using disk diagrams (see [**37**].) We will forgo providing the proof here and view this as a starting point.

The following is also useful.

Exercise 1.6. Show that if G acts cocompactly on a CAT(0) cube complex X, then every hyperplane is acted on cocompactly by its stabilizer.

Hint. Any isometry which sends a cell of a hyperplane to a call of the same hyperplane preserves the entire hyperplane.

3. The pocset structure

The consequence of Theorem 1.1 is that we have a nice collection of sub-CAT(0) complexes which cut up the complex into halfspaces. We want to focus on the combinatorial nature of the collection of halfspaces. Let us first set some notation.

X – CAT(0) cube complex

$\hat{\mathcal{H}}$ – hyperplanes of X

$\mathcal{H}$ – halfspaces of X
$\mathfrak{h}$ – a halfspace in $\mathcal{H}$
$\mathfrak{h}^*$ – the complementary halfpspace of $\mathfrak{h}$
$\hat{\mathfrak{h}}$ – the bounding hyperplane of $\mathfrak{h}$

Note that $\mathcal{H}$ is a poset under inclusion such that

(1) $\mathcal{H}$ has a natural order-reversing involution $\mathfrak{h} \mapsto \mathfrak{h}^*$ so that $\mathfrak{h}$ and $\mathfrak{h}^*$ are incomparable (this is called a *pocset*),
(2) $\mathcal{H}$ is locally finite (meaning that there are finitely many elements between any two given elements),
(3) $\mathcal{H}$ there is a bound on the size of a collection of halfspaces which are not nested after possibly replacing some of them with their complements (follows from finite dimensionality).

Following Roller [**35**], we are thus motivated to define the following notion of a *pocset* ("poset with complementation").

Definition 1.5. A *pocset* is a poset Σ together with an order reversing involution $A \mapsto A^*$ such that

- $A \neq A^*$ and A and A^* are incomparable
- $A < B \Rightarrow B^* < A^*$

Elements $A, B \in \Sigma$ are *nested* if one of $A < B, A < B^*, A^* < B, A^* < B^*$ holds. Otherwise we say that A and B are *transverse*.

If $A < B$ then we let $[A, B] = \{C | A < C < B\}$. This is called the *interval between A and B*.

A pocset is said to be *locally finite* if every interval is finite.

The *width* of a pocset is the size of a maximal transverse subset.

We will focus on locally finite, finite width pocsets. Some of the claims will hold in somewhat broader generality, namely when every transverse subset is finite (see [**21**]).

Example. A *space with walls* is simple a set a pair (S, Σ) is simply a set S together with a collection of subsets Σ closed under complementation (see Haglund and Paulin [**23**]). A space with walls clearly forms a pocset under inclusion. It is said to be *discrete* if for any two elements of $a, b \in S$ the collection of subsets in Σ containing a and not containing b is finite. We will see several examples of spaces with walls later on. A simple example comes from hyperbolic surfaces. Take any finite collection of essential geodesics on a closed surface. The collection of lifts $\mathcal{L}$ of these curves to the universal cover $\mathbf{H}^2$ is a collection of lines. The space $\mathbf{H}^2 - \mathcal{L}$ is now naturally a discrete space with walls, where the walls are given by the lines in $\mathcal{L}$; an element of Σ is a complementary region of a single line in $\mathcal{L}$.

Example. As already mentioned, a particular example of a space with walls is the pocset of halfspaces of a CAT(0) cube complex. This pocset is locally finite and the dimension of the complex is the width of the pocset.

Exercise 1.7. Prove or disprove or salvage if possible. A space walls (S, Σ) is discrete if and only if Σ is locally finite.

LECTURE 2

Cubulations: from pocsets to CAT(0) cube complexes

In this lecture, we will describe the cubulation construction described in [**37**]. There are variants of this construction in various settings (Roller [**35**], Nica [**34**], Chatterji-Niblo [**12**] and Guralnik [**21**]). We will follow most closely Roller's treatment.

1. Ultrafilters

We assume now we have a locally finite pocset Σ. We wish to construct a CAT(0) cube complex $X(\Sigma)$ from Σ. First, we will describe the vertex set.

Definition 2.1. An *ultrafilter* α on Σ is a subset of Σ satisfying

(1) **Choice:** for all pairs $\{A, A^*\}$ in Σ, precisely one of them is in α.
(2) **Consistency:** $A \in \alpha$ and $A < B \implies B \in \alpha$.

The notion of an ultfilter on a pocset is reminiscent from the classical notion of an ultrafilter on the natural numbers, where the pocset is the collection of subsets of $\mathbf{N}$. However, as we shall see, ultrafilters have geometric meaning; they will be the vertices or "vertices at infinity" of a CAT(0) cube complex.

An ultrafilter α is said to satisfy the *Descending Chain Condition (DCC)* if every descending chain of elements terminates.

Example. The first example to think about is the case of a tree. The pocset Σ is the collection of halfspaces of the tree. Here a halfspace is simply a complementary region of the midpoint of an edge. Each edge has two halfspaces associated to it, and an ultrafilter will make a choice of one of these. Thus we can view an the choice condition of an ultrafilter as as a way of putting an arrow on each edge, where the arrow points towards the chosen halfspace. The consistency condition restricts the way the arrows can be oriented, as shown in Figure 1.

Note that this implies that if you choose an orientation at some edge, then for all the arrows in the tail halfspace of that edge, the arrows must point towards the edge.

One type of ultrafilter can be obtained by choosing a vertex v and have all the arrows pointing at it. So the ultrafilter α_v, is given by

$$\alpha_v = \{\mathfrak{h} | v \in \mathfrak{h}\}$$

So each vertex is associated with an ultrafilter. Note that these all satisfy DCC.

Conversely, given an ultrafilter that satisfies DCC, it is not hard to see that there then exists some vertex (and hence a unique vertex) such that all the arrows point at that vertex. Thus the vertices are of the tree are in one-to-one correspondence with DCC-ultrafilters.

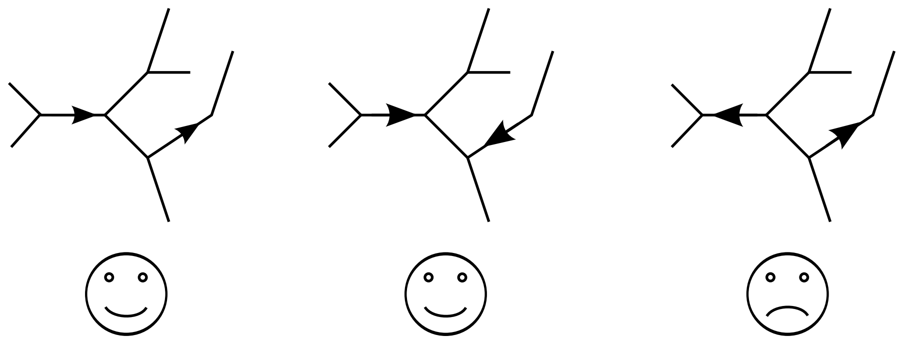

Figure 1. The orientation on the right is not allowed.

Now suppose that α is an ultrafilter which does not satisfy DCC. For every vertex v, the consistency condition ensures that at most one edge adjacent to v is oriented away from v. Moreover, since α does not satisfy DCC, it follows that exactly one edge is oriented away from v. This means that α determines a ray emanating from v. See Figure 2.

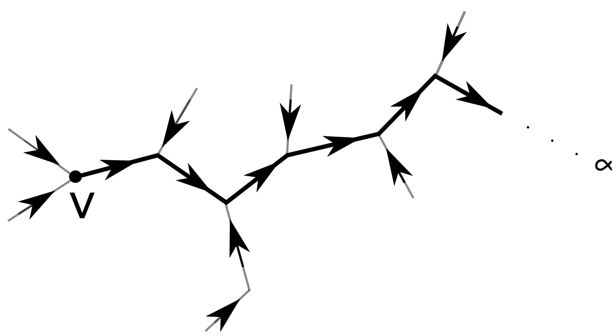

Figure 2. An ultrafilter at infinity.

We thus have a picture of all the ultrafilters. The DCC ultrafilters are the the vertices of the tree and the rest correspond to boundary points of the tree.

Example. Another elementary example that is good to think about is the usual squaring of the plane (Figure 3). As in the case of a tree the DCC ultrafilters correspond to the vertices of the complex. The other ultrafilters are “at infinity” as shown below. There is a line of ultrafilters on all four “sides” of the plane and an ultrafilter for each “corner”

What we see in the above two examples turns out to be the general picture as well: the collection of ultrafilters forms a compactification of the complex. We will not delve into this in these notes; for further discussion, see [**21**] and [**31**]. We just note here that for a pocset coming from the halfspaces of a finite dimensional CAT(0) cube complex, the DCC ultrafilters are the same as the vertices.

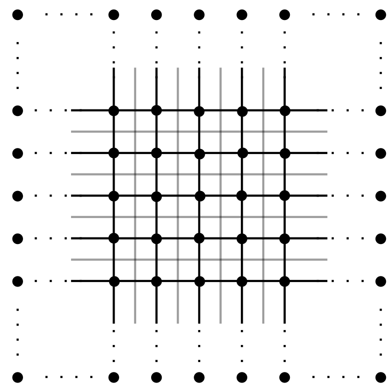

FIGURE 3. The ultrafilters associated to the squaring of the Euclidean plane.

Exercise 2.8. Let X be a finite dimensional CAT(0) cube complex and consider the pocset of halfspaces $\mathcal{H}$. Then the DCC ultrafilters are precisely the ultrafilters associated to vertices. Namely, every DCC ultrafilter is of the form

$$\alpha_v = \{\mathfrak{h} | v \in \mathfrak{h}\}$$

2. Constructing the complex from a pocset

We are now given a discrete, finite width pocset Σ and we wish to construct a CAT(0) cube complex $X = X(\Sigma)$.

Vertices. The vertex set X^0 of X will be the collection of DCC ultrafilters on Σ. From now on we will use letters like v and w to denote vertices of X^0.

Edges. We join two such vertices v and w by an edge if $|v \triangle w| = 2$. That is, there exists $A \in \Sigma$ such that $w = (v - \{A\}) \cup \{A^*\}$.

Exercise 2.9. Let $A \in v$, then $(v - \{A\}) \cup \{A^*\}$ is an ultrafilter if and only if A is minimal in v. Note that $(v - \{A\}) \cup \{A^*\}$ is also a DCC ultrafilter and hence another vertex of X.

Notation. When A is minimal in v, we will use the following notation:

$$(v; A) \equiv (v - \{A\}) \cup \{A^*\}$$

When B is minimal in $(v; A)$, we will use the following notation:

$$(v; A, B) \equiv ((v; A); B)$$

And similarly we use the notation $(v; A_1, \ldots, A_n)$ for multiple elements.

Having constructed X^1, we now need to know that it is connected. This is the content of the next exercise.

Exercise 2.10. If Σ has finite width, then any two DCC ultrafilters are joined by a finite path.

Hint. Consider two DCC ultrafilters v and w. One needs to show that there are finitely many A's in Σ such that $A \in v$ and $A^* \in w$. (Why is this enough?) Suppose there are infinitely many such and use finite width to build a descending sequence of elements in v or w which does not terminate.

Squares. We attach a square to $X^{(1)}$ to every 4-cycle appears in $X^{(1)}$.

Let us look a bit more closely at the vertices of a square. Let's say one of them is v, then by the way we defined the 1-skeleton, there are $A, B \in v$ such that the two vertices of the square adjacent to v are $(v; A)$ and $(v; B)$. Now it is not too difficult to see that the vertex diagonally opposite v in the square is simply $(v; A, B)$. This in turn tells us that A and B are transverse. (The reader should check this.)

Since ultimately X is supposed to be CAT(0), we better have a simply connected 2-skeleton.

Exercise 2.11. $X^{(2)}$ is simply connected.

