

Coarse cubical rigidity

Elia Fioravanti¹  | **Ivan Levcovitz**² | **Michah Sageev**³

¹Max Planck Institute for Mathematics, Bonn, Germany

²Department of Mathematics, Joyce Cummings Center, Tufts University, Medford, Massachusetts, USA

³Mathematics Department, Technion - Israel Institute of Technology, Technion, Haifa, Israel

Correspondence

Elia Fioravanti, Max Planck Institute for Mathematics, Bonn, Germany.
Email: elia.fioravanti@kit.edu

Abstract

We show that for many right-angled Artin and Coxeter groups, all cocompact cubulations coarsely look the same: They induce the same coarse median structure on the group. These are the first examples of non-hyperbolic groups with this property. For all graph products of finite groups and for Coxeter groups with no irreducible affine parabolic subgroups of rank ≥ 3 , we show that all automorphisms preserve the coarse median structure induced, respectively, by the Davis complex and the Niblo–Reeves cubulation. As a consequence, automorphisms of these groups have nice fixed subgroups and satisfy Nielsen realisation.

MSC 2020

20F65, 20F67 (primary), 51F30 (secondary)

Contents

1. INTRODUCTION	2
1.1. Right-angled Artin groups	4
1.2. Right-angled Coxeter groups	5
1.3. Consequences for automorphisms	7
1.4. A word on proofs	7
1.5. Structure of the paper	8
2. PRELIMINARIES	9
2.1. Cube complexes	9
2.2. Convex-cocompactness	10
2.3. Median-cocompactness	13
2.4. Cubical coarse median structures	13
3. FIRST FORMS OF COARSE CUBICAL RIGIDITY	17

3.1. Strongly cellular cubulations	17
3.2. Graph products of finite groups	18
3.3. Virtually abelian groups	20
3.4. Right-angled Artin groups	24
4. UNIFORMLY NON-QUASI-CONVEX RAYS	29
4.1. Terminology, notation and conventions	29
4.2. Quadrants, grids and quasi-convexity	32
4.3. Constructing uniformly non-quasi-convex geodesics.	34
4.4. Conclusion	38
5. THE STANDARD COARSE MEDIAN STRUCTURE OF AN RIGHT-ANGLED COXETER GROUP.	38
5.1. Notation and a preliminary lemma	38
5.2. The main argument	40
5.3. Loose and bonded squares	42
6. GENERAL COXETER GROUPS	45
APPENDIX: COCOMPACTNESS OF INTERSECTIONS	47
ACKNOWLEDGEMENTS.	48
REFERENCES.	48

1 | INTRODUCTION

Cubulating a group G — constructing a proper G -action on a CAT(0) cube complex — is a particularly effective strategy to study many algebraic properties of G . For instance, the existence of a cocompact cubulation implies that G is biautomatic [47], has finite asymptotic dimension [61] and satisfies the Tits alternative [57].

More recently, cubulations have also proven particularly fruitful in the study of group automorphisms. For a right-angled Artin group A_Γ , a certain space of particularly nice cubulations of A_Γ provides a remarkably simple (rational) classifying space for $\text{Out}(A_\Gamma)$ [2, 15, 22], extending the classical construction of Outer Space for $\text{Out}(F_n)$ [23]. In addition, for a general cocompactly cubulated group G , automorphisms that ‘coarsely preserve’ a given cubulation of G always enjoy particularly good properties [31], for instance, their fixed subgroups are finitely generated and undistorted in G .

In light of these results, it is tempting to study spaces of cubulations in greater generality. One promising feature is that cubulations are often completely determined by their length function [3, 4]. However, the space of *all cubulations* of a group can be extremely vast and flexible. For instance, if S is a closed hyperbolic surface, every finite filling collection of simple closed curves on S gives rise to a cocompact cubulation of $\pi_1(S)$ [54], and there are additional cubulations originating from subsurfaces of S . Even right-angled Artin groups always admit many more cubulations than those appearing in the Outer Spaces mentioned above (see, e.g. Example 3.25 below).

This leads to two natural questions, which serve as our main motivation for this work.

- (1) Do all cocompact cubulations of a group G have anything in common?
- (2) Is there a ‘best’ cubulation for G ?

In relation to Question (1), it is natural to wonder whether all cubulations of G give rise to the same notion of ‘convex-cocompactness’. Specifically, we say that a subgroup $H \leq G$ is *convex-cocompact* in a cubulation $G \curvearrowright X$ if there exists an H -invariant convex subcomplex $C \subseteq X$ on which H acts cocompactly.

It is well known that, when G is word-hyperbolic, all cocompact cubulations of G give rise to the same convex-cocompact subgroups; in fact, convex-cocompact subgroups are precisely quasi-convex ones [36, Theorem H]. However, this property can fail rather drastically for non-hyperbolic groups, whose cubulations can induce infinitely many distinct notions of convex-cocompactness.

For instance, let $\mathbb{Z}^2 \curvearrowright \mathbb{R}^2$ be the standard \mathbb{Z}^2 -action on the standard square tiling of \mathbb{R}^2 , with one generator x translating horizontally by 1 and the other generator y translating vertically by 1. The only convex-cocompact maximal cyclic subgroups for this action are $\langle x \rangle$ and $\langle y \rangle$. Precomposing the action $\mathbb{Z}^2 \curvearrowright \mathbb{R}^2$ with an element $A \in \mathrm{SL}_2(\mathbb{Z})$ produces another action of \mathbb{Z}^2 on the standard tiling of \mathbb{R}^2 , where the subgroups $\langle A^{-1}x \rangle$ and $\langle A^{-1}y \rangle$ become convex-cocompact. These two actions have the same convex-cocompact subgroups if and only if A is a signed permutation matrix. Thus, we get infinitely many types of actions in terms of which subgroups are convex-cocompact.

As a matter of fact, both our motivating questions are best phrased in terms of *coarse median structures*. Rather than the general notion introduced by Bowditch [9], the following very particular instance will suffice for most of our purposes in this paper. Recall that Chepoi–Roller duality states that every CAT(0) cube complex X is equipped with a *median operator* $m : X^3 \rightarrow X$ and that, conversely, every discrete median algebra uniquely arises in this way [17, 52]. If $x, y, z \in X$ are vertices, then the median $m(x, y, z)$ is the only vertex of X that lies on a geodesic in the 1-skeleton of X between any two of the three vertices x, y, z .

Definition. A *cubical coarse median* on G is a map $\mu_X : G^3 \rightarrow G$ obtained by pulling back the median operator of a cocompact cubulation $G \curvearrowright X$ via an equivariant quasi-isometry $G \rightarrow X$ (see Subsection 2.4 for details). Two cocompact cubulations are said to induce the same *coarse median structure* if their corresponding cubical coarse medians on G are at uniformly bounded distance from each other.

In other words, two cubulations induce the same coarse median structure on G if they are G -equivariantly quasi-isometric via a map that coarsely respects medians. Additional motivation for studying these ‘coarsely median preserving’ quasi-isometries, even beyond the world of cubulated groups, comes from recent work of Petyt. He showed that mapping class groups are coarsely median preserving quasi-isometric to CAT(0) cube complexes [41, 50], albeit certainly not to any that admit geometric group actions, since mapping class groups are quasi-isometrically rigid (QI-rigid) [6, 39].

The fundamental connection between cubical coarse medians and convex-cocompactness is provided by the following result, which will be one of our main tools. It can be deduced from the first author’s work in [30] and we explain in detail how in Theorem 2.17.

Theorem ([30]). *Let $G \curvearrowright X, Y$ be cocompact cubulations. The following are equivalent:*

- (1) $G \curvearrowright X$ and $G \curvearrowright Y$ induce the same coarse median structure on G ;
- (2) $G \curvearrowright X$ and $G \curvearrowright Y$ have the same convex-cocompact subgroups.

In view of this result and the above questions, we are interested in the following property.

Definition. A group G satisfies *coarse cubical rigidity* if it has a unique cubical coarse median structure. Equivalently, all cocompact cubulations of G induce the same notion of convex-cocompactness.

Hyperbolic groups always satisfy this form of rigidity, provided that they are cocompactly cubulated. More generally, an arbitrary hyperbolic group always admits a unique coarse median structure (not necessarily a cubical one), as shown in [49, Theorem 4.2]. By contrast, as exemplified above, for $n \geq 2$ the group \mathbb{Z}^n admits countably many distinct cubical coarse median structures (see Proposition 3.9 for a classification), and uncountably many non-cubical ones.

Some exceptional groups satisfy even stronger forms of cubical rigidity. For instance, many Burger–Mozes–Wise groups [7, 13, 59] admit a *unique* cocompact cubulation with no free faces [26, Proposition C]. However, these stronger kinds of rigidity appear to be extremely rare in general, and they are likely to never occur for right-angled Artin groups, see Example 3.25.

A weaker form of coarse cubical rigidity is the existence of a coarse median structure which is preserved by any automorphism of the ambient group. Thus, while there is not a unique coarse median structure, there is a special one.

The main goal of this paper is to develop tools for proving coarse cubical rigidity of groups, and to exhibit many new examples of groups with this property. Right-angled Artin and Coxeter groups provide some natural candidates and they will indeed be our main focus, though we will discuss results for graph products and Coxeter groups generally.

1.1 | Right-angled Artin groups

Recall that a right-angled Artin group A_Γ is obtained from a finite simplicial graph Γ via the presentation:

$$A_\Gamma = \langle \Gamma^{(0)} \mid [v, w] = 1 \iff v, w \text{ span an edge in } \Gamma \rangle.$$

We say that A_Γ is *twistless* if there do not exist any vertices $v, w \in \Gamma$ with $\text{st}(v) \subseteq \text{st}(w)$; equivalently, A_Γ does not have any twist automorphisms (see Section 3 for details). As an example, A_Γ is twistless if the outer automorphism group $\text{Out}(A_\Gamma)$ happens to be finite.

In previous work, it was shown that A_Γ is twistless if and only if all automorphisms of A_Γ preserve the coarse median structure induced by the Salvetti complex [31, Proposition A(3)]. Thus, twistlessness is a necessary condition for A_Γ to have a unique cubical coarse median structure. Our first main result is that this condition is also sufficient.

Theorem A. *Let A_Γ be a right-angled Artin group. Then A_Γ satisfies coarse cubical rigidity if and only if it is twistless.*

We emphasise that Theorem A concerns arbitrary cocompact cubulations of A_Γ , not just the ones that make up the Outer Spaces constructed by Charney, Vogtmann and coauthors [2, 15, 22]. In general, A_Γ has many perfectly nice cubulations outside these Outer Spaces, even cubulations with the same dimension as the Salvetti complex and without any free faces. We demonstrate this in the case when Γ is a hexagon in Example 3.25.

Theorem A can be rephrased as follows: Every cocompact cubulation of A_Γ can be obtained from the standard action on (the universal cover of) the Salvetti complex by blowing up to hyperplanes finitely many walls associated with convex-cocompact subgroups of A_Γ , and then possibly collapsing some hyperplanes of the Salvetti complex. See Theorem 2.17(4) for the equivalence with

the above formulation of Theorem A. When A_Γ is not twistless, more complicated procedures are required to reach arbitrary cocompact cubulations (e.g. precomposition with an automorphism of A_Γ), because we need to be able to modify the induced coarse median structure.

We will actually prove a more general version of Theorem A that applies to *all* right-angled Artin groups. This is best described in terms of the product decomposition

$$\text{Aut}(A_\Gamma) = T(A_\Gamma) \cdot U(A_\Gamma),$$

where $T(A_\Gamma)$ and $U(A_\Gamma)$ are, respectively, the *twist subgroup* and the *untwisted subgroup*. The subgroup $U(A_\Gamma)$ consists of all those automorphisms that fix the coarse median structure induced by the Salvetti complex [31]. In Theorem 3.19, we will show that $T(A_\Gamma)$ acts transitively and with finite stabilisers on cubical coarse median structures with ‘decomposable flats’ (Definition 3.15).

Theorem A is a consequence of this last result: When A_Γ is twistless, all cocompact cubulations have decomposable flats and $T(A_\Gamma)$ is trivial. At the opposite end of the spectrum, not all cubulations of \mathbb{Z}^n have decomposable flats and, in fact, there are infinitely many $\text{Out}(\mathbb{Z}^n)$ -orbits of cubical coarse median structures on \mathbb{Z}^n .

1.2 | Right-angled Coxeter groups

Recall that the right-angled Coxeter group W_Γ is the quotient of A_Γ by the normal closure of the set of squares of standard generators $\{v^2 \mid v \in \Gamma^{(0)}\}$.

On the one hand, right-angled Coxeter groups tend to behave more rigidly than general right-angled Artin groups because all their automorphisms preserve the coarse median structure induced by the Davis complex [31, Proposition A(2)]. On the other, Coxeter groups are granted some additional flexibility by the fact that their standard generators have finite order. One exotic example to keep in mind is the ‘ $\frac{\pi}{4}$ -rotated’ action of the product of infinite dihedrals $D_\infty \times D_\infty$ on the square tiling of \mathbb{R}^2 : Each factor preserves one of the two lines through the origin forming an angle of $\frac{\pi}{4}$ or $\frac{3\pi}{4}$ with the coordinate axes, and each reflection axis is parallel to one of these two lines. Note that products of dihedrals can sometimes act in this way even within cubulations of larger, directly irreducible Coxeter groups (Example 5.7).

Our second main result provides two natural conditions that prevent these exotic behaviours from occurring. An action on a cube complex $G \curvearrowright X$ is *strongly cellular* if, whenever an element $g \in G$ preserves a cube of X , it fixes it pointwise.

Theorem B. *Let $W_\Gamma \curvearrowright X$ be a cocompact cubulation of a right-angled Coxeter group. Suppose that it satisfies **at least one** of the following:*

- (1) *the action is strongly cellular;*
- (2) *for all $x, y \in \Gamma^{(0)}$, the subgroup $\langle x, y \rangle$ is convex-cocompact in X .*

Then $W_\Gamma \curvearrowright X$ induces the same coarse median structure on W_Γ as the Davis complex.

Item (1) of Theorem B holds more generally for all graph products of finite groups (Theorem 3.5). These groups also have a preferred cocompact cubulation due to Davis [24], which is




FIGURE 1 The right-angled Coxeter group W_{Λ_2} satisfies coarse cubical rigidity, while W_{Λ_1} does not. Every square in Λ_2 is bonded, while the square in Λ_1 is ‘loose’.

always strongly cellular. We will refer to this cubulation as the *graph-product complex* in this paper, to avoid confusion in the Coxeter case.[†]

Item (2) of Theorem B has non-trivial content only when x and y are not adjacent. Item (2) is the result requiring the most technical argument of the paper, which will occupy the entirety of Sections 4 and 5. Importantly, Item (2) implies that cubical coarse median structures that are standard on maximal virtually abelian subgroups of W_Γ must necessarily coincide with the one induced by the Davis complex.

As a consequence of Theorem B(2) and the cubical flat torus theorem [62], we obtain many examples of right-angled Coxeter groups satisfying coarse cubical rigidity. Say that an induced square $\Delta \subseteq \Gamma$ is *bonded* if there exists another induced square $\Delta' \subseteq \Gamma$ such that the intersection $\Delta \cap \Delta'$ has exactly 3 vertices.

Corollary C. *If every induced square in Γ is bonded, then W_Γ satisfies coarse cubical rigidity.*

Bonded squares are important because they are precisely those squares $\Delta \subseteq \Gamma$ for which the standard coarse median structure is forced upon the virtually abelian subgroup $W_\Delta \leq W_\Gamma$ by its intersection pattern with highest abelian subgroups of W_Γ . On the other hand, when Δ is not bonded, there are no obvious restrictions on the coarse median structure on W_Δ . Thus, it would not be unreasonable to expect the reverse arrow in Corollary C to also hold true, though we were unable to show this.

We emphasise that there are plenty of right-angled Coxeter groups that do not satisfy coarse cubical rigidity. A one-ended, irreducible example is provided by the left-hand graph in Figure 1 (see Example 5.7). However, we do suspect that all right-angled Coxeter groups will satisfy a slightly weaker form of rigidity, namely the following.

Conjecture. *Let W_Γ be a right-angled Coxeter group.*

- There are only finitely many cubical coarse median structures on W_Γ .*
- Each cubical coarse median structure on W_Γ is completely determined by its restriction to maximal parabolic virtually abelian subgroups.*

In view of Theorem B(2), Item (b) of the Conjecture should look rather plausible. We show in Remark 3.12 that (b) \Rightarrow (a).

[†] Viewing a right-angled Coxeter group W_Γ as a graph product of order-2 groups, its ‘graph-product complex’ is the first cubical subdivision of the ‘usual Davis complex’ for W_Γ , which is the CAT(0) cube complex having the standard Cayley graph of W_Γ as its 1-skeleton. Importantly, the W_Γ -action on the Davis complex is *not* strongly cellular, while the action on the graph-product complex is. We have good reason to work with both complexes associated with W_Γ in this paper, depending on the section, so we prefer not to call them both ‘the Davis complex’.

1.3 | Consequences for automorphisms

For a cocompactly cubulated group G , the outer automorphism group $\text{Out}(G)$ naturally permutes cubical coarse median structures on G .

Automorphisms that preserve some cubical coarse median structure on G enjoy exceptionally nice properties. For instance, their fixed subgroups are finitely generated, undistorted and cocompactly cubulated [31, Theorem B], properties that can drastically fail for general automorphisms. In addition, a cubical version of Nielsen realisation works in this context [31, Corollary H].

All automorphisms of right-angled Coxeter groups preserve a cubical coarse median structure (the one of the Davis complex), while this can fail for right-angled Artin groups. Thus, it is natural to wonder if this weaker form of cubical rigidity extends to two larger classes of groups: graph products of finite groups and general Coxeter groups.

We show that this is indeed the case for these graph products and for those Coxeter groups W whose Niblo–Reeves cubulation is cocompact [48]. By [18, 58], the latter happens exactly when W does not have any irreducible affine parabolic subgroups of rank ≥ 3 . Note that, when the Niblo–Reeves cubulation is not cocompact, it is not even known whether W admits a coarse median structure (except in a few cases that turn out to be cocompactly cubulated for other reasons[†]).

Theorem D.

- (1) *Let G be a graph product of finite groups. The coarse median structure on G induced by the graph-product complex is fixed by $\text{Out}(G)$.*
- (2) *Let W be a Coxeter group with cocompact Niblo–Reeves cubulation. The coarse median structure on W induced by the Niblo–Reeves cubulation is fixed by $\text{Out}(W)$.*

It is worth remarking that, if W is not right-angled, the Niblo–Reeves cubulation is never strongly cellular, even after subdivisions. So the techniques used for Theorem B(1) cannot be applied here. We also mention that, when W is 2-spherical (i.e. when there are no ∞ labels in its Coxeter graph), the group $\text{Out}(W)$ is finite by [19, 27]. However, this does not directly imply that these automorphisms fix a coarse median structure on W without going through Theorem D.

1.4 | A word on proofs

Here we briefly sketch the proofs of Theorems A and B.

Let G_Γ be a right-angled Artin/Coxeter group, let $G_\Gamma \curvearrowright \mathcal{X}_\Gamma$ be the standard action on the (universal cover of the) Salvetti/Davis complex and let $G_\Gamma \curvearrowright Y$ be an arbitrary cocompact cubulation. Theorem A and Theorem B(1) can be quickly deduced from the following three ingredients:

- (a) To prove that Y and \mathcal{X}_Γ induce the same coarse median structure, it suffices to show that hyperplane-stabilisers of $G \curvearrowright \mathcal{X}_\Gamma$ are convex-cocompact in Y (Theorem 2.17).
- (b) Hyperplane-stabilisers of $G \curvearrowright \mathcal{X}_\Gamma$ are centralisers of finite-order elements in the Coxeter case, and direct factors of centralisers of infinite-order elements in the Artin case.

[†] For instance, the (2,4,4) Coxeter group is cocompactly cubulated, as it coincides with the automorphism group of the standard square tiling of \mathbb{R}^2 . However, its Niblo–Reeves cubulation is 4-dimensional and it is not cocompact.

- (c) Centralisers of (sufficiently high powers of) convex-cocompact, infinite-order elements are always convex-cocompact (Lemma 2.4(4)). Centralisers of finite-order elements are convex-cocompact in all cocompact cubulations that are also strongly cellular (Lemma 3.1).

The proof of Theorem B(2) is significantly more involved. It relies on the following key facts about a general group G , though in a less obvious way.

- (d) Centralisers of finitely generated subgroups of G are ‘median-cocompact’ in all cocompact cubulations of G (Proposition 2.10). That is, they act cofinitely on a median subalgebra of the cube complex.
- (e) If $G \curvearrowright X$ is a cocompact cubulation and $H < G$ is a median-cocompact subgroup that is not convex-cocompact, then there exists a combinatorial ray $r \subseteq X$ that stays uniformly close to an H -orbit \mathcal{O} , but such that convex hulls of all long subsegments of r go very far from \mathcal{O} , and do so uniformly in their length (Theorem 4.1).
- (f) Let $W_\Delta \leq W_\Gamma$ be right-angled Coxeter groups such that W_Δ is irreducible and has finite centraliser in W_Γ . If a quasi-geodesic $\alpha \subseteq W_\Delta$ spends uniformly bounded time near cosets of proper parabolic subgroups of W_Δ , then α is Morse in W_Γ (Lemma 5.1).

It is worth remarking that, using Chepoi–Roller duality, Fact (d) has the following general consequence, which seems new and interesting. Rather surprisingly, we are not aware of any proofs of this result that do not rely on any abstract median-algebra or coarse-median techniques.

Corollary E. *Let G be a cocompactly cubulated group.*

- (1) *For every finitely generated subgroup $H \leq G$, the centraliser $Z_G(H)$ is cocompactly cubulated.*
- (2) *If $G = G_1 \times G_2$ and G_2 has finite centre, then G_1 is cocompactly cubulated.*

Special cases of this result already appear in the literature. Item (1) was shown by Haettel in [35, Theorem 2.1] when H is a sufficiently deep finite-index subgroup of an abelian subgroup of G , and by Genevois in [33, Theorem 5.1] when $H \cong \mathbb{Z}$. A special case of Item (2) under stronger assumptions on G_2 follows from work of Kropholler–O’Donnell [45, Theorem A].

Regarding Item (1), we emphasise that it is possible that no cubulation of $Z_G(H)$ can be realised as a convex subcomplex of any cubulation of G . We give one such example in Example 3.14. Also note that Item (2) in the corollary can fail if G_2 has infinite centre: If W is the (3,3,3) Coxeter group, then the group $W \times \mathbb{Z}$ is cocompactly cubulated (Example 2.16), while W is not [37].

1.5 | Structure of the paper

Section 2 collects various basic facts on convex-cocompactness, median-cocompactness and (cubical) coarse median structures. Theorem 2.17 is the most important result of the section, as it will be our main tool in all proofs of coarse cubical rigidity. We prove Corollary E in Subsection 2.4.

Section 3 is concerned with the simplest cases of coarse cubical rigidity. We prove Theorem B(1) for general graph products of finite groups in Subsection 3.2; Theorem D(1) immediately follows from that. In Subsection 3.3, we study cubical coarse median structures on virtually abelian groups, also proving that Item (b) in the Conjecture implies Item (a). Finally, we prove Theorem A in Subsection 3.4.

Sections 4 and 5 are devoted to the proof of Theorem B(2). Section 4 is only concerned with median subalgebras of cube complexes, its main goal being Theorem 4.1, showing the existence

of the ‘uniformly non-quasi-convex’ rays mentioned in Item (e) of the above sketch. Then Section 5 applies this result to prove Theorem B(2) and deduce Corollary C.

Finally, Section 6 contains the proof of Theorem D(2). The Appendix proves two basic lemmas about subsets of metric spaces with cocompact stabilisers.

2 | PRELIMINARIES

2.1 | Cube complexes

We will assume a certain familiarity with CAT(0) cube complexes and group actions on them. For an introduction to the topic, the reader can consult [20, 52, 55, 60]. In this subsection, we only fix terminology and notation.

A *cubulation* of a group G is a proper action on a CAT(0) cube complex. We say that G is *cocompactly cubulated* if it admits a cocompact cubulation. All actions on cube complexes are assumed to be cellular (i.e. by cubical automorphisms).

