UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

School of Mathematics

Embeddings of CAT(0) Cube Complexes in Products of Trees

by

Gemma Lauren Holloway

Thesis for the degree of Doctor of Philosophy

August 2007

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING,SCIENCE AND MATHEMATICS

SCHOOL OF MATHEMATICS

Doctor of Philosophy

TITLE OF THESIS

by Gemma Lauren Holloway

In 'Groups acting on connected cubes and Kazhdan's property T', [29], Niblo and Roller showed that any CAT(0) cube complex embeds combinatorially and quasi-isometrically in the Hilbert space $\ell^2(\mathcal{H})$ where $\mathcal{H}$ is the set of hyperplanes. This Hilbert space may be viewed as the completion of an infinite product of trees. In this thesis, we consider the question of the existence of quasi-isometric maps from CAT(0) cube complexes to finite products of trees, restricting our attention to folding maps as used in [29].

Following an overview of the properties of CAT(0) cube complexes, we first prove that there exists CAT(0) square complexes which do not fold into a product of trees with fewer than k factors for arbitrary k, giving examples which admit co-compact proper actions by right-angled Coxeter groups. We also show that there exists a CAT(0) square complex which does not fold into any finite product of trees.

We then identify a class of group actions on CAT(0) cube complexes for which the existence of such an action implies the existence of a quasi-isometric embedding of that group in a finite product of finitely branching trees. We apply this result to surface groups, certain 3-manifold groups and more generally to Coxeter groups which do not contain affine triangle subgroups.

Contents

List of Figures

DECLARATION OF AUTHORSHIP

I, Gemma Lauren Holloway, declare that the thesis entitled
Embeddings of CAT(0) Cube Complexes in Products of Trees
and the work presented in it are my own. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- none of this work has been published before submission.

Signed: ....

Date:...

Acknowledgements

Firstly, many thanks to my supervisor, Graham Niblo, for his help and direction both while working on my PhD and before.

I also want to thank those with whom I shared an office, Jon, Jon and Cormac, for letting me bounce ideas of them and for distractions when I needed them.

Most importantly, I want to thank my Mum and Dad and the rest of my family for their patience and support in so many ways. I couldn't have done this without them.

Chapter 1

Background

1.1 Introduction

In order to study the geometry of a space we can employ one of two general methods, either controlling the geometry of the space intrinsically or extrinsically. Examples of intrinsic controls on the geometry of a space are the CAT(0) condition and the hyperbolicity condition (see [20]). Extrinsic control on the geometry of a space is achieved via an embedding of the space within another space with known geometry. Examples of spaces in which we can usefully embed are Euclidean space $\mathbb{E}^n$, Hyperbolic space $\mathbb{H}^n$, the sphere $\mathbb{S}^n$ and Hilbert space. We also consider embeddings into products of the above spaces.

A notable success of the extrinsic method is the following theorem

Theorem. *[42] If G is a finitely generated group which coarsely embeds in Hilbert space then the strong Novikov conjecture holds for G.*

It is known that hyperbolic groups quasi-isometrically embed in products of $\mathbb{H}^2$ and these coarsely embed in Hilbert space (see [5]). Other examples include groups with Yu's property A (as defined in [42]) (which includes hyperbolic groups (see [39])), which also embed in Hilbert space.

In [16] Dranishnikov and Januszkiewicz noted that Coxeter groups embed in products of trees (we give a proof of this in lemma 1.48) and used this to show that every Coxeter group acts amenably on a compact space and has finite asymptotic dimension (see [41] for the definition of asymptotic dimension).

It is natural to ask what one can say about the products of trees in which a group can embed quasi-isometrically. In particular, when does it quasi-isometrically embed in a finite product of locally finite trees. Dranishnikov and Schroeder answered this for right-angled Coxeter groups in [17].

In chapter 3 of this thesis we explore generalisations of this result showing that their theorem holds for several widely studies classes of groups, including finitely generated Coxeter groups which contain no subgroups isomorphic to the Euclidean triangle groups $\triangle(2,3,6)$, $\triangle(2,4,4)$ or $\triangle(3,3,3)$, surface groups and some 3-mainifold groups.

Given these results one can ask if every CAT(0) cube complex quasi-isometrically embeds in a finite product of trees. Every asymptotically finite dimensional space X is quasi-isomorphic to a subset of a finite product of trees ([15]). Hence the question of whether every CAT(0) cube complex X embeds quasi-isometrically in a product of trees is equivalent to the question of whether every CAT(0) cube complex has finite asymptotic dimension.

This is a delicate question and it is not even known whether every CAT(0) square complex has finite asymptotic dimension. In this context we provide the following interesting result in chapter 2:

Theorem. *For each $k \in \mathbb{N}$ there exists a right-angled Coxeter group W_k and a 2-dimensional CAT(0) cube complex $\mathcal{U}_k$ such that W_k acts isometrically, cocompactly and properly on $\mathcal{U}_k$ and there is no bending map from $\mathcal{U}_k$ to a product of less than k trees.*

It should be noted that this does not in itself settle the question of the asymptotic dimension as we have restricted the class of maps from quasi-isometries to bending maps (see section 2.1 for definition). However this class of maps seems natural, arising in "Groups acting on CAT(0) cube complexes" by Niblo and Reeves ([26]), and does suggest a possible class of counter examples to the conjecture that CAT(0) cube complexes have finite asymptotic dimension.

1.2 Coxeter Groups

We begin by defining Coxeter groups and considering some of their properties. The material in this section is mainly based on the books "Buildings"[6] and "Reflection Groups and Coxeter Groups"[23].

1.2.1 Definitions

A *Coxeter system* (W, S) consists of a group W and a set of generators $S \subset W$, such that the generating relations for W can be written in the form $(ss')^{m_{s,s'}} = 1$ for some $s, s' \in S$ and some $m_{s,s'} \in \mathbb{N} \cup \{\infty\}$, where $m_{s,s'} = 1$ if and only if $s = s'$. Since $m_{s,s'} = 1$ if $s = s'$, each of the generators of W has order 2. In the case where there is no relation between a pair of generators s, s', we define $m_{s,s'} = \infty$.

We call the group W a *Coxeter group*. If no element $s \in S$ can be written as a product of the elements of $S \setminus \{s\}$ then S is a minimal generating set for W. The *rank* of a Coxeter group W is the size $|S|$ of the set S, where S is a minimal generating set for W. For a given group W the generating set S is not necessarily unique, but any two minimal generating sets for W have the same size, hence the rank of a Coxeter group is well defined. A Coxeter group is *finitely generated* if its rank is finite.

Coxeter groups are a generalisation of the Euclidean reflection groups. An element of a Coxeter group W is called a *reflection* if it is a conjugate of an element of the generating set S. It is easy to see that all elements of this form have order 2.

A *Coxeter matrix* $(m_{s,s'})_{s,s' \in S}$ is a symmetric $|S| \times |S|$ matrix with $m_{s,s} = 1$ for all $s \in S$ and with every other entry either an integer greater than 1 or ∞. A Coxeter matrix defines a Coxeter group W generated by the index set S and relations $(ss')^{m_{s,s'}} = 1$ for all $s, s' \in S$.

Let (W, S) be a Coxeter system. The *Coxeter diagram* of this system consists of a vertex for each element of S together with an edge connecting distinct vertices s and t if $m_{s,t} \neq 2$ and the edge is labelled by $m_{s,t}$.[9]

A Coxeter diagram is *irreducible* if it is connected. A *subdiagram* of a Coxeter diagram is a subcomplex with the same labels as the Coxeter diagram. [18]

1.2.2 The length function

Any element $\gamma \in W$ can be represented by a word w with letters in the set S, say $\gamma = w = s_1 s_2 \dots s_r$. The *length* of the word w is r. We define the length of the identity element to be 0. A word representing an element γ is called a *reduced word for* γ if it has minimal length among all words representing γ.

Theorem 1.1. ([6], page 50)

If (W, S) is a Coxeter system and γ is any element of W then every word w in the generators S which represents γ can be transformed to a reduced word representing γ by a finite sequence of operations of the following types:

(i) delete a subword of the form ss, $s \in S$;

or

(ii) given $s, t \in S$ with $m_{s,t} < \infty$ replace a subword of the form $stst\dots$ of length $m_{s,t}$ by a word $tsts\dots$ of length $m_{s,t}$.

If w and w' are reduced words representing the same element $\gamma \in G$ then w can be transformed into w' by a finite sequence of operations of type (ii).

For a proof of this theorem see chapter II, section 3C of [6].

Analysing the effect of these operations on the length of a word gives the following corollary:

Corollary 1.2. *(a) If the words w and w' represent the same element of W then the lengths of w and w' are either both even or both odd.*

(b) If w and w' are reduced representations of the same element $\gamma \in G$ then w and w' have the same length.

Let $\gamma \in W$ and let w be a reduced representation of γ. The length of w is denoted by $\ell(\gamma)$ and called the *norm* of γ. By corollary 1.2 (b) we see that the lengths of any two reduced words for γ are equal and hence $\ell(\gamma)$ is well defined.

Humphreys [23] lists the following properties of the norm function ℓ for any $\gamma,\ \gamma' \in W$:

Lemma 1.3. *Let γ, γ' be any pair of elements in a Coxeter group W with generating set S.*

1. $\ell(\gamma) = \ell(\gamma^{-1})$.

2. $\ell(\gamma) = 1$ *if and only if* $\gamma \in S$.

3. $\ell(\gamma\gamma') \leq \ell(\gamma) + \ell(\gamma')$.

4. $\ell(\gamma\gamma') \geq \ell(\gamma) - \ell(\gamma')$.

5. $\ell(\gamma) - 1 \leq \ell(\gamma s) \leq \ell(\gamma) + 1$, *for* $s \in S$.

Proof. 1. If $w = s_1 s_2 \dots s_n$ is a reduced word for γ then γ^{-1} can be written as $w' = s_r^{-1} \dots s_2^{-1} s_1^{-1} = s_r \dots s_2 s_1$, which is also reduced, hence $\ell(\gamma) = \ell(\gamma^{-1})$.

2. This result is trivial.

3. Let w be a reduced word for γ and w' a reduced word for γ'. Then the element $\gamma\gamma'$ can be represented by the word ww' and hence $\ell(\gamma\gamma') \leq$ length of w + length of w'= $\ell(\gamma) + \ell(\gamma')$.

4. The element γ can be written as $\gamma(\gamma'\gamma'^{-1}) = (\gamma\gamma')(\gamma'^{-1})$. Hence

$$\begin{aligned}
\ell(\gamma) &= \ell((\gamma\gamma')\gamma'^{-1}) \\
&\leq \ell(\gamma\gamma') + \ell(\gamma'^{-1}) \text{ by point 3} \\
&\leq \ell(\gamma\gamma') + \ell(\gamma') \text{ by point 1}
\end{aligned}$$

Rearranging the inequality gives $\ell(\gamma\gamma') \geq \ell(\gamma) - \ell(\gamma')$.

5. By points 3 and 4 we have

$$\ell(\gamma) - \ell(s) \leq \ell(\gamma s) \leq \ell(\gamma) + \ell(s)$$

Then since $s \in S$ point 2 gives

$$\ell(\gamma) - 1 \leq \ell(\gamma s) \leq (\gamma) + 1$$

$\square$

It follows that setting $d_W(\gamma, \beta) = \ell(\gamma^{-1}\beta)$ we obtain a metric on W. A geodesic between two points $\alpha, \beta \in W$ is given by a sequence $\alpha = \gamma_0, \ldots, \gamma_k = \beta$ with $d(\gamma_i, \gamma_j) = |i - j|$. Such a geodesic exists and corresponds to the reduced word $w = s_1 s_2 \ldots s_r$ for $\alpha^{-1}\beta$ by the following rule:

$$\gamma_0 = \alpha, \gamma_1 = \alpha s_1, \gamma_2 = \alpha s_1 s_2, \ldots, \gamma_r = \alpha s_1 s_2 \ldots s_r = \beta$$

Note that subwords of reduced words are reduced, hence for all $i \geq j$,

$$\begin{aligned} d(\gamma_i, \gamma_j) &= \ell(\gamma_i^{-1}\gamma_j) \\ &= \ell((\alpha s_1 \ldots s_i)^{-1} \alpha s_1 \ldots s_j) \\ &= \ell(s_i^{-1} \ldots s_1^{-1} \alpha^{-1} \alpha s_1 \ldots s_j) \\ &= \ell(s_{i+1} \ldots s_j) = |i - j| \end{aligned}$$

Let α, β and γ be elements of G. We say that γ *lies between* α *and* β if $d(\alpha, \gamma) + d(\gamma, \beta) = d(\alpha, \beta)$. If γ lies between α and β then there exists a geodesic from α to β which contains γ.

1.2.3 Coxeter groups are linear

In this section we will show that every Coxeter group is linear by constructing for each Coxeter group W an injective group homomorphism from W to a general linear group. The proof will be in several stages.

Lemma 1.4. *Given any finitely generated Coxeter group W there is a homomorphism σ from W to a finitely generated linear group.*

The proof of this lemma comes from section 5.3 of [23].

Proof. Given a Coxeter group W choose a Coxeter system (W, S). Let V be a vector space with coefficients in $\mathbb{R}$ and a basis $\{\alpha_s | s \in S\}$. Then a general vector λ in V can be written as $(r_1\alpha_{s_1}, \ldots, r_n\alpha_{s_n})$ where $S = \{s_1, \ldots, s_n\}$ and $r_i \in \mathbb{R}$ for all $i \in \{1, \ldots, n\}$. We will construct a homomorphism from W to the general linear group $GL(V) = GL_n(\mathbb{R})$, where n is the rank of W.

For each pair of generators s and t in S we have a relation of the form $(st)^{m_{s,t}} = 1$, where $m_{s,t} \in \mathbb{N} \cup \{\infty\}$. We can realise these relations geometrically as reflections in hyperplanes such that for any pair of elements s and t the corresponding hyperplanes meet at an angle of $\frac{\pi}{m_{s,t}}$ if $m_{s,t} \in \mathbb{N}$, or as reflections in parallel hyperplanes if $m_{s,t} = \infty$.

We define a bilinear form B on V by

$$B(\alpha_s, \alpha_t) = \begin{cases} -\cos\left(\frac{\pi}{m_{s,t}}\right) & m_{s,t} \neq \infty \\ -1 & m_{s,t} = \infty \end{cases}$$

Since B is bilinear, for a general vector $\lambda = (r_1\alpha_{s_1}, \ldots, r_n\alpha_{s_n})$ we have $B(\lambda, \alpha_s) = r_1B(\alpha_{s_1}, \alpha_s) + \ldots + r_nB(\alpha_{s_n}, \alpha_s)$ and $B(\alpha_s, \lambda) = r_1B(\alpha_s, \alpha_{s_1}) + \ldots + r_nB(\alpha_s, \alpha_{s_n})$.

For each $s \in S$ we define a linear transformation $\sigma_s : V \to V$ by $\sigma_s(\lambda) = \lambda - 2B(\alpha_s, \lambda)\alpha_s$. Since each element of W can be written as a product of the elements of S, we can now define a natural homomorphism $\sigma : W \to GL_n(\mathbb{R})$ by taking the product operation on $GL_n(\mathbb{R})$ to be composition. Let γ be any element of W and let $w = t_1t_2 \ldots t_k$ be a word representing γ with $t_i \in S$ for all $i \in \{1, \ldots, k\}$. Then $\sigma(\gamma) = \sigma(t_1)\sigma(t_2) \ldots \sigma(t_k) = \sigma_{t_1}\sigma_{t_2} \ldots \sigma_{t_k}$.

In order to show that σ is well defined we must check that any two words w and w' representing the same element are mapped to the same linear transformation. By Lemma 1.1 w can be transformed to w' by a sequence of operations of types (i) and (ii). Hence it suffices to show that $\sigma(ss)(\lambda) = \lambda$ for all $s \in S$ and all $\lambda \in V$, and that if $m_{s,t} < \infty$ then the product $sts \ldots$ of length $m_{s,t}$ maps to the same linear transformation as the product $tst \ldots$ of length $m_{s,t}$. Equivalently we show that $\sigma_s^2 = e$ for all $s \in S$ and that $(\sigma_s\sigma_t)^{m_{s,t}} = e$ for any pair $s, t \in S$ such that $m_{s,t} < \infty$.

- $m_{s,t} = 1$

Every element in the generating set S is mapped to an element of order two in the linear group.

To see this, first note that for any $s \in S$

$$\begin{aligned}
\sigma_s(\alpha_s) &= \alpha_s - 2B(\alpha_s, \alpha_s)\alpha_s \\
&= \alpha_s - 2\left(-\cos\left(\frac{\pi}{1}\right)\right)\alpha_s \\
&= \alpha_s - 2\alpha_s = -\alpha_s.
\end{aligned}$$

Hence

$$\begin{aligned}
\sigma_s^2(\lambda) &= \sigma_s(\sigma_s(\lambda)) \\
&= \sigma_s(\lambda - 2B(\alpha_s, \lambda)\alpha_s) \\
&= \sigma_s(\lambda) - 2B(\alpha_s, \lambda)\sigma_s(\alpha_s) \\
&= (\lambda - 2B(\alpha_s, \lambda)\alpha_s) - 2B(\alpha_s, \lambda)(-\alpha_s) \\
&= \lambda
\end{aligned}$$

- $1 < m_{s,t} < \infty$

Definition. A bilinear form B on a vector space V is *positive definite* if for all $x \in V$ with $x \neq 0$, $B(x, x) > 0$. B is *non-degenerate* if $B(x, y) = 0$ for all x implies $y = 0$, and $B(x, y) = 0$ for all y implies $x = 0$.

Consider the two-dimensional subspace $V_{s,t} = \mathbb{R}\alpha_s \bigoplus \mathbb{R}\alpha_t$. Then the restriction of B to $V_{s,t}$ is positive definite and non-degenerate. To see this consider a general non-zero vector $\lambda = a\alpha_s + b\alpha_t$ in $V_{s,t}$. Note that

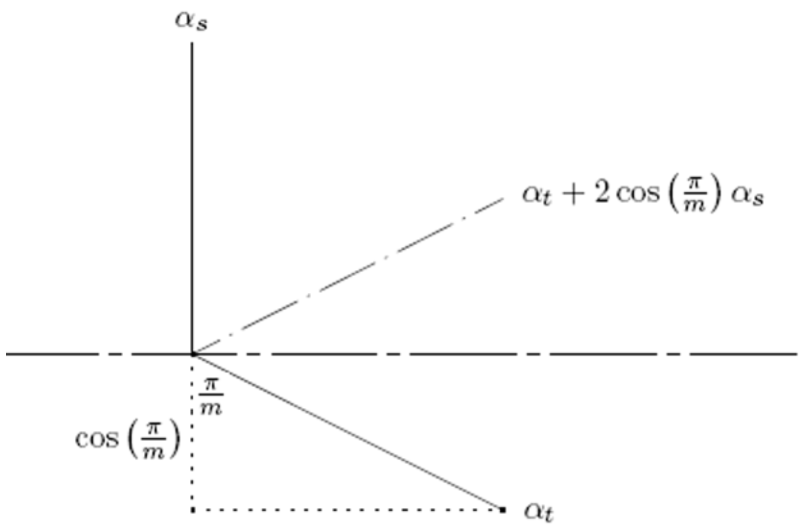

Figure 1.1: The transformation σ_s applied to α_t.

$B(\alpha_s, \alpha_s) = B(\alpha_t, \alpha_t) = -\cos(\pi) = 1$ and let $m = m_{s,t}$. Then

$$\begin{aligned}
B(\lambda, \lambda) &= B(\lambda, a\alpha_s + b\alpha_s) \\
&= aB(\lambda, \alpha_s) + bB(\lambda, \alpha_t) \\
&= a^2 B(\alpha_s, \alpha_s) + 2abB(\alpha_s, \alpha_t) + b^2 B(\alpha_t, \alpha_t) \\
&= a^2 - 2ab\cos\left(\frac{\pi}{m}\right) + b^2 \\
&= \left(a - b\cos\left(\frac{\pi}{m}\right)\right)^2 - b^2\cos^2\left(\frac{\pi}{m}\right) + b^2 \\
&= \left(a - b\cos\left(\frac{\pi}{m}\right)\right)^2 + b^2\sin^2\left(\frac{\pi}{m}\right) \\
&> 0 \qquad \text{when } 1 < m < \infty \text{ and one of } a, b \text{ non-zero}
\end{aligned}$$

Since B is positive definite, we can consider $V_{s,t}$ as the Euclidean plane. Then the transformations $\sigma_s(\lambda) = \lambda - 2B(\alpha_s, \lambda)\alpha_s$ and $\sigma_t(\lambda) = \lambda - 2B(\alpha_t, \lambda)\alpha_t$ are the standard form for orthogonal reflections in two hyperplanes. Since $\sigma_s(\alpha_t) = \alpha_t - 2B(\alpha_s, \alpha_t)\alpha_t = \alpha_t + 2\cos(\frac{\pi}{m})\alpha_t$ we see that the angle between α_s and α_t is $\pi - \frac{\pi}{m}$ (see figure 1.1) and so the angle between the reflecting lines must be $\frac{\pi}{m}$. Comparing this to the case of two reflecting lines in a dihedral group, we see that $\sigma_{s,t}$ must be a rotation by $\frac{2\pi}{m}$ and hence the order of the element $\sigma_s\sigma_t$ in $V_{s,t}$ is m (for a discussion of dihedral groups see section 1.1 of [23]). This tells us that the order of $\sigma_s\sigma_t$ on the entire space V is at least m.

Because the bilinear form B is non-degenerate on $V_{s,t}$, V can be written as the orthogonal direct sum of $V_{s,t}$ and its orthogonal complement $V_{s,t}^{\perp}$ (see lemma 1.2 of [36]). The order of $\sigma_s\sigma_t$ on V will be the least common multiple of its order on $V_{s,t}$ and its order on $V_{s,t}^{\perp}$. Any vector λ in V can be written as the sum of a component in $V_{s,t}$ and a component in the complement, say $\lambda = \lambda_{s,t} \oplus \lambda_{s,t}^{\perp}$. Then

$$\begin{aligned}
\sigma_s(\lambda_{s,t} \oplus \lambda_{s,t}^{\perp}) &= \lambda_{s,t} \oplus \lambda_{s,t}^{\perp} - 2B(\alpha_s, \lambda_{s,t} \oplus \lambda_{s,t}^{\perp})\alpha_s \\
&= \lambda'_{s,t} \oplus \lambda_{s,t}^{\perp}
\end{aligned}$$

for some $\lambda'_{s,t}$ in $V_{s,t}$. This shows that the transformation σ_s fixes $V_{s,t}^{\perp}$ pointwise. Similarly σ_t fixes $V_{s,t}^{\perp}$ pointwise and so $\sigma_s\sigma_t$ acts as the identity on the orthogonal complement of $V_{s,t}$. Hence $\sigma_s\sigma_t$ has order m on the entire vector space V.

In order to check σ is a homomorphism, we will need to show that $(\sigma_s\sigma_t)$ has degree $m_{s,t}$ for all s, t. Since we have already checked the cases $m_{s,t} = 1$ and $1 < m_{s,t} < \infty$ it only remains to check the case when $m_{s,t} = \infty$.

- $m_{s,t} = \infty$

  If $m_{s,t} = \infty$ then $B(\alpha_s, \alpha_t) = B(\alpha_t, \alpha_s) = -1$. We also have $B(\alpha_s, \alpha_s) = B(\alpha_t, \alpha_t) = 1$.

  Consider the vector $\lambda = \alpha_s + \alpha_t$

  $B(\alpha_s, \lambda) = B(\alpha_s, \alpha_s) + B(\alpha_s, \alpha_t) = 1 - 1 = 0$. Similarly $B(\alpha_t, \lambda) = 0$.

  Then

$$\begin{aligned}
\sigma_s(\sigma_t(\lambda)) &= \sigma_s(\lambda - 2B(\alpha_t, \lambda)\alpha_t) \\
&= \sigma_s(\lambda) \\
&= \lambda - 2B(\alpha_s, \lambda)\alpha_s \\
&= \lambda
\end{aligned}$$

Consider the vector α_s in V. Then

$$\begin{aligned}\sigma_s(\sigma_t(\alpha_s)) &= \sigma_s(\alpha_s - 2B(\alpha_t,\alpha_s)\alpha_t) \\ &= \sigma_s(\alpha_s) + 2\sigma_s(\alpha_t) \\ &= \alpha_s - 2B(\alpha_s,\alpha_s)\alpha_s + 2\alpha_t - 4B(\alpha_s,\alpha_t)\alpha_s \\ &= \alpha_s - 2\alpha_s + 2\alpha_t + 4\alpha_s \\ &= 3\alpha_s + 2\alpha_t\end{aligned}$$

So $\sigma_s\sigma_t(\alpha_s) = 2\lambda + \alpha_s$ and by iteration $(\sigma_s\sigma_t)^n(\alpha_s) = 2n\lambda + \alpha_s$. Hence $\sigma_s\sigma_t$ has infinite order on V as required.

Hence the map σ is well defined and is a homomorphism from G to the set of linear transformations on a vector space V.

□

We will show that σ is injective. In order to do this, we need the following lemma from [22]:

Lemma 1.5. *For any $w \in W$ and α_s with $s \in S$ if $\ell(ws) > \ell(w)$ then $\sigma(w)(\alpha_s)$ can be written as $\sum_{s' \in S} \lambda_{s'}\alpha_{s'}$ where $\lambda_{s'} \geq 0$ for all s'.*

Proof. Choose some word $w \in W$ of minimal length such that the theorem fails for some $s \in S$, and choose such an s. Such a w must be non-trivial, hence $\ell(w) \geq 1$ and we can choose some $t \in S\backslash\{s\}$ such that $\ell(wt) = \ell(w) - 1$. We can write w as a reduced word w_1 either of the form $v(st)^k$ or of the form $vt(st)^k$ where k is maximal over all possible reduced words for w and $\ell(vt) > \ell(v)$. Since w_1 is a reduced word and $\ell(ws) > \ell(w)$, we must have $k < \frac{m_{s,t}}{2}$ in the first case and $k + 1 < \frac{m_{s,t}}{2}$ in the second case.

We first consider the case where w_1 is of the form $v(st)^k$:

If $m_{s,t} = \infty$, then we have shown in the proof of the previous lemma that

$$(\sigma_s\sigma_t)^k(\alpha_s) = (2k+1)\alpha_s + 2k\alpha_t.$$

If $m_{s,t} < \infty$ then we have shown that $(\sigma_s\sigma_t)^k$ acts on the space $V_{s,t}$ as a rotation by $\frac{k\pi}{m_{s,t}}$, and on the orthogonal complement of $V_{s,t}$ trivially. Hence

$$(\sigma_s\sigma_t)^k(\alpha_s) = \frac{1}{\sin\left(\frac{\pi}{m_{s,t}}\right)}\left[\sin\left(\frac{(2k+1)\pi}{m_{s,t}}\right)\alpha_s + \sin\left(\frac{2k\pi}{m_{s,t}}\right)\alpha_t\right].$$

Note that this is not the standard form for a rotation due to the fact that the basis for V is orthonormal. Since $k < \frac{m_{s,t}}{2}$ both coefficients are non-negative.

We now consider the case where w_1 is of the form $vt(st)^k$:

If $m_{s,t} = \infty$ then

$$\begin{aligned}
\sigma_t(\sigma_s\sigma_t)^k(\alpha_s) &= \sigma_t((2k+1)\alpha_s + 2k\alpha_t) \\
&= (2k+1)\alpha_s + 2k\alpha_t - 2(2k+1)B(\alpha_t,\alpha_s)\alpha_t - 2(2k)B(\alpha_t,\alpha_t)\alpha_t \\
&= (2k+1)\alpha_s + 2k\alpha_t + 2(2k+1)\alpha_t - 2(2k)\alpha_t \\
&= (2k+1)\alpha_s + (2k+2)\alpha_t
\end{aligned}$$

If $m_{s,t} < \infty$ then

$$\begin{aligned}
\sigma_t(\sigma_s\sigma_t)^k(\alpha_s) &= (\sigma_t\sigma_s)^{k+1}(-\alpha_s) \\
&= -\frac{1}{\sin\left(\frac{\pi}{m_{s,t}}\right)}\left[-\sin\left(\frac{2k+1}{m_{s,t}}\pi\right)\alpha_s - \sin\left(\frac{2k+2}{m_{s,t}}\pi\right)\alpha_t\right]
\end{aligned}$$

and since $k+1 < \frac{m_{s,t}}{2}$ each of these coefficients is non-negative.

Hence $\sigma(w)(\alpha_s) = \sigma(v)(\eta\alpha_s + \mu\alpha_t)$ for some $\lambda, \mu \geq 0$.

Since the word $v(st)^k$ or $vt(st)^k$ was chosen so that k is maximal, it follows that $\ell(vs) > \ell(v)$ and that $\ell(v) < \ell(w)$. Hence by the minimality of w among elements for which the lemma fails, $\sigma(v)(\alpha(s))$ and $\sigma(v)(\alpha_t)$ can be expressed as a linear combination of the basis vectors with non-negative coefficients. Hence $\sigma(w)(\alpha_s) = \eta\sigma(v)(\alpha_s) + \mu\sigma(v)(\alpha_t)$ can be expressed as a linear combination of the basis vectors with non-negative coefficients. This contradicts our choice of w and s as a counterexample to the lemma and completes the proof. □

Lemma 1.6. *Any finitely generated Coxeter group W is a finitely generated linear group.*

Proof. By lemma 1.4, for any finitely generated Coxeter group W there is a homomorphism from W to a finitely generated linear group. In order to complete the proof, we must show that σ is injective, that is for any non-trivial $w \in W$ there exists a $\lambda \in V$ such that $\sigma(w)(\lambda) \neq \lambda$.

Suppose that w is a non-trivial element such that $\sigma(w)$ is trivial. Since w is non-trivial $\ell(w) \geq 1$ and for some $s \in S$ we can write w as $w's$ with $\ell(w') = \ell(w) - 1$. Hence by lemma 1.5 $\sigma(w')(\alpha_s)$ can be written as a linear combination of $\{\alpha_s | s \in S\}$ with non-negative coefficients.

Consider $\lambda = \alpha_s$. Then $\alpha_s = \sigma(w)(\alpha_s) = \sigma(w's)(\alpha_s) = \sigma(w')\sigma_s(\alpha_s) = \sigma(w')(-\alpha_s) = -\sigma(w')(\alpha_s)$. Thus α_s can be expressed as a linear combination of $\{\alpha_s | s \in S\}$ in which every coefficient is non-positive. This contradicts the linear independence of the basis $\{\alpha_s | s \in S\}$ of V. Hence there is no such w and σ is an injective homomorphism. Hence W is a linear group generated by the set $\{\sigma_s | s \in S\}$. □

1.2.4 The Coxeter complex

Let W be a Coxeter group and S be a generating set for W. We call a subgroup $V < W$ a *special subgroup* if $V = \langle S' \rangle$ for some (possibly empty) generating set $S' \subset S$. Consider the cosets of the form wV where $w \in W$ and V is a special subgroup. We define the partial order $\leq$ on the set of cosets by setting $H \leq H'$ if and only if $H' \subset H$ as subsets of W. The partially ordered set $\Sigma(W, S) = \{w \langle S' \rangle | w \in W, \ S' \subset S\}$ is called the *Coxeter complex* of W. We realise this complex as a simplicial complex by identifying each coset of a special subgroup containing k elements with a $k-1$ simplex, mapping the elements of the coset to the vertices of the simplex (see chapter III section 1 of [6]). W acts on $\Sigma(W, S)$.

