On CAT(0) cube complexes

Graham A. Niblo

University of Southampton

May 2011

The Mobius action of $PGL_2(\mathbb{Z})$ on the upper half plane

$PGL_2(\mathbb{Z})$ acts by Mobius transformations on the upper half plane model of the hyperbolic plane so that

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}(z) = \begin{cases} \frac{az+b}{cz+d} & \text{if } ad - bc = 1, \\ \frac{a\bar{z}+b}{c\bar{z}+d} & \text{if } ad - bc = -1. \end{cases}$$

Matrix	$S_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	$S_2 = \begin{bmatrix} -1 & 1 \\ 0 & 1 \end{bmatrix}$	$S_3 = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$
Isometry	$s_1 : z \mapsto 1/\bar{z}$	$s_2 : z \mapsto 1 - \bar{z}$	$s_3 = z \mapsto -\bar{z}$
Fixed set	$\{z \mid z = 1\}$	$\{z \mid Re(z) = 1/2\}$	$\{z \mid Re(z) = 0\}$

Picturing the action

Now we can examine the action of $PGL_2(\mathbb{Z})$ by studying the group of isometries generated by the three reflections s_1, s_2, s_3.

1. The mirror lines ℓ_1, ℓ_2 intersect at the point $P = (1 + \sqrt{3}i)/2$ at an angle of $2\pi/6$, so the two reflections s_1, s_2 generate the finite dihedral group D_6.
2. The reflection lines ℓ_1, ℓ_3 intersect at right angles at the point $Q = i$ and therefore generate the finite dihedral group D_4.
3. The lines ℓ_2, ℓ_3 only share the ideal point ∞ and therefore the isometries s_2, s_3 generate a parabolic subgroup isomorphic to the infinite dihedral group D_∞.
4. The three lines ℓ_1, ℓ_2, ℓ_3 cut the plane into six regions, one of which, labelled C in the following diagram, is bounded by all three lines. This region forms a fundamental domain for the action of $PGL_2(\mathbb{Z})$.

The Coxeter complex

The tiling of the plane by images of C is well known and appears in a paper by Fricke which appeared in 1890.

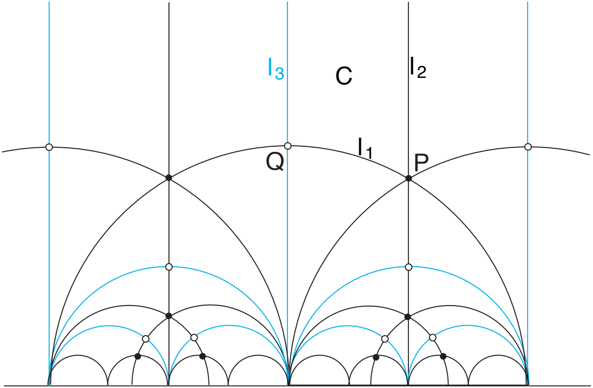

The cubing associated to the Coxeter complex

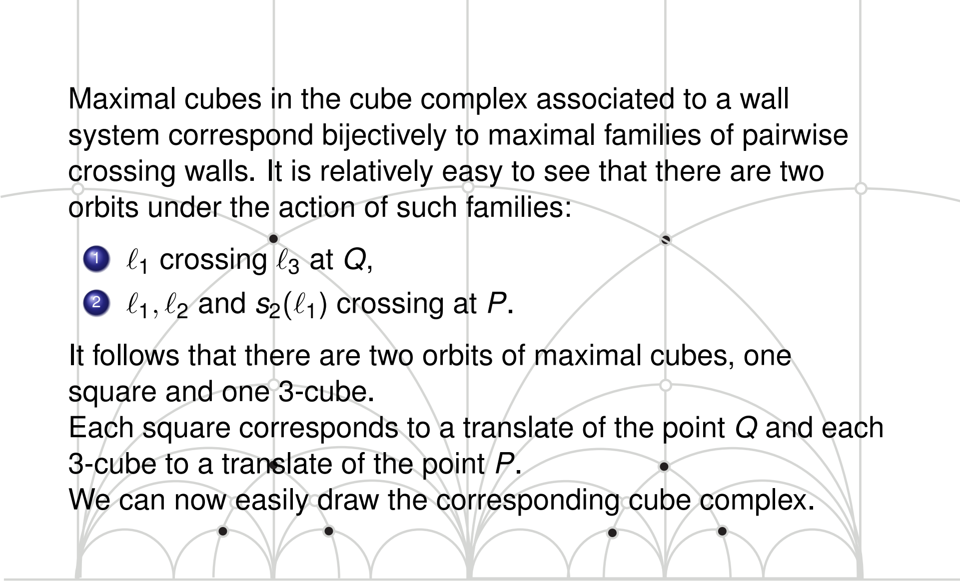

Maximal cubes in the cube complex associated to a wall system correspond bijectively to maximal families of pairwise crossing walls. It is relatively easy to see that there are two orbits under the action of such families:

1. ℓ_1 crossing ℓ_3 at Q,
2. ℓ_1, ℓ_2 and $s_2(\ell_1)$ crossing at P.

It follows that there are two orbits of maximal cubes, one square and one 3-cube.

Each square corresponds to a translate of the point Q and each 3-cube to a translate of the point P.

We can now easily draw the corresponding cube complex.

The cubing associated to $PGL_2(\mathbb{Z})$

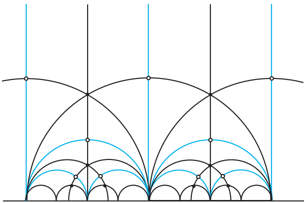

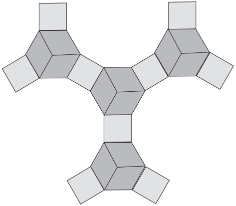

The group acts with two orbits of hyperplanes.

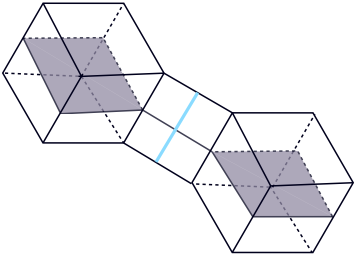

The cubing associated to the blue walls

Note that this is not the Bass Serre tree of a splitting since the group acts with involutions on the edges.

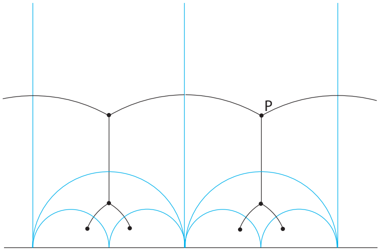

The tree dual to the configuration of lines

The Bass-Serre tree associated to the blue walls

Taking the barycentric subdivision repairs this and we obtain a splitting of $PGL_2(\mathbb{Z})$ as the amalgamated free product of the stabilisers of P and Q over their intersection.

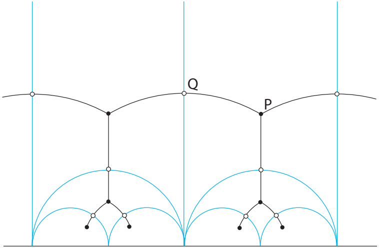

The tree dual to the configuration of lines

$$PGL_2(\mathbb{Z}) = D_4 \underset{\mathbb{Z}_2}{*} D_6 = \langle S_3, S_1 \rangle \underset{\langle S_1 \rangle}{*} \langle S_1, S_2 \rangle.$$

The cubing associated to the black walls

Note that while this cubing is not a tree it is quasi-isometric to the Bass-Serre tree.

