Classifying virtually special tubular groups

Daniel J. Woodhouse

Abstract. A group is tubular if it acts on a tree with $\mathbb{Z}^2$ vertex stabilizers and $\mathbb{Z}$ edge stabilizers. We prove that a tubular group being virtually special is equivalent to it acting freely on either a locally finite or finite dimensional CAT(0) cube complex. Furthermore, we prove that if a tubular group acts freely on a finite dimensional CAT(0) cube complex, then it virtually acts freely on a three dimensional CAT(0) cube complex.

1. Introduction

A *tubular group G* splits as a graph of groups with $\mathbb{Z}^2$ vertex groups and $\mathbb{Z}$ edge groups. Equivalently, G is the fundamental group of a graph of spaces, denoted by X, with each vertex space homeomorphic to a torus and each edge space homeomorphic to $S^1 \times [-1, 1]$. The graph of spaces X is a *tubular space*. In this paper all tubular groups will be finitely generated and therefore have compact tubular spaces.

Tubular groups have been studied from various perspectives: Brady and Bridson provided tubular groups with isoperimetric function n^α for all α in a dense subset of $[2, \infty)$ in [1]. Cashen determined when two tubular groups are quasi-isometric [5]. Wise determined whether or not a tubular group acts freely on a CAT(0) cube complex [12], and classified which tubular groups are cocompactly cubulated. The author determined a criterion for finite dimensional cubulations [13]. Button has proven that all free-by-cyclic groups that are also tubular groups act freely on finite dimensional cube complexes [3]. The main theorem of this paper is the following one.

Theorem 1.1. *Let G be a tubular group. The following are equivalent*:
 (1) *G acts freely on a locally finite* CAT(0) *cube complex*;
 (2) *G acts freely on a finite dimensional* CAT(0) *cube complex*;
 (3) *G is virtually special.*

Haglund and Wise introduced special cube complexes in [7]. The main consequence of a group being special is that it embeds in a right angled Artin group (see [10] or [11] for a full outline of Wise's program). Note that if G' is a finite index, special subgroup of G, then G' acts freely on the universal cover of the Salvetti complex of some finitely generated right angled Artin group. The Salvetti complex is finite dimensional and locally finite, and G acts freely on $X^{[G:G']}$, so we can deduce that (3) implies (1) and (2). The proof of Theorem 1.1 will therefore be completed by proving that (1) implies (3) in Theorem 5.10, and (2) implies (1) in Corollary 6.4.

1.1. Structure of the Paper. In [12] Wise obtained free actions of tubular groups on CAT(0) cube complexes by first finding *equitable sets* that allow the construction of *immersed walls*. Such a set of immersed walls determines a *wallspace* $(\widetilde{X}, \mathcal{W})$ which yields a dual cube complex $C(\widetilde{X}, \mathcal{W})$ which G acts freely on. Wallspaces were first introduced by Haglund and Paulin [6], and the dual cube complex construction was first developed by Sageev [9]. In [13] the author defined a criterion, called *dilation*, that determines if an immersed wall produces infinite or finite dimensional cubulations. More precisely, if the immersed walls are *non-dilated*, then $C(\widetilde{X}, \mathcal{W})$ is finite dimensional. We recall the relevant definitions and background in Section 2.

Section 3 establishes a technical result using techniques from [13]. It is shown that that immersed walls can be replaced with *primitive* immersed walls without losing the finite dimensionality or local finiteness of the associated dual cube complex. The reader is encouraged to either read this section alongside [13], or skip it on a first reading.

In Section 4 we analyze $C(\widetilde{X}, \mathcal{W})$ 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 $\widetilde{X}$. A further criterion, the notion of a *fortified* immersed wall, determines when $C(\widetilde{X}, \mathcal{W})$ 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 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 result.

Proposition 5.9. *Let G be a tubular group acting freely on a* CAT(0) *cube complex $\widetilde{Y}$. There is a tubular space X and a finite set of immersed walls in X such that, 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 5.10 which is the equivalence of (1) and (3) in Theorem 1.1.

In Section 6 we further exploit the results obtained in Section 4 to obtain the following, demonstrating that the cubical dimension of tubular groups with finite dimensional cubulations are virtually within 1 of their cohomological dimension.

Corollary 6.3. *If a tubular group G acts freely on a finite dimensional* CAT(0) *cube complex, then G virtually acts freely on a* CAT(0) *cube complex of dimension* 3.

The techniques in this section will also allow us to prove Corollary 6.4 that (2) implies (1) in Theorem 1.1.

2. Background tubular groups and their cubulations

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 $\varphi_e^-\colon X_e^- \to X_{-e}$, and $\varphi_e^+\colon X_e^+ \to X_{+e}$. 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 $\widetilde{X}_{\widetilde{v}}$ and $\widetilde{X}_{\widetilde{e}}$ 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 nonpositively curved geodesic metric space and that attaching maps φ_e^- and φ_e^+ define locally geodesic curves in X_{-e} and X_{+e}.

2.1. Equitable Sets and Intersection Numbers. Given a pair of closed curves in a torus $\alpha, \beta\colon S^1 \to T$, 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* $\#[[\alpha], [\beta]]$ is the minimal number of intersection points realized by a pair of representatives from the respective classes. This number is realized 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 $\#[[\alpha], [\beta]] = \det[[\alpha], [\beta]]$. 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 disjoint from the attaching maps of adjacent edge spaces,

such that S_v generate a finite index subgroup of $\pi_1 X_v = G_v$, and $\#[\varphi_e^-, S_{-e}] = \#[\varphi_e^+, S_{+e}]$. 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.

