Geometry & Topology

Volume 28 (2024)

Coarse-median preserving automorphisms

Elia Fioravanti

msp

Geometry & Topology 28:1 (2024) 161–266
DOI: 10.2140/gt.2024.28.161
Published: 27 February 2024

Coarse-median preserving automorphisms

ELIA FIORAVANTI

This paper has three main goals.

First, we study fixed subgroups of automorphisms of right-angled Artin and Coxeter groups. If φ is an untwisted automorphism of a RAAG, or an arbitrary automorphism of a RACG, we prove that $\mathrm{Fix}\,\varphi$ is finitely generated and undistorted. Up to replacing φ with a power, we show that $\mathrm{Fix}\,\varphi$ is quasiconvex with respect to the standard word metric. This implies that $\mathrm{Fix}\,\varphi$ is a virtual retract and a special group in the sense of Haglund and Wise.

By contrast, there exist "twisted" automorphisms of RAAGs for which $\mathrm{Fix}\,\varphi$ is undistorted but not of type F (hence not special), of type F but distorted, or even infinitely generated.

Secondly, we introduce the notion of "coarse-median preserving" automorphism of a coarse median group, which plays a key role in the above results. We show that automorphisms of RAAGs are coarse-median preserving if and only if they are untwisted. On the other hand, all automorphisms of Gromov-hyperbolic groups and right-angled Coxeter groups are coarse-median preserving. These facts also yield new or more elementary proofs of Nielsen realisation for RAAGs and RACGs.

Finally, we show that, for every special group G (in the sense of Haglund and Wise), every infinite-order, coarse-median preserving outer automorphism of G can be realised as a homothety of a finite-rank median space X equipped with a "moderate" isometric G–action. This generalises the classical result, due to Paulin, that every infinite-order outer automorphism of a hyperbolic group H projectively stabilises a small H–tree.

20F65, 20F67; 20E36, 20F28, 20F34, 20F36, 20F55

1. Introduction 162
2. Preliminaries 174
3. Cubical convex-cocompactness 189
4. Fixed subgroups of CMP automorphisms 203
5. Invariant splittings of RAAGs 223
6. Projectively invariant metrics on finite-rank median algebras 230
7. Ultralimits and coarse-median preserving automorphisms 240
Appendix. Measurable partitions of halfspace-intervals 260
References 262

1 Introduction

This paper is inspired by the following, at first sight unrelated, questions.

Question 1 Given a finitely generated group G and $\varphi \in \operatorname{Aut} G$, what is the structure of the subgroup of fixed points $\operatorname{Fix}\varphi \leq G$?

Question 2 Given a finitely generated group G and $\varphi \in \operatorname{Aut} G$, when can we realise φ as a homothety of a nonpositively curved metric space X equipped with a "nice" G–action by isometries?

Our motivation comes from the theory of automorphisms of free groups. When $G = F_n$, a complete answer to Question 1 was first conjectured by Peter Scott in 1978, and later proved — after work by Dyer and Scott [46], Jaco and Shalen [72], Gersten [56; 57], Culler [36], Goldstein and Turner [58], Cooper [33] and Cohen and Lustig [32], among others — by Bestvina and Handel [12]:

> *For every* $\varphi \in \operatorname{Aut} F_n$*, the fixed subgroup* $\operatorname{Fix}\varphi \leq F_n$ *is generated by at most* n *elements.*

In particular, $\operatorname{Fix}\varphi$ is finitely generated, free, and quasiconvex in F_n.

Bestvina and Handel's proof is based on the extension of several ideas of Nielsen–Thurston theory from surfaces to graphs. Specifically, every homotopy equivalence between finite graphs is homotopic to a (*relative*) *train track map* [12; 11]. This result is also a key ingredient in providing the following answer to Question 2, by Gaboriau, Jaeger, Levitt and Lustig [53]:

> *For every* $\varphi \in \operatorname{Aut} F_n$*, there exists an action by homotheties* $F_n \rtimes_\varphi \mathbb{Z} \curvearrowright T$*, where* T *is an* $\mathbb{R}$*–tree and the restriction* $F_n \curvearrowright T$ *is isometric, minimal, and has trivial arc-stabilisers.*

If φ is exponentially growing, then $F_n \curvearrowright T$ has dense orbits and $\operatorname{Fix}\varphi$ is elliptic.

We are interested in Question 2 because of its connections to Question 1. Indeed, if one admits the existence of an F_n–tree as above, it is possible to give more elementary proofs of the Scott conjecture, which are completely independent of the complicated machinery of train tracks and instead rely on an "index theory" for F_n–trees; see Gaboriau, Levitt and Lustig [55] and Gaboriau and Levitt [54].

More generally, a satisfactory answer to Question 2 was obtained by Paulin [88] for all *Gromov-hyperbolic groups* G. If $\phi \in \operatorname{Out} G$ has infinite order, then it can be similarly realised as a homothety of a *small* G–tree, ie an $\mathbb{R}$–tree with a minimal isometric G–action such that no G–stabiliser of an arc contains a copy of the free group F_2.

Paulin's proof is abstract in nature, but his result can be pictured quite concretely in the case when $G = \pi_1(S)$ for a closed surface S: Thurston [97] showed that ϕ is induced by a homeomorphism of S that preserves a projective measured singular foliation on S; the $\mathbb{R}$–tree T can then be constructed by lifting this singular foliation to the universal cover $\widetilde{S}$ and considering its leaf space.

It is natural to wonder if the above discussion is specific to hyperbolic groups. This might be suggested by the fact that automorphism groups of one-ended hyperbolic groups can essentially be understood in terms of mapping class groups of finite-type surfaces (see Levitt [78] and Sela [93]), for which Nielsen–Thurston theory is available.

In recent years, the study of outer automorphisms of groups other than $\pi_1(S)$ and F_n has gained significant traction. The groups $\operatorname{Out}\mathcal{A}_\Gamma$ — where $\mathcal{A}_\Gamma$ is a right-angled Artin group (RAAG) — are particularly appealing in this context, as they can exhibit a variety of interesting behaviours ranging between the extremal cases of $\operatorname{Out} F_n$ and $\operatorname{Out}\mathbb{Z}^n = \mathrm{GL}_n\mathbb{Z}$.

One may look at the large body of work on $\operatorname{Out} F_n$ hoping to extract a blueprint that will direct the study of the groups $\operatorname{Out}\mathcal{A}_\Gamma$. This has proved a successful approach in some cases, remarkably with the definition of analogues of Outer Space (see Bregman, Charney and Vogtmann [20] and Charney, Stambaugh and Vogtmann [25]) and its consequences for the study of homological properties. However, there are limits to such analogies: in practice, techniques that are tailored to general RAAGs and based on induction on the complexity of the graph Γ seem to provide the most effective approach to many problems; see for instance Charney and Vogtmann [27; 28], Day and Wade [43], Day, Sale and Wade [42] and Guirardel and Sale [61].

Our aim is to investigate Questions 1 and 2 when G is a RAAG or, more generally, a cocompactly cubulated group. These are just two of the many basic questions that have been fully solved for $\operatorname{Out} F_n$, but have so far remained out of the limelight for the groups $\operatorname{Out}\mathcal{A}_\Gamma$.

One quickly realises that it is necessary to impose some restrictions on $\varphi \in \operatorname{Aut}\mathcal{A}_\Gamma$ if the two questions are to be fruitfully addressed. To begin with, it is not hard to construct automorphisms of $F_2 \times \mathbb{Z}$ whose fixed subgroup is infinitely generated (Example 4.13), which would prevent us from relying on the tools of geometric group theory in relation to Question 1. In addition, when $G = \mathbb{Z}^n$, it should heuristically always be possible to equivariantly collapse the space X in Question 2 to a copy of $\mathbb{R}$, which forces $\varphi \in \mathrm{GL}_n\mathbb{Z}$ to have a positive eigenvalue.

We choose to consider the subgroup of *untwisted automorphisms* $U(\mathcal{A}_\Gamma) \leq \operatorname{Aut}\mathcal{A}_\Gamma$, which was introduced by Day in [41] (with the name of "long-range automorphisms") and further studied by Charney, Stambaugh and Vogtmann [25] and Hensel and Kielak [69]. This can be defined as the subgroup generated by a certain subset of the Laurence–Servatius generators for $\operatorname{Aut}\mathcal{A}_\Gamma$ (see Laurence [75] and Servatius [94]), excluding generators that "resemble" too closely elements of $\mathrm{GL}_n\mathbb{Z}$.

The subgroup $U(\mathcal{A}_\Gamma) \leq \operatorname{Aut}\mathcal{A}_\Gamma$ displays stronger similarities to $\operatorname{Aut} F_n$ and often makes up a large portion of the entire group $\operatorname{Aut}\mathcal{A}_\Gamma$. For instance, $U(F_n) = \operatorname{Aut} F_n$ and $U(\mathcal{A}_\Gamma)$ always contains the kernel of the homomorphism $\operatorname{Aut}\mathcal{A}_\Gamma \to \mathrm{GL}_n\mathbb{Z}$ induced by the $(\operatorname{Aut}\mathcal{A}_\Gamma)$–action on the abelianisation of $\mathcal{A}_\Gamma$.

Our first result is a novel, *coarse geometric* characterisation of untwisted automorphisms. This will play a fundamental role in addressing both Questions 1 and 2 in the rest of the paper.

Recall that every right-angled Artin group $\mathcal{A}_\Gamma$ is equipped with a median operator $\mu\colon \mathcal{A}_\Gamma^3 \to \mathcal{A}_\Gamma$ coming from the fact that $\mathcal{A}_\Gamma$ is naturally identified with the 0–skeleton of a CAT(0) cube complex (the universal cover of its Salvetti complex); see Chepoi [31]. Thus, one can consider those automorphisms of $\mathcal{A}_\Gamma$ with respect to which μ is coarsely equivariant.

More generally, it makes sense to study such automorphisms for any *coarse median group* (G,μ). This remarkably broad class of groups was introduced by Bowditch in [15] and contains all Gromov-hyperbolic groups, as well as all groups admitting a geometric action on a CAT(0) cube complex, and all *hierarchically hyperbolic groups* in the sense of Behrstock, Hagen and Sisto [6, Definition 1.21].

Definition An automorphism φ of a coarse median group (G,μ) is *coarse-median preserving*[1] (*CMP*) if there exists a constant $C \geq 0$ such that

$$\varphi(\mu(g_1,g_2,g_3)) \approx_C \mu(\varphi(g_1),\varphi(g_2),\varphi(g_3)) \quad \text{for all } g_1,g_2,g_3 \in G,$$

where $x \approx_C y$ means $d(x,y) \leq C$ with respect to some fixed word metric d on G.

It is easy to see that CMP automorphisms form a subgroup of $\operatorname{Aut} G$ containing all inner automorphisms.[2] Thus, it makes sense to speak of CMP *outer* automorphisms, as this property does not depend on the specific lift to $\operatorname{Aut} G$.

It turns out that, in the setting of right-angled Artin groups, CMP automorphisms coincide with untwisted automorphisms, perhaps explaining the closer analogy between $U(\mathcal{A}_\Gamma)$ and $\operatorname{Aut} F_n$. In particular, every element of $\operatorname{Aut} F_n$ is CMP, while only a finite subgroup of $\operatorname{Aut} \mathbb{Z}^n$ is CMP.

More precisely, we have the following. We endow right-angled Artin/Coxeter groups with the coarse median structure induced by the action on the universal cover of the Salvetti/Davis complex.

Proposition A (1) *All automorphisms of hyperbolic groups are CMP.*

(2) *All automorphisms of right-angled Coxeter groups are CMP.*

(3) *Automorphisms of right-angled Artin groups are CMP if and only if they are untwisted.*

Part (1) is due to the fact that hyperbolic groups admit a unique coarse median structure, which was shown in [83]; see Example 2.28 below. That CMP automorphisms of RAAGs are untwisted can be easily deduced from the proof, due to Laurence [75], that elementary automorphisms generate the automorphism group. We prove the rest of Proposition A in Section 3.4.

[1]This terminology is motivated in Section 2.6; see Remark 2.25.

[2]Here it is important that our definition of *coarse median group* (Definition 2.24) is slightly stronger than Bowditch's original definition [15], in that we require μ to be coarsely G–equivariant. The difference between the two notions is analogous to the distinction between *hierarchically hyperbolic groups* and groups that are just a *hierarchically hyperbolic space*.

Our first result on Question 1 applies to all CMP automorphisms of *cocompactly cubulated* groups, ie those groups that admit a proper cocompact action on a CAT(0) cube complex.

We remark that, in addition to Proposition A, examples of CMP automorphisms of cubulated groups are provided by [52, Theorem E], which characterises when a generalised Dehn twist preserves the coarse median structure induced by the cubulation.

Theorem B *Let G be a cocompactly cubulated group, with the induced coarse median structure. If $\varphi \in \operatorname{Aut} G$ is coarse-median preserving, then:*

(1) $\operatorname{Fix}\varphi$ *is finitely generated and undistorted in G.*

(2) $\operatorname{Fix}\varphi$ *is itself cocompactly cubulated.*

Both parts of this result fail badly for "twisted" automorphisms of right-angled Artin groups. For every finite graph Γ, there exist automorphisms $\psi \in \operatorname{Aut}(\mathcal{A}_\Gamma \times \mathbb{Z})$ with $\operatorname{Fix}\psi = BB_\Gamma \times \mathbb{Z}$, where $BB_\Gamma \leq \mathcal{A}_\Gamma$ denotes the Bestvina–Brady subgroup [8]; see Example 4.13. When finitely generated, BB_Γ is quadratically distorted in $\mathcal{A}_\Gamma$ as soon as $\mathcal{A}_\Gamma$ is directly irreducible and noncyclic; see Tran [98]. Even when $\operatorname{Fix}\psi$ is finitely generated and undistorted, one can ensure that $\operatorname{Fix}\psi$ not be of type F, which implies that $\operatorname{Fix}\psi$ is not cocompactly cubulated. These examples can be easily extended to RAAGs that do not split as products.

We emphasise that the cubulation of $\operatorname{Fix}\varphi$ provided by Theorem B does not arise from a *convex* subcomplex of the cubulation of G in general, but just from a *median subalgebra* of it; see Section 2.2 for a definition. In fact, the subgroup $\operatorname{Fix}\varphi$ need not be *quasiconvex* in G, as can be observed for the automorphism $\varphi \in \operatorname{Aut}\mathbb{Z}^2$ that swaps the standard generators, where $\operatorname{Fix}\varphi$ is the diagonal subgroup of $\mathbb{Z}^2$.

Nevertheless, in many situations, $\operatorname{Fix}\varphi$ does turn out to be quasiconvex in the ambient group. We prove this fact in the context of right-angled Artin and Coxeter groups, where it has the remarkable consequence that $\operatorname{Fix}\varphi$ is a retract of a finite-index subgroup of the ambient group; see Haglund and Wise [68, Section 6].

Theorem C *Consider the right-angled Artin group $\mathcal{A}_\Gamma$ or the right-angled Coxeter group $\mathcal{W}_\Gamma$. There are finite-index subgroups $U_0(\mathcal{A}_\Gamma) \leq U(\mathcal{A}_\Gamma)$ and $\operatorname{Aut}_0 \mathcal{W}_\Gamma \leq \operatorname{Aut}\mathcal{W}_\Gamma$ such that, for any automorphism φ lying in either of these subgroups:*

(1) $\operatorname{Fix}\varphi$ *is quasiconvex in $\mathcal{A}_\Gamma$ or $\mathcal{W}_\Gamma$ with respect to their standard word metric, ie geodesics in their standard Cayley graph with endpoints in* $\operatorname{Fix}\varphi$ *stay uniformly close to* $\operatorname{Fix}\varphi$.

(2) *In particular,* $\operatorname{Fix}\varphi$ *is a virtual retract and it is a special group in the Haglund–Wise sense.*

For the experts, the finite-index subgroups in Theorem C are generated by the elementary automorphisms known as inversions, folds and partial conjugations; see Section 3.4 and Remark 3.27. Quasiconvexity of $\operatorname{Fix}\varphi$ can alternatively be characterised saying that $\operatorname{Fix}\varphi$ acts properly and cocompactly on a convex subcomplex of the universal cover of the Salvetti/Davis complex, or, again, in coarse median terms; see Definition 2.30, Remark 2.31 and Lemma 3.2.

In light of Theorem C, it is only natural to wonder what isomorphism types of special groups can arise as $\operatorname{Fix}\varphi$, and whether their complexity can be bounded in any way in terms of the ambient group, in the spirit of Scott's conjecture. We only provide a very partial result on these questions (Corollary E), leaving a more detailed treatment for later work. The main proof ingredient, which we believe is of independent interest, is the following construction of $U_0(\mathcal{A}_\Gamma)$–invariant Bass–Serre trees for most right-angled Artin groups.

Proposition D *Let $\mathcal{A}_\Gamma$ be directly irreducible, freely irreducible and noncyclic. Then there exists an amalgamated product splitting $\mathcal{A}_\Gamma = \mathcal{A}_+ *_{\mathcal{A}_0} \mathcal{A}_-$, with $\mathcal{A}_\pm$ and $\mathcal{A}_0$ parabolic subgroups of $\mathcal{A}_\Gamma$, such that the corresponding Bass–Serre tree $\mathcal{A}_\Gamma \curvearrowright T$ is $U_0(\mathcal{A}_\Gamma)$–invariant. That is: for every $\varphi \in U_0(\mathcal{A}_\Gamma)$, there exists an isometry $f\colon T \to T$ satisfying $f \circ g = \varphi(g) \circ f$ for all $g \in \mathcal{A}_\Gamma$.*

Corollary E *Consider a right-angled Artin group $\mathcal{A}_\Gamma$ and $\varphi \in U_0(\mathcal{A}_\Gamma)$.*

1. *If $\mathcal{A}_\Gamma$ splits as a direct product $\mathcal{A}_1 \times \mathcal{A}_2$, then $\varphi(\mathcal{A}_i) = \mathcal{A}_i$ and $\operatorname{Fix}\varphi = \operatorname{Fix}\varphi|_{\mathcal{A}_1} \times \operatorname{Fix}\varphi|_{\mathcal{A}_2}$.*
2. *If $\mathcal{A}_\Gamma$ is directly irreducible, then the subgroup $\operatorname{Fix}\varphi \leq \mathcal{A}_\Gamma$ splits as a (possibly trivial) finite graph of groups with vertex and edge groups of the form $\operatorname{Fix}\varphi|_P$, for proper parabolic subgroups $P \leq \mathcal{A}_\Gamma$ with $\varphi(P) = P$ and $\varphi|_P \in U_0(P)$.*

The same two results hold for right-angled Coxeter groups $\mathcal{W}_\Gamma$ and automorphisms $\varphi \in \operatorname{Aut}_0 \mathcal{W}_\Gamma$.

We now turn to Question 2, which is the second main focus of the paper. Recall that Paulin [88] showed that, for every Gromov-hyperbolic group G, every infinite-order element of $\operatorname{Out} G$ can be realised as a homothety of a small, isometric G–tree.

Our main result on Question 2, generalises Paulin's theorem to CMP automorphisms of *special groups* G, in the Haglund–Wise sense [68; 90]. This is a broad class of groups including right-angled Artin groups, finite-index subgroups of right-angled Coxeter groups, as well as free and surface groups and a number of other hyperbolic examples.

Note that *small* G–actions on $\mathbb{R}$–trees are not the right notion to consider in this context. Indeed, if a special group G has a small action on an $\mathbb{R}$–tree T, then every arc stabiliser is free abelian and the work of Rips and Bestvina–Feighn implies that G splits over an abelian subgroup; see Bestvina and Feighn [10, Theorem 9.5]. However, there exist special groups that admit an infinite-order CMP outer automorphism, but do not split over any abelian subgroup (eg the RAAG $\mathcal{A}_\Gamma$ with Γ as in Figure 1, by Groves and Hull [59]).

In fact, due to the lack of hyperbolicity, it is reasonable to expect that $\mathbb{R}$–trees will need to be replaced by higher-dimensional analogues.

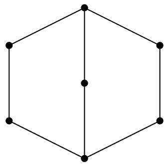

Figure 1

The correct setting seems to be provided by the simultaneous generalisation of $\mathbb{R}$–trees and CAT(0) cube complexes known as *median spaces*. These are those metric spaces (X,d) such that, for all $x_1,x_2,x_3\in X$, there exists a unique point $m(x_1,x_2,x_3)$ (known as their *median*) satisfying

$$d(x_i,x_j)=d(x_i,m(x_1,x_2,x_3))+d(m(x_1,x_2,x_3),x_j)\quad \text{for all } 1\leq i<j\leq 3.$$

A connected median space X is said to have *rank* $\leq r$ if all its locally compact subsets have topological dimension $\leq r$. Rank-1 connected median spaces are precisely $\mathbb{R}$–trees.

The following is our main result on Question 2 (a more general statement for infinite abelian subgroups of $\operatorname{Out} G$ is Theorem 7.25). Note that, although higher-rank median spaces are never nonpositively curved, they always admit a canonical, bi-Lipschitz equivalent CAT(0) metric;[3] see Bowditch [17].

Theorem F *Let G be the fundamental group of a compact special cube complex. Suppose G has trivial centre. Let $\phi\in\operatorname{Out} G$ be infinite-order and coarse-median preserving. Then*:

1. *There is a geodesic, finite-rank median space X and an action by homotheties $G\rtimes_\phi\mathbb{Z}\curvearrowright X$.*
2. *The restriction $G\curvearrowright X$ is isometric, minimal, with unbounded orbits, and "moderate".*
3. *If $\varphi\in\operatorname{Aut} G$ represents ϕ, then the subgroup $\operatorname{Fix}\varphi\leq G$ fixes a point of X.*
4. *If ϕ and ϕ^{-1} are subexponentially growing, then the action $G\rtimes_\phi\mathbb{Z}\curvearrowright X$ is isometric.*

As for actions on $\mathbb{R}$–trees, we say that $G\curvearrowright X$ is *minimal* if X does not contain any proper, G–invariant convex subsets. We propose the notion of "moderate" action on a median space as a higher-rank generalisation of the notion of small action on an $\mathbb{R}$–tree.

Definition (moderate actions) Let G be a group and X be a median space.

1. A *k–cube* in X is a median subalgebra $C\subseteq X$ isomorphic to the product $\{0,1\}^k$.
2. An isometric action $G\curvearrowright X$ is *moderate* if, for every $k\geq 1$ and every k–cube $C\subseteq X$, the subgroup of G fixing C pointwise contains a copy of $\mathbb{Z}^k$ in its centraliser.

Any 2–element subset of X is a 1–cube. Thus, if G is hyperbolic and $G\curvearrowright X$ is moderate, the intersection of any two point-stabilisers must be virtually cyclic. In particular, if G is torsionfree hyperbolic and T is an $\mathbb{R}$–tree, then the action $G\curvearrowright T$ is moderate if and only if it is small. We remark that, when G is hyperbolic, the space X provided by Theorem F is indeed an $\mathbb{R}$–tree.