Hint. Consider the shortest possible loop γ in X^1 which is non-trivial in $\pi_1(X^2)$. Let v be some vertex along γ. Then we will have along γ a sequence of vertices v, $(v; A_1), \ldots, (v; A_1, \ldots, A_n), v(A_1, \ldots, A_n, A_k^*)$ such that

(1) For all i, j with $1 \leq i < j \leq n$, we have $A_i \neq A_j$ and $A_i \neq A_j^*$.
(2) $1 \leq k \leq n$

If $k = n$, then we have backtracking along γ and it is not the shortest non-trivial loop. Now argue that A_n and A_k are transverse. Use this to produce a new loop which has along it the following sequence of vertices:

$$v, (v; A_1), \ldots, (v; A_1, \ldots, A_{n-1}, A_k^*), (v; A_1, \ldots, A_{n-1}, A_k^*, A_n)$$

Proceed until backtracking is produced.

Higher dimensional cubes. We now construct the n skeleton inductively. Simply add and n-cube whenver the boundary of one appears in the $n-1$ skeleton. It is again instructive to think about what happens locally when v is a vertex of an n-cube σ. The neighboring vertices in the n-cube are of the form $(v; A_1), \ldots, (v; A_n)$. Since any pair of such edges spans a square, we have that $\{A_1, \ldots, A_n\}$ is a collection of pairwise transverse elements. We then see that all the vertices of σ are of the form

$$(v; A_{i_1}, \ldots, A_{i_k})$$

for some distinct collection of indices $i_j \in \{1, \ldots, n\}$.

Conversely, a collection $A_1, \ldots, A_n$ of pairwise transverse elements of some vertex v, defines for us a n-cube. These observations allow us to establish the following.

Exercise 2.12. X satisfies the Gromov flag link condition.

Not that the dimension of X is equal to the width of the original pocset Σ. We call this construction of a cube complex from a pocset a *cubulation*.

Exercise 2.13. Explore this construction when the pocset is not of finite width. For example, suppose the pocset is completely un-nested: no two elements are comparable. Is the complex connected? What do the components look like? etc.

Group actions. Let Σ be a pocset and suppose that a group G acts on Σ in an order preserving manner and without inversions (for no $g \in G$ and $A \in \Sigma$ do we have $gA = A^*$). Then we obtain an action on $X(\Sigma)^0$:

$$g\alpha = \{g\mathfrak{h} | \mathfrak{h} \in \alpha\}$$

It is easy to check that this extends to an action on $X(\Sigma)$ by cellular isometries.

3. Examples of cubulations

3.1. The pocset of halfspaces

Suppose that X is a finite dimensional CAT(0) cube complex. As already noted the collection of halfspaces $\mathcal{H}(X)$ of X is a pocset.

Proposition 2.1. *The cubulation obtained from $\mathcal{H}(X)$ is X.*

Proof. Let $Y = X(\mathcal{H}(X))$, namely the cubulation obtained from $\mathcal{H}(X)$. First, note that every vertex in $v \in X^0$ determines a DCC ultrafilter

$$\alpha_v = \{\mathfrak{h} | v \in \mathfrak{h}\}$$

And so we have a map

$$\Phi : X^0 \to Y^0$$
$$v \mapsto \alpha_v$$

It is easy to see that Φ is injective, since any two vertices are separated by some hyperplanes of X and hence determine different ultrafilters.

For any two adjacent vertices $v, w \in X$, they are separated by a unique hyperplane $\hat{\mathfrak{h}}$ transverse to the edge of which v and w are endpoints. Thus, by construction, it follows that α_v and α_w will be joined by an edge in the Y. We thus can extend Φ to the 1-skeleton X^1.

Once this is done, it is not hard to see that the map can be extended to the higher dimensional skeleta, since by construction, a cube is attached to Y for every 1-skeleton of a cube that appears in Y^1.

Finally, it remains to show that Φ is onto. Suppose that α is a vertex of Y, namely a DCC ultrafilter on the pocset $\mathcal{H}$. Let α_v be some vertex of Y which is in the image of Φ and let $\mathfrak{h}_1, \dots, \mathfrak{h}_n \in \alpha_v$ be the halfspaces so that

$$\alpha = (\alpha_v; \mathfrak{h}_1, \dots, \mathfrak{h}_n)$$

Now $\mathfrak{h}_1$ is minimal in α_v. This means that if we consider the vertex v in X, the hyperplane $\hat{\mathfrak{h}}_1$ is transverse to an edge adjacent to v; let $[v, w]$ be that edge. Now we observe that

$$\Phi(w) = (\alpha_v; \mathfrak{h}_1)$$

We continue in this manner finding the vertices that get mapped by Φ to $(\alpha_v; \mathfrak{h}_1, \dots, \mathfrak{h}_i)$ for each $i \geq 1$. We thus find a vertex mapped by Φ to α, as required. $\square$

3.2. Lines in the plane

A collection of lines in the plane is called *discrete* if there is a lower bound to the distance between two parallel lines in the collection.

Exercise 2.14. Let $\mathcal{L}$ be a discrete collection of lines which has finitely many parallelism classes (for example, think of the plane triangluated by unit isosceles triangles). Consider the set $S = \mathbf{R}^2 - \mathcal{L}$. Then the pocset associated to the space with walls (S, Ω) is a finite width, locally finite pocset. What is the cube complex associated to this?

3.3. Small cancellation groups (Wise)

We give here a (very) brief account of Wise's cubulation of small cancellation groups. For a complete discussion, see [**41**].

Let $G = \langle S|R \rangle$ be a finitely presented group, where S is closed under taking inverses. We consider the presentation 2-complex K associated to this presentation. The universal cover $\tilde{K}$ is a 2-complex, called the *Cayley complex* of the presentation, whose 2-cells we call *relator polygons*. The 1-skeleton of $\tilde{K}$ can be identified with the Cayley graph of the presentation and so that the each edge of $\tilde{K}$ is labeled by an element of S. A *piece* of the presentation is a reduced word in S that appears as the label of a path in $\tilde{K}^1$ which is contained in the boundary of more than one 2-cells of $\tilde{K}$.

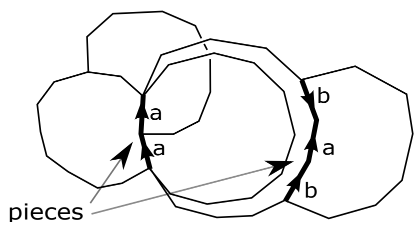

FIGURE 4. Part of the Cayley complex and some pieces.

The presentation is said to be a C$'$(1/n) *presentation* for G if the length of a piece is always less than $1/n$'th the length of the boundary of a relator polygon in which it appears. There are many more small cancellation conditions which come up in small cancellation theory (see [**29**]).

Small cancellation groups are studied via disk diagrams which are disks (actually something slightly more general called a multi-disk) tiled by relator polygons. A multi-disk is a connected, simply connected union of disks and finite trees in the plane, where each of the disks and trees meet along their boundaries. Given a trivial word w in G, we can represent the triviality of w by drawing a labeled multi-disk, as seen in Figure 5, and a tiling of the disk regions by relator polygons.

The original word w is obtained by reading the labels around the outside of the multi-disk. The multi-disk together with the labeling of the edges by generators and the tiling by relator polygons is called a *disk diagram*. For small cancellation groups there is a fundamental lemma [**42**] which gives a trichotomy regarding disk diagrams. This goes back at least to Greendlinger [**19**] and even earlier to Dehn [**15**].

Lemma 2.1 (Fundamental Lemma)**.** *A disk diagram for a word in a* C$'$(1/6)*-group is of one of the following types.*

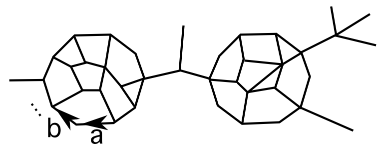

FIGURE 5. A multi-disk. The edges are are labeled by generators and the cells by relator polygons.

(1) *A single polygon.*
(2) *A ladder: this means a diagram formed by attaching polygons and/or edges "end-to-end" in a linear fashion. The two endpoints of the ladder are either* shells *or* spurs. *A shell is a polygon which is attached to the diagram along an arc whose length is less than half of the length of its boundary. A spur is simply an edge which is attached to the diagram at one endpoint and free at the other endpoint.*

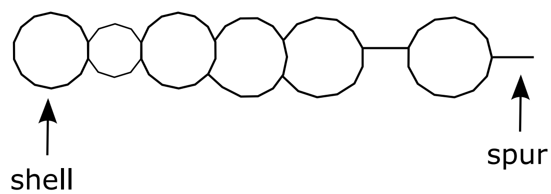

(3) *A diagram with at least 3 shells and/or spurs. A typical diagram might look like this.*

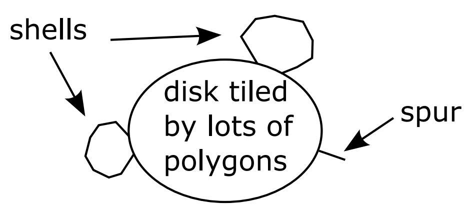

From this Fundamental Lemma one can deduce that $C'(1/6)$ have a linear isoperimetric inequality and hence such groups are Gromov hyperbolic (see [**20**]).

We now describe how one finds walls in $\tilde{K}$. First, if necessary, subdivide the boundary of the relator polygons so that each relator polygon has an even number of sides. We now describe certain types of tracks (a la Dunwoody) called *wisetracks*; these tracks will serve as the walls and the complement of the their union will serve as the space with walls.

We build a graph Δ as follows. We have a vertex $v_e \in \Delta$ for each edge e in $\tilde{K}^1$. Two vertices v_e and v_f are joined by an edge in Δ if e and f are opposite edges in some relator polygon of $\tilde{K}$. We then have a natural map

$$\eta : \Delta \to \tilde{K}$$

Which maps each vertex v_e of Δ to the midpoint of edge e in $\tilde{K}^1$, and each edge $[v_e, v_f]$ of Δ to a straight arc in the appropriate relator polygon joining the midpoints of the edges e and f. (Notice that the $C'(1/6)$ property insures that there is unique such relator polygon.)

The image of a connected component of Δ is called a *wisetrack*.

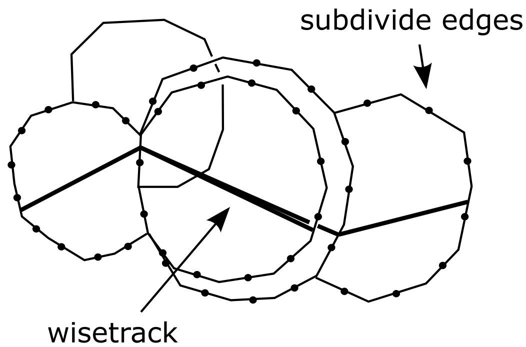

Figure 6. Some wisetracks in a small cancellation complex.

Using small cancellation theory Wise then shows that each wisetrack is embedded. The idea is to use the fundamental lemma. If a wisetrack crosses itself, one sees a sequence of relator polygons as below, and this in turn gives rise to a a disk diagram without 3 shells or spurs, contradicting the Fundamental Lemma.

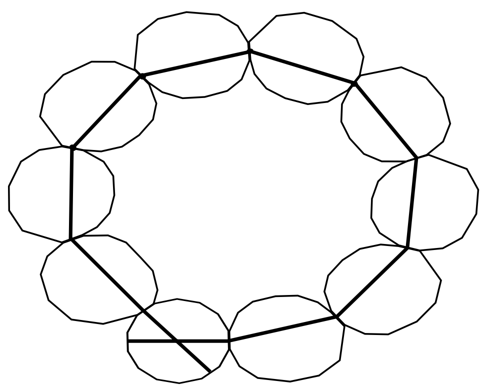

Figure 7. A self-intersection leads to a diagram which violates the Fundamental Lemma.

A result of Dunwoody then tells us that each of these wisetracks separate $\tilde{K}^2$. So we take the space with walls to be the complement of the union of all wisetracks. The original group acts on $\tilde{K}$ and therefore acts on the cube complex. We will say more about the complex and the action once we say some more about actions.

3.4. Coxeter groups

In a similar vain, Niblo and Reeves [**32**] cubulated Coxeter groups. The space with walls can be described in terms of tracks in an appropriate presentation.