2.2. Immersed Walls From Equitable Sets.

Immersed walls are constructed from *circles* and *arcs*. For each $\alpha \in S_v$, let S_α^1 be the domain of α. The disjoint union $\bigsqcup S_\alpha^1$ 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 $(p^-, q^-) \in S_{\alpha^-}^1 \times X_e^-$ and $(p^+, q^+) \in S_{\alpha^+}^1$ be corresponding intersection points between $\alpha^\pm \in S_{\pm e}$ and $\varphi_e^\pm$. 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*. Each graph Λ_i has its own graph of groups structure with infinite cyclic vertex groups and trivial edge groups. *As in* [13], *all "immersed walls" in this paper are immersed walls constructed from equitable sets as above.* This means that we are free use the results obtained in [13].

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 $\alpha_e \colon S^1 \to X_e$ 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 horizontal and vertical walls gives a *wallspace* $(\widetilde{X}, \mathcal{W})$ where the G-action on $\widetilde{X}$ also gives an action on $\mathcal{W}$. The main theorem of [12] is that a tubular group acts freely on a CAT(0) cube complex if and only if there exists an equitable set.

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.

Let $\widetilde{\Lambda}$ and $\widetilde{\Lambda}'$ be horizontal walls in $\widetilde{X}$. An point $x \in \widetilde{\Lambda} \cap \widetilde{\Lambda}$ is a *regular intersection point* if it lies in a vertex space $\widetilde{X}_{\widetilde{v}}$, and the lines $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ and $\widetilde{\Lambda}' \cap \widetilde{X}_{\widetilde{v}}$ are non-parallel. Otherwise, a point $x \in \widetilde{\Lambda} \cap \widetilde{\Lambda}'$ is a *non-regular intersection point*, and either $x \in \widetilde{X}_{\widetilde{v}}$ where $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}} = \widetilde{\Lambda}' \cap \widetilde{X}_{\widetilde{v}}$, or $x \in \widetilde{X}_{\widetilde{e}}$.

An *infinite cube* in a CAT(0) cube complex $\widetilde{Y}$, is an sequence of $c_0, c_1, \ldots, c_n, \ldots$ such that c_n is an n-cube in $\widetilde{Y}$, and c_n is a face of c_{n+1}. In [13] a *dilation function* is constructed for each immersed wall $R \colon \pi_1 \Lambda \to \mathbb{Q}^*$, and an immersed

wall is said to be *dilated* if R has infinite image. The following is Theorem 1.2 from that paper.

Theorem 2.1. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The following conditions 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*.

The following result is also obtained from [13] by combining Thm 1.2, Prop 4.6, and Prop 3.4. The last part follows from the last paragraph of the proof of Prop 3.4.

Proposition 2.2. *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.*

3. Primitive immersed walls

The following result uses the techniques in Section 5. of [13] to compute the dilation function. Let Λ be an immersed wall in X, and let $R: \pi_1(\Lambda) \to \mathbb{Q}^*$ be its dilation function. If R has finite image then Λ is non-dilated. Let $q: \Lambda \to \Omega$ be the quotient map obtained by crushing each circle to a vertex. Note that the arcs in Ω correspond to the arcs in Λ. The dilation function R factors through $q_*: \pi_1\Lambda \to \pi_1\Omega$, so there exists a function $\widehat{R}: \pi_1\Omega \to \mathbb{Q}^*$ such that $R = \widehat{R} \circ q_*$. We can therefore determine if Λ is dilated by computing the function $\widehat{R}$.

We orient each arc in Λ so that all arcs embedded in the same edge space are oriented in the same direction. We orient the arcs in Ω accordingly. We define a weighting $\omega: E(\Omega) \to \mathbb{Q}^*$. Let X_e be an edge space in X, and let σ be an arc mapped into X_e connecting the circles C^- and C^+. Let $\alpha^-: C^- \to X_{-e}$ and $\alpha^+: C^+ \to X_{+e}$ be the corresponding elements in the equitable set. Then

$$\omega(\sigma) = \frac{\#[\varphi_e^-, \alpha^-]}{\#[\varphi_e^+, \alpha^+]}.$$

If $\gamma = \sigma_1^{\epsilon_1} \cdots \sigma_n^{\epsilon_n}$ is an edge path in Ω where σ_i is an oriented arc in Ω, and $\epsilon \in \pm 1$, then $\widehat{R}(\gamma) = \omega(\sigma_1)^{\epsilon_1} \cdots \omega(\sigma_n)^{\epsilon_n}$.

Lemma 3.1. *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 $S^1_{g^n}$ with $S^1_1, \dots S^1_n$ and reattaching the arcs that were attached to the intersection points in S^1_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_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 $\widehat{R}_i : \pi_1 \Omega_i \to \mathbb{Q}^*$ and $\widehat{R}_{ij} : \pi_1 \Omega_{ij} \to \mathbb{Q}^*$ be the unique maps such that $R_i = \widehat{R}_i \circ q_i$ and $R_{ij} = \widehat{R}_{ij} \circ q_{ij}$. Let ω_i and ω_{ij} be the respective weightings of the arcs in Ω_i and Ω_{ij}. By assumption, R_i and $\widehat{R}_i$ 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 $\widehat{R}_{ij}(\gamma) = \widehat{R}_i(\rho_i \circ \gamma)$.

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 that 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_ι^1, 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, $\widehat{R}_{ij}(\gamma) = \widehat{R}_i(\rho_i \circ \gamma)$.

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. $\square$

4. 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 2.1, 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 $\widetilde{X}_{\tilde{e}}$.

We refer to [8] for full background on the dual cube complex construction. A 0-cube z in $\widetilde{Z}$ is a choice of halfspace $z[\widetilde{\Lambda}]$ of $\widetilde{\Lambda}$ for each $\widetilde{\Lambda} \in \mathcal{W}$ such that

(1) if $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2 \in \mathcal{W}$, then $z[\widetilde{\Lambda}_1] \cap z[\widetilde{\Lambda}_2] \neq \emptyset$;

(2) if $x \in X$, then there are only finitely many $\widetilde{\Lambda} \in \mathcal{W}$ such that $x \notin z[\widetilde{\Lambda}]$.

