CAT(0) CUBE COMPLEXES, MEDIAN GRAPHS, AND CUBULATING GROUPS

MARK HAGEN

Abstract. Notes from a mini-course at *Into the Forest*, 28 July–2 August 2019.

Contents

1. Lecture 1: CAT(0) cube complexes, hyperplanes, convexity 1
1.1. What is being skipped/ignored in these notes 1
1.2. The cubical viewpoint 1
1.3. The median graph viewpoint 4
1.4. Equivalence of the two viewpoints 5
1.5. Bridges, convexity, and the Helly property 8
2. Lecture 2: Cubulating wallspaces: enforcement of the Helly property 9
3. Lecture 3: Cubulating groups 11
3.1. Consequences of cubulating 11
3.2. Obtaining properness and cocompactness in the hyperbolic case 12
More exercises 15
References 16

1. Lecture 1: CAT(0) cube complexes, hyperplanes, convexity

There are (at least) two seemingly different viewpoints on CAT(0) cube complexes, which we'll call the *cubical* viewpoint and the *median graph* viewpoint. Since they are closely related, and there are situations where each of them is more convenient than the other, we'll see both.

1.1. What is being skipped/ignored in these notes.

In the interest of time, I am not going to talk about disc diagrams in CAT(0) cube complexes. This means that we won't be able to prove certain things from the "cubical" viewpoint. We will be able to do things more formally from the median viewpoint, which is a bit less technical. So some faith will be required before we get to the median stuff. If you want to know about disc diagrams, you can read about them e.g. here [Wis18] or here [Wis12].

1.2. The cubical viewpoint.

Definition 1.1 (Cube, cube complex)**.** Given $d \geq 0$, a *d–cube* is a copy of $[-\frac{1}{2}, \frac{1}{2}]^d$, equipped with the ℓ^2 (Euclidean) or ℓ^1 metric (according to convenience). Its *dimension* is d. A *codimension–k face* is a subspace obtained by restricting k of the coordinates to $\pm\frac{1}{2}$. Note that each face is a $(d-k)$–cube. A *midcube* is a subspace obtained by restricting exactly one coordinate to 0. Midcubes are $(d-1)$–cubes, but are not subcomplexes of X when X is given the obvious (product) cell structure in which each 1–cube is viewed as a graph with two vertices and one edge.

Date: July 26, 2019.

A *cube complex* is a CW complex X such that

- each cell is a d–cube for some d;
- each cell is attached using an isometry of some face.

The *dimension* of X is the supremum of the set of dimensions of cubes of X.

Definition 1.2 (Link, nonpositively-curved cube complex, CAT(0) cube complex)**.** Let X be a cube complex. Given $v \in X^{(0)}$, the *link* $\mathrm{Lk}(v)$ is the simplex complex with a $(d-1)$–simplex for each d–cube of X containing v; a simplex σ is a face of a simplex τ if the cube corresponding to σ is a face of the cube corresponding to τ. The complex X is *nonpositively curved* if $\mathrm{Lk}(v)$ is *flag complex* for each $v \in X^{(0)}$, i.e.

- $\mathrm{Lk}(v)$ is a simplicial complex – there are no loops or multiple edges in the 1–skeleton and more generally, any simplex is completely determined by its 0–skeleton;
- if $\sigma_0, \ldots, \sigma_k$ are pairwise-adjacent 0–simplices in $\mathrm{Lk}(v)$, then they span a k–simplex.

Intuitively, the second condition says that if X contains the "corner" of a cube of dimension at least 3, then it contains the cube.

A simply connected nonpositively-curved cube complex is a *CAT(0) cube complex*.

Remark 1.3 (The CAT(0) metric, the ℓ^1 metric, and the graph metric)**.** A result of Bridson [Bri91] implies (in the finite-dimensional case) that, if X is a CAT(0) cube complex, then by regarding each cube as a Euclidean unit cube, we can equip X with a geodesic metric making it a CAT(0) space; this was later shown to hold in the infinite-dimensional case by Leary [Lea13]. We will denote by d_2 the CAT(0) metric on X, but we're not going to do much CAT(0) geometry. There's too much other structure that's usually (although not always) more useful.

We denote by d_1 the usual graph metric on the 1–skeleton $X^{(1)}$. This is the metric we'll always use (in fact, we're just going to think about the resulting subspace metric on the 0–skeleton). There is a way to extend it to an ℓ^1 metric on X so that the 1–skeleton is isometrically embedded, and each cube, equipped with the usual ℓ^1 metric, is also isometrically embedded. But for now, we just need the graph metric.

Example 1.4. Here are some CAT(0) cube complexes:

- Every tree is a 1–dimensional CAT(0) cube complex, and vice versa.
- The product of two CAT(0) cube complexes is again a CAT(0) cube complex.
- (Salvetti complex.) The *right-angled Artin group* $A(\Gamma)$ presented by the finite simplicial graph Γ has a generator for each vertex, and two generators commute if and only if their corresponding vertices are joined by an edge. The *Salvetti complex* is formed as follows. Start with the presentation complex of the presentation just described (so, a square complex). Now, add any cube whose 1–skeleton appears. This is a nonpositively curved cube complex with fundamental group $A(\Gamma)$. So, $A(\Gamma)$ acts freely and cocompactly on a CAT(0) cube complex. One can do something similar for right-angled Coxeter groups.
- Let X decompose as a tree of spaces so that each vertex and edge space is a CAT(0) cube complex, and each edge space embeds in the incident vertex spaces as a convex subcomplex (in the sense defined below). Then X is a CAT(0) cube complex.

The most important feature of a CAT(0) cube complex is the family of *hyperplanes*:

Definition 1.5 (Hyperplane)**.** Let X be a CAT(0) cube complex. A *hyperplane* $H \subset X$ is a subspace such that for each (closed) cube c of X, either $H \cap c = \emptyset$ or $H \cap c$ is a single midcube of c. Note that for each edge e, the midpoint of e is contained in a unique hyperplane, which we call the hyperplane *dual* to e. The *carrier* $\mathcal{N}(H)$ of the hyperplane H is the union of all closed cubes c such that $H \cap c \neq \emptyset$.

The following important theorem is due to Sageev [Sag95]. We won't prove it here, but we'll see some of the ideas from the median viewpoint.

Theorem 1.6. *Let X be a CAT(0) cube complex. Let H be a hyperplane of X. Then:*

- *H is a CAT(0) cube complex, where the cubes are the midcubes of the cubes of $\mathcal{N}(H)$.*
- *If $\dim X < \infty$, then $\dim H < \dim X$.*
- *H is convex with respect to the CAT(0) metric d_2.*
- *$X - H$ has exactly two components, called* halfspaces*.*
- *$\mathcal{N}(H)$ is isometric to $H \times [-\frac{1}{2}, \frac{1}{2}]$, and $H \times \{\pm 12\}$ lie in different components of $X - H$.*

Moreover, if γ is an edge-path in $X^{(1)}$, then γ is a geodesic (with respect to the graph metric d_1 on $X^{(1)}$) if and only if it contains at most one edge dual to each hyperplane.

Remark 1.7 (Separation and distance)**.** The last clause has a very useful consequence. Given $x, y \in X$, we say that the hyperplane H *separates* x, y if x, y lie in different components of $X - H$ (i.e. in different halfspaces). So, if x, y are vertices of X, then $\mathsf{d}_1(x, y)$ is exactly the number of hyperplanes separating x from y.

At this point, we're taking as read a very important theorem about CAT(0) cube complexes. We're going to see some of the ideas underlying this theorem, in a slightly different context, a bit later, when we see why CAT(0) cube complexes and *median graphs* are the same thing.

Remark 1.8. Let X be a CAT(0) cube complex. Then each midcube in X is contained in a unique hyperplane. For each d–cube c of X, there are exactly d hyperplanes intersecting c, and they all contain the barycentre of c.

Next to hyperplanes, the most important feature of a CAT(0) cube complex is its collection of *convex subcomplexes*.

Definition 1.9 (Convex subcomplex, crossing)**.** A subcomplex Y of a CAT(0) cube complex X is *convex* if the following holds. Let $\mathcal{H}_Y$ be the set of all halfspaces in X (i.e. components of complements of hyperplanes) containing Y. Then every cube in $\bigcap_{H \in \mathcal{H}} H_Y$ belongs to Y. The *convex hull* of an arbitrary subspace Y is the intersection of all convex subcomplexes containing Y.

Given a hyperplane H and a subspace $Y \subset X$ which is either a convex subcomplex or a hyperplane, we say that H *crosses* Y if $H \cap Y \neq \emptyset$. The hyperplane H crosses the hyperplane H' if and only if each of the halfspaces of $X - H$ intersects each of the halfspaces of $X - H'$.