[3]The reader should keep in mind the case of $\mathbb{R}^n$, where the ℓ^1 metric is median and the Euclidean metric is CAT(0).

We would like to emphasise that Theorem F does not provide any *lower* bounds to the rank of the median space X. In particular, we still do not have an answer to the following:

Question 3 (1) Can we always take the median space X in Theorem F to be an $\mathbb{R}$–tree?

(2) If G is a directly and freely irreducible RAAG, can we even take X to be a simplicial tree?

We have seen that, when $\mathcal{A}_\Gamma$ is directly and freely irreducible, Proposition D yields a $U_0(\mathcal{A}_\Gamma)$–invariant simplicial $\mathcal{A}_\Gamma$–tree. However, it remains unclear if such a simplicial tree can always be taken to be moderate and, more importantly, if it can be constructed so that $\mathrm{Fix}\,\varphi$ is elliptic.

We conclude this overview by highlighting two more results. These fall outside the main purpose of this text, but they are almost immediate consequences of the techniques used in this paper and we find them of independent interest. We prove them at the end of Section 4.2.

Recall that the property of being cocompactly cubulated does not, in general, pass to finite-index *overgroups*. Many examples of this are provided by crystallographic groups (see Hagen [62]): for instance, the $(3,3,3)$ triangle group has $\mathbb{Z}^2$ as a finite-index subgroup, but it is not itself cocompactly cubulated.

The following is a criterion for cubulating finite-index overgroups. Its proof is loosely inspired by the idea of Guirardel cores (see Guirardel [60] and Hagen and Wilton [65]), but it requires none of the technical machinery. Instead, it is a simple consequence of Proposition 4.1 (or the earlier result of Bowditch [18, Proposition 4.1]).

Corollary G *Let G be a group with a cocompactly cubulated finite-index subgroup H. Suppose that the coarse median structure on G induced by the cubulation of H is G–invariant (it is automatically H–invariant). Then G is cocompactly cubulated.*

Along with Proposition A, the previous corollary implies the following version of Nielsen realisation for automorphisms of right-angled Artin and Coxeter groups.

Corollary H (Nielsen realisation for RA$*$Gs) *Consider one of the following two settings:*

(1) *A centreless right-angled Artin group $G=\mathcal{A}_\Gamma$ and a finite subgroup $F\le \mathrm{Out}\,\mathcal{A}_\Gamma$ contained in the projection to outer automorphisms of the untwisted subgroup $U(\mathcal{A}_\Gamma)\le \mathrm{Aut}\,\mathcal{A}_\Gamma$.*

(2) *A centreless right-angled Coxeter group $G=\mathcal{W}_\Gamma$ and any finite subgroup $F\le \mathrm{Out}\,\mathcal{W}_\Gamma$.*

In either case, F can be realised as a group of automorphisms of a compact, nonpositively curved, cube (orbi)complex Q with $G=\pi_1 Q$.

Part (2) is new, while part (1) is originally due to Hensel and Kielak [69]. When $F\le U_0(\mathcal{A}_\Gamma)$, they constructed Q quite explicitly via a glueing construction, ensuring that $\dim Q=\dim \mathcal{X}_\Gamma$. By comparison, our approach does not offer much control on dimension (except $\dim Q\le \#F\cdot \dim \mathcal{X}_\Gamma$), but it provides a much more elementary proof of the existence of some Q.

We expect our complex Q to be special, but this would require additional arguments in the proof (the only delicate point being lack of interosculations). We also think it should be possible to "trim" Q into having the optimal dimension $\dim \mathcal{X}_\Gamma$ by relying on the "panel collapse" procedure of Hagen and Touikan [64] (or small variations thereof), but the details seem too technical to be discussed here.

1.1 On the proof of Theorems B and C

The two theorems are proved in Section 4 under the aliases of Theorem 4.10 and Corollaries 4.34 and 4.35.

Regarding Theorem B, the starting observation is that $\mathrm{Fix}\,\varphi$ is an *approximate median subalgebra* of the group G; see Definition 2.33 and Lemma 2.35. Fixing a proper cocompact action on a CAT(0) cube complex $G \curvearrowright \mathcal{Z}$, the proof then takes place in three steps.

(1) If a subgroup $H \leq G$ is an approximate median subalgebra, H is finitely generated (Proposition 4.11). We prove this by relying on a straightforward adaptation of an argument due to Paulin [86] in the context of hyperbolic groups. Paulin's argument is itself a generalisation of Cooper's proof [33] in the case when the group G is free (a result originally due to Gersten [57] from the early 80s).

(2) Approximate median subalgebras of CAT(0) cube complexes are always at finite Hausdorff distance from actual median subalgebras (Proposition 4.1 or [18, Proposition 4.1]).

(3) Applying the previous step to H–orbits in $\mathcal{Z}$, we obtain an H–invariant median subalgebra $M \subseteq \mathcal{Z}^{(0)}$ such that $H \curvearrowright M$ is cofinite. Along with the fact that H is finitely generated, this yields a cocompact cubulation that quasi-isometrically embeds into $\mathcal{Z}$ (Lemma 4.12), though not necessarily as a convex subcomplex.

A similar strategy gives a new proof of W Neumann's result [82] that fixed subgroups of automorphisms of hyperbolic groups are quasiconvex; see also Minasyan and Osin [81]. Indeed, recall that, although not all hyperbolic groups are cocompactly cubulated, they are all coarse median, and all their automorphisms φ are CMP by Proposition A. It is easy to see that all coarsely connected, approximate median subalgebras of hyperbolic spaces are quasiconvex. As above, this implies that $\mathrm{Fix}\,\varphi$ is quasiconvex.

When dealing with nonhyperbolic groups, quasiconvexity is significantly harder to ensure and the proof of Theorem C requires additional work. Namely, assuming that $\varphi \in U_0(\mathcal{A}_\Gamma)$ or $\varphi \in \mathrm{Aut}_0\,\mathcal{W}_\Gamma$, we need to show that $(\mathrm{Fix}\,\varphi)$–orbits in the Salvetti complex $\mathcal{X}_\Gamma$ or Davis complex $\mathcal{Y}_\Gamma$ are quasiconvex (in the coarse median sense; see Definition 2.30, Remark 2.31 and Lemma 3.2).

The proof of this is based on a quasiconvexity criterion for median subalgebras of CAT(0) cube complexes (Proposition 4.25). The most important ingredients are the fact that $\mathcal{X}_\Gamma$ and $\mathcal{Y}_\Gamma$ do not contain "infinite staircases" (Section 4.3), and certain properties that distinguish elements of $U_0(\mathcal{A}_\Gamma)$ and $\mathrm{Aut}_0\,\mathcal{W}_\Gamma$ from more general CMP automorphisms in $U(\mathcal{A}_\Gamma)$ and $\mathrm{Aut}\,\mathcal{W}_\Gamma$ (Lemmas 4.30 and 4.32).

We conclude by mentioning that other important tools for the study of undistortion and quasiconvexity of subgroups of cubulated groups were recently developed by Beeker and Lazarovich in [2] and [3, Theorem 1.2(2)], and Dani and Levcovitz [38, Theorem A], based on extensions of the classical machinery of *Stallings folds* [95; 96] from graphs to higher-dimensional cube complexes. These techniques play no role in our arguments, but it is possible that they can be used to give alternative proofs of certain special cases of Theorems B and C.

1.2 On the proof of Theorem F

Keeping the case of $\mathrm{Out}\, F_n$ in mind, as described eg in Gaboriau, Jaeger, Levitt and Lustig [53, Section 2], there are two main obstacles to overcome:

(a) No good analogue of (*relative*) *train track maps* is available to represent homotopy equivalences between nonpositively curved cube complexes.

(b) It is not known if (isometric) actions on finite-rank median spaces are completely determined by their length function. There are results of this type for actions on $\mathbb{R}$–trees (see Culler and Morgan [37]) and cube complexes (see Beyrer and Fioravanti [13; 14]), but their extension to a general median setting would require some significantly new ideas.

The proof of Theorem F is made up of two main steps, which we now describe. In this sketch, we restrict our attention to the construction of the homothetic action $G \rtimes_\phi \mathbb{Z} \curvearrowright X$ (parts (1) and (2) of the theorem). Parts (3) and (4) follow, respectively, from parts (1) and (2) of Remark 7.27.

Let G be a special group, let $\mathcal{Z}$ be a CAT(0) cube complex, and let $\rho\colon G \to \mathrm{Aut}\,\mathcal{Z}$ be the homomorphism corresponding to a proper, cocompact, cospecial action $G \curvearrowright \mathcal{Z}$. Equip G with the coarse median structure arising from $\mathcal{Z}$. Let $\varphi \in \mathrm{Aut}\, G$ be a coarse-median preserving automorphism projecting to an infinite-order element of $\mathrm{Out}\, G$.

Step 1 *There exist a finite-rank median space X, an isometric action $G \curvearrowright X$ with unbounded orbits, and a homeomorphism $H\colon X \to X$ satisfying $H \circ g = \varphi(g) \circ H$ for all $g \in G$.*

In order to prove this, we consider the sequence of homomorphisms $\rho_n := \rho \circ \varphi^n$ and the sequence of G–actions on cube complexes $G \curvearrowright \mathcal{Z}_n$ that they induce. We then fix a nonprincipal ultrafilter ω, choose basepoints $p_n \in \mathcal{Z}_n$ and scaling factors $\epsilon_n > 0$, and consider the ultralimit

$$(X, p) := \lim_\omega (\epsilon_n \mathcal{Z}_n, p_n).$$

This is easily seen to be a finite-rank median space and, for a suitable choice of p_n and λ_n, the actions $G \curvearrowright \mathcal{Z}_n$ converge to an isometric action $G \curvearrowright X$ with unbounded orbits.

So far this is just a classical Bestvina–Paulin construction; see Bestvina [7] and Paulin [85]. The actual subtleties lie in the definition of the map $H\colon X \to X$. By the Milnor–Schwarz lemma, there exists a quasi-isometry $h\colon \mathcal{Z} \to \mathcal{Z}$ satisfying $h \circ g = \varphi(g) \circ h$ for all $g \in G$. We would like to define H as the ultralimit of the corresponding sequence of quasi-isometries $\mathcal{Z}_n \to \mathcal{Z}_n$, but this might displace the basepoint $p \in X$ by an infinite amount.

In order to rule out this eventuality, we rely on an argument similar to the one used in Paulin [88] for hyperbolic groups. On closer inspection, Paulin's argument only requires the following property, which is satisfied by nonelementary hyperbolic groups.

Definition Let G be a infinite group with a (fixed) Cayley graph $(\mathcal{G}, d)$. We say that G is *uniformly nonelementary* (*UNE*) if there exists a constant $c > 0$ with the following property. For every finite generating set $S \subseteq G$ and for all $x, y \in \mathcal{G}$, we have

$$d(x, y) \le c \cdot \max_{s \in S}[d(x, sx) + d(y, sy)].$$

The important part of this definition is that the constant c *does not* depend on the generating set S. Note that the UNE property is independent of the specific choice of $\mathcal{G}$; cf Definition 2.36.

Our main contribution to Step 1 is the proof of the following fact (Corollary 7.23), which is potentially of independent interest.

Theorem I *Let G be the fundamental group of a compact special cube complex. If G has trivial centre, then G is uniformly nonelementary.*

Now, let $m\colon X^3 \to X$ denote the median operator of the median space X. The fact that $\varphi \in \operatorname{Aut} G$ is coarse-median preserving easily implies that the homeomorphism $H\colon X \to X$ arising from the above construction satisfies $H(m(x, y, z)) = m(H(x), H(y), H(z))$ for all $x, y, z \in X$. However, H need not be a homothety at this stage.

Step 2 *There exists a G–invariant (pseudo)metric $\eta\colon X \times X \to [0, +\infty)$ such that (X, η) is a median space with the same median operator m, and H is a homothety with respect to η.*

Since $H : X \to X$ preserves the median operator m, there is an action of H on the space of all G–invariant median pseudometrics on X that induce m. More precisely, we show that H gives a homeomorphism of a certain space of (projectivised) median pseudometrics on X, and that the latter is a compact *absolute retract* (*AR*). The existence of the required pseudometric η then follows from the Lefschetz fixed point theorem for homeomorphisms of compact ANRs. This is discussed mainly in Sections 6.2 and 7.4; see especially Corollaries 6.23 and 7.24.

Once the pseudometric η is obtained, we can pass to the quotient metric space to obtain a genuine median space.

1.3 Further questions

We would like to highlight four questions raised by our results.

As mentioned earlier, every hyperbolic group admits a unique *coarse median structure* (Definition 2.22). At the opposite end of the spectrum, any RAAG for which $U(\mathcal{A}_\Gamma)$ has infinite index in $\mathrm{Aut}\,\mathcal{A}_\Gamma$ will admit infinitely many $\mathcal{A}_\Gamma$–invariant coarse median structures.

Right-angled Coxeter groups $\mathcal{W}_\Gamma$ seem to place themselves in between these two extremal situations: they can admit infinitely many distinct coarse median structures — eg because every RAAG is a finite-index subgroup of a RACG; see Davis and Januszkiewicz [40] — but it is not clear which of these structures are $\mathcal{W}_\Gamma$–invariant. For instance, Proposition A(2) implies that all Coxeter generating sets of $\mathcal{W}_\Gamma$ give rise to the same coarse median structure (which fails for Artin generating sets of $\mathcal{A}_\Gamma$).

Question 4 Does each RACG $\mathcal{W}_\Gamma$ have only *finitely many* $\mathcal{W}_\Gamma$–invariant coarse median structures?

As an example of why one might expect this kind of rigidity, we suggest looking at the difference between the RAAG $\mathbb{Z}^n$ and the RACG $(D_\infty)^n$, where D_∞ is the infinite dihedral group. The space of $\mathbb{Z}^n$–invariant coarse median structures on $\mathbb{Z}^n$ (equivalently, on $\mathbb{R}^n$) is uncountable, simply because it is endowed with a natural $GL_n\mathbb{R}$–action and we can consider the orbit of the standard structure. However, of the structures in this orbit, only finitely many are $(D_\infty)^n$–invariant.

The second question naturally arises from Theorem C and was already mentioned above:

Question 5 Consider $\varphi \in U_0(\mathcal{A}_\Gamma)$ or $\varphi \in \mathrm{Aut}_0\,\mathcal{W}_\Gamma$.

(1) What isomorphism types of special groups can arise as $\mathrm{Fix}\,\varphi$ for some choice of φ and Γ? When $\varphi \in U_0(\mathcal{A}_\Gamma)$, is $\mathrm{Fix}\,\varphi$ itself a right-angled Artin group?

(2) Can we bound the "complexity" of $\mathrm{Fix}\,\varphi$ in terms of $\#\Gamma^{(0)}$, in the spirit of Scott's conjecture?

Regarding part (1) of Question 5, note that every RAAG can arise as the fixed subgroup of some element of $U_0(\mathcal{A}_\Gamma)$, simply because we can always take $\varphi = \mathrm{id}$. One can easily construct more elaborate examples using this observation as a starting point.

One can also wonder about fixed subgroups of automorphisms of general coarse median groups G. By Lemma 2.35, this reduces to understanding subgroups that are *approximate median subalgebras* (Definition 2.33). We study these subgroups when G is cocompactly cubulated (Theorem 4.10), but some of our arguments should work more generally (especially the proof of Proposition 4.11).

Question 6 Let (G,μ) be a finite-rank coarse median group. Let a subgroup $H \leq G$ be an approximate median subalgebra.

(1) Is H finitely generated?

(2) Is H undistorted? Which properties of G does H retain?

For instance, when G is hierarchically hyperbolic, I do not know if H must be finitely generated. However, assuming that it is, the second part of the question has a positive answer: H is undistorted and hierarchically hyperbolic. This is evident from Bowditch's axioms (B1)–(B10) for (weak) hierarchically hyperbolic spaces [18, Section 7] and the coarse median characterisation of hierarchy paths [18, Theorem 1.1].

We emphasise that our definition of *coarse median group* (Definition 2.24) is slightly stronger than Bowditch's original definition [15], in that we require μ to be coarsely G–equivariant.

Our last question regards UNE groups. It is clear that UNE groups have finite centre, and it is not hard to show that nonelementary hyperbolic groups are UNE. All other examples of UNE groups that we are aware of are provided by Theorem I.

Are there other interesting examples or nonexamples of UNE groups? Given the proof of Theorem I, a positive answer to the following seems likely:

Question 7 Are hierarchically hyperbolic groups with finite centre UNE?

Outline of the paper

Section 2 mostly contains background material on median algebras, cube complexes and coarse median groups. An exception is Section 2.4, which reviews some of the results of [51]. The latter will be helpful, mostly in Sections 6 and 7, for some of the more technical arguments in the proof of Theorem F.

In Section 3, we consider cocompactly cubulated groups G and study a notion of *convex-cocompactness* for subgroups of G, which is a special instance of quasiconvexity in coarse median spaces (Definition 2.30). Section 3.2 studies cyclic, convex-cocompact subgroups of RAAGs (whose generators we call *label-irreducible*). Section 3.4 contains the proof of Proposition A.

Section 4 is concerned with fixed subgroups of CMP automorphisms. First, Sections 4.1 and 4.2 are devoted to the proof of Theorem B. Then Section 4.3 studies staircases in cube complexes, allowing us to formulate a quasiconvexity criterion for median subalgebras in Section 4.4. Finally, Section 4.5 restricts to Salvetti and Davis complexes, proving Theorem C.

Section 5 is completely independent from the subsequent part of the paper and can be safely skipped. It only contains the proof of Proposition D and Corollary E.

Finally, Sections 6 and 7 are the most technical parts of the paper and they contain the bulk of the proof of Theorem F. In Section 6, we consider group actions on finite-rank median algebras and develop a criterion for the existence of a (projectively) invariant metric (as required for Step 2 of the proof sketch for Theorem F). In Section 7, we study ultralimits of actions on Salvetti complexes, in order to obtain the properties needed to apply the results of Section 6. Theorems F and I are proved in Section 7.4.

Acknowledgements I would like to thank Mark Hagen for noticing that a result similar to Proposition 4.1 already appeared in [18], Camille Horbez for mentioning that ARs have good fixed-point properties, Robert Kropholler for pointing me to [98] and [62], Alessandro Sisto for discussions related to Definition 2.36, and Ric Wade for suggesting a major simplification to my original argument for Proposition A in the case of RAAGs. I thank Jason Behrstock, Corey Bregman, Ruth Charney, Ashot Minasyan and Karen Vogtmann for other interesting conversations, and the referee for many useful comments.

I am grateful to Ursula Hamenstädt and the Max Planck Institute for Mathematics in Bonn for their hospitality and financial support while part of this work was being completed.

2 Preliminaries

2.1 Frequent notation and identities

Throughout the paper, all groups will be equipped with the discrete topology. Thus, we will refer to *properly discontinuous* actions on topological spaces simply as *proper* actions.

If G is a group and $F \subseteq G$ is a subset, we denote by $\langle F\rangle$ the subgroup of G generated by F. We denote by $Z_G(F)$ the *centraliser* of the subset F, ie the subgroup of elements of G commuting with all elements of F.

If (X,d) is a metric space, $A \subseteq X$ is a subset, and $R \geq 0$ is a real number, we denote by $\mathcal{N}_R(A)$ the closed R–neighbourhood of A. If $x,y \in X$, we write $x \approx_R y$ with the meaning of $d(x,y) \leq R$.

Consider a group action on a set $G \curvearrowright X$. If η is a G–invariant pseudometric on X, we write, for every $x \in X$, $g \in G$, and $F \subseteq G$,

$$\ell(g,\eta) = \inf_{x\in X} \eta(x,gx), \quad \tau_F^\eta(x) = \max_{f\in F} \eta(x,fx), \quad \overline{\tau}_F^\eta = \inf_{x\in X} \tau_F^\eta(x).$$

When X is a metric space and we do not name its metric explicitly, we also write $\ell(g,X)$, τ_F^X and $\overline{\tau}_F^X$. If X is equipped with several G–actions originating from homomorphisms $\rho_n\colon G \to \operatorname{Isom} X$, we will write $\ell(g,\rho_n)$, $\tau_F^{\rho_n}$, $\overline{\tau}_F^{\rho_n}$ in order to avoid confusion.

If $S \subseteq G$ is a finite generating set, we denote by $|\cdot|_S$ and $\|\cdot\|_S$ the associated *word length* and *conjugacy length*, respectively:

$$|g|_S = \inf\{k \mid g = s_1 \cdot \dots \cdot s_k,\ s_i \in S^{\pm}\} \quad \text{and} \quad \|g\|_S = \inf_{h\in G} |hgh^{-1}|_S.$$

The following useful identities will be repeatedly used in this text. We consider a G–action on a set X, a G–invariant pseudometric η, a point $x \in X$, and finite generating sets $S, S_1, S_2 \subseteq G$. We have

$$\eta(x,gx) \leq |g|_S \cdot \tau_S^\eta(x), \quad \ell(g,\eta) \leq \|g\|_S \cdot \overline{\tau}_S^\eta, \quad \tau_{S_1}^\eta(x) \leq |S_1|_{S_2} \cdot \tau_{S_2}^\eta(x),$$

where we have defined $|S_1|_{S_2} := \max_{s\in S_1} |s|_{S_2}$.

2.2 Median algebras

In this and the next section, we only fix notation and prove a few simple facts that do not appear elsewhere in the literature. For a comprehensive introduction to median algebras and median spaces, the reader can consult [29, Sections (2)–(4)], [15, Sections (4)–(6)] and [50, Section 2].

A median algebra is a pair (M, m), where M is a set and $m\colon M^3 \to M$ is a map satisfying, for all $a, b, c, x \in M$,

$$m(a,a,b)=a, \quad m(a,b,c)=m(b,c,a)=m(b,a,c), \quad m(m(a,x,b),x,c)=m(a,x,m(b,x,c)).$$

The third identity, usually known as the 4–*point condition*, is sometimes replaced by a different identity involving 5 points (for instance, in [89; 29; 15; 50]). The equivalence of the two conditions [74; 1] is quite nontrivial, but not required in the rest of the paper.

A map $\phi\colon M \to N$ between median algebras is a *median morphism* if, for all $x, y, z \in M$, we have $\phi(m(x,y,z)) = m(\phi(x), \phi(y), \phi(z))$. We denote by $\operatorname{Aut} M$ the group of median automorphisms of M. Throughout the paper, all group actions on median algebras will be by (median) automorphisms, unless stated otherwise.