If A, B are cosets of special subsets with $A \leq B$ then we say that A is a *face* of B. We call a maximal element of the poset Σ a *chamber*. Since $\{e\} \subset V$ for any subgroup V of W, the set of chambers is the set of single element subsets of W. The subgroup $\{e\}$ is the *fundamental chamber* of Σ. The group W acts simply transitively on the left on the set of chambers, that is given any pair of chambers C and C' there is a unique element $w \in W$

such that $C = wC'$.

Given a reflection $r \in W$ we define the *wall* $\mathcal{H}_r$ to be the fixed set of the action of r, $\mathcal{H}_r = \{x \in \Sigma(W,S) | r.x = x\}$. The *stabiliser* of a wall $\mathcal{H}_r$, denoted $\mathrm{stab}_W(\mathcal{H}_r)$, is the set $\{\gamma \in W | \gamma\mathcal{H}_r = \mathcal{H}_r\}$. The *centraliser* of the element r, denoted $C_W(r)$ is the set $\{\gamma \in W | \gamma r = r\gamma\}$.

Any wall $\mathcal{H}_r$ in $\Sigma(W,S)$ separates the complex into two components, called *half-apartments* ([28], p.5). We denote the half-apartments associated to the wall $\mathcal{H}_r$ by X_r and X_r^*. Then $*$ is an involution which exchanges X_r and X_r^*.

Two chambers are *adjacent* in $\Sigma(W,S)$ if they have a common codimension 1 face. Given any pair of adjacent chambers C and C' there is a unique automorphism s of the Coxeter complex of order 2 which exchanges C and C' while fixing $C \cap C'$, and $C \cap C'$ lies in the unique wall $\mathcal{H}_s$ (see [9] pp.2-3).

Lemma 1.7. [9] *The stabiliser of the wall $\mathcal{H}_r$ is the centraliser of the reflection r.*

Proof. Given a reflection r, there exists a pair of adjacent chambers C and C' such that r exchanges C and C' and fixes their common face $C \cap C'$. To see this, note that any reflection r can be written as $\gamma s \gamma^{-1}$ for some $\gamma \in W$ and $s \in S$. Then the pair of chambers $C = \gamma\{e\}$ and $C' = r\gamma\{e\} = \gamma s\{e\}$ satisfies $rC = C'$ and $C \cap C' = \{\gamma, \gamma s\} = \gamma \langle s \rangle$ is a common codimension 1 face of C and C'. Choose such a pair C, C'.

Suppose that $\gamma \in \mathrm{stab}_W(\mathcal{H}_r)$. Since $C \cap C'$ is fixed by r, $C \cap C'$ is in $\mathcal{H}_r$. Then γ maps $C \cap C'$ to another chamber D in $\mathcal{H}_r$ such that r fixes D. Hence both r and $\gamma^{-1} r \gamma$ fix $C \cap C'$ pointwise and exchange γC with $\gamma C'$. Since W acts simply transitively on the set of chambers we must have $\gamma^{-1} r \gamma = r$. Hence $\gamma \in \mathrm{stab}_W(\mathcal{H}_r)$ implies $\gamma \in C_W(r)$.

Now suppose that $\gamma \in C_W(r)$, that is $r\gamma = \gamma r$. Hence $r\gamma(\mathcal{H}_r) = \gamma r(\mathcal{H}_r) = \gamma(\mathcal{H}_r)$. As r fixes the unique wall $\mathcal{H}_r$, $r(\gamma(\mathcal{H}_r)) = \gamma(\mathcal{H}_r)$ means we must have $\gamma\mathcal{H}_r = \mathcal{H}_r$ and hence $\gamma \in \mathrm{stab}_W(\mathcal{H}_r)$. $\square$

1.2.5 Dranishnikov and Schroeder's construction

Definition. [17] A *right-angled Coxeter group* is a Coxeter group W with generating set S and all relations of the form $(st)^{m_{s,t}} = 1$ such that $m_{s,t} \in \{1, 2, \infty\}$ for any pair $s, t \in S$.

Definition. Let W be a right-angled Coxeter group with generating set S. Then a *colouring map* for W is a map $c : S \to \{1, 2, \ldots, n\}$ which satisfies the condition $c(s) \neq c(t)$ if $m_{s,t} = 2$. The minimum value of n for which such a map is possible is called the *chromatic number* of W.

If W is finitely generated then a colouring map exists. If the rank of W is k then the chromatic number of W is less than or equal to k.

Dranishnikov and Schroeder proved the following theorem on embeddings of Coxeter groups in products of trees in [17]. In this theorem the metric on the group W is the word length metric with respect to the generating set S.

Theorem 1.8. *Let W be a finitely generated right-angled Coxeter group with chromatic number n. Then W admits an equivariant isometric embedding into a product of n simplicial trees.*

In general these trees are locally infinite. Dranishnikov and Schroeder also prove the following theorem on embeddings in products of locally finite trees.

Theorem 1.9. *Let W be a finitely generated right-angled Coxeter group with chromatic number n and let T be an exponentially branching locally compact tree. Then for every $r > 0$ there exists a bilipschitz embedding $\psi : W \to T \times \ldots \times T$ (n-factors) such that ψ restricted to every ball of radius r is isometric.*

1.3 Separability Properties

1.3.1 Definitions

Definition. [34] A group G is *residually finite* (RF) if for every $g \in G \setminus \{e\}$ there exists a finite index subgroup G_g of G such that $g \notin G_g$.

Definition. [34] A group G is *locally extended residually finite* (LERF) if for every finitely generated subgroup $S < G$ and every $g \in G \setminus S$ there exists a finite index subgroup G_g of G such that $S < G_g$ and $g \notin G_g$.

Definition. A subgroup $S < G$ is *separable* if there exists a (possibly infinite) set of finite index subgroups of G, which we denote by $\{H_i,\ i \in I\}$, such that $S = \bigcap_{i \in I} H_i$.

Remark 1.10. A group G is residually finite if and only if $\{e\}$ is separable in G. G is locally extended residually finite if and only if every finitely generated $S < G$ is separable in G.

Lemma 1.11. *A group G is residually finite if and only if for every $g \in G \setminus \{e\}$ there exists a finite group F_g and a homomorphism $\phi_g : G \to F_g$ such that $\phi_g(g) \neq e$.*

Proof. Let G be an RF group. Then for each $g \in G$ there exists a finite index subgroup G_g of G such that $g \notin G_g$. Let $\overline{G_g} = \bigcap_{\gamma \in G} \gamma G_g \gamma^{-1}$. Since G_g is finite index, $\overline{G_g}$ is finite index, and since $g \notin G_g$, $g \notin \overline{G_g}$. Then the map $\phi_g : G \to G/\overline{G_g}$ is a homomorphism and $G/\overline{G_g}$ is a finite group F_g. Since $g \notin \overline{G_g}$, $\phi(g) \neq e_{F_g}$.

For every $g \in G \setminus \{e\}$ suppose there exists a finite group F_g and a homomorphism $\phi_g : G \to F_g$ such that $\phi_g(g) \neq e_{F_g}$. Then the kernel of ϕ_g is a finite index (normal) subgroup of G and $g \notin ker(\phi_g)$. $\square$

Lemma 1.12. *G is locally extended residually finite if and only if for every finitely generated subgroup $S < G$ and every $g \in G \setminus S$ there exists a finite group F_g and a homomorphism $\phi_g : G \to F_g$ such that $\phi_g(g) \notin \phi_g(S)$.*

Proof. Suppose that G is a LERF group. Then for any finitely generated subgroup S of G and any $g \in G \setminus S$ there exists a finite index subgroup G_g of G such that $g \notin G_g$ and $S \subset G_g$. Let $\overline{G_g} = \bigcap_{\gamma \in G} \gamma G_g \gamma^{-1}$. Then $\overline{G_g}$ is a finite index normal subgroup of G and $\overline{G_g} \subset G_g$. Let ϕ_g be the homomorphism $\phi_g : G \to G/\overline{G_g}$. Let F_1 be the subset $\phi(S)$ of $G/\overline{G_g}$. Then $\phi_g^{-1}(F_1) = S\overline{G_g}$. Since $S \subset G_g$ and $\overline{G_g} \subset G_g$, $S\overline{G_g} \subset G_g$ and hence $g \notin S\overline{G_g}$, that is $\phi_g(g) \notin \phi_g(S)$.

Suppose that for every finitely generated subset S of G and every $g \in G \setminus \{S\}$ there exists a finite group F_g and a homomorphism $\phi_g : G \to F_g$ such that $\phi_g(g) \notin \phi_g(S)$. Then the kernel of ϕ_g is a finite index (normal) subgroup of G. Let F_1 denote the set $\phi_g(S)$. Then $\phi_g^{-1}(F_1)$ is a union of cosets of $ker(\phi_g)$, and hence is a finite index subgroup of G containing S. Since $\phi_g(g) \notin \phi_g(S)$, $g \notin \phi_g^{-1}(F_1)$ and G is locally extended residually finite. □

Lemma 1.13. [34], [35] *If G is RF or LERF, then any subgroup of G has the same property and so does any group K which contains G as a subgroup of finite index.*

Proof. We first consider the case where G is residually finite. Let H be a subgroup of G and choose any $h \in H \setminus \{e\}$. Then $h \in G \setminus \{e\}$ and by definition there exists a finite index subgroup G_h of G such that $h \notin G_h$. Then $H_h = G_h \cap H$ is a finite index subgroup of H and $h \notin H_h$.

Now let K be a group containing G as a subgroup of finite index. If G is not normal in K, replace G by $G' = \bigcap_{k \in K} kGk^{-1}$. Since G is finite index in K, there are finitely many conjugates of G and hence G' is a finite index normal subgroup of K. Since $G' < G$ we know from the previous paragraph that G' is RF. Suppose $k \in K \setminus \{e\}$. If $k \in G$ then since G is RF there exists a subgroup $G_k < G$ which is a finite index subgroup of K not containing k. If $k \notin G$ then we take the subgroup G itself as the finite index subgroup.

We now consider the case where G is locally extended residually finite. Let H be a subgroup of G, let S_H be a finitely generated subgroup of H and suppose $h \in H \setminus S_H$. Then $h \in G \setminus S_H$ and by definition there exists a finite index subgroup G_h of G such that $S_H < G_h$ and $h \notin G_h$. Then $H_h = G_h \cap H$ is a finite index subgroup of H such that $S_H < H_h$ and $h \notin H_h$.

Now let K be a group containing G as a subgroup of finite index. If G is not normal in K, replace G by $G' = \bigcap_{k \in K} kGk^{-1}$. Since $G' < G$ we know G' is LERF. Let F be the finite quotient group K/G and let $p : K \to F$ be the natural projection map.

Let S_K be a finitely generated subgroup of K and suppose $k \in K \setminus S_K$. Since G is normal in K, $S_G = S_K \cap G$ is a finitely generated normal subgroup of S_K. Suppose g_1 and g_2 are elements of S_K with $p(g_1) = p(g_2)$, then there exists some $g \in G$ for which $g_1 = gg_2$. Since g_1 and g_2 lie in the subgroup S_K, so does $g_1 g_2^{-1} = g$ and hence g lies in $S_K \cap G = S_G$. Hence for any pair g_1, g_2 of elements of S_K, $p(g_1) = p(g_2)$ if and only if $g_1 = gg_2$ for some $g \in G \cap S_K$ and there exists an isomorphism from the quotient group S_K/S_G to a subgroup F_1 of F.

Consider $K_1 = p^{-1}(F_1)$. If $k \notin K_1$ then we can take K_1 to be the required finite index subgroup of K containing S_K. If $k \in K_1$ then we can write $k = gs$ where $g \in G$ and $s \in S_K$. Since $k \notin S_K$, we know $g \notin S_K$ and hence $g \notin S_G$. Then, since G is LERF there exists a finite index subgroup G_g of G (and hence of K) such that $S_G < G_g$ and $g \notin G_g$. Let $\overline{G_g}$ be the subset of G_g consisting of all elements normalised by S_K, $\overline{G_g} = \bigcap_{s \in S_K} sG_g s^{-1}$. Since we already know that S_G is normal in S_K, we have $S_G < \overline{G_g}$. Clearly we have $g \notin \overline{G_g}$. Let K_2 be the subgroup generated by $\overline{G_g}$ and S_K. Since K_2 contains $\overline{G_g}$, a finite index subgroup of K, K_2 is finite index in K, and by construction K_2 contains S_K and does not contain k. Hence K is LERF. $\square$

1.3.2 Examples of RF and LERF groups.

Examples of RF groups include all finite groups, free groups, surface groups and the fundamental groups of Haken 3-manifolds ([21]). In "On faithful representations of infinite groups of matrices", [25], Malcev proved the following:

Theorem 1.14. *Let G be a finitely generated linear group. For every finite set $\{g_1, \ldots, g_k\}$ of elements of G there exists a finite group H and a homomorphism Φ from G to H such that if $g_i \neq g_j$ then $\Phi(g_i) \neq \Phi(g_j)$. Hence G is residually finite.*

Proof. Let G be a finitely generated linear group. Then G can be faithfully represented as a subgroup of $GL(n,\mathbb{R})$. Choose a set $M_1,\dots,M_m$ of matrices such that $M_j \in GL(n,\mathbb{R})$ for all j and every element of G can be written as a product of non-negative powers of these M_j. Since G is finitely generated, this can be done by selecting the matrices representing the generators of G and their inverses. We write the matrix M_j as $(a^j_{\alpha\beta})_{\alpha,\beta\in 1,\dots,n}$.

We have a set of relations S between the generators of the group, written as $S_i(M_1,\dots,M_m)-I_n = 0_n$ where S_i is a polynomial with integer coefficients and variables in the set $M_1,\dots M_m$.

In general the set S contains an infinite number of equations. By the Hilbert basis theorem, we can replace S by a finite set of equations such that all the equations in the original set are a consequence of these.

In each of these equations the elements of the matrix on the left hand side are polynomials in $x^j_{\alpha\beta}$ with integer coefficients which give the required relation when set equal to 0. There are a finite number of such polynomials and a set $a^j_{\alpha\beta}$ with elements in the field $\mathbb{R}$ such that setting $x^j_{\alpha\beta} = a^j_{\alpha\beta}$ satisfies the set of polynomials. Since there are a finite number of polynomials which these $x^j_{\alpha\beta}$ must satisfy, there is a finite degree extension of $\mathbb{Z}$ in which such a choice of $x^j_{\alpha\beta}$ is possible. Hence each of the generators $M_1,\dots,M_m$ and each element of G can be written as an element of $GL(n,R)$ where R is a finite degree extension of $\mathbb{Z}$. Hence G is a subgroup of $GL(n,R)$.

Consider a finite set of elements $\{g_1,g_2,\dots,g_k\} \subset G$. For each g_i let $P_i(M_1,\dots,M_m)$ be an expression for g_i in terms of the generating set. For each pair g_i, g_h of elements in the set $\{g_1,\dots,g_k\}$ with $g_i \neq g_h$, we have the non-equality $P_i(M_1,\dots,M_m) - P_h(M_1,\dots,M_m) \neq 0_n$. In each of these relations the entries of the matrix on the left hand side are polynomials in $x^j_{\alpha\beta}$. The $x^j_{\alpha\beta}$'s must be chosen so that at least one entry in each matrix is non-zero.

Now consider the finite set S^* of relations in the generators given by the set of equalities S together with the set of non-equalities of the form $P_i(M_1,\dots,M_m)-P_h(M_1,\dots,M_m) \neq 0_n$. As our chosen matrices $M_1,\dots,M_m$ in $GL(n,R)$ generate G and the expressions P_i and P_h represent distinct group elements g_i and g_h, setting $x^j_{\alpha\beta} = a^j_{\alpha\beta}$ as above satisfies the set of relations S^*.

With these choices of $x^j_{\alpha\beta}$ each of the polynomial entries in the matrices $P_i(M_1,\dots,M_m) - P_h(M_1,\dots,M_m)$ takes a value $p^{ih}_{\alpha\beta} \in R$, and for each pair i,h at least one of the $p^{ih}_{\alpha\beta}$ is non-zero. Choose a prime p such that $p > \max\{a^j_{\alpha\beta}, p^{ih}_{\alpha\beta} | j = 1,\dots,m;\ i,h = 1,\dots,k, i \neq h, \alpha,\beta = 1,\dots,n\}$. Take K to be the field $R/p\mathbb{Z}$. There is a natural map $\phi : R \to K$. The set of relations S^* is soluble in the field K, specifically we can take $x^j_{\alpha\beta} = \phi(a^j_{\alpha\beta})$ as a set of solutions. Since R is a finite degree extension of $\mathbb{Z}$ the field K is finite.

Let $N_j = (x^j_{\alpha\beta})$ where $x^j_{\alpha\beta} = \phi(a^j_{\alpha\beta})$ for all $j = 1,\dots,k,\ \alpha,\beta = 1,\dots,n$. This choice of matrices satisfies the set of relations S. For each $g \in G$, g can be written as a product $P_g(M_1,\dots,M_m)$ of the generators. Define $\Phi : G \to H$ by $g \mapsto P_g(N_1,\dots N_m)$. By our choice of prime p, $P_{g_i}(N_1,\dots,N_m) - P_{g_h}(N_1,\dots N_m)$ is non-zero for all $i,h \in 1,\dots,k$ and hence $\Phi(g_i) \neq \Phi(g_h)$ for all g_i, g_h in our chosen finite set.

Since K is a finite field, the group $GL(n,K)$ must also be finite. The group H is generated by a finite set of elements of $GL(n,K)$ so $H < GL(n,K)$ and H is finite. □

Corollary 1.15. *Let G be a finitely generated Coxeter group. Then G is residually finite.*

Proof. By Lemma 1.6 every finitely generated Coxeter group is a finitely generated linear group. Theorem 1.14 completes the proof. □

Lemma 1.16. *Let G be a residually finite group and α be an automorphism of G. Then the subgroup $\mathrm{Fix}(\alpha) = \{g \in G | \alpha(g) = g\}$ is separable in G.*

Proof. Let $\alpha : G \to G$ be an automorphism of a residually finite group. Choose an element $\gamma \in G \setminus \mathrm{Fix}(\alpha)$. Then, by the definition of $\mathrm{Fix}(\alpha)$, we have $\gamma^{-1}\alpha(\gamma) \neq e_G$. Denote $\gamma^{-1}\alpha(\gamma)$ by γ'. Since G is a residually finite group, there exists a finite group $F_{\gamma'}$ and a homomorphism $\phi_{\gamma'} : G \to F_{\gamma'}$ such that $\phi_{\gamma'}(\gamma^{-1}\alpha(\gamma)) \neq e_{F_{\gamma'}}$.

Construct a map $\Phi : G \to F_{\gamma'} \times F_{\gamma'}$ defined by $\Phi(g) = (\phi_{\gamma'}(g), \phi_{\gamma'}(\alpha(g)))$. Note that $\phi_{\gamma'}(\gamma^{-1}\alpha(\gamma)) \neq e_{F_\gamma}$ implies that $\phi_{\gamma'}(\gamma) \neq \phi_{\gamma'}(\alpha(\gamma))$ and so $\Phi(\gamma) \notin \{(f,f) | f \in F_{\gamma'}\}$, the diagonal subgroup of $F_{\gamma'} \times F_{\gamma'}$.

Clearly, if $g \in \mathrm{Fix}(\alpha)$ then $g = \alpha(g)$ and so $\Phi(g) \in \{(f,f)|f \in F_\gamma\}$. Hence for any $\gamma \in G \setminus \mathrm{Fix}(\alpha)$ the pullback $\Phi^{-1}(\{(f,f)|f \in F_\gamma\})$ is a finite index subgroup of G containing $Fix(\alpha)$ but not containing γ, and so $\mathrm{Fix}(\alpha)$ is separable in G. □

Corollary 1.17. *Let G be a residually finite group. Then $C_G(g)$, the centraliser of g in G, is separable for all $g \in G$.*

Proof. Consider the homomorphism $\alpha : G \to G$ given by $r \mapsto g^{-1}rg$. Then $\mathrm{Fix}(\alpha) = C_G(r)$ and Lemma 1.16 gives the result. □

Corollary 1.18. *Let W be a finitely generated Coxeter group. Let $\mathcal{H}_r$ be a wall in the Coxeter complex $\Sigma(W,S)$. Then $\mathrm{stab}_W(\mathcal{H}_r)$ is separable.*

Proof. By Corollary 1.15 W is residually finite and by Lemma 1.7, $\mathrm{stab}_W(\mathcal{H}_r) = C_W(r)$. Hence by Corollary 1.17 $\mathrm{stab}_W(\mathcal{H}_r)$ is separable. □

Lemma 1.19. [24] *Let G be a RF group and H a maximal abelian subgroup of G. Then H is separable in G.*

Proof. Suppose H is a maximal abelian subgroup of G. Let $f(x_1,x_2) = x_1x_2x_1^{-1}x_2^{-1}$. Then H is maximal subject to the condition $f(h_j,h_k) = e$ for all $h_j,h_k \in H$. Since G is residually finite we can choose a set of finite index normal subgroups N_i in G such that $\bigcap_{i\in I} N_i = \{e\}$. We take $M_i = HN_i$ and will show that $H = \bigcap_{i\in I} M_i$. Since N_i is finite index in G and $N_i < HN_i = M_i$, each M_i is finite index in G.

Note that

$$\bigcap_{i\in I} M_i = \bigcap_{i\in I} HN_i \supset H\{\bigcap_{i\in I} N_i\} = H\{e\} = H.$$

For each choice of N_i we have

$$\begin{aligned} f(h_jN_i, h_kN_i) &= (h_jN_i)(h_kN_i)(N_ih_j^{-1})(N_ih_k^{-1}) \\ &= h_jh_kh_j^{-1}h_k^{-1}N_i \\ &= N_i \end{aligned}$$

and so $f(h_jN_i, h_kN_i) \in N_i$. Hence $\bigcap_{i\in I} f(HN_i, HN_i) < \bigcap_{i\in I} N_i = \{e\}$, so that any pair x_1, x_2 in the subgroup $\bigcap_{i\in I} HN_i$ satisfy $f(x_1, x_2) = e$. Since H was chosen to be maximal with respect to this property, it follows that $H > \bigcap_{i\in I} HN_i$ and so $H = \bigcap_{i\in I} M_i$. $\square$

The following results are due to Scott and Hall respectively.

Theorem 1.20. [34, 35] *Every surface group is LERF.*

Corollary 1.21. *Every finitely generated free group is LERF.*

Proof. Every finitely generated free group is a subgroup of the fundamental group of the genus 2 orientable surface. Hence by lemma 1.13 and theorem 1.20, every finitely generated free group is LERF. $\square$

Corollary 1.22. *Every Fuchsian group is LERF.*

Proof. Any Fuchsian group has a subgroup of finite index which is a surface group. Hence by lemma 1.13 and theorem 1.20, every Fuchsian group is LERF. $\square$

1.3.3 Scott's results

Definition. [2] A topological space X is called *Hausdorff* if for any pair of distinct points $u, v \in X$ we can choose open neighbourhoods U and V such that $u \in U$, $v \in V$ and $U \cap V = \emptyset$.

In"Subgroups of surface groups are almost geometric", [34], and a later correction, [35], Peter Scott gives the following characterisation of residual finiteness.

Lemma 1.23. ([34],page 557) *Let X be a Hausdorff topological space with a regular covering $\tilde{X}$ and covering group G. Then the following conditions are equivalent*

(a) G is RF,

(b) If C is a compact set in $\tilde{X}$ then G has a subgroup G_1 of finite index such that $gC \cap C$ is empty for every non-trivial element g of G_1,

(c) If C is a compact set in $\tilde{X}$ then the projection map $\tilde{X} \to X$ factors through a finite covering X_1 of X such that C projects by a homeomorphism into X_1.

This generalises to the following characterisation of locally extended residual finiteness.

Lemma 1.24. ([34],page 557) *Let X be a Hausdorff topological space with a regular covering $\tilde{X}$ and covering group G. Then G is LERF if and only if given any f.g. subgroup S of G and a compact subset C of $\tilde{X}/S$ there is a finite covering X_1 of X such that the projection $\tilde{X}/S \to X$ factors through X_1 and C projects homeomorphically into X_1.*

Definition. ([34], page 555) Let F be a surface and X be a compact subsurface of F. X is incompressible if no component of the closure of $F \setminus X$ is a 2-disc whose boundary is contained in the boundary of X.

Definition. ([34], page 555) Let F be a surface. A subgroup S of $\pi_1(F)$ is *geometric* if $S = \pi_1(X)$ for some incompressible subsurface X of F.

Scott gave the following theorem about embedded loops in surfaces, which implies that the fundamental group of a surface is LERF.

Theorem 1.25. ([34], page 561) *Let F be a surface, let S be a finitely generated subgroup of $\pi_1(F)$ and let $g \in \pi_1(F) \setminus S$. Then there is a finite covering F_1 of F such that $\pi_1(F_1)$ contains S but not g and S is geometric in F_1.*

Scott also uses this approach to give a proof that the fundamental group of a compact Seifert fibre space is LERF.

1.4 Cube Complexes and Products of Trees

1.4.1 Definitions

Definition. ([4], p. 111) The *unit n-cube* I^n is the n-fold product $[0,1]^n$, isometric to a cube in $\mathbb{E}^n$ with edges of length one. I^0 is defined to be a point. Faces of the cube are defined in the usual way as lower dimensional cubes embedded in the boundary of the cube. We denote the boundary of the cube σ by $\partial\sigma$.

Definition. ([4], p.112) A *cube complex* K is the quotient of a disjoint union of cubes $X = \coprod_{\lambda\in\Lambda} I^{n_\lambda}$ by an equivalence relation $\sim$. The restrictions $p_\lambda : I^{n_\lambda} \to K$ of the natural projection $p : X \to K = X/\sim$ to a cube are required to satisfy:

1. for every $\lambda \in \Lambda$ the map p_λ is injective;
2. if $p_\lambda(I^{n_\lambda}) \cap p_{\lambda'}(I^{n_{\lambda'}}) \neq \emptyset$ then there is an isometry $h_{\lambda,\lambda'}$ from a face $T_\lambda \subset I^{n_\lambda}$ onto a face $T_{\lambda'} \subset I^{n_{\lambda'}}$ such that $p_\lambda(x) = p_{\lambda'}(x')$ if and only if $x' = h_{\lambda,\lambda'}(x)$.

The *dimension* of a cube complex X is equal to the dimension of the highest dimension cube in X. The set of 0-dimensional cubes are called the vertices of X. The set of 1-dimensional cubes are called the edges of X.

Throughout this thesis, we will assume that all cube complexes are finite dimensional and locally finite, that is only finitely many cubes meet any point.

We will consider two different metrics on the set of vertices of a cube complex X. The first is the natural path metric inherited from $\mathbb{E}^n$, denoted by d_2. This metric can be defined on the whole of X. We will also consider the edge metric d_1 defined on the set of vertices of X by $d_1(u,v) = \min\{\text{length of } p | p \text{ is a edge path joining } u \text{ and } v\}$ where u and v are any pair of vertices of X.

Definition. A *graph* G is a collection of vertices V and a set of edges E which join pairs of vertices. We can denote the edge e which joins the vertices u

and v by the pair (u, v). We say an edge e is *incident* with a vertex v if e joins u and v for some vertex u in G.

A *path* in G is a sequence of edges $(u_1, v_1), (u_2, v_2), \ldots, (u_j, v_j)$ such that $v_i = u_{i+1}$ for all $1 \leq i < j$. The *length of a path* is the number of edges in the sequence. A *cycle* in G is a path in G with the same initial and terminal vertices, in which no edge or vertex occurs more than once. The *length of a cycle* is the number of edges in the cycle. A *tree* is a graph containing no cycles.

A *geodesic* is a path whose length is minimal among the set of lengths of paths between its end points. In a general metric space there may be more than one geodesic with the same endpoints.

Given a cube complex X denote by $X^{(i)}$ the *i-skeleton of* X, that is the subcomplex formed from all cubes of X of dimension less than or equal to i. The 1-skeleton $X^{(1)}$ of a cube complex is a graph. An edge path in a CAT(0) cube complex X is entirely contained in its 1-skeleton $X^{(1)}$.

Any tree is a 1-dimensional CAT(0) cube complex.

Definition. ([4], p.34) Let $\triangle$ be a triangle in a metric space (X, d_X) with vertices p, q and r and edge lengths $d_X(p, q)$, $d_X(q, r)$ and $d_X(r, p)$. A *comparison triangle* for $\triangle$ in Euclidean space is a Euclidean triangle $\bar{\triangle}$ with vertices $\bar{p}$, $\bar{q}$, $\bar{r}$ such that $d(\bar{p}, \bar{q}) = d_X(p, q)$, $d(\bar{q}, \bar{r}) = d_X(q, r)$ and $d(\bar{r}, \bar{p}) = d_X(r, p)$.

A point $\bar{x} \in [\bar{q}, \bar{r}]$ is called a *comparison point* for $x \in [q, r]$ if $d(\bar{q}, \bar{x}) = d_X(q, x)$. Comparison points for points on $[p, q]$ and $[p, r]$ are defined in the same way.

Definition. ([4], p.158) A metric space X is *CAT(0) with respect to the metric* d_X if for every geodesic triangle $\triangle$ in X and comparison triangle $\bar{\triangle}$ in Euclidean space, every pair of points $x, y \in \triangle$ with comparison points $\bar{x}, \bar{y} \in \bar{\triangle}$ satisfies $d_X(x, y) \leq d(\bar{x}, \bar{y})$. Where the metric is clear from the context, we will say that the space is CAT(0).

A metric space X is *locally CAT(0)* if for every $x \in X$ there exists a ball $B(x, r_x)$ which is CAT(0) with respect to the induced metric.

We say that a cube complex is CAT(0) if it is CAT(0) with respect to the inherited metric d_2.

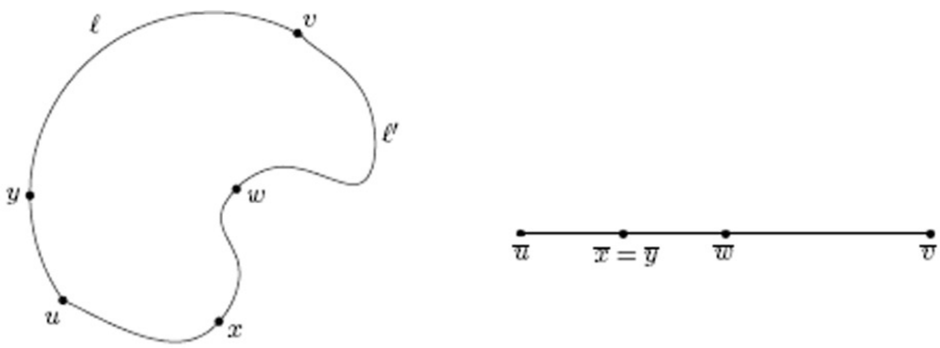

Figure 1.2: A comparison triangle for a pair of geodesics

Lemma 1.26. *Geodesics in CAT(0) spaces are unique.*

Proof. Let (X, d) be a metric space which is CAT(0) with respect to the metric d. Let u and v be two points in X and suppose that l and l' are two geodesics joining u and v. Choose a point w on l. Then the comparison triangle for the triangle with vertices u, v, w is a line segment with length equal to the length of l, as in figure 1.2. Let x and y be points on l and l' respectively with $d_X(u, x) = d_X(u, y)$. Then we have $\overline{x} = \overline{y}$ and, since X is CAT(0), $d_X(x, y) \leq d(\overline{x}, \overline{y}) = 0$ and by the definition of a metric must have $x = y$. This is true for any pair of corresponding points, and hence we must have $l = l'$.