Recall that a Coxeter group has a presentation of the following form

$$G = \langle S | s_i^2 = 1, (s_i s_j)^{m_{ij}} \rangle$$

Where $S = \{s_1, \ldots, s_n\}$ is a finite set, $m_{ij} = m_{ji}$ and $2 \leq m_{ij} \leq \infty$.

We let K denote the presentation 2-complex for this and $\tilde{K}$ its universal cover. Note that the presentation complex for this presentation has relator polygons with an even number of sides. Also note that in $\tilde{K}$, for each pair i, j with $m_{ij} < \infty$, and each polygon P which reads $(s_i s_j)^{m_{ij}}$ along its boundary, there are a total of m_{ij} other polygons which share the same boundary as P. We can construct a quotient of $\tilde{K}$ in which each of these "pillows" is collapsed to a single polygon. Now we build tracks as in the small cancellation case.

Niblo and Reeves then check that the resulting complex is finite dimensional and the action on it is proper. Caprace then showed that for the action is cocompact unless the original Coxeter group contains a Euclidean triangle group. For details, see [**32**] and [**9**].

3.5. Codimension 1 subgroups

A general situation where the above construction is applicable is when G is a finitely generated group and H is a subgroup that "separates" the Cayley graph of G. More precisely, the subgroup H is said to be a *codimension 1* subgroup if the coset graph G/H has more than one end. In the Cayley graph, this translates to the following: there exists a number R, such that the R-neighborhood of H separates the Cayley graph into two *deep* components. If we choose one such component $A \subset G$, we see that the translates under G of A and its complement A^*, form a collection of walls Σ on the set G. With a bit of work, one can show that this is a discrete collection of walls. Moreover, the action of G on the resulting cube complex is fixed-point free. See [**37**], [**33**], [**16**] for more details.

4. Cocompactness and properness

In the examples given in Sections 3.3 and 3.4, one cubulates using a collection of walls and then one would like to know that the action is proper and better yet, proper and cocompact. It turns out that cocompactness is assured by hyperbolicity and quasiconvexity, and properness by a kind of "filling" condition. Before stating some general theorems, let us look at a simple example to highlight the ideas.

4.1. An example: curves on surfaces

As discussed in Section 3, an example of a space with walls can be obtained by considering a finite collection of simple closed geodesics on a closed hyperbolic surface S. The universal cover of S is identifiied with the hyperbolic plane $\mathbf{H}^2$ and the curves on S lift to a collection of lines $\mathcal{L}$ in the univeral cover. The space with walls is $\mathbf{H}^2 - \bigcup_{\ell \in \mathcal{L}} \mathcal{L}$ and the walls are the halfspaces defined by the lines in $\mathcal{L}$. We let $G = \pi_1(S)$.

A collection of lines in $\mathcal{L}$ that pairwise intersect is called *transverse*. If we review the cubulation construction in which a cube complex X is contructed from this space with walls, we see that cocompactness is implied by the following two claims.

Claim 2.1. *For each $k > 0$, there are finitely many G-orbits of transverse collections of k lines in $\mathcal{L}$*

Claim 2.2. *There is a bound on the size of a transverse collection of lines in $\mathcal{L}$.*

Claim 2.1 ensures that in X/G, there are finitely many cubes in each dimension. Claim 2.2 ensures that the complex X is finite dimensional. To prove these claims we will need the following, which we leave as an exercise.

Exercise 2.15. Let L be a transverse collection of n lines in $\mathbf{H}^2$, with $n > 1$. Then, there exists a number $R = R(L) > 0$ such any line intersecting all the lines in L, intersects the ball of radius R about the origin.

Proof of Claim 2.1. For $k = 1$, the statement is that there are finitely many conjugacy classes of lines, which is simply the fact that there are finitely curves in the quotient of $\mathbf{H}^2$ under the action of $G = \pi_1(S)$.

For $k = 2$, any two transverse lines have a point of intersection, which by cocompactness, can be translated into some fixed fundamental domain D for the action. Since only finitely many lines intersect D, there are only finitely many points of intersection in D.

We now proceed by induction. Let $L = \{\ell_1, \ldots, \ell_{k+1}\}$ denote a transverse collection of lines in L. By induction, we can translate L so the $\{\ell_1, \ldots, \ell_k\}$ is one of finitely many transverse collections. Now by the exercise, ℓ_{k+1} meets the ball of radius R about the origin. By the discreteness of the pattern, there are only finitely many choices for ℓ_{k+1}. □

Exercise 2.16. Prove Claim 2.2.

Thus, we see that the cube complex construction yields a cocompact action by G.

Now let us think about properness. Each element of the group G is a hyperbolic isometry of the hyperbolic plane and hence has an axis. We say that the pattern of lines $\mathcal{L}$ is *filling* if for every $g \in G$, there exists a line $\ell \in \mathcal{L}$, such that the axis ℓ_g of g crosses ℓ. It is not hard see that this corresponds to each complementary regions of the union of lines in $\mathcal{L}$ being bounded. Each of these complementary regions corresponds to a vertex in the resulting cube complex and it is also then not hard to see that the orbit of such a vertex is unbounded. In fact, since G is torsion free, this tells us that the action of G on X is not only proper, but free.

4.2. Hyperbolic groups, quasi-convex subgroups and hyperbolic 3-manifolds
We imagine that we are in a more general situation in which G is a hyperbolic group and H is a quasi-convex codimension-1subgroup. We recall that a hyperbolic group G has a natural visual boundary ∂G and that H has a limit set $\Lambda(H) \subset \partial G$. We refer the reader to any a reference on hyperbolic groups ([**20**] or [**17**], for example.) As in Section 3.5, we obtain a pocset Σ from which we can obtain an action of G on a CAT(0) cube complex X.

Regarding cocompactness, we have the following theorem [**18**].

Theorem 2.1 (Gitik-Mitra-Rips-S)**.** *If G is a hyperbolic and H is quasiconvex, then the action of G on X is cocompact.*

Remark. The same proof applies if we apply the cubulation construction to a finite collection of codimension-1 subgroups.

Regarding properness, thinking along the same lines as the example above leads to the following.

Theorem 2.2 (Bergeron-Wise)**.** *If G is hyperbolic and H is quasi-convex, such that for ever element $g \in G$, there exists a conjugate of H whose limit set $\Lambda(H)$ separates the endpoints of the axis of g, then the action of G on X is proper.*

Bergeron and Wise [**5**] actually prove a more applicable result.

Theorem 2.3 (Bergeron-Wise). *Let G be a hyperbolic group. Suppose that for every pair of points a, b in ∂G, there exists a quasiconvex codimension-1 subgroup H whose limit set separates a and b. Then there exists a finite collection of codimension-1, quasiconvex subgroups such that the action of G on the resultying cube complex is proper,*

A particular application of this theorem is in the setting of 3-manifold groups, in light of the following deep result of Kahn and Markovic [**27**].

Theorem 2.4 (Kahn-Markovic). *Let $M = \mathbf{H}^3/G$ be a closed hyperbolic 3-manifold. Then every great circle in $S^2 = \partial \mathbf{H}^3$ is a limit of quasicircles which are limit sets of quasi-fuchsian subgroups. In particular, every pair of points in S^2 is separated by the limit set of a quasiconvex surface subgroup.*

Putting the Kahn-Markovic theorem together with the above theorems on cubulations, we obtain

Corollary 2.1. *Every hyperbolic 3-manifold acts properly and cocompactly on a* CAT(0) *cube complex.*

5. Roller duality

5.1. Statement of duality

We have seen two construction in this lecture:

$$\text{cube complex } X \rightsquigarrow \text{pocset of halfspaces } \mathcal{H}(X)$$

$$\text{pocset } \Sigma \rightsquigarrow \text{ cube complex } X(\Sigma)$$

Exercise 2.17 (Roller Duality). These constructions are dual to one another:

(1) Given a finite width locally finite pocset, Σ, then $\mathcal{H}(X(\Sigma)) \equiv \Sigma$.
(2) Given a finite dimensional cube complex X, $X(\mathcal{H}(X)) = X$.

Proof. □

5.2. Applications

Subpocsets and collapsing. If Σ is pocset and $\Delta \subset \Sigma$ is a *subpocset* (i.e. a subset closed under involution), then there is a natural map $\rho_\Delta : \mathcal{U}(\Sigma) \to \mathcal{U}(\Delta)$, defined by

$$\rho_\Delta(\alpha) = \alpha \cap \Delta$$

It is elementary to check that $\rho_\Delta(\alpha)$ is indeed an ultrafilter. It is also easy to see that this maps sends DCC ultrafilters to DCC ultrafilters. Moreover it then extends to the cubes of X and we obtain a map $\rho_\Delta : X(\Sigma) \to X(\Delta)$.

To see what the map ρ_Δ looks like, note that the collection of cubes meeting a hyperplane $\hat{\mathfrak{h}}$ is isometric to $\hat{\mathfrak{h}} \times I$. We call this the *carrier* of $\hat{\mathfrak{h}}$. Since the pocset of halfspaces of the cube complex $X(\Delta)$ is simply Δ, the map ρ_Δ collapses the carrier of every hyperplane associated to half-spaces which are not in Δ in the I direction.

We now give some examples of this construction.

Orbit quotients. Suppose that a group G acts on a CAT(0) cubical complex X. Then given a hyperplane $\hat{\mathfrak{h}}$, we can look at the orbit of $\hat{\mathfrak{h}}$ under G. We then get the pocset $G(\mathfrak{h} \cup \mathfrak{h}^*)$, which of course is a subpocset of $\mathcal{H}$. By the collapsing construction above, we get a new CAT(0) cubical complex $X(G, \hat{\mathfrak{h}})$. We call this the orbit quotient of X associated to $\hat{\mathfrak{h}}$. This quotient has the property that there is a single orbit of hyperplanes, which is sometimes useful.

Exercise 2.18. Consider $\mathbf{Z} \times \mathbf{Z}$ acting on the standard squaring of the plane. What are the orbit quotients?

Exercise 2.19. Consider the standard description of the surface of genus two given as the quotient of the octagon whose edges are identified $ab\bar{a}\bar{b}cd\bar{c}\bar{d}$. Square the surface by putting a vertex in the middle and joining this vertex to the midpoint of each edge. Let X be the universal cover of this surface acted on by the fundamental group of the surface G.

(1) What are the orbit quotients? Are they locally finite?
(2) Are the actions on the orbit quotients proper?
(3) G acts on the product of the orbit quotients. Is this action proper? Is it cocompact?

Products.

Corollary 2.2 (Recognizing Products)**.** *Let X be a* CAT(0) *cube complex and $\hat{\mathcal{H}}$ its collection of hyperplanes. Then a decomposition of X into a product $X = X_1 \times X_2$, corresponds to a decomposition of $\hat{\mathcal{H}}$ as a disjoint union $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \hat{\mathcal{H}}_2$ where every hyperplane in $\hat{\mathcal{H}}_1$ crosses every hyperplane in $\hat{\mathcal{H}}_2$.*

Exercise 2.20. Prove this corollary

Hint. The direction that has not been discussed before is the one in which one is given a decomposition of the hyperplanes as a disjoint union $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \hat{\mathcal{H}}_2$. Build the cube complexes $X(\mathcal{H}_1)$ and $X(\mathcal{H}_2)$. Show that $\mathcal{H}(X)$ has the same pocset structure as $\mathcal{H}(X(\mathcal{H}_1) \times X(\mathcal{H}_2))$.