A subset $S \subseteq M$ is a *median subalgebra* if $m(S\times S\times S) \subseteq S$. A subset $C \subseteq M$ is *convex* if $m(C\times C\times M) \subseteq C$. Helly's lemma states that any finite family of pairwise-intersecting convex subsets of M has nonempty intersection [89, Theorem 2.2]. We say that C is *gate-convex* if it admits a *gate-projection*, ie a map $\pi_C\colon M \to C$ with the property that $m(z, \pi_C(z), x) = \pi_C(z)$ for all $x \in C$ and $z \in M$. Gate-convex subsets are convex, and convex subsets are median subalgebras. Each gate-convex subset admits a unique gate-projection, and gate-projections are median morphisms.

The *interval* $I(x, y)$ between points $x, y \in M$ is defined as the set $\{z \in M \mid m(x,y,z) = z\}$. Note that $I(x, y)$ is gate-convex with projection given by the map $z \mapsto m(x, y, z)$. Intervals can be used to give an alternative description of convexity: a subset $C \subseteq M$ is convex if and only if $I(x, y) \subseteq C$ for all $x, y \in C$.

A *halfspace* is a subset $\mathfrak{h} \subseteq M$ such that both $\mathfrak{h}$ and $\mathfrak{h}^* := M \setminus \mathfrak{h}$ are convex and nonempty. A *wall* is a set of the form $\mathfrak{w} = \{\mathfrak{h}, \mathfrak{h}^*\}$, where $\mathfrak{h}$ and $\mathfrak{h}^*$ are halfspaces. We say that $\mathfrak{w}$ is the wall *bounding* $\mathfrak{h}$, and that $\mathfrak{h}$ and $\mathfrak{h}^*$ are the halfspaces *associated* to $\mathfrak{w}$.

Two halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are *transverse* if all four intersections $\mathfrak{h}_1 \cap \mathfrak{h}_2$, $\mathfrak{h}_1^* \cap \mathfrak{h}_2$, $\mathfrak{h}_1 \cap \mathfrak{h}_2^*$ and $\mathfrak{h}_1^* \cap \mathfrak{h}_2^*$ are nonempty. If $\mathfrak{w}_1$ and $\mathfrak{w}_2$ are the walls bounding $\mathfrak{h}_1$ and $\mathfrak{h}_2$, we also say that $\mathfrak{w}_1$ is *transverse* to $\mathfrak{w}_2$ and $\mathfrak{h}_2$. If $\mathcal{U}$ and $\mathcal{V}$ are sets of walls or halfspaces, we say that $\mathcal{U}$ and $\mathcal{V}$ are *transverse* if every element of $\mathcal{U}$ is transverse to every element of $\mathcal{V}$. If $\mathcal{H}$ is a set of halfspaces, we write $\mathcal{H}^* := \{\mathfrak{h}^* \mid \mathfrak{h} \in \mathcal{H}\}$.

We denote by $\mathscr{W}(M)$ and $\mathscr{H}(M)$, respectively, the set of all walls and all halfspaces of M. Given subsets $A, B \subseteq M$, we write

$$\mathscr{H}(A|B) = \{\mathfrak{h} \in \mathscr{H}(M) \mid A \subseteq \mathfrak{h}^*, B \subseteq \mathfrak{h}\}, \quad \mathscr{W}(A|B) = \{\mathfrak{w} \in \mathscr{W}(M) \mid \mathfrak{w} \cap \mathscr{H}(A|B) \neq \varnothing\}.$$

If $\mathfrak{w}_1$ and $\mathfrak{w}_2$ are walls bounding disjoint halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$, we set

$$\mathscr{W}(\mathfrak{w}_1|\mathfrak{w}_2) := \mathscr{W}(\mathfrak{h}_1|\mathfrak{h}_2) \setminus \{\mathfrak{w}_1, \mathfrak{w}_2\}.$$

If $A, B \subseteq M$ are nonempty, then $\mathscr{H}(A|B)$ admits minimal elements under inclusion. This follows from Zorn's lemma since, for every totally ordered subset $\mathscr{C} \subseteq \mathscr{H}(A|B)$, the intersection of all halfspaces in $\mathscr{C}$ is again a halfspace in $\mathscr{H}(A|B)$. Note that any two minimal elements $\mathfrak{h}_1, \mathfrak{h}_2 \in \mathscr{H}(A|B)$ are transverse, since $\mathfrak{h}_1 \cap \mathfrak{h}_2$ and $\mathfrak{h}_1^* \cap \mathfrak{h}_2^*$ are nonempty and there is no inclusion relation between $\mathfrak{h}_1$ and $\mathfrak{h}_2$.

If $\mathfrak{w} \in \mathscr{W}(A|B)$, we say that the wall $\mathfrak{w}$ *separates* A and B. Any two disjoint convex subsets of M are separated by at least one wall [89, Theorem 2.8]; in particular, distinct points of M are always separated by a wall.

Given a subset $A \subseteq M$, we also introduce

$$\mathscr{H}_A(M) := \{\mathfrak{h} \in \mathscr{H}(M) \mid \mathfrak{h} \cap A \neq \varnothing, \mathfrak{h}^* \cap A \neq \varnothing\}, \quad \mathscr{W}_A(M) := \{\mathfrak{w} \in \mathscr{W}(M) \mid \mathfrak{w} \subseteq \mathscr{H}_A(M)\}.$$

Equivalently, a wall $\mathfrak{w}$ lies in $\mathscr{W}_A(M)$ if and only if it separates two points of A.

Remark 2.1 If $\mathcal{U} \subseteq \mathscr{H}(M)$ and $\mathcal{V} \subseteq \mathscr{H}(N)$ are subsets, we say that a map $\phi: \mathcal{U} \to \mathcal{V}$ is a *morphism of pocsets* if, for all $\mathfrak{h}, \mathfrak{k} \in \mathcal{U}$ with $\mathfrak{h} \subseteq \mathfrak{k}$, we have $\phi(\mathfrak{h}) \subseteq \phi(\mathfrak{k})$ and $\phi(\mathfrak{h}^*) = \phi(\mathfrak{h})^*$.

Every median morphism $\phi: M \to N$ induces a morphism of pocsets $\phi^*: \mathscr{H}_{\phi(M)}(N) \to \mathscr{H}(M)$ defined by $\phi^*(\mathfrak{h}) = \phi^{-1}(\mathfrak{h})$. When $\phi: M \to N$ is surjective, we obtain a map $\phi^*: \mathscr{H}(N) \to \mathscr{H}(M)$ that is injective and preserves transversality.

Remark 2.2 (1) If $S \subseteq M$ is a subalgebra, we have a map $\mathrm{res}_C: \mathscr{H}_C(M) \to \mathscr{H}(C)$ given by $\mathrm{res}_C(\mathfrak{h}) = \mathfrak{h} \cap C$. This is a morphism of pocsets and, by [15, Lemma 6.5], it is a surjection.

(2) If $C \subseteq M$ is convex, then the map res_C is also injective and it preserves transversality. In particular, the sets $\mathscr{H}(C)$ and $\mathscr{H}_C(M)$ are naturally identified in this case.
Indeed, if $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}_C(M)$ are intersecting halfspaces, Helly's lemma guarantees that $\mathfrak{h} \cap C$ and $\mathfrak{k} \cap C$ intersect too. Moreover, we have $\mathfrak{h} = \mathfrak{k}$ if and only if $\mathfrak{h} \cap \mathfrak{k}^*$ and $\mathfrak{h}^* \cap \mathfrak{k}$ are empty.

(3) If C is gate-convex with projection π_C, then $\mathrm{res}_C \circ \pi_C^* = \mathrm{id}_{\mathscr{H}(C)}$ and $\pi_C^* \circ \mathrm{res}_C = \mathrm{id}_{\mathscr{H}_C(M)}$.

If $C_1, C_2 \subseteq M$ are gate-convex subsets with gate-projections π_1, π_2, then $\mathscr{H}(x|C_i) = \mathscr{H}(x|\pi_i(x))$ for all $x \in M$. We say that $x_1 \in C_1$ and $x_2 \in C_2$ are a *pair of gates* if $\pi_2(x_1) = x_2$ and $\pi_1(x_2) = x_1$. Pairs of gates always exist and satisfy $\mathscr{H}(x_1|x_2) = \mathscr{H}(C_1|C_2)$.

The *standard k–cube* is the finite set $\{0,1\}^k$ equipped with the median operator m determined by a majority vote on each coordinate. A subset $S \subseteq M$ is a *k–cube* if it is a median subalgebra isomorphic to the standard k–cube. In particular, any subset of M with cardinality 2 is a 1–cube.

Remark 2.3 An important example of median algebra is provided by the 0–skeleton of any CAT(0) cube complex X; see [31]. The vertex set of any k–cell of X is a k–cube in the above sense, but the converse does not hold. For instance, in the standard tiling of $\mathbb{R}^n$, every set of the form $\{a_1, b_1\} \times \cdots \times \{a_n, b_n\}$ with $a_i < b_i$ is a k–cube according to the above notion. To avoid confusion, when dealing with cube complexes we will refer to k–cubes in $X^{(0)}$ as *generalised k–cubes*.

The *rank* of M, denoted by $\operatorname{rk} M$, is the largest cardinality of a set of pairwise-transverse walls of M. Equivalently, $\operatorname{rk} M$ is the supremum of the integers k such that M contains a k–cube (assuming $\operatorname{rk} M$ is at most countable); see [15, Proposition 6.2]. We will be exclusively interested in median algebras of finite rank.

We will need the following criterion, which summarises Lemmas 2.9 and 2.11 in [51]. If $\mathcal{H} \subseteq \mathscr{H}(M)$, we denote by $\bigcap \mathcal{H} \subseteq M$ the intersection of all halfspaces in $\mathcal{H}$.

Lemma 2.4 *Let M be a finite-rank median algebra. Partially order $\mathscr{H}(M)$ by inclusion.*

(1) *Let $\mathcal{H} \subseteq \mathscr{H}(M)$ be a set of pairwise intersecting halfspaces. Suppose that every chain in $\mathcal{H}$ admits a lower bound in $\mathcal{H}$. Then $\bigcap \mathcal{H}$ is a nonempty convex subset of M.*

(2) *A convex subset $C \subseteq M$ is gate-convex if and only if there does not exist a chain $\mathscr{C} \subseteq \mathscr{H}_C(M)$ such that $\bigcap \mathscr{C}$ is nonempty and disjoint from C.*

If $A \subseteq M$ is a subset, we denote by $\langle A \rangle$ the median subalgebra generated by A, ie the smallest subalgebra of M containing A. We also denote by $\operatorname{Hull} A$ the smallest convex subset of M that contains A; this coincides with the intersection of all halfspaces of M that contain A.

The sets $\langle A \rangle$ and $\operatorname{Hull} A$ are best understood in terms of the following operators:

$$\mathcal{M}(A) = \mathcal{M}^1(A) := m(A \times A \times A), \qquad \mathcal{M}^{n+1}(A) := \mathcal{M}(\mathcal{M}^n(A)),$$
$$\mathcal{J}(A) = \mathcal{J}^1(A) := m(A \times A \times M) = \bigcup_{x,y \in A} I(x,y), \qquad \mathcal{J}^{n+1}(A) := \mathcal{J}(\mathcal{J}^n(A)).$$

It is clear that $\operatorname{Hull} A = \bigcup_{n \geq 1} \mathcal{J}^n(A)$ and $\langle A \rangle = \bigcup_{n \geq 1} \mathcal{M}^n(A)$.

Remark 2.5 When $\operatorname{rk} M = r$ is finite, [15, Lemma 6.4] shows that already $\mathcal{J}^r(A) = \operatorname{Hull} A$. A similar result holds for $\langle A \rangle$ and the operator $\mathcal{M}$ (see Proposition 4.2 below), but its proof will require considerable work.

If M_1 and M_2 are median algebras, we denote by $M_1 \times M_2$ their *product*. This is the median algebra with underlying set $M_1 \times M_2$ and the only median operator for which both coordinate projections are median morphisms.

The set $\mathscr{W}(M_1 \times M_2)$ is naturally partitioned into two transverse subsets $\mathscr{W}_1$ and $\mathscr{W}_2$. A wall lies in $\mathscr{W}_1$ if and only if it separates two points in one (equivalently, every) fibre $M_1 \times \{*\}$; halfspaces associated to walls in $\mathscr{W}_1$ are unions of fibres $\{*\} \times M_2$. The set $\mathscr{W}_2$ is defined similarly, swapping the roles played by the two indices. Since all fibres are gate-convex in $M_1 \times M_2$, Remark 2.2 gives natural identifications between $\mathscr{W}_i$ and $\mathscr{W}(M_i)$.

In finite rank, product splittings can be completely characterised in terms of walls. The following is [51, Lemma 2.12]; also see [23, Lemma 2.5] in the special case of cube complexes.

Lemma 2.6 *For a finite-rank median algebra M, the following are equivalent:*

(1) *M splits as a product of median algebras $M_1 \times M_2$, where neither M_i is a singleton.*

(2) *There exists a partition $\mathscr{W}(M) = \mathscr{W}_1 \sqcup \mathscr{W}_2$, where the $\mathscr{W}_i$ are nonempty and transverse.*

When this happens, the set $\mathscr{W}_i$ is identified with $\mathscr{W}(M_i)$ as described above.

2.3 Compatible metrics on median algebras

A metric space (X, d) is a *median space* if, for all $x_1, x_2, x_3 \in X$, there exists a unique point $m(x_1, x_2, x_3)$ in X such that

$$d(x_i, x_j) = d(x_i, m(x_1, x_2, x_3)) + d(m(x_1, x_2, x_3), x_j)$$

for all $1 \leq i < j \leq 3$. In this case, the map $m\colon X^3 \to X$ gives a median algebra (X, m).

Remark 2.7 (rank of median spaces) We define the *rank* of X as the rank of the underlying median algebra (X, m). If X is a connected median space, then this notion of rank coincides with the supremum of the topological dimensions of the locally compact subsets of X. The latter is the definition of rank that we used in the introduction. One inequality follows from Theorem 2.2 and Lemma 7.6 in [15], while the other from [17, Proposition 5.6].

For the purposes of this paper, it is convenient to think of median spaces in terms of the following notion. Let M be a median algebra.

Definition 2.8 A pseudometric $\eta\colon M \times M \to [0, +\infty)$ is *compatible* if, for every $x, y, z \in M$,

$$\eta(x, y) = \eta(x, m(x, y, z)) + \eta(m(x, y, z) + y).$$

Thus, we can equivalently define median spaces as pairs (M, d), where M is a median algebra and d is a compatible metric on M.

We write $\mathcal{D}(M)$ and $\mathcal{PD}(M)$, respectively, for the sets of all compatible metrics and all compatible pseudometrics on M. In the presence of a group action $G \curvearrowright M$, we write $\mathcal{D}^G(M)$ and $\mathcal{PD}^G(M)$ for the subsets of G–invariant (pseudo)metrics (or just $\mathcal{D}^g(M)$ and $\mathcal{PD}^g(M)$ if $G = \langle g \rangle$).

To avoid confusion, we will normally denote compatible *metrics* by the letter δ, and general compatible *pseudometrics* by the letter η.

Consider a gate-convex subset $C \subseteq M$ and its gate-projection $\pi_C \colon M \to C$. For every pseudometric $\eta \in \mathcal{PD}(M)$, the maps $\pi_C \colon M \to C$ and $m \colon M^3 \to M$ are 1–Lipschitz, in the sense that

$$\eta(\pi_C(x), \pi_C(y)) \le \eta(x, y), \quad \eta(m(x, y, z), m(x', y', z')) \le \eta(x, x') + \eta(y, y') + \eta(z, z').$$

This can be proved as in Lemma 2.13 and Corollary 2.15 of [29]. In addition, gate-projections are nearest-point projections, in the sense that $\eta(x, \pi_C(x)) = \eta(x, C)$ for all $x \in M$.

If $\delta \in \mathcal{D}(M)$ and (M, δ) is complete, then a subset $C \subseteq M$ is gate-convex if and only if it is convex and closed in the topology induced by δ; see [29, Lemma 2.13].

If M is the 0–skeleton of a CAT(0) cube complex X, then a natural compatible metric on M is given by the restriction of the *combinatorial metric* on X: this is just the intrinsic path metric of the 1–skeleton of X. All cube complexes in this paper will be implicitly endowed with their combinatorial metric, rather than the CAT(0) metric. All geodesics will be assumed to be *combinatorial geodesics*.

Remark 2.9 A *halfspace-interval* is a set of the form $\mathscr{H}(x|y) \subseteq \mathscr{H}(M)$ for $x, y \in M$. Let $\mathscr{B}(M) \subseteq 2^{\mathscr{H}(M)}$ denote the σ–algebra generated by halfspace-intervals. We say that a subset $\mathcal{H} \subseteq \mathscr{H}(M)$ is $\mathscr{B}$–measurable if it lies in $\mathscr{B}(M)$.

Every $\eta \in \mathcal{PD}(M)$ induces a measure ν_η on $\mathscr{B}(M)$ such that $\nu_\eta(\mathscr{H}(x|y)) = \eta(x, y)$ for all $x, y \in M$; see eg [29, Theorem 5.1]. If $\eta \in \mathcal{PD}^G(M)$, then ν_η is G–invariant.

Lemma 2.10 *Let (X, d) be a median space. Let $A \subseteq X$ be a subset such that $\mathcal{J}(A) \subseteq \mathcal{N}_R(A)$ for some $R \ge 0$. Then, for every $D \ge 0$, we have*

$$\mathcal{J}(\mathcal{N}_D(A)) \subseteq \mathcal{N}_{2D+R}(A).$$

In addition, if $\operatorname{rk} X = r$, *we have* $\operatorname{Hull} A \subseteq \mathcal{N}_{2^r R}(A)$.

Proof If $z \in \mathcal{J}(\mathcal{N}_D(A))$, there exist $x, y \in \mathcal{N}_D(A)$ and $z \in I(x, y)$. Consider points $x', y' \in A$ with $d(x, x'), d(y, y') \le D$. Set $z' = m(x', y', z)$. Since $z' \in \mathcal{J}(A)$, we have $d(z', A) \le R$. Furthermore,

$$d(z, z') = d(m(x, y, z), m(x', y', z)) \le d(x, x') + d(y, y') \le 2D.$$

In conclusion, $d(z, A) \le d(z, z') + d(z', A) \le 2D + R$, as required.

Proceeding by induction, it is straightforward to obtain $\mathcal{J}^i(A) \subseteq \mathcal{N}_{(2^i-1)R}(A)$ for every $i \ge 0$. If $\operatorname{rk} X = r$, we have $\operatorname{Hull} A = \mathcal{J}^r(A)$ by Remark 2.5, hence $\operatorname{Hull} A \subseteq \mathcal{N}_{(2^r-1)R}(A) \subseteq \mathcal{N}_{2^r R}(A)$. □

2.4 Convex cores in median algebras

In this subsection, we collect a few facts proved in [51] extending the notion of "essential core" [23, Section 3] from actions on cube complexes to general actions on finite-rank median algebras (even with no invariant metric or topology). These results will only play a role in the proofs of Theorems F and I (especially in Sections 6 and 7). The reader only interested in the other results mentioned in the introduction can safely read this subsection with CAT(0) cube complexes in mind, just to familiarise themselves with our notation.

Let M be a median algebra of finite rank r.

Definition 2.11 We say that $g \in \operatorname{Aut} M$ acts

(1′) *nontransversely* if there does not exist a wall $\mathfrak{w} \in \mathscr{W}(X)$ such that $\mathfrak{w}$ and $g\mathfrak{w}$ are transverse;

(2′) *stably without inversions* if there do not exist $n \in \mathbb{Z}$ and $\mathfrak{h} \in \mathscr{H}(X)$ with $g^n\mathfrak{h} = \mathfrak{h}^*$.

An action $G \curvearrowright M$ by automorphisms is

(1) *nontransverse* if every $g \in G$ acts nontransversely;

(2) *without wall inversions* if every $g \in G$ acts stably without inversions;

(3) *essential* if, for every $\mathfrak{h} \in \mathscr{H}(M)$, there exists $g \in G$ with $g\mathfrak{h} \subsetneq \mathfrak{h}$.

Remark 2.12 If there exists $\delta \in \mathcal{D}^G(M)$ such that (M,δ) is connected, then $G \curvearrowright M$ is without wall inversions. This follows from [50, Proposition B] when (M,δ) is complete, and from [51, Remark 4.3] in general.

Keeping the notation of [51], each action $G \curvearrowright M$ determines sets of halfspaces

$$\begin{aligned}
\mathcal{H}_1(G) &:= \{\mathfrak{h} \in \mathscr{H}(M) \mid \exists g \in G \text{ such that } g\mathfrak{h} \subsetneq \mathfrak{h}\},\\
\overline{\mathcal{H}}_{1/2}(G) &:= \{\mathfrak{h} \in \mathscr{H}(M) \setminus \mathcal{H}_1(G) \mid \exists g \in G \text{ such that } g\mathfrak{h}^* \cap \mathfrak{h}^* = \varnothing \text{ and } g\mathfrak{h} \neq \mathfrak{h}^*\},\\
\overline{\mathcal{H}}_0(G) &:= \{\mathfrak{h} \in \mathscr{H}(M) \mid \forall g \in G \text{ either } g\mathfrak{h} \in \{\mathfrak{h},\mathfrak{h}^*\} \text{ or } g\mathfrak{h} \text{ and } \mathfrak{h} \text{ are transverse}\}.
\end{aligned}$$

As observed in [51, Section 3.1], we have a G–invariant partition

$$\mathscr{H}(M) = \overline{\mathcal{H}}_0(G) \sqcup \mathcal{H}_1(G) \sqcup \overline{\mathcal{H}}_{1/2}(G) \sqcup \overline{\mathcal{H}}_{1/2}(G)^*.$$

We write $\mathcal{W}_1(G)$ and $\mathcal{W}_0(G)$ for the sets of walls bounding the halfspaces in $\mathcal{H}_1(G)$ and $\overline{\mathcal{H}}_0(G)$.

Definition 2.13 The *reduced core* $\overline{\mathcal{C}}(G)$ is the intersection of all halfspaces lying in $\overline{\mathcal{H}}_{1/2}(G)$.

We adopt the convention that $\overline{\mathcal{C}}(G) = M$ when $\overline{\mathcal{H}}_{1/2}(G)$ is empty. We will write $\overline{\mathcal{C}}(G,M)$ (and $\mathcal{H}_\bullet(G,M)$, $\mathcal{W}_\bullet(G,M)$) if it is necessary to specify the ambient median algebra. We just write $\overline{\mathcal{C}}(g)$ (and $\mathcal{H}_\bullet(g)$, $\mathcal{W}_\bullet(g)$) if $G = \langle g\rangle$.

Theorem 2.14 [51] *Let G be finitely generated and let $G \curvearrowright M$ be without wall inversions.*

(1) *The reduced core $\overline{\mathcal{C}}(G)$ is nonempty, G–invariant and convex.*

Suppose in addition that $\mathcal{D}^G(M) \neq \varnothing$.