An action $G \curvearrowright X$ on a cube complex is *strongly cellular* if the set of fixed points of each element of G is a subcomplex of X (this definition is equivalent to the one given in the Introduction). The action is *non-transverse* if there do not exist a hyperplane \mathfrak{w} and element $g \in G$ such that \mathfrak{w} and $g\mathfrak{w}$ are transverse; when G acts cocompactly, this is equivalent to saying that X ℓ_1 -embeds equivariantly in a finite product of trees. Non-transverse actions are strongly cellular, but the converse does not hold; for instance, since every free action is strongly cellular, one can consider the universal cover of any non-positively curved cube complex with a self-intersecting hyperplane, together with the action by deck transformations.

Every CAT(0) cube complex admits two natural metrics. The ℓ_2 -metric d_2 (usually known as the CAT(0) metric) and the ℓ_1 -metric d_1 , namely the induced path metric obtained by equipping every cube with its ℓ_1 -metric. The restriction of d_1 to the 1-skeleton is also known as the *combinatorial metric*. Whenever we speak of ‘geodesics’ without any prefixes, we will always refer to ℓ_1 -geodesics contained in the 1-skeleton. Note however that, on a few occasions, it will be useful to consider ℓ_1 -geodesics outside the 1-skeleton. For instance, we will need the following observation in Lemma 3.6, though this is not involved in the proof of any of the main results of the paper.

Lemma 2.1. *Let X be a CAT(0) cube complex. Every ℓ_2 -geodesic is an ℓ_1 -geodesic.*

Proof. Recall that (X, d_1) is a finite-rank median space. Let $\alpha \subseteq X$ be an ℓ_2 -geodesic. If α were not an ℓ_1 -geodesic, it would cross a median wall \mathfrak{w} at least twice. Every median wall has the form $H \times \{t\}$, where H is a hyperplane with carrier $H \times [-\frac{1}{2}, \frac{1}{2}]$ and $t \in [-\frac{1}{2}, \frac{1}{2}]$. In particular, \mathfrak{w} is ℓ_2 -convex, so it cannot be crossed twice by α . \square

Every CAT(0) cube complex admits a median operator $m : X^3 \rightarrow X$. The median $m(x, y, z)$ is the only point lying on an ℓ_1 -geodesic between any two of the three points x, y, z . With the exception of the proof of Lemma 3.6, we will only need the restriction of m to the 0-skeleton of X . If x, y, z are vertices, then $m(x, y, z)$ is the only vertex with the property that no hyperplane of X separates $m(x, y, z)$ from two of the vertices x, y, z .

A subset $A \subseteq X$ is a *median subalgebra* if $m(A, A, A) \subseteq A$. A full subcomplex $C \subseteq X$ is *convex* if its 0-skeleton satisfies the stronger property $m(C^{(0)}, C^{(0)}, X^{(0)}) \subseteq C^{(0)}$. This notion is equivalent to

convexity with respect to either the ℓ_1 - or the ℓ_2 -metric [38]. If $A \subseteq X^{(0)}$, we speak of the *convex hull* of A , referring to the vertex set of the smallest convex subcomplex of X containing A . We also speak of the median subalgebra *generated* by A , referring to the smallest median subalgebra of X containing A . We will need the following lemma stating that convex hulls and generated subalgebras can be constructed by taking medians a bounded number of times.

Lemma 2.2. *Let X be a CAT(0) cube complex of dimension δ and let $A \subseteq X^{(0)}$ be a subset.*

- (1) *The convex hull of A is equal to the set $\mathcal{J}^\delta(A)$, where the operators $\mathcal{J}^i(\cdot)$ are defined inductively by $\mathcal{J}^0(A) = A$ and $\mathcal{J}^{i+1}(A) = m(\mathcal{J}^i(A), \mathcal{J}^i(A), X^{(0)})$.*
- (2) *The median subalgebra generated by A is equal to $\mathcal{M}^{2\delta}(A)$, where the operators $\mathcal{M}^i(\cdot)$ are defined inductively by $\mathcal{M}^0(A) = A$ and $\mathcal{M}^{i+1}(A) = m(\mathcal{M}^i(A), \mathcal{M}^i(A), \mathcal{M}^i(A))$.*

Part (1) of the lemma is [9, Lemma 5.5]. The bound mentioned in part (2) was obtained in Bowditch’s book [11, Proposition 8.2.4]; also see [10, Proposition 4.1] or [31, Proposition 4.2] for previous results of the same type.

2.2 | Convex-cocompactness

Let X be a CAT(0) cube complex.

Definition 2.3. Let $H \curvearrowright X$ be an action. We say that H is *convex-cocompact* in X if there exists an H -invariant convex subcomplex $C \subseteq X$ that is acted upon cocompactly by H .

We define the *rank* of a finitely generated, virtually abelian group as the rank of its finite-index free abelian subgroups. If G is a group, we say that a finitely generated, virtually abelian subgroup $A \leq G$ is *highest* if it is not virtually contained in a free abelian subgroup of strictly higher rank.

The following lemma collects the main properties of convex-cocompactness from [31, 62] and [30].

Lemma 2.4. *Let $G \curvearrowright X$ be a cocompact cubulation.*

- (1) *Finite intersections of convex-cocompact subgroups of G are convex-cocompact in X .*
- (2) *Convex-cocompactness in X is a commensurability invariant for subgroups of G .*
- (3) *All highest virtually abelian subgroups of G are convex-cocompact in X .*
- (4) *If $H \leq G$ is convex-cocompact in X and the action $H \curvearrowright X$ is non-transverse, then the normaliser $N_G(H)$ is convex-cocompact in X .*

Proof. Part (1) is [30, Lemma 2.7]. Part (2) is immediate from the equivalence between convex-cocompactness and coarse median quasi-convexity (Proposition 2.14 below). Part (3) follows from the cubical flat torus theorem [62, Theorem 3.6]. Part (4) is [30, Corollary 2.12(3)]. □

We caution the reader that the non-transversality assumption really is necessary for Lemma 2.4(4) to hold; see [30, Example 2.11].

Remark 2.5. Consider a product of CAT(0) cube complexes $X_1 \times X_2$ and an element $g \in \text{Aut}(X_1 \times X_2)$ preserving each factor of the splitting. Then $\langle g \rangle$ is convex-cocompact in $X_1 \times X_2$ if and only if it is elliptic in one of X_1, X_2 and convex-cocompact in the other.

Indeed, every $\langle g \rangle$ -invariant convex subcomplex $C \subseteq X_1 \times X_2$ splits as a product $C_1 \times C_2$ of $\langle g \rangle$ -invariant convex subcomplexes $C_i \subseteq X_i$. If g were loxodromic in both X_i , then both C_i would contain bi-infinite geodesics, hence C would contain a copy of \mathbb{R}^2 , barring cocompactness.

The following technical lemma is an important ingredient in the proof of Theorems A and B(2).

Lemma 2.6. *Consider a cocompact cubulation $G \curvearrowright X$ satisfying all of the following conditions:*

- $G = H \times K$ and $H \simeq \mathbb{Z}$;
- H is convex-cocompact in X ;
- K is generated by k_1, \dots, k_n , where each $\langle k_i \rangle$ is convex-cocompact in X .

Then K is convex-cocompact in X .

Proof. In the Caprace–Sageev terminology [20, Section 3.3], every hyperplane of X is either H -essential, H -half-essential or H -trivial. Of the two halfspaces of X associated with an H -half-essential hyperplane, one is H -deep and the other is H -shallow.

Since H is normal in G , the set of H -deep halfspaces bounded by H -half-essential hyperplanes is G -invariant. The intersection Y of all these halfspaces is non-empty, as it contains any axis for $H \simeq \mathbb{Z}$. Thus, $Y \subseteq X$ is a G -invariant convex subcomplex on which G acts properly and cocompactly. Replacing X with Y , we can safely assume that there are no H -half-essential hyperplanes.

Now, observing that every H -essential hyperplane is transverse to every H -trivial hyperplane, we obtain a splitting of X as a product of cube complexes $X_1 \times X_0$ [20, Lemma 2.5]. Every hyperplane of X_1 is H -essential, while every hyperplane of X_0 is H -trivial. Since G normalises H , it takes H -essential hyperplanes to H -essential hyperplanes. It follows that the G -action respects the splitting $X = X_1 \times X_0$.

We only need three observations to conclude the proof.

Claim 1. X_1 is a quasi-line.

Since all hyperplanes of X_1 are H -essential, every H -invariant convex subcomplex of X is a union of fibres $X_1 \times \{*\}$. Thus, since H is convex-cocompact in X , it must act cocompactly on X_1 . Since all hyperplanes of X_0 are H -trivial, the action $H \curvearrowright X_0$ has a global fixed point (e.g. by [16, Proposition B.8]), hence H must act properly on X_1 . In conclusion, the action $H \curvearrowright X_1$ is proper and cocompact, proving that X_1 is a quasi-line.

Claim 2. There exists a G -invariant ℓ_2 -geodesic line $\alpha \subseteq X_1$.

Indeed, the union of all H -axes in X_1 (for the ℓ_2 -metric) splits as a product of CAT(0) spaces $\mathbb{R} \times Z$ (e.g. by [5, Theorem II.7.1]). This subset, as well as its splitting, is G -invariant. Since X_1 is a quasi-line, the CAT(0) space Z is bounded, hence the induced G -action on Z has a global fixed point z . The fibre $\mathbb{R} \times \{z\}$ is the required axis α .

Claim 3. Each generator $k_i \in K$ is elliptic in X_1 .

This is clear if k_i has finite order; suppose instead that k_i has infinite order. Note that k_i must be elliptic in either X_0 or X_1 by Remark 2.5, since $\langle k_i \rangle$ is convex-cocompact in X by hypothesis. However, no infinite-order k_i can be elliptic in X_0 . Otherwise, since H is elliptic in X_0 , the sub-

group $H \times \langle k_i \rangle$ would also be elliptic in X_0 . This would force $H \times \langle k_i \rangle$ to act properly on X_1 , which cannot happen, since $H \times \langle k_i \rangle \simeq \mathbb{Z}^2$ while X_1 is a quasi-line.

We now finish the proof using the above claims. Since each k_i commutes with H , preserves the H -axis α and is elliptic in X_1 , we conclude that each k_i fixes α pointwise. This implies that the entire K is elliptic in X_1 . Since $G = H \times K$ acts cocompactly on X , we deduce that K must act cocompactly on X_0 . Since X_0 embeds K -equivariantly in X as a convex subcomplex, this shows that K is convex-cocompact in X . \square

We emphasise that it really is important in Lemma 2.6 that K be generated by convex-cocompact elements. Indeed, any epimorphism $\alpha : K \rightarrow \mathbb{Z}$ gives an automorphism of $G = H \times K$ fixing $H \cong \mathbb{Z}$ pointwise and taking K to the subgroup $K' = \{(\alpha(k), k) \mid k \in K\}$. So G can have many product splittings $G = H \times K'$ where K' is a subgroup isomorphic to K , though only one of these subgroups will be convex-cocompact in X .

The following lemma is needed in the proof of Theorem B(2).

Lemma 2.7. *Let X be a finite-dimensional CAT(0) cube complex. Consider an action $G \curvearrowright X$ where $G = H \times K$ is finitely generated. If H and K are convex-cocompact in X , then so is G .*

Proof. We will only use this result in the case when X is locally finite and G acts properly, where one could actually give a more direct proof using [20] instead of [29]. Nevertheless, we choose to prove the general statement, as this will likely be useful elsewhere.

Up to subdividing X , we can assume that G acts without hyperplane-inversions. By [29, Theorem 3.16 and Remark 3.17], the reduced core $\overline{C}(H, X)$ is a non-empty, G -invariant, convex subcomplex of X . In addition, there is a G -invariant splitting $\overline{C}(H, X) = C_1 \times C_0$, where every hyperplane of C_1 is skewed by an element of H , while the action $H \curvearrowright C_0$ factors through a finite group (see [29, Lemma 3.22, Corollary 4.6 and Remark 3.9]). The fact that $H \curvearrowright C_0$ factors through a finite group, rather than simply being elliptic, is the key difference from [20].

If H is convex-cocompact in X , the action $H \curvearrowright C_1$ is cocompact. Indeed, if $Y \subseteq X$ is an H -cocompact, convex subcomplex, we can project Y first to $\overline{C}(H, X)$ and then to C_1 . The result is an H -cocompact, convex subcomplex $Y' \subseteq C_1$, which must be the entire C_1 , since every hyperplane of C_1 is H -skewed. Thus, $H \curvearrowright C_1$ is cocompact.

Since the splitting $C_1 \times C_0$ is G -invariant, we have an action $G \curvearrowright C_0$. With another application of the above results from [29], we obtain a G -invariant splitting $\overline{C}(K, C_0) = C_{01} \times C_{00}$. Here $C_{01} \subseteq C_0$ is a G -invariant, convex subcomplex all whose hyperplanes are skewed by elements of K . Again, if K is convex-cocompact in X , the action $K \curvearrowright C_{01}$ is cocompact.

In conclusion, we have found a G -invariant, convex subcomplex $C_1 \times C_{01} \subseteq X$ such that the actions $H \curvearrowright C_1$ and $K \curvearrowright C_{01}$ are cocompact, and the action $H \curvearrowright C_{01}$ factors through a finite group.

Now, let $A \subseteq C_1$ and $B \subseteq C_{01}$ be compact fundamental domains for the actions of H and K , respectively. Denote by $H \cdot B$ the union of all H -translates of B ; this is again a compact set, since the action $H \curvearrowright C_{01}$ factors through a finite group. Finally, the set $A \times (H \cdot B)$ is a compact fundamental domain for the action $G \curvearrowright C_1 \times C_{01}$, showing that G is convex-cocompact in X . \square

Remark 2.8. Although we will not be needing this, it is interesting to note that the proof of Lemma 2.7 is actually showing a bit more. The group G acts cocompactly on a convex subcomplex of X with a G -invariant product splitting $C = C_1 \times C_2$ such that the actions $H \curvearrowright C_1$ and $K \curvearrowright C_2$ are cocompact and essential, while the action $H \curvearrowright C_2$ factors through a finite group. Without

stronger assumptions, we cannot say more, for instance, the action $K \curvearrowright C_1$ need not even be elliptic (the lack of symmetry between H and K is due to the fact that the splitting of C is not canonical). Indeed, we might well have $H \cong K$ and the G -action might be given by a homomorphism $\rho : H \times H \rightarrow \text{Aut}(X)$ with $\rho(h, 1) = \rho(1, h)$ for all $h \in H$, in which case, $K \curvearrowright C_1$ is essential and C_2 is a single point (not having any hyperplanes, C_2 is nevertheless K -essential as well).

2.3 | Median-cocompactness

Let X be a CAT(0) cube complex.

In some cases, one can prove convex-cocompactness of a subgroup $H \leq \text{Aut}(X)$ by first establishing a significantly weaker property — *median-cocompactness* — and then promoting this to genuine convex-cocompactness.

Definition 2.9. Let $H \curvearrowright X$ be an action. We say that H is *median-cocompact* in X if there exists an H -invariant median subalgebra $M \subseteq X^{(0)}$ that is acted upon cofinitely by H .

As an example, if X is the standard square tiling of \mathbb{R}^2 , the diagonal $\{(n, n) \mid n \in \mathbb{Z}\}$ is a median subalgebra, although the only convex subcomplex containing it is the whole X . Thus, considering the standard action $\mathbb{Z}^2 \curvearrowright \mathbb{R}^2$, the diagonal subgroup $\langle (1, 1) \rangle \simeq \mathbb{Z}$ is median-cocompact, but not convex-cocompact.

If $G \curvearrowright X$ is a cocompact cubulation and $H \leq G$ is a median-cocompact subgroup, then H is necessarily finitely generated, undistorted and cocompactly cubulated [31, Lemma 4.12]. However, the cubulation of H is ‘abstractly’ provided by Chepoi–Roller duality and it cannot be realised as a convex subcomplex of X in general.

A fundamental observation is that median-cocompactness always comes for free for *centralisers*, while this is not true of convex-cocompactness without strong additional assumptions (compare Lemma 2.4(4) above and Lemma 3.1 below).

Proposition 2.10. Let $G \curvearrowright X$ be a cocompact cubulation. For every finitely generated subgroup $H \leq G$, the centraliser $Z_G(H)$ is median-cocompact in X .

The proof of this result is not hard, but it is best phrased in terms of coarse median structures, so we postpone it until the next subsection (Lemma 2.15). Corollary E will quickly follow from Proposition 2.10, as we also explain in the next subsection.

2.4 | Cubical coarse median structures

Let G be a finitely generated group equipped with a word metric. The specific choice of word metric is inconsequential in the following discussion.

A *coarse median* is a map $\mu : G^3 \rightarrow G$ for which there exists a constant $C \geq 0$ such that, for all $a, b, c, x \in G$, the following hold:

- (1) $\mu(a, a, b) = a$ and $\mu(a, b, c) = \mu(b, c, a) = \mu(b, a, c)$;
- (2) $\mu(\mu(a, x, b), x, c)$ and $\mu(a, x, \mu(b, x, c))$ are at distance $\leq C$;
- (3) $d(\mu(a, b, c), \mu(x, b, c)) \leq Cd(a, x) + C$;
- (4) $x\mu(a, b, c)$ and $\mu(xa, xb, xc)$ are at distance $\leq C$.

Coarse medians were introduced by Bowditch [9], while the above formulation is due to Niblo–Wright–Zhang [49]. Condition (4) is sometimes omitted, but it is important for our purposes.

Two coarse medians μ, ν are *at bounded distance* from each other if the points $\mu(a, b, c)$ and $\nu(a, b, c)$ are at uniformly bounded distance for $a, b, c \in G$. A *coarse median structure* on G is an equivalence class $[\mu]$ of coarse medians pairwise at bounded distance.

Every cocompact cubulation $G \curvearrowright X$ induces a specific coarse median structure $[\mu]$ on G . This consists of the maps $\mu : G^3 \rightarrow G$ for which there exists a constant C such that:

$$d(\mu(g_1, g_2, g_3)v, m(g_1v, g_2v, g_3v)) \leq C, \quad \forall g_1, g_2, g_3 \in G, \quad \forall v \in X^{(0)}.$$

That such maps μ exist is a straightforward consequence of the Milnor–Schwarz lemma. Coarse median structures arising this way are much better behaved than arbitrary ones, so they deserve a special name.

Definition 2.11. A coarse median structure on G is *cubical* if induced by a cocompact cubulation. Let $\mathcal{CM}(G)$ be the set of coarse median structures on G , and $\mathcal{CM}_{\square}(G)$ the subset of cubical ones.

If $[\mu] \in \mathcal{CM}(G)$ and $\varphi \in \text{Aut}(G)$, then the map

$$(a, b, c) \mapsto \varphi(\mu(\varphi^{-1}(a), \varphi^{-1}(b), \varphi^{-1}(c)))$$

is also a coarse median. If φ is an inner automorphism, then this procedure does not alter the coarse median structure $[\mu]$. For this to hold, it is important that we included condition (4) in the above definition of coarse median, and also that we identify coarse medians at bounded distance from each other (indeed, conjugation by an element $g \in G$ displaces the actual coarse median μ by a linear function of the word length of g). In conclusion, the group $\text{Aut}(G)$ acts on the set of coarse medians on G , and the subgroup of inner automorphisms fixes pointwise the quotient of coarse median structures $\mathcal{CM}(G)$. Thus, we obtain a natural $\text{Out}(G)$ -action on $\mathcal{CM}(G)$ and $\mathcal{CM}_{\square}(G)$.

Example 2.12.

- (1) If G is word-hyperbolic, $\mathcal{CM}(G)$ is a singleton by [49, Theorem 4.2]. In particular, all cocompact cubulations of G induce the same (cubical) coarse median structure.
- (2) Instead, $\mathcal{CM}_{\square}(\mathbb{Z}^2)$ is countably infinite (see Proposition 3.9) and $\mathcal{CM}(\mathbb{Z}^2)$ is uncountable.

A fundamental fact is that the information of which subgroups of G are convex-cocompact or median-cocompact in a given cubulation is entirely encoded in the induced coarse median structure.

Definition 2.13. Consider $[\mu] \in \mathcal{CM}(G)$ and a subgroup $H \leq G$.

- (1) H is $[\mu]$ -*quasi-convex* if $\mu(H, H, G)$ is at finite Hausdorff distance from H .
- (2) H is $[\mu]$ -*quasi-submedian* if $\mu(H, H, H)$ is at finite Hausdorff distance from H .

When the coarse median structure is understood, we will sometimes simply speak of *quasi-convex* or *quasi-submedian* subgroups. We emphasise that there is no connection between quasi-submedian subgroups and the ‘quasi-median graphs’ studied, for example, in [8, 14, 32].

Proposition 2.14. *Let $G \curvearrowright X$ be a cocompact cubulation, inducing $[\mu_X] \in \mathcal{CM}_\square(G)$. Then:*

- (1) $H \leq G$ is convex-cocompact in X if and only if it is $[\mu_X]$ -quasi-convex;
- (2) $H \leq G$ is median-cocompact in X if and only if it is $[\mu_X]$ -quasi-submedian.

Proof. Both parts follow from Lemma 2.2, which guarantees that both the convex hull of an H -orbit in X and the median subalgebra generated by it can be constructed by taking medians a bounded number of times. More details on the argument can be found, for example, in [31, Lemma 3.2] or in the proof of [31, Theorem 4.10]. (Note that, in [31], quasi-submedian subgroups are referred to as ‘approximate median subalgebras’.) □

In view of Proposition 2.14, the fact that centralisers of finitely generated subgroups are median-cocompact in all cocompact cubulations (Proposition 2.10) becomes an immediate consequence of the following observation about general coarse median structures.

Lemma 2.15. *If $[\mu] \in \mathcal{CM}(G)$ and $H \leq G$ is finitely generated, then $Z_G(H)$ is $[\mu]$ -quasi-submedian.*

Proof. For simplicity, if $x, y \in G$ are at distance $\leq D$ in the chosen word metric, we write $x \approx_D y$. We also write $|x|$ for the word length of x . Finally, let C be the constant for which μ satisfies the four conditions in the definition of coarse medians at the beginning of this subsection.

Set $Z := Z_G(H)$. Consider a point $x \in \mu(Z, Z, Z)$, say $x = \mu(z_1, z_2, z_3)$ with $z_1, z_2, z_3 \in Z$. If $h \in H$, using Conditions (4) and (3) in this order, we obtain:

$$hx = h\mu(z_1, z_2, z_3) \approx_C \mu(hz_1, hz_2, hz_3) = \mu(z_1h, z_2h, z_3h) \approx_{3C(1+|h|)} \mu(z_1, z_2, z_3) = x.$$

It follows that $x^{-1}hx$ lies in the finite subset of G with word length at most $4C + 3C|h|$. Hence x lies in a finite union of right cosets of $Z_G(h)$, call it $R(h)$.

Since x was arbitrary, we have $\mu(Z, Z, Z) \subseteq R(h)$ for every $h \in H$. If h_1, \dots, h_n are a finite generating set for H , the intersection $R(h_1) \cap \dots \cap R(h_n)$ is a finite union of right cosets of $Z_G(h_1) \cap \dots \cap Z_G(h_n) = Z$, and it contains $\mu(Z, Z, Z)$. This shows that $\mu(Z, Z, Z)$ is contained in a bounded neighbourhood of Z . Since $\mu(Z, Z, Z)$ trivially contains Z , it is at finite Hausdorff distance from Z , proving the lemma. □

Now that the proof of Proposition 2.10 is completed, we take the chance to quickly deduce Corollary E from it.

Proof of Corollary E. Let $G \curvearrowright X$ be a cocompact cubulation. If $H \leq G$ is finitely generated, Proposition 2.10 guarantees that $Z_G(H)$ acts cofinitely on a median subalgebra $M \subseteq X^{(0)}$. By Chepoi–Roller duality, there exists an action on a CAT(0) cube complex $Z_G(H) \curvearrowright Y$ such that M and $Y^{(0)}$ are equivariantly isomorphic as median algebras. It follows that the action $Z_G(H) \curvearrowright Y$ is a cocompact cubulation. This proves part (1) of the corollary.

If $G = G_1 \times G_2$ and G_2 has finite centre, then G_1 is a finite-index subgroup of the centraliser $Z_G(G_2)$, which is cocompactly cubulated by part (1). Hence G_1 is also cocompactly cubulated, proving part (2). □

The proof of Corollary E also shows that, if a product $G_1 \times G_2$ is cocompactly cubulated, then $G_1 \times \mathbb{Z}^n$ is cocompactly cubulated for some $n \geq 1$. In general, however, this does not imply that G_1 is itself cocompactly cubulated, as the next example demonstrates.

Example 2.16. Consider the (3,3,3) Coxeter group W , which is not cocompactly cubulated by [37]. Let us see that, instead, $W \times \mathbb{Z}$ acts properly and cocompactly on the cubical tiling of \mathbb{R}^3 .