A CAT(0) cube complex is called irreducible if it is not a product of two complexes. Applying Corollary 2.2 we obtain the following

Corollary 2.3. *Let X be a finite dimensional* CAT(0) *cube complex then X admits a canonical decomposition as a product of finitely many irreducible factors (up to permutation of factors).*

Proof. Consider a maximal decomposition of X as a product $X = \prod_{i=1}^n X_i$. Note than n is bounded by the dimension of X, so that each of the X_i is irreducible. Now suppose that $X = \prod_{j=1}^m Y_i$ is another decomposition of X into irreducibles. We then obtain transverse disjoint decompositions of the collection of hyperplanes of X

$$\hat{\mathcal{H}} = \bigcup_{i=1}^n \hat{\mathcal{H}}_i = \bigcup_{j=1}^m \hat{\mathcal{K}}_j$$

Since each $\hat{\mathcal{H}}_i$ does not admit a disjoint transverse decomposition, we have that for each i, there exists j such that $\hat{\mathcal{H}}_i \subset \hat{\mathcal{K}}_j$. Similarly, for each $\hat{\mathcal{K}}_j$, there exists some $\hat{\mathcal{H}}_i$ such that $\hat{\mathcal{K}}_j \subset \hat{\mathcal{H}}_i$. Putting these two facts together and the fact that these are disjoint decompositions of $\hat{\mathcal{H}}$, we obtain that for each i there exists j, such that $\hat{\mathcal{H}}_i = \hat{\mathcal{K}}_j$, and we are done. $\square$

LECTURE 3

Rank Rigidity

In this lecture, we will see that under mild conditions, CAT(0) cube complexes are products of irreducible complexes that are either "line-like" or exhibit hyperbolic-like behavior. In the course of sketching a proof of this theorem, we will discuss some useful features of CAT(0) cube complexes, including the notion of an essential core and the interaction of isometries with hyperplanes.

A hyperbolic isometry of a CAT(0) space is called *rank 1* if an (any) axis for the isometry does not bound a half-flat (an isometrically embedded half Euclidean plane).

Example. Let X be a Gromov hyperbolic CAT(0) space. Then every isometry is rank 1.

The reason here is simply that there are no half-flats in X: was a half flat, there would be large triangles that are not δ-thin, for any $\delta > 0$.

Exercise 3.21. Suppose that X is a product of two infinite locally finite trees. Then no isometry of X is rank 1. Is this still true when we allow locally infinite trees?

Remark. The above generalizes to products of CAT(0) spaces with extendible geodesics.

Exercise 3.22. Let Γ be the pentagon graph and let X be the universal cover of the Salvetti complex associated the the RAAG $A(\Gamma)$. Show that there are rank 1 and non rank 1 isometries of X.

Problem 3.1. Is there a CAT(0) space with a proper cocompact group action and an isometry with an axis that bounds a half-flat, but for which the axis is not a bounded distance from an isometrically embedded flat.

It turns out that for symmetric spaces of higher rank and Euclidean buildings, there are no rank 1 elements. Together with what we saw above about products, Ballmann and Buyalo [**4**] were lead to the following conjecture.

Rank Rigidity Conjecture. *Let G act properly cocompactly on a* CAT(0) *space X with extendible geodesics. Then one of the following three possibilities holds.*

(1) G contains a rank 1 element
(2) X is a non-trivial product
(3) X is a higher rank symmetric space or a Euclidean building

This conjecture was originally proven in the setting of non-positively curved manifolds by Ballmann [**2**] and was subsequently generalized by others (see [**3**] for further discussion.) The goal of this lecture, is the following [**10**]

Theorem 3.1. *Let G act properly and cocompactly on a* CAT(0) *cube complex X. Then one of the following two possibilites holds.*

1. *G contains a rank 1 element*
2. *X contains a convex invariant subcomplex which splits as a product*

1. Essential cores

Just as when a group acts on a tree, one can often reduce to an action on a minimal invariant subtree by removing edges with valence 1 vertices, one has a similar notion for cube complexes.

Definition 3.1. A hyperplane is said to be *essential* if both of its halfspaces contain points arbitrarily far away from it. The hyperplane is said to be *inessential* if some neighborhood of it is the whole complex. If the hyperplane is neither essential not inessential, we say that it is *half-essential*. This means that only one of the halfspaces it defines contains points arbitrarily far away from it.

Lemma 3.1. *Let X be a cocompact* CAT(0) *cube complex. Then there exists an* $\mathrm{Aut}(X)$*-invariant subcomplex Y such that Y decomposes as a product $Y = Z \times C$, where Z is essential and C is finite.*

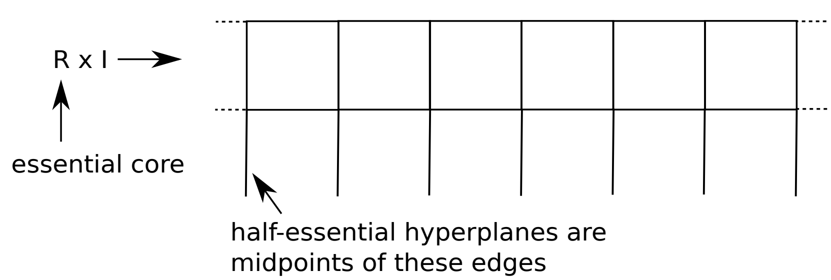

FIGURE 1. A complex whose essential core is the real line.

SKETCH OF PROOF. We consider first the half-essential hyperplanes. Note that if there are half-essential hyperplanes, then there exists one which is *extremal*, meaning that on one side of it, all the vertices are endpoints of edges transverse to the hyperplane. Recall that each hyperplane has a carrier, which is the union of the closed cells meeting it. If a hyperplane is extremal, then its carrier is of the form $\hat{\mathfrak{h}} \times [0,1]$, where one of its boundaries, say $\hat{\mathfrak{h}} \times \{0\}$, is "free" in the sense that every cell meeting it is contained in the carrier. Thus we can remove $\hat{\mathfrak{h}} \times [0,1)$ from the complex and remain with a connected subcomplex. If we do this to all of the extremal hyerplanes at once, then one obtains a new complex X', invariant under $\mathrm{Aut}(X)$, with fewer orbits of hyperplanes. It is easy to check that X' is indeed CAT(0) and convex. In fact, there is a deformation retraction from X to X'. We then continue this process, eliminating at each stage orbits of half-essential hyperplanes. Since there are finitely many orbits of hyperplanes, we end up with an invariant CAT(0) subcomplex Y with no half-essential hyperplanes.

Now we consider Y and observe that every essential hyperplane crosses every inessential hyerplane. For suppose that $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ are disjoint hyperplane with $\hat{\mathfrak{h}}$ essential. Then up to renaming the halfspaces associate to $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$, we have that $\mathfrak{h} \subset \mathfrak{k}$. But $\hat{\mathfrak{h}}$ is essential, which means there are points in $\mathfrak{h}$ arbitrarily far from $\hat{\mathfrak{h}}$ and these points are then arbitrarily far from $\hat{\mathfrak{k}}$.

Thus, by Corollary 2.2, we have that Y decomposes as a product $Y = Z \times C$, where the hyperplanes associated to Z are the essential hyperplanes and the hyperplanes associated to C are the inessential ones. Since C has only inessential hyperplanes, it follows that it is finite.

We already know that Y is $\mathrm{Aut}(X)$-invariant. Since the notions of essential and inessential are $\mathrm{Aut}(X)$-invariant, so is the decomposition $Y = Z \times C$. □

The conclusion of this lemma is that whenever we have an action of a proper, cocompact action of a group on a CAT(0) cube complex, we can pass to a proper cocompact action on an essential one, by passing to the complex Z in the lemma.

One example of an essential CAT(0) cube complex is one with extendible geodesics. In order to simplify parts of the proof, from here on in, we will restrict to this situation:

STANDING ASSUMPTION. For the rest of this lecture, we will assume that our CAT(0) cube complex has extendible geodesics.

2. Skewering

In this section, we examine the ways that that automorphisms act behave with respect to hyperplanes.

Definition 3.2. An automorphism $g \in \mathrm{Aut}(X)$ is said to *skewer* a half-space $\mathfrak{h}$ if $g\mathfrak{h} \subset \mathfrak{h}$.

The term "skewer" becomes clear from the following exercise.

Exercise 3.23. Let $g \in \mathrm{Aut}(X)$ and $\mathfrak{h}$ a half-space.

(1) If g skewers $\mathfrak{h}$, then g is hyperbolic and any axis for g crosses $\hat{\mathfrak{h}}$.
(2) If g is hyperbolic and the axis of g crosses h, then for some $n \in \mathbf{Z}$, we have that g^n skewers $\mathfrak{h}$.

Note that g skewers $\mathfrak{h}$ if and only if g^{-1} skewers $\mathfrak{h}^*$, so that it makes sense to speak of g skewering the hyperplane $\hat{\mathfrak{h}}$ whenever g skewers $\mathfrak{h}$ or $\mathfrak{h}^*$.

3. Single Skewering

We start by showing that hyperplanes are skewered.

Proposition 3.1 (Single Skewering Lemma)**.** *Let G act cocompactly on X, then ever hyperplane is skewered by some element of G.*

Proof. Consider a hyperplane $\hat{\mathfrak{h}} \in X$. Let $X(G, \hat{\mathfrak{h}})$ be the orbit quotient and $p : X \to X(G, \hat{\mathfrak{h}})$ the G-equivariant quotient map. There are two possibilities depending on whether the diameter of $X(G, \hat{\mathfrak{h}})$ is bounded or not.

If the diameter of $X(G, \hat{\mathfrak{h}})$ is unbounded, then there exists a 1-skeleton geodesic α of length larger than the ramsey number $R(dim(X(G, \hat{\mathfrak{h}}), 3)$. This means that the are three disjoint hyperplanes crossing α. Since these are all in the same orbit of $\hat{\mathfrak{k}} = p(\hat{\mathfrak{h}})$, we can label then $a\hat{\mathfrak{k}}, b\hat{\mathfrak{k}}, c\hat{\mathfrak{k}}$. Now if we put transverse orientations on these three hyperplanes, we see that two of them must be oriented in the same direction along α. This means that one of the elements ab^{-1}, bc^{-1}, ac^{-1} skewers one of the three hyperplanes.

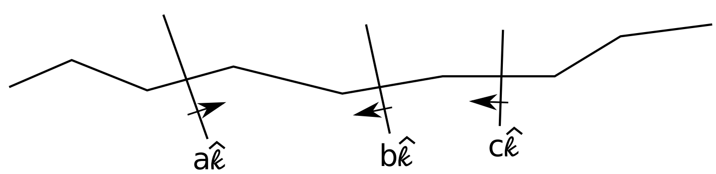

FIGURE 2. The element bc^{-1} carries $c\mathfrak{k}$ into $b\mathfrak{k}$ and hence skewers $c\hat{\mathfrak{k}}$.

Since the action on the hyperplanes of $X(G, \hat{\mathfrak{h}})$ is transitive, it follows that the hyperplane $\hat{\mathfrak{k}}$ is also skewered in $X(G, \hat{\mathfrak{h}})$. Lifting the action to X, we see that the same element that skewers $\hat{\mathfrak{k}}$ in the action of G on $X(G, \hat{\mathfrak{h}})$ skewers $\hat{\mathfrak{h}}$ in the action of G on X.

The second case is that the diameter of $X(G, \hat{\mathfrak{h}})$ is bounded. This means that there exists a fixed point for the action of G on $X(G, \hat{\mathfrak{h}})$. After perhaps passing to a finite index subgroup of G, we may assume that there is a fixed vertex $v \in X(G, \hat{\mathfrak{h}})$ for the action. Now lifting to X, we see that the collection of vertices $p^{-1}(v)$ are stabilized by the action of G on X. All of these vertices lie to one side of some hyperplane in X. This contradicts the fact that X is essential and the action of G on X is cocompact. $\square$

4. Flipping

An alternative to a hyperbolic element skewering a hyperplane is the following.

Definition 3.3. A hyperbolic isometry g of X is said to *flip* a half-space $\mathfrak{h}$ if $g\mathfrak{h} \subset \mathfrak{h}^*$.

Exercise 3.24 (Trichotomy)**.** Let g be a hyperbolic isometry of X and let $\mathfrak{h}$ be a halfspace. Then one of the following holds.

(1) Some power of g skewers $\mathfrak{h}$.
(2) Some power of g flips $\mathfrak{h}$ or $\mathfrak{h}^*$.
(3) Some power of g stabilizes $\hat{\mathfrak{h}}$.

The first of the above possiblities is when the axis of g meets $\hat{\mathfrak{h}}$ and the last is when the axis for g lies in a bounded neighborhood of $\hat{\mathfrak{h}}$.