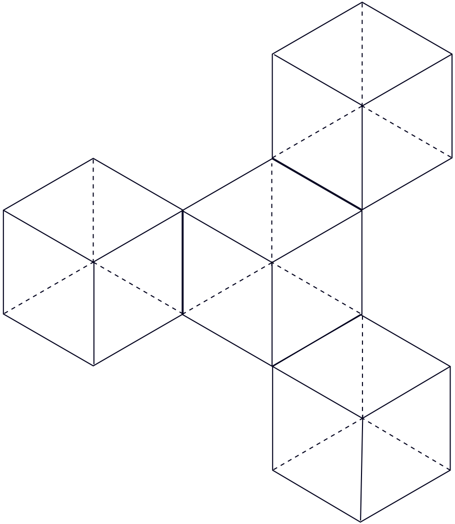

An alternative way to obtain a cubing - The Charney-Davis method

The complex of groups picture for PGL(2, Z) and the (3,3,3,) triangle group

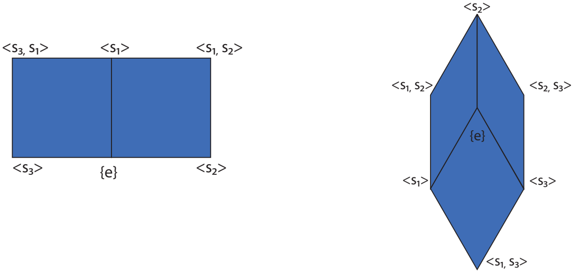

The Charney-Davis squaring for $PGL(2,\mathbb{Z})$

The Charney-Davis CAT(0) squaring for PGL(2, Z)

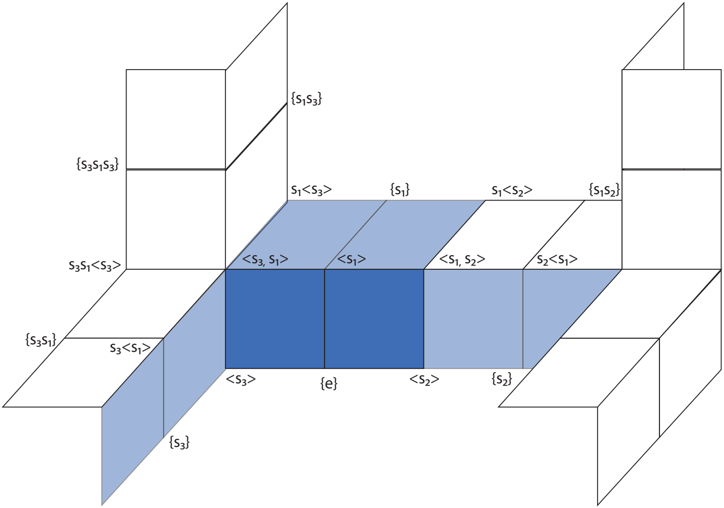

The Charney-Davis cubing is not always CAT(0)

Positive, negative and 0 curvature in the Charney-Davis CAT(0) squaring for the (3,3,3) triangle group

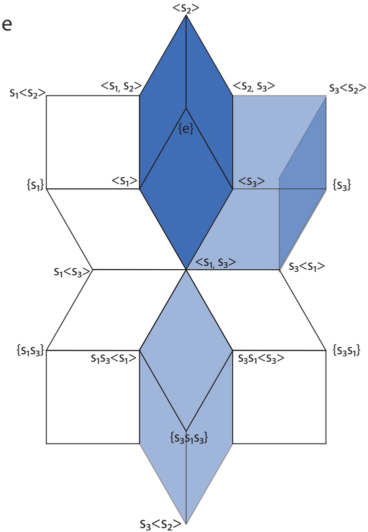

The dark blue block is a fundamental domain for the action and the two pale blue blocks are translates of it by the elements s_3 and $s_3 s_1 s_3$.

Let (W, R) be a finite rank Coxeter system and for generators $r, s \in R$ let $m(r, s)$ denote the order of the product rs in W, where by convention $m(r, r) = 1$ and $m(r, s) = \infty$ when the product has infinite order. The standard presentation for W is given by:

$$W \cong \langle R \mid (rs)^{m(r,s)} = 1, \forall r, s \in R \rangle.$$

The linear representation of the Coxeter group W

Let V denote a real vector space with basis $\{\vec{r} \mid r \in R\}$ in one to one correspondence with the elements of the generating set R. A symmetric bilinear form on V is defined by setting

$$\langle \vec{r}, \vec{s} \rangle = -\cos\left(\frac{\pi}{m(r,s)}\right) \quad \text{for all } r, s \in R.$$

The action of W on V is defined by

$$s\vec{r} = \vec{r} - 2\langle \vec{s}, \vec{r} \rangle \vec{s} \quad \text{for all } r, s \in R.$$

Definition

The ***roots*** of the Coxeter system (W, R) are the elements of the set $\Phi := \{w\vec{r} \mid w \in W, r \in R\} \subseteq V$. The elements, $\vec{r}$, of the basis for V are called ***simple roots***.

Positive roots

A root $w\vec{r}$ may be expressed uniquely as a real linear combination of the simple roots, and the coordinates will either all be positive, or all be negative.

We say that the root is positive (respectively negative) according to which of these alternatives occurs.
We denote the set of positive roots by Φ^+, and the set of negative roots by Φ^-. The set of roots Φ is preserved under the involution $-I : V \to V$ which sends v to $-v$.

Key fact

For $w \in W, r \in R$, $w\vec{r} \in \Phi^+$ if and only if $\ell(wr) > \ell(w)$.

Brink and Howlett's dominance relation (adapted)

Definition

Let $\alpha, \beta \in \Phi$. We say that α dominates β (and write $\alpha \,\mathrm{dom}\, \beta$) if

1. $\langle \alpha, \beta \rangle \geq 1$, and
2. there exists $g \in W$ such that $g\alpha \in \Phi^+$ and $g\beta \in \Phi^-$.

This relation is a partial order on Φ which is reversed by the involution $-I$. Let $\preceq$ denote the reverse partial order, that is $\alpha \preceq \beta$ if and only if $\beta \,\mathrm{dom}\, \alpha$.

The triple $(\Phi, \preceq, -I)$ is a pocset.

Halfspaces in W

Definition

For adjacent vertices u, ur in the natural Cayley graph Γ_W, define
$$\mathfrak{H}(u, ur) = \{w \in W \mid d(w, u) < d(w, ur)\}.$$

It is a standard (non-trivial) fact that there are no odd length loops in Γ_W, so $\mathfrak{H}(ur, u) = \mathfrak{H}(u, ur)^*$.

Definition

$$\mathcal{H} = \{\mathfrak{h}(u, ur) \mid u \in W, r \in R\}.$$

The triple $(\mathcal{H}, \subseteq, *)$ is also a pocset.

The triples $(\mathcal{H}, \subseteq, *)$ and $(\Phi, \preceq, -I)$ are isomorphic

Definition

Given a root $w\vec{r} \in \Phi$, set $\mathfrak{h}(w\vec{r}) = \{u \in W \mid u^{-1}w\vec{r} \in \Phi^+\}$.

Lemma

For all $w \in W, r \in R$, we have $\mathfrak{h}(w\vec{r}) = \mathfrak{h}(w, wr)$.

