Classifying Cubulations of Tubular Groups

Daniel James Woodhouse

Doctor of Philosophy

Mathematics

McGill University

Montreal,Quebec

2016-06-01

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy.

DEDICATION

For Mum and Dad.

ACKNOWLEDGEMENTS

I am grateful to my parents for sending me to school. I would like to thank Dani Wise for being very generous with his time and ideas, and Stefan Friedl for encouraging me to consider graduate school in the first place. I consider myself lucky to have met Jake Langham, Tom Felix Hamilton, Richard Brooker, Callan McGill, Julian Brough, Liz Buckingham-Jeffery, Monty West, and many other people I did mathematics with at Warwick. During my time at McGill I benefited from the company of Andrew Fiori, Kael Dixon, Mathilde Gerbelli-Gauthier, Kasia Jankiewicz, Alice Pozzi, Rob Graham, Leta Montopoli, Mark Hagen, Janine Hagen, Adam Wilks, Mathieu Carette, Will Cavendish, Hadi Bigdely, (the second) Ben Smith, Joj Helfer, Krista Reimer, Adam Alcolado and many others. I would like to thank Emily Stark again for motivating me to find Example 4.4.1. These kind of things don't happen without friends listening to you complain, so I would also like to thank (the first) Ben Smith, Mary Finnegan, and Sam Brown. Matthew Taylor also gets a special mention since he was probably the first person I talked to about mathematics.

ABSTRACT

We consider groups acting on CAT(0) cube complexes. Typically, an action of a group on a CAT(0) cube complex is obtained using Sageev's *dual cube complex* construction. In the setting of hyperbolic groups, if there are sufficiently many quasiconvex codimension-1 subgroups one can obtain an action of the group on a CAT(0) cube complex that is proper and cocompact. In which case, by a theorem of Agol that builds upon a program of Wise, the quotient will be *virtually special*. In general, however, the dual cube complex might be locally infinite and possibly even infinite dimensional. The quotient may not even be compact or virtually special. The aim of this thesis is to study a natural family of groups that demonstrates the range of potential diversity. A *tubular group* is a group that splits as a graph of groups with $\mathbb{Z}^2$ vertex groups and $\mathbb{Z}$ edge groups. Wise formulated a necessary and sufficient criterion for a tubular group to act freely on a CAT(0) cube complex, but speculated that many of these cubulations would be infinite dimensional. We develop a *dilation* criterion for the walls used to cubulate a tubular group that explicitly determines if the dual cube complex is finite or infinite dimensional. A cubulation of a tubular group will be finite dimensional if none of the walls used to cubulate it are dilated. Otherwise it will be infinite dimensional. Building upon the characterization of finite dimensionality, we analyse the dual cube complex in the finite dimensional case more closely to prove that a tubular group acts freely on a locally finite CAT(0) cube complex if and only if it is virtually special.

ABRÉGÉ

Nous considérons les groupes agissant sur les complexes cubiques CHAT(0) en toute généralité. L'action d'un groupe sur un complexe cubique CHAT(0) est typiquement obtenue par la construction du complexe cubique dual de Sageev. Pour les groupes hyperboliques, en trouvant suffisamment de sous-groupes quasiconvexes de codimension-1, on peut obtenir une action propre et cocompacte du groupe sur un complexe cubique CHAT(0). De plus, d'après un théorème d'Agol, le quotient sera virtuellement spécial. Cependant, en général, le complexe cubique dual peut être de dimension infinie. Le but de cette thèse est d'étudier une famille naturelle de groupes qui présente toute l'étendue de la diversité potentielle des situations. Un groupe tubulaire est un groupe qui se scinde en un graphe de groupes dont les groupes de sommets sont $\mathbb{Z}^2$, et les groupes d'arêtes sont $\mathbb{Z}$. Wise formula une condition nécessaire et suffisante pour qu'un groupe tubulaire agisse librement sur un complexe cubique CHAT(0), mais conjectura que beaucoup de ces cubulations seraient de dimension infinie. Nous développons un critère de dilatation pour les murs utilisés pour cubuler un groupe tubulaire qui détermine explicitement si le complexe cubique dual est de dimension finie ou infinie. La cubulation d'un groupe tubulaire sera de dimension finie si aucun de murs utilisés lors de la cubulation ne sont dilatés. Dans le cas contraire, la cubulation sera de dimension infini. Puis, nous prouvons qu'un groupe tubulaire agit sur un complexe cubique CHAT(0) localement fini si et seulement si celui-ci est virtuellement spécial.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
Introduction

Geometric group theory studies finitely generated groups by considering their actions on metric spaces, especially those with geometry resembling non-positive curvature. Given an action of a group on a space, statements about the geometry of the space and the quality of the action are then treated as insightful comments about the group itself. Conversely, the algebra of a group will either permit or prohibit certain actions, allowing conclusions about the "geometry" of the group.

The conversation usually begins with a finitely generated group and its Cayley graph, but quickly moves on to the action of a fundamental group on the universal cover, general actions on trees, $\mathbb{R}$-trees, buildings, and possible boundaries of all of the above. In Gromov's seminal 1987 essay *Hyperbolic Groups* [16] two geometric notions were discussed, generalising the idea of non-positive and negative curvature to the setting of geodesic metric spaces. The first, the notion of a *δ-hyperbolic geodesic metric space*, was due to Gromov himself. The second, the notion of a *CAT(0) space*, Gromov credited to Cartan, Aleksandrov and Toponogov. This thesis will consider a specialized family of CAT(0) spaces first given by Gromov as simple examples: *CAT(0) cube complexes*.

A CAT(0) space is a geodesic metric space that satisfies the comparison triangle criterion. Basic examples include trees and hyperbolic manifolds, but Gromov provided other examples constructed by gluing together various polyhedrons from either Euclidean or Hyperbolic space. The advantage of such constructions is that verifying a space is CAT(0) reduces to checking a local non-positive curvature criterion and passing to the universal cover to ensure the space is simply connected. In the special case where all the polygons are Euclidean cubes with unit side lengths, the local non-positive condition becomes a purely combinatorial criterion on the *link* of each 0-cube in the complex. These spaces became known as CAT(0) cube complexes.

Since their initial introduction, it has become clear that CAT(0) cube complexes have their own rich and unique geometry that should be studied aside from the broader class of CAT(0) spaces. At the heart of this geometry is the role played by *hyperplanes*, the codimension-1 subspaces bisecting the cubes contained in a CAT(0) cube complex. The fundamental development in this direction were the results developed in the thesis of Sageev which showed that these codimension-1 subspaces correspond to *codimension-1 subgroups* [33]. Sageev showed that a group acting essentially on a CAT(0) cube complex contains a codimension-1 subgroup commensurable to the stabilizer of a hyperplane in the cube complex. Conversely, Sageev gave an explicit construction called the *dual cube complex* that takes as input a group and a collection of codimension-1 subgroups, and outputs a CAT(0) cube complex with a group action such that the codimension-1 subgroup virtually stabilizes a hyperplane. This result, strengthened by Gerasimov [14], and Niblo and Roller [26] makes it possible to say that the existence of an action on a CAT(0) cube complex, the quality of the action, and the qualities of the cube complex itself are in fact statements about the existence and structure of codimension-1 subgroups.

A great deal of progress has been made in recent years *cubulating groups*; that is to say, constructing actions of groups on CAT(0) cube complexes. This progress has been particularly strong for families of Gromov hyperbolic groups. Most of these groups have been cubulated using Sageev's construction in one form or another. They include Coxeter groups [25, 21], random groups [28], limit groups [36], small cancellation groups, some one relator groups [38], diagram groups [12], hyperbolic free-by-cyclic groups [18, 17], hyperbolic three manifold groups [2], and many non-hyperbolic 3-manifold groups [30, 29, 23]. Most impressively, perhaps, is the result of Agol, confirming the conjecture of Wise, that a compact non-positively curved cube complex with hyperbolic fundamental group is virtually special [1]. This result was particularly significant as it positively resolved Thurston's virtual Haken conjecture.

The information given about a group acting on a CAT(0) cube complex is not restricted to the codimension-1 subgroups. In the broader theory of CAT(0) spaces, the *Flat Torus Theorem* [6] shows that for groups acting properly and cocompactly on CAT(0) spaces, virtually $\mathbb{Z}^n$ subgroups stabilize a convex subspace isometric to $\mathbb{E}^n$. Among many other consequences, this means that subgroups with distorted cyclic subgroups such as non-trivial Baumslag-Solitar subgroups cannot act properly and cocompactly on CAT(0) cube complexes.

The Flat Torus Theorem does not hold for actions on CAT(0) space in general. The primary obstruction is guaranteeing that the isometries on the space are *semi-simple*. When a group acts properly and cocompactly on a CAT(0) space, then it is immediate that the group acts semisimply. In fact, by a result of Bridson, all isometries of a finite dimensional CAT(0) cube complex are semisimple [5]. This leaves a problem, since Sageev's construction in general will produce infinite dimensional cube complexes.

Motivated by this issue, Haglund considered isometries of CAT(0) cube complexes via their combinatorial geometry [19]. Although a fixed point or geodesic axis in the CAT(0) metric does not necessarily exist for an isometry, a combinatorial axis exists that is isometrically embedded in the combinatorial metric. Typically, the combinatorial axis is not a convex subcomplex. The existence of such an axis implies that groups that act freely on a CAT(0) cube complex don't have distorted cyclic subgroups. The first set of results in this thesis, which appear in [42], generalize Haglund's combinatorial axis to groups of isometries that are virtually $\mathbb{Z}^n$. We construct an invariant subcomplex that is isometrically embedded in the combinatorial metric and is the finite product of CAT(0) cube complexes quasi-isometric to $\mathbb{R}$. As in the case of Haglund's axis, this subcomplex is not convex, so it is not a Flat Torus Theorem, since the torus is not flat.

One benefit of studying hyperbolic groups is that cocompact cubulations can be guaranteed by producing a collection of quasiconvex codimension-1 subgroups [33]. Outside of this setting, cubulations are expected to be more complicated. A *tubular group* is a group that splits as a finite graph of groups with $\mathbb{Z}^2$ vertex groups, and $\mathbb{Z}$ edge groups. Alternatively, a tubular group can be considered as the fundamental group of a *tubular space*: a graph of spaces with tori as vertex spaces, and cylinders as edge space. Tubular groups have already been studied by Brady and Bridson for their interesting isoperimetric qualities [4], and Cashen determined when two tubular groups are quasi-isometric [10]. Wise [40] considered which tubular groups act freely on a CAT(0) cube complex, and provided a necessary and sufficient criterion for such an action to exist. This criterion is the existence of *equitable sets*: a finite set of elements from each vertex group that allows the construction of *immersed walls* in the tubular space. Sageev's construction can then be applied to obtain a CAT(0) cube complex. Wise showed that a very limited class of such groups could be cocompactly cubulated, and speculated that many of the cubulations would be infinite dimensional.

The second collection of results in this thesis, which appear in [44], determine when the cubulations are infinite dimensional. Let X be a tubular space and let $G = \pi_1 X$. Given an immersed wall $\Lambda \looparrowright X$ obtained from an equitable set, we construct a homomorphism called the *dilation map* $\pi_1\Lambda \to \mathbb{Q}^*$. The dilation map is analogous to the holonomy of a leaf in a foliation. When the dilation map has finite image we say that the immersed wall is *non-dilated*. If all the immersed walls are non-dilated, then the resulting dual cube complex is finite dimensional. Conversely, when the dilation map has infinite image, the immersed wall is *dilated*. If an immersed wall is dilated then the dual cube complex is infinite dimensional. Moreover, we prove that a dilated immersed wall is covered by infinitely many pairwise intersecting walls in the universal cover.

This result should be compared with the results of Rubinstein and Wang, who showed that there exists π_1-injective surfaces in graph manifolds whose lifts to the universal cover are pairwise intersecting. Rubinstein and Wang's interest was in showing that the corresponding subgroups were not separable and indeed the same conclusion can be made about the codimension-1 subgroups corresponding to dilated immersed walls. Yi Liu has also generalised the work of Rubinstein and Wang to closed aspherical 3-manifolds [24].

We can then conclude that tubular groups provide a rich diversity of cubulations: cocompact, finite dimensional, and infinite dimensional. One might ask what possible benefit there might be to knowing that a group has a non-cocompact, finite dimensional cubulation? For tubular groups, the existence of a finite dimensional cubulation will imply that the group has very nice properties that are encapsulated in the notion of being virtually special.

A group is *virtually special* if it has a finite index subgroup that acts on a CAT(0) cube complex, such that the quotient is a *special cube complex*.

Special cube complexes were introduced by Haglund and Wise [20], and they are characterized by the fact that they π_1-injectively map into the Salvetti complex of a right angled Artin group. Much is known about right angled Artin groups and their subgroups (see [11]), so showing a group is virtually special has many immediate corollaries. For example, any virtually special group is a subgroup of $SL_n(\mathbb{Z})$ for some n.

The final collection of results in this thesis, which appear in [43], show that a tubular group is virtually special if and only if it acts freely on a locally finite CAT(0) cube complex. Moreover, we show that a tubular group acts freely on a finite dimensional CAT(0) cube complex if and only if it acts freely on a locally finite CAT(0) cube complex. As a consequence we deduce that if a tubular group acts freely on a finite dimensional or locally finite CAT(0) cube complex, then it virtually acts freely on an n-dimensional CAT(0) cube complex for $n \geq 3$.

CHAPTER 2
CAT(0) Geometry and Groups

In this section we review basic definitions and results in the theory of CAT(0) spaces and CAT(0) cube complexes. We refer the reader to [6] for a full account of the theory of CAT(0) spaces, to [3, 8, 13] for further background on geometric group theory, and [34] for an account of the elementary theory of CAT(0) cube complexes and [39] for an overview of Wise's program.

Let X be a geodesic metric space. Let $x, y, z \in X$, and let $\triangle$ be a geodesic triangle with corners at x, y, z. As $\triangle$ satisfies the triangle inequality, there exists, up to isometry, a *comparison triangle* $\triangle' \subseteq \mathbb{E}^2$ with the same side lengths as $\triangle$. We say that $\triangle$ satisfies the *CAT(0) comparison triangle criterion* if for $a, b \in \triangle$, and a', b' the corresponding points in $\triangle'$, then $\mathsf{d}_X(a, b) \leq \mathsf{d}_{\mathbb{E}}(a', b')$. A geodesic metric space is *CAT(0)*, if a geodesic triangle for all $x, y, z \in X$ satisfies the CAT(0) comparison triangle criterion.

Example 2.0.1. A metric tree must satisfy the CAT(0) comparison triangle criterion since all triangle in a tree are either trivial points or lines, or a tripod. As the product of CAT(0) spaces is CAT(0) we can deduce that the product of trees is similarly CAT(0). In the case of manifolds, it is immediate that $\mathbb{E}^n$ are CAT(0), and an elementary fact for $\mathbb{H}^n$.

We refer to [6] for the following basic results on CAT(0) spaces.

Proposition 2.0.2. *Let X be a CAT(0) space.*

1. *There is a unique geodesic joining any two points in X.*
2. *X is contractible.*

A group G is a *CAT(0) group* if G acts metrically properly and cocompactly on a CAT(0) space X.

Example 2.0.3. Realizing the free group $\mathbb{F}_n$ as the fundamental group of a wedge of n-circles, we obtain an action of $\mathbb{F}_n$ on the universal cover which is a tree. Thus $\mathbb{F}_n$ is CAT(0). By realising a closed surface with genus > 1 as the cocompact quotient of $\mathbb{H}^2$ we deduce that the fundamental group is CAT(0).

Showing that particular groups may or may not be CAT(0) became a source of great interest in geometric group theory. Studying the isometries of CAT(0) spaces has been particularly profitable for understanding which groups are CAT(0). Perhaps the most famous result is the Flat Torus Theorem [6]. A subspace $Y \subseteq X$ is *convex*, for all $x, y \in Y$, the geodesic joining x and y is contained in Y.

Theorem 2.0.4. *Let G be a group acting metrically properly and cocompactly on a CAT(0) space X. Let $A \subseteq G$ be a virtually abelian group, commensurable to $\mathbb{Z}^n$. Then there is a convex subspace $Y \subseteq X$ stabilized by A, such that $Y \cong \mathbb{E}^n$. Moreover, Y realises the minimal translation length of the elements in A.*

This allows us to produce examples of non-CAT(0) groups.

Example 2.0.5. Let G be the Baumslag-Solitar group $BS(1,2) = \langle a, t \mid a = ta^2t^{-1} \rangle$. The subgroup $\langle a \rangle \leq BS(1,2)$ is a distorted cyclic subgroup of G. Suppose that G were a CAT(0) group, the G would act properly and cocompactly on a CAT(0) space X. By the Švarc-Milnor Lemma [6, I.8.19], G is quasi-isometric to X. But, by Theorem 2.0.4 $\langle a \rangle$ stabilizes a convex subspace $L \cong \mathbb{E}^1$ in X, that realises the minimal translation length of a. This would imply that $\langle a \rangle$ is not a distorted subgroup of G.

2.1 CAT(0) Cube Complexes

An n-cube is a geodesic metric space isometric to $[0,1]^n$. A 0-cube is a singleton $\{0\}$. A *face* of an n-cube $C \cong [0,1]^n$ is the subspace obtained from restricting finitely many coordinates to 0 or 1. A *cube complex* is a metric space that is the union of subspaces $\mathcal{C} = \{C_\alpha\}$, where each C_α is isometric to an n-cube, for some n, and the intersection of $C_1, C_2 \in \mathcal{C}$ is a face of C_1 and C_2, that itself belongs in $\mathcal{C}$. A cube complex is essentially a combinatorial construction; we can give X the structure of a CW-complex by letting the n-cubes determine the n-skeleton.

The *link* of a 0-cube $x \in X$ is a simplicial complex such that each n-cube C containing x corresponds to an $(n-1)$-simplex, and if C' is the face of C, then the simplex corresponding to C' is contained in the simplex corresponding to C. Equivalently, the link at $x \in X$ is the simplicial complex obtained by taking the subspace of all points distance $1/2$ from x. The simplicial structure is obtained from its intersection with X, viewed as a CW-complex. Intuitively, we imagine that the link of a 0-cube is its combinatorial unit tangent space. Indeed, if a cube complex is homeomorphic to an n-manifold, then the link of each 0-cube is homeomorphic to an $(n-1)$-sphere.

A simplicial complex is *flag* if every set of n distinct 0-simplicies that are pairwise contained in a 1-simplex, is contained in an $(n-1)$-simplex. A cube complex X is *non-positively curved* if the link of every 0-cube is flag. A *CAT(0) cube complex* is a simply connected, non-positively curved cube complex such that the link of every 0-cube is a flag.

We can obtain a CAT(0) cube complex from a non-positively curved cube complex X by taking the universal cover $\widetilde{X}$. This provides a simple means to construct CAT(0) cube complexes. The following example will prove to be central in this thesis:

Example 2.1.1. The presentation complex of $G = \langle a, b, t \mid [a, b] = 1, a = tbt^{-1} \rangle$ is a non-positively curved cube complex. See Figure 2.1.1 to see the presentation complex alongside the link.

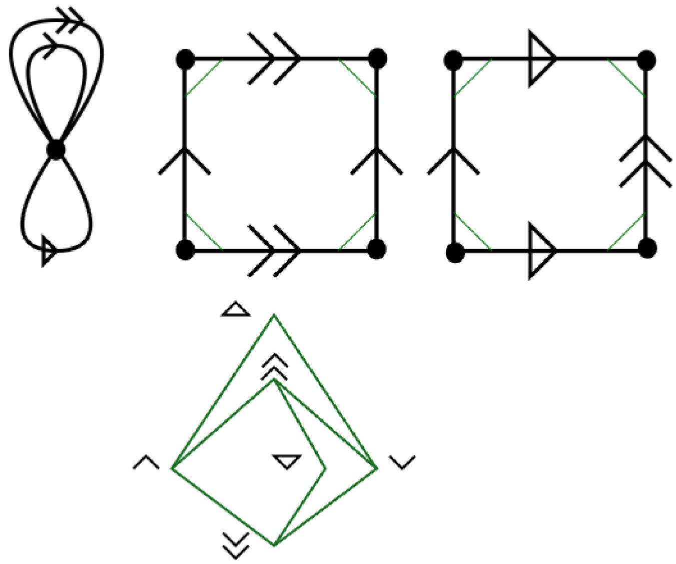

Figure 2–1: The presentation complex (above) and the link (below) of G in Example 2.1.1. The vertices of the link are marked to indicate which edge and orientation it corresponds to.

Although our definition of CAT(0) cube complexes does not involve any comparison triangle criterion, a CAT(0) cube complex is a CAT(0) geodesic metric space. In general, CAT(0) spaces may be constructed from polyhedral cell complexes, satisfying a more sophisticated local non-positive curvature criterion. See [6] for full details. The metric on a CAT(0) cube complex X will be referred to as the *CAT(0) metric* d_X.

As above definition suggests, CAT(0) cube complexes are often best understood as combinatorial constructions. The key to understanding a CAT(0) cube complex combinatorially is via its *hyperplanes*.

Let X be a CAT(0) cube complex. Let $C \cong [0,1]^n$ be an n-cube. The *midcubes* of C are the subspaces obtained by restricting a single coordinate in C to $1/2$. A *hyperplane* Λ in a non-positively curved cube complex X is a minimal, non-empty, subspace that restricts to a set of mid-cubes in each n-cube. A hyperplane in X is *embedded* if its intersection with each cube in X is either a single mid-cube or empty. A hyperplane in X *self-intersects* if it is not embedded. Alternatively, a hyperplane in a CAT(0) cube complex is the subspace of points equidistant in the CAT(0) metric from two adjacent 0-cubes.

Proposition 2.1.2. *If X is a CAT(0) cube complex, then the hyperplanes in X do not self-intersect.*

As hyperplanes in non-positively curved cube complexes may self intersect, we will refer to them as *immersed hyperplanes*.

There exists an alternative metric on the 0-cubes of X, that we will refer to as the *combinatorial metric* d^c_X, sometimes referred to as the *ℓ^1-metric*. The *combinatorial distance* between two 0-cubes in X is the length of the shortest combinatorial path in X joining them. A *combinatorial path* $\gamma : [a,b] \to X$ is a continuous map with $a, b \in \mathbb{Z}$ such that α restricted to $[n, n+1]$ is an isometry with 1-cube in X. A *combinatorial geodesic* $\gamma : [a,b] \to X$ is a combinatorial path such that if $n < m$ are integers in the interval $[a,b]$, then $\mathsf{d}^c_X(\gamma(n), \gamma(m)) = m - n$. The following lemma shows that the combinatorial distance between two 0-cubes is the number of hyperplanes in X separating them.

Lemma 2.1.3. *A combinatorial path γ in X is a combinatorial geodesic if γ intersects each hyperplane at most once.*

We will always assume that a group G acting on a CAT(0) cube complex preserves its cell structure and maps cubes isometrically to cubes. A group G acts without *inversions* if the stabilizer of a hyperplane also stabilizes each complementary component. The requirement that the action be without inversions is not a serious restriction as G acts without inversions on the cubical subdivision.

Alongside hyperplanes are their carriers and halfspaces. The *carrier* $N(\Lambda)$ of a hyperplane Λ is the minimal closed subcomplex of X containing Λ. Note that a hyperplane $\Lambda \subseteq X$ is itself a CAT(0) cube complex.

Lemma 2.1.4. *Let X be a CAT(0) cube complex, and $\Lambda \subseteq X$ a hyperplane. The carrier $N(\Lambda)$ is isometric to $\Lambda \times [0,1]$. Moreover, $N(\Lambda)$ is a convex subcomplex.*

Lemma 2.1.5. *There are precisely two connected components in $X - \Lambda$.*

As the complement $X - \Lambda$ has two components we can partition $X = \overleftarrow{\Lambda} \sqcup \overrightarrow{\Lambda}$ such that $\overleftarrow{\Lambda}$ is the first component which we refer to as the *open halfspace*, and $\overrightarrow{\Lambda}$ is the union of the second component with Λ, which we refer to as the *closed halfspace*. Note that $\overleftarrow{\Lambda} \sqcup \overrightarrow{\Lambda} = X$. This determines what will later be called a *wall* in X. Let $\mathcal{W}$ be the set of walls determined by the hyperplanes in X. Let $\mathcal{L}(\Lambda)$ and $\mathcal{R}(\Lambda)$ denote the minimal subcomplexes containing $\overleftarrow{\Lambda}$ and $\overrightarrow{\Lambda}$ respectively. Let $\mathcal{L}^w(\Lambda)$ and $\mathcal{R}^w(\Lambda)$ denote the maximal subcomplexes contained in $\overleftarrow{\Lambda}$ and $\overrightarrow{\Lambda}$ respectively. Note that $\mathcal{L}(\Lambda), \mathcal{R}(\Lambda), \mathcal{L}^w(\Lambda), \mathcal{R}^w$ are convex subcomplexes.

Let S be a non-empty subset of X. The *combinatorial hull* of S is

$$\mathsf{hull}(S) = \bigcap_{S \subseteq \mathcal{L}(\Lambda)} \mathcal{L}(\Lambda) \bigcap_{S \subseteq \mathcal{R}(\Lambda)} \mathcal{R}(\Lambda).$$

Note that $S \subseteq \mathsf{hull}(S)$, and $\mathsf{hull}(S)$ is a convex subcomplex as it is the intersection of convex subcomplexes. If a group G acts on X, and S is G-invariant, then so is $\mathsf{hull}(S)$.

2.2 Special Cube Complexes

We refer to [20] for full background on special cube complexes.

Let Γ be a simplicial graph. The *right angled Artin group* $G(\Gamma)$ is a group with presentation

$$\langle g_v : v \in V\Gamma \mid [g_v, g_w] = 1 : (v, w) \in \Gamma \rangle.$$

Example 2.2.1. If Γ is the totally disconnected graph with n vertices, then $R(\Gamma)$ is the group with presentation containing n generators and no relations, so $R(\Gamma) \cong \mathbb{F}_n$. If Γ is the complete graph K_n, then all the generators pairwise commute so $R(\Gamma) \cong \mathbb{Z}^n$. If Γ is the complete bipartite graph $K_{n,m}$, then $R(\Gamma) \cong \mathbb{F}_n \times \mathbb{F}_m$.

There is a non-positively curved cube complex $R(\Gamma)$ called the *Salvetti complex* that has $\pi_1 R(\Gamma) = G(\Gamma)$. We obtain $R(\Gamma)$ by taking the presentation complex associated to the above presentation of $G(\Gamma)$ and inserting an n-cube for every n-clique in Γ.

Let X, Y be non-positively curved cube complexes. A cubical map $X \to Y$ is a map that restricts for each n-cube C in X to an isometry from C to an n-cube in Y. A cubical map $\phi : X \to Y$ induces a map $\phi_x : \text{link}(x) \to \text{link}(\phi(x))$ for each 0-cube x in X. A simplicial map has *no missing edges* if two adjacent 0-simplicies in the image also have adjacent preimage. A cubical map $\phi : X \to Y$ is a *local isometry* if ϕ_x has no missing edges for all 0-cubes x in X.

Proposition 2.2.2. *Let $\phi : X \to Y$ be a local isometry, then $\phi_* : \pi_1 X \to \pi_1 Y$ is an injection.*

A non-positively curved cube complex is *special* if it maps by a local isometry into the Salvetti complex of a finitely generated right angled Artin group. A non-positively curved cube complex X is *virtually special* if a finite index cover of X maps by local isometry into a Salvetti complex of a finitely generated right angled Artin group. By Proposition 2.2.2, a non-positively curved cube complex X being virtually special implies that $\pi_1 X$ virtually

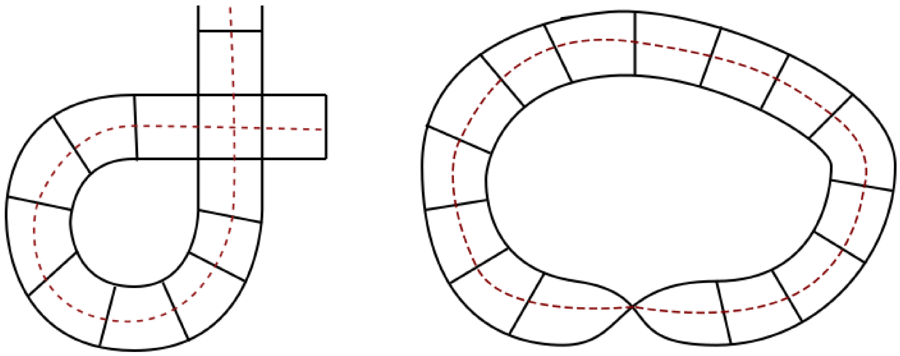

Figure 2–2: An immersed hyperplane self-intersection on the left, and a one sided immersed hyperplane on the right.

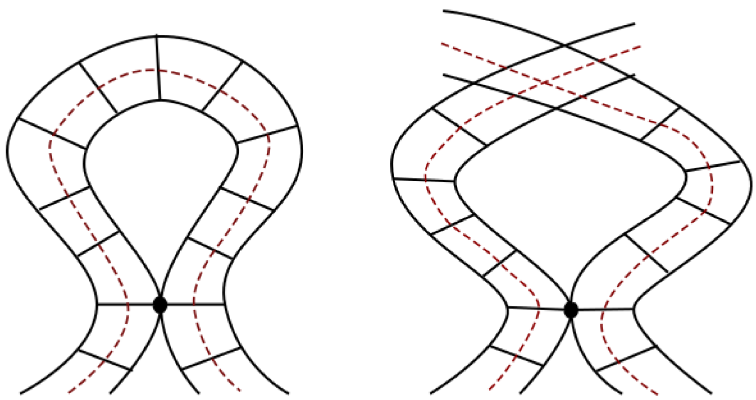

Figure 2–3: An immersed hyperplane self-osculating on the left, and a pair of immersed hyperplanes inter-osculating on the right. The vertex at which the osculation is happening is in bold. Note that if the adjacent 1-cubes to this vertex belong to a common 2-cube, then it is *not* a point of osculation.

embeds inside a finitely generated right angled Artin group. In particular this means that $\pi_1 X \leqslant SL_n(\mathbb{Z})$.

In practice, a non-positively curved cube complex X may be verified as being special by verifying that the immersed hyperplanes are *embedded* (see the left diagram in Figure 2.2) and avoid the following pathologies:

1. Let $\Lambda \subseteq X$ be an embedded hyperplane in X. We say Λ is *2-sided* if the embedding $\Lambda \hookrightarrow X$ can be extended to an embedding $\Lambda \times (-\epsilon, \epsilon) \hookrightarrow X$. See the right diagram in Figure 2.2.
2. If Λ is 2-sided, embedded hyperplane then the 1-cubes intersecting Λ may be consistently oriented. We say that Λ *directly self-osculates* if

there is 0-cube that is either the initial vertex for two distinct edges dual to Λ that aren't contained in a 2-cube, or it is the terminal vertex for two distinct edges dual to Λ that aren't contained in a 2-cube. See the left diagram in Figure 2–3.