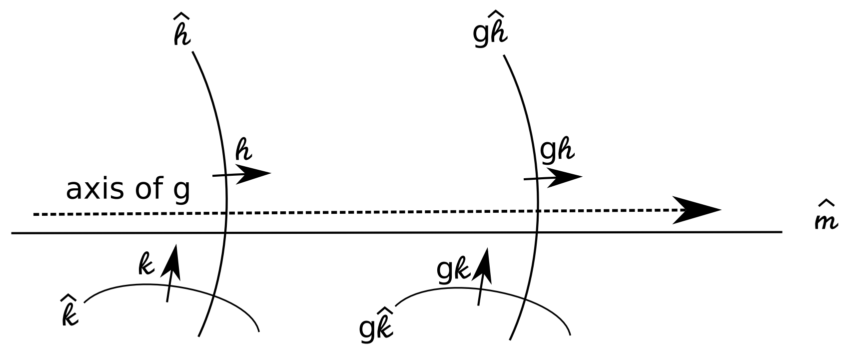

FIGURE 3. The three possibilities. g skewers $\hat{\mathfrak{h}}$, flips $\hat{\mathfrak{k}}^*$ and stabilizes $\hat{\mathfrak{m}}$.

We say that a halfspace is *unflippable* if there does not exist any $g \in G$ flipping it. A key lemma is then the following.

Lemma 3.2 (Flipping Lemma)**.** *Let G act on X properly and cocompactly. Let $\mathfrak{h}$ be an unflippable hyperplane. Then X decomposes as a product $X = Y \times \mathbf{R}$ and $\hat{\mathfrak{h}}$ appears as the preimage of a point in $\mathbf{R}$ under the natural projection $X \to \mathbf{R}$.*

Before sketching a proof of this lemma, we recall an elementary lemma, whose proof we leave as an exercise.

Exercise 3.25 (Endometry Lemma)**.** Let X be a proper metric space with a cocompact isometry group. Let $f : X \to X$ be an isometric map. Then f is bijective.

Hint. Injectivity is obvious. For surjectivity, use the fact that for every $R > 0$ and $\epsilon > 0$, there exists a number N such that any ball of radius R, the size of an ϵ-separated collection of points in the ball is at most N.

SKETCH OF PROOF OF FLIPPING LEMMA. The idea will be to break up the hyperplanes of X into those that intersect $\hat{\mathfrak{h}}$ and those that do not. It will turn out that this is a transverse decomposition of the collection of hyperplanes. We will thus aim to prove the following.

Claim. Show that for every hyperplane $\hat{\mathfrak{k}}$ disoint from $\hat{\mathfrak{h}}$, the hyperplanes that cross $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ are the same. We do this in several steps.

Step 1. When $\hat{\mathfrak{k}} \subset \mathfrak{h}^*$, show that every hyperplane crossing $\hat{\mathfrak{k}}$ crosses $\hat{\mathfrak{h}}$. Suppose that $\hat{\mathfrak{m}}$ is a hyperplane meeting $\hat{\mathfrak{k}}$ and is disjoint from $\hat{\mathfrak{h}}$. The hyperplane $\hat{\mathfrak{k}}$ is a CAT(0) cube complex and itself has extendible geodesics. By the Single Skewering Lemma 3.1, there exists an element $g \in \mathrm{Stab}(\hat{\mathfrak{k}})$ which skewers the hyperplane $\hat{\mathfrak{k}} \cap \hat{\mathfrak{m}}$. Now the element g has an axis which is disjoint from $\hat{\mathfrak{h}}$, so it cannot skewer $\hat{\mathfrak{h}}$. Moreover, since g skewers a hyperplane which is disoint from $\hat{\mathfrak{h}}$, no power of g can stablizer $\hat{\mathfrak{h}}$. So by Trichotomy, it follows that some power of g flips $\mathfrak{h}$, a contradiction.

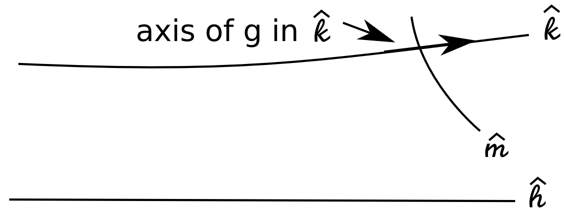

FIGURE 4. A hyperplane meeting $\hat{\mathfrak{k}}$ but not $\hat{\mathfrak{h}}$.

Step. 2. Step 1 yields an embedding of $\hat{\mathfrak{k}}$ into $\hat{\mathfrak{h}}$ as follows. We wish to define a map $f : \hat{\mathfrak{k}} \to \hat{\mathfrak{h}}$.

We first define f on the vertices of $\hat{\mathfrak{k}}$. Let v be a vertex of $\hat{\mathfrak{k}}$. For every hyperplane $\hat{\mathfrak{m}}$ meeting $\hat{\mathfrak{h}}$, we need to chose a halfspace bounded $\hat{\mathfrak{m}}$. Simply choose the side that contains v. This gives us an ultrafilter on the halfspaces of $\hat{\mathfrak{h}}$. It is easy to see that this satisfies DCC, so we have a vertex $f(v)$.

The conclusion of Step 1 now tells us that the map f is injective on vertices. For if v and w are two vertices of $\hat{\mathfrak{k}}$, separated by a hyperplane $\hat{\mathfrak{m}}$, then $\hat{\mathfrak{m}}$ also intersects $\hat{\mathfrak{h}}$ and therefore separates $f(v)$ and $f(w)$.

Finally, one checks that f extends to the cubes of $\hat{\mathfrak{k}}$. We leave this to the reader. This yields the desired embedding.

Step 3. For every translated hyperplane $g\hat{\mathfrak{h}} \subset \mathfrak{h}^*$, a hyperplane crosses $\hat{\mathfrak{h}}$ if and only if crosses $g\hat{\mathfrak{h}}$.

We already know by Step 1 that every hyperplane meeting $g\hat{\mathfrak{h}}$ also meets $\hat{\mathfrak{h}}$. Step 2 gives us an isometric embedding of $g\hat{\mathfrak{h}}$ into $\hat{\mathfrak{h}}$. But now the Endometry Lemma, tells us that this embedding is surjective. It then easily follows that every hyperplane meeting $\hat{\mathfrak{h}}$ also meets $g\hat{\mathfrak{h}}$. For suppose there was a hyperplane $\hat{\mathfrak{m}}$ meeting $\hat{\mathfrak{h}}$ and not meeting $g\hat{\mathfrak{h}}$. Then any vertex in $\hat{\mathfrak{h}}$ separated from $g\hat{\mathfrak{h}}$ by $\hat{\mathfrak{m}}$ whould not be in the image of the map f defined in Step 2.

Step 4. For every hyperplane $\hat{\mathfrak{k}} \subset \mathfrak{h}^*$, a hyperplane crosses $\hat{\mathfrak{h}}$ if and only if it crosses $\hat{\mathfrak{k}}$.

Let $\mathfrak{k}$ denote the halfspace of $\hat{\mathfrak{k}}$ which contains the hyperplane $\hat{\mathfrak{h}}$. We will show that there exists a translate of $\hat{\mathfrak{h}}$ lying in $\mathfrak{k}^*$. By the Single Skewering Lemma, there exists $g \in G$ such that $g\mathfrak{k}^* \subset \mathfrak{k}^*$. Since higher powers of g move points deeper and deeper into $\mathfrak{k}^*$, there exists n such that $g^n\hat{\mathfrak{h}} \cap \mathfrak{k}^* \neq \emptyset$. We can also choose n such that $g^n\hat{\mathfrak{h}} \cap \hat{\mathfrak{h}} = \emptyset$. By Step 1, since $g^n\hat{\mathfrak{h}}$ is disjoint from $\hat{\mathfrak{h}}$, it must also be disjoint from $\hat{\mathfrak{k}}$ and therefore must be contained in $\mathfrak{k}^*$. Now by Step 3, the hyperplanes that cross $\hat{\mathfrak{h}}$ are precisely those that cross $g\hat{\mathfrak{h}}$, so that all the hyperplanes crossing $\hat{\mathfrak{h}}$ must cross $\hat{\mathfrak{k}}$ as well.

Step 5. For every hyperplane $\hat{\mathfrak{k}} \subset \mathfrak{h}$, a hyperplane crosses $\hat{\mathfrak{h}}$ if and only if it crosse $\hat{\mathfrak{k}}$.

We leave this to the reader. We proceed as in Steps 1-4, but one needs to take care in proving the last step in this case. This completes the proof of the claim.

We now break up the collection of hyperplanes $\hat{\mathcal{H}}$ of X into a disjoint union $\hat{\mathcal{H}} = \hat{\mathcal{H}}_{\|} \cup \hat{\mathcal{H}}_{\perp}$, where

$$\hat{\mathcal{H}}_{\|} = \{\hat{\mathfrak{k}} \in \hat{\mathcal{H}} | \hat{\mathfrak{k}} \cap \hat{\mathfrak{h}} = \emptyset\}$$

$$\hat{\mathcal{H}}_\perp = \{\hat{\mathfrak{k}} \in \hat{\mathcal{H}} | \hat{\mathfrak{k}} \cap \hat{\mathfrak{h}} \neq \emptyset\}$$

The claim tells us that every hyperplane in $\hat{\mathcal{H}}_\|$ intersects every hyperplane in $\hat{\mathcal{H}}_\perp$ and this gives us a product decomposition $X = X_\| \times X_\perp$.

Observe now that since no two hyperplanes in $\hat{\mathcal{H}}_\|$ can intersect, the space $X_\|$ is a tree T. There is a copy of T which appears in X as a maximal intersection of hyperplanes meeting $\hat{\mathfrak{h}}$. The stabilizer of T acts properly and cocompactly on T with an unflippable hyperplane. We leave it to the reader to check that this means that T is a line. □

5. Double Skewering

We now seek another property which tells us more about how automorphisms of X interact with hyperplanes.

Definition 3.4. An automorphism $g \in \mathrm{Aut}(X)$ is said to *double skewer* two nested halfspaces $\mathfrak{h} \subset \mathfrak{k}$ if $g\mathfrak{k} \subset \mathfrak{h}$.

Proposition 3.2 (Double Skewering Lemma)**.** *Let G act on X cocompactly, then for any two nested halfspaces $\mathfrak{h} \subset \mathfrak{k}$ there exists $g \in G$ double skewering $\mathfrak{h}$ and k.*

Proof. We will employ the Flipping Lemma. There are two cases. First, suppose that either $\mathfrak{h}$,$\mathfrak{h}^*$,$\mathfrak{k}$ or $\mathfrak{k}^*$ is unflippable. Then by the Flipping Lemma, we have that X decomposes as a product $X = Y \times \mathbf{R}$ and that the hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ are preimages of points $p, q \in \mathbf{R}$. Choose g skewering $\mathfrak{k}$ so that $g\mathfrak{k} \subset \mathfrak{k}$. Since the hyperplanes that intersect $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ are the same, it follows that $g\hat{\mathfrak{k}} \cap \hat{\mathfrak{h}} = \emptyset$. Now there are only finitely many hyperplanes between $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$, it follows that for a sufficiently large power of g, we have $g\mathfrak{k} \subset \mathfrak{h}$, as required.

Otherwise, all the halfspaces associated to $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ are flippable. We do two flips to obtain double skewering (see Figure 5 below).

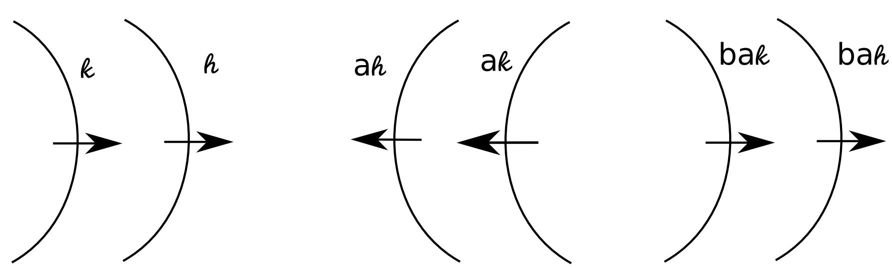

Figure 5. Double flipping leads to double skewering.