(2) *There is a G–fixed point in M if and only if $\mathcal{H}_1(G) = \varnothing$.*

(3) *The sets $\mathcal{W}_1(G)$ and $\mathcal{W}_0(G)$ are transverse and $\mathscr{W}_{\overline{\mathcal{C}}(G)}(M) = \mathcal{W}_0(G) \sqcup \mathcal{W}_1(G)$.*

(4) *The resulting partition of $\mathscr{W}(\overline{\mathcal{C}}(G))$ gives a product splitting $\overline{\mathcal{C}}(G) = \overline{\mathcal{C}}_0(G) \times \overline{\mathcal{C}}_1(G)$. The normaliser of the image of G in $\operatorname{Aut} M$ leaves $\overline{\mathcal{C}}(G)$ invariant, preserving the two factors. The action $G \curvearrowright \overline{\mathcal{C}}_1(G)$ is essential, while $G \curvearrowright \overline{\mathcal{C}}_0(G)$ fixes a point.*

Proof We just refer the reader to the relevant statements in [51]. Part (1) follows from Theorem 3.17(2). The two implications in part (2) are obtained from Proposition 3.23(2) and Lemma 4.5(1), respectively. Part (3) is a consequence of Lemma 4.5 and Lemma 3.22(2). Finally, part (4) follows from Remark 3.16 and the previous parts. □

Remark 2.15 If G acts on a CAT(0) cube complex X and $M = X^{(0)}$, then the action $G \curvearrowright \overline{\mathcal{C}}_1(G)$ in Theorem 2.14(4) is easily identified as the *G–essential core* of Caprace and Sageev; cf [23, Section 3.3]. In particular, note that Theorem 2.14 strengthens [23, Proposition 3.5], showing that the G–essential core always embeds G–equivariantly as a convex subcomplex of X.

Theorem 2.16 *If $g \in \operatorname{Aut} M$ acts nontransversely and stably without inversions, then*

(1) *the reduced core $\overline{\mathcal{C}}(g)$ is gate-convex, and*

(2) *for every $x \in M$ and every $\eta \in \mathcal{PD}^g(M)$, we have $\eta(x, gx) = \ell(g, \eta) + 2\eta(x, \overline{\mathcal{C}}(g))$.*

Proof Part (1) is [51, Proposition 3.36] and part (2) is [51, Proposition 4.9(3)]. □

Note that $\overline{\mathcal{C}}(G)$ is not gate-convex in general, even when $G \curvearrowright M$ is an isometric action of a finitely generated free group on a complete $\mathbb{R}$–tree. See [51, Example 3.37].

Remark 2.17 Part (2) of Theorem 2.16 implies that, if $\delta \in \mathcal{D}^g(M)$ and (M, δ) is a geodesic space, then g is *semisimple*: either g fixes a point of M or g translates along a $\langle g \rangle$–invariant geodesic.

The next two remarks will only be needed in Section 7.

Remark 2.18 Let $g \in \operatorname{Aut} M$ act nontransversely and stably without inversions, with $\mathcal{D}^g(M) \neq \varnothing$.

(1) Each $\mathfrak{h} \in \mathcal{H}_1(g)$ satisfies $\bigcap_{n \in \mathbb{Z}} g^n \mathfrak{h} = \varnothing$; see [51, Lemma 4.5(1)].

(2) A halfspace $\mathfrak{h}$ lies in $\mathfrak{h} \in \overline{\mathcal{H}}_0(g)$ if and only if $g\mathfrak{h} = \mathfrak{h}$, and it lies in $\mathcal{H}_1(g)$ if and only if either $g\mathfrak{h} \subsetneq \mathfrak{h}$ or $g\mathfrak{h} \supsetneq \mathfrak{h}$. This follows from Remarks 3.33 and 3.34 in [51], after observing that $\mathcal{H}_1(g) \subseteq \mathscr{H}_{\overline{C}(g)}(M)$ (eg by part (1) of this remark).

(3) Let $N \subseteq M$ be a $\langle g\rangle$–invariant median subalgebra. By Remark 2.2, intersecting the halfspaces of M with N, we obtain a surjective restriction map $\mathrm{res}_N : \mathscr{H}_N(M) \to \mathscr{H}(N)$. Parts (1) and (2) show that:

- If $\mathfrak{h} \in \overline{\mathcal{H}}_0(g, M) \cap \mathscr{H}_N(M)$, then $g \cdot \mathrm{res}_N(\mathfrak{h}) = \mathrm{res}_N(\mathfrak{h})$ and $\mathrm{res}_N(\mathfrak{h}) \in \overline{\mathcal{H}}_0(g, N)$.
- If $\mathfrak{h} \in \overline{\mathcal{H}}_{1/2}(g, M) \cap \mathscr{H}_N(M)$, then either $\mathrm{res}_N(\mathfrak{h}) \in \overline{\mathcal{H}}_{1/2}(g, N)$ or $g \cdot \mathrm{res}_N(\mathfrak{h}) = \mathrm{res}_N(\mathfrak{h})^*$.
- We have $\mathcal{H}_1(g, M) \subseteq \mathscr{H}_N(M)$ and $\mathrm{res}_N(\mathcal{H}_1(g, M)) = \mathcal{H}_1(g, N)$.

Remark 2.19 Let $g \in \mathrm{Aut}\, M$ act nontransversely and stably without inversions. Let ν_η be the measure introduced in Remark 2.9. Part (2) of Theorem 2.16 shows that $\ell(g, \eta) = \nu_\eta(\mathscr{H}(x|gx))$ for any $x \in \overline{C}(g)$. In view of parts (1) and (2) of Remark 2.18, the set $\mathscr{H}(x|gx) \sqcup \mathscr{H}(gx|x)$ is a $\mathscr{B}$–measurable fundamental domain for the action $\langle g\rangle \curvearrowright \mathcal{H}_1(g)$. It follows that, for *any* fundamental domain $\Omega \in \mathscr{B}(M)$ for the action $\langle g\rangle \curvearrowright \mathcal{H}_1(g)$, we have $\ell(g, \eta) = \frac{1}{2}\nu_\eta(\Omega)$.

2.5 Two constructions involving cube complexes

2.5.1 Restriction quotients Restriction quotients of CAT(0) cube complexes were originally introduced in [23, page 860]. Our interest is due to the fact that the Salvetti blowups and collapses from [25] are a particular instance of this construction, which can actually be phrased purely in median-algebra terms. This is mainly needed in the proof of Proposition A(3) in Section 3.4, though it will also be useful in Sections 3.1 and 7.3.

A map $f : X \to Y$ between cube complexes is said to be *cubical* if, on every cube $c \subseteq X$, it factors as a projection of c onto one of its faces, followed by an isomorphism onto a cube of Y.

Let X be a CAT(0) cube complex. The *carrier* of a hyperplane $\mathfrak{w} \in \mathscr{W}(X)$ is the smallest convex subcomplex of X that contains all edges crossing $\mathfrak{w}$. It naturally splits as a product $C \times [0, 1]$, where $C \times \{0\}$ and $C \times \{1\}$ are convex subcomplexes of X on the two sides of $\mathfrak{w}$.

Given a hyperplane $\mathfrak{w} \in \mathscr{W}(X)$, we can construct a new CAT(0) cube complex Y by *collapsing* $\mathfrak{w}$: we remove from X the interior of the carrier $C \times (0, 1)$ and we identify the isomorphic subcomplexes $C \times \{0\}$ and $C \times \{1\}$. The natural collapse map $X \to Y$ is a cubical map.

Now, consider a set of hyperplanes $\mathcal{U} \subseteq \mathscr{W}(X)$. The *restriction quotient* of X determined by $\mathcal{U}$ is the CAT(0) cube complex $X(\mathcal{U})$ obtained by collapsing all hyperplanes in $\mathscr{W}(X) \setminus \mathcal{U}$ (which usually involves infinitely many collapses). It has one vertex for every connected component of the complement in X of the union of the hyperplanes in $\mathcal{U}$, with two vertices joined by an edge exactly when the corresponding components are separated by a single element of $\mathcal{U}$. Let $\pi_{\mathcal{U}}\colon X \to X(\mathcal{U})$ be the natural collapse, which is again a cubical map.

If $G \curvearrowright X$ is an action and the subset $\mathcal{U} \subseteq \mathscr{W}(X)$ is G–invariant, then the restriction quotient $X(\mathcal{U})$ is also equipped with a natural G–action and the collapse map $\pi_{\mathcal{U}}$ is G–equivariant.

Proposition 2.20 *Consider* CAT(0) *cube complexes* X, Y *and a surjective cubical map* $\pi : X \to Y$. *Then the following are equivalent*:

(1) *There exists a subset* $\mathcal{U} \subseteq \mathscr{W}(X)$ *and an isomorphism* $Y \cong X(\mathcal{U})$ *with respect to which* π *corresponds to the natural collapse* $\pi_{\mathcal{U}}\colon X \to X(\mathcal{U})$.

(2) *For every vertex* $v \in Y$, *the preimage* $\pi^{-1}(v)$ *is a convex subcomplex of* X.

(3) *The restriction* $\pi : X^{(0)} \to Y^{(0)}$ *is a median morphism.*

If X *and* Y *are equipped with* G*–actions and* π *is* G*–equivariant*, *then the set* $\mathcal{U}$ *is* G*–invariant.*

Proof The equivalence of (1) and (2) was shown in [71, Theorem 4.4]. Fibres of median morphisms between median algebras are always convex, so (3) implies (2). Finally, (1) $\Longrightarrow$ (3) can be shown by observing that single hyperplane-collapses are median morphisms. □

2.5.2 Roller boundaries In two proofs (Proposition 4.11 and, briefly, Lemma 3.13), we will need the notion of Roller boundary of a CAT(0) cube complex X, denoted by ∂X. We list here the (well-known) properties that we will use.

The 0–skeleton of any CAT(0) cube complex X has a natural structure of median algebra; see for instance [31, Theorem 6.1] and [89, Theorem 10.3]. The ℓ^1 metric on X, denoted by d, is a compatible metric in the sense of Definition 2.8. Thus, the pair $(X^{(0)}, d)$ is a median space. The notions of "halfspace" and "wall" coincide with the usual notion of halfspace and hyperplane in CAT(0) cube complexes. Thus, we write $\mathscr{W}(X)$ and $\mathscr{H}(X)$ with the meaning of $\mathscr{W}(X^{(0)})$ and $\mathscr{H}(X^{(0)})$.

We can embed $X^{(0)} \hookrightarrow 2^{\mathscr{H}(X)}$ by mapping each vertex v to the subset $\sigma_v \subseteq \mathscr{H}(X)$ of halfspaces that contain it. This is a median morphism if we endow $2^{\mathscr{H}(X)}$ with the structure of median algebra given by

$$m(\sigma_1, \sigma_2, \sigma_3) = (\sigma_1 \cap \sigma_2) \cup (\sigma_2 \cap \sigma_3) \cup (\sigma_3 \cap \sigma_1).$$

The space $2^{\mathscr{H}(X)}$ is compact with the product topology, and we can consider the closure $\overline{X}$ of $X^{(0)}$ inside it. We define the *Roller boundary* ∂X as the set $\overline{X} \setminus X^{(0)}$.

For us, the only important facts will be:

(1) The subset $\overline{X} = X \sqcup \partial X \subseteq 2^{\mathscr{H}(X)}$ is a median subalgebra and $X^{(0)}$ is convex in $\overline{X}$.

(2) The median $m\colon \overline{X}^3 \to \overline{X}$ is continuous with respect to the topology that $\overline{X}$ inherits from $2^{\mathscr{H}(X)}$. With this topology, $\overline{X}$ is compact and totally disconnected. If X is locally finite, the subset $X^{(0)} \subseteq \overline{X}$ is discrete.

(3) If $\mathfrak{h} \in \mathscr{H}(X)$, its closure $\overline{\mathfrak{h}}$ inside $\overline{X}$ is gate-convex. In fact, $\overline{\mathfrak{h}}$ and $\overline{\mathfrak{h}^*}$ are complementary halfspaces of the median algebra $\overline{X}$. The gate-projection $\pi_{\mathfrak{h}}\colon \overline{X} \to \overline{\mathfrak{h}}$ takes $X^{(0)}$ to $\mathfrak{h}$.

(4) Two halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}(X)$ are said to be *strongly separated* if $\mathfrak{h} \cap \mathfrak{k} = \varnothing$ and no halfspace of X is transverse to both $\mathfrak{h}$ and $\mathfrak{k}$; see [4]. If $\mathfrak{h}$ and $\mathfrak{k}$ are strongly separated, then the gate-projection $\pi_{\mathfrak{h}}\colon \overline{X} \to \overline{\mathfrak{h}}$ maps $\overline{\mathfrak{k}}$ to a single point.

The reader can consult [47, Sections 2.3–2.4] and [50, Theorem 4.14] for more details on facts (1)–(3). Fact (4) follows, for example, from Corollary 2.22 and Lemma 2.23 in [49].

2.6 Coarse median structures

Coarse median spaces were introduced by Bowditch in [15]. We present the following equivalent definition from [83].

Definition 2.21 Let X be a metric space. A *coarse median* on X is a map $\mu\colon X^3 \to X$ for which there exists a constant $C \geq 0$ such that, for all $a, b, c, x \in X$, we have

(1) $\mu(a,a,b) = a$ and $\mu(a,b,c) = \mu(b,c,a) = \mu(b,a,c)$,

(2) $\mu(\mu(a,x,b),x,c) \approx_C \mu(a,x,\mu(b,x,c))$,

(3) $d(\mu(a,b,c),\mu(x,b,c)) \leq C d(a,x) + C$.

Note that part (2) of the definition is an approximate version of the 4–point condition, from our definition of median algebras at the beginning of Section 2.2.

There is an appropriate notion of *rank* also for coarse median spaces. Since this notion will play no significant role in our paper (except when we briefly mention it at the end of Section 7.1), we simply refer the reader to [15; 83; 84] for more details.

The following notion of *coarse median structure* is different from the one in [84, Definition 2.8], but it is hard to imagine this being cause for confusion.

Definition 2.22 Two coarse medians $\mu_1, \mu_2\colon X^3 \to X$ are *at bounded distance* if there exists a constant $C \geq 0$ such that $\mu_1(x,y,z) \approx_C \mu_2(x,y,z)$ for all $x, y, z \in X$. A *coarse median structure* on X is an equivalence class $[\mu]$ of coarse medians pairwise at bounded distance. A *coarse median space* is a pair $(X, [\mu])$ where X is a metric space and $[\mu]$ is a coarse median structure on it.

Remark 2.23 Let $f\colon X \to Y$ be a quasi-isometry with a coarse inverse denoted by $f^{-1}\colon Y \to X$. If $\mu\colon X^3 \to X$ is a coarse median on X, then

$$(f_*\mu)(x,y,z) := f\big(\mu(f^{-1}(x), f^{-1}(y), f^{-1}(z))\big)$$

is a coarse median on Y. If $[\mu_1] = [\mu_2]$, then $[f_*\mu_1] = [f_*\mu_2]$.

If $\mathrm{QI}(X)$ is the group of quasi-isometries $X \to X$ up to bounded distance (as defined for example in [44, Definition 8.22]), the above defines a natural left action of $\mathrm{QI}(X)$ on the set of coarse median structures on X.

Definition 2.24 A *coarse median* group is a pair $(G, [\mu])$ where G is a finitely generated group equipped with a word metric and $[\mu]$ is a G–invariant coarse median structure on G.

The requirement that $[\mu]$ be G–invariant can be equivalently stated as follows: for each $g \in G$, there exists a constant $C(g) \geq 0$ such that $g\mu(g_1, g_2, g_3) \approx_{C(g)} \mu(gg_1, gg_2, gg_3)$ for all $g_1, g_2, g_3 \in G$.

Note that Definition 2.24 is stronger than Bowditch's original definition from [15], which did not ask for $[\mu]$ to be G–invariant. Definition 2.24 is better suited to our needs in this paper, but it is not QI–invariant or even commensurability-invariant (unlike Bowditch's).

These two definitions of coarse median group parallel the notions of HHS and HHG from [5; 6]. Namely, every hierarchically hyperbolic group is a coarse median group in the sense of Definition 2.24, while any group that admits a structure of hierarchically hyperbolic space is coarse median in the sense of Bowditch [19] (we will simply refer to these as "groups with a coarse median structure").

Remark 2.25 If G is finitely generated, any group automorphism $\varphi\colon G \to G$ is bi-Lipschitz with respect to any word metric on G. The resulting homomorphism $\mathrm{Aut}\, G \to \mathrm{QI}(G)$ defines an $(\mathrm{Aut}\, G)$–action on the set of coarse median structures on G that takes G–invariant structures to G–invariant structures. If $(G, [\mu])$ is a coarse median group, then every inner automorphism of G fixes $[\mu]$, and we obtain an action of $\mathrm{Out}\, G$ on the $(\mathrm{Aut}\, G)$–orbit of $[\mu]$.

Definition 2.26 Let $(G, [\mu])$ be a coarse median group. We say that $\phi \in \mathrm{Out}\, G$ (or $\varphi \in \mathrm{Aut}\, G$) is *coarse-median preserving* if it fixes $[\mu]$. We denote by $\mathrm{Out}(G, [\mu]) \leq \mathrm{Out}\, G$ and $\mathrm{Aut}(G, [\mu]) \leq \mathrm{Aut}\, G$ the subgroups of coarse-median preserving automorphisms.

Thus $\varphi \in \mathrm{Aut}\, G$ is coarse-median preserving exactly when, fixing a word metric on G, there exists a constant $C \geq 0$ such that, for all $g_i \in G$,

$$\varphi(\mu(g_1, g_2, g_3)) \approx_C \mu(\varphi(g_1), \varphi(g_2), \varphi(g_3)).$$

Remark 2.27 Let $G \curvearrowright X$ be a proper cocompact action on a CAT(0) cube complex. Any orbit map $o\colon G \to X$ is a quasi-isometry that can be used to pull back the median operator $m\colon X^3 \to X$ to a coarse median structure $[\mu_X] := o_*^{-1}[m]$ on G. It is straightforward to check that $[\mu_X]$ is independent of all choices involved (though the notation is slightly improper, as $[\mu_X]$ does depend on the specific G–action on X). We refer to $[\mu_X]$ as the *coarse median structure induced by $G \curvearrowright X$*.

Let us write gx for the action of $g \in G$ on $x \in X$ according to $G \curvearrowright X$. Then, every $\varphi \in \operatorname{Aut} G$ gives rise to a twisted G–action on X, which we denote by $G \curvearrowright X^{\varphi}$, and is defined as $g \cdot x = \varphi^{-1}(g)x$. Note that $\varphi_*[\mu_X] = [\mu_{X^\varphi}]$ and thus $\varphi \operatorname{Out}(G, [\mu_X])\varphi^{-1} = \operatorname{Out}(G, [\mu_{X^\varphi}])$.

Each of the structures $[\mu_{X^\varphi}]$ is G–invariant. In particular, $(G, [\mu_X])$ is a coarse median group.

Example 2.28 Every geodesic Gromov-hyperbolic space X is equipped with a natural coarse median structure $[\mu]$ represented by the operators μ that map each triple (x, y, z) to an approximate incentre for a geodesic triangle with vertices x, y, z; cf [15, Section 3]. In fact, by [83, Theorem 4.2], this is the only coarse median structure that X can be endowed with. It follows that $[\mu]$ is preserved by every quasi-isometry of X.

In particular, all automorphisms of Gromov-hyperbolic groups are coarse-median preserving. Alternatively, it is not hard to prove this last fact directly, relying on the Morse lemma and the observation that group automorphisms are quasi-isometries with respect to any word metric.

Example 2.29 Equipping $\mathbb{Z}^n$ with the median operator μ associated to its ℓ^1 metric, we obtain a coarse median group $(\mathbb{Z}^n, [\mu])$. An automorphism $\varphi \in \operatorname{Aut} \mathbb{Z}^n = \mathrm{GL}_n\mathbb{Z}$ is coarse-median preserving if and only if it lies in the signed permutation group $O(n, \mathbb{Z}) \leq \mathrm{GL}_n\mathbb{Z}$, ie if it can be realised as an automorphism of the standard tiling of $\mathbb{R}^n$ by unit cubes. This will follow from Proposition A(3) once we prove it in Section 3.4 (though it also is easily shown by hand).

We end this subsection with the definitions of *quasiconvex* subsets and *approximate median subalgebras*, which will play an important role in Sections 3 and 4.

Definition 2.30 Let $(X, [\mu])$ be a coarse median space. A subset $A \subseteq X$ is *quasiconvex* if there exists $R \geq 0$ such that $\mu(A \times A \times X) \subseteq \mathcal{N}_R(A)$.

This notion is clearly independent of the chosen representative μ of the structure $[\mu]$. Moreover, by Definition 2.21(3), if subsets A and B have finite Hausdorff distance, then A is quasiconvex if and only if B is.

By Example 2.28, Definition 2.30 extends the usual notion of quasiconvexity in hyperbolic spaces. The next remark shows that this is also the notion of quasiconvexity appearing in the statement of Theorem C. We will discuss in Section 3.1 other equivalent notions of quasiconvexity in (nonhyperbolic) cube complexes.

Remark 2.31 Let G be a right-angled Artin/Coxeter group. Let $G \curvearrowright X$ be the action on the universal cover of the Salvetti/Davis complex and let $[\mu_X]$ be the induced coarse median structure on G, as in Remark 2.27. Recall that, for a subset $A \subseteq X^{(0)}$, the set $\mathcal{J}(A) = \mu_X(A \times A \times X)$ is the union of all geodesics joining points of A.

Since the standard Cayley graph of G is precisely the 1–skeleton of X, a subgroup $H \leq G$ is quasiconvex as defined in the statement of Theorem C if and only if we have $\mathcal{J}(H\cdot x) \subseteq \mathcal{N}_R(H\cdot x)$ for some $x \in X$ and $R \geq 0$. This is clearly equivalent to quasiconvexity of H with respect to the coarse median structure $[\mu_X]$.

Remark 2.32 If X is a finite-rank median space, then a subset $A \subseteq X$ is quasiconvex if and only if $d_{\mathrm{Haus}}(A, \mathrm{Hull}\, A) < +\infty$. This follows from Lemma 2.10.

A similar, weaker notion is that of *approximate median subalgebra*.

Definition 2.33 Let $(X, [\mu])$ be a coarse median space. A subset $A \subseteq X$ is an *approximate median subalgebra* if there exists $R \geq 0$ such that $\mu(A \times A \times A) \subseteq \mathcal{N}_R(A)$.

Again, the definition only depends on the structure $[\mu]$ and passes on to all subsets of X at finite Hausdorff distance from A. An analogue of Remark 2.32 also holds, but it is more complicated and will be discussed in Section 4.1.