3. If Λ, Λ' are hyperplanes in X, then they *inter-osculate* if they intersect in X, and there exists a 0-cube in X that is incident to a 1-cube e dual to Λ, and a 1-cube e' incident to Λ' such that e and e' aren't contained in the same 2-cube. See the right diagram in Figure 2–3.

Theorem 2.2.3. *Let X be a non-positively curved cube complex containing finitely many hyperplanes. Then X is special if and only if*

1. *Each hyperplane embeds in X.*
2. *Each hyperplane is 2-sided.*
3. *No hyperplane directly self osculates.*
4. *No two hyperplanes inter-osculate.*

We say that a group is *special* if it is the fundamental group of a non-positively curved cube complex, and is *virtually special* if it has a finite index subgroup that is special.

2.3 Dual Cube Complexes

We refer to [32] and [22] for full background.

Let S be a set. A *wall* $\Lambda = \{\overleftarrow{\Lambda}, \overrightarrow{\Lambda}\}$ in S is a partition of S into two disjoint, nonempty subsets. The subsets $\overleftarrow{\Lambda}, \overrightarrow{\Lambda}$ are the *halfspaces* of Λ. A wall Λ *separates* $x, y \in S$ if they belong to distinct halfspaces of Λ. Let $K \subseteq S$. A wall Λ *intersects* K if K nontrivially intersects both $\overleftarrow{\Lambda}$ and $\overrightarrow{\Lambda}$. Let $\mathcal{W}$ be a set of walls in S, then $(S, \mathcal{W})$ is a wallspace if for all $x, y \in S$, the number of walls separating x and y is finite. If Λ intersects K, then the *restriction of* Λ *to* K, is the wall in K determined by $\Lambda\big|_K = \{\overleftarrow{\Lambda} \cap K, \overrightarrow{\Lambda} \cap K\}$. We will not permit duplicate walls in $\mathcal{W}$. Let $\mathcal{H}$ be the set of halfspaces determined by $\mathcal{W}$.

Example 2.3.1. Let X be a CAT(0) cube complex, and let $\Lambda \subseteq X$ be a hyperplane in X. The complement $X - \Lambda$ has two components, therefore defining a wall in X such that $\overleftarrow{\Lambda}$ is an open halfspace not containing Λ, and $\overrightarrow{\Lambda}$ is a closed halfspace containing Λ. Note that $\overleftarrow{\Lambda} \sqcup \overrightarrow{\Lambda} = X$. Let $\mathcal{W}$ be the set of walls determined by the hyperplanes in X. Then $(X, \mathcal{W})$ is the wallspace associated to X. Note that we are using Λ to denote both the hyperplane and the wall corresponding to the hyperplane.

A function $c : \mathcal{W} \to \mathcal{H}$ is a 0-*cube* if $c[\Lambda] \in \{\overleftarrow{\Lambda}, \overrightarrow{\Lambda}\}$ and the following two conditions are satisfied:

1. For all $\Lambda_1, \Lambda_2 \in \mathcal{W}$ the intersection $c[\Lambda_1] \cap c[\Lambda_2]$ is nonempty.
2. For all $x \in S$, the set $\{\Lambda \in \mathcal{W} \mid x \notin c[\Lambda]\}$ is finite.

The *dual cube complex* $C(S, \mathcal{W})$ is the connected CAT(0) cube complex obtained as follows: Let the union of all 0-cubes be the 0-skeleton. Two 0-cubes $c_1 \neq c_2$ are endpoints of a 1-cube if $c_1[\Lambda] = c_2[\Lambda]$ for all but precisely one $\Lambda \in \mathcal{W}$. An n-cube is then inserted wherever there is the 1-skeleton of an n-cube. The hyperplanes in $C(S, \mathcal{W})$ are identified naturally with the walls in $\mathcal{W}$. A proof of the fact that $C(S, \mathcal{W})$ is in fact a CAT(0) cube complex can be found in [32].

A point $x \in S$ determines a 0-cube c_x defined such that $x \in c_x[\Lambda]$ for all $\Lambda \in \mathcal{W}$. Condition (1) holds immediately since $x \in c_x[\Lambda]$ for all $\Lambda \in \mathcal{W}$. Condition (2) holds for c_x, since if $y \in S$ a wall Λ does not separate x and y, we can deduce that $y \in c_x[\Lambda]$, hence all but finitely many Λ satisfy $y \in c_x[\Lambda]$. Such 0-cubes are called the *canonical* 0-*cubes*.

CHAPTER 3
Generalizing Haglund's Axis Theorem

3.1 A Preliminary Discussion

A connected CAT(0) cube complex X is a *quasiline* if it is quasiisometric to $\mathbb{R}$. The *rank* of a virtually abelian group commensurable to $\mathbb{Z}^n$ is n. The goal of this chapter will be the following theorem:

Theorem 3.4.3. Let G be virtually $\mathbb{Z}^n$. Suppose G acts properly and without inversions on a CAT(0) cube complex X. Then G stabilizes a finite dimensional subcomplex $Y \subseteq X$ that is isometrically embedded in the combinatorial metric, and $Y \cong \prod_{i=1}^{m} C_i$, where each C_i is a cubical quasiline and $m \geq n$. Moreover, $\mathrm{Stab}_G(\Lambda)$ is a codimension-1 subgroup for each hyperplane Λ in Y.

Note that Y might not be a convex subcomplex, nor even isometrically embedded in the CAT(0) metric.

Corollary 3.1.1. *Let A be a finitely generated virtually abelian group acting properly on a CAT(0) cube complex X. Then A acts metrically properly on X.*

Let g be an isometry of X, and let $x \in X$. The *displacement of g at x*, denoted $\tau_x(g)$, is the distance $\mathsf{d}_X(x, gx)$. The *translation length* of g, denoted $\tau(g)$, is $\inf\{\tau_x(g) \mid x \in X\}$. Similarly, if x is a 0-cube of X, we can define the *combinatorial displacement of g at x*, denoted $\tau_x^c(g)$, as $\mathsf{d}_X^c(x, gx)$ and the *combinatorial translation length*, denoted $\tau^c(g)$, is $\inf\{\tau_x^c(g) \mid x \in X\}$. Note that τ, and τ^c are conjugacy invariant. An isometry g of a CAT(0) space is *semisimple* if $\tau_x(g) = \tau(g)$ for some $x \in X$, and G acts *semisimply* on a CAT(0) space X if each $g \in G$ is semisimple.

If a virtually $\mathbb{Z}^n$ group G acts metrically properly by semisimple isometries on a CAT(0) space X, then the Flat Torus Theorem [6] provides a G-invariant, convex, flat $\mathbb{E}^n \subseteq X$. A virtually abelian subgroup is *highest* if it is not virtually contained in a higher rank abelian subgroup. If G is a highest virtually abelian subgroup of a group acting properly and cocompactly on a CAT(0) cube complex X, then G cocompactly stabilizes a convex subcomplex Y which is a product of quasilines, as above [41]. However, this theorem fails without the highest hypothesis. Moreover, most actions do not arise in the above fashion.

Despite the fact that the flat torus theorem will not hold under the hypotheses of Theorem 3.4.3, we can deduce the following:

Corollary 3.4.4. Let G be virtually $\mathbb{Z}^n$. Suppose G acts properly and without inversions on a CAT(0) cube complex X. Then G cocompactly stabilizes a subspace $F \subseteq X$ homeomorphic to $\mathbb{R}^n$ such that for each hyperplane $\Lambda \subseteq X$, the intersection $\Lambda \cap F$ is either empty or homeomorphic to $\mathbb{R}^{n-1}$.

The initial motivation for Theorem 3.4.3 and Corollary 3.4.4 was to resolve the following question posed by Wise. Although we have not found a combinatorial flat, Corollary 3.4.4 is perhaps better suited to applications (see [43]).

Problem 3.1.2. Let $\mathbb{Z}^2$ act freely on a CAT(0) cube complex Y. Does there exist a $\mathbb{Z}^2$-equivariant map $F \to Y$ where F is a square 2-complex homeomorphic to $\mathbb{R}^2$, and such that no two hyperplanes of F map to the same hyperplane in Y?

A *combinatorial geodesic axis for g* is a g-invariant, isometrically embedded, subcomplex $\gamma \subseteq X$ with $\gamma \cong \mathbb{R}$. Note that γ realizes the minimal combinatorial translation length of g. Theorem 3.4.3 is a high dimensional generalization of Haglund's combinatorial geodesic axis theorem. Haglund's proof involved an argument by contradiction, exploiting the geometry of hyperplanes. We reprove the result in Section 3.5 by using the dual cube complex construction of Sageev. The results are further support for Haglund's slogan "in CAT(0) cube complexes the combinatorial geometry is as nice as the CAT(0) geometry".

The following is an application of Theorem 3.4.3.

Corollary 3.1.3. *Let H be virtually $\mathbb{Z}^n$, and let $\phi : H \to H$ be an injection with $\phi \neq \phi^i$ for all $i > 1$. Then $G = \langle H, t \mid t^{-1}ht = \phi(h) : \ h \in H\rangle$ cannot act properly on a CAT(0) cube complex.*

Proof. Suppose that G acts properly on a CAT(0) cube complex X. After subdividing X we can assume that G acts without inversions. As H is finitely generated, there exists an a in the finite generating set such that $\phi^i(a) \neq a$ for all $i \in \mathbb{N}$, otherwise $\phi^i = \phi$ for some i, contradicting our hypothesis. Thus, $|\{\phi^i(a)\}| = \infty$. By Theorem 3.4.3 there is an H-equivariant isometrically embedded subcomplex $Y \subseteq X$ such that $Y \cong \prod_{i=1}^m C_i$ where each C_i is a cubical quasiline.

As Y is isometrically embedded in X in the combinatorial metric, the combinatorial translation length $\tau^c(\phi^i(a))$ is the same in Y as it is in X. The set $\{\tau^c(\phi^i(a))\}_{i\in\mathbb{N}}$ must be unbounded since the action of H on Y is proper and Y is locally finite. However, since τ^c is conjugacy invariant in G, we conclude that $\tau^c(\phi^i(a)) = \tau^c(\phi^j(a))$ for all $i, j \in \mathbb{N}$. Thus, we arrive at the contradiction that $\{\tau^c(\phi^i(a))\}_{i\in\mathbb{N}}$ is both contant and unbounded. $\square$

The above argument is inspired by the solvable subgroup theorem [6, II.7.8]. However, we have the following example of a solvable group which does act freely on a CAT(0) cube complex.

Example 3.1.4. Let $H = \langle a_1, a_2, \dots \mid [a_i, a_j] : i \neq j\rangle$. Note that H is the fundamental group of the non-positively curved cube complex Y obtained from a 0-cube v, and 1-cubes $e_1, e_2, e_3 \dots$ with n-cubes inserted for every cardinality n collection of 1-cubes to create an n-torus. One should think of Y as an infinite cubical torus. The oriented loop e_i represents the element a_i.

Let $\phi : H \to H$ be the monomorphism such that $\phi(a_i) = a_{i+1}$. Let $G = H*_\phi = \langle t, a_1, a_2, \dots \mid [a_i, a_j] : i \neq j \ , t^{-1}a_i t = a_{i+1} \rangle$ be the associated ascending HNN extension. Note that G is generated by a_1 and t. There is a graph of spaces X obtained by letting Y be the vertex space and $Y \times [0, 1]$ be the edge space and identifying $(v, 1)$ and $(v, 0)$ with v, and the 1-cube $e_i \times \{1\}$ with e_i and $e_i \times \{0\}$ with e_{i+1}. Note that X is nonpositively curved, and therefore $G = \pi_1 X$ acts freely on the CAT(0) cube complex $\widetilde{X}$, the universal cover of X.

3.2 Technical Results Relating to Dual Cube Complexes

Lemma 3.2.1. *Let X be a CAT(0) cube complex. Let $\mathcal{W}$ be a set of walls obtained from the hyperplanes in X. Let Z be a connected subcomplex of X, and let $\mathcal{W}_Z \subseteq \mathcal{W}$ be the subset of walls intersecting Z. Let $\mathcal{V}$ be walls in $\mathcal{W}_Z$ restricted to Z. Then $(Z, \mathcal{V})$ is a wallspace and $C(Z, \mathcal{V})$ embeds in $C(X, \mathcal{W})$ isometrically in the combinatorial metric.*

Proof. We first claim that the map $\mathcal{W}_Z \to \mathcal{V}$ is an injection. Suppose that $\Lambda_1, \Lambda_2 \in \mathcal{W}_Z$ are distinct walls. As Λ_1, Λ_2 intersects Z, and since Z is connected, there are 1-cubes e_1, e_2 in Z that are dual to the hyperplanes corresponding to Λ_1, Λ_2. Therefore, both 0-cubes in e_1 belong in a single halfspace of $\Lambda_2\big|_Z$, so $\Lambda_1\big|_Z \neq \Lambda_2\big|_Z$.

We construct a map $\phi : C(Z, \mathcal{V}) \to C(X, \mathcal{W})$ on the 0-skeleton first. Let c be a 0-cube in $C(Z, \mathcal{V})$. We let $\phi(c) \in C(X, \mathcal{W})$ be the uniquely defined 0-cube such that $\phi(c)[\Lambda] \supseteq c[\Lambda\big|_Z]$ for $\Lambda\big|_Z \in \mathcal{V}$, and $\phi(c)[\Lambda] \supseteq Z$ for $\Lambda \in \mathcal{W} - \mathcal{W}_Z$. To verify that $\phi(c)$ is a 0-cube, first observe that $\phi(c)[\Lambda_1] \cap \phi(c)[\Lambda_2]$ is nonempty since $\Lambda_1\big|_Z \cap \Lambda_2\big|_Z \subseteq X$. Secondly, if $x \in X$ we need to show that $x \in \phi(c)[\Lambda]$

for all but finitely many $\Lambda \in \mathcal{W}$. Choose $z \in Z$, then $z \in c[\Lambda\big|_Z]$ for all $\Lambda\big|_Z \in \mathcal{V} - \{\Lambda_1\big|_Z, \ldots, \Lambda_k\big|_Z\}$, hence $z \in \phi(c)[\Lambda]$ for all $\Lambda \in \mathcal{W}_Z - \{\Lambda_1, \ldots, \Lambda_k\}$. Let $\{\Lambda_{k+1}, \ldots, \Lambda_{k+\ell}\}$ be the set of walls in $\mathcal{W}$ separating x and z. Then $x \in \phi(c)[\Lambda]$ for all $\Lambda \in \mathcal{W} - \{\Lambda_1, \ldots \Lambda_{k+\ell}\}$.

The 0-cubes are embedded since if $c_1 \neq c_2$, there exists $\Lambda\big|_Z \in \mathcal{V}$ such that $c_1[\Lambda\big|_Z] \neq c_2[\Lambda\big|_Z]$, hence $\phi(c_1)[\Lambda] \neq \phi(c_2)[\Lambda]$. If c_1, c_2 are adjacent 0-cubes in $C(Z, \mathcal{V})$, then $c_1[\Lambda\big|_Z] = c_2[\Lambda\big|_Z]$ for all $\Lambda\big|_Z \in \mathcal{V}$, with the exception of precisely one wall $\hat{\Lambda}\big|_Z$. Therefore, we can deduce that $\phi(c_1)[\Lambda] = \phi(c_2)[\Lambda]$ for all walls in $\mathcal{W}$, with the precise exception of $\hat{\Lambda}$. Therefore, the 1-skeleton of $C(Z, \mathcal{V})$ embeds in $C(X, \mathcal{W})$, which is sufficient for ϕ to extend to an embedding of the entire cube complex.

Consider $C(Z, \mathcal{V})$ as a subcomplex of $C(X, \mathcal{W})$. The set of hyperplanes in $C(Z, \mathcal{V})$ embeds into the set of hyperplanes in $C(X, \mathcal{W})$. To see that $C(Z, \mathcal{V})$ is an isometrically embedded subcomplex, let z_1, z_2 be 0-cubes in Z and γ be a geodesic combinatorial path in $C(Z, \mathcal{V})$ joining them. Each hyperplane dual to γ in $C(Z, \mathcal{V})$ intersects γ precisely once, and since the hyperplanes in $C(Z, \mathcal{V})$ inject to hyperplanes in $C(X, \mathcal{W})$, it is geodesic there as well. $\square$

Lemma 3.2.2. *Let S be a set and let $\mathcal{W}$ be a set of walls of S. Let G be a group acting on $(S, \mathcal{W})$. Let $\mathcal{V} \subseteq \mathcal{W}$ be a G-invariant subset. Then there is a G-equivariant function $\phi : C(S, \mathcal{W})^0 \to C(S, \mathcal{V})^0$. Moreover, $\phi^{-1}(z)$ is nonempty for all 0-cubes z in $C(S, \mathcal{V})$.*

Proof. Let c be a 0-cube in $C(S, \mathcal{W})$. Let $\phi(c)[\Lambda] = c[\Lambda]$ for $\Lambda \in \mathcal{V}$. It is immediate that ϕ is G-equivariant.

To verify $\phi(c)[\Lambda]$ is a 0-cube in $C(S, \mathcal{V})$ first note that $\phi(c_1)[\Lambda_1] \cap \phi(c_2)[\Lambda_2] \neq \emptyset$ for all $\Lambda_1, \Lambda_2 \in \mathcal{V}$, since $c_1[\Lambda_1] \cap c_2[\Lambda_2] \neq \emptyset$ for all $\Lambda_1, \Lambda_2 \in \mathcal{W}$. Secondly, for all $x \in S$ observe that $x \in \phi(c)[\Lambda]$ for all but finitely many $\Lambda \in \mathcal{V}$. Indeed, this is true for all but finitely many $\Lambda \in \mathcal{W}$.

To see that $\phi^{-1}(z)$ is non-empty for all 0-cubes z in $C(S, \mathcal{V})$ we determine a 0-cube x in $C(S, \mathcal{W})$ such that $\phi(x) = z$. Fix $s \in S$. Let $x[\Lambda] = z[\Lambda]$ for $\Lambda \in \mathcal{V}$. Suppose that $\Lambda \in \mathcal{W} - \mathcal{V}$. If $\overrightarrow{\Lambda} \supseteq z[\Lambda']$ for some $\Lambda' \in \mathcal{V}$ let $x[\Lambda] = \overrightarrow{\Lambda}$. Similarly if $\overleftarrow{\Lambda} \supseteq z[\Lambda']$. Otherwise, if Λ intersects $z[\Lambda']$ for all $\Lambda' \in \mathcal{V}$ then let $s \in x[\Lambda]$.

To verify that x is a 0-cube, consider the following cases to show $x[\Lambda_1] \cap x[\Lambda_2] \neq \emptyset$ for $\Lambda_1, \Lambda_2 \in \mathcal{W}$. If $\Lambda_1, \Lambda_2 \in \mathcal{V}$ then $x[\Lambda_1] \cap x[\Lambda_2] = z[\Lambda_1] \cap z[\Lambda_2] \neq \emptyset$. Suppose that $\Lambda_1 \in \mathcal{W} - \mathcal{V}$ and $x[\Lambda_1] \subseteq z[\Lambda_1']$ for some $\Lambda_1' \in \mathcal{V}$. If $\Lambda_2 \in \mathcal{V}$, then $x[\Lambda_1] \cap x[\Lambda_2] \supseteq z[\Lambda_1'] \cap z[\Lambda_2] \neq \emptyset$. If $\Lambda_2 \in \mathcal{W} - \mathcal{V}$ and $x[\Lambda_2] \subseteq z[\Lambda_2']$ for some $\Lambda_2' \in \mathcal{V}$ then $x[\Lambda_1] \cap x[\Lambda_2] \subseteq z[\Lambda_1'] \cap z[\Lambda_2'] \neq \emptyset$. If Λ_2 intersects $z[\Lambda]$ for all $\Lambda \in \mathcal{V}$, then $x[\Lambda_1] \cap x[\Lambda_2] \supseteq z[\Lambda_1'] \cap x[\Lambda_2] \neq \emptyset$. Finally if both $s \in x[\Lambda_1]$ and $x[\Lambda_2]$, then their intersection will contain at least s.

Finally, we verify that for $s' \in S$ there are only finitely many $\Lambda \in \mathcal{W}$ such that $s' \notin x[\Lambda]$. Suppose, by way of contradiction, that there is an infinite subset of walls $\{\Lambda_1, \Lambda_2, \ldots\} \subseteq \mathcal{W}$ such that $s' \notin x[\Lambda_i]$ for all $i \in \mathbb{N}$. We can assume, by excluding at most finitely many walls, that each $\Lambda_i \in \mathcal{W} - \mathcal{V}$. Similarly, by excluding finitely many walls, we can assume that Λ_i does not separate s and s'. Therefore, $s \notin x[\Lambda_i]$ for $i \in \mathbb{N}$. Therefore, by construction of x, there exist $\Lambda_i' \in \mathcal{V}$ such that $z[\Lambda_i'] \subseteq x[\Lambda_i]$, which implies that $s' \notin z[\Lambda_i']$. There are infinitely many distinct Λ_i', as otherwise there is a $\Lambda' \in \mathcal{V}$ such that $z[\Lambda'] \subseteq x[\Lambda_i]$ for infinitely many i, which would imply that infinitely many Λ_i separate s' from an element in the complement of $z[\Lambda']$. Therefore, infinitely many distinct walls $\Lambda_i' \in \mathcal{V}$ have $s' \notin z[\Lambda_i']$, contradicting that z is a 0-cube in $C(S, \mathcal{V})$. □

3.3 Minimal $\mathbb{Z}^n$-invariant convex subcomplexes

The following Theorem is found in [14] (or less explicitly in [26]).

Theorem 3.3.1. *Let G be a finitely generated group that acts on a CAT(0) cube complex X without a fixed point or inversions. Then there is a hyperplane in X that is stabilized by a codimension-1 subgroup of G.*

The goal of this section is to prove the following:

Lemma 3.3.2. *Let G be a finitely generated group acting without fixed point or inversions on a CAT(0) cube complex X. There exists a minimal, G-invariant, convex subcomplex $X_o \subseteq X$ such that X_o contains only finitely many hyperplane orbits, and every X_o hyperplane stabilizer is a codimension-1 subgroup of G.*

Proof. Since G is finitely generated, by taking the convex hull of a G-orbit we obtain a G-invariant convex subcomplex $X_o \subseteq X$ containing finitely many G-orbits of hyperplanes. Assume that X_o is a minimal such subcomplex in terms of the number of hyperplane orbits.

Let $(X, \mathcal{W})$ be the wallspace obtained from the hyperplanes in X. Suppose that $\mathrm{Stab}_G(\Lambda)$ is not a codimension-1 subgroup of G for some $\Lambda \in \mathcal{W}$. Let $G\Lambda \subseteq \mathcal{W}$ be the G-orbit of Λ. By Lemma 3.2.2 there is an G-invariant map $\phi : X_o^0 \to C(X_o, G\Lambda)^0$. Since $\mathrm{Stab}_G(\Lambda)$ is not commensurable to a codimension-1 subgroup, Theorem 3.3.1 implies that there is a fixed 0-cube x in $C(X_o, G\Lambda)$. Lemma 3.2.2 then implies that $\phi^{-1}(x)$ is non-empty. Assuming that $\phi^{-1}(x) \subseteq \overleftarrow{\Lambda}$, then the intersection $\bigcap_{g \in G} gL(\Lambda)$ contains a proper, convex, G-invariant subcomplex of X_o, with one less hyperplane orbit. This contradicts the minimality of X_o. $\square$

The following Corollary follows since all codimension-1 subgroups of a rank n virtually abelian group are of rank $(n-1)$.

Corollary 3.3.3. *Let G be a rank n, virtually abelian group acting without fixed point or inversions on a CAT(0) cube complex X. Then there exists a minimal, G-invariant, convex subcomplex $X_o \subseteq X$ such that X_o contains* *only finitely many hyperplane orbits, and every hyperplane stabilizer is a rank $(n-1)$ subgroup of G.*

3.4 Proof of Main Theorem

Definition 3.4.1. Regard $\mathbb{R}$ as a CAT(0) cube complex whose 0-skeleton is $\mathbb{Z}$. Let g be an isometry of X. A *geodesic combinatorial axis* for g is a g-invariant subcomplex homeomorphic to $\mathbb{R}$ that embeds isometrically in X.

Definition 3.4.2. Let (M, d) be a metric space. The subspaces $N_1, N_2 \subseteq M$ are *coarsely equivalent* if each lies in an r-neighbourhood of the other for some $r > 0$.

Theorem 3.4.3. *Let G be virtually $\mathbb{Z}^n$. Suppose G acts properly and without inversions on a CAT(0) cube complex X. Then G stabilizes a finite dimensional subcomplex $Y \subseteq X$ that is isometrically embedded in the combinatorial metric, and $Y \cong \prod_{i=1}^{m} C_i$, where each C_i is a cubical quasiline and $m \geq n$. Moreover,* $\mathrm{Stab}_G(\Lambda)$ *is a codimension-1 subgroup for each hyperplane Λ in Y.*

Proof. By Corollary 3.3.3 there is a minimal, non-empty, convex subcomplex $X_o \subseteq X$ stabilized by G, containing finitely many hyperplane orbits, and $\mathrm{Stabilizer}_G(\Lambda)$ is a rank $(n-1)$ subgroup of G, for each hyperplane $\Lambda \subseteq X_o$.

Let $S = \{g_1, \ldots, g_r\}$ be a generating set for G. Let $x \in X_o$ be a 0-cube. Let Υ be the Cayley graph of G with respect to S. Let $\phi : \Upsilon \to X_o$ be a G-equivariant map that sends vertices to vertices, and edges to combinatorial paths or vertices in X_o. Let $Q = \phi(\Upsilon)$. As G acts properly on X, and cocompactly on Υ, the graph Q is quasiisometric to G. Let $\mathcal{W}_Q$ be the set of hyperplanes intersecting Q, and let $(Q, \mathcal{W}_Q)$ be the associated wallspace. By Lemma 3.2.1 we know that $C(Q, \mathcal{W}_Q)$ is an isometrically embedded subcomplex of X_o. Fix a proper action of G on $\mathbb{R}^n$, and let $q : Q \to \mathbb{R}^n$ be a G-equivariant quasiisometry. Note that $\mathrm{Stabilizer}_G(\Lambda)$ is a quasiisometrically embedded subgroup of G, for all $\Lambda \in \mathcal{W}_Q$. Thus $q(\Lambda \cap Q)$ is coarsely equivalent to a codimension-1 affine subspace $H \subseteq \mathbb{R}^n$. Moreover, $q(\overleftarrow{\Lambda} \cap Q)$ and $q(\overrightarrow{\Lambda} \cap Q)$ are coarsely equivalent to the halfspaces of H.

Let $n > 0$. Since there are finitely many orbits of hyperplanes in X_o, there are only finitely many commensurability classes of stabilizers. Therefore, we may partition $\mathcal{W}_Q$ as the disjoint union $\bigsqcup_{i=1}^m \mathcal{W}_i$ where each $\mathcal{W}_i$ contains all walls with commensurable stabilizers. For each $\Lambda_i \in \mathcal{W}_i$ let $q(\Lambda_i \cap Q)$ be coarsely equivalent to a codimension-1 affine subspace $H_i \subseteq \mathbb{R}^n$, stabilized by $\text{Stabilizer}_G(\Lambda_i)$. If $i \neq j$ then H_i and H_j are nonparallel affine subspaces, and therefore Λ_i and Λ_j will intersect in Q. Therefore, every wall in $\mathcal{W}_i$ intersects every wall in $\mathcal{W}_j$ if $i \neq j$, and thus $C(Q, \mathcal{W}_Q) \cong \prod_{i=1}^m C(Q, \mathcal{W}_i)$.

Finally, we show that $C(Q, \mathcal{W}_i)$ is a quasiline for each $1 \leq i \leq m$. As G permutes the factors in $\prod_{i=1}^m C(Q, \mathcal{W}_i)$, there is a finite index subgroup $G' \leqslant G$ that preserves each factor. For each i, the stabilizers $\text{Stab}_G(\Lambda)$ are commensurable for all $\Lambda \in \mathcal{W}_i$. Therefore, there is a cyclic subgroup Z_i that is not virtually contained in any $\text{Stab}_G(\Lambda)$ and thus acts freely on $C(Q, \mathcal{W}_i)$. As the stabilizers of $\Lambda \in \mathcal{W}_i$ are commensurable, all $q(\Lambda \cap Q)$ will be quasi-equivalent to parallel codimension-1 affine subspaces of $\mathbb{R}^n$, which implies that only finitely many Z_i-translates of Λ can pairwise intersect. As there are finitely many Z_i-orbits of Λ in $\mathcal{W}_i$, there is an upper bound on the number of pairwise intersecting hyperplanes in $\mathcal{W}_i$. Thus, there are finitely many Z_i-orbits of maximal cubes in $C(Q, \mathcal{W}_i)$, which implies that $C(Q, \mathcal{W}_i)$ is CAT(0) cube complex quasiisometric to $\mathbb{R}$. $\square$

We can now prove Corollary 3.4.4.

Corollary 3.4.4. *Let G be virtually $\mathbb{Z}^n$. Suppose G acts properly and without inversions on a CAT(0) cube complex X. Then G cocompactly stabilizes a* *subspace $F \subseteq X$ homeomorphic to $\mathbb{R}^n$ such that for each hyperplane $\Lambda \subseteq X$, the intersection $\Lambda \cap F$ is either empty or homeomorphic to $\mathbb{R}^{n-1}$.*

Proof. By Theorem 3.4.3 there is a G-equivariant, isometrically embedded, subcomplex $Y \subseteq X$, such that $Y = \prod_{i=1}^{m} C_i$, where each C_i is a quasiline, and $\mathrm{Stab}_G(\Lambda)$ is a codimension-1 subgroup. Considering Y with the CAT(0) metric, note that Y is a complete CAT(0) metric space in its own right, and G acts semisimply on Y. By the Flat Torus Theorem [6] there is an isometrically embedded flat $F \subseteq Y$. Note that $F \subseteq X$ is not isometrically embedded. As $\mathrm{Stab}_G(\Lambda)$ is a codimension-1 subgroup of G for each hyperplane Λ in X, the intersection $\Lambda \cap F = (\Lambda \cap Y) \cap F$ is either empty or, as $F \subseteq Y$ is isometrically embedded, the hyperplane intersection is an isometrically embedded copy of $\mathbb{R}^{n-1}$. $\square$

3.5 Haglund's Axis

The goal of this section is to reprove the following result of Haglund as a consequence of Corollary 3.4.4.