Consider the elements $\rho, \sigma \in O(3, \mathbb{Z})$ and the translations T_0, T_1 defined as follows:

$$\begin{aligned}\rho(x, y, z) &= (x, z, y), & T_0(x, y, z) &= (x + 1, y + 1, z + 1), \\ \sigma(x, y, z) &= (y, z, x), & T_1(x, y, z) &= (x - 2, y + 1, z + 1).\end{aligned}$$

All four elements preserve the standard cubical tiling of \mathbb{R}^3 , and T_0 commutes with ρ, σ, T_1 . The plane $P = \{(x, y, z) \mid x + y + z = 0\}$ is orthogonal to the axes of T_0 and it is preserved by the reflection ρ and the rotation σ . In particular, $\sigma|_P$ is an order-3 rotation around the origin and $\rho|_P$ is a reflection in a line through the origin that is also an axis for T_1 . This shows that $\langle \rho, \sigma \rangle$ is an order-6 dihedral group, and the reflection axes of the restrictions to P of $\rho, \sigma\rho\sigma^{-1}, (T_1\sigma^2)\rho(T_1\sigma^2)^{-1}$ intersect forming an equilateral triangle.

In conclusion, $\langle \rho, \sigma\rho\sigma^{-1}, (T_1\sigma^2)\rho(T_1\sigma^2)^{-1} \rangle$ is a copy of W acting properly and cocompactly on the plane P . This group commutes with T_1 , which translates perpendicularly to P , providing the required geometric action $W \times \mathbb{Z} \curvearrowright \mathbb{R}^3$.

We emphasise that, although this example is probably well known, the action $W \curvearrowright \mathbb{R}^3$ considered here is *not* the Niblo–Reeves cubulation of W . The latter is also an action of W on the cubical tiling of \mathbb{R}^3 preserving the same exact plane P , but the two actions differ in the following key way.

- In the current example, each order-2 element of W acts as a reflection in a plane orthogonal to P . In particular, the W -action preserves each of the two connected components of $\mathbb{R}^3 - P$.
- In the Niblo–Reeves cubulation, each order-2 element of W acts as a rotation by π around a line contained in P . Each of these elements swaps the two connected components of $\mathbb{R}^3 - P$.

As a consequence, the Niblo–Reeves cubulation $W \curvearrowright \mathbb{R}^3$ *cannot* be extended to an action $W \times \mathbb{Z} \curvearrowright \mathbb{R}^3$: In order to commute with all order-2 elements of W , the \mathbb{Z} -factor would have to translate in a direction that is parallel to all their fixed lines, but no such direction exists.

We now turn to an important tool to check that two cubulations of a group G induce the same coarse median structure, namely Theorem 2.17 below. Before stating the theorem, we need to recall a classical construction.

If $G \curvearrowright X$ is an action on a CAT(0) cube complex and $\mathcal{U} \subseteq \mathcal{W}(X)$ is a G -invariant set of hyperplanes, we can collapse the hyperplanes in \mathcal{U} and obtain a new action on a CAT(0) cube complex $G \curvearrowright X(\mathcal{U})$, along with a G -equivariant, surjective, cubical map $\pi_{\mathcal{U}} : X \rightarrow X(\mathcal{U})$. The action $G \curvearrowright X(\mathcal{U})$ is known as a *restriction quotient* of $G \curvearrowright X$, see [20, section 2.3] for details. Note that restriction quotients of cocompact actions will remain cocompact, but restriction quotients of proper actions may well stop being proper.

A map $f : M \rightarrow N$ between median algebras (in our case, subalgebras of cube complexes) is a *median morphism* if, for all $x, y, z \in M$, we have $f(m(x, y, z)) = m(f(x), f(y), f(z))$. Consider two actions $G \curvearrowright X$ and $G \curvearrowright Y$ on CAT(0) cube complexes. As observed in [31, Proposition 2.20] (crucially exploiting [40, Theorem 4.4]), a G -equivariant surjective cubical map $f : X \rightarrow Y$ corresponds to a restriction quotient if and only if f is a median morphism on the 0-skeleton. In particular, if a cocompact cubulation is a restriction quotient of another cocompact cubulation, then the two induced coarse median structures on G will coincide.

The next result will be our main tool to show that two cubulations of G induce the same coarse median structure. It reduces the problem to checking that finitely many subgroups are convex-cocompact. Note that Item (3) in Theorem 2.17 is deliberately asymmetric in X and Y .

Theorem 2.17. *Let $G \curvearrowright X$ and $G \curvearrowright Y$ be cocompact cubulations. The following are equivalent:*

- (1) $G \curvearrowright X$ and $G \curvearrowright Y$ induce the same coarse median structure on G ;
- (2) $G \curvearrowright X$ and $G \curvearrowright Y$ have the same convex-cocompact subgroups;
- (3) every G -stabiliser of a hyperplane of X is convex-cocompact in Y ;
- (4) there exists a third cocompact cubulation $G \curvearrowright Z$ of which both $G \curvearrowright X$ and $G \curvearrowright Y$ are restriction quotients.

Proof. The implication (1) \Rightarrow (2) follows from Proposition 2.14, while the implication (2) \Rightarrow (3) is obvious. The implication (4) \Rightarrow (1) is also clear, since then Z and X induce the same coarse median structure on G , and so do Z and Y .

We are left to prove that (3) \Rightarrow (4), which is the main content of the theorem. Suppose that all stabilisers of hyperplanes of X are convex-cocompact in Y . By [30, Proposition 7.9], there exists a non-empty, G -invariant, G -cofinite median subalgebra $N \subseteq X^{(0)} \times Y^{(0)}$.

Let $K_X \subseteq X$ and $K_Y \subseteq Y$ be finite subcomplexes with $G \cdot K_X = X$ and $G \cdot K_Y = Y$. Consider the diagonal action $G \curvearrowright X \times Y$ and the finite subcomplex $K := K_X \times K_Y \subseteq X \times Y$.

Let us show that there exists a G -invariant, G -cofinite median subalgebra $M \subseteq X \times Y$ such that $K \subseteq M$. To begin with, since K is compact, the set $G \cdot K$ is at finite Hausdorff distance from N . Since the median operator m is 1-Lipschitz, it follows that the set of medians $m(G \cdot K, G \cdot K, G \cdot K)$ is at finite Hausdorff distance from $m(N, N, N) = N$. Since the median subalgebra M generated by $G \cdot K$ is obtained by taking medians a bounded number of times (Lemma 2.2), this implies that M is at finite Hausdorff distance from N too. It is clear that M is G -invariant and contains K . Finally, since $X \times Y$ is locally finite and N is G -cofinite, M is also G -cofinite.

Now, by Chepoi–Roller duality, the median algebra $M \cap (X^{(0)} \times Y^{(0)})$ is the 0-skeleton of a CAT(0) cube complex Z equipped with a proper and cocompact action $G \curvearrowright Z$. The restriction to M of the coordinate projection $X \times Y \rightarrow X$ gives a G -equivariant cubical map $\pi_X : Z \rightarrow X$ that is a median morphism on the 0-skeleton; in addition, since $K \subseteq M$ and $G \cdot K_X = X$, the map π_X is surjective. As mentioned before the statement of the theorem, this implies that π_X corresponds to a restriction quotient. The same argument shows that $G \curvearrowright Y$ is a restriction quotient of $G \curvearrowright Z$, concluding the proof. \square

3 | FIRST FORMS OF COARSE CUBICAL RIGIDITY

In this section, we prove Theorem A (Theorem 3.19) and a general version of Theorem B(1) for all graph products of finite groups (Theorem 3.5). Theorem D(1) immediately follows from the latter. In addition, in Subsection 3.3, we study cubical coarse medians on virtually abelian groups.

3.1 | Strongly cellular cubulations

In this subsection, we obtain a criterion guaranteeing that all strongly cellular, cocompact cubulations of certain groups induce the same coarse median structure (Proposition 3.2).

We first need the following simple lemma. This should be compared to Proposition 2.10 stating that centralisers of finite sets are median-cocompact in all cocompact cubulations. Here we exploit stronger assumptions on the centraliser and the action to obtain the stronger conclusion of convex-cocompactness.

Lemma 3.1. *Let G be a group and let $Z \leq G$ be the centraliser of a finite set of finite-order elements. Then Z is convex-cocompact in every strongly cellular, cocompact cubulation of G .*

Proof. Let $G \curvearrowright Y$ be a strongly cellular, cocompact cubulation of G . Let $f \in G$ have finite order. The subset $\text{Fix}(f) \subseteq Y$ is non-empty and, since the ℓ_2 -metric on Y is uniquely geodesic, it is ℓ_2 -convex. By definition of strongly cellular action, the subset $\text{Fix}(f)$ is also a subcomplex. We conclude that $\text{Fix}(f)$ is a convex subcomplex (recall that, for subcomplexes, ℓ_2 -convexity is equivalent to ℓ_1 -convexity [38]).

Observe that the action $Z_G(f) \curvearrowright \text{Fix}(f)$ is cocompact. In order to see this, consider a point $x \in \text{Fix}(f)$ and let $G_x \leq G$ be the finite subgroup fixing x . If $gx \in \text{Fix}(f)$ for some $g \in G$, then $g^{-1}fg \in G_x$. The set $\{g \in G \mid g^{-1}fg \in G_x\}$ is a finite union of right cosets of $Z_G(f)$. Hence $Z_G(f)$ acts cofinitely on every intersection between $\text{Fix}(f)$ and a G -orbit. Since there are only finitely many G -orbits of vertices in Y , it follows that $Z_G(f)$ acts cocompactly on $\text{Fix}(f)$.

In conclusion, for every finite-order element $f \in G$, the centraliser $Z_G(f)$ is convex-cocompact in Y . By Lemma 2.4, finite intersections of these subgroups are also convex-cocompact. \square

Combining the previous lemma with Theorem 2.17, we immediately obtain the following criterion. We emphasise that the action $G \curvearrowright X$ in Proposition 3.2 is *not* required to be itself strongly cellular.

Proposition 3.2. *Let G be a group with a cocompact cubulation $G \curvearrowright X$ where every hyperplane-stabiliser is commensurable to the centraliser of a finite set of finite-order elements of G . Then all strongly cellular, cocompact cubulations of G induce the same coarse median structure as $G \curvearrowright X$.*

Proof. Let $G \curvearrowright X$ be a cocompact cubulation where hyperplane-stabilisers are commensurable to centralisers of finite sets of finite-order elements. Let $G \curvearrowright Y$ be any strongly cellular, cocompact cubulation. Since convex-cocompactness is a commensurability invariant (Lemma 2.4), Lemma 3.1 implies that all stabilisers of hyperplanes of X are convex-cocompact in Y . Now, Theorem 2.17 implies that X and Y induce the same coarse median structure on G . \square

3.2 | Graph products of finite groups

In this subsection, we use Proposition 3.2 to deduce that all strongly cellular, cocompact cubulations of graph products of finite groups induce the same coarse median structure (Theorem 3.5). In particular, this will prove Theorem B(1) and Theorem D(1).

Every graph product of finite groups admits a particularly nice cocompact cubulation: its *graph-product complex*. This was shown in [24] (also see [53] and [34, Theorem 2.27]), but we briefly recall here the construction.[†]

[†] For graph products of order-2 groups (i.e. right-angled Coxeter groups), the complex that we are about to describe is the first cubical subdivision of what is usually known as the Davis complex.

Let G be the graph product determined by the data $(\Gamma, \{F_v\}_{v \in \Gamma})$. Here Γ is a finite simplicial graph and $\{F_v\}_{v \in \Gamma}$ is a collection of finite groups indexed by the vertices of Γ . The group G is the quotient of the free product of the F_v by the normal subgroup generated by the commutator sets $[F_w, F_{w'}]$ such that $[w, w']$ is an edge of Γ . If $\Lambda \subseteq \Gamma$ is a subgraph, we denote by G_Λ the subgroup of G generated by the union of the F_v with $v \in \Lambda$. If Λ is the empty set, then we let G_Λ denote the trivial subgroup of G .

The 0-skeleton of the graph-product complex \mathcal{D} is identified with the collection of cosets gG_c , where $g \in G$ and $c \subseteq \Gamma$ is a clique (possibly empty). We add edges $[gG_c, gG_{c'}]$ when $c \subseteq c'$ and $\#c = \#c' - 1$. Fixing $g \in G$, if $c \subseteq c'$ are cliques with $\#c = \#c' - k$, then the subgraph of \mathcal{D} spanned by the vertices $gG_{c''}$ with $c \subseteq c'' \subseteq c'$ is isomorphic to the 1-skeleton of a k -cube. We complete the construction of \mathcal{D} by glueing a cube of the appropriate dimension to each such subgraph, with obvious identifications ensuring that each cube is uniquely determined by its vertex set.

Note that G permutes left cosets by left multiplication. This gives a G -action on the 0-skeleton of \mathcal{D} , which naturally extends to a cellular action on the whole \mathcal{D} .

Proposition 3.3. *Let G be a graph product of finite groups.*

- (1) *The graph-product complex \mathcal{D} is a CAT(0) cube complex.*
- (2) *The action $G \curvearrowright \mathcal{D}$ is proper, cocompact and strongly cellular.*
- (3) *Hyperplane-stabilisers are precisely conjugates of the subgroups $G_{\text{lk}(v)}$ for $v \in \Gamma^{(0)}$.*

Proof. Checking (1) and (2) is straightforward and we leave it to the reader. We prove (3).

Let $\mathfrak{m} \subseteq \mathcal{D}$ be the hyperplane dual to the edge $e = [gG_c, gG_cF_v]$, where $g \in G$ is an element and $c \cup \{v\}$ is a clique. Suppose $f \subseteq \mathcal{D}$ is an edge opposite to e in a square containing e . Then f takes one of the following two forms:

- $f = [gG_cF_w, gG_cF_vF_w]$, where $c \cup \{v, w\}$ is a clique of Γ ;
- $f = [ghG_{c-\{w\}}, ghG_{c-\{w\}}F_v]$, where $w \in c$ and $h \in F_w$.

From this, we deduce that the edges of \mathcal{D} dual to \mathfrak{m} are precisely those of the form:

$$[g'G_{c'}, g'G_{c'}F_v],$$

where $c' \sqcup \{v\}$ is a clique in Γ and the element $g^{-1}g'$ is a product of elements of those F_w for which $[v, w]$ is an edge of Γ .

Now, if $h \in G$, we have $h\mathfrak{m} = \mathfrak{m}$ if and only if the edge $[hgG_c, hgG_cF_v]$ is of the above form. This happens exactly when $g^{-1}hg$ lies in the subgroup $G_{\text{lk}(v)}$, as required. □

Remark 3.4. For every vertex $v \in \Gamma$ and every $g \in F_v - \{1\}$, we have $G_{\text{lk}(v)} \leq Z_G(g) \leq G_{\text{st}(v)}$ (see e.g. [1, Theorem 32]). Since $G_{\text{st}(v)} = F_v \times G_{\text{lk}(v)}$, this shows that $G_{\text{lk}(v)}$ is always commensurable to the centraliser of a finite-order element.

Theorem 3.5. *Let G be a graph product of finite groups. All strongly cellular, cocompact cubulations of G induce the same coarse median structure on G .*

Proof. It suffices to apply Proposition 3.2 to the graph-product complex of G . Proposition 3.3(3) and Remark 3.4 ensure that hyperplane-stabilisers have the required form. □

Theorem D(1) immediately follows from Theorem 3.5 and the fact that the action on the graph-product complex is strongly cellular.

3.3 | Virtually abelian groups

In this subsection, we completely classify cubical coarse median structures on free abelian groups \mathbb{Z}^n (Proposition 3.9) and on products of infinite dihedrals D_∞^n (Proposition 3.11). While \mathbb{Z}^n has many cubical coarse medians — infinitely many orbits up to the natural $\text{Aut}(\mathbb{Z}^n)$ -action — there are only finitely many cubical coarse median structures on D_∞^n .

The first result is used in the next subsection to study cubical coarse medians on right-angled Artin groups and prove Theorem 3.19, which implies Theorem A. The second result is needed to deduce Corollary C from Theorem B and it also shows that, in the Conjecture from the Introduction, Item (b) implies Item (a) (Remark 3.12).

Lemma 3.6. *Let A be a virtually abelian group. Every cubical coarse median structure on A is induced by a proper cocompact A -action on the standard cubical tiling of \mathbb{R}^n , for some $n \geq 0$.*

Proof. Let $A \curvearrowright X$ be a proper cocompact action on a CAT(0) cube complex. Up to passing to an invariant convex subcomplex (which does not alter the induced coarse median structure), the cubical flat torus theorem [62, Theorem 3.6] allows us to assume that X is a product of quasi-lines $L_1 \times \dots \times L_n$.

By the CAT(0) flat torus theorem [5, Corollary II.7.2], there exists an A -invariant ℓ_2 -convex subspace $F \subseteq X$ that is ℓ_2 -isometric to a Euclidean flat. The projection of A to each factor L_i is an ℓ_2 -geodesic line α_i . Each α_i is also an ℓ_1 -geodesic by Lemma 2.1, hence it is a median subalgebra of L_i . It follows that $F = \alpha_1 \times \dots \times \alpha_n$ is an A -invariant median subalgebra of X .

In conclusion, the actions $A \curvearrowright X$ and $A \curvearrowright F$ induce the same coarse median structure on A . In addition, equipped with the restriction of the ℓ_1 -metric on X , the flat F is isometric to \mathbb{R}^n with the ℓ_1 -metric. Finally, since A acts on F discretely and permuting the orthogonal directions $\alpha_1, \dots, \alpha_n$, it preserves a tiling of \mathbb{R}^n by cuboids. Scaling to 1 all edge-lengths of these cuboids, we obtain the standard cubical tiling of \mathbb{R}^n without altering the coarse median structure. \square

As we are about to show, cubical coarse medians on \mathbb{Z}^n are parameterised by the following objects.

Definition 3.7. A *virtual basis* of \mathbb{Z}^n is a set $\{C_1, \dots, C_n\}$, where each $C_i \leq \mathbb{Z}^n$ is a maximal cyclic subgroup and $\langle C_1 \cup \dots \cup C_n \rangle$ has finite index in \mathbb{Z}^n . We denote by $\mathcal{VB}(\mathbb{Z}^n)$ the set of virtual bases.

An example of a virtual basis that is not a basis is given by $C_1 = \langle (1, 1) \rangle, C_2 = \langle (1, -1) \rangle$ in \mathbb{Z}^2 .

Remark 3.8. If $A \in \text{GL}_n(\mathbb{Q})$ is a matrix with integer entries and $\{C_1, \dots, C_n\}$ is a virtual basis of \mathbb{Z}^n , we obtain a new virtual basis $\{C'_1, \dots, C'_n\}$ by taking C'_i to be the maximal cyclic subgroup of \mathbb{Z}^n containing $A(C_i)$. Multiplying A by a constant does not alter the C'_i , so we obtain a natural transitive action:

$$\text{PGL}_n(\mathbb{Q}) \curvearrowright \mathcal{VB}(\mathbb{Z}^n).$$

As another way of making sense of this action, note that $\mathcal{VB}(\mathbb{Z}^n)$ can be equivalently defined as the collection of all cardinality- n general-position subsets of $\mathbb{Q}\mathbb{P}^{n-1}$. The action $\mathrm{PGL}_n(\mathbb{Q}) \curvearrowright \mathcal{VB}(\mathbb{Z}^n)$ is then directly induced by the standard action $\mathrm{PGL}_n(\mathbb{Q}) \curvearrowright \mathbb{Q}\mathbb{P}^{n-1}$ by projective automorphisms.

The stabiliser of the standard basis of \mathbb{Z}^n is the subgroup $N \leq \mathrm{PGL}_n(\mathbb{Q})$ generated by diagonal matrices and permutation matrices (i.e. the normaliser of the diagonal subgroup). This shows that we can naturally identify:

$$\mathcal{VB}(\mathbb{Z}^n) \cong \mathrm{PGL}_n(\mathbb{Q})/N \cong \mathrm{GL}_n(\mathbb{Q})/\overline{N},$$

where $\overline{N} \leq \mathrm{GL}_n(\mathbb{Q})$ is again the normaliser of the diagonal subgroup.

Proposition 3.9. *The sets $\mathcal{CM}_{\square}(\mathbb{Z}^n)$ and $\mathcal{VB}(\mathbb{Z}^n)$ are in 1-to-1 correspondence, equivariantly with respect to the action $\mathrm{Out}(\mathbb{Z}^n) \curvearrowright \mathcal{CM}_{\square}(\mathbb{Z}^n)$ and the left-multiplication action $\mathrm{GL}_n(\mathbb{Z}) \curvearrowright \mathrm{GL}_n(\mathbb{Q})/\overline{N}$.*

The correspondence pairs each coarse median structure $[\mu]$ with the set $\{C_1, \dots, C_n\}$ of maximal, cyclic, $[\mu]$ -quasi-convex subgroups of \mathbb{Z}^n .

Proof. By Lemma 3.6, every cubical coarse median on \mathbb{Z}^n is induced by a \mathbb{Z}^n -action on the standard cubical tiling of \mathbb{R}^n , which we denote by S . We have $\mathrm{Aut}(S) = T \rtimes \mathrm{O}(n, \mathbb{Z})$, where $T \simeq \mathbb{Z}^n$ is the translation subgroup, and $\mathrm{O}(n, \mathbb{Z})$ is the group of signed permutation matrices.

Up to passing to a finite-index subgroup of \mathbb{Z}^n (for instance, the intersection of all subgroups of index $\leq \#\mathrm{O}(n, \mathbb{Z})$), we can assume that all our actions $\mathbb{Z}^n \curvearrowright S$ factor through T .

We equip T with its standard basis t_1, \dots, t_n (where each t_i translates in only one coordinate direction), as well as the coordinate-wise median operator m associated with this basis. Note that $[m]$ is precisely the coarse median structure on T induced by the action $T \curvearrowright S$.

The above discussion shows that every cubical coarse median structure $[\mu]$ on \mathbb{Z}^n can be obtained by pulling back $[m]$ via an embedding $\iota : \mathbb{Z}^n \hookrightarrow T$. To each such structure $[\mu]$, we associate the virtual basis $\{\iota^{-1}(\langle t_1 \rangle), \dots, \iota^{-1}(\langle t_n \rangle)\} \in \mathcal{VB}(\mathbb{Z}^n)$, which is precisely the set of maximal, cyclic, $[\mu]$ -quasi-convex subgroups of \mathbb{Z}^n .

It is straightforward to see that $[\mu]$ is completely determined by this element of $\mathcal{VB}(\mathbb{Z}^n)$. Finally, every element of $\mathcal{VB}(\mathbb{Z}^n)$ arises in this way, as every action by translations $H \curvearrowright \mathbb{R}^n$ of a finite-index subgroup $H \leq \mathbb{Z}^n$ can be extended to a \mathbb{Z}^n -action by translations, which will preserve a tiling of \mathbb{R}^n by sufficiently fine cubes. This concludes the proof. □

Having classified cubical coarse median structures on \mathbb{Z}^n , we move on to the case of D_{∞}^n .

Remark 3.10. There are two obvious cubical coarse median structures on D_{∞}^2 . The first (or ‘standard’) structure is induced by the action on the Davis complex: It is the action on the standard square tiling of \mathbb{R}^2 with each reflection axis parallel either to the x -axis or to the y -axis. The second (or ‘ $\frac{\pi}{4}$ -rotated’) structure is induced by the action on the standard tiling of \mathbb{R}^2 with all reflection axes meeting the x - and y -axes at an angle of $\frac{\pi}{4}$ or $\frac{3\pi}{4}$. This is just the first action rotated by $\frac{\pi}{4}$ and, unlike it, it is not strongly cellular. Indeed, recall from Theorem 3.5 that all strongly cellular cubulations induce the standard coarse median structure.

The next result shows that there are no other cubical coarse median structures on D_∞^2 . Similarly, all cubical coarse medians on D_∞^n are simply obtained by choosing the ‘ $\frac{\pi}{4}$ -rotated’ structure on some D_∞^2 -factors and the ‘standard’ one on the rest.

