

Automorphisms of Contact Graphs of CAT(0) Cube Complexes

Elia Fioravanti*

Max Planck Institute for Mathematics, Vivatsgasse 7, 53111, Bonn,
Germany

*Correspondence to be sent to: e-mail: fioravanti@mpim-bonn.mpg.de

Abstract

We show that, under weak assumptions, the automorphism group of a CAT(0) cube complex X coincides with the automorphism group of Hagen's contact graph $\mathcal{C}(X)$. The result holds, in particular, for universal covers of Salvetti complexes, where it provides an analogue of Ivanov's theorem on curve graphs of non-sporadic surfaces. This highlights a contrast between contact graphs and Kim–Koberda extension graphs, which have much larger automorphism group. We also study contact graphs associated with Davis complexes of right-angled Coxeter groups. We show that these contact graphs are less well behaved and describe exactly when they have more automorphisms than the universal cover of the Davis complex.

1 Introduction

The curve graph associated with a finite-type surface is a fundamental object in the study of mapping class groups [17, 18]. It has arguably been at the heart of the monumental developments in our understanding of mapping class groups and Kleinian groups that were initiated by the work of Masur and Minsky [31, 32] and have ultimately led to the solution of Thurston's ending lamination conjecture [1, 30].

Received February 6, 2020; Revised September 15, 2020; Accepted September 24, 2020

An important property of curve graphs is that they are *rigid*: Ivanov showed that every automorphism of the curve graph of a non-sporadic surface is induced by an element of the extended mapping class group [25, 26]. This was a central ingredient first in the computation of the abstract commensurator [26], and later in the proof that mapping class groups are quasi-isometrically rigid [8, 16]. Due to the classical result of Tits that the automorphism group of a building coincides with the associated algebraic group [34], Ivanov's theorem also strengthens the analogy between mapping class groups and arithmetic groups.

It was recently shown by Behrstock, Hagen and Sisto [4, 5] that the Masur–Minsky machinery can be applied to a much vaster class of spaces, which they name *hierarchically hyperbolic spaces*. This has proved a fruitful approach to a number of problems [3, 12, 13], notably allowing these authors to obtain a particularly strong rigidity result for quasi-flats [6].

It is natural to wonder if analogues of Ivanov's theorem hold for other hierarchically hyperbolic spaces. In this note, we address this question for CAT(0) cube complexes.

Many CAT(0) cube complexes were shown to be hierarchically hyperbolic in [4, 21], with Hagen's *contact graph* playing the role of a curve complex. In particular, all virtually special groups [24] are hierarchically hyperbolic, and these include all right-angled Artin and Coxeter groups. In the case of right-angled Artin groups, a parallel between curve graphs and *extension graphs* (a relative of contact graphs) had already been observed by Kim and Koberda [27, 28].

For a general CAT(0) cube complex X , there are actually at least three (closely related) graphs that can claim some analogy with curve graphs:

- (i) Hagen's *contact graph* $\mathcal{C}(X)$ [15] mentioned above. Vertices are hyperplanes of X and edges connect pairs of hyperplanes that are not separated by a 3rd.
- (ii) The *crossing graph* $\mathcal{C}_\hbar(X)$. This is the subgraph of $\mathcal{C}(X)$ with full vertex set and edges only joining pairs of transverse hyperplanes.
- (iii) The *reduced crossing graph* $\mathcal{C}_r(X)$. This is the simplicial graph obtained by identifying vertices of $\mathcal{C}_\hbar(X)$ with the same link (see Definition 1 for a more precise description).

When X_Γ is the universal cover of a Salvetti complex, and if no two vertices of the graph Γ have the same link, then $\mathcal{C}_r(X_\Gamma)$ coincides with the *extension graph* Γ^e introduced by Kim and Koberda [27].

If X is uniformly locally finite and has no cut vertices, then these three graphs are quasi-isometric to each other (see e.g., [14, Appendix A]). By [15, Theorem 4.1], they are quasi-trees, hence, in particular, they are δ -hyperbolic. When, in addition, X is hierarchically hyperbolic, a number of further analogies with curve graphs and mapping class groups is discussed in [4], including acylindricity of actions, existence of hierarchy paths, and a Masur–Minsky-style distance formula. We also refer the reader to [28] for other analogies in the case of right-angled Artin groups.

If $\mathcal{G}(X)$ denotes any of the three graphs above, we have a natural homomorphism $\text{Aut } X \rightarrow \text{Aut } \mathcal{G}(X)$. The closest we can get to an analogue of Ivanov’s theorem is if one of these homomorphisms is an isomorphism. It is important to remark that the group $\text{Aut } X$ will often be *uncountable* and, in fact, this is essentially always the case for universal covers of Salvetti complexes (see Remark 1 below).

It can be deduced from the work of Huang [22] that the images of the homomorphisms $\text{Aut } X \rightarrow \text{Aut } \mathcal{C}_\cap(X)$ and $\text{Aut } X \rightarrow \text{Aut } \mathcal{C}_r(X)$ have uncountable index for most universal covers of Salvetti complexes (see Remark 2 below). This dashes any hopes that these two maps be isomorphisms, even for relatively harmless cube complexes like Salvetti complexes.

As we are about to see, things are a lot better behaved for the contact graph $\mathcal{C}(X)$. We only draw the reader’s attention to Example 7.1 in [4], which shows that even the homomorphism $\iota: \text{Aut } X \rightarrow \text{Aut } \mathcal{C}(X)$ cannot be an isomorphism in complete generality.

We say that a vertex $v \in X^{(0)}$ is *extremal* if its link is a cone (cf. [2, Definition 2.2]). In other words, v lies in the carrier of some hyperplane \mathfrak{w} that is transverse to all other hyperplanes containing v in their carrier. We stress that any vertex that belongs to a single edge of X is extremal.

Every hyperplane $\mathfrak{w} \subseteq X$ inherits a structure of CAT(0) cube complex from X , where cubes of \mathfrak{w} are intersections with cubes of X . In particular, vertices of \mathfrak{w} are in one-to-one correspondence with edges of X crossing \mathfrak{w} . It thus makes sense to speak of *extremal vertices of \mathfrak{w}* .

Our main result is the following (its two parts will be proved in Corollary 1 and Theorem 3.1, respectively). We denote by $S(X^{(0)})$ the group of permutations of the vertex set of X .

Theorem. Let X be a uniformly locally finite CAT(0) cube complex with no extremal vertices. Let $\iota: \text{Aut } X \rightarrow \text{Aut } \mathcal{C}(X)$ be as above.

1. There exists a homomorphism $\rho: \text{Aut } \mathcal{C}(X) \rightarrow S(X^{(0)})$ such that $\rho \circ \iota = \text{id}_{\text{Aut } X}$. In particular, the homomorphism ι is injective.

Suppose in addition that no hyperplane of X has extremal vertices. Then:

2 ι is a group isomorphism with inverse ρ .

The main idea in the proof of the Theorem 1 is that, when X has no extremal vertices, vertices of X are in one-to-one correspondence with maximal cliques in $\mathcal{C}(X)$. If, in addition, the hyperplanes of X have no extremal vertices, edges of X correspond to pairs of maximal cliques of $\mathcal{C}(X)$ with the largest possible intersection.

Cube complexes with no free faces never have extremal vertices (see e.g., [2, Remark 2.5]), and neither do their hyperplanes. Hence:

Corollary. Let X be a uniformly locally finite CAT(0) cube complex with no free faces. Then $\text{Aut } X \cong \text{Aut } \mathcal{C}(X)$ via the map ι .

The Corollary applies in particular to all universal covers of Salvetti complexes associated with right-angled Artin groups.

It would be nice to use this result to study abstract commensurators of right-angled Artin groups, or to expand the known results on their quasi-isometry classification [7, 10, 22, 29] and their quasi-isometric rigidity properties [9, 19, 23]. Unfortunately, it appears that these applications would rather require the *extension graph*, which, as discussed above, is highly non-rigid.

For instance, Huang showed that every quasi-isometry between right-angled Artin groups with finite outer automorphism groups induces an isomorphism of their extension graphs [22, Lemma 4.5]. In the same context, quasi-isometries can fail to induce isomorphisms of contact graphs (see e.g., Remark 3).

