Cubulating spaces and groups, lecture notes
(working draft – February 23, 2024)

Jason F. Manning

Contents

Preface

These notes are a work in progress, based partly on a Fall 2014 course at Cornell. Thanks very much to all the participants in that course. Almost nothing (correct) in this manuscript is original, but right now the references are extremely incomplete. If you see a mistake or missing reference please let me know about it! Thanks to Pallavi Dani, Daniel Groves, and Chaitanya Tappu for pointing out errors in earlier versions. All errors remaining are due to me.

Introduction: Subgroup separability

CHAPTER 1

Outline and conventions

1. Dependence of chapters written so far

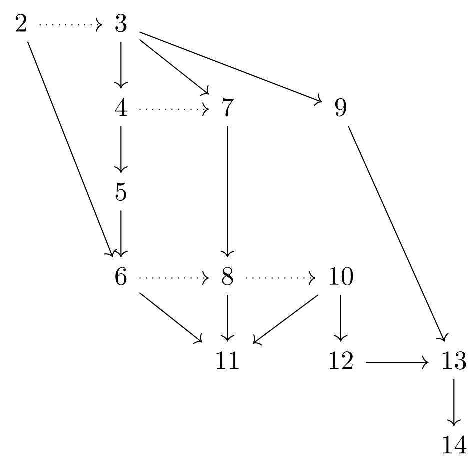

2. Things this text covers or should eventually cover

- Residual properties of groups. (Chapter 2.)
- NPC cube complexes. (Part I, starting with Chapter 3.)
- RAAGs and special cube complexes. (Chapters 4,5.)
- Geometry of CAT(0) cube complexes. (Chapter 7.)
- Hyperbolic groups. (Part II, starting with Chapter 10)
- Cubulating with codimension 1 subgroups. (Chapter 9, and much of Part II.)
- Hierarchies and (special) combination theorems.
- MSQT and Dehn filling.
- Agol's theorem.

Some highlights of the part written already are Haglund and Wise's virtual specialness criterion for cubulated hyperbolic groups (Chapter 11) and the finiteness criteria of Sageev and Bergeron–Wise for hyperbolic groups acting on cube complexes given in Chapter 14.

3. Conventions

$H \dot{<} G$ means H is a *finite index* subgroup of G. $H \dot{\triangleleft} G$ means H is a finite index normal subgroup of G. All metrics are written $d(\cdot,\cdot)$ unless there is a chance of ambiguity about the ambient metric space X, in which case the metric is written $d_X(\cdot,\cdot)$.

CHAPTER 2

Subgroup separability in free and surface groups

The purpose of this section is to prove some profinite statements about free and surface groups using the geometric methods of Stallings and Scott.

1. Residual finiteness

DEFINITION 2.1. A group G is *residually finite* if for every $g \in G \smallsetminus \{1\}$, there is a finite Q and a homomorphism $\phi\colon G \to Q$ so that $\phi(g) \neq 1$.

This basic notion has a number of equivalent formulations. One is in terms of the *profinite topology* on a group, which is the topology generated by finite index subgroups and their cosets. We collect a few here:

LEMMA 2.2. *Let G be a group. The following conditions are equivalent:*

(1) *G is residually finite.*
(2) *G is* fully *residually finite: For any finite set $F \subseteq G$, there is a finite quotient $\phi\colon G \to Q$ so that $\phi|F$ is injective.*
(3) $\bigcap\{H \dot{<} G\} = \{1\}$.
(4) *The profinite topology on G is Hausdorff.*

Verification is left to the reader.

One could argue that all the characterizations in Lemma 2.2 are essentially algebraic. Here is a topological characterization from Scott [**Sco78**].

PROPOSITION 2.3. [**Sco78**, Lemma 1.3] *Let K be a connected CW-complex, with $G = \pi_1 K$, and let $\pi\colon \tilde{K} \to K$ be the universal cover.*[1] *The following are equivalent:*

(1) *G is residually finite.*
(2) *For any compact $C \subseteq \tilde{K}$ there is a $G_0 \dot{<} G$ with $gC \cap C = \emptyset$ for all $g \in G_0 \smallsetminus \{1\}$.*
(3) *For any compact $C \subseteq \tilde{K}$ there is a finite-sheeted cover $K_C \to K$ so that the natural covering map $\tilde{K} \to K_C$ restricts to an embedding of C.*

Condition (3) is saying that the green part of the following diagram can be filled in, where all maps not from C are covering maps (the one from K_C to K

[1] Scott more generally allows $\tilde{K}$ to be any Hausdorff space on which G acts freely and properly discontinuously.

being finite sheeted).

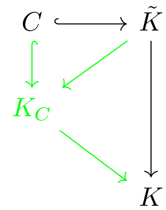

Proof. We fix a basepoint $p \in K$ and a lift $\tilde{p} \in \tilde{K}$, and suppose all covers of K come with a basepoint which is the image of this $\tilde{p}$.

(2) $\iff$ (3): Here we are simply using the correspondence between subgroups of $\pi_1 K$ and covers of K. We have (for (2)$\Rightarrow$(3)) $K_C = G_0 \backslash \tilde{K}$ and (for (3)$\Rightarrow$(2)) $G_0 = \pi_1 K_C$.

(1)$\Rightarrow$(2): Suppose G is RF. Let $T = \{g \mid gC \cap C \neq \emptyset\}$. This set is finite by proper discontinuity of the action. By Lemma 2.2.(2) there is a finite Q and a homomorphism $\phi\colon G \to Q$ which is injective on T. Let $G_0 = \ker \phi$.

(2)$\Rightarrow$(1): Suppose the condition about compact sets holds, and let $g \in G \setminus \{1\}$. Let $C = \{\tilde{p}, g\tilde{p}\} \subseteq \tilde{K}$, and let K_C be a finite cover of K in which this C embeds. If γ is a loop based at p representing $g \in \pi_1 K$, then γ doesn't lift to K_C, so $\gamma \notin \pi_1 K_C \dot{<} G$. $\square$

As an example, we give a topological proof that free groups are residually finite.

Remark 2.4. (This can also be seen using Mal′cev's theorem that linear groups are residually finite, after verifying the existence of free linear groups, for example

$$\left\langle \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \right\rangle \cong F_2.$$

In this case we have an embedding into $SL(2,\mathbb{Z})$, so it suffices to note that every element survives in $SL(2,\mathbb{Z}/p)$ for some p.)

Some lemmas, to be proved by the reader:

Lemma 2.5. *If finitely generated free groups are RF, then all free groups are RF.*

Lemma 2.6. *Suppose $H \dot{<} G$ (ie H is finite index in G). H is RF $\iff$ G is RF.*

Lemma 2.7. *A free group of rank 2 has finite index subgroups of all finite ranks bigger than 2.*

Theorem 2.8. *Free groups are residually finite.*

Proof. By the above lemmas we really only need to prove $F = \langle a, b \rangle$ is RF. We have $F = \pi_1 K$ where K is a rose with two petals (wedge of two circles). The universal cover is a 4–valent tree. It can be identified with the Cayley graph Γ of F with respect to the generating set $\{a, b\}$. Thus the edges can be given orientations and labeled by the generators a and b.

Let C be a compact subset of $\Gamma = \tilde{K}$. Let D be a connected subgraph of Γ containing C. The cover $\tilde{K} \to K$ restricts to an immersion of D, which fails to be a cover because of some missing edges. We correct this as follows: For each

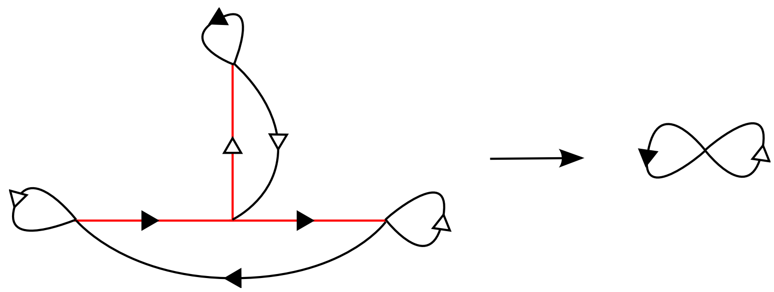

FIGURE 1. Completing the red graph D to a cover of the rose. The two generators are indicated by black and white arrow markers.

generator $x \in \{a, b\}$, let γ_x be the loop of the rose corresponding to x, and let L be a maximal component of $\pi^{-1}(\gamma_x)$. This component is an interval, to which we can add a single edge e so that $L \cup e$ is a finite cover of γ_x. After doing this to each such component, we have embedded D (and hence C) into a finite-sheeted cover of the rose. □

Notice that we didn't really use that D was a subset of a tree, but just that it had some immersion to the rose. This suggests that there is something stronger we could have proven!

2. Subgroup separability

DEFINITION 2.9. Let $H < G$. We say H is *separable* if for every $g \in G \setminus H$, there is a finite group Q and a homomorphism $\phi\colon G \to Q$ so that $\phi(g) \notin \phi(H)$.

Again, there are a number of group-theoretic equivalences:

LEMMA 2.10. *Let $H < G$. The following are equivalent:*

(1) *H is separable.*
(2) $\bigcap\{K \mid H < K \dot{<} G\} = H$.
(3) *H is closed in the profinite topology.*

Some easy consequences:

LEMMA 2.11. *G is RF if and only if $\{1\}$ is separable in G.*

LEMMA 2.12. *Let $G_0 \dot{<} G$, $H < G$, and let $H_0 = H \cap G_0$.*
H_0 is separable in G_0 $\iff$ H_0 is separable in G $\iff$ H is separable in G.

Here is a way to generate examples of separable subgroups of RF groups.

DEFINITION 2.13. Say $H < G$ is a *virtual retract* of G if there is some $G_0 < G$ containing H which retracts to H.

LEMMA 2.14. *Let G be RF, and suppose R is a virtual retract of G. Then R is separable in G.*

PROOF. By Lemma 2.12 it suffices to consider the case that G retracts to R. Let $\rho\colon G \to R$ be a retraction (i.e. $\rho|R$ is the identity), and suppose that $g \in G \setminus R$. It follows that $\rho(g) \neq g$. Since G is residually finite, G is fully residually finite, so there is a finite group Q and a $\phi\colon G \to Q$ so that $\phi(g) \neq \phi(\rho(g))$. Now consider the map

$$\Phi = (\phi, \phi \circ \rho)\colon G \to Q \times Q.$$

We have $\Phi(R)$ contained in the diagonal subgroup of $Q \times Q$, but $\Phi(g)$ outside it. □

Again, there is a topological criterion for separability:

THEOREM 2.15. [**Sco78**, 1.4] *Let K be a CW-complex with $\pi_1 K = G$, and let $H < G$. Let $\tilde{K}_H$ be the cover of K corresponding to H. The following are equivalent.*

(1) *$H < G$ is separable.*
(2) *For any compact $C \subseteq \tilde{K}_H$, there is a finite-sheeted intermediate cover $K_C \to K$ so that the natural covering map $\tilde{K}_H \to K_C$ restricts to an embedding of C.*

As with the topological criterion for residual finiteness, it may be helpful to draw a diagram of condition (2), with the part to be filled in in green.

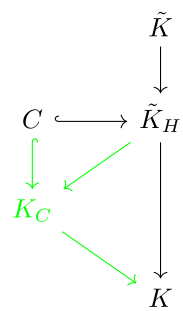

PROOF. Again, we fix basepoints $p \in K$, $\tilde{p} \in \tilde{K}$ to pin down the correspondence between subgroups of G and covers of K. Let $\pi_H\colon \tilde{K} \to \tilde{K}_H$ be the cover $x \mapsto Hx$.

(2)⇒(1): For $g \in G \smallsetminus H$, we have $H\tilde{p} \neq Hg\tilde{p}$, so $\pi_H(\tilde{p}) \neq \pi_H(g\tilde{p})$. Let $C = \{\pi_H(\tilde{p}), \pi_H(g\tilde{p})\}$. The cover K_C provided by (2) has the property that no based loop representing g lifts to it, so $G_0 = \pi_1 K_C$ doesn't contain g. But since K_C is covered by $\tilde{K}_H$, G does contain H. Since g was arbitrary, H is separable.

(1)⇒(2): Let C be a compact subset of $\tilde{K}_H$. There is a finite subcomplex D containing C. Lifting the open cells of D one by one, we can find a $\tilde{D} \subseteq \tilde{K}$, composed of finitely many open cells, so that π_H maps $\tilde{D}$ bijectively to C. This $\tilde{D}$ is contained in a finite subcomplex $E \subseteq \tilde{K}$. The set $T_0 = \{g \in G \mid gE \cap E \neq \emptyset\}$ is finite, since G is acting properly discontinuously. Let $T = T_0 \smallsetminus H$.

Since H is separable, there is a finite index A in G so that $H < A$, but $H \cap T = \emptyset$. Let K_C be the cover corresponding to A, and suppose by way of contradiction that C doesn't embed. Then D doesn't embed. In particular there is some $g \in G \smallsetminus A$ so that $g\tilde{D} \cap \tilde{D} \neq \emptyset$. But $G \smallsetminus A \subseteq G \smallsetminus H$, so this $g \in T$, a contradiction. □

3. Stallings folds and covers of the rose

In [**Sta83**], Stallings gives a powerful method for understanding finitely generated subgroups of a free group. In particular, here is an algorithm, given n words $\{w_1, \dots, w_n\}$ in a free group $F = \langle x_1, \dots, x_k \rangle$, to build a core for the cover of the rose corresponding to the group $H < F$ generated by the words:

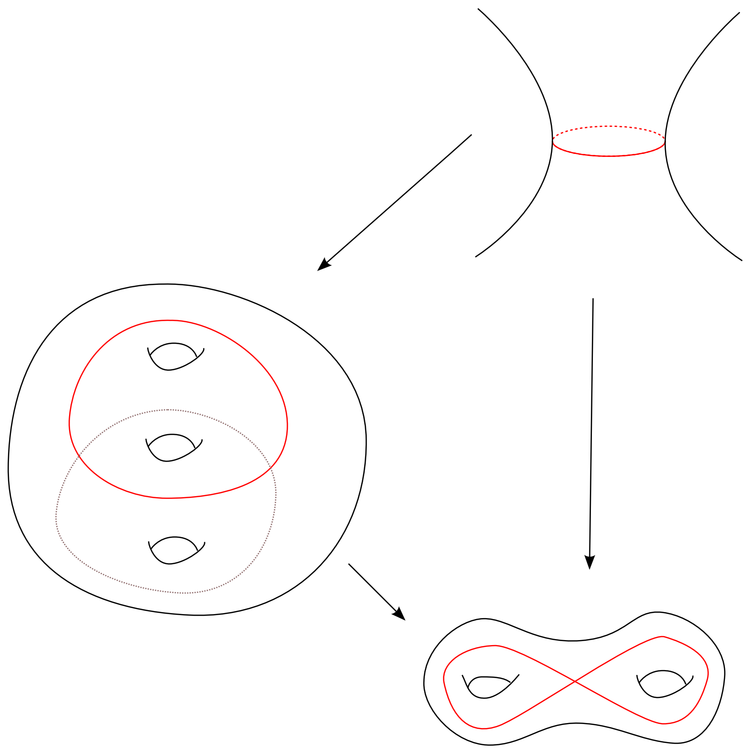

FIGURE 2. An illustration of the topological criterion for separability. The compact set C is the red circle, embedding in the intermediate cover at left. The dashed circle in the cover is some other elevation of the immersed circle in the surface K.

(1) We start with a map of roses $R_n \to R_k$ representing the map from $F_n \to F_k$ sending the ith generator to w_i. The ith petal of R_n can be subdivided into $|w_i|$ edges so that each edge goes to a constant speed loop around some letter. Label and direct each edge accordingly, obtaining a graph Γ_0, which we now modify inductively.
(2) Given Γ_i, we check to see whether some vertex has two adjacent edges with the same direction and label. If so, Γ_i is *foldable*, and we obtain Γ_{i+1} from Γ_i by identifying these two edges. If there is more than one choice, make one at random and check again. If Γ_i is not foldable, then set $\Gamma_H = \Gamma_i$.

Since each fold decreases the number of edges, the process above must terminate.

Note that each of the Γ_i constructed above is still a directed graph labeled by basis elements of F_k, so it comes equipped with a canonical map to R_k. For Γ_H, call this map η_H.

LEMMA 2.16. *The map $\eta_H \colon \Gamma_H \to R_k$ is an immersion (locally injective map).*

LEMMA 2.17. *Any immersion of connected 1–complexes $i \colon A \to B$ can be extended to a covering by attaching trees to A.*

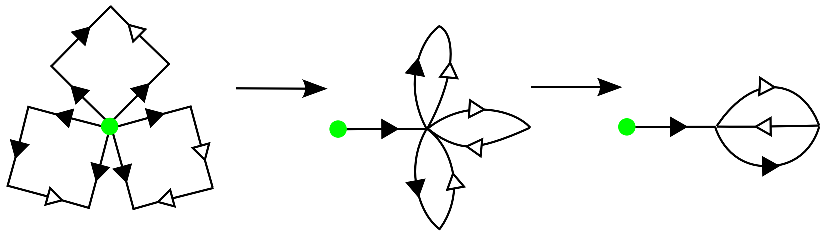

FIGURE 3. Folding the graph representing $H = \langle aaBA, abbA, aBAA\rangle$. Note that each step in the picture is several steps in the textual description. The resulting graph proves that H has rank 2 and infinite index in $\langle a, b\rangle$.

COROLLARY 2.18. *If $i\colon A \to B$ is an immersion of 1–complexes, then $i_*\colon \pi_1 A \to \pi_1 B$ is injective.*

Before discussing separability, let's record two consequences of Stallings' construction:

PROPOSITION 2.19. *There is an algorithm which takes as input a finite collection $w_1, \ldots, w_n \subseteq F_k$ and outputs the rank and index of the subgroup they generate.*

PROOF. Letting $H = \langle w_1, \ldots, w_n\rangle$ as above, the rank of H is $1 - \chi(\Gamma_H)$. If $\Gamma_H \to R_k$ is a covering map, then the index is the number of vertices in Γ_H. Otherwise, the index is infinite. □

DEFINITION 2.20. A group G is *LERF* (locally extended residually finite) if every finitely generated subgroup of G is separable.

THEOREM 2.21. *Free groups are LERF.*

PROOF. Again it suffices to consider finitely generated free groups, so fix such a free group F of rank k. We have $F = \pi_1 R_k$, where R_k is the rose with k petals. Let $H = \langle w_1, \ldots, w_n\rangle$ be a finitely generated subgroup of F. We've shown how to describe the cover $\tilde{R}_H$ corresponding to H, together with a compact core $\Gamma_H \subseteq \tilde{R}_H$. Let $C \subseteq \Gamma_H$ be some compact subset, and let D be a connected subcomplex of $\tilde{R}_H$ containing both C and Γ_H. The covering map $\tilde{R}_H \to R_k$ restricts to an immersion of D, which we can complete to a cover in exactly the same way as we completed D to a cover in the proof of Theorem 2.8. □

In fact the above proof finds a finite cover of the rose containing Γ_H as a subcomplex.

LEMMA 2.22. *Let $A \subseteq B$ be an inclusion of connected 1–complexes. Then $\pi_1 A$ is a free factor of $\pi_1 B$.*

COROLLARY 2.23 (Marshall Hall's Theorem). *Let $H < F$ be finitely generated, where F is free. Then there is a finite index $F' \dot{<} F$ containing H, so that $F = H * K$ for some K.*

4. Surface groups are LERF

Our aim here is to prove the following theorem of Scott.

THEOREM 2.24. *Let Σ be a surface. Then $\pi_1\Sigma$ is LERF.*

We more or less follow Scott's proof from [**Sco78, Sco85**]. A key insight there is to note that all closed hyperbolic surface groups are abstractly commensurable to a certain reflection group acting on the hyperbolic plane. Before getting to that, we deal with some simple situations.

(1) Suppose Σ is not closed. Then $\pi_1\Sigma$ is free, hence LERF by Theorem 2.21.
(2) Suppose Σ is closed, but $\chi(\Sigma) > 0$. Then $\pi_1\Sigma \in \{\{1\}, \mathbb{Z}/2\}$ is finite, hence LERF.
(3) Suppose Σ is closed and $\chi(\Sigma) = 0$. Then $\pi_1\Sigma$ is either $\mathbb{Z} \oplus \mathbb{Z}$ (if Σ is a torus) or contains $\mathbb{Z} \oplus \mathbb{Z}$ as an index 2 subgroup (if Σ is a Klein bottle). It's easy to show $\mathbb{Z} \oplus \mathbb{Z}$ is LERF.

We're left with the situation that Σ is closed and $\chi(\Sigma) < 0$. In such a case, Σ finitely covers Σ_{-1}, the nonorientable closed surface with Euler characteristic -1. Moreover, we'll see that $\pi_1\Sigma_{-1}$ is finite index in a reflection group.

Let ⬠ be a right-angled regular pentagon in $\mathbb{H}^2$. Let P be the group of isometries of $\mathbb{H}^2$ generated by reflections in the lines bounding ⬠. There is a finite index subgroup P_0 which is torsion-free, so that $P_0 \backslash \mathbb{H}^2$ is a hyperbolic surface.

The group P preserves a family of lines $\mathcal{L}$, which cut $\mathbb{H}^2$ into pentagons which are translates of ⬠. Each line in $\mathcal{L}$ determines two convex halfspaces which are unions of pentagons. Call these the *combinatorial halfspaces* determined by P.

If $C \subset \mathbb{H}^2$, we define the *combinatorial hull* of C to be the intersection of the combinatorial halfspaces containing C.

Here's a lemma about hyperbolic geometry which will be used a couple of times to control the size of combinatorial hulls.

LEMMA 2.25. *Let C be a closed convex subset of $\mathbb{H}^2$, and let $\gamma\colon [0,\infty) \to \mathbb{H}^2$ be a geodesic ray in $\mathbb{H}^2 \smallsetminus C$. Define $\alpha(t)$ to be the visual angle subtended by C, as seen from $\gamma(t)$. Then $\lim_{t\to\infty} \alpha(t) = 0$.*

Let's warm up with the following (which is also a consequence of Mal′cev's theorem).

PROPOSITION 2.26. *Hyperbolic surface groups are RF.*

PROOF. It suffices to show $P_0 = \pi_1\Sigma_0$ is RF. We'll use the topological criterion. Let $C \subseteq \mathbb{H}^2$ be a compact set, and let D be the combinatorial hull of C.

CLAIM. *D is compact (a union of finitely many pentagons).*

PROOF OF CLAIM. Let $l \in \mathcal{L}$. If l does not meet C, then D lies entirely on one side of l. It follows that if l meets the interior of D, then l meets C. Since C is bounded, there are only finitely many lines l meeting the interior of D. If each only meets D in a bounded set, there can only be finitely many pentagons in D.

Let $l \in \mathcal{L}$ be some line which intersects the interior of D, and therefore hits C. Since C is bounded, there are two unbounded components of $l \smallsetminus C$. Each of these is a geodesic ray to which we can apply Lemma 2.25. Moreover, this ray crosses infinitely many perpendicular lines $k_1, k_2, \ldots$ from $\mathcal{L}$. Let $p_i = k_i \cap l$, and let α_i be the visual angle subtended by C at p_i. Lemma 2.25 implies that $\lim_{i\to\infty} \alpha_i = 0$. In particular it is eventually less than $\pi/2$, so the k_i must eventually miss C. The part of l separated from C by k_i cannot be part of D, so $l \cap D$ is bounded. □

Let Γ be the group generated by reflections in the lines bounding D. Since $\Gamma\backslash\mathbb{H}^2$ is compact, the group Γ is finite index in P. It follows that $H = \Gamma \cap P_0$ is also finite index in P_0, so that if $\Sigma = H\backslash\mathbb{H}^2$, then Σ is a finite-sheeted cover of Σ_0. But since C embeds in $\Gamma\backslash\mathbb{H}^2$, it also embeds in Σ. □

We'll show:

Theorem 2.27. *For every finitely generated $H < P_0$, there is a finite index $P' < P$ containing H as a retract.*

Proof. If $H \dot{<} P_0$, there is nothing to show, so assume that H is infinite index in P_0. We can also assume that $H \neq 1$. It follows that $\Sigma = H\backslash\mathbb{H}^2$ is a noncompact hyperbolic surface. Since P doesn't contain any parabolics, this surface has no cusps, so its convex core C is compact, bounded by finitely many simple geodesic loops. (Or possibly C consists of a single simple geodesic loop.) Let $\pi_H\colon \mathbb{H}^2 \to \Sigma$ be the covering map, and let $\tilde{C} = \pi_H^{-1}(C)$. This is some convex subset of $\mathbb{H}^2$. Let Y be the combinatorial convex core of $\tilde{C}$.

Claim. *$\bar{Y} = \pi_H(Y)$ is compact, consisting of finitely many pentagons.*

Given the claim, let $\tilde{R}$ be the subgroup of P generated by reflections in the faces of Y, and let $P' = \langle \tilde{R}, H\rangle$. The quotient of $\mathbb{H}^2$ by P' is exactly $\bar{Y}$, so $P' \dot{<} P$. We observe

(1) H normalizes $\tilde{R}$. (Since H preserves Y, it conjugates generating reflections of $\tilde{R}$ to other generating reflections.)
(2) $H \cap \tilde{R} = \{1\}$. (Every element of H preserves Y, while no nontrivial element of $\tilde{R}$ does.)

These two facts together imply that $P' = \tilde{R} \rtimes H$, so P' retracts to H.

Proof of claim. The proof of the claim consists of two parts. First, we note that there are finitely many H–orbits of lines l of the pentagonal tiling which meet the interior of Y. Indeed such a line l must meet $\tilde{C}$, and so $\pi_H(l)$ meets C. Since C is compact, there are only finitely many lines of the tiling meeting C.

Second we show that each such line has compact intersection with $\bar{Y}$. Note first that if $l \cap \tilde{C}$ is noncompact, then $l \subseteq \tilde{C}$, so $\pi_H(l)$ is a closed (hence compact) curve. So we may suppose that $l \cap \tilde{C}$ is compact, and so $l \smallsetminus \tilde{C}$ is a pair of rays. Let $\gamma\colon [0,\infty) \to \mathbb{H}^2$ be a geodesic ray in $l \smallsetminus \tilde{C}$. Lemma 2.25 implies that the visual angle of $\tilde{C}$ as seen from $\gamma(t)$ is eventually less than $\pi/2$, say for $t \geq t_0$. Let $k \in \mathcal{L}$ be a line crossed by $\gamma(t)$ for $t \geq t_0$. Then k separates $\tilde{C}$ from $\gamma([t,\infty))$, so $\gamma(t)$ is not in the interior of Y. We've shown that $l \cap Y$ is compact, so $\pi_H(l) \cap \bar{Y}$ is compact.

It follows that there are only finitely many pentagons in $\pi_H(Y)$, since each such pentagon must meet the interior of $\pi_H(Y)$. □

□

We've shown that surface groups are RF (2.26) and that finitely generated subgroups are virtual retracts (2.27). By Lemma 2.14, we have the following Corollary:

Corollary 2.28. *Surface groups are LERF.*

Part I

Non-positively curved cube complexes

CHAPTER 3

Introduction to cube complexes

1. Nonpositive curvature

An *n–cube* is a copy of $I^n = [0,1]^n$ metrized as a subset of Euclidean space. A k–dimensional *face* of I^n is a subset in which all but k of the coordinates are held constant at either 0 or 1. A *cube complex* is built from a disjoint union of cubes of various dimensions, glued together by isometries of faces. The 0–cubes will also be referred to as *vertices*; the 1–cubes as *edges*.

Definition 3.1. The $\frac{1}{3}$–neighborhood of a vertex v in a cube complex inherits the structure of a Δ–complex from the cube-complex structure: Each n–cube incident to v contributes an $(n-1)$–simplex, and simplices are glued together along faces exactly when the cubes are glued along a face incident to v. This Δ–complex is called the *link* of v, or $\mathrm{lk}(v)$.

Definition 3.2. A cube complex is *non-positively curved* or *NPC* if every link of a vertex is a flag simplicial complex.

Example 3.3. A square complex (2–dimensional cube complex) is NPC if and only if there is no cycle of length less than 4 in any link.

Example 3.4. Let $K \subseteq S^3$ be a knot, and consider a (generic) projection of that knot K to an equatorial sphere E. Let N and S be the north and south poles. For each region R of $E \smallsetminus K$, choose a geodesic arc from N to S through that region, and label it with the region R. For each crossing we attach a square with labels given by the following rule:

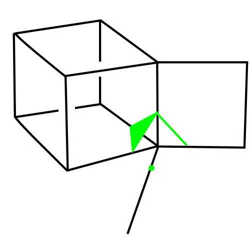

Figure 1. Picture of the link of a vertex in a simple cube complex

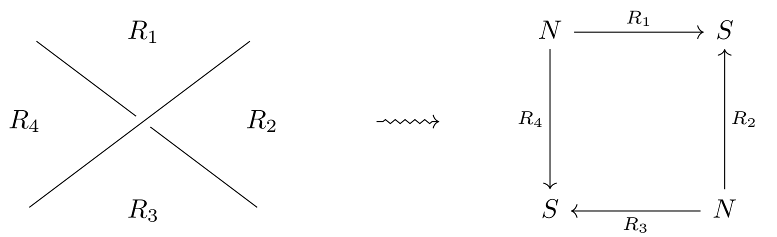

It's not too hard to show that the resulting square complex is NPC if and only if the link projection was alternating.

2. The cube complex associated to a right angled Artin group

Let Γ be an unoriented simplicial graph, with vertex set V, and let $E \subseteq V \times V$ be the edge set. We define the *right angled Artin group* (or *RAAG*) based on Γ to be the group:

$$A(\Gamma) = \langle V \mid vw = wv, \text{ for } (v, w) \in E \rangle .$$

Some important examples:

(1) If Γ has no edges, then $A(\Gamma)$ is free of rank $\|\Gamma^{(0)}\|$.
(2) If $\Gamma = K_n$, the complete graph on n vertices, then $A(\Gamma) \cong \mathbb{Z}^n$.
(3) If $\Gamma = K_{p,q}$ is a complete bipartite graph, then $A(\Gamma) \cong F_p \times F_q$, a product of free groups.
(4) If Γ is a segment of length 2, then $A(\Gamma)$ is the fundamental group of the complement of a certain 3–component link (see Figure).

The last (3–manifold) example can be generalized. Droms [**Dro87**] showed that $A(\Gamma)$ is a 3–manifold group if and only if Γ is a disjoint union of trees and triangles. One direction is a straightforward construction. To show the others are not 3–manifold groups, one either embeds $\mathbb{Z}^4$ (if there is a K_4) or shows the groups are incoherent, contradicting the Scott–Shalen Core Theorem.

Remark 3.5. RAAGs are not LERF in general. In fact even $F_2 \times F_2$ is not LERF since it has unsolvable membership problem [**?**].

Definition 3.6. If V is a finite set, one can form the torus $T_V = (S^1)^V$. This torus has a nice CW-complex structure with the set of cells in one-to-one correspondence with 2^V. In this correspondence, the empty set corresponds to the unique 0–cell, the singletons to 1–cells, 2–cells to pairs, etc. Notice that this CW complex is also a NPC cube complex.

If Γ is a graph with vertex V, then the *Salvetti complex* $S(\Gamma)$ is the subcomplex of T_V consisting of those cubes corresponding to cliques in Γ.

Lemma 3.7. *Let Γ be a finite simplicial graph. The Salvetti complex $S(\Gamma)$ is a NPC cube complex with $\pi_1 S(\Gamma) \cong A(\Gamma)$.*

To prove the lemma, we need a couple of definitions

Definition 3.8. Let Λ be a simplicial graph. The *flagification* $\mathrm{Flag}(\Lambda)$ is the unique flag complex whose 1–skeleton is Λ.

Definition 3.9. Let K be a simplicial complex with vertex set V. The *double* $D(K)$ is the complex with vertex set $V^+ \sqcup V^-$ (two disjoint copies of V), and so that vertices $\{v_0^{\epsilon_0}, \ldots, v_n^{\epsilon_n}\}$ span a simplex if and only if $\{v_0, \ldots, v_n\}$ do.

The proof of Lemma 3.7 thus boils down to the following:

EXERCISE 1. (1) The double of a flag complex is flag.
(2) The link of the vertex of $S(\Gamma)$ is the double of $\mathrm{Flag}(\Gamma)$.

Right-angled Artin groups turn out to be absolutely central to the subject of these notes. In particular, finding geometrically nice embeddings of groups into RAAGs turns out to be very useful. One reason for this is that RAAGs are *linear*, meaning they admit faithful representations into $GL(n,\mathbb{C})$ for some n. This follows from the fact that they are abstractly commensurable to right-angled Coxeter groups, as we now explain.