More precisely, since $\mathfrak{h}$ is flippable, there exists $a \in G$ such that $a\mathfrak{h}^* \subset \mathfrak{h}$. Since $\mathfrak{h} \subset \mathfrak{k}$, we have that $\mathfrak{k}^* \subset \mathfrak{h}^*$, which implies that $a\mathfrak{k}^* \subset a\mathfrak{h}^* \subset \mathfrak{h}$. Now since $\mathfrak{k}$ is flippable, so is $a\mathfrak{k}$. This means there exists $b \in G$ such that $ba\mathfrak{k} \subset a\mathfrak{k}^*$. Since $\mathfrak{h} \subset \mathfrak{k}$, we have $ba\mathfrak{h} \subset ba\mathfrak{k}$. So we then obtain:

$$ba\mathfrak{h} \subset ba\mathfrak{k} \subset a\mathfrak{k}^* \subset a\mathfrak{h}^* \subset \mathfrak{h}$$

as required. □

6. Hyperplanes in sectors

Consider n intersecting hyperplanes. They divide X into 2^n regions which we call *sectors*. In this section, we will only consider sectors determined by two intersecting hyperplanes. If X were, say a product of two trees, then none of these regions would contain hyperplanes since every hyperplane intersects one of the original pair. However, if one images a CAT(0) cube square complex obtained by taking the universal cover of the squaring of a hyperbolic surface described in Example 3. In this cube complex, the hyperplanes are quasigeodesics in a hyperbolic space. Thus, it is easy to see that each sector contains a hyperplane (in fact, infinitely many). The following proposition tells us that this is a general phenomenon.

Proposition 3.3 (Sector Lemma)**.** *Let X be an irreducible cocompact* CAT(0) *cube complex with extendible geodesics that is not the real line. Let $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ be two intersecting hyperplanes in X. Then each of the four sectors defined by X contains a hyperplane.*

Sketch of proof. First, we show that there exists some hyperplane disjoint from $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$. Assume that this is not the case. We now seek a decomposition of the collection of hyperplanes into disjoint transverse subsets, which will contradict the fact that X is not a product. Recall the $\hat{\mathcal{H}}$ denotes the collection of hyperplanes of X. We now focus on the following collections of hyperplanes.

$\hat{\mathcal{H}}_{\mathfrak{h}} = \{\text{hyperplanes disjoint from } \hat{\mathfrak{h}}\}$
$\hat{\mathcal{H}}_{\mathfrak{k}} = \{\text{hyperplanes disjoint from } \hat{\mathfrak{k}}\}$
$\hat{\mathcal{H}}'_{\mathfrak{h}} = \{\text{hyperplanes disjoint from some hyperplane in } \hat{\mathcal{H}}_{\mathfrak{h}}\}$
$\hat{\mathcal{H}}'_{\mathfrak{k}} = \{\text{hyperplanes disjoint from some hyperplane in } \hat{\mathcal{H}}_{\mathfrak{k}}\}$
$\hat{\mathcal{R}} = \hat{\mathcal{H}} - (\hat{\mathcal{H}}' \cup \hat{\mathcal{K}}')$

We leave it to the reader to check that $\hat{\mathcal{H}} = \hat{\mathcal{H}}'_{\mathfrak{h}} \cup \hat{\mathcal{H}}'_{\mathfrak{k}} \cup \hat{\mathcal{R}}$ is a transverse decomposition of $\hat{\mathcal{H}}$.

Thus there exists some hyperplane disjoint from $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$. Say that $\hat{\mathfrak{m}}$ is such a hyperplane, so that $\mathfrak{m} \subset \mathfrak{h} \cap \mathfrak{k}$. We now need to show all the other sectors contain hyperplanes as well. The double skewering lemma applied to the hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{m}}$ gives us an element $g \in G$ such that $g\mathfrak{h} \subset \mathfrak{m}$. By applying a sufficiently high power of g, we have that $g^n\mathfrak{h} \subset \mathfrak{m}$ and $g^n\hat{\mathfrak{k}} \cap \hat{\mathfrak{k}} = \emptyset$. It follows that $\hat{\mathfrak{k}} \subset g^n\mathfrak{h}^* \cap g^n\mathfrak{k}$ or $\hat{\mathfrak{k}} \subset g^n\mathfrak{h}^* \cap g^n\mathfrak{k}^*$. In either case, by applying g^{-n} we conclude that there exists a hyperplane in one of the two sectors $\mathfrak{h}^* \cap \mathfrak{k}$ or $\mathfrak{h}^* \cap \mathfrak{k}^*$.

By reversing the roles of $\hat{\mathfrak{k}}$ and $\hat{\mathfrak{m}}$ we get that there exists a hyperplane in one of the two sectors $\mathfrak{k}^* \cap \mathfrak{h}$ or $\mathfrak{k}^* \cap \mathfrak{h}^*$. This gives that there are hyperplanes in diagonally opposite sectors. Without loss of generality, let us assume that there is a hyperplane in $\mathfrak{h} \cap \mathfrak{k}$ and $\mathfrak{h}^* \cap \mathfrak{k}^*$.

We now consider the hyperplane $\hat{\mathfrak{h}}_1 = \hat{\mathfrak{k}} \cap \hat{\mathfrak{h}}$ as a hyperplane in $\hat{\mathfrak{h}}$. If both of the halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_1^*$ are flippable as halfspaces in $\hat{\mathfrak{h}}$, then we could these flipping elements and obtain hyperplanes in the remaining two sectors $\mathfrak{h} \cap \mathfrak{k}^*$ and $\mathfrak{h}^* \cap \mathfrak{k}$. So we can assume that one of the halfspaces $\mathfrak{h}_1$ or $\mathfrak{h}_1^*$ in $\hat{\mathfrak{h}}$ is unflippable. Similarly, one of the hyperplanes associated to $\hat{\mathfrak{h}} \cap \hat{\mathfrak{k}}$ in $\hat{\mathfrak{k}}$ is unflippable.

We now find a group element $g \in \mathrm{Stab}(\hat{\mathfrak{h}})$ skewering $\hat{\mathfrak{h}} \cap \hat{\mathfrak{k}}$. Let $\hat{\mathfrak{m}}$ be the hyperplane containing in $\mathfrak{h} \cap \mathfrak{k}$. We leave it as an exercise to the reader to show that some power of g carries $\hat{\mathfrak{m}}$ into $\mathfrak{h} \cap \mathfrak{k}^*$. Similarly some power of an element in $\mathrm{Stab}(\hat{\mathfrak{k}})$ carries $\hat{\mathfrak{m}}$ into $\mathfrak{h}^* \cap \mathfrak{k}$. $\square$

Exercise 3.26 (The Tits Alternative)**.** Use Proposition 3.3 to prove every group which acts cocompactly on an irreducible CAT(0) cube complex with extendible geodesics has a free subgroup of rank 2, unless the complex is a real line.

Hint. Suppose that X is as in Proposition 3.3, then we obtain a pattern of four disjoint halfspaces $\mathfrak{h}_1, \mathfrak{h}_2, \mathfrak{h}_3, \mathfrak{h}_4$ such that $\mathfrak{h}_i \subset \mathfrak{h}_j^*$, for each $i \neq j$. Play ping-pong.

7. Proving rank rigidity

We sketch now how the elements in the previous sectiongs provide a proof of rank rigidity. Recall that we are assuming that X is an irreducible CAT(0) cube complex with extendible geodesics and that G is acting properly and cocompactly on X. We wish to show that there exist rank 1 elements in G. So we assume that no element is rank 1.

In the previous section, we considered sectors which were the intersection of two halfspaces. Consider a maximal collection of intersecting hyperplanes which contain hyperplanes in "diagonally opposite" sectors, More precisely, let $\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n$ be a maximal collection of hyperplanes such that $\bigcap_i \hat{\mathfrak{h}}_i \neq \emptyset$ and the sectors $\bigcap_i \mathfrak{h}_i$ and $\bigcap_i \mathfrak{h}_i^*$ contain hyperplanes. We know such collections exist, when $n = 1$. We will show that if X does not have rank 1 elements, there always exists a larger collection of hyperplanes with these properties. Since the dimension is bounded, this is a contradiction.

Let $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ be the hyperplanes contained in the sectors: $\mathfrak{h}_i \in \bigcap_i \mathfrak{h}_i$ and $\mathfrak{k}_i^* \bigcap \mathfrak{h}_i^*$. The goal will now be to find a hyperplane $\hat{\mathfrak{h}}_{n+1}$ that intersects both $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$. If we can do this, then $\hat{\mathfrak{h}}_n$ will intersect the hyperplanes $\hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_n$ and the Sector Lemma applied to $\hat{\mathfrak{h}}_{n+1}, \hat{\mathfrak{h}}$ and $\hat{\mathfrak{h}}_{n+1}$ and $\hat{\mathfrak{k}}$ will tell us that there exist hyperplanes in diagonally opposite sectors of the collection $\{\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_{n+1}\}$, and the proof will be complete.

By the Double Skewering Lemma, there exists $g \in G$ such that $g\mathfrak{k} \subset \mathfrak{h}$. Since g is not rank 1, there exists a half-flat F bounding an axis ℓ for g. The intersection of F with the hyperplanes $\hat{\mathfrak{h}}_i, \hat{\mathfrak{k}}$ and $g\hat{\mathfrak{k}}$ is a collection of rays meeting ℓ in points. We also consider the intersection of ℓ with $g^{-1}\hat{\mathfrak{h}}_i$, and $g^{-1}\hat{\mathfrak{k}}$, as in Figure 6. We observe that by discreteness, there exist finitely many points of intersection of ℓ and the hyperplanes of X between any two given points of ℓ.

Let R be the ray of intersection $R = \hat{\mathfrak{k}} \cap F$. Since there infinitely many hyperplanes crossed by R and there are finitely many hyperplanes crossing ℓ between $\mathfrak{g}^{-1}\hat{\mathfrak{k}} \cap \ell$ and $g\hat{\mathfrak{k}} \cap \ell$, there exists some hyperplane $\hat{\mathfrak{m}}$ intersecting R that which does not intersect ℓ between $\mathfrak{g}^{-1}\hat{\mathfrak{k}} \cap \ell$ and $g\hat{\mathfrak{k}} \cap \ell$. So either $\hat{\mathfrak{m}}$ or $g\hat{\mathfrak{m}}$ is the required hyperplane $\hat{\mathfrak{h}}_{n+1}$.

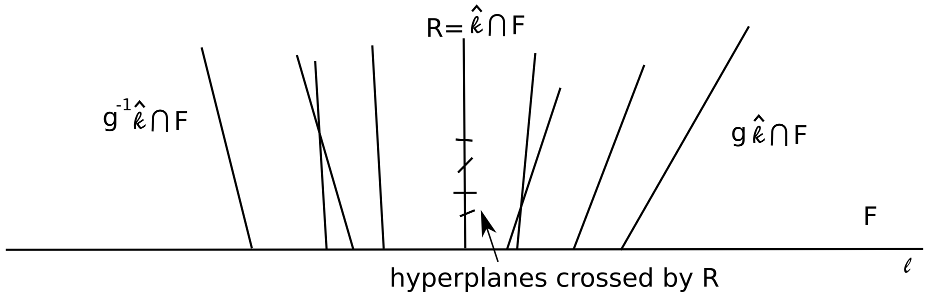

FIGURE 6. There are infinitely many hyperplanes crossing R. One of these must cross $g^{-1}\hat{\mathfrak{k}}$ or $g\hat{\mathfrak{k}}$.

LECTURE 4

Special cube complexes

In this lecture, we give a very brief introduction to special cube complexes and the notion of canonical completion and retraction. This should give you an idea of why special cube complexes have anything to do with subgroup separability. The material in this lecture follows closely some of the material in [**24**].

1. Subgroup Separability

We first recall some basic notions regarding subgroup separability

Definition 4.1. Let G be a group and $H < G$. We say that H is *separable* if for every $g \in G - H$, there exists a finite index subgroup $K < G$ such that $H < K$ and $g \notin K$. The group G is said to be *residually finite* if the trivial subgroup is separable.

Exercise 4.27. Show that $H < G$ is separable if for every $g \in G - H$, there exists a homomorphism to a finite group $\phi : G \to F$, such that $\phi(g) \notin \phi(H)$.