As pointed out by one of the anonymous referees, it is even true that every right-angled Artin group G with $\#\text{Out}(G) < +\infty$ admits a proper cobounded quasi-action of some group H that fails to induce an H -action by automorphisms on the contact graph associated with the Salvetti complex of G (see the proof of [19, Theorem 6.10] along with our Theorem 1). It remains unclear to me whether there can be such examples where this issue cannot be resolved by passing to a finite-index subgroup of H .

In relation to this, the following question was suggested by the same anonymous referee (several special cases have been studied in [23]):

Question 1. Let a group H be quasi-isometric to a right-angled Artin group G with $\#\text{Out}(G) < +\infty$. Does a finite-index subgroup of H act properly and cocompactly on the universal cover of the Salvetti complex of G ?

Finally, it is reasonable to wonder whether it really is necessary to require that hyperplanes of X have no extremal vertices in the Theorem 1. Davis complexes provide a nice class of counterexamples when we drop this hypothesis.

Recall that, for a graph Γ and a vertex $a \in \Gamma^{(0)}$, the *star* $\text{st } a \subseteq \Gamma^{(0)}$ is the set of vertices that are either equal to a or joined to a by an edge. The following will be proved in Proposition 1.

Proposition 1. Let W_Γ be a right-angled Coxeter group with no finite direct factors. Let Y_Γ denote the universal cover of its Davis complex. Then:

1. Y_Γ has no extremal vertices, so $\iota: \text{Aut } Y_\Gamma \rightarrow \text{Aut } \mathcal{C}(Y_\Gamma)$ is injective.
2. Y_Γ has a hyperplane with extremal vertices if and only if there exist distinct vertices $a, b \in \Gamma^{(0)}$ with $\text{st } a \subseteq \text{st } b$. In this case, the subgroup $\iota(\text{Aut } Y_\Gamma) < \text{Aut } \mathcal{C}(Y_\Gamma)$ has infinite index.
3. The homomorphism $\rho: \text{Aut } \mathcal{C}(Y_\Gamma) \rightarrow S(Y_\Gamma^{(0)})$ is injective if and only if there do not exist vertices $a, b \in \Gamma^{(0)}$ with $\text{st } a = \text{st } b$. When ρ is not injective, its kernel is an uncountable torsion subgroup.

2 Preliminaries and Examples

All proofs are elementary, but we assume a certain familiarity with the basics of CAT(0) cube complexes. See for instance [33, 35] for an introduction.

Given a graph \mathcal{G} , we say that two vertices are *adjacent* if they belong to a common edge. Given $x \in \mathcal{G}^{(0)}$, the *link* $\text{lk } x \subseteq \mathcal{G}^{(0)}$ is the subset of vertices adjacent to x . The *star* of x is the set $\text{st } x = \text{lk } x \cup \{x\}$.

An *automorphism* of \mathcal{G} is a self-bijection of $\mathcal{G}^{(0)}$ that preserves adjacency of vertices. We denote the automorphism group of \mathcal{G} by $\text{Aut } \mathcal{G}$. If every vertex in \mathcal{G} has finite degree, $\text{Aut } \mathcal{G}$ is a 2nd-countable, locally compact topological group with the compact-open topology,

If \mathcal{G}_1 and \mathcal{G}_2 are graphs, their *join* $\mathcal{G}_1 * \mathcal{G}_2$ is obtained by taking the disjoint union $\mathcal{G}_1 \sqcup \mathcal{G}_2$ and joining by an edge each vertex of \mathcal{G}_1 to each vertex of \mathcal{G}_2 . We say that the graph \mathcal{G} is a *cone* if it is the join of some other graph \mathcal{G}' and a singleton.

Let X be a CAT(0) cube complex. The group of (*cubical*) *automorphisms* of X coincides the automorphism group of the 1-skeleton $X^{(1)}$; we denote it by $\text{Aut } X$. A path $\gamma \subseteq X$ is a *combinatorial geodesic* if it is contained in $X^{(1)}$ and it is a geodesic for the graph metric on $X^{(1)}$.

We denote by $\mathscr{W}(X)$ and $\mathscr{H}(X)$, respectively, the set of all hyperplanes and all halfspaces of X . Given a hyperplane $\mathfrak{w} \in \mathscr{W}(X)$, we refer to the two halfspaces $\mathfrak{h}, \mathfrak{h}^*$ bounded by \mathfrak{w} as its *sides*. Given subsets $A, B \subseteq X$, we denote by $\mathscr{W}(A|B)$ the set of hyperplanes \mathfrak{w} such that A and B are contained in opposite sides of \mathfrak{w} .

Given $\mathfrak{w} \in \mathscr{W}(X)$, the union of all cubes of X that intersect \mathfrak{w} forms a subcomplex $C(\mathfrak{w}) \subseteq X$ known as the *carrier* of \mathfrak{w} . We say that \mathfrak{w} is *adjacent* to a vertex $v \in X^{(0)}$ if \mathfrak{w} intersects an edge of X incident to v (equivalently, if v belongs to the carrier of \mathfrak{w}). We denote by $\mathscr{W}_v \subseteq \mathscr{W}(X)$ the set of hyperplanes adjacent to the vertex $v \in X^{(0)}$.

We say that a hyperplane \mathfrak{u} *contacts* another hyperplane \mathfrak{w} if their carriers intersect. Equivalent conditions are that the set $\mathscr{W}(\mathfrak{u}|\mathfrak{w})$ is empty, or that there exists a vertex $v \in X^{(0)}$ such that $\{\mathfrak{u}, \mathfrak{w}\} \subseteq \mathscr{W}_v$.

For each $v \in X^{(0)}$, we redefine the link $\text{lk } v$ as follows.¹ The graph $\text{lk } v$ has a vertex for every edge of X incident to v , and an edge joining two vertices of $\text{lk } v$ if and only if the corresponding edges of X span a square. Equivalently, the vertex set of $\text{lk } v$ is \mathscr{W}_v , and we join two hyperplanes by an edge when they are transverse. We say that v is *extremal* if $\text{lk } v$ is a cone.

The intersection between a hyperplane $\mathfrak{w} \subseteq X$ and a cube $c \subseteq X$ is always either empty or a mid-cube in c . It follows that \mathfrak{w} inherits a decomposition into cubes $\mathfrak{w} \cap c$, where c ranges through all cubes in the carrier $C(\mathfrak{w})$. This gives \mathfrak{w} a structure of CAT(0) cube complex. We are thus allowed to speak of “vertices of \mathfrak{w} ” (which are midpoints of edges of X crossing \mathfrak{w}) and of their link in the CAT(0) cube complex \mathfrak{w} .

A subcomplex $C \subseteq X$ is said to be *convex* if every combinatorial geodesic joining two vertices of C is entirely contained in C . The carrier of every hyperplane is a convex subcomplex, and so is the full subcomplex spanned by the vertices contained in any given halfspace. If C_1, \dots, C_k are pairwise-intersecting convex subcomplexes of X , we have $C_1 \cap \dots \cap C_k \neq \emptyset$. This fact normally goes by the name of Helly’s lemma.

Before we go any further, let us give a less ambiguous definition of the *reduced crossing graph* $\mathcal{C}_r(X)$, which was sketched in the introduction.

Definition 1. Vertices of $\mathcal{C}_r(X)$ are maximal subsets $\mathcal{W} \subseteq \mathscr{W}(X)$ with the property that, for any two $\mathfrak{u}, \mathfrak{v} \in \mathcal{W}$, each hyperplane of X is transverse to \mathfrak{u} if and only if it is transverse to \mathfrak{v} . Vertices of $\mathcal{C}_r(X)$ corresponding to subsets $\mathcal{U}, \mathcal{V} \subseteq \mathscr{W}(X)$ are joined by an edge of $\mathcal{C}_r(X)$ if and only if there exist transverse hyperplanes $\mathfrak{u} \in \mathcal{U}$ and $\mathfrak{v} \in \mathcal{V}$ (in which case every hyperplane in \mathcal{U} is transverse to every hyperplane in \mathcal{V}).