DEFINITION 3.10. Let Γ be a graph with vertex set V and edge set E. The *right-angled Coxeter group* $C(\Gamma)$ based on Γ is the group

$$C(\Gamma) = \langle V \mid v^2 = 1 \text{ for } v \in V; (vw)^2 = 1 \text{ for } (v,w) \in E\rangle.$$

(More general Coxeter groups are also generated by involutions, but the relations $(vw)^2 = 1$ are replaced by $(vw)^{m(v,w)} = 1$ for some collection of $m(v,w) \in \{2, 3, \dots, \infty\}$

THEOREM 3.11. *(Tits)* [**Dav08**, Appendix D] *Each Coxeter group embeds into $SL(n,\mathbb{Z})$ for some n. In particular, Coxeter groups are linear.*

PROOF. (Idea) We won't really prove this; just give the representation in the right-angled case. The proof of faithfulness can be found in Davis' book (also in Bourbaki). Let $C(\Gamma)$ is the right-angled Coxeter group based on Γ, with vertex set V and edge set $E \subseteq V \times V$. We describe an action of $C(\Gamma)$ on $\mathbb{R}^V$ which preserves a certain quadratic form. Let $v \mapsto e_v$ be a bijection between V and the basis of $\mathbb{R}^V$. For $v, w \in V$, define

$$\langle e_v, e_w\rangle = \begin{cases} 0, & \text{if } (v,w) \in E \\ -1, & \text{otherwise.} \end{cases} \tag{1}$$

Now we describe an action $C(\Gamma) \curvearrowright \mathbb{R}^V$ by

$$v(x) = x - 2\langle e_v, x\rangle e_v.$$

Clearly this representation has image in $GL(n,\mathbb{Z})$ where $n = |V|$. But $GL(n,\mathbb{Z})$ embeds into $SL(n+1,\mathbb{Z})$. □

EXERCISE 2. Suppose Γ consists of three vertices v_1, v_2, v_3, where v_1 and v_2 are connected by an edge, and v_3 is isolated. Compute the representation into $GL(3,\mathbb{Z})$. What's the signature of the form described in equation (1)?

There is an obvious surjection $A(\Gamma) \to C(\Gamma)$, but the kernel of this map is infinite, so it doesn't give abstract commensurability. We have to choose a different graph.

THEOREM 3.12. *(Davis-Januszkiewicz)*[**DJ00**] *Let Γ be a finite graph. Then there is another graph Γ' and an injective homomorphism $A(\Gamma) \to C(\Gamma')$ with finite index image.*

PROOF. (Sketch) Let V be the vertex set of Γ and let $E \subseteq V \times V$ be the edge set. Davis and Januszkiewicz describe Γ' in the following way: The vertex set V' is equal to two copies of V, the vertex set of Γ. Decorate the elements of V by hats and checks so $V' = \hat{V} \sqcup \check{V}$. The edge set E' of Γ' is given by three rules:

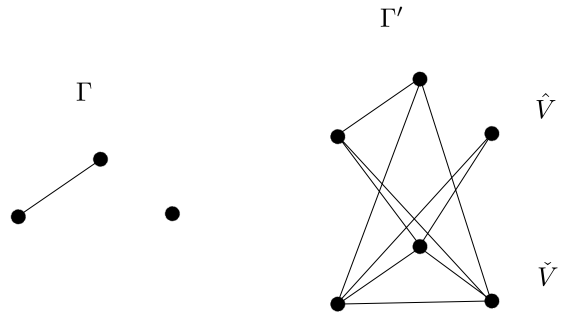

FIGURE 2. Davis–Januszkiewicz' construction.

(1) Every pair of vertices in $\check{V}$ is connected by an edge, so $\check{V}$ spans a complete graph.
(2) Vertices $\hat{v}$ and $\hat{w}$ are connected by an edge if and only if $(v, w) \in E$.
(3) Connect $\hat{v}$ to $\check{w}$ if and only if $v \neq w$.

An example is shown in Figure 2. Notice that $\langle \check{v}, \hat{v} \rangle \cong \mathbb{Z}/2 * \mathbb{Z}/2$ is virtually infinite cyclic, and that $\check{v}\hat{v}$ is infinite order. The embedding $\beta\colon A(\Gamma) \to C(\Gamma')$ is given by $v \mapsto \check{v}\hat{v}$. Davis and Januszkiewicz show that $C(\Gamma')$ has a proper and cocompact action on the universal cover X of the Salvetti complex for $A(\Gamma)$ which agrees (via β) with the usual action of $A(\Gamma)$ on X. □

Using Mal′cev's theorem that linear groups are residually finite (actually very easy in this case), we obtain the corollary:

COROLLARY 3.13. *For any finite graph* Γ, $A(\Gamma)$ *is residually finite.*

But as we will see below, there is a *geometric* proof of residual finiteness along the lines of the proof for free groups given in Theorem 2.8. Moreover, though $A(\Gamma)$ is often not LERF, we will be able to use a geometric argument as in 2.21 to show many subgroups are separable.

CHAPTER 4

Special cube complexes

In this section we will meet *special* cube complexes for the first time, as cube complexes which lack certain "hyperplane pathologies." We'll also see the connection with RAAGs, which Haglund and Wise only discovered after noticing how useful the notion was for geometric separability arguments[**HW08**].

1. Special via hyperplanes

A cube $I^n = [0,1]^n$ has one *midcubes* for each dimension: $M_i = \{(x_1, \ldots, x_n) \in I^n \mid x_i = \frac{1}{2}\}$. Some midcubes are pictured in Figure 1.

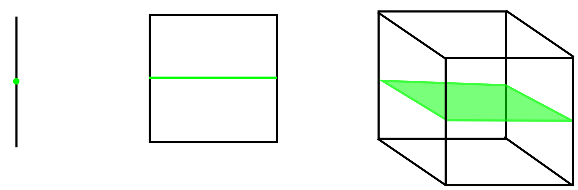

FIGURE 1. Some midcubes

Let X be a cube complex. The *midcube complex of* X, $M(X)$ is a cube complex whose cubes are in one-to-one correspondence with midcubes of cubes of X. Whenever one of the face-identifications of X identifies two faces of midcubes, we identify those faces in $M(X)$. A component H of $M(X)$ is called a *hyperplane*. It comes equipped with an immersion $m_H \colon H \to X$. An example is shown in Figure 2. A hyperplane is *embedded* if this immersion is an embedding. Otherwise we say the hyperplane *self-intersects*.

Each cube of X can be thought of as an I–bundle over any of its midcubes. We can therefore pull back an I–bundle over H, for any hyperplane. The hyperplane is said to be 2*–sided* if this bundle is trivial; otherwise it is 1*–sided*.

Two (unoriented) edges e_1, e_2 *corner a square* if there is a square of the form:

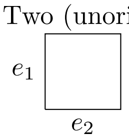

Suppose H is a 2–sided hyperplane in X, so that $m_H \colon H \to X$ extends to a cubical immersion $\bar{m}_H \colon H \times [0,1] \to X$. Suppose there are distinct vertices v_1, v_2 of H so that $\bar{m}_H(v_1, 0) = \bar{m}_H(v_2, 0)$ or $\bar{m}_H(v_1, 1) = \bar{m}_H(v_2, 1)$, and suppose that $\bar{m}_H(v_1 \times I)$ and $\bar{m}_H(v_2 \times I)$ don't corner a square. Then H is said to *self-osculate* (see Figure 3 for an example).

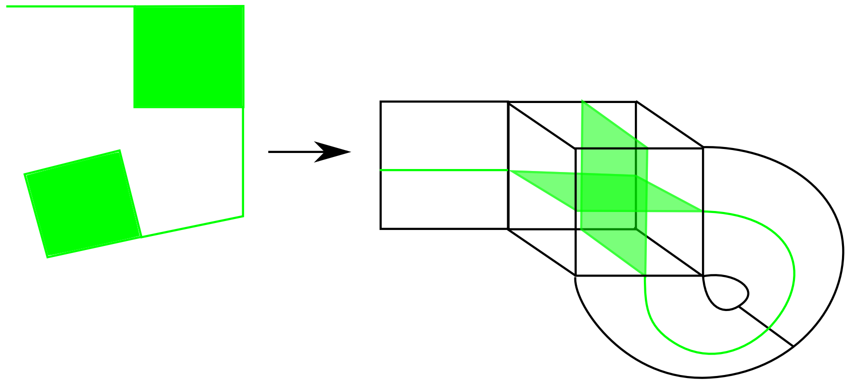

FIGURE 2. An immersed hyperplane. The cube complex shown has four other hyperplanes, each consisting of a single midcube.

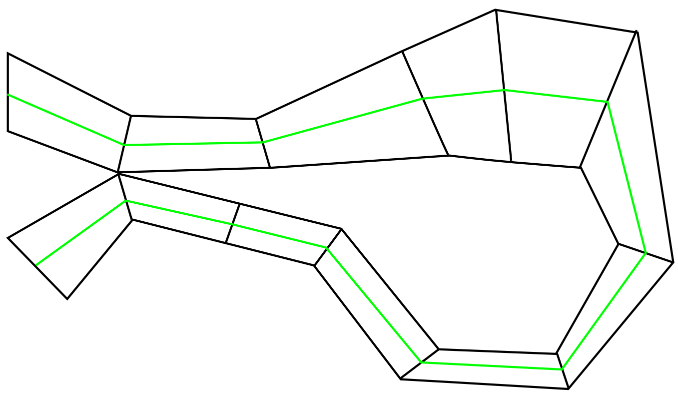

FIGURE 3. A self-osculating hyperplane

Let H_1 and H_2 be distinct 2–sided hyperplanes in X. It's not hard to see that $H_1 \cap H_2$ is nonempty if and only if there are vertices $v_i \in H_i$ so the edges $\bar{m}_{H_i}(v_i \times I)$ corner a square.

Two hyperplanes H_1 and H_2 are said to *osculate* if there are vertices $v_i \in H_i$ so that $\bar{m}_{H_1}(v_1 \times \partial I) \cap \bar{m}_{H_2}(v_2 \times \partial I)$ is nonempty, but the edges $\bar{m}_{H_i}(v_i \times I)$ do *not* corner a square.

The hyperplanes H_1 and H_2 *interosculate* if they both cross and osculate. See Figure 4 for an example.

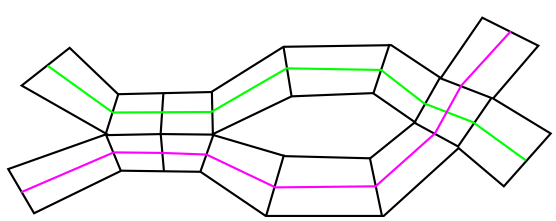

FIGURE 4. Two inter-osculating hyperplanes

Definition 4.1. A cube complex is *special* if all of the following hold:

(1) No hyperplane self-intersects.
(2) No hyperplane is 1–sided.
(3) No hyperplane self-osculates.
(4) No two hyperplanes inter-osculate.

2. Parallelism of edges

The definitions of the hyperplane pathologies can all be phrased in terms of parallelism classes of edges. This point of view is important for some proofs. All the proofs in this section are left to the reader.

Fix X a cube complex, and let $\vec{E}$ be the set of oriented edges of X. To any hyperplane H in X we associate a set $E(H) \subseteq \vec{E}$ of edges which are *dual* to H in the sense that their midpoints are 0–cubes of H. Two oriented edges e_1, e_2 *corner a square* if there is a square of the form:

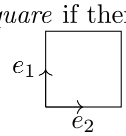

Lemma 4.2. *H self-intersects if and only if there are two edges $e_1, e_2 \in E(H)$ which corner a square.*

Say that e and $e' \in \vec{E}$ are *elementary parallel* if there is a square of the form:

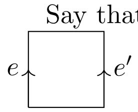

The equivalence relation $\|$ of *parallelism* on $\vec{E}$ is generated by elementary parallelism.

Lemma 4.3. *Let H be an i–sided hyperplane for $i \in \{1, 2\}$. Then $E(H)$ contains i parallelism classes.*

A *co-oriented* hyperplane $\vec{H}$ is a hyperplane H together with a choice of parallelism class $\vec{E}(H) \subsetneq E(H)$. We refer to $\vec{E}(H)$ as a *co-orientation* of H.

Lemma 4.4. *Let H be a 2–sided hyperplane. H self-osculates if and only if there is a co-orientation $\vec{E}(H)$ containing two edges e_1, e_2 which do not corner a square, but which have the same origin.*

Lemma 4.5. *Let H_1 and H_2 be 2–sided hyperplanes. Then H_1 and H_2 interosculate if and only if there are co-orientations $\vec{E}(H_1), \vec{E}(H_2)$, and edges $e_i, f_i \in \vec{E}(H_i)$ so that:*

(1) *e_1, e_2 have a common origin but don't corner a square, and*
(2) *f_1, f_2 corner a square.*

Exercise 3. Draw some square complexes and see what hyperplane pathologies occur. Can you make a special cube complex homeomorphic to a closed surface of negative Euler characteristic?

CHAPTER 5

Special cube complexes and RAAGs

In this section we prove Haglund and Wise's characterization of special cube complexes as those which locally isometrically embed in the Salvetti complex of some RAAG [**HW08**]. First note the following, which can be easily proved using the lemmas from the last section.

PROPOSITION 5.1. *Let* Γ *be a finite graph, and let* $S(\Gamma)$ *be the Salvetti complex based on* Γ*. Then* $S(\Gamma)$ *is special.*

1. Kinds of maps between cube complexes

We'll deal exclusively with *combinatorial* maps of cube complexes. To put it somewhat formally, we require that if $\phi\colon X \to Y$ is such a map, and $x\colon I^k \to X$ is the characteristic map of some cube of X, then $\phi \circ x|\mathring{I}^k$ is a homeomorphism onto the interior of some cube of Y. Moreover if $y\colon I^k \to Y$ is the characteristic map of the target cube, then $y^{-1} \circ \phi \circ x|\mathring{I}^k$ is an isometry.

If $f\colon X \to Y$ is a map of cube complexes, and v is a vertex of X, then f induces a map of links

$$\mathrm{lk}(v) \xrightarrow{f_v} \mathrm{lk}(f(v)).$$

Recall that an *immersion* is a locally injective map. We can detect whether f is an immersion by looking at the induced maps on links.

LEMMA 5.2. *Let* $f\colon X \to Y$ *be a map of cube complexes. Then* f *is an immersion if and only if* f_v *is injective for each vertex* $v \in X$*.*

A *full subcomplex* S of a simplicial complex K contains every simplex in K whose vertices are contained in S.

DEFINITION 5.3. $f\colon X \to Y$ an immersion of cube complexes is a *local isometry* if $f_v(\mathrm{lk}(v))$ is a full subcomplex of $\mathrm{lk}(f(v))$ for all vertices $v \in X$.

EXAMPLE 5.4. Consider any cube Y of dimension at least 2, and let X be the subcomplex consisting of two adjacent codimension one faces. Then X is immersed in Y, but the inclusion is not a local isometry.

REMARK 5.5. If X and Y are both NPC cube complexes, then $f\colon X \to Y$ is a local isometry if and only if no two non-adjacent vertices in a link are sent to adjacent vertices in a target link.

PROPOSITION 5.6. *Let* $f\colon X \to Y$ *be a local isometry of NPC cube complexes, where* Y *is special. Then* X *is special.*

PROOF. Without loss of generality, we may suppose that X and Y are connected. Since edges go to edges and squares to squares, there is a well-defined induced map:

$$\vec{E}(X)\Big/\| \longrightarrow \vec{E}(Y)\Big/\|. \tag{2}$$

We suppose X is not special, and show that Y cannot be special either.

We consider the hyperplane pathologies in turn, according to their characterizations in terms of edge parallelism given in Section 2. Suppose first there is a self-intersecting hyperplane in X. Then there are two edges $a \parallel b$ which corner a square of X. But it follows that $f(a) \parallel f(b)$ also corner a square, so Y is not special.

Suppose that X contains some one-sided hyperplane. Then some oriented a is parallel to $-a$, the edge with the opposite orientation. But then the same must hold in Y using (2).

Similarly, if X contains a self-osculating hyperplane, then there is a pair of oriented edges $a \parallel b$ with the same source, but which don't corner a square. The images $f(a) \neq f(b)$, since f is an immersion. The map from (2) gives $f(a) \parallel f(b)$. But since f is a local isometry, $f(a)$ and $f(b)$ don't corner a square in Y. Thus Y contains a self-osculating hyperplane.

If X contains a pair of interosculating hyperplanes, X contains edges $e_1 \parallel f_1$ and $e_2 \parallel f_2$ so that $e_1 \nparallel e_2$, exhibiting the interosculation. Namely, e_1 and e_2 have a common origin but don't corner a square, but f_1 and f_2 corner a square. As before, the same properties must hold of their images in Y. □

As a special case of the previous proposition, any NPC cube complex which locally isometrically immerses to a Salvetti complex is special. Haglund and Wise proved a remarkable converse to this fact, which we prove in the next subsection.

2. Special cube complexes embed in RAAGs

To state the result properly, we need another definition. Let X be a cube complex, and let Γ_X be the *hyperplane graph of* X: The vertices of Γ_X correspond to the immersed hyperplanes of X, and two vertices are connected to one another if the corresponding hyperplanes cross. If Γ_X is a finite graph, we can form the Salvetti complex $S(\Gamma_X)$ (whose fundamental group is a RAAG) as in Section 2.

THEOREM 5.7. [**HW08**] *Let X be a special cube complex with finitely many hyperplanes. Then there is a locally isometric immersion $\phi\colon X \to S(\Gamma_X)$.*

Before we prove the theorem, we note a corollary.

COROLLARY 5.8. *Let G be the fundamental group of a special cube complex with finitely many hyperplanes. Then G is a subgroup of some RAAG.*

PROOF OF 5.7. Let X be a special cube complex with finitely many hyperplanes. For each hyperplane H, fix a co-orientation $\vec{E}(H)$. For each hyperplane H, there is a corresponding (oriented) 1–cell e_H in $S(\Gamma_X)$, and we define $\phi|e\colon e \to e_H$ to be the orientation-preserving (combinatorial) map for each $e \in \vec{E}(H)$.

More generally, a k–cube C of X has k different hyperplanes passing through it and crossing one another. (They're different because hyperplanes of X are embedded.) Since they cross one another, these hyperplanes correspond to the vertices of a clique in Γ_X, which corresponds to a cube D in $S(\Gamma_X)$. The map has already been defined on the 1–skeleton of C, and this definition extends combinatorially uniquely to a map $\phi|C\colon C \to D$. It's not hard to see that definitions on cubes which share a face are consistent, so we've defined a combinatorial map $\phi\colon X \to S(\Gamma_X)$, just using the fact that hyperplanes are embedded and two-sided.

To see the map is an immersion, we have to use the fact that there are no self-osculating hyperplanes. Indeed, since the source and target are both NPC, ϕ is an immersion so long as it doesn't identify any two oriented edges with the same origin. If $\phi(e_1) = \phi(e_2)$ as oriented edges, then we must have had $e_1 \parallel e_2$ in X. If e_1 and e_2 originate at the same point, then the corresponding hyperplane must self-osculate, contradicting specialness.

Finally, we use the lack of inter-osculation to see that ϕ is a local isometry. Again using the fact that X and Y are NPC, ϕ can only fail to be a local isometry if there are vertices a and b in some $\mathrm{lk}(v)$ which are not connected by an edge, but $\phi_v(a)$ and $\phi_v(b)$ are connected by an edge. Let e_a, e_b be the oriented edges corresponding to a and b. If there is an edge connecting $\phi_v(a)$ to $\phi_v(b)$, then the hyperplanes dual to e_a and e_b must cross somewhere. In other words there are $f_a \parallel e_a$ and $f_b \parallel e_b$ which corner a square. Thus the hyperplanes dual to e_a and e_b interosculate. □

REMARK 5.9. The graph Γ_X is not necessarily the smallest graph Γ so that X locally isometrically immerses in $S(\Gamma)$. One can often get a smaller Γ by considering the crossing graph of a collection of not-necessarily connected "hyperplanes," each of which is a disjoint union of non-crossing hyperplanes. One has to be careful of course that these "hyperplanes" don't self-osculate or inter-osculate. In the context of graphs, this means coloring and orienting the edges, so that no two edges of the same color have the same origin or the same terminus. The corresponding Γ has vertex set equal to the set of colors.

CHAPTER 6

Canonical completion and retraction, take 1

The canonical completion and retraction allows us to prove separability of subgroups of special cube complexes, in case the subgroup is represented by a locally isometrically immersion of cube complexes. In this section we deal only with the case where the target is a Salvetti complex. The "completion" step is essentially the same as the cover described in the proof that free groups are LERF in Section 3. The *canonical retraction* is new.

Goal: From $i\colon X \to S = S(\Gamma)$ a locally isometric immersion, produce a finite cover (the completion) $p\colon \mathsf{C} \to S$ with $X \subseteq \mathsf{C}$, and so C retracts to X:

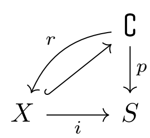

If Γ has no edges (so $A(\Gamma)$ is free) we've basically seen how to build C already in Section 3: For each maximal non-closed segment mapping of X mapping to a petal of the rose $S(\Gamma)$, we add an additional edge mapping to the same petal, completing the segment to a circle. The retraction works by mapping this additional edge continuously onto the segment it was added to. (See Figure 1.) If a new edge is

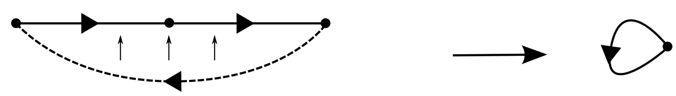

FIGURE 1. An immersion of a segment to a rose with a single petal. Complete by adding an edge; retract by projecting that edge to the preexisting segment.

attached to a single vertex (a "length 0 segment") then the retraction r maps that new edge to the attaching vertex.

The 1–skeleton $\mathsf{C}^{(1)}$ is produced from $X^{(1)} \to S^{(1)}$ in exactly the same way. We then need to check that squares and higher-dimensional cells can be added in a consistent way. For example let's complete the immersion shown in Figure 2. The procedure already described gives a graph covering the 1–skeleton of $S(\Gamma)$. One now checks that the boundary of the square in $S(\Gamma)$ lifts to a path beginning at any of the four vertices. Gluing in squares to these lifts gives a cover of $S(\Gamma)$, as in Figure 3.

1. Definition of the completion

We recall and name the completion and retraction, which we already described informally.

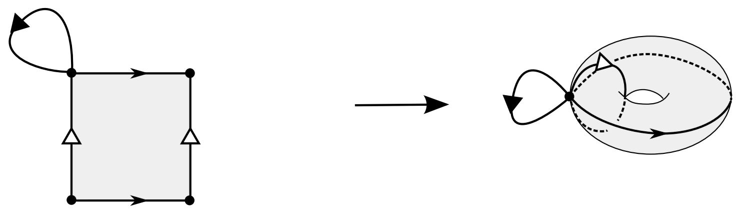

Figure 2. An immersion of the wedge of a square and circle into the wedge of a torus and circle. The target is a Salvetti complex, so we should be able to build a canonical completion.

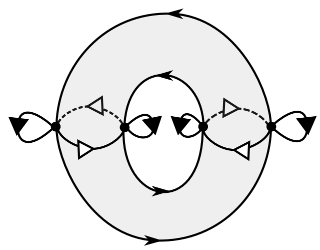

Figure 3. Here we show the completion of the map from Figure 2. The original complex X consists of the square on the lower part of the torus, in the front, together with the circle attached to the inner rim of the torus, on the left.

Definition 6.1. Let $\phi\colon K \to R$ be a combinatorial immersion from a finite graph to a rose, so the petals of the rose are $\{P_1, \ldots, P_n\}$. For each i, let K_i be the preimage of the petal P_i. This K_i is a disjoint union of points, circles, and segments. We attach edges to K_i to get a cover $\bar{K}_i \to P_i$: To each isolated point, attach a copy of P_i, and to each segment, attach a single edge joining the endpoints of the segment. The union of these K_i is a graph $\mathsf{C}_{K\to R}$ (the *completion*), with covering map

$$p\colon \mathsf{C}_{K\to R} \longrightarrow R$$

extending ϕ. Each K_i retracts to P_i, giving a retraction

$$r\colon \mathsf{C}_{K\to R} \longrightarrow K.$$

These are called the *canonical completion and retraction of* $\phi\colon K \to R$.

The following lemma will imply we can extend the preceding construction to 2–complexes.

Lemma 6.2. *(cf.* [**BRHP15**, Section 2]*) Let $K \to S$ be a locally isometric immersion from an NPC cube complex K to a Salvetti complex $S = S(\Gamma)$ with 1–skeleton R. Let $p\colon \mathsf{C}_{K^{(1)}\to R} \to R$ and $r\colon \mathsf{C}_{K^{(1)}\to R} \to K^{(1)}$ be the cover and retraction defined in 6.1.*

Let $\sigma\colon [0,4] \to S$ be the boundary map of a square of S, and let v be a vertex in $p^{-1}(\sigma(0))$. Then there is a lift of σ to a loop of length 4 based at v.

Proof. Let $\Gamma = (V, E)$ be the graph on which $S(\Gamma)$ is based, so $A(\Gamma) = \pi_1 S(\Gamma)$ is the corresponding RAAG. The path σ has label $aba^{-1}b^{-1}$ for some a and b in $V \sqcup \bar{V}$ which commute in $A(\Gamma)$.

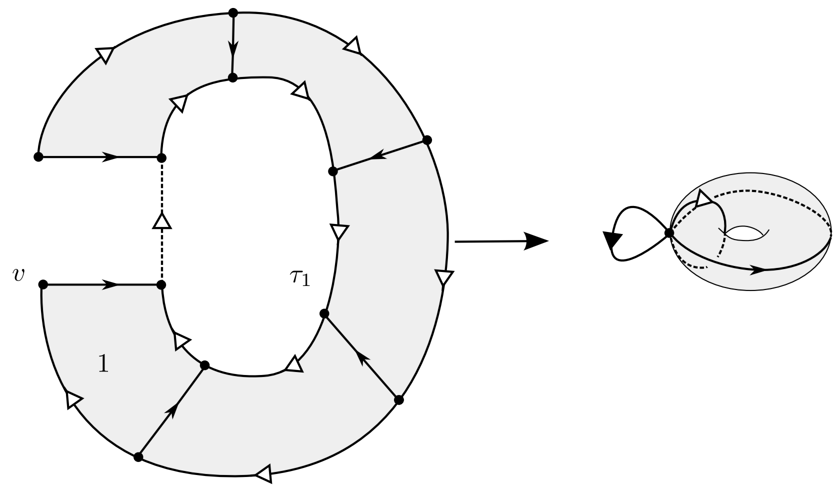

FIGURE 4. Case 1. Open arrowheads correspond to b, closed to a. The inner path τ_1 is the one completed by the new edge which is the second edge of $\tilde{\sigma}$.

Let $C = \mathtt{C}_{K^{(1)}\to R}$ be the canonical completion of the map on $K^{(1)}$. Since C contains $K^{(1)}$, we can distinguish between *old* and *new* edges of C. Let $\tilde{\sigma}$ be a lift of σ starting at v. We must show $\tilde{\sigma}(0) = \tilde{\sigma}(4)$. We'll assume that $\tilde{\sigma}$ is "non-degenerate" in the sense that every edge of $\tilde{\sigma}$ has distinct beginning and end. The "degenerate" cases will be left as an exercise.

If $\tilde{\sigma}$ passes through two consecutive old edges, then the local isometry assumption implies that those two edges span a square in K. The entire boundary of that square must be equal to the image of $\tilde{\sigma}$.

We can therefore assume that at least two of the edges of $\tilde{\sigma}$ are new edges, including one of the first two.

CASE 1. *Some edge of $\tilde{\sigma}$ is old.*

We consider only the (sub)case that the *first* edge is old. The other cases are very similar.

Since the second edge is new and non-degenerate it must complete some segment τ_1 labeled b^n terminating at $\sigma(1)$. In particular, there is a path of old edges labeled ab^{-1} starting at v. Since $K \to S$ is a locally isometric immersion, the two edges in this path corner a square (square 1 in Figure 4). In fact there must be a rectangle of squares all along τ_1. The opposite side, τ_0, of this chain of squares is also labeled by b^n, and terminates at $v = \tilde{\sigma}(0)$. Opposite the rectangle from $\tilde{\sigma}[0,1]$ is another edge labeled b, and we see that $\tilde{\sigma}|[2,3]$ must run along this (old) edge from τ_1 to τ_0. It follows that $\tilde{\sigma}[3,4]$ is a new edge. In fact τ_1 must be another maximal segment labeled b^n, and $\tilde{\sigma}[3,4]$ the new edge which completes it. Thus $\tilde{\sigma}(4) = v$ as required.

CASE 2. *All the edges of $\tilde{\sigma}$ are new edges.*

Again assuming non-degeneracy, this implies that there is a segment of old edges labeled ba going through $\tilde{\sigma}(1)$ and a segment of old edges labeled $a^{-1}b$ passing through $\tilde{\sigma}(2)$. These segments must corner squares of K (the grey squares from Figure 5), which must be part of a strip of squares joining the a–edge issuing from

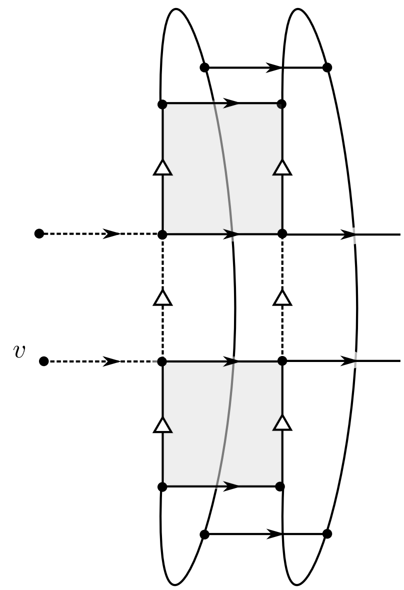

FIGURE 5. Case 2. Open arrowheads correspond to b, closed to a. Just the beginning of the rectangle (the two grey squares) is shown.

$\tilde{\sigma}(1)$ to the a–edge issuing from $\tilde{\sigma}(2)$. In particular, we get another new b–edge parallel to $\tilde{\sigma}[1,2]$. Continuing around the paths (labeled a^n) which $\tilde{\sigma}[0,1]$ and $\tilde{\sigma}[2,3]$ complete, we find a whole rectangle of squares, whose corners are exactly $\tilde{\sigma}\{0,1,2,3\}$. In the end, we discover that $\tilde{\sigma}(0)$ and $\tilde{\sigma}(3)$ are the extremities of a segment labeled b^k for some k, which can only be completed by $\tilde{\sigma}[3,4]$. Again we've shown that $\tilde{\sigma}(4) = v$.

□

EXERCISE 4. What happens in the degenerate cases we omitted from the above proof?

The following is a corollary of the preceding proof and exercise:

COROLLARY 6.3. *Let $K \to S$ be a locally isometric immersion of an NPC cube complex to a Salvetti complex, with $R = S^{(1)}$. Let $\mathsf{C}_{K^{(1)}\to R}$ be the canonical completion of the map on 1–skeleta, let $r\colon \mathsf{C}_{K^{(1)}\to R} \to K^{(1)}$ be the canonical retraction, and let $\tilde{\sigma}$ be some lift of a square boundary. Then $r \circ \tilde{\sigma}$ bounds a rectangle of squares in K.*

In particular $r \circ \tilde{\sigma}$ is null-homotopic.

We now can prove the main result of this section:

THEOREM 6.4. *For any locally isometric immersion $\phi\colon K \to S$ where K is an NPC cube complex and S is a Salvetti complex, K embeds in a finite-sheeted cover $K \subseteq \mathsf{C} \xrightarrow{p} S$, with a retraction $r\colon \mathsf{C} \to K$. Moreover, $\mathsf{C}^{(1)}$ is the completion $\mathsf{C}_{K^{(1)}\to S^{(1)}}$ described in Definition 6.1, and the maps p, r extend the maps described there.*

PROOF. Lemma 6.2 tells us how to build the 2–skeleton of C: Starting with $K^{(2)} \cup \mathsf{C}^{(1)}$, attach a 2–cell to every lift to $\mathsf{C}^{(1)}$ of the boundary of a square in S, which doesn't already bound a square in K. We obtain thereby a covering space $p\colon \mathsf{C}^{(2)} \to S^{(2)}$. Corollary 6.3 tells us we can extend the retraction to $r\colon \mathsf{C}^{(2)} \to K$.

Inductively suppose we have built $\mathsf{C}^{(n-1)}$ for some $n \geq 3$, and that we have defined a covering map $p\colon \mathsf{C}^{(n-1)} \to S^{(n-1)}$ and retraction $r\colon \mathsf{C}^{(n-1)} \to K^{(n-1)}$. We build $\mathsf{C}^{(n)}$ by attaching n–cubes to $K^{(n)} \cup \mathsf{C}^{(n-1)}$. Boundaries of n–cells in S are simply connected, so they always lift to $\mathsf{C}^{(n-1)}$, and we attach n–cubes to all lifts not already bounding cubes in K, to get a covering space $p\colon \mathsf{C}^{(n)} \to S^{(n)}$. For $\sigma\colon \partial I^n \to \mathsf{C}^{(n-1)}$ such a lift, we note that $r \circ \sigma$ has image in the NPC cube complex K, so it is contractible in K (actually in $K^{(n)}$), so we can use this contraction to extend the retraction r to the n–skeleton. □

Definition 6.5. We'll denote the cover from 6.4 either by $\mathsf{C}_{K\to S}$ or by C_ϕ, depending on whether we want to emphasize the map or the complexes.

2. Geometric separability

We note some immediate corollaries:

Corollary 6.6. *If X is a compact special cube complex, and $G = \pi_1 X$, then G is a virtual retract of some RAAG.*

Corollary 6.7. *Theorem 5.7 gives a locally isometric immersion $\phi\colon X \to S(\Gamma_X)$, so $\phi_*\colon G \to A(\Gamma_X)$ is injective. Moreover, Theorem 6.4 gives a finite-sheeted cover C of $S(\Gamma_X)$ which retracts to X. But then $\pi_1\mathsf{C} \dot{<} A(\Gamma_X)$ retracts to G.*

The following is notable in that it doesn't mention RAAGs or Salvetti complexes at all.

