msp

Algebraic & Geometric Topology 17 (2017) 2737–2751

A generalized axis theorem for cube complexes

DANIEL J WOODHOUSE

We consider a finitely generated virtually abelian group G acting properly and without inversions on a CAT(0) cube complex X. We prove that G stabilizes a finite-dimensional CAT(0) subcomplex $Y \subseteq X$ that is isometrically embedded in the combinatorial metric. Moreover, we show that Y is a product of finitely many quasilines. The result represents a higher-dimensional generalization of Haglund's axis theorem.

20F65

1 Introduction

A CAT(0) *cube complex* X is a cell complex that satisfies two properties: it is a geodesic metric space satisfying the CAT(0) comparison triangle condition, and each n–cell is isometric to $[0,1]^n$. We will call this metric the CAT(0) *metric* d_X and refer to Bridson and Haefliger [2] for a comprehensive account. A *hyperplane* $\Lambda \subseteq X$ is the subset of points equidistant between two adjacent vertices. Despite the brevity of this definition, hyperplanes are better understood via their combinatorial definition, and the reader is urged to consult the literature; see Sageev [10], Haglund [6] and Wise [12] for the required background. There also exists an alternative metric on the 0–cubes of X, which we will refer to as the *combinatorial metric* d_X^c, sometimes referred to as the ℓ^1*–metric*. The combinatorial distance between two 0–cubes is the length of the shortest combinatorial path in X joining the 0–cubes. Equivalently, the combinatorial distance between two 0–cubes is the number of hyperplanes in X separating them. 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.

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 paper will be the following theorem:

Published: 19 September 2017 DOI: 10.2140/agt.2017.17.2737

Theorem 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 will not in general be a convex subcomplex.

Corollary 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.*

Corollary 1.2 *Let G be a finitely generated group acting properly on a* CAT(0) *cube complex X. Then virtually $\mathbb{Z}^n$ subgroups are undistorted in G.*

Let g be an isometry of X, and let $x \in X$. The *displacement of g at x*, denoted by $\tau_x(g)$, is the distance $\mathsf{d}_X(x, gx)$. The *translation length* of g, denoted by $\tau(g)$, 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 by $\tau_x^c(g)$, as $\mathsf{d}_X^c(x, gx)$ and the *combinatorial translation length*, denoted by $\tau^c(g)$, 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 of Bridson and Haefliger [2] provides a G–invariant, convex, flat $\mathbb{E}^n \subseteq X$. A group acting on a CAT(0) cube complex does not, in general, have to do so semisimply. See Algom-Kfir, Wajnryb and Witowicz [1] for examples of nonsemisimple isometries in Thompson's group F acting on an infinite-dimensional CAT(0) cube complex. Alternatively, in Gersten [5], a free-by-cyclic group G is shown not to permit a semisimple action on a CAT(0) space. Yet in Wise [13] it is shown that G does act freely on a CAT(0) cube complex. Thus Theorem 4.3 can be applied to such actions, whereas the classical flat torus theorem cannot.

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; see Wise and Woodhouse [14]. 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 4.3, we can deduce the following:

Corollary 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 4.3 and Corollary 4.4 was to resolve the following question, posed by Wise. Although we have not found a combinatorial flat, Corollary 4.4 is perhaps better suited to applications (see Woodhouse [15]).

Problem 1.3 Let $\mathbb{Z}^2$ act freely on a CAT(0) cube complex Y. Does there exists 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 in the combinatorial metric, subcomplex $\gamma \subseteq X$ with $\gamma \cong \mathbb{R}$. Note that γ realizes the minimal combinatorial translation length of g. Theorem 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 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 4.3, and the argument is inspired by the solvable subgroup theorem of Bridson and Haefliger [2]. Note that since we do not require that the action of G on a CAT(0) cube complex be semisimple the following is not covered by the solvable subgroup theorem.

Corollary 1.4 *Let H be virtually $\mathbb{Z}^n$, and let $\phi\colon H \to H$ 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 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 bounded and unbounded. □

However, we have the following example of a solvable group which does act freely on a CAT(0) cube complex.