Two 0-cubes z_1, z_2 are adjacent if $z_1[\widetilde{\Lambda}] = z_2[\widetilde{\Lambda}]$ for all but precisely one hyperplane $\widetilde{\Lambda}_{12} \in \mathcal{W}$. The 1-cube joining z_1 and z_2 is dual to the hyperplane corresponding to $\widetilde{\Lambda}_{12}$. An n-cube is then present wherever the 1-skeleton of an n-cube appears. We say that two disjoint walls $\widetilde{\Lambda}, \widetilde{\Lambda}'$ *face* each other in z if $z[\widetilde{\Lambda}]$ is not contained in $z[\widetilde{\Lambda}']$ and vice versa.

Proposition 4.1. *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 $\widetilde{\Lambda}_{\tilde{e}} \in \mathcal{W}_v$.*

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 $\widetilde{\Lambda} = \widetilde{\Lambda}_e \in \mathcal{W}_v$ 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 $\widetilde{\Lambda}_e$. □

Proposition 4.1 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 following proposition, which collects the principal consequences of finite dimensionality, is Proposition 4.12 in [13].

Proposition 4.2. *Let X be tubular space with geodesic attaching maps, and let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a finite set of immersed walls in X. If the dual cube complex $C(\widetilde{X}, \mathcal{W})$ is finite dimensional, then the horizontal walls in $\mathcal{W}$ can be partitioned into a collection $\mathcal{P}$ of subsets such that the following conditions are satisfied.*

(1) *The partition $\mathcal{P}$ is preserved by G.*

(2) *For each $A \in \mathcal{P}$, the walls in A are pairwise non-intersecting.*

(3) *Let $\widetilde{\Lambda} \in A \in \mathcal{P}$ be a wall intersecting $\widetilde{X}_{\tilde{v}}$. There exists $h \in G_{\tilde{v}}$ stabilizing an axis in $\widetilde{X}_{\tilde{v}}$ perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\tilde{v}}$ such that $A = \{h^r \widetilde{\Lambda}\}_{r \in \mathbb{Z}}$.*

Any partition of the horizontal walls in $\mathcal{W}$ satisfying conditions (1)–(3) in Proposition 4.2 will be called a *stable partition*.

Lemma 4.3. *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 $\widetilde{X}_{\tilde{v}}$. □

The immersed walls $\Lambda_1, \ldots, \Lambda_k$ are non-dilated, and therefore $\widetilde{Z}$ is finite dimensional, so by Proposition 4.2 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 4.3, 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 $A_i = \{h_i^r \tilde{\Lambda}_i\}_{r \in \mathbb{Z}}$ such that $h_i \in G_{\tilde{v}}$ stabilizes an axis in $\tilde{X}_{\tilde{v}}$ perpendicular to $\tilde{\Lambda}_i \cap \tilde{X}_{\tilde{v}}$. 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 $g \cdot h_i^{\alpha_i} \tilde{\Lambda}_i = h_j^{\beta_j} \tilde{\Lambda}_j$. 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 $g \cdot h_i^{\alpha_i} \tilde{\Lambda}_i = h_j^{\beta_j} \tilde{\Lambda}_j$, then necessarily $g \cdot h_i^{\alpha_i + 1} \tilde{\Lambda}_i = h_j^{\beta_j \pm 1} \tilde{\Lambda}_j$, 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 $\tilde{Z}_{\tilde{v}}$. Since $G_{\tilde{v}}$ acts freely on $\tilde{Z}_{\tilde{v}}$ 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 $\phi_{\tilde{v}} : \tilde{Z}_{\tilde{v}} \to \mathbf{R}^{d_{\tilde{v}}}$. If z is a 0-cube in $\tilde{Z}_{\tilde{v}}$ then every wall $\tilde{\Lambda}$ that is either vertical or not in contained in the subpartition $\mathcal{P}_{\tilde{v}}$ has $\tilde{X}_{\tilde{v}} \subseteq z[\tilde{\Lambda}]$. Therefore z is entirely determined by $z[\tilde{\Lambda}]$ for $\tilde{\Lambda}$ in $\mathcal{P}_{\tilde{v}}$. For $1 \leq i \leq d\tilde{v}$ the set $\{h_i^r \tilde{\Lambda}_i \cap \tilde{X}_{\tilde{v}}\}_{r \in \mathbb{Z}}$ is an infinite collection of disjoint parallel lines in $\tilde{X}_{\tilde{v}}$. As all the walls in A_i are disjoint in $\tilde{X}$, for each 0-cube z in $\tilde{Z}_{\tilde{v}}$ there exists a unique $\alpha_i \in \mathbb{Z}$ such that $h_i^{\alpha_i} \tilde{\Lambda}_i$ and $h_i^{\alpha_i + 1} \tilde{\Lambda}_i$ 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 4.4. *The embedding $\phi_{\tilde{v}} : \tilde{Z}_{\tilde{v}} \to \mathbf{R}^{d_{\tilde{v}}}$ 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 $(gz)[h_i^{\alpha_i} \tilde{\Lambda}_i] = z[gh_i^{\alpha_i} \tilde{\Lambda}_i] = z[h_j^{\beta_j} \tilde{\Lambda}_j]$, which implies that $\phi_{\tilde{v}}(gz) = (\beta_1, \ldots, \beta_{d_{\tilde{v}}}) = g\phi_{\tilde{v}}(z)$. $\square$