Corollary 6.8. *Let $f\colon X \to Y$ be a locally isometric immersion of compact special cube complexes, $G = \pi_1 Y$, and $H = f_*\pi_1 X$. Then H is separable in Y.*

Proof. Let $\phi\colon Y \to S(\Gamma_Y)$ be the locally isometric immersion from Theorem 5.7. Then $\phi \circ f\colon X \to S(\Gamma_Y)$ is a locally isometric immersion to a Salvetti complex, so we can form the canonical completion $\mathsf{C}_{\phi\circ f}$. The retraction gives a finite index $A_0 \dot{<} A(\Gamma_Y)$ which retracts onto H. But since $H < G$, we have that $G_0 = A_0 \cap G$ also retracts onto H. Since $G < A(\Gamma_Y)$, it is residually finite, and so (using Lemma 2.14) H is separable in G. □

3. What's canonical about it? (Deferring the question)

If G is the fundamental group of a special cube complex, it makes sense to call G *special*. If moreover $H < G$ is represented by a locally isometric immersion of cube complexes, we can say that H is a *geometric* subgroup of G. For many purposes, all we need to know is: *Geometric subgroups of special groups are separable.*

To know the above (ie to prove Corollary 6.8) we just needed to build *some* finite-sheeted cover to which a given immersed subcomplex lifts. But it is important later that this cover be built "canonically". We'll come back later to precisely what this means. It's more relevant when we complete maps between NPC cube complexes, neither of which is a Salvetti complex.

CHAPTER 7

Geometry of CAT(0) cube complexes

In this section we take a combinatorial approach to the geometry of CAT(0) cube complexes, very much like that discussed in Chapter 3 of Wise's CBMS notes [**Wis12**] and in Sageev's paper [**Sag95**].

1. Finding disk diagrams for null-homotopic loops

We fix a CAT(0) (meaning simply connected and NPC) cube complex X. Since X is simply connected, any combinatorial loop γ has a null-homotopy $h\colon D^2 \to X$ with $h|\partial D = \gamma$. Cellular approximation tells us that D^2 has a cell structure for which this map can be assumed cellular. In fact the map can be improved still further, in that the cell structure can be (nearly) a cube complex structure on D^2, so that h is combinatorial. The possibility that some subset of the disk with nonempty interior is forced to have 1–dimensional image means that this isn't exactly true, but still h can be chosen to factor through a combinatorial map from a planar 2–dimensional cube complex V, called a *disk diagram*.

$$\begin{array}{ccc} D^2 & \xrightarrow{\quad h \quad} & X \\ & {\scriptstyle\beta}\searrow \quad \nearrow{\scriptstyle\phi} & \\ & V & \end{array}$$

(See Figure 1 for an example.) The fact that this can be done is essentially *van*

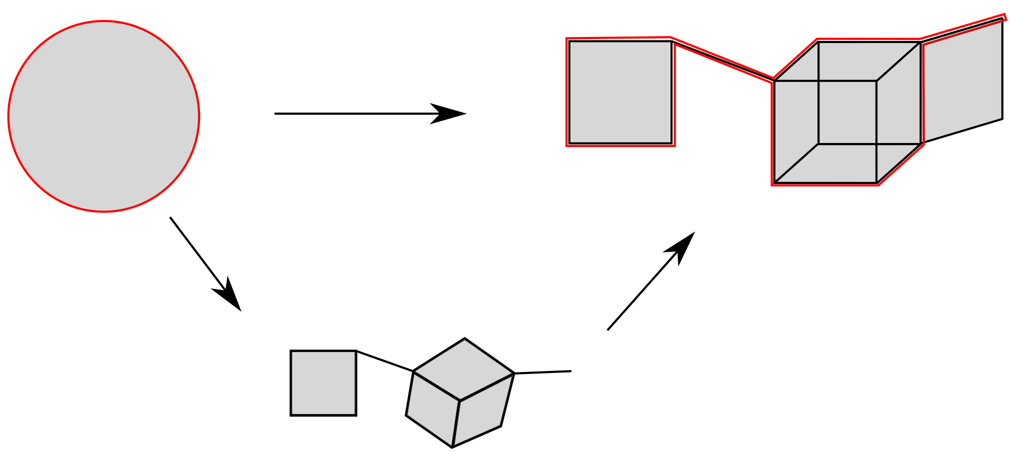

Figure 1. A disk diagram filling a combinatorial loop. The loop is not meant to be injective, but is shown as if it is for clarity.

Kampen's Lemma (see [**Bri02**]). The disk diagram is not unique. For example in Figure 1 we could have just as well used the diagram in Figure 2. Formally, we

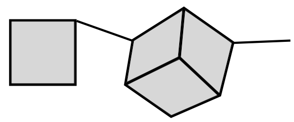

FIGURE 2. Another disk diagram filling the same combinatorial loop.

define a disk diagram over $\gamma\colon S^1 \to X$ as follows:

DEFINITION 7.1. Let X be a cube complex and let $\gamma\colon S^1$ to X be map which is combinatorial for some subdivision of S^1. A *disk diagram over* γ is a 2–dimensional square complex $V \subseteq \mathbb{E}^2$, together with maps $\beta\colon D^2 \to V$ and $\phi\colon V \to X$ so that:

(1) $(\phi \circ \beta)|\partial D^2 = \gamma$;
(2) ϕ is combinatorial;
(3) β extends to a small neighborhood $N_\epsilon(D^2)$, and β restricted to $N_\epsilon(D^2) \smallsetminus D^2$ is an orientation-preserving homeomorphism onto $N_\epsilon(V) \smallsetminus V$.

The last requirement is so that "reading" ϕ counterclockwise around V gives the same sequence of edges as reading γ counterclockwise around S^1. Sometimes we'll omit mention of the maps ϕ and β and just refer to V as a disk diagram. By the *boundary* ∂V of a diagram, we'll mean the curve $\beta|\partial D^2$.

The van Kampen Lemma implies that any combinatorial loop in a CAT(0) cube complex has a disk diagram as above. Since a disk diagram is a square complex, it has hyperplanes. These are either arcs, loops, or single points (if there is an edge which doesn't meet the interior of V). Extending these hyperplanes a bit to separate a regular neighborhood $N(V)$, we get a system of *dual curves* to V (See left hand side of Figure 3.) It's often convenient just to work with these dual

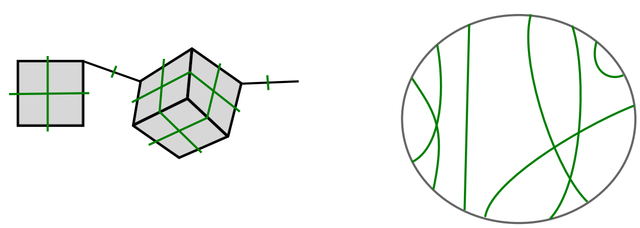

FIGURE 3. The dual curves for the disk diagram in Figure 2.

curves. The neighborhood $N(V)$ is a disk, and we can just draw the dual curves on a disk, as at the right hand side of Figure 3. The reader should check that V can be recovered from this pattern of dual curves, so we haven't really lost any information.

DEFINITION 7.2. A disk with a system of dual curves coming from a disk diagram will be called a *dual curve diagram*.

2. Features of disk diagrams

Let V be a disk diagram. Two edges of V are V*–equivalent* if they are parallel, which is the same as to say they cross the same dual curve in the dual curve diagram.

DEFINITION 7.3. Let $\gamma\colon S^1 \to X$ be a loop, where X is NPC, and S^1 has been divided into edges to make γ combinatorial. Let

$$D^2 \xrightarrow{\beta} V \xrightarrow{\phi} X \text{ and } D^2 \xrightarrow{\beta'} V' \xrightarrow{\phi'} X$$

be two disk diagrams over γ (so $\phi \circ \beta | \partial D^2 = \phi' \circ \beta' | \partial D^2 = \gamma$).

The two diagrams V and V' induce equivalence relations on the edges of S^1. We say the disk diagrams are ∂*–equivalent*, if these equivalence relations are the same.

The dual curves of V are either arcs or loops, which cross transversely. Crossings in the dual curve diagram are in one-to-one correspondence with squares of the disk diagram.

DEFINITION 7.4 (n–gons). A dual curve which is an embedded loop is called a 0*–gon*. Let $n \geq 1$, and suppose that σ is a circle made up of n consecutive arcs of dual curves $\sigma_i\colon [0,1] \to V$ so that (mod n) $\sigma_i(1) = \sigma_{i+1}(0)$ is a crossing point of dual curves. The circle σ bounds a disk D in V, and we suppose that the angle made by σ_i and σ_{i+1} inside D is $\frac{\pi}{2}$. Then σ is called an n*–gon*. See Figure 4.

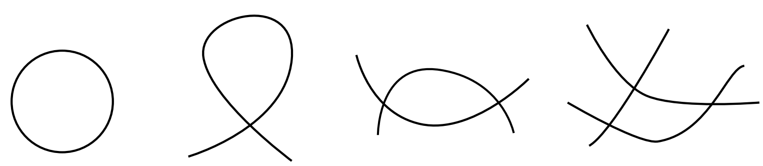

FIGURE 4. From left to right, a 0–gon, 1–gon, 2–gon, and 3–gon.

A 1–gon is also called a *monogon*; a 2–gon is called a *bigon*. An n–gon needn't be isolated; other arcs of dual curves can pass through it. If none do, we say the n–gon is *empty*. The *area* of an n–gon is the number of crossings occurring in the region bounded by the n–gon, including on the n–gon itself. See Figure 5 for some examples. In particular the area of an empty n–gon is n.

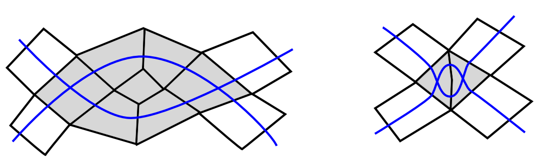

FIGURE 5. Gray squares contribute to the area of the bigon. The left hand bigon has area 7, and the right one has the minimal area possible, 2.

The following is clear from the construction:

LEMMA 7.5. *No disk diagram contains an empty* 0*–gon.*

Next we deal with minimal area monogons:

LEMMA 7.6. *Let V be a disk diagram over $\gamma\colon S^1 \to X$ where X is a CAT(0) cube complex. Then V contains no empty 1–gon.*

PROOF. Figure 6 shows the only way an empty monogon could arise in a (general) disk diagram; from a square with two incoming edges at some vertex identified. But the existence of such a square implies that the target complex X must have a

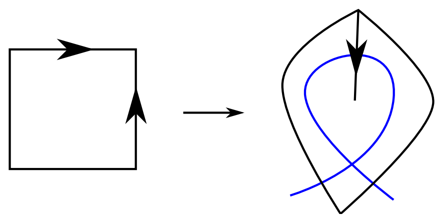

FIGURE 6. A degenerate monogon coming from a single square two of whose sides have been glued together.

link which is not simplicial. In particular it isn't flag, so X is not NPC, a contradiction. □

We'll next give some lemmas which act like "Reidemeister moves" and can sometimes be used to simplify a diagram. We start with removing a minimal 0–gon; by Lemma 7.5, this minimal area is 2.

LEMMA 7.7 (minimal 0–gon removal). *Let $\gamma\colon S^1 \to X$ where X is a CAT(0) cube complex. If a disk diagram V over γ contains a 0–gon of area 2, then there is a ∂–equivalent diagram V' with $\mathrm{Area}(V') = \mathrm{Area}(V) - 2$. The dual curve diagram of V' is obtained from the dual curve diagram of V by removing the 0–gon of area 2.*

PROOF. A minimal area monogon corresponds to a pair of squares with $\frac{3}{4}$ of their boundary identified (see Figure 7). Since the target is NPC, these squares

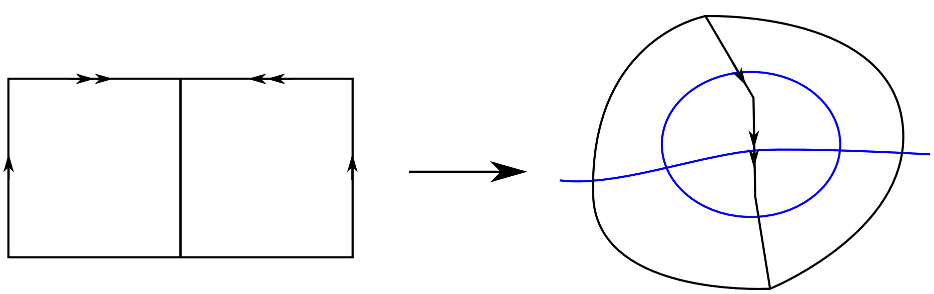

FIGURE 7. A minimal 0–gon coming from two squares which have been glued together.

must actually go to the same square, and we can obtain a new diagram V' by excising the two squares, and identifying the resulting pair of free edges. The only effect on the dual curve diagram is to remove the 0–gon, so the new diagram is ∂–equivalent to the old one. □

The next two lemmas give analogues of the Reidemeister moves for knot diagrams. See Figure 8. Note that there is no analog of Reidemeister I, since Lemma

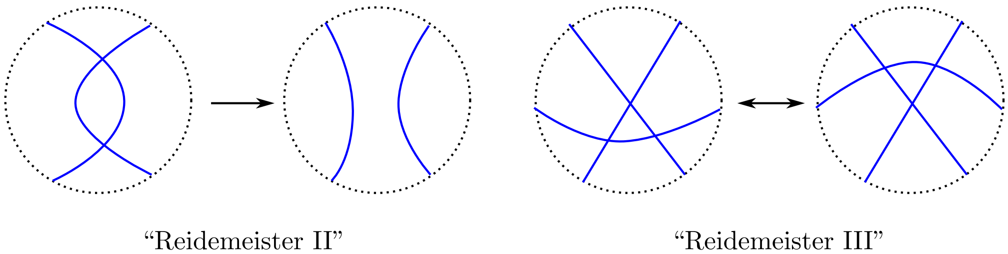

FIGURE 8. Lemmas 7.8 and 7.9 say that we can make local changes to a disk diagram which induce the pictured changes on dual curve diagrams.

7.6 rules out empty monogons.

The next lemma gives a way to remove empty bigons.

LEMMA 7.8 (minimal 2–gon removal). *Let γ be a loop in an CAT(0) cube complex X, and suppose that V is a disk diagram over γ which contains an empty 2–gon σ. Then there is a ∂–equivalent diagram V' with*

1. $\text{Area}(V') = \text{Area}(V) - 2$.
2. *If σ is not part of a minimal 0–gon, the dual curve diagrams of V' and V differ by a Reidemeister II move at σ.*

PROOF. If the diagram V contains a minimal area 0–gon, then we can appeal to the minimal 0–gon removal Lemma 7.7. So suppose that there are no minimal area 0–gons in V, but that there is an empty 2–gon, as in the right-hand side of figure 5. Because the target X is NPC, the two squares in such a bigon must map to the same square of X, and they can therefore be removed from the diagram (see Figure 9) to get a new diagram over γ. This new diagram is clearly ∂–equivalent

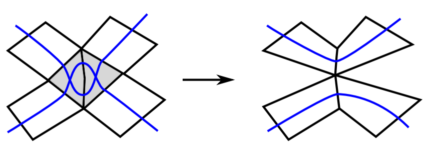

FIGURE 9. Removing a minimal bigon.

to the old one. □

Finally, there is a version of the Reidemeister III move, called a *hexagon move*:

LEMMA 7.9. *(hexagon move) Suppose γ is a loop in an CAT(0) cube complex, and that V is a disk diagram over γ containing an empty 3–gon σ. Then there is another ∂–equivalent diagram V' with*

1. $\text{Area}(V') = \text{Area}(V)$
2. *The dual curve diagrams of V and V' differ by a Reidemeister III move at σ.*

PROOF. Let $D \subset V$ be the union of the three squares corresponding to the crossings of the empty 3–gon σ. Let v be the vertex corresponding to the empty 3–gon, and let x be the vertex $\beta(v) \in X$, where β is the combinatorial map associated to the disk diagram. The subdiagram D gives a triangle T in $\mathrm{lk}(v)$. Because β is combinatorial, $\beta_v(T)$ is a path of length 3 in $\mathrm{lk}(x)$. Since $\mathrm{lk}(x)$ is simplicial, this path can only be a triangle. Since $\mathrm{lk}(x)$ is flag, this triangle is filled in with a 2–simplex, which corresponds to a 3–cube in X. There is a homotopy of β across this 3–cube to a new map β'. (This homotopy fixes all points not in the interior of D.) The new map β' is not any more combinatorial, but D can be recubulated (yielding a new, homeomorphic complex V') to make it combinatorial. Figure 10 indicates how to perform this recubulation of D: The resulting diagram V' clearly

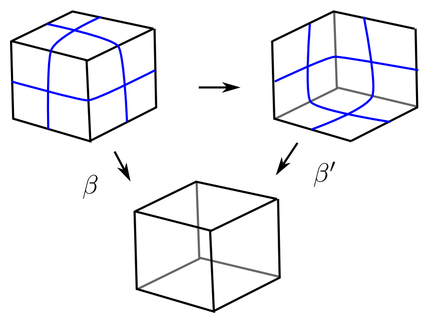

FIGURE 10. Performing a hexagon move.

has the same area and is related to the old diagram by a Reidemeister III move. ☐

We now apply these lemmas to prove that minimal area diagrams don't contain n–gons (empty or not) for $n \leq 2$:

CASSON LEMMA. *Suppose that V is a disk diagram over $\gamma\colon S^1 \to X$, where X is a CAT(0) cube complex. Suppose that D contains either*

- *an n–gon with $n \leq 2$, or*
- *a pair of adjacent V–equivalent edges.*

Then there is another disk diagram V' over γ, so that

1. $\mathrm{Area}(V') \leq \mathrm{Area}(V) - 2$, *and*
2. *V' is ∂–equivalent to V.*

PROOF. **Step 1: Find a bigon.** Any 0–gon, 1–gon, or pair of adjacent V–equivalent edges gives rise to a bigon. (See Figure 11.) Note that an adjacent pair of V–equivalent edges gives rise to a dual curve which "self-osculates" either directly or indirectly, as in one of the two pictures at the right of Figure 11. Let σ_1 be either this dual curve or the 0–gon or 1–gon. In the figure, this curve is shown in blue. The cases illustrated have plenty of squares to see what is going on. In particular, there is at least one dual curve σ_2 (shown in green) crossing σ_1, which must form a bigon with σ_1. The only way we could fail to have such an additional dual curve is that σ_1 is an empty monogon, which is forbidden by Lemma 7.6.
Step 2: Let bigons be bygones. If there is an empty 2–gon, we may finish by applying Lemma 7.8. We choose a least area bigon $\sigma = \sigma_1 \cup \sigma_2$ in V, cutting off a disk D in the dual curve diagram. We suppose $\mathrm{Area}(\sigma)$ is strictly bigger than 2,

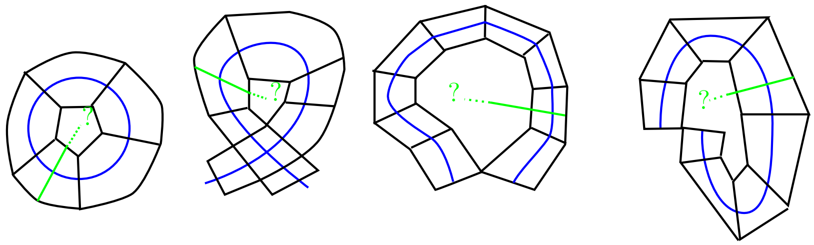

FIGURE 11. Finding bigons.

and show how to use hexagon moves to find a diagram with the same area but a smaller area bigon.

Since Area(σ) > 2, some other arc of a dual curve crosses σ_1 or σ_2. Let τ be a maximal arc of this dual curve which meets D. If both endpoints of τ are on σ_1 or on σ_2, then τ forms a bigon with that arc, which is necessarily of smaller area than σ. Thus each such arc crosses from σ_1 to σ_2. We therefore have a picture something like the left hand side of Figure 12. Let $\{\tau_1, \ldots, \tau_k\}$ be the curves crossing from σ_1

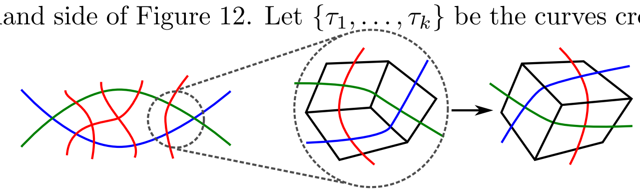

FIGURE 12. On the left, a bigon is formed by a blue and green dual curve. Red dual curves cross the bigon. If there is a red curve which forms an empty 3–gon with the dual curves, then we can do a hexagon move to decrease the area of the bigon.

to σ_2. If some τ_i forms an empty 3–gon with subarcs of σ_1 and σ_2, then we can apply Lemma 7.9 to get a diagram with the same area, but with a smaller area bigon, as in Figure 12.

Suppose at least some of the τ_i cross each other, and suppose there is an empty 3–gon involving σ_1. If so, we again have an available hexagon move (see Figure 13) to produce a new diagram V'' with a smaller bigon.

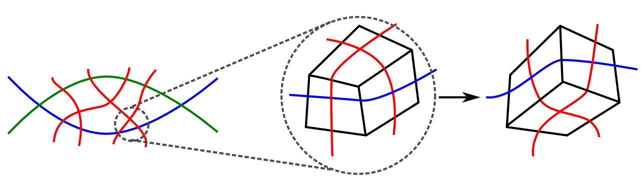

FIGURE 13. A hexagon move coming from an empty 3–gon using σ_1 but not σ_2.

It therefore suffices to prove the following claim:

CLAIM. *If $\sigma = \sigma_1 \cup \sigma_2$ is a least area bigon in the diagram V, and* Area(σ) > 2, *then there is an empty 3–gon inside σ using a subarc of σ_1.*

PROOF. The arcs crossing σ can be ordered from left to right according to where they cross σ_1. Let $\tau_1, \ldots, \tau_k$ be the ordered collection of such arcs which cross some other such arc, ignoring those which only cross σ_1 and σ_2. See the left hand side of Figure 14. Now each τ_i crosses some other τ_j after crossing σ_1. Let

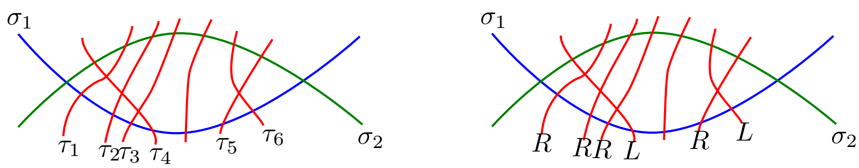

FIGURE 14. Finding an empty 3–gon.

$\tau_{j(i)}$ be the arc which τ_i crosses first. We mark τ_i by the letter R if $j(i) > i$, and by L if $j(i) < i$. See the right-hand side of Figure 14. If we had some $j(i) = i$, this would mean τ_i formed a monogon, and hence we could find a smaller area bigon than σ.

We thus obtain some sequence of R's and L's. In the example we have $RRRLRL$. Note that the first letter of this sequence must be R, and the last letter must be L. It follows that there is some initial sequence R^kL. It is then easy to see that τ_k, τ_{k+1} and σ_1 form an empty 3–gon. □

□

COROLLARY 7.10. *Let $\gamma\colon S^1 \to X$ be a combinatorial loop in a CAT(0) cube complex, and let V be a minimal area disk diagram over γ. Then V contains no n–gon for $n \le 2$, and no pair of adjacent, V–equivalent edges.*

EXERCISE 5. Show that if V is a disk diagram for $\gamma\colon S^1 \to X$ with X a CAT(0) cube complex, then V has no monogons. (Hint: use the Reidemeister moves to get an empty monogon)

3. Geodesics and hyperplanes in CAT(0) cube complexes

In this section we fix a CAT(0) cube complex X. It's important to distinguish between geodesics in X with respect to the CAT(0) metric, which need not be (and usually aren't) combinatorial, and *combinatorial* geodesics which are really geodesics in the 1–skeleton. We'll almost never talk about the first kind of geodesic in these notes. The next proposition tells us how a combinatorial geodesic can interact with a hyperplane.

PROPOSITION 7.11. *Let H be a hyperplane of X, and let γ be a combinatorial geodesic. Then H contains at most one edge dual to H.*

PROOF. Let γ be a shortest counterexample to the Proposition. Then γ begins and ends with edges e_1, e_2 which are parallel. Let n be the length of a shortest gallery of squares exhibiting the fact that e_1 and e_2 are parallel. There are two possible pictures, as shown in Figure 15. In either case one has a planar diagram V_0 consisting of a segment of length length(γ) together with n squares, and a combinatorial map of this diagram into X. The inner boundary of this diagram goes to a loop in X. Since X is simply connected, we can fill the boundary with

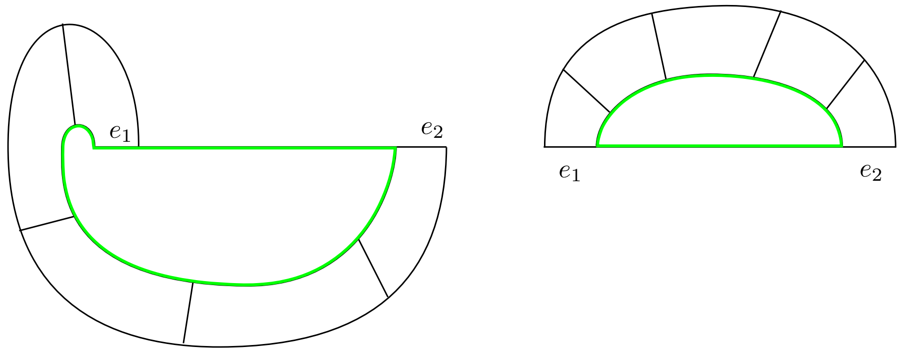

FIGURE 15. A geodesic containing parallel edges, either oriented together (as at left) or oppositely (as at right). In both pictures $n = 5$. The green curve should be filled in with a disk diagram.

some disk diagram V_1 realizing this null-homotopy. We may suppose this disk diagram is minimal area. Let V be the disk diagram $V_0 \cup V_1$.

If the edges e_1 and e_2 are oriented together along γ, as on the left hand side of Figure 15, we see that V must contain a monogon, which violates NPC as in Exercise 5.

We must therefore have that e_1 and e_2 are oriented oppositely, as in the right hand side of Figure 15. Let α be the dual curve to the gallery in V_0. Suppose first that there is another dual curve β making a bigon with α. We may suppose that this bigon B is least area among those involving α. If the bigon is nonempty, then we can argue as in the proof of the Casson Lemma that there is some empty 3–gon in B meeting α, and so we can perform a hexagon move on V to obtain a new diagram as in Figure 16. We can delete a square from this diagram to get

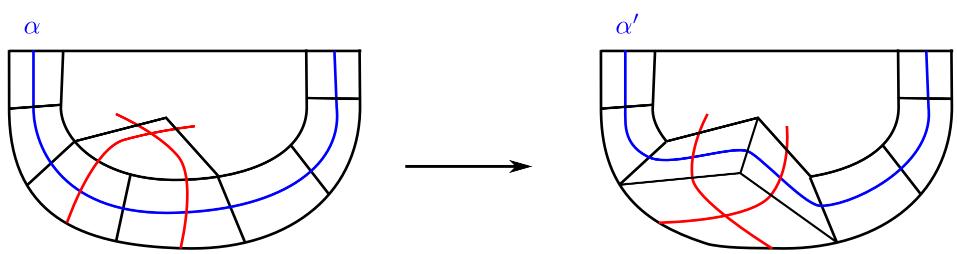

FIGURE 16. Finding a diagram with a smaller bigon, but the same length gallery connecting e_1 to e_2.

a new diagram V' of the same type as V: The new diagram again has an outer "gallery" exhibiting the fact that e_1 is parallel to e_2, and dual curve α' running through this gallery. However, the smallest area bigon involving α' is smaller area than the smallest one involving α. Eventually we obtain an empty bigon using α. Removing the two squares involved in this bigon leaves a shorter gallery connecting e_1 to e_2, contradicting our assumption that the gallery was shortest to begin with. We deduce that there were no bigons involving α.

But this implies that every dual curve crossing α also crosses the part of ∂V mapping to γ. Thus the length of that part of ∂V on the outside of the gallery is at most $|\gamma| - 2$, contradicting the assumption that γ was geodesic. □

EXERCISE 6. (1) Any square in a CAT(0) cube complex is embedded.
(2) Any (combinatorial) loop in a CAT(0) cube complex has even length.

We note a corollary of the exercise (remember that we only allow combinatorial paths).

COROLLARY 7.12. *Any path of length* 2 *in a CAT*(0) *cube complex is geodesic.*

PROPOSITION 7.13. *Let* X *be a CAT*(0) *cube complex. Then* X *is special.*

PROOF. (Sketch) We rule out the hyperplane pathologies in turn.

Let H be a hyperplane of X. If H were 1–sided, we could build a nontrivial element of $H^1(X, \mathbb{Z}/2)$, contradicting the simple connectedness of X.

If H self-osculates or self-intersects, there are a pair of distinct dual edges to H which meet at a point. These give a path of length two in X, which is geodesic by Corollary 7.12, and crosses H twice, contradicting Proposition 7.11.

Suppose H_1 and H_2 interosculate. Then there is a combinatorial map into X of a diagram like either the left or right side of Figure 17. In either case, we

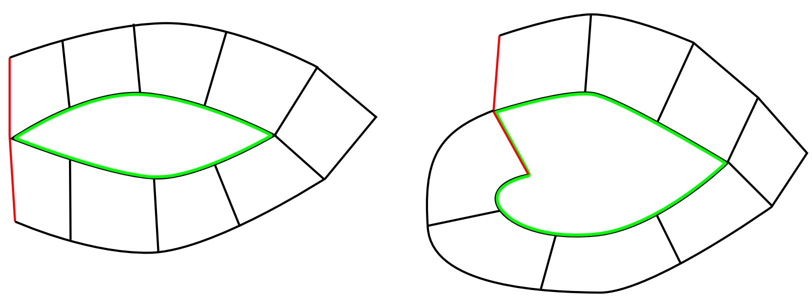

FIGURE 17. Possible interosculations. The red edges represent an osculation. Fill the green loop with a least area filling.

can fill in the diagram with some least area disk diagram. In case the diagram is like the one on the right, this filled-in diagram will contain either a monogon, contradicting Exercise 5, or a subdiagram like the left-hand interosculation picture. We can therefore assume the picture is like the one at left. In this case we'll be able to find some hexagon move reducing the area of the filled-in part. But if the filled-in part has no squares, there is a hexagon move available at the right, and we can decrease the size of the diagram. Eventually we discover that the edges which form the "osculation" actually corner a square, and so there was no inter-osculation to begin with. □

4. π_1–injectivity of locally isometrically immersed subcomplexes

One way to show that locally immersed subcomplexes of NPC cube complexes are π_1–injective would be to invoke the Cartan-Hadamard Theorem (see [**BH99**, II.4.14]). Since we are avoiding explicit use of CAT(0) geometry in these notes, we sketch a different proof. This proof also allows us to introduce Wise's *cornsquares*.

DEFINITION 7.14. Let V be a disk diagram over $\gamma\colon S^1 \to X$, where X is CAT(0). A *corner* of V is a pair of consecutive edges of ∂V which corner a square in V.

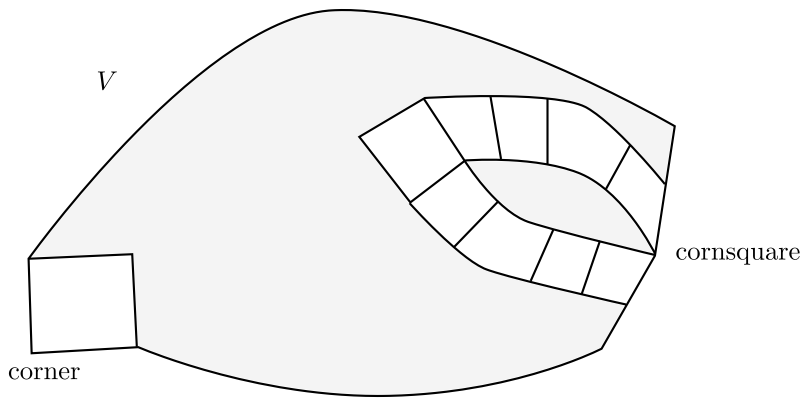

FIGURE 18. A corner and a cornsquare.

A *cornsquare* of V is a pair of consecutive edges e_1 and e_2 of ∂V which are V–parallel to edges e_1' and e_2' which corner a square.

The next lemma says that, in least area diagrams, cornsquares can be improved to corners.