□

In the case of a cube complex X, we have an alternative method of checking whether X is CAT(0). This method is based on examining the set of directions at each point in X, and results in a combinatorial condition which must be satisfied for X to be locally CAT(0).

Definition. ([20], pp.102-103, 108) Let X be a cube complex, and let x be some point in X. The *open star* of x, denoted $st(x)$, is the union of the interiors of the cells that contain x. Given $y, y' \in st(x) \setminus \{x\}$, we say that the geodesic segments $[x, y]$ and $[x, y']$ define the same *direction* at x if one of them is contained in the other. The *geometric link* of x is the set $Lk_X(x)$ of directions at x. For any two directions u and u' contained in the same cube of

$st(x)$, we define the distance between u and u' to be the angle between u and u'. We hence define a distance function on $Lk_X(x)$ by taking the distance between u and u' to be the length of the shortest path $u = u_0, u_1, \ldots, u_m = u'$ where for each $1 \leq i \leq m$ u_{i-1} and u_i lie in the same simplex of $st(x)$.

Lemma 1.27. *(Theorem 7.16 of [4]) Let X be a cube complex and x be a vertex in X. Then $Lk_X(x)$ is isometric to $\partial(B(x,1) \cap X)$, that is the intersection of the boundary of a ball of radius 1 with X.*

Definition. [33] The *link of a m-cube σ* in X, denoted $link_X(\sigma)$ is a simplicial complex whose n-skeletons are defined inductively as follows: The set of vertices is the set $\{\tau \in X^{(n+1)} : \sigma \in \partial\tau\}$, and the set $\{\tau_0, \ldots, \tau_n\} \subset link_X(\sigma)^{(0)}$ spans an n-simplex if there exists a cube C such that every τ_i lies in the boundary of C. Where the space X is clear from the context, we denote $link_X(\sigma)$ by $link(\sigma)$.

We note that in the case where σ is a vertex in X the link of the 0-cube σ, $link_X(\sigma)$, is isometric to the geometric link $Lk_X(\sigma)$ of the point σ in X, where the metric on $link_X(\sigma)$ is defined by setting each edge in $link_X(\sigma)$ to have length $\frac{\pi}{2}$.

Definition. [27] A cube complex X is *non-positively curved* if for any cube C the following conditions are satisfied:

(i) (no bigons) For each pair of vertices in $link_X(C)$ there is at most one edge containing them.

(ii) (no triangles) Every edge cycle of length three in $link_X(C)$ is contained in a 2-simplex of $link(C)$.

Definition. [2] Let X be a topological space. A *path* in X is a continuous function $\gamma : [0,1] \to X$. A *loop* in X is a path $\gamma : [0,1] \to X$ such that $\gamma(0) = \gamma(1)$.

X is *path-connected* if any pair a, b of its points can be joined by a path, that is there exists a path $\gamma : [0,1] \to X$ such that $\gamma(0) = a$ and $\gamma(1) = b$.

For any subset A of X, let $cl(A)$ denote the closure of A. X is *connected* if whenever it is decomposed as the union $A \cup B$ of two non-empty subsets then either $cl(A) \cap B \neq \emptyset$ or $A \cap cl(B) \neq \emptyset$.

X is *simply connected* if it is path-connected and has trivial fundamental group, that is every loop in X is homotopic to the constant loop.

Remark 1.28. If X is a metric space then X is simply connected if any two points in X can be joined by a unique geodesic path.

Lemma 1.29. *[20] X is locally CAT(0) if and only if it is non-positively curved, and it is CAT(0) if and only if it is simply connected and non-positively curved.*

Definition. Given a cell complex X and a cell in x, the *star of* x, denoted $star(x)$, is the union of all cells of X which contain x in their boundaries, $star(x) = \bigcup_{C \in X, x \in \partial C} C$.

Definition. [33] An *automorphism* of a cube complex X is a homeomorphism of the underlying space which restricted to each cube is a linear isomorphism onto a cube of X.

An *action* of G on X is a homomorphism ϕ from G to the automorphism group of X. We will write $g(x)$ for $\phi(g)(x)$ when the action is clear from the context.

Definition. We say the action of a group G on a space X is *cocompact* if there exists a compact subspace C of X such that $GC = X$.

Definition. We say the action of a group G on a space X is *free* if for any $g \in G$ and $x \in X$ $g(x) = x$ implies $g = e$ where e is the identity.

Definition. We say the action of a group G on a space X is *properly discontinuous* (or *proper*) if for every compact subspace K of X the set $\{g \in G | gK \cap K \neq \emptyset\}$ is finite.

Lemma 1.30. *Let G be a group and X be a CAT(0) cube complex on which G acts properly and cocompactly. Then G is a finitely generated group.*

Proof. If G acts properly and cocompactly on X, there is some compact region $C \subset X$ such that $GC = X$. Let $S = \{g \in G | gC \cap C \neq \emptyset\}$. Since that action of G on X is properly discontinuous, S is finite. We claim S is a generating set for X. Suppose $\gamma \in G$. Then $\gamma C \in X$ and since X is path-connected we can choose a path p from a point in eC to a point in γC. Let $g_0 C = eC, g_1 C, \ldots, g_n C = \gamma C$ denote the images of C through which p passes.

Since each $g_i C$ is compact and hence closed, it follows that $g_i C \cap g_{i+1} C \neq \emptyset$. Hence $g_i^{-1} g_i C \cap g_i^{-1} g_{i+1} C = C \cap g_i^{-1} g_{i+1} C \neq \emptyset$ and by definition $g_i^{-1} g_{i+1} \in S$. This is true for each choice of i and hence $g_n = \gamma \in \langle S \rangle$. Hence G is generated by the finite set S. □

Suppose G is a group with finite generating set S. Then there is a natural length function on the elements of G given by defining $\ell_w(g)$ to be the minimum possible length of γ where γ is a word for g in the elements of S and their inverses. Using this length function, we define a metric on G by setting $d_W(g_1, g_2)$ to be $\ell_w(g_1 g_2^{-1})$ (see section 1.1.2 for a discussion of this metric in the case of finitely generated Coxeter groups). We call this metric the word length metric.

Definition. Let G be a group with finite generating set S. We construct the *Cayley graph* $\Gamma(G; S)$ of G as follows: Let $V(\Gamma) = G$. We connect vertices $g, h \in G$ by an edge in Γ if $g^{-1}h \in S$, i.e. For all $g \in G$, $s \in S$ we have an edge connecting g to gs.

We can use any action a of a group G on a space X to define a map α from G into X as follows. Choose any point $x \in X$ and define $\alpha(e) = x$ where e is the identity element of G. Then for any element g of G we define $\alpha(g) = g(x)$.

If the action of G on X is free then the map α will be an embedding of G in X, that is $\alpha(g) = \alpha(h) \implies g = h$.

Definition. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \to Y$ is an *isometric embedding* if for every pair of points $u, v \in X$, $d_Y(f(u), f(v)) = d_X(u, v)$. The map f is an *isometry* if it is also surjective. If there is an isometric map $f : X \to Y$ then we say X and Y are *isometric.*

For example, there exists a isometric embedding $f : \mathbb{Z} \to \mathbb{R}$ given by mapping the integer z to the point z on the real line. However, this map is not surjective, and the spaces $\mathbb{Z}$ and $\mathbb{R}$ are not isometric.

Suppose X and Y are a pair of cube complexes and let d_1 be the edge metric. A map $f : X \to Y$ is *d_1-isometric* if for any pair of points $u, v \in X$, $d_1(f(u), f(v)) = d_1(u, v)$.

Note that metric space formed by the group G with finite generating set S with the word length metric is isometric to the Cayley graph $\Gamma(G; S)$ with the edge path metric.

Definition. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \to Y$ is a *quasi-isometry* if there exists some $k_1, k_2 \in \mathbb{R}$ with $k_1, k_2 > 0$ such that for all $u, v \in X$ $k_1 d_X(u, v) \leq d_Y(f(u), f(v)) \leq k_2 d_X(u, v)$.

If there is a quasi-isometry $f : X \to Y$ then we say X and Y are *quasi-isometric.*

Lemma 1.31. ([3], p.24) *If S and S' are finite generating sets for G then $\Gamma(G; S)$ is quasi-isometric to $\Gamma(G; S')$.*

Lemma 1.32. ([3], p.26) *Suppose that Γ is a finitely generated group and that $G \leq \Gamma$ is finite index. Then G is quasi-isometric to Γ.*

We say that the action of a group G on a space X is isometric if for any $g \in G$, $x, y \in X$ $d_X(x, y) = d_X(g(x), g(y))$. The action is quasi-isometric if there exists some k_1, k_2 such that $k_1 d_X(x, y) \leq d_X(g(x), g(y)) \leq k_2 d_X(x, y)$.

Note that the metric d_2 and the edge metric d_1 on a finite dimensional CAT(0) cube complex X are quasi-isometric, that is for any $x, y \in X^{(0)}$ there exists some k_1, k_2 such that $k_1 d_1(x, y) \leq d_2(x, y) \leq k_2 d_1(x, y)$. In fact, we can see that $k_2 = 1$ and $k_1 = \sqrt{n}$ where n is the maximum dimension of a cube in X.

1.4.2 Hyperplanes

Definition. [19] Consider the cube $[0, 1]^n$. Then for each $i \in \{1, 2, \dots, n\}$ there is a *midplane* M of $[0, 1]^n$ given by setting the coordinate x_i to be $\frac{1}{2}$, that is $M = \{(\alpha_1, \dots, \alpha_{i-1}, \frac{1}{2}, \alpha_{i+1}, \dots, \alpha_n) | \alpha_j \in [0, 1]\}$.

Given a CAT(0) cube complex X, take the set of all midplanes of cubes in X. Two midplanes M and N in X are said to be *hyperplane equivalent* if there is a sequence of midplanes $M = M_1, \dots, M_n = N$ such that $M_i \cap M_{i+1}$ is a midplane for all $i \in \{1, \dots, n-1\}$. We call such a sequence a *chain of midplanes*.

Given an equivalence class of midplanes, we build a complex by identifying the faces of midplanes in the class along their common intersections. We will use the term *hyperplane* for both the equivalence class of a midplane and the associated complex.

Lemma 1.33. [27] *The hyperplanes of a CAT(0) cube complex X are also CAT(0) cube complexes.*

Proof. By definition, any hyperplane $\mathfrak{h}$ in a CAT(0) cube complex is a connected cube complex. In order to show that $\mathfrak{h}$ is CAT(0), we will show that there is a global isometry from $\mathfrak{h}$ to X. Hence any triangle in $\mathfrak{h}$ is isometric to a triangle in X, and so $\mathfrak{h}$ satisfies the definition of a CAT(0) space.

The following proof is due to Niblo and Reeves in [27].

Let M be a midplane in $\mathfrak{h}$ and let C_M denote the unique lowest dimension cube of X containing M as a midplane. Let $\phi : \mathfrak{h} \to X$ denote the natural inclusion of $\mathfrak{h}$ in X, which maps cubes to midplanes. We show that for all $x \in \mathfrak{h}$ there is a neighbourhood U of x such that $\phi|_U : \mathfrak{h} \to X$ is an isometry. For any M in $\mathfrak{h}$, $\phi|_M$ is an isometry and so when x lies in a single cube M in $\mathfrak{h}$, we can choose as U any neighbourhood of x contained in M.

Suppose $x \in \mathfrak{h}$ lies in more than one cube of $\mathfrak{h}$. For each midplane $M \subset X$ let $\rho_M : C_M \to C_M$ denote reflection in M. Let $St_{\mathfrak{h}}(x)$ denote the set of cubes in $\mathfrak{h}$ containing x and $St_X(\phi(x))$ denote the set of cubes in X containing $\phi(x)$. We can define a map $\rho : St_X(\phi(x)) \to St_X(\phi(x))$ by setting $\rho|_{C_{\phi(M)}}$ to be $\rho_{\phi(M)}$ for each $M \in St_{\mathfrak{h}}(x)$. To see that this function is well defined we note that if M_1 and M_2 lie in a hyperplane and both contain x, then $\rho_{\phi(M_1)}$ and $\rho_{\phi(M_2)}$ agree on $C_{\phi(M_1 \cap M_2)}$. We also check that ρ is defined on every cube of $St_X(C_{\phi(x)})$, which follows from the fact that the minimal cube in X containing $\phi(x)$ is $C(M)$ where M is the minimal cube of $\mathfrak{h}$ containing x.

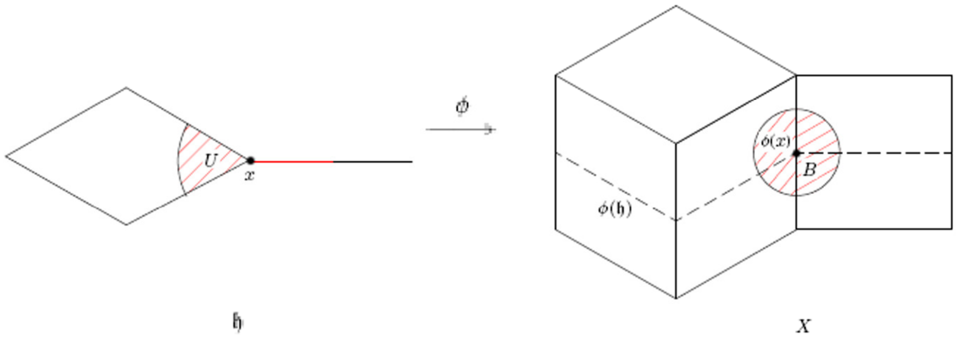

Figure 1.3: An isometry from a hyperplane to the CAT(0) cube complex

Given any $x \in \mathfrak{h}$, there exists a ball of B of radius $\epsilon > 0$ centered at $\phi(x)$ such that B is convex and is contained in $St_X(\phi(x))$ (see fingure 1.3. Hence the map ρ preserves B and is an isometry with fixed point set $\phi(St_{\mathfrak{h}}(x)) \cap B$. Since $\phi(St_{\mathfrak{h}}(x)) \cap B$ is the fixed set of an isometry, it is a totally geodesic subspace of X. Hence the set $U = \phi^{-1}(\phi(St_{\mathfrak{h}}(x)) \cap B)$ is a set containing x such that $\phi|_U : \mathfrak{h} \to X$ is an isometry. Hence for each $x \in \mathfrak{h}$ we can choose a neighbourhood U of x such that $\phi|_U : \mathfrak{h} \to X$ is an isometry, and so ϕ is a local isometry.

Hence since X is CAT(0), by the Cartan-Hadamard theorem (section 4 [20]), the local isometry ϕ must be a global isometry. Hence the cube complex $\mathfrak{h}$ is CAT(0).

□

Remark 1.34. If a group G acts on a CAT(0) cube complex X, then G acts on the set of hyperplanes of X. Since the action of G on X preserves incidence of cubes, it preserves incidence of midplanes and hence extends to an action on hyperplanes.

Each hyperplane $\mathfrak{h}$ divides X into two simply connected pieces called halfspaces (see [33] p.610-611). We denote the two halfspaces of X with respect to $\mathfrak{h}$ by $X_{\mathfrak{h}}$ and $X^*_{\mathfrak{h}}$. Each hyperplane is *geodesically convex*, that is for any hyperplane $\mathfrak{h}$ every geodesic with initial and end points both in $\mathfrak{h}$ is entirely contained in $\mathfrak{h}$ (see [27] p.624).

Suppose that $X_{\mathfrak{h}}$ is not geodesically convex. Then there exists a pair of points u, v in $X_{\mathfrak{h}}$ and a geodesic ℓ from u to v which has non-trivial intersection with $X \setminus X_{\mathfrak{h}}$. This geodesic must intersect the boundary $\mathfrak{h}$ of $X_{\mathfrak{h}}$ in at least two points. Let u' and v' be the points of $\mathfrak{h} \cap \ell$ closest to u and v respectively. Since $\mathfrak{h}$ is geodesically convex, any geodesic between u' and v' must be entirely contained in $\mathfrak{h}$, hence the intersection of ℓ with $X \setminus X_{\mathfrak{h}}$ is trivial. It follows that for every $\mathfrak{h}$ the halfspaces $X_{\mathfrak{h}}$ and $X^*_{\mathfrak{h}}$ are geodesically convex.

We say that two hyperplanes $\mathfrak{h}_1$ and $\mathfrak{h}_2$ *intersect* if each of the sets

$$X_{\mathfrak{h}_1} \cap X_{\mathfrak{h}_2},\ X^*_{\mathfrak{h}_1} \cap X_{\mathfrak{h}_2},\ X_{\mathfrak{h}_1} \cap X^*_{\mathfrak{h}_2},\ X^*_{\mathfrak{h}_1} \cap X^*_{\mathfrak{h}_2},$$

is non empty. Equivalently, $\mathfrak{h}_1$ and $\mathfrak{h}_2$ intersect if and only if $\mathfrak{h}_1 \neq \mathfrak{h}_2$ and there exists midplanes M_1 of $\mathfrak{h}_1$ and M_2 of $\mathfrak{h}_2$ such that $M_1 \cap M_2 \neq \emptyset$.

Suppose there exists a pair of vertices u, v in X which lie in the same set of halfspaces. Since X is connected, there exists a shortest edge path from u to v. Since each edge contains a midpoint which is a midplane in some hyperplane, the endpoints of any edge lie in different halfspaces. Since halfspace are geodesically convex, the geodesic from u to v crosses between halfspaces $X_{\mathfrak{h}}$ and $X^*_{\mathfrak{h}}$ at most once for each hyperplane $\mathfrak{h}$,. Hence any geodesic of length > 1 has endpoints lying in different halfspaces for at least one hyperplane, and so $d(u, v) = 0$, that is $u = v$. Hence every vertex in X is uniquely determined by the set of halfspaces containing it.

We say a hyperplane $\mathfrak{h}$ *separates* the vertices u and v if u lies in $X_{\mathfrak{h}}$ and v lies in $X^*_{\mathfrak{h}}$, or vice versa. We say the edge e and the hyperplane $\mathfrak{h}$ intersect if the midpoint of e lies in a midplane in the equivalence class of midplanes $\mathfrak{h}$. Let u and v be a pair of vertices in X and let p be a geodesic edge path between them. For each edge in the path p, p intersects a single hyperplane in X. Since halfspaces in X are geodesically convex, the path p intersects each hyperplane at most once and the length of the path p is equal to the number of hyperplanes intersecting edges in p. Using the fact that halfspaces are geodesically convex for a second time, we note that if the path p intersects a hyperplane $\mathfrak{h}$, then u and v lie in different halfspaces $X_{\mathfrak{h}}$ and $X^*_{\mathfrak{h}}$. Hence the length of any geodesic path from u to v, denoted by $d_1(u, v)$, is equal to the number of hyperplanes separating u and v.

For any hyperplane $\mathfrak{h}$ in X, let $\mathrm{stab}_G(\mathfrak{h})$ denote the set of elements of G which act on X to map $\mathfrak{h}$ to itself.

1.4.3 Almost Residual Finiteness

Definition. A group G is *almost residually finite* if there exists a finite subgroup H of G such that H is separable.

Lemma 1.35. *Let G be a group and suppose that G has a proper action on a CAT(0) cube complex X such that for every hyperplane $\mathfrak{h}$ in X $\mathrm{stab}_G(\mathfrak{h})$ is separable. Then G is almost residually finite.*

Proof. Choose a vertex v in X. Let $\mathcal{H}$ denote the set of hyperplanes in X and let $\overline{\mathcal{H}}$ be the set of hyperplanes adjacent to v, that is hyperplanes which intersect an edge incident with v. Since X is locally finite, the set $\overline{\mathcal{H}}$ is finite.

For some choice of halfspace $X_{\mathfrak{h}}^{(*)}$ in $\{X_{\mathfrak{h}}, X_{\mathfrak{h}}^*\}$ for each $\mathfrak{h}$, v is the unique vertex contained in the intersection $\cap_{\mathfrak{h}\in\mathcal{H}} X_{\mathfrak{h}}^{(*)}$. Since the set $\overline{\mathcal{H}}$ contains all hyperplanes adjacent to v, v is the unique vertex contained in $\bigcap_{\mathfrak{h}\in\overline{\mathcal{H}}} X_{\mathfrak{h}}^{(*)}$.

Consider the group $I = \bigcap_{\mathfrak{h}\in\overline{\mathcal{H}}} \mathrm{stab}_G(\mathfrak{h})$. Since each subgroup $\mathrm{stab}_G(\mathfrak{h})$ is separable, the group I is separable. Then I acts on X and preserves each hyperplane in the set $\overline{\mathcal{H}}$, and so we can say that I acts on the set of components of $X \setminus \overline{\mathcal{H}}$.

Since the set $\overline{\mathcal{H}}$ is finite and each $\mathfrak{h} \in \overline{\mathcal{H}}$ divides X into two components, it follows that $X \setminus \overline{\mathcal{H}}$ has only finitely many components, say n. Then I permutes the components of $X \setminus \{\overline{\mathcal{H}}\}$ and is homomorphic to a subgroup of the symmetric group S_n. Since S_n is finite, it follows that for each component of $X\setminus\overline{\mathcal{H}}$, the subgroup of I preserving that component is finite index in I. Let H_v be the subgroup fixing the component containing v. Since v is uniquely determined by the set $X_{\mathfrak{h}}^{(*)}$ of halfspaces containing it, it follows that H_v fixes v.

Since the action of G on X is proper, $\{g \in G | gv = v\}$ is finite. H_v is a subgroup of this set and hence if finite. H_v is a finite index subgroup of the separable group I, so H_v is also separable. Hence G is almost residually finite. $\square$

1.4.4 Coxeter groups act on CAT(0) cube complexes

The following construction is due to Niblo and Reeves and can be found in [28].

Definition. ([28], p.2) We consider the triple $(H, \leq, *)$ where H is a set, $\leq$ is a partial order on H and $*$ is an order reversing involution on H, denoted by $X \mapsto X^*$.

Suppose $(H, \leq, *)$ satisfies the following conditions:

(1) Given any two elements X_1 and X_2 of H, there exist only finitely many elements $X_3 \in H$ such that $X_1 \leq X_3 \leq X_2$.

(2) Given any pair of elements X_1 and X_2 of H, at most one of the following holds:
$$X_1 \leq X_2,\ X_1 \leq X_2^*,\ X_1^* \leq X_2,\ X_1^* \leq X_2^*$$

Then $(H, \leq, *)$ is a *halfspace system*. The elements of the set H are called *halfspaces*.

A pair of halfspaces X_1 and X_2 in H are said to be *transverse* if none of the conditions in part (2) of the definition above hold.

For any set of halfspaces $K \subset H$ a halfspace $X \in K$ is said to be a *minimal halfspace in the set* K if for every $X' \in K$, $X' \nleq X$.

For a general halfspace system $(H, \leq, *)$, we define an equivalence relation $\sim$ on H by $X_1 \sim X_2$ if and only if $X_1 = X_2$ or $X_1 = X_2^*$. We denote the equivalence class containing X by $[X] = \{X, X^*\}$. The boundary map ∂ is a map $\partial : H \to H/\sim$ defined by $X \mapsto [X]$. We call the equivalence class $[X]$ the *boundary* of X.

A pair of boundaries $[X_1], [X_2]$ are said to *intersect* if the halfspaces X_1 and X_2 are transverse.

Given a CAT(0) cube complex, there is an obvious triple given by the set of halfspaces H associated to the set of hyperplanes in the complex with the partial order $\leq$ induced by inclusion and with the order reversing involution $*$.

Lemma 1.36. *Let X be a CAT(0) cube complex. Then the associated triple $(H, \leq, *)$ is a halfspace system*

Proof. Given any pair of halfspaces $X_{\mathfrak{h}_1}$ and $X_{\mathfrak{h}_2}$, let $I(X_{\mathfrak{h}_1}, X_{\mathfrak{h}_2})$ denote the set $\{X_{\mathfrak{h}_3} \in H | X_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_3} \leq X_{\mathfrak{h}_2}\}$. If $X_{\mathfrak{h}_1} \not\leq X_{\mathfrak{h}_2}$ then $I(X_{\mathfrak{h}_1}, X_{\mathfrak{h}_2})$ is empty. Suppose $X_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_2}$, then there is a finite length edge path p in the cube complex X joining the hyperplanes $\mathfrak{h}_1$ and $\mathfrak{h}_2$, that is a path whose initial and final edges have midpoints contained in midplanes in the hyperplanes $\mathfrak{h}_1$ and $\mathfrak{h}_2$ respectively.

If $X_{\mathfrak{h}_3}$ satisfies $X_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_3} \leq X_{\mathfrak{h}_2}$ then $\mathfrak{h}_3$ must cross the path p, that is some edge of p must have midpoint in $\mathfrak{h}_3$. The midpoint of each edge in p is in exactly one hyperplane. Since p has finite length, there must be a finite number of $\mathfrak{h}_3$ crossing p and hence finitely many $X_{\mathfrak{h}_3}$ with $X_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_3} \leq X_{ch_2}$.

Given any pair of halfspaces $X_{\mathfrak{h}_1}$ and $X_{\mathfrak{h}_2}$ with $\mathfrak{h}_1 \neq \mathfrak{h}_2$, we show that at most one of the inequalities listed in part 2 of the definition of a halfspace system. Suppose $X_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_2}$, and consider in turn each of the other inequalities.

- If $X_{\mathfrak{h}_1} \leq X^*_{\mathfrak{h}_2}$, then $X_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_2} \cap X^*_{\mathfrak{h}_2} = \emptyset$.
- If $X^*_{\mathfrak{h}_1} \leq X_{\mathfrak{h}_2}$ then $X_{\mathfrak{h}_1} \cup X^*_{\mathfrak{h}_1} = X \leq X_{\mathfrak{h}_2}$.
- If $X^*_{\mathfrak{h}_1} \leq X^*_{\mathfrak{h}_2}$ then, applying the order reversing involution $*$, $X_{\mathfrak{h}_2} \leq X_{\mathfrak{h}_1}$ and we have $X_{\mathfrak{h}_1} = X_{\mathfrak{h}_2}$.

In each case we have a contradiction. Similarly, taking any of $X_1 \leq X_2^*$, $X_1^* \leq X_2$, $X_1^* \leq X_2^*$ to be true, we show that assuming any other inequality in the list to also be true leads to a contradiction. □

We will consider a vertex of the cube complex X as defined by the set of halfspaces containing it, and examine the properties of these sets of halfspaces. Each vertex will lie in exactly one of X_1, X_1^* for every X_1. Suppose a halfspace X_1 is contained in a halfspace X_2. If the vertex v in X lies in

X_1 then it also lies in X_2. Two vertices u, v in X which are joined by an edge will be separated by a single hyperplanes $\mathfrak{h}$, the hyperplane equivalence class containing the midplane of the edge joiing u and v. Since u and v are separated by a single hyperplane, the set of halfspaces containing u and v will be identical with the exception of u being contained in $X_{\mathfrak{h}}$ and v being contained in $X_{\mathfrak{h}}^*$. We note that for a halfspace system arising from hyperplanes in a CAT(0) cube complex the definition of intersection of boundaries in a halfspace system is equivalent to the definition of intersection of hyperplanes given in section 1.3.2.

Lemma 1.37. ([28],p.3) *Given any halfspace system $(H, \subseteq, *)$, we can construct a CAT(0) cube complex whose hyperplanes form the halfspace system $(H, \subseteq, *)$*

In the following paragraphs, we outline the construction of a CAT(0) cube complex from a halfspace system. Our definitions will be motivated by the properties of sets of halfspaces in cube complexes containing vertices, discussed above.

Definition. A map $v : H/\sim \longrightarrow H$ is a *section* for ∂ if $\partial v([X]) = [X],\ \forall X$.

A section is interpreted as an orientation on a boundary. We view the section v as a list of the halfspaces in which a vertex lies. We saw that if $X_1 \leq X_2$ and the vertex v in X lies in X_1 then v should also lie in X_2. Hence for any pair of hyperplanes X_1 and X_2 $v([X_1]) \nleq v([X_2])^*$ (see figure 1.4). We take the set of vertices for the CAT(0) cube complex to be all sections v for ∂ such that $v([X_1]) \nleq v([X_2])^*$ for all $X_1, X_2 \in H$.

We say two vertices u and v are *separated* by the boundary $[X]$ if $u([X]) \neq v([X])$. We join two vertices u and v by an edge if and only if u and v are separated by exactly one boundary $[X]$. Not every halfspace $u([X])$ can be replaced by its complement to give a new vertex. Suppose two vertices u and v are joined by an edge. Then for some element $[X]$ of $H/\sim$, $u([X]) = v([X])^*$. By the definition of a vertex, for any halfspace Y we have $u([Y]) = v([Y]) \nleq v([X])^* = u([X])$. Hence $u([X])$ must be minimal among the set of halfspaces in the image of $H/\sim$ under u. The set of vertices adjacent to

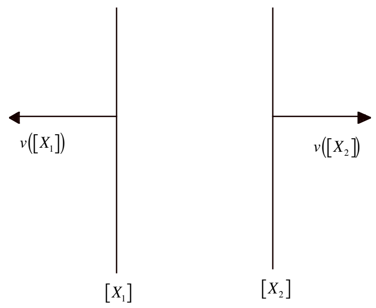

Figure 1.4: Vertices must satisfy $v([X_1]) \nleq v([X_2])^*$

u are those obtained by replacing $u([X])$ by $u([X])^*$ where $[X]$ is a minimal halfspace in $u(H/\sim)$.

Now consider conditions necessary for a set of vertices to be the vertices of a square. If the vertices u and v are to lie at opposite corners of the square, then the values of u and v must differ on precisely two elements of $H/\sim$, call them $[X_i],\ i \in \{1,2\}$. For $i = 1,2$ we must have $u([X_i])$ minimal in $u(H/\sim)$.

The pair u, v has the following properties:

$$\begin{aligned}
u([X_1]) &\nleq u([X_2]) && \text{by the minimality of } u([X_1]) \\
u([X_1])^* &\nleq u([X_2]) && \text{by the minimality of } u([X_2]) \\
u([X_1]) &\nleq u([X_2])^* && \text{since } u \text{ is a vertex} \\
v([X_1]) = u([X_1])^* &\nleq u([X_2]) = v([X_2])^* && \text{since } v \text{ is a vertex}
\end{aligned}$$

So $u([X_1])$ and $u([X_2])$ are transverse. For $i = 1,2$ we replace either or both of the $u([X_i])$ with $u([X_i])^*$ to get the remaining three vertices of the square.

In general, if we have a set $[X_1], [X_2], \ldots, [X_n]$ of boundaries such that each $u([X_i])$ is minimal in $u(H/\sim)$ and any pair $u([X_i]), u([X_j])$ are transverse, then by replacing the $u([X_i])$ with $u([X_j])^*$, we form the vertices of an n-dimensional cube. We construct a cube complex by filling in the cubes which occur in this way.

The cube complex constructed in this way may have more than one component. For example, consider the halfspace system $(\mathcal{H}, \leq, *)$ where

$\mathcal{H} = \{\mathcal{H}_i, \mathcal{H}_i^* | i \in \mathbb{Z}\}$ and $\mathcal{H}_i \leq \mathcal{H}_j$ if and only if $i \leq j$. Let $\mathfrak{h}_i = [\mathcal{H}_i, \mathcal{H}_i^*]$. The CAT(0) cube complex constructed from this halfspace system as defined by lemma 1.37 is isometric to the real line with additional points $\{+\infty, -\infty\}$. This cube complex has three connected components, $\{-\infty\}, \mathbb{R}$ and $\{+\infty\}$. The vertex corresponding to a integer point z can be defined by the section

$$z(\mathfrak{h}_i) = \begin{cases} \mathcal{H}_i & \text{if } i < z \\ \mathcal{H}_i^* & \text{if } i \geq z. \end{cases}$$

The points ∞ and $-\infty$ are defined by the sections

$$\begin{aligned} -\infty(\mathfrak{h}_i) &= \mathcal{H}_i \quad \text{for all } i \in \mathbb{Z} \\ +\infty(\mathfrak{h}_i) &= \mathcal{H}_i^* \quad \text{for all } i \in \mathbb{Z} \end{aligned}$$

To see that the cube complex has three connected components, suppose ∞ and $\mathbb{R}$ lie in the same component. Then, by the definition of a connected set, for any decomposition of $\infty \cup \mathbb{R}$ as two subsets A and B, at least one of the sets $cl(A) \cap B$, $A \cap cl(B)$ is non-empty.