Example 1.5 Let $H = \langle a_1, a_2, \dots \mid [a_i, a_j] \text{ for } i \neq j\rangle$. Note that H is the fundamental group of the nonpositively 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 $\phi\colon H \to H$ be the monomorphism such that $\phi(a_i) = a_{i+1}$. Let $G = H*_\phi = \langle t, a_1, a_2, \dots \mid [a_i, a_j] \text{ for } i \neq j,\ t^{-1}a_it = a_{i+1}\rangle$ 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.

Acknowledgements I would like to thank Daniel T Wise, Mark F Hagen, Jack Button, Piotr Przytycki and Dan Guralnik.

2 Dual cube complexes

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 $\overleftarrow{\Lambda}$ and $\overrightarrow{\Lambda}$ 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 $\Lambda|_K = \{\overleftarrow{\Lambda} \cap K, \overrightarrow{\Lambda} \cap K\}$.

In this paper, duplicate walls are not permitted in $\mathcal{W}$. Let $\mathcal{H}$ be the set of all halfspaces corresponding to the walls in $\mathcal{W}$.

Example 2.1 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 $L(\Lambda)$ and $R(\Lambda)$ denote the maximal subcomplexes contained in $\overleftarrow{\Lambda}$ and $\overrightarrow{\Lambda}$, respectively. Note that $L(\Lambda)$ and $R(\Lambda)$ are convex subcomplexes. 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 $c\colon \mathcal{W} \to \mathcal{H}$ 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 by letting 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 [9].

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*.

Lemma 2.2 *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 the set of 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 and Λ_2 intersect Z, and since Z is connected, there are 1–cubes e_1 and e_2 in Z that are dual to the hyperplanes corresponding to Λ_1 and Λ_2. Therefore, both 0–cubes in e_1 belong in a single halfspace of $\Lambda_2|_Z$, so $\Lambda_1|_Z \neq \Lambda_2|_Z$.

We construct a map $\phi\colon C(Z, \mathcal{V}) \to C(X, \mathcal{W})$ 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 $\phi(c)[\Lambda] \supseteq c[\Lambda|_Z]$ for $\Lambda|_Z \in \mathcal{V}$, 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 $\Lambda_1|_Z \cap \Lambda_2|_Z \subseteq X$. 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 $z \in c[\Lambda|_Z]$ for all $\Lambda|_Z \in \mathcal{V} - \{\Lambda_1|_Z, \dots, \Lambda_k|_Z\}$, hence $z \in \phi(c)[\Lambda]$ for all $\Lambda \in \mathcal{W}_Z - \{\Lambda_1, \dots, \Lambda_k\}$. Let $\{\Lambda_{k+1}, \dots, \Lambda_{k+\ell}\}$ be the set of walls in $\mathcal{W}$ separating x and z. Then $x \in \phi(c)[\Lambda]$ for all $\Lambda \in \mathcal{W} - \{\Lambda_1, \dots \Lambda_{k+\ell}\}$.

The 0–cubes are embedded since if $c_1 \neq c_2$, there exists $\Lambda|_Z \in \mathcal{V}$ such that $c_1[\Lambda|_Z] \neq c_2[\Lambda|_Z]$, hence $\phi(c_1)[\Lambda] \neq \phi(c_2)[\Lambda]$. If c_1 and c_2 are adjacent 0–cubes in $C(Z, \mathcal{V})$, then $c_1[\Lambda|_Z] = c_2[\Lambda|_Z]$ for all $\Lambda|_Z \in \mathcal{V}$, with the exception of precisely one wall $\widehat{\Lambda}|_Z$. Therefore, we can deduce that $\phi(c_1)[\Lambda] = \phi(c_2)[\Lambda]$ for all walls in $\mathcal{W}$, with the precise exception of $\widehat{\Lambda}$. 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 and 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. □

Given a wall Λ associated to a hyperplane in X we let $N(\Lambda)$ denote the *carrier* of Λ, by which we mean the union of all cubes intersected by Λ.

The following lemma describes what is called the *restriction quotient* in [3].