Exercise 1. *Let X be a CAT(0) cube complex. Let $Y \subset X$ be a convex subcomplex. Prove that Y is again a CAT(0) cube complex.*

Remark 1.10 (Examples of convex subcomplexes)**.** Each cube is a convex subcomplex. For each hyperplane H, the carrier $\mathcal{N}(H)$ is a convex subcomplex. Similarly, let H be a hyperplane, let $\bar{H}_0$ be one of the associated halfspaces, and let $\bar{H}$ be the subcomplex of X spanned by the vertices in $\bar{H}_0$. Then $\bar{H}$ is a convex subcomplex. (We will call $\bar{H}$ a *combinatorial halfspace*.) Each hyperplane determines two disjoint combinatorial halfspaces whose union contains all of $X^{(0)}$.

A key property of CAT(0) cube complexes is that each 0–cube is completely determined by the set of combinatorial halfspaces that contain it. Let's explore this:

Definition 1.11 (Consistent orientation)**.** A *consistent orientation* is a collection σ of combinatorial halfspaces such that:

- for each hyperplane H, *exactly one* of the two associated combinatorial halfspaces belongs to σ;
- if $\bar{H}, \bar{H}' \in \sigma$, then $\bar{H} \cap \bar{H}' \neq \emptyset$;
- for any $x \in X^{(0)}$, the set of $\bar{H} \in \sigma$ such that $x \notin \bar{H}$ is finite.

Exercise 2 (Principal orientations)**.** *Show that if $y \in X^{(0)}$, then the set of combinatorial halfspaces containing y is a consistent orientation.*

Lemma 1.12 (Every consistent orientation is principal)**.** *Let σ be a consistent orientation. Then there exists a unique 0–cube y such that $y \in \bigcap_{\bar{H} \in \sigma} \bar{H}$.*

Proof. Uniqueness of y is clear from Remark 1.7. Indeed, if y, y' are as in the statement, then no hyperplane separates them, so $\mathsf{d}_1(y, y') = 0$.

Now we have to prove existence. Fix $x_0 \in X^{(0)}$ and let σ_0 be the principal orientation corresponding to x_0, provided by Exercise 2. By the definition of consistency, there are finitely many halfspaces in $\sigma \triangle \sigma_0$. More precisely, we have hyperplanes $H_1, \ldots, H_k$ such that the following holds. For each H_i, let $\overleftarrow{H}_i, \overrightarrow{H}_i$ be the associated combinatorial halfspaces. Then (after relabelling), $x_0 \in \overleftarrow{H}_i$ for all i, and $\overrightarrow{H}_i \in \sigma$. For any other hyperplane H, the associated halfspace $\overleftarrow{H}$ belonging to σ is the one containing x_0.

Let the *size* of σ be the smallest $\ell \in \mathbb{N}$ such that there exists $z \in X^{(0)}$ with the property that $\sigma_z \triangle \sigma$ is a set of halfspaces defined by ℓ hyperplanes. We've just seen that the size of σ is at most k; by choosing x_0 to minimise the size, we can assume that σ has size k.

If $k = 0$, then $\sigma = \sigma_0$ is principal, and we are done with $y = x_0$. So, suppose that $k \geq 1$, which amounts to supposing that σ is not principal.

The halfspaces $\overrightarrow{H}_i$ are partially ordered by inclusion. Up to relabeling, we have that $\overrightarrow{H}_i \subset \overrightarrow{H}_j$ implies $i \leq j$. Let $1, \ldots, m$ be such that $\overrightarrow{H}_1, \ldots, \overrightarrow{H}_m$ are the maximal halfspaces in the partial order, where $m \leq k$.

Suppose some hyperplane H separates x_0 from $\mathcal{N}(H_1)$. Then $x_0 \in \overleftarrow{H}$ and $\mathcal{N}(H_1) \subset \overrightarrow{H}$. Hence $\overrightarrow{H}_1 \subseteq \overrightarrow{H}$. This implies that $\overrightarrow{H} \in \sigma$ (by consistency). On the other hand, if $H \neq H_i$ for all i, then $\overleftarrow{H} \in \sigma$, a contradiction. Hence $H = H_i$ for some $i \neq 1$, contradicting the maximality of $\overrightarrow{H}_1$. Thus no hyperplane separates x_0 from $\mathcal{N}(H_1)$.

We have just shown that for each halfspace $\overrightarrow{H}$ containing $\mathcal{N}(H_1)$, we have $x_0 \in \overrightarrow{H}$. So, by convexity of $\mathcal{N}(H_1)$, we have $x_0 \in \mathcal{N}(H_1)$. Since $\mathcal{N}(H_1) \cong H_1 \times [-\frac{1}{2}, \frac{1}{2}]$, there is a 0–cube $x_0' \in \mathcal{N}(H_1)$ such that the only hyperplane separating x_0, x_0' is H_1 itself. Let σ_0' be the principal orientation for x_0'. Then the set of halfspaces in σ but not σ_0' is $\{\overrightarrow{H}_2, \ldots, \overrightarrow{H}_k\}$. In other words, σ has size $k - 1$ with respect to x_0', contradicting our choice of k. So, σ is principal. $\square$

The preceding fact will be used later to find medians. It also prefigures the idea of "cubulating a wallspace", which we will see in the next lecture.

1.3. **The median graph viewpoint.**

Definition 1.13 (Median graph)**.** Let Γ be a connected graph, equipped with the combinatorial metric d in which each edge has length 1. We say Γ is a *median graph* if there exists $\mu : (\Gamma^{(0)})^3 \to \Gamma^{(0)}$ such that for all vertices $x_1, x_2, x_3 \in \Gamma$, the vertex $\mu = \mu(x_1, x_2, x_3)$ satisfies

$$\mathsf{d}(x_i, x_j) = \mathsf{d}(x_i, \mu) + \mathsf{d}(\mu, x_j)$$

for all distinct $i, j \in \{1, 2, 3\}$ and μ is the unique vertex with that property.

Definition 1.14 (Median-convex)**.** Let Γ be a median graph. The (induced) subgraph $\Lambda \subset \Gamma$ is *median-convex* if for all $x, y \in \Lambda$ and $z \in \Gamma$, the median $\mu(x, y, z) \in \Lambda$.

Exercise 3. *Let $\Lambda \subset \Gamma$ be median-convex. Show that Λ is convex in the metric sense: any geodesic of Γ with endpoints in Λ lies in Λ. Also prove the converse.*

Definition 1.15 (Gate)**.** Let Γ be a median graph, and let $\Lambda \subset \Gamma$ be a median-convex subgraph. Let $x \in \Gamma$ be a vertex. Just because Γ is a graph, there exists a vertex $y \in \Lambda$ such that $\mathsf{d}(y, x) = \mathsf{d}(\Lambda, x)$. Suppose that y' is another such vertex. Then by median-convexity, $\mu = \mu(x, y, y') \in \Lambda$. By the definition of a median, μ lies on a geodesic from x to y and from x to y'. Our choice of y, y' implies $\mu = y = y'$. Hence Λ has a unique closest vertex to x, called the *gate* of x in Λ.

Let $\mathfrak{g}_\Lambda : \Gamma^{(0)} \to \Lambda^{(0)}$ be the map sending each vertex to its gate. (When Λ is clear, we will just write $\mathfrak{g}$ for the gate map.)

Exercise 4. *Let Γ be a median graph and let Λ be a median-convex subgraph. Suppose that $x, y \in \Gamma$ are adjacent. Show that $\mathfrak{g}_\Lambda(x), \mathfrak{g}_\Lambda(y)$ are adjacent or equal. Using this, extend the gate map over edges to get a 1–lipschitz retraction $\mathfrak{g}_\Lambda : \Gamma \to \Lambda$.*

Median graphs also have hyperplanes, which can be defined in a very natural way:

Definition 1.16 (Hyperplane, convex split, carrier)**.** Let Γ be a median graph and let e be an edge. Observe that e is median-convex. Let $\mathfrak{g} : \Gamma \to e$ be the gate map. Let $v, w \in e$ be the vertices.

Observe that $\Gamma^{(0)} = (\mathfrak{g}^{-1}(v) \cap \Gamma^{(0)}) \sqcup (\mathfrak{g}^{-1}(w) \cap \Gamma^{(0)})$. Each of the halves of this bipartition is a *vertex-halfspace* associate to e, and we let $\overleftarrow{e}, \overrightarrow{e}$ denote the vertex-halfspaces.

Exercise 5. *The subgraphs spanned by $\overleftarrow{e}$ and $\overrightarrow{e}$ are median-convex.*

A partition $\{\overleftarrow{e}, \overrightarrow{e}\}$ of $\Gamma^{(0)}$ of the above type is the *convex split dual to* the edge e.