Proposition 3.11. *The set $\mathcal{CM}_\square(D_\infty^n)$ is finite. More precisely, for every $[\mu] \in \mathcal{CM}_\square(D_\infty^n)$, there is a product splitting $D_\infty^n = A_1 \times \dots \times A_k \times B_1 \times \dots \times B_{n-2k}$ for some $0 \leq k \leq \frac{n}{2}$, where all A_i and B_j are $[\mu]$ -quasi-convex, we have $A_i \simeq D_\infty^2$ and $B_j \simeq D_\infty$, and the restriction of $[\mu]$ to each A_i is the $\frac{\pi}{4}$ -rotated structure.*

Proof. Let again S denote the standard cubical tiling of \mathbb{R}^n . By Lemma 3.6, every cubical coarse median on D_∞^n arises from a proper cocompact action $D_\infty^n \curvearrowright S$. The latter corresponds to a homomorphism $\rho : D_\infty^n \rightarrow \text{Aut}(S) = T \rtimes O(n, \mathbb{Z})$ with finite-index image, where $T \simeq \mathbb{Z}^n$.

Note that ρ must be faithful. Indeed, the kernel of ρ is necessarily finite, and D_∞ (and hence D_∞^n) does not have any non-trivial finite normal subgroups.

Now, choose a reflection r_i in each factor of D_∞^n and consider the subgroup $R := \langle r_1, \dots, r_n \rangle \simeq (\mathbb{Z}/2\mathbb{Z})^n$. Let $\pi : \text{Aut}(S) \rightarrow O(n, \mathbb{Z})$ be the projection with $\ker \pi = T$. Since T is torsion-free and ρ is injective, the homomorphism $\pi\rho$ must be injective on R . Thus, $\pi\rho(R)$ is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^n$ and it is a subgroup of $O(n, \mathbb{Z})$, the group of signed permutation matrices.

Every element of $\pi\rho(R)$ has order 2, hence it is $O(n, \mathbb{Z})$ -conjugate to a block-diagonal matrix, with each block chosen from:

$$(\pm 1), \quad \pm \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Since $\pi\rho(R)$ is abelian, its elements can be simultaneously block-diagonalised. Since $\pi\rho(R) \simeq (\mathbb{Z}/2\mathbb{Z})^n$, it follows that $\pi\rho(R)$ is $O(n, \mathbb{Z})$ -conjugate to a subgroup of the form:

$$\left(\begin{array}{ccccccc} R_1 & & & & & & \\ & \ddots & & & & & \\ & & R_k & & & & \\ & & & \pm 1 & & & \\ & & & & \ddots & & \\ & & & & & \ddots & \\ & & & & & & \pm 1 \end{array} \right),$$

where $0 \leq k \leq n/2$ and each $R_\ell \simeq (\mathbb{Z}/2\mathbb{Z})^2$ is the group of 2×2 matrices $\left\{ \pm \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \pm \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right\}$.

Finally, the coarse median structure that ρ induces on D_∞^n is completely determined by the datum of which pairs of indices $1 \leq i, j \leq n$ are such that $\langle r_i, r_j \rangle$ corresponds to one of the 2×2 blocks R_ℓ in the above decomposition. Indeed, this datum determines which infinite cyclic subgroups of D_∞^n translate in exactly one coordinate direction in $S \simeq \mathbb{R}^n$ and are thus convex-cocompact. And the latter determines the coarse median structure by applying Proposition 3.9 to a free abelian finite-index subgroup of D_∞^n . □

Remark 3.12. Item (b) implies Item (a) in the Conjecture from the Introduction.

Indeed, if W_Γ is a right-angled Coxeter group, W_Γ has only finitely many maximal virtually abelian parabolic subgroups P_1, \dots, P_k up to conjugacy. Each of them has a finite-index subgroup $P'_i \simeq D_\infty^{m_i}$ for some $m_i \geq 0$. By Lemma 2.4(3) and Proposition 2.14, all P_i and P'_i are $[\mu]$ -quasi-convex with respect to each $[\mu] \in \mathcal{CM}_\square(W_\Gamma)$, so they inherit cubical coarse median structures from $[\mu]$.

By Proposition 3.11, there are only finitely many possible restrictions of $[\mu]$ to the P_i . Thus, if these restrictions completely determine $[\mu]$, then the set $\mathcal{CM}_\square(W_\Gamma)$ is finite.

Remark 3.13. The techniques used in this subsection can be similarly used to completely describe the set of cubical coarse medians $\mathcal{CM}_\square(A)$ for any virtually abelian group A .

Indeed, passing to a normal free abelian subgroup, we can always write

$$1 \longrightarrow \mathbb{Z}^n \longrightarrow A \longrightarrow F \longrightarrow 1,$$

for some $n \geq 1$ and some finite group F . Conjugation gives a monodromy homomorphism $\rho : F \rightarrow \text{Out}(\mathbb{Z}^n) = \text{GL}_n(\mathbb{Z})$. In particular, the subgroup $\rho(F) \leq \text{GL}_n(\mathbb{Z})$ acts on $\mathcal{CM}_\square(\mathbb{Z}^n)$, which is identified with $\mathcal{VB}(\mathbb{Z}^n) \cong \text{GL}_n(\mathbb{Q})/\overline{N}$ (recalling Proposition 3.9 and Remark 3.8).

Now, the space $\mathcal{CM}_\square(A)$ is naturally identified with the set of $\rho(F)$ -fixed points in $\mathcal{CM}_\square(\mathbb{Z}^n)$. In one direction, every cocompact cubulation of A determines such a fixed point by restriction to the finite-index subgroup \mathbb{Z}^n . Conversely, if a cubical coarse median structure on a normal finite-index subgroup is preserved by the conjugacy action of the whole group, then it comes from a cocompact cubulation of the whole group; this was shown in [31, Corollary G] using very similar ideas to those in the proof of Theorem 2.17 above.

Thus, the determination of $\mathcal{CM}_\square(A)$ boils down to describing the set of fixed points of a finite subgroup of $\text{GL}_n(\mathbb{Z})$ acting on the coset space $\text{GL}_n(\mathbb{Q})/\overline{N}$, where $\overline{N} = D \rtimes \text{O}(n, \mathbb{Z})$ is the normaliser of the subgroup D of diagonal matrices, and $\text{O}(n, \mathbb{Z})$ is the group of signed permutation matrices.

As a special case of this, we see that A is cocompactly cubulated if and only if $\rho(F)$ fixes at least one point in $\text{GL}_n(\mathbb{Q})/\overline{N}$. Equivalently, $\rho(F) \leq \text{GL}_n(\mathbb{Z})$ can be conjugated into \overline{N} by an element of $\text{GL}_n(\mathbb{Q})$, in which case, $\rho(F)$ can even be conjugated into $\text{O}(n, \mathbb{Z})$. This recovers Hagen’s characterisation of cocompactly cubulated crystallographic groups from [37].

The following example, mentioned in the Introduction, shows that Proposition 2.10 cannot be improved. In general, we cannot expect centralisers to be convex-cocompact even if we have the freedom to choose the cubulation.

Example 3.14. Consider $G = \mathbb{Z}^n \rtimes \mathfrak{S}_n$, where the symmetric group \mathfrak{S}_n acts by permuting the standard basis of \mathbb{Z}^n . For $n = 3$ or $n \geq 5$, the centraliser $Z_G(\mathfrak{S}_n)$ is not convex-cocompact in any cocompact cubulation of G .

Indeed, $Z_G(\mathfrak{S}_n)$ is the infinite cyclic subgroup generated by the element $v := (1, 1, \dots, 1)$. If $\langle v \rangle$ were convex-cocompact in a cocompact cubulation of G , then, by Proposition 3.9, there would exist a virtual basis $\langle v \rangle, C_1, \dots, C_{n-1}$ of \mathbb{Z}^n that is permuted by \mathfrak{S}_n . In particular, the subgroups C_1, \dots, C_{n-1} would be permuted by \mathfrak{S}_n , since v is fixed. However, for every maximal infinite cyclic subgroup $C < \mathbb{Z}^n$ with $C \neq \langle v \rangle$, the orbit $\mathfrak{S}_n \cdot C$ contains at least n distinct subgroups.

Let us prove this last statement. Suppose that C is generated by an element (x_1, \dots, x_n) . If the absolute values $|x_i|$ are not all equal, then, without loss of generality, there exist an integer $a > 0$ and an index $1 \leq k < n$ such that $|x_i| = a$ for $1 \leq i \leq k$ and $|x_i| < a$ for $k < i \leq n$. Since \mathfrak{S}_n acts k -transitively on $\{1, \dots, n\}$, it follows that the orbit $\mathfrak{S}_n \cdot C$ contains at least $\binom{n}{k} \geq n$ distinct subgroups. If instead all the $|x_i|$ are equal, we can assume that there exists $1 \leq k < n$ such that $x_i = 1$ for $1 \leq i \leq k$ and $x_i = -1$ for $k < i \leq n$. In this case, the orbit $\mathfrak{S}_n \cdot C$ has cardinality $\binom{n}{k}$ if $n \neq 2k$ and $\frac{1}{2} \cdot \binom{n}{k}$ if $n = 2k$. When $n \geq 5$, even this last quantity is $\geq n$, completing the proof.

3.4 | Right-angled Artin groups

Let A_Γ be a right-angled Artin group. In this subsection, we study the set $\mathcal{CM}_\square(A_\Gamma)$. The main result is Theorem 3.19, which implies Theorem A from the Introduction, but also concerns right-angled Artin groups that are not twistless.

Let $A_\Gamma \curvearrowright \mathcal{X}_\Gamma$ be the standard action on the universal cover of the Salvetti complex. Let $[\mu_\Gamma] \in \mathcal{CM}_\square(A_\Gamma)$ be the induced coarse median structure on A_Γ . We refer to $[\mu_\Gamma]$ as the *standard* coarse median structure on A_Γ .

We will obtain Theorem 3.19 by adapting the proof of Proposition 3.2. The main differences are that: (1) we cannot exploit torsion and (2) stabilisers of hyperplanes of the Salvetti complex are not commensurable to centralisers. On the other hand, no analogue of the ‘strongly cellular’ assumption will be required. The following notion will play an important role.

Definition 3.15. A cocompact cubulation $A_\Gamma \curvearrowright X$ has *decomposable flats* if every maximal abelian subgroup $A \leq A_\Gamma$ admits a basis a_1, \dots, a_k with each $\langle a_i \rangle$ convex-cocompact in X .

An example of a cubulation that fails to have this property is provided by the \mathbb{Z}^2 -action on the standard tiling of \mathbb{R}^2 where the standard generators of \mathbb{Z}^2 translate by (1,1) and (1, -1), respectively.

Remark 3.16. Here is some motivation for considering the above condition.

- (1) The standard action $A_\Gamma \curvearrowright \mathcal{X}_\Gamma$ on the universal cover of the Salvetti complex has decomposable flats, since each $v \in \Gamma$ generates a convex-cocompact subgroup.
- (2) Whether or not a given cocompact cubulation $A_\Gamma \curvearrowright X$ has decomposable flats only depends on the induced coarse median structure on A_Γ , since this is true of convex-cocompactness.
- (3) The action $\text{Out}(A_\Gamma) \curvearrowright \mathcal{CM}_\square(A_\Gamma)$ preserves the property of having decomposable flats.

Indeed, let $A_\Gamma \curvearrowright X$ be a cocompact cubulation with decomposable flats. Consider $\varphi \in \text{Aut}(A_\Gamma)$ and denote by $A_\Gamma \curvearrowright X^\varphi$ the action on X precomposed with φ . If $A \leq A_\Gamma$ is a maximal abelian subgroup, then so is $\varphi(A)$, which then admits a basis a_i of convex-cocompact elements for the action $A_\Gamma \curvearrowright X$. It follows that the elements $\varphi^{-1}(a_i)$ form a basis of A and are convex-cocompact for $A_\Gamma \curvearrowright X^\varphi$. Thus, $A_\Gamma \curvearrowright X^\varphi$ has decomposable flats.

- (4) For a general cocompact cubulation $G \curvearrowright X$ of a general group and a convex-cocompact abelian subgroup $A \leq G$, it is only possible to find a basis a_1, \dots, a_k of a *finite-index* subgroup of A such that each $\langle a_i \rangle$ is convex-cocompact in X . Compare Proposition 3.9.

Lemma 3.17. *Let $A_\Gamma \curvearrowright X$ be a cocompact cubulation with decomposable flats. Let $A \leq A_\Gamma$ be a (not necessarily maximal) abelian subgroup that is both convex-cocompact in X and closed under taking roots. Then A admits a basis a_1, \dots, a_k with each $\langle a_i \rangle$ convex-cocompact in X .*

Proof. Let $A' \leq A_\Gamma$ be a maximal abelian subgroup containing A and let x_1, \dots, x_n be a basis of A' such that each $\langle x_i \rangle$ is convex-cocompact. If an element $x_{i_1}^{a_1} \cdot \dots \cdot x_{i_s}^{a_s}$ with $a_1, \dots, a_s \neq 0$ lies in A , then the fact that A is convex-cocompact implies that A must contain non-trivial powers of x_{i_1}, \dots, x_{i_s} (for instance, by the argument for [31, Lemma 3.16]). Since A is closed under taking roots, it must then contain x_{i_1}, \dots, x_{i_s} themselves.

This shows that A contains all x_i required to write any of its elements. They provide the required basis of A . □

The automorphism group $\text{Aut}(A_\Gamma)$ is generated by elementary automorphisms described by Laurence and Servatius [46, 56]. In the terminology of [22, Section 2.2], these are known as *graph automorphisms*, *inversions*, *partial conjugations* and *transvections*; the latter are in turn divided into *folds* and *twists*. The only automorphisms that will be important for us are twists.

Definition 3.18. If $v, w \in \Gamma$ are distinct vertices with $\text{st}(v) \subseteq \text{st}(w)$, there is a well-defined automorphism of A_Γ that fixes all generators in $\Gamma - \{v\}$ and maps $v \mapsto vw$. We denote this automorphism by $\tau_{v,w}$ and refer to it as a *twist*. The *twist subgroup* $T(A_\Gamma) \leq \text{Aut}(A_\Gamma)$ is the subgroup generated by all twists.

We say that A_Γ is *twistless* if the twist subgroup $T(A_\Gamma)$ is trivial. Equivalently, we never have $\text{st}(v) \subseteq \text{st}(w)$ for distinct vertices $v, w \in \Gamma$.

The right-angled Artin group A_Γ is twistless precisely when the entire group $\text{Out}(A_\Gamma)$ fixes the standard coarse median structure $[\mu_\Gamma]$ [31, Proposition A]. This is always the case when $\text{Out}(A_\Gamma)$ is finite.

We are now ready to state the main result of this subsection.

Theorem 3.19. *Let A_Γ be a right-angled Artin group.*

- (1) *The twist group $T(A_\Gamma)$ acts transitively and with finite stabilisers on the subset of $\mathcal{CM}_\square(A_\Gamma)$ corresponding to cubulations with decomposable flats.*
- (2) *If A_Γ is twistless, all cocompact cubulations of A_Γ have decomposable flats.*

Remark 3.20. Let $U(A_\Gamma) \leq \text{Aut}(A_\Gamma)$ be the subgroup generated by all Laurence–Servatius generators except for twists, that is, graph automorphisms, inversions, folds and partial conjugations (including inner automorphisms). This is known as the *untwisted subgroup* [22]. As shown in [31], $U(A_\Gamma)$ is precisely the stabiliser of the standard coarse median structure on A_Γ .

It is clear that $U(A_\Gamma)$ and $T(A_\Gamma)$ generate $\text{Aut}(A_\Gamma)$, but Theorem 3.19(1) implies that there is even a product decomposition $\text{Aut}(A_\Gamma) = T(A_\Gamma) \cdot U(A_\Gamma)$ and that the intersection $T(A_\Gamma) \cap U(A_\Gamma)$ is finite. Neither $U(A_\Gamma)$ nor $T(A_\Gamma)$ is normal in general, so this is not (virtually) a semi-direct product. The existence of this splitting also follows from [2, Corollary 7.12].

We begin the proof of Theorem 3.19 with a few elementary observations on the twist group.

Remark 3.21. We have $\text{Aut}(\mathbb{Z}^n) = \langle T(\mathbb{Z}^n), \text{O}(n, \mathbb{Z}) \rangle$, where $T(\mathbb{Z}^n)$ is the twist subgroup, while $\text{O}(n, \mathbb{Z}) \simeq (\mathbb{Z}/2\mathbb{Z})^n \rtimes \mathfrak{S}_n$ is the group of signed permutation matrices (here \mathfrak{S}_n denotes the symmetric group). The finite subgroup $\text{O}(n, \mathbb{Z})$ permutes the elements of the standard basis of \mathbb{Z}^n , possibly inverting some of them. Since $\text{O}(n, \mathbb{Z})$ normalises $T(\mathbb{Z}^n)$, we have $\text{Aut}(\mathbb{Z}^n) = T(\mathbb{Z}^n) \cdot \text{O}(n, \mathbb{Z})$. The intersection $T(\mathbb{Z}^n) \cap \text{O}(n, \mathbb{Z})$ is non-trivial for $n \geq 2$.

For every vertex $v \in \Gamma$, denote by $\kappa(v) \subseteq \Gamma$ the intersection of all maximal cliques of Γ that contain v . Equivalently, it is easy to see that we have:

$$\kappa(v) = \left\{ w \in \Gamma^{(0)} \mid \text{st}(v) \subseteq \text{st}(w) \right\}. \tag{3.1}$$

The action $T(A_\Gamma) \curvearrowright A_\Gamma$ can be completely described in terms of the behaviour on the free abelian subgroups $A_{\kappa(v)} \leq A_\Gamma$. Indeed, by Equation (3.1) and the definition of twists, we have

$\varphi(A_{\kappa(v)}) = A_{\kappa(v)}$ for every $v \in \Gamma$ and every $\varphi \in T(A_\Gamma)$. Clearly, φ is the trivial automorphism if and only if it has trivial restriction to all the free abelian subgroups $A_{\kappa(v)}$, so $T(A_\Gamma)$ embeds in the finite product $\prod_v T(A_{\kappa(v)})$. It follows that the $T(A_\Gamma)$ -stabiliser of the standard coarse median structure on A_Γ is finite, since this is true in the free abelian case (Remark 3.21).

We now prove a few lemmas from which Theorem 3.19 can be quickly deduced.

Lemma 3.22. *Let $A_\Gamma \curvearrowright X$ be a cocompact cubulation with decomposable flats. Let $v \in \Gamma$ be a vertex such that the subgroup $\langle v \rangle$ is not convex-cocompact in $A_\Gamma \curvearrowright X$ and such that the star $\text{st}(v) \subseteq \Gamma$ is maximal among stars of such vertices. Then there exists an automorphism $\varphi_v \in T(A_\Gamma)$ such that:*

- $\langle \varphi_v(v) \rangle$ is convex-cocompact in $A_\Gamma \curvearrowright X$;
- $\varphi_v(w) = w$ for all $w \in \Gamma$ such that $\langle w \rangle$ is convex-cocompact in $A_\Gamma \curvearrowright X$.

Proof. If $c \subseteq \Gamma$ is a maximal clique, then A_c is convex-cocompact in X by Lemma 2.4(3). Since $\kappa(v)$ is the intersection of all maximal cliques containing v , it follows from Lemma 2.4(1) that $A_{\kappa(v)}$ is convex-cocompact in X .

Write $\kappa(v) = \{v\} \sqcup \{w_1, \dots, w_k\}$ and choose $0 \leq s \leq k$ so that $\langle w_{s+1} \rangle, \dots, \langle w_k \rangle$ are convex-cocompact in X , while $\langle w_1 \rangle, \dots, \langle w_s \rangle$ are not. Also set $\bar{\kappa}(v) := \{v\} \sqcup \{w_1, \dots, w_s\}$ and note that, by maximality of $\text{st}(v)$ and Equation (3.1), all vertices in $\bar{\kappa}(v)$ have the same star.

Since the action $A_\Gamma \curvearrowright X$ has decomposable flats, Lemma 3.17 guarantees that there exists a basis x_0, \dots, x_k of $A_{\kappa(v)}$ such that all $\langle x_i \rangle$ are convex-cocompact in X . Up to permuting and inverting the x_i , we can assume $x_i = w_i$ for $s + 1 \leq i \leq k$, since $A_{\kappa(v)} \simeq \mathbb{Z}^{k+1}$ has exactly $k + 1$ convex-cocompact directions by Proposition 3.9.

Let $\psi \in \text{Aut}(A_{\kappa(v)})$ map the basis $\{v, w_1, \dots, w_k\}$ to the basis $\{x_0, \dots, x_k\}$ while fixing the elements w_{s+1}, \dots, w_k . Recalling that the twist subgroup of a free abelian group acts transitively on bases (Remark 3.21), we can take $\psi \in T(A_{\kappa(v)})$. Moreover, ψ is a product of twists $\tau_{a,b}$ with $a \in \bar{\kappa}(v)$ and $b \in \kappa(v)$, in the notation of Definition 3.18. Since all elements of $\bar{\kappa}(v)$ have the same star in Γ , which is contained in the star of each element of $\kappa(v)$, each of the twists $\tau_{a,b}$ extends to a twist of the whole A_Γ .

In conclusion, the product of these twists of A_Γ is an automorphism $\varphi_v \in T(A_\Gamma)$ fixing $\Gamma - \{v, w_1, \dots, w_s\}$ and mapping v to the convex-cocompact element $\psi(v) = x_0$. This is the required automorphism, concluding the proof. □

Lemma 3.23. *Let $A_\Gamma \curvearrowright X$ be a cocompact cubulation with decomposable flats. Then there exists $\varphi \in T(A_\Gamma)$ such that each subgroup $\langle \varphi(v) \rangle$ with $v \in \Gamma$ is convex-cocompact in $A_\Gamma \curvearrowright X$.*

Proof. We prove the lemma by induction on the number $N \geq 0$ of standard generators that fail to be convex-cocompact in $A_\Gamma \curvearrowright X$. The base step $N = 0$ is trivial, taking $\varphi = \text{id}_{A_\Gamma}$.

For the inductive step, suppose that the lemma has been proved for cubulations in which at most $N - 1$ standard generators are not convex-cocompact. Pick a generator $v \in \Gamma$ such that $\langle v \rangle$ is not convex-cocompact in $A_\Gamma \curvearrowright X$ and such that $\text{st}(v)$ is maximal among stars of vertices of Γ with this property. Let $A_\Gamma \curvearrowright Y$ denote the action $A_\Gamma \curvearrowright X$ precomposed with the automorphism $\varphi_v \in T(A_\Gamma)$ provided by Lemma 3.22.

By Remark 3.16(3), $A_\Gamma \curvearrowright Y$ is again a cocompact cubulation with decomposable flats. In addition, by the choice of φ_v , there are at most $N - 1$ standard generators that are not convex-cocompact in $A_\Gamma \curvearrowright Y$. By the inductive hypothesis, there exists an automorphism $\psi \in T(A_\Gamma)$ such

that each element $\psi(x)$ with $x \in \Gamma$ is convex-cocompact in $A_\Gamma \curvearrowright Y$. In other words, the elements $\varphi_v \psi(x)$ are all convex-cocompact in $A_\Gamma \curvearrowright X$. We conclude by setting $\varphi := \varphi_v \psi$. \square

Lemma 3.24. *Let $A_\Gamma \curvearrowright X$ be a cocompact cubulation. If $\langle v \rangle$ is convex-cocompact in X for every $v \in \Gamma^{(0)}$, then the induced coarse median structure on A_Γ is the standard one.*

Proof. Hyperplane-stabilisers of the standard A_Γ -action on the universal cover of the Salvetti complex are of the form $A_{\text{lk}(v)}$ with $v \in \Gamma^{(0)}$, so our goal is to show that these subgroups are all convex-cocompact for $A_\Gamma \curvearrowright X$. Then we can apply Theorem 2.17.

First, we show that all centralisers $A_{\text{st}(v)} = Z_{A_\Gamma}(v)$ are convex-cocompact in X . In order to see this, note that v^n must act non-transversely on X for some $n \geq 1$ (e.g. by [3, Proposition 2.7(5)]). The subgroup $\langle v \rangle$ is convex-cocompact by hypothesis, so Lemma 2.4(4) implies that $Z_{A_\Gamma}(v^n)$ is convex-cocompact in X (since it is commensurable to $N_{A_\Gamma}(v^n)$). Finally, since we are in a right-angled Artin group, we have $Z_{A_\Gamma}(v^n) = Z_{A_\Gamma}(v)$.