Theorem 3.5.1 (Haglund [19])**.** *Let G be a group acting on a CAT(0) cube complex without inversions. Every element $g \in G$ either fixes a 0-cube of G, or stabilizes a combinatorial geodesic axis.*

Proof. As finite groups don't contain codimension-1 subgroups, Theorem 3.3.1 implies that if g is finite order then it fixes a 0-cube. Suppose that G does not fix a 0-cube, then $\langle g \rangle$ must act properly on X. By Corollary 3.4.4, there is a line $L \subset X$ stabilized by $\langle g \rangle$, that intersects each hyperplane at most once at a single point in L. Let $\mathcal{W}_L$ be the set of hyperplanes intersecting L. Note that the intersection points of the walls in $\mathcal{W}_L$ with L is locally finite subset.

Fix a basepoint $p \in L$ that doesn't belong to a hyperplane intersecting L, and let x be the canonical 0-cube corresponding to p. Let $\Lambda_1, \ldots, \Lambda_k$ be the set of hyperplanes separating p and gp, and assume that $p \in \overleftarrow{\Lambda}_i$. Reindex the hyperplanes such that $\overleftarrow{\Lambda}_1 \cap L \subseteq \overleftarrow{\Lambda}_2 \cap L \subseteq \cdots \subseteq \overleftarrow{\Lambda}_k \cap L$. The ordering of the hyperplanes separating p and gp determines a combinatorial geodesic joining x and gx of length k, where the i-th edge is a 1-cube dual to Λ_i. This can be extended $\langle g \rangle$-equivariantly, to obtain a combinatorial geodesic axis L_c, since each hyperplanes intersects L_c at most once. $\square$

CHAPTER 4
Tubular Groups and their Finite Dimensional Cubulations

4.1 A Preliminary Discussion

A *tubular group* G is a group which splits as a graph of groups with $\mathbb{Z}^2$ vertex groups and $\mathbb{Z}$ edge groups. A tubular group is the fundamental group of a graph of spaces X with each vertex space homeomorphic to a torus and each edge space homeomorphic to a cylinder. The graph of spaces X is a *tubular space*. Note that in this chapter X will denote a tubular space, and *not* a CAT(0) cube complex as in Chapter 3. Moreover, we will let $\widetilde{X}$ denote the universal cover.

In [40] Wise reduces the existence of cubulations to a combinatorial criterion called *equitable sets*. Given an equitable set one can construct a finite set of *immersed walls*. An immersed wall is a graph Λ immersed π_1-injectively in X, such that $\widetilde{\Lambda}$ lifts to a 2-sided embedding $\widetilde{\Lambda} \to \widetilde{X}$. By 2-sided we mean that the image of $\widetilde{\Lambda}$ in $\widetilde{X}$ is contained in a neighbourhood homeomorphic to $\widetilde{\Lambda} \times [-1, 1]$. The set of all such lifts gives a G-invariant set $\mathcal{W}$ of *walls*. The pair $(\widetilde{X}, \mathcal{W})$ is a *wallspace*. In this chapter, all references to "immersed walls" will be in reference to immersed walls obtained from an equitable set. **Note that the theorems referring to immersed walls will not apply to any other kind of immersed walls.** Section 4.6 defines the notion of a *dilating wall* which can be recognized through a combinatorial criterion. As explained in Proposition 4.6.13, a wall not being dilated means its G-translates intersecting a vertex space in $\widetilde{X}$ can be partitioned into finitely many sets of pairwise nonintersecting walls.

Definition 4.1.1. An *infinite cube* in a CAT(0) cube complex is the union of an ascending sequence of n-cubes c_n of $\widetilde{X}$ such that c_n is a subcube of c_{n+1} for each n.

The following is the main goal of this chapter.

Theorem 4.6.10. Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The following are equivalent:

1. The dual cube complex $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional.
2. The dual cube complex $C(\widetilde{X}, \mathcal{W})$ contains an infinite cube.
3. One of the immersed walls is dilated.

Corollary 4.6.11 states that it is decidable if a given immersed wall is dilated. However, in contrast to cocompactness, there is no known simple criterion to determine whether or not a given tubular group acts on a finite dimensional CAT(0) cube complex. See Example 4.7.1 for an example of a direct proof that a specific tubular group does not posses an equitable set that can produce non-dilated immersed walls.

A group G is *separable* if every finitely generated subgroup $H \leqslant G$ is the intersection of all finite index subgroups containing H. Burns, Karrass, and Solitar gave the first example of a non-separable 3-manifold group [7]. Niblo and Wise reproved this result [27]. Rubinstein and Wang produced an example of a non-embedded immersed surface S in a graph manifold M such that $\widetilde{S} \to \widetilde{M}$ is injective and any two $\pi_1 M$ translates of $\widetilde{S}$ intersect [31]. It follows that $\pi_1 S$ is not separable in $\pi_1 M$. This is an application of the geometric interpretation for separability given by Scott [35]. Theorem 4.6.10 carries information about the separability of the associated codimension-1 subgroups of G, and provides a new proof of the non-subgroup separability of certain 3-manifold groups that conceptually unifies the result of Burns, Karrass, and Solitar with the geometric proof of Rubinstein and Wang (see Example 4.4.1).

Finally, the following Theorem is a consequence of Theorem 4.6.10 together with a study of the geometry of the walls.

Theorem 4.1.2. *Let the tubular space X associated to the tubular group G have a finite set of immersed walls. If the corresponding wallspace $(\widetilde{X}, \mathcal{W})$ has quasi-isometrically embedded walls, then the dual CAT(0) cube complex $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Example 4.8.2 demonstrates that quasi-isometrically embedded walls are not a necessary condition for finite dimensionality.

4.2 Equitable Sets and Immersed Walls

Let G be a tubular group with associated tubular space X and underlying graph Γ. Given an edge e in a graph we will let $-e$ and $+e$ respectively denote the initial and terminal vertices of e. Let X_v and X_e denote vertex and edge spaces in this graph of spaces. Let X_e^- and X_e^+ be the boundary circles of X_e, and denote the attaching maps by $\varphi_e^- : X_e^- \to X_{-e}$, and $\varphi_e^+ : X_e^+ \to X_{+e}$. Note that φ_e^- and φ_e^+ respectively represent generators of G_e in G_{-e} and G_{+e}. We will let $\widetilde{X}$ denote the universal cover of X. Let $\widetilde{X}_{\tilde{v}}$ and $\widetilde{X}_{\tilde{e}}$ denote vertex and edge spaces in the universal cover $\widetilde{X}$, and let $\widetilde{\Gamma}$ denote the Bass-Serre tree. We will assume that each vertex space has the structure of a non-positively curved geodesic metric space and that attaching maps φ_e^- and φ_e^+ define locally geodesic curves in X_{-e} and X_{+e}.

4.2.1 Equitable Sets and Intersection Numbers

Given a pair of closed curves in a torus $\alpha, \beta : S^1 \to T$, the *intersection points* are the elements $(p, q) \in S^1 \times S^1$ such that $\alpha(p) = \beta(q)$. For a pair of homotopy classes $[\alpha], [\beta]$ of closed curves in a torus T, their *geometric intersection number* $\#\big[[\alpha], [\beta]\big]$ is the minimal number of intersection points realised by a pair of representatives from the respective classes. This number is realised by any pair of geodesic representatives of the classes. If $B = \{[\beta_i]\}$ is a finite set of homotopy classes of curves in T, then $\#[\alpha, B] := \sum_i \#[\alpha, \beta_i]$. Viewing $[\alpha]$ and $[\beta]$ as elements of $\pi_1 T = \mathbb{Z}^2$, we can compute that $\#\Big[[\alpha], [\beta]\Big] = \Big|\det\Big[[\alpha], [\beta]\Big]\Big|$. Given an identification of $\mathbb{Z}^2$ with $\pi_1 T$, the elements of $\mathbb{Z}^2$ are identified with homotopy classes of curves in T, so it makes sense to consider their geometric intersection number. An *equitable set* for a tubular group G is a collection of sets $\{S_v\}_{v\in\Gamma}$, where S_v is a finite set of distinct geodesic curves in X_v that are either disjoint or transverse each other and to the attaching maps of adjacent edge spaces, such that S_v generate a finite index subgroup of $\pi_1 X_v = G_v$, and $\#\Big[\varphi_e^-, S_{-e}\Big] = \#\Big[\varphi_e^+, S_{+e}]$. Note that equitable sets can also be given with S_v a finite subset of G_v that generates a finite index subgroup of G_v and satisfies the corresponding equality for intersection numbers. This is how Wise formulates equitable sets, and its equivalence follows from exchanging elements of $G_v = \pi_1 X_v$ with geodesic closed curves in X_v that represent the corresponding elements. An equitable set is *fortified* if for each edge e in Γ, there exists $\alpha_e^+ \in S_{+e}$ and $\alpha_e^- \in S_{-e}$ such that $\#[\alpha_e^+, \varphi_e^+] = \#[\alpha_e^-, \varphi_e^-] = 0$. An equitable set is *primitive* if every element $\alpha \in S_v$ represents a primitive element in G_v.

4.2.2 Immersed Walls From Equitable Sets

The main theorem proven in [40] is

Theorem 4.2.1 (Wise)**.** *A tubular group G acts freely on a CAT(0) cube complex if and only if there is an equitable set for G.*

As we repeatedly use the construction given in the proof, here is an outline of how an equitable set provides a CAT(0) cube complex with a free G-action.

Immersed walls are constructed from *circles* and *arcs*. For each $\alpha \in S_v$, let S^1_α be the domain of α. The disjoint union $\bigsqcup S^1_\alpha$ over all $\alpha \in S_v$ and $v \in V\Gamma$ are the *circles*. Since $\#[\varphi_e^-, S_{-e}] = \#[\varphi_e^+, S_{+e}]$, there exists a bijection from the intersection points between curves in S_{-e} and φ_e^-, and the intersection points between curves in S_{+e} and φ_e^+. Let $(p^-, q^-) \in S^1_{-e} \times X_e^-$ and

$(p^+, q^+) \in S^1_{+e} \times X^+_e$ be corresponding intersection points between $\alpha^\pm \in S_{\pm e}$ and $\varphi^\pm_e$. Then an *arc* $a \cong [0,1]$ has its endpoints attached to p^- and p^+. The endpoints of a are mapped into $X_{-e} \cap X_e$ and $X_{+e} \cap X_e$, so the interior of a can be embedded in X_e. After attaching an arc for each pair of corresponding intersection points, we obtain a set $\{\Lambda_1, \ldots, \Lambda_n\}$ of connected graphs that map into X, called *immersed walls.* The fundamental group of Λ_i has its own graph of groups structure with infinite cyclic vertex groups and trivial edge groups. **All "immersed walls" in this Chapter and Chapter 5 are immersed walls constructed from equitable sets as above.**

A lift of $\widetilde{\Lambda}_i \to \widetilde{X}$ is a two sided embedding in $\widetilde{X}$, separating $\widetilde{X}$ into two *halfspaces.* The images of the lifts of $\widetilde{\Lambda}_i$ to $\widetilde{X}$ are *horizontal walls* $\mathcal{W}_\mathsf{h}$. The *vertical* walls $\mathcal{W}_\mathsf{v}$ are obtained from the lifts of curves $\alpha_e : S^1 \to X_e$ given by the inclusion $S^1 \times \{0\} \hookrightarrow S^1 \times [-1,1]$. The set $\mathcal{W} = \mathcal{W}_\mathsf{h} \sqcup \mathcal{W}_\mathsf{v}$ of all horizonal and vertical walls gives a *wallspace* $(\widetilde{X}, \mathcal{W})$ where the G-action on $\widetilde{X}$ also gives an action on $\mathcal{W}$.

A set of immersed walls is *fortified* if they are obtained from a fortified equitable set. A set of immersed walls is *primitive* if they are obtained from a primitive equitable set.

Remark 4.2.2. In the above construction, the resulting dual CAT(0) cube complex depends on the choice of correspondence in each edge space, as well as how the interior of each arc embeds in the cylinder. Example 4.7.3 shows finite dimensionality of the cubulation is dependent on the choice of correspondence between the intersection points at either end of the edge space. By contrast, Theorem 4.6.10 and the definition of a dilated immersed wall imply finite dimensionality is invariant of the subsequent choice of immersion of the interior of the arcs in the edge spaces.

4.3 Examples of tubular groups

Tubular groups include examples from many other classes of interesting groups within geometric group theory. For example, there are free-by-cyclic groups, right angled Artin groups, and the fundamental groups of graph manifolds that are also tubular groups. By way of motivation we will give some key examples that illustrate the variety available.

Example 4.3.1 (Gersten's Group [15])**.** Let $\mathbb{F}_3 = \langle a, b, c\rangle$ and let ϕ be the automorphism of $\mathbb{F}_3$ determined by mapping $a \mapsto a, b \mapsto ba$, and $c \mapsto ca^2$. Let G be the Free-by-Cyclic group constructed from $\mathbb{F}_3$ and ϕ. $G = \langle a, b, c, t \mid a = tat^{-1}, b = tbat^{-1}, c = tca^2t^{-1}\rangle$. Note that $\langle a, t\rangle$ is a rank 2 free abelian group.

Proposition 4.3.2 ([15])**.** *G is not CAT(0).*

Proof. Suppose that G were to act metrically properly and cocompactly on a CAT(0) space X. Then $\langle a, t\rangle$ would stabilize a convex subspace $Y \subseteq X$ such that $Y \cong \mathbb{E}^2$ and such that the minimal translation lengths of a and t would be realized on Y. Rewriting the second relator as $b^{-1}tb = ta^{-1}$ we can deduce that t is conjugate to ta^{-1}. Similarly, by rewriting the third relator as $c^{-1}tc = ta^{-2}$ we can deduce that t is conjugate to ta^{-2}.

The minimal translation length in X of an element of G is invariant under conjugation. Therefore the translation lengths of t, ta^{-1}, and ta^{-2} coincide. Therefore if $y \in Y$, then $ty, ta^{-1}y$, and $ta^{-2}y$ are three distinct points of a circle centered at y. This contradicts the fact that $ty, ta^{-1}y$, and $ta^{-2}y$ must also be colinear points in Y, since they are each $\langle a\rangle$-translates of each other. $\square$

Although G is not CAT(0), as a consequence of Wise's criterion for cubulation we can see that G does act freely on a CAT(0) cube complex. This action will necessarily be non-semisimple. The presentation may be rewritten as $\langle x, y, s, t \mid [x, y] = 1, y = s^{-1}xys, x^{-1}y = t^{-1}xyt\rangle$, by sending $a \mapsto x, t \mapsto xy, b \mapsto s, c \mapsto t$. Note that this presentation makes it explicit that G is a tubular group,decomposing with a single vertex group $\langle x, y\rangle \cong \mathbb{Z}^2$, and two $\mathbb{Z}$ edge groups with stable letters s and t. It can then be verified that $S = \{xy, xy^{-1}\}$ is an equitable set since $\#[xy, S] = 2 = \#[xy^{-1}, S] = \#[y, S]$.

Example 4.3.3 (Non-Hopfian [37])**.** Let $G = \langle a, b, s, t \mid [a, b] = 1, sas^{-1} = (ab)^2, tbt^{-1} = (ab)^2\rangle$. It is clear from the presentation that G splits as a graph of groups with a single vertex group $\langle a, b\rangle \cong \mathbb{Z}^2$, and two $\mathbb{Z}$ edge groups with stable letters s and t.

In [37], Wise constructed an explicit non-positively curved 2-complex with fundamental group G. Moreover, he showed that G is *non-Hopfian*: there exists a surjection from G to itself that is not injective. As a consequence this implies that G is not residually finite.

We define a homomorphism $\phi : G \to G$, by mapping $a \mapsto a^2, b \mapsto b^2, s \mapsto s$, and $t \mapsto t$. Since $\phi(ab) = (ab)^2$, we can then conjugate by either s or t to obtain a and b in the image of ϕ. Therefore ϕ is surjective.

To see that ϕ is not injective we consider the element $g = [s^{-1}(ab)s, t^{-1}(ab)t]$ and verify using the normal form theorem for graphs of groups that $g \neq 1$. Then we calculate that $\phi(g) = [s^{-1}(ab)^2s, t^{-1}(ab)^2t] = [a, b] = 1$, which implies that ϕ is not injective.

It can be shown that G has an equitable set, and therefore that G acts freely on a CAT(0) cube complex. Furthermore, Thm 1.2 in [40] states that all CAT(0) tubular groups act freely on a CAT(0) cube complex, although not necessarily a finite dimensional CAT(0) cube complex.

4.4 The Motivating Example

The next example motivates the definitions and theorems given in this chapter that identify which immersed walls give infinite cubulations.

Example 4.4.1. Consider the tubular group $G = \langle a, b, t \mid [a, b] = 1, tat^{-1} = b\rangle$. This is a cyclic HNN extension of $\mathbb{Z}^2 = \langle a, b\rangle$. The associated graph of

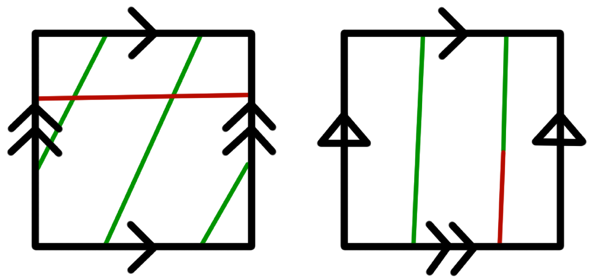

Figure 4–1: The graph of spaces X, with an immersed wall given by the equitable set.

spaces has a torus vertex space X_v, and a cylindrical edge space X_e that is attached along geodesic paths associated to a and b. Consider the immersed wall given by the equitable set $\{a, ab^2\}$ and the choice of arcs shown at right in Figure 4–1.

An infinite set of pairwise crossing walls in $(\widetilde{X}, \mathcal{W})$ will be exhibited to demonstrate that $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional. The immersed wall in this example is the prototype for a general dilated immersed wall, which will be discussed in Section 4.6. Choose a vertex space $\widetilde{X}_{\tilde{v}_1}$ in the universal cover. The circle ab^2 in X_v is covered by an infinite set of parallel lines in $\widetilde{X}_{\tilde{v}_1}$. Index these lines consecutively by the elements of $\mathbb{Z}$. See the top of Figure 4–2.

Choose a vertex space $\widetilde{X}_{\tilde{v}_2}$ adjacent to $\widetilde{X}_{\tilde{v}_1}$ via an edge space attached along an a-line in $\widetilde{X}_{\tilde{v}_1}$. The walls indexed $2n$ in $\widetilde{X}_{\widetilde{v_1}}$ travel through to $\widetilde{X}_{\tilde{v}_2}$ where they either cover ab^2 or a. By adding $+1$ to each index if necessary, assume that they cover ab^2. Then the walls indexed $2n+1$ also travel through to $\widetilde{X}_{\tilde{v}_2}$ where they cover a. Observe the wall indexed $2n$ in $\widetilde{X}_{\tilde{v}_1}$ can be indexed n in $\widetilde{X}_{\tilde{v}_2}$ and the indexing of the lines remains consecutive. This follows from how the immersed wall in the base space determines the walls in the universal cover (see Figure 4–2).

The odd indexed walls in $\widetilde{X}_{\tilde{v}_1}$ are now parallel to a in $\widetilde{X}_{\tilde{v}_2}$. Thus the even walls in $\widetilde{X}_{\tilde{v}_1}$ cross the odd walls in $\widetilde{X}_{\tilde{v}_1}$.

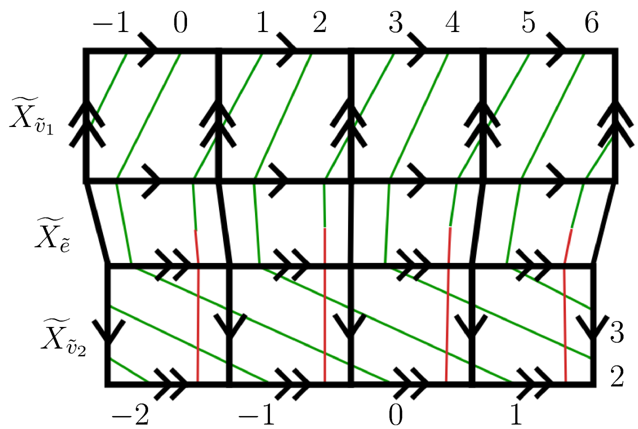

Figure 4–2: A portion of the universal cover $\widetilde{X}$, consisting of parts of $\widetilde{X}_{\tilde{v}_1}$, $\widetilde{X}_{\tilde{e}}$, and $\widetilde{X}_{\tilde{v}_2}$. The numbers at the top show the indexing of the walls in $\widetilde{X}_{\tilde{v}_1}$ and the numbers at the bottom show the indexing of the walls in $\widetilde{X}_{\tilde{v}_2}$.

We will now show that the walls indexed $1, 2, 4, 8, \ldots, 2^n, \ldots$ in $\widetilde{X}_{\tilde{v}_1}$ are pairwise intersecting. Indeed, as shown above, the wall numbered 1 intersects all the even walls in $\widetilde{X}_{\tilde{v}_1}$. To repeat the above argument, consecutively index the walls intersecting $\widetilde{X}_{\tilde{v}_2}$ as ab^2-lines, by the elements of $\mathbb{Z}$, such that the walls indexed $2, 4, 8, \ldots, 2^n, \ldots$ in $\widetilde{X}_{\tilde{v}_1}$ are then indexed $1, 2, 4, \ldots, 2^{n-1}, \ldots$ in $\widetilde{X}_{\tilde{v}_2}$. Arguing inductively, given the vertex space $\widetilde{X}_{\tilde{v}_i}$ containing walls parallel to ab^2 indexed $1, 2, 4, 8, \ldots, 2^n, \ldots$, there exists an adjacent vertex space $\widetilde{X}_{\tilde{v}_{i+1}}$ where the wall indexed 1 will cross all the other walls which can be re-indexed $1, 2, 4, 8, \ldots, 2^n, \ldots$ in $\widetilde{X}_{\tilde{v}_{i+1}}$. Therefore there is an infinite family of pairwise crossing walls intersecting $\widetilde{X}_{\widetilde{v_1}}$. It will be shown in Proposition 4.5.4 that such a family determines an infinite cube.

Remark 4.4.2. The existence of an infinite set of pairwise crossing walls in $\widetilde{X}$ implies the non-separability of the associated codimension-1 subgroup by Scott's criterion [35]. The tubular group of Example 4.4.1 is the 3-manifold group that Burns, Karrass and Solitar proved was not subgroup separable in [7]. The group that they used to prove non subgroup separability is a subgroup of the stabilizer of this wall. Niblo and Wise reproved this result in

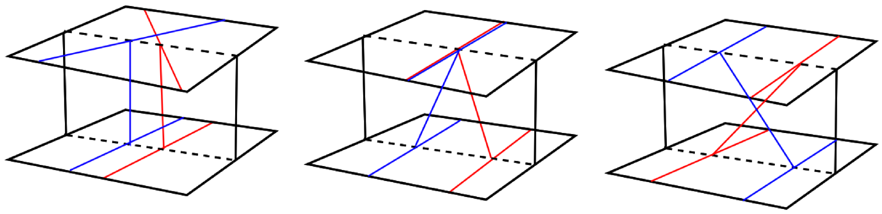

Figure 4–3: A regular intersection on the left and non-regular intersections on the right and center.

[27] by finding a family of subgroups that are not contained in any finite index subgroup. The stabilizer of the wall belongs to that family.

4.5 Infinite Dimensional Cubulations Must Contain Infinite Cubes

The goal of this section is to prove Proposition 4.5.4, a weaker version of implication (1)⇒(2) from Theorem 4.6.10.

Remark 4.5.1. A distinction will be made between two types of intersections occurring between a pair of walls in $(\widetilde{X}, \mathcal{W})$ (see Figure 4–3).

Regular An intersection in a vertex space as between non-parallel lines in $\mathbb{R}^2$.

Non-Regular An intersection in a vertex space as overlapping parallel lines in $\mathbb{R}^2$, or an intersection occurring in an edge space.

Note that walls may intersect more than once, and a given pair of walls may have multiple regular and non-regular intersections. Note that if $\widetilde{\Lambda}$ and $\widetilde{\Lambda}'$ regularly intersect in $\widetilde{X}_{\tilde{v}}$, then $\widetilde{\Lambda}$ and $g\widetilde{\Lambda}'$ also regularly intersect in $\widetilde{X}_{\tilde{v}}$ for all $g \in G_{\tilde{v}}$. Two walls $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2$ that intersect a vertex space $\widetilde{X}_{\tilde{v}}$, but do not intersect regularly are *parallel in* $\widetilde{X}_{\tilde{v}}$.

The following short lemma is used extensively throughout the paper.

Lemma 4.5.2. *Let $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ be walls in $(\widetilde{X}, \mathcal{W})$ that intersect $\widetilde{X}_{\tilde{v}}$. Let $\widetilde{Y} \subset \widetilde{X}$ be a subtree of spaces containing $\widetilde{X}_{\tilde{v}}$ such that $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2$ have no regular intersections in $\widetilde{Y}$. Then every vertex or edge space in $\widetilde{Y}$ that intersects $\widetilde{\Lambda}_1$ also intersects $\widetilde{\Lambda}_2$.*

Proof. Let $\widetilde{X}_{\widetilde{u}}$ be a vertex space in $\widetilde{Y}$. If $\widetilde{\Lambda}_1 \cap \widetilde{X}_{\widetilde{u}}$ and $\widetilde{\Lambda}_2 \cap \widetilde{X}_{\widetilde{u}}$ are a pair of parallel (possibly overlapping) lines in $\widetilde{X}_{\widetilde{u}}$ then $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2$ intersect an identical set of edge spaces adjacent to $\widetilde{X}_{\widetilde{u}}$, and thus an identical set of vertex spaces adjacent to $\widetilde{X}_{\widetilde{u}}$. Therefore it can be deduced inductively, starting with $\widetilde{X}_{\tilde{v}}$, that $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ must intersect an identical set of vertex and edge spaces in $\widetilde{Y}$. $\square$

Lemma 4.5.3. *Let S be a finite set of walls in $(\widetilde{X}, \mathcal{W})$ that pairwise intersect. There exists a vertex space $\widetilde{X}_{\widetilde{v}}$ that intersects every wall in S.*

Proof. Each wall in S maps to a connected subset of the Bass-Serre tree. These subsets pairwise intersect, therefore by application of Helly's Theorem for trees, there exists some vertex in the total intersection. Therefore every wall in S intersects the corresponding vertex space. $\square$

A set of horizontal walls in $(\widetilde{X}, \mathcal{W})$ that pairwise regularly intersect is *crossing*.

Proposition 4.5.4. *If there are crossing sets of arbitrary finite cardinality in $(\widetilde{X}, \mathcal{W})$, then there exists an infinite cube in $C(\widetilde{X}, \mathcal{W})$. Moreover, the infinite cube contains a canonical 0-cube.*

Proof of Proposition 4.5.4. Let $\{K_i\}$ be a sequence of crossing sets such that $|K_i| \to \infty$ as $i \to \infty$. There are only finitely many wall orbits so assume that each K_i consists exclusively of translates of a single wall $\widetilde{\Lambda}$. By Lemma 4.5.3, for every i there exists a vertex $\tilde{v}_i$ in $\widetilde{\Gamma}$ such that every wall in K_i intersects $\widetilde{X}_{\tilde{v}_i}$. Since G acts cocompactly on $\widetilde{\Gamma}$, after passing to a subsequence of $\{K_i\}$, it can be assumed that the $\{\tilde{v}_i\}$ lie in a single G-orbit, $G\tilde{v}$. Choose g_i such that $g_i\tilde{v}_i = \tilde{v}$ and replace K_i with $g_i^{-1}K_i$. Therefore all walls in K_i intersect the vertex space $\widetilde{X}_{\widetilde{v}}$ for each i.

There are only finitely many parallelism classes of lines in $\widetilde{X}_{\widetilde{v}}$ belonging to walls in $\mathcal{W}$, so there must exist Q_0, a set of walls parallel in $\widetilde{X}_{\widetilde{v}}$, with the

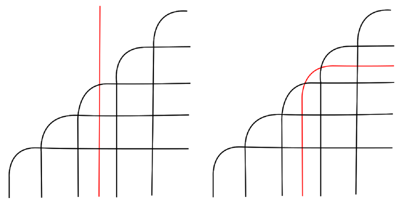

Figure 4–4: The potential obstruction to an infinite cube on the left, and the concluding situation on the right. Compare with the picture in [22].

following property: *there is no upper bound on the size of crossing subsets of* Q_0. Given Q_0, there must exist a crossing set $\{\widetilde{\Lambda}_1, \widetilde{\Lambda}_1'\} \subset Q_0$. The regular intersection points between $\widetilde{\Lambda}_1, \widetilde{\Lambda}_1'$ must be in vertex spaces other than $\widetilde{X}_{\tilde{v}}$. Choose $\widetilde{\Lambda}_1, \widetilde{\Lambda}_1'$ such that they intersect in $\widetilde{X}_{\widetilde{v_1}}$, and the distance from $\tilde{v}$ to $\widetilde{v_1}$ in $\widetilde{\Gamma}$ is minimized over all such choices of intersection points and choices of $\widetilde{\Lambda}_1, \widetilde{\Lambda}_1'$.

By Lemma 4.5.2, setting $\widetilde{Y} \subset \widetilde{X}$ to be the union of vertex and edge spaces between and including $\tilde{v}$ and $\tilde{v}_1$, every other wall in Q_0 also intersects $\widetilde{X}_{\widetilde{v_1}}$. Repeating the above argument there exists a set Q_1, of walls parallel in $\widetilde{X}_{\widetilde{v_1}}$, such that there is no upper bound on the size of crossing subsets. Either $\widetilde{\Lambda}_1$ or $\widetilde{\Lambda}_1'$ will therefore intersect every wall in Q_1, so assume that it is $\widetilde{\Lambda}_1$. Iterate this argument to obtain a sequence of walls $\{\widetilde{\Lambda}_i\}_{i=1}^{\infty}$ that pairwise intersect.