We say that e, f are *equivalent* if they induce the same convex split (up to exchanging the vertex-halfspaces). An equivalence class of edges is a *(median graph) hyperplane.*

1.4. **Equivalence of the two viewpoints.** The two viewpoints are united by the following theorem. The fact that median graphs determine CAT(0) cube complexes was proved by Chepoi [Che00].

Theorem 1.17 (Cube complexes and median graphs)**.** *Let X be a CAT(0) cube complex. Then*

- *$X^{(1)}$ is a median graph.*
- *If $Y \subset X$ is a cubically convex subcomplex, then $Y^{(1)}$ is a median-convex subgraph of $X^{(1)}$.*
- *If $Y \subset X$ is a full subcomplex of X (i.e. every cube of X with 0–skeleton in Y is in Y) and $Y^{(1)}$ is median-convex in $X^{(1)}$, then Y is a convex subcomplex.*
- *For each hyperplane H of X, dual to an edge e, the convex split $\overleftarrow{e}, \overrightarrow{e}$ is exactly the partition of $X^{(0)}$ induced by the two components of $X - H$.*

Conversely:

- *If Γ is a median graph, then there exists a unique (up to cubical isomorphism) CAT(0) cube complex X such that $X^{(1)} \cong \Gamma$.*
- *If $\Lambda \subset \Gamma$ is a median-convex subgraph, then the CAT(0) cube complex Y with 1–skeleton Λ is a convex subcomplex of X (spanned by Λ).*

First, we show that 1–skeleta of CAT(0) cube complexes are median.

Lemma 1.18 ($X^{(1)}$ is a median graph)**.** *Let X be a CAT(0) cube complex. Then $X^{(1)}$ is a median graph.*

Proof. Let $x, y, z \in X$ be vertices. Let H be a hyperplane. Let $\overrightarrow{H}, \overleftarrow{H}$ be the associated halfspaces. Then one of the halfspaces — say $\overleftarrow{H}$ — contains at least 2 of the vertices $\{x, y, z\}$.

Let σ be the set of halfspaces $\overleftarrow{H}$ defined above, as H varies over all the hyperplanes.

Observe that σ contains exactly one halfspace associated to each hyperplane.

If H, H' are hyperplanes, then $\overleftarrow{H}, \overleftarrow{H}'$ both contain at least two of the points $\{x, y, z\}$, so they contain a common point, i.e. $\overleftarrow{H} \cap \overleftarrow{H}' \neq \emptyset$.

Finally, let σ_x be the set of halfspaces containing x. Then any hyperplane H such that $x \notin \overleftarrow{H}$ separates x from both y and z, so $|\sigma - \sigma_x| \leq \mathsf{d}_1(x, y) + \mathsf{d}_1(x, z)$.

Hence σ is a consistent orientation. By Lemma 1.12, there is a unique $m \in X^{(0)}$ such that $\sigma_m = \sigma$.

If H is a hyperplane separating x from y, then let $\overleftarrow{H}$ be the associated halfspace containing two of $\{x,y,z\}$. So, $m \in \overleftarrow{H}$. If $x \in \overleftarrow{H}$, then $y \in \overrightarrow{H}$, since H separates x, y. So, H separates y, m. Otherwise, if $x \in \overrightarrow{H}$, then similarly H separates x, m. Now, H cannot separate x, y from m, for otherwise we would have $x, y \in \overrightarrow{H}$, contradicting how $\overleftarrow{H}$ was defined. Thus $\mathsf{d}_1(x,y) = \mathsf{d}_1(x,m) + \mathsf{d}_1(y,m)$. A similar equality holds for x, z and z, y, so m is a median.

Exercise 6. *Prove that m is the unique vertex of X satisfying the requirements in Definition 1.13.*

Hence $X^{(1)}$ is a median graph. □

Next, we show that convexity and median-convexity coincide.

Lemma 1.19. *Let $Y \subset X$ be a full subcomplex of X. Then Y is convex if and only if $Y^{(1)}$ is median-convex in $X^{(1)}$.*

Proof. Suppose that Y is convex and let $x, y \in Y$ be vertices of Y. Let $z \in X$ be a vertex. Let m be the median of x, y, z. By convexity of Y, to show that $m \in Y$, we need to show that no hyperplane separates m from Y. If H is such a hyperplane, then the halfspaces $\overleftarrow{H}, \overrightarrow{H}$ satisfy $m \in \overleftarrow{H}$ and $Y \subset \overrightarrow{H}$. Thus $x, y \in \overrightarrow{H}$, so $\overleftarrow{H}$ cannot contain two of x, y, z. But the construction of m ensures that the halfspace containing m contains two of x, y, z, a contradiction. Hence $m \in Y$, so $Y^{(1)}$ is median-convex.

Conversely, suppose that $Y^{(1)}$ is median-convex, and let $\mathfrak{g} : X^{(1)} \to Y^{(1)}$ be the gate map (the existence of which requires median-convexity). For each hyperplane H, let $\overleftarrow{H}$ be the combinatorial halfspace containing Y, let $Z = \bigcap_H \overleftarrow{H}$. We need to show that $Y = Z$.

For a contradiction, suppose that there exists a 0–cube $x \in Z - Y$, and let $y = \mathfrak{g}(x) \in Y$. Our assumption implies that $x \neq y$. So, there exists a hyperplane H separating x, y. Since $x \in Z$, the hyperplane H cannot separate x from Y, so H crosses Y. This contradicts Exercise 7 below:

Exercise 7. *Let $Y \subset X$ be a subcomplex with $Y^{(1)}$ median-convex. Let $x \in X$ be a vertex and let H be a hyperplane. Then H separates x from Y if and only if H separates x from $\mathfrak{g}_Y(x)$, where $\mathfrak{g}_Y : X^{(1)} \to Y^{(1)}$ is the gate map. (Hint: one direction is obvious. For the other, suppose H separates x from $\mathfrak{g}_Y(x)$ but not from Y, and deduce that there exists $z \in Y$ such that H separates y, z. Show that y lies on a geodesic from x to z, and reach a contradiction with Theorem 1.6.)*

□

Next, we show that the two notions of "hyperplane" coincide:

Lemma 1.20. *Let X be a CAT(0) cube complex and let H be the hyperplane dual to an edge e. Let $\overleftarrow{H}, \overrightarrow{H}$ be the associated combinatorial halfspaces. Then (up to relabelling), $\overleftarrow{e} = \overleftarrow{H}^{(0)}$ and $\overrightarrow{e} = \overrightarrow{H}^{(0)}$.*

Sketch. Recall that e, f are equivalent edges if $\overleftarrow{e} = \overleftarrow{f}$ and $\overrightarrow{e} = \overrightarrow{f}$. Denote this eqivalence relation $\sim$. Define an equivalence relation $\sim'$ by $e \sim' f$ if e, f are dual to the same (CAT(0) cubical) hyperplane H.

Note that the edges f with $e \sim' f$ are precisely the edges with one endpoint in $\overleftarrow{H}$ and one in $\overrightarrow{H}$.

The edges f with $e \sim f$ are precisely the edges with one endpoints in $\overleftarrow{e}$ and one in $\overrightarrow{e}$, by the definition of $\sim$. Let $\mathcal{F}'$ denote the former set of edges and $\mathcal{F}$ denote the latter set.

Let $f \in \mathcal{F}'$. Let m_f, m_e be the endpoints of e, f respectively. Let γ be a (combinatorial) geodesic in H from m_e to m_f (use that H is a cubical complex). So, each edge of γ is a midcube of a square of X, and the product structure of $\mathcal{N}(H)$ gives us a strip $\gamma \times [-\frac{1}{2}, \frac{1}{2}]$ in X, where

$\gamma \times \{\pm\frac{1}{2}\}$ is a d_1–geodesic and $m_f \times [-\frac{1}{2}, \frac{1}{2}] = f$ and $m_e \times [-\frac{1}{2}, \frac{1}{2}] = e$. Let u, v be the endpoints of f in $H \times \{\frac{1}{2}\}, H \times \{-\frac{1}{2}\}$ respectively. Then $\mathfrak{g}_e(u)$ and $\mathfrak{g}_e(v)$ are the endpoints of $\gamma \times \{\frac{1}{2}\}$ and $\gamma \times \{-\frac{1}{2}\}$ respectively. So, $f \in \mathcal{F}$.