LEMMA 7.15. *Let V be a least area diagram over $\gamma\colon S^1 \to X$ where X is CAT(0). If V has a cornsquare at e_1, e_2 in ∂V, then there is another least area diagram with a corner at e_1, e_2.*

PROOF. In our picture of a cornsquare (Figure 18), the galleries leading from the boundary to the edges which corner a square bound a nonempty region in the diagram V. If that is the case, though, we can apply hexagon moves to V to decrease the area of that region, eventually obtaining a diagram V' with the same area, but so the cornsquare looks like the one in Figure 19. But now we can apply

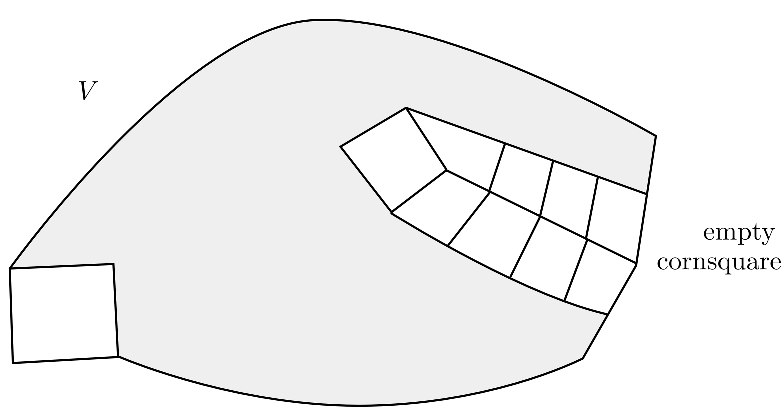

FIGURE 19. A diagram with an empty cornsquare.

further hexagon moves to decrease the length of the galleries in the cornsquare, eventually converting it to a corner. □

PROPOSITION 7.16. *If $\phi\colon Y \to X$ is a locally isometric immersion of NPC cube complexes, then $\phi_*\colon \pi_1 Y \to \pi_1 X$ is injective.*

PROOF. If ϕ is not π_1–injective, then there is some combinatorial loop σ in Y so that $\phi \circ \sigma$ is null-homotopic, so admits a finite-area disk diagram. Fix some σ which is least area among shortest such loops.[1] Now let V be a least area disk diagram over $\phi \circ \sigma$.

CLAIM. *V has a cornsquare.*

Given the claim we can replace V by a diagram V' with the same area, but with a corner. By the local isometry property, this corner comes from a square in Y, and we can homotope σ across that square to get a loop σ' of the same length as σ, but with $\mathrm{Area}(\phi \circ \sigma') < \mathrm{Area}(\phi \circ \sigma)$. This contradicts our choice of σ.

PROOF OF CLAIM. We argue using the dual curves to V, which are properly embedded arcs in $D = N_\epsilon(V)$. Note that a corner in V corresponds to a pair of dual curves which have adjacent endpoints in the boundary, and which immediately cross, forming an "empty 3–gon" together with ∂D. A cornsquare in V corresponds to a pair of dual curves which have adjacent endpoints in the boundary and *eventually* cross each other.

Choose some point $p \in \partial D$ which is not an endpoint of a dual curve. The orientation of D gives an ordering on points of $\partial D \smallsetminus \{p\}$, which restricts to an ordering of the endpoints of dual curves. Let $\mathcal{P}$ be the set of pairs of endpoints (f_1, f_2) so that $f_1 < f_2$ in this ordering, and either:

(1) f_1 and f_2 are the endpoints of a single dual curve, or
(2) the dual curves issuing from f_1 and f_2 cross.

Let (e_1, e_2) be an *innermost* element of $\mathcal{P}$: this means that if $(f_1, f_2) \in \mathcal{P}$ and $e_1 \le f_1 < f_2 \le e_2$, then $f_1 = e_1$ and $f_2 = e_2$. One quickly sees that no endpoint f of any dual curve satisfies $e_1 < f < e_2$.

If (e_1, e_2) is of type (1), so that e_1 and e_2 are the endpoints of a single dual curve α, then this curve α must correspond to a backtrack in $\phi \circ \sigma$. Since ϕ is a local isometry, there must also be a backtrack in σ, which implies that σ is not shortest, a contradiction.

If (e_1, e_2) is of type (2), then we have found a pair of adjacent dual curves which cross. □

□

The techniques of this section can also be used to establish the following:

PROPOSITION 7.17. *Let $\phi\colon Y \to X$ be a locally isometric immersion of NPC cube complexes. Then the lift $\tilde{\phi}\colon \tilde{Y} \to \tilde{X}$ to universal covers is injective, with convex image.*

EXERCISE 7. Prove Proposition 7.17.

[1] You might want to reread that sentence until it makes sense.

CHAPTER 8

Quasiconvex subcomplexes

In this subsection we prove Haglund's theorem that quasiconvex subgroups of cubulated groups can be represented by locally isometric immersed cube complexes. We should emphasize that which subgroups of a cubulated group are quasiconvex depends on the cubulation (unless the cubulated group is hyperbolic, as we'll see in Chapter 10).

DEFINITION 8.1. Let X be a geodesic space, $K \geq 0$. The subset $A \subseteq X$ is *K–quasiconvex* if every geodesic with endpoints in A lies in $N_K(A)$.

EXERCISE 8. Let Γ be the Cayley graph of $\mathbb{Z} \oplus \mathbb{Z}$ with the standard generating set. A subgroup is also a subset of Γ, so we can ask which subgroups are quasiconvex. Show:

(1) The subgroup generated by $(1, 0)$ is quasiconvex in X.
(2) The subgroup generated by $(1, 1)$ is not quasiconvex in X.

Since the two subgroups in the exercise are related by an automorphism of $\mathbb{Z} \oplus \mathbb{Z}$, this makes it seem like quasiconvexity is a pretty unstable notion. But we will continue on with it anyway. Our goal is the following theorem of Haglund:

THEOREM 8.2. *Let X be a NPC cube complex, $G = \pi_1 X$, $\tilde{X} \to X$ the universal cover. Let $H < G$ be such that some orbit $Hx_0 \subseteq \tilde{X}$ is quasiconvex.*

Then there is a compact NPC cube complex Y with $\pi_1 Y = H$ and a locally isometric immersion of cube complexes $\phi\colon Y \to X$ so that ϕ_* is the inclusion $H < G$.

This theorem combines with Corollary 6.8 to give a powerful tool for finding separable subgroups of fundamental groups of special cube complexes.

1. Median spaces

Trees are distinguished by the fact that every geodesic triangle is 0–thin (a tripod). Geodesic median spaces are those in which every triple of points is the vertex set of *some* 0–thin geodesic triangle.

DEFINITION 8.3. M a metric space is *median* if for all x, y, z in M, there is a unique m (called the *median of x, y, z*) so that

$$d(x,y) = d(x,m)+d(m,y),\ d(y,z) = d(y,m)+d(m,z), \text{ and } d(x,z) = d(x,m)+d(m,z).$$

We'll show below (Gerasimov Lemma) that the 1–skeleton of a CAT(0) cube complex is median. We first need the converse to Proposition 7.11.

LEMMA 8.4. *Let γ be a combinatorial path in the CAT(0) cube complex X. Then γ is geodesic if and only if γ crosses each hyperplane at most once.*

PROOF. Proposition 7.11 told us that geodesics cross each hyperplane at most once.

Suppose that γ is any combinatorial path which crosses each hyperplane at most once and let σ be a geodesic with the same endpoints as γ. Fill in the loop $\gamma \cdot \sigma^{-1}$ with a minimal area disk diagram, and consider dual curves starting on γ. Since γ crosses each hyperplane at most once, every such dual curve must end on σ. But this implies that the length of γ is at most the length of σ, so γ is geodesic. □

EXERCISE 9. Let H be a hyperplane in a CAT(0) cube complex, and let $N(H) \cong H \times [-\frac{1}{2}, \frac{1}{2}]$ be the union of cubes intersecting H (also called the *carrier of* H). Write $\partial N(H)$ for the subset $H \times \{-\frac{1}{2}, \frac{1}{2}\}$. This has two connected components. Show that the intersection of each with $X^{(1)}$ is convex.

We consider the possible ways a path γ could fail to be geodesic.

DEFINITION 8.5. A *backtrack* in γ is a length 2 subsegment e_1e_2, so that e_2 is just e_1 with the opposite orientation. By an *elementary shortening* of γ, we mean a subsegment of γ of the form $e_1\sigma e_2$, where e_1 is parallel to e_2 with the opposite orientation, and σ is geodesic.

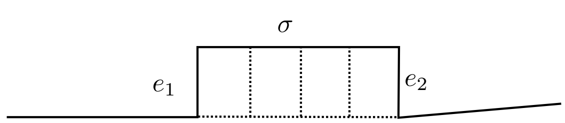

FIGURE 1. An elementary shortening.

REMARK 8.6. If γ has an *elementary shortening* $e_1\sigma e_2$, with e_1 and e_2 both dual to the hyperplane H, the endpoints of σ lie in the same component of $\partial N(H)$. By Exercise 9, σ is contained entirely in one component of $\partial N(H)$. If $\hat{\sigma}$ is the corresponding path in the other component of $\partial N(H)$, then γ can be modified to a shorter path by just replacing $e_1\sigma e_2$ with $\hat{\sigma}$.

The next lemma says that any non-geodesic in a CAT(0) cube complex is non-geodesic for one of these two simple reasons.

LEMMA 8.7. *Let X be a CAT(0) cube complex, and γ a combinatorial path. If γ is not geodesic, then γ contains either a backtrack or an elementary shortening.*

PROOF. If γ is not geodesic, then by Lemma 8.4, γ must cross some hyperplane twice. Let $e_1\sigma e_2$ be a shortest subsegment of γ which crosses a hyperplane twice. If σ is empty, then e_1e_2 is a backtrack. Otherwise σ crosses no hyperplane twice, so it is geodesic, again using 8.4. It follows that $e_1\sigma e_2$ is an elementary shortening. □

GERASIMOV LEMMA. *The 0–skeleton of any CAT(0) cube complex is median (using the metric induced from the path metric on the 1–skeleton).*

PROOF. We first prove uniqueness. Suppose there are two median points m_1 and m_2 for some triple x, y, z. Consider a hyperplane H separating m_1 from m_2. Two of x, y, z must lie on one side of H. Without loss of generality, say x, y lie on the same side as m_1. But then there is a geodesic from x to y through m_1 which crosses H twice. This contradicts Lemma 8.4.

We next establish existence. Let x, y, z be vertices of X. Define the set

$$I(x,y) = \{p \mid d(x,p) + d(p,y) = d(x,y)\},$$

and let m be a point of $I(x,y)$ closest to z. We claim that m is a median point for x, y, z. Indeed, choose geodesics $[x,m]$, $[y,m]$ and $[m,z]$. If m is not a median, then either $[x,m] \cup [m,z]$ or $[y,m] \cup [m,z]$ is not geodesic. We may as well suppose it is $\gamma = [x,m] \cup [m,z]$, and apply Lemma 8.7 to find a backtrack or elementary shortening. If there is a backtrack, it must occur at m, and we find a point in $I(x,y)$ which is closer to z. If on the other hand there is an elementary shortening, then we have some picture like Figure 2. In particular, the path γ crosses a hyperplane

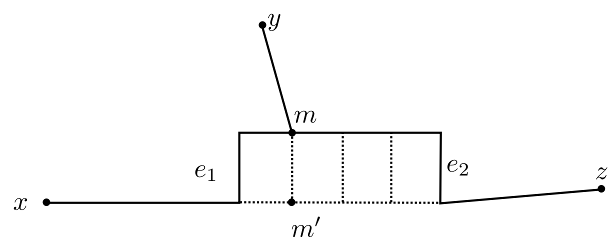

FIGURE 2. If m is not a median, we find some $m' \in I(x,y)$ closer to z.

H twice at edges e_1 and e_2, and a subsegment of the path γ looks like $e_1 \sigma e_2$, where σ is a geodesic in $\partial N(H)$. Since $[x,m]$ and $[m,z]$ are geodesic, m must occur in σ. We claim the reflection m' of m across H is

(1) in $I(x,y)$ and
(2) closer to z than m is.

The path σ can be broken into two parts, $\sigma_1 \subseteq [x,m]$, and $\sigma_2 \subseteq [m,z]$. For $i \in \{1,2\}$, let σ_i' be the path obtained by reflecting σ_i across H. Let e_m be the edge from m' to m. The path $[x,m] \cup [m,z]$ can be decomposed as $\sigma_x e_1 \sigma e_2 \sigma_z$. Notice that the path $\sigma_x \sigma_1' e_m [m,y]$ passes through m' and has the same length as the geodesic $[x,m] \cup [m,y] = \sigma_x e_1 \sigma_1 [m,y]$, so $m' \in I(x,y)$.

For the second claim, note that the path $\sigma_2' \sigma_z$ from m' to z has length one less than $[m,z]$. $\square$

2. Combinatorial hulls

DEFINITION 8.8. Let X be a CAT(0) cube complex, and $H \subseteq X$ a hyperplane. The *carrier* $N(H)$ is the union of cubes in X intersecting H. Denote the interior of $N(H)$ by $\dot{N}(H)$. Then $X \smallsetminus \dot{N}(H)$ has precisely two components, called the *halfspaces determined by* H. If we co-orient H, then we can distinguish between the *positive* halfspace H^+ (into which the edges dual to H point) and the *negative* halfspace H^-.

Let X be a CAT(0) cube complex, and let $A \subseteq X$. The *combinatorial convex hull of* A is

$$\mathrm{Hull}(A) = \bigcap \{B \mid B \text{ is a halfspace containing } A\}.$$

REMARK 8.9. By exercise 9, half-spaces are convex. (This means the 1–skeleton of a half-space is convex in $X^{(1)}$.)

We now state a theorem about combinatorial hulls which will be used to prove Theorem 8.2. A cube complex is *uniformly locally finite* if there is some n so that every vertex lies in at most n cubes.

THEOREM 8.10. *Let X be a uniformly locally finite CAT(0) cube complex. For all K, there exists an L so that: If Q is a K–quasiconvex subset of X,* Hull(Q) *is contained in the L–neighborhood of Q.*

PROOF. As a reminder, we work entirely in the 1–skeleton of X. Let Q be a K–quasiconvex subset. We will show that if $d(v, Q)$ is too large, then $v \notin \mathrm{Hull}(Q)$.

Precisely, let L be the maximum number of hyperplanes meeting a K–ball in X, and suppose that $d(v, Q) > L$. We must find a hyperplane separating v from Q.

Choose some vertex $w \in Q$ closest to v, and a geodesic γ from v to w. Let H be a co-oriented hyperplane crossed by γ, with $v \in H^+$, and suppose that H does *not* separate v from Q. In particular, there is some $u \in Q \cap H^+$. Let m be the median point of u, v, w. Since H^+ is convex, the median m lies in H^+. See Figure 3.

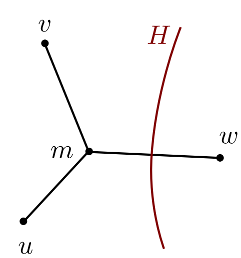

FIGURE 3. If H doesn't separate v from Q, it must come close to $w \in Q$ realizing $d(w, Q)$.

Since m lies on a geodesic from w to u, and Q is K–quasiconvex, $d(m, Q) \leq K$. Since w is closest in Q to v, we have $d(m, w) \leq K$. The hyperplane H must cross γ between m and w, so H meets a K–ball about w.

Since there are at most L hyperplanes meeting the K–ball about w, and $d(v, w) > L$, the geodesic γ crosses at least one hyperplane which *does* separate v from Q, so $v \notin \mathrm{Hull}(Q)$. ☐

Theorem 8.2 now follows from Theorem 8.10 and the following exercise:

EXERCISE 10. Let X be a CAT(0) cube complex, with $G \curvearrowright X$ freely and cocompactly. If $Q \subseteq X$, then the inclusion of Hull(Q) in X is a local isometric embedding. Thus if $x \in X$, we get $H \curvearrowright \mathrm{Hull}\, Hx$, and after taking quotients

$$H \backslash \mathrm{Hull}(Hx) \to G \backslash X$$

is a locally isometric immersion of compact NPC cube complexes inducing $H < G$ on the level of fundamental groups.

CHAPTER 9

Finding cubes: codimension one subgroups and pocsets

In this section we explain where cube complexes actually come from. The key ideas here are due to Michah Sageev, and the following account owes a lot to his Park City notes [**Sag14**].

The name "codimension one" refers to immersions of manifolds in one another. But as we will see below, not every π_1–injective codimension one submanifold has codimension one image. There is a subtle issue to do with one-or-two-sidedness.

Let G be generated by the finite set S, and let $H < G$. The *Schreier coset graph* $\mathrm{Sch}(G, H, S)$ is a graph whose vertex set is $H \backslash G$, and an edge joining Hg to Hgs whenever $g \in G$, $s \in S$.

Observe that $\mathrm{Sch}(G, H, S)$ is equal to the quotient $H \backslash \Gamma(G, S)$, under the natural isometric action.

Definition 9.1. If G is finitely generated and $H < G$, then H is *codimension one* in G if $\mathrm{Sch}(G, H, S)$ has at least two ends.

Here's an easy exercise and a harder one.

Exercise 11. Whether $H < G$ is codimension one does not depend on the finite generating set S.

Exercise 12. Let $G = \pi_1 M$ where M is a compact aspherical n–manifold, and let $\phi\colon N \to M$ be a π_1–injective immersion of some connected aspherical $(n-1)$–manifold. Then $H = \phi_*(\pi_1 N)$ is codimension one if and only if the immersion is 2–sided. (A codimension one submanifold is 2*–sided* if its normal bundle is trivial.)

1. Pocsets

The notion of a "system of half-spaces" studied by Sageev was reformalized by Roller [**Rol98**] as a *pocset*, or *poset-with-complementation.*

Definition 9.2. A *pocset* is a poset $(P, \leq)$ together with an involution $A \mapsto A^*$ satisfying:

(1) A and A^* are incomparable, and
(2) $A \leq B \implies B^* \leq A^*$.

Example 9.3. Let X be a set, and let $P \subseteq 2^X \setminus \{X, \emptyset\}$ be closed under complementation. Then P is a pocset, with involution $A^* = X \setminus A$. A pocset of this form is also sometimes called a *space with walls.* The *walls* are the pairs $\{A, A^*\}$. Usually it is also required that, for $x, y \in X$, the set $\{A \in P \mid x \in A, y \notin A\}$ is finite. (Note that Hruska–Wise's notion of a *wallspace* is different [**HW14**, Section 2].)

EXAMPLE 9.4. Let X be a cube complex, and let P be the collection of combinatorial halfspaces. Any hyperplane determines two such halfspaces, and the involution exchanges them. Order by inclusion.

Later, we'll see how to get a pocset with a G–action from a codimension one subgroup $H < G$. For now we continue with the general theory of how to turn a pocset into a cube complex.

We need a few more definitions:

DEFINITION 9.5. Let $(P, \leq)$ be a pocset, and let A, B be distinct elements of P. Say A, B are *nested* if one of the following holds:

$$A \leq B,\ A \leq B^*,\ A^* \leq B, \text{ or } A^* \leq B^*.$$

Otherwise they are *transverse*.

If $(P, \leq)$ is a space with walls, then A, B are transverse if and only if all four intersections $A \cap B, A \cap B^*, A^* \cap B, A^* \cap B^*$ are nonempty.

DEFINITION 9.6. The *width* of a pocset is the maximum number of pairwise transverse elements. If there is no such maximum number, the width is ∞.

For a pocset P coming from a finite dimensional cube complex as in 9.4, the width of P is equal to the dimension of the cube complex.

For a general pocset, we must reconstruct the cube complex. The vertices of the cube complex will be ultrafilters: "consistent" choices of A or A^* for every pair $\{A, A^*\} \subset P$.

DEFINITION 9.7. An *ultrafilter* on a pocset P is a subset $\alpha \subset P$ satisfying:

(1) (Completeness) For every $A \in P$, exactly one of $\{A, A^*\}$ is in α.
(2) (Consistency) If $A \in \alpha$, and $A \leq B$, then $B \in \alpha$.

EXAMPLE 9.8. (Principal ultrafilters) Suppose P is the pocset associated to a space with walls X. Let $x \in X$. Then $\alpha_x = \{A \in P \mid x \in A\}$ is an ultrafilter.

In "nicely behaved" spaces with walls, principal ultrafilters also satisfy the following:

DESCENDING CHAIN CONDITION. *Every sequence $A_1 > A_2 > \cdots$ of elements of α terminates.*

We say an ultrafilter is *DCC* if it satisfies this condition.

2. The cube complex associated to a pocset

Our cube complex will have 0–cells corresponding to DCC ultrafilters, and edges corresponding to "flips" $A \leftrightarrow A^*$. We use the following notation, when ω is a DCC ultrafilter on the pocset P, and $A \in \omega$:

$$(\omega; A) = (\omega \smallsetminus \{A\}) \cup \{A^*\}$$

LEMMA 9.9. *If ω is a DCC ultrafilter, the following are equivalent:*

(1) *$(\omega; A)$ is a DCC ultrafilter;*
(2) *A is minimal in ω (with respect to the order on P).*

PROOF. (1) $\implies$ (2)**.** Suppose ω and $\omega' = (\omega; A)$ are both DCC ultrafilters, and that $B < A$ is in ω. Then $A^* \in \omega'$, but $B^* \notin \omega'$. Since $A^* < B^*$, this contradicts consistency.
(2) $\implies$ (1)**.** Let $A \in \omega$ be minimal, and let $\omega' = (\omega; A)$. It is clear that ω' is complete and DCC, so we just need to check consistency. Suppose by contradiction that for some $B < C$, $B \in \omega'$ but $C \notin \omega'$. Since ω is consistent, one of B^* or C is A.

In case $B = A^* < C \notin \omega'$, we have $C^* < A$ and $C^* \in \omega$ ($C \neq A$ since A and A^* are incomparable). But then A is not minimal in ω.

In case $C = A$, we have $B < A$ in ω, and again A is not minimal. □

DEFINITION 9.10. Let P be a pocset. We define a cube complex $X = X(P)$ as follows:

- Let $X^{(0)}$ be the set of DCC ultrafilters on P.
- Connect ω to ω' by an edge if and only if $\omega' = (\omega; A)$ for a minimal $A \in \omega$.
- Inductively glue in n–cubes for $n \geq 2$, wherever the $(n-1)$–skeleton appears.

The following lemma follows from the construction:

LEMMA 9.11. *Let $\omega \in X^{(0)}$, and let $k \geq 1$. There is a one-to-one correspondence between k–cubes incident to ω and $(k-1)$–tuples of pairwise transverse minimal elements of ω.*

In particular, the link of a vertex is flag, so we have:

COROLLARY 9.12. *For P a pocset, the cube complex $X(P)$ is NPC.*

Connectedness is not guaranteed in general, but we do have the following:

LEMMA 9.13. *If P is a pocset of finite width, then $X(P)$ is connected.*

PROOF. Let ω and η be DCC ultrafilters on P. We want to find a path from ω to η. This boils down to two claims:

CLAIM. *If A is a minimal element of $\omega \smallsetminus \eta$, A is minimal in ω.*

PROOF. Exercise. □

CLAIM. *If P has finite width, then $\omega \smallsetminus \eta$ is finite.*

PROOF. Suppose that $\delta = \omega \smallsetminus \eta$ is infinite. Ramsey's theorem implies there is either an infinite collection $\mathcal{A}$ of pairwise nested or pairwise transverse elements of δ. Since P has finite width, this collection must be pairwise nested. For any two $A, B \in \mathcal{A}$, we claim that either $A < B$ or $B < A$. Otherwise either (i) $A^* < B$, or (ii) $B < A^*$. In case (i), $A^* \in \eta$ but $B \notin \eta$ contradicts consistency of η. In case (ii), $B \in \omega$ but $A^* \notin \omega$ contradicts consistency of ω.

But since $A < B$ or $B < A$ for every pair of elements of δ, we can construct an infinitely long chain of elements of δ, of one of the two forms:

$$A_1 < A_2 < A_3 < \cdots \text{ or } A_1 > A_2 > A_3 > \cdots .$$

The first kind of chain contradicts DCC of η; the second contradicts DCC of ω. □

Given the two claims, we can always choose $A \in \omega \setminus \eta$ so that $(\omega; A)$ is a DCC ultrafilter closer to η than ω was. □

REMARK 9.14. In some natural situations, we don't necessarily know finite width, but have some other way to pick out a "canonical" component of $X(P)$.

LEMMA 9.15. *Let P be a pocset. Then any component of $X(P)$ is simply connected.*

PROOF. Consider a shortest non-contractible edge loop $\sigma = e_1 \cdots e_n$. Each edge e_i connects some ω_{i-1} to $\omega_i = (\omega_{i-1}; A_i)$. Call A_i the *label* of e_i. Since $\omega_n = \omega_0$, for each i there is at least one j so that $A_j = A_i^*$. Choose such a pair i, j with $i < j$ so that $j - i$ is minimal. If $j = i + 1$, then there is a backtrack, contradicting the minimality of σ. Suppose then that there is some k between i and j.

CLAIM 9.15.1. *A_k and A_i are transverse.*

PROOF OF CLAIM. By minimality, there is no edge between e_i and e_j labeled by A_i, A_i^* or A_k^*, and only e_k is labeled by A_k. We thus have

$$\begin{aligned}
\{A_i, A_k\} &\subset \omega_{i-1} \\
\{A_i^*, A_k\} &\subset \omega_i, \omega_{k-1} \\
\{A_i^*, A_k^*\} &\subset \omega_k, \omega_{j-1} \\
\{A_i, A_k^*\} &\subset \omega_j
\end{aligned}$$

In particular, none of the possible pairs violate the consistency condition of an ultrafilter, so A_i, A_k must be transverse. □

Applying the claim with $k = i+1$, we can homotop the subpath $e_i \cdot e_{i+1} \cdots e_j$ to a path $e_i' \cdot e_{i+1}' \cdots e_j$ so that e_i' joins ω_i to $\omega_{i+1}' = (\omega_i; A_{i+1})$ and e_{i+1}' joins ω_{i+1}' to $\omega_{i+2} = (\omega_{i+1}'; A_i)$. In this way we obtain a loop of the same length but so there are edges labeled by A_i and A_i^* which are closer together than the ones in σ. Repeating this process we eventually obtain a loop with a backtrack, which again contradicts the assumption that σ has minimal length. □

To summarize:

THEOREM 9.16. *Let P be a pocset. Then any component of $X(P)$ is a CAT(0) cube complex. If P has finite width, then $X(P)$ is connected.*

We remark that if a group G acts on a pocset, this action naturally induces an action on the cube complex $X(P)$.

The following is a motivating example:

EXAMPLE 9.17. Let M be a compact hyperbolic manifold, and let F be an immersed totally geodesic submanifold of codimension one. Then any elevation $\tilde{F}$ to the universal cover $\tilde{M} \to M$ divides $\tilde{M}$ into two half-spaces. Let $P = P(M, F)$ be the pocset of such half-spaces, with order given by inclusion, and involution given by switching half-spaces.

The width of the pocset P is equal to the maximum number of pairwise crossing elevations of F, and is finite. In particular $X(P)$ is finite dimensional and connected. The fundamental group of M acts cocompactly on $X(P)$. An example is shown in Figure 1 of an immersed curve in a surface. The vertices are (in this case) in bijective correspondence with components of the complement of the preimage in the universal cover. (This is because at most two lifts pairwise cross.) There's one orbit of square, and the quotient of the cube complex by the group action is a sphere with three points identified.

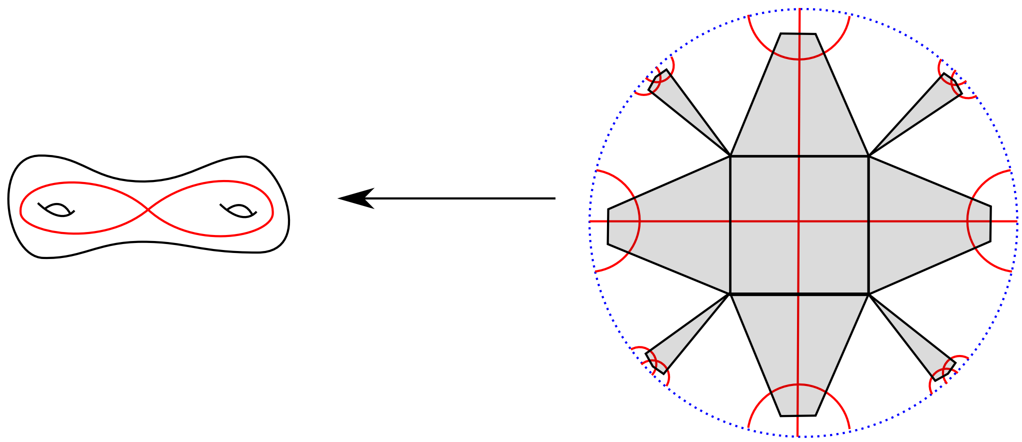

FIGURE 1. An immersed curve and part of the corresponding cube complex.

3. Codimension one subgroups

The main source of examples is a "coarsening" of Example 9.17: codimension one subgroups as defined in Definition 9.1. Remember a subgroup $H < G$ is said to be *codimension one* if the Schreier coset graph $\mathrm{Sch}(G, H, S)$ has at least two ends. In this subsection we construct a pocset from a collection of codimension one subgroups, and pick out a canonical component of the associated cube complex. This component is preserved by the natural G–action.

DEFINITION 9.18. Let $H < G$. A subset $A \subseteq G$ is *H–finite* if it is a union of finitely many right cosets of H. It is *H–infinite* if it is a union of right cosets of H, but isn't H–finite.

Recall that for two sets the *symmetric difference* $A \triangle B = (A \setminus B) \cup (B \setminus A)$.

LEMMA 9.19. *Let $H < G$ be codimension one. Then there is an $A \subset G$ satisfying the following:*

(1) *A and $G \setminus A$ are both H–infinite; and*
(2) *for every $g \in G$, $Ag \triangle A$ is H–finite.*

PROOF. Note that the right cosets of H are canonically identified with the vertices of any Schreier coset graph $\mathrm{Sch}(G, H, S)$. Fix S, and let $\Gamma = \mathrm{Sch}(G, H, S)$. Since H is codimension one, it has more than one end. Choose a neighborhood of some end, and let $E \subset \Gamma^{(0)}$ be the vertices in that neighborhood. The set A is the union of the right cosets corresponding to vertices of E. Since there is more than one end, both A and $G \setminus A$ are H–infinite.

To prove statement (2), it suffices to consider $g = s \in S$. Then As is the union of the cosets corresponding to

$$Es := \{v \mid \exists w \in E \text{ and a directed edge labeled } s \text{ from } w \text{ to } v\}.$$

The symmetric difference between E and Es consists of vertices which are either in E and connected by an edge to a vertex in $\Gamma \setminus E$, or in $\Gamma \setminus E$ and connected by an edge to a vertex in E. In particular, this set must be finite, since E is the set of vertices in the neighborhood of an end. It follows that the corresponding union of cosets $As \triangle A$ is H–finite. □

DEFINITION 9.20. If A is a subset of G satisfying the conclusions of Lemma 9.19, we will call A an *H–halfspace in G* and call the pair $\{A, G \setminus A\}$ an *H–wall in G*.

DEFINITION 9.21. Let $H_1, \dots, H_n$ be a collection of codimension one subgroup of G. For each i, let A_i be an H_i–halfspace in G, and let $\mathcal{A} = \{A_1, \dots, A_n\}$. We suppose all the sets A_i are distinct from each other and from each others complements $A_i^c := G \setminus A_i$. We define a pocset $(P_{\mathcal{A}}, \leq)$ to be

$$P_{\mathcal{A}} = \{gA \mid g \in G, A \in \mathcal{A}\} \cup \{gA^c \mid g \in G, A \in \mathcal{A}\},$$

with order given by inclusion, and involution given by complementation.

Note that the pocset defined in 9.21 gives G the structure of a *space with walls* (see 9.3). In particular, the notion of a *principal* ultrafilter $\omega_g = \{B \in P \mid g \in B\}$ (see 9.8) makes sense.

LEMMA 9.22. *Let $P_{\mathcal{A}}$ be as in Definition 9.21. Any principal ultrafilter on $P_{\mathcal{A}}$ is DCC, and any two principal ultrafilters lie in the same component of the associated cube complex $X = X(P_{\mathcal{A}})$.*

PROOF. Let $p, q \in G$, and define $\#(p, q)$ to be the number of pairs $\{B, B^*\}$ so that $p \in B$ and $q \in B^*$.

CLAIM. $\#(p, q) < \infty$.

Given the claim, we prove the lemma as follows. Let α_p, α_q be the principal ultrafilters associated to p, q, respectively. Let $D = \{B \mid p \in B, q \notin B\}$, and let $n = \#(p, q) = \#D$. We want to flip the elements of D, one at a time. Formally, let $\omega_0 = \alpha_p$. Inductively for $i \in 1, \dots n$, choose B_i minimal in $\omega_{i-1} \cap D$, and set $\omega_i = (\omega_{i-1}; B_i)$. Thus $\omega_n = \alpha_q$, and we just need to check that each ω_i is a DCC ultrafilter. By Lemma 9.9, this is true so long as B_i is always minimal in ω_{i-1}.

To show B_i is minimal in ω_{i-1}, we argue as follows. Let $B' \in \omega_{i-1}$ satisfy $B_i \subseteq B'$. Since $B_i \in D$, we have $q \notin B'$, so B' must already have been in α_p. In particular $p \in B'$, so $B' \in D \setminus \{B_1, \dots, B_{i-1}\}$. Since B_i is minimal in $\omega_{i-1} \cap D$, we must have $B' = B_i$. Thus B_i is minimal.

We now establish the claim. It suffices to consider a single set A as in the conclusion of Lemma 9.19, associated to a single codimension one subgroup H. We'll show that $\Psi = \{gA \mid p \in gA, q \in gA^c\}$ is finite; a similar argument shows that $\{gA \mid p \in gA^c, q \in gA\}$ is finite.

Let $g \in \cup\Psi$. Since $p \in gA$, we have $g^{-1} \in Ap^{-1}$. Similarly since $q \in gA^c$, we have $g^{-1} \in A^c q^{-1}$. In particular, we have g^{-1} contained in $Ap^{-1} \cap A^c q^{-1}$, which is a finite union of right cosets $Ht_1 \cup \cdots \cup Ht_s$. Equivalently, g is contained in the finite union of left cosets $t_1^{-1}H \cup \cdots t_s^{-1}H$. Since $HA = A$ this implies the finiteness of the set Ψ. □

Recall that if a pocset has finite width, the associated cube complex is connected. If the width is infinite, there may be many components. The lemma we just proved says that in the special case of a pocset coming from a collection of codimension one subgroups, there is nonetheless a distinguished component preserved by the group action.