We now obtain Remarks 1, 2, and 3, which were promised in the introduction. Throughout this discussion, we denote by X_Γ the universal cover of the Salvetti complex associated with a right-angled Artin group A_Γ .

Recall that the 1-skeleton of X_Γ coincides with the standard Cayley graph of A_Γ . It follows that each edge of X_Γ is labelled by a standard generator of A_Γ , that is, a vertex of the graph Γ . Opposite edges in a square of X_Γ always have the same label, hence all edges crossing a given hyperplane $\mathfrak{w} \in \mathscr{W}(X_\Gamma)$ have the same label. We will refer to this element of $\Gamma^{(0)}$ as the *label* of \mathfrak{w} .

Remark 1. It was pointed out to me by Nir Lazarovich that the group $\text{Aut } X_\Gamma$ is always uncountable, except when $X_\Gamma \cong \mathbb{R}^n$. The argument is essentially the one in Theorem 5.12 of [20], but I briefly recall it here.

Assume that the graph Γ is not complete and pick vertices $x, y \in \Gamma^{(0)}$ that are not joined by an edge. Let ϕ be the automorphism of the group A_Γ that fixes each standard generator except for the one corresponding to y , which is taken to its inverse. Identifying A_Γ with $X_\Gamma^{(0)}$, it is clear that ϕ determines a cubical automorphism of X_Γ , which we also denote by ϕ .

Let $v \in X_\Gamma^{(0)}$ be the vertex corresponding to the identity of A_Γ . There are two hyperplanes adjacent to v that are labelled by $x \in \Gamma^{(0)}$. Let \mathfrak{w} be one of them and let $\mathfrak{h}, \mathfrak{h}^*$ denote its two sides. Observe that ϕ fixes the carrier $C(\mathfrak{w})$ pointwise and therefore leaves invariant \mathfrak{h} and \mathfrak{h}^* .

Let the map $\psi: X_\Gamma \rightarrow X_\Gamma$ be defined as the identity on $\mathfrak{h}^* \cup C(\mathfrak{w})$, and as ϕ on $\mathfrak{h} \cup C(\mathfrak{w})$. It follows from the previous observation that ψ is well defined. Since every edge of X_Γ is contained in either $\mathfrak{h}^* \cup C(\mathfrak{w})$ or $\mathfrak{h} \cup C(\mathfrak{w})$, this is an automorphism of X_Γ . It is clear that $\psi \neq \text{id}_{X_\Gamma}$.

Now, given any finite set F of vertices of X_Γ , there exists $g \in A_\Gamma$ with $gF \subseteq \mathfrak{h}^*$. Hence the automorphism $g^{-1}\psi g \in \text{Aut } X_\Gamma - \{\text{id}_{X_\Gamma}\}$ fixes F pointwise. We conclude that the locally compact group $\text{Aut } X_\Gamma$ is non-discrete.

The fact that $\text{Aut } X_\Gamma$ is uncountable follows from Baire's theorem (see e.g., [11, Remark 2.A.18]).

We say that a combinatorial geodesic $\gamma \subseteq X_\Gamma$ is *standard* if all edges of γ have the same label. Two bi-infinite standard geodesics are at finite Hausdorff distance if and only if they cross the same hyperplanes; in this case, we say that they are *parallel*. Given that $X_\Gamma^{(1)}$ coincides with the usual Cayley graph of A_Γ , we will also speak of standard geodesics in A_Γ .

The *extension graph* Γ^e [22, 27, 28] has a vertex for every parallelism class of standard geodesics in X_Γ ; the vertices determined by standard geodesics γ_1 and γ_2 are joined by an edge of Γ^e if and only if the hyperplanes crossed by γ_1 are transverse to the hyperplanes crossed by γ_2 .

Note that, in general, we do not have a homomorphism $\text{Aut } X_\Gamma \rightarrow \text{Aut } \Gamma^e$, but only a homomorphism $A_\Gamma \rightarrow \text{Aut } \Gamma^e$.

Let d_w denote the usual word metric on A_Γ , which coincides with the graph metric on $X_\Gamma^{(1)}$ under the identification $A_\Gamma = X_\Gamma^{(0)}$. Let d_r be the *syllable metric* on A_Γ , as defined for example, in [28, Section 5.2] and [22, Section 4.3]. More precisely, d_r is the largest metric on A_Γ satisfying $d_r(x, y) = 1$ for all distinct $x, y \in A_\Gamma$ that are joined by a standard geodesic.

In order to make Remark 2 below, we will need the following lemma, which can in large part be deduced from the work of Huang [22].

Lemma 1. Let Γ be a finite graph.

1. There is a natural homomorphism $\text{Isom}(A_\Gamma, d_r) \rightarrow \text{Aut } \Gamma^e$. This is injective if and only if Γ is not a cone.²
2. If no two vertices of Γ have the same link, every isometry of (A_Γ, d_w) is an isometry of (A_Γ, d_r) . Moreover, $\Gamma^e = \mathcal{C}_r(X_\Gamma)$ in this case.

Proof. We begin with part (2). Consider a vertex $v \in X_\Gamma^{(0)}$, a vertex $x \in \Gamma^{(0)}$, and the two vertices $x^\pm \in (\text{lk } v)^{(0)}$ determined by x . If no two vertices of Γ have the same link, no vertex of $\text{lk } v - \{x^\pm\}$ can have the same link as x^+ and x^- . It follows that every element of $\text{Aut } X_\Gamma$ takes standard geodesics to standard geodesics. Identifying $\text{Isom}(A_\Gamma, d_w)$ with $\text{Aut } X_\Gamma$, we deduce that every $\phi \in \text{Isom}(A_\Gamma, d_w)$ is 1-Lipschitz with respect to d_r . Since ϕ^{-1} must also be 1-Lipschitz for d_r , this shows that $\phi \in \text{Isom}(A_\Gamma, d_r)$.

Finally, it is easy to see that, since no two vertices of Γ have the same link, the projection $\mathcal{C}_\cap(X_\Gamma) \rightarrow \mathcal{C}_r(X_\Gamma)$ identifies two vertices if and only if there exists a standard geodesic crossing the corresponding hyperplanes. It follows that $\mathcal{C}_r(X_\Gamma)$ coincides with Γ^e in this case.

Regarding part (1), Huang showed that every element of $\text{Isom}(A_\Gamma, d_r)$ takes standard geodesics to standard geodesics, as unparametrized sets (see Remark 4.16 and the proof of Corollary 4.15 in [22]). Since vertices of Γ^e correspond to families of hyperplanes transverse to standard geodesics, every isometry of the syllable metric induces a permutation of the vertices of Γ^e . Huang also showed (*loc. cit.*) that such

permutations preserve adjacency of vertices of Γ^e . We thus obtain a homomorphism $\text{Isom}(A_\Gamma, d_r) \rightarrow \text{Aut } \Gamma^e$.

If Γ is the cone over a subgraph Δ , we have $A_\Gamma = A_\Delta \times \mathbb{Z}$. For any permutation $\sigma: \mathbb{Z} \rightarrow \mathbb{Z}$, the map $(g, n) \mapsto (g, \sigma(n))$ is an isometry of the syllable metric on $A_\Delta \times \mathbb{Z}$, but it maps to the identity in $\text{Aut } \Gamma^e$.

Conversely, let us show that, when Γ is not a cone, the homomorphism $\text{Isom}(A_\Gamma, d_r) \rightarrow \text{Aut } \Gamma^e$ is injective. In other words, given $\phi \in \text{Isom}(A_\Gamma, d_r)$ taking each standard geodesic to a standard geodesic at finite Hausdorff distance, we need to show that ϕ fixes each element of A_Γ .