If φ is a coarse-median preserving automorphism of a coarse median group $(G, [\mu])$, the fixed subgroup $\mathrm{Fix}\,\varphi \leq G$ is in general not quasiconvex (for instance, consider the automorphism of $\mathbb{Z}^2$ that swaps the standard generators). However, it is always an approximate median subalgebra, as the next two lemmas show. This will be important in the proof of Theorem B.

Lemma 2.34 *Let G be a finitely generated group and let d be a word metric on G. For every $\varphi \in \mathrm{Aut}\, G$, there exist functions $\zeta_1, \zeta_2 \colon \mathbb{N} \to \mathbb{R}_{>0}$, with ζ_1 linear, such that, for every $g \in G$,*

$$\zeta_1\big(d(g, \varphi(g))\big) \leq d(g, \mathrm{Fix}\,\varphi) \leq \zeta_2\big(d(g, \varphi(g))\big).$$

Proof For the first inequality, note that $\varphi \colon G \to G$ is C–bi-Lipschitz with respect to d, for some constant $C \geq 0$. If $g' \in \mathrm{Fix}\,\varphi$ is an element closest to g, we have

$$d(g, \varphi(g)) \leq d(g, g') + d(\varphi(g'), \varphi(g)) \leq (1+C)\cdot d(g, g') = (1+C)\cdot d(g, \mathrm{Fix}\,\varphi).$$

Thus, we can take $\zeta_1(t) := t/(1+C)$.

Regarding the second inequality, suppose for the sake of contradiction that there does not exist a function ζ_2 so that it is satisfied. Then, there exist elements $g_n \in G$ with $d(g_n, \mathrm{Fix}\,\varphi) \to +\infty$, but $d(g_n, \varphi(g_n)) \leq D$ for some $D \geq 0$. Passing to a subsequence, we can assume that $\varphi(g_n) = g_n x$ for some $x \in G$ and all n. Thus $g_n g_m^{-1} \in \mathrm{Fix}\,\varphi$, hence $d(g_n, \mathrm{Fix}\,\varphi) = d(g_m, \mathrm{Fix}\,\varphi)$ for all $n, m \geq 0$, contradicting the fact that the distances $d(g_n, \mathrm{Fix}\,\varphi)$ diverge. □

Lemma 2.35 *Let $(G, [\mu])$ be a coarse median group. If $\varphi \in \mathrm{Aut}(G, [\mu])$, then $\mathrm{Fix}\,\varphi \leq G$ is an approximate median subalgebra.*

Proof Since $\varphi \in \operatorname{Aut}(G, [\mu])$, there is a constant C such that

$$\varphi(\mu(x, y, z)) \approx_C \mu(\varphi(x), \varphi(y), \varphi(z)) \quad \text{for all } x, y, z \in G.$$

Thus, if $x, y, z \in \operatorname{Fix}\varphi$, we have $\varphi(\mu(x, y, z)) \approx_C \mu(x, y, z)$. Lemma 2.34 gives a constant C' such that $d(\mu(x, y, z), \operatorname{Fix}\varphi) \le C'$ for all $x, y, z \in \operatorname{Fix}\varphi$, as required. □

2.7 UNE actions and groups

The following (seemingly novel) notion will play an important role in the proof of Theorem F, especially in Sections 6.2, 7.1 and 7.4.

Definition 2.36 Let G be a finitely generated group and let (X, d) be a (pseudo)metric space.

(1) An isometric action $G \curvearrowright X$ is *uniformly nonelementary* (*UNE*) if there exists a constant $c > 0$ with the following property. For every finite generating set $S \subseteq G$ and for all $x, y \in X$,

$$d(x, y) \le c \cdot [\tau_S^d(x) + \tau_S^d(y)].$$

We say that $G \curvearrowright X$ is *c–uniformly nonelementary* (*c–UNE*) when we need to specify c.

(2) An infinite group G is *UNE* if it admits a UNE, proper, cocompact action on a geodesic metric space.

The previous definition differs slightly from the one given in the introduction, but it is easily seen to be equivalent.

Remark 2.37 If G is infinite and an action $G \curvearrowright X$ is proper and cocompact, then there exists $\epsilon > 0$ such that, for every generating set $S \subseteq G$ and every $x \in X$, we have $\tau_S^d(x) \ge \epsilon$.

Along with the Milnor–Schwarz lemma, this can be used to show that a group is UNE if and only if *every* proper, cocompact action on a geodesic space is UNE. Equivalently, if the action of G on its locally finite Cayley graphs is UNE.

Example 2.38 (1) Nonelementary hyperbolic groups are UNE (for instance, this is implicitly shown in the last two paragraphs of the proof of [88, Lemme 3.1]).

(2) Fundamental groups of compact special cube complexes with finite centre are UNE. We will obtain this in Corollary 7.23.

(3) UNE groups have finite centre.

3 Cubical convex-cocompactness

This section is devoted to *convex-cocompact* subgroups of cocompactly cubulated groups (Definition 3.1). First, in Section 3.1, we discuss the relationship between convex-cocompactness and coarse median quasiconvexity. Then, Section 3.2 discusses basic properties of *cyclic*, convex-cocompact subgroups of RAAGs. Finally, Proposition A is proved in Section 3.4.

The reader who is not interested in the proofs of Theorems F and I can safely skip Section 3.3, which is devoted to some of the finer properties of convex-cocompact subgroups of RAAGs and is more technical. Its results will only be needed in Section 7.

3.1 Cubical convex-cocompactness in general

Let $G \curvearrowright X$ be a proper cocompact action on a CAT(0) cube complex. In particular, X is finite-dimensional and locally finite.

Definition 3.1 A subgroup $H \leq G$ is *convex-cocompact* in $G \curvearrowright X$ if there exists an H–invariant, convex subcomplex $C \subseteq X$ that is acted upon cocompactly by H.

Despite the similarity in terminology, we emphasise that the above is much weaker than the notion of "boundary convex-cocompactness" due to Cordes and Durham [34]. For instance, all convex-cocompact subgroups of RAAGs are free if we consider the Cordes–Durham notion [73], whereas every special group is a convex-cocompact subgroup of some RAAG acting on its Salvetti complex according to Definition 3.1; see [68].

Let $[\mu_X]$ be the coarse median structure on G induced by $G \curvearrowright X$ as in Remark 2.27. Recall that quasiconvex subsets of coarse median spaces were introduced in Definition 2.30. For the notion of H–essential core, see Remark 2.15 or [23, Section 3.3].

The following is just a restating of some well-known facts. The equivalence of the first two parts is due to Haglund; see [66, Theorem H] and [91].

Lemma 3.2 *The following are equivalent for a subgroup $H \leq G$:*

1. *H is convex-cocompact in $G \curvearrowright X$.*
2. *H is quasiconvex in $(G, [\mu_X])$.*
3. *H is finitely generated and acts cocompactly on the H–essential core of $H \curvearrowright X$.*

Proof Let us begin with the equivalence of (1) and (2). Picking a vertex $v \in X$, condition (2) holds if and only if there exists a constant R' such that $m(H \cdot v, H \cdot v, G \cdot v) \subseteq \mathcal{N}_{R'}(H \cdot v)$. Since G acts cocompactly and m is 1–Lipschitz in each component, this is equivalent to the existence of R'' with

$$\mathcal{J}(H \cdot v) = m(H \cdot v, H \cdot v, X) \subseteq \mathcal{N}_{R''}(H \cdot v).$$

It is clear that this holds when (1) is satisfied, so (1) $\Longrightarrow$ (2).

Conversely, if (2) holds, then $H \cdot v$ is quasiconvex in X and Remark 2.32 implies that $\mathrm{Hull}(H \cdot v)$ is at finite Hausdorff distance from $H \cdot v$. Since X is locally finite, this means that H acts cocompactly on $\mathrm{Hull}(H \cdot v)$, hence H is convex-cocompact.

We now show the equivalence of (1) and (3). First, if $C \subseteq X$ is convex and H–invariant, the H–essential core of $H \curvearrowright X$ is a restriction quotient of C (as defined in Section 2.5). Thus, if H acts cocompactly on C, it also acts cocompactly on the H–essential core. Moreover, the action $H \curvearrowright C$ is proper and cocompact, which implies that H is finitely generated. This proves (1) $\Longrightarrow$ (3).

Conversely, let X' be the cubical subdivision. Since H is finitely generated and $H \curvearrowright X'$ has no inversions, the essential core of $H \curvearrowright X'$ embeds H–equivariantly as a convex subcomplex of X'; see Remark 2.15. This shows that (3) $\Longrightarrow$ (1). □

Recalling that automorphisms of G are bi-Lipschitz with respect to word metrics on G, the equivalence of (1) and (2) in Lemma 3.2 has the following straightforward consequence.

Corollary 3.3 *If $\varphi \in \mathrm{Aut}(G, [\mu_X])$, then a subgroup $H \le G$ is convex-cocompact in $G \curvearrowright X$ if and only if $\varphi(H)$ is.*

Example 3.4 If G is Gromov-hyperbolic, then a subgroup $H \le G$ is convex-cocompact in $G \curvearrowright X$ if and only if H is quasiconvex in G (again since (1) $\Longleftrightarrow$ (2) in Lemma 3.2). In particular, the notion of convex-cocompactness is independent of the chosen cubulation of G in this case. A quick look at the standard cubulation of $\mathbb{Z}^2$ immediately shows that the latter does not hold in general.

3.2 Label-irreducible elements in RAAGs

This subsection studies convex-cocompact *cyclic* subgroups of right-angled Artin groups. Let Γ be a finite simplicial graph. Let $\mathcal{A} = \mathcal{A}_\Gamma$ be a RAAG and $\mathcal{X} = \mathcal{X}_\Gamma$ the universal cover of its Salvetti complex. Set $r = \dim \mathcal{X}$.

The Cayley graph of $\mathcal{A}$ corresponding to the standard generating set $\Gamma^{(0)}$ is naturally identified with the 1–skeleton of the CAT(0) cube complex $\mathcal{X}$. Thus, every edge of $\mathcal{X}$ is labelled by a vertex of Γ. Observing that edges crossing the same hyperplane have the same label, we obtain a map $\gamma \colon \mathscr{W}(\mathcal{X}) \to \Gamma^{(0)}$.

We can apply the discussion in Section 2.4 to the standard action $\mathcal{A} \curvearrowright \mathcal{X}$ (or, to be precise, the action on the 0–skeleton of $\mathcal{X}$). Every element of $\mathcal{A}$ acts nontransversely and stably without inversions. For every $g \in \mathcal{A} \setminus \{1\}$, the reduced core $\overline{\mathcal{C}}(g)$ is the union of all axes of g.

A hyperplane of $\mathcal{X}$ lies in $\mathcal{W}_1(g)$ if and only if it is crossed by one (equivalently, all) axis of g. Hyperplanes lie in $\mathcal{W}_0(g)$ when they are preserved by g; equivalently, when they are transverse to all elements of $\mathcal{W}_1(g)$, or, again, when they separate two axes of g.

The factor $\overline{\mathcal{C}}_1(g)$ is $\langle g\rangle$–equivariantly isomorphic to the convex hull in $\mathcal{X}$ of any axis of g. The factor $\overline{\mathcal{C}}_0(g)$ is fixed pointwise by g and it is isomorphic to $\mathcal{X}_\Lambda$, where $\Lambda\subseteq\Gamma$ is the maximal subgraph all of whose vertices are joined by an edge to all vertices in $\gamma(\mathcal{W}_1(g))$.

For a simplicial graph Δ, we denote by Δ^o the *opposite* of Δ. This the graph that has the same vertex set as Δ and an edge between two vertices exactly when they are not connected by an edge in Δ.

Definition 3.5 Consider $g\in\mathcal{A}\setminus\{1\}$.

(1) We define $\Gamma(g):=\gamma(\mathcal{W}_1(g))\subseteq\Gamma^{(0)}$. These are precisely the standard generators of $\mathcal{A}$ that appear in the cyclically reduced words representing elements conjugate to g.

(2) We say that g is *label-irreducible* if the full subgraph of Γ spanned by $\Gamma(g)$ does not split as a nontrivial join (ie its opposite graph is connected). Equivalently, g is contracting [26] within a parabolic subgroup of $\mathcal{A}$.

Two label-irreducible elements $g,h\in\mathcal{A}$ are *independent* if $\langle g,h\rangle\not\simeq\mathbb{Z}$. If g,h are independent and commute, then $\langle g,h\rangle\simeq\mathbb{Z}^2$. We will also use the following result of Servatius; see eg [94, Proposition III.1].

Lemma 3.6 *If $g,h\in\mathcal{A}$ are commuting, independent, label-irreducible elements, then every vertex of $\Gamma(g)$ is joined to every vertex of $\Gamma(h)$ by an edge of Γ.*

To each element $g\in\mathcal{A}$, we can associate a canonical collection of label-irreducible elements $g_1,\dots,g_k$, called the *label-irreducible components* of g, as shown in the next result.

Lemma 3.7 (label-irreducible components) *For every element $g\in\mathcal{A}$, the following hold.*

(1) *We can write $g=g_1\cdot\dots\cdot g_k$ for pairwise-commuting, pairwise-independent label-irreducibles $g_i\in\mathcal{A}$. In addition, $0\le k\le r$ and the g_i are unique up to permutation.*

(2) *The sets $\mathcal{W}_1(g_i)$ are transverse to each other and $\mathcal{W}_1(g_i)\subseteq\mathcal{W}_0(g_j)$ for $i\neq j$. In addition,*
$$\mathcal{W}_1(g)=\mathcal{W}_1(g_1)\sqcup\dots\sqcup\mathcal{W}_1(g_k),\quad \ell(g,\mathcal{X})=\ell(g_1,\mathcal{X})+\dots+\ell(g_k,\mathcal{X}),$$
$$\overline{\mathcal{C}}_1(g)\simeq\overline{\mathcal{C}}_1(g_1)\times\dots\times\overline{\mathcal{C}}_1(g_k),\qquad \overline{\mathcal{C}}(g)=\overline{\mathcal{C}}(g_1)\cap\dots\cap\overline{\mathcal{C}}(g_k).$$

(3) *Centralisers satisfy $Z_{\mathcal{A}}(g)=Z_{\mathcal{A}}(g_1)\cap\dots\cap Z_{\mathcal{A}}(g_k)$. Moreover, $Z_{\mathcal{A}}(g)$ splits as the direct product of a parabolic subgroup of $\mathcal{A}$ and a copy of $\mathbb{Z}^k$ freely generated by roots of $g_1,\dots,g_k$.*

Proof Since label-irreducibility is invariant under taking conjugates, we assume throughout the proof that g is cyclically reduced. If g is the identity, we can simply take $k=0$ and the entire lemma holds trivially. Suppose instead that $g\neq 1$.

We begin with part (1). Let $\Lambda_1,\dots,\Lambda_k$ be the connected components of the subgraph of Γ^o spanned by $\Gamma(g)$. In Γ, every vertex of Λ_i is joined by an edge to every vertex of Λ_j with $j\neq i$. Thus, permuting the letters in a word representing g, we can write $g = g_1 \cdot \dots \cdot g_k$, where each g_i is cyclically reduced and $\Gamma(g_i) = \Lambda_i^{(0)}$. The elements g_i commute pairwise and, since each Λ_i is connected, they are all label-irreducible. It is clear that g_i and g_j are independent for $i \neq j$.

Uniqueness of the g_i up to permutations follows from the fact that, by Lemma 3.6, $\Gamma(g_1), \dots, \Gamma(g_k)$ must coincide with the vertex sets of $\Lambda_1, \dots, \Lambda_k$ in any such decomposition of g. Furthermore, choosing a vertex from each Λ_i, we obtain a k–clique in Γ, so $k \leq r$. This proves part (1).

We now prove part (2). Since g_i is cyclically reduced, there exists a (combinatorial) axis $\alpha_i \subseteq \mathcal{X}_{\Gamma(g_i)}$ through the identity. Note that the product $\mathcal{X}_{\Gamma(g_1)} \times \dots \times \mathcal{X}_{\Gamma(g_k)} \subseteq \mathcal{X}$ is preserved by all g_i and that each g_i leaves invariant every hyperplane of $\mathcal{X}_{\Gamma(g_j)}$ for all $j \neq i$. Thus, the sets $\mathcal{W}_1(g_i)$ are transverse to each other and $\mathcal{W}_1(g_i) \subseteq \mathcal{W}_0(g_j)$ for $i \neq j$. The equality $\mathcal{W}_1(g) = \mathcal{W}_1(g_1) \sqcup \dots \sqcup \mathcal{W}_1(g_k)$ now follows by observing that $\alpha_1 \times \dots \times \alpha_k$ contains an axis of g. The product splitting of $\bar{\mathcal{C}}_1(g)$ can be deduced from the transverse partition of $\mathcal{W}_1(g)$ using Lemma 2.6.

If Ω_i is a fundamental domain for the $\langle g_i \rangle$–action on $\mathcal{W}_1(g_i)$, the previous paragraph implies that $\Omega_1 \sqcup \dots \sqcup \Omega_k$ is a fundamental domain for the $\langle g \rangle$–action on $\mathcal{W}_1(g)$. Taking cardinalities, this shows that $\ell(g, \mathcal{X}) = \ell(g_1, \mathcal{X}) + \dots + \ell(g_k, \mathcal{X})$. Finally, the characterisation of $\bar{\mathcal{C}}(g)$ can be deduced from the fact that this is the set of points of $\mathcal{X}$ that realise the translation length $\ell(g, \mathcal{X})$.

We conclude with part (3). The inclusion $Z_{\mathcal{A}}(g_1) \cap \dots \cap Z_{\mathcal{A}}(g_k) \leq Z_{\mathcal{A}}(g)$ is clear. Conversely, if $h \in \mathcal{A}$ commutes with g, uniqueness in part (1) implies that the elements hg_ih^{-1} coincide with the g_i up to permutation. Since $\Gamma(hg_ih^{-1}) = \Gamma(g_i)$, it follows that $hg_ih^{-1} = g_i$ for each i. Hence $h \in Z_{\mathcal{A}}(g_1) \cap \dots \cap Z_{\mathcal{A}}(g_k)$, as required. The last statement is Servatius' centraliser theorem from [94, Section III]. □

Remark 3.8 For every $H \leq \mathcal{A}$, there exists a finite subset $F \subseteq H$ such that $Z_{\mathcal{A}}(H) = Z_{\mathcal{A}}(F)$.

Indeed, we have observed in Lemma 3.7(3) that the centraliser of every element of $\mathcal{A}$ splits as a product of a free abelian group and a parabolic subgroup of $\mathcal{A}$. It follows that every descending chain of centralisers of subsets of $\mathcal{A}$ eventually stabilises, since this is true of chains of parabolics.

We conclude this subsection by showing that label-irreducibles are precisely those elements $g \in \mathcal{A}$ such that the subgroup $\langle g \rangle$ is convex-cocompact in $\mathcal{A} \curvearrowright \mathcal{X}$. After a couple of preliminary results, this is shown below in Lemma 3.11.

Lemma 3.9 *Every connected full subgraph $\Lambda \subseteq \Gamma^o$ has diameter $\leq 2r - 1$.*

Proof Suppose towards a contradiction that there exist vertices $x, y \in \Lambda$ and a shortest path $\sigma \subseteq \Lambda$ joining them, with σ made up of $2r$ edges. Let σ_i be the i^{th} vertex of Γ^o met by σ, with $\sigma_0 = x$ and $\sigma_{2r} = y$. Since σ is shortest and Λ is full, no two of the $r + 1$ vertices $\sigma_0, \sigma_2, \dots, \sigma_{2r}$ are joined by an edge of Γ^o. Thus, they form an $(r + 1)$–clique in Γ, contradicting the fact that $r = \dim \mathcal{X}$. □

Lemma 3.10 *Let $g \in \mathcal{A}$ be label-irreducible. Then, for every $\mathfrak{u} \in \mathcal{W}_1(g)$, there exists a point $x \in \bar{\mathcal{C}}(g)$ such that $\mathscr{W}(x|gx) \subseteq \mathscr{W}(\mathfrak{u}|g^{4r-2}\mathfrak{u})$. In particular, $\gamma(\mathscr{W}(\mathfrak{u}|g^{4r-2}\mathfrak{u})) = \Gamma(g)$.*

Proof Pick a point y on an axis of g so that $\mathfrak{u} \in \mathscr{W}(y|gy)$. Set $x = g^{2r-1}y$ and consider a hyperplane $\mathfrak{w} \in \mathscr{W}(x|gx)$. Since g is label-irreducible, the full subgraph of Γ^o spanned by $\Gamma(g)$ is connected. By Lemma 3.9, there exists a sequence $\sigma_0 = \gamma(\mathfrak{u}), \sigma_1, \dots, \sigma_k = \gamma(\mathfrak{w})$ of vertices in $\Gamma(g)$ such that $k \leq 2r-1$ and consecutive σ_i are not joined by an edge of Γ. Set $\sigma_j = \sigma_k$ for $k < j \leq 2r-1$.

For $0 \leq i \leq 2r-1$, pick a hyperplane $\mathfrak{w}_i \in \mathscr{W}(g^i y|g^{i+1}y)$ with $\gamma(\mathfrak{w}_i) = \sigma_i$, making sure that $\mathfrak{w}_0 = \mathfrak{u}$ and $\mathfrak{w}_{2r-1} = \mathfrak{w}$. Since σ_i and σ_{i+1} are not joined by an edge, the hyperplanes $\mathfrak{w}_i$ and $\mathfrak{w}_{i+1}$ are not transverse. Since these hyperplanes are all crossed by an axis of g, we conclude that each $\mathfrak{w}_i$ separates the $\mathfrak{w}_j$ with $j < i$ from those with $j > i$. In particular, $\mathfrak{u}$ and $\mathfrak{w}$ are not transverse.

The same argument shows that $\mathfrak{w}$ and $g^{4r-2}\mathfrak{u}$ are not transverse, hence $\mathfrak{w} \in \mathscr{W}(\mathfrak{u}|g^{4r-2}\mathfrak{u})$. Since $\mathfrak{w} \in \mathscr{W}(x|gx)$ was arbitrary, we have shown that $\mathscr{W}(x|gx) \subseteq \mathscr{W}(\mathfrak{u}|g^{4r-2}\mathfrak{u})$. □

Lemma 3.11 (1) *If g is label-irreducible and $\alpha \subseteq \mathcal{X}$ is an axis, then $d_{\mathrm{Haus}}(\alpha, \mathrm{Hull}\,\alpha) \leq (8r-4)\ell(g, \mathcal{X})$.*

(2) *An element $g \in \mathcal{A} \setminus \{1\}$ is label-irreducible if and only if $\langle g \rangle$ is convex-cocompact in $\mathcal{A} \curvearrowright \mathcal{X}$.*

Proof Assuming part (1), we first prove part (2). Using the third characterisation of convex-cocompactness in Lemma 3.2 and the fact that $\bar{\mathcal{C}}_1(g)$ is equivariantly isomorphic to $\mathrm{Hull}\,\alpha$, part (1) shows that label-irreducible elements are convex-cocompact. Conversely, if g is not label-irreducible, the nontrivial splitting of $\bar{\mathcal{C}}_1(g)$ provided by Lemma 3.7(2) implies that $\langle g \rangle$ cannot act cocompactly on $\bar{\mathcal{C}}_1(g)$.