Proof.

$$\begin{aligned} g \in \mathfrak{h}(w\vec{r}) &\Leftrightarrow g^{-1}w\vec{r} \in \Phi^+ \\ &\Leftrightarrow \ell(g^{-1}wr) > \ell(g^{-1}w) \quad \text{Key fact above} \\ &\Leftrightarrow d(1, g^{-1}wr) > d(1, g^{-1}w) \\ &\Leftrightarrow d(g, wr) > d(g, w) \\ &\Leftrightarrow g \in \mathfrak{h}(w, wr) \end{aligned}$$

□

Given a word $w = r_1 \dots r_n \in R^*$ we write $\vec{w_i} = r_1 \dots r_{i-1}\vec{r_i}$, and set $\vec{w} = \vec{w_n}$.

Lemma

Let $\rho, \sigma, \tau \in \Phi$, with $\rho = \vec{w}$.

1. $\mathfrak{h}(\rho)^* = \mathfrak{h}(-\rho)$.
2. *If $\rho, \sigma \in \Phi^+$ and $\mathfrak{h}(\sigma) \subseteq \mathfrak{h}(\rho)$, then $\sigma = \vec{w_{i+1}}$ for some $1 \leq i \leq \ell(w)$.*
3. *If $\mathfrak{h}(\tau) \subseteq \mathfrak{h}(\rho)$, then there are only finitely many $\sigma \in \Phi$ with $\mathfrak{h}(\tau) \subseteq \mathfrak{h}(\sigma) \subseteq \mathfrak{h}(\rho)$.*
4. *If $\rho \neq \sigma$ then $\mathfrak{h}(\rho) \neq \mathfrak{h}(\sigma)$.*

Lemma

The map f which sends $w\vec{r} \in \Phi$ *to* $\mathfrak{h}(w\vec{r}) = \mathfrak{h}(w, wr)$ *defines an isomorphism between the partially ordered sets with involution* $(\Phi, \preceq, -I)$ *and* $(\mathcal{H}, \subseteq, *)$.

Proof.

The map f is clearly surjective.
Injectivity follows from part 4 of the previous Lemma.
That $\mathfrak{h}(\alpha) \subseteq \mathfrak{h}(\beta)$ if and only if $\alpha \preceq \beta$ for all $\alpha, \beta \in \Phi$ follows from Lemma 2.3 of Brink and Howlett.
$\mathfrak{h}(-\alpha) = W - \mathfrak{h}(\alpha)$ so $f \circ (-I) = * \circ f$.

□

Corollary

1. *The triple* $(\mathcal{H}, \subseteq, *)$ *is a discrete pocset. Hence, there is a cube complex associated to the triple* $(\mathcal{H}, \subseteq, ^*)$ *and each of the components is CAT(0).*
2. *The Cayley graph* Γ_W *maps equivariantly into a component of the cube complex, so G preserves a component.*

The cube complex is finite dimensional

Lemma

Given a finitely generated Coxeter system (W, R), there is a number $N = N(W, R)$, such that any collection of more than N halfspaces contains a nested pair.

Proof.

It suffices to show that there is a positive integer N such that, given any subset $S \subseteq \Phi$ of cardinality greater than or equal to N there are roots $\alpha, \beta \in S$ such that $|\langle \alpha, \beta \rangle| \geq 1$.
If not, let $B \subseteq S$ be a basis for the subspace of V spanned by S. So $|B| \leq |R|$, the rank of W. If $|S| > N'|B|$, then there will be two roots $\alpha, \beta \in S$ such that $\langle \alpha, \gamma \rangle = \langle \beta, \gamma \rangle$ for all $\gamma \in B$. Write $\alpha = \sum_{\gamma \in B} \alpha_\gamma \gamma$ and $\beta = \sum_{\gamma \in B} \beta_\gamma \gamma$. Then,
$1 = \langle \alpha, \alpha \rangle = \sum_{\gamma \in B} \alpha_\gamma (\gamma, \alpha) = \sum_{\gamma \in B} \alpha_\gamma (\gamma, \beta) = \langle \alpha, \beta \rangle$ which gives a contradiction. Hence we may put $N = N'|B|$. □

Local finiteness

The parallel walls theorem of Brink and Howlett showed that there is a bound on the number of walls in a collection that do not dominate some other wall in the collection. This implies local finiteness.
The strong parallel walls conjecture gives a similar global bound on the size of a link in the cubing associated to a given Coxeter group. It was proved by Caprace and others.

The cubing associated to $PGL_2(\mathbb{Z})$

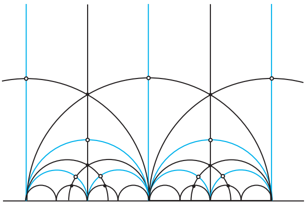

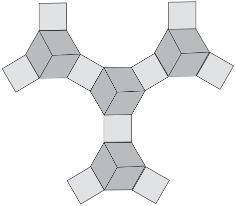

The group acts with two orbits of hyperplanes.

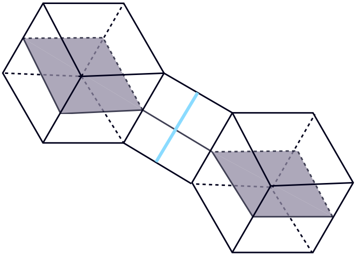

The pocset splits as a disjoint union of three families of half spaces, where each family is linearly (totally) ordered and each hyperplane from one family is transvers to each hyperplane in the other two families.

It follows that the induced cubing for each of the three families is a line and that the three lines form a direct product. So the induced cubing is the natural integer lattice cubing on $\mathbb{R}^3$.

The cubing associated to the $3, 3, 3$ triangle group is $\mathbb{R}^3$ with its usual cubing.

In this case the action is NOT co-compact.

The Coxeter complex can be seen as a slice of this picture, orthogonal to the line $x = y = z$.

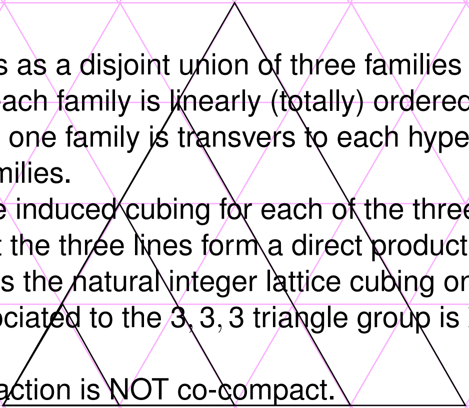

Co-compactness of the action

According to Williams the orbits of 3-cells correspond to conjugacy classes of 3-generator reflection subgroups such that the exponents of the rotations are finite.
According to Caprace there are finitely many such classes if and only if the Coxeter group does not contain a Euclidean crystallographic triangle group.

Certain Artin groups can also be cubed.
In the case of right-angled Artin groups the procedure is elementary and the group acts freely on the CAT(0) cubing.
For Artin groups of type FC Charney and Davis showed how to construct a CAT(0) cubing on which the group acts isometrically, but with large stabilisers (the special subgroups of finite type).

Given a subset B of a group G with $e \in B$ we may build a G-pocset from it as follows:

Let $\mathcal{H} = \{gB \mid g \in G\} \cup \{gB^* \mid g \in G\}$. The involution $*$ denotes complementation.

Given an element $g \in G$ there is a natural vertex v_g in the cubing defined by setting $v_g(\{\mathfrak{h}, \mathfrak{h}^*\}$ equal to the halfspace containing it, and so there is a map $g \to v_g$ from G to X

When is there an invariant component?

Let (X, d) denote the component of the cubing containing v_e, equipped with the edge metric d on its vertex set.

Lemma

Let H be the left stabiliser of the pair $\{B, B^*\}$ in G. Then $d(v_g, v_e) < \infty$ if and only if the symmetric difference $B + Bg^{-1}$ is contained in finitely many cosets HF. (We say that B is H-almost invariant in G.)