Let

$$A = \bigcap_{i \in \mathbb{Z}} X^*_{\mathfrak{h}_i}$$

and

$$B = X - A = X - \bigcap_{i \in \mathbb{Z}} X^*_{\mathfrak{h}_i} = \bigcup_{i \in \mathbb{Z}} X_{\mathfrak{h}_i}.$$

Now B is a union of open sets, and hence is open. Hence A is closed and $cl(A) \cap B = A \cap B$. Any point $x \in B$ lies inside $X_{\mathfrak{h}_j}$ for some j, and hence $x \notin X^*_{\mathfrak{h}_j}$ and $x \notin \bigcap_{i \in \mathbb{Z}} X^*_{\mathfrak{h}_i} = A$. Hence $cl(A) \cap B = \emptyset$.

We now consider the closure of B, $cl(B) = B \cup \partial B \subset \bigcup_{i \in \mathbb{Z}} X_{\mathfrak{h}_i} \cup \bigcup_{i \in Z} \partial X_{\mathfrak{h}_i}$. For any $i \in \mathbb{Z}$, $(X_{\mathfrak{h}_i} \cup \partial X_{\mathfrak{h}_i}) \subset X_{\mathfrak{h}_{i+1}}$ and hence $(X_{\mathfrak{h}_i} \cup \partial X_{\mathfrak{h}_i}) \cap X^*_{\mathfrak{h}_{i+1}} = \emptyset$. Suppose $cl(B) \cap A \neq \emptyset$. Then for some $j \in \mathbb{Z}$

$$(X_{\mathfrak{h}_j} \cup \partial X_{\mathfrak{h}_j}) \cap \left(\bigcap_{i \in \mathbb{Z}} X^*_{\mathfrak{h}_i} \right) \neq \emptyset$$

Hence for each $k \in \mathbb{Z}$

$$(X_{\mathfrak{h}_j} \cup \partial X_{\mathfrak{h}_j}) \cap X_{\mathfrak{h}_k} \neq \emptyset$$

which is false for $k = j+1$. Hence $cl(B) \cap A = \emptyset$ and $+\infty$ and $\mathbb{R}$ are separate components of the cube complex.

In a similar fashion, we can construct a halfspace system corresponding to the Euclidean plane. The cube complex constructed using this halfspace system has 9 components, $\mathbb{E}^2$, four points corresponding to $(\pm\infty, \pm\infty)$ and four copies of the real line, corresponding to $(\pm\infty, \mathbb{R})$ and $(\mathbb{R}, \pm\infty)$.

For a proof that the components of the space we just constructed are CAT(0) cube complexes, see [33].

Theorem 1.38. ([28],p.1) *If W is a finitely generated Coxeter group then there exists a locally finite, finite dimensional CAT(0) cube complex X on which W acts properly discontinuously by isometries, and in which there is an isometric embedding of W.*

Given a finitely generated Coxeter group W, we construct a cube complex X as follows: We begin by defining a halfspace system for the Coxeter group W. We then use the general method of constructing a cube complex given in lemma 1.37 on this halfspace system. Finally, we show that if a halfspace system is constructed from a Coxeter group W then W embeds quasi-isometrically in a component of the corresponding cube complex.

There are three ways of defining a halfspace system for a Coxeter group W. We will describe two of them. For the third definition of a halfspace system associated to the Coxeter group W, see section 2.1 of [28]. We choose a generating set S for W.

Definition. Given a Coxeter system (W, S) construct the Coxeter complex $\Sigma(W, S)$. Define the set of halfspaces H_W to be the half-apartments of $\Sigma(W, S)$, $H_W = \{X_r,\ X_r^* | r \text{ is a reflection in } W\}$. Then the triple $(H_W, \subseteq, *)$ where $\subseteq$ is inclusion of half-apartments is a halfspace system. In this case the map ∂ is equivalent to the map $X_r \mapsto \mathcal{H}_r$ which takes the halfspace corresponding to the reflection r to the wall corresponding to r.

Definition. Let Γ_W denote the Cayley graph of (W, S). Let u, v be adjacent vertices in Γ_W and define $H(u, v) = \{w \in W | d(w, u) < d(w, v)\}$. Since u and v are adjacent, $u = vs$ for some $s \in S$. Recalling the definition of the metric on Γ_W we have $d(g, h) = \ell(g^{-1}h)$, where $\ell(\gamma)$ is the minimum length of a word for γ in the generators. By corollary 1.2

$$d(w, u) = \ell(w^{-1}u) = \ell(w^{-1}vs) \neq \ell(w^{-1}v) = d(w, v)$$

and hence $H(u, v) \cap H(v, u)$ is empty. For a fixed u and v, W is the disjoint union of $H(u, v)$ and $H(v, u)$. Let $H'_W = \{H(u, v) | u, v \in W, d(u, v) = 1\}$ be the collection of all such sets, and denote by $* : H'_W \to H'_W$ the involution which sends $H(u, v)$ to $H(v, u)$. Let $\leq$ be the natural order given by inclusion as subsets. Then $(H_W, \leq, *)$ is a halfspace system.

Lemma 1.39. *For any Coxeter group W the halfspace systems $(H_W, \subseteq, *)$ and $(H'_W, \leq, *)$ are isomorphic as partially ordered sets.*

Proof. For the proof of this lemma, see proposition 1 of [28] and proposition 2.6 of [32]. □

We have seen that the cube complex X corresponding to a halfspace system may have more than one component. We define an embedding of W into a single component of X as follows:

For each edge (u, v) in Γ_W, the pair $H(u, v), H(v, u)$ represents a subdivision of Γ_W into two components, with exactly one of the two halfspaces $H(u, v), H(v, u)$ containing the vertex e. We define a section v_e to ∂ by setting $v_e([H(u, v), H(v, u)])$ to be the halfspace containing e. By definition, the intersection of any two halfspaces in the image of v_e contains the vertex e and so is non-empty. Hence for any pair of halfspaces $v_e(H(u, v))$ and $v_e(H(x, y))$ in the image, $v_e(H(u, v)) \nleq v_e(H(x, y))^*$ and $v_e(H(x, y)) \nleq v_e(H(u, v))^*$. Hence by definition v_e is a vertex in the CAT(0) cube complex..

Similarly, for any $g \in \Gamma_W$ a section v_g is defined by choosing the halfspace of Γ_W containing the vertex g for each pair $H(u, v), H(v, u)$ where $\{u, v\}$ is an edge.

Vertices of Γ_W are labelled by elements of W by the construction of Γ_W. By the definition of the metric on W, the embedding of W in Γ_W defined by this labelling is an isometry. Hence it is sufficient to show that Γ_W embeds isometrically in the cube complex. We define the embedding to be the map which takes g to v_g for all $g \in \Gamma_W$

Let g, h be a pair of vertices in Γ_W and let $g = \gamma_0, \gamma_1, \ldots, \gamma_n = h$ be a geodesic from g to h in Γ_W. Then for each edge $\{\gamma_i, \gamma_{i+1}\}$ in the geodesic, $g \in H(\gamma_i, \gamma_{i+1})$ and $h \in H(\gamma_{i+1}, \gamma_i)$. Hence the value of the sections v_g and v_h differ on at least n boundaries and the distance $d_1(v_g, v_h)$ between v_g and v_h in X satisfies $d_1(v_g, v_h) \geq d_W(g, h)$. In fact, there is a bijection from the set of boundaries in Γ_W which separate g and h to the set of hyperplanes in X which separate v_g and v_h. Hence we have $d_1(v_g, v_h) = d_{\Gamma_W}(g, h)$ and so the map taking g to v_g for all $g \in W$ is an isometry from (W, d_W) to (X, d_1).

See [28] for the proof that X is locally finite (p.9) and finite dimensional (p.8), and for the proof that W acts properly discontinuously (by isometries) on X (p.9).

We note that the CAT(0) cube complex constructed in this way has the property that for any hyperplane $\mathfrak{h}$ in the complex, $\mathrm{stab}_W(\mathfrak{h}) = \mathrm{stab}_W(\mathcal{H}_r)$ for some $\mathcal{H}_r$ in the Coxeter complex $\Sigma(W, S)$.

The action of W on X as defined by Niblo and Reeves is not necessarily cocompact. Williams gives the following theorem.

Theorem 1.40. [40] *Let W be a Coxeter group acting on a CAT(0) cube complex X as defined in the proof of theorem 1.38. Then the action of W on X is cocompact if and only if for any triple p, q, r of positive integers, W contains only finitely many conjugacy classes of subgroups isomorphic to the p, q, r triangle group $< a, b, c | a^2 = b^2 = c^2 = (ab)^p = (bc)^q = (ac)^r = 1 >$.*

Niblo and Reeves conjectured that the action of W on X will be cocompact if and only if W contains no subgroups isomorphic to Euclidean triangle groups, that is to either $\triangle(2, 3, 6)$, $\triangle(2, 4, 4)$ or $\triangle(3, 3, 3)$. This result was later proved by Caprace and Mühlherr in [7].

Lemma 1.41. *([7], p.468) Let (W, S) be a Coxeter system of finite rank. The following statements are equivalent:*

(i) there are only finitely many conjugacy classes of reflection triangles,

(ii) the Coxeter diagram of (W, S) has no irreducible affine subdiagrams of rank ≥ 3.

Considering the Coxeter diagrams of the affine groups (see for example [38]), we see that the irreducible affine diagrams of rank 3 are those corresponding to the groups $\triangle(2,3,6)$, $\triangle(2,4,4)$ and $\triangle(3,3,3)$. One of these diagrams appears as a subdiagrams of the Coxeter diagram for a group (W, S) if and only if W contains a subgroup isomorphic to the corresponding Euclidean triangle group.

1.4.5 Actions and maps on CAT(0) cube complexes

Lemma 1.42. *Let G be a finitely generated group which acts freely, isometrically, cocompactly and properly on a CAT(0) cube complex X. Then there is a quasi-isometric map from G to X.*

Proof. Let S be a generating set for G. Choose any point $x \in X$ and use the action of G on X to define a map p from G to X by setting $p(g) = g(x)$ for all $g \in G$. Denote by k_1 the maximum distance between x and $s(x)$ for any $s \in S$, that is $k_1 = \max\{d_X(x, s(x)) | s \in S\}$. Both the map p and the specific value of k_1 are dependant on our choice of x. Since the action of G on X is isometric, it follows that for any $g \in G$ and $s \in S$, $d_X(g(x), (gs)(x)) \leq k_1$. Given any $g, h \in G$ there is a word $s_1 \dots s_n,\ s_i \in S$ representing $g^{-1}h$ such that $n = d_G(g, h)$. It follows that there is a path in X from $p(g)$ to $p(h)$ defined by the vertices $g(x), gs_1(x), \dots, gs_1 \dots s_n(x) = (gg^{-1}h)(x) = h(x)$. Since $d_X(g(x), gs(x)) \leq k_1$ for all $g \in G, s \in S$ it follows that $d_X(p(g), p(h)) \leq k_1 n = k_1 d_G(g, h)$.

In order to prove the map p is quasi-isometric it remains to show that there exists some k_2 such that $k_2 d_G(g, h) \leq d_X(p(g), p(h))$ holds for any pair g, h. We consider an alternative generating set for the group G. Since the action of G on X is cocompact we can choose a compact region $C \subset X$ containing x such that $GC = X$. Then there is some $r \in \mathbb{R}$, the diameter of the fundamental region C, such that for every $y \in X$ $\exists g \in G$ such that $d_X(y, g(x)) \leq r$.

Let $k = max\{2r+1, k_1\}$. Let $S' = \{g \in G \setminus \{1\} | d_X(x, g(x)) \leq k\}$. Then $S \subset S'$ and S' is a generating set for G. We construct the graph Δ with vertex set $V(\Delta) = G$ by connecting $g, h \in G$ by an edge if $d_X(g(x), h(x)) \leq k$. Then Δ is the Cayley graph of G with generating set S'. By lemma 1.31 the metric d_G on the group G with generating set S is quasi-isometric to the metric d_Δ on the Cayley graph for G with generating set S'. Hence it is sufficient to prove that, for all $g, h \in G$, $k_2 d_\Delta(g, h) \leq d_X(p(g), p(h))$.

Given any $g, h \in G$, let $\alpha \subset X$ be a geodesic connecting $g(x)$ to $h(x)$. Choose a sequence of points $g(x) = x_0, x_1, \ldots, x_n = h(x)$ along α, such that $d(x_i, x_{i+1}) \leq 1$ for all i and such that $d_X(g(x), h(x)) > n - 1$.

For each x_i, we can choose some g_i such that $d(g_i(x), x_i) < r$. Hence we can construct a path $g(x) = g_0(x), g_1(x), \ldots, g_n(x) = h(x)$ in X such that $d_X(g_i(x), g_{i+1}(x)) \leq r + 1 + r \leq k$ for all i. Hence by the definition of Δ for each i we have $d_\Delta(g_i, g_{i+1}) = 1$, and there is a path in Δ from g to h with length n. Hence $d_\Delta(g, h) \leq n \leq d_X(g(x), h(x)) + 1$.

Since the action of G on X is free and properly discontinuous, we can choose some $k' > 0$ in $\mathbb{R}$ such that for all $g, h \in X$ $d_X(g(x), h(x)) > k'$. Then

$$\begin{aligned} \frac{d_X(g(x), h(x))}{k'} &> 1 \\ \Rightarrow \frac{k'+1}{k'} d_X(g(x), h(x)) &> d_X(g(x), h(x)) + 1 \\ &> d_\Delta(g, h) \end{aligned}$$

and the map $p : G \to X$ is a quasi-isometric embedding.

□

We can define a quasi-inverse to p, that is a quasi-isometric map $q : X \to G$ such that $p \circ q$ and $q \circ p$ are a bounded distance from the identity maps. We have seen that for each point $y \in X$ there exists a $g \in G$ such that $d_X(y, g(x)) \leq r$. Let q be a map which takes each y to some element g of G satisfying $d_X(y, g(x)) \leq r$.

Consider $p \circ q(y)$. We have defined $q(y)$ to be some $g \in G$ such that $d_X(y, g(x)) \leq r$ and hence $d_X(y, p \circ q(y)) = d_X(y, g(x)) \leq r$.

Let C' be the compact subset of X defined by $C' = \{y | d_X(y, g(x)) \leq r\}$. Since the action of G on X is properly discontinuous, there are finitely many $g_i \in G$ such that $g_i C' \cap C' \neq \emptyset$. Let $m = max\{d(e, g_i) | g_i C' \cap C' \neq \emptyset\}$. Consider $q \circ p(g) = q(g(x))$. Then $q \circ p(g)$ is some $g' \in G$ for which $d_X(g'(x), g(x)) < r$. Since the action of G on X is isometric $d_X(g^{-1}g'(x), x) < r$ and hence $g^{-1}g' \in \{g_i | g_i C' \cap C' \neq \emptyset\}$ and $\ell(g^{-1}g') = d_G(g, g') \leq m$.

1.4.6 Products of CAT(0) cube complexes

Definition. The *Cartesian product* of a pair of sets X and Y is the set of all possible ordered pairs whose first component is a member of X and whose second component is a member of Y

$$X \times Y = \{(x, y) | x \in X, y \in Y\}$$

Definition. The product a pair of cube complexes X and Y, denoted $X \times Y$ is the cartesian product $X \times Y$ together with the cubical structure inherited from X and Y, that is each pair of cubes C_X in X and C_Y in Y with dimensions m and n respectively gives rise to a $(m+n)$-cube $C_X \times C_Y$ in $X \times Y$, $C_X \times C_Y = \{(x, y) | x \in C_x, y \in C_y\}$.

We wish to show that the product of a pair of CAT(0) cube complexes is $CAT(0)$. We will do this by considering the links of the vertices in the product.

Lemma 1.43. *A CAT(0) cube complex X is CAT(0) if X is simply connected and for every vertex v in X the link $Lk_X(v)$ is a CAT(1) space.*

Proof. See definition 5.1, theorem 5.2 and theorem 5.4 of [4]. □

Lemma 1.44. *$Lk(v)$ is CAT(1) if and only if every pair x and y of point in $Lk(v)$ with $d_{Lk(v)}(u, v) \leq \pi$, x and y are joined by at most one geodesic.*

Proof. See 4.2.B of [20]. Note that $Lk_X(v)$ is isomorphic to the intersection of the boundary of the ball $B(v, 1)$ with X and hence has curvature 1. □

Definition. ([4], p.63) Let (X, d_X) and (Y, d_Y) be two metric spaces. As a set, their *spherical join* $X * Y$ is $[0, \frac{\pi}{2}] \times X \times Y$ modulo the equivalence relation $\sim$ where $(\theta, x, y) \sim (\theta', x', y')$ whenever

- $\theta = \theta' = 0$ and $x = x'$ or
- $\theta = \theta' = \frac{\pi}{2}$ and $y = y'$ or
- $\theta = \theta' \notin \{0, \frac{\pi}{2}\}$ and $x = x', y = y'$.

We define a metric d on $X * Y$ by requiring that the distance between the points $u = (\theta, x, y)$ and $u' = (\theta', x', y')$ be at most π and that d satisfy the formula $\cos(d(u, u')) = \cos\theta \cos\theta' \cos(d_X(x, x')) + \sin\theta \sin\theta' \cos(d_Y(y, y'))$.

We think of $X * Y$ as the product of $X \times Y$ with the interval $[0, \frac{\pi}{2}]$, identifying points to 'collapse' the ends so that $(0, X, Y) = \{(0, x, y) | x \in X, y \in Y\}$ is isometric to X and $(\frac{\pi}{2}, X, Y) = \{(\frac{\pi}{2}, x, y) | x \in X, y \in Y\}$ is isometric to Y.

Lemma 1.45. ([4], p.284) *Let X and Y be two complete CAT(0) spaces. Then $\partial(X \times Y)$ with the angular metric $\angle$ (as defined in definition 23 of [14]) is isometric to the spherical join $\partial X * \partial Y$ of $(\partial X, \angle)$ and $(\partial Y, \angle)$. More specifically, given $\psi = (\theta, x, y)$ and $\psi' = (\theta', x', y')$ in $\partial(X \times Y)$, we have*

$$\cos(\angle(\psi, \psi')) = \cos\theta \cos\theta' \cos(\angle(x, x')) + \sin\theta \sin\theta' \cos(\angle(y, y'))$$

Lemma 1.46. *Let X and Y be cube complexes and let $u \in X$, $v \in Y$ be vertices of X and Y respectively. Then the geometric link of the vertex $u \times v$ in $X \times Y$ satisfies $Lk_{X \times Y}(u \times v) = Lk_X(u) * Lk_Y(v)$*

Proof. The link of a vertex in a CAT(0) cube complex is isometric to the boundary of the ball of radius 1 centered at that vertex with the angular metric. Given a pair of vertices u, v in X, Y respectively, let X' denote the sphere of radius 1 in X centered at u and let Y' denote the sphere of radius 1 in Y centered at v. Then the sphere of radius 1 in $X \times Y$ centered at (u, v)

is isometric to the direct product $X' \times Y'$. Hence by lemma 1.45

$$Lk_{X\times Y}((u,v)) = \partial(X' \times Y') = \partial(X') * \partial(Y') = Lk_X(u) * Lk_Y(v).$$

$\square$

Lemma 1.47. *The product of a pair of CAT(0) cube complexes is a CAT(0) cube complex.*

Proof. Let X and Y be a pair of CAT(0) cube complexes. By definition, the product $X \times Y$ is a cube complex. By lemma 1.43, $X \times Y$ is CAT(0) if and only if for every vertex v in $X \times Y$ $Lk_{X\times Y}(v)$ satisfies the conditions in lemma 1.44. Every vertex v in x is of the form (x, y) for some vertex x in X and some vertex y in Y. By lemma 1.46, $Lk_{X\times Y}(v) = Lk_X(x) * Lk_Y(y)$. By the definition of the metric on the spherical product, any geodesic between two points has length at most $\frac{\pi}{2}$, and so $Lk_{X\times Y}(v)$ is CAT(1) if and only if no two points in $Lk_{X\times Y}(v)$ are joined by more than one geodesic, that is if and only if $Lk_{X\times Y}(v)$ is simply connected.

Suppose ℓ is a non-trivial loop in $Lk_{X\times Y}(v)$, then we can homotope the loop ℓ to a loop contained in the subspace $(0, Lk_X(x), Lk_Y(y))$ by a continuous change in the first coordinate. This loop is then homotopic to any loop in $(0, Lk_X(x), t)$ for fixed t, since points $(0, s, t)$ and $(0, s', t')$ are identified under the equivalence if $s = s'$. We can then homotope to a loop in the subspace $(\frac{\pi}{2}, Lk_X(x), Lk_Y(y))$, which since t is fixed and $(0, s, t)$ and $(0, s', t')$ are identified under the equivalence if $t = t'$ is a point . Hence the product of a pair of CAT(0) cube complexes is a CAT(0) cube complex. $\square$

A product of trees $T_1 \times T_2 \times \ldots \times T_n$ is a CAT(0) cube complex of dimension n. The product of CAT(0) cube complexes X_1 and X_2 of dimensions n and m will have dimension $n+m$. We consider the d_1 metric on products of CAT(0) cube complexes. This choice of metric has the useful property that for any pair of vertices $u = (u_1, u_2 \ldots, u_n)$ and $v = (v_1, v_2, \ldots, v_n)$ in $X_1 \times \ldots \times X_n$ we have $d_1(u, v) = \sum_{i=1}^{n} d_1(u_i, v_i)$, where $d_1(u_i, v_i)$ is the distance between u_i and v_i in X_i.

The Niblo-Roller construction found in [29] may be regarded as showing that any CAT(0) cube complex can be d_1-isometrically embedded in $[0,1]^\infty$, which can be viewed as an infinite product of trees.

Definition. A *colouring map* for a graph G is a map $c : V \to \{1, 2, \ldots, n\}$ such that $c(v) = c(u) \implies u$ and v are not joined by any edge of G.

The *chromatic number* of a graph G is the smallest n such that there exists a colouring map $c : V \to \{1, 2, \ldots, n\}$.

Definition. ([31]) The *transversality graph* of a halfspace system $(H, \leq, *)$ is the graph $T(H)$ with vertex set $H/\sim$, where $\sim$ is the equivalence relation $X_1 \sim X_2$ if and only if $X_1 = X_2$ or $X_1 = X_2^*$. We denote the equivalence class containing X_1 by $[X_1]$. Two vertices $[X_1]$ and $[X_2]$ are connected by an edge in $T(H)$ if and only if X_1 and X_2 are transverse in H.

Definition. The *hyperplane chromatic number* of a CAT(0) cube complex X is the chromatic number of the transversality graph $T(H)$ of the halfspace system $(H, \leq, *)$ associated to X. Note that the chromatic number of a cube complex may be infinite.

For any finitely generated Coxeter group W Dranishnikov and Schroeder [17] give a construction of a CAT(0) cube complex in which W embeds and a proof that the hyperplane chromatic number of X is finite. In the proof of lemma 2.9, we will give a construction for a CAT(0) cube complex with hyperplanes chromatic number k for any $k \in \mathbb{N}$.

The following lemma is proved using a generalisation of the methods of Niblo and Roller in [29].

Lemma 1.48. *Let X be a CAT(0) cube complex with hyperplane chromatic number n. Then X can be embedded d_1-isometrically in a product of n trees.*

Proof. Suppose X has hyperplane chromatic number n and let H denote the set of hyperplanes in X. We can find a map $c : H/\sim \to \{1, \ldots, n\}$ such that for all $[X_1], [X_2]$ in H $c([X_1]) = c([X_2]) \implies [X_1]$ and $[X_2]$ are not joined by an edge in $T(H)$, that is X_1 and X_2 are not transverse in X. We rewrite the set H as a disjoint union of sets $H = \coprod H_i$ where $H_i = \{[X_j] \in H | c([X_j]) = i\}$.

The construction outlined here is based on the construction given in the proof of lemma 1.37, and is discussed in detail in section 3.1.2.

For each set H_i we can build a tree T_i using the construction given in lemma 1.37 and the halfspace system $(\{X_h, X_h^* | h \in H_i\}, \leq, *)$ where $\leq$ is the order defined by inclusion and $*$ exchanges X_h and X_h^* for each $h \in H_i$. $X \setminus H_i$ consists of a number of connected components. For each component there is a vertex v in T_i defined by the section which maps each boundary $[X_j]$ to the halfspace X_j or X_j^* containing that component. We join two vertices by an edge if and only if the corresponding components are separated by exactly one hyperplane. By definition no pair of hyperplanes in H_i are transverse. Hence since n-cubes arise from sets of n boundaries which are pairwise transverse, T_i contains no cubes of dimension greater than 1 and is a graph. Since the components of the complex constructed from $(H_i, \leq, *)$ are CAT(0), T_i is simply connected and hence must be a tree.

For each H_i and T_i we define a map $\sigma_i : X^{(0)} \to T_i$ by sending each vertex v to the vertex of T_i corresponding to the component of $X \setminus H_i$ which contains v. We note that this map is not injective, two vertices may lie in the same component and hence be mapped to the same vertex of T_i. Two vertices in T_i are joined by an edge if and only if the corresponding components of $X \setminus H_i$ are separated by a single hyperplane. Hence for any two components in $X \setminus H_i$ separated by k hyperplanes, there is a path of length k between the corresponding vertices in T_i. Since T_i is a tree, there is no shorter path, and the distance between components in the number of hyperplanes separating them. Hence for any pair of vertices u, v the distance between $\sigma_i(u)$ and $\sigma_i(v)$ is the number of hyperplanes in the set H_i separating them.

We define the map $\sigma : X^{(0)} \to T_1 \times T_2 \times \ldots \times T_n$ for every vertex $v \in X$ by $\sigma(v) = (\sigma_1(v), \sigma_2(v), \ldots, \sigma_n(v))$. Since every edge in the CAT(0) cube complex intersects exactly one hyperplane, the distance $d_1(u, v)$ between two vertices u, v in the CAT(0) cube complex is precisely the number of hyperplanes separating them. Each hyperplane lies in H_i for exactly one i in $\{1, \ldots, n\}$, hence taking the product metric on the product of trees $d_1(u, v)$

is exactly the distance between $\sigma(u)$ and $\sigma(v)$ in $T_1 \times T_2 \times \ldots \times T_n$ and the map σ is an d_1-isometry. $\square$

Chapter 2

Cube complexes which do not embed in finite products of trees

We will prove the following result:

Theorem 2.1. *For each $k \in \mathbb{N}$ there exists a right-angled Coxeter group W_k and a 2-dimensional CAT(0) cube complex $\mathcal{U}_k$ such that W_k acts isometrically, cocompactly and properly on $\mathcal{U}_k$ and there is no bending map from $\mathcal{U}_k$ to a product of less than k trees.*

2.1 Preliminaries

Definition. A *hyperplane colouring map* for a CAT(0) cube complex X with hyperplane set $\mathcal{H}$ is a map $c : \mathcal{H} \to \{1, 2, \ldots, n\}$ such that for all $\mathfrak{h}$, $\mathfrak{h}' \in \mathcal{H}$, $c(\mathfrak{h}) = c(\mathfrak{h}') \implies \mathfrak{h}$ and $\mathfrak{h}'$ do not intersect.

Note that a hyperplane colouring map corresponds to a colouring of the transversality graph of the corresponding halfspace system.

The *hyperplane chromatic number* of a CAT(0) cube complex X is the smallest n such that there exists a hyperplane colouring map $c : \mathcal{H} \to \{1, 2, \ldots, n\}$ for X.

Remark 2.2. Let X be a CAT(0) cube complex and $\mathcal{H}$ be the set of hyperplanes in X. Let k be the hyperplane chromatic number of X and $c : \mathcal{H} \to \{1, \ldots, k\}$ be a hyperplane colouring map for X. Let $\overline{\mathcal{H}} \subset \mathcal{H}$ be a subset of the hyperplanes of X. Then the restriction of c to the set $\mathcal{H}$ has the property that for all $\mathfrak{h},\ \mathfrak{h}' \in \overline{\mathcal{H}}$, $c(\mathfrak{h}) = c(\mathfrak{h}') \implies \mathfrak{h}$ and $\mathfrak{h}'$ do not intersect. We make use of this fact in later in this chapter.

Recall the definition of a hyperplane as an equivalence class of midplanes. For any CAT(0) cube complex there is an canonical map m from the set of edges to the set of hyperplanes given by inclusion of midpoints of edges in these classes.

Definition. Let X be a CAT(0) cube complex and $e, \overline{e}$ a pair of edges in X. A *square-path* from e to $\overline{e}$ in X is a sequence of 2-dimensional cubes $C_1, \ldots, C_n$ in X such that $e \cap C_1 = e$, $C_n \cap \overline{e} = \overline{e}$ and $C_i \cap C_{i+1}$ is an edge for all $1 \leq i \leq n-1$.

A *straight square-path* from e to $\overline{e}$ in X is a square-path $C_1, \ldots, C_n$ from e to $\overline{e}$ which in addition satisfies $e \cap C_1 \cap C_2 = \emptyset, C_{n-1} \cap C_n \cap \overline{e} = \emptyset$ and $C_i \cap C_{i+1} \cap C_{i+2} = \emptyset$ for all $1 \leq i \leq n-2$. This additional condition ensures that for all i the cubes C_i and C_{i+2} meet C_{i+1} at opposite edges.

Lemma 2.3. *Let X be a CAT(0) cube complex and $e, \overline{e}$ be edges in X. The edges e and $\overline{e}$ are mapped by m to the same hyperplane $\mathfrak{h}$ if and only if there is a straight square-path from e to $\overline{e}$.*

Proof. Suppose that there is a straight square-path $C_1, \ldots, C_n$ from e to $\overline{e}$. Denote e by C_0 and $\overline{e}$ by C_{n+1}. Let M_0 be the midpoint of e and M_{n+1} be the midpoint of $\overline{e}$. Then for all $0 \leq i \leq n$, $C_i \cap C_{i+1}$ is an edge, call this edge e_i. Since $C_i \cap C_{i+1} \cap C_{i+2} = \emptyset$, the edges e_i and e_{i+1} must be opposite faces of the square C_{i+1}. Let M_{i+1} be the midplane of C_{i+1} which cuts both e_i and e_{i+1}. Then both M_{i+1} and M_{i+2} cut the edge e_{i+1}, and $M_{i+1} \cap M_{i+2}$ is a midplane - the midpoint of the edge e_{i+1}. Thus the sequence $M_0, M_1, \ldots, M_n, M_{n+1}$ is a chain of midplanes and $M_0 = m(e)$ and $M_{n+1} = m(\overline{e})$ lie in the same hyperplane.