DEFINITION 9.23 (Sageev construction). Let $\mathcal{H} = \{H_1, \dots, H_n\}$ be a collection of codimension one subgroups of a group G, let $\mathcal{A} = \{A_1, \dots, A_n\}$ and $P_{\mathcal{A}}$ be as in Definition 9.21, and let X be the component of the associated cube complex containing the principal ultrafilters. Then X is said to be obtained from $(G, \mathcal{H})$ *via the Sageev construction*.

Next we'll want to find criteria for properness or cocompactness of the action. In order to formulate such criteria, we'll need a bit of hyperbolic geometry.

Part II

Hyperbolic geometry and cube complexes

CHAPTER 10

Quasi-Isometries and Hyperbolicity

The technology of cube complexes really starts to shine when applied to *hyperbolic* (or at least relatively hyperbolic) groups. As we'll see below, these groups have a robust notion of quasiconvex subgroup, allowing us to get more mileage out of Theorem 8.2.

We'll just review the basics of Gromov hyperbolicity here. A good source for more details is [**BH99**, III.H and III.Γ]. One key feature of hyperbolicity (not shared by NPC-ness) is that it is *coarse*, in the sense that quasi-isometric groups are either both hyperbolic or both non-hyperbolic.

1. Coarse geometry

We first recall the idea of a quasi-isometry.

Definition 10.1. Let X, Y be metric spaces, $K \geq 1$, $C \geq 0$. A (not necessarily continuous) function $f\colon X \to Y$ is a (K,C)*–quasi-isometric embedding* if, for all $a, b \in X$,

$$\frac{1}{K}d(a,b) - C \leq d(f(a), f(b)) \leq K d(a,b) + C.$$

If in addition, every point $y \in Y$ lies within C of $f(x)$ for some x, then f is a (K,C)*–quasi-isometry.*

As usual for this kind of terminology, f is a *quasi-isometric embedding* if it is a (K,C)–quasi-isometric embedding for *some* K, C, and so on.

Exercise 13. If this terminology is new, you should convince yourself of the following:

(1) A composition of quasi-isometries is a quasi-isometry.
(2) Any quasi-isometry $f\colon X \to Y$ has a *quasi-inverse*; a quasi-isometry $g\colon Y \to X$ so that $f \circ g$ and $g \circ f$ are bounded distance from $\mathbf{1}_Y$ and $\mathbf{1}_X$, respectively.
(3) A metric space is quasi-isometric to a point if and only if it is bounded.
(4) A quasi-isometry of complete locally compact geodesic metric spaces induces a bijection on the set of topological ends. (So $\mathbb{R}$ is not quasi-isometric to $\mathbb{R}^2$, for example.)

Note that the first two parts of this exercise imply that quasi-isometry determines an *equivalence relation* on the class of all metric spaces.

The Schwarz–Milnor Lemma gives the fundamental connection between group theory and coarse (qi) geometry. See [**BH99**, I.8.18] for a proof.

Schwarz–Milnor Lemma. *Let X be a proper geodesic metric space on which a group G acts properly, cocompactly, by isometries. Then*

(1) *G is finitely generated by some set S, and*
(2) *any orbit map $g \mapsto gx$ is a quasi-isometry from (G, d_S) to X.*

(Here d_S is the word metric on G given by the generating set S.)

We note some Corollaries/Exercises.

Corollary 10.2. *Let G be finitely generated.*

(1) *If A and B are two finite generating sets for G, then the Cayley graphs $\Gamma(G, A)$ and $\Gamma(G, B)$ are quasi-isometric. (So G determines a unique quasi-isometry type.)*
(2) *If $H \dot{<} G$, then H is quasi-isometric to G.*
(3) *If $N \lhd G$ is finite, then G/N is quasi-isometric to G.*
(4) *If G is the fundamental group of a closed hyperbolic n–manifold, then G is quasi-isometric to $\mathbb{H}^n$.*
(5) *If G is a finitely generated free group, then G is quasi-isometric to F_2, the free group on 2 letters.*

2. Hyperbolic metric spaces

By a *triangle* we always mean a geodesic triangle, which is the union of three geodesics, the *sides* of the triangle. A triangle T is *δ–slim* if each side is contained in the δ–neighborhood of the union of the other two sides. Any particular triangle is δ–slim for δ equal to the diameter of the triangle, for example. In the euclidean plane, for any fixed δ there are triangles which fail to be δ–slim. Hyperbolic space, on the other hand, has the property that *all* triangles are δ–slim for some universal δ. It turns out that a great deal of geometry can be done using this fact alone.

Definition 10.3. Let $\delta \geq 0$. A geodesic space X is *δ–hyperbolic* if every triangle in X is δ–slim.[1] A space is said to be *Gromov hyperbolic* (or just *hyperbolic*) if it is δ–hyperbolic for some δ.

Exercise 14. (1) If $\delta_1 < \delta_2$, then every δ_1–hyperbolic space is δ_2–hyperbolic.
(2) A tree is 0–hyperbolic.
(3) For any $n \geq 2$, hyperbolic space $\mathbb{H}^n$ is δ–hyperbolic for some δ. Find a δ which works. (Hint: think about ideal triangles)
(4) Let $n \geq 4$. Any geodesic n–gon in a δ–hyperbolic space is $(n-2)\delta$–slim: any side is contained in the $(n-2)\delta$–neighborhood of the union of the other sides.

We'll see that being (Gromov) hyperbolic is a quasi-isometry invariant, though the particular δ is not. To prove this we have to understand the images of geodesics under a quasi-isometry. But this is the same as understanding quasi-isometric embeddings of intervals.

Definition 10.4. Let X be a metric space. Let $I \subseteq R$ be closed and connected. A (K, C)–quasi-isometric embedding of $\sigma \colon I \to X$ is called a *(K, C)–quasi-geodesic*. (Sometimes the image of such a map is referred to as a (K, C)–quasi-geodesic.) If $I = [a, b]$ we say the quasi-geodesic is *from $\sigma(a)$ to $\sigma(b)$*.

[1] But if you are skipping around in these notes be sure to read Remark 10.11.

Quasigeodesics in general metric spaces can be quite badly behaved. (For example $t \mapsto te^{i(1+\ln(t))}$ gives a qi embedding of $[0, \infty)$ into $\mathbb{C} = \mathbb{E}^2$ [**BH99**, Exercise I.8.23].)

However, we have the following nice statement in hyperbolic spaces (by the Hausdorff distance between two paths, we mean the Hausdorff distance between their images).

THEOREM 10.5 (Quasigeodesic Stability). [**BH99**, III.H.1.7] *Given any K, C, δ, there is an R so that:*

If σ is a (K, C)–quasigeodesic from x to y in a δ–hyperbolic space X, and γ is a geodesic from x to y, then the Hausdorff distance between σ and γ is at most R.

We refer to Bridson–Haefliger for the proof. Here are some easy consequences.

COROLLARY 10.6. *Let $g\colon Y \to X$ be a quasi-isometric embedding of geodesic metric spaces, where X is hyperbolic. Then Y is hyperbolic.*

PROOF. Fix δ so that X is δ–hyperbolic.

For some $K \geq 1$, $C \geq 0$ there is a (K, C)-quasi-isometry $g\colon Y \to X$. Let T be a geodesic triangle in Y. Then $g(T)$ is a (K, C)–quasi-geodesic triangle in X. By Quasi-geodesic stability, each side is Hausdorff distance at most R from a geodesic with the same endpoints, where $R = R(K, C, \delta)$. These geodesics form a δ–slim triangle. For y on the triangle T, there is therefore some y' on another side of T so $d(g(y'), g(y)) \leq 2R + \delta$. But then since g is a (K, C)–quasi-isometry, $d(y', y) \leq K(2R + \delta) + KC$.

Thus Y is δ'–hyperbolic for $\delta' = K(2R + \delta) + KC$. □

Similar arguments give the following.

COROLLARY 10.7. *Let $f\colon Y \to X$ be a quasi-isometric embedding of hyperbolic spaces, and let $Q \subseteq Y$. Then Q is quasi-convex if and only if $f(Q)$ is quasi-convex.*

We can therefore make the following definitions:

DEFINITION 10.8. Let G be finitely generated. G is *hyperbolic* if some (equivalently every) Cayley graph is Gromov hyperbolic.

If G is hyperbolic and $H < G$, say that H is *quasiconvex* if it is quasiconvex as a subset of some (equivalently every) Cayley graph of G.

We also note the following important observation:

LEMMA 10.9. *Let X be hyperbolic, and let A, B be subsets of X which are a finite Hausdorff distance apart. Then A is quasiconvex if and only if B is quasiconvex.*

3. Gromov products and reformulating hyperbolicity

The property of being (Gromov) hyperbolic has a number of useful reformulations. It's useful to have at least a few of these at our fingertips. (Lots more can be found, for example in Bridson–Haefliger.) We'll start with some terminology.

A *tripod* is a geodesic space which is the union of three (possibly degenerate intervals), wedged together at a point:

$$[x_1, y_1] \sqcup [x_2, y_2] \sqcup [x_3, y_3]/x_1 \sim x_2 \sim x_3. \tag{3}$$

Given three points p_1, p_2, p_3 in a metric space M, the triangle inequality implies that there is always a *comparison tripod*, i.e. a tripod T as in (3) so that $d(p_i, p_j) = (y_j - x_j) + (y_i - y_j)$ for any $i \neq j$. If M is a geodesic space, and Δ is a geodesic triangle with the points p_1, p_2, p_3 as vertices, there is always a comparison map $c\colon \Delta \twoheadrightarrow T$ which restricts to an isometry on each side. (Remember that a geodesic

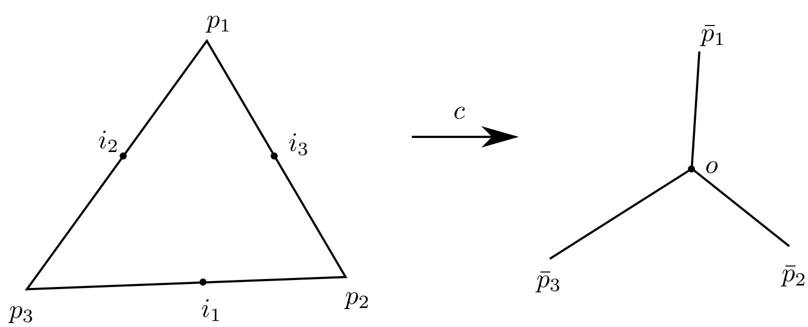

FIGURE 1. Comparison tripod and comparison map.

triangle is just a union of three geodesic segments — the triangle is not necessarily "filled in".) For each j, we refer to the point y_j as $\bar{p}_j$. We refer to the common image of the points x_j as o (see Figure 1). The points $\{i_1, i_2, i_3\} = c^{-1}(o)$ are called the *internal points* of the triangle Δ. For each $j \in \{1, 2, 3\}$, we have (considering subscripts mod 3)

$$\begin{aligned} (p_{j+1} \,|\, p_{j+2})_{p_j} &:= d(\bar{p}_j, o) = d(p_j, i_{j+1}) = d(p_j, i_{j+2}) \\ &= \frac{1}{2}\left(d(p_j, p_{j+1}) + d(p_j, p_{j+2}) - d(p_{j+1}, p_{j+2})\right). \end{aligned}$$

This quantity is called the *Gromov product of p_{j+1} and p_{j+2} with respect to p_j*.

With notation as in the previous paragraph, the *insize* of Δ is the diameter of $c^{-1}(o)$. If $\delta \geq 0$, then Δ is *δ–thin* (not to be confused with δ–slim) if $\sup\{\mathrm{diam}(c^{-1}(z)) \mid z \in T\} \leq \delta$.

With this language we can give a number of equivalent formulations of hyperbolicity.

PROPOSITION 10.10. *Let X be a geodesic space. The following are equivalent:*

(1) $\exists \delta_1 \geq 0$ *so that all triangles in X are δ_1–thin.*
(2) $\exists \delta_2 \geq 0$ *so that all triangles in X are δ_2–slim.*
(3) $\exists \delta_3 \geq 0$ *so that all triangles in X have insize at most δ_3.*

The (elementary) proof can be found in [**BH99**, III.H.1.16]. There is some slight worsening of constants as one moves from one formulation to the next. For a fixed δ, the strongest statement is (1), that all triangles are δ–thin.

REMARK 10.11. (WARNING) We will henceforth mean, by *δ–hyperbolic*, the statement that all triangles are δ–thin, not just δ–slim.

Another important interpretation of the Gromov product in hyperbolic spaces is as the approximate distance from a basepoint to a geodesic. The following is an exercise.

LEMMA 10.12. *Let $x, y, z \in X$ where X is δ–hyperbolic, and let σ be a geodesic joining x to y. Then $d(z, \sigma) - \delta \leq (x \,|\, y)_z \leq d(z, \sigma)$.*

4. Infinite hyperbolic groups have elements of infinite order

The result in the title of this section is a special application of the theory of *regular languages* to studying hyperbolic groups. That these ideas were relevant was realized by Cannon – much more information can be found in the book [**ECH⁺92**].

DEFINITION 10.13. Suppose G is finitely generated by S, and $g \in G$. The *cone type* $C(g)$ is the collection of words w in the free group on S so that

$$d_S(1, gw) = d_S(1, g) + |w|.$$

In other words, given any geodesic γ from 1 to g in the Cayley graph $\Gamma(G, S)$, the cone type $C(g)$ is the collection of paths which can be appended to γ to produce a new geodesic.

EXERCISE 15. If a generator s is in $C(g)$, then $C(gs)$ depends only on $C(g)$ and s. Associated to (G, S) there is therefore a directed labeled graph with vertices equal to possible cone types in G, and edges labeled by elements of $S \cup S^{-1}$. Draw this graph for G a free group or a free abelian group of rank 2, with S the standard generators.

THEOREM 10.14 (Cannon's cone types theorem). *Let G be hyperbolic, generated by the finite set S. There is a finite collection of cone types $\mathcal{C} = \{C_1, \ldots, C_k\}$ so that, for every $g \in G$, $C(g) \in \mathcal{C}$.*

PROOF. The idea is that the cone type of g is determined by the shape of the set of nearby elements which are *closer* to the identity than g. Let $\delta \geq 0$ be some constant so that all triangles in the Cayley graph are δ–thin.

Let $g \in G$. Define the *tail* of g to be the set of $t \in G$ satisfying both

(1) $d_S(gt, 1) < d_S(g, 1)$ and
(2) $d_S(1, t) \leq 2\delta + 3$.

CLAIM. *If g, h have the same tail, they have the same cone type.*

PROOF. We'll induct on the length of a word v in $C(g)$, showing it must also lie in $C(h)$. The base case is that v is the empty word. So we must show that if $v \in C(g) \cap C(h)$, and $s \in S$, that $vs \in C(g)$ implies $vs \in C(h)$.

Suppose not. Then $vs \in C(g) \setminus C(h)$. In particular, $d_S(1, hvs) < d_S(1, h) + |v|_S + 1$. Let γ be a geodesic from 1 to hvs, and let w be the word labeling γ. Since $v \in C(h)$, there is a geodesic σ from 1 to hv, passing through h, labeled by some word $v_h v$. Together with an edge labeled s, the geodesics γ and σ form a geodesic triangle shown in Figure 2. Write $w = w_1 w_2$, where $|w_1| = d_S(1, h) - 1$. This implies $|w_2| \leq |v| + 1$. We notice that $(hv \,|\, hvs)_1 \geq d_S(1, h) + |v| - 1$. Using δ–thin-ness, w_1 represents a group element $a \in \gamma$ with $d_S(a, h) \leq \delta + 1$. Since a is closer to 1 than h is, the group element $t = h^{-1}a$ is in the tail of h. In G we have the equality $t^{-1}vs = w_2$. The tail of g is equal to the tail of h, so $d_S(gt, 1) < d_S(g, 1)$. We note now that

$$gvs = (gt)t^{-1}vs = gtw_2.$$

Using the inequalities established already, we get

$$d_S(1, gvs) \leq d_s(1, gt) + |w_2| < d_S(g, 1) + |v| + 1,$$

contradicting the assertion that $vs \in C(g)$. □

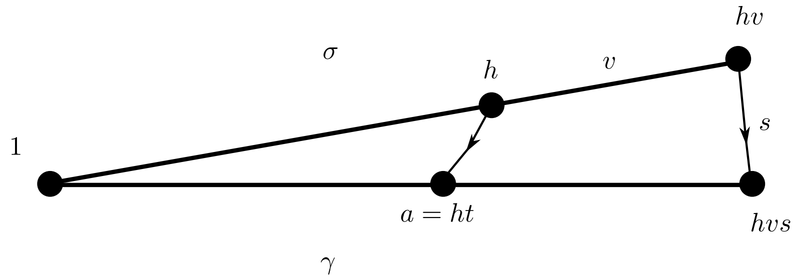

FIGURE 2. The difference between h and a lies in the tail of h.

□

Now let's prove the corollary promised in the title of the section.

COROLLARY 10.15. *If G is an infinite hyperbolic group, G contains an infinite order element.*

PROOF. This is essentially the "Pumping Lemma" from automata theory. Let k be the number of cone types with respect to some generating set S, and let $g \in G$ satisfy $d(1, g) = n > k$.

Choose a geodesic word $w = s_1 \cdots s_n$ representing g. Each prefix $w_i = s_1 \cdots s_i$ of w is also a geodesic representative of some group element g_i. There must be some $0 \leq i < j \leq n$ so that $C(g_i) = C(g_j)$. Divide up the word w into three subwords $u_1u_2u_3$ so that $u_1 = w_i$, and $u_1u_2 = w_j$. (See Figure 3.) Now clearly $u_2u_3 \in C(g_i)$.

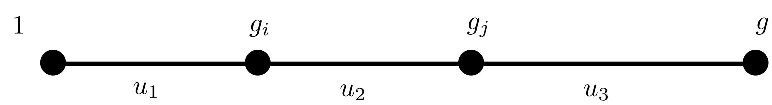

FIGURE 3. The middle word begins and ends with the same cone type.

Since $C(g_i) = C(g_j)$, we have $u_2u_3 \in C(g_j)$. But this means $u_2^2u_3 \in C(g_i) = C(g_j)$. Inductively we see that $u_2^nu_3 \in C(g_i)$ for all $i > 0$. In particular u_2^n is geodesic for every $n > 0$. In particular u_2^n is never a loop, for positive n, so the group element $g_i^{-1}g_j$ represented by u_2 must be infinite order. □

5. Quasiconvexity in cube complexes with hyperbolic π_1

THEOREM 10.16. *Let X be a compact NPC cube complex, and suppose that $G = \pi_1 X$ is hyperbolic. The following are equivalent, for $H < G$ a subgroup:*

(1) *H is quasiconvex in G.*
(2) *There is a compact NPC cube complex Y and a locally isometric immersion $\phi\colon Y \to X$ so that $\phi_*(\pi_1 Y) = H$.*

PROOF. Let $\tilde{X}$ be the universal cover of X.

(1) $\Longrightarrow$ (2): Choose $x \in \tilde{X}$ and a finite generating set S for G. The Schwarz–Milnor Lemma implies that the orbit map $g \mapsto gx$ is a quasi-isometry from (G, d_S) to $\tilde{X}$. Corollary 10.7 implies that Hx is κ–quasiconvex for some κ. Haglund's Hull theorem 8.2 implies that there is a locally isometric immersion of cube complexes as specified.

(2) $\Longrightarrow$ (1): Lift ϕ to a map $\tilde{\phi}\colon \tilde{Y} \to \tilde{X}$ of universal covers. Proposition 7.17 says that this map is one-to-one, with convex image. If $x \in \tilde{X}$, then Hx is finite Hausdorff distance from $\tilde{\phi}(\tilde{Y})$, so Hx is quasiconvex by Lemma 10.9. But this implies that $H < G$ is quasiconvex by Corollary 10.7. □

REMARK 10.17. In particular *hyperplane subgroups* (π_1–images of hyperplanes) are quasiconvex. This can be seen either by subdividing the cube complex, or by using the carrier of a hyperplane.

6. Stability of paths built from geodesic segments

In this section we prove that local geodesics and "broken geodesics" are close to geodesics. It is also possible to show such paths are quasi-geodesics (we'll do this later for broken geodesics) so one could deduce they are close to geodesics using the Quasi-geodesic Stability Theorem 10.5. However it is possible to get better bounds by attacking the question directly.

LEMMA 10.18. *Let c be a path in a δ–hyperbolic space, and suppose that $c \subset N_R(\gamma)$, where γ is a geodesic connecting the endpoints of c. Then $\gamma \subset N_{R+\delta}(c)$.*

PROOF. Suppose the paths c and γ go from p to q. If $a, b \in \gamma$ we'll write $[a, b]$ for the subsegment of γ joining them. Let $r \in \gamma$. The image of c is contained in the union of the R–neighborhoods of $[p, r]$ and $[r, q]$. Since the image is closed and connected, there is some x on c which is contained in $N_R([p, r]) \cap N_R([r, q])$. Let $p' \in [p, r]$ and $q' \in [r, q]$ be points within R of x, and consider the comparison tripod for the triangle with vertices $\{x, p', q'\}$. The image of r on this tripod is distance at most R from the image of x, so $d(r, x) \leq R + \delta$. Since $r \in \gamma$ was arbitrary, the Lemma is proved. □

LEMMA 10.19. *Let γ, σ be geodesic segments in a δ–hyperbolic space. Let x be a point of σ so that $d(x, \gamma) \geq d(y, \gamma)$ for all $y \in \sigma$. Then either*

(1) *σ is contained in a 2δ–neighborhood of γ, or*
(2) *some endpoint of σ is within 2δ of x.*

PROOF. Let x be the point on σ which is farthest from γ, and let $x' \in \gamma$ be closest to x. Join the endpoints of σ to x' and consider the comparison tripod T_1 for the resulting triangle. Let $\bar{x}$ be the image of x in this tripod. Let y be an endpoint of σ whose corresponding leg contains $\bar{x}$. (It might be the central point, in which case y can be either endpoint.) Let x'' be the point on $[y, x']$ in the preimage of $\bar{x}$, so $d(x'', y) = d(x, y)$.

Let y' be the closest point on γ to y, and consider also the triangle with vertices $\{y, y', x'\}$ and its comparison tripod T_2. The image $\bar{x'}$ of x' in T_2 is either in the leg corresponding to y or the leg corresponding to x'. If it is in the leg corresponding to x', then $d(x', \gamma) \leq \delta$, and so $d(x, \gamma) \leq 2\delta$; we are in case (1).

If $\bar{x'}$ is contained in the leg corresponding to y, then

$$d(x, y') \leq 2\delta + d(y, y') - d(x, y).$$

(Consider the tripods.) Since x is farthest from γ,

$$0 \leq d(x, y') - d(y, y') \leq 2\delta - d(x, y),$$

so $d(x, y) \leq 2\delta$ and we are in case (2). □

DEFINITION 10.20. Let $K > 0$. A K–*local geodesic* is a unit speed path so that every subpath of length at most K is geodesic.

The next statement is immediate from Lemmas 10.18 and 10.19.

COROLLARY 10.21. *Let $\epsilon > 0$ and suppose c is a $(4\delta + \epsilon)$–local geodesic. If γ is a geodesic with the same endpoints,*

(1) *c is contained in a 2δ–neighborhood of γ.*
(2) *γ is contained in a 3δ–neighborhood of c.*

The above statement is sometimes useful, but it is more common to be given a broken geodesic than a local geodesic. The following proposition gives a similar statement about broken geodesics.

PROPOSITION 10.22. *Let X be a δ–hyperbolic, geodesic space, and let $l \geq 0$. Let $c = c_1 \cdots c_n$ be a concatenation of geodesics $c_i = [p_{i-1}, p_i]$ so that*

- *(Gromov products are small) for each i, $(p_{i-1}, p_{i+1})_{p_i} \leq l$; and*
- *(segments are long) for each $i \notin \{1, n\}$, $|c_i| > 2l + 8\delta$.*

Then:

(1) *c is contained in a $(l + 3\delta)$–neighborhood of γ; and*
(2) *γ is contained in a $(l + 4\delta)$–neighborhood of c.*

PROOF. If $n \leq 2$, the Proposition follows easily from slimness of triangles, so we assume $n \geq 3$.

Item (2) follows from item (1) and Lemma 10.18, so we only need to prove item (1). Let x be the farthest point from γ on c, and let $M = d(x, \gamma)$. Without loss of generality, we suppose that $M > 2\delta$. Then Lemma 10.19 implies that x is within 2δ of some breakpoint p_i. Since $M > 2\delta$, the breakpoint p_i cannot be either endpoint of the geodesic γ; in particular $i \notin \{0, n\}$. There are two cases, depending on whether or not $i \in \{1, n-1\}$.

Suppose first that $i \notin \{i, n-1\}$. By the assumption that segments are long, $d(x, \{p_{i\pm 1}\}) > 2l + 6\delta$. Choose a geodesic σ joining p_{i-1} to p_{i+1}, and note that $d(x, \sigma) \leq l + \delta$, by the assumption on Gromov products. Let y be a closest point to p_{i-1} in γ, and let z be a closest point to p_{i+1} in γ. Choose geodesics $[y, z] \subseteq \gamma$, $[p_{i-1}, y]$, and $[p_{i+1}, z]$. The point x lies within $l + 3\delta$ of some point w on the union of these three geodesics. We claim that $w \in [y, z]$, so we have $M \leq l + 3\delta$.

Indeed, suppose that $w \in [p_{i-1}, y]$ (the case $w \in [p_{i+1}, z]$ being identical). Now we have

$$
\begin{aligned}
0 \leq d(x, y) - d(p_{i-1}, y) &\leq d(x, w) + d(w, y) - (d(p_{i-1}, w) + d(w, y)) \\
&= d(x, w) - d(p_{i-1}, w) \\
&\leq d(x, w) - (d(x, p_{i-1}) - d(x, w)) \\
&= 2d(x, w) - d(x, p_{i-1}) \\
&\leq 2(l + 3\delta) - d(x, p_{i-1}) < 0
\end{aligned}
$$

a contradiction. We have established item (1) in case $i \notin \{1, n-1\}$.

Now suppose $i \in \{1, n-1\}$. Reversing the indices if necessary, we can assume that $i = 1$. Let z be the point on γ closest to p_2. Clearly we have $d(p_2, z) \leq M$. Choose a geodesic $[p_0, p_2]$; the point x is within $l + \delta$ of a point x' on $[p_0, p_2]$. The point x' is within δ of a point w either on $[p_0, z] \subset \gamma$, or on a geodesic $[p_2, z]$. If $w \in \gamma$, we have $d(x, \gamma) \leq l + 2\delta$, and we are finished.

Suppose that $w \in [p_2, z]$. We have

$$\begin{aligned}
0 \leq d(x,\gamma) - d(p_2,z) &\leq d(x,z) - d(p_2,z) \\
&\leq d(x,w) + d(w,z) - (d(p_2,w) + d(w,z)) \\
&= d(x,w) - d(p_2,w) \\
&\leq d(x,w) - (d(p_2,x) - d(x,w)) \\
&= 2d(x,w) - d(p_2,x) \\
&< 2(l+2\delta) - (2l + 8\delta - 2\delta) \leq -2\delta,
\end{aligned}$$

a contradiction. □

CHAPTER 11

Characterizing virtually special in terms of separability

We have already seen that quasiconvex subgroups of virtually special hyperbolic groups are separable. In this chapter we see how separability can be used to remove the hyperplane pathologies. This brings us tantalizingly close to proving that if $G = \pi_1 X$ is hyperbolic, and X is a compact NPC cube complex, then X is virtually special if and only if the hyperplane subgroups are separable in G. This statement is true, but seems to require the Malnormal Quasiconvex Hierarchy Theorem of Hsu–Wise and Haglund–Wise.

We will be able to show:

Theorem 11.1. [**HW08**] *Let $G = \pi_1 X$ be hyperbolic, where X is a compact NPC cube complex. The following are equivalent:*

1. *X is virtually special.*
2. *Every quasiconvex subgroup of G is separable. ("G is QCERF.")*

A crucially important corollary of this statement is that, for G hyperbolic, virtual specialness is a property of the *group* G, and not of any particular cube complex whose fundamental group is isomorphic to G.

Corollary 11.2. *Let $\pi_1 X \cong \pi_1 Y \cong G$, where G is hyperbolic, and X, Y are compact NPC cube complexes. If X is virtually special, then so is Y.*

1. Resolving the "easy" pathologies

Recall the four hyperplane pathologies: one-sidedness, self-intersection, self-osculation, and inter-osculation. The first is easiest to resolve (assuming there are no self-intersections), and doesn't really have anything to do with separability:

Lemma 11.3. *Let X be a NPC cube complex with finitely many hyperplanes, all of which are embedded. Then there is a finite-sheeted cover $\ddot{X} \to X$ so that no hyperplane of $\ddot{X}$ is one-sided.*

No finite-sheeted cover of $\ddot{X}$ contains a one-sided hyperplane.

Proof. If H is a hyperplane of X, let $\omega_H \colon \pi_1 X \to \mathbb{Z}/2\mathbb{Z}$ measure the intersection mod 2 with H. Notice that if H is one-sided, then ω_H can't be zero, since there must be a loop in the closed carrier of H (as in Figure 1) which crosses H an odd number of times. If $\mathcal{H}$ is the set of hyperplanes, we can put all the ω_H together to get a map:

$$\Omega \colon \pi_1 X \to (\mathbb{Z}/2\mathbb{Z})^{\mathcal{H}}. \qquad (4)$$

Since X has finitely many hyperplanes, the kernel of Ω is finite index in $\pi_1 X$, and so there is a finite-sheeted cover $\ddot{X} \to X$ with $\pi_1 \ddot{X} = \ker \Omega$. If Ω contained a

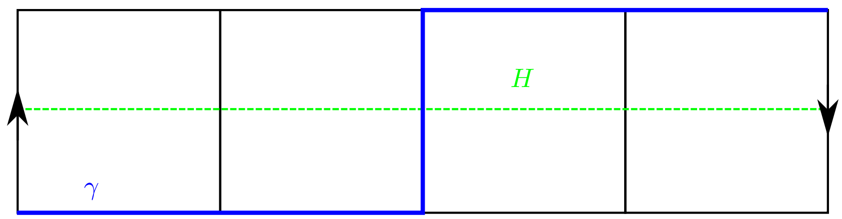

FIGURE 1. The blue loop γ witnesses the one-sidedness of H and shows that ω_H is nontrivial.

one-sided hyperplane $\tilde{H}$, there would be a loop in the carrier of $\tilde{H}$ as in Figure 1 witnessing that fact. The image of this loop in X would also cross some hyperplane an odd number of times, contradicting the fact that it lifts to $\ddot{X}$. □

REMARK 11.4. The "double-dot" cover of a cube complex just described as corresponding to the kernel of the map in equation (4) has a lot of nice properties, and will be useful for other things later.

We will remove self-intersection and self-osculation in a fairly straightforward way, using the topological characterization of separability. In separability gives us quite a bit more – we can lift hyperplanes to ones with "large embedded neighborhoods." Let's make this precise:

DEFINITION 11.5. Let H be a hyperplane of a NPC cube complex X, and let $\tilde{H} \subset \tilde{X}$ be an elevation of H to the universal cover of X. Let $N = N(\tilde{H})$ be the (closed) carrier of $\tilde{H}$ in $\tilde{X}$. The *self-interaction radius* of H, written $\mathrm{selfint}(H)$, is the smallest length of a combinatorial path joining N to some γN with $\gamma \notin \mathrm{Stab}(\tilde{H})$.

Notice that if H self-intersects or self-osculates, then $\mathrm{selfint}(H) = 0$. Observe that self-interaction radius is monotone under covers in the following sense:

LEMMA 11.6. *Let H be a hyperplane of an NPC cube complex X, and let $\hat{X} \to X$ be a cover. If $\hat{H}$ is any elevation of H to $\hat{X}$, then* $\mathrm{selfint}(\hat{H}) \geq \mathrm{selfint}(H)$.

The following says essentially that if $\pi_1 H$ is separable, then we can increase $\mathrm{selfint}(H)$ in finite covers as much as we want.

LEMMA 11.7. *Let X be a compact NPC cube complex, and let H be a compact hyperplane so that $\pi_1 H$ is separable in $\pi_1 X$. Let $n \geq 0$. Then there is a finite-sheeted regular cover $\hat{X} \to X$ so that any elevation of H to $\hat{X}$ has* $\mathrm{selfint}(H) > n$.

PROOF. Let $G = \pi_1 X$, $W = \pi_1 H$, and fix an elevation $\tilde{H}$ to the universal cover $\tilde{X}$ of X. Let S be a finite generating set for G, and endow G with the word metric coming from S. Choose x a vertex of the closed carrier N of $\tilde{H}$. The Schwarz–Milnor Lemma tells us that there is a (λ, ϵ)–quasi-isometry from G to Gx for some λ, ϵ.