Conversely, suppose $f \in \mathcal{F}$. Let u, v be the endpoints of f, and let $\mathfrak{g}(u), \mathfrak{g}(v)$ be the (distinct, since $f \in \mathcal{F}$) endpoints of e. Let γ be a geodesic from u to $\mathfrak{g}(u)$, so $\gamma \cdot e$ is a geodesic from u to $\mathfrak{g}(v)$. For each $w \in \gamma$, let $m(w)$ be the median of $v, w, \mathfrak{g}(v)$. If U is a hyperplane separating $m(w)$ from w, then one halfspace $\overleftarrow{U}$ contains $v, \mathfrak{g}(v), m(w)$, and $\overrightarrow{U}$ contains w. Suppose that $U \neq H$. Then U does not separate u, v or $\mathfrak{g}(u), \mathfrak{g}(v)$ and we have seen that U does not separate $v, \mathfrak{g}(v)$. So $u, v, \mathfrak{g}(u), \mathfrak{g}(v) \in \overleftarrow{U}$. Since γ is a geodesic, it contains at most one edge dual to U, so since $w \in \gamma$, we have $w \in \overleftarrow{U}$. This is a contradiction, so $U = H$.

On the other hand, if $U = H$, then $u, \mathfrak{g}(u) \in \overrightarrow{U}$ and $v, \mathfrak{g}(v) \in \overleftarrow{U}$. As above, $w \in \overrightarrow{H}$. On the other hand, $m(w)$ lies on a geodesic from v to $g(v)$, by the definition of a median, so arguing as above shows $m(w) \in \overleftarrow{H}$. Hence H separates $w, m(w)$.

Since there is a unique hyperplane separating them, $\mathsf{d}_1(w, m(w)) = 1$, i.e. w and $m(w)$ are joined by an edge $e(w)$.

Exercise 8. *X contains a strip $\gamma \times [-\frac{1}{2}, \frac{1}{2}]$ where $\gamma = \gamma \times \{-\frac{1}{2}\}$, and $\gamma \times \{\frac{1}{2}\}$ is a geodesic from v to $\mathfrak{g}(v)$, and the edges of the form $\gamma(i) \times [-\frac{1}{2}, \frac{1}{2}]$ in the strip are exactly the edges $e(w), w \in \gamma^{(0)}$. Hence $e \sim' f$.*

We have shown that $\mathcal{F} = \mathcal{F}'$. In other words, an edge f has endpoints in different halfspaces $\overleftarrow{H}, \overrightarrow{H}$ if and only if f has endpoints in different median-halfspaces $\overleftarrow{e}, \overrightarrow{e}$.

To conclude from here is an exercise. □

Now we do the other direction. Fix a median graph Γ, and let X be the cube complex obtained as follows: for each induced subgraph of Γ isomorphic to the 1–skeleton of an n–cube, $n \geq 2$, add the n–cube (identifying faces in the obvious way).

Lemma 1.21. *X is simply connected.*

Proof. Let σ be a closed edge-path in X. Recall that edges e, f of Γ are equivalent, denoted $e \sim f$, if $\overleftarrow{e} = \overleftarrow{f}$ and $\overrightarrow{e} = \overrightarrow{f}$ (up to relabelling).

Since σ is a closed path, we can choose e, f to be distinct edges on σ whose images in X are equivalent. Choose them to be *innermost*, in the sense that, if σ' is the subpath of σ strictly between e and f, then no two edges of σ' are equivalent, and no edge on σ' is equivalent to e.

If e, f are consecutive, then let u be the initial vertex of e, let v be a vertex common to e and f, and let $w \neq u$ be the terminal vertex of f. Then $u, w \in \overleftarrow{e}$ (say), and $v \in \overrightarrow{e} = \overrightarrow{f}$. The median of u, v, w lies in $\overleftarrow{e}$, since the latter space is median-convex. On the other hand, the median lies on a geodesic from u to v and from w to v, so the median is v. This is a contradiction, since $v \notin \overleftarrow{e}$. Hence e, f are not consecutive.

Let $d = |\sigma'|$.

Let h be an edge of σ'. Since e, f were an innermost pair, $(\overleftarrow{h}, \overrightarrow{h})$ and $(\overleftarrow{e}, \overrightarrow{e})$ are different convex splits. Without loss of generality, $e \subset \overleftarrow{h}$ and $f \subset \overrightarrow{h}$. Assume h is the edge on σ' immediately after e. Let u, v be the vertices of e, with $v \in h$. Let v', u' be the vertices of f, with $u' \in \overleftarrow{e}$ and $u \in \overleftarrow{e}$. Let m be the median of u, u', w, where w is the other vertex of h. Then m lies on a geodesic $[u, w]$. Since Γ is median, it contains no 3–cycle, so $|[u, w]| = 2$. Since m is on a geodesic from $u \in \overleftarrow{e}$ to $u' \in \overleftarrow{e}$, we have $m \in \overleftarrow{e}$, and in particular $m \neq v$. So u, v, w, m are the vertices of a square s in X. So replacing the path (u, v, w) with (u, m, w) amounts to homotoping σ across a square to produce a closed path τ. Now, the edge $[m, w]$ is equivalent to e, and the distance from $[m, w]$ to f in τ is strictly less than d. So, by induction, τ is nullhomotopic, so σ was as well. Thus X is simply connected. □

Exercise 9. *Show that X is nonpositively curved (or see the first paragraph of the proof of Theorem 6.1 in this paper by Chepoi:).*

So, we know that Γ is the 1–skeleton of a CAT(0) cube complex X. We also know that Γ has a median coming from the hyperplanes of X, and uniqueness of medians shows that these two coincide. So, we're finished.

From now on, we will work in a CAT(0) cube complex X, equipped with (geometric) hyperplanes, but we can freely use either the median or the halfspace viewpoint interchangeably when studying convexity.

1.5. **Bridges, convexity, and the Helly property.** Here are two very useful facts about convexity that are very helpful in geometric arguments.

Theorem 1.22 (Bridge theorem)**.** *Let X be a CAT(0) cube complex and let $A, B \subset X$ be convex subcomplexes. Let $\mathfrak{g}_A, \mathfrak{g}_B : X \to A, B$ be the gate maps. Then:*

- *A hyperplane H crosses $\mathfrak{g}_A(B)$ if and only if H crosses A and B. A hyperplane H separates A, B if and only if H separates $\mathfrak{g}_A(B)$ from $\mathfrak{g}_B(A)$.*
- *The convex hull of $\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A)$ is a CAT(0) cube complex isomorphic to $I \times \mathfrak{g}_A(B)$, where I is the convex hull of a pair of vertices $a \in \mathfrak{g}_A(B)$ and $b \in \mathfrak{g}_B(A)$. In particular, $\mathfrak{g}_A(B)$ and $\mathfrak{g}_B(A)$ are isomorphic.*

Proof. Suppose that H crosses A and B. Choose $b, b' \in B^{(0)}$ on opposite sides of H. Let U be the combinatorial halfspace associated to H and containing b. Then since H crosses A, we have $U \cap A \neq \emptyset$, and it is easy to see that the intersection of convex subcomplexes is convex, so $U \cap A$ is convex.

Let $a = \mathfrak{g}_A(b)$ and let $a' = \mathfrak{g}_{U\cap A}(b)$. Suppose that H separates a from b. Let $\mu = \mu(b, a, a')$. Then $\mu \in A$, since $a, a' \in A$ and $A^{(1)}$ is median-convex (by convexity of A). Moreover, since $a', b \in U$, and U is median-convex, we have $\mu \in U$. So $\mu \in U \cap A$, and μ lies on a geodesic from b to a. Hence $\mathsf{d}_1(\mu, b) \leq \mathsf{d}_1(a, b)$, whence, since a is the closest vertex of a to b, we have $\mu = a$. Thus $a \in U$, a contradiction. Hence H does not separate b from $\mathfrak{g}_A(b)$. Similarly, H does not separate b' from $\mathfrak{g}_A(b')$. So $\mathfrak{g}_A(b), \mathfrak{g}_A(b')$ are points of $\mathfrak{g}_A(B)$ separated by H. Thus H crosses $\mathfrak{g}_A(B)$. Similarly, H crosses $\mathfrak{g}_B(A)$.

The other direction — H crosses $\mathfrak{g}_A(B)$ and $\mathfrak{g}_B(A)$ implies that H crosses A, B — is obvious.

Now suppose that H separates A, B. Clearly H separates $\mathfrak{g}_A(B), \mathfrak{g}_B(A)$. Conversely, suppose that H separates $\mathfrak{g}_A(B)$ from $\mathfrak{g}_B(A)$. Then either H separates A, B or H crosses, say, A. Let U be the combinatorial halfspace associated to H and not containing $\mathfrak{g}_A(B)$ (but containing some $a \in A$, since H crosses A). By the preceding statement, H cannot cross B, so $B \subset U$.

Let $b = \mathfrak{g}_B(a)$ and let $a' = \mathfrak{g}_A(b)$. Then $b \in U$, since $b \in B$. Hence, by the argument above, $a' \in U$. This is a contradiction, so H cannot cross A. Thus H separates A, B.