Consider an element $g \in A_\Gamma$ and let $\{\gamma_x\}_{x \in \Gamma^{(0)}}$ be the collection of all standard geodesics containing g . Observe that $\bigcap_x \gamma_x = \{g\}$, that each $\phi(\gamma_x)$ is a standard geodesic at finite Hausdorff distance from γ_x , and that $\bigcap_x \phi(\gamma_x) = \{\phi(g)\}$. Let α be a combinatorial geodesic joining g and $\phi(g)$ in X_Γ . Since α joins a point of γ_x to a point of $\phi(\gamma_x)$, every edge crossed by α must be labelled by an element of $\text{st } x \subseteq \Gamma^{(0)}$. However, since Γ is not a cone, we have $\bigcap_{x \in \Gamma^{(0)}} \text{st } x = \emptyset$. In conclusion, α does not cross any edges, and we have $\phi(g) = g$ for all $g \in A_\Gamma$. ■

Remark 2. Suppose that Γ is not a cone and that no two vertices of Γ have the same link. We show here that, in this case, the images of the two natural maps $\iota_\text{th}: \text{Aut } X_\Gamma \rightarrow \text{Aut } \mathcal{C}_\text{th}(X_\Gamma)$ and $\iota_r: \text{Aut } X_\Gamma \rightarrow \text{Aut } \mathcal{C}_r(X_\Gamma)$ have uncountable index.

It follows from Lemma 1 that we have a commutative diagram:

$$\begin{array}{ccccc} \text{Isom}(A_\Gamma, d_w) & \hookrightarrow & \text{Isom}(A_\Gamma, d_r) & \hookrightarrow & \text{Aut } \Gamma^e \\ \parallel & & & & \parallel \\ \text{Aut } X_\Gamma & \xleftarrow{\iota_r} & & \xrightarrow{\iota_r} & \text{Aut } \mathcal{C}_r(X_\Gamma). \end{array}$$

Now, the argument in [22, Example 4.14] shows that the embedding $\text{Isom}(A_\Gamma, d_w) \hookrightarrow \text{Isom}(A_\Gamma, d_r)$ is very far from being surjective. More precisely, for every standard geodesic $\gamma \subseteq X_\Gamma$ and every permutation σ of its vertex set $\gamma^{(0)} \subseteq X_\Gamma^{(0)} = A_\Gamma$, we can construct an element of $\text{Isom}(A_\Gamma, d_r)$ that preserves the set $\gamma^{(0)}$, and acts on it as σ .

On closer inspection, this corresponds to a copy of the infinite symmetric group $S(\mathbb{Z}) < \text{Isom}(A_\Gamma, d_r)$ that intersects the subgroup $\text{Isom}(A_\Gamma, d_w)$ in an infinite dihedral subgroup. Thus, $\text{Isom}(A_\Gamma, d_w) < \text{Isom}(A_\Gamma, d_r)$ has uncountable index, and so does $\text{Aut } X_\Gamma < \text{Aut } \mathcal{C}_r(X_\Gamma)$.

Finally, observe that the quotient projection $\mathcal{C}_{\text{th}}(X_\Gamma) \twoheadrightarrow \mathcal{C}_r(X_\Gamma)$ induces a surjective homomorphism $\pi_r: \text{Aut } \mathcal{C}_{\text{th}}(X_\Gamma) \rightarrow \text{Aut } \mathcal{C}_r(X_\Gamma)$.³ This yields the commutative diagram:

$$\begin{array}{ccc}
 \text{Aut } X_\Gamma & \xleftarrow{\iota_{\text{th}}} & \text{Aut } \mathcal{C}_{\text{th}}(X_\Gamma) \\
 & \searrow \iota_r & \swarrow \pi_r \\
 & \text{Aut } \mathcal{C}_r(X_\Gamma) &
 \end{array}$$

We conclude that $\text{Aut } X_\Gamma < \text{Aut } \mathcal{C}_{\text{th}}(X_\Gamma)$ also has uncountable index.

Remark 3. There are many examples of quasi-isometries of X_Γ that fail to induce an automorphism of the contact graph $\mathcal{C}_\Gamma := \mathcal{C}(X_\Gamma)$.

A 1st example is provided by the argument in Remark 2 (also appearing in [9, Section 11] and [22, Example 4.14]). Under its assumptions, we have seen that every permutation of the vertex set of the standard geodesic γ gives rise to an isometry of (A_Γ, d_r) . Considering bi-Lipschitz permutations of $\gamma^{(0)}$, the resulting isometries of d_r are quasi-isometries of (A_Γ, d_w) and most of them will not induce automorphisms of \mathcal{C}_Γ .

For a 2nd example, suggested by one of the anonymous referees, let Λ be a pentagon. Consider the homomorphism $\tau: A_\Lambda \rightarrow \mathbb{Z}/2\mathbb{Z}$ with $\tau(v) = 1$ for some vertex $v \in \Lambda$ and $\tau(w) = 0$ for all other $w \in \Lambda$. Then $\ker \tau$ is isomorphic to $A_{\Lambda'}$, where Λ' is the graph obtained by glueing two copies of Λ along $st v$. Since $\ker \tau$ has index 2 in A_Λ , there is a natural proper cobounded quasi-action of A_Λ on $X_{\Lambda'}$. It is not hard to see that every hyperplane of $X_{\Lambda'}$ is taken within bounded distance of another hyperplane, so we obtain an action of A_Λ on $\mathcal{C}_{\Lambda'}^{(0)}$ by bijections. However, the bijection associated with $v \in A_\Lambda$ does not extend to an isomorphism of $\mathcal{C}_{\Lambda'}$.

Note that it is not clear if the quasi-isometries in the 1st example can take part in a proper cobounded quasi-action on X_Γ . In the 2nd example, we do have a geometric quasi-action, but $\#\text{Out}(A_{\Lambda'}) = +\infty$. Moreover, the restriction to $\ker \tau < A_\Lambda$ of the quasi-action on $A_{\Lambda'}$ does in fact induce an action by automorphisms on $\mathcal{C}_{\Lambda'}$. As mentioned in the Introduction, other examples are provided by [19, Theorem 6.10], although, again, all issues disappear in a finite-index subgroup.

3 Proof of the Theorem

Let X be a CAT(0) cube complex with contact graph $\mathcal{C} = \mathcal{C}(X)$, as defined in the introduction. We identify subsets of \mathcal{C} with their intersection with $\mathcal{C}^{(0)}$ and with the corresponding subset of $\mathcal{W}(X)$.

Our 1st goal is to establish a correspondence between vertices of X and maximal cliques in \mathcal{C} .

Lemma 2.

1. For every finite clique $C \subseteq \mathcal{C}$, there exists $v \in X^{(0)}$ with $C \subseteq \mathcal{W}_v$.
2. If $C \subseteq \mathcal{C}$ is a maximal finite clique, there is $v \in X^{(0)}$ with $C = \mathcal{W}_v$.
3. If X is uniformly locally finite, the cliques of \mathcal{C} are uniformly finite.

Proof. For every vertex $v \in X^{(0)}$, the subset $\mathcal{W}_v \subseteq \mathcal{C}$ is a clique. Parts (2) and (3) thus follow immediately from part (1), which we now prove.

Let $C \subseteq \mathcal{C}$ be a finite clique. For every hyperplane $w \in C$, at most one side of w can contain a hyperplane in C disjoint from w . Picking this side for every $w \in C$, or just any side if w is transverse to all other hyperplanes in C , we obtain a finite collection of pairwise-intersecting halfspaces $\mathcal{H} \subseteq \mathcal{H}(X)$. By Helly's lemma, there exists $w \in X^{(0)}$ lying in all elements of \mathcal{H} .

Let $d(w)$ denote the sum of the distances from w to the carriers of the hyperplanes in C , using the graph metric of $X^{(1)}$. Thus, $d(w) = 0$ if and only if $C \subseteq \mathcal{W}_w$. If $d(w) > 0$, there exist hyperplanes $u \in C - \mathcal{W}_w$ and $v \in \mathcal{W}_w \cap \mathcal{W}(w|u)$. Let $w' \in X^{(0)}$ be the vertex with $\mathcal{W}(w|w') = \{v\}$. No hyperplane in C can be contained in the side of v that contains w , or they would not contact u . It follows that $d(w') < d(w)$ and, iterating this procedure finitely many times, we obtain a vertex $v \in X^{(0)}$ with $d(v) = 0$. This yields part (1). ■

Remark 4. Locally finite CAT(0) cube complexes need not be ω -dimensional (i.e., they can contain infinite families of pairwise-transverse hyperplanes). In particular, there exist locally finite cube complexes whose contact, crossing, and reduced crossing graphs all contain infinite cliques. Thus, the hypothesis in part (3) of Lemma 2 cannot be weakened.