Let $\{\widetilde{\Lambda}_i\}_{i=1}^{\infty}$ be the (sub)sequence of walls re-indexed so that they are ordered consecutively in $\widetilde{X}_{\tilde{v}}$. An ascending sequence of n-cubes c_n in $C(\widetilde{X}, \mathcal{W})$ will now be constructed. Orient all walls not intersecting $\widetilde{X}_{\tilde{v}}$ towards $\widetilde{X}_{\tilde{v}}$. Let c_0 be the canonical 0-cube given by orienting all walls towards some $x \in \widetilde{\Lambda}_1 \cap \widetilde{X}_{\tilde{v}}$. By induction, there is an ascending sequence of cubes $c_0 \subset \cdots \subset c_n$ dual to the walls $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_{n+1}$. Suppose that for some orientation of $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_{n+1}$ the orientation of $\widetilde{\Lambda}_{n+2}$ could not be reversed (see Figure 4–4). This means that there exists a wall that is parallel to $\widetilde{\Lambda}_{n+2}$, that intersects $\widetilde{X}_{\tilde{v}}$, and that lies between x and $\widetilde{\Lambda}_{n+2}$. There can only exist finitely many such walls and, by Lemma 4.5.2, setting $\widetilde{Y} = \widetilde{X}$ they all intersect $\widetilde{\Lambda}_i$ for $i \neq n+2$. Replace $\widetilde{\Lambda}_{n+2}$ with the closest such wall to x to obtain an $(n+1)$-cube c_{n+1}. $\square$

4.6 Dilating Walls and Infinite Dimensional Cubulations

Let $\widetilde{X}_{\tilde{v}}$ be a vertex space covering X_v. Identify $\widetilde{X}_{\tilde{v}}$ with $\mathbb{R}^2$ so that a $G_{\tilde{v}}$-orbit in $\widetilde{X}_{\tilde{v}}$ is identified with $\mathbb{Z}^2 \subset \mathbb{R}^2$. Let $\widetilde{\Lambda}$ denote a wall covering the immersed wall Λ, intersecting $\widetilde{X}_{\tilde{v}}$. The intersection $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ is a line with rational slope, stabilized by an infinite cyclic subgroup of $G_{\tilde{v}}$. Since the slope is rational there is a maximal infinite cyclic subgroup $H \leqslant G_{\tilde{v}}$ that stabilizes a line perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$. The elements of H are *perpendicular to* $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$. Note that there are finitely many $G_{\tilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\tilde{v}}$. Given vertices $\tilde{u}$, $\tilde{v}$ in $\widetilde{\Gamma}$, the *carrier of* $\tilde{u}, \tilde{v}$, denoted $\mathrm{Carrier}(\tilde{u}, \tilde{v})$, is the union of all vertex and edge spaces corresponding to the vertices and edges on the geodesic in the Bass-Serre tree $\widetilde{\Gamma}$ between and including $\tilde{u}$ and $\tilde{v}$.

The following is elementary:

Lemma 4.6.1. *Let T be a torus and $C \to T$ an immersed geodesic circle, with a choice of lift $\widetilde{C} \to \widetilde{T}$. Given $a, b \in \pi_1 T - \mathrm{Stab}(\widetilde{C})$, there exists $m_a, m_b \in \mathbb{Z} - \{0\}$ such that $a^{m_a}\widetilde{C} = b^{m_b}\widetilde{C}$. Moreover*

$$\frac{m_a}{m_b} = \frac{\#[C, b]}{\#[C, a]}.$$

Lemma 4.6.2. *Let $\widetilde{\Lambda}$ be a wall intersecting the vertex spaces $\widetilde{X}_{\tilde{v}_0}$ and $\widetilde{X}_{\tilde{v}_n}$, where $\tilde{v}_1, \ldots, \tilde{v}_{n-1}$ is the sequence of vertices on the geodesic between $\tilde{v}_0$ and $\tilde{v}_n$ in $\widetilde{\Gamma}$. Let $\tilde{e}_i$ be the edge between $\tilde{v}_{i-1}$ and $\tilde{v}_i$. Let $G_{\tilde{e}_i} = \langle g_i \rangle$ for $1 \leq i \leq n$. Let $\mathrm{Stab}(\widetilde{\Lambda}) \cap G_{\tilde{v}_i} = \langle \rho_i \rangle$. Let $g_0 \in G_{\tilde{v}_0}$ and $g_{n+1} \in G_{\tilde{v}_n}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}_0}$ and $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}_n}$ respectively. Then there exist $\overleftarrow{m}$, $\overrightarrow{m} \in \mathbb{Z} - \{0\}$ such that $g_0^{\overleftarrow{m}}\widetilde{\Lambda} = g_{n+1}^{\overrightarrow{m}}\widetilde{\Lambda}$ and*

$$\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{\overleftarrow{m}_0 \cdots \overleftarrow{m}_n}{\overrightarrow{m}_0 \cdots \overrightarrow{m}_n}$$

where $\overleftarrow{m}_i = \#[\rho_i, g_{i+1}]$ *and* $\overrightarrow{m}_i = \#[\rho_i, g_i]$ *for* $0 \leq i \leq n$. *Moreover the walls in* $\{g_0^{\overleftarrow{m}r}\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ *do not intersect within* $\mathrm{Carrier}(\tilde{v}_0, \tilde{v}_n)$.

Proof. For each $0 \leq i \leq n$ apply Lemma 4.6.1 to $g_i, g_{i+1} \in G_{\tilde{v}_i}$ to deduce that $g_i^{\overleftarrow{m}_i}\widetilde{\Lambda} = g_{i+1}^{\overrightarrow{m}_i}\widetilde{\Lambda}$. Let $\overleftarrow{m} = \overleftarrow{m}_0 \cdots \overleftarrow{m}_n$ and $\overrightarrow{m} = \overrightarrow{m}_0 \cdots \overrightarrow{m}_n$. Observe that for $0 \leq i \leq n$ we have:

$$g_0^{\overleftarrow{m}}\widetilde{\Lambda} = g_0^{\overleftarrow{m}_0\overleftarrow{m}_1\cdots\overleftarrow{m}_n}\widetilde{\Lambda} = g_i^{\overrightarrow{m}_0\cdots\overrightarrow{m}_i\overleftarrow{m}_{i+1}\cdots\overleftarrow{m}_n}\widetilde{\Lambda},$$

so there are no intersections of walls in $\{g_0^{\overleftarrow{m}r}\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ within either $\widetilde{X}_{\tilde{v}_i}$ or $\widetilde{X}_{\tilde{e}_i}$. $\square$

Definition 4.6.3. Let $\widetilde{\Lambda}$ be a wall covering Λ and let $\widetilde{X}_{\tilde{v}}$ intersect $\widetilde{\Lambda}$. Observe that $\widetilde{\Lambda}$ also intersects $\widetilde{X}_{g\tilde{v}}$ for any $g \in \mathrm{Stab}(\widetilde{\Lambda})$. If $h \in G_{\tilde{v}}$ is perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$, then ghg^{-1} is a perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{g\tilde{v}}$. From Lemma 4.6.2 there exist $\overleftarrow{m}, \overrightarrow{m} \in \mathbb{Z} - \{0\}$ such that $h^{\overleftarrow{m}}\widetilde{\Lambda} = gh^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$ and $\{h^{\overleftarrow{m}r}\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ contains no pair intersecting in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$. The pair $\overleftarrow{m}, \overrightarrow{m}$ are *g-shift exponents*. Note that g-shift exponents are not unique. If $\overleftarrow{m}, \overrightarrow{m}$ are g-shift exponents, then so are $\overleftarrow{m}n, \overrightarrow{m}n$ for $n \in \mathbb{Z} - \{0\}$. The *dilation function of* $\widetilde{\Lambda}$ is the map $R : \mathrm{Stab}(\widetilde{\Lambda}) \to \mathbb{Q}^*$ with $g \mapsto \dfrac{\overleftarrow{m}}{\overrightarrow{m}}$, where $\overleftarrow{m}, \overrightarrow{m}$ are g-shift exponents.

Lemma 4.6.4. *The map* $R : \mathrm{Stab}(\widetilde{\Lambda}) \to \mathbb{Q}^*$ *is a homomorphism, and does not depend on the choice of vertex space* $\widetilde{X}_{\tilde{v}}$ *or perpendicular element* $h \in G_{\tilde{v}}$. *Moreover, if* R' *is the dilation function of* $\widetilde{\Lambda}'$, *where* $\widetilde{\Lambda}' = g'\widetilde{\Lambda}$, *then* $R(g'^{-1}gg') = R'(g)$ *for* $g \in \mathrm{Stab}(\widetilde{\Lambda}')$.

Proof. Fix a choice of $\widetilde{X}_{\tilde{v}}$ and let $h \in G_{\tilde{v}}$ be a primitive element perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$. To see that R is well defined take any two choices of g-shift exponents $\overleftarrow{m}, \overrightarrow{m}$, and $\overleftarrow{m}', \overrightarrow{m}'$. Since $h^{\overleftarrow{m}}\widetilde{\Lambda} = gh^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$ and $h^{\overleftarrow{m}'}\widetilde{\Lambda} = gh^{\overrightarrow{m}'}g^{-1}\widetilde{\Lambda}$ we have $gh^{\overleftarrow{m}\overrightarrow{m}'}g^{-1}\widetilde{\Lambda} = h^{\overleftarrow{m}\overleftarrow{m}'}\widetilde{\Lambda} = gh^{\overrightarrow{m}\overleftarrow{m}'}g^{-1}\widetilde{\Lambda}$ which implies that $\dfrac{\overleftarrow{m}}{\overrightarrow{m}} = \dfrac{\overleftarrow{m}'}{\overrightarrow{m}'}$. Now suppose that $h' \in G_{\tilde{v}}$ is any other choice of an element perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$, then $h' = h^n$ with $n \in \mathbb{Z} - \{0\}$. Therefore, $(h')^{\overleftarrow{m}}\widetilde{\Lambda} = h^{n\overleftarrow{m}}\widetilde{\Lambda} =$

$gh^{n\overrightarrow{m}}g^{-1}\widetilde{\Lambda} = g(h')^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$, and $\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{n\overleftarrow{m}}{n\overrightarrow{m}}$, so $\overleftarrow{m}, \overrightarrow{m}$ are also g-shift exponents with respect to h'.

To see that R is independent of the choice of $\widetilde{X}_{\tilde{v}}$ suppose $\widetilde{\Lambda}$ also intersects $\widetilde{X}_{\tilde{w}}$. Let $h_{\tilde{v}} \in G_{\tilde{v}}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$, and $h_{\tilde{w}} \in G_{\tilde{w}}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{w}}$. Then by Lemma 4.6.2 there exist $p, q \in \mathbb{Z} - \{0\}$ such that $h_{\tilde{w}}^{p}\widetilde{\Lambda} = h_{\tilde{v}}^{q}\widetilde{\Lambda}$. Let $g \in \mathrm{Stab}(\widetilde{\Lambda})$ and $\overleftarrow{m}, \overrightarrow{m}$ be g-shift exponents with respect to $\widetilde{X}_{\tilde{u}}$ and deduce that $h_{\tilde{w}}^{\overleftarrow{m}p}\widetilde{\Lambda} = h_{\tilde{v}}^{\overleftarrow{m}q}\widetilde{\Lambda} = gh_{\tilde{v}}^{\overrightarrow{m}q}g^{-1}\widetilde{\Lambda} = gh_{\tilde{w}}^{\overrightarrow{m}p}g^{-1}\widetilde{\Lambda}$. This means that $\overleftarrow{m}p, \overrightarrow{m}p$ are g-shift exponents with respect to $\widetilde{X}_{\tilde{w}}$

Let $\widetilde{\Lambda}' = g'\widetilde{\Lambda}$ with $g' \in G$. Let $g \in \mathrm{Stab}(g'\widetilde{\Lambda})$. Then we know that $g'^{-1}gg' \in \mathrm{Stab}(\widetilde{\Lambda})$ so there exist $g'^{-1}gg'$-shift exponents $\overleftarrow{m}, \overrightarrow{m}$ for $\widetilde{\Lambda}$ with respect to $\widetilde{X}_{\tilde{v}}$ so that $h^{\overleftarrow{m}}\widetilde{\Lambda} = (g'^{-1}gg')h^{\overrightarrow{m}}(g'^{-1}gg')^{-1}\widetilde{\Lambda}$, where h is perpendicular to $\widetilde{\Lambda}$ in $\widetilde{X}_{\tilde{v}}$. This can be rewritten as $(g'hg'^{-1})^{\overleftarrow{m}}(g'\widetilde{\Lambda}) = g(g'hg'^{-1})^{\overrightarrow{m}}g^{-1}(g'\widetilde{\Lambda})$, where $g'hg'^{-1}$ is perpendicular to $\widetilde{\Lambda}'$ in $\widetilde{X}_{g'\tilde{v}}$. This means that $\overleftarrow{m}, \overrightarrow{m}$ are g-shift exponents for $g'\widetilde{\Lambda}$ with respect to $\widetilde{X}_{g'\tilde{v}}$, and $R'(g) = R(g'^{-1}gg')$.

To see that R is a homomorphism, let $g_1, g_2 \in \mathrm{Stab}(\widetilde{\Lambda})$, let $\overleftarrow{m}_1, \overrightarrow{m}_1$ be g_1-shift exponents of $\widetilde{\Lambda}$ with respect to $\widetilde{X}_{\tilde{v}}$ so $h^{\overleftarrow{m}_1}\widetilde{\Lambda} = g_1h^{\overrightarrow{m}_1}g_1^{-1}\widetilde{\Lambda}$, and let $\overleftarrow{m}_2, \overrightarrow{m}_2$ be g_2-shift exponents of $\widetilde{\Lambda}$ with respect to $\widetilde{X}_{g_1\tilde{v}}$ so $g_1h^{\overleftarrow{m}_2}g_1^{-1}\widetilde{\Lambda} = g_2g_1h^{\overrightarrow{m}_2}g_1^{-1}g_2^{-1}\widetilde{\Lambda}$. Then $h^{\overleftarrow{m}_1\overleftarrow{m}_2}\widetilde{\Lambda} = g_2g_1h^{\overrightarrow{m}_1\overrightarrow{m}_2}g_1^{-1}g_2^{-1}\widetilde{\Lambda}$. Therefore $R(g_1g_2) = \frac{\overleftarrow{m}_1\overleftarrow{m}_2}{\overrightarrow{m}_1\overrightarrow{m}_2} = R(g_1)R(g_2)$. $\square$

Definition 4.6.5. If the dilation function R has infinite image then $\widetilde{\Lambda}$ is *dilated*. Otherwise $\widetilde{\Lambda}$ is *non-dilated* (and the image of R is either trivial or $\{-1, 1\}$). An element $g \in \mathrm{Stab}\,\widetilde{\Lambda}$ is *dilated* if $R(g) \neq \pm 1$. Lemma 4.6.4 implies that if $\widetilde{\Lambda}$ is [non-]dilated, then all its G-translates are [non-]dilated. Therefore, it makes sense to say that an immersed wall Λ is [non-]dilated if its associated walls are [non-]dilated walls.

A wall $\widetilde{\Lambda}$ being dilated means there is an infinite family of walls, $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$, intersecting $\widetilde{X}_{\tilde{v}}$, that become closer together while traveling from $\widetilde{X}_{\tilde{v}}$ to $\widetilde{X}_{g\tilde{v}}$.

Considering higher powers of g, some infinite subset of these walls must intersect each other. The proof of the following proposition makes this idea precise.

Proposition 4.6.6. *If $(\widetilde{X}, \mathcal{W})$ contains a dilated wall then there are crossing sets of arbitrary finite cardinality.*

Proof. Let $\widetilde{\Lambda}_0$ be a dilated wall intersecting $\widetilde{X}_{\tilde{v}}$ and let $h \in G_{\tilde{v}}$ be perpendicular to $\widetilde{\Lambda}_0 \cap \widetilde{X}_{\tilde{v}}$. Choose $g \in \mathrm{Stab}(\widetilde{\Lambda}_0)$ such that $R(g) > 1$. By Lemma 4.6.7 the set of all G-translates of $\widetilde{\Lambda}_0$ intersecting $\widetilde{X}_{\tilde{v}}$ that do not regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$ can be partitioned:

$$\mathcal{P} = \Big\{\{h^{\overleftarrow{n}r}\widetilde{\Lambda}_0\}_{r\in\mathbb{Z}}, \ldots, \{h^{\overleftarrow{n}r}\widetilde{\Lambda}_p\}_{r\in\mathbb{Z}}\Big\},$$

where $h^{\overleftarrow{n}}\widetilde{\Lambda}_i = gh^{\overrightarrow{n}}g^{-1}\widetilde{\Lambda}_i$ and $\dfrac{\overleftarrow{n}}{\overrightarrow{n}} = R(g)$.

Let $s \geq 1$ and $\widetilde{M} = h^{\overleftarrow{n}s}\widetilde{\Lambda}_0$. We will show that $Q_s^\infty = \{\widetilde{M}, g^{-1}\widetilde{M}, g^{-2}\widetilde{M}, \ldots, \}$ contains a cardinality s crossing set. All $g^{-i}\widetilde{M}$ are distinct because if $g^{-p}\widetilde{M} = g^{-q}\widetilde{M}$ for $p \neq q$ then $h^{\overleftarrow{n}s}\widetilde{\Lambda}_0 = g^{p-q}h^{\overleftarrow{n^s}}\widetilde{\Lambda}_0 = g^{p-q}h^{\overleftarrow{n^s}}g^{q-p}\widetilde{\Lambda}_0$ which would imply $R(g^{p-q}) = 1$, but R is a homomorphism so this would contradict $R(g) > 1$. This implies that $Q_s^\infty = g(Q_s^\infty - \{\widetilde{M}\})$. Note that none of the walls $\widetilde{M}, \ldots, g^{1-s}\widetilde{M}$ regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$, since

$$g^{-i}\widetilde{M} = g^{-i}h^{\overleftarrow{n}s}\widetilde{\Lambda}_0 = g^{-i}h^{\overleftarrow{n}i\overleftarrow{n}s-i}g^i\widetilde{\Lambda}_0 = h^{\overrightarrow{n}i\overleftarrow{n}s-i}\widetilde{\Lambda}_0,$$

and the right hand side belongs to the set $\{h^{\overleftarrow{n}r}\widetilde{\Lambda}_0\}_{r\in\mathbb{Z}} \in \mathcal{P}$ for $0 \leq i \leq s-1$.

Suppose there exists k such that $g^{-k}\widetilde{M}$ regularly intersects $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$. Let k be the minimal such value and note that $k \geq s$. Thus, $g^{-i}\widetilde{M}$ does not regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$ for $0 \leq i < k$, which implies that $g^{-k}\widetilde{M}$ regularly intersects $g^{-i}\widetilde{M}$ for $0 \leq i < k$. Then $Q_s^k = \{\widetilde{M}, \ldots, g^{-k}\widetilde{M}\}$ is a crossing set of cardinality at least $s+1$, since $g^{-k}\widetilde{M}$ regularly intersects all the other elements of Q_s^k and $g^{-1}(Q_s^k - \{g^{-k}\widetilde{M}\}) = Q_s^k - \{\widetilde{M}\}$.

Suppose that $g^{-i}\widetilde{M}$ does not regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$ for all $i \geq 0$. Therefore each $g^{-i}\widetilde{M}$ belongs to some subset in the partition $\mathcal{P}$. Therefore there exists $r_i \in \mathbb{Z}$ and $\sigma(i) \in \{0, \ldots, p\}$ such that $g^{-i}\widetilde{M} = h^{\overleftarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i)}$. We will obtain a contradiction by showing that $\{r_i\}$ is a bounded sequence, contradicting that Q_s^∞ is an infinite set. From the construction of the sequence we have

$$g^{-1}h^{\overleftarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i)} = g^{-1}(g^{-i}\widetilde{M}) = g^{-(i+1)}\widetilde{M} = h^{\overleftarrow{n} r_{i+1}}\widetilde{\Lambda}_{\sigma(i+1)},$$

and from the properties of $\mathcal{P}$ it can be inferred that

$$h^{\overleftarrow{n} r_{i+1}}\widetilde{\Lambda}_{\sigma(i+1)} = g^{-1}h^{\overleftarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i)} = h^{\overrightarrow{n} r_i}g^{-1}\widetilde{\Lambda}_{\sigma(i)}.$$

Hence

$$h^{\overleftarrow{n} r_{i+1} - \overrightarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i+1)} = g^{-1}\widetilde{\Lambda}_{\sigma(i)},$$

so $a_i = \overleftarrow{n} r_{i+1} - \overrightarrow{n} r_i$ is a bounded integer sequence because the right hand side can only be one of a finite number of walls. Therefore $r_{i+1} = \frac{\overrightarrow{n}}{\overleftarrow{n}} r_i + \frac{a_i}{\overleftarrow{n}}$ is bounded since $\frac{1}{R(g)} = \frac{\overrightarrow{n}}{\overleftarrow{n}} < 1$. $\square$

The following Lemma is used in the proof of Proposition 4.6.6.

Lemma 4.6.7. *Fix $g \in \mathrm{Stab}(\widetilde{\Lambda}_0)$, and $h \in G_{\tilde{v}}$ perpendicular to $\widetilde{\Lambda}_0 \cap X_{\tilde{v}}$. Let U be the set of G-translates of $\widetilde{\Lambda}_0$ that intersect $\widetilde{X}_{\tilde{v}}$ and do not regularly intersect $\widetilde{\Lambda}_0$ in* $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$. *Then there exist $\overleftarrow{n}, \overrightarrow{n} \in \mathbb{Z} - \{0\}$ such that $R(g) = \frac{\overleftarrow{n}}{\overrightarrow{n}}$, and a subset $\{\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_p\} \subset U$ such that there is a partition of U,*

$$\mathcal{P} = \Big\{ \{h^{\overleftarrow{n} r}\widetilde{\Lambda}_0\}_{r\in\mathbb{Z}}, \ldots, \{h^{\overleftarrow{n} r}\widetilde{\Lambda}_p\}_{r\in\mathbb{Z}} \Big\},$$

and such that $h^{\overleftarrow{n}}\widetilde{\Lambda}_i = gh^{\overrightarrow{n}}g^{-1}\widetilde{\Lambda}_i$ *for each* i.

Proof. Let $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$ be the union of the consecutive vertex and edge spaces $\widetilde{X}_{\tilde{v}_0}, \widetilde{X}_{\tilde{e}_1}, \ldots, \widetilde{X}_{\tilde{e}_\ell}, \widetilde{X}_{\tilde{v}_\ell}$. For each i there exists primitive $\rho_i \in G_{\tilde{v}_i}$ such that $\mathrm{Stab}(\widetilde{\Lambda}_0) \cap G_{\tilde{v}_i} = \langle \rho_i^{k_i} \rangle$ for some $k_i \in \mathbb{Z} - \{0\}$. Let $\widetilde{\Lambda}$ be any wall in U. Since $\widetilde{\Lambda}$ doesn't regularly intersect $\widetilde{\Lambda}_0$ in $\widetilde{X}_{\tilde{v}_i}$ observe that $\mathrm{Stab}(\widetilde{\Lambda}) \cap G_{\tilde{v}_i} = \langle \rho_i^{j_i} \rangle$ for some $j_i > 0$.

Let $G_{\tilde{e}_i} = \langle g_i \rangle$ for $0 < i < \ell + 1$, and let $g_0 = h$ and $g_{\ell+1} = ghg^{-1}$. By Lemma 4.6.2 there exists $\overleftarrow{m}, \overrightarrow{m} \in \mathbb{Z} - \{0\}$ such that $h^{\overleftarrow{m}}\widetilde{\Lambda} = gh^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$ and

$$\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{\overleftarrow{m}_0 \cdots \overleftarrow{m}_\ell}{\overrightarrow{m}_0 \cdots \overrightarrow{m}_\ell} = \frac{\#[\rho_0, g_1] \cdots \#[\rho_\ell, g_{\ell+1}]}{\#[\rho_0, g_0] \cdots \#[\rho_\ell, g_\ell]},$$

where the j_i-s have canceled on the right hand side. Hence $\frac{\overleftarrow{m}}{\overrightarrow{m}} = R(g)$ because the final ratio is independent of the values of j_i. There are only finitely many values that j_i can take, so we can obtain $\overleftarrow{n}, \overrightarrow{n}$ such that for any $\widetilde{\Lambda}$, we have $h^{\overleftarrow{n}}\widetilde{\Lambda} = gh^{\overrightarrow{n}}g^{-1}\widetilde{\Lambda}$. Since there are only finitely many $G_{\tilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\tilde{v}}$ the partition is obtained by taking a finite number of representatives $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_p$, such that the sets $\{h^{\overleftarrow{n}r}\widetilde{\Lambda}_i\}_{r \in \mathbb{Z}}$ are a maximal collection of disjoint sets. $\square$

We now consider the non-dilated walls.

Lemma 4.6.8. *Let* $\widetilde{\Lambda}$ *be a non-dilated wall. For any* $\widetilde{X}_{\tilde{v}}$ *intersected by* $\widetilde{\Lambda}$, *there exists* $h \in G_{\tilde{v}}$ *perpendicular to* $\widetilde{X}_{\tilde{v}} \cap \widetilde{\Lambda}$ *such that if* $g \in \mathrm{Stab}(\widetilde{\Lambda})$ *then* $ghg^{-1}\widetilde{\Lambda} = h\widetilde{\Lambda}$, *and the walls in* $\{h^r\widetilde{\Lambda}\}_{r \in \mathbb{Z}}$ *are pairwise disjoint.*

Proof. Choose a generating set $\{f_1, \ldots, f_n\}$ for $\mathrm{Stab}(\widetilde{\Lambda})$, and fix $h_0 \in G_{\tilde{v}}$ perpendicular to $\widetilde{X}_{\tilde{v}} \cap \widetilde{\Lambda}$. By Lemma 4.6.2 and the assumption that $\widetilde{\Lambda}$ is nondilated, each f_i has an s_i such that $h_0^{s_i}\widetilde{\Lambda} = f_i h_0^{\pm s_i} f_i^{-1}\widetilde{\Lambda}$ and the walls $\{h_0^{s_i r}\widetilde{\Lambda}\}$ pairwise do not intersect in $\mathrm{Carrier}(\tilde{v}, f_i\tilde{v})$. Let $s = \mathrm{lcm}\, s_i$ and let $h = h_0^s$, then $h\widetilde{\Lambda} = gh^{\pm 1}g^{-1}\widetilde{\Lambda}$ and $\{h^r\widetilde{\Lambda}\}$ pairwise do not intersect in $\mathrm{Carrier}(\tilde{v}, g\tilde{v})$, for all $g \in \mathrm{Stab}(\widetilde{\Lambda})$, as any g can be expressed as the product of the generating set.

For any other $\widetilde{X}_{\widetilde{w}}$ intersected by $\widetilde{\Lambda}$, there is a geodesic $\gamma : I \to \widetilde{\Lambda}$ with initial point $\widetilde{a} \in \widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ and endpoint in $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}}$. Extend γ to a geodesic γ' such that the endpoint of γ' is $g\widetilde{a}$ for some $g \in \mathrm{Stab}(\widetilde{\Lambda})$. Therefore $h\widetilde{\Lambda} = gh^{\pm 1}g^{-1}\widetilde{\Lambda}$ and $h^r\widetilde{\Lambda} \cap \widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}} = \emptyset$ for all r, as $\widetilde{X}_{\widetilde{w}}$ lies in Carrier$(\tilde{v}, g\tilde{v})$. This implies that the walls in $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ are pairwise disjoint. □

Proposition 4.6.9. *Let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a finite set of non-dilated immersed walls, then $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Proof. Let $\{K_i\}$ be a sequence of finite sets of pairwise intersecting walls such that $\lim |K_i| = \infty$. By Lemma 4.5.3 there exist vertex spaces $\widetilde{X}_{\tilde{v}_i}$ such that each wall in K_i intersects $\widetilde{X}_{\tilde{v}_i}$. Since there are finitely many G-orbits of vertices in the Bass-Serre tree, choose a subsequence of $\{K_i\}$ and find a sequence $\{g_i\} \subset G$ such that all $\{g_iK_i\}$ intersect a fixed $\widetilde{X}_{\tilde{v}}$. Since there are finitely many $G_{\tilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\tilde{v}}$, we may restrict to subsets, still of unbounded cardinality, such that all the walls lie in the same $G_{\tilde{v}}$-orbit.

Fix some $\widetilde{\Lambda}$ in g_1K_1. By Lemma 4.6.8 there exists $h \in G_{\tilde{v}}$ such that $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ consists of pairwise non-intersecting walls. There are only finitely many $G_{\tilde{v}}$-orbits of $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$, so by the pigeonhole principle there must be some g_iK_i that has more than one wall in one of these orbits. This contradicts that the walls in g_iK_i pairwise intersect. □

Theorem 4.6.10. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The following are equivalent:*

1. *$C(\widetilde{X}, \mathcal{W})$ is infinite dimensional.*
2. *$C(\widetilde{X}, \mathcal{W})$ contains an infinite cube.*
3. *Some immersed wall is dilated.*

Proof. (3) $\Rightarrow$ (2) follows from Proposition 4.6.6 and Proposition 4.5.4.

(2) $\Rightarrow$ (1) is immediate from the definition of an infinite cube.

(1) $\Rightarrow$ (3) is the contrapositive of 4.6.9. $\square$

Corollary 4.6.11. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a set of immersed walls in X. It is decidable whether or not $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Proof. For each immersed wall Λ, the corresponding dilation function can be computed by using Lemmas 4.6.1 and 4.6.2 to find their values on a finite generating set of $\pi_1\Lambda$. $\square$

The following is the result of combining Theorem 4.6.10, Proposition 4.6.6, and Proposition 4.5.4.

Proposition 4.6.12. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. If $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional, then $\mathcal{W}$ contains an set of pairwise regularly intersecting walls of infinite cardinality, that correspond to the hyperplanes in an infinite cube in $C(\widetilde{X}, \mathcal{W})$. Moreover, the infinite cube contains a canonical 0-cube.*

The following characterization of non-dilated immersed walls will be used in Chapter 5.

Proposition 4.6.13. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of non-dilated immersed walls in X. The horizontal walls in $\mathcal{W}$ can be partitioned into a collection $\mathcal{A}$ of subsets such that:*

1. *The partition is preserved by G,*
2. *The walls in A are pairwise non-intersecting for each $A \in \mathcal{A}$,*
3. *Let $\widetilde{\Lambda} \in A \in \mathcal{A}$ be a wall intersecting $\widetilde{X}_{\tilde{v}}$. There exists $h \in G_{\tilde{v}}$ perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ such that $A = \{h^r \widetilde{\Lambda}\}_{r \in \mathbb{Z}}$.*

Proof. Let $\widetilde{\Lambda}$ be a horizontal wall in $\mathcal{W}$. There are finitely many $\mathrm{Stab}(\widetilde{\Lambda})$-orbits of lines of the form $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}}$. Let $\{\ell_i\}_{i=1}^n$ be a set of representatives for these orbits and let $\{\tilde{v}_i\}_{i=1}^n$ be vertices in $\widetilde{\Gamma}$ such that $\ell_i = \widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}_i}$. By Proposition 4.6.8, for each i there exist $h_i \in G_{\tilde{v}_i}$ perpendicular to ℓ_i such that $h_i\widetilde{\Lambda} = gh_ig^{-1}\widetilde{\Lambda}$ for all $g \in \mathrm{Stab}(\widetilde{\Lambda})$, and each $A_i = \{h_i^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ is a collection of pairwise disjoint walls. Let $\mathcal{A}_i = \{gA_i \mid g \in G\}$. If $A_i \cap gA_i \neq \emptyset$ then there exist r, s such that $h_i^s\widetilde{\Lambda} = gh_i^r\widetilde{\Lambda}$ which implies $h_i^{-s}gh_i^r \in \mathrm{Stab}(\widetilde{\Lambda})$. Therefore, for all $t \in \mathbb{Z}$

$$h_i^{-s}gh_i^{t+r}\widetilde{\Lambda} = (h_i^{-s}gh_i^r)h_i^t(h_i^{-s}gh_i^r)^{-1}\widetilde{\Lambda} = h_i^{\pm t}\widetilde{\Lambda},$$

so $gh_i^{t+r}\widetilde{\Lambda} = h_i^{\pm t+s}\widetilde{\Lambda}$, and so $gA_i = A_i$. Therefore each $\mathcal{A}_i$ is a partition of the G-orbit of $\widetilde{\Lambda}$ satisfying (1) and (2).