Let $\mathcal{H}$ be the set of hyperplanes crossing A and B, and let $\mathcal{V}$ be the set of hyperplanes separating A and B. Then:

- If $H \in \mathcal{H}$ and $V \in \mathcal{V}$, then H and V cross.
- Let $c \in \mathfrak{g}_A(B)$ and $d \in \mathfrak{g}_B(A)$ be chosen as close as possible. Then every element of $\mathcal{V}$ separates c, d and every V separating c, d belongs to $\mathcal{V}$. So $\mathcal{V}$ is exactly the set of hyperplanes crossing I, the hull of $\{c, d\}$.

Let J be the convex hull of $\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A)$. Then J is a CAT(0) cube complex whose hyperplanes are partitioned into two classes such that every element of one class crosses every element of the other. By e.g. Proposition 2.5 of [CS11], $J \cong J_1 \times J_2$, where J_1, J_2 the hyperplanes of J_1 map to elements of $\mathcal{H}$ and those in J_2 map to elements of $\mathcal{V}$. In particular, $J_2 \cong I$. (This is an exercise in consistent orientations.) ☐

The following theorem lies at the heart of the utility of CAT(0) cube complexes:

Theorem 1.23 (Helly property)**.** *Let X be a CAT(0) cube complex, and let $Y_1, \ldots, Y_k$ be convex subcomplexes such that $Y_i \cap Y_j \neq \emptyset$ for all i, j. Then $\bigcap_{i=1}^{k} Y_i \neq \emptyset$.*

Proof. The cases $k \leq 2$ are tautological. If $k = 3$, then let $y_{ij} \in Y_i \cap Y_j$ for distinct $i, j \in \{1, 2, 3\}$. Let $\mu = \mu(y_{12}, y_{23}, y_{31})$. Then $\mu \in Y_i$ for $i = 1, 2, 3$ by median-convexity.

Now suppose that $k \geq 3$. By induction, $Z = Y_3 \cap \cdots \cap Y_k \neq \emptyset$, and Z is convex, being an intersection of convex subcomplexes. By hypothesis, there exists $y_{12} \in Y_1 \cap Y_2$. By induction, there exists $y_1 \in Y_1 \cap Z$ and $y_2 \in Y_2 \cap Z$. Let $\mu = \mu(y_1, y_2, y_{12})$. Then, again, $\mu \in Y_1 \cap Y_2 \cap Z$ by convexity. □

2. Lecture 2: Cubulating wallspaces: enforcement of the Helly property

Here is the construction that explains why CAT(0) cube complexes are so ubiquitous, and answers the question one often gets: why cubes, instead of simplices or something? The construction goes back to (independently) Gerasimov and Sageev [Ger97, Sag95]; the way it's formalised here is due, independently, to Chatterji–Niblo and Nica [CN05, Nic04].

The viewpoint we take here is: start with a *wallspace* (a totally set-theoretic object) and build a median graph so that the "walls" in the wallspace are mapped bijectively to convex splits; our discussion most closely resembles Nica's treatment.

Definition 2.1 (Wallspace)**.** A *wallspace* is a pair $(S, \mathcal{W})$, where S is a set and $\mathcal{W}$ is a set of *walls*, where a wall is a pair $(\overleftarrow{W}, \overrightarrow{W})$ of disjoint subsets of S whose union is S. The subsets $\overleftarrow{W}, \overrightarrow{W}$ are *halfspaces*. We also require that, for all $s, t \in S$, the set of halfspaces containing exactly one of s, t is finite.

Example 2.2. If X is a CAT(0) cube complex, then each hyperplane H partitions $X^{(0)}$ into two sets, namely the 0–skeleta of the combinatorial halfspaces. So, the hyperplanes endow $X^{(0)}$ with a wallspace structure.

Remark 2.3. In a lot of applications, it's useful to relax the requirement that walls be actual partitions, by allowing the halfspaces to intersect. In practice, S is often a metric space, and we have geometric walls — subspaces with two complementary components, or whose (many) complementary components are grouped into two subsets. For example, hyperplanes in a CAT(0) cube complex are geometric walls. This introduces some technicalities, but the construction in this lecture works in this context, too. See [HW14] for more on this.

We are also often concerned with the situation where some group G acts on S, and for each $(\overleftarrow{W}, \overrightarrow{W}) \in \mathcal{W}$ and each $g \in G$, the bipartition $(g\overleftarrow{W}, g\overrightarrow{W}) \in \mathcal{W}$. This is a *$G$–wallspace.* (For example, if G is acting by cubical automorphisms on the CAT(0) cube complex X, then the wallspace structure coming from the hyperplanes is a G–wallspace.)

We now describe the *CAT(0) cube complex dual to* the wallspace $(S, \mathcal{W})$.

First — familiarly — an *consistent orientation* σ is a set of halfspaces with exactly one from each wall, such that $\overleftarrow{W}, \overleftarrow{W}' \in \sigma$ implies $\overleftarrow{W} \cap \overleftarrow{W}' \neq \emptyset$. When $s \in S$, note that the set σ_s of halfspaces containing s is consistent. An *admissible orientation* is a consistent orientation σ such that for some (hence any) $s \in S$, there are finitely many walls $(\overleftarrow{W}, \overrightarrow{W})$ such that $\overleftarrow{W} \in \sigma, \overrightarrow{W} \in \sigma_s$ or vice versa.

The dual cube complex $X = X(S, \mathcal{W})$ has vertex set consisting of the admissible orientations.

Next, given admissible orientations σ, σ', we say that σ, σ' are adjacent (and join them by a 1–cube in X) if there is exactly one wall W such that one of the associated halfspaces is in σ and the other is in σ'.

Let $X^{(1)}$ be the resulting graph.

Lemma 2.4. *$X^{(1)}$ is a median graph.*

Proof. Let $\sigma_1, \sigma_2, \sigma_3$ be admissible orientations. For each $(\overleftarrow{W}, \overrightarrow{W}) \in \mathcal{W}$, exactly one of $\overleftarrow{W}, \overrightarrow{W}$ is contained in at least two of $\sigma_1, \sigma_2, \sigma_3$. Up to relabelling, assume it's $\overleftarrow{W}$. Let σ be the set of such $\overleftarrow{W}$, as $(\overleftarrow{W}, \overrightarrow{W})$ varies over all walls.

Given $\overleftarrow{W}, \overleftarrow{W}'$, there exists i such that σ_i contains $\overleftarrow{W}$ and $\overleftarrow{W}'$, so by consistency of σ_i, we have $\overleftarrow{W} \cap \overleftarrow{W}' \neq \emptyset$. So σ is a consistent orientation.

Moreover, for all but finitely many walls $(\overleftarrow{W}, \overrightarrow{W})$, all three of $\sigma_1, \sigma_2, \sigma_3$ contain a fixed halfspace, say $\overleftarrow{W}$. Hence σ also contains $\overleftarrow{W}$. So σ differs from, say σ_1 on finitely many walls. Thus σ is an admissible orientation, i.e. a vertex of X.

Let $i, j \in \{1, 2, 3\}$ be distinct. Let $\#(\sigma_i, \sigma_j)$ be the number of walls $(\overleftarrow{W}, \overrightarrow{W})$ such that σ_i contains $\overleftarrow{W}$ and σ_j contains $\overrightarrow{W}$, or vice versa (i.e. the number of walls on which the two orientations differ).

Let $\sigma_i = \tau_0, \tau_1, \ldots, \tau_\ell = \sigma_j$ be a geodesic in $X^{(1)}$ from σ_i to σ_j. Then for each $k \leq \ell$, there is a unique wall $(\overleftarrow{W}_k, \overrightarrow{W}_k)$ on which the orientations τ_{k-1}, τ_k differ, say $\overleftarrow{W}_k \in \tau_{k-1}$ and $\overrightarrow{W}_k \in \tau_k$. So, σ_i and τ_k do not differ on any wall not in $\{(\overleftarrow{W}_s, \overrightarrow{W}_s) : 1 \leq s \leq k\}$. Hence $\#(\sigma_i, \sigma_j) \leq \mathsf{d}_1(\sigma_i, \sigma_j)$.

Now let $\{(\overleftarrow{V}_s, \overrightarrow{V}_s)\}$ be the (necessarily finite) set of walls on which σ_i, σ_j differ. Label the halfspaces so that $\overleftarrow{V}_s \in \sigma_i$ and $\overrightarrow{V}_s \in \sigma_j$. Partially order the walls so that $(\overleftarrow{V}_s, \overrightarrow{V}_s) \prec (\overleftarrow{V}_t, \overrightarrow{V}_t)$ if $\overleftarrow{V}_s \subset \overleftarrow{V}_t$. Let τ_1 be the set of halfspaces obtained from σ_i by replacing $\overleftarrow{V}_s$ by $\overrightarrow{V}_s$ for a single $(\overleftarrow{V}_s, \overrightarrow{V}_s)$ which is $\prec$–minimal.