Proof.

$d(v_g, v_e)$ is the number of translates kB such that exactly one of e, g belongs to kB. This is the number of translates kB such that $k^{-1} \in B + Bg^{-1}$. If this is contained in finitely many right cosets HF then k lies in finitely many left cosets $F^{-1}H$ so there are only finitely many translates $F^{-1}B$ separating v_g from v_e as required.
Reversing the argument gives the converse. □

When is there an interesting invariant component?

Lemma

The orbit Gv_e is bounded if and only if B or B^* is H-finite.

Corollary

If $B \subset G$ has left stabiliser H, is H-almost invariant, and neither B nor B^* is H-finite, then G acts with an unbounded orbit on a CAT(0) cube complex.

Definition

We say that a countable group G is semi-splittable over a subgroup H if it admits a subset $B \subset G$ with left stabiliser H such that B is H-almost invariant, and neither B nor B^* is H-finite.

Generalising Bass-Serre theory

Theorem (Sageev/Roller-N/Gerasimov)

Let G be a countable group, $H < G$. Then G is semi-splittable over H if and only if it admits an action on a CAT(0) cube complex satisfying the following:

- *The action has a single orbit of hyperplanes $G\bar{\mathfrak{h}}$.*
- *The action has no fixed points.*
- *H is the left stabiliser of a hyperplane.*

The geometric meaning of H-almost invariance

Let G be a finitely generated subgroup, Γ a locally finite Cayley graph for G and $B \subset G$ with left stabiliser H

Definition

The coboundary of a set of vertices B in a graph Γ is the set of edges of Γ with precisely one end point in B.

Lemma (Scott-Houghton)

B is H-almost invariant in G if and only if the set of cosets HB has finite coboundary in the quotient $H \setminus \Gamma$. It is proper if and only if neither HB nor HB^* is finite.

Proof.

Bg is the right translate of B by a given G and this consists of points within $|g|$ of B, so H-finiteness of $B + Bg$ asserts that there are only finitely many vertices in $H \setminus \Gamma$ which are not in HB but are within $|g|$ of it. □

An aside

The right action of G on $H \setminus G$ induces a linear action on $\ell^2(H \setminus B)$ by post and premultiplication with projection onto HB.
When B is H-almost invariant we can define an index for each operator ρ_g yielding a homomorphism:

$$G \to \mathbb{Z}$$
$$g \mapsto \text{rk}(\text{Ker})(\rho_g) - \text{rk}(\text{Coker}(\rho_g))$$

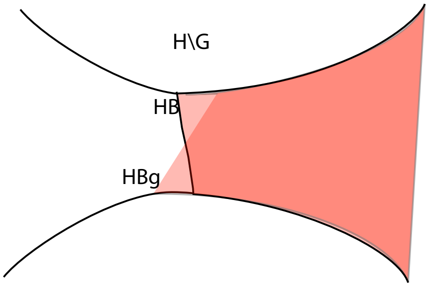

A connection with K-theory

When $G = \mathbb{Z}$ and $B = \mathbb{N}$ the operators ρ_g form the Toeplitz algebra $\mathcal{T}$. The index map induces the Toeplitz extension which plays a key role in the proof of Bott periodicity for algebraic K-theory:

$$1 \to \mathcal{K} \to \mathcal{T} \to C(S^1) \to 1.$$

A generalisation of this by Pimsner, Voicelscu and others allows one to compute the K-theory of the reduced C^* algebra of a graph product of groups in terms of the K-theory of the vertex stabilisers.

Constructing almost invariant sets from actions on cubings

Given an action of G on a CAT(0) cube complex X choose a vertex v adjacent to a halfspace $\mathfrak{h}$.
Now set $B = \{g \in G \mid gv \in \mathfrak{h}\}$.
Then B is H-almost invariant where H is the left stabiliser of the corresponding hyperplane $\overline{\mathfrak{h}}$. B is proper if and only if the orbit Gv is unbounded in the cube complex.

Example

Let G be the free group on generators a, b acting on its standard Cayley graph. Let e denote the edge joining e to a and $\overline{\mathfrak{h}}$ be its midpoint. Then the set $B = \{g \in G \mid ge$ is not separated from e by $\overline{\mathfrak{h}}\}$ is a proper $\{1\}$-almost invariant set. It is equal to all elements whose first syllable is not a positive power of a.

Hyperplanes

The complexity of a hyperplane, measure by its intersections with its translates in some sense measures how far the corresponding almost invariant set is from yielding a splitting of the group. If the hyperplane is a point then the group splits over its stabiliser.

1. Each hyperplane $\mathcal{H}$ is itself a cube complex and is contained in a totally geodesic subspace of the form $\mathcal{H} \times [0, 1]$.
2. It is therefore a CAT(0) cube complex and has hyperplanes of its own.
3. Given a hyperplane $\mathcal{K}$ of $\mathcal{H}$ the subspace $\mathcal{K} \times [0, 1]$ hits each edge of a cube which meets it orthogonally at its midpoint.
4. It follows that the hyperplanes in $\mathcal{H}$ are precisely the intersection of $\mathcal{H}$ with the other hyperplanes of X.

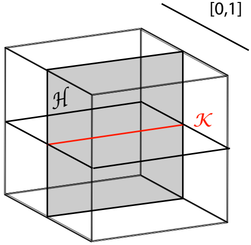

Lemma

Let X be a CAT(0) cube complex and G be a group G acting on X with a single orbit of hyperplanes. Let H be the stabiliser of a hyperplane $\bar{\mathfrak{h}}$ *and let* $Sing_G(\bar{\mathfrak{h}})$ *be the set* $\{g \in G \mid g\bar{\mathfrak{h}}$ *is transverse to* $\bar{\mathfrak{h}}\}$. *Then* $Sing_G(\bar{\mathfrak{h}})$ *is of the form HFH for some subset* $F \subset G$.

Proof.

If $f\bar{\mathfrak{h}}$ is transverse to $\bar{\mathfrak{h}}$ then $fh_1\bar{\mathfrak{h}}$ is transverse to $h_2^{-1}h$ for any $h_1, h_2 \in H$, so $\mathrm{Sing}_G(\bar{\mathfrak{h}})$ is invariant under left and right multiplication by H. Hence $\mathrm{Sing}_G(\bar{\mathfrak{h}}) = HFH$ for some subset $F \subset G$. □

The finite hyperplane theorem

Theorem

Let G be a finitely generated group, and H be a finitely generated subgroup with $e(G, H) \geq 2$. Let X be the Sageev cubing and $\overline{\mathfrak{h}}$ the hyperplane. Then $Sing_G(\overline{\mathfrak{h}})$ is a finite union of double cosets HFH and the hyperplane $\overline{\mathfrak{h}}$ is finite if and only if $Sing_G(\overline{\mathfrak{h}})$ lies in the subgroup

$$Comm_G(H) = \{g \in G \mid H^g \cap H \text{ has finite index in } H \text{ and } H^g\}.$$

Proof of the finite hyperplane theorem - part 1

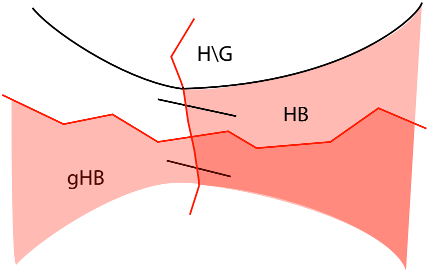

We use the finite generation of H to arrange for the coboundary of HB to be connected so that two translates cross if and only if their coboundaries intersect.
Since the couboundary is H-finite there can be only finitely many translates of B up to the action of H which cross it. The double cosets HfH in $\mathrm{Sing}_G(\mathcal{H})$ count this number.