Lemma 2.3 *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 ϕ: $C(S, \mathcal{W})^0 \to C(S, \mathcal{V})^0$. 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 $\phi(c_1)[\Lambda_1] \cap \phi(c_2)[\Lambda_2] \neq \varnothing$ for all $\Lambda_1, \Lambda_2 \in \mathcal{V}$, since $c_1[\Lambda_1] \cap c_2[\Lambda_2] \neq \varnothing$ 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 nonempty 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 $\vec{\Lambda} \supseteq z[\Lambda']$ for some $\Lambda' \in \mathcal{V}$ let $x[\Lambda] = \vec{\Lambda}$. 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 $x[\Lambda_1] \cap x[\Lambda_2] \neq \varnothing$ for $\Lambda_1, \Lambda_2 \in \mathcal{W}$. If $\Lambda_1, \Lambda_2 \in \mathcal{V}$ then $x[\Lambda_1] \cap x[\Lambda_2] = z[\Lambda_1] \cap z[\Lambda_2] \neq \varnothing$. 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 $x[\Lambda_1] \cap x[\Lambda_2] \supseteq z[\Lambda_1'] \cap z[\Lambda_2] \neq \varnothing$. If $\Lambda_2 \in \mathcal{W} - \mathcal{V}$ and $x[\Lambda_2] \subseteq z[\Lambda_2']$ for some $\Lambda_2' \in \mathcal{V}$ then $x[\Lambda_1] \cap x[\Lambda_2] \subseteq z[\Lambda_1'] \cap z[\Lambda_2'] \neq \varnothing$. If Λ_2 intersects $z[\Lambda]$ for all $\Lambda \in \mathcal{V}$, then $x[\Lambda_1] \cap x[\Lambda_2] \supseteq z[\Lambda_1'] \cap x[\Lambda_2] \neq \varnothing$. Finally if both $s \in x[\Lambda_1]$ and $s \in x[\Lambda_2]$, 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 $\{\Lambda_1, \Lambda_2, \dots\} \subseteq \mathcal{W}$ 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 Minimal $\mathbb{Z}^n$–invariant convex subcomplexes

The following is Theorem 2 from [4]. As this paper is written in Russian, we give a proof in an appendix based on the work in [8] as well as stating the definition of codimension-1.

Theorem 3.1 (Gerasimov [4]) *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 3.2 *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 $(X_o, \mathcal{W})$ be the wallspace obtained from the hyperplanes in X_o. 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 2.3 there is an G–invariant map $\phi\colon X_o^0 \to C(X_o, G\Lambda)^0$. Since $\mathrm{Stab}_G(\Lambda)$ is not commensurable to a codimension-1 subgroup, Theorem 3.1 implies that there is a fixed 0–cube x in $C(X_o, G\Lambda)$. Lemma 2.3 then implies that $\phi^{-1}(x)$ is nonempty. 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. □

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

Corollary 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.*

4 Proof of the main theorem

Definition 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 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–neighborhood of the other for some $r > 0$.

Theorem 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 there is a minimal, nonempty, convex subcomplex $X_o \subseteq X$ stabilized by G, containing finitely many hyperplane orbits, and $\mathrm{Stab}_G(\Lambda)$ is a rank $n-1$ subgroup of G for each hyperplane $\Lambda \subseteq X_o$.

Let $x \in X_o$ be a 0–cube. Let Υ be the Cayley graph of G with respect to a finite generating set S. Let $N = \max\{\mathrm{d}_X^{\mathrm{c}}(x, gx) \mid g \in S\}$. Let $\phi\colon \Upsilon \to X_o$ be a G–equivariant map that maps the vertex corresponding to 1_G to x, and edges to geodesic combinatorial paths in X_o. Note that the image $\phi(e)$ of each edge e in Υ has length at most N and is intersected at most once by each hyperplane. 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 2.2 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 $q\colon Q \to \mathbb{R}^n$ be a G–equivariant quasiisometry. Note that $\mathrm{Stab}_G(\Lambda)$ is a quasiisometrically embedded codimension-1 subgroup of G, for all $\Lambda \in \mathcal{W}_Q$. We claim that $q(\Lambda \cap Q)$ is coarsely equivalent to a codimension-1 affine subspace $H \subseteq \mathbb{R}^n$.