Exercise 10. τ_1 *is an admissible orientation.*

Hence, by induction on $\#(\sigma_i, \sigma_j)$, we have a path from σ_i to σ_j of length $\#(\sigma_i, \sigma_j)$, so $\mathsf{d}_1(\sigma_i, \sigma_j) \leq \#(\sigma_i, \sigma_j)$.

Now let $(\overleftarrow{V}_s, \overrightarrow{V}_s)$ separate σ_1 from both σ_2 and σ_3, so that $\sigma_1 \ni \overleftarrow{V}_s$ and $\sigma \ni \overrightarrow{V}_s$. Let $(\overleftarrow{V}_t, \overrightarrow{V}_t)$ separate σ_1 from σ_2 but not σ_3. Then either $(\overleftarrow{V}_s, \overrightarrow{V}_s) \prec (\overleftarrow{V}_t, \overrightarrow{V}_t)$, or the two walls are incomparable (i.e. they cross). So the path constructed above can be chosen to pass through σ. Repeating for σ_2, σ_3 shows that σ is a median for $\sigma_1, \sigma_2, \sigma_3$.

Uniqueness of σ is an exercise. □

Since $X^{(1)}$ is a median graph, we can now, by Theorem 1.17, throw in all cubes whose 1–skeleta appear, and obtain a CAT(0) cube complex $X = X(S, \mathcal{W})$.

Remark 2.5 (Walls go to walls)**.** The assignment $s \mapsto \sigma_s$ defines a map $f : S \to X^{(0)}$. Let $\overleftarrow{W}$ be a halfspace in S. By construction, $f(\overleftarrow{W}) = \{\sigma_s : s \in \overleftarrow{W}\}$ is contained in the set $F_{\overleftarrow{W}}$ of admissible orientations containing $\overleftarrow{W}$. If $\overrightarrow{W} = S - \overleftarrow{W}$ is the complementary halfspace, then $f(\overrightarrow{W})$ is contained in the set $F_{\overrightarrow{W}}$ of admissible orientations containing $\overrightarrow{W}$, so $f(\overleftarrow{W}) \cap f(\overrightarrow{W}) = \emptyset$.

Let $\sigma, \tau \in F_{\overleftarrow{W}}$ and let $\eta \in X^{(0)}$. Let μ be the median of σ, τ, η. Since $\overleftarrow{W}$ belongs to σ and τ, the definition of μ ensures that $\overleftarrow{W} \in \mu$, so $\mu \in F_{\overleftarrow{W}}$. Hence $F_{\overleftarrow{W}}$ (and similarly $F_{\overrightarrow{W}}$) is median-convex.

Exercise 11. *Let X be a CAT(0) cube complex and let Y, Z be disjoint median-convex subgraphs of $X^{(1)}$. Suppose that $Y^{(0)} \sqcup Z^{(0)} = X^{(0)}$. Show that there exists an edge e such that $Y^{(0)} = \overleftarrow{e}$ and $Z^{(0)} = \overrightarrow{e}$.*

Hence $F_{\overleftarrow{W}}, F_{\overrightarrow{W}}$ is a convex split in X, i.e. it corresponds to a hyperplane. So, we have an induced map $F : \mathcal{W} \to \mathcal{H}$, where $\mathcal{H}$ is the set of hyperplanes in X, and a wall $(\overleftarrow{W}, \overrightarrow{W})$ separates $s, t \in S$ if and only if $F(\{\overleftarrow{W}, \overrightarrow{W}\})$ separates σ_s, σ_t. In particular, walls that *cross* — i.e. all four possible intersections of halfspaces, one from each wall — get taken to hyperplanes that cross.

It is not hard to verify that F is bijective.

Remark 2.6 (G–action on X)**.** Suppose $(S,\mathcal{W})$ is a G–wallspace. Let σ be an admissible orientation and let $g \in G$. Let $g\sigma$ be the set of halfspaces of the form $g\overleftarrow{W}$, where $\overleftarrow{W}$ is a halfspace in σ. More precisely, if $W = (\overleftarrow{W}, \overrightarrow{W})$, then $g\sigma$ contains $\overleftarrow{W}$ if σ contains $g^{-1}\overleftarrow{W}$ and $g\sigma$ contains $\overrightarrow{W}$ otherwise. It's easy to check that this defines a G–action on $X = X(S,\mathcal{W})$ by cubical automorphisms (i.e. an action of G on $X^{(1)}$ by graph automorphisms). So: **a G–action on a wallspace gives a G–action by isometries on a CAT(0) cube complex**.

Remark 2.7 (Duality)**.** Let X be a CAT(0) cube complex. Let $\mathcal{W}$ be the set of walls in $X^{(0)}$ induced by the hyperplanes. We saw before that each $x \in X^{(0)}$ gives rise to a consistent orientation σ_x of the walls, namely the set of halfspaces containing x. On the other hand, each admissible orientation of the walls yields a 0–cube, by Lemma 1.12. So, the cube complex dual to the wallspace $(X^{(0)},\mathcal{W})$ is isomorphic to X.

Here's one important takeaway: a wallspace $(S,\mathcal{W})$ might badly fail to satisfy a "Helly property" for halfspaces, in the sense that there may be admissible orientations (collections of pairwise intersecting halfspaces) σ such that for all $s \in S$, the orientations σ, σ_s are very far apart (i.e. differ on many walls). Passing to the dual cube complex gives a new wallspace (namely, the 0–skeleton with the walls coming from hyperplanes) in which every admissible orientation is principal.

In many natural examples, S is a finitely-generated group G (with a fixed word-metric) and $\mathcal{W}$ is a collection of bipartitions of G (usually each bipartition is stabilised by some interesting subgroup). The failure of the Helly property for halfspaces in G to hold "coarsely" will be reflected in the failure of G to act coboundedly on the dual cube complex.

3. Lecture 3: Cubulating groups

Given a group G and a wallspace $(S,\mathcal{W})$, with a G–action on S preserving the walls in $\mathcal{W}$, we recover an action of G on the dual cube complex $X(S,\mathcal{W})$ by cubical automorphisms (hence by isometries in both the ℓ^1 and ℓ^2 metrics, and by median-preserving isometries of the 1–skeleton).

In this lecture, we discuss:

- What does a (nice) action on a CAT(0) cube complex tell one about G?
- How do proper actions on CAT(0) cube complexes arise in nature?
- How do cocompact actions arise in nature?

For the last two questions, we'll focus on the case of hyperbolic groups.

3.1. Consequences of cubulating.

Fix a finitely-generated group G. Suppose that G acts on the CAT(0) cube complex X. Then, for example:

(1) If G acts without a global fixed point, then G does not have property (T); if the action is metrically proper, G is *a-T-menable* [NR97, NR98b].
(2) If G has bounded finite subgroups and X is finite-dimensional, then G satisfies a Tits alternative [SW05].
(3) If G acts properly and X is finite-dimensional, then G has finite *asymptotic dimension* [Wri12].
(4) If G acts properly and cocompactly, then G acts on a quasi-tree with a *WPD* element [Hag14, BHS17]. In this case, either X contains an invariant subcomplex splitting as a product, or G is *acylindrically hyperbolic*, hence is not simple, has many nontrivial quasimorphisms, etc. [CS11].
(5) If G acts properly and cocompactly, then G is bi-automatic [NR98a], of type F_∞, satisfies a quadratic isoperimetric inequality, etc.
(6) If G acts freely, then G has no distorted cyclic subgroup [Hag07].

A lot of the original motivation comes from the case where G is a hyperbolic group (in particular, the fundamental group of a hyperbolic 3–manifold). In this case, a remarkable theorem of Agol [Ago13] says that, if G acts properly and cocompactly on a CAT(0) cube complex X, then there is a finite-index subgroup $G' \leq G$ such that $G\backslash X$ is a compact *special* cube complex as defined by Haglund-Wise [HW08]. We won't give the exact definition here, but this implies that G' embeds in a right-angled Artin group, and moreover:

- G is linear (over $\mathbb{Z}$) [HW99, DJ00] and hence...
- ...G is residually finite;
- more generally, quasiconvex subgroups of G are *separable* (so $G\backslash X$ has tons of finite covers) [HW08];
- G is either virtually abelian or virtually surjects onto F_2 [Wis18];
- G is *conjugacy separable* [MZ16];
- ...

It's often also interesting to take a (not necessarily hyperbolic) group G and consider not-necessarily-proper, or not-necessarily-cocompact actions on CAT(0) cube complexes; since this is much easier to arrange, one shouldn't be surprised that the class of groups admitting such actions is large and includes wild examples.