Let us now prove part (1). Considering a point $p \in \mathrm{Hull}\,\alpha$, it is enough to obtain the inequality $d(p, \alpha) \leq (8r-4)\ell(g, \mathcal{X})$.

Every element of $\mathscr{H}_{\mathrm{Hull}\,\alpha}(\mathcal{X})$ intersects α in a subray. Let $\mathcal{H}_+$ be the subset of halfspaces intersecting α in a positive subray (ie containing all points $g^n z$ with $n \geq 0$, for a suitable choice of $z \in \alpha$). Any two maximal halfspaces lying in $\mathcal{H}_+$ and not containing p are transverse. It follows that there are only finitely many such maximal halfspaces, which we denote by $\mathfrak{h}_1, \dots, \mathfrak{h}_k$.

A negative subray of α is contained in $\mathfrak{h}_1^* \cap \dots \cap \mathfrak{h}_k^*$, so we can pick a point $x \in \alpha \cap \mathfrak{h}_1^* \cap \dots \cap \mathfrak{h}_k^*$. In particular, x does not lie in any halfspaces of $\mathcal{H}_+$ that do not contain p; hence $\mathscr{H}(x|p) \subseteq \mathcal{H}_+$. Let $y \in \alpha$ be the point with $d(x, p) = d(x, y)$ and $\mathscr{H}(x|y) \subseteq \mathcal{H}_+$. Setting $m = m(x, p, y)$, we note that every $\mathfrak{j} \in \mathscr{H}(m|p)$ is transverse to every $\mathfrak{k} \in \mathscr{H}(m|y)$. Indeed, $m \in \mathfrak{j}^* \cap \mathfrak{k}^*$, $p \in \mathfrak{j} \cap \mathfrak{k}^*$ and $y \in \mathfrak{j}^* \cap \mathfrak{k}$, while $\mathfrak{j} \cap \mathfrak{k}$ is nonempty because $\mathfrak{j}$ and $\mathfrak{k}$ both lie in $\mathcal{H}_+$.

Now, suppose for the sake of contradiction that $d(p, y) > (8r-4)\ell(g, \mathcal{X})$. Since we chose y with $d(x, p) = d(x, y)$, we have $d(p, m) = d(m, y) > (4r-2)\ell(g, \mathcal{X})$. Now $\mathscr{W}(p|m) \subseteq \mathscr{W}_{\mathrm{Hull}\,\alpha}(\mathcal{X}) = \mathcal{W}_1(g)$, a set on which $\langle g^{4r-2} \rangle$ acts with exactly $(4r-2)\ell(g, \mathcal{X})$ orbits. Thus, there exists a hyperplane $\mathfrak{u} \in \mathscr{W}(p|m)$ such that $g^{4r-2}\mathfrak{u} \in \mathscr{W}(p|m)$. Lemma 3.10 implies that $\gamma(\mathscr{W}(p|m)) = \Gamma(g)$. Similarly, we obtain $\gamma(\mathscr{W}(m|y)) = \Gamma(g)$. This contradicts the fact that $\mathscr{W}(p|m)$ is transverse to $\mathscr{W}(m|y)$. □

3.3 More on convex-cocompactness in RAAGs

The results in this subsection will only be used in Section 7 and can be skipped by the reader uninterested in the proof of Theorems F and I.

First, we discuss additional properties of label-irreducible elements of RAAGs. Our aim is obtaining *uniform* control on the extent to which axes of distinct label-irreducibles can track each other. The main result here is Lemma 3.13, along with its direct consequence Corollary 3.14. Both results will be fundamental building blocks in the proof that centreless special groups are UNE.

Then, in the second part of the subsection, we study general convex-cocompact subgroups of RAAGs, proving only a couple of simple properties related to label-irreducible components.

3.3.1 Additional properties of label-irreducible elements We maintain the notation introduced at the beginning of Section 3.2. Recall that $r = \dim \mathcal{X}$.

Recall that the carrier of a hyperplane $\mathfrak{w} \in \mathscr{W}(\mathcal{X})$ is the smallest convex subcomplex of $\mathcal{X}$ that contains all edges crossing $\mathfrak{w}$. A hyperplane of $\mathcal{X}$ separates two points in the carrier of $\mathfrak{w}$ if and only if it is either equal or transverse to $\mathfrak{w}$. If two hyperplanes $\mathfrak{u}$ and $\mathfrak{w}$ have intersecting carriers, then they are transverse if and only if $\gamma(\mathfrak{u})$ and $\gamma(\mathfrak{w})$ are joined by an edge of Γ.

Lemma 3.12 *If $g, h \in \mathcal{A}$ and $\Gamma(g) \subseteq \gamma\big(\mathscr{W}_{\bar{\mathcal{C}}(g)}(\mathcal{X}) \cap \mathscr{W}_{\bar{\mathcal{C}}(h)}(\mathcal{X})\big)$, then $\bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h) \neq \varnothing$.*

Proof Suppose for the sake of contradiction that $\bar{\mathcal{C}}(g)$ and $\bar{\mathcal{C}}(h)$ are disjoint. Then there exists a hyperplane $\mathfrak{v}$ separating them, which we pick so that the carrier of $\mathfrak{v}$ intersects $\bar{\mathcal{C}}(g)$. This guarantees that g admits an axis α that intersects the carrier of $\mathfrak{v}$.

Since $\mathfrak{v}$ separates $\bar{\mathcal{C}}(g)$ and $\bar{\mathcal{C}}(h)$, it is transverse to $\mathscr{W}_{\bar{\mathcal{C}}(g)}(\mathcal{X}) \cap \mathscr{W}_{\bar{\mathcal{C}}(h)}(\mathcal{X})$, so the vertex $\gamma(\mathfrak{v})$ is connected by an edge of Γ to all elements of $\Gamma(g)$. Observing that all hyperplanes crossed by α are labelled by elements of $\Gamma(g)$ and recalling that α intersects the carrier of $\mathfrak{v}$, we deduce that all hyperplanes crossed by α are transverse to $\mathfrak{v}$. In other words, $\mathfrak{v}$ is transverse to $\mathcal{W}_1(g)$, hence $\mathfrak{v} \in \mathcal{W}_0(g) \subseteq \mathscr{W}_{\bar{\mathcal{C}}(g)}(\mathcal{X})$. This is the required contradiction. □

Lemma 3.13 *Let $g, h \in \mathcal{A}$ be label-irreducible. If there exist hyperplanes $\mathfrak{u}, \mathfrak{w} \in \mathscr{W}(\mathcal{X})$ such that $\{\mathfrak{u}, g^{4r}\mathfrak{u}, \mathfrak{w}, h^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h)$, then $\langle g, h\rangle \simeq \mathbb{Z}$.*

Proof The proof will consist of three steps.

Step 1 *We can assume that $1 \in \mathcal{A} \cong \mathcal{X}^{(0)}$ lies in $\bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h)$, and that $\Gamma(g) = \Gamma(h) = \Gamma^{(0)}$.*

Since $\mathcal{W}_1(g) \cap \mathcal{W}_1(h)$ contains any hyperplane separating two of its elements, we have $\mathscr{W}(\mathfrak{u}|g^{4r-2}\mathfrak{u}) \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h)$. Lemma 3.10 yields

$$\Gamma(g) = \gamma(\mathscr{W}(\mathfrak{u}|g^{4r-2}\mathfrak{u})) \subseteq \gamma(\mathcal{W}_1(g) \cap \mathcal{W}_1(h)) \subseteq \Gamma(h).$$

One the one hand, this allows us to apply Lemma 3.12 and deduce that $\bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h) \neq \varnothing$. On the other, this shows that $\Gamma(g) \subseteq \Gamma(h)$ and the inclusion $\Gamma(h) \subseteq \Gamma(g)$ is obtained similarly, so $\Gamma(g) = \Gamma(h)$.

Conjugating g and h by any $x \in \bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h)$, we can assume that $1 \in \bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h)$. Equivalently, g and h are cyclically reduced, so they lie in the parabolic subgroup $\mathcal{A}_{\Gamma(g)} = \mathcal{A}_{\Gamma(h)} \leq \mathcal{A}$. Replacing $\mathcal{A}$ with $\mathcal{A}_{\Gamma(g)}$ does not alter the properties in the statement of the lemma, so we can assume that $\Gamma(g) = \Gamma(h) = \Gamma^{(0)}$.

Step 2 *Assume without loss of generality that $\ell(g, \mathcal{X}) \leq \ell(h, \mathcal{X})$. Possibly replacing g and h with their inverses and conjugating them, there exists a geodesic $\sigma \subseteq \mathcal{X}$ from 1 to g such that*

- *the union $\rho := \bigcup_{i \geq 0} g^i \sigma$ is a ray and contains h and h^2 (viewing $1, g, h, h^2$ as vertices of $\mathcal{X}$), and*
- *if $\tau \subseteq \rho$ is the arc joining 1 to h, then $h \cdot \tau$ is the arc of ρ joining h to h^2.*

Let $\mathfrak{k} \in \mathscr{H}(\mathcal{X})$ be a halfspace bounded by $h^{4r-2}\mathfrak{w} \in \mathcal{W}_1(g) \cap \mathcal{W}_1(h)$. Possibly replacing g and/or h with their inverses, we have $g\mathfrak{k} \subsetneq \mathfrak{k}$ and $h\mathfrak{k} \subsetneq \mathfrak{k}$. Since $\Gamma^{(0)} = \Gamma(h)$, Lemma 3.10 shows that $\mathfrak{w}$ and $h^{4r-2}\mathfrak{w}$ are strongly separated in $\mathcal{X}$.

The subray contained in $\mathfrak{k}^*$ of any (combinatorial) axis of g defines a point ξ in the Roller boundary $\partial\mathcal{X}$ such that $g\xi = \xi$ and $\xi \in h^{-4r+2}\mathfrak{k}^*$ (recall that this halfspace is bounded by $\mathfrak{w}$). Similarly, there exists $\eta \in \partial\mathcal{X}$ with $h\eta = \eta$ and $\eta \in h^{-4r+2}\mathfrak{k}^*$. Since the halfspaces $h^{-4r+2}\mathfrak{k}^*$ and $\mathfrak{k}$ are strongly separated, the gate-projections of ξ and η to $\mathfrak{k}$ coincide and they are a vertex $x \in \bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h)$. Conjugating g and h by x, we can assume that $x = 1$.

Label $\mathfrak{k}_1 \supsetneq \mathfrak{k}_2 \supsetneq \cdots \supsetneq \mathfrak{k}_m$ the elements of $\mathscr{H}(1|h^2)$ bounded by hyperplanes with label $\gamma(\mathfrak{w})$. Set $\mathfrak{k}_0 := \mathfrak{k}$ and observe that $\mathfrak{k}_m = h^2\mathfrak{k}$, which is bounded by $h^{4r}\mathfrak{w} \in \mathcal{W}_1(g) \cap \mathcal{W}_1(h)$. In conclusion,

$$\xi, \eta \notin h^{-4r+2}\mathfrak{k} \supsetneq \mathfrak{k} = \mathfrak{k}_0 \supsetneq \mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_m = h^2\mathfrak{k}.$$

Note that the hyperplanes bounding the $\mathfrak{k}_i$ all lie in $\mathcal{W}_1(g) \cap \mathcal{W}_1(h)$. Since $1 \in \bar{\mathcal{C}}(g) \cap \bar{\mathcal{C}}(h)$, there exist an axis of h and an axis of g each crossing all hyperplanes bounding the $\mathfrak{k}_i$. Hence there exist $1 \leq t \leq s$ such that $g\mathfrak{k}_j = \mathfrak{k}_{j+t}$ for all $0 \leq j \leq m-t$, and $h\mathfrak{k}_i = \mathfrak{k}_{i+s}$ for all $0 \leq i \leq m-s$.

Let x_i be the gate-projection of $x = 1$ to $\mathfrak{k}_i$. Note that this is also the gate-projection to $\mathfrak{k}_i$ of ξ and η. Since $g\xi = \xi$ and $h\eta = \eta$, we must have $gx_j = x_{j+t}$ and $hx_i = x_{i+s}$ for all $1 \leq j \leq m-t$ and $1 \leq i \leq m-s$. In particular, since $x_0 = 1$, we have $h = x_s$, $h^2 = x_{2s} = x_m$ and $g = x_t$.

Observe that each x_i is also the gate-projection to $\mathfrak{k}_i$ of each x_j with $j < i$. Thus, we can construct a (combinatorial) geodesic σ from 1 to g by concatenating arbitrary geodesics σ_j from x_j to x_{j+1} for $0 \leq j < t$. The union $\rho = \bigcup_{i \geq 0} g^i \sigma$ is a ray since $1 \in \bar{\mathcal{C}}(g)$. Let $k, l \geq 1$ be the integers with $0 \leq s - kt < t$ and $0 \leq 2s - lt < t$. Since σ contains the points $g^{-k}h = x_{s-kt}$ and $g^{-l}h^2 = x_{2s-lt}$, it is clear that h and h^2 lie on the ray ρ.

Finally, note that we can choose the geodesics σ_j so that the following compatibility condition is satisfied: *whenever there exist $f \in \mathcal{A}$ and $0 \leq i, j < t$ with $fx_i = x_j$ and $fx_{i+1} = x_{j+1}$, we have $f\sigma_i = \sigma_j$*. This is possible because the action $\mathcal{A} \curvearrowright \mathcal{X}$ is free and so the element f is uniquely determined by i and j (when it exists). Now, given $0 \le j < s$, the arc of the ray ρ joining x_{s+j} to x_{s+j+1} is precisely $g^{a_j}\sigma_{b_j}$, where $s+j = a_j t + b_j$ and $0 \le b_j < t$. The element $g^{-a_j}h$ maps x_j and x_{j+1} to x_{b_j} and x_{b_j+1}, so it takes σ_j to σ_{b_j} by our construction. Thus $h\sigma_j = g^{a_j}\sigma_{b_j}$ is contained in ρ for every $0 \le j < s$. This proves the second condition in the statement of Step 2.

Step 3 *We have $\langle g, h\rangle \simeq \mathbb{Z}$.*

Let $S \cong \Gamma^{(0)}$ be the standard generating set of $\mathcal{A}$. Let $F(S)$ be the free group freely generated by S, and let $\pi\colon F(S) \to \mathcal{A}$ be the surjective homomorphism that takes each generator of $F(S)$ to the corresponding standard generator of $\mathcal{A}$. Let $w_g \in F(S)$ be the word spelled by the labels of the edges met moving from 1 to g along the geodesic σ. Let $w_h \in F(S)$ be the word spelled moving from 1 to h along the ray $\rho = \bigcup_{i\ge 0} g^i\sigma$. It is clear that $\pi(w_g) = g$ and $\pi(w_h) = h$.

From Step 2, we have $w_h = w_g^p a$ for some $p \ge 1$ and an initial subword a of w_g, and $w_h^2 = w_g^{p+1}ab$ for some word b such that $w_g^{p+1}ab$ is reduced in $F(S)$. It follows that $w_g^p a w_g^p a = w_g^{p+1}ab$ in $F(S)$, where both sides of the equality are reduced words. Looking at the first $(p+1)|w_g| + |a|$ letters on the left, we deduce that $aw_g = w_g a$. Hence $\langle w_g, w_h\rangle = \langle w_g, a\rangle$ is a cyclic subgroup of $F(S)$. We conclude that $\langle g, h\rangle = \pi\left(\langle w_g, w_h\rangle\right) \simeq \mathbb{Z}$. □

Corollary 3.14 *Consider two elements $g, h \in \mathcal{A}$. Suppose that g is label-irreducible. Assume in addition that **one** of the following conditions is satisfied*:

(1) *There exists $\mathfrak{w} \in \mathcal{W}_1(g)$ such that h preserves $\mathfrak{w}$ and $g^{4r}\mathfrak{w}$.*

(2) *There exist hyperplanes $\mathfrak{u}, \mathfrak{w} \in \mathscr{W}(\mathcal{X})$ with $\{\mathfrak{u}, \mathfrak{w}, h^{4r}\mathfrak{u}, g^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h)$.*

Then g and h commute in $\mathcal{A}$.

Proof Assume first that there exists $\mathfrak{w} \in \mathcal{W}_1(g)$ such that $\mathfrak{w}$ and $g^{4r}\mathfrak{w}$ are preserved by h. Then $\{\mathfrak{w}, g^{4r}\mathfrak{w}\} = \{\mathfrak{w}, (hgh^{-1})^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(hgh^{-1})$. Since g and hgh^{-1} are label-irreducible, Lemma 3.13 implies that $\langle g, hgh^{-1}\rangle \simeq \mathbb{Z}$. Observing that $\ell(g, \mathcal{X}) = \ell(hgh^{-1}, \mathcal{X})$, we deduce that hgh^{-1} must coincide with either g or g^{-1}. The second option cannot occur in a right-angled Artin group, hence $hgh^{-1} = g$, as required.

Suppose now that there exist hyperplanes $\mathfrak{u}, \mathfrak{w}$ with $\{\mathfrak{u}, \mathfrak{w}, h^{4r}\mathfrak{u}, g^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h)$. In light of Lemma 3.7(2), there exist (possibly equal) irreducible components h_1, h_2 of h, such that $\{\mathfrak{u}, g^{4r}\mathfrak{u}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h_1)$ and $\{\mathfrak{w}, h^{4r}\mathfrak{w}\} = \{\mathfrak{w}, h_2^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h_2)$.

Since g is label-irreducible and $\gamma(\mathscr{W}(\mathfrak{u}|g^{4r}\mathfrak{u})) = \Gamma(g)$ by Lemma 3.10, no element of $\mathcal{W}_1(g)$ can be transverse to both $\mathfrak{u}$ and $g^{4r}\mathfrak{u}$. Hence $h_1 = h_2$, otherwise $\mathcal{W}_1(h_1)$ and $\mathcal{W}_1(h_2)$ would be transverse. Thus $\{\mathfrak{u}, g^{4r}\mathfrak{u}, \mathfrak{w}, h_2^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g) \cap \mathcal{W}_1(h_2)$ and Lemma 3.13 yields $\langle g, h_2\rangle \simeq \mathbb{Z}$. Now, a power of g coincides with a power of h_2, hence it commutes with h. It follows that g and h commute. □

We conclude with the following lemma, which is actually independent from the notion of label-irreducibility and from the discussion in the rest of this subsection, albeit in a similar spirit.

Lemma 3.15 *Consider elements $g_1, g_2 \in \mathcal{A}$ and vertices $x_1, x_2 \in \mathcal{X}$ such that the two sets $\mathscr{W}(x_i|g_i x_i)$ are transverse. Then g_1 and g_2 commute and we have $\mathcal{W}_1(g_i) \subseteq \mathcal{W}_0(g_j)$ for $i \neq j$.*

Proof We begin with the following observation:

Claim *For $\mathfrak{w} \in \mathscr{W}(\mathcal{X})$ and $x, y \in \mathcal{X}$, the hyperplane $\mathfrak{w}$ is transverse to $\mathscr{W}(x|y)$ if and only if every vertex in the set $\gamma(\mathscr{W}(x|y))$ is joined by an edge of Γ to every vertex in the set $\{\gamma(\mathfrak{w})\} \cup \gamma(\mathscr{W}(x|\mathfrak{w}))$.*

Proof Suppose first that $\mathfrak{w}$ is transverse to $\mathscr{W}(x|y)$. Then, since every hyperplane in $\mathscr{W}(\mathfrak{w}|x)$ is disjoint from $\mathfrak{w}$, we have $\mathscr{W}(\mathfrak{w}|x) \subseteq \mathscr{W}(\mathfrak{w}|x, y)$. Since $\mathscr{W}(x|y)$ is transverse to $\mathscr{W}(\mathfrak{w}|x, y)$, it is also transverse to $\{\mathfrak{w}\} \cup \mathscr{W}(\mathfrak{w}|x)$. It follows that every vertex in $\gamma(\mathscr{W}(x|y))$ is joined by an edge to every vertex in $\{\gamma(\mathfrak{w})\} \cup \gamma(\mathscr{W}(x|\mathfrak{w}))$, as required.

Suppose now instead that $\mathfrak{w}$ is disjoint from a hyperplane $\mathfrak{u} \in \mathscr{W}(x|y)$. Choosing $\mathfrak{u}$ so that it is closest to $\mathfrak{w}$, we can assume that no hyperplane of $\mathscr{W}(x|y)$ separates $\mathfrak{w}$ and $\mathfrak{u}$. If the carriers of $\mathfrak{w}$ and $\mathfrak{u}$ intersect, then $\gamma(\mathfrak{w})$ and $\gamma(\mathfrak{u})$ cannot be joined by an edge, as $\mathfrak{u}$ and $\mathfrak{w}$ are disjoint. Otherwise, there exists a hyperplane $\mathfrak{v} \in \mathscr{W}(\mathfrak{u}|\mathfrak{w})$ such that its carrier intersects the carrier of $\mathfrak{u}$; in particular, $\gamma(\mathfrak{u})$ and $\gamma(\mathfrak{v})$ are not joined by an edge. Since $\mathfrak{v}$ does not separate x and y, we must have $\mathfrak{v} \in \mathscr{W}(\mathfrak{w}|x, y)$, hence $\gamma(\mathfrak{v}) \in \gamma(\mathscr{W}(x|\mathfrak{w}))$, as required. ◁

Consider for a moment $g \in \mathcal{A}$, $x \in \mathcal{X}$ and $n \geq 1$. Since $\mathscr{W}(x|g^n x)$ is contained in the union $\mathscr{W}(x|gx) \cup \cdots \cup \mathscr{W}(g^{n-1}x|g^n x)$, we have $\gamma(\mathscr{W}(x|g^n x)) \subseteq \gamma(\mathscr{W}(x|gx))$. Thus, the claim implies that a hyperplane $\mathfrak{w}$ is transverse to $\mathscr{W}(x|gx)$ if and only if it is transverse to $\bigcup_{n\in\mathbb{Z}} \mathscr{W}(x|g^n x)$.