Now, we know that both $\langle v \rangle$ and $A_{\text{st}(v)} = \langle v \rangle \times A_{\text{lk}(v)}$ are convex-cocompact in X . In addition, $A_{\text{lk}(v)}$ is generated by elements $w \in \Gamma$, which all generate convex-cocompact subgroups by hypothesis. Lemma 2.6 thus implies that $A_{\text{lk}(v)}$ is convex-cocompact in X , as required. \square

Proof of Theorem 3.19. The combination of Lemmas 3.23 and 3.24 shows that $T(A_\Gamma)$ acts transitively on the set of coarse median structures induced by cocompact cubulations with decomposable flats. Proving that $T(A_\Gamma)$ acts with finite stabilisers then amounts to proving that the $T(A_\Gamma)$ -stabiliser of the *standard* structure is finite. As observed above, this follows from the fact that $T(A_\Gamma)$ embeds in the finite product of the twist groups of the convex-cocompact free abelian groups $A_{\kappa(v)}$, each of which has only a finite subgroup stabilising the standard coarse median structure on $A_{\kappa(v)}$. This proves part (1).

Regarding part (2), note that, when A_Γ is twistless, all cliques $\kappa(v) \subseteq \Gamma$ are singletons, because of the description of $\kappa(v)$ given in Equation (3.1). We have already observed that Lemma 2.4 implies that $A_{\kappa(v)}$ is convex-cocompact in all cocompact cubulations of A_Γ . Thus, each $\langle v \rangle$ is convex-cocompact in all cocompact cubulations of A_Γ , which implies that all these cubulations have decomposable flats. \square

Theorem A might (wrongly) lead us to believe that, under the right assumptions, right-angled Artin groups have only finitely many distinct cocompact cubulations. A natural guess would be that this holds when $\text{Out}(A_\Gamma)$ is finite, if we restrict to cubulations with no free faces. After all, it was shown in [26, Proposition C] that many Burger–Mozes–Wise groups have a *unique* cubulation with no free faces (while they have many essential ones).

The next example shows that such guesses are incorrect and right-angled Artin groups are way too flexible for this kind of result to hold.

Example 3.25. Let Γ be a hexagon, as in Figure 2. The group $\text{Out}(A_\Gamma)$ is finite and A_Γ satisfies coarse cubical rigidity by Theorem A. Nevertheless, A_Γ has infinitely many 2-dimensional cocompact cubulations with no free faces, as we are about to show.

The rough idea is that there should exist cocompact cubulations of $A_\Gamma \curvearrowright X$ such that the action of each free group $\langle x_{i-1}, x_{i+1} \rangle$ on its own essential core in X is an arbitrary tree in the Outer Space of $\langle x_{i-1}, x_{i+1} \rangle$ (modulo some compatibility conditions). We only prove a special case of this, where four of these trees are standard and the remaining two have a simple form.




FIGURE 2 Top left: the hexagon graph Γ . Right: the $n + 6$ squares making up the square complex C_n . Bottom left: the reduced link of the single vertex of C_n . For clarity, the vertex of $\text{lk}(v)$ labelled e_1 has degree 2, while the vertices labelled e_2, \dots, e_n all have degree 1.

For each $n \geq 0$, let C_n be the finite square complex described on the right-hand side of Figure 2. Note that C_n has a single vertex (call it v) and its 1-skeleton is a rose with $n + 6$ petals, which we name a_1, \dots, a_6 and e_1, \dots, e_n . It is clear that C_n is 2-dimensional and does not have any free faces. In addition, C_0 is simply the Salvetti complex of A_Γ .

Each complex C_n is non-positively curved. In order to see this, it suffices to check that the link of the only vertex v does not contain any 3-cycles. Note that $\text{lk}(v)$ contains a pair of non-adjacent vertices for each edge of C_n . The graph $\overline{\text{lk}}(v)$ obtained from $\text{lk}(v)$ by collapsing these pairs of vertices (and identifying any resulting pairs of edges with the same endpoints) is pictured in Figure 2, below on the left. Any 3-cycle in $\text{lk}(v)$ would give rise to a 3-cycle in $\overline{\text{lk}}(v)$, but the latter does not contain any, proving that C_n is non-positively curved.

The description of C_n given in Figure 2 yields the following presentation for its fundamental group:

$$\pi_1(C_n, v) = \langle a_1, \dots, a_6 \mid [a_1, a_2], [a_2, a_3], [a_3, a_4], [a_4, a_5], [a_5, a_6], [a_6, a_1 a_3^n] \rangle.$$

This can be rewritten as follows, replacing a_1 with $\bar{a}_1 := a_1 a_3^n$:

$$\begin{aligned} \pi_1(C_n, v) &= \langle \bar{a}_1, a_2, \dots, a_6 \mid [\bar{a}_1 a_3^{-n}, a_2], [a_2, a_3], [a_3, a_4], [a_4, a_5], [a_5, a_6], [a_6, \bar{a}_1] \rangle \\ &= \langle \bar{a}_1, a_2, \dots, a_6 \mid [\bar{a}_1, a_2], [a_2, a_3], [a_3, a_4], [a_4, a_5], [a_5, a_6], [a_6, \bar{a}_1] \rangle, \end{aligned}$$

where we have used the fact that $[a_2, a_3]$ is one of the relators. This shows that there exists an isomorphism $\varphi_n : A_\Gamma \rightarrow \pi_1(C_n, v)$ with $\varphi(x_1) = a_1 a_3^n$ and $\varphi(x_i) = a_i$ for $2 \leq i \leq 6$.

The deck-transformation actions $A_\Gamma \curvearrowright \tilde{C}_n$ given by the isomorphisms φ_n are the required 2-dimensional cocompact cubulations of A_Γ with no free faces.

In order to check that these cubulations are truly pairwise distinct, we consider the essential core of the free group $\langle x_1, x_3 \rangle$. First, notice that the loops a_1, a_3 form a *convex* rose in C_n , which lifts to a convex tree in \tilde{C}_n . Thus, the action of the free group $\langle a_1, a_3 \rangle$ on its essential core in \tilde{C}_n coincides with the action of $\langle a_1, a_3 \rangle$ on its standard Cayley graph (corresponding to the generating set $\{a_1, a_3\}$). Now, since $\varphi_n(x_1) = a_1 a_3^n$ and $\varphi_n(x_3) = a_3$, the action of the free group $\langle x_1, x_3 \rangle$ on its essential core in \tilde{C}_n is rather the point of Outer Space obtained by twisting the standard Cayley graph (corresponding to the generating set $\{x_1, x_3\}$) by the outer automorphism $(x_1, x_3) \mapsto (x_1 x_3^n, x_3)$. This completes our example.

One could also wonder about the essential cores of the other free groups $\langle x_{i-1}, x_{i+1} \rangle$ in \tilde{C}_n . For $\langle x_2, x_4 \rangle, \langle x_3, x_5 \rangle, \langle x_4, x_6 \rangle, \langle x_6, x_2 \rangle$, the essential core is just the standard Cayley graph, whereas for $\langle x_5, x_1 \rangle$ it is the subdivision of the standard Cayley graph where one of the two orbits of edges gets subdivided into $n + 1$ orbits of edges. To see the latter, note that a_1, a_3, a_5 also form a convex rose in C_n . Thus, the essential core of $\langle a_1, a_3, a_5 \rangle$ in \tilde{C}_n is a copy of the standard Cayley graph of F_3 , and, in there, the minimal subtree for $\langle a_5, a_1 a_3^n \rangle$ gives the essential core of $\langle x_5, x_1 \rangle$.

Finally, we emphasise that the construction described in this example applies more generally to the case when Γ is a cycle of length at least 6.

4 | UNIFORMLY NON-QUASI-CONVEX RAYS

Having shown Theorem A and Theorem B(1), we now embark on the proof of Theorem B(2), which will occupy the current section and the next. This section is devoted to proving the following result, which is the most important ingredient in the proof of Theorem B(2).

Recall that quasi-convexity was introduced in Definition 2.13. If $A \subseteq X$ is a subset of a CAT(0) cube complex, we denote by $\text{Hull}_X(A)$ the smallest convex subcomplex of X containing A .

Theorem 4.1. *Let $G \curvearrowright X$ be a cocompact cubulation. Let $M \subseteq X^{(0)}$ be a median subalgebra that is preserved and acted upon cofinitely by a subgroup $H \leq G$. If M is not quasi-convex, then there exists a (combinatorial) ray $r : [0, +\infty) \rightarrow X$ such that:*

- (1) r stays at bounded distance from M ;
- (2) there exists a constant $K \geq 0$ such that, for all $t \geq s \geq 0$, the set $\text{Hull}_X(r|_{[s,t]})$ contains points at distance $\geq \lfloor \frac{t-s}{K} \rfloor$ from M .

This result should look rather plausible, but its proof will require significant technical work. To motivate the reader through it and emphasise that things are more delicate than they may seem, we recall the following classical example. It shows that Theorem 4.1 badly fails if M is replaced with just a connected subcomplex of X (or its vertex set); in general, lack of quasi-convexity is not witnessed by a ray, let alone by a ‘uniformly non-quasi-convex’ one. Thus, the assumption that M be a median subalgebra is key and we will have to significantly rely on it.

Example 4.2. Let $G = \pi_1(S) \rtimes \mathbb{Z}$ be the fundamental group of a fibred closed hyperbolic 3-manifold. Let $G \curvearrowright X$ be a cocompact cubulation, whose existence is provided by [12, 44]. Let $H = \pi_1(S)$ be the fibre subgroup and define $M \subseteq X$ as a thickening of an H -orbit, so that M is a connected subcomplex on which H acts cocompactly; this is possible since H is finitely generated. However, M is not a median subalgebra of X , as H is distorted in G .

Since X is hyperbolic, coarse-median quasi-convexity and hyperbolic quasi-convexity are the same notion. In particular, M is not quasi-convex in X , as the subgroup H is not quasi-convex in G . At the same time, we certainly cannot find any non-quasi-convex rays near M : Since X is hyperbolic, every ray in X is quasi-convex.

4.1 | Terminology, notation and conventions

Let X be a finite-dimensional CAT(0) cube complex and $M \subseteq X^{(0)}$ a median subalgebra.

All distances, neighbourhoods and geodesics should always be understood to correspond to the combinatorial metric on the 0-skeleton of X . The notation $d(\cdot, \cdot)$ will always refer to this metric. We write $\mathcal{N}_R(A)$ for the (closed) R -neighbourhood in $X^{(0)}$ of a subset $A \subseteq X^{(0)}$.

Equipped with the restriction of the median operator of X , the subalgebra M is a *median algebra* (see [9, 52] for some background) and we will make use of the corresponding terminology. Specifically, a subset $A \subseteq M$ is said to be *convex* (in M) if $m(A, A, M) \subseteq A$. A subset $\mathfrak{h} \subseteq M$ is a

halfspace if both \mathfrak{h} and $\mathfrak{h}^* := M - \mathfrak{h}$ are convex and non-empty. A wall of M is an unordered pair $\{\mathfrak{h}, \mathfrak{h}^*\}$, where $\mathfrak{h} \subseteq M$ is a halfspace. See Remark 4.4(2) below for the relation between halfspaces of M and X .

We denote by $\mathcal{H}(M)$ and $\mathcal{W}(M)$ the sets of halfspaces and walls of M , respectively. Similarly, $\mathcal{H}(X)$ and $\mathcal{W}(X)$ are the sets of halfspaces and hyperplanes of the cube complex X , or, equivalently, the halfspaces and walls of the median subalgebra $X^{(0)}$. If $A, B \subseteq X$ are subsets, we write:

$$\mathcal{H}(A|B) := \{\mathfrak{h} \in \mathcal{H}(X) \mid A \subseteq \mathfrak{h}^*, B \subseteq \mathfrak{h}\},$$

and denote by $\mathcal{W}(A|B) \subseteq \mathcal{W}(X)$ the set of hyperplanes bounding these halfspaces. Note that $\mathfrak{h} \in \mathcal{H}(A|\mathfrak{h})$ and $\mathfrak{h}^* \in \mathcal{H}(\mathfrak{h}|A)$ for every subset $A \subseteq \mathfrak{h}^*$.

Halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}(M)$ are *transverse* if all four intersections $\mathfrak{h} \cap \mathfrak{k}, \mathfrak{h}^* \cap \mathfrak{k}, \mathfrak{h} \cap \mathfrak{k}^*, \mathfrak{h}^* \cap \mathfrak{k}^*$ are non-empty. In this case, we also say that the corresponding walls are *transverse*. Two sets \mathcal{A}, \mathcal{B} of halfspaces/walls are *transverse* if every element of \mathcal{A} is transverse to every element of \mathcal{B} . If $\mathfrak{h} \subseteq \mathfrak{k}$ or $\mathfrak{k} \subseteq \mathfrak{h}$, we say that \mathfrak{h} and \mathfrak{k} are *nested*.

By Chepoi–Roller duality, the median algebra M is canonically isomorphic to the 0-skeleton of a CAT(0) cube complex $\square(M)$, equipped with its natural median operator. The sets $\mathcal{W}(\square(M))$ and $\mathcal{H}(\square(M))$ are naturally identified with $\mathcal{W}(M)$ and $\mathcal{H}(M)$. In general, $\square(M)$ is not a subcomplex of X , let alone a convex one.[†] Note that M inherits its own intrinsic metric from $\square(M)$, but this will not cause any ambiguity: In our cases of interest, the metric inherited from $\square(M)$ coincides with the restriction of the combinatorial metric of X (see Remark 4.4(3)).

In fact, we will mostly deal with median algebras satisfying the following additional property (studied, e.g. in [31, subsection 4.4.1]).

Definition 4.3. A subset $A \subseteq X^{(0)}$ is *edge-connected* if, for all $x, y \in A$, there exists a sequence of points $x_1, \dots, x_n \in A$ such that $x_1 = x, x_n = y$ and, for all i , the points x_i and x_{i+1} are joined by an edge of X .

Remark 4.4. The following three observations explain our interest in edge-connected subalgebras.

- (1) Suppose $H \leq \text{Aut}(X)$ is a finitely generated subgroup that preserves M and acts cofinitely on it. Suppose further that X is locally finite. Then M is contained in an *edge-connected*, H -invariant, H -cofinite subalgebra $M' \subseteq X^{(0)}$.

Indeed, since H is finitely generated, its orbits in X are coarsely connected. Thus, since M is H -cofinite, there exists $R \geq 0$ such that $\mathcal{N}_R(M)$ is an edge-connected subset of X . By [31, Lemma 4.21], the subalgebra M' generated by this subset is also edge-connected. Finally, M' is at finite Hausdorff distance from M since it is constructed by taking medians a bounded number of times (Lemma 2.2), hence it is H -cofinite because X is locally finite.

- (2) Let $\mathcal{H}_M(X) \subseteq \mathcal{H}(X)$ be the subset of halfspaces \mathfrak{h} such that both $\mathfrak{h} \cap M$ and $\mathfrak{h}^* \cap M$ are non-empty. We have a map:

$$\text{res}_M : \mathcal{H}_M(X) \rightarrow \mathcal{H}(M), \quad \text{res}_M(\mathfrak{h}) := \mathfrak{h} \cap M.$$

[†] For instance, any subset of \mathbb{Z} is a median subalgebra, if we view \mathbb{Z} as the vertex set of the standard cubical structure on the real line.

This map is always *surjective* by [9, Lemma 6.5]. It is also clear that it is a morphism of pocsets, that is, that $\text{res}_M(\mathfrak{h}^*) = \text{res}_M(\mathfrak{h})^*$ and $\mathfrak{h} \subseteq \mathfrak{k} \Rightarrow \text{res}_M(\mathfrak{h}) \subseteq \text{res}_M(\mathfrak{k})$. On the other hand, res_M is *injective* if and only if M is edge-connected [31, Lemma 4.20].

Note that, even when M is edge-connected, res_M does not preserve transversality: Halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}_M(X)$ can be transverse even if $\text{res}_M(\mathfrak{h})$ and $\text{res}_M(\mathfrak{k})$ are not. In fact, this always happens, unless M is the vertex set of a *convex* subcomplex of X .

- (3) Denote for a moment by d_M the metric on M corresponding to the combinatorial metric on $\square(M)$, the cube complex given by Chepoi–Roller duality. Recall that d is instead the restriction of the combinatorial metric of X .

The metrics d and d_M coincide on M if and only if M is edge-connected (in general, we only have $d_M \leq d$). This is an immediate consequence of Item (2) above, recalling that, if $x, y \in M$, the distance $d_M(x, y)$ is the number of walls in $\mathcal{W}(M)$ separating x and y , while $d(x, y)$ is the number of hyperplanes in $\mathcal{W}(X)$ separating them.

In addition, if M is edge-connected, the dual cube complex $\square(M)$ can be realised as a subcomplex of X : It is simply the union of all cubes of X whose 0-skeleton is contained in M . In general, $\square(M)$ is not a convex subcomplex of X . Nevertheless, if an edge path is contained in $\square(M)$, then it is a geodesic in X if and only if it is a geodesic in $\square(M)$.

We conclude this subsection with some more notation. First, the map $\text{res}_M : \mathcal{H}_M(X) \rightarrow \mathcal{H}(M)$ introduced in Remark 4.4(2) has an obvious twin

$$\text{res}_M : \mathcal{W}_M(X) \rightarrow \mathcal{W}(M),$$

where $\mathcal{W}_M(X) \subseteq \mathcal{W}(X)$ is the subset of hyperplanes bounding halfspaces in $\mathcal{H}_M(X)$.

We will speak of X -*transverse* or M -*transverse* hyperplanes/halfspaces depending on the pocset of interest: $\mathcal{H}(X)$ or $\mathcal{H}(M)$. This will be important in order to avoid confusion, since the maps res_M are bijections in our cases of interest and they do not preserve transversality.

For a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$, we denote by $C_X(\mathfrak{w}) \subseteq X$ its X -*carrier*: This is the convex subcomplex of X spanned by all edges crossing \mathfrak{w} . On the other hand, for a wall $\mathfrak{u} \in \mathcal{W}(M)$, we denote by $C_M(\mathfrak{u}) \subseteq M$ its M -*carrier*: This is the intersection between the 0-skeleton of $\square(M)$ and the carrier in $\square(M)$ of the hyperplane of $\square(M)$ determined by \mathfrak{w} . In other words, a point $x \in M$ lies in $C_M(\mathfrak{u})$ when there exists $y \in M$ such that \mathfrak{u} is the only wall of M separating x and y .

Finally, if \mathfrak{h} is a halfspace bounded by a hyperplane/wall \mathfrak{w} , it is convenient to write $C_X(\mathfrak{h})$ and $C_M(\mathfrak{h})$ for the intersections $C_X(\mathfrak{w}) \cap \mathfrak{h}$ and $C_M(\mathfrak{w}) \cap \mathfrak{h}$.

Remark 4.5. For every $\mathfrak{w} \in \mathcal{W}_M(X)$, we have $C_M(\text{res}_M(\mathfrak{w})) \subseteq C_X(\mathfrak{w}) \cap M$, but in general this inclusion can be very far from an equality, even at the level of 0-skeleta.

For instance, let X be the strip $\mathbb{R} \times [0, 1]$ with its standard decomposition into squares, let \mathfrak{w} be the hyperplane corresponding to the $[0,1]$ -factor and let M be the 0-skeleton of the subcomplex $(\mathbb{R} \times \{0\}) \cup (\{0\} \times [0, 1])$. Then $C_X(\mathfrak{w}) \cap M = M$, while $C_M(\text{res}_M(\mathfrak{w})) = \{(0, 0), (0, 1)\}$.

Luckily, when X and M have geometric group actions, the distinction between $C_M(\text{res}_M(\mathfrak{w}))$ and $C_X(\mathfrak{w}) \cap M$ will not be quite so drastic. See Lemma 4.11 below.

We record here also the following observation for later use.

Remark 4.6. Let $A \subseteq X^{(0)}$ be a subset of a CAT(0) cube complex. Set $\delta := \dim X$.

- (1) If $m(A, A, X) \subseteq \mathcal{N}_R(A)$, then $\text{Hull}_X(A) \subseteq \mathcal{N}_{2\delta_R}(A)$. This is a straightforward consequence of the fact that the median operator is 1-Lipschitz, together with the fact that hulls are constructed by taking medians a bounded number of times (Lemma 2.2).
- (2) In particular, $A \subseteq X$ is quasi-convex (in the sense of Definition 2.13) if and only if A is at finite Hausdorff distance from $\text{Hull}_X(A)$.

4.2 | Quadrants, grids and quasi-convexity

Let X be a finite-dimensional CAT(0) cube complex and let $M \subseteq X^{(0)}$ be a median subalgebra.

The goal of this subsection is to rephrase (failure of) quasi-convexity for the subalgebra M in terms of ‘grids’ of hyperplanes of X . The important result in this regard is Corollary 4.9.

Just like a convex subcomplex of X is the complement of the union of *halfspaces* disjoint from it, an (edge-connected) median *subalgebra* of X is the complement of the union of *quadrants* disjoint from it. The next lemma proves a stronger version of this fact.

Recall that the subset $\mathcal{H}_M(X) \subseteq \mathcal{H}(X)$ was introduced in Remark 4.4(2).

Lemma 4.7. *Let $C \subseteq X$ be a convex subcomplex with $C \cap \text{Hull}_X(M) \neq \emptyset$.*

- (1) *If $C \cap M = \emptyset$, there exist distinct halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}_M(X)$ with $C \subseteq \mathfrak{h} \cap \mathfrak{k}$ and $M \subseteq \mathfrak{h}^* \cup \mathfrak{k}^*$.*
- (2) *If, in addition, M is edge-connected, then \mathfrak{h} and \mathfrak{k} are X -transverse.*

Proof. The pocset $\mathcal{H}_M(X) \subseteq \mathcal{H}(X)$ is naturally identified with $\mathcal{H}(\text{Hull}_X(M))$, and a halfspace $\mathfrak{h} \in \mathcal{H}_M(X)$ contains C if and only if it contains $C \cap \text{Hull}_X(M)$. Thus, up to replacing X with $\text{Hull}_X(M)$ and C with $C \cap \text{Hull}_X(M)$, we can assume $\mathcal{H}_M(X) = \mathcal{H}(X)$.

Now, we prove part (1). Consider the subset $\sigma_C \subseteq \mathcal{H}(M)$ defined as follows:

$$\sigma_C := \{\mathfrak{h} \cap M \mid \mathfrak{h} \in \mathcal{H}(X) \text{ and } C \subseteq \mathfrak{h}\}.$$

Note that σ_C cannot contain any infinite descending chains $\mathfrak{h}_1 \cap M \supseteq \mathfrak{h}_2 \cap M \supseteq \dots$. Indeed, the halfspaces $\mathfrak{h}_i \in \mathcal{H}(X)$ would be pairwise distinct, giving infinitely many elements of $\mathcal{H}(p|C)$. This is impossible since $d(p, C)$ is finite.

Also note that the elements of $\sigma_C \subseteq \mathcal{H}(M)$ cannot pairwise intersect. Otherwise, since σ_C does not contain any infinite descending chains, there would be a point $p \in M$ that lies in all elements of σ_C (this is clear viewing M as the 0-skeleton of the cube complex $\square(M)$ via Chepoi–Roller duality). By the definition of σ_C , we would also have $p \in C$, contradicting the assumption that $C \cap M = \emptyset$.

In conclusion, there exist two disjoint elements of σ_C . This means that there exist halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}(X) = \mathcal{H}_M(X)$ such that $C \subseteq \mathfrak{h} \cap \mathfrak{k}$ and $M \cap \mathfrak{h} \cap \mathfrak{k} = \emptyset$. This proves part (1).

Regarding part (2), recall that \mathfrak{h} and \mathfrak{k} intersect M , as they lie in $\mathcal{H}_M(X)$. Since $M \cap \mathfrak{h} \cap \mathfrak{k} = \emptyset$, the intersections $\mathfrak{h} \cap \mathfrak{k}^*$ and $\mathfrak{h}^* \cap \mathfrak{k}$ must intersect M , hence they are non-empty. The intersection $\mathfrak{h} \cap \mathfrak{k}$ is also non-empty, as it contains C .

Finally, if M is edge-connected, it contains the vertex set of an edge-path in X joining a point of $M \cap \mathfrak{h} \cap \mathfrak{k}^*$ to a point of $M \cap \mathfrak{h}^* \cap \mathfrak{k}$. Since this path cannot intersect $\mathfrak{h} \cap \mathfrak{k}$, which is disjoint from M , it must intersect $\mathfrak{h}^* \cap \mathfrak{k}^*$. This proves that $\mathfrak{h}^* \cap \mathfrak{k}^*$ is non-empty, hence \mathfrak{h} and \mathfrak{k} are X -transverse. □

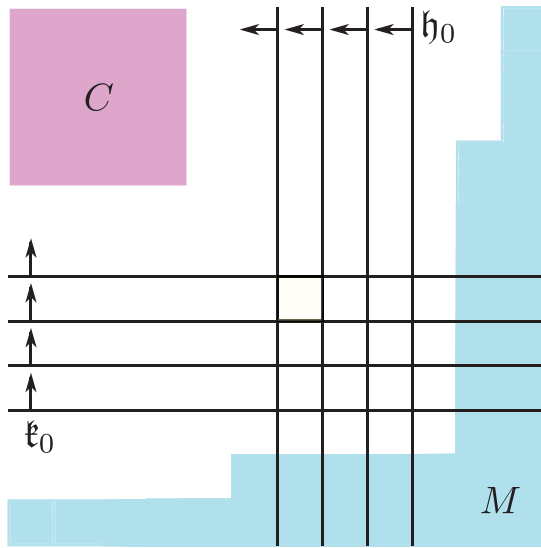


FIGURE 3 Grid of hyperplanes separating M and C in Lemma 4.8.