Suppose e and $\overline{e}$ are mapped by m to the same hyperplane $\mathfrak{h}$. Then there exists a chain of midplanes $M_1, M_2, \ldots, M_n$ where M_1 is the midpoint of e and M_n the midpoint of $\overline{e}$. Given such a chain of midplanes, we can construct a chain of midplanes from M_1 to M_n in which each midplane has dimension less than 2. Suppose M_2 has dimension $d_2 \geq 2$. Then we can choose an edge path $m_2^1, \ldots, m_2^{k_2}$ of length at most d_2 from the vertex $M_1 \cap M_2$ to a vertex in $M_2 \cap M_3$, since the edges and vertices of M_2 are also midplanes in X, $m_2^1, \ldots, m_2^{k_2}$ is a chain of midplanes. Replace the chain $M_1, M_2, \ldots M_n$ with $M_1, m_2^1, \ldots, m_2^{k_2}, M_3, \ldots, M_n$. Repeating this process with the new chain for each M_i of dimension greater than 1, we obtain a chain of midplanes $M_1, m_2^1, \ldots, m_2^{k_2}, \ldots, m_{n-1}^1, \ldots, m_{n-1}^{k_{n-1}}, M_n$ in which every midplane has dimension at most one.

Given a chain of midplanes $m_1, m_2, \ldots, m_n$ we can replace any subsequence m_j, m_{j+1}, m_{j+2} such that $m_j \cap m_{j+1} \cap m_{j+2}$ is a midplane with the subsequence m_j, m_{j+2} to construct a shorter chain of midplanes from m_1 to m_n. Using this process, we remove midplanes from the chain $M_1, m_2^1, \ldots, m_2^{k_2}, \ldots, m_{n-1}^1, \ldots, m_{n-1}^{k_{n-1}}, M_n$ as necessary to produce a chain of midplanes which has no subsequence of length three such that the intersection of the elements in that subsequence is a midplane.

Let C_i^j be the unique square with midplane m_i^j. Let m_j, m_{j+1} be adjacent midplanes in the chain of midplanes. Then $m_j \cap m_{j+1}$ is a midplane of dimension 0 and is the midpoint of an edge e_j. Hence we must have $C_j \cap C_{j+1} = e_j$. Since $m_{j+1} \cap m_{j+2}$ is also a midplane and $m_j \cap m_{j+1} \cap m_{j+1}$ is not a midplane, it follows that $m_{j+1} \cap m_{j+2}$ is the midpoint of the edge of C_{j+1} opposite e_j. Label this edge e_{j+1}. Then $C_{j+1} \cap C_{j+2} = e_{j+1}$, and $C_j \cap C_{j+1} \cap C_{j+2} = e_j \cap e_{j+1} = \emptyset$. Hence $C_2^1, C_2^2, \ldots, C_2^{k_2}, \ldots, C_{n-1}^1, \ldots C_{n-1}^{k_{n-1}}$ is a straight square-path from e to $\overline{e}$. $\square$

Let $T_1 \times \ldots \times \hat{T_i} \times \ldots \times T_k$ denote the product $T_1 \times \ldots \times T_{i-1} \times T_{i+1} \times \ldots \times T_k$. Let p be the map from the set of edges of $T_1, \ldots, T_k$ to sets of edges in $T_1 \times \ldots \times T_k$ given by defining the image of an edge e_i in the tree T_i to be the product of e_i with the set of vertices of $T_1 \times \ldots \times \hat{T_i} \times \ldots \times T_k$. Denote $p(e_i)$ by E_i.

Lemma 2.4. *Let T be a product of trees $T_1 \times \ldots \times T_k$. Suppose that there exists a straight square-path $C_1, \ldots, C_n$ from e to $\overline{e}$ in T. Then both e and $\overline{e}$*

lie in E_i for some $i \in 1, \ldots, k$, and some edge e_i in T_i.

Proof. Every edge in T is of the form $e_i \times \hat{v}_i$ for some unique choice of tree T_i, edge $e_i \in T_i$ and vertex $\hat{v}_i \in T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k$. Similarly, every square in T is of the form $e_i \times \hat{e}_i$ for some T_i and some pair of edges $e_i, \hat{e}_i$ with e_i an edge in T_i and $\hat{e}_i$ an edge in $T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k$.

Consider the straight square-path $C_1, \ldots, C_n$ from e to $\overline{e}$. By the above, e must be of the form $e_i \times \hat{v}_i$ for some T_i and some edge e_i in T_i and vertex $\hat{v}_i$ in $T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k$. Since $e \cap C_1 = e$, C_1 must be of the form $e_i \times \hat{e}_i{}^1$ where $\hat{e}_i{}^1$ is an edge in $T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k$ which has $\hat{v}_i$ as one of its vertices.

Since $C_1 \cap C_2$ is an edge, C_2 must be of the form

$$\begin{array}{lll} e_i \times \hat{e}_i{}^2 & \text{where} & e_i \text{ is an edge in } T_i, \\ & & \hat{e}_i{}^2 \text{ is an edge in } T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k \text{ and} \\ & & \hat{e}_i{}^1 \cap \hat{e}_i{}^2 \text{is a vertex in } T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k \end{array}$$

or

$$\begin{array}{lll} \hat{e}_i{}^1 \times e_i{}^1 & \text{where} & e_i^1 \text{ is an edge in } T_i, \\ & & \hat{e}_i{}^1 \text{ is an edge in } T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k \text{ and} \\ & & e_i^1 \cap e_i \text{ is a vertex in } T_i \end{array}$$

Suppose C_2 is of the form $\hat{e}_i{}^1 \times e_i^1$ then

$$\begin{aligned} e \cap C_1 \cap C_2 &= (e_i \times \hat{v}_i) \cap (e_i \times \hat{e}_i{}^1) \cap (e_i \times \hat{e}_i{}^1) \\ &\subseteq (e_i \cap e_i^1) \times \hat{v}_i \end{aligned}$$

as $\hat{v}_i \in \hat{e}_i{}^1$. Since $e_i \cap e_i^1$ is a vertex, $(e_i \cap e_i^1) \times \hat{v}_i{}^1 \neq \emptyset$ which contradicts the hypothesis that $C_1, \ldots, C_n$ is a straight square-path. Hence C_2 must be of the form $e_i \times \hat{e}_i{}^2$. Similarly, each C_α must be of the form $e_i \times \hat{e}_i{}^\alpha$ for some edge $\hat{e}_i{}^\alpha$ in $T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k$, and $\overline{e}$ must be of the form $e_i \times \overline{v_i}$ for some vertex $\overline{v_i}$ in $T_1 \times \ldots \times \hat{T}_i \times \ldots \times T_k$. Hence e and $\overline{e}$ lie in E_i for some i. $\square$

Definition. An map α from a cube complex X to a cube complex T is called a *bending map* if

1. it is injective and
2. the restriction $\alpha|_S$ of α to the cube S is an isometry from the n-cube S in X to an n-cube $\alpha(S)$ in T for every cube S of X.

Remark 2.5. Let $\alpha : X \to T$ be a bending map. Suppose that a pair of cubes S and S' are adjacent in X, that is $S \cap S'$ is non empty. Then the cubes $\alpha(S)$ and $\alpha(S')$ are adjacent in T. To see this, note that, by the definition of a cube complex, if $S \cap S'$ is non-empty it contains a cube C which lies in both S and S'. Restricting α to S we see that $\alpha(C)$ lies in $\alpha(S)$, and similarly, $\alpha(C)$ lies in $\alpha(S')$, and hence $\alpha(C) \subset \alpha(S) \cap \alpha(S')$ and so $\alpha(S)$ and $\alpha(S')$ are adjacent.

Note that a bending map is not necessarily an isometry. To see this, consider the possible images under a bending map of a pair of adjacent cubes (for example see the map α in figure 2.1), and the distance between a pair of points where one point lies in each of these cubes.

In fact, a bending map is not necessarily a quasi-isometry, since we can choose a pair of cube complexes X and T and a bending map $\alpha : X \to T$ such that for any $k_1 \in \mathbb{R}$ there exists a pair of vertices u and v in X with $d(u, v) > k_1$ for which $d(\alpha(u), \alpha(v)) = 1$. For example, consider a bending map from the infinite tree X in which every vertex has valency two into the 2-dimensional cube complex T isomorphic to the Euclidean plane which maps X to a double spiral in T (see α' in figure 2.1). Then for every $k_1 > 0$ we can choose two vertices in T which are a distance 1 apart and which lie in different 'arms' of the spiral whose preimages are at least k_1 apart in X. To see that this is possible, consider how the distance in the preimage changes as we move away from the central point of the spiral.

Lemma 2.6. *Suppose X and T are CAT(0) cube complexes and that α is a bending map from X to T. Then there is a map h from the set $\mathcal{H}_X$ of hyperplanes in X to a subset of the set $\mathcal{H}_T$ of hyperplanes in T such that if the pair of hyperplanes $\mathfrak{h}, \mathfrak{h}' \in \mathcal{H}_X$ intersect then so do $h(\mathfrak{h})$ and $h(\mathfrak{h}')$.*

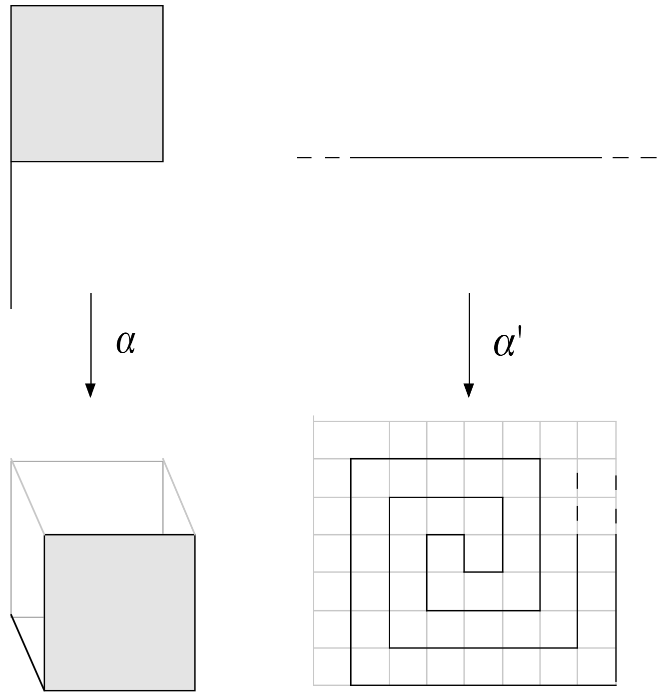

Figure 2.1: Bending maps which are not quasi-isometries

Proof. The bending map α takes n-cells in X isometrically to n-cells in T, and hence induces a map h which takes midplanes of n-cells in X to midplanes of n-cells in T. To see that h extends to a well-defined map on hyperplanes, consider two midplanes M and M' which lie in the same hyperplane in X. Then there exists a sequence $M_1, M_2, \ldots, M_k$ of midplanes in X such that each of $M \cap M_1$, $M_i \cap M_{i+1}$, $i \in \{1, \ldots, k-1\}$ and $M_k \cap M'$ is also a midplane.

Then $h(M)$ and $h(M')$ lie in the same hyperplane in T. To see this, consider the sequence $h(M_1), h(M_2), \ldots, h(M_k)$ of midplanes in T. We know that $M_1 \cap M_2$ is a midplane in X. Hence the n-cells containing M_1 and M_2 must both be adjacent to an $(n-1)$-cell with midplane $M_1 \cap M_2$. Since the map α preserves adjacency, the cells containing $h(M_1)$ and $h(M_2)$ must both be adjacent to an $(n-1)$- cell with midplane $h(M_1) \cap h(M_2)$. Similarly, each of the intersections $h(M) \cap h(M_1)$, $h(M_i) \cap h(M_{i+1})$, $i \in \{1, \ldots, k-1\}$ and $h(M_k) \cap h(M')$ is a midplane in T. Hence the midplanes $h(M)$ and $h(M')$ lie in the same hyperplane as required.

It remains to show that if $\mathfrak{h}$ and $\mathfrak{h}'$ are hyperplanes in X which intersect, then the hyperplanes $h(\mathfrak{h})$ and $h(\mathfrak{h}')$ intersect in T. To see this, note that if $\mathfrak{h}$ and $\mathfrak{h}'$ intersect, then then there must be a pair of midplanes $M \in \mathfrak{h}$ and $M' \in \mathfrak{h}'$ which are midplanes of the same cube C in X and hence intersect within that cube. Restricting α to C we have an isometry from C to $\alpha(C)$, and hence the images $h(M)$ and $h(M')$ intersect in the cube $\alpha(C)$ in T, and hence the hyperplanes $h(\mathfrak{h})$ and $h(\mathfrak{h}')$ intersect. □

Lemma 2.7. *Let X be a CAT(0) cube complex. If there is a bending map from X to a product of k trees then X has hyperplane chromatic number less than or equal to k.*

Proof. Suppose there is a bending map from X to the product of trees $T = T_1 \times \ldots \times T_k$. Let $\mathcal{H}_T$ denote the set of hyperplanes in T and $\mathcal{H}_X$ the hyperplanes in X.

Since there is a bending map from X to T, by lemma 2.6 there is a map h from the set of hyperplanes $\mathcal{H}_X$ to a subset of $\mathcal{H}_T$ such that if $\mathfrak{h}, \mathfrak{h}' \in \mathcal{H}_X$ intersect then so do $h(\mathfrak{h})$ and $h(\mathfrak{h}')$. Hence by remark 2.2 it suffices to show that the product of trees T has hyperplane chromatic number less than or equal to k.

Let m be the canonical map from the set of edges in the CAT(0) cube complex T to the set of hyperplanes $\mathcal{H}_T$. We extend this map to a map $\overline{m}$ from the set of edges in the trees $T_1, \ldots, T_k$ to the set of hyperplanes $\mathcal{H}_T$. For every edge e_i in a tree T_i, choose any edge $\overline{e_i}$ in the complex T which lies in the image of e_i under p. We define $\overline{m}(e_i)$ to be $m(\overline{e_i})$.

To see that this map is well defined, we need to show that $\overline{m}$ does not depend on our choice of edge $\overline{e_i}$ in E_i. Consider an edge e_i in the tree T_i, which has image a set of edges E_i in T. Let e and $\overline{e}$ be a pair of edges in E_i. Then e and $\overline{e}$ are given by $e_i \times v$ and $e_i \times \overline{v}$ respectively where v, $\overline{v}$ are vertices in $T_1 \times \ldots \times \hat{T_i} \times \ldots \times T_k$. Since $T_1 \times \ldots \times \hat{T_i} \times \ldots \times T_k$ is a CAT(0) cube complex, it is connected and contains an edge path p from v to $\overline{v}$. Then $\{e_i\} \times \{p\}$ is a straight square-path from e to $\overline{e}$, and so by lemma 2.3 each image of e_i in T is mapped to the same hyperplane in $\mathcal{H}_T$.

Suppose e, $\overline{e}$ are a pair of edges in T such that $m(e) = m(\overline{e})$. By lemma

2.3 if $m(e) = m(\overline{e})$ then there is a straight square path from e to $\overline{e}$. Suppose e is an edge in $E_i = p(e_i)$ and $\overline{e}$ is an edge in $E_j = p(e_j)$. Then we have a straight square-path in T from the edge e to the edge $\overline{e}$. By lemma 2.4, it follows that $e_i = e_j$. Hence the map $\overline{m}^{-1}$ is also well defined.

Define a map $c_T : \mathcal{H}_T \to \{1, \ldots, k\}$ by taking $c_T(\mathfrak{h})$ to be the index of the tree containing $\overline{m}^{-1}(\mathfrak{h})$. To see that c_T is a hyperplane colouring map, note that for any $e_i, \overline{e_i} \in T_i$ the midpoints of e_i and $\overline{e_i}$ are distinct and hence the hyperplanes $\overline{m}(e_i)$ and $\overline{m}(\overline{e_i})$ do not intersect in T. Hence T has hyperplane chromatic number at most k. □

2.2 2-dimensional cube complexes which do not embed in products of k trees

2.2.1 CAT(0) cube complexes

Lemma 2.8. *([12]) For every integer $k > 0$ there exists a graph G_k with chromatic number k and no cycle of length less than 6.*

Proof. For $k < 3$ the result is trivial. The following construction is due to Blanche Descartes in [12]. We define inductively a sequence $G_3, G_4, \ldots$ of graphs. For each $k \geq 3$ G_k has chromatic number k and contains no cycle of length less than 6.

Let G_3 be a graph with chromatic number 3 and no cycles of length less than 6, for example the cycle of length 7. Let m_i be the number of vertices in G_i, and define $M_i := \begin{pmatrix} im_i - i + 1 \\ m_i \end{pmatrix}$. For all $i \geq 3$, let G_{i+1} be defined as follows: take M_i copies of G_i, and $im_i - i + 1$ additional vertices, which we will refer to as central vertices. Choose a one-to-one map between the set of copies of G_i and the set of subsets of the central vertices with m_i members. Denote the sets of m_i central vertices by $S^1, S^2, \ldots S^{M_i}$ and denote the copy of G_i corresponding to S^j by G_i^j.

For each set S^j of central vertices, we add edges joining each vertex in S^j to a vertex of G_i^j in such a way that no two edges are incident with the same

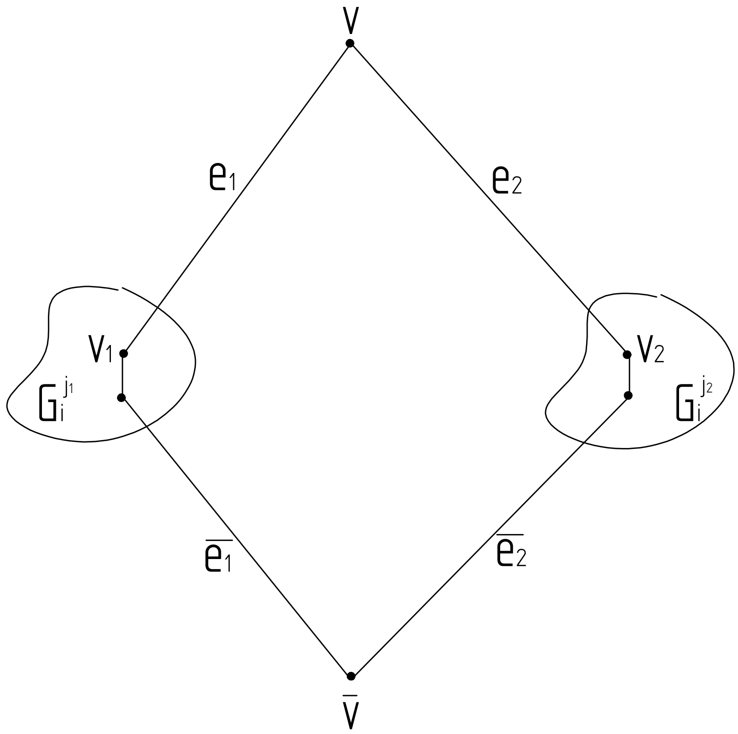

Figure 2.2: Cycles must have length at least six

vertex. This is possible since we have chosen the size of the set S^j to be the number of vertices in G_i^j. The resulting graph is G_{i+1}.

To see that there are no cycles of length less than 6 in the graph G_i we will use induction on i. Suppose that G_i has no cycles of length less than 6. Then any cycle in G_{i+1} with length less than 6 cannot lie entirely within a copy of G_i, and hence must contain a central vertex v and a pair of edges $e_1,\ e_2$ incident with v. Since there is no edge joining two central vertices, e_1 joins v to some vertex v_1 of $G_i^{j_1}$ for some $j_1 \in \{1, \ldots, M_i\}$. Since no two edges from S^{j_1} to $G_i^{j_1}$ have a common end point, the edge e_2 must join v to some vertex v_2 of $G_i^{j_2}$ for some $j_2 \in \{1, \ldots, M_i\} \setminus j_1$.

There are no edges joining vertices in different copies of G_i, hence since $j_1 \neq j_2$ there must be a second central vertex $\overline{v}$ in the cycle. As above, within the cycle each central vertex must be incident with two edges which join it to vertices in distinct copies of G_i. Note that any cycle containing 3 or more central vertices must therefore have length greater than or equal to 6.

Let us assume that only two of the vertices of the cycle are central vertices, v and $\overline{v}$. Then there must be edges in the cycle from $\overline{v}$ to vertices in $G_i^{j_1}$ and $G_i^{j_2}$. Let $\overline{e_1}$ be an edge from $G_i^{j_1}$ to $\overline{v}$. Suppose $\overline{e_1}$ is incident with v_1. Then $\overline{v}$ does not lie in the set S^{j_1} of central vertices since no two edges from S_1^j to $G_i^{j_1}$ can share an endpoint. Hence there is no edge from $\overline{v}$ to a vertex in $G_i^{j_1}$, which is a contradiction. Hence $\overline{e_1}$ must be incident with a vertex in $G_i^{j_i}$ other than v_1. Since we are interested in the shortest possible circuit, let us assume that this vertex is joined to v_1 by an edge.

Similarly, there must be an edge $\overline{e_2}$ from $\overline{v}$ to a vertex of $G_i^{j_2}$ other than v_2. Hence it follows that any cycle in G_{i+1} has length at least six. See figure 2.2.

Since G_3 contains no cycles of length less than 6, it follows that for all $k \geq 3$ G_k contains no cycle of length less than 6.

To see that for each k the chromatic number of G_k is k we again use induction on i. Suppose that G_i has chromatic number i. Since no pair of central vertices are joined by an edge, and no pair of vertices in different copies of G_i are joined by an edge, we can colour G_{i+1} with the colours $1, \ldots, i, i+1$ by colouring the vertices in each copy of G_i using a colouring map for G_i, and the central vertices with colour $i+1$. Hence G_{i+1} has chromatic number at most $i+1$.

Suppose G_{i+1} can be coloured with i colours. Then the $im_i - i + 1$ central vertices can be coloured with i colours, and for some $d \in 1, \ldots, i$ there are at least m_i vertices with colour d. Let S^j denote a set of m_i central vertices in which every vertex has colour d. Then every vertex in G_i^j is joined by an edge to a central vertex which has colour d, and so G_i^j must be coloured in $i-1$ colours, which contradicts the fact that G_i has chromatic number i. Hence G_{i+1} has chromatic number $i+1$. G_3 is a cycle with odd length, and hence has chromatic number 3. Hence for each $k \geq 3$ G_k has chromatic number k.

□

Lemma 2.9. *For every integer $k > 0$ there exists a compact 2-dimensional CAT(0) cube-complex X_k such that there is no bending map from X_k to a product of less than k trees.*

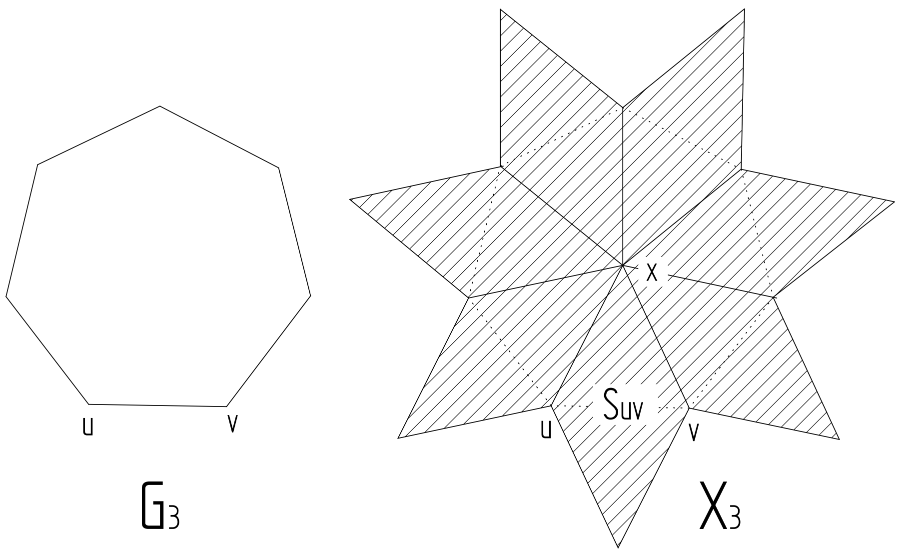

Figure 2.3: The construction of X_3 from G_3

Proof. By Lemma 2.8 we can construct a graph G_k which has chromatic number k and contains no cycles of length less than six. We construct X_k as follows: begin with a vertex x. For each vertex v in G_k we add a vertex v and an edge $e_v = (x, v)$ to X_k. If the vertices u and v are joined by an edge in G_k then we attach a square S_{uv} to X_k by identifying the edges e_u and e_v with two adjacent faces of S_{uv}. For an example of this construction in the case $k = 3$, see figure 2.3.

Since X_k is constructed by identifying the faces of cubes of dimension less than or equal to 2, X_k is a 2-dimensional cube complex. For all $k > 1$ the cube complex X_k is locally CAT(0). In order to see this, consider the link of each vertex in X_k. By Lemma 1.29 X_k is CAT(0) if the link of every vertex in X_k contains no cycle of length less than 4.

We consider the vertices of X_k in three sets. By definition, the link of the vertex x is the graph G_k which was chosen to contain no cycles of length less than 6. Consider the vertices of X_k which correspond to vertices in G_k. If v is such a vertex then the set of edges in $link(v)$ corresponds to the set of squares in X_k which contain v in their boundaries. Every such square is also incident with the edge e_v, but since G_k contains no cycles of length two, no two squares in X_k share more than one edge. It follows that $link(v)$ contains no cycles. The remaining set of vertices are those lying on the boundary of squares of the form S_{uv} which are not joined to the vertex x by an edge. Each of these vertices is in the boundary of a single square, and hence its link contains a single edge. Hence no vertex has a link which contains a cycle of less than 4 edges, and it follows that X_k is locally CAT(0).

In order to complete the proof that X_k is CAT(0), we need to show that X_k is simply connected. Considering the definition of X_k, we observe that each square in X_k lies in the star of the vertex x, that each edge of X_k lies either in $star(x)$ or in the boundary of a square in $star(x)$, and that each vertex lies either in $star(x)$ or in the boundary of some higher dimensional cube of $star(x)$. Hence X_k is equal to the closure of the star of x in X_k, and X_k is path-connected.

Suppose there exists a loop ℓ in X_k which is not homotopic to the constant loop. Then ℓ is homotopic to a loop in the 1-skeleton of X_k. Suppose this loop contains edges which do not lie in the star of x. Each such edge lies in the boundary of a square in $star(x)$, and so we can construct a homotopy from ℓ to a loop which lies in the 1-skeleton of $star(x)$. Hence the existance of a loop ℓ in X_k which is not homotopic to the constant loop implies the existence of a non-trivial loop in the 1-skeleton of $star(x)$. By definition, $star(x)^{(1)}$ is a tree and hence contains no non-trivial loops. Hence X_k is simply connected.

We claim that X_k has hyperplane chromatic number at least k. Consider the set $\mathcal{H}_x$ of hyperplanes in X_k which contain midpoints of edges which are incident with the vertex x. Since the vertex x has finite valency, the set $\mathcal{H}_x$ is finite and hence a hyperplane colouring of the set $\mathcal{H}_x$ with finitely many colours exists.

Let VG_k denote the set of vertices of G_k. We define a map g from VG_k to the set of edges incident with x by setting $g(v) = e_v$. Let m be the map from the set of edges in a cube complex to the set of hyperplanes, as defined in the proof of lemma 2.7. Suppose $c : \mathcal{H}_x \to \{1, 2, \ldots, \kappa\}$ is a hyperplane colouring map on the set $\mathcal{H}_x$.

Then $c \circ m \circ g : VG_k \to \{1, 2, \ldots, \kappa\}$ is a colouring map for G_k. To see this, suppose u and v are joined by an edge in G_k. Then there is a square S_{uv} in X_k with the edges e_u and e_v in its boundary. The hyperplanes $m(e_u)$ and $m(e_v)$ intersect in the square S_{uv} and it follows that $c(m(e_u)) \neq c(m(e_v))$. Hence $c \circ m \circ g(u) \neq c \circ m \circ g(v)$ as required.

Suppose $\kappa < k$. Then the chromatic number of G_k is less than k, which is a contradiction. So the hyperplane chromatic number of X_k is greater than or equal to k. Hence by lemma 2.7, there is no bending map from X_k to a product of less than k trees.

□

Lemma 2.10. *There exists a 2-dimensional CAT(0) cube complex X_∞ such that there is no bending map from X_∞ to a finite product of trees.*

Proof. By lemma 2.9, for every $k \in \mathbb{N}$ we can choose a 2-dimensional CAT(0) cube complex X_k such there is no bending map from X_k to a product of less than k trees. For each $k \in \mathbb{N} \setminus \{1\}$, take a copy of X_k and choose two vertices on the boundary of X_k which are not joined by an edge to the vertex x. To see that this is possible, consider the proof of lemma 2.9. For each $k > 1$ we can choose X_k so that each X_k contains more than one square, and each square contains a vertex not joined by an edge to the vertex x. Label the chosen vertices as v_k^- and v_k^+.

We define a CAT(0) cube complex X_∞ as follows: Consider the union $\bigcup_{k \in \mathbb{N} \setminus \{1\}} X_k$ and define the equivalence relation $\sim$ by $v_{k-1}^- \sim v_k^+$ for all $k > 2$. Then $X_\infty = \bigcup_{k \in \mathbb{N} \setminus \{1\}} X_k / \sim$.

Clearly X_∞ is a 2-dimensional cube complex. To see that X_∞ is CAT(0), we consider the link of each vertex in X_∞. Since we know that each X_k is CAT(0), we need only consider the link of those vertices given by the identification of v_{k-1}^- and v_k^+ for some $k > 2$. We saw in the proof of lemma 2.9 that each of these vertices has link consisting of a single edge. Since no higher dimensional cubes are identified under $\sim$, the link of the vertex in X_∞ corresponding to the equivalence class $[v_{k-1}^-, v_k^+]$ is a pair of disjoint arcs, and contains no cycles. Hence X_∞ is CAT(0).

Suppose there is a bending map from X_∞ to a product of K trees for some K. Then by lemma 2.7 there is a hyperplane colouring map $c : X_\infty \to \{1, 2, \ldots, K\}$ and by remark 2.2, the restriction of c to the set $\mathcal{H}_{X_{K+1}}$ is a hyperplane colouring map, where $\mathcal{H}_{X_{K+1}}$ is the set of hyperplanes in X_{K+1}. Then X_{K+1} has hyperplane chromatic number less than or equal to K and this contradicts lemma 2.9, hence there exists no bending map from X_∞ to a finite product of trees. $\square$

2.2.2 Hyperbolic cube complexes

In order to prove that the cube complexes constructed in the previous section are CAT(0), we made use of the fact that the link of any vertex in the complex contains no vertex whose link contains a cycle of less than 4 edges. In fact, we have constructed a cube complex which contains no vertex whose link contains a cycle of less than 7 edges. This allows us to prove a stronger result, that for any $k > 0$ we can choose a metric $d_{\mathbb{H}}$ on X_k such that $(X_k, d_{\mathbb{H}})$ is hyperbolic. In order to show this, we need the following definition and lemmas.

Definition. ([20],page 119) The *model space* (M, χ_0) is the complete simply connected manifold of constant curvature χ_0. A (M, χ_0)-*simplicial space* is a simplicial complex in which each simplex is isometric to a simplex in the model space (M, χ_0).

Let X be a cube complex in which each n-cube is isometric to the Euclidean n-cube with side length 1. Let $sub(X)$ be a subdivision of X such that every cell in $sub(X)$ is a simplex. Such a subdivision is always possible, for example take the barycentric subdivision of X. Then $sub(X)$ is a $(M, 0)$-simplicial space.

Definition. We say two cell complexes P_1 and P_2 are combinatorially equivalent if there is a bijective map f from the vertex set of P_1 to the vertex set of P_2 such that if $u_1, \ldots, u_j$ are vertices lying in the boundary of an i-dimensional cell of P_1 then $f(u_1), \ldots, f(u_j)$ lie in the boundary of an i-dimensional cell of P_2.