We'll mostly focus on the hyperbolic setting, and mostly focus on proper, cocompact actions. To give some (hyperbolic) examples, G acts properly and cocompactly on a CAT(0) cube complex whenever G is:

- the fundamental group of a closed hyperbolic 3–manifold [KM12, BW12];
- a $C'(\frac{1}{6})$ small-cancellation group [Wis04];
- a hyperbolic Coxeter group [NR03];
- a hyperbolic group of the form $F_k \rtimes \mathbb{Z}$ [HW15];
- a random group at density $< \frac{5}{24}$ in Gromov's density model [OW11a, OW11b].

3.2. **Obtaining properness and cocompactness in the hyperbolic case.**

Definition 3.1 (Codimension–1 subgroup)**.** Let G be a finitely generated group, with a fixed Cayley graph Γ. A finitely-generated subgroup $H \leq G$ is *codimension*–1 if there exists $r \geq 0$ such that $\mathcal{N}_r^\Gamma(H)$ is connected and $\Gamma - \mathcal{N}_r^\Gamma(H)$ has at least 2 distinct H–orbits of *deep* components, where a component C is deep if it contains points arbitrarily far from H.

If H is a codimension–1 subgroup, then we can build a (non-canonical) wall $(\overleftarrow{H}, \overrightarrow{H})$ in G as follows. Let C be a deep component of $\Gamma - \mathcal{N}_r^\Gamma(H)$, and let $\overleftarrow{H} = H \cdot C$. Let $\overrightarrow{H} = \Gamma - \overleftarrow{H}$. Then $(\overleftarrow{H}, \overrightarrow{H})$ is an H–invariant wall. For each $g \in G$, the bipartition $(g\overleftarrow{H}, g\overrightarrow{H})$ is a gHg^{-1}–invariant wall. Letting $\mathcal{W}$ be any set of such walls, we see that $(G, \mathcal{W})$ is a G–wallspace, and hence G acts by isometries on the dual cube complex $X = X(G, \mathcal{W})$. More generally, any finite collection of codimension–1 subgroups of G gives rise to an action on a CAT(0) cube complex in the same way.

In this very general setting, one obtains a cocomapct action on the dual cube complex provided the walls satisfy a "coarse Helly property" (which is very hard to satisfy without some kind of hyperbolicity assumption on G and some kind of quasiconvexity assumption on the codimension–1 subgroups):

Lemma 3.2 (Coarse Helly implies cocompact)**.** *Let G be a finitely generated group, let Γ be a Cayley graph and d the associated word-metric. Let $H_1, \ldots, H_k \subset G$ be codimension–1 subgroups. Suppose that for all $D \geq 0$, there exists R such that the following holds: let $A_1, \ldots, A_n \in \{gH_i : g \in G, i \leq k\}$ satisfy $\mathsf{d}(A_i, A_j) \leq D$ for all $i, j \leq n$. Then there exists $x \in G$ such that $\mathsf{d}(x, A_i) \leq R$ for all $i \leq n$.*

Given the above, the action of G on any dual cube complex X constructed as above from $H_1, \ldots, H_k$ is cocompact.

Proof. Let c be a maximal cube of X, viewed as an admissible orientation. Let $U_1, \ldots, U_n$ be the hyperplanes that intersect c. For each $i \leq n$, let $(\overleftarrow{W}_i, \overrightarrow{W}_i)$ be the wall corresponding to U_i. Then $\overleftarrow{W}_i$ is a $g_i H_{j_i} g_i^{-1}$–orbit of components of $\Gamma - \mathcal{N}^{\Gamma}(g_i H_{j_i})$ for some $g_i \in G$ and H_{j_i} among our codimension–1 subgroups.

Exercise 12. *For all $i \neq j$, the walls $(\overleftarrow{W}_i, \overrightarrow{W}_i)$ and $(\overleftarrow{W}_j, \overrightarrow{W}_j)$* cross, *i.e. $\overleftarrow{W}_i \cap \overleftarrow{W}_j, \overleftarrow{W}_i \cap \overrightarrow{W}_j, \overrightarrow{W}_i \cap \overleftarrow{W}_j, \overrightarrow{W}_i \cap \overrightarrow{W}_j$ are all nonempty. (This is just a dual cube complex/wallspace thing: hyperplanes cross if and only if the corresponding walls cross. No need to mention any groups.)*

By the exercise, $g_i H_{j_i}$ and $g_\ell H_{j_\ell}$ come D–close for all i, ℓ, where D depends only on r. Hence, by hypothesis, up to enlarging R uniformly, there exists $x \in G$ such that x is R–close to each $\overleftarrow{W}_i$ and $\overrightarrow{W}_i$.

By translating, we can assume that $x = 1$. But the R–ball in Γ intersects only boundedly many left cosets of $H_1, \ldots, H_k$, so there are boundedly many families $\{g_i H_{j_i}\}$ as above (up to the action of G), so there are finitely many orbits of maximal cubes in X. Hence the action of G on X is cocompact. □

The hypothesis in the previous lemma was very artificial. How does it arise in practice?

Suppose that G is word-hyperbolic (so Γ is δ–hyperbolic for some δ depending on the generating set) and each H_i is quasiconvex (so, for some κ depending on the generating set and $\{H_i\}$, each H_i is κ–quasiconvex in Γ). In this situation, we get:

Theorem 3.3. *Let G be a hyperbolic group and let $H_1, \ldots, H_k$ be quasiconvex codimension–1 subgroups. Then G acts cocompactly on any associated dual CAT(0) cube complex.*

Remark 3.4. The word "any" is because the passage from codimension–1 subgroups to walls is not in general canonical.

The above theorem follows from Lemma 3.2 once we prove the following two things:

Lemma 3.5 (Coarse Helly property for hyperbolic spaces)**.** *Let Γ be a δ–hyperbolic geodesic space. Then for all D, κ, N there exists R such that the following holds. Let $H_1, \ldots, H_N$ be κ–quasiconvex subspaces such that $\mathsf{d}_\Gamma(H_i, H_j) \leq D$ for all i, j. Then there exists $x \in \Gamma$ such that $\mathsf{d}_\Gamma(x, H_i) \leq R$ for all i.*

Proof. This is an exercise in δ–hyperbolic geometry. For each distinct i, j, choose some y_{ij} that is D–close to H_i and H_j. A standard fact about hyperbolic spaces is that there is a tree T and a $(1, C)$–quasi-isometric embedding $T \to \Gamma$ (where C depends on δ and N) such that the leaves of T get sent to points in $\{y_{ij}\}$.

For each i, let T_i be the subtree of T spanned by points of the form y_{ij}, $1 \leq j \leq N$. Then for all i, i', the point $y_{ii'} \in T_i \cap T_{i'}$, i.e. the trees T_i pairwise intersect. By the Helly property (applied to the convex subcomplexes T_i of the 1–dimensional CAT(0) cube complex T), there exists $z \in \bigcap_i T_i$. Let $\bar{z} \in \Gamma$ be the image of z under $T \to \Gamma$. Then $\bar{z}$ lies in the image of each T_i. By quasiconvexity, the image of T_i lies uniformly close to H_i, so taking $x = \bar{z}$ completes the proof. □

Lemma 3.6 (Bounded packing)**.** *Let $H_1, \ldots, H_k$ be κ–quasiconvex subgroups of the δ–hyperbolic group G. Then for each D, there exists N such that any family of pairwise D–close left cosets of the H_i has cardinality at most N.*

Sketch. This argument is due to Hruska and Wise, see [HW14] for details. The key point is that the H_i have finite *height*, i.e. there exists n so that any collection of $n+1$ distinct conjugates of H_i has finite intersection. By induction on the height, Hruska-Wise show that for each D, there exists N_i such that any family of cosets of H_i that are pairwise D–close has cardinality at most N_i. So, any family $\{g_j H_{i_j}\}$ of cardinality at least $k \max_j N_j + 1$ contains at least $N_j + 1$ cosets of some H_j, and hence cannot be pairwise D–close. □

So, cocompactness is automatic when we cubulate a hyperbolic group, as long as we use quasiconvex codimension–1 subgroups to get our walls.

Now we turn our attention to the following question. Let G be a hyperbolic group. When can we find "enough" walls in G (coming from quasiconvex codimension–1 subgroups) to get a proper action on the dual cube complex?