Proof of the finite hyperplane theorem - part 2

Recall that hyperplanes transverse to $\bar{\mathfrak{h}}$ correspond bijectively with cosets gH in HFH.
Since $HFH \subset \mathrm{Comm}_G(H)$ each double coset HfH decomposes as finitely many left cosets of H:

$$HfH = (ff^{-1})HfH = f(H^f H) = fS_f(H \cap H^f)H = fS_f H.$$

Applications of the finite hyperplane theorem

If G is finitely generated and has a finitely generated subgroup H with $e(G, H) \geq 2$ and satisfying any of the following conditions then the hyperplanes in the Sageev cubing are finite:

1. H is finite,
2. H is normal in G,
3. The commensurator of H in G is G.

The Plan

When G acts on a CAT(0) cube complex with finite hyperplanes then the complex is sufficiently tree-like for us to obtain an action on a tree. In particular we obtain:

Theorem (Stallings' theorem)

Let G be a finitely generated group. Then G splits over a finite subgroup if and only if $e(G) \geq 2$.

Outline proof of Stallings' theorem

1. $e(G) > 1 \Rightarrow G$ has an almost invariant set B with finite stabiliser.
2. Therefore G acts on a CAT(0) cube complex with finite hyperplane stabiliser.
3. The hyperplanes are compact and the cube complex is contractible.
4. So G acts on a simply connected 2-complex X which contains a finite essential track.
5. So X contains a G invariant pattern of minimal tracks.
6. So G acts on a tree.

- Steps 1 and 2 were outlined earlier on.
- Step 3 The cubing is contractible because it is CAT(0) so geodesics are unique and vary continuously with their endpoints. Hyperplanes are finite because $\mathrm{Comm}_G(H) = G$.
- Step 4 asserts that G acts on a simply connected 2-complex. For this we take a suitably refined triangulation of the 2-skeleton of the cubing.

Dunwoody patterns

Definition

A ***pattern*** in the 2-complex consists of a closed subset $\mathcal{P}$ in the complement of the 0-skeleton of X such that $\mathcal{P}$ meets each closed 1-simplex γ in a finite union of points all lying in the interior of γ, and each closed 2-simplex in a finite union of disjoint closed line segments, each joining two points in the boundary edges.

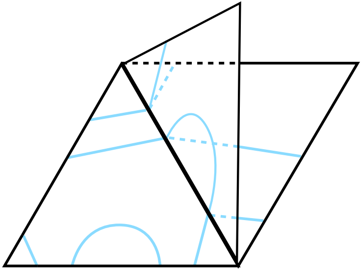

Dunwoody patterns

Definition

A ***pattern*** in the 2-complex consists of a closed subset $\mathcal{P}$ in the complement of the 0-skeleton of X such that $\mathcal{P}$ meets each closed 1-simplex γ in a finite union of points all lying in the interior of γ, and each closed 2-simplex in a finite union of disjoint closed line segments, each joining two points in the boundary edges.

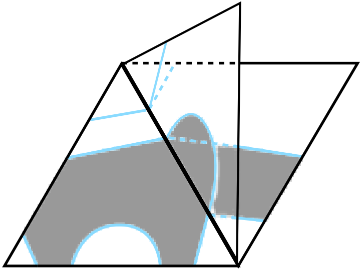

Note that patterns are locally separating.

Notes on patterns

1. This definition does not quite agree with that given by Dunwoody since we allow the line segments to join two points in the same edge of a 2-simplex.
2. A connected pattern is called a ***track***

Theorem

If X is a simply connected 2-complex then any track is a separating set.

Proof

Given a track τ and an edge e define $w(e) = |\tau \cap e|$. Reducing mod 2, we obtain a $\mathbb{Z}_2$-valued 1-cochain z.

The weight function w and cochain z

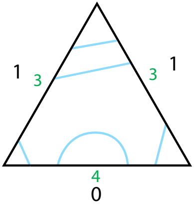

Given a track τ and an edge e define $w(e) = |\tau \cap e|$. Reducing mod 2, we obtain a $\mathbb{Z}_2$-valued 1-cochain z.

The weight function w and cochain z

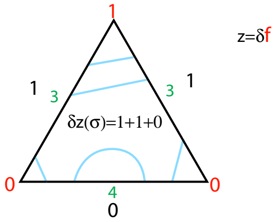

- $\delta z(\sigma) = 0$ for each 2-simplex σ.
- The corresponding class in $H^1(X, \mathbb{Z}_2)$ is zero if and only if the pattern separates X
- If $\pi_1(X) = \{0\}$ then $H^1(X, \mathbb{Z}_2) = 0$ so every track separates.

Least weight tracks

Now assume that X is a simply connected, triangulated 2-complex (so every track separates)

Definition

A pattern $\mathcal{P}$ (or track) is said to be finite if it intersects only finitely many 1-simplices, and in this case we assign it a ***weight***, $\|\mathcal{P}\| = |\mathcal{P} \cap X^{(1)}|$. A finite pattern is said to be ***essential*** if at least two of its complementary components are unbounded. It is said to be least weight if it is essential and has least weight among all essential tracks.

The Big Idea: “Least weight tracks do not intersect”

Example

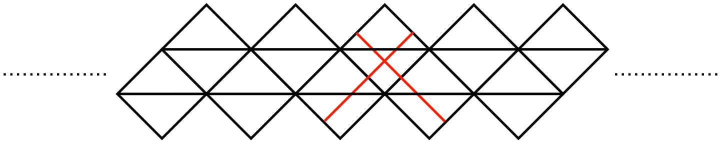

. . . unfortunately least weight tracks can (and often do) intersect.

Cut and paste

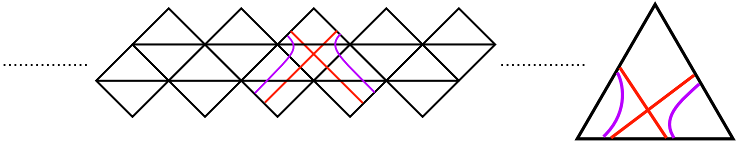

Lemma

Any least weight pattern on X consists of a single track which intersects each edge of X in at most one point and each 2-simplex in at most a single arc joining distinct edges of the simplex.

Given a pattern $\mathcal{P}$ and an edge e define $w(e) = |\mathcal{P} \cap e|$. If $\mathcal{P}$ has more than one component they each have weight less than $\mathcal{P}$ so they are all inessential. They cut the space in a tree pattern, and, by hypothesis two of the complementary components (corresponding to vertices of this tree) are unbounded. Choose an edge in the tree separating them This is a track in the pattern and its two complementary components are both unbounded. This contradicts the minimality of $\mathcal{P}$.

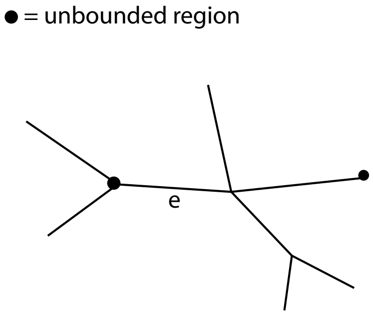

Reducing the characteristic function of $\mathcal{P}$ mod 2 we obtain a new function with values 0 and 1, from which we can build a new pattern $\mathcal{Q}$.

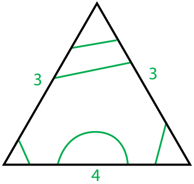

Reducing the characteristic function of $\mathcal{P}$ mod 2 we obtain a new function with values 0 and 1, from which we can build a new pattern $\mathcal{Q}$.