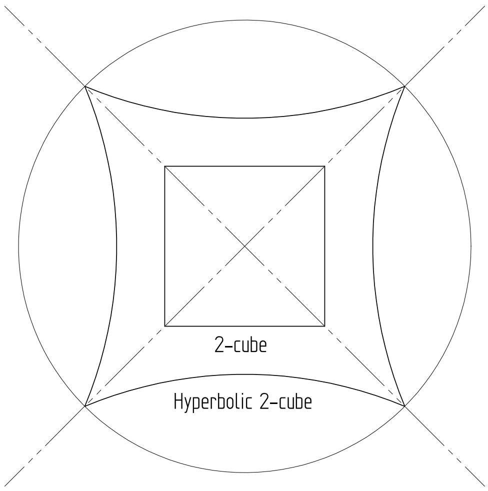

Figure 2.4: Constructing a “square” in the $(M, -1)$ model space.

For each n, we define a n-dimensional polyhedron in the model space $(M, -1)$ which is combinatorially equivalent to the Euclidean n-cube and in which every edge has length 1. Call this the *hyperbolic n-cube*.

To see that this is possible, consider embedding the Euclidean n-cube in the model space $(M, -1)$. We can construct a geodesic from the center of the cube to each vertex. Extending these geodesics, we can take the intersection of these geodesics with a n-sphere centered at the centre of the cube, and construct a polyhedron whose vertices are these points and which is combinatorially equivalent to the Euclidean n-cube. For an illustration of the 2-dimensional case, see figure 2.4.

By varying the radius of the sphere, we can construct such a polyhedron in which the lengths of the edges are any non-zero length, and hence can choose the edge length to be 1. This polyhedron is the hyperbolic n-cube. Note that the angle between adjacent edges will be less than $\frac{\pi}{2}$ in such a polyhedron.

Define a piecewise hyperbolic metric on the cube complex X by taking each n-cube in X to be isometric to the hyperbolic n-cube. For clarity, we will denote X with this metric by $\overline{X}$. Let $sub(\overline{X})$ denote a subdivision of $\overline{X}$ such that every cell in $sub(\overline{X})$ is a simplex. Since each cube in $\overline{X}$ was isometric to a cube in the $(M,-1)$ model space, $sub(X)$ is a $(M,-1)$-simplicial space.

The following lemmas are due to Gromov.

Lemma 2.11. ([20], page 120) *An (M,χ_0)-simplicial space L satisfies CAT(1) if and only if the curvature of L is less than or equal to 1 and for every two points l_1 and l_2 in L with distance between them less than π, l_1 and l_2 can be joined by at most one geodesic segment in L.*

Lemma 2.12. ([20], page 120) *An (M,χ_0)-simplicial space X has curvature less than or equal to χ if and only if $\chi \geq \chi_0$ and the link of each cell is CAT(1).*

Lemma 2.13. *For every integer $k > 0$ there exists a compact, 2-dimensional hyperbolic cube-complex $\overline{X_k}$ such that there is no bending map from $\overline{X_k}$ to $\overline{T}$ where $\overline{T}$ is a product of less than k trees with the piecewise hyperbolic metric.*

Proof. By Lemma 2.9, there exists a 2-dimensional CAT(0) cube complex X_k such that there is no bending map from X_k to a product of less than k trees with the standard piecewise Euclidean metric.

Replace each of the euclidean cubes in X_k with a hyperbolic n-cube in the model space $(M,-1)$ as described above. Hence form the piecewise hyperbolic space $\overline{X_k}$. Suppose that there is a bending map from $\overline{X_k}$ to a product $\overline{T} = T_1 \times \ldots \times T_n$where $n < k$ and where $\overline{T}$ has the piecewise hyperbolic metric. Then there is a bending map from X_k to the product $T = T_1 \times \ldots \times T_n$ with the piecewise Euclidean metric, which is a contradiction.

It remains to show that $\overline{X_k}$ is hyperbolic. Since lemmas 2.11 and 2.12 apply to (M,χ_0)-simplicial complexes, we need to subdivide $\overline{X_k}$ in such a way that every cell is a simplex. Since $\overline{X_k}$ is 2-dimensional, we can do this by placing a vertex at the center of each square, and joining that vertex to each corner of the square. We call this complex $sub(\overline{X_k})$, and note that with the metric inherited from $\overline{X_k}$ it is a $(M,-1)$-simplicial space. Then by lemma 2.12, $sub(\overline{X_k})$ (and hence $\overline{X_k}$) has curvature -1 if and only if the link of every cell in $sub(\overline{X_k})$ is CAT(1).

Since the dimension of $\overline{X_k}$ is 2, we need only consider the links of the vertices, which will be graphs whose edges are given length equal to the corresponding angle of the triangle in $\overline{X_k}$. We first consider the vertices of $sub(\overline{X_k})$ which were not vertices of $\overline{X_k}$. Each of these vertices lies within a hyperbolic space, namely the hyperbolic 2-cube of which it is the center. Hence by lemmas 2.11 and 2.12, the links of these vertices contain no cycles of length less than 2π.

We now consider those vertices of $sub(\overline{X_k})$ which are also vertices of $\overline{X_k}$. Since the length of edges in the link of a vertex is given by the size of the corresponding angles, taking $sub(\overline{X_k})$ preserves the total length of a cycle in the link of a vertex under this metric, while doubling the number of edges. Using the 2nd hyperbolic cosine rule (see Hyperbolic Geometry, [1]), we calculate that the angle α at a corner of a regular hyperbolic 4-gon with sides of length 1 is approximately 1.36 (2 d.p). Note that $4 \times \alpha < 2\pi < 5 \times \alpha$. Hence if the link of a vertex in $\overline{X_k}$ contains no cycles of less than 5 edges, it contains no cycles of length less than 2π and hence any two points in $link(v)$ joined by more than one geodesic have distance between them greater that π. Hence by lemma 2.11 the link of the corresponding vertex in X is CAT(1).

By lemma 2.8 and the proof of lemma 2.9, $\overline{X_k}$ satisfies the condition that the link of any vertex contains no cycle with less than 5 edges. Hence $\overline{X_k}$ has curvature -1, and is hyperbolic. □

2.3 Cube complexes with isometric group actions

2.3.1 The space $\mathcal{U}$

The following definitions and lemmas are from "The Geometry and Topology of Coxeter groups" by Michael Davis ([10]).

Definition. (page 59, [10]) A *mirror structure* on a space Y is an index set Q and a family $(Y_q)_{q \in Q}$ of closed subspaces of Y. Each Y_q is called a *mirror* of Y. For any subset $P \subset Q$, define $Y_P = \bigcap_{p \in P} Y_p$. Let $Q(y)$ denote the set $\{q \in Q | y \in Y_q\}$.

Definition. (page 59, [10]) A *family of groups* with index set Q consists of a group Γ, a subgroup B of Γ and a family $(\Gamma_q)_{q\in Q}$ of subgroups of Γ such that each Γ_q contains B. For any non empty subset $P \subset Q$, define Γ_P to be the group generated by $\{\Gamma_p | p \in P\}$. Let $\Gamma_\emptyset = \emptyset$.

In this thesis we will only consider families of groups where the subgroup B of Γ is trivial. In this case, any set $(\Gamma_q)_{q\in Q}$ of subgroups of Γ is a family of groups.

Definition. (page 60, [10]) Given a space Y with mirror structure $(Y_q)_{q\in Q}$ and a group Γ with family of groups $(\Gamma_q)_{q\in Q}$ we can define a space $\mathcal{U}(\Gamma, Y) = \mathcal{U}$ on which there is a Γ action with fundamental region homeomorphic to Y.

$$\mathcal{U}(\Gamma, Y) = \Gamma \times Y / \sim$$

where $\sim$ is the equivalence relation on points of $\Gamma \times Y$ given by $(\gamma_1, y_1) \sim (\gamma_2, y_2)$ if and only if $y_1 = y_2$ and $\gamma_1^{-1}\gamma_2 \in \Gamma_{Q(y_1)}$

There is a natural action of Γ on $\Gamma \times Y$ given by $g(\gamma, y) = (g\gamma, y)$ for all $g \in \Gamma$, $(\gamma, y) \in \Gamma \times Y$. Suppose (γ_1, y_1) and (γ_2, y_2) are a pair of points in $\Gamma \times Y$ such that $(\gamma_1, y_1) \sim (\gamma_2, y_2)$. Then $g(\gamma_1, y_1) \sim g(\gamma_2, y_2)$ since $(g\gamma_1)^{-1}g\gamma_2 = \gamma_1^{-1}(g^{-1}g)\gamma_2 = \gamma_1^{-1}\gamma_2 \in \Gamma_{Q(y_1)}$. Hence the action of Γ on $\Gamma \times Y$ descends to an action on $\mathcal{U}$.

Let $i : Y \to \mathcal{U}$ be the map defined by $y \mapsto (e, y)$ where e is the identity element of Γ. Then $i(Y)$ is an embedded copy of Y in $\mathcal{U}$, and is a fundamental region for the action of Γ on $\mathcal{U}$. For each $\gamma \in \Gamma$ and subspace $\overline{Y}$ of Y let $(\gamma, \overline{Y})$ denote the subspace $\{(\gamma, y) | y \in \overline{Y}\}$.

Suppose that there is a metric d on the space Y. If $x, y \in Y$ then for any γ in Γ, we define a metric on (γ, Y) by $d((\gamma, x), (\gamma, y)) = d(x, y)$. We define a *piecewise geodesic path* in $\mathcal{U}$ to be a path p such that for any $\gamma \in \Gamma$ the intersection of p with (γ, Y) is a geodesic with respect to the metric d in the subspace (γ, Y). Suppose $(g, x), (h, y)$ are points in $\mathcal{U}$. The space $\mathcal{U}$ inherits a metric from Y by taking $d((g, x), (h, y))$ to be the length of the shortest piecewise geodesic path from (g, x) to (h, y).

Definition. (page 61, [10]) A mirror structure on Y is Γ-finite (with respect to a family of subgroups for Γ) if $X_P = \emptyset$ for any finite subset $P \subset Q$ such that Γ_P is infinite.

Lemma 2.14. (page 61, [10]) *Given a group Γ and a space Y with associated mirror structure and family of groups, the Γ-action on $\mathcal{U}(\Gamma, Y)$ is properly discontinuous if and only if the following conditions hold:*

(a) Y is Hausdorff

(b) The mirror structure is Γ-finite.

Lemma 2.15. (page 62, [10]) *Suppose (W, S) is a Coxeter system. Define a family of subgroups indexed by S by taking, for each $s \in S$, W_s to be the subgroup of order 2 generated by s. Then $\mathcal{U}(W, Y)$ is connected (resp. path-connected) if the following two conditions hold:*

(a) Y is connected (resp. path connected) and

(b) Y_s is nonempty for each $s \in S$.

Definition. (page 113, [10]) Let (W, S) be a Coxeter system. A subset $T \subset S$ is *spherical* if the subgroup generated by T is finite.

Lemma 2.16. (page 151, [10]) *Let (Γ, S) be a Coxeter system and let Y be a connected cell complex withh associated family of groups and mirror structure indexed by S. Then $\mathcal{U}(\Gamma, Y)$ is simply connected if and only if the following three conditions hold:*

(a) Y is simply connected.

(b) For each $s \in S$, Y_s is nonempty and path connected.

(c) For each spherical subset $\{s, t\} \in S^{(2)}$, $Y_s \cap Y_t$ is nonempty.

2.3.2 The cube complex $\mathcal{U}_k$

For each $k \in \mathbb{N}$, let G_k be the graph as defined in lemma 2.8. Let W_k be the Coxeter group $\langle S_k | s^2 = 1 \forall s \in S_k \rangle$ where S_k is in one-to-one correspondence with the set of edges of G_k.

We define a family of groups with respect to W_k with index set S_k by taking $W_{k\,s}$ to be the order two subgroup of W_k generated by s.

Let X_k be the two-dimensional cube complex as defined in lemma 2.9. We define a mirror structure on X_k with index set S_k as follows: For each $s \in S_k$ there is a corresponding edge $\{u, v\}$ in G_k. For each such edge, we have a square S_{uv} of X_k. Let $X_{k\,s}$ be the vertex of S_{uv} opposite the vertex x.

For each $k \in \mathbb{N}$, we define

$$\mathcal{U}_k = \mathcal{U}(W_k, X_k)$$

Recall that $\mathcal{U}(W_k, X_k) = W_k \times X_k / \sim$. $W_k \times X_k$ is a collection of 2-dimensional cube complexes. The equivalence relation $\sim$ leads to the identification of a pair of points only if those points are vertices on the boundaries of distinct copies of X_k. It is clear from this that the resulting complex $\mathcal{U}_k$ is a 2-dimensional cube complex.

Let $i(X_k) = \{(e, x) | x \in X_k\}$. Then $i(X_k)$ is an embedded copy of X_k in $\mathcal{U}_k$. $i(X_k)$ is a fundamental region for the action of W_k on $\mathcal{U}_k$, hence since X_k is compact the action is cocompact.

Corollary 2.17. *$\mathcal{U}_k$ is connected and path connected.*

Proof. X_k is connected and path connected, and W_k is a Coxeter group. Each $X_{k\ s}$ is non-empty, hence by lemma 2.15 $\mathcal{U}_k$ is connected and path connected. □

Lemma 2.18. *The action of W_k on $\mathcal{U}_k$ is isometric with the inherited metric.*

Proof. Let (γ_1, x_1), (γ_2, x_2) be a pair of distinct points in $\mathcal{U}_k$ and let p be a shortest geodesic path from (γ_1, x_1) to (γ_2, x_2). Let

$$(\gamma_1, X_k) = (g_1, X_k), (g_2, X_k), \ldots, (g_n, X_k) = (\gamma_2, X_k)$$

denote the sequence of copies of X_k in $\mathcal{U}_k$ through which p passes, and let $p_i = (g_i, X_k) \cap p$ denote the geodesic segment of p_i in (g_i, X_k).

Since p is a path, $p_i \cap p_{i+1}$ must be a point $(g_i, x_i) \sim (g_{i+1}, x_i)$ and by the definition of $\mathcal{U}_k$ x_i must lie in a mirror $X_{k\ s_i}$ of X_k and we must have $g_i^{-1} g_{i+1} \in W_{k\ s_i}$.

Now consider the pair of points $g(\gamma_1, x_1) = (g\gamma_1, x_1)$, $g(\gamma_2, x_2) = (g\gamma_2, x_2)$. Then the set $g(p) = \{g(x)|x \in p\}$ is a piecewise geodesic path from $g(\gamma_1, x_1)$ to $g(\gamma_2, x_2)$. To see this, note that $g(p_i)$ contains the point (gg_i, x_i) and $g(p_{i+1})$ the point (gg_{i+1}, x_i) and that $(gg_i)^{-1}(gg_{i+1}) = g_i^{-1}(g^{-1}g)g_{i+1} = g_i^{-1}g_{i+1}$. Hence $g(p)$ is a path. For each i, $g(p_i)$ is a geodesic path in $(g\gamma_i, X_k)$ with length equal to the length of p_i, and hence $g(p)$ is a piecewise geodesic path. Hence $d(g(\gamma_i, x_1), g(\gamma_2, x_2)) \leq d((\gamma_1, x_1)(\gamma_2, x_2))$.

Suppose there exists a piecewise geodesic path from $g(\gamma_1, x_1)$ to $g(\gamma_2, x_2)$ with length less than p. Then

$$\begin{aligned} d(g(\gamma_1, x_1), g(\gamma_2, x_2) &\geq d(g^{-1}(g(\gamma_1, x_1)), g^{-1}(g(\gamma_2, x_2)) \\ &= d((\gamma_1, x_1), (\gamma_2, x_2)) \end{aligned}$$

Hence $d(g(\gamma_1, x_1), g(\gamma_2, x_2) = d((\gamma_1, x_1), (\gamma_2, x_2))$ and the action of W_k on $\mathcal{U}_k$ is isometric. □

Corollary 2.19. *The action of W_k on $\mathcal{U}_k$ is proper.*

Proof. Suppose M is a subset of S_k such that $W_{k\,M}$ is infinite. Then M must contain more than one element. But $X_{k\,s} \cap X_{k\,t} = \emptyset$ for any choice s, t of distinct elements of S_k, hence the mirror structure on X_k is W_k-finite. X_k is Hausdorff. Hence by lemma 2.14 the action of W_k on $\mathcal{U}_k$ is proper. □

Lemma 2.20. *$\mathcal{U}_k$ is non-positively curved.*

Proof. We consider the links of vertices of $\mathcal{U}_k$. Let $p : W_k \times X_k \to \mathcal{U}_k$ denote the canonical map onto $\mathcal{U}_k$, and let $p^{-1}(x)$ denote the preimage of the point x. The mirror structure on X_k was chosen so that each point which lies in a mirror is a vertex of X_k. The relation $\sim$ identifies points (γ_1, x_1) and (γ_2, x_2) of $W_k \times X_k$ only if x_1 and x_2 lie in mirrors, and so $\sim$ identifies only those points which are vertices. It follows that the link of a vertex v is isomorphic to the disjoint union of the links of the vertices $p^{-1}(v)$. As in the proof of lemma 2.9 the link of any vertex in X_k contains no cycle of less than 4 edges. Hence by lemma 1.29 $\mathcal{U}_k$ is non-positively curved. $\square$

Lemma 2.21. *$\mathcal{U}_k$ is simply connected.*

Proof. We saw in the proof of lemma 2.9 that the space X_k is simply connected. Each $X_{k\,s}$ contains a single point, and hence is nonempty and path connected. For any $s, t \in S$ the group generated by s and t is infinite, hence the Coxeter system (W_k, S_k) has no 2-element spherical subsets. Hence by lemma 2.16 $\mathcal{U}_k$ is simply connected. $\square$

Theorem 2.1. *For each $k \in \mathbb{N}$ there exists a right-angled Coxeter group W_k and a 2-dimensional CAT(0) cube complex $\mathcal{U}_k$ such that W_k acts isometrically, cocompactly and properly on $\mathcal{U}_k$ and there is no bending map from $\mathcal{U}_k$ to a product of less than k trees.*

Proof. Let W_k and $\mathcal{U}_k$ be as defined at the beginning of the section. Then by lemmas 2.20 and 2.21, $\mathcal{U}_k$ satisfies the conditions of lemma 1.29 and is a CAT(0) cube complex. By the definition of $\mathcal{U}_k$, since X_k is a 2-dimensional cube complex $\mathcal{U}_k$ is also a 2-dimensional cube complex.

By lemmas 2.19 and 2.18 and the definition of $\mathcal{U}_k$, W_k acts isometrically, cocompactly and properly on $\mathcal{U}_k$. Suppose there is a bending map α from $\mathcal{U}_k$ to a product of less than k trees. $i(X) = (e, X_k)$ is an embedded copy of X_k in $\mathcal{U}_k$, so the restriction $\alpha|_{i(X)}$ is a bending map from X_k to a product of less than k trees. This is a contradiction of lemma 2.9, hence there is no bending map from $\mathcal{U}_k$ to a product of less than k trees. $\square$

Chapter 3

Embeddings in finite products of trees

We prove the following result:

Theorem 3.1. *Let G be a group which acts isometrically, properly, and cocompactly on a finite dimensional, locally finite CAT(0) cube complex X in such a way that* $\mathrm{stab}_G(\mathfrak{h})$ *is separable for each hyperplane $\mathfrak{h}$ of X. Then there is a quasi-isometric embedding of X in a finite product of locally finite trees.*

If in addition to satisfying the conditions of theorem 3.1, the action of G on X is free, by lemma 1.42 there is a quasi-isometric embedding of the group in the cube complex X. Hence we have the following corollary to theorem 3.1:

Corollary 3.2. *Let G be a group which acts freely, isometrically, properly, and cocompactly on a finite dimensional, locally finite CAT(0) cube complex X in such a way that* $\mathrm{stab}_G(\mathfrak{h})$ *is separable for each hyperplane $\mathfrak{h}$ of X. Then there is a quasi-isometric embedding of G in a finite product of locally finite trees.*

3.1 Embeddings of the CAT(0) Cube Complex in a Product of Trees.

3.1.1 Choosing N-orbits of $\mathfrak{h}$ which do not cross

Let a group G act properly and cocompactly on a CAT(0) cube complex X. Then there exists a compact subset C of X such that $GC = X$. Since C is compact, there is a finite set of hyperplanes in X which intersect the subset C. Denote these hyperplanes by $\mathfrak{h}_1, \mathfrak{h}_2, \ldots, \mathfrak{h}_n$. Then the set of all hyperplanes in X is given by $\{G\mathfrak{h}_1, G\mathfrak{h}_2, \ldots, G\mathfrak{h}_n\}$, where $G\mathfrak{h}_i = \{g\mathfrak{h}_i | g \in G\}$.

In general, there may be some $g \in G$ and $\mathfrak{h}_i \in \{\mathfrak{h}_1 \ldots, \mathfrak{h}_n\}$ such that $g\mathfrak{h}_i \cap \mathfrak{h}_i \neq \mathfrak{h}_i$ and $g\mathfrak{h}_i \cap \mathfrak{h}_i \neq \emptyset$. In this case we say that $g\mathfrak{h}_i$ *crosses* $\mathfrak{h}_i$.

For a given hyperplane $\mathfrak{h} \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_n\}$, let $L_\mathfrak{h}$ denote the set $\{g \in G | g\mathfrak{h} \text{ crosses } \mathfrak{h}\}$ and $H_\mathfrak{h}$ the group $\text{stab}_G(\mathfrak{h}) = \{g \in G | g\mathfrak{h} = \mathfrak{h}\}$.

Lemma 3.3. *Let G be a group which acts properly and cocompactly on a CAT(0) cube complex X. Then for any hyperplane $\mathfrak{h} \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_n\}$, $L_\mathfrak{h} = \text{stab}_G(\mathfrak{h}) F_\mathfrak{h} \text{stab}_G(\mathfrak{h})$ for some finite set $F_\mathfrak{h} \subset G$.*

Proof. Suppose G acts properly on X and let $\mathfrak{h}$ be any hyperplane of X. Any compact subset K of the hyperplane $\mathfrak{h}$ is a compact subset of X, hence $\{g \in \text{stab}_G(\mathfrak{h}) | gK \cap K \neq \emptyset\} \subset \{g \in G | gK \cap JK \neq \emptyset\}$ is a finite set, and the action of $\text{stab}_G(\mathfrak{h})$ on $\mathfrak{h}$ is proper.

Suppose G acts cocompactly on X, and let $\mathfrak{h}$ be any hyperplane in X. Let C be a compact subset of X such that $GC = X$. Any midplane in the hyperplane equivalence class $\mathfrak{h}$ must be the image of a midplane in C. Let M be the set of midplanes in C which are mapped by some $g \in G$ to a midplane of $\mathfrak{h}$. For each $M_i \in M$ choose a $g_i \in GH$ such that $g_i M_i \in \mathfrak{h}$. Then the set $C' = \bigcup_{M_i \in M} g_i M_i$ is the union of a finite set of midplanes of $\mathfrak{h}$ and hence is a compact subset of $\mathfrak{h}$. We claim $H_\mathfrak{h} C' = \mathfrak{h}$.

Suppose M' is a midplane in $\mathfrak{h}$. Then $M' = gM_i$ for some $g \in G$ and some $M_i \in M$. Then $g_i(g^{-1}M) = g_i M_i \in C'$ and $g_i g^{-1} \in H_\mathfrak{h}$ since any element of the group which maps a midplane in $\mathfrak{h}$ to a midplane in $\mathfrak{h}$ stabilises $\mathfrak{h}$. Hence $H_\mathfrak{h} C' = \mathfrak{h}$ and the action of $H_\mathfrak{h}$ on $\mathfrak{h}$ is cocompact.

If $g \in L_{\mathfrak{h}}$ then $g\mathfrak{h}$ crosses $\mathfrak{h}$. Hence for some $h_1, h_2 \in H_{\mathfrak{h}}$, gh_1C' and h_2C' intersect but are not equal. We will say that subsets which intersect but are not equal cross.

Since gh_1C' and h_2C' cross, so do $h_2^{-1}gh_1C'$ and C'. Since G acts properly on X and C' is a compact subset of X, the set $F_{\mathfrak{h}} = \{f \in G | fC' \text{ crosses } C'\} \subseteq \{f \in G | fC' \cap C' \neq \emptyset\}$ is finite. We have shown that if $g \in L_{\mathfrak{h}}$, then $g \in H_{\mathfrak{h}}F_{\mathfrak{h}}H_{\mathfrak{h}}$, that is $L_{\mathfrak{h}} \subseteq H_{\mathfrak{h}}F_{\mathfrak{h}}H_{\mathfrak{h}}$.

Suppose $g \in H_{\mathfrak{h}}F_{\mathfrak{h}}H_{\mathfrak{h}}$. Then for some $h_1, h_2 \in H_{\mathfrak{h}}$ and some $f \in F_{\mathfrak{h}}$, $g = h_1fh_2$. Then $g\mathfrak{h}$ crosses $\mathfrak{h}$ if and only if $h_1fh_2\mathfrak{h} = h_1f\mathfrak{h}$ crosses $\mathfrak{h}$. Since $h_1 \in \text{stab}_G(\mathfrak{h})$, $h_1^{-1} \in \text{stab}_G(\mathfrak{h})$ and $h_1f\mathfrak{h}$ crosses $\mathfrak{h}$ if and only if $f\mathfrak{h}$ crosses $\mathfrak{h}$. By the definition of the set $F_{\mathfrak{h}}$, $f\mathfrak{h}$ crosses $\mathfrak{h}$, hence $g\mathfrak{h}$ crosses $\mathfrak{h}$ and $g \in L_{\mathfrak{h}}$. Hence $H_{\mathfrak{h}}F_{\mathfrak{h}}H_{\mathfrak{h}} \subset L_{\mathfrak{h}}$ and we have $L_{\mathfrak{h}} = H_{\mathfrak{h}}F_{\mathfrak{h}}H_{\mathfrak{h}} = \text{stab}_G(\mathfrak{h})F_{\mathfrak{h}}\text{stab}_G(\mathfrak{h})$ as required. □

Lemma 3.4. *Let G be a group which acts properly and cocompactly on a CAT(0) cube complex such that* $\text{stab}_G(\mathfrak{h})$ *is separable for each hyperplane* $\mathfrak{h}$*. Then for each $\mathfrak{h} \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_n\}$ there exists a finite index subgroup $K_{\mathfrak{h}}$ of G containing* $\text{stab}_G(\mathfrak{h})$ *such that for all $k \in K_{\mathfrak{h}}$, $k\mathfrak{h}$ does not cross $\mathfrak{h}$.*

Proof. By the hypothesis $\text{stab}_G(\mathfrak{h}) = H_{\mathfrak{h}}$ is separable for any $\mathfrak{h}$, so we have $H_{\mathfrak{h}} = \bigcap H_j$ where, for each j, H_j is a finite index subset of G. Hence for all $g \in G \setminus H_{\mathfrak{h}}$, there exists a H_j with $g \notin H_j$.

As $F_{\mathfrak{h}} = \{f \in G | fC' \text{ crosses } C'\}$ if $f \in F_{\mathfrak{h}}$ then $f\mathfrak{h}$ crosses $\mathfrak{h}$, and so if $f \in F_{\mathfrak{h}}$, f does not stabilise $\mathfrak{h}$. Hence $F_{\mathfrak{h}} \cap H_{\mathfrak{h}} = \emptyset$ and for each $f \in F_{\mathfrak{h}}$ we can choose a finite index H_j not containing f. We intersect these to form a subgroup $K_{\mathfrak{h}}$ which contains $H_{\mathfrak{h}} = \text{stab}_G(\mathfrak{h})$ and contains no element of $F_{\mathfrak{h}}$. Since $L_{\mathfrak{h}} = H_{\mathfrak{h}}F_{\mathfrak{h}}H_{\mathfrak{h}}$, it follows that $K_{\mathfrak{h}}$ contains no element of $L_{\mathfrak{h}}$, and hence for all $k \in K_{\mathfrak{h}}$, $k\mathfrak{h}$ does not cross $\mathfrak{h}$. Since $F_{\mathfrak{h}}$ is finite and each H_j is finite index in G, it follows that $K_{\mathfrak{h}}$ is a finite index subgroup of G. □

Lemma 3.5. *Let G be a group which acts properly and cocompactly on a CAT(0) cube complex X such that* $\text{stab}_G(\mathfrak{h})$ *is separable for any hyperplane $\mathfrak{h}$. Then we can find a finite index normal subgroup N of G such that $m\mathfrak{h}_i$ does not cross $\mathfrak{h}_i$ for any $m \in N$, $\mathfrak{h}_i \in \{\mathfrak{h}_1, \mathfrak{h}_2, \ldots, \mathfrak{h}_n\}$.*

Proof. Taking $M = \bigcap_i K_{\mathfrak{h}_i}$ gives a subgroup of G such that $h\mathfrak{h}_i$ does not cross $\mathfrak{h}_i$ for any $h \in M, \mathfrak{h}_i \in \{\mathfrak{h}_1, \mathfrak{h}_2, \ldots, \mathfrak{h}_n\}$. M is an intersection of a finite number of finite index subgroups of G, and hence is finite index in G.

Let $N = \bigcap_{g\in G} M^g$. Since M is finite index in G, there are a finite number of subgroups conjugate to M in G. Hence the intersection $\bigcap_{g\in G} M^g$ is an intersection of a finite number of finite index subgroups, and so N is a finite index normal subgroup of G. By definition, for all $m \in N$, $\mathfrak{h}_i \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_n\}$, $m\mathfrak{h}_i$ does not cross $\mathfrak{h}_i$.

□

Remark 3.6. Since N is finite index in G, we can choose a finite set of coset representatives for N in G, $\{\gamma_1, \ldots, \gamma_l\}$. Let $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ denote the finite set of hyperplanes $\{\gamma_i\mathfrak{h}_j | i \in 1, \ldots, l, j \in 1, \ldots, n\}$. Then the set $\{N\mathfrak{h}_1, \ldots, N\mathfrak{h}_k\}$ includes all the hyperplanes of the cube complex X.

Note that for any $m \in N$ and any hyperplane $\mathfrak{h}$ in X $m\mathfrak{h}$ does not cross $\mathfrak{h}$. This follows from the fact that N is a normal subgroup: For any $\mathfrak{h}$ we can write $\mathfrak{h} = g\mathfrak{h}_i$ for some $g \in G$ and $\mathfrak{h}_i \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_n\}$ and so for any $m \in N$

$$m\mathfrak{h} \cap \mathfrak{h} = mg\mathfrak{h}_i \cap g\mathfrak{h}_i = g(m'\mathfrak{h}_i) \cap g\mathfrak{h}_i = g(m'\mathfrak{h}_i \cap \mathfrak{h}_i)$$

for some $m' \in N$, which by lemma 3.5 is either the empty set or the hyperplane $g\mathfrak{h}_i$.

3.1.2 Using N to construct a product of trees

Given a group G with a proper, cocompact action on a CAT(0) cube complex X such that $\mathrm{stab}_G(\mathfrak{h})$ is separable for every hyperplane $\mathfrak{h}$, we will construct an embedding of X in a product of trees, by first constructing for each $\mathfrak{h}_i$ a tree T_i by considering $N\mathfrak{h}_i$. We will use the method described in the proof of lemma 1.38. We begin by defining a halfspace system.