By Lemma 4.6.2, there exists p_i, q_i such that $h_1^{p_i}\widetilde{\Lambda} = h_i^{q_i}\widetilde{\Lambda}$. Let $p = p_1 \cdots p_n$, and $\hat{q}_i = p_1 \cdots p_{i-1}q_ip_{i+1}\cdots p_n$. Let $h = h_1^p$. Let $A = \{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ and $\mathcal{A} = \{gA \mid g \in G\}$. Note that $\mathcal{A}$ is a common refinement of the partitions $\mathcal{A}_i$, still satisfying (1) and (2) since $\{h_i^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}} \supseteq \{h_i^{\hat{q}_ir}\widetilde{\Lambda}\}_{r\in\mathbb{Z}} = \{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$. Condition (3) holds since if $\widetilde{\Lambda}$ intersects $\widetilde{X}_{\widetilde{w}}$ then there exists i and $f \in \mathrm{Stab}(\widetilde{\Lambda})$ such that $\widetilde{w} = f\tilde{v}_i$ so $A = \{fh_i^{\widehat{q_i}r}f^{-1}\widetilde{\Lambda}\}$. $\square$

Any partition of the horizontal walls in $\mathcal{W}$ satisfying conditions (1)-(3) in Proposition 4.6.13 will be called a *stable partition.*

Lemma 4.6.14. *Let X be a tubular space and $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a finite set of immersed walls in X. Let $\mathcal{P}$ be a stable partition of the horizontal walls in $\mathcal{W}$. Then for each $\tilde{v} \in \widetilde{\Gamma}$, only finitely many $A \in \mathcal{P}$ contain walls intersecting $\widetilde{X}_{\tilde{v}}$.*

Proof. Suppose that $\widetilde{\Lambda}$ is a wall intersecting $\widetilde{X}_{\tilde{v}}$, then, by condition (3) of a stable partition, there exists some $h \in G_{\tilde{v}}$ that is perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ such that $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}} \in \mathcal{P}$. By G-invariance we can deduce that each of the

$G_{\tilde{v}}$-translates of $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ is also in $\mathcal{P}$. There are only finitely many such translates, therefore each $G_{\tilde{v}}$-orbit of a wall in $\widetilde{X}_{\tilde{v}}$ is contained in finitely many elements of $\mathcal{P}$. The claim then follows from the fact that there are only finitely many $G_{\tilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\tilde{v}}$. □

4.7 Computing the Dilation Function

In this section the results of Section 4.6 are applied to concrete examples to determine whether or not they are finite dimensional. In this section we assume that all immersed walls are non-trivial, in the sense that they do not consist of a single immersed circle. This allows us to identify $\mathrm{Stab}(\widetilde{\Lambda})$ with $\pi_1\Lambda$. The main focus will be on computing the dilation function $R : \pi_1\Lambda \to \mathbb{Q}^*$ of an immersed wall. Let $q : \Lambda \to \Omega$ be the quotient map obtained by quotienting each circle in the equitable set to a vertex. This simplifies the computation since R factors through $q_* : \pi_1\Lambda \to \pi_1\Omega$.

Regard Ω as a directed graph by fixing orientations of its edges $E(\Omega)$, choosing the orientation of each edge σ to be consistent with all other edges mapping into the same edge space of X. Recall that φ_e^- and φ_e^+ denote the attaching maps of the edge space X_e in X. Define a *weighting* $\omega : E(\Omega) \to \mathbb{Q}^*$ as follows: for each directed edge $\sigma \in E(\Omega)$ let $\omega(\sigma) = \frac{\#[C_\iota:\varphi_e^-]}{\#[C_\tau:\varphi_e^+]}$, where X_e is the edge space σ maps into, and C_ι, C_τ are the elements in the equitable set attached to the initial and terminal ends of σ. Let γ be a combinatorial path in Λ representing an element $[\gamma] \in \pi_1\Lambda$, and let $\overline{\gamma}$ be the combinatorial path in Ω obtained by quotienting the circle-edges of γ to vertices, then $q\circ\gamma$ is homotopic to $\overline{\gamma}$. If $\overline{\gamma} = \sigma_1^{\epsilon_1}\cdots\sigma_n^{\epsilon_n}$ with $\epsilon \in \{\pm 1\}$, then $R([\gamma]) = \omega(\sigma_1)^{\epsilon_1}\cdots\omega(\sigma_n)^{\epsilon_n}$. Note that this does not depend on the choice of representative γ. See Figure 4–6 for examples of this quotient and edge weighting.

Example 4.7.1. Let $G = \langle a, b, s, t \mid [a, b], s^{-1}abs = a^2, t^{-1}abt = b^2\rangle$. Let X be the tubular space such that $\pi_1X = G$. Note that there is a single vertex group

$G_v = \pi_1 X_v$. The group G is a free-by-cyclic tubular group that acts freely on a CAT(0) cube complex, but cannot act on a finite dimensional CAT(0) cube complex obtained from an equitable set. The proof uses the dilation function.

The equitable set $S_v = \{a, b\}$ demonstrates that G acts freely on a CAT(0) cube complex, but will only produce dilated walls so the dual cube complex will be infinite dimensional. It is possible to construct non-dilated immersed walls in X, by taking any collection of disjoint curves in the vertex space X_v that represent powers of ab^{-1}. However, such a collection is not sufficient for an equitable set because such elements will not generate a finite index subgroup of G_v. We therefore claim that any immersed wall containing a circle that doesn't represent a power of ab^{-1} is dilated since any equitable set for G will produce such an immersed wall.

Suppose there is an immersed wall Λ obtained from a subset of curves $S = \{v_1, \ldots, v_n\}$ contained in a equitable set, where $v_i = a^{x_i} b^{y_i}$ where we assume $v_1 \neq a^n b^{-n}$. The equitable set S must satisfy the equations

$$2\sum_{i=1}^{n} |x_i| = 2\sum_{i=1}^{n} |y_i| = \sum_{i=1}^{n} |x_i - y_i|. \tag{$*$}$$

The claim follows by finding a closed path γ in Ω such that $R([\gamma]) \neq \pm 1$. Direct the edges of Ω such that edges exiting the vertex space via the attaching maps ab are the initial ends. Therefore the number of edges leaving the vertex corresponding to v_i is $2|x_i - y_i|$, while the number of vertices arriving at that vertex is $2|x_i| + 2|y_i|$. If some $|x_i| + |y_i| > |x_i - y_i|$, then since each $|x_i| + |y_i| \geq |x_i - y_i|$ we have

$$\sum_{i=1}^{n} |x_i| + |y_i| > \sum_{i=1}^{n} |x_i - y_i|,$$

which would contradict $(*)$. Therefore the number of edges entering and exiting each vertex in Ω are equal, so there exists a *directed Eulerian trail* γ, which is a cycle traversing each edge in Ω precisely once. Therefore

$$|R([\gamma])| = \prod_{i=1}^{n} \frac{|x_i - y_i|^{2|x_i - y_i|}}{(2|x_i|)^{2|x_i|}(2|y_i|)^{2|y_i|}},$$

and the claim will be proven by showing that $|R([\gamma])| < 1$. Considering each factor in the product separately we want to show that

$$\frac{|x-y|^{2|x-y|}}{(2|x|)^{2|x|}(2|y|)^{2|y|}} \leq 1$$

with equality only when $x = -y$. Since $v_1 \neq a^n b^{-n}$, the first term in the product is strictly less than one, so $|R([\gamma])| < 1$.

This is immediate in the trivial case when $|x - y| = 0, |x| = 0$, or $|y| = 0$. After taking the square root and applying the logarithm, the non-trivial case is equivalent to showing that

$$|x-y| \log |x-y| \leq |x| \log(2|x|) + |y| \log(2|y|), \tag{$**$}$$

with equality only when $x = -y$. As $z \log(z)$ is strictly convex for $z > 0$, the following inequality holds with equality when $p = q$;

$$\frac{|p+q|}{2} \log\left(\frac{|p+q|}{2}\right) \leq \frac{|p|}{2} \log|p| + \frac{|q|}{2} \log|q|.$$

Thus $(**)$ holds by letting $p = 2x$ and $q = -2y$.

Example 4.7.2. The Right Angled Artin Group $A = \langle a, b, c, d \mid [a, b] = [b, c] = [c, d] = 1 \rangle$ is not subgroup separable (see [27]). From the alternative presentation $A = \langle a_i, b_i \mid [a_i, b_i], a_1 = a_2, b_2 = b_3; 1 \leq i \leq 3 \rangle$ it is clear that A is fundamental group of a tubular group given by three tori and two cylinders. Consider an immersed wall with the following equitable set: $\{a_1 b_1^2, b_1\}, \{a_2 b_2^{-1}, a_2 b_2^2\}$ and $\{a_3^2\}$, and the arcs given in Figure 4–5. The

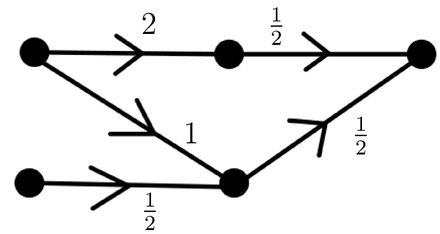

Figure 4–5: The graph Ω from Example 4.7.2.

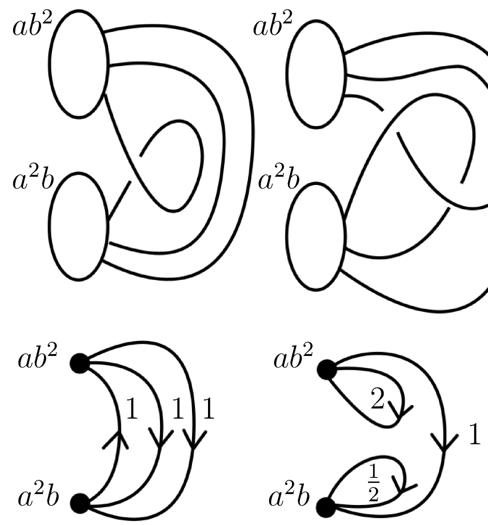

Figure 4–6: Example 4.7.3. The immersed walls Λ_1, Λ_2 are shown on the top, and the associated quotient Ω_i is shown below with its orientation and weighting. The immersed wall on the left is non-dilated and the immersed wall on the right is dilated.

diagram shows the underlying graph Ω with the associated edge weighting. There is a single simple closed path with $R([\gamma]) = (2) \cdot (1/2) \cdot (1/2)^{-1} \cdot (1) = 2$, therefore $\widetilde{\Lambda}$ is dilated and $\mathrm{Stab}(\widetilde{\Lambda})$ is not separable (see discussion at the start of this chapter). Furthermore by Theorem 4.6.10 there must exist an infinite cube in the associated cubulation.

Example 4.7.3. The group $G = \langle a, b, t \mid [a, b], t^{-1}at = b \rangle$ from Example 4.4.1 demonstrates that the choice of arcs connecting circles can change whether or not a cubulation is finite dimensional. Consider the equitable set $\{a^2b, ab^2\}$ which extends to both a dilating wall and a non-dilating wall shown in Figure 4–6.

The following lemma will be made use of in Chapter 5

Lemma 4.7.4. *Let X be a tubular space, and let $G = \pi_1 X$. Let $\Lambda_1, \dots \Lambda_k$ be a set of immersed walls in X obtained from an equitable set $\{S_v\}_{v \in \Gamma}$. Then there exists a set of primitive immersed walls $\Lambda'_1, \dots, \Lambda'_\ell$ in X obtained from an equitable set $\{S'_v\}_{v \in \Gamma}$. Moreover:*

1. *If $\Lambda_1, \dots, \Lambda_k$ are non-dilated, then so are $\Lambda'_1, \dots \Lambda'_\ell$.*
2. *If $\Lambda_1, \dots \Lambda_k$, are fortified, then so are $\Lambda'_1, \dots \Lambda'_\ell$.*

Proof. Each Λ_i decomposes as the union of disjoint circles ,which are the domain of locally geodesic closed paths in the equitable set, and arcs. Suppose that $\alpha^n \in S_v$, where $[\alpha] \in \pi_1 X_v = G_v$ is primitive. Let Λ_i be the immersed wall containing the circle $S^1_{\alpha^n}$ corresponding to α^n.

A new equitable set is obtained by replacing α^n in S_v with n locally geodesic curves $\{\alpha_i : S^1_i \to X_v\}_{i=1}^n$ with disjoint images in X_v that are isotopic to α in $\widetilde{X}_v$. This remains an equitable set since $\#[\alpha^n, \gamma] = n\#[\alpha, \gamma] = \sum_{i=1}^n \#[\alpha_i, \gamma]$ for any locally geodesic curve γ in X_v. New immersed walls are obtained from Λ_i by replacing $S^1_{\alpha^n}$ with $S^1_1, \dots S^1_n$ and reattaching the arcs that were attached to the intersection points in $S^1_{\alpha^n}$ to the corresponding intersection points on $S^1_1, \dots S^1_n$. Let $\Lambda_{i1}, \dots, \Lambda_{i\ell}$ be the new set of immersed walls obtained in this way. Note that each arc in $\Lambda_{i1}, \dots, \Lambda_{i\ell}$, corresponds to a unique arc in Λ_i.

Assume that Λ_i is non-dilated. We claim that the new immersed walls $\Lambda_{i1}, \dots, \Lambda_{i\ell}$ are also non-dilated. Let $q_i : \Lambda_i \to \Omega_i$ and $q_{ij} : \Lambda_{ij} \to \Omega_{ij}$ be the quotient maps obtained by crushing the circles to vertices. Let u be the vertex in Ω_i corresponding to $S^1_{\alpha^n}$. Let $R_i : \pi_1 \Lambda_i \to \mathbb{Q}^*$ and $R_{ij} : \pi_1 \Lambda_{ij} \to \mathbb{Q}^*$ be the dilation functions. Let $\hat{R}_i : \pi_1 \Omega_i \to \mathbb{Q}^*$ and $\hat{R}_{ij} : \pi_1 \Omega_{ij} \to \mathbb{Q}^*$ be the unique maps such that $R_i = \hat{R}_i \circ q_i$ and $R_{ij} = \hat{R}_{ij} \circ q_{ij}$. Let ω_i and ω_{ij} be the respective weighting of the arcs in Ω_i and Ω_{ij}. By assumption, R_i and $\hat{R}_i$ have finite image. As the arcs in Λ_{ij} correspond to arcs in Λ_i, there is a map

$\rho_i : \Omega_{ij} \to \Omega_i$. We show Λ_{ij} is non-dilated by showing that if $[\gamma] \in \pi_1\Omega_{ij}$ then $\hat{R}_{ij}([\gamma]) = \hat{R}_i([\rho_i \circ \gamma])$.

Let σ be an oriented arc in Ω_{ij}. The edge $q_{ij}^{-1}(\sigma)$ embeds in an edge space X_e. If the vertices of σ^ϵ are disjoint from $\rho^{-1}(u)$, then $\omega_{ij}(\sigma^\epsilon) = \omega_i(\rho_{ij} \circ \sigma^\epsilon)$. If the endpoints of σ^ϵ are contained in $\rho_{ij}^{-1}(v)$, and correspond to the circles S^1_ι and S^1_τ then

$$\omega_{ij}(\sigma^{\pm 1}) = \frac{\#[\varphi_e^{\mp}, \alpha_\iota]}{\#[\varphi_e^{\pm}, \alpha_\tau]} = \frac{n\#[\varphi_e^{\mp}, \alpha_\iota]}{n\#[\varphi_e^{\pm}, \alpha_\tau]} = \frac{\#[\varphi_e^{\mp}, \alpha^n]}{\#[\varphi_e^{\pm}, \alpha^n]} = \omega(\rho_{ij} \circ \sigma^{\pm 1}).$$

Suppose that exactly one endpoint of σ^ϵ is contained in $\rho_{ij}^{-1}(u)$. If σ^ϵ terminates a vertex in $\rho_{ij}^{-1}(u)$ corresponding to S^1_τ, and the initial vertex corresponds to a circle which is the domain of a locally geodesic curve β then

$$\omega_{ij}(\sigma^{\pm 1}) = \frac{\#[\varphi_e^{\mp}, \beta]}{\#[\varphi_e^{\pm}, \alpha_\tau]} = \frac{n\#[\varphi_e^{\mp}, \beta]}{n\#[\varphi_e^{\pm}, \alpha_\tau]} = \frac{n\#[\varphi_e^{\mp}, \beta]}{\#[\varphi_e^{\pm}, \alpha^n]} = n\omega(\rho_{ij} \circ \sigma^{\pm 1}).$$

If σ^ϵ starts at a vertex in $\rho_{ij}^{-1}(u)$ corresponding to S^1_ι, and the terminal vertex correspond to a circle that is the domain of a locally geodesic curve β then

$$\omega_{ij}(\sigma^{\pm 1}) = \frac{\#[\varphi_e^{\mp}, \alpha_\iota]}{\#[\varphi_e^{\pm}, \beta]} = \frac{n\#[\varphi_e^{\mp}, \alpha_\iota]}{n\#[\varphi_e^{\pm}, \beta]} = \frac{\#[\varphi_e^{\mp}, \alpha^n]}{n\#[\varphi_e^{\pm}, \beta]} = \frac{1}{n}\omega(\rho_{ij} \circ \sigma^{\pm 1}).$$

Therefore, given an edge path γ in Ω_{ij}, since the number of edges exiting vertices in $\rho_{ij}^{-1}(v)$ is the same as the number of vertices entering, $\hat{R}_{ij}(\gamma) = \hat{R}_i(\rho_i \circ \gamma)$.

This procedure produces immersed walls with one fewer non-primitive element in the equitable set. Repeating this procedure for each non-primitive element in the equitable set produces a primitive set of non-dilated immersed walls. It is also clear, that if $\Lambda_1, \ldots, \Lambda_k$ are fortified, then so are the new immersed walls. □

4.8 Dilated Walls are not Quasi-Isometrically Embedded

Using Theorem 4.6.10, Theorem 4.1.2 can be restated in the following form.

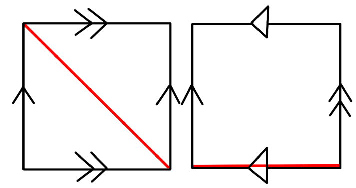

Figure 4–7: The tubular space and immersed wall in Example 4.8.2

Theorem 4.8.1. *If* $\widetilde{\Lambda}$ *is dilated then* $\mathrm{Stab}(\widetilde{\Lambda}) \leqslant \pi_1 X$ *is not quasi-isometrically embedded.*

The following example motivates the proof of Theorem 4.8.1, and illustrates that a non-dilated wall can fail to be quasi-isometrically embedded.

Example 4.8.2. Consider the tubular group $G = \langle a, b, t \mid [a, b], t^{-1}at = b \rangle$ from Example 4.4.1. Let Λ be the immersed wall containing a single circle representing $a^{-1}b$, with a single arc whose endpoints are attached to the intersection point at either end of the edge space (see Figure 4–7). Observe that Λ is embedded, and therefore non-dilated. Let $H = \mathrm{Stab}(\widetilde{\Lambda}) \leqslant G$. Note that H is a rank 2 free group, with basis $\{a^{-1}b, t\}$. We will show that H is quadratically distorted in G.

Define $\alpha_0 = 1$ and for each $n \geq 1$ define $\alpha_n = t(a^{-1}b)^n \alpha_{n-1}$. This sequence of elements has the property that $|\alpha_n|_H = \sum_{i=1}^{n}(i+1) = \frac{1}{2}n(n+1) + n$. The element $\beta_n = a^n(ta^{-1})^n \in G$ satisfies $|\beta|_G \leq 3n$. It can be verified by induction that $\alpha_n = \beta_n$. Therefore H is a quadratically distorted subgroup of G. Figure 4–8 illustrates the paths these elements correspond to in the universal cover: α_n spirals outwards increasing in length quadratically, while β_n cuts through inside, increasing linearly in length before traveling out by another linear factor.

Before embarking on the proof of Theorem 4.8.1 we outline the strategy employed.

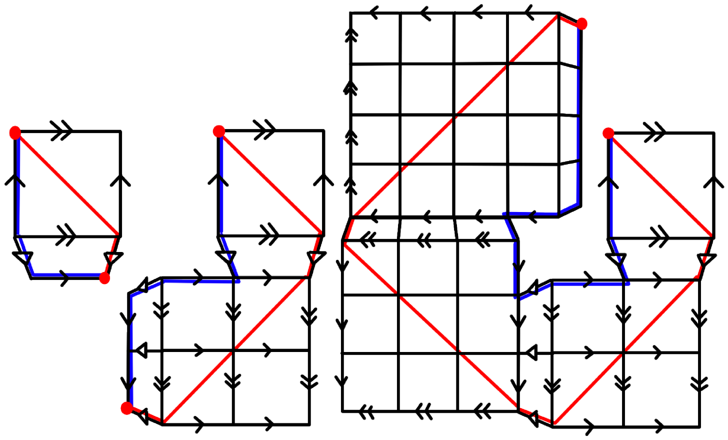

Figure 4–8: The paths corresponding to α_1, β_1, and α_2, β_2, and α_4, β_4

1. Inspired by the geometric interpretation of Example 4.8.2, construct a sequence of spiraling paths γ_m in Λ determined by a sequence of parameters.
2. Choose the parameters of γ_m such that the length of γ_m in Λ is a quadratic function of m.
3. Double up γ_m, as illustrated in Figure 4–10, to produce a second sequence of paths ρ_m that spiral out before spiraling back in, with the length growing quadratically in Λ. This stage is necessary because in Example 4.8.2 the rate at which γ_m spiraled outwards was linear.
4. Specify a path homotopy from ρ_m to ρ'_m, where the length of ρ'_m in X is a linear function of m.

Proof of Theorem 4.8.1. Let Λ be a dilated immersed wall in X. Let $z \in \pi_1\Lambda$ satisfy $|R(z)| \neq 1$, and let γ be a closed immersed path in Λ representing z.

Stage 1: Decompose γ as

$$\gamma = x_0 \cdot y_1 \cdot x_1 \cdot y_2 \cdots x_{\ell-1} \cdot y_\ell$$

where each x_i is a path in a circle of Λ immersed in a vertex space X_{u_i}, and y_i is an arc of Λ embedded in an edge space X_{e_i}. As γ is a closed path, the terminal point of y_ℓ is the initial point of x_0. The subpath x_i is *diagonal* if $X^+_{e_i}$ and $X^-_{e_{i+1}}$ are not parallel in X_{u_i}, and the subpath $y_i \cdot x_i \cdots x_{i+r} \cdot y_{i+r+1}$ is *straight* if $x_i, x_{i+1}, \ldots, x_{i+r}$ are not diagonal. The subscripts are considered modulo ℓ so that subpaths of γ containing the initial point are considered. Thus, after cyclically parameterizing, γ decomposes as

$$\gamma = a_0 \cdot b_1 \cdot a_1 \cdot b_2 \cdots a_{n-1} \cdot b_n$$

where each a_i is diagonal and each b_i is straight. For notational purposes later these subscripts are considered modulo n.

Claim 1. There is at least one diagonal subpath in γ.

Proof. Let $\widetilde{\Lambda}$ be a wall in $\widetilde{X}$ stabilized by $\pi_1\Lambda$, and $\widetilde{X}_{\widetilde{u}_0}$be a vertex space intersected by $\widetilde{\Lambda}$. Let $\widetilde{u}_1, \ldots, \widetilde{u}_{\ell-1}$ be the vertices in $\widetilde{\Gamma}$ on the geodesic between $\widetilde{u}_0$ and $\widetilde{u}_\ell = z\widetilde{u}_0$. Let $\tilde{e}_i$ be the edge between $\widetilde{u}_{i-1}$ and $\widetilde{u}_i$. Let $\langle g_i \rangle = G_{\tilde{e}_i}$ for $1 \leq i \leq \ell$, let $g_0 \in G_{\widetilde{u}_i}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{u}_0}$ and $g_{\ell+1} = zg_0z^{-1}$. Let $\langle \rho_i \rangle = \mathrm{Stab}(\widetilde{\Lambda}) \cap G_{\widetilde{u}_i}$ for $0 \leq i \leq \ell$. Suppose that the claim is false. This would imply that exists a primitive $h_i \in G_{\widetilde{u}_i}$ and $p_i, q_i \in \mathbb{Z} - \{0\}$ such that $g_i = h_i^{p_i}$ and $g_{i+1} = h_i^{q_i}$ for $1 \leq i \leq \ell - 1$ and $g_\ell = zh_0^{p_\ell}z^{-1}$ and $g_1 = h_0^{q_\ell}$. Since G acts freely on a CAT(0) cube complex it cannot contain a subgroup isomorphic to $\langle r, t \mid t^{-1}r^ntr^m \rangle$ where $n \neq \pm m$ [19]. Therefore $\frac{p_1 \cdots p_\ell}{q_1 \cdots q_\ell} = \pm 1$, so Lemma 4.6.2 says

$$R(z) = \prod_{i=0}^{\ell} \frac{\#[\rho_i, g_{i+1}]}{\#[\rho_i, g_i]} = \frac{\#[\rho_0, h_0^{q_\ell}]}{\#[\rho_0, g_0]} \frac{\#[\rho_\ell, zg_0z^{-1}]}{\#[\rho_\ell, zh_0^{p_\ell}z^{-1}]} \prod_{i=1}^{\ell-1} \frac{\#[\rho_i, h_i^{q_i}]}{\#[\rho_i, h_i^{p_i}]} = \frac{q_1 \cdots q_\ell}{p_1 \cdots p_\ell} = \pm 1,$$

since $\#[\rho_0, h_0] = \#[\rho_\ell, zh_0z^{-1}]$ and $\#[\rho_0, g_0] = \#[\rho_\ell, zg_0z^{-1}]$. This contradicts our choice of γ. $\square$

Parameterise the circle in Λ containing a_i as an immersed path c_i with $c_i(0) = c_i(1) = a_i(1)$ and such that if $a_i(0) \neq a_i(1)$ there is an immersed path

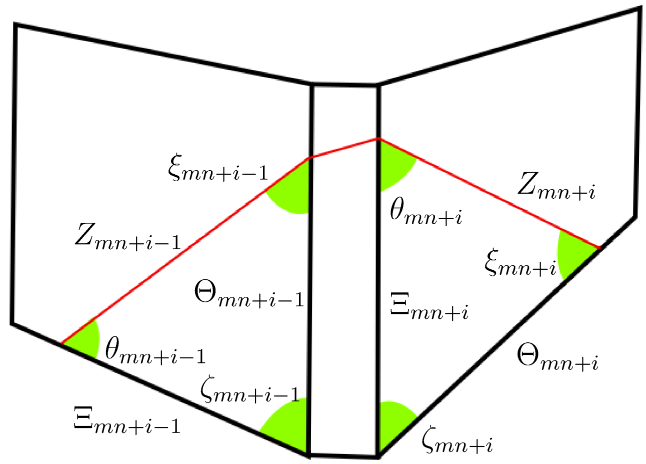

Figure 4–9: A portion of γ_m traversing $X_{v_{i-1}}$, Y_i, and X_{v_i} is illustrated schematically with the corresponding side lengths labeled.

a'_i satisfying $c_i^{k_i} = (a'_i)^{-1} \cdot a_i$ for some $k_i \geq 1$. If $a_i(0) = a_i(1)$ then let $c_i^{k_i} = a_i$, and let $a'_i = a_i^{-1}$. Note that $a_i(0) = a'_i(0)$ and $a_i(1) = a'_i(1)$. Define a sequence of immersed paths $\{\gamma_m : I \to \Lambda\}$ inductively by setting $\gamma_{-1} = a_0(0)$, and

$$\gamma_m = \gamma_{m-1} \cdot \prod_{i=0}^{n-1} \hat{a}_{mn+i} \cdot c_i^{s(mn+i)} \cdot b_{i+1}$$

where $\hat{a}_{mn+i} \in \{a_i, a'_i\}$ and $s(mn+i) \in \mathbb{Z} - \{0\}$. Both $\hat{a}_{mn+i}$ and $s(mn+i)$ will be defined inductively to replicate the spiralling effect in Example 4.8.2. Note that for γ_m to be an immersed path, if $\hat{a}_{mn+i} = a_i$ then $s(mn+i) > 0$, and if $\hat{a}_{mn+i} = a'_i$ then $s(mn+i) < 0$.

Stage 2: Let X_{v_i} be the vertex space containing the image of a_i. Observe that for each b_i there is a cylinder $Y_i \looparrowright X$, that b_i factors through, where Y_i is a union of cylinders that cover edge spaces and cylinders that immerse in the vertex spaces. Let Y_i^- and Y_i^+ denote the boundary components of Y_i with $f_i^- : Y_i^- \to X_{v_{i-1}}$ and $f_i^+ : Y_i^+ \to X_{v_i}$ denoting the corresponding restrictions of the immersion. When γ traverses every vertex space diagonally then each Y_i is an edge space.

We choose $\hat{a}_{mn+i} \in \{a_i, a_i'\}$ inductively by first setting $\hat{a}_1 = a_1$. Assuming that $\hat{a}_{mn+i-1}$ has been chosen we specify that $s(mn+i-1) \in \mathbb{Z}-\{0\}$ is positive or negative so that $\hat{a}_{mn+i-1} \cdot c_i^{s(mn+i-1)}$ is a locally geodesic path in X_{v_i}. Choose $\hat{a}_{mn+i} \in \{a_i, a_i'\}$ such that the terminal point of $\hat{a}_{mn+i-1} \cdot c_i^{s(mn+i-1)}$ meets Y_i^- with an acute angle on the same side as the initial point of $\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}$ leaves Y_i^+. This ensures that γ_m spirals in a consistent direction.