The next lemma shows that, if a convex subcomplex and a median subalgebra of X are far from each other, then they are separated by a large grid of hyperplanes of X . This configuration is depicted in Figure 3.

Lemma 4.8. *Let M be edge-connected. Let $C \subseteq X$ be a convex subcomplex with $C \cap \text{Hull}_X(M) \neq \emptyset$. Set $\delta := \dim X$. If $d(C, M) > (2^\delta + 1)\delta n$ for some $n \geq 0$, then there exist halfspaces $\mathfrak{h}_0 \supseteq \dots \supseteq \mathfrak{h}_n$ and $\mathfrak{k}_0 \supseteq \dots \supseteq \mathfrak{k}_n$ in $\mathcal{H}_M(X)$ such that each \mathfrak{h}_i is X -transverse to every \mathfrak{k}_j , and $C \subseteq \mathfrak{h}_n \cap \mathfrak{k}_n$ and $M \subseteq \mathfrak{h}_0^* \cup \mathfrak{k}_0^*$.*

Proof. Set $C' := \text{Hull}_X(\mathcal{N}_{\delta n}(C))$. Since the median operator is 1-Lipschitz in each coordinate, we have $m(\mathcal{N}_{\delta n}(C), \mathcal{N}_{\delta n}(C), X) \subseteq \mathcal{N}_{2\delta n}(C)$, hence Remark 4.6 guarantees that $C' \subseteq \mathcal{N}_{(2^\delta + 1)\delta n}(C)$. It follows that $C' \cap M = \emptyset$.

By Lemma 4.7, there exist X -transverse halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}_M(X)$ such that $C' \subseteq \mathfrak{h} \cap \mathfrak{k}$ and $M \subseteq \mathfrak{h}^* \cup \mathfrak{k}^*$. Since $\mathcal{N}_{\delta n}(C) \subseteq C'$, we have $d(C, \mathfrak{h}^*) > \delta n$ and $d(C, \mathfrak{k}^*) > \delta n$.

By Dilworth’s lemma, the sets $\mathcal{H}(\mathfrak{h}^*|C)$ and $\mathcal{H}(\mathfrak{k}^*|C)$ each contain a chain of halfspaces of length $n + 1$. Also note that $\mathcal{H}(\mathfrak{h}^*|C)$ and $\mathcal{H}(\mathfrak{k}^*|C)$ are contained in $\mathcal{H}_M(X)$, since $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}_M(X)$ and $C \cap \text{Hull}_X(M) \neq \emptyset$. Finally, each $\mathfrak{j} \in \mathcal{H}(\mathfrak{h}^*|C)$ is X -transverse to each $\mathfrak{j}' \in \mathcal{H}(\mathfrak{k}^*|C)$: The proof of this fact is identical to that of Lemma 4.7(2). This concludes the proof of the lemma. \square

Corollary 4.9. *Let M be edge-connected. Then M fails to be quasi-convex if and only if, for every $n \geq 0$, there exist halfspaces $\mathfrak{h}_0 \supseteq \dots \supseteq \mathfrak{h}_n$ and $\mathfrak{k}_0 \supseteq \dots \supseteq \mathfrak{k}_n$ in $\mathcal{H}_M(X)$ such that each \mathfrak{h}_i is X -transverse to every \mathfrak{k}_j , and $M \subseteq \mathfrak{h}_0^* \cup \mathfrak{k}_0^*$.*

Proof. Recall from Remark 4.6 that M is quasi-convex if and only if it is at finite Hausdorff distance from $\text{Hull}_X(M)$. Now, if $\text{Hull}_X(M)$ contains points arbitrarily far from M , then Lemma 4.8 yields the required hyperplane grids. Conversely, if there exist halfspaces $\mathfrak{h}_i, \mathfrak{k}_j$ as in the statement of the corollary, then $\text{Hull}_X(M) \cap \mathfrak{h}_n \cap \mathfrak{k}_n \neq \emptyset$ by Helly’s lemma and $d(\mathfrak{h}_n \cap \mathfrak{k}_n, M) > n$. Hence $\text{Hull}_X(M)$ contains points arbitrarily far from M , showing that M is not quasi-convex. \square

4.3 | Constructing uniformly non-quasi-convex geodesics

Let X be a CAT(0) cube complex with a geometric action $G \curvearrowright X$. Let $M \subseteq X^{(0)}$ be a median subalgebra. We assume that M is edge-connected and that there exists a subgroup $H \leq G$ acting on M with finitely many orbits.

The following is the main result of this subsection. We will quickly deduce Theorem 4.1 from it in the next subsection.

Recall that $\square(M)$ is canonically realised as a (non-convex) subcomplex of X , since M is edge-connected (Remark 4.4(3)). Hyperplanes of $\square(M)$ are precisely intersections with $\square(M)$ of hyperplanes of X (though this is not true of carriers). Also recall that we defined the carrier of a halfspace as the intersection of the halfspace with the carrier of the corresponding wall/hyperplane.

Proposition 4.10. *There exists a constant $K' \geq 0$ such that the following holds. For all X -transverse halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}_M(X)$ with $M \cap \mathfrak{h} \cap \mathfrak{k} = \emptyset$, there exists a geodesic $\alpha \subseteq \square(M) \subseteq X$ from a vertex in the M -carrier $C_M(\text{res}_M(\mathfrak{k}^*))$ to a vertex in $C_M(\text{res}_M(\mathfrak{h}^*))$ with the following property. For all integers $t \geq s \geq 0$, the set $\text{Hull}_X(\alpha|_{[s,t]})$ contains points at distance $> \lfloor \frac{t-s}{K'} \rfloor$ from M .*

Before embarking on the proof of Proposition 4.10, we need to make the following key observation. It exploits cocompactness of M and X to deduce that M -carriers are not too different from intersections of X -carriers with M (recall Remark 4.5). We denote by $I(x, y)$ the interval in X with endpoints x and y , that is, the union of all combinatorial geodesics in X joining x and y .

Lemma 4.11. *There exists a constant $K \geq 0$ with the following property. Consider a halfspace $\mathfrak{h} \in \mathcal{H}_M(X)$ and a geodesic $\beta \subseteq \square(M) \subseteq X$ connecting a point $x \in C_M(\text{res}_M(\mathfrak{h}))$ to a point $y \in M \cap C_X(\mathfrak{h})$. If $d(x, y) > K$, then $I(x, y) \cap C_M(\text{res}_M(\mathfrak{h}))$ contains a vertex other than x .*

Proof. For every halfspace $\mathfrak{h} \in \mathcal{H}_M(X)$, its stabiliser $G(\mathfrak{h}) \leq G$ acts cocompactly on $C_X(\mathfrak{h})$. Since H acts cofinitely on M , Lemma A.2 shows that the intersection $H \cap G(\mathfrak{h})$ acts cofinitely on the set $M \cap C_X(\mathfrak{h})$. Note that there are only finitely many H -orbits of halfspaces in $\mathcal{H}_M(X)$, since this set is equivariantly in bijection with $\mathcal{H}(M)$. Thus, there exists a constant $K \geq 0$ such that, for every $\mathfrak{h} \in \mathcal{H}_M(X)$, all orbits of the action $H \cap G(\mathfrak{h}) \curvearrowright M \cap C_X(\mathfrak{h})$ are K -dense.

The set $M \cap C_X(\mathfrak{h})$ contains the subset $C_M(\text{res}_M(\mathfrak{h}))$, which is $(H \cap G(\mathfrak{h}))$ -invariant and non-empty. Hence, for every $\mathfrak{h} \in \mathcal{H}_M(X)$, the subset $C_M(\text{res}_M(\mathfrak{h}))$ is K -dense in $M \cap C_X(\mathfrak{h})$.

Now, consider points $x \in C_M(\text{res}_M(\mathfrak{h}))$ and $y \in M \cap C_X(\mathfrak{h})$ with $d(x, y) > K$. By the previous paragraph, there exists a point $z \in C_M(\text{res}_M(\mathfrak{h}))$ with $d(z, y) \leq K$. Setting $w := m(x, y, z)$, we have $w \in M \cap I(x, y)$ and $w \neq x$. We are left to show that w lies in $C_M(\text{res}_M(\mathfrak{h}))$.

Since $x, z \in C_M(\text{res}_M(\mathfrak{h}))$, there exist $x', z' \in C_M(\text{res}_M(\mathfrak{h}^*))$ with $d(x, x') = d(z, z') = 1$. Consider the point $w' = m(x', y, z')$. It is clear that $w' \in M \cap \mathfrak{h}^*$ and that $d(w, w') = 1$. This shows $w \in C_M(\text{res}_M(\mathfrak{h}))$, concluding the proof. \square

Now, we begin by describing how to construct the geodesic α appearing in Proposition 4.10. Then we will show that it satisfies the required properties, which will take a few lemmas. The reader might find Figure 4 helpful while working through Construction 4.12 and the subsequent Lemma 4.13.

We say that a segment $\beta \subseteq \alpha$ is a \parallel -segment (resp. a \perp -segment) if all halfspaces entered by β lie in Ω_{\parallel} (resp. in Ω_{\perp}). The next lemma collects the key properties of the geodesic α .

Lemma 4.13. *The oriented geodesic α obtained in Construction 4.12 satisfies the following.*

- (1) *If some $j_{\parallel} \in \Omega_{\parallel}$ is entered by α before some $j_{\perp} \in \Omega_{\perp}$, then j_{\parallel} is X -transverse to j_{\perp} .*
- (2) *All \parallel -segments and \perp -segments of α have length $\leq K$, with K as in Lemma 4.11.*
- (3) *If α contains a \parallel -segment β_{\parallel} immediately followed by a \perp -segment β_{\perp} , and if j_{\parallel} and j_{\perp} are any halfspaces entered, respectively, by β_{\parallel} and β_{\perp} , then $M \cap j_{\parallel}^* \cap j_{\perp} = \emptyset$.*

Proof. Property (1) is almost immediate. For every $j_{\parallel} \in \Omega_{\parallel}$ and $j_{\perp} \in \Omega_{\perp}$, the intersection $j_{\parallel}^* \cap j_{\perp}$ is non-empty, as it contains $\mathfrak{k} \cap j_{\perp} \neq \emptyset$. Thus, if α enters j_{\parallel} before j_{\perp} , these two halfspaces must be X -transverse.

We now prove Property (2), beginning with some preliminary remarks.

In Construction 4.12, we introduced an equivalence relation \sim on $\mathcal{G}_M(p, q)$ and a total order $<$ on its set of equivalence classes. What matters for these relations is the function $t \mapsto d(\beta(t), \mathfrak{k})$, where $\beta \in \mathcal{G}_M(p, q)$. In turn, this function is completely determined by the order in which β enters elements of Ω_{\parallel} and Ω_{\perp} : Entering some $j \in \Omega_{\perp}$ does not change the value of $d(\beta(t), \mathfrak{k})$, since j is X -transverse to \mathfrak{k} ; on the other hand, entering some $j \in \Omega_{\parallel}$ increases $d(\beta(t), \mathfrak{k})$ by 1, since $j \cap \mathfrak{k} = \emptyset$.

If $\beta \in \mathcal{G}_M(p, q)$ enters halfspaces j_1, j_2 consecutively and if j_1, j_2 are M -transverse, then there exists another geodesic $\beta' \in \mathcal{G}_M(p, q)$ only differing from β in the fact that it enters j_2 before j_1 . If j_1, j_2 both lie in Ω_{\parallel} , or both lie in Ω_{\perp} , then $\beta' \sim \beta$. Instead, if $j_1 \in \Omega_{\parallel}$ and $j_2 \in \Omega_{\perp}$, then $\beta' < \beta$.

With these observations in hand, we complete the proof of Property (2) by the following two claims. Let $\alpha_{\parallel} \subseteq \alpha$ and $\alpha_{\perp} \subseteq \alpha$ be, respectively, a maximal \parallel -segment and a maximal \perp -segment. Let K be the constant provided by Lemma 4.11.

Claim 1. The segment α_{\perp} has length $\leq K$.

Proof of Claim 1. First, suppose that α_{\perp} is an initial segment of α . Since all elements of Ω_{\perp} are X -transverse to \mathfrak{k} , we have $\alpha_{\perp} \subseteq C_X(\mathfrak{k}^*)$. Every halfspace j entered by α_{\perp} satisfies $\text{res}_M(j) \cap \text{res}_M(\mathfrak{k}) = \emptyset$, since α joins the points p and q in the bridge \mathcal{B} . Thus, every point of $\text{Hull}_X(\alpha_{\perp})$ (other than p) lies in the difference $C_X(\mathfrak{k}^*) - C_M(\text{res}_M(\mathfrak{k}^*))$. Lemma 4.11 then shows that α_{\perp} has length $\leq K$.

Suppose now instead that α_{\perp} is not an initial segment of α . Let $\mathfrak{w} \in \mathcal{W}(X)$ be the last hyperplane crossed by α before the start of α_{\perp} ; let $e \subseteq \alpha$ be the edge crossing \mathfrak{w} . The hyperplane \mathfrak{w} bounds an element of Ω_{\parallel} , by maximality of α_{\perp} . So, by part (1), \mathfrak{w} is X -transverse to every hyperplane crossed by α_{\perp} and hence $\alpha_{\perp} \subseteq C_X(\mathfrak{w})$. If α_{\perp} had length $> K$, Lemma 4.11 would imply the existence of a point $x \in \text{Hull}_X(\alpha_{\perp}) \cap C_M(\text{res}_M(\mathfrak{w}))$ other than the initial vertex of α_{\perp} . Let $x' \in M$ be the point with $\mathcal{W}(x|x') = \{\mathfrak{w}\}$. Then, replacing the segment $e \cup \alpha_{\perp} \subseteq \alpha$ with a geodesic in $\square(M) \subseteq X$ passing through x' and x , we would find an element $\alpha' \in \mathcal{G}_M(p, q)$ with $\alpha' < \alpha$, a contradiction. □

Claim 2. The segment α_{\parallel} has length $\leq K$.

Proof of Claim 2. This is entirely analogous to the previous proof.

First, suppose that α_{\parallel} is a terminal segment of α . Note that all elements of Ω_{\parallel} are X -transverse to \mathfrak{h} (by part (1), since $\mathfrak{h} \in \Omega_{\perp}$). So we have $\alpha_{\parallel} \subseteq C_X(\mathfrak{h}^*)$. Again, by the properties of bridges, every halfspace j entered by α_{\parallel} satisfies $\text{res}_M(j^*) \cap \text{res}_M(\mathfrak{h}) = \emptyset$. As in Claim 1, every point of $\text{Hull}_X(\alpha_{\parallel})$

(other than q) lies in the difference $C_X(\mathfrak{h}^*) - C_M(\text{res}_M(\mathfrak{h}^*))$ and Lemma 4.11 shows that α_{\parallel} has length $\leq K$.

Suppose instead that α_{\parallel} is not a terminal segment of α . Let $\mathfrak{w} \in \mathcal{W}(X)$ be the first hyperplane crossed by α after the end of α_{\parallel} ; let $e \subseteq \alpha$ be the edge crossing \mathfrak{w} . The hyperplane \mathfrak{w} bounds an element of Ω_{\perp} , so, by part (1), it is X -transverse to every hyperplane crossed by α_{\parallel} and $\alpha_{\parallel} \subseteq C_X(\mathfrak{w})$. If α_{\parallel} had length $> K$, Lemma 4.11 would imply the existence of a point $x \in \text{Hull}_X(\alpha_{\parallel}) \cap C_M(\text{res}_M(\mathfrak{w}))$ other than the terminal endpoint of α_{\parallel} . If $x' \in M$ is the point with $\mathcal{W}(x|x') = \{\mathfrak{w}\}$, we can form a new geodesic $\alpha' \in \mathcal{G}_M(p, q)$ by replacing the segment $e \cup \alpha_{\parallel} \subseteq \alpha$ with a geodesic in $\square(M) \subseteq X$ that follows α_{\parallel} up to x , then crosses \mathfrak{w} to reach x' and finally moves to the vertex in $e - \alpha_{\parallel}$. As before, we have $\alpha' < \alpha$, a contradiction. \square

Finally, we prove Property (3). Consider $\beta_{\parallel}, \beta_{\perp}$ and $\mathfrak{j}_{\parallel}, \mathfrak{j}_{\perp}$ as in the statement. Let x be the point where β_{\parallel} and β_{\perp} meet. Suppose for the sake of contradiction that $M \cap \mathfrak{j}_{\parallel}^* \cap \mathfrak{j}_{\perp} \neq \emptyset$. Then \mathfrak{j}_{\parallel} and \mathfrak{j}_{\perp} are M -transverse (since \mathfrak{j}_{\parallel} is entered before \mathfrak{j}_{\perp} by a geodesic contained in $\square(M)$).

Without loss of generality, suppose $x \in C_M(\text{res}_M(\mathfrak{j}_{\parallel})) \cap C_M(\text{res}_M(\mathfrak{j}_{\perp}^*))$; otherwise it suffices to replace \mathfrak{j}_{\parallel} and \mathfrak{j}_{\perp} by a different pair of M -transverse halfspaces entered by $\beta_{\parallel}, \beta_{\perp}$, corresponding to walls in $\mathcal{W}(M)$ that are closer to x . This does not affect the assumption that \mathfrak{j}_{\parallel} and \mathfrak{j}_{\perp} are M -transverse.

Up to changing the order in which β_{\parallel} and β_{\perp} enter the respective halfspaces (which yields $\alpha' \sim \alpha$), we can also assume that \mathfrak{j}_{\parallel} is the last halfspace entered before x and that \mathfrak{j}_{\perp} is the first entered after it. Now, since \mathfrak{j}_{\parallel} and \mathfrak{j}_{\perp} are M -transverse, we can swap the order in which these halfspaces are entered to produce $\alpha'' < \alpha$, a contradiction. \square

We only need one last simple combinatorial lemma before proving Proposition 4.10.

Lemma 4.14. *Suppose some geodesic in X enters halfspaces $\mathfrak{h}_1, \dots, \mathfrak{h}_N$ in this order (not necessarily consecutively), and also halfspaces $\mathfrak{k}_1, \dots, \mathfrak{k}_N$ in this order. Consider $n \geq 0$. If $N > 2n \dim X$, there exist an index $1 \leq k \leq N$ and indices $1 \leq i_1 < \dots < i_n < k$ and $k < j_1 < \dots < j_n \leq N$ such that:*

$$\mathfrak{h}_{i_1} \supset \mathfrak{h}_{i_2} \supset \dots \supset \mathfrak{h}_{i_n} \supset \mathfrak{h}_k, \quad \mathfrak{k}_k \supset \mathfrak{k}_{j_1} \supset \mathfrak{k}_{j_2} \supset \dots \supset \mathfrak{k}_{j_n}.$$

Proof. Define functions $f, g : \{1, \dots, N\} \rightarrow \{0, \dots, N\}$ as follows. The integer $f(k)$ is the longest length of a chain of \mathfrak{h}_i all strictly containing \mathfrak{h}_k , while $g(k)$ is the longest length of a chain of \mathfrak{k}_i all strictly contained in \mathfrak{k}_k . We need an index k so that $f(k) \geq n$ and $g(k) \geq n$ hold simultaneously.

For an integer i , consider the halfspaces \mathfrak{h}_j with $f(j) = i$. It is clear that they must be pairwise transverse. Hence $\#f^{-1}(i) \leq \dim X$ and, similarly, $\#g^{-1}(i) \leq \dim X$ for every $1 \leq i \leq N$. It follows that $\#f^{-1}([0, n - 1]) \leq n \dim X$ and $\#g^{-1}([0, n - 1]) \leq n \dim X$, which implies the lemma. \square

The next result immediately implies Proposition 4.10.

Corollary 4.15. *Let α be the geodesic obtained in Construction 4.12. Consider $n \geq 0$. For all integers $0 \leq s \leq t \leq d(p, q)$ with $t - s \geq 2K(2n \dim X + 2)$, the segment $\alpha|_{[s,t]}$ enters halfspaces $\mathfrak{h}_0, \dots, \mathfrak{h}_n \in \Omega_{\perp}$ and $\mathfrak{k}_0, \dots, \mathfrak{k}_n \in \Omega_{\parallel}$ such that, for all i, j , we have $M \cap \mathfrak{h}_i \cap \mathfrak{k}_j^* = \emptyset$ and \mathfrak{h}_i is X -transverse to \mathfrak{k}_j . In particular, $\text{Hull}_X(\alpha|_{[s,t]})$ contains points at distance $> n$ from M .*

Proof. After an initial segment of length $\leq K$, the geodesic $\alpha|_{[s,t]}$ contains a \parallel -segment β_1 followed by a \perp -segment γ_1 , and so on up to a \parallel -segment β_N and a \perp -segment γ_N , for some $N \geq 0$. Each β_i and γ_i has length $\leq K$ by Lemma 4.13(2) and the sum of their lengths is $\geq t - s - 2K$. Thus, $N \geq \lfloor \frac{t-s-2K}{2K} \rfloor > 2n \dim X$.

Let $\mathfrak{h}_i \in \Omega_{\parallel}$ and $\mathfrak{k}_i \in \Omega_{\perp}$ be arbitrary halfspaces entered by β_i and γ_i , respectively. By Lemma 4.13(3), we have $M \cap \mathfrak{h}_i^* \cap \mathfrak{k}_i = \emptyset$, hence $\text{res}_M(\mathfrak{k}_i) \subseteq \text{res}_M(\mathfrak{h}_i)$. Lemma 4.14 (applied to the cube complex $\square(M)$) ensures the existence of a chain:

$$\text{res}_M(\mathfrak{k}_{j_1}) \subset \dots \subset \text{res}_M(\mathfrak{k}_{j_n}) \subset \text{res}_M(\mathfrak{k}_k) \subset \text{res}_M(\mathfrak{h}_k) \subset \text{res}_M(\mathfrak{h}_{i_n}) \subset \dots \subset \text{res}_M(\mathfrak{h}_{i_1}).$$

It follows that $M \cap \mathfrak{h}_{i_a}^* \cap \mathfrak{k}_{j_b} = \emptyset$ for all indices a, b , while Lemma 4.13(1) guarantees that \mathfrak{h}_{i_a} and \mathfrak{k}_{j_b} are X -transverse. These are the required halfspaces.

Finally, by Helly’s lemma, there exists a point $z \in \text{Hull}_X(\alpha|_{[s,t]})$ lying in all $\mathfrak{h}_{i_a}^*$ and all \mathfrak{k}_{j_b} (as well as \mathfrak{h}_k^* and \mathfrak{k}_k). We clearly have $d(z, M) \geq n + 1$, concluding the proof. \square

4.4 | Conclusion

We finally prove Theorem 4.1 by combining Proposition 4.10 and Corollary 4.9.

Proof of Theorem 4.1. Since H acts cofinitely on the subalgebra M , [31, Lemma 4.12] shows that H is finitely generated. Thus, we can argue as in Remark 4.4(1) and thicken M to an edge-connected subalgebra $M \subseteq M' \subseteq X^{(0)}$ that is still H -invariant and H -cofinite.

Since M' is at finite Hausdorff distance from M , it is not quasi-convex. Hence Corollary 4.9 guarantees that, for every $n \geq 0$, there exist X -transverse halfspaces $\mathfrak{h}_n, \mathfrak{k}_n \in \mathcal{H}_{M'}(X)$ with $d(\text{res}_{M'}(\mathfrak{h}_n), \text{res}_{M'}(\mathfrak{k}_n)) > n$.

Proposition 4.10 yields geodesics $\alpha_n \subseteq \square(M')$ from $C_{M'}(\text{res}_{M'}(\mathfrak{h}_n^*))$ to $C_{M'}(\text{res}_{M'}(\mathfrak{k}_n^*))$ with the property that, for all $t \geq s \geq 0$, the set $\text{Hull}_X(\alpha_n|_{[s,t]})$ contains points at distance $> \lfloor \frac{t-s}{K'} \rfloor$ from M' . Note that the length of the α_n diverges with n , since $d(\text{res}_{M'}(\mathfrak{h}_n), \text{res}_{M'}(\mathfrak{k}_n)) > n$.