Definition. Given any hyperplane $\mathfrak{h}_i \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ consider the set of hyperplanes $N\mathfrak{h}_i$. For each $n \in N$, $n\mathfrak{h}_i$ separates X into two connected components which we denote by $X_{n\mathfrak{h}_i}$ and $X^*_{n\mathfrak{h}_i}$. We denote the set of halfspaces obtained in this way by $H_i = \{X_{n\mathfrak{h}_i}, X^*_{n\mathfrak{h}_i} | n \in N\}$. We consider the triple

$(H_i, \leq, *)$ where $\leq$ is the order given by inclusion of halfspaces and $*$ is the order reversing involution given by interchanging the two halfspaces defined by any hyperplane.

We define the boundary map $\partial : H_i \to H_i / \sim$ to be the map which takes halfspaces to their boundaries, i.e. $\partial(X_{n\mathfrak{h}_i}) = \{X_{n\mathfrak{h}_i}, X^*_{n\mathfrak{h}_i}\} = \partial(X^*_{n\mathfrak{h}_i})$. For simplicity of notation, we equate the equivalence class $\{X_{n\mathfrak{h}_i}, X^*_{n\mathfrak{h}_i}\}$ with the hyperplane $n\mathfrak{h}_i$.

Lemma 3.7. *$(H_i, \leq, *)$ is a halfspace system.*

Proof. By lemma 1.36 the set of halfspaces of a CAT(0) cube complex form a halfspce system. H_i is a subset of the halfspaces H of X, with the partial order $\leq$ and the involution $*$ on H_i agreeing with the partial order and involution on H. Suppose X_1 and X_2 are any two elements of H_i. Since $H_i \subset H$ the set $\{X_3 \in H_i | X_1 \leq X_3 \leq X_2\}$ is contained in the set $\{X_3 \in H | X_1 \leq X_3 \leq X_2\}$ and hence is finite. Similarly, since the partial order on H_i agrees with the partial order on H, at most one of the inequalities $X_1 \leq X_2,\ X_1 \leq X_2^*,\ X_1^* \leq X_2,\ X_1^* \leq X_2^*$ holds. These two observations on the properties of triple $(H, \leq, *)$ show that it is a halfspace system. $\square$

Lemma 3.8. *For each $i \in \{1, \ldots, k\}$ let $(H_i, \leq, *)$ be the halfspace system defined above. Then the components of the cube complex corresponding to $(H_i, \leq, *)$ (as defined in lemma 1.37) are trees. There is an injective map ξ_1 from the set of components of $X \setminus N\mathfrak{h}_i$ into the vertex set of one of these trees.*

Proof. For any $i \in \{1, \ldots, k\}$ we construct a CAT(0) cube complex C_i using $(H_i, \leq, *)$ as follows: Take the set of vertices to be all sections for ∂ such that $v(n_1\mathfrak{h}) \not\leq v(n_2\mathfrak{h})^*$ for any $n_1\mathfrak{h},\ n_2\mathfrak{h} \in H_i$.

Let $\mathfrak{h} = \mathfrak{h}_i$. We define a map ξ_i from the set of components of $X \setminus N\mathfrak{h}$ in the set of vertices by mapping each component D of $X \setminus N\mathfrak{h}$ to the section $\xi_i(D)$ defined by setting $\xi_i(D)(n\mathfrak{h})$ to be the halfspace $X_{n\mathfrak{h}}$ or $X^*_{n\mathfrak{h}}$ containing D. Since D is non-empty, the section $\xi_i(D)$ will satisfy $\xi_i(D)(n_1\mathfrak{h}) \not\leq \xi_i(D)(n_2\mathfrak{h})^*$ for all $n_1\mathfrak{h}, n_2\mathfrak{h} \in N$ and hence is a vertex of C_i

We join two vertices u and v by an edge if and only if the values of the sections u and v differ on exactly one hyperplane. If two components D and

D' are adjacent in $X \setminus N\mathfrak{h}$ then they are separated by exactly one hyperplane $\mathfrak{h}$. Hence the values of the sections $\xi_i(D)$ and $\xi_i(D')$ differ on exactly one boundary $\{X_{\mathfrak{h}}, X_{\mathfrak{h}}^*\}$, and so by definition the vertices $\xi_i(D)$ and $\xi_i(D')$ are adjacent in C'.

Choose any component of $X \setminus N\mathfrak{h}$ and denote it by E. Let T_i be the component of C_i containing $\xi_i(E)$, the vertex corresponding to the component E. X is connected, and any two vertices are separated by at most finitely many hyperplanes. Hence for any component D of $X \setminus N\mathfrak{h}$, the vertex $\xi_i(D)$ also lies in T_i and so the canonical map from $X \setminus N\mathfrak{h}$ to C_i gives an embedding in a single component of C_i.

To see that T_i is a tree, suppose that T_i contains a cycle. Then there is a finite set of hyperplanes $H' = \{n_1\mathfrak{h}, \ldots, n_k\mathfrak{h}\}$ such that for each $n_i\mathfrak{h}$, there is a pair of vertices v, v' in the cycle such that the values of v and v' on $n_i\mathfrak{h}$ differ.

Consider the finite set of halfspaces $\{X_{n_i\mathfrak{h}}, X_{n_i\mathfrak{h}}^* | n_i\mathfrak{h} \in H'\}$. Since $\leq$ is an order on the set of all halfspaces from this set, we can choose a minimal halfspace in this set, i.e a halfspace $X_{n\mathfrak{h}}$ with $n\mathfrak{h} \notin H'$ such that for all $n_i\mathfrak{h} \in H'$ $X_{n_i\mathfrak{h}} \nleq X_{n\mathfrak{h}}$ and $X_{n_i\mathfrak{h}}^* \nleq X_{n\mathfrak{h}}$. Without loss of generality, let this halfspace be $X_{n_1\mathfrak{h}}$. By the definition of the set H', there exists a vertex v in the cycle such that $v(n_1\mathfrak{h}) = X_{n_1\mathfrak{h}}$. Suppose that the vertex u is adjacent to v and lies in the cycle. The values of u and v differ on precisely one hyperplane, call it $n\mathfrak{h}$. Then $u(n\mathfrak{h}) = v(n\mathfrak{h})^*$.

Suppose that $n\mathfrak{h} \neq n_1\mathfrak{h}$. As $n\mathfrak{h}$ lies in the set H' we know that $X_{n\mathfrak{h}} \nleq X_{n_1\mathfrak{h}}$ and $X_{n\mathfrak{h}}^* \nleq X_{n_1\mathfrak{h}}$. Suppose that $v(n\mathfrak{h}) = X_{n\mathfrak{h}}$. By the definition of a vertex we have $X_{n\mathfrak{h}} = v(n\mathfrak{h}) \nleq v(n_1\mathfrak{h})^* = X_{n_1\mathfrak{h}}^*$ and $X_{n\mathfrak{h}}^* = v(n\mathfrak{h})^* = u(n\mathfrak{h}) \nleq u(n_1\mathfrak{h})^* = X_{n_1\mathfrak{h}}^*$. Similarly, taking $v(n\mathfrak{h}) = X_{n\mathfrak{h}}^*$ yields the same relations. However, at least one of the relations $X_{n\mathfrak{h}} \leq X_{n_1\mathfrak{h}}$, $X_{n\mathfrak{h}} \leq X_{n_1\mathfrak{h}}^*$, $X_{n\mathfrak{h}}^* \leq X_{n_1\mathfrak{h}}$, $X_{n\mathfrak{h}}^* \leq X_{n_1\mathfrak{h}}^*$ must hold, and hence we must have $n\mathfrak{h} = n_1\mathfrak{h}$. This means that the vertex defined by $u(n_1\mathfrak{h}) = v(n_1\mathfrak{h})^*, u(n_i\mathfrak{h}) = v(n_i\mathfrak{h}) \forall n_i\mathfrak{h} \in H' \setminus \{n_1\mathfrak{h}\}$ is the only vertex adjacent to u which lies in the cycle. This contradicts the existence of such a cycle, hence T_i is a tree. $\square$

By our choice of subgroup N the action of N on X maps hyperplanes in

$N\mathfrak{h}_i$ to hyperplanes in $N\mathfrak{h}_i$. Hence the action of N on the set $\{X_{n\mathfrak{h}_i}, X^*_{n\mathfrak{h}_i} | n \in N\}$ maps halfspaces to halfspaces in such a way that, for any $m \in N$, $m(X_{n\mathfrak{h}_i})$ is either $X_{(mn)\mathfrak{h}_i}$ or $X^*_{(mn)\mathfrak{h}_i}$. We define an action of N on T_i by taking the value of $m(v)$ on the hyperplane $n\mathfrak{h}_i$ to be $m(v(n\mathfrak{h}_i))$ for each $m, n \in N$ and each vertex v in T_i.

We check that the resulting section $m(v)$ is a vertex as follows:

$$\begin{aligned} v(n_i\mathfrak{h}) \not\leq v(n_2\mathfrak{h})^* \;&\Rightarrow\; v(n_1\mathfrak{h}) \cap v(n_2\mathfrak{h}) \neq \emptyset \\ &\Rightarrow\; m(v(n_1\mathfrak{h}) \cap v(n_2\mathfrak{h})) \neq \emptyset \\ &\Rightarrow\; m(v(n_1\mathfrak{h})) \cap m(v(n_2\mathfrak{h})) \neq \emptyset \\ &=\; m(v(n_1\mathfrak{h})) \not\leq m(v(n_2))^* \end{aligned}$$

Hence $m(v)$ is a vertex of T_i. Similarly, we can show that if $\xi_i(D)$ is the image in T_i of the component D of $X \setminus N\mathfrak{h}_i$ then $m(\xi_i(D)) = \xi_i(m(D))$. Let Y denote the product of trees $T_1 \times \ldots \times T_n$. Then there is a natural action of N on Y resulting from the action of N on T_i.

Lemma 3.9. *Suppose N is a normal subgroup of a finitely generated group G with index n and that N acts on the left on a product of trees Y. Then G acts on the product of n copies of Y.*

Proof. In order to show this, we use Serre's construction as described in "Groups acting on graphs", [13]. We have N normal in G with finite index n. Let $Y_1, \ldots, Y_n$ be copies of Y. We have an action of N on each Y_i. Let $x_1, \ldots, x_n$ be left coset representatives for N in G. G acts on the left on the set of cosets $\{x_1N, \ldots, \ldots, x_nN\}$. Define an action of G on the index set $\{1, \ldots . n\}$ by $gi = j$ if and only if $g(x_iN) = x_jN$.

Consider any $g \in G$. For each $i \in \{1, \ldots, n\}$ there is a unique expression $gx_i = x_{gi}h_i$ where $gi \in \{1, \ldots, n\}$ and $h_i \in N$. Let $(w_1, \ldots, w_n)$, $w_i \in Y_i$ represent a general point of $Y_1 \times \ldots \times Y_n$. We define

$$g(w_1, \ldots, w_n) = (h_{g^{-1}1}w_{g^{-1}1}, \ldots, h_{g^{-1}n}w_{g^{-1}n})$$

Clearly $e(w_1, \ldots, w_n) = (w_1, \ldots, w_n)$. To see that this is an action, it remains to check that $f(g(w_1,\ldots,w_n)) = (fg)(w_1,\ldots,w_n)$.

For each $i \in \{1,\ldots,n\}$ there is a unique expression $gx_i = x_{gi}h_i$ and a unique expression $fx_i = x_{fi}k_i$ where $gi,\ fi \in \{1,\ldots,n\}$ and $h_i,\ k_i \in N$. Then $(fg)x_i = f(gx_i) = f(x_{gi}h_i) = (fx_{gi})h_i = x_{fgi}k_{gi}h_i = x_{fgi}j_i$ since multiplication in the group is associative, and we have the expression $j_i = k_{gi}h_i$. In order to check that we have an action, we will need the following expression for $j_{(fg)^{-1}i}$;

$$j_{(fg)^{-1}i} = j_{g^{-1}f^{-1}i} = k_{g(g^{-1}f^{-1}i)}h_{g^{-1}f^{-1}i} = k_{(gg^{-1})f^{-1}i}h_{g^{-1}f^{-1}i} = k_{f^{-1}i}h_{g^{-1}f^{-1}i}.$$

Then

$$\begin{aligned}
f(g(w_1,\ldots,w_n)) &= f(h_{g^{-1}1}w_{g^{-1}1},\ldots,h_{g^{-1}n}w_{g^{-1}n}) \\
&= f(y_1,\ldots,y_n) \text{ where } y_i = h_{g^{-1}i}w_{g^{-1}i} \\
&= (k_{f^{-1}1}y_{f^{-1}1},\ldots,k_{f^{-1}n}y_{f^{-1}n}) \\
&= (k_{f^{-1}1}h_{g^{-1}f^{-1}1}w_{g^{-1}f^{-1}1},\ldots,k_{f^{-1}n}h_{g^{-1}f^{-1}n}w_{g^{-1}f^{-1}n}) \\
&= (j_{g^{-1}f^{-1}1}w_{g^{-1}f^{-1}1},\ldots,j_{g^{-1}f^{-1}n}w_{g^{-1}f^{-1}n}) \\
(fg)(w_1,\ldots,w_n) &= (j_{(fg)^{-1}1}w_{(fg)^{-1}1},\ldots,j_{(fg)^{-1}n}w_{(fg)^{-1}n})
\end{aligned}$$

Hence G acts on the finite product $Y \times \ldots \times Y$ as required. □

Lemma 3.10. *Let G be a group which acts properly and cocompactly by isometries on a finite dimensional, locally finite CAT(0) cube complex X such that* $\mathrm{stab}_G(\mathfrak{h})$ *is separable for any hyperplane $\mathfrak{h}$. Then for some $k \in \mathbb{N}$ there is an isometric map from X to a product of trees $T_1 \times \ldots \times T_k$*

Proof. By lemma 3.5 and remark 3.6 there exists a finite index subgroup N and a finite set of hyperplanes $\mathfrak{h}_1,\ldots,\mathfrak{h}_k$ such that the action of N on the set of hyperplanes generates all hyperplanes of X and such that, for all $m \in N$ and any hyperplane $\mathfrak{h} \in X$, $m\mathfrak{h}$ does not cross $\mathfrak{h}$.

Each vertex x in X is uniquely defined by the set of half spaces containing it. Since each hyperplane in X is the image under the action of N of some unique hyperplane in the set $\{\mathfrak{h}_1,\ldots,\mathfrak{h}_k\}$, each pair of halfspaces $X_{\mathfrak{h}}, X^*_{\mathfrak{h}}$ is contained in the set H_i for some unique i. Hence for each vertex x the set of halfspaces $\{X_{\mathfrak{h}}^{(*)}|x \in X_{\mathfrak{h}}^{(*)}\}$ can be decomposed into the disjoint union of the sets $\{X_{\mathfrak{h}}^{(*)} \in H_i|x \in X_{\mathfrak{h}}^{(*)}\}$, $i = 1, \ldots, k$. Hence each vertex in X is uniquely defined by the list of components $D_1, \ldots, D_k$ containing it, where $D_i \in X \setminus N\mathfrak{h}_i$.

Define the map from the vertex set of the cube complex X to the vertex set of the product of trees $T_1 \times \ldots \times T_k$ defined as follows: for each vertex x in X and each $i \in \{1, \ldots, k\}$ let $D_i(x)$ be the component of $X \setminus N\mathfrak{h}_i$ containing x. As defined in lemma 3.8 for each i there is an injective map $\xi_i : X \setminus N\mathfrak{h}_i \mapsto T_i\xi$. We define $\xi : X^(0) \mapsto T_1 \times \ldots \times T_k$ by $\xi(x) = \xi_1(D_1(x)), \ldots, \xi_k(D_k(x))$.

If u and u' are adjacent vertices in X then they are separated by exactly one hyperplane. Hence they map to vertices in $T_1 \times \ldots \times T_k$ which differ in only one co-ordinate and are adjacent in that co-ordinate tree. Extending this to a shortest edge path between any two vertices in X, we see that the distance between two vertices is precisely the number of hyperplanes separating them. Similarly, we have one edge in the shortest edge path between corresponding vertices in the product of trees for each of these hyperplanes. Hence on the level of edge metrics distance is preserved and so $X^{(1)}$ embeds d_1-isometrically in the one skeleton of $T_1 \times \ldots \times T_k$.

Since both the cube complex X and the product of trees $T_1 \times \ldots \times T_k$ are CAT(0), the map ξ extends to a map from n-cubes to n-cubes for all n which is d_2-isometric. To see this, note the map ξ preserves incidence of edges and consider the image under ξ of the 1-skeleton of a n-cube in X where $n \geq 2$. If $n = 2$ then, since $T_1 \times \ldots \times T_k$ is CAT(0) and hence simply connected, the image of the 1-skeleton of every 2-cube in X must be the 1-skeleton of a 2-cube in $T_1 \times \ldots \times T_k$. If $n > 2$, then since $T_1 \times \ldots \times T_k$ is non-positively curved the image of the 1-skeleton of every n-cube in X is the 1-skeleton of an n-cube in $T_1 \times \ldots \times T_k$. Similarly, we can show that for any pair of vertices $\xi(u)$ and $\xi(v)$ in the image of X in $T_1 \times \ldots \times T_k$, if n is the minimal value for which $\xi(u)$ and $\xi(v)$ are vertices in the boundary of an n-cube in $T_1 \times \ldots \times T_k$, then u and v lie in the boundary of an n-cube in X.

□

Corollary 3.11. *Let G be a group which acts freely, properly and cocompactly by isometries on a finite dimensional, locally finite CAT(0) cube complex X*

such that $\mathrm{stab}_G(\mathfrak{h})$ *is separable for each hyperplane* $\mathfrak{h}$ *of* X*. Then* G *embeds quasi-isometrically in a finite product of trees.*

Proof. By lemma 1.42, for any such G and X there is a quasi-isometric embedding of G in X. By lemma 3.10 there is an isometric map from X to a product of trees $T_1 \times \ldots \times T_k$. The composition of these two maps gives a quasi-isometric map from the group G to the product of trees $T_1 \times \ldots \times T_k$. $\square$

3.2 Embeddings in a product of finitely branching trees

We now want to construct a product of locally finite trees into which X will embed quasi-isometrically. The following construction is based on the work of Dranishnikov and Schroeder in [17].

We begin by looking at trees in a different way to the previous section, and consider rooted trees. Let $(Q) = Q_1, Q_2, \ldots$ be a sequence of non-empty sets. We associate to (Q) the *rooted simplicial tree* $T_{(Q)}$ as follows: The set of vertices is the set of finite sequences $(q_1, \ldots, q_\alpha)$ with $q_i \in Q_i$. The empty sequence defines the root vertex and is denoted by $v_\emptyset$. We denote the vertex given by $(q_1, \ldots, q_\alpha)$ as $v_{(q_1,\ldots,q_\alpha)}$. Two vertices are connected by an edge in $T_{(Q)}$ if their lengths as sequences differ by one and the shorter can be obtained by erasing the last term of the longer.

Let $v = v_{(q_1,\ldots,q_\alpha)}$ and $u = v_{(q'_1,\ldots,q'_{\alpha'})}$. Then there exists a unique integer r, such that $q_i = q'_i$ for all $i \leq r$, and such that $q_{r+1} \neq q'_{r+1}$. Then $d(v, u) = (\alpha' - r) + (\alpha - r)$.

The vertex $v_\emptyset$ has $|Q_1|$ neighbours and every vertex of distance $i > 0$ from $v_\emptyset$ has $|Q_{i+1}| + 1$ neighbours. The tree $T_{(Q)}$ is *locally compact* if and only if Q_i is finite for all i.

Throughout this section we will assume that G acts properly and co-compactly by isometries on a CAT(0) cube complex X in such a way that $\mathrm{stab}_G(\mathfrak{h})$ is separable for every hyperplane. Let N denote a finite index normal subgroup of G such that, for all $n \in N$, $n\mathfrak{h}$ does not intersect $\mathfrak{h}$, and let $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ be a set of hyperplanes such that for any hyperplane $\mathfrak{h}$ in X $\mathfrak{h} = n\mathfrak{h}_i$ for some $n \in N$ and some unique $\mathfrak{h}_i \in \{\mathfrak{h}_i, \ldots, \mathfrak{h}_k\}$.

For ease of notation, we will refer to vertices in X by elements of the group G, as if there was an embedding of G in X. In the case where there is no such embedding, or where a vertex is not the image of an element of G, the same arguments apply with the "identity"being a chosen vertex in X and γ, γ' being any pair of vertices.

Define the *length of the vertex* γ to be the number of hyperplanes crossed by a geodesic from the identity in X to γ. We denote the length of γ by $l(\gamma)$.

Lemma 3.12. *If $\mathfrak{h} = n\mathfrak{h}_i$ for some $n \in N$ then any two shortest edge paths from the identity to $\mathfrak{h}$ in X must cross the same set of images of $\mathfrak{h}_i$ under N.*

Proof. First note that a shortest edge path crosses each hyperplane at most once. This follows from the properties of CAT(0) cube complexes, in which geodesics are unique and hyperplanes are geodesically convex.

Let p be a shortest edge path from the identity to $n\mathfrak{h}_i$, then p crosses a set of images of $\mathfrak{h}_i$ which can be listed as $n_1\mathfrak{h}_i, n_2\mathfrak{h}_i, \ldots, n_\alpha\mathfrak{h}_i$ with $n_j \in N \backslash \{n\}$. Any hyperplane $n_j\mathfrak{h}_i$, $n_j \in N$ in X separates the cube complex into two connected components which we will denote by $X_{n_j\mathfrak{h}_i}$ and $X^*_{n_j\mathfrak{h}_i}$. Without loss of generality, label the components so that the identity lies in $X_{n_j\mathfrak{h}_i}$. By our choice of N, there is no $n_j \in N$ such that $n_j\mathfrak{h}_i$ intersects $n\mathfrak{h}_i$, hence if the path p crosses $n_j\mathfrak{h}_i$ we can say that $n_j\mathfrak{h}_i$ separates the identity from $n\mathfrak{h}_i$, that is the identity lies in $X_{n_j\mathfrak{h}_i}$ and $n\mathfrak{h}_i$ lies in$X^*_{n_j\mathfrak{h}_i}$. Hence any shortest path from the identity to $n\mathfrak{h}_i$ must cross each of the hyperplanes $n_1\mathfrak{h}_i, n_2\mathfrak{h}_i, \ldots, n_\alpha\mathfrak{h}_i$.

Suppose a path p' crosses a hyperplane $m\mathfrak{h}_i,\ m \in N\backslash\{n\}$ not crossed by p. Since the path p from the identity e to the hyperplane $n\mathfrak{h}_i$ does not cross $m\mathfrak{h}_i$, both e and some midplane in the hyperplane equivalence class $n\mathfrak{h}_i$ must lie in $X_{m\mathfrak{h}_i}$. By our choice of subgroup N in which n and m are contained, $n\mathfrak{h}_i$ and $m\mathfrak{h}_i$ do not cross, and hence every midplane in $n\mathfrak{h}_i$ lies in $X_{m\mathfrak{h}_i}$. Since $X_{m\mathfrak{h}_i}$ is geodesically convex, any shortest path from e to $n\mathfrak{h}_i$ must be entirely contained in $X_{m\mathfrak{h}_i}$. Hence since the path p' crossed $m\mathfrak{h}_i$, it is not a shortest path from e to $n\mathfrak{h}_i$.

Hence any two shortest paths between the identity vertex and the hyperplane $\mathfrak{h}$ cross the same set of images of $\mathfrak{h}_i$ under N. $\square$

Every hyperplane in X is the image under an element of N of exactly one hyperplane in the set $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$. For each $i \in \{1, \ldots, k\}$ and each vertex γ we define $l_i(\gamma)$ to be the number of times a geodesic from the identity to γ crosses a hyperplane which is in the orbit of $\mathfrak{h}_i$ under N. By lemma 3.12, l_i is well defined. Note that $l(\gamma) = \sum_{i=1}^{k} l_i(\gamma)$.

If $\mathfrak{h}$ is a hyperplane then there is a unique $\mathfrak{h}_i \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ such that $\mathfrak{h}$ is an image of $\mathfrak{h}_i$ under the action of N, and the *level* of $\mathfrak{h}$, denoted $lev(\mathfrak{h})$, is the number of images of that unique $\mathfrak{h}_i$ intersected by a shortest length path from the origin to $\mathfrak{h}$. Lemma 3.12 shows that $lev(\mathfrak{h})$ is well defined.

Denote by $\mathcal{H}^i$ the hyperplanes which are images of $\mathfrak{h}_i$ under N, and by $\mathcal{H}^i_m$ the set of hyperplanes in $\mathcal{H}^i$ with level m. We consider the tree $T_{\mathcal{H}^i}$ belonging to the sequence $(\mathcal{H}^i) = \mathcal{H}^i_1, \mathcal{H}^i_2, \ldots$. In general the set $\mathcal{H}^i_m$ will be infinite, so $T_{(\mathcal{H}^i)}$ will be an infinitely branching tree.

For each i we define a map $\phi^i : X^{(0)} \to T_{(\mathcal{H}^i)}$ by $\phi^i(\gamma) = v_{(g_1\mathfrak{h}_i, g_2\mathfrak{h}_i, \ldots, g_{l_i(\gamma)}\mathfrak{h}_i)}$ where $g_k \in N$ and $(g_1\mathfrak{h}_i, g_2\mathfrak{h}_i, \ldots, g_{l_i(\gamma)}\mathfrak{h}_i)$ is the ordered sequence of translates of $\mathfrak{h}_i$ through which the geodesic from the origin to γ passes. Note that for all i, the identity vertex is mapped to the root vertex of $T_{(\mathcal{H}_i)}$.

By construction, $\phi^i(\gamma) \in T_{(\mathcal{H}^i)}$ and $lev(g_\alpha\mathfrak{h}_i) = \alpha$, the distance from the root vertex $v_\emptyset$ to the vertex determined by the sequence $(g_1\mathfrak{h}_i, g_2\mathfrak{h}_i \ldots, g_\alpha\mathfrak{h}_i)$. As we would hope, the distance from $v_\emptyset$ to $\phi_i(\gamma)$ in $T_{(\mathcal{H}^i)}$ is equal to $l_i(\gamma)$.

Definition. Let $g_1\mathfrak{h}_i$ and $g_2\mathfrak{h}_i$ be images of the hyperplane $\mathfrak{h}_i$. Then the distance between them is defined to be the number of images of $\mathfrak{h}_i$ separating them, that is the number of hyperplanes $n\mathfrak{h}_i$ such that either $g_1\mathfrak{h}_i \in X_{n\mathfrak{h}_i}$ and $g_2\mathfrak{h}_i \in X^*_{n\mathfrak{h}_i}$ or $g_1\mathfrak{h}_i \in X^*_{n\mathfrak{h}_i}$ and $g_2\mathfrak{h}_i \in X_{n\mathfrak{h}_i}$.

Lemma 3.13. *For any hyperplane $\mathfrak{h}_i$ in $\{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ and any $m \in \mathbb{N}$ there exists a map $fin^i_m : \mathcal{H}^i_m \to F^i_m$, where F^i_m is a finite set, such that $fin^i_m(g_1\mathfrak{h}_i) = fin^i_m(g_2\mathfrak{h}_i)$ only if either $g_1\mathfrak{h}_i = g_2\mathfrak{h}_i$ or $d(g_1\mathfrak{h}_i, g_2\mathfrak{h}_i) \geq 4nm$.*

Proof. By hypothesis, $\mathrm{stab}_G(\mathfrak{h}_i)$ is separable in G, that is $\mathrm{stab}_G(\mathfrak{h}_i)$ can be written as an intersection of finite index subgroups of G.

Let $\nu = 4nm$. Then there exists a finite set of hyperplanes which are images of $\mathfrak{h}_i$ under N and which are at a distance of less than ν from $\mathfrak{h}_i$. We choose a finite set of elements $\{n_1, \ldots, n_\alpha\}$ in N such that $\{n_1\mathfrak{h}_i, \ldots, n_\alpha\mathfrak{h}_i\}$ is a list of these hyperplanes (not including $\mathfrak{h}_i$ itself). Note that the choice of these n_i is not unique.

For each $n_j \in \{n_1, \ldots, n_\alpha\}$ $n_j\mathfrak{h}_i \neq \mathfrak{h}_i$, that is $n_j \notin \mathrm{stab}_G(\mathfrak{h}_i)$ and we can choose a finite index subgroup H_j of N such that $\mathrm{stab}_N(\mathfrak{h}_i) \subset H_j$ and $n_j \notin H_j$. Hence there exists a finite group F_j and a homomorphism $\sigma_j : N \to F_j$ satisfying $\sigma_j(n_j) \notin \sigma_j(\mathrm{stab}_N(\mathfrak{h}_i))$.

We define $\sigma_{\mathfrak{h}_i} : N \to F_1 \times \ldots \times F_\alpha,\ n \mapsto (\sigma_1(n), \ldots, \sigma_\alpha(n))$. Then $\forall n_j \in \{n_1, \ldots, n_\alpha\},\ \sigma_{\mathfrak{h}_i}(n_j) \notin \sigma(\mathrm{stab}_N(\mathfrak{h}_i))$.

We define $\mathrm{fin}^i_m : \mathcal{H}^i_m \to F^i_m$ by $fin^i_m(n\mathfrak{h}_i) = \sigma_{\mathfrak{h}_i}(n)$.

Choose any $g_1, g_2 \in N$ with $d(g_1\mathfrak{h}_i, g_2\mathfrak{h}_i) < 4nm$ and $g_1\mathfrak{h}_i \neq g_2\mathfrak{h}_i$. Then $d(g_2^{-1}g_1\mathfrak{h}_i, \mathfrak{h}_i) < 4nm$, and $g_2^{-1}g_1 \in \{n_1, \ldots, n_\alpha\}$.

Suppose $\mathrm{fin}^i_m(g_1\mathfrak{h}_i) = \mathrm{fin}^i_m(g_2\mathfrak{h}_i)$. Let $e_{F^i_m}$ denote the identity element of the group F^i_m. Then since $\sigma_{\mathfrak{h}_i}$ is a homomorphism

$$\begin{aligned}
& \sigma_{\mathfrak{h}_i}(g_1) = \sigma_{\mathfrak{h}_i}(g_2) \\
\implies \quad & \sigma_{\mathfrak{h}_i}(g_2^{-1})\sigma_{\mathfrak{h}_i}(g_1) = e_{F^i_m} \\
\implies \quad & \sigma_{\mathfrak{h}_i}(g_2^{-1}g_1) = e_{F^i_m} \\
\implies \quad & \sigma_{\mathfrak{h}_i}(g_2^{-1}g_1) \in \sigma_{\mathfrak{h}_i}(\mathrm{stab}_N(\mathfrak{h}_i)).
\end{aligned}$$

which is a contradiction. Hence we must have either $g_1\mathfrak{h}_i = g_2\mathfrak{h}_i$ or $d(g_1\mathfrak{h}_i, g_2\mathfrak{h}_i) \geq 4nm$.

□

For simplicity we use the notation fin instead of fin^i_m if the indices are clear from the context. For each hyperplane $\mathfrak{h}_i \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ we consider the locally compact tree $T_{(F^i)}$ coming from the sequence $(F^i) = F^i_1, F^i_2, \ldots$. We consider the map $\psi_i : X^{(0)} \to T_{(F^i)}$ defined by

$$\psi_i(\gamma) = v_{(\mathrm{fin}(g_1\mathfrak{h}_i), \ldots, \mathrm{fin}(g_{l_i(\gamma)}\mathfrak{h}_i))}$$

where

$$\phi^i(\gamma) = v_{(g_1\mathfrak{h}_i,\ldots,g_{l_i(\gamma)}\mathfrak{h}_i)}.$$

We define

$$\psi = \prod_{i=1}^{k} \psi_i : X^{(0)} \to \prod_{i=1}^{k} T_{(F^i)}$$

We now show that this map is a quasi-isometry. First recall that $d(\gamma, \gamma')$ is the number of hyperplanes crossed by a geodesic from γ to γ'. The function ϕ^i gives a list of hyperplane images of $\mathfrak{h}_i$ crossed by this geodesic, so we have $d(\gamma, \gamma') = \sum_{i=1}^{k} d(\phi^i(\gamma), \phi^i(\gamma'))$. Applying the map fin^i_m may identify two hyperplanes, but does not create any new hyperplanes. Hence we have $d(\psi_i(\gamma), \psi_i(\gamma')) \leq d(\phi^i(\gamma), \phi^i(\gamma')$ for every i and $d(\gamma, \gamma') \geq \sum_{i=1}^{k} d(\psi_i(\gamma), \psi_i(\gamma')) = d(\psi(\gamma), \psi(\gamma'))$.