As G is virtually $\mathbb{Z}^n$ and $\mathrm{Stab}_G(\Lambda)$ is a codimension-1, there exists $g \in S$ such that $\langle g\rangle$ is not virtually contained in $\mathrm{Stab}_G(\Lambda)$. There are finitely many $\mathrm{Stab}_G(\Lambda)$–orbits of vertices in $\Upsilon/\langle g\rangle$, so let $A = \{a_0, \dots a_k\}$ be representatives in Υ such that Λ separates $\phi(a_i)$ and $\phi(ga_i)$. Let γ_i be the biinfinite geodesic in Υ containing $\langle g\rangle a_i$. Then $\Lambda \cap \phi(\gamma_i)$ is contained in the $N(N+1)$ neighborhood of $\phi(a_i)$ in $\phi(\gamma_i)$, since otherwise Λ would intersect a pair of 1–cubes in $\phi(\gamma_i)$ that lie in the same $\langle g\rangle$–orbit, implying that $\langle g\rangle$ is virtually contained in $\mathrm{Stab}_G(\Lambda)$. Thus, $\Lambda \cap h\phi(\gamma_i)$ is contained in the $N(N+1)$ neighborhood of $h\phi(a_i)$ in $h\phi(\gamma_i)$ for all $h \in \mathrm{Stab}_G(\Lambda)$.

Now suppose that Λ intersects Q outside of the $N(N+2)$ neighborhood of $\mathrm{Stab}_G(\Lambda)A$. Then Λ must intersect $\phi(e)$, where e is an edge connecting $h_1\gamma_i$ and $h_2\gamma_j$ for some $h_1, h_2 \in \mathrm{Stab}_G(\Lambda)$. Up to taking the inverse of g, we can assume that ge is further away from h_1a_i and h_2a_j than e. Then Λ must intersect $\phi(ge)$ since Λ is 2–sided, intersects $\phi(e)$ precisely once, and cannot intersect the intervals in $h_1\gamma_i$ and $h_2\gamma_j$ that lie between e and ge. Similarly, Λ intersects $\phi(g^ne)$ for all $n > 0$ implying that Λ intersects a pair of 1–cubes in the same $\langle g\rangle$–orbit, further implying that $\langle g\rangle$ is virtually contained in $\mathrm{Stab}_G(\Lambda)$ and contradicting the initial assumption on g. Thus, Λ cannot intersect Q outside of the $N(N+2)$ neighborhood of $\mathrm{Stab}_G(\Lambda)A$. 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 $\mathrm{Stab}_G(\Lambda_i)$. 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 $1 \le i \le m$. 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 $\mathrm{Stab}_G(\Lambda)$ are commensurable for all $\Lambda \in \mathcal{W}_i$.

Therefore, there is a cyclic subgroup Z_i that is not virtually contained in any $\mathrm{Stab}_G(\Lambda)$ 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 quasiequivalent 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 $\mathbb{R}$. □

We can now prove Corollary 4.4.

Corollary 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 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 [2] 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 $\mathbb{R}^{n-1}$. □

5 Haglund's axis

The goal of this section is to reprove the following result of Haglund as a consequence of Corollary 4.4.

Theorem 5.1 (Haglund [6]) *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 3.1 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 4.4, there is a line $L \subset X$ stabilized by G, 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 a 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 $\Lambda_1, \dots, \Lambda_k$ 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. □

Appendix: Codimension-1 subgroups

Definition A.1 Let G be a finitely generated group. Let Υ denote the Cayley graph of G with respect to some finite generating set. A subgroup $H \leqslant G$ is *codimension-*1 if K/Υ has more than one end.

Let $\oplus$ denote the operation of symmetric difference. A subset $A \subseteq G$ is *H–finite* if $A \subseteq HF$ where F is some finite subset of G. We will use the following equivalent formulation (see [11]) of codimension-1: A subgroup $H \leq G$ is a codimension-1 subgroup if there exists some $A \subseteq G$ such that:

(1) $A = HA$.

(2) A is *H–almost invariant*, that is to say that $A \oplus Ag$ is H–finite for any $g \in G$.

(3) A is *H–proper*, that is to say that neither A nor $G - A$ is H–finite.