Recall that the *profinite* topology on G is the topology whose basic open sets are the cosets of finite index subgroups of G. If you have not seen the profinite topology before, you should check that it is indeed a topology We then have the following exercise.

Exercise 4.28. A subgroup $H < G$ is separable if and only if it is closed in the profinite topology on G.

Recall that a *retraction* $\phi : G \to H$ is simply a homomorphism which is the identity on H. We say that H is a *retract* of G.

Exercise 4.29. Let G be a residually finite group, and let $H < G$ be a retract. Then H is separable.

Hint. A retract of a Hausdorff space is closed.

2. Warmup - Stallings proof of Marshall Hall's Theorem

Marshall Hall [**25**] proved back in 1949 that every finitely generated subgroup of a finitely generated free group is a virtually a free factor. John Stallings came up with a nice, very elementary graph-theoretic proof of this fact [**39**]. Roughly speaking, the idea is to represent the subgroup $H < G$ as an immersion of graphs $\Delta \to \Gamma$ where Γ has fundamental group G. One then seeks a finite cover in which Δ lifts to an embedding. This then tells you that H is a virtual retract and a virtual free factor. By Exercise 4.28, we then have that H is separable. We review this a bit more closely.

Theorem 4.1. *Every finitely generated subgroup of a finitely generated free group is a virtual retract.*

Proof. This will be a proof by example. Consider the free group $G = \langle a, b\rangle$ and consider the subgroup $H = \langle ab^2a^{-1}, abab^2\rangle$. First we represent $H < G$ as a map between bouquets of circles $\Delta \to \Gamma$.

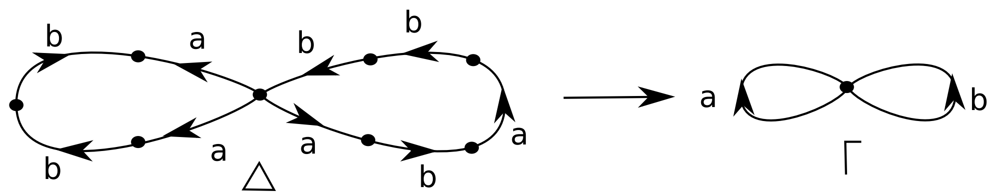

Figure 1. Representing a subgroup by a map between bouquets of circles.

On the level of fundamental group, this map represents our subgroup H. Note that the map is not necessarily an immersion (i.e. a local embedding.) In this case, for example there are is more than one edge on the left graph labeled a and pointing out of the central vertex. To remedy this, we "fold" two such edges in the graph on the left as shown in the diagram below to obtain a new graph with fewer edges. The map still represents the same subgroup, and the map is closer to an immersion. We continue to fold, reducing the number of edges each time, until we obtain an immersion.

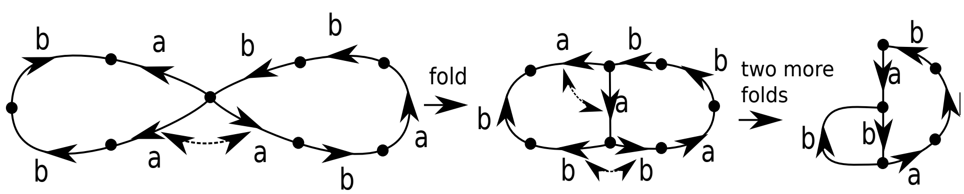

Figure 2. Folding the source graph.

Note that on the graph on the right there is no more folding that can be done, so that we now have an immersion of graphs representing our subgroup. We now replace Δ by this new folded graph and we have a new map $\Delta \to \Gamma$ which is an immersion.

Now that we have an immersion, the next thing to do is to complete. There are many ways to do this. We describe what Haglund and Wise call "canonical completion". The idea is to add edges to the graph Δ until the map becomes a covering space. Consider the a-loop in Γ. The fact that the map $\Delta \to \Gamma$ is an immersion tells us that each component of the preimage of the a-loop in Γ is either a cycle of a's, an arc of a's, or a single vertex. If it is a loop of a's we do nothing. If it is a loop of a's we add an edge labeled with an a to complete the arc of a's to a loop of a's. Finally, if it is a vertex, then it means that we have a vertex with only b-edges adjacent to it. In this case, we simply attach an a-loop at that vertex. See Figure 2.

We then do the same for the b's. It is now easy to check that the resulting map from the completed graph $\overline{\Gamma}$ to Γ is a covering space.

Now note that not only do we have a finite covering space $\overline{\Gamma}$ of the original bouquet of circles Γ, The original map $\Delta \to \Gamma$ factors through an embedding $\Delta \to \overline{\Gamma}$ and there is natural retraction map $\overline{\Gamma} \to \Delta$. Simply map each added a-edge to the arc of a-edges it completed and each a-loop to the vertex it is attached to; and do the same for the b's.

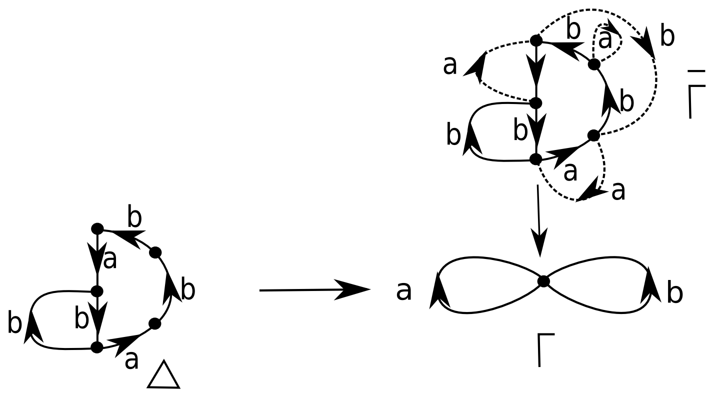

FIGURE 3. Completing an immersion of graphs. The dotted edges are the edges we add to complete.

We thus obtain that the subgroup H is a retract. It is also easy to see from this picture that H is virtually a free factor.

$\square$

Exercise 4.30. Use the proof of the above theorem to show that a finitely generated free group is residually finite.

Since the free group is residually finite, Exercise 4.28 and Theorem 4.1 now tell us that every finitely generated subgroup of a finitely generated free group is separable.

3. Special Cube Complexes

One of Wise's main goals was to seek a more general setting in which the technique for of the previous section for graphs can be made to work. This led to the notion of special cube complexes which we now describe.

First we describe "osculation". Two hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ in an NPC complex are said to *osculate* at a vertex v if there exists edges e and f with endpoint v such that e and f are not on the boundary of a square and e and f are transverse to the hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$. If $\hat{\mathfrak{h}} = \hat{\mathfrak{k}}$ we say that $\hat{\mathfrak{h}}$ *self-osculates*. If $\hat{\mathfrak{h}} \neq \hat{\mathfrak{k}}$ and $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ also intersect, we say that $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ *interosculate*.

A hyperplane is said to be *2-sided* if it separates its carrier. This is the same as saying that the carrier is a product of the hyperplane with an interval.

Figure 4. Self osculation and interosculation.

Definition 4.2. Let X be an NPC cube complex. We say that X is *special* if every hyperplane is embedded, 2-sided and does not self-osculate, and no two hyperplanes interosculate.

A group is said to be *special* if it is the fundamental group of a compact special cube complex. The group is said to be *virtually special* if it has a finite index subgroup which is special.

Exercise 4.31. Examine the examples of NPC cube complexes discussed in Lecture 1 and decide which is special.

Exercise 4.32. Let $X \to Y$ be a locally isometric map between NPC cube complexes and suppose that Y is special. Then X is special.

Recall now the example discussed in Lecture 1, Example 5, namely the Salvetti complex associated to a right-angled Artin group. We call such a complex a *RAAG complex* for short. It is easy to see that these cube complexes are indeed special (the reader should check this). The reason that these examples are so important is because of the following proposition.

Proposition 4.1. *Let X be a compact NPC cube complex. Then X is special if and only if there exists a locally isometric embedding to a RAAG complex. In particular, the fundamental group of a compact special cube complex is a subgroup of a right-angled Artin group.*

Sketch of proof. If there exists a locally isometric embedding $X \to R$, then we know by Exercise 4.31 that since R is special, so is X.

On the other hand, suppose that X is special. We let Γ be the graph whose vertices correspond to hyperplanes of X and where two vertices are joined by an edge if and only if the corresponding hyperplanes in X intersect. We consider the RAAG-complex $R = R(\Gamma)$. We now construct a map $X^{(1)} \to R$ in the natural way: vertices of X get mapped to the unique vertex of R, edges of X get sent to edges that cross the corresponding hyperplane in R. Now one checks that this map can be extended over the all cubes of X to a cubical map and that the map is a local isometry. □

We wil focus on compact special cube complexes, although it is possible to discuss natters in the context of complexes with finitely many hyperplanes. A final remark is that by a result proved independentally by Davis-Januszkiewicz [**14**] and

Hsu-Wise [**26**], RAAGs are linear. Thus, we know immediately that every virtually special group is residually finite. Therefore, in order to show that a subgroup of virtually special group is separable, we just need to show that it is a virtual retract.

4. Canonical Completion and Retraction

We now wish to generalize the the canonical completion and retraction construction from the world of graphs to special cube complexes. First, we consider an immersion $X \to R$ from an NPC complex X to a RAAG complex R. We will build a covering space of R which we call $C(X, R)$. An instructive example is seen in Figure 5.

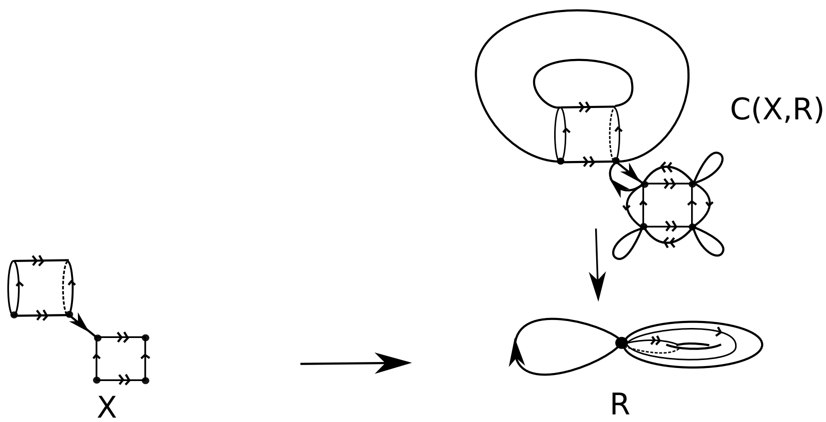

FIGURE 5. Canonical completion and retraction for a map to a RAAG complex.

First we focus on the 1-skeleton of X as a map to the 1-skeleton of R, which is a bouquet of circles. We canonical complete as in the case of graphs. This is the 1-skeleton of $C(X, R)$. Now one checks that one can add squares and higher dimensional cubes wherever they "should be". For example, on the top left side of X, there is an annulus. We added an arc with a double arrow to complete the original double arrow arc to a loop. Now one simply glues in another annulus so that the original annulus is completed to a torus. The original square on the lower left of X has four extra arcs attached and four extra squares. That this all works requires a bit of thought and checking some cases. We leave it to the reader.

One has then built a covering space of X and there is indeed a retraction $C(X, R) \to X$ extending the one described earlier for graphs.

Generally, for a locally isometric immersion between two special cube complexes $X \to Y$, where Y is not necessarily a RAAG complex, one uses a fiber product construction. Given two cubical maps $X \to Y$ and $Z \to Y$ between cube complexes, one can construct a complex denoted $X \otimes_Y Z$ called the *fiber product of X and Z over Y*, which is a subspace of $X \times Y$ consisting of products $\sigma \times \tau$ where σ and τ get mapped to the same cell in Z. If the original maps $X \to Y$ and $Z \to Y$ were covering spaces, this would corresponds to the usual common covering space.

If you are not familiar with this notion, you should first draw some simple examples of fiber products using graphs. Figure 6 shows one such example.