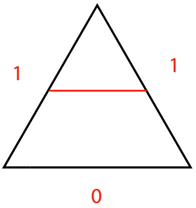

If $\mathcal{P}$ crossed some edge at least twice then the norm of $\mathcal{Q}$ is strictly less than $||\mathcal{P}||$, but their characteristic functions give identical elements of $H^1_{bdd}(X, \mathbb{Z}_2)$ by construction, so $\mathcal{Q}$ is also a least weight pattern, contradicting the minimality of $\mathcal{P}$. So $\mathcal{P}$ is a single track which crosses each 2-simplex in a single arc and which separates X into two unbounded components.

Theorem

Let σ and τ be least weight tracks in a simply connected 2-complex X which intersect transversely in the interior of the 2-cells of X. Then there are disjoint least weight tracks σ' and τ' in X such that $(\sigma \cup \tau) \cap X^{(1)} = (\sigma' \cup \tau') \cap X^{(1)}$.

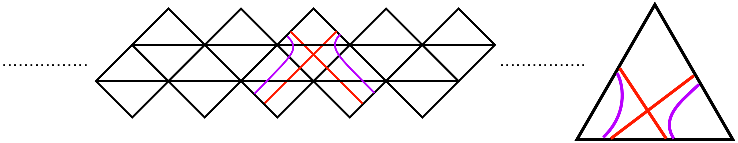

Proof

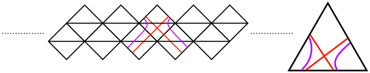

- Let $\|\sigma\| = \|\tau\| = n$ (every least weight pattern is a track of weight n).
- Taking the boundary of a small regular neighbourhood of the union of the two tracks gives a new pattern $\mathcal{P}$ with weight $2n$. There are “opposite corners” which are both unbounded so at least two of the tracks in $\mathcal{P}$ are essential.
- Hence they each have weight n and they are least weight.
- They are given by the canonical cut and paste sketched above.

Least length tracks

Now we "hyperbolise" our complex by uniformizing each simplex as an ideal hyperbolic triangle so that the midpoint of each edge is identified with the natural midpoint of each edge. We can replace each arc of our pattern with the corresponding geodesic arc in the hyperbolic metric and define the length of the pattern to be the sum of these lengths.
We say that a pattern is ***minimal*** if it is a least weight pattern of least length among all least weight patterns. (In particular it is a track.)

Theorem

Minimal tracks do not cross.

The Meeks Yau rounding trick

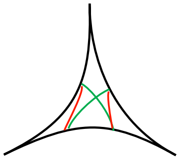

Cutting and pasting reduces length

How do we know minimal tracks exist?

We don't and they might not.
If the complex has a separating vertex which cuts it into at least two unbounded components then we obtain least weight tracks in the neighbourhood of that vertex. By moving them out to infinity we can make them arbitrarily short AND arbitrarily close to the separating vertex.

Shortening horocyclic tracks

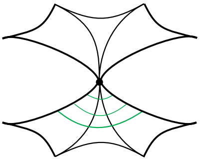

We cannot arbitrarily shorten non-horocyclic tracks

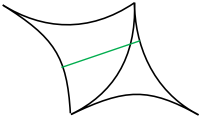

Since non-horocyclic tracks have to have length at least $\log(1 + \sqrt{2})$ they cannot be arbitrarily shortened. We can apply the Arzela-Ascoli theorem to prove that minimal tracks exist in this case.

Theorem

The Arzela-Ascoli theorem If C is a compact metric space and Γ is a separable metric space, then every sequence of equicontinuous maps $f_n : \Gamma \longrightarrow C$ has a subsequence that converges uniformly on compact subsets to a continuous map $f : \Gamma \longrightarrow C$.

Defining the compact metric space C

- Let $\ell = \inf\{\text{lengths of least weight tracks}\}$.
- If there are no horocyclic tracks $\ell \geq \log(1+\sqrt{2})$.
- Choose $\epsilon > 0$ and consider only least weight tracks of length $< \ell + \epsilon$.
- There is a neighbourhood of the ideal vertices which no such track can enter since it is too far from opposite sides of triangles.
- So we can remove an open neighbourhood of all ideal vertices and obtain a 2-complex Y which contains all sufficiently short least weight tracks.
- This 2-complex is locally finite since all infinite branching in the cube complex occurs at the vertices. This follows from the fact that hyperplanes are finite.

Now take a sequence of least weight tracks with length converging to ℓ and all of length $< \ell + \epsilon$.

- Since there is one hyperplane orbit and only finitely many edges on each hyperplane there are only finitely many edge orbits, so, translating by elements of *G* we can assume all the tracks in the sequence intersect some given edge *e*.
- The ball *C* of radius $\ell + \epsilon$ around *e* in *Y* is compact and contains all the tracks of length $< \ell + \epsilon$ which meet the edge.
- There are only finitely many possibilities for the topology of a least weight track in *C*. So we can take a subsequence consisting of tracks which are all homeomorphic to some graph Γ.

We now have a sequence of maps

$$f_n : \Gamma \longrightarrow C$$

defining tracks with lengths converging to ℓ.

Theorem

The Arzela-Ascoli theorem If C is a compact metric space and Γ is a separable metric space, then every sequence of equicontinuous maps $f_n : \Gamma \longrightarrow C$ has a subsequence that converges uniformly on compact subsets to a continuous map $f : \Gamma \longrightarrow C$.

Since length varies continuously with the tracks the map f defines a track of length ℓ as required.

Generalising Stallings' Theorem

All we needed to make this argument work is an essential action of G on a CAT(0) cube complex so that each hyperplane is finite.

Theorem

If G is a finitely generated group, H is a finitely generated subgroup with $e(G, H) \geq 2$ and $G = Comm_G(H)$ then G splits over a subgroup commensurable with H.

Example - The cubing of $PGL(2, \mathbb{Z})$ associated to the black walls

Note that while this cubing is not a tree it is quasi-isometric to the Bass-Serre tree.

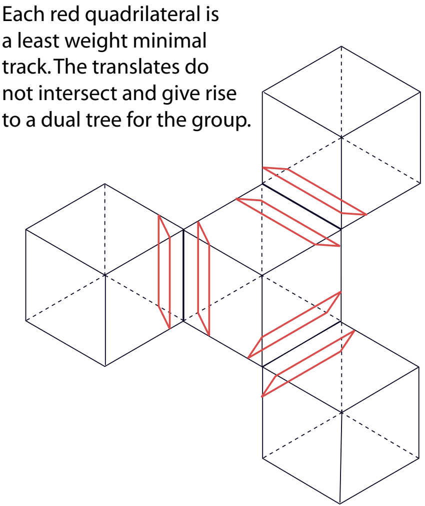

Maps between cubings

Let X and Y be CAT(0) cube complexes and $f : X \to Y$ be a cellular immersion.

The map induces a map on the set of half spaces which respects inclusion and complement, i.e., it induces a map on the dual pocset.

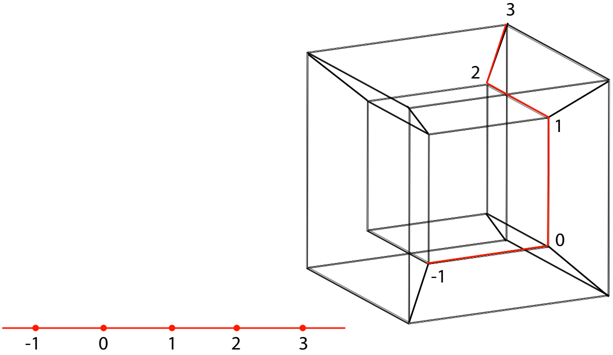

The Serre embedding

- Given a countable cube complex X choose a vertex v and for each vertex u define a $0, 1$ valued function on the hyperplanes as follows: $\chi_u(\overline{\mathfrak{h}}) = 1$ if and only if $\overline{\mathfrak{h}}$ separates u from v.
- This defines an embedding of X in the infinite unit cube, so the cube is a terminal object in the category.
- The inclusion induces a map between the pocset of half spaces - this is a covariant functor which "forgets" the partial order.