Inside the torus X_{v_i} the paths f_i^+, f_{i+1}^- and c_i determine a triangle with angles θ_i between the f_i^+ and c_i sides, ξ_i between the f_{i+1}^- and c_i sides and ζ_i between the f_i^+ and f_{i+1}^- sides. The rule of sines states

$$\frac{Z_i}{\sin \zeta_i} = \frac{\Xi_i}{\sin \xi_i} = \frac{\Theta_i}{\sin \theta_i}$$

where the value of the numerator is the length of the side opposite the angle in the denominator (see Figure 4–9). Let $Z_{mn+i} = |\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|$. The subscripts of ζ_i, ξ_i, and θ_i are considered modulo n, while the subscripts of Z_i, Ξ_i, and Θ_i are not. Note that we chose $\hat{a}_{mn+i}$ and the sign of $s(mn+i)$ so that $0 < \xi_i, \theta_i < \pi/2$.

Let $r_i = \frac{1}{|f_i^-|}\frac{\sin \xi_i}{\sin \zeta_i}$, and $t_i = \frac{1}{|f_i^+|}\frac{\sin \theta_i}{\sin \zeta_i}$ with $r_{i+n} = r_i$ and $t_{i+n} = t_i$. Define $s(mn+i) \in \mathbb{Z} - \{0\}$ inductively by setting $s(0) = 1$, and assuming that $s(mn+i-1)$ is defined choose $|s(mn+i)|$ to be large enough that

$$\begin{aligned} 1 \leq\ & r_i Z_{mn+i} - t_{i-1} Z_{mn+i-1} \\ &= r_i\,(|\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|) -\ t_{i-1}(|\hat{a}_{mn+i-1} \cdot c_{i-1}^{s(mn+i-1)}|) \end{aligned} \tag{†}$$

and small enough that

$$r_i\,(|\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|) -\ t_{i-1}(|\hat{a}_{mn+i-1} \cdot c_{i-1}^{s(mn+i-1)}|) \leq 1 + r_i |c_i^{1+k_i}|. \tag{‡}$$

Combining (†) and (‡), replacing $|\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|$ with Z_{mn+i}, and applying the rule of sines produces the following inequality:

$$1 \leq \frac{\Xi_{mn+i}}{|f_i^-|} - \frac{\Theta_{mn+i-1}}{|f_i^+|} \leq 1 + r_i|c_i^{1+k_i}|. \tag{$\star$}$$

Geometrically, this means that $s(mn+i)$ is defined so that consecutive triangles in the sequence have adjacent sides increasing at a bounded rate (see Figure 4–9). The $|f_i^-|$ and $|f_i^+|$ accounts for the different lengths of the attaching maps of Y_i. To estimate a lower bound on the length of γ_m, first rewrite (†) as

$$Z_{mn+i} \geq r_i^{-1}\Big[1 + t_{i-1}Z_{mn+i-1}\Big].$$

Without loss of generality $\prod_{i=0}^{n-1} r_i^{-1}t_i = \prod_{i=0}^{n-1} \frac{\sin\theta_i}{\sin\xi_i}\frac{|f_i^-|}{|f_i^+|} \geq 1$, otherwise replace γ with γ^{-1}, which switches θ_i for ξ_i and f_i^- for f_i^+. Applying this formula recursively produces

$$Z_{mn+i} \geq r_i^{-1}\Big[1 + t_{i-1}r_{i-1}^{-1}\Big[1 + t_{i-2}r_{i-2}^{-1}\Bigg[\cdots\Bigg[1 + t_{i-n}Z_{mn+i-n}\Bigg]\cdots\Bigg]\Big]\Big],$$

hence

$$Z_{mn+i} \geq \left[\prod_{i=0}^{n-1} r_i^{-1}t_i\right] Z_{m(n-1)+i} + \cdots \;\geq\; Z_{m(n-1)+i} + D,$$

where $D > 0$ is a lower bound on the lower order terms, which are independent of any value of Z_{mn+i}, and the value of i. Therefore there is a lower bound on the length of γ_m:

$$|\gamma_m| \geq \sum_{j=0}^{m(n+1)-1} Z_j \geq \sum_{j=0}^{m} jD \geq \frac{D}{2}m(m+1).$$

Hence the length of γ_m grows at least quadratically.

Stage 3: Define $\rho_m := \gamma_m \cdot a_0 \cdot c_0 \cdot a_0^{-1} \cdot \gamma_m^{-1}$, an immersed path with length growing quadratically with m. See Figure 4–10.

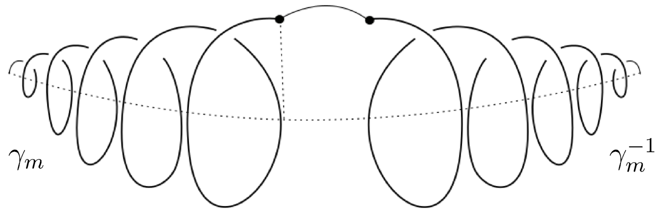

Figure 4–10: The spiral γ_m is doubled up.

Stage 4: A path homotopy of ρ_m in X will be described in two stages, and the resulting curve will have length a linear function of m. Compare the following description to Figure 4–11.

Step 1: Path homotope each diagonal segment $\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}$ in γ_m across the vertex space X_{v_i} to a path $p_{mn+i} \cdot q_{mn+i}$ where p_{mn+i} is a geodesic with image in $X_{v_i} \cap Y_i$ and q_{mn+i} is a geodesic with image in $X_{v_i} \cap Y_{i+1}$ (see the top right diagram in Figure 4–11). These paths respectively have lengths Ξ_{mn+i} and Θ_{mn+i}. The same path homotopies are made to γ_m^{-1}.

Step 2: Path homotope each $q_{mn+i-1} \cdot b_i \cdot p_{mn+i}$ factoring through Y_i, and the corresponding inverse path, to a local geodesic factoring through Y_i joining the initial and terminal points (see the bottom right diagram in Figure 4–11). From the inequality $(\star)$, we can deduce that the resulting path has length at most $|b_{i-1}| + |f_i^+|\left(1 + r_i|c_i^{1+k_i}|\right)$. The segment $q_{mn-1} \cdot b_n \cdot a_0 \cdot c_0 \cdot a_0^{-1} \cdot b_n^{-1} \cdot q_{mn-1}^{-1}$ can similarly be homotoped since $b_n \cdot a_0 \cdot c_0 \cdot a_0^{-1} \cdot b_n^{-1}$ factors through a cylinder $Y \looparrowright X$.

Therefore the length of the resulting curve ρ'_m is less than

$$2mn \max_{0 \le i < n} \left\{ |b_{i-1}| + |f_i^+|\left(1 + r_i|c_i^{1+k_i}|\right) \right\} + |a_0 \cdot c_0 \cdot a_0^{-1}|$$

which is a linear function of m. Thus $\widetilde{\Lambda}$ is distorted. $\square$

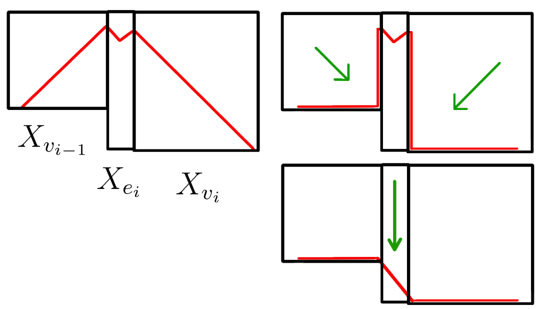

Figure 4–11: The left diagram shows $X_{v_{i-1}}$, Y_i, and X_{v_i} containing $a_{i-1} \cdot c_{i-1}^{s(mn+i-1)} \cdot b_i \cdot a_i \cdot c_i^{s(mn+i)}$. The top right diagram shows the first step of the homotopy, and the bottom right diagram shows the second step.

CHAPTER 5
Virtually Special Tubular Groups

The goal of this chapter is the following theorem:

Theorem 5.0.3. *A tubular group G acts freely on a locally finite CAT(0) cube complex if and only if G is virtually special.*

In Section 5.1 we analyse $C(\widetilde{X}, \mathcal{W})$ in the finite dimensional case to establish a set of conditions that imply that $G \backslash C(\widetilde{X}, \mathcal{W})$ is virtually special. We decompose $C(\widetilde{X}, \mathcal{W})$ as a tree of spaces, with the same underlying tree as $\widetilde{X}$ and then, under the assumption that the walls are *primitive*, we show that $C(\widetilde{X}, \mathcal{W})$ maps G-equivariantly into $\mathbf{R}^d \times \widetilde{\Gamma}$, where $\mathbf{R}$ is the standard cubulation of $\mathbb{R}$, and $\widetilde{\Gamma}$ is the underlying graph of $\widetilde{X}$. We will show that when the immersed walls are fortified and non-dilated $C(\widetilde{X}, \mathcal{W})$ is locally finite. Combining these results allow us to give criterion for $G \backslash C(\widetilde{X}, \mathcal{W})$ to be virtually special.

In Section 5.2 we consider a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. We show that we can obtain from such an action immersed walls that preserve the important properties of $\widetilde{Y}$. More precisely, we prove the following:

Proposition 5.2.8. Let G be a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. Then there is a tubular space X and a finite set of immersed walls in X. Moreover, if $(\widetilde{X}, \mathcal{W})$ is the associated wallspace, then

1. G acts freely on $C(\widetilde{X}, \mathcal{W})$.
2. $C(\widetilde{X}, \mathcal{W})$ is finite dimensional if $\widetilde{Y}$ is finite dimensional.
3. $C(\widetilde{X}, \mathcal{W})$ is finite dimensional and locally finite if $\widetilde{Y}$ is locally finite.

This Proposition is sufficient to allow us to prove Theorem 5.2.9.

In Section 5.3 we further exploit the results obtained in Section 5.1 to prove the following converse to (3) Proposition 5.2.8.

Corllary 5.3.3. Let X be a tubular space. If the tubular group $G = \pi_1 X$ acts freely on a finite dimensional CAT(0) cube complex, then G acts on a locally finite, finite dimensional CAT(0) cube complex

5.1 Finite Dimensional Dual Cube Complexes

Let X be a tubular space and let $G = \pi_1 X$. Let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a set of non-dilated immersed walls $\Lambda_1, \ldots \Lambda_k$ constructed from an equitable set, and a vertical immersed wall in each edge space. **We emphasize that in this section all immersed walls are assumed to be non-dilated, even when it is not explicitly stated.** Let $\widetilde{Z} = C(\widetilde{X}, \mathcal{W})$ and let $Z = G \backslash \widetilde{Z}$. By Theorem 4.6.10, the immersed walls being non-dilated is equivalent to $\widetilde{Z}$ being finite dimensional. For each edge $\tilde{e}$ in $\widetilde{\Gamma}$ let $\widetilde{\Lambda}_{\tilde{e}}$ denote the vertical wall in $\widetilde{X}_{\tilde{e}}$.

Proposition 5.1.1. *There is a G-equivariant map $f : \widetilde{Z} \to \widetilde{\Gamma}$. Therefore $\widetilde{Z}$ decomposes as a tree of spaces with $\widetilde{Z}_{\tilde{v}} = f^{-1}(\tilde{v})$, and $\widetilde{Z}_{\tilde{e}} = f^{-1}(\tilde{e})$ is the carrier of the hyperplane corresponding to $\widetilde{\Lambda}_{\tilde{e}} \in \mathcal{W}_{\mathsf{v}}$.*

By $f^{-1}(e)$ we mean the union of all cubes c in $\widetilde{Z}$ such that $f(c) = e$.

Proof. As there is a vertical wall in each edge space, and since the vertical walls are all disjoint we can identify $C(\widetilde{X}, \mathcal{W}_v)$ with the Bass-Serre tree $\widetilde{\Gamma}$ of $\widetilde{X}$. We define a map $f : \widetilde{Z} \to \widetilde{\Gamma}$: let z be a 0-cube in Z, then define $f(z)$ by letting $f(z)[\widetilde{\Lambda}_e] = z[\widetilde{\Lambda}_e]$. If z_1, z_2 are adjacent 0-cubes, then $z_1[\widetilde{\Lambda}] \neq z_2[\widetilde{\Lambda}]$ for precisely one wall $\widetilde{\Lambda} \in \mathcal{W}$. If $\widetilde{\Lambda}$ is a horizontal wall then $f(z_1) = f(z_2)$ and the 1-cube joining them is also mapped to the same vertex. If $\widetilde{\Lambda} = \widetilde{\Lambda}_e \in \mathcal{W}_{\mathsf{v}}$ then $f(z_1)$ and $f(z_2)$ are adjacent in $\widetilde{\Gamma}$, so the 1-cube joining z_1 and z_2 maps to the edge joining $f(z_1)$ and $f(z_2)$. As f is defined on the 1-skeleton, the map extends uniquely to the entire cube complex $\widetilde{Z}$. Then $\widetilde{Z}_{\tilde{v}} = f^{-1}(\tilde{v})$ and $\widetilde{Z}_{\tilde{e}} = f^{-1}(\tilde{e})$ is the carrier of the hyperplane corresponding to $\widetilde{\Lambda}_e$. $\square$

Proposition 5.1.1 implies that Z decomposes as a graph of spaces with vertex spaces Z_v, edge spaces Z_e, and underlying graph $\Gamma = G\backslash\widetilde{\Gamma}$.

The immersed walls $\Lambda_1, \ldots, \Lambda_k$ are non-dilated, and therefore $\widetilde{Z}$ is finite dimensional, so by Proposition 4.6.13 there exists a stable partition $\mathcal{P}$ of the horizontal walls in $\mathcal{W}$. Let $\mathcal{P}_{\tilde{v}}$ be the subpartition containing walls intersecting $\widetilde{X}_{\tilde{v}}$. Let $\mathcal{P}_{\tilde{e}}$ be the subpartition of walls intersecting $\widetilde{X}_{\tilde{e}}$. By Lemma 4.6.14, both $\mathcal{P}_{\tilde{v}}$ and $\mathcal{P}_{\tilde{e}}$ are finite subpartitions. If $\tilde{e}$ is incident to the vertex $\tilde{v}$, then $\mathcal{P}_{\tilde{e}} \subseteq \mathcal{P}_{\tilde{v}}$. Let $\mathcal{P}_{\tilde{v}} = \{A_1, \ldots, A_{d_{\tilde{v}}}\}$. By criterion (3) of a stable partition $A_i = \{h_i^r \widetilde{\Lambda}_i\}_{r \in \mathbb{Z}}$ such that $h_i \in G_{\tilde{v}}$ stabilizes an axis in $\widetilde{X}_{\tilde{v}}$ perpendicular to $\widetilde{\Lambda}_i \cap \widetilde{X}_{\tilde{v}}$. The action of $G_{\tilde{v}}$ preserves both the partition $\mathcal{P}_{\tilde{v}}$ and the ordering of the walls in each A_i.

Let $\mathbf{R}$ denote the cubulation of $\mathbb{R}$ with a vertex for each integer and an edge joining consecutive integers. Therefore, each 0-cube in $\mathbf{R}^d$ is an element of $\mathbb{Z}^d$. We construct a free action of $G_{\tilde{v}}$ on $\mathbf{R}^{d_{\tilde{v}}}$. Let $g \in G_{\tilde{v}}$ and let $(\alpha_1, \ldots, \alpha_{d_{\tilde{v}}})$ be a 0-cube in $\mathbf{R}^{d_{\tilde{v}}}$. Define the map $g \cdot (\alpha_1, \ldots, \alpha_{d_{\tilde{v}}}) = (\beta_1, \ldots, \beta_{d_{\tilde{v}}})$ such that $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$. As g permutes the walls in $\mathcal{P}_{\tilde{v}}$, the map g is a bijection on the 0-cubes in $\mathbf{R}^{d_{\tilde{v}}}$. If $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$, then necessarily $g \cdot h_i^{\alpha_i+1} \widetilde{\Lambda}_i = h_j^{\beta_j \pm 1} \widetilde{\Lambda}_j$, so adjacent 0-cubes are mapped to adjacent 0-cubes and the map extends to an isomorphism of $\mathbf{R}^{d_{\tilde{v}}}$. If $g \cdot (\alpha_1, \ldots, \alpha_{d_{\tilde{v}}}) = (\alpha_1, \ldots, \alpha_{d_{\tilde{v}}})$ then g would stabilize all the walls in $\mathcal{P}_{\tilde{v}}$, which would imply that it fixed every 0-cube in $\widetilde{Z}_{\tilde{v}}$. Since $G_{\tilde{v}}$ acts freely on $\widetilde{Z}_{\tilde{v}}$ this would imply that $g = 1_G$, and hence $G_{\tilde{v}}$ acts freely on $\mathbf{R}^{d_{\tilde{v}}}$.

We also define an embedding $\phi_{\tilde{v}} : \widetilde{Z}_{\tilde{v}} \to \mathbf{R}^{d_{\tilde{v}}}$. If z is a 0-cube in $\widetilde{Z}_{\tilde{v}}$ then every wall $\widetilde{\Lambda}$ that is either vertical or not in contained in the subpartition $\mathcal{P}_{\tilde{v}}$ has $\widetilde{X}_{\tilde{v}} \subseteq z[\widetilde{\Lambda}]$. Therefore z is entirely determined by $z[\widetilde{\Lambda}]$ for $\widetilde{\Lambda}$ in $\mathcal{P}_{\tilde{v}}$. For

$1 \leq i \leq d\tilde{v}$ the set $\{h_i^r \widetilde{\Lambda}_i \cap \widetilde{X}_{\tilde{v}}\}_{r \in \mathbb{Z}}$ is an infinite collection of disjoint parallel lines in $\widetilde{X}_{\tilde{v}}$. As all the walls in A_i are disjoint in $\widetilde{X}$, for each 0-cube z in $\widetilde{Z}_{\tilde{v}}$ there exists a unique $\alpha_i \in \mathbb{Z}$ such that $h_i^{\alpha_i} \widetilde{\Lambda}_i$ and $h_i^{\alpha_i+1} \widetilde{\Lambda}_i$ face each other in z. Let $\phi_{\tilde{v}}(z) = (\alpha_1, \ldots, \alpha_{d_{\tilde{v}}})$. Note that the map is injective and sends adjacent 0-cubes to adjacent 0-cubes, so the map on the 0-cubes extends to an embedding of the entire cube complex.

Lemma 5.1.2. *The embedding $\phi_{\tilde{v}} : \widetilde{Z}_{\tilde{v}} \to \mathbf{R}^{d_{\tilde{v}}}$ is $G_{\tilde{v}}$-equivariant.*

Proof. Let $g \in G_{\tilde{v}}$. If $\phi_{\tilde{v}}(z) = (\alpha_1, \ldots, \alpha_{d_{\tilde{v}}})$ and $g \cdot \phi_{\tilde{v}}(z) = (\beta_1, \ldots, \beta_{d_{\tilde{v}}})$, then $(gz)[h_i^{\alpha_i} \widetilde{\Lambda}_i] = z[g h_i^{\alpha_i} \widetilde{\Lambda}_i] = z[h_j^{\beta_j} \widetilde{\Lambda}_j]$, which implies that $\phi_{\tilde{v}}(gz) = (\beta_1, \ldots, \beta_{d_{\tilde{v}}}) = g\phi_{\tilde{v}}(z)$. □

Let $\tilde{e}$ be an edge adjacent to $\tilde{v}$. Then either $+\tilde{e} = \tilde{v}$ or $-\tilde{e} = \tilde{v}$. We define a free action of $G_{\tilde{e}}$ on $\mathbf{R}^{d_{\tilde{e}}} \times [-1, 1]$. After reindexing, let $\mathcal{P}_{\tilde{e}} = \{A_1, \ldots, A_{d_{\tilde{e}}}\} \subseteq \mathcal{P}_{\tilde{v}}$ where $d_{\tilde{e}} \leq d_{\tilde{v}}$. Let $(\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}, \pm 1)$ be a 0-cube in $\mathbf{R}^d \times [-1, 1]$ and let $g \in G_{\tilde{e}}$. Then $g \cdot (\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}, \pm 1) = (\beta_1, \ldots, \beta_{d_{\tilde{e}}}, \pm 1)$ such that $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$. As in the case of vertex spaces, this map extends to an isomorphism of $\mathbf{R}^{d_{\tilde{e}}} \times [-1, 1]$.

As with the vertex spaces, there is a $G_{\tilde{e}}$-equivariant embedding $\phi_{\tilde{e}} : \widetilde{Z}_{\tilde{e}} \to \mathbf{R}^{d_{\tilde{e}}} \times [-1, 1]$. Let z be a 0-cube in $\widetilde{Z}_{\tilde{e}}$. Then for each $1 \leq i \leq d_{\tilde{e}}$ there exists a unique α_i such that $h_i^{\alpha_i} \widetilde{\Lambda}_i$ faces $h_i^{\alpha_i+1} \widetilde{\Lambda}_i$ in z, and $\widetilde{X}_{\pm \tilde{e}} \subseteq z[\widetilde{\Lambda}_{\tilde{e}}]$. Define $\phi_{\tilde{e}}(z) = (\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}, \pm 1)$.

Let $\tilde{v} = \pm \tilde{e}$. The free action of $G_{\tilde{v}}$ on $\mathbf{R}^{d_{\tilde{v}}}$ restricts to a free action of $G_{\tilde{e}}$. We claim that we can embed $\mathbf{R}^{d_{\tilde{e}}} \times \{\pm 1\}$ into $\mathbf{R}^{d_{\tilde{v}}}$ in a $G_{\tilde{e}}$-equivariant way. Let $H_{\tilde{e}} \subseteq \widetilde{Z}$ be the hyperplane corresponding to $\widetilde{\Lambda}_e$. As $\widetilde{Z}_{\tilde{e}}$ is the carrier of $H_{\tilde{e}}$, we can identify $\widetilde{Z}_{\tilde{e}}$ with $H_{\tilde{e}} \times [-1, 1]$. Note that $H_{\tilde{e}} \times \{\pm 1\}$ embeds as a subspace in $\widetilde{Z}_{\tilde{v}}$, and $\phi_{\tilde{e}}$ restricts to an embedding $\phi_{\tilde{e}}^{\pm} : H_{\tilde{e}} \times \{\pm 1\} \to \mathbf{R}^{d_{\tilde{e}}}$, where $\tilde{v} = \pm \tilde{e}$.

We construct an embedding $\Psi_{\tilde{e}}^{\pm} : \mathbf{R}^{d_{\tilde{e}}} \to \mathbf{R}^{d_{\tilde{v}}}$. Recall that $\mathcal{P}_{\tilde{e}} = \{A_1, \ldots, A_{d_{\tilde{e}}}\} \subseteq \mathcal{P}_{\tilde{v}} = \{A_1, \ldots, A_{d_{\tilde{v}}}\}$. For $d_{\tilde{e}} < j \leq d_{\tilde{v}}$ if $h_j^r \widetilde{\Lambda}_j \in A_j$, then $\widetilde{X}_{\tilde{e}} \subseteq z[h_j^r \widetilde{\Lambda}_j]$ for all 0-cubes z in $\widetilde{Z}_{\tilde{e}}$. Therefore, there is a unique $\alpha_j^{\tilde{e}} \in \mathbb{Z}$ such that $h_j^{\alpha_j^{\tilde{e}}} \widetilde{\Lambda}_j$ faces $h_j^{\alpha_j^{\tilde{e}}+1} \widetilde{\Lambda}_j$ for every 0-cube z in $\widetilde{Z}_{\tilde{e}}$ and $d_{\tilde{e}} < j \leq d_{\tilde{v}}$. Thus we define

$$\Psi_{\tilde{e}}^{\pm}(\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}) = (\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}, \alpha_{d_{\tilde{e}}+1}^{\tilde{e}}, \ldots, \alpha_{d_{\tilde{v}}}^{\tilde{e}}).$$

The $G_{\tilde{e}}$-equivariance of $\Psi_{\tilde{e}}^{\pm}$ will require a further assumption:

Lemma 5.1.3. *The following commutative square is $G_{\tilde{e}}$-equivariant provided the immersed walls are primitive.*

$$\begin{tikzcd}
H_{\tilde{e}} \times \{\pm 1\} \arrow[r, hook, "\phi_{\tilde{e}}^{\pm}"] \arrow[d, hook] & \mathbf{R}^{d_{\tilde{e}}} \arrow[d, "\Psi_{\tilde{e}}^{\pm}"] \\
\widetilde{Z}_{\tilde{v}} \arrow[r, hook, "\phi_{\tilde{v}}"] & \mathbf{R}^{d_{\tilde{v}}}
\end{tikzcd}$$

Moreover, $\Psi_{\tilde{e}}^{\pm}$ is a $G_{\tilde{e}}$-equivariant inclusion that is equivalent to extending the $G_{\tilde{e}}$-action on $\mathbf{R}^{d_{\tilde{e}}}$ by a trivial action on $\mathbf{R}^{d_{\tilde{v}}-d_{\tilde{e}}}$.

Proof. Let z be a 0-cube in $H_{\tilde{e}} \times \{\pm 1\}$. Then by construction

$$\Psi_{\tilde{e}}^{\pm} \circ \phi_{\tilde{e}}^{\pm}(z) = (\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}, \alpha_{d_{\tilde{e}}+1}^{\tilde{e}}, \ldots, \alpha_{d_{\tilde{v}}}^{\tilde{e}}) = \phi_{\tilde{v}}(z).$$

To verify that $\Psi_{\tilde{e}}$ is $G_{\tilde{e}}$-equivariant, let $g \in G_{\tilde{e}}$. For $1 \leq i \leq d_{\tilde{e}}$ there exists $1 \leq j \leq d_{\tilde{e}}$ and β_j be such that $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$. For $d_{\tilde{e}} < i \leq d_{\tilde{v}}$ the intersection $\widetilde{\Lambda}_i \cap \widetilde{X}_{\tilde{v}}$ is a geodesic line parallel to $\widetilde{X}_{\tilde{e}} \cap \widetilde{X}_{\tilde{v}}$. Thus, $G_{\tilde{e}}$ stabilizes $\widetilde{\Lambda}_i \cap \widetilde{X}_{\tilde{v}}$. As the immersed walls are primitive we can deduce that $G_{\tilde{e}}$ stabilizes

$\widetilde{\Lambda}_i$. For $d_{\tilde{e}} < i \leq d_{\tilde{v}}$ we deduce that $g \cdot h_i^{\alpha}\widetilde{\Lambda}_i = h_i^{\alpha}\widetilde{\Lambda}_i$ for $\alpha \in \mathbb{Z}$, and conclude:

$$\begin{aligned} g \cdot \Psi_{\tilde{e}}(\alpha_1, \ldots, \alpha_{\tilde{e}}) &= g \cdot (\alpha_1, \ldots, \alpha_{d_{\tilde{e}}}, \alpha^{\tilde{e}}_{d_{\tilde{e}}+1}, \ldots, \alpha^{\tilde{e}}_{d_{\tilde{v}}}) \\ &= (\beta_1, \ldots, \beta_{d_{\tilde{e}}}, \alpha^{\tilde{e}}_{d_{\tilde{e}}+1}, \ldots, \alpha^{\tilde{e}}_{d_{\tilde{v}}}) \\ &= \Psi_{\tilde{e}}(\beta_1, \ldots, \beta_{\tilde{e}}) \\ &= \Psi_{\tilde{e}}(g \cdot (\alpha_1, \ldots, \alpha_{\tilde{e}})). \end{aligned}$$

Observe that $G_{\tilde{e}}$ acts trivially on the last $d_{\tilde{v}} - d_{\tilde{e}}$ coordinates. □

Let $d = \max\{|\mathcal{P}_{\tilde{v}}| \mid \tilde{v} \in V\widetilde{\Gamma}\}$, which is finite, since there are only finitely many vertex orbits.

Proposition 5.1.4. *If the immersed walls $\Lambda_1, \ldots, \Lambda_k$ are primitive, then G acts freely on $\mathbf{R}^d \times \widetilde{\Gamma}$ such that the action on the $\widetilde{\Gamma}$ factor is the action of G on the Bass-Serre tree. Moreover, there is a G-equivariant embedding $\phi : \widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$.*

Proof. The $G_{\tilde{v}}$ and $G_{\tilde{e}}$-actions on $\mathbf{R}^{d_{\tilde{v}}}$ and $\mathbf{R}^{d_{\tilde{e}}}$ can be equivariantly extended to actions on $\mathbf{R}^d = \mathbf{R}^{d_{\tilde{v}}} \times \mathbf{R}^{d-d_{\tilde{v}}} = \mathbf{R}^{d_{\tilde{e}}} \times \mathbf{R}^{d-d_{\tilde{e}}}$ such that $G_{\tilde{v}}$ and $G_{\tilde{e}}$ act trivially on the additional factors. Therefore, the $G_{\tilde{e}}$-commutative square in Lemma 5.1.3 can be extended:

$$\begin{tikzcd} H_{\tilde{e}} \times \{\pm 1\} \arrow[r, hook, "\phi^{\pm}_{\tilde{e}}"] \arrow[d, hook] & \mathbf{R}^{d_{\tilde{e}}} \arrow[r] \arrow[d, "\Psi^{\pm}_{\tilde{e}}"] & \mathbf{R}^d \arrow[d] \\ \widetilde{Z}_{\tilde{v}} \arrow[r, hook, "\phi_{\tilde{v}}"] & \mathbf{R}^{d_{\tilde{v}}} \arrow[r] & \mathbf{R}^d \end{tikzcd}$$

The right square commutes and is $G_{\tilde{e}}$-equivariant since by Lemma 5.1.3, the $G_{\tilde{e}}$-equivariant inclusion $\Psi^{\pm}_{\tilde{e}}$ is equivalent to extending $\mathbf{R}^{d_{\tilde{e}}}$ by the trivial action on $\mathbf{R}^{d_{\tilde{v}}-d_{\tilde{e}}}$.

The decomposition of $\widetilde{Z}$ into a tree of spaces with underlying tree $\widetilde{\Gamma}$ gives a decomposition of Z as a graph of spaces with underlying graph Γ. By taking the quotient by $G_{\tilde{e}}$ of the top row of the above diagram, the bottom row by

$G_{\tilde{v}}$, and forgetting the middle column we obtain the following diagram:

$$\begin{tikzcd} G_{\tilde{e}}\backslash H_{\tilde{e}} \times \{\pm 1\} \arrow[r] \arrow[d] & G_{\tilde{e}}\backslash \mathbf{R}^d \arrow[d] \\ G_{\tilde{v}}\backslash Z_{\tilde{v}} \arrow[r] & G_{\tilde{v}}\backslash \mathbf{R}^d \end{tikzcd}$$

The vertical maps on the left give the attaching maps of the edge spaces in Z. The vertical maps on the right can be used as attaching maps of edge spaces for a new graph of spaces $\mathcal{Z}$ with underlying graph Γ and $\pi_1\mathcal{Z} = G$. The universal cover of $\mathcal{Z}$ will be $\widetilde{\mathcal{Z}} = \mathbf{R}^d \times \widetilde{\Gamma}$, and there is a π_1-isomorphic map $Z \to \mathcal{Z}$, determined by the horizontal maps on the vertex and edge spaces in the above diagram, that lifts to a G-equivariant inclusion of the universal covers $\widetilde{Z} \to \widetilde{\mathcal{Z}} = \mathbf{R}^d \times \widetilde{\Gamma}$.