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 $g \cdot h_i^{\alpha_i} \tilde{\Lambda}_i = h_j^{\beta_j} \tilde{\Lambda}_j$. 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 $\phi_{\tilde{e}} : \tilde{Z}_{\tilde{e}} \to \mathbf{R}^{d_{\tilde{e}}} \times [-1, 1]$. Let z be a 0-cube in $\tilde{Z}_{\tilde{e}}$. Then for each $1 \leq i \leq d_{\tilde{e}}$ there exists a unique α_i such that $h_i^{\alpha_i} \tilde{\Lambda}_i$ faces $h_i^{\alpha_i + 1} \tilde{\Lambda}_i$ in z, and $\tilde{X}_{\pm \tilde{e}} \subseteq z[\tilde{\Lambda}_{\tilde{e}}]$. 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 $H_{\tilde{e}} \subseteq \tilde{Z}$ be the hyperplane corresponding to $\tilde{\Lambda}_e$. As $\tilde{Z}_{\tilde{e}}$ is the carrier of $H_{\tilde{e}}$,
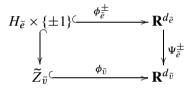
we can identify $\widetilde{Z}_{\tilde{e}}$ with $H_{\tilde{e}} \times [-1, 1]$. Note that $H_{\tilde{e}} \times \{\pm 1\}$ embeds as a subspace in $\widetilde{Z}_{\tilde{v}}$, and $\phi_{\tilde{e}}$ restricts to an embedding $\phi_{\tilde{e}}^{\pm} \colon H_{\tilde{e}} \times \{\pm 1\} \to \mathbf{R}^{d_{\tilde{e}}}$, where $\tilde{v} = \pm \tilde{e}$.

We construct an embedding $\Psi_{\tilde{e}}^{\pm} \colon \mathbf{R}^{d_{\tilde{e}}} \to \mathbf{R}^{d_{\tilde{v}}}$. 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 4.5. *The following commutative square is $G_{\tilde{e}}$-equivariant provided the immersed walls are primitive.*

$$
\begin{array}{ccc}
H_{\tilde{e}} \times \{\pm 1\} & \overset{\phi_{\tilde{e}}^{\pm}}{\hookrightarrow} & \mathbf{R}^{d_{\tilde{e}}} \\
\downarrow & & \downarrow{\scriptstyle \Psi_{\tilde{e}}^{\pm}} \\
\widetilde{Z}_{\tilde{v}} & \overset{\phi_{\tilde{v}}}{\hookrightarrow} & \mathbf{R}^{d_{\tilde{v}}}
\end{array}
$$

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

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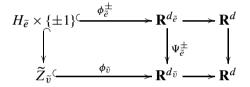$$

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

Let $d = \max\{|\mathcal{P}_{\tilde{v}}| \mid \tilde{v} \in V\tilde{\Gamma}\}$, which is finite, since there are only finitely many vertex orbits.
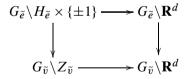
Proposition 4.6. *If the immersed walls $\Lambda_1, \ldots, \Lambda_k$ are primitive, then G acts freely on $\mathbf{R}^d \times \tilde{\Gamma}$ such that the action on the $\tilde{\Gamma}$ factor is the action of G on the Bass-Serre tree. Moreover, there is a G-equivariant embedding $\phi \colon \tilde{Z} \to \mathbf{R}^d \times \tilde{\Gamma}$.*

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 4.5 can be extended:

$$
\begin{array}{ccccc}
H_{\tilde{e}} \times \{\pm 1\} & \overset{\phi_{\tilde{e}}^{\pm}}{\hookrightarrow} & \mathbf{R}^{d_{\tilde{e}}} & \longrightarrow & \mathbf{R}^d \\
\downarrow & & \downarrow{\scriptstyle \Psi_{\tilde{e}}^{\pm}} & & \downarrow \\
\tilde{Z}_{\tilde{v}} & \underset{\phi_{\tilde{v}}}{\hookrightarrow} & \mathbf{R}^{d_{\tilde{v}}} & \longrightarrow & \mathbf{R}^d
\end{array}
$$

The right square commutes and is $G_{\tilde{e}}$-equivariant since by Lemma 4.5, the $G_{\tilde{e}}$-equivariant inclusion $\Psi_{\tilde{e}}^{\pm}$ 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 $\tilde{Z}$ into a tree of spaces with underlying tree $\tilde{\Gamma}$ 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:

$$
\begin{array}{ccc}
G_{\tilde{e}} \backslash H_{\tilde{e}} \times \{\pm 1\} & \longrightarrow & G_{\tilde{e}} \backslash \mathbf{R}^d \\
\downarrow & & \downarrow \\
G_{\tilde{v}} \backslash Z_{\tilde{v}} & \longrightarrow & G_{\tilde{v}} \backslash \mathbf{R}^d
\end{array}
$$

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 $\tilde{\mathcal{Z}} = \mathbf{R}^d \times \tilde{\Gamma}$, 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 $\tilde{Z} \to \tilde{\mathcal{Z}} = \mathbf{R}^d \times \tilde{\Gamma}$.

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

Proposition 4.7. *$\tilde{Z}$ 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 $\mathcal{P}_{\tilde{e}} = \mathcal{P}_{\tilde{v}}$. 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 $\phi_{\tilde{v}} \colon \widetilde{Z}_{\tilde{v}} \to \mathbf{R}^{d_{\tilde{v}}}$ 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 4.8. *If $\Lambda_1, \ldots, \Lambda_k$ are primitive, fortified, non-dilated immersed walls, then G is virtually special.*

Proof. By Proposition 4.6, there is a free action of G on $\mathbf{R}^d \times \widetilde{\Gamma}$, so G is a subgroup of $\mathrm{Isom}(\mathbf{R}^d \times \widetilde{\Gamma}) \cong (\mathbb{Z}^d \rtimes \mathrm{Aut}([-1, 1])^d) \times \mathrm{Aut}(\widetilde{\Gamma})$. Therefore, there is a projection $\rho \colon G \to \mathbb{Z}^d \rtimes \mathrm{Aut}([-1, 1]^d)$. Each vertex group $G_{\tilde{v}}$ embeds in $\mathbb{Z}^d$, and the mapping is invariant under conjugation. As there are only finitely many 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 $G' = \rho^{-1}((D\mathbb{Z})^d)$. Then $G' \leqslant G$ is a finite index subgroup that embeds in $(D\mathbb{Z})^d \times \mathrm{Aut}(\widetilde{\Gamma})$ such that each edge group is generated by an element that is primitive in the adjacent vertex groups. By Proposition 4.6 there is a G-equivariant embedding $\phi \colon \widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$. As G' does not permute the factors

of $\mathbf{R}^d$ we can deduce that the hyperplanes in $(\mathbf{R}^d \times \widetilde{\Gamma})/G$ 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 $gh_{d_{\tilde{v}}}^r \widetilde{\Lambda}_{d_{\tilde{v}}} \neq h_{d_{\tilde{v}}}^r \widetilde{\Lambda}_{d_{\tilde{v}}}$ 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. $\qquad\square$

We note that the requirement in Proposition 4.8 that the immersed walls are fortified is necessary, as the following example demonstrates.