Exploiting H -cocompactness of M' , we can assume that the initial vertices of the α_n all lie in a given finite subset of M' . Passing to a subsequence, the α_n converge to a combinatorial ray $r \subseteq \square(M')$. Every segment of r is a segment of some α_n for large n , so it is still true that $\text{Hull}_X(r|_{[s,t]})$ contains points at distance $> \lfloor \frac{t-s}{K'} \rfloor$ from M' .

Finally, since $r \subseteq \square(M')$ and M' is at finite Hausdorff distance from M , the ray r stays at bounded distance from M . This concludes the proof. \square

5 | THE STANDARD COARSE MEDIAN STRUCTURE OF AN RIGHT-ANGLED COXETER GROUP

Subsections 5.1 and 5.2 are devoted to the proof of Theorem B(2), which is Theorem 5.3 below. Then, in Subsection 5.3, we deduce Corollary C.

5.1 | Notation and a preliminary lemma

Let W_{Γ} be a right-angled Coxeter group.

In this case, the graph-product complex \mathcal{D} defined in Subsection 3.2 is the cubical subdivision of a simpler CAT(0) cube complex, which is usually known as the *Davis complex* and which we denote by \mathcal{D}_Γ . The 1-skeleton of \mathcal{D}_Γ is naturally identified with the Cayley graph of W_Γ with respect to the generating set $\Gamma^{(0)}$. Thus, every edge and every hyperplane of \mathcal{D}_Γ is labelled by a vertex of Γ .

Let $[\mu_\Gamma] \in \mathcal{CM}_\square(W_\Gamma)$ be the coarse median structure induced on W_Γ by \mathcal{D}_Γ . We refer to $[\mu_\Gamma]$ as the *standard* coarse median structure on W_Γ .

We say that a graph Δ is *irreducible* if it is not a join of two proper subgraphs; equivalently, W_Δ is not a direct product of proper subgroups. For an (induced) subgraph $\Delta \subseteq \Gamma$, we write:

$$\Delta^\perp := \{v \in \Gamma^{(0)} \mid \Delta \subseteq \text{lk}(v)\}.$$

The following is fairly classical, but we were not able to find a proof in the literature.

Lemma 5.1. *Let $\Delta \subseteq \Gamma$ be an irreducible induced subgraph such that Δ^\perp spans a (possibly empty) clique. Let $\alpha : [0, +\infty) \rightarrow \mathcal{D}_\Delta$ be an infinite edge path in $\mathcal{D}_\Delta \subseteq \mathcal{D}_\Gamma$. If we have*

$$K := \sup_{g \in W_\Delta, v \in \Delta} \text{diam } \alpha^{-1}(g\mathcal{D}_{\Delta-\{v\}}) < +\infty,$$

then α is a Morse quasi-geodesic in \mathcal{D}_Γ .

Proof. For every (oriented) edge $e \subseteq \alpha$, denote by $\gamma(e) \in \Delta$ its label, by $\mathfrak{w}(e) \in \mathcal{W}(\mathcal{D}_\Gamma)$ the hyperplane it crosses and by $\mathfrak{h}(e) \in \mathcal{H}(\mathcal{D}_\Gamma)$ the halfspace it enters. We say that e is *good* if $\alpha(0)$ lies in $\mathfrak{h}(e)^*$ and the unbounded connected component of $\alpha - e$ is entirely contained in $\mathfrak{h}(e)$.

Furthermore, for $n \geq 0$, let e_n be the edge connecting $\alpha(n)$ and $\alpha(n + 1)$.

Claim 1. For every vertex $w \in \Delta$ and every sub-path $\alpha_0 \subseteq \alpha$ of length $> K$, there exists a good edge $e \subseteq \alpha_0$ with $\gamma(e) = w$.

Proof of Claim 1. Recall that the hyperplanes of \mathcal{D}_Δ labelled by w are pairwise disjoint, and that the connected components of the complement of their union are precisely the translates $g\mathcal{D}_{\Delta-\{w\}}$ with $g \in W_\Delta$ (with some shreds of cubes attached).

Let $[m, n] \subseteq [0, +\infty)$ be a maximal interval such that the vertices $\alpha(m)$ and $\alpha(n)$ lie in the same translate of $\mathcal{D}_{\Delta-\{w\}}$. Let $g\mathcal{D}_{\Delta-\{w\}}$ be this translate, where $g \in W_\Delta$. Then $\gamma(e_n) = w$ and $\gamma(t) \in \mathfrak{h}(e_n)$ for all $t \geq n + 1$. If $m \neq 0$, we also have $\gamma(e_{m-1}) = w$ and $\mathfrak{w}(e_{m-1}) \neq \mathfrak{w}(e_n)$; otherwise, $\alpha(m - 1)$ and $\alpha(n + 1)$ would lie in some other translate $g'\mathcal{D}_{\Delta-\{w\}}$, contradicting maximality of $[m, n]$. This shows that either $m = 0$ or $\alpha(0) \in \mathfrak{h}(e_m)^* \subseteq \mathfrak{h}(e_n)^*$, hence e_n is always a good edge.

Now, let $\alpha(t)$ be the initial vertex of α_0 and let $[m, n]$ be an interval containing t that is maximal in the above sense. By definition of K , such an interval exists and we have $n - m \leq K$. Hence $n \leq t + K$ and e_n is contained in α_0 . This is the required good edge, proving the claim. \square

Now, fix a vertex $v \in \Delta$. Since Δ is irreducible, its complement graph Δ^c (where two vertices are adjacent if and only if they are not adjacent in Δ) is connected. Let D be its diameter. Choose a sequence of integers $(n_k)_{k \geq 0}$ such that each edge e_{n_k} is good, with $\gamma(e_{n_k}) = v$ and:

$$(2D - 1)(K + 1) < n_{k+1} - n_k \leq 2D(K + 1).$$

This is possible by Claim 1. Set $N := \#\Delta^\perp$.

Recall that two hyperplanes of a cube complex are said to be L -separated, for some $L \geq 0$, if they are disjoint and at most L hyperplanes are transverse to both.

Claim 2. For $k \neq k'$, the hyperplanes $\mathfrak{w}(e_{n_k})$ and $\mathfrak{w}(e_{n_{k'}})$ are N -separated in D_Γ .

Proof of Claim 2. We have $\mathfrak{h}(e_{n_0}) \supseteq \mathfrak{h}(e_{n_1}) \supseteq \dots$ by construction. For every $w \in \Delta$ and $k \geq 0$, it is standard to show that the hyperplanes $\mathfrak{w}(e_{n_k})$ and $\mathfrak{w}(e_{n_{k+1}})$ are separated by a hyperplane labelled by w (using Claim 2 and the inequality $n_{k+1} - n_k > (2D - 1)(K + 1)$, where $D = \text{diam } \Delta^c$).

Hence, if some $\mathbf{u} \in \mathcal{W}(D_\Gamma)$ is transverse to $\mathfrak{w}(e_{n_k})$ and $\mathfrak{w}(e_{n_{k+1}})$, then \mathbf{u} is transverse to hyperplanes labelled by all vertices of Δ , hence \mathbf{u} must be labelled by a vertex of Δ^\perp . Since Δ^\perp is a clique, it follows that at most N hyperplanes are transverse to $\mathfrak{w}(e_{n_k})$ and $\mathfrak{w}(e_{n_{k+1}})$, proving the claim. \square

For each $k \geq 0$, let p_k be the gate-projection of $\alpha(0)$ to the carrier of $\mathfrak{w}(e_{n_k})$. Let $\beta \subseteq D_\Delta$ be a geodesic ray obtained by concatenating geodesics from $\alpha(0)$ to p_0 and from each p_i to p_{i+1} .

Claim 3. We have $d(p_k, \alpha(n_k)) \leq N + 2D(K + 1)$ for all $k \geq 0$.

Proof of Claim 3. Denote by C_k the carrier of $\mathfrak{w}(e_{n_k})$ and by P_k the gate-projection of the halfspace $\mathfrak{h}(e_{n_{k-1}})^*$ to C_k . Both p_k and $\alpha(n_k)$ lie in C_k , with p_k in fact lying in P_k . Since $\mathfrak{w}(e_{n_k})$ and $\mathfrak{w}(e_{n_{k-1}})$ are N -separated, the convex subcomplex P_k is crossed by at most N hyperplanes, and so it has diameter $\leq N$. By the properties of gate-projections, we have:

$$d(\alpha(n_k), P_k) \leq d(\alpha(n_k), \mathfrak{h}(e_{n_{k-1}})^*) \leq d(\alpha(n_k), \alpha(n_{k-1})) \leq 2D(K + 1).$$

Recalling that $p_k \in P_k$ and $\text{diam}(P_k) \leq N$, we obtain the claim. \square

Finally, Claim 3 shows that the distances $d(p_k, p_{k+1})$ are uniformly bounded, so [21, Theorem 4.2 and Theorem 2.14] imply that β is Morse. It is clear that α and β are at finite Hausdorff distance from each other, so it follows that α is a Morse quasi-geodesic. \square

5.2 | The main argument

Lemma 5.1 allows us to translate Theorem 4.1 into the following practical result, which will quickly yield Theorem B(2).

Proposition 5.2. *Let $W_\Gamma \curvearrowright X$ be a cocompact cubulation. Suppose that there exists an irreducible subgraph $\Delta \subseteq \Gamma$ such that W_Δ is median-cocompact in X . Suppose further that, for every $x \in \Delta$, the subgroup $W_{\Delta - \{x\}}$ is convex-cocompact in X . Then:*

- (1) *either W_Δ is itself convex-cocompact in X ,*
- (2) *or the centraliser $Z_{W_\Gamma}(W_\Delta) = W_{\Delta^\perp}$ is infinite.*

Proof. Let $M \subseteq X^{(0)}$ be a median subalgebra on which W_Δ acts cofinitely. Suppose that W_Δ is not convex-cocompact in X , that is, that M is not quasi-convex in X (Proposition 2.14).

Theorem 4.1 gives us a ray $r \subseteq X$ at bounded distance from M , and a constant K such that, for every segment $\sigma \subseteq r$ of length ℓ , the set $\text{Hull}_X(\sigma)$ contains points at distance $\geq \lfloor \frac{\ell}{K} \rfloor$ from M .

Fix a basepoint $p \in M$ and, for every $x \in \Delta$, denote $\mathcal{O}_x := W_{\Delta - \{x\}} \cdot p$. Since $W_{\Delta - \{x\}}$ is convex-cocompact in X , Proposition 2.14 guarantees that there exists a constant $R \geq 0$ such that $m(\mathcal{O}_x, \mathcal{O}_x, X) \subseteq \mathcal{N}_R(\mathcal{O}_x)$ for every $x \in \Delta$. Setting $\delta := \dim X$, Remark 4.6 implies $\text{Hull}_X(\mathcal{O}_x) \subseteq \mathcal{N}_{2^\delta R}(\mathcal{O}_x) \subseteq \mathcal{N}_{2^\delta R}(M)$, hence $\text{Hull}_X(g\mathcal{O}_x) \subseteq \mathcal{N}_{2^\delta R}(M)$ for every $g \in W_\Delta$.

It follows that, for every constant $C \geq 0$, there exists a constant $C' \geq 0$ such that the intersection between r and the C -neighbourhood of any $g\mathcal{O}_x$ with $g \in W_\Delta$ has diameter at most C' .

Now, by the Milnor-Schwarz lemma, there exists a W_Γ -equivariant quasi-isometry $q : D_\Gamma \rightarrow X$ with $q(D_\Delta) \subseteq M$. Let $\alpha \subseteq D_\Delta$ be a quasi-geodesic edge path such that $d_{\text{Haus}}(q(\alpha), r) < +\infty$. By the previous paragraph and the fact that α is a quasi-geodesic, we see that α satisfies the hypothesis of Lemma 5.1.

If Δ^\perp were a clique, then Lemma 5.1 would show that α , and hence r , is Morse. However, this would imply $d_{\text{Haus}}(r, \text{Hull}_X(r)) < +\infty$, contradicting the fact that $\text{Hull}_X(r)$ contains points arbitrarily far from M .

Thus, there must exist vertices $x, y \in \Delta^\perp$ that are not connected by an edge. The subgroup $\langle x, y \rangle$ is infinite and it is contained in $Z_{W_\Gamma}(W_\Delta) = W_{\Delta^\perp}$, proving the proposition. \square

We are finally ready to prove Theorem B(2), which is the following result.

Theorem 5.3. *Let $W_\Gamma \curvearrowright X$ be a cocompact cubulation where, for all $x, y \in \Gamma$, the subgroup $\langle x, y \rangle$ is convex-cocompact. Then $W_\Gamma \curvearrowright X$ induces the standard coarse median structure on W_Γ .*

Proof. We prove the statement by induction on $\#\Gamma^{(0)}$. The base case where Γ is a singleton is obvious. We now discuss the inductive step.

Hyperplane-stabilisers for the action $W_\Gamma \curvearrowright D_\Gamma$ are subgroups of the form $W_{\text{lk}(v)}$ with $v \in \Gamma$. In view of Theorem 2.17, it suffices to show that these subgroups are all convex-cocompact in X .

Fix a vertex $v \in \Gamma$.

Claim 1. For every $\Delta \subseteq \text{lk}(v)$, the subgroup W_Δ is median-cocompact in X .

Proof of Claim 1. Since $W_{\text{lk}(v)}$ has index 2 in the centraliser of v , Proposition 2.10 guarantees that it is median-cocompact in X . Let $M \subseteq X$ be a median subalgebra on which $W_{\text{lk}(v)}$ acts cofinitely. By Chepoi-Roller duality, there exists a cocompact cubulation $W_{\text{lk}(v)} \curvearrowright Y$ such that $Y^{(0)}$ is equivariantly isomorphic to M as a median algebra. For every $x, y \in \text{lk}(v)$, the fact that $\langle x, y \rangle$ is convex-cocompact in X implies that it is also convex-cocompact in Y .

Since $\text{lk}(v)$ is a proper subgraph of Γ , the inductive hypothesis implies that Y induces the standard coarse median structure on $W_{\text{lk}(v)}$. Using Proposition 2.14, this implies that, for every $\Delta \subseteq \text{lk}(v)$, the subgroup W_Δ is median-cocompact in Y , and hence in X . \square

Now, suppose for the sake of contradiction that $W_{\text{lk}(v)}$ is not convex-cocompact in X . Let $\Delta_0 \subset \text{lk}(v)$ be a minimal subgraph such that W_{Δ_0} is not convex-cocompact in X . Note that Δ_0 exists and has at least 3 vertices, by our assumptions.

By Claim 1, W_{Δ_0} is median-cocompact in X and, by minimality of Δ_0 , all subgroups $W_{\Delta_0 - \{x\}}$ with $x \in \Delta_0$ are convex-cocompact in X . Minimality also implies that Δ_0 is irreducible, because of Lemma 2.7. Thus, Proposition 5.2 guarantees that $W_{\Delta_0^\perp}$ is infinite, that is, there exist $z, z' \in \Gamma$ such that $\langle z, z' \rangle \simeq D_\infty$ and $\Delta_0 \subseteq \text{lk}(z) \cap \text{lk}(z')$.

Claim 2. The subgroup $W_{\text{lk}(z) \cap \text{lk}(z')}$ is convex-cocompact in X .

Proof of Claim 2. The proof of this fact is almost identical to that of Lemma 3.24.

We have $Z_{W_\Gamma}(zz') = \langle zz' \rangle \times W_{\text{lk}(z) \cap \text{lk}(z')}$. The subgroup $\langle zz' \rangle \simeq \mathbb{Z}$ is convex-cocompact in X by our assumptions, since it has finite index in $\langle z, z' \rangle$. Choosing $n \geq 1$ such that $(zz')^n$ acts non-transversely on X , Lemma 2.4(4) implies that $Z_{W_\Gamma}((zz')^n)$ is convex-cocompact in X . Note that $Z_{W_\Gamma}((zz')^n) = Z_{W_\Gamma}(zz')$. In addition, the index-2 subgroup of $W_{\text{lk}(z) \cap \text{lk}(z')}$ consisting of words of even length is generated by elements uu' with $u, u' \in \text{lk}(z) \cap \text{lk}(z')$, which all generate infinite cyclic subgroups that are convex-cocompact in X , by assumption. Finally, Lemma 2.6 implies that $W_{\text{lk}(z) \cap \text{lk}(z')}$ is convex-cocompact in X . \square

Now, let $C \subseteq X$ be a convex subcomplex on which $W_{\text{lk}(z) \cap \text{lk}(z')}$ acts cocompactly. Again, since $\text{lk}(z) \cap \text{lk}(z')$ is a proper subgraph of Γ , the inductive hypothesis implies that $W_{\text{lk}(z) \cap \text{lk}(z')}$ inherits the standard coarse median structure from its action on C . Using Proposition 2.14, it follows that W_{Δ_0} is convex-cocompact in C , and hence in X . This is the required contradiction. \square

Now that the proof of Theorem B is complete, it is interesting to discuss why the Conjecture from the Introduction is harder.

Remark 5.4. Say that we have an ‘exotic’ cocompact cubulation $W_\Gamma \curvearrowright X$, that is, one that does not induce the coarse median structure of the Davis complex. Suppose that we are interested in showing that a different cubulation $W_\Gamma \curvearrowright Y$ induces the same coarse median structure, under the assumption that the two cubulations give the same coarse median structure on all abelian subgroups.

The only tool we have for this at the moment is Theorem 2.17, which would require us to show that the hyperplane-stabilisers of the action $W_\Gamma \curvearrowright X$ are convex-cocompact with respect to the cocompact cubulation $W_\Gamma \curvearrowright Y$.

In the ‘standard’ case, we were able to exploit the fact that hyperplane-stabilisers of the Davis complex are centralisers, and the fact that centralisers are universally median-cocompact (Proposition 2.10). While still far from convex-cocompactness, median-cocompactness gives us powerful information that puts Theorem 4.1 in motion (recall the cautionary Example 4.2 without this assumption).

For exotic cubulations there is no guarantee that hyperplane-stabilisers will behave as nicely. We can follow this blueprint only if we can find cubulations with ‘good’ hyperplane-stabilisers representing our exotic coarse median structures. Here ‘good’ would need to mean ‘universally median-cocompact’ (i.e. median-cocompact in every cocompact cubulation of W_Γ), or significantly new ideas would be required.

5.3 | Loose and bonded squares

In this subsection, we prove Corollary C by combining Theorem B(2) with the cubical flat torus theorem [62] and our study of cubical coarse medians on products of dihedrals (Proposition 3.11). Rather than with the *bonded squares* mentioned in the Introduction, it will be convenient to work with their opposite: *loose squares* (Definition 5.6 below).

Let Γ be a finite simplicial graph. For simplicity, we say that a *square* is an induced 4-cycle $\Delta \subseteq \Gamma$. A *hyperoctahedron* is an induced subgraph $\Lambda \subseteq \Gamma$ whose opposite graph is a union of pairwise disjoint edges (in other words, $W_\Lambda \simeq D_\infty^n$ for some $n \geq 0$).

It is worth remarking that maximal virtually abelian subgroups of the right-angled Coxeter group W_Γ are not always *highest* in the sense of Subsection 2.2. As a consequence, they might not be convex-cocompact in all cocompact cubulations of W_Γ . Instead, if $\Lambda \subseteq \Gamma$ is a maximal hyperoctahedron, then the virtually abelian subgroup $W_\Lambda \simeq D_\infty^n$ is necessarily highest in W_Γ .

The following equivalent conditions characterise loose squares $\Delta \subseteq \Gamma$.

Lemma 5.5. *For a square $\Delta \subseteq \Gamma$, the following conditions are equivalent.*

- (1) *For every maximal induced subgraph $\Lambda \subseteq \Gamma$ such that W_Λ is virtually abelian, either $\Delta \subseteq \Lambda$ or $W_{\Delta \cap \Lambda}$ is finite.*
- (2) *For every maximal hyperoctahedron $\Lambda \subseteq \Gamma$, either $\Delta \subseteq \Lambda$ or $W_{\Delta \cap \Lambda}$ is finite.*
- (3) *For every square $\Delta' \subseteq \Gamma$, either $W_{\Delta \cup \Delta'}$ is virtually abelian or $W_{\Delta \cap \Delta'}$ is finite.*
- (4) *There does not exist a square $\Delta' \subseteq \Gamma$ such that $\Delta \cap \Delta'$ has exactly 3 vertices.*

Proof. We begin with (1) \Rightarrow (2). If $\Lambda \subseteq \Gamma$ is a maximal hyperoctahedron and $\Lambda' \subseteq \Gamma$ is a maximal subgraph such that $W_{\Lambda'}$ is virtually abelian and $\Lambda \subseteq \Lambda'$, then $W_{\Lambda'} = W_\Lambda \times (\mathbb{Z}/2\mathbb{Z})^m$ for some $m \geq 0$. By (1), either $\Delta \subseteq \Lambda'$ or $W_{\Delta \cap \Lambda'}$ is finite, which yields the analogous statement for Λ .

Let us prove (2) \Rightarrow (3), or rather $\neg(3) \Rightarrow \neg(2)$. Let $\Delta' \subseteq \Gamma$ be a square such that $W_{\Delta \cup \Delta'}$ is not virtually abelian and $W_{\Delta \cap \Delta'}$ is infinite. Choose a maximal hyperoctahedron $\Lambda \supseteq \Delta'$. Then $\Delta \not\subseteq \Lambda$, since $W_{\Delta \cup \Delta'}$ is not virtually abelian, unlike W_Λ . In addition, $W_{\Delta \cap \Lambda}$ is infinite, since it contains $W_{\Delta \cap \Delta'}$. Thus, Λ is a hyperoctahedron witnessing $\neg(2)$.

Regarding the implication (3) \Rightarrow (4), cyclically label by a, b, c, d the vertices of Δ . If there existed a square Δ' with vertices a, b, c, x and $x \neq d$, the subgroup $W_{\Delta \cap \Delta'}$ would be infinite. Moreover, we would have $W_{\Delta \cup \Delta'} = \langle a, c \rangle \times \langle b, d, x \rangle$, with $\langle b, d, x \rangle$ isomorphic to either $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ or $\mathbb{Z}/2\mathbb{Z} * (\mathbb{Z}/2\mathbb{Z})^2$. In particular, $W_{\Delta \cup \Delta'}$ would not be virtually abelian, contradicting (3).

Finally, we prove (4) \Rightarrow (1). Suppose $\Lambda \subseteq \Gamma$ is a maximal induced subgraph such that W_Λ is virtually abelian. Suppose further that $W_{\Delta \cap \Lambda}$ is a proper, infinite subgroup of W_Δ . Cyclically labelling the vertices of Δ by a, b, c, d , we can assume $a, c \in \Lambda$ and $b \notin \Lambda$. Moreover, we have a splitting $W_\Lambda = \langle a, c \rangle \times D_\infty^n \times (\mathbb{Z}/2\mathbb{Z})^m$ with $n + m \geq 1$, by maximality of Λ .

For a vertex $v \in \Lambda - \{a, c\}$, neither $\{a, b, c, v\}$ nor $\{a, v, c, d\}$ can be the vertex set of square distinct from Δ , because of (4). This implies that either $v = d$ or v commutes with both b and d . It follows that $W_{\Delta \cup \Lambda}$ is virtually abelian, which contradicts maximality of Λ , since $b \notin \Lambda$.

This completes the proof of the lemma. □

Definition 5.6. A square $\Delta \subseteq \Gamma$ is *loose* if it satisfies the equivalent conditions in Lemma 5.5. A square is *bonded* if it is not loose.

This is equivalent to the definition given in the Introduction: A square $\Delta \subseteq \Gamma$ is bonded if there exists another square $\Delta' \subseteq \Gamma$ such that the intersection $\Delta \cap \Delta'$ has exactly 3 vertices.

Corollary C claims that W_Γ satisfies coarse cubical rigidity when the intersection pattern of squares in Γ is sufficiently intricate, namely when every square is bonded. We are finally ready to prove this statement.

Proof of Corollary C. Suppose Γ has no loose squares. Let $W_\Gamma \curvearrowright Y$ be a cocompact cubulation.




FIGURE 5 The square $abcd$ is loose.