We will reprove the following theorem from [4] using techniques from [8].

Theorem A.2 *Let G be a finitely generated group acting on a* CAT(0) *cube complex X without edge inversions or fixing a 0–cube. Then the stabilizer of some hyperplane in X is a codimension-*1 *subgroup of G.*

Proof Suppose that no hyperplane stabilizer is a codimension-1 subgroup of G. We will find a 0–cube fixed by G.

Let $\mathcal{H}$ denote the set of hyperplanes in X. We can assume that X has finitely many G–orbits of hyperplanes after possibly passing to the convex hull of a single 0–cube orbit in X. If x and y are 0–cubes in X, then let $\Delta(x, y) \subseteq \mathcal{H}$ denote the set of hyperplanes separating x and y. Note that

$$\mathsf{d}_X^c(x, y) = |\Delta(x, y)|.$$

Let $\Lambda_1, \dots, \Lambda_n$ be a minimal set of representatives of the orbits of hyperplanes. Let x_0 be some fixed choice of 0–cube in X. Let

$$H_i = \mathrm{Stab}_G(\Lambda_i) \quad \text{and} \quad A_i = \{g \in G \mid gx_0 \in \overleftarrow{\Lambda}_i\}.$$

We can verify that A_i satisfies the first two criteria in Definition A.1.

(1) It is immediate that $A_i = H_i A_i$, as G doesn't invert the hyperplanes in X.

(2) Let *xor* denote the exclusive or. For $f \in G$ we can deduce that $A_i \oplus A_i f$ is H_i–finite:

$$\begin{aligned} g \in A_i \oplus A_i f \quad &\iff \quad gx_0 \in \overleftarrow{\Lambda}_i \text{ xor } gf^{-1}x_0 \in \overleftarrow{\Lambda}_i \\ &\iff \quad x_0 \in g^{-1}\overleftarrow{\Lambda}_i \text{ xor } f^{-1}x_0 \in g^{-1}\overleftarrow{\Lambda}_i \\ &\iff \quad g \in G \text{ is such that } g^{-1}\Lambda_i \text{ separates } x_0 \text{ and } f^{-1}x_0. \end{aligned}$$

As $(X, \mathcal{H})$ is a wallspace, there are only finitely many $g \in G$ such that $g^{-1}\Lambda_i$ separates x_0 and $f^{-1}x_0$. If $g_1\Lambda_i, \dots, g_k\Lambda_i$ are the translates then

$$A_i \oplus A_i f = \{g_1, \dots, g_k\} H_i,$$

which implies almost H_i–invariance.

Therefore, A_i cannot be H_i–proper for any i, as we have assumed that none of the H_i are codimension-1. This means that either A_i or $G - A_i$ is H_i–finite. After possibly reversing the orientation of Λ_i we can assume that A_i is H_i–finite, so $A_i \subseteq H_i F_i$ where $F_i \subseteq G$ is finite.

Claim $\mathsf{d}_X(x_0, fx_0) < 2\max_i(|F_i|)$ *for all* $f \in G$.

Proof

$$\begin{aligned} g\Lambda_i \in \Delta(x_0, fx_0) \quad &\iff \quad x_0[g\Lambda_i] \neq fx_0[g\Lambda_i] \\ &\iff \quad x_0 \in g\overleftarrow{\Lambda}_i \text{ xor } fx_0 \in g\overleftarrow{\Lambda}_i \\ &\iff \quad g^{-1}x_0 \in \overleftarrow{\Lambda}_i \text{ xor } g^{-1}fx_0 \in \overleftarrow{\Lambda}_i \\ &\iff \quad g^{-1} \in A_i \text{ xor } g^{-1} \in Af^{-1} \\ &\iff \quad g^{-1} \in A_i \oplus A_i f^{-1}. \end{aligned}$$

As the final set is covered by $2|F_i|$ translates of H_i, we can deduce that there are at most $2|F_i|$ hyperplanes in $\Delta(x_0, fx_0)$. □

Thus, we can conclude that the G–orbit of x_0 is a bounded set. If G has a finite orbit in X, then the convex hull of the orbit is a compact, finite-dimensional, complete CAT(0) cube complex, and we can apply Corollary II.2.8(1) from [2] to find a fixed point p. If p is a 0–cube then we are done. Otherwise, p is in the interior of some n–cube that is fixed by G, and since G doesn't invert hyperplanes we can deduce that G fixes a 0–cube in that cube. If the G–orbits in X are infinite, then their convex hull may not be complete, so the above argument will not hold. Instead, we will follow the strategy of [8] and embed the cube complex into a Hilbert space.