Example 4.9. 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. 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 can 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.

5. Revisiting equitable sets

Although Wise proved in [12] that a tubular group 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 2, no relationship was established between $\widetilde{Y}$ and the
resulting dual $C(\widetilde{X}, \mathcal{W})$. Proposition 5.9 gives the relationship required to reduce
Theorem 1.1 to considering cubulations obtained from equitable sets.

This section will apply the following theorem from [14]. A *cubical quasiline*
is a CAT(0) cube complex quasi-isomorphic to $\mathbb{R}$.

Theorem 5.1. *Let G be virtually $\mathbb{Z}^n$. Suppose G acts properly and without
inversions on a CAT(0) cube complex $\widetilde{Y}$. Then G stabilizes a finite dimensional
subcomplex $\widetilde{Z} \subseteq \widetilde{Y}$ that is isometrically embedded in the combinatorial metric,
and $\widetilde{Z} \cong \prod_{i=1}^{m} C_i$, 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 $\widetilde{Z}$.*

Theorem 5.1 allows us to prove the following result.

Lemma 5.2. *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-equivariant subspace $\widetilde{X}_v \subseteq \widetilde{Y}$
homeomorphic to $\mathbb{R}^2$. Moreover, $\widetilde{X}_{\widetilde{v}}$ has a metric such that the intersection of a
hyperplane in $\widetilde{Y}$ with $\widetilde{X}_v$ is either empty, or a geodesic line.*

Proof. By Theorem 5.1, 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 [2], 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 $\widetilde{Y}_v$. $\qquad\square$

If S is a subset of a CAT(0) cube complex $\widetilde{Y}$, then let $\mathrm{hull}(S)$ denote the
combinatorial convex hull of S. The combinatorial convex hull of S is the minimal
convex subcomplex containing S. Equivalently, $\mathrm{hull}(S)$ is the intersection of all
closed halfspaces containing S.

Definition 5.3. Let X be a tubular space and let Y be a nonpositively curved cube
complex. A map $f \colon X \to Y$ is an *amicable immersion* if

(1) $f_* \colon \pi_1 X \to \pi_1 Y$ is an isomorphism;

(2) the G-equivariant map $\tilde{f} \colon \widetilde{X} \to \widetilde{Y}$ embeds each vertex space $\widetilde{X}_{\widetilde{v}}$ in $\widetilde{Y}$;

(3) each $\widetilde{X}_{\widetilde{v}}$ has a Euclidean metric such that if $H \subseteq \widetilde{Y}$ is a hyperplane, then the
intersection $H \cap \widetilde{X}_{\widetilde{v}}$ is either the empty set, or a single geodesic line in $\widetilde{X}_{\widetilde{v}}$;

(4) each edge space $\widetilde{X}_{\tilde{e}}$ is embedded transverse to the hyperplanes;

(5) each $\widetilde{X}_{\tilde{e}}$ is contained in $\mathsf{hull}\big(\tilde{f}\big(\bigcup_{\tilde{v}\in V\widetilde{\Gamma}}\widetilde{X}_{\tilde{v}}\big)\big)$.

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

Lemma 5.4. *Let X be a tubular space and let Y be a nonpositively 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 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 mapped into $\widetilde{Y}$ transverse to the hyperplanes in $\widetilde{Y}$ and 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 5.5. *Let $f:X \to Y$ be an amicable immersion, where Y is finite dimensional. If $\tilde{v}$ is a vertex in $\widetilde{\Gamma}$, then $\mathsf{hull}(\tilde{f}(\widetilde{X}_{\tilde{v}}))$ embeds as a subcomplex of $\mathbf{R}^d$ for some d.*

Proof. Let $G = \pi_1 X$. Let y be a 0-cube in $\mathsf{hull}(\tilde{f}(\widetilde{X}_{\tilde{v}}))$. 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 $\tilde{f}(\widetilde{X}_{\tilde{v}})$, then $y[H]$ is the halfspace containing $\widetilde{X}_{\tilde{v}}$, and therefore $y[H]$ is fixed for all 0-cubes y in $\mathsf{hull}(\tilde{f}(\widetilde{X}_{\tilde{v}}))$.

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 $\{g_i^r H_i \cap \widetilde{E}\}_{r\in\mathbb{Z}}$ 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 $\phi:\mathsf{hull}(\tilde{f}(\widetilde{X}_{\tilde{v}})) \to \mathbf{R}^d$ by letting $\phi(y) = (y_1, \ldots, y_d)$ for each 0-cube y.