Now, consider the situation in the statement of the lemma. If x_i' is the gate-projection of x_i to $\bar{\mathcal{C}}(g_i)$, we have $\mathscr{W}(x_i'|g_i x_i') \subseteq \mathscr{W}(x_i|g_i x_i)$ and $\mathcal{W}_1(g_i) = \bigcup_{n\in\mathbb{Z}} \mathscr{W}(x_i'|g_i^n x_i')$. It follows that the sets $\mathcal{W}_1(g_1)$ and $\mathcal{W}_1(g_2)$ are transverse, or, equivalently, $\mathcal{W}_1(g_i) \subseteq \mathcal{W}_0(g_j)$ for $i \neq j$. This implies that g_1 and g_2 commute (for instance, by decomposing g_i into label-irreducible components as in Lemma 3.7 and applying Corollary 3.14). □

3.3.2 Convex-cocompact subgroups of RAAGs Again, we keep the notation from Section 3.2. We will simply say that a subgroup $G \leq \mathcal{A}$ is *convex-cocompact* when G is convex-cocompact for the action $\mathcal{A} \curvearrowright \mathcal{X}$ (in the sense of Definition 3.1).

Lemma 3.16 *Let $G \leq \mathcal{A}$ be convex-cocompact. If $g \in G$ and $g = a_1 \cdot \dots \cdot a_k$ is its decomposition into label-irreducible components $a_i \in \mathcal{A}$, then there exists $m \geq 1$ such that all a_i^m lie in G.*

Proof Let $A \leq G$ be a free abelian subgroup containing a power of g, such that no finite-index subgroup of A is contained in a free abelian subgroup of G of higher rank. Since G is convex-cocompact, Theorem 3.6 in [100] shows that there exists a convex, A–invariant, A–cocompact subcomplex $Y \subseteq \mathcal{X}$ that splits as a product $L_1 \times \cdots \times L_p$, where $A \simeq \mathbb{Z}^p$ and each L_i is a quasiline. Replacing each L_i with a subcomplex, we can assume that all quasilines are A–essential.

Note that Y must contain an axis of g in $\mathcal{X}$, hence its convex hull, which is isomorphic to:

$$\bar{\mathcal{C}}_1(g) = \bar{\mathcal{C}}_1(a_1) \times \cdots \times \bar{\mathcal{C}}_1(a_k).$$

Since each a_i is label-irreducible, Lemma 3.11 shows that $\bar{\mathcal{C}}_1(a_i)$ is an irreducible quasiline. Up to permuting the factors of Y, we can thus assume that $L_i \simeq \bar{\mathcal{C}}_1(a_i)$ for $1 \leq i \leq k$, where $k \leq p$.

Since the L_i are locally finite, none of the groups $\operatorname{Aut} L_i$ contains subgroups isomorphic to $\mathbb{Z}^2$. It follows that every projection of $\mathbb{Z}^p \simeq A \leq \prod_i \operatorname{Aut} L_i$ to a product of $(p-1)$ factors must have nontrivial kernel. Equivalently, there exist elements $h_i \in A$ such that h_i acts loxodromically on L_i, and fixes pointwise each L_j with $j \neq i$. For each $1 \leq i \leq k$, the elements h_i and a_i stabilise a common copy of $L_i \simeq \bar{\mathcal{C}}_1(a_i)$ inside Y, and act freely and cocompactly on it. It follows that h_i and a_i are commensurable, hence a power of a_i lies in $A \leq G$. This concludes the proof. □

The exponent m in Lemma 3.16 can be chosen independently of $g \in G$ due to the following.

Remark 3.17 Suppose that $G \leq \mathcal{A}$ is convex-cocompact and, more precisely, that there exists a G–invariant, convex subcomplex $Y \subseteq \mathcal{X}$ such that the action $G \curvearrowright Y^{(0)}$ has q orbits. Then, for every $g \in \mathcal{A}$ such that $\langle g \rangle \cap G \neq \{1\}$, there exists $1 \leq k \leq q$ such that $g^k \in G$.

Indeed, consider $N \geq 1$ such that $g^N \in G$. Since Y is G–invariant and acted upon without inversions, it contains an axis α for g^N; see [67]. Every axis of a power of g is, in fact, also an axis of g (this property is specific to the action $\mathcal{A} \curvearrowright \mathcal{X}$). Thus, picking any $x \in \alpha$, we have $g^i x \in Y$ for all $i \in \mathbb{Z}$. Hence there exist $0 \leq i < j \leq q$ such that $g^i x$ and $g^j x$ are in the same G–orbit. Since $\mathcal{A}$ acts freely on $\mathcal{X}$, we have $g^{j-i} \in G$ and $0 < j - i \leq q$.

3.4 CMP automorphisms of right-angled groups

This subsection is devoted to the proof of Proposition A. Automorphisms of hyperbolic groups were already discussed in Example 2.28, so we are only concerned with right-angled Artin/Coxeter groups.

Let Γ be a finite simplicial graph. Let $\mathcal{A} = \mathcal{A}_\Gamma$ and $\mathcal{W} = \mathcal{W}_\Gamma$ be, respectively, the right-angled Artin group and the right-angled Coxeter group defined by Γ.

We identify with $\Gamma^{(0)}$ the standard generating sets of $\mathcal{A}$ and $\mathcal{W}$. The standard Cayley graphs of $\mathcal{A}$ and $\mathcal{W}$ are 1–skeleta of CAT(0) cube complexes: the universal covers of the Salvetti and Davis complex, respectively. Thus, $\mathcal{A}$ and $\mathcal{W}$ are each endowed with a natural median operator μ_Γ.

Remark 3.18 We have $g \cdot \mu_\Gamma(x,y,z) = \mu_\Gamma(gx,gy,gz)$ for all elements g,x,y,z in $\mathcal{A}$ or $\mathcal{W}$. This implies that $(\mathcal{A},[\mu_\Gamma])$ and $(\mathcal{W},[\mu_\Gamma])$ are coarse median groups, in the sense of Definition 2.24.

Unlike hyperbolic groups, $\mathcal{A}$ and $\mathcal{W}$ can admit infinitely many different coarse median structures. For this reason, we will never omit the subscript in μ_Γ, in order to emphasise that this is the coarse median structure provided by the *standard* generating set of $\mathcal{A}$ or $\mathcal{W}$. Other Artin/Coxeter generating sets can a priori give different coarse median structures; it will be a consequence of Proposition A(2) that this does not actually happen in the Coxeter case.

It was shown by Laurence [75], Servatius [94] and Corredor and Gutierrez [35] that $\operatorname{Aut}\mathcal{A}$ and $\operatorname{Aut}\mathcal{W}$ are generated by finitely many *elementary automorphisms*. These take the same form in both cases.

- **Graph automorphisms** Every automorphism of the graph Γ gives a permutation of the standard generating sets that defines an automorphism of $\mathcal{A}$ and $\mathcal{W}$.
- **Inversions** ι_v for each $v \in \Gamma^{(0)}$. We have $\iota_v(v) = v^{-1}$ and $\iota_v(u) = u$ for all $u \in \Gamma^{(0)} \setminus \{v\}$.
- **Partial conjugations** $\kappa_{w,C}$ for $w \in \Gamma^{(0)}$ and a connected component C of $\Gamma \setminus \operatorname{st} w$. We have $\kappa_{w,C}(u) = w^{-1}uw$ if $u \in C^{(0)}$ and $\kappa_{w,C}(u) = u$ if $u \in \Gamma^{(0)} \setminus C$.
- **Transvections** $\tau_{v,w}$ for $v,w \in \Gamma^{(0)}$ with $\operatorname{lk} v \subseteq \operatorname{st} w$. They are defined by $\tau_{v,w}(v) = vw$ and $\tau_{v,w}(u) = u$ for all $u \in \Gamma^{(0)} \setminus \{v\}$.
  We refer to $\tau_{v,w}$ as a *fold* if v and w are not joined by an edge (equivalently, $\operatorname{lk} v \subseteq \operatorname{lk} w$), and as a *twist* if v and w are joined by an edge (equivalently, $\operatorname{st} v \subseteq \operatorname{st} w$).

Remark 3.19 Graph automorphisms and inversions can be realised as automorphisms of the Salvetti/Davis complex, so they preserve the operator μ_Γ (hence the coarse median structure $[\mu_\Gamma]$).

In the case of right-angled Artin groups, the following class of automorphisms was introduced by Day [41] and Charney, Stambaugh and Vogtmann [25].

Definition 3.20 An automorphism $\varphi \in \operatorname{Aut}\mathcal{A}$ is *untwisted* if it lies in the subgroup $U(\mathcal{A}) \leq \operatorname{Aut}\mathcal{A}$ generated by graph automorphisms, inversions, partial conjugations and folds.

We now proceed to prove parts (2) and (3) of Proposition A. We will treat separately the Coxeter and Artin cases, as the simplest arguments appear to be quite different in spirit. Still, in both situations, the following basic observation is important.

Remark 3.21 Let a group G act properly and cocompactly on two CAT(0) cube complexes X and Y. Let $[\mu_X]$ and $[\mu_Y]$ be the induced coarse median structures on G. If there exists an equivariant restriction-quotient map $\pi: X \to Y$, then $[\mu_X] = [\mu_Y]$. This is immediate from the third characterisation of restriction quotients in Proposition 2.20.

3.4.1 The Coxeter case Here our aim is to prove that all elements of $\mathrm{Aut}\,\mathcal{W}$ preserve the coarse median structure $[\mu_\Gamma]$. We will achieve this by showing that all elementary automorphisms of $\mathcal{W}$ restrict to graph automorphisms on certain finite-index Coxeter subgroups of $\mathcal{W}$. This guarantees that they are all coarse-median preserving.

Given a vertex $w \in \Gamma$, we denote by $\Delta(\Gamma, w)$ the *double* of $\Gamma \setminus \{w\}$ along $\mathrm{lk}\, w$. More precisely, $\Delta(\Gamma, w)$ is obtained from two disjoint copies of the graph $\Gamma \setminus \{w\}$ by identifying the two subgraphs corresponding to $\mathrm{lk}\, w$. We continue to denote by $\mathrm{lk}\, w$ the resulting subgraph of $\Delta(\Gamma, w)$, even though w does not appear in $\Delta(\Gamma, w)$ and so this is not the link of any vertex of $\Delta(\Gamma, w)$.

Let $\alpha_w : \mathcal{W} \to \mathbb{Z}/2\mathbb{Z}$ be the homomorphism that maps w to the generator of $\mathbb{Z}/2\mathbb{Z}$, and all other standard generators of $\mathcal{W}$ to the identity.

Lemma 3.22 *Consider a vertex $w \in \Gamma$. Then*:

(1) $\ker \alpha_w$ *is generated by* $\{x \mid x \in \Gamma \setminus \{w\}\} \sqcup \{w^{-1} y w \mid y \in \Gamma \setminus \mathrm{st}\, w\}$.

(2) *This is a Coxeter generating set giving an isomorphism between* $\ker \alpha_w$ *and* $\mathcal{W}_{\Delta(\Gamma,w)}$.

(3) *The coarse median structure* $[\mu_{\Delta(\Gamma,w)}]$ *induced on* $\ker \alpha_w$ *by this isomorphism with* $\mathcal{W}_{\Delta(\Gamma,w)}$ *coincides with the restriction of the coarse median structure* $[\mu_\Gamma]$ *on* $\mathcal{W}$.

Proof The first two parts are a straightforward application of the normal form for words in Coxeter groups; see eg [39, Chapter 3.4]. We instead focus on part (3).

Let $\mathcal{Y}_\Gamma$ and $\mathcal{Y}_{\Delta(\Gamma,w)}$ be the universal covers of the Davis complexes of $\mathcal{W}$ and $\mathcal{W}_{\Delta(\Gamma,w)}$. We aim to show that, under the above identification between $\ker \alpha_w$ and $\mathcal{W}_{\Delta(\Gamma,w)}$, the standard action $\mathcal{W}_{\Delta(\Gamma,w)} \curvearrowright \mathcal{Y}_{\Delta(\Gamma,w)}$ is a restriction quotient of the action $\ker \alpha_w \curvearrowright \mathcal{Y}_\Gamma$. This proves the lemma, since, by Remark 3.21, the two actions then induce the same coarse median structure on $\ker \alpha_w$.

First, if $\Omega \subseteq \mathcal{Y}_\Gamma$ is a fundamental domain for the $\mathcal{W}$–action, note that $\Omega \cup w\Omega$ is a fundamental domain for $\ker \alpha_w \curvearrowright \mathcal{Y}_\Gamma$. In addition, observe that $\ker \alpha_w$ contains the entire $\mathcal{W}$–stabiliser of a hyperplane $\mathfrak{u} \in \mathscr{W}(\mathcal{Y}_\Gamma)$ precisely when $\gamma(\mathfrak{u}) \notin \mathrm{st}\, w$. Thus, the orbit $\mathcal{W} \cdot \mathfrak{u}$ is made up of two $(\ker \alpha_w)$–orbits of hyperplanes when $\gamma(\mathfrak{u}) \notin \mathrm{st}\, w$, while it is a single $(\ker \alpha_w)$–orbit when $\gamma(\mathfrak{u}) \in \mathrm{st}\, w$. Combining these two observations, the reader should convince themselves that, starting with the action $\ker \alpha_w \curvearrowright \mathcal{Y}_\Gamma$ and collapsing the single orbit of hyperplanes $\mathfrak{u}$ with $\gamma(\mathfrak{u}) = w$, we obtain precisely the action $\mathcal{W}_{\Delta(\alpha,w)} \curvearrowright \mathcal{Y}_{\Delta(\alpha,w)}$, as required. □

Proposition 3.23 *All automorphisms of $\mathcal{W}$ preserve the coarse median structure* $[\mu_\Gamma]$.

Proof Recall that $\mathrm{Aut}\,\mathcal{W}$ is generated by graph automorphisms, partial conjugations and transvections, as defined above. We have already noticed in Remark 3.19 that graph automorphisms are coarse-median preserving. We make two additional observations.

- *Every partial conjugation* $\kappa_{w,C}$ *preserves the subgroup* $\ker\alpha_w \leq \mathcal{W}$. *The restriction of* $\kappa_{w,C}$ *to* $\ker\alpha_w$ *is a graph automorphism with respect to the identification* $\ker\alpha_w \simeq \mathcal{W}_{\Delta(\Gamma,w)}$ *constructed in Lemma 3.22.*

Indeed, the connected component $C \subseteq \Gamma \setminus \mathrm{st}\, w$ gets doubled to two connected components C', C'' of the graph $\Delta(\Gamma, w) \setminus \mathrm{lk}\, w$. These two subgraphs correspond to the sets of generators $C^{(0)}$ and $w^{-1}C^{(0)}w$ for $\ker\alpha_w$. The automorphism $\kappa_{w,C}$ swaps these two sets of generators, while fixing all generators of $\ker\alpha_w$ corresponding to vertices of $\Delta(\Gamma, w) \setminus (C' \cup C'')$. This is realised by an automorphism of the graph $\Delta(\Gamma, w)$.

- *Every transvection* $\tau_{v,w}$ *preserves* $\ker\alpha_v \leq \mathcal{W}$. *The restriction of* $\tau_{v,w}$ *to* $\ker\alpha_v$ *is a product of partial conjugations with respect to the identification* $\ker\alpha_v \simeq \mathcal{W}_{\Delta(\Gamma,v)}$.

Indeed, if $x \in \Gamma \setminus \{v\}$, we have $\tau_{v,w}(x) = x$ and $\tau_{v,w}(v^{-1}xv) = w^{-1}(v^{-1}xv)w$. Setting $A := \Gamma \setminus \mathrm{st}\, v \subseteq \Gamma$, we have $\Delta(\Gamma, v) = \mathrm{lk}\, v \sqcup A' \sqcup A''$, where the subgraphs A', A'' correspond, respectively, to the subsets $A^{(0)}, v^{-1}A^{(0)}v \subseteq \mathcal{W}$. Let $w' \in A'$ be the vertex originating from $w \in \Gamma$. The set A'' is a union of connected components of $\Delta(\Gamma, v) \setminus \mathrm{lk}\, v$. Since $\mathrm{lk}\, v \subseteq \mathrm{st}\, w$ in Γ, we have $\mathrm{lk}\, v \subseteq \mathrm{st}\, w'$ in $\Delta(\Gamma, v)$, hence A'' is also a union of connected components of $\Delta(\Gamma, v) \setminus \mathrm{st}\, w'$. The composition of the partial conjugations $\kappa_{w',C} \in \mathrm{Aut}\, \mathcal{W}_{\Delta(\Gamma,v)}$, as C ranges through these connected components, is precisely the restriction of $\tau_{v,w}$ to $\ker\alpha_v$.

Now, in view of Lemma 3.22 and Remark 3.19, partial conjugations and transvections each preserve the restriction of the coarse median structure $[\mu_\Gamma]$ to a finite-index subgroup of $\mathcal{W}$. Since finite-index subgroups are coarsely dense in $\mathcal{W}$, this implies that these automorphisms actually preserve $[\mu_\Gamma]$ itself, proving the proposition. □

3.4.2 The Artin case We begin by showing that *untwisted* automorphisms of $\mathcal{A}$ preserve the coarse median structure $[\mu_\Gamma]$. I present a proof that was suggested to me by Ric Wade, as it is much simpler than my original brute-force argument.

The main ingredient is the construction of Salvetti blowups from the work of Charney, Stambaugh and Vogtmann [25], which we record in the following lemma. Restriction quotients were discussed in Section 2.5.

Lemma 3.24 *Let* $\varphi \in U(\mathcal{A}_\Gamma)$ *be a fold or partial conjugation. Then there exists a proper cocompact action on a* CAT(0) *cube complex* $\mathcal{A}_\Gamma \curvearrowright Z$ *and two restriction-quotient maps* $\pi_1, \pi_2 \colon Z \to \mathcal{X}_\Gamma$ *such that, for all* $g \in \mathcal{A}_\Gamma$, *we have* $\pi_1 \circ g = g \circ \pi_1$ *and* $\pi_2 \circ g = \varphi(g) \circ \pi_2$.

Proof This holds more generally when φ is a Γ–*Whitehead automorphism*, as defined at the beginning of [25, Section 2.3]. Our statement is a straightforward rephrasing of [25, Lemma 3.2] in terms of universal covers. □

Corollary 3.25 *Automorphisms in* $U(\mathcal{A})$ *preserve the coarse median structure* $[\mu_\Gamma]$.

Proof Let $\varphi \in U(\mathcal{A}_\Gamma)$ be a fold or partial conjugation. Let the action $\mathcal{A}_\Gamma \curvearrowright Z$ and maps $\pi_1, \pi_2 : Z \to \mathcal{X}_\Gamma$ be as provided by Lemma 3.24, and let $[\mu_Z]$ be the coarse median structure on $\mathcal{A}_\Gamma$ induced by Z. Since π_1 is $\mathcal{A}_\Gamma$–equivariant, Remark 3.21 guarantees that $[\mu_\Gamma] = [\mu_Z]$.

On the other hand, π_2 becomes $\mathcal{A}_\Gamma$–equivariant if we endow $\mathcal{X}_\Gamma$ with the φ–twisted action: using the notation from Remark 2.27, this corresponds to replacing $\mathcal{X}_\Gamma$ with $\mathcal{X}_\Gamma^{\varphi^{-1}}$, which induces the coarse median structure $(\varphi^{-1})_*[\mu_\Gamma]$ on $\mathcal{A}_\Gamma$. Thus, another application of Remark 3.21 yields $(\varphi^{-1})_*[\mu_\Gamma] = [\mu_Z]$. We conclude that $\varphi_*[\mu_\Gamma] = [\mu_\Gamma]$.

This shows that all folds and partial conjugations preserve $[\mu_\Gamma]$. Graph automorphisms and inversions are also coarse-median preserving, by Remark 3.19. Since these four types of elementary automorphisms generate $U(\mathcal{A})$, this proves the corollary. □

In order to complete the proof of Proposition A(3), we are left to show that all coarse-median preserving automorphisms of $\mathcal{A}$ are untwisted. This can be easily deduced from the work of Laurence [75], as we now describe.

Proposition 3.26 *If $\varphi \in \operatorname{Aut} \mathcal{A}$ preserves the coarse median structure $[\mu_\Gamma]$, then $\varphi \in U(\mathcal{A})$.*

Proof In the terminology of [75, Section 2], an automorphism $\varphi \in \operatorname{Aut} \mathcal{A}$ is *conjugating* if it preserves the conjugacy class of each standard generator $v \in \Gamma$. More generally, φ is *simple* if, for every $v \in \Gamma$, the image $\varphi(v)$ is label-irreducible and $v \in \Gamma(\varphi(v))$; compare [75, Definition 5.3] and Definition 3.5 in our paper.

Consider a coarse-median preserving automorphism $\varphi = \varphi_0$. By [75, Corollary to Lemma 4.5], there exists a graph automorphism ψ_1 such that, setting $\varphi_1 := \psi_1\varphi$, we have $v \in \Gamma(\varphi_1(v))$ for every generator $v \in \Gamma$. Since graph automorphisms are coarse-median preserving, φ_1 is again coarse-median preserving. By Corollary 3.3 and Lemma 3.11(2), the element $\varphi_1(v)$ is label-irreducible for every $v \in \Gamma$. Thus, φ_1 is simple.

By the proofs of [75, Lemma 6.8] and [75, Corollary to Lemma 6.6], there exists a product of inversions, folds and partial conjugations ψ_2 such that $\varphi_2 := \varphi_1\psi_2$ is conjugating. Finally, by [75, Theorem 2.2], the automorphism φ_2 is a product of partial conjugations. This shows that $\varphi \in U(\mathcal{A})$, as required. □

3.4.3 Pure automorphisms We end this subsection by introducing the subgroups $U_0(\mathcal{A}) \leq U(\mathcal{A})$ and $\operatorname{Aut}_0 \mathcal{W} \leq \operatorname{Aut} \mathcal{W}$ generated by inversions, folds and partial conjugations (no graph automorphisms or twists, in both cases). These are the subgroups appearing in the statements of Theorem C and Proposition D, and we will study them further in Sections 4.5 and 5. For the time being, we limit ourselves to a few quick observations.

Remark 3.27 The subgroups $U_0(\mathcal{A}) \leq U(\mathcal{A})$ and $\mathrm{Aut}_0 \mathcal{W} \leq \mathrm{Aut}\, \mathcal{W}$ have finite index. In the Coxeter case, see eg [92, Proposition 1.2]. In the Artin case, it suffices to observe that $U_0(\mathcal{A})$ is normalised by all graph automorphisms, and that the latter generate a finite subgroup of $U(\mathcal{A})$.

Remark 3.28 Although they do not appear in our chosen generating set for $U_0(\mathcal{A})$, graph automorphisms of $\mathcal{A}$ can still lie in $U_0(\mathcal{A})$. Indeed, confusing $\sigma \in \mathrm{Aut}\, \Gamma$ with the induced $\sigma \in \mathrm{Aut}\, \mathcal{A}$, we have $\sigma \in U_0(\mathcal{A})$ if and only if $\mathrm{lk}\, \sigma(v) = \mathrm{lk}\, v$ for every $v \in \Gamma$.

The "only if" part follows from Lemma 4.30. For the "if" part, it suffices to show that $\sigma \in U_0(\mathcal{A})$ whenever σ swaps two vertices of Γ with the same link and fixes the rest of Γ. In this case, σ is a product of 3 folds and 3 inversions, as described at the end of the proof of [43, Proposition 3.3].