Why is it a question of having "enough" walls? Well, if $g \in G$ is an infinite-order element, then $\langle g \rangle$ is a bi-infinite quasigeodesic in G. If $H_1, \ldots, H_k$ are quasiconvex codimension–1 subgroups, and $(\overleftarrow{W}_1, \overrightarrow{W}_1), \ldots, (\overleftarrow{W}_k, \overrightarrow{W}_k)$ are associated walls in G (with $(\overleftarrow{W}_i, \overrightarrow{W}_i)$ constructed from H_i as above), we need to know that we can find $h \in G$ and $i \leq k$ such that $g^n \in h\overleftarrow{W}_i$ for all $n \geq N$, and $g^n \in h\overrightarrow{W}_i$ for all $n < N$, for some N. Indeed, if we don't achieve this, then for each wall, exactly one associated halfspace contains all of $\langle g \rangle$. Hence, by choosing, for each wall, the halfspace $\langle g \rangle$, we obtain an admissible orientation, so in the dual cube complex, the corresponding vertex is stabilised by the infinite subgroup $\langle g \rangle$. This means that the action can't possibly be proper. So, we need to find a finite collection of walls such that each infinite-order element is "cut" by a translate of one of these walls.

The following remarkable theorem of Bergeron-Wise gives a practical criterion for doing this:

Theorem 3.7 (Bergeron-Wise)**.** *Let G be a hyperbolic group and let ∂G be its Gromov boundary. Suppose that for all distinct $p, q \in \partial G$, there exists a quasiconvex codimension–1 subgroup $H \leq G$ such that p, q lie in H–distinct components of $\partial G - \partial H$. Then G acts properly and cocompactly on a CAT(0) cube complex.*

We won't go through the proof; instead we'll do an example and discuss how to apply the theorem.

Remark 3.8. The key point in the proof is that G acts as a *uniform convergence group* on ∂G. This means the following: let T be the subspace of ∂G^3 consisting of tuples (p, q, r) with p, q, r all distinct. Equip T with the subspace topology coming from ∂G^3. Then the action of G on T is proper and cocompact (this amazing fact *characterises* hyperbolic groups; see [Bow98]).

The idea for why this is true: the points p, q, r determine an ideal geodesic triangle in G. By hyperbolicity, this triangle has a coarsely well-defined "median" in G...

Remark 3.9 (How and why to use the Bergeron-Wise theorem)**.** Let G be a hyperbolic group. Bergeron-Wise won't help you find codimension–1 subgroups. For that, you have to appeal to the specific nature of the group. So, the main task is to come up with a general procedure for constructing quasiconvex walls in G.

Now, for each bi-infinite geodesic γ, you need to use your procedure to build a quasiconvex wall W such that for all sufficiently large positive n, $\gamma(n)$ is on the "right side" of W, and $\gamma(-n)$ is on the "left side". The advantage is that you get to build a new wall for each γ, without worrying about only building finitely many orbits of walls.

Now the theorem gives you the desired cubulation. Many of the examples mentioned earlier (hyperbolic 3–manifold groups, free-by-cyclic groups, etc.) have been cubulated in this way.

More exercises

(1) Tile a closed surface S of genus $g \geq 2$ by squares, so that links (which are all circles) have length at least 4. This gives a nonpositively curved square complex X homeomorphic to S. Do this in such a way that (1) no hyperplane in X crosses itself; (2) more generally, if e, f are edges dual to the same hyperplane, then e, f don't have a common vertex. How many squares do you need, in terms of g?

(2) Let X be a CAT(0) cube complex, and let H, H' be crossing hyperplanes. Let $\overleftarrow{H}, \overleftarrow{H}'$ be associated (necessarily intersecting) halfspaces. Form a new cube complex Y by passing to the largest subcomplex of $X - \overleftarrow{H} \cap \overleftarrow{H}'$. Show that Y is CAT(0).

(3) Let X be a CAT(0) cube complex and let $\mathcal{H}$ be a subset of the hyperplanes of X. Show that the CAT(0) cube complex Y dual to the wallspace $(X^{(0)}, \mathcal{H})$ is a quotient of X that can be realised topologically by collapsing each $\mathcal{N}(H), H \in \mathcal{H}$, which we view as $H \times [-\frac{1}{2}, \frac{1}{2}]$: fix a contraction of $[-\frac{1}{2}, \frac{1}{2}]$ to a point and perform it fibrewise to collapse $\mathcal{N}(H)$. Show that for any convex subcomplex Z of Y, the preimage of Z under $X \to Y$ is convex, and that the preimage of any hyperplane is a hyperplane.

(4) Let γ be a 1–skeleton geodesic in a D–dimensional CAT(0) cube complex, and let $\ell \geq 0$. Suppose that $|\gamma| \geq R(\ell, D+1)$. Show that γ is crossed by at least ℓ disjoint hyperplanes. (Here $R(\ell, D+1)$ denotes the Ramsey number associated to $\ell, D+1$.)

(5) Let X be a CAT(0) cube complex. Prove that if $\dim X < \infty$, then (X, d_2) is quasi-isometric to $(X^{(1)}, \mathsf{d}_1)$.

(6) Let X be a CAT(0) cube complex. Prove that $X^{(1)}$ is hyperbolic if and only if the following holds. There exists a constant k such that whenever there is an embedding $[0, p] \times [0, q] \to X$ which is an isometric embedding on 1–skeleta, we have $\min\{p, q\} \leq k$. (Here $p, q \in \mathbb{N}$ an $[0, p]$ and $[0, q]$ are given the obvious 1–dimensional cubical structures.) Hint: since $X^{(1)}$ is median, to prove hyperbolicity, you need to show that if γ, γ' are geodesics with common endpoints, then they are (uniformly) Hausdorff-close.

(7) Let X be a CAT(0) cube complex and let $\mathcal{H}$ be the set of hyperplanes. Suppose that there is a finite set F and a map $c : \mathcal{H} \to F$ such that, for all $f \in F$, the hyperplanes in $c^{-1}(f)$ are all disjoint. Prove that $X^{(1)}$ embeds isometrically in the product of $|F|$ trees.

(8) Prove the following fact, mentioned earlier in the notes: let X be a CAT(0) cube complex and let $\mathcal{H}$ be the set of hyperplanes. Suppose we can write $\mathcal{H} = \mathcal{A} \sqcup \mathcal{B}$, where every hyperplane in $\mathcal{A}$ crosses every hyperplane in $\mathcal{B}$. Then $X \cong A \times B$, where A, B are CAT(0) cube complexes. Moreover, each hyperplane in $\mathcal{A}$ has the form $H \times B$, where H is a hyperplane of A (and a similar description holds for $\mathcal{B}$).

(9) Let X be a CAT(0) cube complex with $|X^{(0)}| < \infty$. Let G be a group acting on X by cubical automorphisms. Prove that G fixes a point in X (ideally without using the CAT(0) metric!). Deduce that, if Y is a proper CAT(0) cube complex on which the group G acts properly and cocompactly by cubical automorphisms, then G contains finitely many conjugacy classes of finite subgroups.

References

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

[BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[Bow98] Brian H. Bowditch. A topological characterisation of hyperbolic groups. *J. Amer. Math. Soc.*, 11(3):643–667, 1998.

[Bri91] Martin R. Bridson. Geodesics and curvature in metric simplicial complexes. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 373–463. World Sci. Publ., River Edge, NJ, 1991.

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

[Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.

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

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21(4):851–891, 2011.

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

[Ger97] V. N. Gerasimov. Semi-splittings of groups and actions on cubings. In *Algebra, geometry, analysis and mathematical physics (Russian) (Novosibirsk, 1996)*, pages 91–109, 190. Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk, 1997.

[Hag07] Frédéric Haglund. Isometries of cat (0) cube complexes are semi-simple. *arXiv preprint arXiv:0705.3386*, 2007.

[Hag14] Mark F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.

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

[HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[HW14] G. C. Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. *Compos. Math.*, 150(3):453–506, 2014.

[HW15] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the general case. *Geom. Funct. Anal.*, 25(1):134–179, 2015.

[KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

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

[MZ16] Ashot Minasyan and Pavel Zalesskii. Virtually compact special hyperbolic groups are conjugacy separable. *Comment. Math. Helv.*, 91(4):609–627, 2016.

[Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004.

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

[NR98a] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.

[NR98b] Graham A. Niblo and Martin A. Roller. Groups acting on cubes and Kazhdan's property (T). *Proc. Amer. Math. Soc.*, 126(3):693–699, 1998.

[NR03] G. A. Niblo and L. D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.

[OW11a] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than 1/6. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.

[OW11b] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than 1/6. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.

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

[SW05] Michah Sageev and Daniel T. Wise. The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.*, 37(5):706–710, 2005.

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

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

[Wis18] Daniel T Wise. The structure of groups with a quasiconvex hierarchy. *Preprint*, 2018.

[Wri12] Nick Wright. Finite asymptotic dimension for CAT(0) cube complexes. *Geom. Topol.*, 16(1):527–554, 2012.