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

Let $\tilde{f}\colon \widetilde{X} \to \widetilde{Y}$ be the lift to the universal cover of an amicable immersion $f\colon X \to Y$. 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 5.6. *Let $f\colon X \to Y$ 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 intersects $\widetilde{X}_{\tilde{e}}$ transversely, $H \cap \widetilde{X}_{\tilde{e}}$ contains a curve γ with basepoint p that is the point of intersection $H \cap \widetilde{X}_{\tilde{e}}$. As H is 2-sided and only intersects $\widetilde{X}_{-\tilde{e}} \cap \widetilde{X}_{\tilde{e}}$ at p, if γ is a closed arc in $\widetilde{X}_{\tilde{e}}$, then γ must have an endpoint in $\widetilde{X}_{+\tilde{e}} \cap \widetilde{X}_{\tilde{e}}$. In which case H intersects $\widetilde{X}_{+\tilde{e}}$ non-parallel to $\widetilde{X}_{\tilde{e}}$. Otherwise, if γ is not a closed arc, then γ is a ray going to infinity in $\widetilde{X}_{\tilde{e}}$. As the $G_{\tilde{e}}$-translates of γ are all distinct, this would imply that a point in $\widetilde{X}_{+\tilde{e}}$ and a point in $\widetilde{X}_{-\tilde{e}}$ are separated by infinitely many hyperplanes in $\widetilde{Y}$, which is impossible. $\qquad\square$

Lemma 5.7. *Let $f\colon X \to Y$ be an amicable immersion, where Y is a finite dimensional, locally finite, nonpositively 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 $\widetilde{Y} = \mathrm{hull}(\tilde{f}(\widetilde{X}))$. 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 $H \in \mathcal{H}_{\tilde{v}}$ intersects $\widetilde{X}_{\tilde{v}}$ non-parallel to an incident edge space $X_{\tilde{e}}$, it is non-parallel to all edge spaces incident to $\widetilde{X}_{\tilde{v}}$. By Lemma 5.6 H must then intersect each vertex space adjacent to $\widetilde{X}_{\tilde{v}}$ non-parallel to their own incident edge spaces. Repeating this argument we can conclude that H intersects every vertex and edge space in $\widetilde{X}$. Therefore, the hyperplanes in $\mathcal{H}$ can be partitioned into two sets: the set $\mathcal{H}_1$ of all hyperplanes that intersect every vertex and edge space in $\widetilde{X}$, and the set $\mathcal{H}_2$ of hyperplanes that only intersect vertex spaces parallel to the incident edge spaces. Note that every hyperplane in $\mathcal{H}_1$ intersects every hyperplane in $\mathcal{H}_2$, so $\widetilde{Y} = C(\widetilde{Y}, \mathcal{H}_1) \times C(\widetilde{Y}, \mathcal{H}_2)$ (see [4, Lem 2.5]). The action of G

preserves the partition $\mathcal{H} = \mathcal{H}_1 \sqcup \mathcal{H}_2$ and decomposes as a product of actions on each $C(\widetilde{Y}, \mathcal{H}_i)$ respectively.

Suppose that there exists a vertex $\tilde{v}$ in $\widetilde{\Gamma}$ such that no hyperplane in $\mathcal{H}_{\tilde{v}}$ is parallel to the incident edge spaces. Then $\mathcal{H}_{\tilde{v}} = \mathcal{H}_1$ so $\text{hull}(\tilde{f}(\widetilde{X}_{\tilde{v}})) \cong C(\widetilde{Y}, \mathcal{H}_1) \times C$ where C is some n-cube. Therefore $C(\widetilde{Y}, \mathcal{H}_1)$ will embed in $\mathbf{R}^d$ by Lemma 5.5. Note that G_v fixes a 0-cube c_1 in $C(\widetilde{Y}, \mathcal{H}_2)$ determined by orienting all the hyperplanes in $\mathcal{H}_2$ towards $\widetilde{X}_{\tilde{v}}$.

Let $\tilde{u}$ be a vertex in the same G-orbit as $\tilde{v}$. Then $\mathcal{H}_{\tilde{u}} = \mathcal{H}_1$ and the same arguments apply so there exists a 0-cube c_2 in $C(\widetilde{Y}, \mathcal{H}_2)$ fixed by $G_{\tilde{u}}$.

Let $g_1 \in G_{\tilde{v}}$ and $g_2 \in G_{\tilde{u}}$ be elements such that $\langle g_1, g_2 \rangle$ is a rank two free subgroup. As $C(\widetilde{Y}, \mathcal{H}_2)$ is locally finite there exists n, m such that g_1^n and g_2^m fix both c_1 and c_2. Thus $\langle g_1^n, g_2^m \rangle$ is a rank 2 free group that acts freely on $C(\widetilde{Y}, \mathcal{H}_1)$, which is impossible since $C(\widetilde{Y}, \mathcal{H}_1)$ embeds in $\mathbf{R}^d$. $\qquad\square$

Lemma 5.7 is a special case of the following more general statement.

Corollary 5.8. *Let $f \colon X \to Y$ be an amicable immersion, where Y is a finite dimensional, locally finite, nonpositively 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 $\widetilde{X}'_{\tilde{e}} = \widetilde{X}_{\tilde{e}}$ and $\widetilde{X}'_{\pm\tilde{e}} = \widetilde{X}_{\pm\tilde{e}}$. Therefore, by Lemma 5.7, there is a hyperplane intersecting $\widetilde{X}_{-\tilde{e}}$ parallel to $\widetilde{X}_{\tilde{e}}$, and similarly for $\widetilde{X}_{+\tilde{e}}$. $\qquad\square$

The following proposition is a strengthening of one direction of Theorem 1.1 in [12]. Let $f_1 \colon A \to C$ and $f_2 \colon B \to C$ 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 $p_1 \colon A \otimes_C B \to A$ and $p_2 \colon A \otimes_C B \to B$.