Lemma 3. Consider a vertex $v \in X^{(0)}$ with $\#\mathcal{W}_v < +\infty$.

1. There exists $w \neq v$ with $\mathcal{W}_v \subseteq \mathcal{W}_w$ if and only if $\text{lk } v$ is a cone.
2. In particular, if $\text{lk } v$ is not a cone, the clique $\mathcal{W}_v \subseteq \mathcal{C}$ is maximal and there does not exist another vertex w with $\mathcal{W}_v = \mathcal{W}_w$.

Proof. By part (1) of Lemma 2, \mathcal{W}_v is a maximal clique if and only if there does not exist $w \in X^{(0)}$ with $\mathcal{W}_v \subsetneq \mathcal{W}_w$. Part (2) thus follows from part (1).

Suppose that $w \neq v$ is a vertex with $\mathcal{W}_v \subseteq \mathcal{W}_w$. Let $\mathfrak{w} \in \mathcal{W}_v$ be a hyperplane separating w and v . Since \mathfrak{w} cannot separate w from any element of \mathcal{W}_v , it must be transverse to all elements of \mathcal{W}_v . Hence $\text{lk } v$ is a cone.

Conversely, if $\text{lk } v$ is a cone, there exists a hyperplane $\mathfrak{w} \in \mathcal{W}_v$ that is transverse to all other hyperplanes adjacent to v . Denoting by $w \in X^{(0)}$ the vertex with $\mathcal{W}(v|w) = \{\mathfrak{w}\}$, we have $\mathcal{W}_v \subseteq \mathcal{W}_w$. ■

Recall from the introduction that the action $\text{Aut } X \curvearrowright \mathcal{W}(X)$ results in a natural homomorphism $\iota : \text{Aut } X \rightarrow \text{Aut } \mathcal{C}$. Lemmas 2 and 3 immediately yield the 1st part of the Theorem 1:

Corollary 1. Let X be uniformly locally finite, with no extremal vertices. There exists a natural one-to-one correspondence between vertices of X and maximal cliques of \mathcal{C} . This induces a homomorphism $\rho : \text{Aut } \mathcal{C} \rightarrow S(X^{(0)})$ satisfying $\rho \circ \iota = \text{id}_{\text{Aut } X}$.

We now proceed to study when the homomorphism ρ is injective.

Definition 2. For a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$, we denote by $\mathcal{I}(\mathfrak{w})$ the intersection of all sets \mathcal{W}_v that contain \mathfrak{w} . Let moreover $\mathcal{I}^0(\mathfrak{w}) \subseteq \mathcal{I}(\mathfrak{w})$ be the subset of those hyperplanes $u \in \mathcal{I}(\mathfrak{w})$ for which $\mathfrak{w} \in \mathcal{I}(u)$.

Note that $\mathfrak{w} \in \mathcal{I}^0(\mathfrak{w}) \subseteq \mathcal{I}(\mathfrak{w})$. Recall that a vertex $v \in X^{(0)}$ lies in the carrier of \mathfrak{w} if and only if $\mathfrak{w} \in \mathcal{W}_v$. Thus, we have $u \in \mathcal{I}(\mathfrak{w})$ if and only if the carrier of \mathfrak{w} is contained in the carrier of u . In particular, $u \in \mathcal{I}^0(\mathfrak{w})$ if and only if u and \mathfrak{w} have the same carrier.

Remark 5. Let X be uniformly locally finite, with no extremal vertices. Since the subsets $\mathcal{W}_v \subseteq \mathcal{W}(X)$ are exactly the maximal cliques of \mathcal{C} , we have $\phi(\mathcal{I}(\mathfrak{w})) = \mathcal{I}(\phi(\mathfrak{w}))$ for all $\mathfrak{w} \in \mathcal{W}(X)$ and $\phi \in \text{Aut } \mathcal{C}$. Moreover:

$$\begin{aligned} u \in \mathcal{I}^0(\mathfrak{w}) &\Leftrightarrow u \in \mathcal{I}(\mathfrak{w}) \ \& \ \mathfrak{w} \in \mathcal{I}(u) \\ &\Leftrightarrow \phi(u) \in \phi(\mathcal{I}(\mathfrak{w})) \ \& \ \phi(\mathfrak{w}) \in \phi(\mathcal{I}(u)) \\ &\Leftrightarrow \phi(u) \in \mathcal{I}(\phi(\mathfrak{w})) \ \& \ \phi(\mathfrak{w}) \in \mathcal{I}(\phi(u)) \Leftrightarrow \phi(u) \in \mathcal{I}^0(\phi(\mathfrak{w})). \end{aligned}$$

Hence $\phi(\mathcal{I}^0(\mathfrak{w})) = \mathcal{I}^0(\phi(\mathfrak{w}))$ as well.

Lemma 4. Consider $\mathfrak{w}, u \in \mathcal{W}(X)$.

1. We have $u \in \mathcal{I}(\mathfrak{w})$ if and only if u is transverse to \mathfrak{w} and to all other hyperplanes transverse to \mathfrak{w} . In particular, $\#\mathcal{I}(\mathfrak{w}) \leq \dim X$.

2. If $u \in \mathcal{I}^0(\mathfrak{w})$, then u and \mathfrak{w} have exactly the same star in \mathcal{C} .
3. We have $u \in \mathcal{I}^0(\mathfrak{w})$ if and only if $\mathcal{I}^0(u) = \mathcal{I}^0(\mathfrak{w})$. In particular, the sets $\mathcal{I}^0(\mathfrak{w})$ provide a partition of $\mathcal{C}^{(0)}$.

Proof. We begin with part (1). If $u \in \mathcal{I}(\mathfrak{w})$, the carrier of \mathfrak{w} is contained in the carrier of u , and it is clear that u is transverse to all other hyperplanes that intersect the carrier of \mathfrak{w} (i.e., \mathfrak{w} and the hyperplanes transverse to \mathfrak{w} other than u). Conversely, suppose that $u \notin \mathcal{I}(\mathfrak{w})$. Then there exists $v \in X^{(0)}$ adjacent to \mathfrak{w} , but not to u . There exists $\mathfrak{v} \in \mathscr{W}_v$ separating u and v . If u is transverse to \mathfrak{w} , so must be \mathfrak{v} . In conclusion, either u is not transverse to \mathfrak{w} , or \mathfrak{v} is transverse to \mathfrak{w} and u is not transverse to \mathfrak{v} .

Finally, $\#\mathcal{I}(\mathfrak{w}) \leq \dim X$ follows from the observation that the elements of $\mathcal{I}(\mathfrak{w})$ are pairwise transverse. This completes the proof of part (1).

Recall that we have $u \in \mathcal{I}^0(\mathfrak{w})$ if and only if u and \mathfrak{w} have the same carrier. Part (3) is immediate from the fact that this is an equivalence relation. Part (2) follows from the additional observation that edges of \mathcal{C} join exactly those pairs of hyperplanes that have intersecting carriers. ■

Corollary 2. Let X be uniformly locally finite, with no extremal vertices. Let $N \leq S(\mathcal{C}^{(0)})$ be the subgroup leaving each subset $\mathcal{I}^0(\mathfrak{w}) \subseteq \mathcal{C}$ invariant.

1. N is contained in $\text{Aut } \mathcal{C}$.
2. N coincides with the kernel of $\rho: \text{Aut } \mathcal{C} \rightarrow S(X^{(0)})$.
3. N is isomorphic to the direct product of the permutation groups $S(\mathcal{I}^0(\mathfrak{w}))$. In particular, ρ is injective if and only if each $\mathcal{I}^0(\mathfrak{w})$ is a singleton.

Proof. Part (1) is immediate from part (2) of Lemma 4. The 1st half of part (3) is immediate from part (3) of Lemma 4. The 2nd half of part (3) will follow from part (2).

Let us then conclude by proving part (2). Recall that an element $\phi \in \text{Aut } \mathcal{C}$ lies in $\ker \rho$ if and only if ϕ leaves invariant each maximal clique of \mathcal{C} . Again by part (2) of Lemma 4, every maximal clique in \mathcal{C} is a union of sets of the form $\mathcal{I}^0(\mathfrak{w})$. Hence $N \leq \ker \rho$.