Therefore, we obtain a G-action on $\mathbf{R}^d \times \widetilde{\Gamma}$ and a G-equivariant embedding of the tree of spaces $\widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$. $\square$

Proposition 5.1.5. *$\widetilde{Z}$ is locally finite if and only if $\Lambda_1, \ldots, \Lambda_k$ are fortified.*

Proof. If $\Lambda_1, \ldots, \Lambda_k$ is not fortified, then there exists a vertex space $\widetilde{X}_{\tilde{v}}$ and an adjacent edge space $\widetilde{X}_{\tilde{e}}$ such that every horizontal wall $\widetilde{\Lambda}$ in $\mathcal{P}_{\tilde{v}}$ intersects $\widetilde{X}_{\tilde{v}}$ as a line $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ that intersects $\widetilde{X}_{\tilde{v}} \cap \widetilde{X}_{\tilde{e}}$. Therefore, every horizontal wall intersecting $\widetilde{X}_{\tilde{v}}$ intersects $\widetilde{X}_{\tilde{e}}$, so $\mathcal{P}_{\widetilde{e}} = \mathcal{P}_{\tilde{v}}$. Let $\tilde{e}_1, \ldots, \tilde{e}_i, \ldots$ be an enumeration of the $G_{\tilde{v}}$-orbit of $\tilde{e}$. Then $\mathcal{P}_{\tilde{e}_i} = \mathcal{P}_{\tilde{v}}$ and $\widetilde{\Lambda}_{\tilde{e}_i}$ intersects all the horizontal walls in $\mathcal{P}_{\tilde{v}}$.

Let z be a 0-cube in $\widetilde{Z}_{\tilde{v}}$. There is a 0-cube z_i such that $z_i[\widetilde{\Lambda}] = z[\widetilde{\Lambda}]$ for $\widetilde{\Lambda} \neq \widetilde{\Lambda}_{\tilde{e}_i}$, and $z_i[\widetilde{\Lambda}_{\tilde{e}_i}] \neq z[\widetilde{\Lambda}_{\tilde{e}_i}]$. To verify z_i is a 0-cube note that every wall in $\mathcal{P}_{\tilde{v}}$ intersects $\widetilde{\Lambda}_{\tilde{e}}$, and every other wall $\widetilde{\Lambda}$ that is not $\widetilde{\Lambda}_e$ has $\widetilde{X}_{\tilde{e}} \subseteq z_i[\widetilde{\Lambda}]$. Therefore $z_i[\widetilde{\Lambda}_{\tilde{e}}] \cap z_i[\widetilde{\Lambda}] \neq \emptyset$ for all $\widetilde{\Lambda} \in \mathcal{W} - \{\widetilde{\Lambda}_{\tilde{e}}\}$. For any walls $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2 \in \mathcal{W} - \{\widetilde{\Lambda}_{\tilde{e}}\}$ the intersection $z_i[\widetilde{\Lambda}_1] \cap z_i[\widetilde{\Lambda}_2] = z[\widetilde{\Lambda}_1] \cap z[\widetilde{\Lambda}_2] \neq \emptyset$. Finally, if $x \in \widetilde{X}$, then $x \in z_i[\widetilde{\Lambda}]$ for all but finitely many $\widetilde{\Lambda} \in \mathcal{W}$, because it is true for z, which differs from

z_i on precisely one wall. Each z_i is adjacent to z since they differ on precisely one wall, so $z_1, \ldots, z_i, \ldots$ is an infinite collection of distinct 0-cubes adjacent to z, and $\widetilde{Z}$ is not locally finite.

To show the converse, we first observe that the embedding $\phi_{\tilde{v}} : \widetilde{Z}_{\tilde{v}} \to \mathbf{R}^{d_{\tilde{v}}}$ proves that $\widetilde{Z}_{\tilde{v}}$ is always locally finite, irrespective of whether the immersed walls are fortified. Let z be a 0-cube in $\widetilde{Z}_{\tilde{v}}$, and let $\tilde{u}$ be adjacent to $\tilde{v}$ in $\widetilde{\Gamma}$ via an edge $\tilde{e}$. Then z can be adjacent to at most one 0-cube $z_{\tilde{e}}$ in $\widetilde{Z}_{\tilde{u}}$ such that $z[\widetilde{\Lambda}] = z_{\tilde{e}}[\widetilde{\Lambda}]$ for all $\widetilde{\Lambda} \in \mathcal{W}$ except $\widetilde{\Lambda}_e$. This $z_{\tilde{e}}$ may not always define a 0-cube however. Let $\tilde{e}$ be an edge adjacent to $\tilde{v}$. As the immersed walls are fortified there exists $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}} \in \mathcal{P}_{\tilde{v}}$ such that $\{h^r\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}\}_{r\in\mathbb{Z}}$ is an infinite set of lines parallel to $\widetilde{X}_{\tilde{e}} \cap \widetilde{X}_{\tilde{v}}$. As $\{h^r\widetilde{\Lambda}\}$ is a set of disjoint walls, there exists r such that $h^r\widetilde{\Lambda}$ and $h^{r+1}\widetilde{\Lambda}$ are facing in z. There are only finitely many edges $g_1\tilde{e}, \ldots, g_m\tilde{e} \in G_{\tilde{v}}\tilde{e}$ such that $\widetilde{X}_{g_i\tilde{e}} \subseteq z[h^r\widetilde{\Lambda}] \cap z[h^{r+1}\widetilde{\Lambda}]$. If $g\tilde{e}$ is an edge such that $\widetilde{X}_{g_i\tilde{e}}$ is not contained in $z[h^r\widetilde{\Lambda}] \cap z[h^{r+1}\widetilde{\Lambda}]$, then either $z_{g\tilde{e}}[h^r\widetilde{\Lambda}] \cap z_{g\tilde{e}}[\widetilde{\Lambda}_e] = \emptyset$ or $z_{g\tilde{e}}[h^{r+1}\widetilde{\Lambda}] \cap z_{g\tilde{e}}[\widetilde{\Lambda}_e] = \emptyset$ so $z_{g\tilde{e}}$ is not a 0-cube. As there are only finitely many $G_{\tilde{v}}$-orbits of edges incident to $\tilde{v}$ we conclude that $\tilde{z}_{\tilde{e}}$ is a 0-cube for finitely many edges $\tilde{e}$ incident to $\tilde{v}$. $\square$

Proposition 5.1.6. *If $\Lambda_1, \ldots, \Lambda_k$ are primitive, fortified, non-dilated immersed walls, then G is virtually special.*

Proof. By Proposition 5.1.4, there is a free action of G on $\mathbf{R}^d \times \widetilde{\Gamma}$, so G is a subgroup of $\text{Isom}(\mathbf{R}^d \times \widetilde{\Gamma}) \cong \Big(\mathbb{Z}^d \rtimes \text{Aut}([-1,1])^d\Big) \times \text{Aut}(\widetilde{\Gamma})$. Therefore, there is a projection $\rho : G \to \mathbb{Z}^d \rtimes \text{Aut}([-1,1]^d)$. Each vertex group $G_{\tilde{v}}$ embeds in $\mathbb{Z}^d$, and the mapping is invariant under conjugation. As there are only finitely orbits of vertices in $\widetilde{\Gamma}$, there exists a finite index subgroup $(D\mathbb{Z})^d \leqslant \mathbb{Z}^d$ such that if $\tilde{e}$ is incident to $\tilde{v}$ then $G_{\tilde{e}} \cap (D\mathbb{Z})^d$ is generated by a primitive element in $G_{\tilde{v}} \cap (D\mathbb{Z})^d$. Let $G' = \rho^{-1}\Big((D\mathbb{Z})^d\Big)$. Then $G' \leqslant G$ is a finite index subgroup that embeds in $(D\mathbb{Z})^d \times \text{Aut}(\widetilde{\Gamma})$ such that each edge group is generated by an element that is primitive in the adjacent vertex groups. By Proposition 5.1.4 there is a G-equivariant embedding $\phi : \widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$. As G' does not permute the factors of $\mathbf{R}^d$ we can deduce that the hyperplanes in $\left(\mathbf{R}^d \times \widetilde{\Gamma}\right)/G$ do not self-intersect, so neither do the hyperplanes in $\widetilde{Z}/G'$. Indeed, they are also 2-sided and cannot inter-osculate.

Let G'' be a finite index subgroup such that the underlying graph Γ'' has girth at least 2. Let $\widetilde{Z}_{\tilde{e}}^{+}$ and $\widetilde{Z}_{\tilde{e}}^{-}$ denote the subcomplexes $H_{\tilde{e}} \times \{-1\}$ and $H_{\tilde{e}} \times \{+1\}$ respectively. Let $\tilde{e}$ be an edge such that $+\tilde{e} = \tilde{v}$. As $\Lambda_1, \ldots, \Lambda_k$ are fortified, we conclude that $d_{\tilde{v}} > d_{\tilde{e}}$ and $\widetilde{Z}_{\tilde{e}}^{+}$ is a proper subcomplex of $\widetilde{Z}_{\tilde{v}}$. As $G''_{\tilde{e}}$ is primitive, if $g \in G''_{\tilde{v}} - G''_{\tilde{e}}$ then $gh^r_{d_{\tilde{v}}}\widetilde{\Lambda}_{d_{\tilde{v}}} \neq h^r_{d_{\tilde{v}}}\widetilde{\Lambda}_{d_{\tilde{v}}}$ as g acts by translation on $\widetilde{X}_{\tilde{v}}$ in a direction non-parallel to $\widetilde{\Lambda}_{d_{\tilde{v}}} \cap \widetilde{X}_{\tilde{v}}$. Thus, $\widetilde{Z}_{\tilde{e}}^{\pm}$ is not stabilized by g, so we can deduce that $\mathrm{Stab}_{G''_{\tilde{v}}}(\widetilde{Z}_{\tilde{e}}^{+}) = G''_{\tilde{e}}$. Therefore, $G''_{\tilde{e}} \backslash \widetilde{Z}_{\tilde{e}}^{+}$ embeds in $G''_{\tilde{v}} \backslash \widetilde{Z}_{\tilde{v}}$.

Let $Z'' = G'' \backslash \widetilde{Z}$. Let z be a 0-cube in Z''_v. Let H_e be the vertical hyperplane contained in Z''_e, and dual to an edge incident to v. As the attaching maps of Z''_e are embeddings, and Γ'' has girth at least 2, we deduce that z can only be incident to one end of a single 1-cube intersected by H_e. Therefore H_e does not self-osculate.

Let σ be a 1-cube in $\mathbf{R}^d \times \widetilde{\Gamma}$ that projects to a 1-cube in $\mathbf{R}^d$. The G''-orbit of σ is a set of 1-cubes, that all project to the same factor of $\mathbf{R}^d$, since G'' does not permute the factors of $\mathbf{R}^d$. As G'' does not invert hyperplanes, after subdividing $\mathbf{R}^d$ we can assume that the G''-orbit of σ is a disjoint set of 1-cubes. Therefore, after the corresponding subdivision, we conclude that the horizontal hyperplanes in Z'' don't self-osculate. $\square$

We note that the requirement in Proposition 5.1.6 that the immersed walls are fortified is necessary, as the following example demonstrates.

Example 5.1.7. Let $G = \langle a, b, t \mid [a, b] = 1, tat^{-1} = a\rangle$. We can decompose G as the cyclic HNN extension of the vertex group $G_v = \langle a, b\rangle$ with stable letter t. Thus, G is a tubular group. Note that $G \cong \mathbb{F}_2 \times \mathbb{Z}$ which is a right angled Artin group, and therefore special. Let X be the corresponding tubular space with a single vertex space X_v and edge space X_e. There is an equitable set $\{\alpha_1, \alpha_2\}$ where α_1 is a geodesic curve in X_v representing $ab \in G_v$, and α_2 is a geodesic curve in X_v representing $ab^{-1} \in G_v$. Note that each attaching map φ_e^+ and φ_e^- intersects each curve in the equitable set precisely once. Therefore, we obtain a pair of embedded immersed horizontal walls Λ_1 and Λ_2, by connecting respective intersection points with φ_e^+ and φ_e^- by an arc. A vertical wall Λ_e is also embedded in X_e.

In the wallspace $(\widetilde{X}, \mathcal{W})$ we can decompose $\mathcal{W}$ into three sets of disjoint walls: the walls $\mathcal{W}_1$ that cover Λ_1, the walls $\mathcal{W}_2$ that cover Λ_2, and the walls $\mathcal{W}_e$ that cover Λ_e. These walls are disjoint since the immersed walls are embedded. Furthermore, the walls in different sets pairwise intersect. Therefore we conclude that $C(\widetilde{X}, \mathcal{W}) = \mathbf{R}^2 \times \widetilde{\Gamma}$. As this is not locally finite, $G \backslash C(\widetilde{X}, \mathcal{W})$ cannot be virtually special, although G itself is special.

5.2 Revisiting Equitable Sets

Although Wise proved in [40] that acting freely on a CAT(0) cube complex $\widetilde{Y}$ implied the existence of an equitable set, and thus a system of immersed walls as in Section 4.2, no relationship was established between $\widetilde{Y}$ and the resulting dual $C(\widetilde{X}, \mathcal{W})$. Proposition 5.2.8 gives the relationship required to reduce Theorem 5.2.9 to considering cubulations obtained from equitable sets.

This section will apply the main result from Chapter 3.

Lemma 5.2.1. *Let G be a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. Let G_v be a vertex group in G, then there exists a G_v-invariant subspace $\widetilde{X}_{\tilde{v}} \subseteq \widetilde{Y}$ homeomorphic to $\mathbb{R}^2$. Moreover, $\widetilde{X}_{\tilde{v}}$ has a metric such that* *the intersection of a hyperplane in $\widetilde{Y}$ with $\widetilde{X}_{\tilde{v}}$ is either empty, or a geodesic line.*

Proof. By Theorem 3.4.3, there exists a G_v-equivariant subcomplex $\widetilde{Y}_v \subseteq \widetilde{Y}$ that isometrically embeds in the combinatorial metric, and such that $\widetilde{Y}_v \cong \prod_{i=1}^{m} C_i$ where each C_i is a cubical quasiline. By the flat torus theorem [6], G_v stabilizes a flat $\widetilde{X}_v \subseteq \widetilde{Y}_v$, that is a convex subset in the CAT(0) metric of $\widetilde{Y}_v$. As the stabilizers of hyperplanes in $\widetilde{Y}_v$ are codimension-1 subgroups of G_v, the intersection of a hyperplane in $\widetilde{Y}$ with $\widetilde{X}_v$ is either empty or a geodesic line in the CAT(0) metric inherited from $\widetilde{Y}_v$. □

Definition 5.2.2. Let X be a tubular space and let Y be a non-positively curved cube complex. A map $f : X \to Y$ is an *amicable immersion* if:

1. $f_* : \pi_1 X \to \pi_1 Y$ is an isomorphism.
2. The G-equivariant map $\widetilde{f} : \widetilde{X} \to \widetilde{Y}$ embeds each vertex space $\widetilde{X}_{\tilde{v}}$ in $\widetilde{Y}$.
3. Each $\widetilde{X}_{\tilde{v}}$ has a Euclidean metric such that if $H \subseteq \widetilde{Y}$ is a hyperplane, then the intersection $H \cap \widetilde{X}_{\tilde{v}}$ is either the empty set, or a single geodesic line in $\widetilde{X}_{\tilde{v}}$.
4. Each edge space $\widetilde{X}_{\tilde{e}}$ is embedded transverse to the hyperplanes.
5. Each $\widetilde{X}_{\tilde{e}}$ is contained in $\mathsf{hull}\Big(\bigcup_{\tilde{v} \in V\widetilde{\Gamma}} \widetilde{X}_{\tilde{v}}\Big)$.

Note that the Euclidean metric on each $\widetilde{X}_{\tilde{v}}$ is not the subspace metric induced from $\widetilde{Y}$.

Lemma 5.2.3. *Let X be a tubular space and let Y be a non-positively curved cube complex. Let $F : \pi_1 X \to \pi_1 Y$ be an isomorphism. Then there is an amicable immersion $f : X \to Y$ such that $f_* = F$.*

Proof. Use F to identify $G = \pi_1 X$ with $\pi_1 Y$. The claim is proven by constructing a G-equivariant map between the tree of spaces $\widetilde{X} \to \widetilde{Y}$. By Lemma 5.2.1 for each $\tilde{v} \in V\widetilde{\Gamma}$, we can $G_{\tilde{v}}$-equivariantly embed a Euclidean flat $\widetilde{X}_{\tilde{v}}$ in $\widetilde{Y}$ such that if $H \subseteq \widetilde{Y}$ is a hyperplane, then the intersection $H \cap \widetilde{X}_{\tilde{v}}$ is either the empty set, or a single geodesic line in $\widetilde{X}_{\tilde{v}}$. Moreover, we can ensure that $\bigcup_{\tilde{v} \in V\widetilde{\Gamma}} \widetilde{X}_{\tilde{v}}$ is G-equivariant. The edges spaces $\widetilde{X}_{\tilde{e}}$ can then be inserted transverse to the hyperplanes in $\widetilde{Y}$ so that $\widetilde{X}_{\tilde{e}}$ is contained inside $\mathsf{hull}\Big(\bigcup_{\tilde{v} \in V\widetilde{\Gamma}} \widetilde{X}_{\tilde{v}}\Big)$. □

Lemma 5.2.4. *Let $X \to Y$ be an amicable immersion, where Y is finite dimensional. If $\tilde{v}$ is a vertex in $\widetilde{\Gamma}$, then $\mathsf{hull}(\widetilde{X}_{\tilde{v}})$ embeds as a subcomplex of $\mathbf{R}^d$ for some d.*

Proof. Let $G = \pi_1 X$. Let y be a 0-cube in $\mathsf{hull}(\widetilde{X}_{\tilde{v}})$. If H is a hyperplane in $\widetilde{Y}$, let $y[H]$ denote the halfspace of H containing y. Each 0-cube is determined by the halfspace containing it for each hyperplane. If H is a hyperplane that doesn't intersect $\widetilde{X}_{\tilde{v}}$, then $y[H]$ is the halfspace containing $\widetilde{X}_{\tilde{v}}$, and therefore $y[H]$ is fixed for all 0-cubes y in $\mathsf{hull}(\widetilde{X}_{\tilde{v}})$.

Let $\mathcal{H}_{\tilde{v}}$ denote the hyperplanes intersecting $\widetilde{X}_{\tilde{v}}$. Let $H \in \mathcal{H}_{\tilde{v}}$. The intersection $\widetilde{X}_{\tilde{v}} \cap H$ is a geodesic line in $\widetilde{X}_{\tilde{v}}$. Let $g \in G_{\tilde{v}}$ be an isometry that stabilizes an axis in $\widetilde{X}_{\tilde{v}}$ that is not parallel to $\widetilde{X}_{\tilde{v}} \cap H$. Then $\{g^r H\}_{r \in \mathbb{Z}}$ is an infinite family of hyperplanes such that $\{g^r H \cap \widetilde{X}_{\tilde{v}}\}_{r \in \mathbb{Z}}$ is a set of disjoint parallel lines in $\widetilde{X}_{\tilde{v}}$. As $\widetilde{Y}$ is finite dimensional, there exists an N such that H and $g^N H$ do not intersect. Otherwise $\{g^r H\}_{r \in \mathbb{Z}}$ would be an infinite set of pairwise intersecting hyperplanes, which would imply that there are cubes of arbitrary dimension in $\widetilde{Y}$.

Therefore, as there are only finitely many $G_{\tilde{v}}$-orbits of hyperplanes intersecting $\widetilde{X}_{\tilde{v}}$, there exists a finite set of hyperplanes $H_1, \ldots, H_d \in \mathcal{H}_{\tilde{v}}$ and $g_1, \ldots, g_d \in G$ such that $\mathcal{H}_{\tilde{v}} = \{g_1^r H_1, \ldots, g_d^r H_d\}_{r \in \mathbb{Z}}$, and each $\{g_i^r H_i\}_{r \in \mathbb{Z}}$ is a disjoint set of hyperplanes in $\widetilde{Y}$. Therefore $\{g_i^r H_i \cap \tilde{e}\}_{r \in \mathbb{Z}}$ is a set of disjoint geodesic lines in $\widetilde{X}_{\tilde{v}}$. Thus, given a 0-cube y, there exists a unique $y_i \in \mathbb{Z}$ such that $y[g_i^{y_i} H_i]$ and $y[g_i^{y_i+1} H_i]$ properly intersect each other. Therefore, construct $\phi : \mathsf{hull}(\widetilde{X}_{\tilde{v}}) \to \mathbf{R}^d$ by letting $\phi(y) = (y_1, \ldots, y_d)$ for each 0-cube

y. The map ϕ extends to the 1-skeleton of $\mathsf{hull}(\widetilde{X}_{\tilde{v}})$ since adjacent 0-cubes lie on the opposite sides of precisely one hyperplane. Therefore ϕ extends to the higher dimensional cubes, and thus $\mathsf{hull}(\widetilde{X}_{\tilde{v}})$. $\square$

Let $\widetilde{X} \to \widetilde{Y}$ be the lift to the universal cover of an amicable immersion $X \to Y$. Let $\widetilde{X}_{\tilde{e}}$ be an edge space adjacent to a vertex space $\widetilde{X}_{\tilde{v}}$. A hyperplane H in $\widetilde{Y}$ intersects $\widetilde{X}_{\tilde{v}}$ *parallel* to $\widetilde{X}_{\tilde{e}}$ if $H \cap \widetilde{X}_{\tilde{v}}$ is a geodesic line parallel to $\widetilde{X}_{\tilde{e}} \cap \widetilde{X}_{\tilde{v}}$. Otherwise, if $H \cap \widetilde{X}_{\tilde{v}}$ is a geodesic line that is not parallel to $\widetilde{X}_{\tilde{e}} \cap \widetilde{X}_{\tilde{v}}$, then we say H intersects $\widetilde{X}_{\tilde{v}}$ *non-parallel* to $\widetilde{X}_{\tilde{e}}$.

Lemma 5.2.5. *Let $X \to Y$ be an amicable immersion. Let $\tilde{e}$ be an edge in $\widetilde{\Gamma}$. Suppose that $H \subseteq \widetilde{Y}$ is a hyperplane intersecting $\widetilde{X}_{-\tilde{e}}$ non-parallel to $\widetilde{X}_{\tilde{e}}$, then H intersects $\widetilde{X}_{+\tilde{e}}$ non-parallel to $\widetilde{X}_{\tilde{e}}$. Moreover, there is an arc in $H \cap \widetilde{X}_{\tilde{e}}$ joining $H \cap \widetilde{X}_{-\tilde{e}}$ to $H \cap \widetilde{X}_{+\tilde{e}}$.*

Proof. Let $G = \pi_1 X$. The geodesic lines $H \cap \widetilde{X}_{-\tilde{e}}$ and $\widetilde{X}_{\tilde{e}} \cap \widetilde{X}_{-\tilde{e}}$ are non parallel in $\widetilde{X}_{-\tilde{e}}$, and therefore intersect in a single point $p \in \widetilde{X}_{\tilde{e}} \cap \widetilde{X}_{-\tilde{e}}$. As H is two sided in $\widetilde{Y}$ and X the vertex and edge spaces are transverse to H, the intersection of H with X is also locally two sided in $\widetilde{X}$. Therefore, p is contained inside a curve in $H \cap \widetilde{X}_{\tilde{e}}$. As $\widetilde{X}_{\tilde{e}}$ is $G_{\tilde{e}}$-invariant and only finitely many hyperplanes separate any two points in $\widetilde{X}$, we can deduce that p is an endpoint of a compact curve in $H \cap \widetilde{X}_{\tilde{e}}$ with its other endpoint contained in $\widetilde{X}_{+\tilde{e}} \cap \widetilde{X}_{\tilde{e}}$. Thus, H must also intersect $\widetilde{X}_{+\tilde{e}}$ non-parallel to $\widetilde{X}_{\tilde{e}}$. $\square$

Lemma 5.2.6. *Let $X \to Y$ be an amicable immersion, where Y is a finite dimensional, locally finite, non-positively curved cube complex. If $\pi_1 X \cong \mathbb{Z}^2 *_{\mathbb{Z}} \mathbb{Z}^2$, then for every vertex space $\widetilde{X}_{\tilde{v}}$ and adjacent edge space $\widetilde{X}_{\tilde{e}}$ there is a hyperplane H in $\widetilde{Y}$ that intersects $\widetilde{X}_{\tilde{v}}$ parallel to $\widetilde{X}_{\tilde{e}}$.*

Proof. Let $G = \pi_1 X$. There are precisely two vertex orbits and one edge orbit in $\widetilde{\Gamma}$. Assume that $\widetilde{Y} = \mathsf{hull}(\widetilde{X})$. Let $\mathcal{H}$ denote the set of all hyperplanes in $\widetilde{Y}$ intersecting $\widetilde{X}$. Let $\mathcal{H}_{\tilde{v}}$ denote the set of all hyperplanes intersecting $\widetilde{X}_{\tilde{v}}$.

For each vertex $\tilde{v}$, there is precisely one $G_{\tilde{v}}$-orbit of adjacent edges $\tilde{e}$. Therefore, if $H \in \mathcal{H}_{\tilde{v}}$ is non-parallel to an adjacent edge space to $\widetilde{X}_{\tilde{v}}$, it is non-parallel to all adjacent edge spaces, and by Lemma 5.2.5 it must intersect all adjacent vertex spaces non-parallel to all adjacent edge spaces. Therefore, we deduce that any hyperplane that doesn't intersect every vertex space in $\widetilde{X}$ will either intersect a vertex space parallel to its adjacent edge spaces, or its intersection will be a line contained in an edge space.

Suppose that there exists a vertex space $\widetilde{X}_{\tilde{v}_1}$ such that no hyperplane in $\mathcal{H}_{\tilde{v}_1}$ intersects $\widetilde{X}_{\tilde{v}_1}$ parallel to the adjacent edge spaces. Let $\hat{\mathcal{H}} = \mathcal{H} - \mathcal{H}_{\tilde{v}_1}$. Every hyperplane in $\mathcal{H} - \mathcal{H}_{\tilde{v}_1}$ must intersect each wall in $\mathcal{H}_{\tilde{v}_1}$ so we deduce that $\widetilde{Y} = \mathsf{hull}(\widetilde{X}_{\tilde{v}_1}) \times C(\widetilde{Y}, \hat{\mathcal{H}})$ (see [9, Lem 2.5]). Furthermore, $\mathcal{H}_{g\tilde{v}_1} = \mathcal{H}_{\tilde{v}_1}$ for all $g \in G$. By Lemma 5.2.4, $\mathsf{hull}(\widetilde{X}_{\tilde{v}})$ embeds in $\mathbf{R}^d$. Since $\widetilde{X}_{\tilde{v}_1}$ is contained inside some subcomplex $\mathsf{hull}(\widetilde{X}_{\tilde{v}_1}) \times C \subseteq \widetilde{Y}$, where C is the 0-cube determined by orienting all hyperplanes towards $\widetilde{X}_{\tilde{v}_1}$, we can conclude that only finitely many hyperplanes in $\hat{\mathcal{H}}$ intersect the r-neighborhood of $\widetilde{X}_{\tilde{v}_1}$.

Let $\widetilde{X}_{\tilde{u}}$ be a vertex space adjacent to $\widetilde{X}_{\tilde{v}_1}$, and let $\widetilde{X}_{\tilde{e}_1}$ be the edge space connecting them. Let $\widetilde{X}_{\tilde{v}_2}$ be another vertex space adjacent to $\widetilde{X}_{\tilde{u}}$ and let $\widetilde{X}_{\tilde{e}_2}$ be the edge space connecting them. Note that $\widetilde{X}_{\tilde{v}_1}$ and $\widetilde{X}_{\tilde{v}_2}$ are in the same $G_{\tilde{u}}$-orbit. The geodesic lines $\widetilde{X}_{\tilde{u}} \cap \widetilde{X}_{\tilde{e}_1}$ and $\widetilde{X}_{\tilde{u}} \cap \widetilde{X}_{\tilde{e}_2}$ are parallel in $\widetilde{X}_{\tilde{u}}$. Let $D \subseteq \widetilde{X}_{\tilde{u}}$ be the subspace isometric to $\mathbb{R} \times [a, b]$ bounded by these parallel lines. Let $U = D \cup \widetilde{X}_{\tilde{e}_1} \cup \widetilde{X}_{\tilde{e}_2}$. Finitely many hyperplanes in $\hat{\mathcal{H}}$ intersect D.