Proposition 5.9. *Let G be a tubular group acting freely on a $\mathrm{CAT}(0)$ cube complex $\widetilde{Y}$. There is a tubular space X with a finite set of immersed walls such that, 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.*

Proof. Let y be a 0-cube in $\widetilde{Y}$. By possibly replacing $\widetilde{Y}$ with $\text{hull}(Gy)$ 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 5.4 we can find an amicable immersion $f \colon X \to Y$. 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 $\tilde{f}(\widetilde{X})$, 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 $h \to Y$. Each component of $X \otimes_Y h_i$ 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 5.6 each hyperplane $H \subseteq \widetilde{Y}$ that intersects a vertex space $\widetilde{X}_{\tilde{v}}$ 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 geodesics into vertex spaces, and arcs that map into edges spaces X_e with an endpoint in each of 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 $\{\Lambda_i^p\}_{i=1}^k$ that map into X_e. As G acts freely on $\widetilde{Y}$, there must be hyperplanes intersecting each vertex space $\widetilde{X}_{\tilde{v}}$ 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 have images in X_v that coincide with each other or the image of an attaching map into X_v. Suppose that $\alpha_1, \ldots \alpha_m \in S_v$ is 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. For $\epsilon > 0$ there is a homotopy of $\bigsqcup_{i=1}^k \Lambda_i^p \to X$ that is the identity outside of $\bigsqcup_{i=1}^k \Lambda_i^p \cap \mathcal{N}_\epsilon(\alpha_1)$ such that $\alpha_1, \ldots, \alpha_m$ are homotoped to a disjoint set of geodesic curves in $X_v \cap \mathcal{N}_\epsilon(\alpha_1)$ transverse or disjoint from all the other curves in S_v^p and the attaching maps into X_v. By choosing $\epsilon < \frac{1}{3}$ small enough we can perform such a homotopy $\Phi \colon \bigsqcup_{i=1}^k \Lambda_i^p \times [0, 1] \to X$ such that all

sets of overlapping curves in $\{S_v^p\}_{v \in V\Gamma}$ are homotoped to a disjoint set of geodesic curves that are either disjoint or transverse to the attaching maps, 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 $\widetilde{\Lambda}^p$.

To prove (2), suppose that $C(\widetilde{X}, \mathcal{W})$ were infinite dimensional, then by Proposition 2.2, 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 fact.

Claim. *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 Lemma 2.2 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 $\widetilde{\Lambda}_1^p, \ldots, \widetilde{\Lambda}_n^p, \ldots$ 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 [7, Lemma 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 $\mathrm{d}_{\widetilde{Y}}(x, H) \geq N$ and $\mathrm{d}_{\widetilde{Y}}(H, H_i) \geq 1$. 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 5.8 to each edge group in G to deduce that $\{\Lambda_i\}_{i=1}^k$ are fortified. Therefore, by Proposition 4.7 we deduce that $C(\widetilde{X}, \mathcal{W})$ is locally finite.   □

We can now prove the equivalence of (1) and (3) in Theorem 1.1.

Theorem 5.10. *A tubular group G acts freely on a locally finite* CAT(0) *cube complex if and only if G is virtually special.*

Proof. As explained in the introduction, G being virtually special immediately implies that G acts on a locally finite CAT(0) cube complex.

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 5.9 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 2.1 and Proposition 4.7 this is equivalent to saying that the immersed walls are non-dilated and fortified. By Lemma 3.1 we can assume that the immersed walls are also primitive. Therefore, by Proposition 4.8, G is virtually special.   □

6. 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 6.1. *Let X be a tubular space and $G = \pi_1 X$. Suppose that the natural maps $H_1(G_v) \to H_1(G)$ 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 $a_{1v}, \ldots, a_{nv} \in G_v$ with $n > 1$, the corresponding set of immersed walls yields an $(n+1)$-dimensional CAT(0) cube complex with a free G-action.

Proof. By our hypothesis, there are inclusion maps $\iota_v \colon G_v \to \mathsf{H}_1^v(G)$, and projection maps $p_v \colon \mathsf{H}_1^v(G) \to G_v$ for every vertex $v \in V\Gamma$. Let g_e generate the edge group G_e and recall that $\varphi_e^{\pm} \colon X_e^{\pm} \to X_{\pm e}$ 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 $a = \iota_v(a_v) \in \mathsf{H}_1^v(G)$.

Suppose that $\mathsf{H}_1^v(G) \cong \mathbb{Z}^2$, then ι_v and p_v are isomorphisms for all $v \in V\Gamma$. Let $S_u = \{a_u = p_u(a)\}$. Note that $\{S_u\}_{u \in V\Gamma}$ is not an equitable set for G as each 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 $A_e = p_{+e} \circ \iota_{-e} \colon G_{-e} \to G_{+e}$ 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 $A \in \mathrm{GL}_2(\mathbb{Z})$. As the elements of $\mathrm{GL}_2(\mathbb{Z})$ have determinant ± 1, the geometric intersection number $\#[\cdot, \cdot]$ 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 $\mathsf{H}_1^v(G)$ as ι_v is an isomorphism, and thus $a_u = p_u(a)$ is also primitive since p_u is an isomorphism.

If $\mathsf{H}_1^v(G) \cong \mathbb{Z}^d$ with $d > 2$, there exist vertex groups G_u and G_v that embed into $\mathsf{H}_1^v(G)$ 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 $\mathsf{H}_1^v(G)$, and that g_v is disjoint from the image of G_u in $\mathsf{H}_1^v(G)$. 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 $G' = \pi_1 X'$ has $\mathsf{H}_1^v(G') \cong \mathbb{Z}^{d-1}$ and $H_1(G_v')$ embeds as a summand in $H_1(G')$. By repeating this process we can obtain a tubular group $\widehat{G}$ that is the fundamental group of a tubular space $\widehat{X}$ such that $\mathsf{H}_1^v(\widehat{G}) \cong \mathbb{Z}^2$, the homology $H_1(\widehat{G}_v)$ embeds as a summand in $H_1(\widehat{G})$, and $X \subseteq \widehat{X}$ is a sub-graph of spaces.