Extending the action

If G acts on X we obtain an induced action on the infinite cube: For a vertex χ of the infinite cube, a hyperplane $\overline{\mathfrak{h}}$ and $g \in G$ define

$$g\chi(\overline{\mathfrak{h}}) := \begin{cases} \chi(g^{-1}\overline{\mathfrak{h}}), & \text{if } \overline{\mathfrak{h}} \text{ does not separate } v, gv, \\ 1 - \chi(g^{-1}\overline{\mathfrak{h}}), & \text{if } \overline{\mathfrak{h}} \text{ does separate } v, gv. \end{cases}$$

Extending the action to Hilbert space

The unit ℓ^1 cube maps naturally into the Hilbert space of ℓ^2-functions on the hyperplanes and the action extends:

$$g\chi(\overline{\mathfrak{h}}) := \begin{cases} \chi(g^{-1}\overline{\mathfrak{h}}), & \text{if } \overline{\mathfrak{h}} \text{ does not separate } v, gv, \\ 1 - \chi(g^{-1}\overline{\mathfrak{h}}), & \text{if } \overline{\mathfrak{h}} \text{ does separate } v, gv. \end{cases}$$

This is affine since each element is acting by a unitary map composed with a translation, and the translation function is a cocycle:

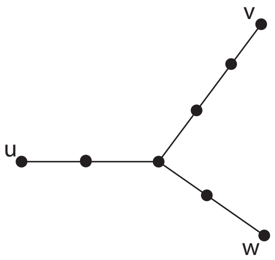

The embedding theorem

Theorem (GAN, Reeves, Roller)

Let G be a group acting on a CAT(0) cube complex X. For each vertex $v \in X^{(0)}$ there is a G-equivariant embedding from X to a Hilbert space on which G acts by affine isometries. The action on Hilbert space will have a global fixed point if and only if the action on the cube complex has a bounded orbit, and will be proper if and only if the action on the cube complex is proper.

Proof.

Equipping the vertices of the unit cube in Hilbert space with the ℓ^1-metric the embedding is an isometry on the 1-skeleton. □

The embedding is NOT an isometry in the geodesic metrics.

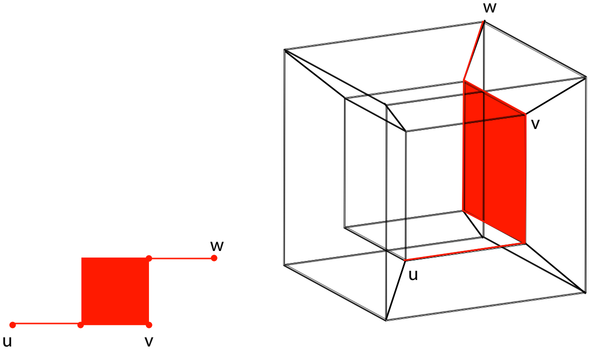

The geodesic distance between the vertices labelled *v* and *w* is 2 in the cube complex but only $\sqrt{2}$ in the Hilbert space.

Uniform embeddings

Recall that a map $f : X \to Y$ between metric spaces $(X, d_X), (Y, d_Y)$ is said to be a uniform embedding if there are increasing functions $S, T : \mathbb{R}^+ \to \mathbb{R}^+$ such that

$$d_Y(f(x), f(x')) \leq S(d_X(x, x'))$$
$$d_X(x, x') \leq T(d_Y(f(x), f(x')))$$

Our embedding of a cube complex in Hilbert space is contracting so we may choose S to be the identity. Since the cube complex is geodesic this shows that f_0 is a coarse Lipschitz map. Pythagoras tells us that the compression is of the order of the square root so it is a uniform embedding.

Fixed points from bounded orbits

Theorem

Let G be a group acting with a single, infinite orbit of hyperplanes on a CAT(0) cube complex X with a bounded orbit. Then G has a unique fixed point and it is a vertex of X.

If X is complete then the metric centre of any bounded orbit provides the required fixed point.

An infinite dimensional cube complex is not complete so we cannot apply this argument. There are two approaches we can try to use to remedy this:

Use the metric completion

- Take the metric completion $\overline{X}$ of X. This is a complete CAT(0) space.
- The G-action extends to $\overline{X}$ and so has a bounded orbit there.
- Since $\overline{X}$ is complete this action has a fixed point p.

Unfortunately we do not it a priori expect p to lie in the subspace X.

Use the Hilbert space

- Embed the cube complex in Hilbert space and extend the action.
- Since G has a bounded orbit in X it has a bounded orbit in Hilbert space.
- Since Hilbert space is complete G has a fixed point s.
- Since G acts transitively on the hyperplanes it acts transitively on the coordinates of s and the action of any given element converts some coordinate r to either r or $1 - r$.

- It follows that an invariant function takes only two values.
- Such a vector is not ℓ^2 unless every coordinate of the fixed vector is 0 or 1 and only finitely many of them are non-zero. Hence any fixed point represents a vertex of the unit cube.
- The fixed set is convex and it consists of vertices of the cube in Hilbert space.
- Since the set of vertices is disconnected this shows that G has a unique fixed point in the Hilbert space and this is a vertex s.

It remains to show that this vertex comes from a vertex of the original cube complex.

Now we put the two constructions together.
Since X embeds in Hilbert space and Hilbert space is complete we can extend the embedding equivariantly to all of $\overline{X}$.

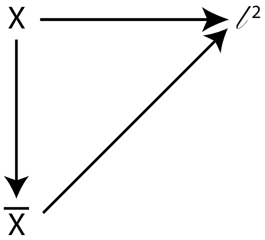

Any fixed point p in $\overline{X}$ must map to the fixed vertex s so it follows that s is the limit of a sequence of points in the image of X.

But each point in the sequence lies in the image of a closed cube of X and the distance from a vertex s of the unit cube to a face not containing it is 1.

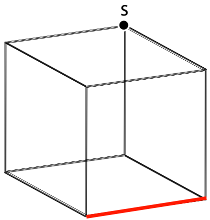

Hence almost all points in the sequence lie in closed cubes containing s. It follows that s lies in a closed cube from X as required.

Definition

A countable discrete group G has property (T) if and only if every affine isometric action on a separable Hilbert space has a global fixed point.

Example

$SL_3(\mathbb{Z})$ has Kazhdan's property (T).

Corollary

Every action of a property T group on a CAT(0) cube complex has a global fixed point.

The Haagerup property (a-T-menability)

Definition (The Haagerup property for geometric group theorists)

A countable discrete group G is said to be a-T-menable if it admits a metrically proper, affine isometric action on a separable Hilbert space.

Recall that an action is metrically proper if for any ball B of finite radius the set $\{g \in G \mid gB \cap B \neq \emptyset\}$ is finite.
It follows that any group acting properly on a CAT(0) cube complex is a-T-menable.