Let $I = \{\gamma \mid d(N, \gamma N) \leq n\}$. We claim that I is a finite union of double cosets of the form WgW for $g \in G$. Indeed, W acts cocompactly on N. Let K be a compact fundamental domain for the action containing x. If $\gamma \in I$, then $d(N, \gamma N) \leq n$. Choose $n_1, n_2 \in N$ so that $d(n_1, \gamma n_2) = d(N, \gamma N)$. For $i \in \{1, 2\}$ there is a $w_i \in W$ so that $n_i \in w_i K$. But then we have that $d(w_1^{-1}\gamma w_2 K, K) \leq n$, and so $d(w_1^{-1}\gamma w_2 x, x) \leq n + 2\,\mathrm{diam}(K)$. This implies in G, that $d(w_1^{-1}\gamma w_2, 1) \leq$

$C(n) := \lambda(n+2\operatorname{diam}(K))+\lambda\epsilon$. Only finitely many cosets $W\gamma W$ intersect the $C(n)$ ball about the identity in G. List these cosets: $\{W, W\gamma_1 W, \ldots, W\gamma_m W\}$, choosing representatives γ_i so that $d(x, \gamma_i x)$ is minimal.

Since W is separable, there is a $G_0 \dot{<} G$ containing H but not containing any of the γ_is. Let $X_0 \to X$ be the corresponding finite-sheeted cover. Then H lifts to a hyperplane H_0 with $\operatorname{selfint}(H_0) > n$. There may be other elevations of H which do not have this property, but after passing to a regular finite-sheeted cover, all elevations will have selfint $> n$. ☐

Combining the above with Lemma 11.3, we obtain the following:

COROLLARY 11.8. *Let X be a compact NPC whose hyperplane subgroups are all separable. Then there is a finite-sheeted $\hat{X} \to X$ with every hyperplane embedded, 2–sided, and non-self-osculating.*

We'll see later that Lemma 11.7 has a lot of other applications besides dealing with hyperplane pathologies.

In order to deal with interosculation, we'll need to use separability of some subgroups obtained by amalgamating hyperplanes together. To show quasi-convexity of these subgroups, we'll need some version of a combination theorem, which we prove in the next subsection.

2. Quasiconvex combination and resolving interosculations

EXERCISE 16. Let A, B, C be NPC cube complexes, and suppose there are combinatorial isometric embeddings $\phi_A\colon C \to A$ and $\phi_B\colon B \to A$. Let X be obtained by gluing A and B together along C:

$$X = A \sqcup B / \phi_A(c) \sim \phi_B(c),\ \forall c \in C.$$

X is a NPC cube complex.

(Hint: Work out the links, and show that gluing two flag complexes along a full subcomplex yields a flag complex.)

In particular, we consider two hyperplanes $H_1, H_2 \subset X$ where X is a NPC cube complex, and suppose $\operatorname{selfint}(H_i) > 0$ for each i. This implies that the closed carriers N_1 and N_2 are embedded in X. Fix a component C of $N_1 \cap N_2$. Then $Y = N_1 \sqcup_C N_2$ is a NPC cube complex with Y immersed (not necessarily locally isometrically) in X. We can choose an elevation $\tilde{C}$ to the universal cover of X, and elevations $\tilde{N}_i$ of N_i so that $\tilde{N}_1 \cap \tilde{N}_2 = \tilde{C}$.

LEMMA 11.9. *With the notation in the previous paragraph, if γ is a shortest path between vertices p and q in $\tilde{N}_1 \cup \tilde{N}_2$, then γ is a geodesic segment in $\tilde{X}$.*

PROOF. (Recall that paths and metrics are assumed combinatorial unless otherwise stated.) If p and q are both in $\tilde{N}_1$ or $\tilde{N}_2$, there is nothing to prove, since these sets are convex. We can therefore assume $p \in \tilde{N}_1$ and $q \in \tilde{N}_2$.

We'll use the median property of the one-skeleton.

Choose some $r \in \gamma \cap \tilde{C}$, and consider the median point m of p, q, r in $\tilde{X}$. Choose particular geodesics $[p, m]$, $[m, q]$, and $[r, m]$. Since the $\tilde{N}_i$ and $\tilde{C}$ are all convex, $[r, m]$ lies in $\tilde{N}_1 \cap \tilde{N}_2 = \tilde{C}$, $[p, m]$ lies in $\tilde{N}_1$, and $[q, m]$ lies in $\tilde{N}_2$. If $r \neq m$, then $[p, m] \cup [m, q]$ is a shorter path than γ in $\tilde{N}_1 \cup \tilde{N}_2$, from p to q. Thus $r = m$, and $d_{\tilde{X}}(p, q) = d(p, r) + d(r, q) = |\gamma|$. ☐

The lemma we just proved shows that even if $\tilde{N}_1 \cup \tilde{N}_2$ is not locally isometrically embedded, it is *isometrically embedded*, in the sense that the metric induced from X agrees with its intrinsic metric.

We now use this lemma to prove a Hyperplane Combination Theorem, which we'll use (together with quasiconvex separability) to resolve interosculations.

THEOREM 11.10. *Let X be a compact NPC cube complex whose universal cover $\tilde{X}$ is δ–hyperbolic. Let H_1, H_2 be hyperplanes of X whose self-interaction radii are at least $100(\delta+1)$, and let N_1, N_2 be the closed carriers of these hyperplanes. Let C be a component of the intersection $N_1 \cap N_2$, and choose a basepoint in C. With respect to this basepoint, let $W_1 = \pi_1 N_1$, $W_2 = \pi_1 N_2$, and $Z = \pi_1 C$.*

(1) *$\langle W_1, W_2\rangle$ is 3δ–quasiconvex in $G = \pi_1 X$.*
(2) *$\langle W_1, W_2\rangle \cong W_1 *_Z W_2$.*

PROOF. **Showing that $\langle W_1, W_2\rangle$ is an amalgam:** Let $K = N_1 \sqcup_C N_2$. Seifert–van Kampen tells us that $\pi_1 K = W_1 *_Z W_2$. There is a canonical immersion $\phi\colon K \to X$. We lift this to a map $\tilde{\phi}\colon \tilde{K} \to \tilde{X}$ of universal covers If $\phi_*\colon \pi_1 K \to G$ is not one-to-one, this means that $\tilde{\phi}$ is not one-to-one. So we may choose any nontrivial geodesic σ in $\tilde{K}$ whose endpoints are identified by $\tilde{\phi}$. This geodesic decomposes a concatenation $\sigma = \sigma_1 \cdots \sigma_n$ of geodesics σ_i each of which lies in an elevation of N_1 or of N_2, alternating between elevations of the two carriers.

We claim that $n \geq 3$. Indeed, $n \neq 1$, since the copies of the $\tilde{N}_i$ embed in $\tilde{X}$. And $n \neq 2$ because the copies of the $\tilde{N}_i$ are convex.

Because of the assumption on self-interaction radius, the geodesics σ_i for $1 < i < n$ must each have length at least $100(\delta+1)$. In particular $|\sigma| > 100(\delta+1)$. Lemma 11.9 implies that $\tilde{\phi}\sigma$ is a $100(\delta+1)$–local geodesic. Corollary 10.21 implies that $\tilde{\phi}\sigma$ lies in a 2δ–neighborhood of any geodesic connecting its endpoints. But $\tilde{\phi}\sigma$ is a loop, so it must actually lie in a ball of radius 2δ. Since it contains a geodesic subsegment of length at least $100(\delta+1)$, this is a contradiction.

Showing that $\langle W_1, W_2\rangle$ is quasiconvex: Let p, q be two vertices of $\tilde{K}$, which we now identify with its image in $\tilde{X}$. Let γ be an $\tilde{X}$–geodesic joining p to q. They are also joined by a $\tilde{K}$–geodesic $\sigma = \sigma_1 \cdots \sigma_n$ as above, composed of segments σ_i each contained in an elevation of N_1 or N_2. As before, Lemma 11.9 shows the path σ is a $100(\delta+1)$–local geodesic, so Corollary 10.21 implies that γ is contained in a 3δ–neighborhood of σ. In particular γ lies in a 3δ–neighborhood of $\tilde{K}$, as required. □

We want to rule out inter-osculation. The following lemma can be proved using hexagon moves:

LEMMA 11.11. *Let X be a NPC cube complex, and let H_1, H_2 be hyperplanes of X. Suppose* $\operatorname{selfint}(H_1)$ *and* $\operatorname{selfint}(H_2)$ *are positive, and let N_1, N_2 be the closed carriers of the hyperplanes. If $N_1 \cap N_2$ is connected, then H_1 and H_2 do not inter-osculate.*

EXERCISE 17. Assume all the hypotheses of Theorem 11.10. Let $A = \langle W_1, W_2\rangle$, and let X_A be the cover of X corresponding to A. The complex $K = N_1 \sqcup_C N_2$ embeds in this cover, so the hyperplanes H_1 and H_2 lift in a canonical way to this cover. Show that these lifts do not interosculate. (If they did, there would be an osculation of lifts of N_1 and N_2 in the image of $\tilde{K}$ in $\tilde{X}$.)

EXERCISE 18. Let X be a NPC cube complex, and let H_1, H_2 be two embedded hyperplanes which don't interosculate. If $\hat{X} \to X$ is a finite-sheeted cover, and $\hat{H}_i$ is an elevation of H_i for each i, then $\hat{H}_1$ and $\hat{H}_2$ don't interosculate.

Applying the Scott's topological characterization of separability to the cover in Exercise 17, we can obtain the following:

COROLLARY 11.12. *Assuming the hypotheses of Theorem 11.10, if Z is separable in $\pi_1 X$, then there is a finite cover $\hat{X} \to X$ in which no elevations of H_1 and H_2 interosculate.*

PROOF. The exercise says that $K = N_1 \sqcup_C N_2$ embeds in the cover corresponding to $\pi_1 K = \langle W_1, W_2 \rangle$. Scott's criterion says that we can then embed K in a finite-sheeted cover $X_0 \to X$. Let $\hat{X} \to X$ be the regular cover corresponding to the normal core of $\pi_1 X_0 < \pi_1 X$.

Then $K \subset X_0$ contains elevations of H_1 and H_2 which cross but do not osculate. Passing to the regular cover $\hat{X}$, we have that *any* crossing pair of elevations fail to osculate. (Note that two elevations of H_1, say, cannot cross or osculate because of the assumption of large self-interaction radius.) □

We've now seen how to resolve all the hyperplane pathologies in finite-sheeted covers, using separability of quasiconvex subgroups. This completes the proof of Theorem 11.1.

CHAPTER 12

Reformulations of hyperbolicity, loxodromic isometries

In this chapter we return to our general discussion of hyperbolic spaces.

1. Four-point reformulations of hyperbolicity

The hyperbolicity condition can also be formulated entirely in terms of Gromov products. We remark (see Figure 1) that if x, y, z, w are any four points in a tree,

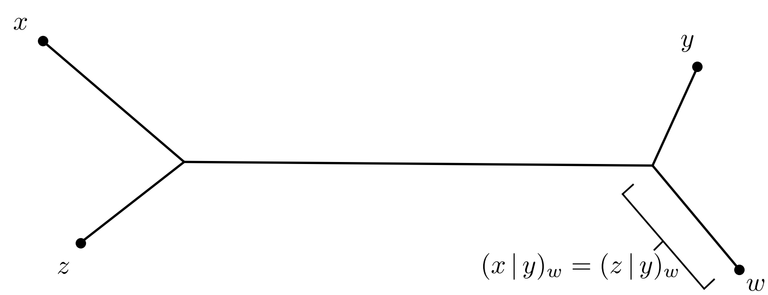

FIGURE 1. the smallest two pairwise Gromov products among three points must coincide

then the smallest two of the quantities

$$(x \mid y)_w, (x \mid z)_w, (y \mid z)_w$$

must be the same. Another way to say this is that in a tree T we have

$$\forall x, y, z, w \in T,\ (x \mid y)_w \geq \min\{(x \mid z)_w, (z \mid y)_w\}. \tag{5}$$

EXERCISE 19. Convince yourself that this is really a reformulation.

Since δ–hyperbolic spaces are "treelike", it should come as no surprise that the above statement is nearly true there. In a δ–hyperbolic space, the following holds:

$$\text{For } x, y, z, w \in X \text{ the two smallest of } \left\{ \begin{array}{l} (x \mid y)_w, \\ (x \mid z)_w, \\ (y \mid z)_w \end{array} \right\} \text{ differ by at most } \delta. \tag{6}$$

This is usually formulated in the following way:

$$(x \mid y)_w \geq \min\{(x \mid z)_w, (z \mid y)_w\} - \delta,\ \forall x, y, z, w \in X. \tag{7}$$

In trees, one also has the following more symmetric "four-point" condition:

$$\forall x, y, z, w \in T,\ d(x, w) + d(y, z) \leq \max\{d(x, y) + d(z, w), d(x, z) + d(y, w)\}. \tag{8}$$

This corresponds to the fact that, among all the ways of adding up pairs of distances between four points without repeating any points, the two *largest* are the same (See Figure 2). An equivalent condition to the hyperbolicity condition (7) is the *four-*

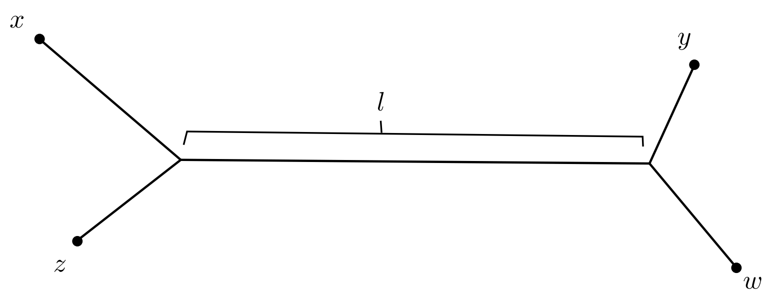

FIGURE 2. Here $d(x,y)+d(z,w)=d(x,w)+d(z,y)$ and $d(x,z)+d(y,w)$ is smaller by $2l$.

point condition: For $x,y,z,w \in X$,

$$\text{(9)} \qquad \text{The two largest of } \left\{ \begin{array}{c} d(x,y)+d(z,w), \\ d(x,z)+d(y,w), \\ d(x,w)+d(y,z) \end{array} \right\} \text{ differ by at most } 2\delta.$$

Again this is usually phrased more obscurely as:

$$\text{(10)} \qquad d(x,w)+d(y,z) \le \max \left\{ \begin{array}{c} d(x,y)+d(z,w), \\ d(x,z)+d(y,w) \end{array} \right\} + 2\delta, \ \forall x,y,z,w \in X.$$

We won't need this, but Equation (7) can be used as a *definition* of hyperbolicity, and it makes sense for metric spaces, not just for geodesic metric spaces. Using this expanded notion of hyperbolic space, any subset of a hyperbolic space, with the restricted metric, is itself a hyperbolic space. (For example one can consider the 0–skeleton of a graph.)

EXERCISE 20. Show that (6),(7),(9) and (10) hold in any δ–hyperbolic space, and that they are equivalent to each other in general.

2. Broken geodesics are quasi-geodesics

The "stability" results 10.21 and 10.22 suggest that local geodesics and broken geodesics with small Gromov products are like quasi-geodesics. We next see that they are in fact quasi-geodesics. For local geodesics it is possible to get a little nicer quantitative statement, see [**BH99**, III.H.1.13].

PROPOSITION 12.1. *Let X be a δ–hyperbolic, geodesic space, and let $l \ge 0$. Let $c = c_1 \cdots c_n$ be a concatenation of geodesics $c_i = [p_{i-1}, p_i]$ so that*

- *(Gromov products are small) for each i, $(p_{i-1}, p_{i+1})_{p_i} \le l$; and*
- *(segments are long) for each $i \notin \{1, n\}$, $|c_i| \ge R > 2(l+\delta)$.*

Then c is a quasi-geodesic whose quality depends only on l, δ, R. More precisely, for any s,t in the domain of c,

$$\text{(11)} \qquad d(c(s), c(t)) \ge \left(1 - \frac{2(l+\delta)}{R}\right)|s-t| - 4(l+\delta).$$

PROOF. Our proof is based on the following claim which will be proved inductively.

CLAIM. *For each* $i \geq 1$, $d(p_i, p_0) \geq d(p_{i-1}, p_0) + l(c_i) - 2(l + \delta)$.

Now assuming the bound, we derive the quasi-geodesic inequality (11). First we prove the inequality for $t = 0$ and s equal to the length of c:

$$\begin{aligned} d(p,q) = d(p_n, p_0) &\geq \sum_i (l(c_i) - 2(l+\delta)) \\ &= c_1 + c_n - 4(l+\delta) + \sum_{i=2}^{n-1} (l(c_i) - 2(l+\delta)) \\ &\geq c_1 + c_n - 4(l+\delta) + \sum_{i=2}^{n-1} \left(1 - \frac{2(l+\delta)}{R}\right) l(c_i) \\ &\geq -4(l+\delta) + \sum_i \left(1 - \frac{2(l+\delta)}{R}\right) l(c_i) \\ &= \left(1 - \frac{2(l+\delta)}{R}\right) l(c) - 4(l+\delta). \end{aligned}$$

Any subsegment of c also satisfies all the hypotheses, so we have the inequality for arbitrary s, t.

PROOF OF CLAIM. It may be helpful to refer to Figure 3. We argue inductively. The base case $i = 1$ is trivial. We suppose the claim has been proved for $1, \ldots, i$ where $i < n$ and prove the claim for $i + 1$.

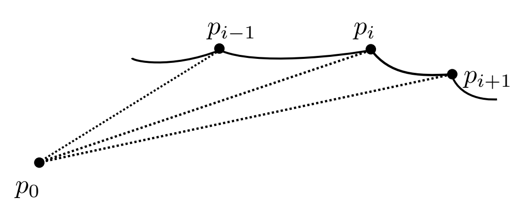

FIGURE 3. A part of the broken geodesic c.

We first assert $(p_0 \,|\, p_{i-1})_{p_i} > l + \delta$. Indeed, using the inductive hypothesis we have

$$\begin{aligned} (p_0 \,|\, p_{i-1})_{p_i} &= \frac{1}{2}(d(p_i, p_0) + l(c_i) - d(p_{i-1}, p_0)) \\ &\geq \frac{1}{2}(d(p_{i-1}, p_0) + l(c_i) - 2(l+\delta) + l(c_i) - d(p_{i-1}, p_0)) \\ &= l(c_i) - (l+\delta) \;\geq\; R - (l+\delta) \;>\; l+\delta. \end{aligned}$$

Now the two smallest of $(p_{i-1} \,|\, p_0)_{p_i}, (p_{i-1} \,|\, p_{i+1})_{p_i}, (p_0 \,|\, p_{i+1})_{p_i}$ must be within δ of one another (see Exercise 20). By hypothesis we have $(p_{i-1} \,|\, p_{i+1})_{p_i} \leq l$ so we must have $(p_0 \,|\, p_{i+1})_{p_i} \leq l + \delta$. This implies

$$d(p_0, p_i) + d(p_{i+1}, p_i) - d(p_{i+1}, p_0) \leq 2(l+\delta),$$

which can be rearranged (using $c_{i+1} = d(p_{i+1}, p_i)$) to give the inductive statement

$$d(p_{i+1}, p_0) \geq d(p_i, p_0) + c_{i+1} - 2(l+\delta).$$

□

□

REMARK 12.2. It follows that c is a (λ, ϵ)–quasi-geodesic for some λ, ϵ which can be worked out from the statement. Plugging these numbers into a generic proof of quasi-geodesic stability (eg [**BH99**, III.H.1.7]) gives some R so that c is contained in the R–neighborhood of any geodesic joining its endpoints. Note however that the R obtained in this way will be much worse than the constant $l + 3\delta$ given by Proposition 10.22.

In the next chapter we will classify isometries of hyperbolic spaces as *elliptic*, *parabolic*, and *loxodromic*. We defer the definitions of parabolic and elliptic until later.

DEFINITION 12.3. An isometry g of a hyperbolic space X is called *loxodromic* if the map $n \mapsto g^n x$ is a quasi-isometric embedding of $\mathbb{Z}$ into X for some (equivalently any) point $x \in X$.

COROLLARY 12.4. *Let g be an isometry of the δ–hyperbolic space X, and suppose that there is a point x so that*

$$d(g^2x, x) > d(gx, x) + 2\delta.$$

Then g is loxodromic.

PROOF. We may suppose $d(g^2x, x) \geq d(gx, x) + C$ where $C > 2\delta$. Consider a bi-infinite broken geodesic γ made of geodesic segments $c_i = [g^{i-1}x, g^ix]$. Each segment has length $R = d(x, gx)$. The Gromov products between any two adjacent segments are

$$\frac{1}{2}\left(d(gx, x) + d(gx, x) - d(g^2x, x)\right) \leq \frac{1}{2}\left(d(gx, x) - C\right).$$

(In particular the right hand side is non-negative.) If we set $l = \frac{1}{2}\left(d(gx, x) - C\right)$, and $R = 2l + C$, then the path γ satisfies the hypotheses of Proposition 12.1. (Or rather any subpath between g^ix and g^jx does, which is enough to show quasi-geodesicity, since the constants are independent of the subpath.) □

3. Finding loxodromic isometries

Corollary 12.4 gives a criterion for an element to be loxodromic, but it isn't obvious how to apply it. The following gives a way to use enough *non-loxodromic* elements to find a loxodromic one.

LEMMA 12.5. [**CDP90**, Chapitre 9, Lemme 2.3] *Suppose X is δ–hyperbolic, with $\delta > 0$. Let g, h be non-loxodromic isometries satisfying*

$$(*) \qquad d(gx, x) \geq 2(gx \,|\, hx)_x + 6\delta, \text{ and } d(hx, x) \geq 2(gx \,|\, hx)_x + 6\delta.$$

Then gh is a loxodromic isometry of X.

PROOF. To simplify notation, we write $|\gamma|$ for $d(x, \gamma x)$ when γ is an isometry of X. We likewise write $|\gamma - \gamma'|$ for $d(\gamma x, \gamma' x)$. The argument consists in applying the four-point inequality (8) to three carefully chosen quadrilaterals.

The isometries g, h are assumed not to be loxodromic. Corollary 12.4 implies

$$(\wr) \qquad |g^2| \leq |g| + 2\delta, \text{ and } |h^2| \leq |h| + 2\delta.$$

The assumptions $(*)$ can be rewritten (expanding the Gromov product and cancelling terms):

$$(**) \qquad |h| + 6\delta \leq |g - h|, \text{ and } |g| + 6\delta \leq |g - h|.$$

Now consider the four points x, gx, ghx, g^2x. The four-point inequality says that

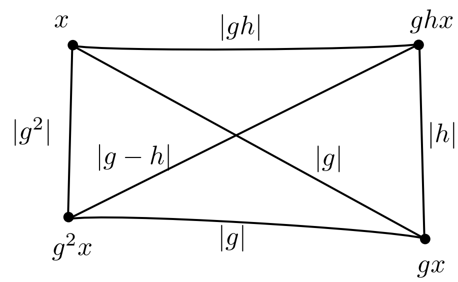

the two largest of the sums

$$|h| + |g^2|, \quad |gh| + |g|, \quad |g| + |g - h|$$

must differ from one another by at most 2δ. The inequalities $(**)$ and $(\wr)$ imply

$$|h| + |g^2| \leq |g - h| + |g| - 4\delta,$$

so the first of the three is smallest. Since the second and third sums are at most 2δ apart, $||gh| - |g - h|| \leq 2\delta$. A symmetric argument gives $||hg| - |g - h|| \leq 2\delta$.

Combining these bounds with the inequalities $(**)$ we obtain

$$(\dagger) \qquad \max\{|g|, |h|\} + 4\delta \leq \min\{|hg|, |gh|\}$$

The second quadrilateral to consider has corners $x, gx, ghx, ghgx$. Again, the

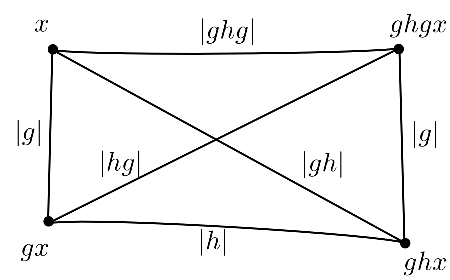

four-point inequality tells us that the two larger of the following three sums must differ by at most 2δ:

$$|g| + |g|, \quad |hg| + |gh|, \quad |h| + |ghg|.$$

Using $(\dagger)$, we get $|g| + |g| \leq |hg| + |gh| - 8\delta$, so again the first sum is smallest and the second and third differ by at most 2δ. In particular

$$|ghg| + |h| \geq |hg| + |gh| - 2\delta.$$

Since $|hg| - |h| \geq 4\delta$ (using $(\dagger)$ again), we have

$$(⬠) \qquad |ghg| \geq |gh| + 2\delta.$$

The third and final quadrilateral has corners $x, ghx, ghgx, ghghx$. The four-

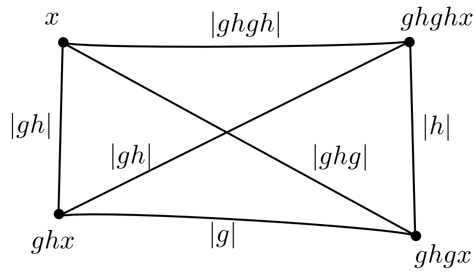

point inequality tells us the two larger of these sums must differ by at most 2δ:

$$|gh| + |h|, \quad |gh| + |ghg|, \quad |g| + |(gh)^2|.$$

Again we claim the first of the three is smallest; indeed, using $(**)$, $(\dagger)$, and $(\pentagon)$ in turn we have

$$|h| + 6\delta \le |g - h| \le |gh| + 2\delta \le |ghg|.$$

Since the second and third sums differ by at most 2δ, we deduce (using $(\pentagon)$ and then $(\dagger)$):

$$\begin{aligned} |(gh)^2| &\ge |gh| + |ghg| - |g| - 2\delta \\ &\ge |gh| + |gh| - |g| \\ &\ge |gh| + 4\delta. \end{aligned}$$

We can therefore apply Corollary 12.4 to conclude that gh is a loxodromic isometry. $\square$

CHAPTER 13

Boundaries of hyperbolic metric spaces

In this chapter we study the boundary of a hyperbolic group, and use it to study quasi-convex subgroups. In particular we show that they have the properties of finite height, finite width, and bounded packing, all of which are useful for proving things about cube complexes.

1. The boundary at infinity

Let X be a proper, geodesic, δ–hyperbolic space. Then X can be compactified by its *Gromov boundary* ∂X. We'll see that any quasiisometry $X \to Y$ of such spaces extends to a continuous map $\partial X \to \partial Y$. In particular, if X is the Cayley graph of a group G, then G acts on ∂X by homeomorphisms.[1]

1.1. The boundary as equivalence classes of rays. Points of ∂X are equivalence classes of geodesic rays $\gamma\colon [0, \infty) \to X$; the equivalence relation is that $\gamma \sim \gamma'$ if the Hausdorff distance between γ and γ' is finite.

We topologize ∂X by describing when sequences converge. First we need a lemma, which is left as an exercise:

LEMMA 13.1. *Fix $p \in X$. Then every point in ∂X is represented by a ray starting at p.*

DEFINITION 13.2 (Topology on Gromov boundary). Let $[\gamma]$ be a point in ∂X, and let $\{\gamma_i\}_{i\in\mathbb{N}}$ be a sequence of rays with the same initial point as γ. We have $\{[\gamma_i]\}_{i\in\mathbb{N}} \to [\gamma]$ in case:

There is some $K > 0$ and a sequence of positive numbers $\{t_i\}_{i\in\mathbb{N}} \to \infty$ so that, for all i, $d_{\text{Haus}}(\gamma_i|_{[0,t_i]}, \gamma|_{[0,t_i]}) \le K$.

The idea of this definition is that $[\gamma_i]$ converge to $[\gamma]$ if the rays γ_i have longer and longer initial segments which fellow travel γ. A sometimes-useful fact is that one can always take K to be a small multiple of δ, independent of the sequence.

REMARK 13.3. Definition 13.2 can be extended to describe the compactification $\bar{X} = X \cup \partial X$ if you think of a point in X as an equivalence class of maps $\gamma\colon [0, \infty) \to X$ which is geodesic on some initial subsegment, and then constant thereafter. Two such maps are equivalent if the eventually constant values coincide. Convergence is then defined in exactly the same way.

The following is a consequence of quasi-geodesic stability (Theorem 10.5).

PROPOSITION 13.4. *If $\phi\colon Y \to X$ is a quasi-isometric embedding of hyperbolic spaces, then ϕ induces a continuous embedding $\hat{\phi}\colon \partial Y \to \partial X$.*

[1] In fact the Gromov boundary makes sense even when X is not proper or geodesic; see [**BH99**, III.H] for details.

PROOF. (Exercise.) □

DEFINITION 13.5. Let $\gamma\colon \mathbb{R} \to X$ be a bi-infinite geodesic (i.e. γ is an isometric embedding). Define $\gamma^-, \gamma^+\colon [0,\infty) \to X$ by $\gamma^-(t) = \gamma(-t)$, and $\gamma^+(t) = \gamma(t)$. If $[\gamma^-] = \eta$, and $[\gamma^+] = \xi$ say γ *connects* η *to* ξ. The points η and ξ are said to be the *endpoints at infinity of* γ.

The next lemma says that such geodesics form slim bigons.

LEMMA 13.6. *Let γ_1, γ_2 be two bi-infinite geodesics with the same endpoints at infinity. Then the Hausdorff distance between γ_1 and γ_2 is at most 2δ.*

PROOF. Let x lie on the image of γ_1. Without loss of generality we may suppose $x = \gamma_1(0)$. Let $R > 0$ be any number so that $d_{\mathrm{Haus}}(\gamma_1^+, \gamma_2^+), d_{\mathrm{Haus}}(\gamma_1^-, \gamma_2^-) < R$. Choose points $a_1^\pm \in \gamma_1^\pm$ so that $d(x, a_1^\pm) > R + 2\delta$. Choose points $a_2^\pm$ on γ_2 which are within R of $a_1^\pm$, and consider a quadrilateral with cornels $a_1^-, a_1^+, a_2^+, a_2^-$ so the side containing the points $a_i^\pm$ is a subsegment of γ_i. This quadrilateral is 2δ–slim, and the sides $[a_1^+, a_2^+]$ and $[a_1^-, a_2^-]$ are further than 2δ from x, so there must be a point x' on the side which is a subsegment of γ_2 which is at most 2δ from x. □

They also form slim triangles (called *ideal triangles*), with coarsely well-defined centers. Say that a point is a *K–center* for a triangle if it is contained in the intersection of the K–neighborhoods of the sides.

LEMMA 13.7. *If $\gamma_1, \gamma_2, \gamma_3$ are bi-infinite geodesics, each pair of which shares a single endpoint at infinity, then every point on γ_1 is within 5δ of a point on γ_1 or γ_2. For any $K \geq 5\delta$, the set of K–centers is nonempty, with diameter bounded above by a constant which depends only on δ and K.*

PROOF. The proof of the first assertion is much like the proof of Lemma 13.6, but uses a hexagon instead of a quadrilateral. Letting x be on one of the sides, say γ_1, we choose three sides of the hexagon to be segments, very far away from x, joining γ_1 to γ_2, etc. The other three sides are subsegments of the sides of the ideal triangle. The point x must be within 5δ of some other side of the hexagon, but the only sides close enough are subsegments of γ_2 and γ_3.

Now let $K \geq 5\delta$. The side γ_1 is contained in the 5δ–neighborhood of $\gamma_2 \cup \gamma_3$. Since each pair of γ_i shares only one point at infinity, there must be points of γ_1 which are 5δ–close to both γ_2 and γ_3. The set of 5δ–centers is nonempty, and the set C of K–centers contains the set of 5δ–centers.

Let $a, b \in C$. We must bound $d(a,b)$ from above. For $1 \leq i < j \leq 3$ choose a point x_{ij} so that $d(x_{ij}, \gamma_i), d(x_{ij}, \gamma_j) \leq 5\delta$, but $d(x_{ij}, \gamma_k) \geq 100\delta + 2K$ for $k \notin \{i,j\}$. Each pair of these points is within 5δ of just one of the γ_i; call the geodesic segment joining them σ_i. Using 2δ–slimness of quadrilaterals, there must be points $a_i, b_i \in \sigma_i$ so that $d(a, a_i), d(b, b_i) \leq K + 2\delta$ for each i. Let T be the comparison tripod for the triangle Δ made of the segments σ_i. The set $\{a_i\}$ must project to at least two legs, and can have diameter at most $2K + 2\delta$. It follows that the set $\{\overline{a}_i\}$ is contained in the $2K + 3\delta$–neighborhood of the central point. In particular, each a_i is within $2K + 4\delta$ of any internal point z of the triangle Δ, so $d(a,z) \leq 3K + 4\delta$. The same argument shows $d(b,z) \leq 3K + 4\delta$, so $d(a,b) \leq 6K + 8\delta$. □

EXERCISE 21. Let X be proper and hyperbolic. Then any two points at infinity are connected by a bi-infinite geodesic. (Hint: pick representatives α, β whose distance from one another is always fairly large. Then use thin-ness of quadrilaterals to show that geodesic segments $[\alpha(n), \beta(n)]$ all pass through some compact set, and properness to show these sub-converge to a bi-infinite geodesic.)

1.2. Gromov product definition and limit sets. There is another way to define the boundary at infinity, which works whenever X is δ–hyperbolic, even if X isn't proper. Namely, fix a basepoint p (which in the end won't matter). Say the sequence $\{a_i\}_{i\in\mathbb{N}}$ *converges to infinity* if $\lim_{i,j\to\infty}(a_i \,|\, a_j)_p = \infty$. Now define points in the *(sequential) Gromov boundary* $\partial_s X$ to be equivalence classes of sequences which converge to infinity, under the following equivalence relation: $\{a_i\}_{i\in\mathbb{N}} \sim \{b_i\}_{i\in\mathbb{N}}$ whenever $\lim_{i,j\to\infty}(a_i \,|\, b_j)_p = \infty$. If $p = [\{a_i\}]$ we can write $\lim_{i\to\infty} a_i = p$. We are going to topologize $X \cup \partial_s X$ so this really is a convergent sequence.