In order to establish an upper bound for the distance in the image, we need the following.

Definition. Let X be a CAT(0) cube complex and $X^{(0)}$ the vertex set of X. For $x, y \in X^{(0)}$ the *median* of x and y, denoted $[x, y]$, is given by $[x, y] = \{z \in X^{(0)} | d(x, y) = d(x, z) + d(z, y)\}$.

If $z \in [x, y]$ then we say z is *between* x and y.

Lemma 3.14. ([30], [8]) *If X is a CAT(0) cube complex then given any 3 vertices $x, y, z \in X^{(0)}$ there is a unique vertex m such that $m \in [x, y] \cap [x, z] \cap [y, z]$*

Lemma 3.15. *Let $\gamma, \gamma' \in N$ and let $d(\psi(\gamma), \psi(\gamma')) = r$ then $d(\gamma, \gamma') \leq 8nr$.*

Proof. By lemma 3.14 there exists an element α between γ and γ' such that α also lies between both 1 and γ and 1 and γ', as shown in figure 3.1. We now consider a geodesic $\alpha = \alpha_0, \ldots, \alpha_\tau = \gamma$ from α to γ and a geodesic $\alpha = \alpha'_0, \ldots, \alpha'_{\tau'} = \gamma'$ from α to γ'. Since α lies between γ and γ', $d(\gamma, \gamma') = d(\gamma, \alpha) + d(\alpha, \gamma') = \tau + \tau'$.

The edges $e_i = [\alpha_i, \alpha_{i-1}]$ and $e'_i = [\alpha'_{i-1}, \alpha'_i]$ are oriented edges of the cube complex and the path $e_\tau, \ldots, e_1, e'_1, \ldots, e'_{\tau'}$ is a geodesic from γ to γ'.

We can assume without loss of generality that $\tau > \tau'$. Let τ_i be the number of hyperplanes in the orbit of $\mathfrak{h}_i$ under N which are crossed by the

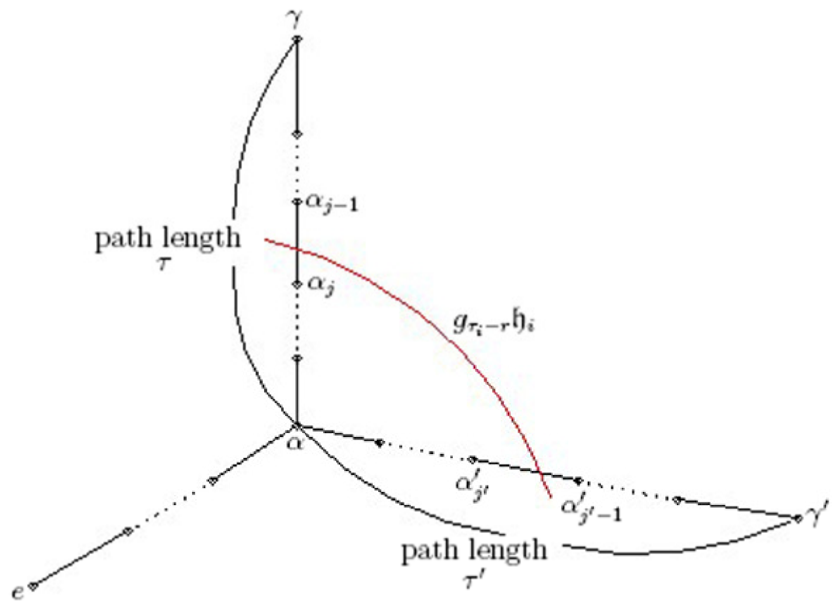

Figure 3.1: Geodesics between e, γ and γ'

geodesic path $e_1, \ldots, e_\tau$. Choose the hyperplane $\mathfrak{h}_i \in \{\mathfrak{h}_1, \ldots, \mathfrak{h}_k\}$ in such a way that τ_i is maximal. If $\tau_i \leq 4r$ then $d(\gamma, \gamma') = \tau + \tau' \leq 2n\tau_i \leq 8nr$ and we are done.

Thus we can assume $\tau_i > 4r$. Consider the images of γ and γ' under the map ψ. Since $r = d(\psi(\gamma), \psi(\gamma'))$, we must have $r \geq d(\psi_i(\gamma), \psi_i(\gamma'))$.

Now, $\psi_i(\gamma) = (\text{fin}(a_1), \ldots, \text{fin}(a_p), \text{fin}(g_1), \ldots, \text{fin}(g_{\tau_i}))$ where $(a_1, \ldots, a_p)$ is the ordered list of images of $\mathfrak{h}_i$ crossed by the geodesic from e to α and $(g_1, \ldots, g_{\tau_i})$ the ordered list of images of $\mathfrak{h}_i$ crossed by the geodesic from α to γ.

Similarly, $\psi_i(\gamma') = (\text{fin}(a_1), \ldots, \text{fin}(a_p), \text{fin}(g'_1), \ldots, \text{fin}(g'_{\tau'}))$ where $(g'_1, \ldots, g'_{\tau'})$ is a ordered list of the images of $\mathfrak{h}_i$ crossed by the geodesic from α to γ'.

Since $r \geq d(\psi_i(\gamma), \psi_i(\gamma'))$, there must be a subsequence $(\text{fin}(a_1), \ldots, \text{fin}(a_p), \text{fin}(g_1), \ldots, \text{fin}(g_\beta))$ of $\psi_i(\gamma)$ such that $\text{fin}(g_i) = \text{fin}(g'_i)\ \forall i \leq \beta$ and such that $(\tau_i - \beta) + (\tau'_i - \beta) \leq r$. It follows that $\beta \geq \tau_i - r$, and hence $\text{fin}(g_{\tau_i - r}) = \text{fin}(g'_{\tau_i - r})$.

We claim that $g_{\tau_i - r} = g'_{\tau_i - r}$. Note that $lev(g_{\tau_i - r}) = p + \tau_i - r \geq \tau_i - r$ and recall that $d(g_{\tau_i-r}\mathfrak{h}_i, g'_{\tau_i-r}\mathfrak{h}_i) \leq d(\gamma,\gamma') < \tau + \tau' \leq 2n\tau_i$. If $g_{\tau_i-r}\mathfrak{h}_i \neq g'_{\tau_i-r}\mathfrak{h}_i$ then by proposition 3.13 $d(g_{\tau_i-r}\mathfrak{h}_i, g'_{\tau_i-r}\mathfrak{h}_i) \geq 4nm$ where m is the level of the hyperplanes.

Hence

$$d(g_{\tau_i-r}\mathfrak{h}_i, g'_{\tau_i-r}\mathfrak{h}_i) \geq 4nm \geq 4n(\tau_i - r) \geq 4n\left(\tau_i - \frac{\tau_i}{4}\right) = 3n\tau_i$$

which contradicts $d(g_{\tau_i-r}, g'_{\tau_i-r}) < 2n\tau_i$.

Thus $g_{\tau_i-r} = g'_{\tau_i-r}$, and there exists a pair of edges $e_j = [\alpha_{j-1}, \alpha_j] \in \{e_1, \ldots, e_\tau\}$ and $e'_{j'} = [\alpha'_{j'-1}, \alpha'_{j'}]$ both intersecting the hyperplane g_{τ_i-r}, (see figure 3.1). We know that the hyperplane g_{τ_i-r} cuts the complex into two totally convex pieces, U and U^*. Without loss of generality we assume α_{j-1} and $\alpha'_{j'-1}$ are contained in U. Then the geodesic $\alpha_{j-1}, \ldots, \alpha_0 = \alpha'_0, \ldots, \alpha'_{j'-1}$ is completely contained in U, and in particular $\alpha \in U$. Now α_j and $\alpha'_{j'}$ are contained in U^*, and by the same argument the complete geodesic $\alpha_j, \ldots, \alpha_0 = \alpha'_0, \ldots, \alpha'_{j'}$ is contained in U^*. Hence $\alpha \in U \cap U^* = \emptyset$. This is a contradiction, so we must have $\tau_i \leq 4r$. We have seen that if $\tau_i \leq 4r$ then $d(\gamma,\gamma') \leq 8nr$, hence the proof is complete. □

We have observed that since the map fin^i_m does not create any hyperplanes, $d(\gamma,\gamma') \geq d(\psi(\gamma),\psi(\gamma'))$. Combining this observation with lemma 3.15, we have shown that the map ψ is quasi-isometric.

We can now prove the theorem.

Theorem 3.1. *Let G be a group which acts, isometrically, properly and cocompactly on a finite dimensional locally finite CAT(0) cube complex X such that* $\mathrm{stab}_G(\mathfrak{h})$ *is separable for each hyperplane $\mathfrak{h}$ of X. Then X embeds quasi-isometrically in a finite product of locally finite trees.*

Proof. We can construct a map ψ from $X^{(0)}$ to the vertex set of a finite product of finitely branching trees. By lemmas 3.13 and 3.15 the map ψ is a quasi-isometry. Quasi-isometric maps between the vertex sets of finite dimensional CAT(0) cube complexes naturally extend to quasi-isometric maps on the entire complex, hence we have a quasi-isometric embedding of X in a finite product of finitely branching trees. □

Given in addition that the action of the group on X is free, we have the following corollary:

Corollary 3.2. *Let G be a group which acts freely, isometrically, properly and cocompactly on a finite dimensional locally finite CAT(0) cube complex X such that* $\mathrm{stab}_G(\mathfrak{h})$ *is separable for each hyperplane $\mathfrak{h}$ of X. Then G embeds quasi-isometrically in a finite product of locally finite trees.*

Proof. By Lemma 1.42 there is a quasi-isometry from the group G to the CAT(0) cube complex X. By Theorem 3.1 we can construct a quasi-isometric map ψ from X to a finite product of finitely branching trees. Hence by composition of maps we have a quasi-isometric embedding of G in a finite product of finitely branching trees. $\square$

Chapter 4

Groups with embeddings in a product of trees

In this chapter we give some examples of groups which satisfy the conditions of corollary 3.2, and hence have quasi-isometric embeddings into a finite product of finitely branching trees. The example of surface groups and the 3-manifold groups mentioned were suggested by the examples of LERF groups in [34].

4.1 Coxeter groups

Let G be a finitely generated Coxeter group. Suppose that for every triple p, q, r of natural numbers G contains only finitely many conjugacy classes of subgroups isomorphic to the (p, q, r)-triangle group . Then G acts isometrically, properly and cocompactly on a CAT(0) cube complex X ([28] ,[40]). Caprace and Mühlherr, [7], showed that G contains only finitely many conjugacy classes of subgroups isomorphic to the (p, q, r)-triangle group if and only if it contains no subgroups isomorphic to the Euclidean triangle groups $\triangle(2, 3, 6)$, $\triangle(2, 4, 4)$ or $\triangle(3, 3, 3)$. By construction, if stab($\mathfrak{h}$) is the stabiliser of a wall in X, then it is equal to the stabiliser of some wall in the Coxeter complex of X. By Corollary 1.17, wall stabilisers of CAT(0) cube complexes are separable.

By lemma 1.38 there is a quasi-isometric embedding of G in X. By theorem 3.1, there is a quasi-isometric embedding of X in a finite product of finitely branching trees. Hence we have:

Corollary 4.1. *Let G be a finitely generated Coxeter group which contains no subgroup isomorphic to a Euclidean triangle group. Then G embeds quasi-isometrically in a finite product of locally finite trees.*

Note that G need not be hyperbolic, since we do not exclude the possibility that G contains affine Coxeter groups, only that it contains affine triangle groups.

4.2 Surface groups

Let G be the fundamental group of a compact, orientable surface F, and X the universal covering space for F. Let g be the genus of the surface F.

Suppose $g = 0$. If F is a compact, orientable surface without boundary, then G is the trivial group. If F is a compact, orientable surface with $n \geq 1$ boundary components then $\pi_1(F) = F_{n-1}$, the free group of rank $n-1$. Hence if F is a genus 0 surface then either G is trivial or G acts on the tree in which each vertex has valency $2n$.

If F is a compact orientable surface of genus 1 without boundary then the covering space X of F is the Euclidean plane. We can choose a Euclidean square in X which is a fundamental region for the action of G on X. This gives a tessellation of X by squares and hence there is a natural isometry from X to a CAT(0) cube complex, on which G acts properly and cocompactly by isometries.

If F is a compact orientable surface with $g \geq 1$ and $s \geq 1$ boundary components then the universal covering space X for F will be the hyperbolic plane. Following the construction of Denvir and Mackay given in [11], we can choose a set of paired geodesics in X in such a way that the compact region bounded by these geodesics is a fundamental region for the action of G on X. In fact, since P lies in the hyperbolic plane P can be chosen to be a regular polygon with at least 5 edges in which all angles are right angles.

The fundamental group G of F is generated by the isometries between the geodesics in the set bounding P. Since P is compact, the action of G on X is cocompact. By the definition of G, if $g_iP \cap P \neq \emptyset$ for some $g_i \in G$ then g_i maps some edge e of P to some edge e' of P. Since there are finitely many edges in the polygon P there are finitely many such g_i in G. Let $I = \{g \in G | gP \cap P \neq \emptyset\}$. Any compact set K in X can be contained in the union $\overline{K}$ of a finite set of copies of P, say $\overline{K} = \bigcup_{g \in \overline{G}} gP$. Let $\{h_1, \ldots, h_k\}$ denote the finite set of elements of G which map g_iP to g_jP for some g_i, g_j in P. If g is such that $g\overline{K} \cap \overline{K} \neq \emptyset$ then g maps some $g_iP \in \overline{K}$ to intersect some $g_jP \in \overline{K}$ and can be written as $h_jg_jig_j^{-1}$ for some $h_j \in H,\ g_j \in \overline{G}$ and $i \in I$. Hence for any compact set $K \subset X$ the set $\{g_i \in G | g_iK \cap K \neq \emptyset\}$ is finite.

We now show that G acts isometrically, properly and cocompactly on a cube complex. Consider the polygonal region P. We divide P into cubes by adding a vertex at the centre of each edge of P and joining each of these to a vertex in the centre of P. We subdivide each copy gP of P in X, and denote the resulting cube complex by $\tilde{X}$. We define a metric on $\tilde{X}$ by defining the length of each edge to be 1 and each square to be isomorphic to a unit square. The metric on $\tilde{X}$ is quasi-isometric to the natural metric on X.

To see that $\tilde{X}$ is CAT(0) we consider the combinatorial link condition on each cell of $\tilde{X}$ (see lemma 1.29). Since the covering space X is either Euclidean or hyperbolic and the angles at the vertices of P are right angles, any cycle in the link of a vertex in $\tilde{X}$ which is the image of a vertex of X contains no cycles of less than 4 edges. Any vertex corresponding to the midpoint of a side in X will also contain no cycles of less than 4 edges in its link. Since P has at least 5 sides , the link of any vertex of $\tilde{X}$ which corresponds to the centre of a copy of P in X has no cycles of less than 5 edges. Hence G acts freely, properly, cocompactly, and isometrically on a CAT(0) cube complex $\tilde{X}$.

By Theorem 1.20 every surface group is locally extended residually finite. Since the G acts properly and cocompactly on $\tilde{X}$ it follows that for each hyperplane $\mathfrak{h}$ in $\tilde{X}$ $\mathrm{stab}_G(\mathfrak{h})$ acts properly and cocompactly on $\mathfrak{h}$ (see the proof of lemma 3.3 for a proof of this fact). By lemma 1.30 $\mathrm{stab}_G(\mathfrak{h})$ is finitely generated for each $\mathfrak{h}$, and since G is LERF it follows that $\mathrm{stab}_G(\mathfrak{h})$ is separable.

Hence by applying corollary 3.2 when $g \geq 1$ we have

Corollary 4.2. *Let G be the fundamental group of a compact, orientable surface. Then G embeds quasi-isometrically in a finite product of locally finite trees.*

4.3 3-manifold groups

Lemma 4.3. *Let G be the fundamental group of the complement in S^3 of the Borromean rings. Then G embeds quasi-isometrically in a finite product of locally finite trees.*

Proof. G is the fundamental group of the complement in S^3 of the Borromean rings. In [37], Thurston showed that the complement of the Borromean ring link can be given a hyperbolic structure coming from a gluing of two ideal octahedra.

Let Γ be the group generated by reflections in the faces of P, where P is a regular octahedron in $\mathbb{H}^3$ all of whose dihedral angles are $\frac{\pi}{2}$. Then G is a subgroup of index 2 of the reflection group Γ.

Γ is a finitely generated right-angled Coxeter group (as defined in section 1.2.5), hence by lemma 1.38 Γ acts properly discontinuously by isometries on a CAT(0) cube complex X. Since Γ is right-angled, any edge in its Coxeter diagram must be labelled by ∞, and hence the Coxeter diagram of Γ contains no affine subdiagram of rank 3. Hence by lemma 1.41 Γ contains only finitely many conjugacy classes of reflection triangles, and by lemma 1.40 Γ acts cocompactly on X. Scott ([34], [35]) proved that Γ is LERF. Hyperplane stabilisers in X are finitely generated, and hence for all hyperplanes $\mathfrak{h}$ $\mathrm{stab}_G(\mathfrak{h})$ is separable.

Applying theorem 3.1 gives a quasi-isometric embedding of Γ in a finite product of finitely branching trees. Since Γ is finitely generated, by lemma 1.32 any finite index subgroup of Γ is quasi-isometric to Γ, hence by composition of quasi-isometries, we have a quasi-isometric map from G to a finite product of finitely branching trees. $\square$

Let A denote the fundamental group of a compact orientable surface F with boundary. We saw in section 4.2 that A acts freely, isometrically, properly and cocompactly on a locally finite CAT(0) cube complex X in such a way that $\mathrm{stab}_A(\mathfrak{h}_X)$ is separable for hyperplane $\mathfrak{h}_X \in X$. Let G be a central extension of an infinite cyclic group $J =< j >$ by A. Since F is a surface with boundary, the group A is free. Hence G is a split central extension of J by A, and G can be written as a direct product $J \times A$. G is the fundamental group of a S^1-bundle over F, a compact Seifert fibre space. We will show that G embeds quasi-isometrically in a finite product of finitely branching trees.

Lemma 4.4. *Let A and B be groups which act freely, isometrically, properly and cocompactly on locally finite CAT(0) cube complexes X and Y respectively, in such a way that $\mathrm{stab}_A(\mathfrak{h}_X)$ and $\mathrm{stab}_B(\mathfrak{h}_Y)$ are separable for every pair of hyperplanes $\mathfrak{h}_X \in X$ and $\mathfrak{h}_Y \in Y$.*

Let G be the direct product $A \times B$. Then G acts freely, isometrically, properly and cocompactly on a locally finite CAT(0) cube complex in such a way that $\mathrm{stab}_G(\mathfrak{h})$ is separable for every hyperplane $\mathfrak{h}$ in the new cube complex.

Proof. For any $g \in G$ we can write g uniquely as (a, b) for some $a \in A$, $b \in B$.

We define the action of G on $X \times Y$ by $g(x, y) = (a, b)(x, y) = (a(x), b(y))$ for any point $(x, y) \in X \times Y$.

Given metrics d_A on A and d_B on B we can define the metric on the group G by $d((a_1, b_1), (a_2, b_2)) = d_A(a_1, a_2) + d_B(b_1, b_2)$. If we take the product metric on $X \times Y$ then the isometric actions of A and B on X and Y respectively give us an isometric action of G on $X \times Y$.

If $K_1 \subset X$ and $K_2 \subset Y$ are compact fundamental regions for the actions of A and B on X and Y respectively, then the region $K_1 \times K_2$ is a fundamental region for the action of G on $X \times Y$, and so we can see that the action of G is cocompact.

Suppose K is a compact subspace of $X \times Y$. Then we can choose compact subspaces $K_1 \subset X$ and $K_2 \subset Y$ such that $K \subset K_1 \times K_2$. Then if $g \in G$

maps K to intersect itself, it follows that g maps $K_1 \times K_2$ to intersect itself. Hence to show that the action of G on $X \times Y$ is properly discontinuous, it is sufficient to show that for any pair of compact subspaces $K_1 \subset X$ and $K_2 \subset Y$, the set $\{g \in G | g(K_1 \times K_2) \cap (K_1 \times K_2) \neq \emptyset\}$ is finite.

Writing g as $(a, b) \in A \times B$,

$$\begin{aligned} & \{(a,b) \in G | (a,b)(K_1 \times K_2) \cap (K_1 \times K_2) \neq \emptyset\} \\ = & \{(a,b) \in G | (a(K_1) \times b(K_2)) \cap (K_1 \times K_2) \neq \emptyset\} \\ = & \{(a,b) \in G | a(K_1) \cap K_1 \neq \emptyset \text{ and } b(K_2) \cap K_2 \neq \emptyset\} \end{aligned}$$

which is finite.

It remains to show that the stabiliser of every hyperplane in $X \times Y$ is separable. Each hyperplane $\mathfrak{h}$ in $X \times Y$ is either of the form $\mathfrak{h}_X \times Y$ where $\mathfrak{h}_X$ is a hyperplane in X or of the form $X \times \mathfrak{h}_Y$ where $\mathfrak{h}_Y$ is a hyperplane in Y.

Consider the case where $\mathfrak{h}$ is $\mathfrak{h}_X \times Y$. The case $X \times \mathfrak{h}_Y$ will follow by similar reasoning. Then

$$\begin{aligned} \mathrm{stab}_G(\mathfrak{h}) &= \{(a,b) | (a,b)(\mathfrak{h}_X \times Y) = \mathfrak{h}_X \times Y\} \\ &= \{(a,b) | a(\mathfrak{h}_X) \times b(Y) = \mathfrak{h}_X \times Y\} \\ &= \{(a,b) | a(\mathfrak{h}_X) = \mathfrak{h}_X \text{ and } b(Y) = Y\} \\ &= \{(a,b) | a \in \mathrm{stab}_A(\mathfrak{h}_X),\ b \in B\} \end{aligned}$$

By hypothesis, $\mathrm{stab}_A(\mathfrak{h}_X)$ is separable in A, that is there exists a collection $\{H_i | i \in I\}$ of finite index subsets of A such that $\mathrm{stab}_A(\mathfrak{h}_X) = \bigcap_{i \in I} H_i$, and hence we have $\mathrm{stab}_G(\mathfrak{h}) = \bigcap_{i \in I} H_i \times B$. $H_i \times B$ is finite index in G, and hence $\mathrm{stab}_G(\mathfrak{h})$ is separable in G.

If $(a, b)(x, y) = (x, y)$ then $a(x) = x$ and $b(y) = y$, hence if the actions of A and B on X and Y respectively are free, we must have $a = e_A$ and $b = e_B$. Hence the action of G on $X \times Y$ is free. $\square$

Corollary 4.5. *Let A and B be groups which act freely, isometrically, properly and cocompactly on locally finite CAT(0) cube complexes X and Y re-*

spectively, in such a way that $\mathrm{stab}_A(\mathfrak{h}_X)$ *and* $\mathrm{stab}_B(\mathfrak{h}_Y)$ *are separable for every pair of hyperplanes* $\mathfrak{h}_X \in X$ *and* $\mathfrak{h}_Y \in Y$*. Let* G *be the direct product* $A \times B$*. Then* G *embeds quasi-isometrically in a finite product of finitely branching trees.*

Proof. Applying corollary 3.2 to the action of G on the CAT(0) cube complex $X \times Y$ as constructed in lemma 4.4, we have the result. $\square$

Lemma 4.6. *Let* G *be a split extension of* $\mathbb{Z}$ *by a finitely generated group* F *which acts isometrically, properly and cocompactly on a cube complex* X *in such a way that* $\mathrm{stab}_F(\mathfrak{h}_X)$ *is separable in* F *for every hyperplane* $\mathfrak{h}_X$ *in* X*. Then* G *embeds quasi-isometrically in a finite product of finitely branching trees.*

Proof. Since G is a split extension of $\mathbb{Z}$ by F, there is a map $f : F \to Aut(\mathbb{Z})$ defined by the extension

$$1 \to \mathbb{Z} \to G \to F \to 1$$

The automorphism group of $\mathbb{Z}$ is $\mathbb{Z}_2$, and hence the kernel of f, which we will denote by K, has index at most 2 in F.

Since F acts isometrically, properly and cocompactly on a cube complex X, there is a isometric, proper cocompact action of K on X. We have $stab_K(\mathfrak{h}_X) = \mathrm{stab}_F(\mathfrak{h}_X) \cap K$ for each $\mathfrak{h}_X$ in X, and so since K is finite index in F $\mathrm{stab}_K(\mathfrak{h}_X)$ is separable in K.

G has an index 2 subgroup G' defined by the extension

$$1 \to \mathbb{Z} \to G' \to K \to 1$$

in which K acts trivially on $\mathbb{Z}$, so G' is a split, central extension $\mathbb{Z} \times K$. There is a tree T on which $\mathbb{Z}$ acts isometrically, properly and cocompactly with $stab_{\mathbb{Z}}(\mathfrak{h}_T)$ separable for every hyperplane $\mathfrak{h}_T$ in T. Hence by lemma 4.5 G' embeds quasi-isometrically in a finite product of finitely branching trees, T. Since G' is finite index in G, by lemma 1.32 G' is quasi-isometric to G and by composition of quasi-isometric embeddings there exists a quasi-isometric embedding of G in a finite product of finitely branching trees. $\square$

List of References

[1] J. W. Anderson. *Hyperbolic Geometry.* Undergraduate Mathematics Series. Springer-Verlag, London, 1999.

[2] M. A. Armstrong. *Basic Topology.* Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1983.

[3] B. H. Bowditch. *A Course on Geometric Group Theory*, volume 16 of *MSJ Memoirs.* Mathematical Society of Japan, Japan, 2006.

[4] M. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *A Series of Comprehensive Studies in Mathematics.* Springer-Verlag, Berlin Heidelberg, 1999.

[5] N. Brodskiy and D. Sonkin. Compression of uniforms embeddings into Hilbert spaces. *Preprint on arxiv:math/5009108v1*, 2005.

[6] K. S. Brown. *Buildings.* Springer-Verlag New York, New York, 1989.

[7] P. Caprace and B. Mühlherr. Reflection triangles in Coxeter groups and biautomaticity. *Journal of Group Theory*, 8:467–489, 2005.

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

[9] D. Long D. Cooper and A. Reid. Infinite Coxeter groups are virtually indicable. *Proc. Edinburgh Math. Soc.*, (2) 41:303–313, 1997.

[10] M. Davis. *The Geometry and Topology of Coxeter Groups.* www.math.ohio-state.edu/∼ mdavis/davisbook.pdf. Not yet Published, 2006.

[11] J. Denvir and R. Mackay. Consequences of contractible geodesics on surfaces. *Transactions of the American mathematical society*, 350(11):4553–4568, 1998.

[12] B. Descartes and P. Ungar. Solution to advanced problem 4526, proposed by Peter Ungar. *The American Mathematical Monthly*, 61:352–353, 1954.

[13] W. Dicks and M. Dunwoody. *Groups acting on graphs*, volume 17 of *Cambridge studies in advanced mathematics*. Cambridge University Press, Cambridge, 1989.

[14] W. Dison. Mostow rigidity. *www.ma.ic.ac.uk/juniorgeometry/Mostow.pdf*, 2005.

[15] A. Dranishnikov. On large scale properties of manifolds. *Preprint on arXiv:math.GT/9912062*, 1999.

[16] A. Dranishnikov and T. Januszkiewicz. Every Coxeter group acts amenable on a compact space. *Topology Proceedings*, 24:135–141, 1999.

[17] A. Dranishnikov and V. Schroeder. Embedding of Coxeter groups in a product of trees. *Preprint on arXiv:math.GR/0402398v1*, 2004.

[18] A. Felikson and P. Tumarkin. On hyperbolic Coxeter polytopes with mutually intersecting facets. *Preprint on arxiv:math/0604248v3*, 2007.

[19] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Advances in Applied Mathematics*, 38:302–323, 2007.

[20] M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in Group Theory*, pages 75–265. Springer-Verlag, New York, 1987.

[21] J. Hempel. *Residual finiteness for 3-manifolds*, volume 111 of *Ann. of Math. Stud.* Princeton Univ. Press, Princeton, NJ, 1987.

[22] R. Howlett. Miscellaneous facts about Coxeter groups. www.maths.usyd.edu.au/res/Algebra/How/anucox.html, 1993. Lectures given at the ANU Group Actions Workshop.

[23] J. E. Humphreys. *Reflection groups and Coxeter groups*, volume 29 of *Cambridge studies in advanced mathematics*. Cambridge University Press, Cambridge, 1990.

[24] D. Long. Immersions and embeddings of totally geodesic surfaces. *Bulletin of the London Mathematical Society*, 19:481–484, 1987.

[25] A. Malcev. On faithful representations of infinite groups of matrices. *Amer. Math. Soc. Transl. Ser. (English translation from Mat. Sb. 8 (1940), 405-422 (Russian))*, (2) 45:1–18, 1965.

[26] G. Niblo and L. Reeves. Groups acting on cat(0) cube complexes. *Geometry and Topology*, 1:1–7, 1997.

[27] G. Niblo and L. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37:621–633, 1998.

[28] G. Niblo and L. Reeves. Coxeter groups act on CAT(0) cube complexes. *Journal of Group theory*, 6:399–413, 2003.

[29] G. Niblo and M. Roller. Groups acting on connected cubes and Kazhdan's property T. *Proc. Amer. Math. Soc.*, 126:693–699, 1998.

[30] B. Nica. Cubulating spaces with walls. *Algebraic and geometric topology*, 4:297–309, 2004.

[31] M. Roller. Poc Sets, Median Algebras and Group Actions. An extended study of Dunwoody's construction and Sageev's Theorem. *Habilitationeschrift, Regensberg*, 1998.

[32] M. Ronan. *Lectures on Buildings*. Academic Press Inc., Boston, Ma, 1989.

[33] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical society (3)*, 71:585–617, 1995.

[34] P. Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc*, (2), 17:555–565, 1978.

[35] P. Scott. Correction to "Subgroups of surface groups are almost geometric". *J. London. Math. Soc*, (2) 32:217–220, 1985.

[36] A. Singh. Bilinear forms. *http://anupamk18.googlepages.com/lecture2.pdf*, 2007.

[37] P. Thurston. *Three-dimensional Geometry and Topology, Volume 1.* Princeton University Press, Princeton, NJ, 1997.

[38] Wikipedia. Coxeter group — wikipedia, the free encyclopedia, 2007. Online; accessed 4 January 2008.

[39] R. Willett. Some note on property A. *Preprint on arXiv:math/0612492v1*, 2007.

[40] B. T. Williams. *Two topics in geometric group theory.* PhD thesis, University of Southampton, 1999.

[41] G. Yu. The Novikov conjecture for groups with finite asymptotic dimension. *Annals of Mathematics*, 147:325–355, 1998.

[42] G. Yu. Baum-Connes conjecture for spaces which admit a uniform embedding into Hilbert space. *Inventiones*, 139:201–240, 2000.