Let $\mathcal{C}(\mathcal{H})$ denote the *connected cube*, a graph with vertices given by functions $c\colon \mathcal{H} \to \{0,1\}$ with finite support, and edges that join a pair of distinct vertices if and only if they differ on precisely one hyperplane.

Fix a 0–cube x_0. Then there is an embedding

$$\phi\colon X^1 \hookrightarrow \mathcal{C}(\mathcal{H})$$

that maps the 0–cube x to c_x, where

$$c_x(\Lambda) = \begin{cases} 1 & \text{if } x[\Lambda] \neq x_0[\Lambda], \\ 0 & \text{if } x[\Lambda] = x_0[\Lambda]. \end{cases}$$

A hyperplane $\Lambda \in \mathcal{H}$ *separates* two vertices c_1 and c_2 in $\mathcal{C}(\mathcal{H})$ if $c_1(\Lambda) \neq c_2(\Lambda)$. Note that Λ separates 0–cubes x, y in X if and only if it separates $\phi(x)$ and $\phi(y)$. Therefore, we can define $\Delta(c_1, c_2)$ for vertices in $\mathcal{C}(\mathcal{H})$ and conclude that if x, y are 0–cubes in X then $\Delta(x, y) = \Delta(\phi(x), \phi(y))$. This implies that ϕ is an isometric embedding in the combinatorial metric.

We will show that a bounded orbit in X implies there is a fixed 0–cube in $\mathcal{C}(\mathcal{H})$ and then argue that we can go one step further and find a fixed 0–cube in X.

Let $\ell^2(\mathcal{H})$ be the Hilbert space of square summable functions $s\colon \mathcal{H} \to \mathbb{R}$. There is an embedding $\rho\colon \mathcal{C}(\mathcal{H}) \to \ell^2(\mathcal{H})$ given by

$$\rho(c)(\Lambda) = c[\Lambda].$$

It is straightforward to verify that $\|\rho(c_1) - \rho(c_2)\|^2 = \mathrm{d}_{\mathcal{C}(\mathcal{H})}(c_1, c_2)$. There is a G–action on $\ell^2(\mathcal{H})$ such that if $s \in \ell^2(\mathcal{H})$, $\Lambda \in \mathcal{H}$ and $g \in G$, then

$$gs(\Lambda) = \begin{cases} s(g^{-1}\Lambda) & \text{if } c_{x_0}(g^{-1}\Lambda) = c_{x_0}(\Lambda), \\ 1 - s(g^{-1}\Lambda) & \text{if } c_{x_0}(g^{-1}\Lambda) \neq c_{x_0}(\Lambda). \end{cases}$$

It is again straightforward to verify that this action is by isometries, and that ρ is G–equivariant.

As Gx_0 is bounded, so is $G(\rho \circ \phi(x_0))$. It then follows that G has a fixed point in $\ell^2(\mathcal{H})$ (a proof is in [8], which also cites Lemma 3.8 in [7]). Let $s\colon \mathcal{H} \to \mathbb{R}$ be the fixed point. For all $g \in G$ we can deduce that $s(g\Lambda)$ is either $s(\Lambda)$ or $1 - s(\Lambda)$. Therefore s can only take two values on the hyperplanes in a single G–orbit. As s has to be square summable the two values have to be 0 and 1, and s can only take the value 1 on finitely many hyperplanes. Thus, s is the image of a point c in $\mathcal{C}(\mathcal{S})$.

Let $c \in \mathcal{C}(\mathcal{S})$ be a G–invariant vertex which minimizes the distance to the image of X^1 in $\mathcal{C}(\mathcal{S})$. Let Z be a G–orbit of 0–cubes in X such that $\phi(Z)$ realizes the minimal distance from c.