First we extend the Gromov product to infinity; if $\eta, \xi \in X \cup \partial_s X$, define

$$(\eta \,|\, \xi)_p = \sup \left\{ \liminf_{i,j\to\infty}(a_i \,|\, b_j)_p \;\middle|\; \lim_{i\to\infty} a_i = \eta, \lim_{i\to\infty} b_i = \xi \right\}. \tag{12}$$

Ok, this is kind of ugly, but what you want to think about is that *if* you could draw geodesic rays from p to the points η and ξ, then $(\eta \,|\, \xi)_p$ measures (up to an additive constant) the length of maximal initial subsegments of η and ξ which δ–fellow travel. Rays which fellow-travel for more time should be considered closer, so we say that $\lim_{k\to\infty} \eta_k = \xi$ in $\partial_s X$ exactly if $\lim_{k\to\infty}(\eta_k \,|\, \xi)_p = \infty$. Points $\{a_i\}$ in X converge to a point in $\partial_s X$ exactly when $\lim_{i,j\to\infty}(a_i \,|\, a_j)_p = \infty$, in which case they converge to $[\{a_i\}]$. We thus get a topology on all of $X \cup \partial_s X$.

PROPOSITION 13.8. *If X is proper and δ–hyperbolic, then $X \cup \partial X$ is homeomorphic to $X \cup \partial_s X$ by a homeomorphism which restricts to the identity on X.*

PROOF. Sketch: Define a map from ∂X to $\partial_s X$ by sending the equivalence class of the ray γ to the equivalence class of the sequence $\{\gamma(i)\}_{i\in\mathbb{N}}$. This map doesn't need properness. To define a map in the other direction, fix a basepoint p. If $\{x_i\}_{i\in\mathbb{N}}$ represents a point of $\partial_s X$, choose geodesic segments $[p, x_i]$. Arzela-Ascoli (properness is used here) can be used to show these segments subconverge to a geodesic ray. For more see [**BH99**, III.H.3]. □

DEFINITION 13.9. If $Z \subset X$ is any subset of the δ–hyperbolic space X, then the *limit set* $\Lambda(Z)$ is that part of $\partial_s X$ which can be represented by sequences $\{z_i\}_{i\in bN}$ of points in Z. If $G \curvearrowright X$ is a group of isometries of X, we define the limit set $\Lambda(G)$ to be the limit set of any orbit of G.

One final remark about the Gromov product at infinity: The supremum is somewhat arbitrary, and could be replaced by an infimum without changing anything essential. This is because of the following.

LEMMA 13.10. *Let $\eta, \xi \in \partial X$, and suppose $\{a'_i\}_{i\in\mathbb{N}}, \{b'_i\}_{i\in\mathbb{N}}$ be arbitrary sequences representing η and ξ, respectively. Then*

$$\liminf_{i,j\to\infty}(a'_i \,|\, b'_j)_p \geq (\eta \,|\, \xi)_p - 4\delta$$

PROOF. Let $\epsilon > 0$, and let $\{a_i\}_{i\in\mathbb{N}}, \{b_i\}_{i\in\mathbb{N}}$ be sequences so that $\liminf_{i,j\to\infty}(a_i \,|\, b_j)_p > (\eta \,|\, \xi)_p - \epsilon$.

Choose large i, j, i', j', we have $(a_i \,|\, a'_{i'})_p$ and $(b_j \,|\, b'_{j'})_p$ much larger than $(\eta \,|\, \xi)_p$, and so that $(a_i \,|\, b_i)_p$ is within ϵ of $(\eta \,|\, \xi)_p$. Let σ be a geodesic joining x_i to y_j and let σ' be a geodesic joining $x'_{i'}$ to $y'_{i'}$. It is not too hard to see under these conditions that $d(p, \sigma)$ and $d(p, \sigma')$ differ by at most 2δ. Thus the Gromov products $(x_i \,|\, y_i)_p$ and $(x'_{i'} \,|\, y'_{i'})_p$ differ by at most 4δ (See Lemma 10.12). Since the indices i', j' were essentially arbitrary, we get

$$\liminf_{i,j\to\infty}(x'_i \,|\, y'_j)_p \geq \liminf_{i,j\to\infty}(x_i \,|\, y_j)_p - 4\delta,$$

and the result follows. □

There is an analog of Lemma 13.7 for triples of points in the sequential boundary. Let p, q, r be distinct points in $\partial_s X$. Say that $x \in X$ is a *K–center* for the triple p, q, r if the following holds: For all sequences $\{p_i\}_{i\in\mathbb{N}}, \{q_i\}_{i\in\mathbb{N}}, \{r_i\}_{i\in\mathbb{N}}$ there is an N so that if $i, j, k \geq N$, then x is in the K–neighborhood of any geodesic between p_i, q_j, r_k. A slightly more complicated version of the argument in Lemma 13.7 gives the statement:

Lemma 13.11. *Let p, q, r be distinct points in $\partial_s X$, and let $K \geq 7\delta$. The set of K–centers of p, q, r is nonempty, and has diameter bounded in terms of δ and K.*

2. Isometries of hyperbolic spaces

There are two main results in this section. Theorem 13.24 classifies isometries of hyperbolic spaces into elliptic, parabolic, and loxodromic (see Definitions 12.3, 13.20). For this result we *don't* assume that the space in question is proper. In particular, when we talk about ∂X in this section, we always mean the *sequential boundary* defined in the last section.

Theorem 13.25 says that any infinite subgroup of a hyperbolic group contains a loxodromic element. Before getting to the proofs, we introduce some machinery which makes them a bit easier to articulate.

2.1. Quasimorphisms. For much more on quasimorphisms, including proofs of the below statements, see [**Cal09**, 2.2]. A *quasimorphism* on a group G is a function $\phi\colon G \to \mathbb{R}$ which is "almost a homomorphism" in the sense that, for some $D(\phi) \geq 0$ (the *defect*), and all $g, h \in G$

$$|\phi(gh) - \phi(g) - \phi(h)| \leq D(\phi).$$

A quasimorphism is called *homogeneous* if it is a homomorphism when restricted to any cyclic subgroup, in other words $\phi(g^n) = n\phi(g)$ for all $g \in G, n \in \mathbb{Z}$. Any quasimorphism can be "homogenized"; given a quasimorphism ϕ, the function

$$\bar{\phi}(g) = \lim_{n\to\infty} \frac{\phi(g^n)}{n}$$

is a homogeneous quasimorphism satisfying

$$D(\bar{\phi}) \leq 2D(\phi), \text{ and } |\bar{\phi}(g) - \phi(g)| \leq D(\phi),\ \forall g \in G. \tag{13}$$

(See [**Cal09**, Lemma 2.21 and Corollary 2.59].) Notice that if two quasimorphisms differ by a bounded amount, their homogenizations are the same.

At first glance you may wonder whether all quasimorphisms are just bounded perturbations of homomorphisms to $\mathbb{R}$. Here's an example of something different.

Example 13.12. Let G be the free group on $\{a, b\}$, and embed the Cayley graph of G in the plane in the standard way. For $g \in G$, let σ_g be the unique embedded edge-path from the identity to g. Let $\phi(g)$ be the number of left turns this edge-path takes minus the number of right turns it takes. Then $\phi((ab)^n) = 2n - 1$ for

$n > 0$, so ϕ is unbounded. On the other hand $\phi(a^n) = \phi(b^n) = 0$ for all n, so ϕ is not boundedly different from a homomorphism.

EXERCISE 22. Show that ϕ defined in the above example is a quasimorphism. What is its defect?

2.2. A little nonstandard analysis. The book [**Gol98**] is a nice readable introduction to the subject of nonstandard analysis. For the purposes of these notes, we will only really need a few facts. One can think of nonstandard analysis as a set-theoretic trick which avoids the cumbersome nature of arguments involving passing repeatedly to subsequences. We will define a gadget $\lim_\omega$ which can be applied to any bounded sequence and always consistently pick out an accumulation point of the sequence. How does this work? Here are the relevant definitions. (Despite the eerie similarity in terminology, try not to confuse this notion of ultrafilter with the one defined for pocsets!)

DEFINITION 13.13. An *ultrafilter* on $\mathbb{N}$ is a subset ω of $2^{\mathbb{N}} \setminus \{\emptyset\}$ which satisfies:

1. (Consistency I) If $A, B \in \omega$, then $A \cap B \in \omega$.
2. (Consistency II) If $A \subseteq B$ and $A \in \omega$, then $B \in \omega$.
3. (Completeness) For any $A \subseteq \mathbb{N}$, exactly one of A, A^c is in ω.

A *principal* ultrafilter on $\mathbb{N}$ is one of the form $\omega_n = \{A \subseteq \mathbb{N} \mid n \in A\}$.

EXERCISE 23. Suppose ω is an ultrafilter on $\mathbb{N}$. Let $\mathbb{N} = A_1 \sqcup \cdots \sqcup A_n$ be a finite partition of $\mathbb{N}$. Show $A_i \in \omega$ for exactly one $i \in \{1, \ldots, n\}$.

EXERCISE 24. Say a subset of $2^{\mathbb{N}} \setminus \{\emptyset\}$ is a *filter* on $\mathbb{N}$ if it satisfies the Consistency requirements 13.13.(1) and 13.13.(2).

1. Show that the "cofinite sets" $F = \{A \subseteq \mathbb{N} \mid \#(\mathbb{N} \setminus A) < \infty\}$ form a filter. Then use Zorn's Lemma to show that there is an ultrafilter on $\mathbb{N}$ which isn't principal.
2. Conversely, show that every nonprincipal ultrafilter contains F.

This gadget (a nonprincipal ultrafilter on $\mathbb{N}$) exists, but can't be described directly; the use of Zorn's Lemma is essential. There is something a little unsettling about this, but we press on anyway, **fixing a nonprincipal ultrafilter ω for the rest of the text.**

DEFINITION 13.14. Let $\mathbf{x} = \{x_i\}_{i \in \mathbb{N}}$ be a sequence of points in a metric space M. Say that $\lim_\omega x_i = x$ if for every $\epsilon > 0$ the set $\{i \mid d(x_i, x) \leq \epsilon\}$ is in ω. We say the point x is the *ω–ultralimit* of the sequence $\mathbf{x}$.

LEMMA 13.15. *If M is a compact metric space, then every infinite sequence has a unique ω–ultralimit among its accumulation points.*

PROOF. (Sketch) This is an application of Exercise 23 above. Since M is a compact metric space, for any $N > 0$, it can be partitioned into finitely many sets $M_{N,1}, \ldots, M_{N,k_N}$ of diameter $\frac{1}{N}$. One can also arrange that each set $M_{N,j}$ is contained in one of the sets $M_{N-1,j'}$.

Fix a sequence $\{x_i\}_{i \in \mathbb{N}}$. The sets $A_{N,j} = \{i \mid x_i \in M_{N,j}\}$ give a finite partition of $\mathbb{N}$, exactly one of whose elements $A_{N,j(N)}$ is in ω. The sets $M_{N,j(N)}$ form a nested sequence of subsets of M whose diameter is going to zero. Each $M_{N,j(N)}$ contains an infinite subsequence of $\{x_i\}_{i \in \mathbb{N}}$, and the unique point in $\bigcap_N M_{N,j(N)}$ is the ultralimit $\lim_\omega x_i$. □

Here is the construction which will be used in the next section.

EXERCISE 25. Let X be a metric space, and suppose $\{f_i\}_{i\in\mathbb{N}}$ is a sequence of K–lipschitz functions which are uniformly bounded on any bounded set. Then $f_\omega(x) = \lim_\omega f_i(x)$ is a K–lipschitz function on X.

2.3. A quasimorphism on the stabilizer of a point at infinity. In this subsection, we fix a δ–hyperbolic space X, and a point $\xi \in \partial X$, and a group G acting by isometries of X. We're not assuming our hyperbolic spaces are proper, so the point ξ should be thought of as an equivalence class of *sequences* not rays, as in Section 1.2.

We are going to define a quasimorphism β_ξ on $G_\xi = \mathrm{Stab}_G(\xi)$ which measures the extent to which a group element pushes elements "towards" or "away from" ξ. (We follow [**CCMT15**] for this definition.) Let $\mathbf{x} = \{x_i\}_{i\in\mathbb{N}}$ be a sequence representing ξ, and define the *Busemann horokernel* on $X \times X$ to be

$$h_{\mathbf{x}}(x, y) = \lim_\omega d(x, x_n) - d(y, x_n).$$

The functions $h^i(x, y) = d(x, x_i) - d(y, x_i)$ are uniformly lipschitz and uniformly bounded on compact sets, so Exercise 25 implies that $h_{\mathbf{x}}$ is a lipschitz function on $X \times X$.

EXERCISE 26. Let $x \in X$, and suppose the sequence $\mathbf{x}$ tends to $\xi \in \partial X$.

(1) Suppose that γ is a geodesic ray in X tending to ξ. Show

$$\lim_{t\to\infty} h_{\mathbf{x}}(x, \gamma(t)) = \infty.$$

(2) Suppose that $\{y_i\}_{i\in\mathbb{N}}$ limits to some point in $\partial X \setminus \{\xi\}$. Show

$$\lim_{i\to\infty} h_{\mathbf{x}}(x, y_i) = -\infty.$$

LEMMA 13.16. *Let $\mathbf{x} = \{x_i\}_{i\in\mathbb{N}}$ and $\mathbf{y} = \{y_i\}_{i\in\mathbb{N}}$ both represent ξ. Then the difference $|h_{\mathbf{x}} - h_{\mathbf{y}}|$ is bounded by 2δ.*

PROOF. Let $x, y \in X$.For $\epsilon > 0$, consider the sets of indices

$$A = \{i \mid |h_{\mathbf{x}}(x, y) - (d(x, x_n) - d(y, x_n))| \le \epsilon\},$$
$$B = \{i \mid |h_{\mathbf{y}}(x, y) - (d(x, y_n) - d(y, y_n))| \le \epsilon\}.$$

Both sets are elements of ω, so their intersection $A \cap B$ is also in ω. In particular $A \cap B$ is infinite.

We have $\lim_{i,j\to\infty}(x_i \,|\, y_j)_x = \infty$, since the sequences $\mathbf{x}$ and $\mathbf{y}$ represent the same point at infinity. The quantities $(y \,|\, x_n)_x$ or $(y \,|\, y_n)_x$ are bounded by $d(x, y)$, and for large n we have $(x_n \,|\, y_n)_x > d(x, y)$. Fix some such large n contained in $A \cap B$. The Gromov product inequality (7) implies that $|(y \,|\, x_n)_x - (y \,|\, y_n)_x| \le \delta$. A computation shows

$$\begin{aligned} |h_{\mathbf{x}}(x, y) - h_{\mathbf{y}}(x, y)| &\le |d(x, x_n) - d(y, x_n) - d(x, y_n) + d(y, y_n)| + 2\epsilon \\ &= 2|(y \,|\, x_n)_x - (y \,|\, y_n)_x| + 2\epsilon \\ &\le 2\delta + 2\epsilon. \end{aligned}$$

Since ϵ, x, y were arbitrary, we see that $h_{\mathbf{x}}$ and $h_{\mathbf{y}}$ differ by at most 2δ. □

Now we use the Busemann kernel to define a quasimorphism.

DEFINITION 13.17. Choose a basepoint $p \in X$. Let $\alpha_{\mathbf{x},p}(g) = h_{\mathbf{x}}(p, gp)$.

LEMMA 13.18. *The function* $\alpha_{\mathbf{x},p}\colon G_\xi \to \mathbb{R}$ *is a quasimorphism of defect at most* 2δ. *Its homogenization* β_ξ *is independent of the choice of sequence* $\mathbf{x}$ *and basepoint* p.

PROOF. We first show that $\alpha_{\mathbf{x},p}$ is a quasimorphism. Let $g_1, g_2 \in G_\xi$, and let $D(g_2, g_2) = \alpha_{\mathbf{x},p}(g_1g_2) - \alpha_{\mathbf{x},p}(g_1) - \alpha_{\mathbf{x},p}(g_2)$. Then using the fact that ultralimits commute with addition and that g_1 is an isometry, we can write

$$\begin{aligned} D(g_1, g_2) &= \lim_\omega [d(p, x_n) - d(g_1g_2p, x_n) - d(p, x_n) + d(g_1p, x_n) - d(p, x_n) + d(g_2p, x_n)] \\ &= \lim_\omega [d(g_1p, x_n) - d(g_1g_2p, x_n) - (d(p, x_n) - d(g_2p, x_n))] \\ &= \lim_\omega [d(p, g_1^{-1}x_n) - d(g_2p, g_1^{-1}x_n)] - \lim_\omega [d(p, x_n) - d(g_2p, x_n)] \\ &= h_{g_1^{-1}\mathbf{x}}(p, g_2p) - h_{\mathbf{x}}(p, g_2p). \end{aligned}$$

Since g_1 fixes ξ, the sequence $g_1^{-1}\mathbf{x}$ also represents ξ. Lemma 13.16 says the difference is at most 2δ.

It is straightforward to see from the definitions and from Lemma 13.16 that changing the basepoint p or the sequence $\mathbf{x}$ only changes $\alpha_{\mathbf{x},p}$ by a bounded amount, so it doesn't affect the homogenization β. □

We will connect this quasimorphism to the detection of loxodromic isometries in the next subsection.

2.4. Classifying isometries of hyperbolic spaces. We have already defined an isometry g of a hyperbolic space X to be *loxodromic* if, for some $p \in X$, the map $n \mapsto g^n p$ gives a quasi-isometric embedding of $\mathbb{Z}$ into X. Note that for a loxodromic isometry the limit set $\Lambda(\langle g \rangle)$ contains two points: $g^{+\infty} = [\{g^i p\}_{i\in\mathbb{N}}]$ and $g^{-\infty} = [\{g^{-i} p\}_{i\in\mathbb{N}}]$. Both of these points are fixed by g. More generally we have the following.

LEMMA 13.19. *Let* g *be an isometry of a hyperbolic space* X. *Then* g *fixes every point in* $\Lambda(\langle g \rangle)$.

PROOF. Fix a base point $p \in X$.

If ξ is in $\Lambda(\langle g \rangle)$, there is some sequence $\{k_i\}$ with $\{g^{k_i} p\}$ tending to ξ. The sequence $\{g^{k_i+1} p\}$ therefore tends to $g\xi$. But $d(g^{k_i+1}p, g^{k_i}p) = d(gp, p)$ is constant, and so the two sequences must tend to the same point at infinity. Thus $g\xi = \xi$. □

Here are two other types of isometries.

DEFINITION 13.20. Let X be hyperbolic. An isometry $g\colon X \to X$ is *elliptic* if some orbit $\langle g \rangle p$ is bounded.

It is *parabolic* if it is not elliptic, and $\#\Lambda(\langle g \rangle) = 1$.

Notice that g is loxodromic/parabolic/elliptic if and only if all its nonzero powers are.

We want to classify isometries with unbounded orbits into loxodromic and parabolic. We will do so according to their fixed point sets in ∂X. Using Lemma 13.11 it is not hard to see the following.

LEMMA 13.21. *Suppose* g *is an isometry of a hyperbolic space* X *which fixes three points in* ∂X. *Then* g *is elliptic.*

We will see that for a non-elliptic isometry g, the fixed point set at infinity is equal to the limit set of the cyclic group generated by g. First we show this limit set is always nonempty. (When X is proper this is automatic, since ∂X compactifies X in that case.)

LEMMA 13.22. *Let g be a non-elliptic isometry of a hyperbolic space X. Then $\Lambda(\langle g\rangle)$ is non-empty.*

PROOF. We argue by contradiction, assuming $\Lambda(\langle g\rangle)$ is empty. If g were loxodromic $\Lambda(\langle g\rangle)$ would contain the points $g^{\pm\infty}$ defined above, so g is not loxodromic.

Fix $p \in X$. Since the limit set is empty, the sequence $\{g^{k_i}p\}_{i\in\mathbb{N}}$ fails to converge to infinity, no matter what indices k_i are chosen. Since g is non-elliptic, we may choose such a sequence so that the distances $d(g^{k_i}p, p)$ are monotone increasing. Since the sequence fails to converge to infinity, there is a constant C so there are arbitrarily large pairs k_i, k_j so that $(g^{k_i}p \,|\, g^{k_j}p)_p \le C$. In particular, we can fix a pair k_i, k_j so that $d(g^{k_i}p, p), d(g^{k_i}p, p) > C + 6\delta$. Applying Lemma 12.5 we see that $g^{k_i - k_j}$ is loxodromic, contradicting the assumption that g was non-loxodromic. □

LEMMA 13.23. *Let g be an isometry of the hyperbolic space X, fixing $\xi \in \partial X$. Then g is loxodromic if and only if $\beta_\xi(g) \neq 0$.*

PROOF. Suppose first that g is loxodromic. The limit set of $\langle g\rangle$ contains the points $g^{\pm\infty}$. The isometry g fixes these points (Lemma 13.19), and isn't elliptic, so it can fix no others (Lemma 13.21). So, possibly replacing g by its inverse, we can suppose that $\xi = g^\infty$. In particular, fixing some $x \in X$, we can use the sequence $\mathbf{x} = \{g^i x\}_{i\in\mathbb{N}}$ to define a horokernel $h_{\mathbf{x}}$ and a quasimorphism $\alpha_{\mathbf{x},x}(k) = h_{\mathbf{x}}(x, kx)$ whose homogenization is β_ξ, as in Subsection 2.3. Exercise 26 implies that $\lim_{i\to\infty} \alpha_{\mathbf{x},x}(g^{-i}) = -\infty$. Since the difference between β_ξ and $\alpha_{\mathbf{x},x}$ is bounded β_ξ cannot vanish on g.

Conversely, suppose that $\beta_\xi(g) \neq 0$. We may suppose $\beta_\xi(g) > 0$. Let $\mathbf{x} = \{x_i\}_{i\in\mathbb{N}}$ be a sequence representing ξ, and let $x \in X$ be some base point. There is some constant C so that $|h_{\mathbf{x}}(x, g^n x) - \beta_\xi(g^n)| \le C$, independent of n. Fix $n \in \mathbb{Z}$ and $\epsilon > 0$ and choose some large k so that, for $a, b \in \{g^n x, x\}$,

$$|h_{\mathbf{x}}(a, b) - (d(a, x_k) - d(b, x_k))| < \epsilon.$$

Now we have

$$\begin{aligned} d(g^n x, x) &\ge d(x, x_k) - d(g^n x, x_k) \\ &\ge h_{\mathbf{x}}(x, g^n x) - \epsilon \\ &\ge \beta_\xi(g^n) - C - \epsilon \\ &\ge n\beta_\xi(g) - C - \epsilon. \end{aligned}$$

Since $d(g^p x, g^q x) = d(g^{p-q}x, x)$, this shows that $n \mapsto g^n x$ is a quasi-isometric embedding, and so g is loxodromic. □

THEOREM 13.24. *Every isometry of a hyperbolic space is either elliptic, parabolic, or loxodromic.*

PROOF. Let g be an isometry of a hyperbolic space X, which we assume is not elliptic, and fix a base point $x \in X$. Since it is not elliptic, Lemma 13.22 implies that $\langle g\rangle$ has non-empty limit set. Let ξ be a point of this limit set. We claim that either g is loxodromic or ξ is the only point. If there is another point η, there is a sequence of powers k_i so that $g^{k_i}x$ tends to η. Exercise 26 implies that $\lim_{i\to\infty} h_{\mathbf{x}}(x, g^{x_i}x) = -\infty$. In particular, it must be the case that $\beta_\xi(g) \neq 0$, so Lemma 13.23 implies that g is loxodromic.

Otherwise, ξ is the only limit point, and g is parabolic. □

2.5. Torsion subgroups of hyperbolic groups are finite. We saw before that infinite hyperbolic groups always contain infinite order elements. The next result says this is true even for infinite *subgroups* of hyperbolic groups. We use an argument adapted from [**GdlH90**, pp. 156–157].

THEOREM 13.25. *Let G be hyperbolic, and let $H < G$ be infinite. Then H contains a loxodromic element.*

PROOF. We fix a δ–hyperbolic Cayley graph X for G. Since H is infinite, the limit set of H in ∂X is nonempty. If it has two points a, b, choose sequences $\alpha_i \to a, \beta_i \to b$ in H. The Gromov products $(\alpha_i \mid \beta_j)_1$ are bounded, so there is some pair α, β with $\min\{|\alpha|, |\beta|\} \geq 2(\alpha \mid \beta)_1 + 6\delta$. By Lemma 12.5, one of α, β, or $\alpha\beta$ is loxodromic.

Now we suppose for a contradiction that H has precisely one limit point, a, so that $H < \mathrm{Stab}(a)$. Let $\beta_a\colon \mathrm{Stab}(a) \to \mathbb{R}$ be the quasimorphism from Lemma 13.18. We will bound the cardinality of the set $K = \{g \in \mathrm{Stab}(a) \mid \beta_a(g) = 0\}$. Since H is infinite there is some $h \in H$ with $\beta_a(h) \neq 0$. By Lemma 13.23, this h is loxodromic. In particular H has at least two limit points, a contradiction.

It remains to bound the size of K. Let $\gamma\colon [0, \infty) \to X$ be a unit speed geodesic ray starting at the identity and representing a. Let $\mathbf{x}$ be the sequence $\{\gamma(i)\}$, let $h_{\mathbf{x}}$ be the corresponding Busemann horokernel, and let $\alpha = \alpha_{\mathbf{x},1}$ be the corresponding quasimorphism. Since β_a is the homogenization of α, which has defect at most 2δ, we have $|\beta_a(g) - \alpha(g)| \leq 2\delta$ for all g. In particular, $|\alpha(g)| \leq 2\delta$ for every $g \in K$.

Let $K_0 < K$ be any finite subset. For each $g \in K_0$, there is some $N_0 = N_0(g)$ so that $g\gamma|_{[N,\infty)}$ lies in a 2δ–neighborhood of γ and $\gamma|_{[N,\infty)}$ lies in a 2δ–neighborhood of $g\gamma$. Choose an integer N so that $N > N_0(g)$ for every $g \in K_0$. For such N, the quantity $h_{\mathbf{x}}(1, g)$ differs by at most δ from $d(1, \gamma(N)) - d(g, \gamma(N))$, and the quantity $h_{g\mathbf{x}}(1, g)$ differs by at most δ from $d(1, g\gamma(N)) - d(g, g\gamma(n))$. By Lemma 13.16 the quantities $h_{\mathbf{x}}(1, g)$ and $h_{g\mathbf{x}}(1, g)$ differ by at most 2δ from each other.

For $g \in K_0$, we define $\eta(g)$ to be any integer so that $d(\gamma(N + \eta(g)), g\gamma(N)) \leq 2\delta + 1$; we claim that $\alpha(g)$ is approximately $\eta(g)$. Indeed

$$\begin{aligned}\alpha(g) = h_{\mathbf{x}}(1, g) &\geq d(1, \gamma(N + \eta(g))) - d(g, \gamma(N + \eta(g))) - \delta \\ &\geq N + \eta(g) - (N + 2\delta + 1) - \delta \\ &= \eta(g) - (3\delta + 1),\end{aligned}$$

and

$$\begin{aligned}\alpha(g) &\leq h_{g\mathbf{x}}(1, g) + 2\delta \\ &\leq d(1, g\gamma(N)) - d(g, g\gamma(N)) + 3\delta \\ &\leq N + \eta(g) + 2\delta + 1 - N + 3\delta \\ &= \eta(g) + 5\delta + 1.\end{aligned}$$

Since $|\alpha(g)|$ is bounded above by 2δ, we have $|\eta(g)| \leq 7\delta + 1$ on K_0. It follows that $d(\gamma(N), g\gamma(N)) \leq 9\delta + 2$, for every $g \in K_0$, and so the cardinality of K_0 is at most the cardinality of a $(9\delta + 2)$–ball in the Cayley graph of G. Since K_0 was an arbitrary finite subset of K, this shows that K is finite, and that $H \backslash K$ is nonempty, as desired. □

Corollary 13.26. *If G is hyperbolic, it contains no parabolic element.*

3. Quasiconvex subgroups of hyperbolic groups

In this section we develop some more nice properties of quasiconvex subgroups of hyperbolic groups.

Lemma 13.27. *If G is a hyperbolic group, and $H < G$ is quasiconvex, then H is finitely generated and quasi-isometrically embedded in G. In particular H is hyperbolic.*

Proof. Fix Γ a Cayley graph for G, and suppose that Γ is δ–hyperbolic, and $H \subset \Gamma$ is K–quasiconvex. The reader can verify that a closed $(K + 10\delta)$–neighborhood N of H in Γ is quasi-isometrically embedded. Corollary 10.6 implies that N is Gromov hyperbolic. Schwarz–Milnor implies H is finitely generated and $H \hookrightarrow N$ is a quasi-isometric embedding. Thus $H \hookrightarrow G$ is a composition of quasi-isometric embeddings. □

Lemma 13.28. *Any finite intersection of quasiconvex subgroups of a hyperbolic group is quasiconvex.*

Proof. (We follow the proof of [**GMRS98**, Lemma 2.7] cf. [**Sho91**, Proposition 3].) It suffices to consider two quasiconvex subgroups A, B of a hyperbolic group G. We argue by contradiction, supposing that $C = A \cap B$ is *not* quasiconvex. Thus there is a sequence of elements $c_i \in C$, and geodesics γ_i joining 1 to c_i, which contain points y_i satisfying $d(y_i, C) \to \infty$.

However there is a fixed quasiconvexity constant λ and elements $a_i \in A$, $b_i \in B$, so that $d(y_i, a_i)$ and $d(y_i, b_i)$ are both bounded by λ, for all i. The distances $d(a_i, b_i)$ are all bounded by 2λ, and there are only finitely many elements in the ball of radius 2λ around 1. Therefore, we may pass to a subsequence in which $a_i = b_i g$ for some fixed g of length $\leq 2\lambda$. Now note that

$$a_i a_1^{-1} = b_i g \cdot g^{-1} b_1^{-1} = b_i b_1^{-1}$$

is in C for all i. Thus $d(y_i, C) \leq d(y_i, a_i) + d(a_i, a_i a_1^{-1}) \leq \lambda + d(1, a_1^{-1})$ is bounded over all i, a contradiction since the y_i are supposed to be getting further and further from C. □

4. Height and width of quasiconvex subgroups

A great many arguments involving quasiconvex subgroups are easier in case the subgroup H involved is *almost malnormal*, meaning that $H \cap gHg^{-1}$ is finite whenever $g \notin H$.[2] The height and width of a subgroup are different measurements of how far a subgroup is from being almost malnormal.

Definition 13.29. Let G be a group, $H < G$ a subgroup. The *height* of H in G is the largest number n so that there are distinct cosets $g_1 H, \ldots, g_n H$ with $\#(g_1 H g_1^{-1} \cap \cdots \cap g_n H g_n^{-1}) = \infty$.

The *width* of H is the largest number n so that there are distinct cosets $g_1 H, \ldots, g_n H$ with $\#(g_i H g_i^{-1} \cap g_j H g_j^{-1}) = \infty$ for all i, j.

[2] H is *malnormal* if $H \cap gHg^{-1} = \{1\}$ whenever $g \notin H$.

If $H < G$ is finite its height and width are 0. If H is infinite and almost malnormal, its height and width are 1.

There's no real reason to restrict to a single subgroup; the above definitions generalize to a collection of subgroups $\mathcal{H}$. Here's another way to think about height and width, in terms of a certain simplicial complex $\mathcal{D} = \mathcal{D}(\mathcal{H})$. Define the zero-skeleton as the disjoint union of coset spaces

$$\mathcal{D}^{(0)} = \sqcup\{G/H \mid H \in \mathcal{H}, \#H = \infty\}.$$

Vertices $g_0H_0, \ldots, g_nH_n$ span a simplex in $\mathcal{D}$ if

$$\#\left(g_0Hg_0^{-1} \cap \cdots g_nHg_n^{-1}\right) = \infty.$$

Let $\mathcal{F}$ be the flag complex with the same one-skeleton as $\mathcal{D}$. The *height* of $\mathcal{H}$ is $\dim(\mathcal{D}) + 1$, and the *width* is $\dim(\mathcal{F}) + 1$. (The empty set has dimension -1.)

In this section we'll prove that quasiconvex subgroups of hyperbolic groups have finite height and width. Essentially the same proofs apply to finite collections of quasiconvex subgroups. These theorems appeared first in [**GMRS98**].