Therefore, by applying the case $d = 2$, for $a_v \in \hat{G}_v = G_v$ we can find a set of non-dilated, primitive embedded immersed walls $\{\hat{\Lambda}_i^a\}$ 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 $a_{1v}, \ldots, a_{nv} \in G_v$ be incommensurate primitive elements with $d \geq 2$, then $\langle a_{1v}, \ldots, a_{nv} \rangle$ 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. Applying the first part of the Proposition to each a_{iv} yields a collection of 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 $(\tilde{X}, \mathcal{W})$ can contain at most $(n + 1)$ walls: a vertical wall, and a single wall corresponding to each element in $\{a_{1v}, \ldots, a_{nv}\}$. Therefore $C(\tilde{X}, \mathcal{W})$ is of dimension at most $(n + 1)$, and since it was obtained from an equitable set, G acts freely on $C(\tilde{X}, \mathcal{W})$. $\qquad\square$

Lemma 6.2. *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 $G' \leq G$ 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 $\tilde{\Gamma}$ be the Bass-Serre tree. By Proposition 4.6, since there are immersed walls that are primitive and non-dilated, G acts freely on $\mathbf{R}^d \times \tilde{\Gamma}$ such that $G_{\tilde{v}}$ fixes the vertex $\tilde{v}$ in $\tilde{\Gamma}$. Therefore G is a subgroup of $\mathrm{Aut}(\mathbf{R}^d \times \tilde{\Gamma}) \cong \mathbb{Z}^d \rtimes \mathrm{Aut}([-1, 1]^d) \times \mathrm{Aut}(\tilde{\Gamma})$. Let K be the kernel of the quotient map $G \to \mathrm{Aut}([-1, 1]^d)$. Note that $K \leq G$ is a finite index subgroup that embeds in $\mathbb{Z}^d \times \mathrm{Aut}(\tilde{\Gamma})$. Let $p \colon K \to \mathbb{Z}^d$ 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 $A_v \leq \mathbb{Z}^d$ such that $p(K_v) \cap A_v$ is a summand of A_v. Let $A = \bigcap_v A_v$ and $G' = p^{-1}(A)$. Note that $G' \leq K \leq G$ 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 $\mathsf{H}_1^{\mathsf{v}}(G')$, so each vertex group survives as a retract in $\mathsf{H}_1(G')$. Each vertex group in G' therefore survives as a summand in the first homology. $\qquad\square$

Corollary 6.3. *If a tubular group G acts freely on a finite dimensional* CAT(0) *cube complex, then G virtually acts freely on a* CAT(0) *cube complex of dimension* 3.

Proof. If G acts freely on a finite dimensional CAT(0) cube complex, then by Proposition 5.9, Theorem 2.1 and Proposition 3.1 there is an equitable set that produces primitive, non-dilated immersed walls. Thus, by Lemma 6.2, 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'. By making a choice of two elements generating some vertex group G_v and applying Proposition 6.1 we obtain a free action of G' on a CAT(0) cube complex of dimension 3. $\square$

The following Corollary completes the proof of Theorem 1.1

Corollary 6.4. *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 5.9, Theorem 2.1 and Proposition 3.1 there is an equitable set that produces primitive, non-dilated immersed walls. By Lemma 6.2 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 6.1, 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 $(\varphi_e^{\pm})_*(g_e)$. If the circles in the immersed walls $\{\Lambda_e^{\pm}\}_{e \in E\Gamma}$ don't constitute an equitable set, use Proposition 6.1, 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 2.1, Proposition 4.8, and Theorem 1.1 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] N. Brady and M. R. Bridson, There is only one gap in the isoperimetric spectrum. *Geom. Funct. Anal.* **10** (2000), no. 5, 1053–1070. Zbl 0971.20019 MR 1800063

[2] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften, 319. Springer-Verlag, Berlin, 1999. Zbl 0988.53001 MR 1744486

[3] J. O. Button, Tubular free by cyclic groups and the strongest Tits alternative. Preprint 2015. arXiv:1510.05842 [math.GR]

[4] P.-E. Caprace and M. Sageev, Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.* **21** (2011), no. 4, 851–891. Zbl 1266.20054 MR 2827012

[5] C. H. Cashen, Quasi-isometries between tubular groups. *Groups Geom. Dyn.* **4** (2010), no. 3, 473–516. Zbl 1243.20052 MR 2653972

[6] F. Haglund and F. Paulin, Simplicité de groupes d'automorphismes d'espaces à courbure négative. In I. Rivin, C. Rourke, and C. Series (eds.), *The Epstein birthday schrift.* Geometry & Topology Monographs, 1. Geometry & Topology Publications, Coventry, 1998, 181–248. Zbl 0916.51019 MR 1668359

[7] F. Haglund and D. T. Wise, Special cube complexes. *Geom. Funct. Anal.* **17** (2008), no. 5, 1551–1620. Zbl 1155.53025 MR 2377497

[8] C. Hruska and D. Wise, Finiteness properties of cubulated groups. *Compos. Math.* **150** (2014), no. 3, 453–506. Zbl 1335.20043 MR 3187627

[9] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3) **71** (1995), no. 3, 585–617. Zbl 0861.20041 MR 1347406

[10] D. T. Wise, Research announcement: the structure of groups with a quasiconvex hierarchy. *Electron. Res. Announc. Math. Sci.* **16** (2009), 44–55. Zbl 1183.20043 MR 2558631

[11] D. T. Wise, *From riches to raags*: *3-manifolds, right-angled Artin groups, and cubical geometry*. CBMS Regional Conference Series in Mathematics, 117. American Mathematical Society, Providence, R.I., 2012. Zbl 1278.20055 MR 2986461

[12] D. T. Wise, Cubular tubular groups. *Trans. Amer. Math. Soc.* **366** (2014), no. 10, 5503–5521. Zbl 1346.20058 MR 3240932

[13] D. Woodhouse, Classifying finite dimensional cubulations of tubular groups. *Michigan Math. J.* **65** (2016), no. 3, 511–532. Zbl 06645291 MR 3542763

[14] D. Woodhouse, A generalized axis theorem for cube complexes. *Algebr. Geom. Topol.* **17** (2017), no. 5, 2737–2751. Zbl 06791381 MR 3704240

Daniel James Woodhouse, Mathematics Department, Technion,
Israel Institute of Technology, Haifa, 32000, Israel

e-mail: daniel.woodhouse@mail.mcgill.ca