Let $\mathcal{V}$ be the set of hyperplanes that intersect $\{c\} \cup \mathcal{V}$. Every hyperplane in $\mathcal{V}$ must intersect Z, otherwise if $\mathcal{F} \subseteq \mathcal{V}$ is the finite, G–invariant subset of hyperplanes separating c from Z then we can define a 0–cube c' such that

$$c'(\Lambda) = \begin{cases} c(\Lambda) & \text{if } \Lambda \notin \mathcal{F}, \\ 1 - c(\Lambda') & \text{if } \Lambda \in \mathcal{F}, \end{cases}$$

and deduce that c' is G–invariant and is $|\mathcal{F}|$ closer to Z than c.

Let $z_0, z_1, z_2, \ldots$ be an enumeration of 0–cubes in Z. Each hyperplane separating z_0 and z_1 must lie in either $\Delta(z_0, c)$ or $\Delta(z_1, c)$. As z_0 is minimal distance in X from c, the edges in X incident to z_0 must be dual to hyperplanes not in $\Delta(z_0, c)$, and instead belongs to $\Delta(z_1, c)$. Therefore, the hyperplane $\Lambda_0 \in \mathcal{V}$ dual to the first edge in a combinatorial geodesic joining z_0 to z_1 must lie in $\Delta(z_1, c)$. Similarly, there exists a hyperplane Λ_1 dual to the first edge of the combinatorial geodesic in X joining z_1 to z_2 that belongs to $\Delta(z_2, c)$ but not $\Delta(z_1, c)$. Note that Λ_1 cannot intersect Λ_0 in X, otherwise Λ_0 would be dual to an edge incident to z_1, which would imply that there exists a 0–cube in X adjacent to z_1 that is closer to c. Therefore Λ_0, Λ_1 separates z_0 from z_2 in X. Iterating this argument produces a sequence of disjoint hyperplanes $\Lambda_0, \Lambda_2, \Lambda_3, \ldots$ such that $\Lambda_0, \ldots, \Lambda_k$ separates z_0 from z_{k+1} in X. This contradicts the hypothesis that Z is a bounded set in X. □

References

[1] **Y Algom-Kfir**, **B Wajnryb**, **P Witowicz**, *A parabolic action on a proper,* CAT(0) *cube complex*, J. Group Theory 16 (2013) 965–984 MR

[2] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999) MR

[3] **P-E Caprace**, **M Sageev**, *Rank rigidity for* CAT(0) *cube complexes*, Geom. Funct. Anal. 21 (2011) 851–891 MR

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

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

[6] **F Haglund**, *Isometries of* CAT(0) *cube complexes are semi-simple*, preprint (2007) arXiv

[7] **P de la Harpe**, **A Valette**, *La propriété* (T) *de Kazhdan pour les groupes localement compacts*, Astérisque 175, Soc. Math. France, Paris (1989) MR

[8] **G A Niblo**, **M A Roller**, *Groups acting on cubes and Kazhdan's property* (T), Proc. Amer. Math. Soc. 126 (1998) 693–699 MR

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

[10] **M Sageev**, *Codimension-*1 *subgroups and splittings of groups*, J. Algebra 189 (1997) 377–389 MR

[11] **P Scott**, *Ends of pairs of groups*, J. Pure Appl. Algebra 11 (1977) 179–198 MR

[12] **D T Wise**, *From riches to raags:* 3*–manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics 117, Amer. Math. Soc., Providence, RI (2012) MR

[13] **D T Wise**, *Cubular tubular groups*, Trans. Amer. Math. Soc. 366 (2014) 5503–5521 MR

[14] **D T Wise**, **D J Woodhouse**, *A cubical flat torus theorem and the bounded packing property*, Israel J. Math. 217 (2017) 263–281 MR

[15] **D J Woodhouse**, *Classifying virtually special tubular groups*, preprint (2016) arXiv To appear in Groups Geom. Dyn.

Mathematics Department, Technion – Israel Institute of Technology
Haifa, Israel

woodhouse.da@technion.ac.il

Received: 3 February 2016 Revised: 9 May 2017