Here's a lemma:

LEMMA 13.30. *Let Q be a λ–quasiconvex subset of the δ–hyperbolic space X, and let γ be a bi-infinite geodesic with endpoints in the limit set of Q. Then $\gamma \subseteq N_{\lambda+2\delta}(Q)$.*

PROOF. Let $p \in \gamma$. Reparametrize γ so that $\gamma(0) = p$. Let $\{a_i\}_{i\in\mathbb{N}}$ and $\{b_i\}_{i\in\mathbb{N}}$ be representative sequences of points in Q which limit to the two endpoints of γ at infinity. Choose N so that $(a_N \,|\, \gamma(-N))_p > R$, and $(b_N \,|\, \gamma(N))_p > R$. Draw a picture and convince yourself that p is within 2δ of a geodesic from a_N to b_N. Since both a_N and b_N are in Q, which is λ–quasiconvex, p is within $2\delta + \lambda$ of some point in Q. □

We use the above lemma to prove quasiconvex subgroups have finite height:

PROPOSITION 13.31. *Let G be hyperbolic, and H quasiconvex. Then the height of H in G is finite.*

PROOF. Fix a generating set S for G, and let δ be a constant of hyperbolicity for $\Gamma = \Gamma(G, S)$. Let λ be the quasiconvexity constant for $H \subset \Gamma$. Note that every coset gH is also a λ–quasiconvex set. Suppose that $I = g_1Hg_1^{-1} \cap \cdots g_nHg_n^{-1}$ is infinite. By Lemma 13.28, I is also a quasiconvex subgroup. In particular it is infinite hyperbolic (Lemma 13.27), so it contains an element of infinite order (Corollary 10.15). Let γ be a bi-infinite geodesic joining distinct points in the limit set of I. The endpoints of γ are also in $\Lambda(g_iH) = \Lambda(g_iHg_i^{-1})$ for each i. Since each g_iH is λ–quasiconvex, Lemma 13.30 implies that $\gamma \subset \bigcap_{i=1}^n N_{2\delta+\lambda}(g_iH)$. In particular, for any vertex g on γ, every coset g_iH intersects the $2\delta+\lambda$–neighborhood about g. These cosets are all disjoint, so there are at most $\#(B_{2\delta+\lambda}(g))$ of them. But this ball is isomorphic to the same radius ball around the identity, so we get that the height of H is at most $\#(B_{2\delta+\lambda}(1)) \leq (2\#S)^{2\delta+\lambda+1}$. □

4.1. Bounded packing and finite width. Next we turn to finite width. Finiteness of width for quasi-convex subgroups will be a consequence of a somewhat more useful notion, bounded packing. This notion was first introduced by Hruska and Wise in [**HW09**] and then applied by them to the study of cube complexes in [**HW14**].

Definition 13.32. Let G be finitely generated, and let $H < G$. Let S be a finite generating set for G, with respect to which we measure distance. The subgroup H has *bounded packing* if for any $D > 0$, there exists $N \geq 2$ so that among any tuple of distinct cosets $\{g_1H, \ldots, g_NH\}$, there must be two of distance at least D.

Exercise 27. Show that if H has bounded packing with respect to the generating set S, then it has bounded packing with respect to any other finite generating set for G.

Exercise 28. Finite subgroups and normal subgroups always have bounded packing.

Proposition 13.34 below shows that bounded packing of a codimension one subgroup forces finite dimensionality of the dual cube complex (via the Sageev construction). Part of the argument will be used again later, so we break it out as a lemma:

Lemma 13.33. *Let G be a finitely generated group with Cayley graph X. For $i \in \{1,2\}$, let H_i be a finitely generated codimension one subgroup, with associated H_i–wall W_i. Then there is a $D > 0$ so that whenever g_1W_1 and g_2W_2 are transverse, then the D–neighborhoods of H_1 and H_2 in X intersect.*

Proof. Each of the walls W_i is really a pair of H_i–halfspaces $\{A_i, A_i^c\}$, and there are finitely many H_i–orbits of edges joining A_i to A_i^c. Thus for some $D > 0$, no edge in the complement of $N_i = N_D(H_i)$ connects A_i to A_i^c. Since each H_i is finitely generated, we can suppose D is large enough that the N_i are connected.

Now suppose that $g_1N_1 \cap g_2N_2$ is empty. We want to show that g_1W_1 and g_2W_2 are nested. It suffices to take $g_1 = 1, g_2 = g$. By exchanging A_1 with A_1^c, we may suppose that $gN_2^{(0)} \subset A_1$. If W_1 and gW_2 are transverse, then in particular the sets $gA_2^c \cap A_1^c$ and $gA_2 \cap A_1^c$ must be nonempty. So let $x \in gA_2^c \cap A_1^c$, and $y \in gA_2 \cap A_1^c$. Now every edge connecting a point of A_1^c to a point of A_1 is contained in N_1, and N_1 is connected, so there is a path σ connecting x to y using only vertices in A_1^c or N_1. The path σ connects a vertex of gA_2^c to one of gA_2, so it must contain a vertex z of gN_2. Since $gN_2^{(0)} \subset A_1$, this vertex z must like in N_1. So $z \in gN_2 \cap N_1$, contradicting the assumption that it is empty. □

Proposition 13.34. *Let G be finitely generated. Let $H < G$ be a finitely generated codimension one subgroup which has bounded packing, and suppose that X is a cube complex obtained from $(G, \{H\})$ via the Sageev construction. Then C is finite dimensional.*

Proof. Fix a finite generating set S for G, and let X be the Cayley graph of G with respect to S. Let A be the H–halfspace on which the Sageev construction is based (See Lemma 9.19 for its properties), and let W be the corresponding H–wall. By Lemma 13.33, there is some D so that if $N = N_D(H)$, and $g_1N \cap g_2N = \emptyset$, then the walls g_1W, g_2W are nested.

The subgroup H is assumed to have bounded packing, so let M be some number so that if $\{g_1H, \ldots, g_MH\}$ are distinct cosets, then two must be at distance at least $2D+1$. Thus there are at most $M-1$ pairwise transverse walls, and the dimension of the dual cube complex is at most $M-1$. □

THEOREM 13.35. *Let G be hyperbolic and $H < G$ quasiconvex. Then H has bounded packing.*

Before we prove the theorem, let's look at a couple of consequences. The first follows directly from Proposition 13.34 above.

COROLLARY 13.36. *Let G be hyperbolic and let $H < G$ be quasiconvex and codimension one. If A is an H–almost invariant subset, and C the corresponding cube complex, then C is finite dimensional.*

The second requires a little geometry.

COROLLARY 13.37. *Let G be hyperbolic and $H < G$ quasiconvex. Then H has finite width.*

PROOF. Let Γ be a δ–hyperbolic Cayley graph for G, and suppose H is λ–quasiconvex in Γ. Let $D = \lambda + 2\delta$. By Theorem 13.35, H has bounded packing in G. Let $N = N(D)$ be the constant from the definition of bounded packing.

Suppose that $g_1H, \ldots, g_kH$ are distinct cosets so that all intersections $g_iHg_i^{-1} \cap g_jHg_j^{-1}$ are infinite. Then we may argue exactly as in the proof of Proposition 13.31 that $d(g_iH, g_jH) \leq \lambda + 2\delta = D$ for each i, j. By the definition of bounded packing, k (and hence the width of H) is at most N. □

The proof of Theorem 13.35 follows Hruska and Wise's proof in [**HW09**] and is likewise based on the following "magic trick" [**HW09**, Lemma 4.5]

LEMMA 13.38. *Let G be a discrete group with a proper left-invariant metric d_G. Let $A, B < G$, and let xA, yB be cosets, and let $L > 0$. Then there is an L' so that*

$$N_L(xA) \cap N_L(yB) \subset N_{L'}\left(xAx^{-1} \cap yBy^{-1}\right).$$

PROOF. Suppose the lemma is false. Then there is a sequence of group elements $z_i \in N_L(xA) \cap N_L(yB)$ so that $d_G(z_i, xAx^{-1} \cap yBy^{-1})$ tends to infinity.

For each i, there are elements p_i, q_i in the L–ball about 1 and $a_i \in A$, $b_i \in B$ so that $z_i = xa_ip_i = yb_iq_i$. The metric on G is proper, so by passing to a subsequence we may assume that p_i and q_i are constant, i.e. $p_i = p$ and $q_i = q$ for all i.

But then $z_iz_1^{-1} = xa_ia_1^{-1}x^{-1} = yb_ib_1^{-1}y^{-1}$, so $d_G(z_i, xAx^{-1} \cap yBy^{-1})$ is bounded above by $d_G(1, z_1)$ for all i. This contradicts our initial choice of sequence $\{z_i\}$. □

Why do I call this a magic trick? As the reader can see, there are essentially no hypotheses, but the conclusion is weaker than it might at first appear. The constant L' depends on all the data given.

PROOF OF THEOREM 13.35. Let H, G constitute a counterexample, chosen to minimize $\text{height}_G(H)$ among all possible counterexamples. Since finite subgroups have bounded packing, this height must be positive. We will show that we can find another counterexample with smaller height, contradicting our initial choice.

We fix a word metric d_G on G. Since H does not have bounded packing in G, there is some L so that for every N there is a collection

$$\mathcal{H}_N = \{H, g_{N,1}H, \ldots, g_{N,N}H\}$$

so that $d_G(A, A') \leq L$ for each $A, A' \in \mathcal{H}_N$. Note that if $d_G(gH, H) \leq L$, then HgH meets $B_L(1)$ in G, and there are finitely many such double cosets.

Refining our sequence of collections $\mathcal{H}_N$, we may therefore assume that they all consist of H together with a collection of cosets of H lying in some *fixed* double coset $HgH \neq H$.

Let $K = H \cap gHg^{-1}$. Since K is an intersection of quasiconvex subgroups it is quasiconvex as well, and moreover is quasiconvex in the hyperbolic group H. We will show first that K has lower height in H than H has in G, and then that K also fails to have bounded packing in H, completing the contradiction.

CLAIM 13.38.1. *The height of K in H is strictly less than the height of H in G.*

PROOF OF CLAIM 13.38.1. Whenever $h_1K, \ldots, h_pK$ are distinct in H/K then the reader may verify that $H, h_1gH, \ldots, h_pgH$ are distinct in G/H, and also

$$\bigcap_{i=1}^{p} h_i K h_i^{-1} \subset H \cap \bigcap_{i=1}^{p} h_i g H (h_i g)^{-1}.$$

If the left-hand side is infinite, so is the right, and so the height of H in G is at least one more than the height of K in H. ☐

In order to show K fails to have bounded packing we first show:

CLAIM 13.38.2. *There is an $R > 0$ so that if*

$$\max\{d_G(agH, bgH), d_G(H, agH), d_G(H, bgH)\} \leq D,$$

then $d_G(aK, bK) \leq R$.

PROOF OF CLAIM 13.38.2. Let x be a point of distance at most $D/2$ from both agH and H, let y be within $D/2$ of both bgH and H, and let z be within $D/2$ of both agH and bgH. Choosing geodesics $[x, y]$, $[y, z]$, and $[z, x]$, there is a number D' (depending on δ, D, and the quasi-convexity constant of H) so that

$$[x, y] \subset N_{D'}(H), [y, z] \subset N_{D'}(bgH), [z, x] \subset N_{D'}(agH).$$

There is a point w which is at distance at most δ from all three sides of this triangle, and thus for $D_1 = D' + \delta$,

$$N_{D_1}(agH) \cap N_{D_1}(bgH) \cap N_{D_1}(H) \neq \emptyset.$$

The magic trick Lemma 13.38 gives us D_2 so that $N_{D_1}(H) \cap N_{D_1}(gH) \subset N_{D_2}(K)$. Translating by a and b gives:

$$N_{D_1}(H) \cap N_{D_1}(agH) \subset N_{D_2}(aK)$$
$$N_{D_1}(H) \cap N_{D_1}(bgH) \subset N_{D_2}(bK)$$

The point w is in both of these, so we have $d(aK, bK) \leq 2D_2$, and we can take $R = 2D_2$. ☐

The following uses no geometry, and we leave it as an exercise:

CLAIM 13.38.3. *The assignment $\psi(hgH) = hK$ gives a well-defined bijection from HgH/H to H/K.*

Now let $\mathcal{K}_N = \psi(\mathcal{H}_N \setminus \{H\})$. This is a sequence of larger and larger collections of cosets of K in H. By Claim 13.38.2, these cosets are uniformly pairwise close in G. Since H is quasi-convex, Lemma 13.27 shows it is quasi-isometrically embedded in G, and so these cosets are uniformly close in H as well. In other words, K does not have bounded packing in H. This completes our argument. ☐

It seems worth remarking that whereas our bound on the height of a quasiconvex subgroup of a hyperbolic group is constructive, the proof of bounded packing just given is highly non-constructive.

QUESTION 13.39. *Is there a constructive proof of bounded packing? Such a proof might be expected to give N as a function of D, the quasiconvexity constant of H, the number of generators of G, and the hyperbolicity constant of G.*

EXERCISE 29. What is the right notion of bounded packing for collections of subgroups? Show that versions of Proposition 13.34 and Theorem 13.35 hold for finite collections of subgroups.

CHAPTER 14

Hyperbolic groups acting on cube complexes: finiteness

In this section we give some tools to show finiteness properties of a cube complex coming from applying the Sageev construction to a hyperbolic group.

1. A cocompactness criterion

Our goal in this section is to prove that whenever the Sageev construction is applied to a finite collection of quasi-convex subgroups of a hyperbolic group G, the G–action on the resulting cube complex is cocompact.

We will need to show that if a finite collection of quasi-convex sets in a hyperbolic space are pairwise close, there is a point which is close to all of them. There are a couple of ways to prove this. A slick proof based on asymptotic cones and the topology of $\mathbb{R}$–trees is given by Calegari. We will instead base the proof on approximation of weak hulls by trees.

Definition 14.1. Let X be a δ–hyperbolic space, and let $S \subset X$. Let $\mathrm{WH}(S)$ be the union of all the geodesics of X joining points of S. The set $\mathrm{WH}(S)$ is called the *weak hull of* S.

Exercise 30. Show the weak hull of any set is λ–quasi-convex, where λ depends only on δ.

Tree Approximation Lemma. *For every n, δ, there is an ϵ satisfying the following: For any n–point set Y in a δ–hyperbolic space, there is a metric tree T_Y and a $(1,\epsilon)$–quasi-isometry from $\mathrm{WH}(S)$ to T_Y so that every leaf of T_Y is the image of a point of Y.*

Proof. For now we refer to Chapter 8 of [**CDP90**], but we'll insert at least a sketch into these notes later. □

Lemma 14.2. *Let $n \in \mathbb{N}$, and let $\delta, \lambda, D \geq 0$. Then there is a $D' \geq 0$ so that: For any collection $S_1, \ldots, S_n$ of pairwise D–close λ–quasi-convex sets in a δ–hyperbolic space X, there is a point p so that*

$$\max\{d_X(p, S_i) \mid i = 1, \ldots, n\} \leq D'.$$

Proof. Since the subsets S_i are pairwise D–close there are $\binom{n}{2}$ points $p_{i,j}$, so that $d(p_{i,j}, S_i)$ and $d(p_{i,j}, S_j)$ are at most D. Since S_i is λ–quasi-convex, any geodesic $[p_{i,j_1}, p_{i,j_2}]$ lies in a $2\delta + D + \lambda$–neighborhood of S_i.

Let C be the coarse hull of the points $p_{i,j}$, i.e. the union of all the geodesics of the form $[p_{i,j}, p_{i',j'}]$. By the Tree Approximation Lemma, there is an $\mathbb{R}$–tree T and a $(1,\epsilon)$–quasi-isometry $\phi\colon C \to T$, where the constant ϵ depends only on the numbers n and δ. Possibly increasing ϵ by a small amount (still depending only on δ and n), we can suppose that ϕ has an ϵ–quasi-inverse ψ, which is also a $(1, \epsilon)$–quasi-isometry.

For each fixed i let T_i be the convex hull in T of the points $\{\overline{p}_{i,j} \mid j \neq i\}$. For each i, j, the intersection $T_i \cap T_j$ contains the point $\overline{p}_{i,j}$, so it is nonempty. It follows that the intersection $\bigcap_i T_i$ is nonempty. Let z be a point in that intersection, and let $\tilde{z} = \psi(z)$. Fix i, and note that since z is in T_i, it lies on a geodesic γ_i joining two points $\overline{p}_{i,j}$ and $\overline{p}_{i,j'}$. Thus the point $\tilde{z}$ is on the $(1, \epsilon)$–quasi-geodesic $\tilde{\gamma}_i = \psi(\gamma_i)$ joining $p_{i,j}$ and $p_{i,j'}$. Let R be the constant of quasi-geodesic stability from Theorem 10.5 applied to $(1, \epsilon)$–quasi-geodesics in δ–hyperbolic spaces.

Then $\tilde{z}$ lies in a $2\delta + D + \lambda + R$–neighborhood of S_i, and we can take $D' = 2\delta + D + \lambda + R$. □

From this lemma we easily obtain the cocompactness criterion.

THEOREM 14.3. [**Sag97**, Theorem 3.1] *Let G be hyperbolic, let $\mathcal{H}$ be a finite collection of quasi-convex subgroups, and suppose that X is a cube complex obtained from $(G, \mathcal{H})$ via the Sageev construction. Then the action of G on X is cocompact.*

PROOF. To each element H_i of $\mathcal{H}$ is associated an H_i–halfspace $A_i \subset G$, and the collection of translates of these forms the space with walls $P_{\mathcal{A}}$. Let σ be a maximal cube of X, corresponding to a collection of transverse walls $g_1W_1, \ldots, g_nW_n$. (Each $W_i = \{A_{j_i}, A^c_{j_i}\}$.) We want to show there are finitely many such cubes, up to the G–action. So we can assume that $g_1 = 1$, and fix $H = H_{j_1}$.

Since the walls are transverse, Lemma 13.33 implies that the D–neighborhoods of the cosets $\{H\} \cup \{g_iH_{j_i}\}_{i=2}^n$ must intersect pairwise, where D depends only on the collection $\mathcal{H}$. By Lemma 14.2, there is a D' depending only on n, D, and $\mathcal{H}$ so that some fixed point p lies within D' of all these cosets. Up to the H–action, there are only finitely many choices for such a p. All the other cosets must intersect the D'–neighborhood of p, so there are only finitely many choices of n–tuple, once p has been fixed. Since n is bounded (See Proposition 13.34, Theorem 13.35 and Exercise 29), there are only finitely many cubes σ up to the action of G. □

2. A properness criterion

We start with a general observation.

LEMMA 14.4. *Suppose G contains no infinite torsion subgroup. Suppose G acts on the pocset $(P, \leq)$ in such a way that, for every infinite order g, there is an $A \in P$ and an $n > 0$ so that $A \supsetneq g^nA$. Then the action of G on the cube complex $X(P)$ has finite vertex stabilizers.*

PROOF. Let ω be a vertex of $X(P)$, i.e., a DCC ultrafilter on P. We suppose that $\mathrm{Stab}(\omega)$ is infinite. Since G has no infinite torsion subgroup, there is some infinite order g with $g\omega = \omega$. There is some $A \in P$ and $n > 0$ so that $A \supsetneq g^nA$. Either $A \in \omega$ or $A^* \in \omega$. In case $A \in \omega$, we must also have $g^{kn}A \in \omega$ for all positive k; if $A^* \in \omega$, we have $g^{-kn}A^* \in \omega$ for all positive k. In either case we obtain an infinite descending sequence in ω, contradicting the assumption that ω is DCC. □

We will see that this lemma in particular applies to hyperbolic groups. We first need to understand a little better how an H–wall shows up at the boundary.

LEMMA 14.5. *Let G be hyperbolic, $H < G$ quasi-convex and codimension-one, and let $\{A, A^c\}$ be an H–wall. Then the following hold:*

(1) $\Lambda(A) \cup \Lambda(A^c) = \partial G$.
(2) $\Lambda(A) \cap \Lambda(A^c) = \Lambda(H)$.
(3) $\Lambda(A) \setminus \Lambda(H)$ *is open in* ∂G.

PROOF. Let $N = N_R(H)$ be a neighborhood of H in a Cayley graph for G so that every edge joining a point of A to a point of A^c is in N. There is a constant D so that $N \subset N_D(A) \cap N_D(A^c)$.

Let $\mathbf{x} = \{x_i\}_{i \in \mathbb{N}}$ represent a point p of ∂G. Passing to a subsequence, we can assume either $\mathbf{x} \subset A$ or $\mathbf{x} \subset A^c$. So either $p \in \Lambda(A)$ or $p \in \Lambda(A^c)$.

Suppose that $p \in \Lambda(A) \cap \Lambda(A^c)$, and consider representative sequences $\mathbf{x} = \{x_i\}_{i \in \mathbb{N}}$ in A and $\mathbf{y} = \{y_i\}_{i \in \mathbb{N}}$ in A^c. For each i, any geodesic $[x_i, y_i]$ passes through N. Let z_i be any vertex of N on $[x_i, y_i]$, and note $\mathbf{z} = \{z_i\}_{i \in \mathbb{N}}$ also represents p. The limit set of N is equal to the limit set of H, so $p \in \Lambda(H)$.

Since limit sets are closed, the last item follows from the first two. □

DEFINITION 14.6. Let $W = \{A, A^c\}$ be a partition of a hyperbolic group G into two subsets, and let $a, b \in \partial G$. Then W *separates* a *from* b if (possibly after exchanging A with A^c) $a \in \Lambda(A) \setminus \Lambda(A^c)$ and $b \in \Lambda(A^c) \setminus \Lambda(A)$.

The following then shows us how to apply Lemma 14.4 to hyperbolic groups.

LEMMA 14.7. *Suppose* G *is hyperbolic,* $g \in G$ *is infinite order, and that* H *is a quasi-convex codimension-one subgroup of* G *so that the fixed points of* g *in* ∂G *are separated by an* H*–wall* $\{A, A^c\}$*. Then for some* $n > 0$ *either* $g^n A \subsetneq A$ *or* $g^n A^c \subsetneq A^c$*.*

PROOF. Let $N = N_R(H)$ be chosen so that every edge joining an element of A to an element of A^c is contained in N. If there were arbitrarily large n so that $g^n N \cap N \neq \emptyset$, then the limit sets of $\langle g \rangle$ and H would intersect. Since this isn't the case, we can choose n so that $g^n N \cap N = \emptyset$. Since N contains all the edges connecting A to A^c, the walls $\{A, A^c\}$ and $\{g^n A, g^n A^c\}$ must be nested. Thus one of the four sets

$$g^n A \cap A, \quad g^n A^c \cap A^c, \quad g^n A \cap A^c, \quad g^n A^c \cap A$$

must be empty.

The sets $g^n A \cap A$ and $g^n A^c \cap A$ are not empty, as they contain all sufficiently large positive and negative powers of g, respectively. If $g^n A \cap A^c = \emptyset$, then $g^n A \subsetneq A$, and if $g^n A^c \cap A = \emptyset$, then $g^n A^c \subsetneq A^c$. □

3. Hyperbolic groups as convergence groups

Isometric group actions on hyperbolic spaces induce topological actions on their Gromov boundaries. Somewhat surprisingly, basically all the relevant information is preserved. The *proper* actions, for example, are exactly those for which the action on the boundary has the following "convergence" property.

DEFINITION 14.8. Let G act on a perfect metrizable compact space M, and let $T = T(M)$ be the space of distinct triples of points in M. The action $G \curvearrowright M$ is a *convergence action* if the induced action on T is properly discontinuous.

REMARK 14.9. It is not obvious why the word "convergence" is used here. This is because of an alternative (and older) formulation: A *convergence sequence* in $G \curvearrowright M$ is a sequence of distinct elements $g_n \in G$ so that there exists a pair of points

a, b so that $g_n z \to a$ uniformly in the complement of b. An equivalent definition to 14.8 is that every sequence of distinct elements of G contains a subsequence which is a convergence sequence. See [**Bow99**] for details and generalizations.

LEMMA 14.10. *Let X be proper and Gromov hyperbolic. If G acts properly by isometries on X, then the induced action of G on ∂X is a convergence action.*

PROOF. Sketch: Define a map $\phi\colon T(\partial X) \to X$, taking a triple (x, y, z) to any point which is a 7δ–center for the triple. There is a uniform bound over triples for the diameter of the set of 7δ–centers, using either Lemma 13.11 or Lemma 13.7. An argument along the same lines as those shows that the image of any compact set in T is also a bounded set. Only finitely many elements of G fail to take a given bounded set off of itself, so the same must be true of a compact set in T. In particular, the action $G \curvearrowright T$ is properly discontinuous. $\square$

REMARK 14.11. In this proof we used a coarsely defined map from triples of points in ∂X to X. If X is nice enough we can be more explicit and even get a continuous map. For example if $X = \partial\mathbb{H}^2$, and $(a, b, c) \in T$, one can construct ϕ as follows: Let γ be the bi-infinite geodesic joining a to b, and let $\phi(a, b, c)$ be the beginning of a perpendicular geodesic ray terminating at c. In other words the triple space can be identified with the unit tangent bundle of $\mathbb{H}^2$.

EXERCISE 31. What is $T(\partial X)$ when X is $\mathbb{H}^3$? How about a tree?

DEFINITION 14.12. A convergence action $G \curvearrowright M$ is *uniform* if the induced action on $T(M)$ is cocompact.

PROPOSITION 14.13. *If G is a hyperbolic group, it acts as a uniform convergence group on ∂G.*

PROOF. (Sketch) Lemma 14.10 implies that G acts as a convergence group on ∂G. Choose an arbitrary basepoint $(a, b, c) \in T = T(M)$. The identity $1 \in G$ is an R–center for this triple, for some $R \geq 7\delta$. Let $C \subset T$ be the set of triples for which 1 is an $(R + 1)$–center. This is relatively compact, and since every triple has a 7δ–center, the translates of C cover all of T. $\square$

REMARK 14.14. There is a remarkable converse to Proposition 14.13, due to Bowditch [**Bow98**]: If M is any compact metrizable space without isolated points, and G acts as a uniform convergence group on M, then G is hyperbolic and M is equivariantly homeomorphic to ∂G.

4. Bergeron–Wise's properness criterion

In this section we explain the Bergeron–Wise criterion for a hyperbolic group to be cubulated.

There is also a version of this criterion for relatively hyperbolic groups, see [**BW12**, Theorem 5.1]. The idea here is to combine a compactness argument with the results of Section 2.

DEFINITION 14.15. Say a hyperbolic group G has *enough codimension-one quasi-convex subgroups* if, for every pair of distinct points $a, b \in \partial G$, there is a quasi-convex subgroup H and an H–wall which separates a from b (in the sense of Definition 14.6).

THEOREM 14.16. *Let G be a hyperbolic group with enough codimension-one quasi-convex subgroups. Then G is cubulated, in the sense that G acts properly cocompactly on a $CAT(0)$ cube complex.*

PROOF. We use the assumption of enough codimension-one quasi-convex subgroups to find a particular open cover $\mathcal{U}$ of T. Let u, v, w be distinct. There is some quasiconvex subgroup $H_{u,v}$ and an $H_{u,v}$–wall $\{A, A^c\}$ separating u from v. We may suppose $u \in \Lambda(A) \setminus \Lambda(A^c) = \Lambda(A) \setminus \Lambda(H_{u,v})$. Choose open neighborhoods $U \subset \Lambda(A) \setminus (\Lambda(H_{u,v}) \cup \{w\})$ of u and $V \subset \Lambda(A^c) \setminus (\Lambda(H_{u,v}) \cup \{w\})$ of v, and an open neighborhood W of w disjoint fom $U \cup V$. The product $M_{u,v,w} = U \times V \times W$ gives an open neighborhood of (u, v, w) in T.

Since G acts cocompactly on T (Proposition 14.13), we only need finitely many $M_{u_i,v_i,w_i} = U_i \times V_i \times W_i$ so that their G–translates cover all of T. For each i let H_i be the associated quasi-convex codimension one subgroup, and let $W_i = \{A_i, A_i^c\}$ be the associated H_i–wall. We claim that the wall-space consisting of these finitely many walls and their G–translates gives a cube complex X with a proper cocompact G–action. We do this by verifying the hypotheses of Lemma 14.4.

First, since G is hyperbolic it contains no infinite torsion subgroup (Theorem 13.25. Now let g be an infinite order element. By Corollary 13.26, the element g is loxodromic, so it has two fixed points $g^{\pm\infty}$ in ∂X. Let w be any third point of ∂X. Then there is some i, and some h so that $(g^{\infty}, g^{-\infty}, w) \in hM_i$). Let $K = hH_ih^{-1}$ and $W = \{A, A^c\} = \{hA_i, hA_i^c\}$. Note that W is a K–wall which separates g^{∞} from $g^{-\infty}$. Applying Lemma 14.7, there is an $n > 0$ so that either $A \supsetneq g^n A$ or $A^c \supsetneq g^n A^c$.

But this implies that g cannot preserve any DCC ultrafilter on the pocset P. □

Part III

Miscellany

CHAPTER 15

(Counter)-examples

1. Non-RF groups

Not every group is residually finite. Here is an easy example:

$$BS(2,3) = \langle a, t \mid t^{-1}a^2t = a^3 \rangle.$$

DEFINITION 15.1. G is *Hopfian* if every epimorphism $\phi\colon G \to G$ is an isomorphism. (G is *co-Hopfian* if every monomorphism $\phi\colon G \to G$ is an isomorphism.)

LEMMA 15.2. *If G is RF, then G is Hopfian.*

PROOF. Suppose $\phi\colon G \to G$ is a surjection, and let $k \neq 1$ be an element of the kernel. We will show that, for any n, k is contained in every subgroup of index n. Indeed, let $S_n = \{H_1, \ldots, H_k\}$ be the set of subgroups of index n. For each i, $\phi^{-1}(H_i)$ is also a subgroup of index n, so ϕ determines a bijection $S_n \to S_n$ via $H_i \mapsto \phi^{-1}(H_i)$. But each $\phi^{-1}(H_i)$ contains k, so k is contained in every subgroup of index n. $\square$

EXERCISE 32. The assignments $t \mapsto t$, $a \mapsto a^2$ determine an epimorphism from $BS(2,3)$ to itself which is not an isomorphism.

In fact, there are fundamental groups of NPC square complexes which fail to be RF [**Wis07**]. Even stranger, they can fail to have any nontrivial normal subgroups at all [**BM00**]!

2. Non-LERF RAAGs

The fundamental example of Burns–Karrass–Solitar is the presentation complex of the group:

$$K = \langle a, b, t \mid aba^{-1}b^{-1} = 1, t^{-1}at = b \rangle.$$

This is just a torus T with a cylinder attached, one end to the meridian, and one end to the longitude of T. It's not hard to see this is NPC by drawing the link of the vertex. But Burns–Karrass–Solitar show (using a slightly different presentation) that the element $[a, t^{-1}bt]$ can't be separated from the subgroup $\langle t, ab^{-1} \rangle$ in any finite quotient [**BKS87**]. Later, Niblo and Wise note that K is abstractly commensurable to $A(\Gamma)$ where Γ is a segment of length 3 [**NW01**].

Bibliography

[BH99] Martin R. Bridson and André Haefliger. *Metric Spaces of Non–Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer–Verlag, Berlin, 1999.

[BKS87] R. G. Burns, A. Karrass, and D. Solitar. A note on groups with separable finitely generated subgroups. *Bull. Austral. Math. Soc.*, 36(1):153–160, 1987.

[BM00] Marc Burger and Shahar Mozes. Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.*, (92):151–194 (2001), 2000.

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

[Bow99] B. H. Bowditch. Convergence groups and configuration spaces. In *Geometric group theory down under (Canberra, 1996)*, pages 23–54. de Gruyter, Berlin, 1999.

[BRHP15] Khalid Bou-Rabee, Mark F. Hagen, and Priyam Patel. Residual finiteness growths of virtually special groups. *Mathematische Zeitschrift*, 279(1-2):297–310, 2015.

[Bri02] Martin R. Bridson. The geometry of the word problem. In *Invitations to geometry and topology*, volume 7 of *Oxf. Grad. Texts Math.*, pages 29–91. Oxford Univ. Press, Oxford, 2002.

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

[Cal09] Danny Calegari. *scl*, volume 20 of *MSJ Memoirs*. Mathematical Society of Japan, Tokyo, 2009.

[CCMT15] Pierre-Emmanuel Caprace, Yves Cornulier, Nicolas Monod, and Romain Tessera. Amenable hyperbolic groups. *J. Eur. Math. Soc. (JEMS)*, 17(11):2903–2947, 2015.

[CDP90] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes: Les groupes hyperboliques de Gromov*, volume 1441 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1990.

[Dav08] Michael W. Davis. *The geometry and topology of Coxeter groups*, volume 32 of *London Mathematical Society Monographs Series*. Princeton University Press, Princeton, NJ, 2008.

[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.

[Dro87] Carl Droms. Graph groups, coherence, and three-manifolds. *J. Algebra*, 106(2):484–489, 1987.

[ECH+92] David B. A. Epstein, James W. Cannon, Derek F. Holt, Silvio V. F. Levy, Michael S. Paterson, and William P. Thurston. *Word Processing in Groups*. Jones and Bartlett Publishers, Boston, 1992.

[GdlH90] É. Ghys and P. de la Harpe, editors. *Sur les groupes hyperboliques d'après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1990. Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988.

[GMRS98] Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.

[Gol98] Robert Goldblatt. *Lectures on the hyperreals*, volume 188 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998. An introduction to nonstandard analysis.

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

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

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

[NW01] Graham A. Niblo and Daniel T. Wise. Subgroup separability, knot groups and graph manifolds. *Proc. Amer. Math. Soc.*, 129(3):685–693, 2001.

[Rol98] Martin Roller. *Poc sets, median algebras and group actions. An extended study of Dunwoody's construction and Sageev's theorem.* habilitation, Universität Regensburg, 1998.

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

[Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.

[Sag14] Michah Sageev. CAT(0) cube complexes and groups. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 7–54. Amer. Math. Soc., Providence, RI, 2014.

[Sco78] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.

[Sco85] Peter Scott. Correction to: "Subgroups of surface groups are almost geometric" [J. London Math. Soc. (2) **17** (1978), no. 3, 555–565; MR0494062 (58 #12996)]. *J. London Math. Soc. (2)*, 32(2):217–220, 1985.

[Sho91] Hamish Short. Quasiconvexity and a theorem of Howson's. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 168–176. World Sci. Publ., River Edge, NJ, 1991.

[Sta83] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[Wis07] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.

[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.