Let $\Lambda \subseteq \Gamma$ be a maximal hyperoctahedron. Thus, $W_\Lambda \simeq D_\infty^n$ for some $n \geq 1$ and W_Λ is a highest virtually abelian subgroup of W_Γ , as defined in Subsection 2.2. By Lemma 2.4(3), W_Λ is convex-cocompact in Y . By Proposition 3.11, either W_Λ inherits the standard coarse median structure from Y (i.e. that of the Davis complex for W_Λ), or there exists a square $\Delta = \{x_1, x_2, x_3, x_4\} \subseteq \Lambda$ such that W_Δ is convex-cocompact in Y , but the subgroups $\langle x_1, x_3 \rangle \simeq \langle x_2, x_4 \rangle \simeq D_\infty$ are not.

However, since the square Δ cannot be loose, Lemma 5.5(2) guarantees the existence of a maximal hyperoctahedron $\Lambda' \subseteq \Gamma$ such that $W_\Delta \cap W_{\Lambda'}$ is a proper, infinite subgroup of W_Δ . By Lemma 2.4(3), the subgroup $W_{\Lambda'}$ is convex-cocompact in Y and, by Lemma 2.4(1), so is the intersection $W_\Delta \cap W_{\Lambda'}$. Now, this intersection is commensurable to either $\langle x_1, x_3 \rangle$ or $\langle x_2, x_4 \rangle$, showing that at least one of them is convex-cocompact in Y .

Combined with the previous paragraph, this proves that, for every maximal hyperoctahedron $\Lambda \subseteq \Gamma$, the action $W_\Gamma \curvearrowright Y$ induces the standard coarse median structure on the subgroup W_Λ . If $x, y \in \Gamma$ are vertices with $\langle x, y \rangle \simeq D_\infty$, we certainly have $\{x, y\} \subseteq \Lambda$ for some maximal hyperoctahedron Λ , so the subgroup $\langle x, y \rangle$ is convex-cocompact in Y .

Now, Theorem 5.3 implies that Y induces the standard coarse median structure on W_Γ , proving the corollary. □

We conclude the section by giving an example of a right-angled Coxeter group W_Γ that fails to satisfy coarse cubical rigidity. Of course, we have seen that $W_\Gamma = D_\infty^n$ is such a group for $n \geq 2$, and it is also easy to come up with examples splitting over finite subgroups.

Instead, the group we are about to exhibit is one-ended and directly irreducible.

Example 5.7. Consider the graph Γ from Figure 5. Let Γ_1 be the subgraph spanned by the vertices $\{a, b, c, d\}$, let Γ_2 be the subgraph spanned by the vertices $\{a, e, f, c\}$, and let Γ_3 spanned by the vertices $\{a, c\}$. We have the following amalgamated product decomposition $W_\Gamma = W_{\Gamma_1} *_{W_{\Gamma_3}} W_{\Gamma_2}$.

Let X_1 be the usual square tiling of \mathbb{E}^2 . See Figure 6. Each vertex $x \in \Gamma_1$ acts by reflection about the line l_x . Note that the subcomplex l is preserved by W_{Γ_3} .

Let C be a 2-cube with opposite pairs of vertices identified, and let X_2 be the universal cover of C (see Figure 7). We define an action of W_{Γ_2} on X_2 by letting e and f act by reflections about the lines l_e and l_f respectively. Let a (resp. c) act as a reflection about the line perpendicular to l_e (resp. l_f) and through the vertex v_a (resp. v_c). Note that there is an infinite subcomplex l' (shown in blue) that is preserved by W_{Γ_3} .

Let T be the Bass–Serre tree associated to $W_\Gamma = W_{\Gamma_1} *_{W_{\Gamma_3}} W_{\Gamma_2}$. We define a blowup X of T . We blow up each vertex v of T corresponding to a coset of W_{Γ_1} (resp. W_{Γ_2}) to a copy X_v of the complex X_1 (resp. X_2). For each edge $[u, v]$ of T corresponding to a coset gW_{Γ_3} , we glue $gl \subset X_u$ to $gl' \subset X_v$ where, up to relabelling, we are assuming that X_u is a copy of X_1 and X_v is a copy of X_2 . Even though we are glueing over non-convex subcomplexes, it is not hard to see that the cube complex X is CAT(0). The result is a cocompact cubulation $W_\Gamma \curvearrowright X$ that is not strongly cellular.

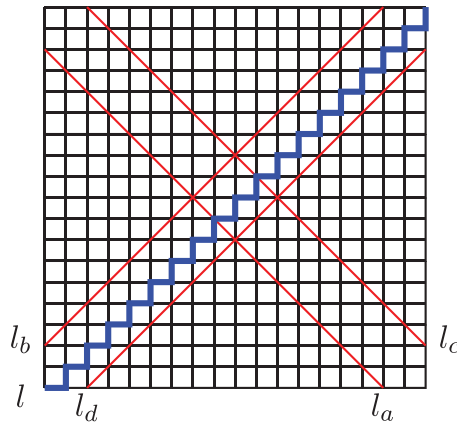


FIGURE 6 Action of W_{Γ_1} on X_1 .




FIGURE 7 Action of W_{Γ_2} on X_2 .

The induced coarse median structure on W_Γ is not the standard one, since $\langle a, c \rangle$ is not convex-cocompact in X , as is evident from Figure 6.

The previous example suggests that it should always be possible to exploit loose squares in Γ to produce cocompact cubulations of W_Γ inducing a coarse median structure other than that of the Davis complex. If this is indeed the case, then Corollary C is sharp.

6 | GENERAL COXETER GROUPS

In this section, we prove Theorem D(2).

Let (W, S) be a general Coxeter system. A subgroup of W is *special* if it is generated by a subset of S , and it is *parabolic* if it is conjugate to a special subgroup. An element of W is an *involution* if it has order 2, and it is a *reflection* if it is conjugate to an element of S .

Involutions in Coxeter groups were fully classified by Richardson [51]. We record his results in the following lemma. For a more recent account, the reader can also consult [43, sections 27.2–27.4].

Lemma 6.1.

- (1) Let P be a finite irreducible Coxeter group. The centre $Z(P)$ is non-trivial if and only if $Z(P) \simeq \mathbb{Z}/2\mathbb{Z}$. In this case, $Z(P) = \langle w_p \rangle$ for an element $w_p \in P$ that is the longest element of P with respect to any Coxeter generating set.
- (2) Let W be a general Coxeter group. Every involution in W is the longest element w_p of a parabolic subgroup $P \leq W$ of the form $P = P_1 \times \dots \times P_k$, where each P_i is a finite irreducible Coxeter group with non-trivial centre. In particular, we have $w_p = w_{P_1} \cdot \dots \cdot w_{P_k}$.

Remark 6.2. Let (W, S) be a Coxeter system with W finite. Let $w_o \in W$ be the longest element with respect to S . Since w_o is longest, any reduced expression for w_o will involve all elements of S , so w_o does not lie in any proper special subgroup of W . In general, however, w_o can lie in a proper parabolic subgroup of W . For instance, this happens when W is the dihedral group with 6 elements.

Things are different if we suppose $W = P$, where P has the form from Lemma 6.1(2). In this case, the longest element $w_p \in P$ is not contained in any proper parabolic subgroup of P . Indeed, since w_p is central, it is contained in a proper parabolic subgroup if and only if it is contained in a proper special subgroup of W (keeping a Coxeter generating set fixed).

Lemma 6.3. Let W be a Coxeter group and let $W \curvearrowright \mathcal{X}$ denote its Niblo–Reeves cubulation. Stabilisers of hyperplanes of \mathcal{X} are precisely conjugates of centralisers $Z_W(r)$ with $r \in W$ a reflection.

Proof. Let Σ be the presentation 2-complex of W associated with the standard presentation. For all $s, t \in S$ with $(st)^m = 1$, there is a W -orbit of $2m$ -gons in Σ with edges alternately labelled by s and t . There are some natural walls in Σ , uniquely determined by the following property: If \mathfrak{w} is a wall and $C \subseteq \Sigma$ is a 2-cell, then $\mathfrak{w} \cap C$ is either empty or a segment joining midpoints of opposite edges of C . The Niblo–Reeves cubulation of W is precisely the CAT(0) cube complex associated with this collection of walls [48].

For each reflection $r \in S$, there is a unique wall $\mathfrak{w}_r \subseteq \Sigma$ intersecting the edge $[1, r]$ in its midpoint. Every wall in Σ is in the W -orbit of one of the walls \mathfrak{w}_r . Thus, it suffices to show that the W -stabiliser of the wall $\mathfrak{w}_r \subseteq \Sigma$ is precisely the centraliser $Z_W(r)$.

Note that r fixes the wall \mathfrak{w}_r pointwise and swaps its two sides. In fact, r is the only element of W with this property, since W acts freely on the 0-skeleton of Σ . Now, if $g \in W$, we have $g\mathfrak{w}_r = \mathfrak{w}_r$ if and only if grg^{-1} again fixes \mathfrak{w}_r pointwise and swaps its two sides. Thus, $g\mathfrak{w}_r = \mathfrak{w}_r$ if and only if $grg^{-1} = r$, that is, $g \in Z_W(r)$ as required. □

Lemma 6.4. Let W be Coxeter group with cocompact Niblo–Reeves cubulation $W \curvearrowright \mathcal{X}$. Then, for every involution $\sigma \in W$, the centraliser $Z_W(\sigma)$ is convex-compact in \mathcal{X} .

Proof. In view of Lemma 6.1, we have $\sigma = w_p$ for a finite, parabolic subgroup $P = P_1 \times \dots \times P_k$, where each P_i is finite and irreducible.

We begin by showing that the centraliser $Z_W(w_p)$ is contained in the normaliser $N_W(P)$. Indeed, if $g \in W$ commutes with w_p , then $w_p \in P \cap gPg^{-1}$. Note that $P \cap gPg^{-1}$ is a parabolic subgroup

of P by [25, Lemma 5.3.6], but it cannot be a *proper* parabolic subgroup because of Remark 6.2. We conclude that $P \leq gPg^{-1}$ and, since P and gPg^{-1} are finite groups of the same cardinality, we must have $gPg^{-1} = P$. This shows $g \in N_W(P)$, as required.

Now, we have a chain of inclusions $Z_W(P) \leq Z_W(w_P) \leq N_W(P)$. Since P is finite, $Z_W(P)$ has finite index in $N_W(P)$, hence $Z_W(P)$ has finite index in $Z_W(w_P)$ as well.

Finally, observe that $Z_W(P)$ is convex-cocompact in \mathcal{X} . Indeed, P is generated by finitely many reflections r_1, \dots, r_k . Each centraliser $Z_W(r_i)$ is the stabiliser of a hyperplane of \mathcal{X} by Lemma 6.3, hence it is convex-cocompact in \mathcal{X} . Thus, $Z_W(P) = \bigcap_i Z_W(r_i)$ is convex-cocompact in \mathcal{X} by Lemma 2.4(1). Lemma 2.4(2) implies that $Z_W(w_P)$ is convex-cocompact, completing the proof. \square

Proof of Theorem D(2). Let $W \curvearrowright \mathcal{X}$ be the Niblo–Reeves cubulation of W , which is cocompact by the assumptions of the theorem. Consider an automorphism $\varphi \in \text{Aut}(W)$ and let $W \curvearrowright \mathcal{X}^\varphi$ denote the standard action on the Niblo–Reeves cubulation precomposed with φ .

We want to show that \mathcal{X} and \mathcal{X}^φ induce the same coarse median structure on W by invoking Theorem 2.17. By Lemma 6.3, up to conjugacy, stabilisers of hyperplanes of \mathcal{X}^φ are of the form $\varphi^{-1}(Z_W(r))$, with $r \in W$ a reflection. Observing that $\varphi^{-1}(Z_W(r)) = Z_W(r')$ for the involution $r' := \varphi^{-1}(r)$, Lemma 6.4 shows that all these subgroups are convex-cocompact in \mathcal{X} , as required for Theorem 2.17. \square

APPENDIX: COCOMPACTNESS OF INTERSECTIONS

A collection \mathcal{C} of subsets of a metric space X is *locally finite* if every ball in X intersects only finitely many elements of \mathcal{C} .

Lemma A.1. *Let $G \curvearrowright X$ be a proper cocompact action on a metric space. A closed subset $A \subseteq X$ is acted upon cocompactly by its G -stabiliser if and only if the orbit $G \cdot A$ is locally finite.*

Proof. The backward arrow is [42, Lemma 2.3], since X is a proper metric space (as it admits a geometric group action).

For the forward arrow, denote by $G_A \leq G$ the stabiliser of A and assume that the action $G_A \curvearrowright A$ is cocompact. Suppose for the sake of contradiction that a ball $B \subseteq X$ intersects infinitely many pairwise distinct translates $g_n A$ with $g_n \in G$. Then the balls $g_n^{-1} B$ all intersect A . Since $G_A \curvearrowright A$ is cocompact and $G \curvearrowright X$ is proper, the elements g_n^{-1} all lie in a product set $G_A \cdot F$ with $F \subseteq G$ finite. Hence all g_n lie in $F^{-1} \cdot G_A$, contradicting that there are infinitely many distinct sets $g_n A$. \square

Lemma A.2 (Cocompact Intersections). *Let $G \curvearrowright X$ be a proper cocompact action on a metric space. Let $A, B \subseteq X$ be closed subsets that are invariant and acted upon cocompactly by subgroups $H, K \leq G$, respectively. Then the action $H \cap K \curvearrowright A \cap B$ is cocompact (possibly, $A \cap B = \emptyset$).*

Proof. Suppose $A \cap B \neq \emptyset$. We first prove that $A \cap B$ is acted upon cocompactly by its G -stabiliser. In view of Lemma A.1, it suffices to show that the orbit $G \cdot (A \cap B)$ is locally finite.

Suppose for the sake of contradiction that $g_n(A \cap B)$ are pairwise distinct translates intersecting a ball $L \subseteq X$. The collections $\{g_n A\}$ and $\{g_n B\}$ also intersect L , so they must be finite by Lemma A.1. Hence there are only finitely many possible intersections $g_n A \cap g_n B$, contradicting our assumption.

Now, we know that $A \cap B$ is acted upon cocompactly by its G -stabiliser $G_{A \cap B}$. A finite-index subgroup of $G_{A \cap B}$ must stabilise A , since $A \cap B$ is contained in all $G_{A \cap B}$ -translates of A , and there are only finitely many such translates of A by local finiteness. Since H acts cocompactly on A , it must have finite index in the G -stabiliser of A , hence a finite-index subgroup of $G_{A \cap B}$ is contained

in H (and, similarly, in K). In conclusion, a finite-index subgroup of $G_{A \cap B}$ acts cocompactly on $A \cap B$ and is contained in $H \cap K$, which concludes the proof. \square

ACKNOWLEDGEMENTS

We thank Jason Behrstock and Mark Hagen for interesting comments, Thomas Haettel for mentioning [33, 35] in relation to Corollary E in June 2023 and George Shaji for catching a couple of errors in an earlier version of this preprint. We are also grateful to the anonymous referees for their many helpful suggestions, particularly to one of them for proposing the terminology ‘bonded square’.

Open access funding enabled and organized by Projekt DEAL.

JOURNAL INFORMATION

The *Journal of Topology* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

ORCID

Elia Fioravanti  <https://orcid.org/0000-0001-5038-9876>

REFERENCES

1. D. A. Barkauskas, *Centralizers in graph products of groups*, *J. Algebra* **312** (2007), no. 1, 9–32.
2. C. Bregman, R. Charney, and K. Vogtmann, *Outer space for RAAGs*, *Duke Math. J.* **172** (2023), no. 6, 1033–1108.
3. J. Beyrer and E. Fioravanti, *Cross-ratios on CAT(0) cube complexes and marked length-spectrum rigidity*, *J. Lond. Math. Soc. (2)* **104** (2021), no. 5, 1973–2015.
4. J. Beyrer and E. Fioravanti, *Cross ratios and cubulations of hyperbolic groups*, *Math. Ann.* **384** (2022), no. 3–4, 1547–1592.
5. M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 319, Springer, Berlin, 1999.
6. J. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, *Geometry and rigidity of mapping class groups*, *Geom. Topol.* **16** (2012), no. 2, 781–888.
7. M. Burger and S. Mozes, *Lattices in product of trees*, *Inst. Hautes Études Sci. Publ. Math.* **92** (2000), 151–194 (2001).
8. H.-J. Bandelt, H. M. Mulder, and E. Wilkeit, *Quasi-median graphs and algebras*, *J. Graph Theory* **18** (1994), no. 7, 681–703.
9. B. H. Bowditch, *Coarse median spaces and groups*, *Pacific J. Math.* **261** (2013), no. 1, 53–93.
10. B. H. Bowditch, *Convex hulls in coarse median spaces*, <https://homepages.warwick.ac.uk/masgag/papers/hulls-cms.pdf>, 2018.
11. B. Bowditch, *Median algebras*, <https://homepages.warwick.ac.uk/~masgag/preprints.html>, 2022.
12. N. Bergeron and D. T. Wise, *A boundary criterion for cubulation*, *Amer. J. Math.* **134** (2012), no. 3, 843–859.
13. P.-E. Caprace, *Finite and infinite quotients of discrete and indiscrete groups*, *Groups St Andrews 2017 in Birmingham*, *London Math. Soc. Lecture Note Ser.*, vol. 455, Cambridge Univ. Press, Cambridge, 2019, pp. 16–69.
14. J. Chalopin, V. Chepoi, H. Hirai, and D. Osajda, *Weakly modular graphs and nonpositive curvature*, *Mem. Amer. Math. Soc.* **268** (2020) no. 1309, vi+159.
15. R. Charney, J. Crisp, and K. Vogtmann, *Automorphisms of 2-dimensional right-angled Artin groups*, *Geom. Topol.* **11** (2007), 2227–2264.
16. I. Chatterji, T. Fernós, and A. Iozzi, *The median class and superrigidity of actions on CAT(0) cube complexes*, *J. Topol.* **9** (2016), no. 2, 349–400. With an appendix by Pierre-Emmanuel Caprace.

17. V. Chepoi, *Graphs of some CAT(0) complexes*, Adv. in Appl. Math. **24** (2000), no. 2, 125–179.
18. P.-E. Caprace and B. Mühlherr, *Reflection triangles in Coxeter groups and biautomaticity*, J. Group Theory **8** (2005), no. 4, 467–489.
19. P.-E. Caprace and B. Mühlherr, *Reflection rigidity of 2-spherical Coxeter groups*, Proc. Lond. Math. Soc. (3) **94** (2007), no. 2, 520–542.
20. P.-E. Caprace and M. Sageev, *Rank rigidity for CAT(0) cube complexes*, Geom. Funct. Anal. **21** (2011), no. 4, 851–891.
21. R. Charney and H. Sultan, *Contracting boundaries of CAT(0) spaces*, J. Topol. **8** (2015), no. 1, 93–117.
22. R. Charney, N. Stambaugh, and K. Vogtmann, *Outer space for untwisted automorphisms of right-angled Artin groups*, Geom. Topol. **21** (2017), no. 2, 1131–1178.
23. M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, 91–119.
24. M. W. Davis, *Buildings are CAT(0)*, Geometry and cohomology in group theory (Durham, 1994), London Math. Soc. Lecture Note Ser., vol. 252, Cambridge Univ. Press, Cambridge, 1998, pp. 108–123.
25. M. W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008.
26. E. Fioravanti and M. F. Hagen, *Deforming cubulations of hyperbolic groups*, J. Topol. **14** (2021), no. 3, 877–912.
27. W. N. Franzsen, R. B. Howlett, and B. Mühlherr, *Reflections in abstract Coxeter groups*, Comment. Math. Helv. **81** (2006), no. 3, 665–697.
28. E. Fioravanti, *Superrigidity of actions on finite rank median spaces*, Adv. Math. **352** (2019), 1206–1252.
29. E. Fioravanti, *Convex cores for actions on finite-rank median algebras*, Annales de l’institut Fourier, arXiv:2101.03060v1, 2021.
30. E. Fioravanti, *On automorphisms and splittings of special groups*, Compos. Math. **159** (2023), no. 2, 232–305.
31. E. Fioravanti, *Coarse-median preserving automorphisms*, Geom. Topol. **28** (2024), no. 1, 161–266.
32. A. Genevois, *Cubical-like geometry of quasi-median graphs and applications to geometric group theory*, arXiv:1712.01618, 2017.
33. A. Genevois, *Median sets of isometries in CAT(0) cube complexes and some applications*, Michigan Math. J. **71** (2022), no. 3, 487–532.
34. A. Genevois and A. Martin, *Automorphisms of graph products of groups from a geometric perspective*, Proc. Lond. Math. Soc. (3) **119** (2019), no. 6, 1745–1779.
35. T. Haettel, *Virtually cocompactly cubulated Artin-Tits groups*, Int. Math. Res. Not. IMRN (2021), no. 4, 2919–2961.
36. F. Haglund, *Finite index subgroups of graph products*, Geom. Dedicata **135** (2008), 167–209.
37. M. F. Hagen, *Cocompactly cubulated crystallographic groups*, J. Lond. Math. Soc. (2) **90** (2014), no. 1, 140–166.
38. F. Haglund, *Isometries of CAT(0) cube complexes are semi-simple*, Ann. Math. Qué. **47** (2023), no. 2, 249–261.
39. U. Hamenstädt, *Geometry of the mapping class groups III: Quasi-isometric rigidity*, arXiv:math/0512429, 2007.
40. J. Huang and B. Kleiner, *Groups quasi-isometric to right-angled Artin groups*, Duke Math. J. **167** (2018), no. 3, 537–602.
41. M. F. Hagen and H. Petyt, *Projection complexes and quasimedial maps*, Algebr. Geom. Topol. **22** (2022), no. 7, 3277–3304.
42. M. F. Hagen and T. Susse, *On hierarchical hyperbolicity of cubical groups*, Israel J. Math. **236** (2020), no. 1, 45–89.
43. R. Kane, *Reflection groups and invariant theory*, CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC, vol. 5, Springer, New York, 2001.
44. J. Kahn and V. Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. (2) **175** (2012), no. 3, 1127–1190.
45. R. Kropholler and C. O’Donnell, *Groups with arbitrary cubical dimension gap*, arXiv:1912.05055, 2019.
46. M. R. Laurence, *A generating set for the automorphism group of a graph group*, J. Lond. Math. Soc. (2) **52** (1995), no. 2, 318–334.
47. G. A. Niblo and L. D. Reeves, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology **37** (1998), no. 3, 621–633.
48. G. A. Niblo and L. D. Reeves, *Coxeter groups act on CAT(0) cube complexes*, J. Group Theory **6** (2003), no. 3, 399–413.

49. G. A. Niblo, N. Wright, and J. Zhang, *A four point characterisation for coarse median spaces*, Groups Geom. Dyn. **13** (2019), no. 3, 939–980.
50. H. Petyt, *Mapping class groups are quasicubical*, Amer. J. Math., arXiv:2112.10681, 2022.
51. R. W. Richardson, *Conjugacy classes of involutions in Coxeter groups*, Bull. Aust. Math. Soc. **26** (1982), no. 1, 1–15.
52. M. A. Roller, *Poc sets, median algebras and group actions. An extended study of Dunwoody's construction and Sageev's theorem*, Preprint, University of Southampton, available at arXiv:1607.07747, 1998.
53. K. Ruane and S. Witzel, *CAT(0) cubical complexes for graph products of finitely generated abelian groups*, New York J. Math. **22** (2016), 637–651.
54. M. Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. Lond. Math. Soc. (3) **71** (1995), no. 3, 585–617.
55. M. Sageev, *CAT(0) cube complexes and groups*, Geometric group theory, IAS/Park City Math. Ser., vol. 21, Amer. Math. Soc., Providence, RI, 2014, pp. 7–54.
56. H. Servatius, *Automorphisms of graph groups*, J. Algebra **126** (1989), no. 1, 34–60.
57. M. Sageev and D. T. Wise, *The Tits alternative for CAT(0) cubical complexes*, Bull. Lond. Math. Soc. **37** (2005), no. 5, 706–710.
58. B. T. Williams, *Two topics in geometric group theory*, Ph.D. thesis, University of Southampton, 1998.
59. D. T. Wise, *Non-positively curved squared complexes: aperiodic tilings and non-residually finite groups*, ProQuest LLC, Ann Arbor, MI, 1996. Ph.D. thesis, Princeton University.
60. D. T. Wise, *From riches to RAAGs: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.
61. N. Wright, *Finite asymptotic dimension for CAT(0) cube complexes*, Geom. Topol. **16** (2012), no. 1, 527–554.
62. D. T. Wise and D. J. Woodhouse, *A cubical flat torus theorem and the bounded packing property*, Israel J. Math. **217** (2017), no. 1, 263–281.