Conversely, consider $\phi \in \ker \rho$. Given that each $\mathcal{I}(\mathfrak{w})$ is an intersection of maximal cliques, and ϕ leaves every maximal clique invariant, we have $\phi(\mathcal{I}(\mathfrak{w})) = \mathcal{I}(\mathfrak{w})$ for all $\mathfrak{w} \in \mathscr{W}(X)$. Recalling that $\phi(\mathcal{I}(\mathfrak{w})) = \mathcal{I}(\phi(\mathfrak{w}))$ by Remark 5, we have $\mathcal{I}(\mathfrak{w}) = \mathcal{I}(\phi(\mathfrak{w}))$, hence $\phi(\mathfrak{w}) \in \mathcal{I}^0(\mathfrak{w})$. Thus, part (3) of Lemma 4 yields $\mathcal{I}^0(\phi(\mathfrak{w})) = \mathcal{I}^0(\mathfrak{w})$ and, again by Remark 5, we have $\phi(\mathcal{I}^0(\mathfrak{w})) = \mathcal{I}^0(\mathfrak{w})$ for every $\mathfrak{w} \in \mathscr{W}(X)$. Hence $\phi \in N$ and $\ker \rho \leq N$. ■

Finally, we discuss when the homomorphism ρ takes values within $\text{Aut } X$.

Lemma 5. Suppose that no hyperplane of X has extremal vertices. Then:

1. given a vertex $v \in X^{(0)}$ and transverse hyperplanes $u, w \in \mathcal{W}_v$, there exists $\mathfrak{v} \in \mathcal{W}_v - \{u, w\}$ that is transverse to u , but not to w ;
2. a vertex $w \in X^{(0)}$ is adjacent to $v \in X^{(0)}$ if and only if there does not exist $x \in X^{(0)} - \{v, w\}$ with $\mathcal{W}_v \cap \mathcal{W}_w \subseteq \mathcal{W}_v \cap \mathcal{W}_x$.

Proof. We first prove part (1). Let $v' \in u$ be the projection of v to u (i.e., the midpoint of the only edge of X that contains v and crosses u). This is a vertex of the cubical structure on u and $u \cap w$ is a hyperplane of u adjacent to v' . Since u has no extremal vertices, there exists a hyperplane \mathfrak{v}' of the cube complex u that is adjacent to v' and disjoint from $u \cap w$. If $\mathfrak{v} \in \mathcal{W}(X)$ is the hyperplane with $\mathfrak{v}' = \mathfrak{v} \cap u$, then \mathfrak{v} is adjacent to v , transverse to u , and disjoint from w (disjointness follows, for instance, from the Helly property for hyperplane carriers).

We now prove part (2). If w is not adjacent to v , there exists a vertex $x \in X^{(0)} - \{v, w\}$ that is adjacent to v and lies on a combinatorial geodesic between v and w . By convexity of carriers, we have $\mathcal{W}_v \cap \mathcal{W}_w \subseteq \mathcal{W}_v \cap \mathcal{W}_x$.

Conversely, suppose that v and w are adjacent and let $x \in X^{(0)} - \{v, w\}$ be such that $\mathcal{W}_v \cap \mathcal{W}_w \subseteq \mathcal{W}_v \cap \mathcal{W}_x$. Let w be the only hyperplane separating v and w ; since $w \in \mathcal{W}_v \cap \mathcal{W}_w$, the vertex x must lie in the carrier of w . Let x' and $v' = w'$ be the projections of the vertices x, v, w to the hyperplane w ; since $x \notin \{v, w\}$, we have $x' \neq v'$. Hence there exists a hyperplane u' of the cube complex w such that u' is adjacent to v' and separates v' from x' . Since w has no extremal vertices, there exists a hyperplane \mathfrak{v}' of w such that \mathfrak{v}' is adjacent to v' and disjoint from u' ; in particular, \mathfrak{v}' is not adjacent to x' . Now, let $\mathfrak{v} \in \mathcal{W}(X)$ be the hyperplane with $\mathfrak{v}' = \mathfrak{v} \cap w$. Note that x is not adjacent to \mathfrak{v} , or x would lie in a square of X crossed by \mathfrak{v} and w , in which case x' would be adjacent to \mathfrak{v}' . In conclusion, $\mathfrak{v} \in \mathcal{W}_v \cap \mathcal{W}_w - \mathcal{W}_x$, which contradicts our assumption that $\mathcal{W}_v \cap \mathcal{W}_w \subseteq \mathcal{W}_v \cap \mathcal{W}_x$. ■

Theorem 3.1. Let X be a uniformly locally finite CAT(0) cube complex with no extremal vertices and with no hyperplanes containing extremal vertices. The map $\iota: \text{Aut } X \rightarrow \text{Aut } \mathcal{C}$ is an isomorphism and ρ is its inverse.

Proof. By part (2) of Lemma 5, any permutation of $X^{(0)}$ in the image of $\rho: \text{Aut } \mathcal{C} \rightarrow S(X^{(0)})$ preserves adjacency of vertices. It follows that ρ takes values in $\text{Aut } X$ and we

have already shown in Corollary 1 that $\rho \circ \iota = \text{id}_{\text{Aut } X}$. Finally, part (1) of Lemma 5 and part (1) of Lemma 4 guarantee that $\mathcal{I}(\mathfrak{w}) = \{\mathfrak{w}\}$ for every $\mathfrak{w} \in \mathscr{W}(X)$. Thus, Corollary 2 shows that ρ is injective. ■

4 Proof of the Proposition

In this section, we consider a right-angled Coxeter group W_Γ , the universal cover Y_Γ of its Davis complex, and the contact graph $\mathcal{C}_\Gamma = \mathcal{C}(Y_\Gamma)$.

Let cl_n denote the complete graph on n vertices. We can always split the finite graph Γ as a join $\text{cl}_n * \Gamma'$ for some $n \geq 0$ and a subgraph $\Gamma' \subseteq \Gamma$ that is not a cone. This corresponds to splittings $W_\Gamma = (\mathbb{Z}/2\mathbb{Z})^n \times W_{\Gamma'}$ and

$$\begin{aligned} Y_\Gamma &= [0, 1]^n \times Y_{\Gamma'}, & \text{Aut } Y_\Gamma &= ((\mathbb{Z}/2\mathbb{Z})^n \rtimes S_n) \times \text{Aut } Y_{\Gamma'}, \\ \mathcal{C}_\Gamma &= \text{cl}_n * \mathcal{C}_{\Gamma'}, & \text{Aut } \mathcal{C}_\Gamma &= S_n \times \text{Aut } \mathcal{C}_{\Gamma'}, \end{aligned}$$

where S_n denotes the symmetric group on n elements. The natural map $\iota: \text{Aut } Y_\Gamma \rightarrow \text{Aut } \mathcal{C}_\Gamma$ vanishes on the $(\mathbb{Z}/2\mathbb{Z})^n$ subgroup, and it restricts to the natural map $\iota: \text{Aut } Y_{\Gamma'} \rightarrow \text{Aut } \mathcal{C}_{\Gamma'}$ on the right-hand factors.

Therefore, the general study of \mathcal{C}_Γ and its automorphism group reduces to the case $n = 0$. Since links of vertices of Y_Γ are all isomorphic to the graph Γ , this is equivalent to Y_Γ having no extremal vertices.

Let us write $Y = Y_\Gamma$ and $\mathcal{C} = \mathcal{C}_\Gamma$ for short in the rest of the section. The discussion on Salvetti complexes in the paragraph before Remark 1 readily generalizes to Davis complexes, showing that every hyperplane of Y is labelled by a vertex of Γ . Let $\gamma: \mathscr{W}(Y) \rightarrow \Gamma^{(0)}$ denote the map assigning to each hyperplane its label.

Lemma 6. Given $\mathfrak{w}, u \in \mathscr{W}(Y)$, we have $u \in \mathcal{I}(\mathfrak{w})$ if and only if the carriers of u and \mathfrak{w} intersect and $\text{st } \gamma(\mathfrak{w}) \subseteq \text{st } \gamma(u)$. In this case, we have $u \in \mathcal{I}^0(\mathfrak{w})$ if and only if $\text{st } \gamma(\mathfrak{w}) = \text{st } \gamma(u)$.