Let g_1 be an isometry in $G_{\tilde{v}_1}$ that stabilizes an axis in $\widetilde{X}_{\tilde{v}_1}$ that is non-parallel to the geodesic $\widetilde{X}_{\tilde{v}_1} \cap \widetilde{X}_{\tilde{e}_1}$. Similarly, let g_2 be an isometry in $G_{\tilde{v}_2}$ that stabilizes an axis in $\widetilde{X}_{\tilde{v}_2}$ that is non-parallel to the geodesic $\widetilde{X}_{\tilde{v}_2} \cap \widetilde{X}_{\tilde{e}_2}$. Note that

$F = \langle g_1, g_2 \rangle$ is a free group on two generators. Let r be such that U is contained in the r-neighbourhood of $\widetilde{X}_{\tilde{v}_1}$. As there are only finitely many hyperplanes in $\hat{\mathcal{H}}$ intersecting $\mathcal{N}_r(\widetilde{X}_{\tilde{v}_1})$, there must exist an n such that g_1^n stabilizes those walls. Similarly, since $\widetilde{X}_{\tilde{v}_2}$ is a translate of $\widetilde{X}_{\tilde{v}_1}$, we can deduce that there are only finitely many hyperplanes in $\hat{\mathcal{H}}$ intersecting $\mathcal{N}_r(\widetilde{X}_{\tilde{v}_2})$, and there must exist an m such that g_2^m stabilizes those walls. Let $F' = \langle g_1^n, g_2^m \rangle$. As U lies in both $\mathcal{N}_r(\widetilde{X}_{\tilde{v}_1})$ and $\mathcal{N}_r(\widetilde{X}_{\tilde{v}_2})$ we can deduce that the hyperplanes in $\hat{\mathcal{H}}$ that intersect the F'-translates of U are precisely the hyperplanes intersecting U. Let $\widetilde{Z} = F'\widetilde{X}_{\tilde{v}_1} \cup F'\widetilde{X}_{\tilde{v}_2} \cup F'U$. Then $\mathsf{hull}(\widetilde{Z}) = \mathsf{hull}(\widetilde{X}_{\tilde{v}_1}) \times K \subseteq \widetilde{Y}$, where K is a compact cube complex. Then F acts freely on $\mathsf{hull}(\widetilde{Z})$, which is a contradiction since number of 0-cubes intersecting the r-neighborhood of a 0-cube in $\mathsf{hull}(\widetilde{X}_{\tilde{v}_1}) \times K$ grows polynomially with r, and therefore cannot permit a free F-action. $\square$

Lemma 5.2.6 is a special case of the following more general statement:

Corollary 5.2.7. *Let $X \to Y$ be an amicable immersion, where Y is a finite dimensional, locally finite, non-positively curved cube complex. Then for every vertex space $\widetilde{X}_{\tilde{v}}$ and adjacent edge space $\widetilde{X}_{\tilde{e}}$ there is a hyperplane H in $\widetilde{Y}$ that intersects $\widetilde{X}_{\tilde{v}}$ parallel to $\widetilde{X}_{\tilde{e}}$.*

Proof. For every edge $\tilde{e}$ in $\widetilde{\Gamma}$ there is a subgroup $G' = \langle G_{-\tilde{e}}, G_{+\tilde{e}} \rangle \leq G$ such that $G' \cong \mathbb{Z}^2 *_{\mathbb{Z}} \mathbb{Z}^2$. Let $Y' = G' \backslash \widetilde{Y}$. Then there is an amicable immersion $X' \to Y'$ such that $\widetilde{X}'_{\tilde{e}} = \widetilde{X}$ and $\widetilde{X}'_{\pm\tilde{e}} = \widetilde{X}_{\pm\tilde{e}}$. Therefore, by Lemma 5.2.6, there is a hyperplane intersecting $\widetilde{X}_{-\tilde{e}}$ parallel to $\widetilde{X}_{\tilde{e}}$, and similarly for $\widetilde{X}_{+\tilde{e}}$. $\square$

The following proposition is a strengthening of one direction of Theorem 1.1 in [40]. Let $f_1 : A \to C$ and $f_2 : B \to C$ be maps between topological spaces A, B, C. The *fiber product* $A \otimes_C B = \{(a, b) \in A \times B \mid f_1(a) = f_2(b)\}$. Note that there are natural projections $p_1 : A \otimes_C B \to A$ and $p_2 : A \otimes_C B \to B$.

Proposition 5.2.8. *Let G be a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. Then there is a tubular space X with a finite set of immersed walls such that the associated wallspace $(\widetilde{X}, \mathcal{W})$ has the following properties:*

1. *G acts freely on $C(\widetilde{X}, \mathcal{W})$.*
2. *$C(\widetilde{X}, \mathcal{W})$ is finite dimensional if $\widetilde{Y}$ is finite dimensional.*
3. *$C(\widetilde{X}, \mathcal{W})$ is finite dimensional and locally finite if $\widetilde{Y}$ is locally finite.*

Proof. Let y be a 0-cube in $\widetilde{Y}$. By possibly replacing $\widetilde{Y}$ with $\mathsf{hull}(Gy)$ we can assume that $\widetilde{Y}$ has finitely many G-orbits of hyperplanes. Assume that $\widetilde{Y}$ is minimal in the sense that we cannot replace $\widetilde{Y}$ with a convex G-subcomplex with less hyperplane orbits. Let $Y = G\backslash\widetilde{Y}$. Let $h_1, \ldots, h_m$ be the immersed hyperplanes in Y. By Lemma 5.2.3 we can find an amicable immersion $X \to Y$. Every hyperplane in $\widetilde{Y}$ must intersect $\widetilde{X}$. Otherwise, if h is a hyperplane in $\widetilde{Y}$ that doesn't intersect $\widetilde{X}$, there is a 0-cube y in $\widetilde{Y}$ in the same halfspace of h as $\widetilde{X}$, and $\mathsf{hull}(Gy)$ is a convex G-subcomplex that doesn't contain the G-orbit of h, contradicting the minimality of $\widetilde{Y}$.

Let $h_i \to Y$ be an immersed hyperplane in Y. We obtain horizontal immersed walls in X by considering the components of the fiber product $X \otimes_Y h_i$ of $X \to Y$ and $h_i \to Y$. Each component Λ has a natural map into X. The components of $X \otimes_Y h_i$ that have image in X contained in an edge space are ignored. Let Λ^p be a component of $X \otimes_Y h_i$ whose image in X intersects a vertex space $X_v \subseteq X$. We will show that after a minor adjustment to Λ^p, we obtain a horizontal immersed wall and by considering all such components we obtain a set of horizontal walls in X obtained from an equitable set.

Using the map $\Lambda^p \to X$ we can decompose Λ^p into the components of the preimages of vertex space and edge spaces. As the intersection of each hyperplane $H \subseteq \widetilde{Y}$ with each vertex space $\widetilde{X}_{\tilde{v}}$ is either empty or a geodesic line, the intersection of each h_i with X_v is a set of geodesic curves, so Λ^p

restricted to the preimage of X_v is a set of geodesic curves. By Lemma 5.2.5 each hyperplane $H \subseteq \widetilde{Y}$ that intersects a vertex space $X_{\tilde{v}}$ non-parallel to an adjacent edge space $\widetilde{X}_{\tilde{e}}$ will intersect $\widetilde{X}_{\tilde{e}}$ as an arc with endpoints in $\widetilde{X}_{-\tilde{e}}$ and $\widetilde{X}_{+\tilde{e}}$. Thus the components of the intersection $X_e \cap h_i$ that intersect X_{-e} or X_{+e}, are arcs with endpoints in both X_{-e} and X_{+e}. Therefore, Λ^p decomposes into circles that map as local geodesics into vertex spaces, and arcs that map into edges spaces X_e with an endpoint in each X_{-e} and X_{+e}.

Let $\{\Lambda_1^p, \ldots, \Lambda_k^p\}$ be the set of all such components of $X \otimes_Y h_i$ that intersect vertex spaces. Let S_v^p be the set curves that map the circles in $\{\Lambda_i^p\}_{i=1}^k$ to the vertex space X_v. The elements of S_v^p and the attaching maps $\varphi_e^{\pm}$ of the edge spaces in X are locally geodesic curves, and $\#[\varphi_e^+, S_{e^+}] = \#[\varphi_e^-, S_{e^-}]$ since both sides are equal to the number of arcs in the walls $\{\Lambda_i^p\}_{i=1}^k$ that map into X_e. As G acts freely on $\widetilde{Y}$, there must be hyperplanes intersecting each vertex space $\widetilde{X}_{\tilde{e}}$ as geodesics in at least two parallelism classes. This implies that S_v^p contains curves generating at least two non-commensurable cyclic subgroups of G_v, and therefore S_v^p generates a finite index subgroup of G_v.

S_v^p is almost an equitable set: the images of the curves in S_v^p may not have images in X_v that coincide with each other or the image of an attaching map into X_v. Suppose that $\alpha_1, \ldots \alpha_m \in S_v$ be a maximal set of curves that have identical image in X_v. Let $\mathcal{N}_\epsilon(Q)$ denote the ϵ-neighborhood of a subset Q of either Y or $\widetilde{Y}$ with respect to the CAT(0) metric. Let $\epsilon \in (0, \frac{1}{3})$ be such that the neighbourhood $\mathcal{N}_\epsilon(\alpha_1) \subseteq Y$ only contains the images of $\alpha_1, \ldots, \alpha_m$ and the arcs connected to them. There is a homotopy of $\bigsqcup_{i=1}^k \Lambda_i^p \to X$ that is the identity outside of $\bigsqcup_{i=1}^k \Lambda_i^p \cap N_\epsilon(\alpha_1)$ such that $\alpha_1, \ldots, \alpha_m$ are homotoped to a disjoint set of geodesic curves in $X_v \cap N_\epsilon(\alpha_1)$ transverse or disjoint from all the other curves in S_v^p and the attaching maps into X_v. By choosing ϵ small enough we can perform such a homotopy $\Phi : \bigsqcup_{i=1}^k \Lambda_i^p \times [0,1] \to X$

such that all sets of overlapping curves in $\{S_v^p\}_{v \in V\Gamma}$ become disjoint and such that Φ is the identity map outside of the ϵ-neighborhood of the overlapping curves. The restriction of Φ to $\Lambda_i^p \times \{1\} \to X$ is an immersed wall which we will denote by Λ_i. Thus, the immersed walls $\{\Lambda_i\}_{i=1}^k$ are obtained from an equitable set $\{S_v\}_{v \in V\Gamma}$. We refer to $\{\Lambda_i^p\}_{i=1}^k$ as the *immersed proto-walls* and the lifts $\widetilde{\Lambda}_i^p \to \widetilde{X}$ as the *proto-walls*. Note that proto-walls have regular and non-regular intersections in the same way that walls do.

Let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from the immersed walls $\{\Lambda_i\}_{i=1}^k$ and adding a single vertical wall for each edge space. Each wall $\widetilde{\Lambda} \to \widetilde{X}$ covers an immersed wall $\Lambda \to X$. There exists a homotopy of $\Lambda \to X$ to the corresponding immersed proto-wall $\Lambda^p \to X$. This homotopy lifts to a homotopy from the immersed wall $\widetilde{\Lambda} \to \widetilde{X}$ to a unique proto-wall $\widetilde{\Lambda}^p \to \widetilde{X}$. Note that each wall is contained in the ϵ-neighborhood of its corresponding proto-wall. Each proto-wall corresponds to the intersection of a unique hyperplane in $\widetilde{Y}$ with the image of $\widetilde{X}$ in $\widetilde{Y}$. Therefore, each wall in $\mathcal{W}$ corresponds to a unique hyperplane in $\widetilde{Y}$.

Let $\widetilde{\Lambda}$ be a wall in $\mathcal{W}$, and let $\widetilde{\Lambda}^p$ be the corresponding proto-wall. Note that $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ and $\widetilde{\Lambda}^p \cap \widetilde{X}_{\tilde{v}}$ are either parallel geodesic lines, or both empty intersections. Therefore, if $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2 \in \mathcal{W}$ are a pair of regularly intersecting walls, then they correspond to a pair of regularly intersecting proto-walls, which correspond to a pair of intersecting hyperplanes in $\widetilde{Y}$.

If a pair of proto-walls $\widetilde{\Lambda}_1^p$ and $\widetilde{\Lambda}_2^p$ are disjoint, then the corresponding walls in $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ in $\mathcal{W}$ are also disjoint. Moreover, since $\widetilde{\Lambda}$ is contained in the ϵ-neighborhood of $\widetilde{\Lambda}^p$, a halfspace of $\widetilde{\Lambda}$ determines a halfspace of $\widetilde{\Lambda}^p$ and therefore a halfspace of the hyperplane H corresponding to $\widetilde{\Lambda}^p$.

Note that (1) follows from the the fact that $C(\widetilde{X}, \mathcal{W})$ was obtained from an equitable set for G, and thus by Theorem 4.2.1 G acts freely on $C(\widetilde{X}, \mathcal{W})$.

To prove (2), suppose that $C(\widetilde{X}, \mathcal{W})$ were infinite dimensional, then by Proposition 4.6.12, there would exists an infinite set of pairwise regularly intersecting walls in $\mathcal{W}$, which implies there is an infinite set of pairwise regularly intersecting proto-walls. Therefore, there is an infinite set of pairwise intersecting hyperplanes in $\widetilde{Y}$. This would imply that $\widetilde{Y}$ is an infinite dimensional CAT(0) cube complex. Therefore, if $\widetilde{Y}$ is finite dimensional, then so is $C(\widetilde{X}, \mathcal{W})$.

To prove (3) we first prove the following:

Claim 2. If $\widetilde{Y}$ is locally finite, then $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.

Proof. Suppose that $\widetilde{Y}$ is locally finite. If $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional, then by Proposition 4.6.12 it contains an infinite cube containing a canonical 0-cube z. Let $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_n, \ldots$ be the set of infinite pairwise crossing walls corresponding to the infinite cube. Let $\widetilde{\Lambda}^p_1, \ldots, \widetilde{\Lambda}^p_n, \ldots$ be the corresponding set of infinite pairwise crossing proto-walls, and let $H_1, \ldots, H_n, \ldots$ be the corresponding infinite family of pairwise crossing hyperplanes.

Suppose that Q is a subcomplex in $\widetilde{Y}$. Let $U(Q)$ denote the *cubical neighborhood* of Q, which is the union of all cubes in $U(Q)$ that intersect Q. As $\widetilde{Y}$ is locally finite, if Q is compact, then $U(Q)$ is also compact. By [20, Lem 13.15], if Q is convex, then so is $U(Q)$. Let $U^n(Q)$ denote the cubical neighborhood of $U^{n-1}(Q)$.

Let $x \in \widetilde{X}$ be a point determining the canonical 0-cube z in $C(\widetilde{X}, \mathcal{W})$. Let x be contained in a cube C in $\widetilde{Y}$. As C is compact and convex, $U^n(C)$ is also compact and convex, and therefore can only be intersected by finitely many H_i. Moreover, as $\bigcup_n U^n(C) = \widetilde{Y}$, every hyperplane intersects $U^n(C)$ for n sufficiently large. Thus, there exists $N, i > 1$ such that H_i intersects $U^{N+2}(C)$, but not $U^{N+1}(C)$. As $U^N(C)$ and the carrier of H_i are disjoint convex subcomplexes, there exists a hyperplane H that separates H_i from

$U^N(C)$. Note that $\mathsf{d}_{\widetilde{Y}}(x, H) \geq N$ and $\mathsf{d}_{\widetilde{Y}}(H, H_i) \geq 1$. Let $\widetilde{\Lambda}^p$ be the proto-wall corresponding to H and $\widetilde{\Lambda}$ be the corresponding wall. As $\widetilde{\Lambda}^p$ separates x from $\widetilde{\Lambda}_i^p$, we can conclude $\widetilde{\Lambda}$ separates x from $\widetilde{\Lambda}_i$ since the $\widetilde{\Lambda}$ and $\widetilde{\Lambda}_i$ are respectively contained in the ϵ-neighborhoods of $\widetilde{\Lambda}^p$ and $\widetilde{\Lambda}_i^p$. This contradicts the fact that z is incident to a 1-cube dual to hyperplane corresponding to $\widetilde{\Lambda}_i$. □

As $C(\widetilde{X}, \mathcal{W})$ is finite dimensional, we can apply Corollary 5.2.7 to each edge group in G to deduce that $\{\Lambda_i\}_{i=1}^k$ are fortified. Therefore, by Proposition 5.1.5 we deduce that $C(\widetilde{X}, \mathcal{W})$ is locally finite. □

We can now prove the main theorem of this thesis.

Theorem 5.2.9. *A tubular group G acts freely on a locally finite CAT(0) cube complex if and only if G is virtually special.*

Proof. Suppose that G is virtually special. Then G embeds as the subgroup of a finitely generated right angled Artin group R. The universal cover $\widetilde{Y}$ of the Salvetti complex corresponding to R is a locally finite CAT(0) cube complex. Therefore, as R acts freely on $\widetilde{Y}$, so does G.

Conversely, suppose that G acts freely on a locally finite CAT(0) cube complex. Let X be a tubular space such that $G = \pi_1 X$. By Proposition 5.2.8 there exists a finite set of immersed walls such that the dual of the associated wallspace $C(\widetilde{X}, \mathcal{W})$ is finite dimensional and locally finite. By Theorem 4.6.10 and Proposition 5.1.5 this is equivalent to saying that the immersed walls are non-dilated and fortified. By Lemma 4.7.4 we can assume that the immersed walls are also primitive. Therefore, by Proposition 5.1.6, G is virtually special.

□

5.3 Virtual Cubical Dimension

Let G be a tubular group. The first homology of G can be written as a direct sum of two factors: $\mathsf{H}_1(G) = \mathsf{H}_1^{\mathsf{v}}(G) \times \mathsf{H}_1^{\mathsf{s}}(G)$. The first factor $\mathsf{H}_1^{\mathsf{v}}(G)$ is generated by the image of the vertex groups of G, and the second factor $\mathsf{H}_1^{\mathsf{s}}(G)$ is generated by the stable letters in the presentation corresponding to the graph of groups decomposition.

Proposition 5.3.1. *Let X be a tubular space and $G = \pi_1 X$. Suppose that the natural maps $\mathsf{H}_1(G_v) \to \mathsf{H}_1(G)$ are injections as summands for all $v \in V\Gamma$. Given a primitive element a_v in a vertex group G_v, we can find non-dilated, primitive, embedded, pairwise disjoint immersed walls $\{\Lambda_i^a\}$ such that there is precisely one circle in each vertex space, and the circle in X_v represents a_v in $\pi_1 X_v$.*

Moreover, given primitive, incommensurable elements $a_{1v}, \dots, a_{nv} \in G_v$, the immersed walls $\{\Lambda_{1i}^{a_1}, \dots, \Lambda_{ni}^{a_n}\}$ produce an $(n+1)$-dimensional CAT(0) cube complex with a free G-action.

Proof. By our hypothesis, there are inclusion maps $\iota_v : G_v \to \mathsf{H}_1^{\mathsf{v}}(G)$, and projection maps $p_v : \mathsf{H}_1^{\mathsf{v}}(G) \to G_v$ for every vertex $v \in V\Gamma$. Let g_e generate the edge group G_e and recall that $\varphi_e^{\pm} : X_e^{\pm} \to X_{\pm e}$ denotes the attaching maps, so $(\varphi_e^-)_*$ and $(\varphi_e^+)_*$ denote the inclusions of G_e into G_{-e} and G_{+e} respectively. Let a_v be a primitive element in G_v, and let $a = \iota_v(a_v) \in \mathsf{H}_1^{\mathsf{v}}(G)$.

Suppose that $\mathsf{H}_1^{\mathsf{v}}(G) \cong \mathbb{Z}^2$, then ι_v and p_v are isomorphisms for all $v \in V\Gamma$. Let $S_u = \{a_u = \iota_u(a)\}$. Note that $\{S_u\}_{u \in V\Gamma}$ is not an equitable set for G as S_u contains a single element, but we will show that $\#[S_{-e}, (\varphi_e^-)_*(g_e)] = \#[S_{+e}, (\varphi_e^+)_*(g_e)]$ for all $e \in E\Gamma$ so $\{S_u\}_{u \in V\Gamma}$ can be used to construct an immersed walls in X.

There is an isomorphism $A_e = p_{+e} \circ \iota_{-e} : G_{-e} \to G_{+e}$ that maps $(\varphi_e^-)_*(g_e) \mapsto (\varphi_e^+)_*(g_e)$, and $a_{-e} \mapsto a_{+e}$. By identifying G_{-e} and G_{+e} with $\mathbb{Z}^2$ we can say that $A \in SL_2(\mathbb{Z})$. As the elements of $SL_2(\mathbb{Z})$ preserve the determinant, $\#[\cdot,\cdot]$ is preserved by A. Therefore,

$$\#[a_{-e},(\varphi_e^-)_*(g_e)] = \#[A(a_{-e}),A(\varphi_e^-)_*(g_e)] = \#[a_{+e},(\varphi_e^+)_*(g_e)],$$

As S_v contains precisely one element for each $v \in V\Gamma$, after adding non-intersecting arcs we obtain embedded immersed walls $\{\Lambda_i^a\}$ that are pairwise non-intersecting and have exactly one circle in each vertex space. As they are embedded, each $\Lambda_i^{a_v}$ is non-dilated. By construction, the circle in X_v represents $a_v \in \pi_1 X_v$. Since a_v is primitive in G_v, we deduce that $a = \iota_v(a_v)$ is also primitive in $\mathsf{H}_1^{\mathsf{v}}(G)$ as ι_v is an isomorphism, and thus $a_u = p_u(a)$ is also primitive since p_u is an isomorphism.

If $\mathsf{H}_1^{\mathsf{v}}(G) \cong \mathbb{Z}^d$ with $d > 2$, there exist vertex groups G_u and G_v that embed into $\mathsf{H}_1^{\mathsf{v}}(G)$ as distinct summands. Let $G_u = \langle g_u, g_u' \rangle$ and $G_v = \langle g_v, g_v' \rangle$. We can assume, since they are distinct summands, that g_u is disjoint from the image of G_v in $\mathsf{H}_1^{\mathsf{v}}(G)$, and that g_v is disjoint from the image of G_u in $\mathsf{H}_1^{\mathsf{v}}(G)$. By attaching an edge space to X connecting X_u and X_v, that has attaching maps representing g_u and g_v respectively we obtain a new graph of spaces X'. The resulting tubular group $G' = \pi_1 \hat{X}'$ has $\mathsf{H}_1^{\mathsf{v}}(G') \cong \mathbb{Z}^{d-1}$. By repeating this process we can obtain a tubular group $\hat{G}$ that is the fundamental group of a tubular space $\hat{X}$ such that $X \subseteq \hat{X}$ is a sub-graph of spaces. For $a_v \in \hat{G}_v = G_v$ we can find a set of non-dilated, primitive embedded immersed walls $\{\hat{\Lambda}_i^{a_v}\}$ that are pairwise disjoint, and such that each vertex space $\hat{X}_v$ contains precisely one circle in the collection of immersed walls, and the circle mapping into $\hat{X}_v = X_v$ represents a_v. By restricting these immersed walls to X we prove the first part of the claim.

Let $a_{1v}, \dots, a_{nv} \in G_v$ be incommensurate primitive elements with $d \geq 2$, then $\langle a_{1v}, \dots, a_{nv} \rangle$ is a finite index subgroup of G_v. Similarly, $\langle a_{1u} = p_u \circ$

$\iota_v(a_{1v}), \ldots, a_{nu} = p_u \circ \iota_u(a_{nv})\rangle$ is a finite index subgroup of G_u. Therefore, $\{S_u = \{p_u \circ \iota_v(a_{1v}), \ldots, p_u \circ \iota_u(a_{nv})\}\}_{u \in V\Gamma}$ is an equitable set. The set $\{\Lambda_i^{a_j}\}$ contains pairwise disjoint embedded, non-dilated immersed walls. After including the vertical immersed walls we conclude that any set of pairwise intersecting walls in the associated wallspace $(\widetilde{X}, \mathcal{W})$ can contain at most $(n+1)$ walls: a vertical wall, and a single wall covering an immersed wall in each $\{\Lambda_i^{a_j}\}$. Therefore $C(\widetilde{X}, \mathcal{W})$ is of dimension at most $(n+1)$, and since it was obtained from an equitable set, G acts freely on $C(\widetilde{X}, \mathcal{W})$. □

Lemma 5.3.2. *Let X be a tubular space and $G = \pi_1 X$. Suppose there exists an equitable set that produces primitive, non-dilated immersed walls in X. Then there exists a finite index subgroup $G' \leqslant G$ such that the natural map $\mathsf{H}_1(G'_v) \to \mathsf{H}_1(G')$ is an injection as a summand for each vertex group G'_v of the induced splitting of G'.*

Proof. Let $\widetilde{\Gamma}$ be the Bass-Serre tree. By Proposition 5.1.4, since there are immersed walls that are primitive and non-dilated, G acts freely on $\mathbf{R}^d \times \widetilde{\Gamma}$ such that $G_{\tilde{v}}$ fixes the vertex $\tilde{v}$ in $\widetilde{\Gamma}$. Therefore G is a subgroup of $\mathrm{Aut}(\mathbf{R}^d \times \widetilde{\Gamma}) \cong \mathbb{Z}^d \rtimes \mathrm{Aut}([-1,1]^d) \times \mathrm{Aut}(\widetilde{\Gamma})$. Let K be the kernel of $G \hookrightarrow \mathbb{Z}^d \rtimes \mathrm{Aut}([-1,1]^d) \times \mathrm{Aut}(\widetilde{\Gamma}) \to \mathrm{Aut}([-1,1]^d)$. Note that $K \leqslant G$ is a finite index subgroup that embeds in $\mathbb{Z}^d \times \mathrm{Aut}(\widetilde{\Gamma})$. Let $p : K \to \mathbb{Z}^d$ be the projection onto the first factor. Each vertex group in K survives in the image of p, and therefore we have embedding $\mathsf{H}_1(K_v) \hookrightarrow \mathsf{H}_1^{\mathsf{v}}(G'') \hookrightarrow \mathbb{Z}^d$.

For each vertex group K_v in the graph of groups decomposition of K, there is a finite index subgroup $A_v \leqslant \mathbb{Z}^d$ such that $p(K_v) \cap A_v$ is a summand of A_v. Let $A = \cap_v A_v$ and $G' = p^{-1}(A)$. Note that $G' \leqslant K \leqslant G$ are finite index subgroups. Each vertex group in G' will be a factor in A. As A is free abelian, the map $G' \to A$ will factor through the $\mathsf{H}_1^{\mathsf{v}}(G')$, so we can deduce that each vertex group survives as a retract in $\mathsf{H}_1(G')$. Therefore, each vertex group in G' survives as a summand in the first homology. $\square$

Corollary 5.3.3. *Let X be a tubular space. If the tubular group $G = \pi_1 X$ acts freely on a finite dimensional CAT(0) cube complex, then G acts on a locally finite, finite dimensional CAT(0) cube complex*

Proof. If G acts freely on a finite dimensional CAT(0) cube complex, then by Proposition 5.2.8, Theorem 4.6.10 and Proposition 4.7.4 there is an equitable set that produces primitive, non-dilated immersed walls. By Lemma 5.3.2 there exists $G' \leqslant G$ such that the natural map $\mathsf{H}_1(G'_v) \to \mathsf{H}_1(G')$ is an injection as a summand for each vertex group G'_v of the induced splitting of G'. Therefore, by Proposition 5.3.1, for each edge group $G_e = \langle g_e \rangle$ there is a pair of primitive, non-dilated, immersed walls Λ_e^+ and Λ_e^- such that $\Lambda_e^\pm$ contains a circle that maps into $X_{\pm e}$, representing $(\phi_e^\pm)_*(g_e)$. If the circles in the immersed walls $\{\Lambda_e^\pm\}_{e \in E\Gamma}$ don't constitute an equitable set, use Proposition 5.3.1, to find more primitive, non-dilated walls, so that the circles do constitute an equitable set. By construction the set of immersed walls is fortified.

By Theorem 4.6.10, Proposition 5.1.6, and Theorem 5.2.9 the corresponding dual cube complex $C(\widetilde{X}, \mathcal{W})$ is finite dimensional, locally finite, and admits a free G' action. There is an induced free action of G on $C(\widetilde{X}, \mathcal{W})^{[G:G']}$, which is also a locally finite, finite dimensional CAT(0) cube complex. $\square$

References

[1] Ian Agol. The virtual haken conjecture. 2012. With an appendix by Ian Agol, Daniel Groves, and Jason Manning.

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

[3] Brian H. Bowditch. *A course on geometric group theory*, volume 16 of *MSJ Memoirs*. Mathematical Society of Japan, Tokyo, 2006.

[4] N. Brady and M. R. Bridson. There is only one gap in the isoperimetric spectrum. *Geom. Funct. Anal.*, 10(5):1053–1070, 2000.

[5] Martin R. Bridson. On the semisimplicity of polyhedral isometries. *Proc. Amer. Math. Soc.*, 127(7):2143–2146, 1999.

[6] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

[7] R. G. Burns, A. Karrass, and D. Solitar. A note on groups with separable finitely generated subgroups. *Bull. Austral. Math. Soc.*, 36(1):153–160, 1987.

[8] J. W. Cannon. Geometric group theory. In *Handbook of geometric topology*, pages 261–305. North-Holland, Amsterdam, 2002.

[9] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for cat(0) cube complexes. *Geometric And Functional Analysis*, 21:851–891, 2011. 10.1007/s00039-011-0126-7.

[10] Christopher H. Cashen. Quasi-isometries between tubular groups. *Groups Geom. Dyn.*, 4(3):473–516, 2010.

[11] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata*, 125:141–158, 2007.

[12] Daniel S. Farley. Finiteness and CAT(0) properties of diagram groups. *Topology*, 42(5):1065–1082, 2003.

[13] Ross Geoghegan. *Topological methods in group theory*, volume 243 of *Graduate Texts in Mathematics*. Springer, New York, 2008.

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

[15] S. M. Gersten. The automorphism group of a free group is not a CAT(0) group. *Proc. Amer. Math. Soc.*, 121(4):999–1002, 1994.

[16] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

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

[18] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *Duke Math. J.*, 165(9):1753–1813, 2016.

[19] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. pages 1–17, 2007.

[20] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1 551–1620, 2008.

[21] Frédéric Haglund and Daniel T. Wise. Coxeter groups are virtually special. *Adv. Math.*, 224(5):1890–1903, 2010.

[22] G. Christopher Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. *Compositio Mathematica*. pp. 1–58, to appear.

[23] Yi Liu. Virtual cubulation of nonpositively curved graph manifolds. *J. Topol.*, 6(4):793–822, 2013.

[24] Yi Liu. A characterization of virtually embedded subsurfaces in 3-manifolds. *Trans. Amer. Math. Soc.*, 369(2):1237–1264, 2017.

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

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

[27] Graham A. Niblo and Daniel T. Wise. The engulfing property for 3-manifolds. In *The Epstein birthday schrift*, pages 413–418 (electronic). Geom. Topol., Coventry, 1998.

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

[29] Piotr Przytycki and Daniel T. Wise. Mixed 3-manifolds are virtually special. pages 1–24. Available at arXiv:1205.6742.

[30] Piotr Przytycki and Daniel T. Wise. Graph manifolds with boundary are virtually special. *J. Topol.*, 7(2):419–435, 2014.

[31] J. Hyam Rubinstein and Shicheng Wang. π_1-injective surfaces in graph manifolds. *Comment. Math. Helv.*, 73(4):499–515, 1998.

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

[33] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.

[34] Michah Sageev. CAT(0) cube complexes and groups. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 7–54. Amer. Math. Soc., Providence, RI, 2014.

[35] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.

[36] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. Available at `http://www.math.mcgill.ca/wise/papers`. pp. 1–189. Submitted.

[37] Daniel T. Wise. A non-Hopfian automatic group. *J. Algebra*, 180(3):845–847, 1996.

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

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

[40] Daniel T. Wise. Cubular tubular groups. *Trans. Amer. Math. Soc.*, 366(10):5503–5521, 2014.

[41] Daniel T. Wise and Daniel J. Woodhouse. A cubical flat torus theorem and the bounded packing property. *Accepted Israel J. Math.* Available at http://arxiv.org/abs/1510.00365.

[42] Daniel J. Woodhouse. A generalized axis theorem for cube complexes. Submitted. Available at http://arxiv.org/abs/1602.01952.

[43] Daniel J. Woodhouse. Classifying virtually special tubular groups. 2015. In Preparation.

[44] Daniel J. Woodhouse. Classifying finite dimensional cubulations of tubular groups. *Michigan Math. J.*, 65(3):511–532, 2016.

Index