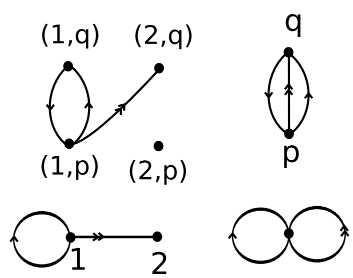

Figure 6. A simple fiber product.

Note that the fiber product need not be connected, but in the situation we will be looking at there will be a natural component to look at.

So now given a local isometric embedding between special cube complexes $X \to Y$, we know there exists a RAAG complex R and a locally isometric embedding $Y \to R$. The composition $X \to Y \to R$ gives a map $X \to R$, and we can form the canonical completion $C(X, R)$ for this map. We also have a natural covering map $C(X, R) \to R$. One then defines $C(X, Y)$ as a fiber product:

$$C(X, Y) \equiv Y \otimes_R C(X, R)$$

We have a retraction which comes from the composition $C(X, Y) \to C(X, R) \to X$, where the second map is the retraction produced before. The canonical completions $C(X, Y)$ becomes difficult to draw for complicated examples (even the one in Figure 5), but Figure 7 displays simple example.

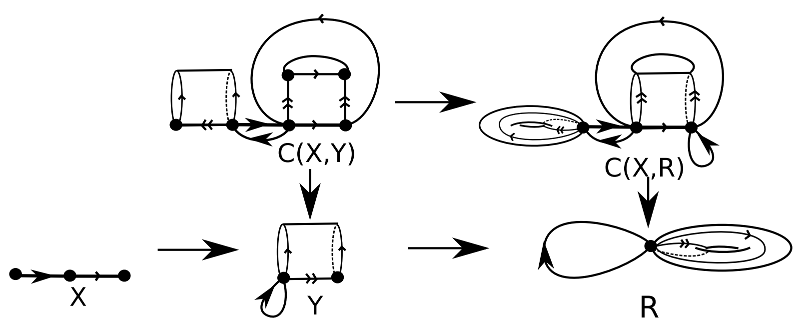

Figure 7. A general completion and retraction.

This particular example is connected, but you should be aware that they need not be. However there is a natural embedding of X in $C(X, Y)$ and one usually focuses on this component.

5. Application: separability of quasiconvex subgroups

In the previous section, we saw that special cube complexes have something to do with subgroup separability. In particular, any subgroup of a special cube complex group which can be represented as an immersion is separable.

A particular application of this is the following result of Haglund and Wise.

Theorem 4.2. *If G is virtually special and Gromov hyperbolic. Then every quasi-convex subgroup is separable.*

We first remark that we may assume that G is itself special. For if we prove that the quasi-convex subgroups of a finite index subgroup of G are separable, then the quasi-convex subgroups of G are separable as well. Secondly, we remark that since G is virtually special, it is linear. To prove the theorem, we first need to discuss the construction of a combinatorial convex core. We give the construction in the form of an exercise. For more details see Haglund [**22**].

Suppose that G is Gromov hyperbolic and the fundamental group of a compact NPC cube complex Y. Then we have G acting on the universal cover X, which is a Gromov hyperbolic CAT(0) cube complex. Now Consider the orbit of some vertex under a quasi-convex subgroup $H < G$. Since H is finitely generated, there exists some neighborhood N of this orbit which is connected. Now using geometry of Gromov hyperbolic spaces, one can prove the following lemma.

Exercise 4.33. There exists some constant C such that if σ is a cube distance at least C from N, one of the hyperplanes meeting σ is disjoint from N.

One now builds a convex hull for H as follows. For each halfspace $\mathfrak{h}$ let $C(\mathfrak{h})$ denote the union of $\mathfrak{h}$ with the carrier of $\hat{\mathfrak{h}}$. This is a convex subcomplex of X. We then set

$$\mathrm{Hull}(H) = \bigcap_{N \subset \mathfrak{h}} C(\mathfrak{h})$$

Some thought and Exercise 4.32 will then tell you that $\mathrm{Hull}(H)$ is contained in some neighborhood of H. In particular this is a convex subcomplex of X on which H acts cocompactly. We call $\mathrm{Hull}(H)$ a combinatorial convex hull for H.

Observe that $\mathrm{Hull}(H)/H$ embedds naturally into X/H and $X/H \to X/G = Y$ is a covering space and hence a locally isometric embedding. We thus have a locally isometric embedding of NPC complexes $\mathrm{Hull}(H)/H \to Y$.

We now apply the canonical completion and retraction construction to this local isometric embedding to obtain a finite covering $\overline{Y}$ of Y and a retraction $\overline{Y} \to \mathrm{Hull}(H)/H$. We thus obtain a retraction $G \to H$ and, so that by Exercise 4.28, H is separable.

6. Hyperbolic cube complexes are virtually special

This lecture was a very cursory introduction to the notion of special cube complexes. Wise, together with Haglund and others, extensively developed the theory of special cube complexes and proved several very deep and difficult theorems about them. For a treatment of much of these theorems the reader should consult Wise's upcoming book [**42**]. Agol [**1**] combined these theorems with an ingenius coloring argument to obtain the following theorem.

Theorem 4.3. *Every hyperbolic group which acts properly and cocompactly on a* CAT(0) *cube complex is virtually special.*

This is a startling theorem which has far reaching implications. For example, it settles the long-standing virtual Haken conjecture by telling us that every hyperbolic manifold has a finite cover which contains an embedded π_1-injective surface.

To see this, first note that Corollary 2.1 tells us that the fundamental group G of a hyperbolic manifold is the fundamental group of a compact NPC cube complex. Since the G is Gromov hyperbolic, we then know by Theorem 4.3 that G is virtually special. In particular it follows that the quasi-convex surface subgroups produced by Kahn and Markovic are separable. But now a theorem of Scott [**38**] tells us that a separable surface subgroups homotops to an embedding in a finite cover.

Bibliography

1. Ian Agol, *The virtual haken conjecture*, Preprint, 2012. With an appendix by I. Agol, D. Groves and J.Manning.
2. Werner Ballmann, *Nonpositively curved manifolds of higher rank*, Ann. of Math. (2) **122** (1985), no. 3, 597–609.
3. ———, *Lectures on spaces of nonpositive curvature*, DMV Seminar, vol. 25, Birkhäuser Verlag, Basel, 1995, With an appendix by Misha Brin.
4. Werner Ballmann and Sergei Buyalo, *Periodic rank one geodesics in Hadamard spaces*, Geometric and probabilistic structures in dynamics, Contemp. Math., vol. 469, Amer. Math. Soc., Providence, RI, 2008, pp. 19–27.
5. Nicolas Bergeron and Daniel T. Wise, *A boundary criterion for cubulation*, Amer. J. of Math. **134** (2023), no. 3, 843–859.
6. Martin Bridson, *Geodesics and curvature in metric simplicial complexes*, Group Theory from a Geometrical Viewpoint, Proc. ICTP Trieste 1990 (A. Haefliger E. Ghys and A. Verjovsky, eds.), World Scientific, Singapore, 1991, pp. 373–464.
7. Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Springer-Verlag, Berlin, 1999. MR 1 744 486
8. Marc Burger and Shahar Mozes, *Lattices in product of trees*, Inst. Hautes Études Sci. Publ. Math. (2000), no. 92, 151–194 (2001). MR 2002i:20042
9. Pierre-Emmanuel Caprace, *Conjugacy of 2-spherical subgroups of Coxeter groups and parallel walls*, Algebr. Geom. Topol. **6** (2006), 1987–2029 (electronic). MR MR2263057
10. Pierre-Emmanuel Caprace and Michah Sageev, *Rank rigidity for cat(0) cube complexes*, Geometric And Functional Analysis **21** (2011), 851–891, 10.1007/s00039-011-0126-7.
11. Indira Chatterji, Cornelia Drutu, and Frederic Haglund, *Kazhdan and haagerup properties from the median viewpoint*, Adv. Math. **225** (2010), 882–921.
12. Indira Chatterji and Graham Niblo, *From wall spaces to* CAT(0) *cube complexes*, Internat. J. Algebra Comput. **15** (2005), no. 5-6, 875–885. MR MR2197811 (2006m:20064)
13. Victor Chepoi, *Graphs of some* CAT(0) *complexes*, Adv. in Appl. Math. **24** (2000), no. 2, 125–179. MR MR1748966 (2001a:57004)
14. Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235.
15. Max Dehn, *Papers on group theory and topology*, Springer-Verlag, New York-Berlin, 1987, Translated from the German and with introductions and an appendix by John Stillwell, With an appendix by Otto Schreier.
16. V. N. Gerasimov, *Semi-splittings of groups and actions on cubings*, Algebra, geometry, analysis and mathematical physics (Russian) (Novosibirsk, 1996), Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk, 1997, pp. 91–109, 190. MR 99c:20049
17. É. Ghys and P. de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov*, Progress in Mathematics, vol. 83, Birkhäuser Boston Inc., Boston, MA, 1990, Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988. MR 92f:53050
18. Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev, *Widths of subgroups*, Trans. Amer. Math. Soc. **350** (1998), no. 1, 321–329.
19. Martin Greendlinger, *On Dehn's algorithms for the conjugacy and word problems, with applications*, Comm. Pure Appl. Math. **13** (1960), 641–677. MR 23 #A2327
20. M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263.
21. D. Guralnik, *Coarse decompositions of boundaries for CAT(0) groups*, Preprint. arXiv:math/0611006 [math.GR], 2006.

22. Frederic Haglund, *Finite index subgroups of graph products*, Geom. Dedicata **135** (2008), 167–209.
23. Frédéric Haglund and Frédéric Paulin, *Simplicité de groupes d'automorphismes d'espaces à courbure négative*, The Epstein birthday schrift, Geom. Topol., Coventry, 1998, pp. 181–248 (electronic). MR 2000b:20034
24. Frédéric Haglund and Daniel T. Wise, *A combination theorem for special cube complexes*, Submitted.
25. Marshall Hall, Jr., *Coset representations in free groups*, Trans. Amer. Math. Soc. **67** (1949), 421–432.
26. Tim Hsu and Daniel T. Wise, *Separating quasiconvex subgroups of right-angled Artin groups*, Math. Z. **240** (2002), no. 3, 521–548. MR 1 924 020
27. Jeremy Kahn and Vladimir Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. **175** (2012), no. 3, 1127–1190.
28. I.J. Leary, *A metric Kan–Thurston theorem*, Preprint. arXiv:1009.1540, to appear in Journal of Topology.
29. Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, Springer-Verlag, Berlin, 1977, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.
30. Gábor Moussong, *Hyperbolic coxeter groups*, Ph.D. thesis, Ohio State University, 1988.
31. Amso Nevo and Michah Sageev, *Poisson boundaries of CAT(0) cube complexes*, to appear in Groups Geom. Dyn.
32. G. A. Niblo and L. D. Reeves, *Coxeter groups act on* CAT(0) *cube complexes*, J. Group Theory **6** (2003), no. 3, 399–413. MR 2004e:20072
33. Graham A. Niblo and Martin A. Roller, *Groups acting on cubes and Kazhdan's property (T)*, Proc. Amer. Math. Soc. **126** (1998), no. 3, 693–699. MR 98k:20058
34. Bogdan Nica, *Cubulating spaces with walls*, Algebr. Geom. Topol. **4** (2004), 297–309 (electronic). MR 2 059 193
35. M.A. Roller, *Poc sets, median algebras and group actions. An extended study of Dunwoody's construction and Sageev's theorem*, Preprint, Univ. of Southampton, 1998.
36. Pascal Rolli, *Notes on cat(0) cube complexes*, Preprint, 2012.
37. Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) **71** (1995), no. 3, 585–617. MR 97a:20062
38. Peter Scott, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2) **17** (1978), no. 3, 555–565.
39. John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565.
40. Daniel T. Wise, *A non-positively curved squared complex with no finite covers*, Preprint, 1995.
41. ———, *Cubulating small cancellation groups*, GAFA, Geom. Funct. Anal. **14** (2004), no. 1, 150–214.
42. Daniel T. Wise, *From riches to raags*, Regional Conference Series in Mathematics, vol. 117, Amer. Math. Soc., Providence, 2012.