Proof. If $u \in \mathcal{I}(\mathfrak{w})$, the carrier of \mathfrak{w} is contained in the carrier of u , and we have $\text{st } \gamma(\mathfrak{w}) \subseteq \text{st } \gamma(u)$ by part (1) of Lemma 4. Conversely, suppose that $\text{st } \gamma(\mathfrak{w}) \subseteq \text{st } \gamma(u)$ and that a vertex $v \in Y^{(0)}$ lies in the carrier of both u and \mathfrak{w} . Any other vertex w in the carrier of \mathfrak{w} is joined to v by a path that only crosses edges labelled by elements of $\text{st } \gamma(\mathfrak{w})$. Since $\text{st } \gamma(\mathfrak{w}) \subseteq \text{st } \gamma(u)$, none of these edges can leave the carrier of u , hence $u \in \mathscr{W}_w$. This shows that $u \in \mathcal{I}(\mathfrak{w})$. The statement about $\mathcal{I}^0(\mathfrak{w})$ follows immediately. ■

We are now ready to prove the Davis intro ?? from the introduction, which we recall here for ease of reference.

Proposition 1. Let W_Γ have no finite direct factors. Then:

1. Y has no extremal vertices, so $\iota: \text{Aut } Y \rightarrow \text{Aut } \mathcal{C}$ is injective.
2. Y has a hyperplane with extremal vertices if and only if there exist distinct vertices $a, b \in \Gamma^{(0)}$ with $\text{st } a \subseteq \text{st } b$. In this case, the subgroup $\iota(\text{Aut } Y) < \text{Aut } \mathcal{C}$ has infinite index.
3. The homomorphism $\rho: \text{Aut } \mathcal{C} \rightarrow S(Y^{(0)})$ is injective if and only if there do not exist vertices $a, b \in \Gamma^{(0)}$ with $\text{st } a = \text{st } b$. When ρ is not injective, its kernel is an uncountable torsion subgroup.

Proof. As already observed, the fact that W_Γ has no finite factors implies that Y has no extremal vertices. Part (1) thus follows from Corollary 1.

Let us now prove part (3). By part (1) of Lemma 4, the cardinality of each set $\#\mathcal{I}^0(\mathfrak{w})$ is bounded above by $\dim Y$. Thus, part (3) of Corollary 2 shows that $\ker \rho$ is a (possibly trivial) torsion group. If no two vertices of Γ have the same star, then Lemma 6 implies that each set $\mathcal{I}^0(\mathfrak{w})$ is a singleton. In this case, part (3) of Corollary 2 shows that ρ is injective.

Conversely, suppose that distinct vertices $a, b \in \Gamma^{(0)}$ have the same star. Since W_Γ has no direct factors, there exist infinitely many $\mathfrak{w}_i \in \mathscr{W}(Y)$ with $\gamma(\mathfrak{w}_i) = a$. The sets $\mathcal{I}^0(\mathfrak{w}_i)$ are pairwise disjoint and each contains at least two elements, by Lemma 6. Part (3) of Corollary 2 yields that $\ker \rho$ is a direct product of countably many non-trivial groups, hence uncountable.

We conclude the proof of the proposition by addressing part (2). For every $\mathfrak{w} \in \mathscr{W}(Y)$, the induced cubical structure on \mathfrak{w} is isomorphic to Y_Λ , where Λ is the full subgraph of Γ with vertex set $\text{lk } \gamma(\mathfrak{w})$. This has extremal vertices if and only if Λ is a cone over some $b \in \Lambda^{(0)}$, that is, if $\text{st } \gamma(\mathfrak{w}) \subseteq \text{st } b$. We deduce that Y contains a hyperplane with an extremal vertex if and only if there exist $a, b \in \Gamma^{(0)}$ with $\text{st } a \subseteq \text{st } b$.

We are only left to show the 2nd half of part (2). Let $a, b \in \Gamma^{(0)}$ be distinct vertices with $\text{st } a \subseteq \text{st } b$, consider $w \in Y^{(0)}$, and let $\mathfrak{a}, \mathfrak{b} \in \mathscr{W}_w$ be labelled by a, b , respectively. Let \mathfrak{a}^\pm be the two halfspaces bounded by \mathfrak{a} . We define a partition $\mathscr{W}(Y) = \mathcal{A}^+ \sqcup \mathcal{A}^- \sqcup \mathcal{T}$, where a hyperplane lies in \mathcal{A}^\pm if it is contained in \mathfrak{a}^\pm , and it lies in \mathcal{T} if it is transverse or equal to \mathfrak{a} .

Let $r_\mathfrak{b} \in W_\Gamma < \text{Aut } Y$ be the reflection in the hyperplane \mathfrak{b} . Consider the map $\phi: \mathcal{C}^{(0)} \rightarrow \mathcal{C}^{(0)}$ defined as the identity on $\mathcal{A}^- \sqcup \mathcal{T}$, and as $\iota(r_\mathfrak{b})$ on $\mathcal{A}^+ \sqcup \mathcal{T}$ (note that $\iota(r_\mathfrak{b})$

coincides with the identity on \mathcal{T}). Since no edge of \mathcal{C} connects an element of \mathcal{A}^+ to an element of \mathcal{A}^- , it is clear that $\phi \in \text{Aut } \mathcal{C}$.

If a vertex $v \in Y^{(0)}$ lies in the halfspace \mathfrak{a}^- , we have $\mathscr{W}_v \subseteq \mathfrak{a}^- \sqcup \mathcal{T}$, and this set is fixed pointwise by ϕ . In this case, we have $\rho(\phi)v = v$. On the other hand, if v lies in \mathfrak{a}^+ , we have $\mathscr{W}_v \subseteq \mathfrak{a}^+ \sqcup \mathcal{T}$, where $\phi = \iota(r_b)$. Hence $\rho(\phi)v = r_b v$. Looking at the action on any square of Y that is crossed by both \mathfrak{a} and \mathfrak{b} , we see that $\rho(\phi) \notin \text{Aut } Y$.

Let us write $\phi_{\mathfrak{a}^+, \mathfrak{b}}$ in the rest of the proof, highlighting the dependence on the choice of \mathfrak{a}^+ and \mathfrak{b} in the definition of ϕ . Pick a sequence $\mathfrak{a}_n \in \mathscr{W}(Y)$ of hyperplanes with $\gamma(\mathfrak{a}_n) = \mathfrak{a}$ such that the distance between \mathfrak{a} and \mathfrak{a}_n diverges.⁴ Let $\mathfrak{b}_n \in \mathscr{W}(Y)$ be (the only) hyperplanes transverse to \mathfrak{a}_n with $\gamma(\mathfrak{b}_n) = \mathfrak{b}$. We choose the side \mathfrak{a}_n^+ so that it is disjoint from \mathfrak{a}^+ and set $\psi_n := \phi_{\mathfrak{a}^+, \mathfrak{b}} \circ \phi_{\mathfrak{a}_n^+, \mathfrak{b}_n}$. By the above discussion, we have $\psi_n \in \text{Aut } Y$.

The permutation $\rho(\psi_n) \in S(Y^{(0)})$ fixes exactly those vertices of Y that do not lie in $\mathfrak{a}^+ \sqcup \mathfrak{a}_n^+$. Since the distance between \mathfrak{a}^+ and \mathfrak{a}_n^+ diverges, the $\rho(\psi_n)$ lie in pairwise distinct cosets of $\text{Aut } Y < \rho(\text{Aut } \mathcal{C})$. Hence, since $\rho \circ \iota = \text{id}_{\text{Aut } Y}$, the ψ_n lie in pairwise distinct cosets of $\iota(\text{Aut } Y) < \text{Aut } \mathcal{C}$.