Examples of a-T-menable groups

- Countable amenable groups (Bekka, Cherix and Valette)
- Proper groups of isometries of real or complex hyperbolic space (e.g. free groups, surface groups etc)
- Coxeter groups (Bozejko, Januszkiewicz, Spatzier)
- Diagram groups, e.g., Thompson's group F (Farley)
- Finitely presented B(4)-T(4) or word-hyperbolic B(6) group (Wise)
- Right angled and special Artin groups (Deligne, Haglund and Wise)

Analytic methods in group theory

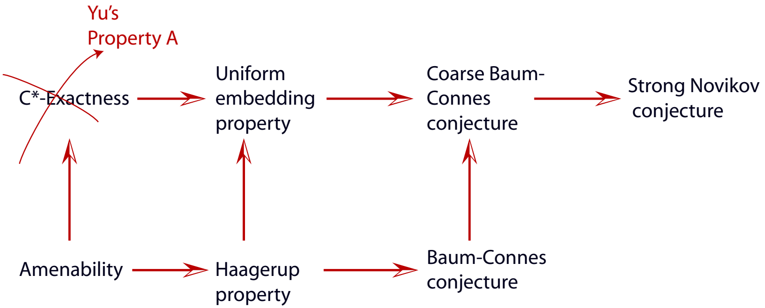

Følner's criterion

A countable discrete group G is amenable if and only if for each n there is a non-empty finite subset $A_n \subset G$ such that for all $R \in \mathbb{R}$

$$\frac{|gA_n + hA_n|}{|A_n|} \to 0$$

uniformly on $\{(g, h) \mid d(g, h) \leq R\}$.

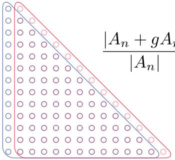

$$\frac{|A_n + gA_n|}{|A_n|} = \frac{4}{(n+1)}$$

A non-equivariant version of Følner's criterion

Yu's property A

A (discrete) metric space X has property A if and only if there exists a sequence $\{S_n\}$ of positive reals such that for each n and each $x \in X$ there is a function $f_{n,x} : X \to \mathbb{N}$ with support contained in $B_{S_n}(x)$ such that for all $R \in \mathbb{R}^+$

$$\frac{\|f_{n,x} - f_{n,x'}\|}{\|f_{n,x}\|} \to 0$$

uniformly on the set $\{(x, x') \mid d(x, x') \leq R\}$.

Property A for groups acting on finite dimensional CAT(0) cube complexes

Theorem (Brodzki, Campbell, Guentner, Niblo, Wright)

Let G be a group acting properly on a CAT(0) cube complex of finite dimension. Then G has property A.

Outline of the proof

We assign binomial weights satisfying certain conditions to intervals in the cube complex:

Definition

An interval in a CAT(0) cube complex is the set of all cubes lying between two vertices x, y.

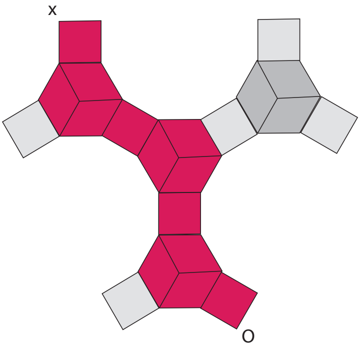

ANNALS OF MATHEMATICS
Vol. 51, No. 1, January, 1950

A DECOMPOSITION THEOREM FOR PARTIALLY ORDERED SETS

BY R. P. DILWORTH

(Received August 23, 1948)

1. Introduction

Let P be a partially ordered set. Two elements a and b of P are *camparable* if either $a \geqq b$ or $b \geqq a$. Otherwise a and b are *non-comparable*. A subset S of P is *independent* if every two distinct elements of S are non-comparable. S is *dependent* if it contains two distinct elements which are comparable. A subset C of P is a *chain* if every two of its elements are comparable.

This paper will be devoted to the proof of the following theorem and some of its applications.

THEOREM 1.1. *Let every set of $k + 1$ elements of a partially ordered set P be dependent while at least one set of k elements is independent. Then P is a set sum of k disjoint chains.*[1]

It should be noted that the first part of the hypothesis of the theorem is also necessary. For if P is a set sum of k chains and S is any subset containing $k + 1$ elements, then at least one pair must belong to the same chain and hence be comparable.

Two different embeddings of the same interval

Two different embeddings of the same interval

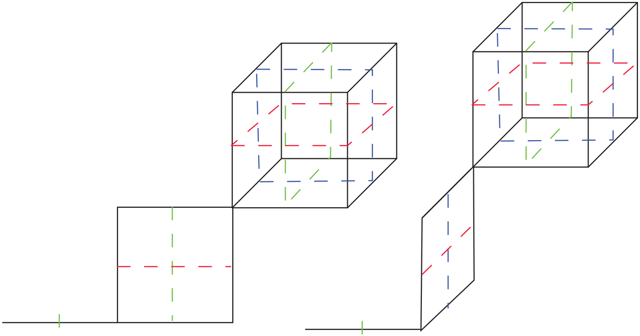

Example

Here we have taken a 1-dimensional CAT(0) cube complex (a tree) and computed the weights with respect to an embedding in the Euclidean plane.

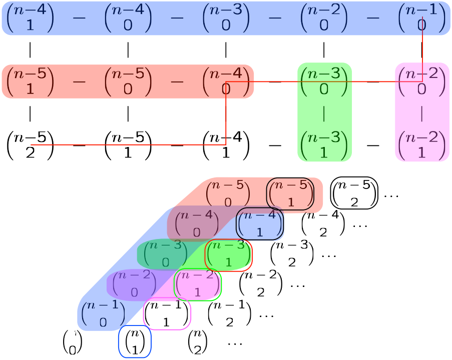

Stabilisers at infinity

Recall that a vertex of the cube complex can be viewed as an ultrafilter on the pocset. The full collection of ultrafilters yields additional vertices "at infinity":

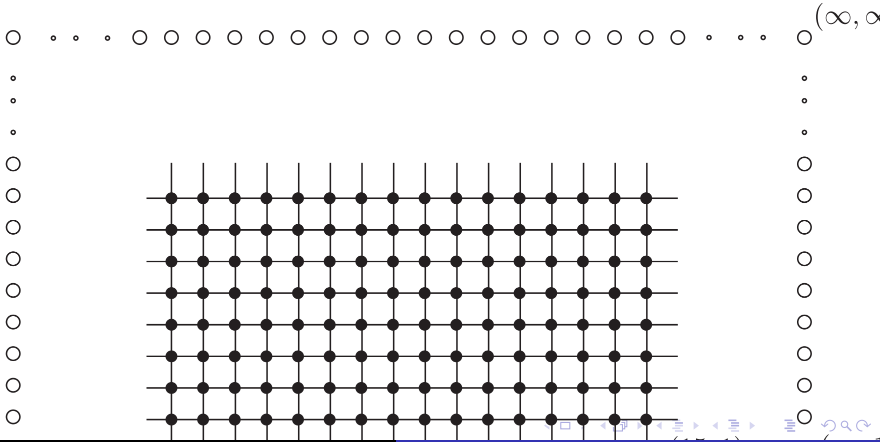

Given a vertex v in the cube complex and a vertex u at infinity we can still define an interval $[v, u]$, and use this to define Yu-style weighting functions for each u at infinity.

But these functions are equivariant under the action of the stabiliser of u, so they actually form a Reiter sequence for the stabiliser, i.e. a weighted form of a Folner sequence.

If G acts properly on a CAT(0) cube complex of finite dimension then stabilisers at infinity are amenable.

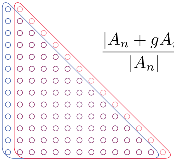

$$\frac{|A_n + gA_n|}{|A_n|} = \frac{4}{(n+1)}$$