This concludes the proof. ■

Notes

1. As a set, $\text{lk } v$ is naturally identified with the link within the graph $X^{(1)}$. However, since v is a vertex of a cube complex, it will be important that $\text{lk } v$ inherits an additional structure of graph.
2. Note that the condition that Γ be not a cone in part (1) of Lemma 1 seems to have been overlooked in Remark 4.16 of [22].
3. This holds for arbitrary CAT(0) cube complexes X , as long as the fibres of the projection $C_{\text{th}}(X) \rightarrow C_r(X)$ all have the same cardinality.
4. One can show the existence of the \mathfrak{a}_n as follows. Since W_Γ has no finite factors, there exists a hyperplane $\mathfrak{c} \in \mathscr{W}(Y)$ disjoint from \mathfrak{a} . Let $r_{\mathfrak{a}}, r_{\mathfrak{c}} \in W_\Gamma$ denote the reflections in \mathfrak{a} and \mathfrak{c} , respectively. Then we can take $\mathfrak{a}_n := (r_{\mathfrak{c}} r_{\mathfrak{a}})^n \mathfrak{a}$.

Acknowledgments

I am grateful to Max Planck Institute for Mathematics in Bonn for its hospitality and financial support. I would like to thank Nir Lazarovich for pointing out that universal covers of Salvetti complexes have uncountable automorphism group (cf. Remark 1) and directing me to [20]. I would also like to thank Jason Behrstock for some comments on an earlier preprint, Mark Hagen for interesting conversations related to this work, and the anonymous referees for their many useful suggestions.

References

- [1] Brock, J. F., R. D. Canary, and Y. N. Minsky. "The classification of Kleinian surface groups, II: the ending lamination conjecture." *Ann. Math. (2)* 176, no. 1 (2012): 1–149.
- [2] Beyrer, J., E. Fioravanti, and M. Incerti-Medici. "CAT(0) cube complexes are determined by their boundary cross ratio." to appear on *Groups Geom. Dyn.* preprint arXiv:1805.08478v4.

- [3] Behrstock, J., M. F. Hagen, and A. Sisto. "Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups." *Proc. Lond. Math. Soc. (3)* 114, no. 5 (2017): 890–926.
- [4] Behrstock, J., M. F. Hagen, and A. Sisto. "Hierarchically hyperbolic spaces, I: curve complexes for cubical groups." *Geom. Topol.* 21, no. 3 (2017): 1731–804.
- [5] Behrstock, J., M. F. Hagen, and A. Sisto. "Hierarchically hyperbolic spaces II: combination theorems and the distance formula." *Pacific J. Math.* 299, no. 2 (2019): 257–338.
- [6] Behrstock, J., M. F. Hagen, and A. Sisto. "Quasiflats in hierarchically hyperbolic spaces." to appear on *Duke Math. J.* preprint arXiv:1704.04271v2.
- [7] Behrstock, J. A., T. Januszkiewicz, and W. D. Neumann. "Quasi-isometric classification of some high dimensional right-angled Artin groups." *Groups Geom. Dyn.* 4, no. 4 (2010): 681–92.
- [8] Behrstock, J., B. Kleiner, Y. Minsky, and L. Mosher. "Geometry and rigidity of mapping class groups." *Geom. Topol.* 16, no. 2 (2012): 781–888.
- [9] Bestvina, M., B. Kleiner, and M. Sageev. "The asymptotic geometry of right-angled Artin groups. I." *Geom. Topol.* 12, no. 3 (2008): 1653–99.
- [10] Behrstock, J. A. and W. D. Neumann. "Quasi-isometric classification of graph manifold groups." *Duke Math. J.* 141, no. 2 (2008): 217–40.
- [11] Cornuier, Y. and P. de la Harpe. "Metric Geometry of Locally Compact Groups," vol. 25 of *EMS Tracts in Mathematics*. Zürich: European Mathematical Society (EMS), 2016. Winner of the 2016 EMS Monograph Award.
- [12] Durham, M. G., M. F. Hagen, and A. Sisto. "Boundaries and automorphisms of hierarchically hyperbolic spaces." *Geom. Topol.* 21, no. 6 (2017): 3659–758.
- [13] Dahmani, F., M. Hagen, and A. Sisto. "Dehn filling Dehn twists." to appear on *Proceedings A of the Royal Society of Edinburgh*. preprint arXiv:1812.09715v1.
- [14] Genevois, A. "Hyperbolicities in CAT(0) cube complexes." *Enseign. Math.* 65 (2020), no. 1–2, 33–100.
- [15] Hagen, M. F. "Weak hyperbolicity of cube complexes and quasi-arboreal groups." *J. Topol.* 7, no. 2 (2014): 385–418.
- [16] Hamenstädt, U. "Geometry of the mapping class groups III: quasi-isometric rigidity." (2007): preprint arXiv:math/0512429.
- [17] Harvey, W. J. "Geometric Structure of Surface Mapping Class Groups." In *Homological Group Theory (Proc. Sympos., Durham, 1977)*, vol. 36 of London Math. Soc. Lecture Note Ser. 255–69. Cambridge–New York: Cambridge Univ. Press, 1979.
- [18] Harvey, W. J. "Boundary Structure of the Modular Group." In *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, vol. 97 of Ann. of Math. Stud. 245–51. Princeton, NJ: Princeton Univ. Press, 1981.
- [19] Huang, J. and B. Kleiner. "Groups quasi-isometric to right-angled Artin groups." *Duke Math. J.* 167, no. 3 (2018): 537–602.

- [20] Haglund, F. and F. Paulin. "Simplicité de Groupes d'Automorphismes d'Espaces à Courbure Négative." In *The Epstein Birthday Schrift*, vol. 1 of *Geom. Topol. Monogr.* 181–248. Coventry: Geom. Topol. Publ, 1998.
- [21] Hagen, M. F. and T. Susse. "On hierarchical hyperbolicity of cubical groups." *Israel J. Math.* 236 (2020), no. 1, 45–89.
- [22] Huang, J. "Quasi-isometric classification of right-angled Artin groups I: the finite out case." *Geom. Topol.* 21, no. 6 (2017): 3467–537.
- [23] Huang, J. "Commensurability of groups quasi-isometric to RAAGs." *Invent. Math.* 213, no. 3 (2018): 1179–247.
- [24] Haglund, F. and D. T. Wise. "Special cube complexes." *Geom. Funct. Anal.* 17, no. 5 (2008): 1551–620.
- [25] Ivanov, N. V. "Automorphism of complexes of curves and of Teichmüller spaces." *Internat. Math. Res. Notices* 14 (1997): 651–66.
- [26] Ivanov, N. V. "Mapping Class Groups." In *Handbook of Geometric Topology*. 523–633. Amsterdam: North-Holland, 2002.
- [27] Kim, S.-H. and T. Koberda. "Embedability between right-angled Artin groups." *Geom. Topol.* 17, no. 1 (2013): 493–530.
- [28] Kim, S.-H. and T. Koberda. "The geometry of the curve graph of a right-angled Artin group." *Internat. J. Algebra Comput.* 24, no. 2 (2014): 121–69.
- [29] Margolis, A. "Quasi-isometry classification of RAAGs that split over cyclic subgroups." to appear on *Groups Geom. Dyn.*, preprint arXiv:1803.05493v2.
- [30] Minsky, Y. "The classification of Kleinian surface groups. I. Models and bounds." *Ann. Math. (2)* 171, no. 1 (2010): 1–107.
- [31] Masur, H. A. and Y. N. Minsky. "Geometry of the complex of curves. I. Hyperbolicity." *Invent. Math.* 138, no. 1 (1999): 103–49.
- [32] Masur, H. A. and Y. N. Minsky. "Geometry of the complex of curves. II. Hierarchical structure." *Geom. Funct. Anal.* 10, no. 4 (2000): 902–74.
- [33] Sageev, M. "CAT(0) Cube Complexes and Groups." In *Geometric Group Theory*, vol. 21 of IAS/Park City Math. Ser. 7–54. Providence, RI: Amer. Math. Soc., 2014.
- [34] Tits, J. "Buildings of Spherical Type and Finite BN-Pairs." *Lecture Notes in Mathematics*, vol. 386. Berlin–New York: Springer, 1974.
- [35] Wise, D. T. "From Riches to Rags: 3-Manifolds, Right-Angled Artin Groups, and Cubical Geometry, vol. 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC. Providence, RI: American Mathematical Society, 2012.