Lemma 3.29 *If $\varphi(\mathcal{A}_\Delta) = \mathcal{A}_\Delta$ for a full subgraph $\Delta \subseteq \Gamma$ and $\varphi \in U_0(\mathcal{A})$, then $\varphi|_{\mathcal{A}_\Delta} \in U_0(\mathcal{A}_\Delta)$.*

Proof We begin with a general observation. As in the proof of Proposition 3.26, we can write $\varphi = \sigma\varphi_1$, where σ is a graph automorphism and φ_1 is a simple automorphism of $\mathcal{A}$. Moreover, simple automorphisms are products of inversions, folds and partial conjugations, so $\varphi_1 \in U_0(\mathcal{A})$. We conclude that $\sigma \in U_0(\mathcal{A})$, and Remark 3.28 shows that $\mathrm{lk}\, \sigma(v) = \mathrm{lk}\, v$ for every $v \in \Gamma$.

If $v \in \Delta$, then $v \in \Gamma(\varphi_1(v))$ because φ_1 is simple. Thus

$$\sigma(v) \in \sigma(\Gamma(\varphi_1(v))) = \Gamma(\sigma\varphi_1(v)) = \Gamma(\varphi(v)) \subseteq \Delta.$$

We deduce that $\sigma(\Delta) = \Delta$, and Remark 3.28 shows that $\sigma|_{\mathcal{A}_\Delta} \in U_0(\mathcal{A}_\Delta)$. Since σ and φ preserve $\mathcal{A}_\Delta$, so does φ_1, and it suffices to show that $\varphi_1|_{\mathcal{A}_\Delta} \in U_0(\mathcal{A}_\Delta)$.

It is clear that $\varphi_1|_{\mathcal{A}_\Delta}$ is a simple automorphism of $\mathcal{A}_\Delta$, so the fact that $\varphi_1|_{\mathcal{A}_\Delta} \in U_0(\mathcal{A}_\Delta)$ follows again from [75] as in the proof of Proposition 3.26. $\square$

4 Fixed subgroups of CMP automorphisms

This section is devoted to fixed subgroups of coarse-median preserving automorphisms of cocompactly cubulated groups. Theorem B is proved in Sections 4.1 and 4.2, where we study the properties of those subgroups of cocompactly cubulated groups that are approximate median subalgebras; see Theorem 4.10. At the end of Section 4.2, we also prove Corollaries G and H.

Then in Sections 4.3 and 4.4, we develop a quasiconvexity criterion for approximate median subalgebras of cube complexes (Proposition 4.25). This is used to prove Theorem C in Section 4.5; see Corollaries 4.34 and 4.35.

The reader interested only in Theorems F and I can just read the proof that $\mathrm{Fix}\, \varphi$ is finitely generated (Proposition 4.11) and skip the rest of this section in its entirety.

4.1 Approximate median subalgebras

The goal of this subsection is to show that approximate median subalgebras (Definition 2.33) of median spaces stay close to actual subalgebras. This is an important ingredient in the proofs of Theorem B and Corollaries G and H, which will be discussed in the next subsection.

Shortly after the first draft of this paper appeared on arXiv, it was pointed out to me by Mark Hagen that a similar result appears in the work of Bowditch [18, Proposition 4.1], which I was not aware of. Although Propositions 4.1 and 4.2 below are more general and our proofs seem different, I want to emphasise that Bowditch's result would suffice for all applications in this paper.

Proposition 4.1 *If X is a finite-rank median space and $A \subseteq X$ is an approximate median subalgebra, then $d_{\mathrm{Haus}}(A, \langle A\rangle) < +\infty$.*

The only focus of this subsection will actually be the next result, which provides an analogue of Remark 2.5. From it, it is straightforward to deduce Proposition 4.1 proceeding as in Lemma 2.10, which we leave to the reader.

Proposition 4.2 *There exists a function $h\colon \mathbb{N} \to \mathbb{N}$ with the following property. If M is a median algebra of rank r and $A \subseteq M$ is a subset, then $\langle A\rangle = \mathcal{M}^{h(r)}(A)$.*

We now obtain a sequence of lemmas leading up to Proposition 4.8, which proves Proposition 4.2.

Let M be a median algebra. We denote by $\mathscr{M}(M)$ the collection of subsets of M of one of these three forms:

- $\mathfrak{h}$, where $\mathfrak{h}$ is a halfspace;
- $\mathfrak{h} \cup \mathfrak{k}$, where $\mathfrak{h}$ and $\mathfrak{k}$ are transverse halfspaces;
- $\mathfrak{h} \cup \mathfrak{k}$, where $\mathfrak{h}$ and $\mathfrak{k}$ are disjoint halfspaces.

Elements of $\mathscr{M}(M)$ are to median subalgebras what halfspaces of M are to convex subsets. More precisely, the following is a straightforward characterisation of the median subalgebra generated by a subset $A \subseteq M$; see for instance [99, II.4.25.7].

Lemma 4.3 *For every subset $A \subseteq M$, the median subalgebra $\langle A\rangle \subseteq M$ is the intersection of all elements of $\mathscr{M}(M)$ containing A.*

We will make repeated use of the following observation, without explicit mention:

Lemma 4.4 *Given points $a, b, c, d \in M$, the three sets $\mathscr{W}(a, b|c, d)$, $\mathscr{W}(a, c|b, d)$ and $\mathscr{W}(a, d|b, c)$ are transverse to each other.*

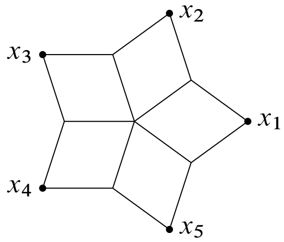

Figure 2: A pentagonal configuration in the 0–skeleton of a CAT(0) square complex.

It is also convenient to give a name to the situation in Figure 2.

Definition 4.5 An ordered 5–tuple $(x_1, x_2, x_3, x_4, x_5) \in M^5$ is a *pentagonal configuration* if the five sets $\mathscr{W}(x_{i-1}, x_i, x_{i+1} | x_{i+2}, x_{i+3})$ are all nonempty (indices are taken mod 5).

This requirement is invariant under cyclic permutations of the 5 points. Also note that, setting $\mathcal{W}_i := \mathscr{W}(x_{i-1}, x_i, x_{i+1} | x_{i+2}, x_{i+3})$, the sets $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ are transverse for all i mod 5.

Lemma 4.6 *Suppose that* $\operatorname{rk} M \leq 2$. *Consider* $x \in M$ *with* $x = m(m(a_1, a_2, a_3), b, c)$ *for points* $a_i, b, c \in M$. *Then one of the following happens*:

- *There exists* $1 \leq i \leq 3$ *such that* $x = m(a_i, b, c)$.
- *There exist* $1 \leq i < j \leq 3$ *such that either* $x = m(a_i, a_j, b)$ *or* $x = m(a_i, a_j, c)$.
- *We have* $x = m(a_1, a_2, a_3)$.
- *The points* a_1, a_2, a_3, b, c *can be ordered to form a pentagonal configuration.*

Proof Set $n = m(a_1, a_2, a_3)$. Consider the projections $\overline{a}_i = m(a_i, b, c)$ to the interval $I(b, c)$. Since gate-projections are median morphisms, we have $x = m(\overline{a}_1, \overline{a}_2, \overline{a}_3)$.

Claim 1 *If we are not in the* 1^{st} *or* 3^{rd} *case, we can assume that the four sets* $\mathscr{W}(x|\overline{a}_1)$, $\mathscr{W}(x|\overline{a}_2)$, $\mathscr{W}(x|\overline{a}_3)$ *and* $\mathscr{W}(a_1, a_2|b, c)$ *are all nonempty, and that* $\mathscr{W}(a_1, c|a_2, b) = \varnothing$.

Proof If one of the sets $\mathscr{W}(x|\overline{a}_i)$ is empty, then $x = \overline{a}_i$ and we are in the 1^{st} case. On the other hand, if the sets $\mathscr{W}(a_i, a_j|b, c)$ are all empty for $i \neq j$, then we are in the 3^{rd} case. Indeed, since $\mathscr{W}(n|b, c) \subseteq \bigcup_{i<j} \mathscr{W}(a_i, a_j|b, c)$, we have $n \in I(b, c)$, hence $x = m(n, b, c) = n = m(a_1, a_2, a_3)$.

Thus, up to permuting the a_i, we can assume that $\mathscr{W}(a_1, a_2|b, c) \neq \varnothing$. Since this is transverse to the transverse sets $\mathscr{W}(a_1, b|a_2, c)$ and $\mathscr{W}(a_1, c|a_2, b)$, one of the latter must be empty. Swapping b and c if necessary, we can assume that it is $\mathscr{W}(a_1, c|a_2, b)$. $\triangleleft$

Claim 2 *If we are not in the* 4^{th} *case either, we can further assume that* $\mathscr{W}(a_1, a_2, b|a_3, c) = \varnothing$.

Proof Note that the assumptions in Claim 1 are left unchanged if we simultaneously swap $b \leftrightarrow c$ and $a_1 \leftrightarrow a_2$. Thus, it suffices to show that we can suppose that at least one of the two sets $\mathscr{W}(b, a_1, a_2|c, a_3)$ and $\mathscr{W}(a_1, a_2, c|a_3, b)$ is empty.

In order to do so, we assume that $\mathscr{W}(b, a_1, a_2|c, a_3)$ and $\mathscr{W}(a_1, a_2, c|a_3, b)$ are both nonempty and show that (b, a_1, a_2, c, a_3) is a pentagonal configuration. This places us in the 4^{th} case.

Since $\mathscr{W}(a_1, c|a_2, b) = \varnothing$ and $x = m(\overline{a}_1, \overline{a}_2, \overline{a}_3)$, where $\overline{a}_i$ is the projection of a_i to $I(b, c)$, we have

$$\mathscr{W}(a_2, c, a_3|b, a_1) \supseteq \mathscr{W}(\overline{a}_2, \overline{a}_3|\overline{a}_1) = \mathscr{W}(x|\overline{a}_1) \neq \varnothing,$$
$$\mathscr{W}(a_3, b, a_1|a_2, c) \supseteq \mathscr{W}(\overline{a}_3, \overline{a}_1|\overline{a}_2) = \mathscr{W}(x|\overline{a}_2) \neq \varnothing.$$

Moreover, since $\mathscr{W}(a_1, a_3|b, c)$ is transverse to the nonempty transverse subsets $\mathscr{W}(b, a_1, a_2|c, a_3)$ and $\mathscr{W}(a_1, a_2, c|a_3, b)$, we have $\mathscr{W}(a_1, a_3|b, c) = \varnothing$. Hence $\mathscr{W}(c, a_3, b|a_1, a_2) = \mathscr{W}(c, b|a_1, a_2) \neq \varnothing$. $\lhd$

Claim 3 *Under these assumptions, we have* $\mathscr{W}(x|m(a_1, a_3, c)) = \mathscr{W}(b, c|a_1, a_3)$.

Proof By the properties of gate-projections, the set $\mathscr{W}(b|c)$ does not intersect any of the sets $\mathscr{W}(a_i|\overline{a}_i)$. Thus, since $x = m(\overline{a}_1, \overline{a}_2, \overline{a}_3)$, we must have

$$\begin{aligned}\mathscr{W}(x|m(a_1, a_3, c)) \cap \mathscr{W}(b|c) &= \mathscr{W}(m(a_1, a_2, a_3)|m(a_1, a_3, c)) \cap \mathscr{W}(b|c)\\ &= \mathscr{W}(a_1|a_3) \cap \mathscr{W}(a_2|c) \cap \mathscr{W}(b|c)\\ &= \mathscr{W}(a_1, a_2, b|a_3, c) \sqcup \mathscr{W}(a_2, a_3, b|a_1, c) = \varnothing,\end{aligned}$$

where we have used Claims 1 and 2 at the very end. Since $x \in I(b, c)$, we have $\mathscr{W}(x|b, c) = \varnothing$. Thus,

$$\mathscr{W}(x|m(a_1, a_3, c)) = \mathscr{W}(x, b, c|m(a_1, a_3, c)) = \mathscr{W}(b, c|a_1, a_3). \qquad \lhd$$

In order to conclude the proof of the lemma, suppose for the sake of contradiction that we are not in the 2^{nd} case, in addition to the assumptions of the claims. Then, Claim 3 implies

$$\varnothing \neq \mathscr{W}(x|m(a_1, a_3, c)) = \mathscr{W}(b, c|a_1, a_3).$$

On the other hand, since $\mathscr{W}(a_1, c|a_2, b)$ and $\mathscr{W}(a_1, a_2, b|a_3, c)$ are empty by Claims 1 and 2,

$$\varnothing \neq \mathscr{W}(x|\overline{a}_1) = \mathscr{W}(\overline{a}_2, \overline{a}_3|\overline{a}_1) = \mathscr{W}(c, a_2, a_3|b, a_1) \subseteq \mathscr{W}(c, a_3|b, a_1),$$
$$\varnothing \neq \mathscr{W}(x|\overline{a}_3) = \mathscr{W}(\overline{a}_1, \overline{a}_2|\overline{a}_3) = \mathscr{W}(a_1, a_2, c|a_3, b) \subseteq \mathscr{W}(a_1, c|a_3, b).$$

Since the three sets $\mathscr{W}(b, c|a_1, a_3)$, $\mathscr{W}(c, a_3|b, a_1)$, $\mathscr{W}(a_1, c|a_3, b)$ are pairwise transverse, this violates the assumption that $\operatorname{rk} M \leq 2$. This proves the lemma. $\square$

Corollary 4.7 *If T_1 and T_2 are rank-1 median algebras, then $\langle A\rangle = \mathcal{M}(A)$ for all $A \subseteq T_1 \times T_2$.*

Proof The product $T_1 \times T_2$ does not contain any pentagonal configurations. Otherwise, there would be walls $\mathfrak{w}_1, \mathfrak{w}_2, \mathfrak{w}_3, \mathfrak{w}_4, \mathfrak{w}_5$ with each $\mathfrak{w}_i$ transverse to $\mathfrak{w}_{i+1}$. If $\mathfrak{w}_1$ originates from the factor T_1, say, then $\mathfrak{w}_2$ must originate from T_2 and, continuing like this, $\mathfrak{w}_5$ again originates from T_1. Since $\mathfrak{w}_5$ and $\mathfrak{w}_1$ are transverse, this would contradict the fact that $\operatorname{rk} T_1 = 1$.

Thus, the 4th case of Lemma 4.6 never occurs, hence $\mathcal{M}^2(A) = \mathcal{M}(A)$ for all $A \subseteq T_1 \times T_2$. □

For the next result, let us consider the functions $f, g, h\colon \mathbb{N} \to \mathbb{N}$ given by

$$f(n) = 2^{2^n}, \quad g(n) = 1 + f\left(\tfrac{1}{2}n(n-1)\right), \quad h(n) = ng(n) + n.$$

Proposition 4.8 *Given a median algebra M and a subset $A \subseteq M$, the following hold*:

(1) *If $\#A \le n$, then $\langle A\rangle = \mathcal{M}^{f(n)}(A)$.*

(2) *If M can be embedded in a product of d rank-1 median algebras, then $\langle A\rangle = \mathcal{M}^{g(d)}(A)$.*

(3) *If $\operatorname{rk} M \le r$, then $\langle A\rangle = \mathcal{M}^{h(r)}(A)$.*

Proof Part (1) is immediate from most constructions of the free median algebra on the set A; for instance, see [15, Lemma 4.2] and the subsequent paragraph.

Regarding part (2), let us fix an injective median morphism $M \hookrightarrow T_1 \times \dots \times T_d$, where the T_i have rank 1. Let $\pi_{ij}\colon M \to T_i \times T_j$ denote the composition with the projection to $T_i \times T_j$. Given $x \in M$, Lemma 4.3 implies that $x \in \langle A\rangle$ if and only if $\pi_{ij}(x) \in \langle \pi_{ij}(A)\rangle$ for all $1 \le i < j \le d$.

Since each π_{ij} is a median morphism, Corollary 4.7 shows that

$$\langle \pi_{ij}(A)\rangle = \mathcal{M}(\pi_{ij}(A)) = \pi_{ij}(\mathcal{M}(A)).$$

Thus, given $x \in \langle A\rangle$, there exist points $m_{ij} \in \mathcal{M}(A)$ such that $\pi_{ij}(x) = \pi_{ij}(m_{ij})$. It follows that

$$x \in \langle \{m_{ij} \mid 1 \le i < j \le d\}\rangle.$$

Since there are at most $\frac{1}{2}d(d-1)$ distinct points m_{ij}, part (1) yields

$$\langle \{m_{ij} \mid 1 \le i < j \le d\}\rangle = \mathcal{M}^{g(d)-1}(\{m_{ij} \mid 1 \le i < j \le d\}) \subseteq \mathcal{M}^{g(d)-1}(\mathcal{M}(A)) = \mathcal{M}^{g(d)}(A).$$

Hence $\langle A\rangle \subseteq \mathcal{M}^{g(d)}(A)$.

Finally, let us prove part (3). Since $\operatorname{rk}\langle A\rangle \le \operatorname{rk} M$, we can safely assume that $M = \langle A\rangle$. Consider two points $a, b \in M$ and recall that the gate-projection $\pi_{ab}\colon M \to I(a,b)$ is given by $\pi_{ab}(x) = m(a,b,x)$. By Dilworth's lemma, the interval $I(a,b) \subseteq M$ can be embedded in a product of r rank-1 median algebras for all $a, b \in M$; cf [16, Proposition 1.4].

If $B \subseteq M$ is a subset with $\langle B\rangle = M$ and $a, b \in B$, then part (2) yields

$$I(a,b) = \pi_{ab}(M) = \pi_{ab}(\langle B\rangle) = \langle \pi_{ab}(B)\rangle = \mathcal{M}^{g(r)}(\pi_{ab}(B)) = \pi_{ab}(\mathcal{M}^{g(r)}(B)) \subseteq \mathcal{M}^{g(r)+1}(B).$$

It follows that $\mathcal{J}(B) \subseteq \mathcal{M}^{g(r)+1}(B)$ for every subset $B \subseteq M$ with $\langle B\rangle = M$. Observing that

$$\mathcal{J}^{k+1}(B) = \mathcal{J}(\mathcal{J}^k(B)) \subseteq \mathcal{M}^{g(r)+1}(\mathcal{J}^k(B)),$$

we inductively obtain $\mathcal{J}^m(B) \subseteq \mathcal{M}^{m(g(r)+1)}(B)$ for all $m \geq 1$. In particular, by Remark 2.5,

$$\langle A\rangle \subseteq \operatorname{Hull} A = \mathcal{J}^r(A) \subseteq \mathcal{M}^{r(g(r)+1)}(A) = \mathcal{M}^{h(r)}(A).$$

This concludes the proof of the proposition. $\square$

Remark 4.9 The bounds provided by Proposition 4.8 are highly nonsharp. For instance, if $\operatorname{rk} M \leq 2$, a slightly more careful use of Lemma 4.6 would show that $\langle A\rangle = \mathcal{M}^2(A)$ for every $A \subseteq M$, while the proposition only gives $\langle A\rangle = \mathcal{M}^{244}(A)$. For the purposes of this paper, we only care that such bounds exist and only depend on the rank of M.

4.2 Approximate subalgebras of cubulated groups

This subsection is devoted to the proof of Theorem B and to a few examples of how this result can fail for automorphisms that do not preserve the coarse median structure (Example 4.13). Towards the end, we use similar techniques to prove Corollaries G and H.

Our main focus will be the following result. Recall that, if φ is a coarse-median preserving automorphism of a cocompactly cubulated group, Lemma 2.35 guarantees that the subgroup $\operatorname{Fix}\varphi$ is an approximate median subalgebra. Thus, Theorem B immediately follows from:

Theorem 4.10 *Let $G \curvearrowright X$ be a proper cocompact action on a* CAT(0) *cube complex. Let $[\mu_X]$ be the induced coarse median structure on G. If a subgroup $H \leq G$ is an approximate median subalgebra of $(G, [\mu_X])$, then*:

(1) *H is finitely generated and undistorted in G.*

(2) *H admits a proper cocompact action on a* CAT(0) *cube complex.*

As a first step, we need to show that the subgroup H in Theorem 4.10 is finitely generated. The proof of this is a straightforward adaptation of an argument due to Cooper and Paulin [33; 86] for fixed subgroups of automorphisms of hyperbolic groups.

Proposition 4.11 *Let $G \curvearrowright X$ be a proper cocompact action on a* CAT(0) *cube complex. If a subgroup $H \leq G$ is an approximate median subalgebra of $(G, [\mu_X])$, then H is finitely generated.*

Proof Fix a base vertex $p \in X$. The main observation is the following:

Claim *If $x_n \in H \cdot p$ is a diverging sequence, then there exists an element $h \in H$ such that $d(p, hx_n) < d(p, x_n)$ holds for infinitely many values of n.*

Proof Since H is an approximate median subalgebra of $(G, [\mu_X])$, there exists $L \geq 0$ such that all medians of points in $H \cdot p$ lie in the L–neighbourhood of $H \cdot p$ in X.

Passing to a subsequence, we can assume that the vertices x_n converge to a point in the Roller boundary $\xi \in \partial X$. Recalling that $X^{(0)}$ is discrete in the Roller compactification and that the median map is continuous, there exist integers $M(n) \geq 0$ such that, for every $m \geq M(n)$, we have

$$m(p, x_n, \xi) = m(p, x_n, x_m).$$

In particular, there exist elements $h_n \in H$ such that $h_n p$ is L–close to $m(p, x_n, \xi)$.

Now, since $x_n \to \xi$, the medians $m(p, x_n, \xi)$ diverge, and so do the points $h_n p$. In particular, there exist indices $i < j$ such that

$$d(p, h_i p) + 2L < d(p, h_j p).$$

Since, for $m \geq M(j)$, the point $h_j p$ is L–close to the median $m(p, x_j, x_m)$, we also have

$$d(p, h_j p) + d(h_j p, x_m) \leq d(p, x_m) + 2L.$$

In conclusion, setting $h := h_i h_j^{-1}$, we obtain, for all $m \geq M(j)$,

$$\begin{aligned} d(p, hx_m) = d(p, h_i h_j^{-1} x_m) &\leq d(p, h_i p) + d(h_i p, h_i h_j^{-1} x_m) = d(p, h_i p) + d(h_j p, x_m) \\ &\leq d(p, h_i p) + d(p, x_m) - d(p, h_j p) + 2L \\ &< d(p, x_m). \end{aligned}$$ ◁

Now, suppose for the sake of contradiction that H is not finitely generated. Write H as the union of an infinite ascending chain of subgroups $H_1 \lneq H_2 \lneq \cdots$, where $H_{n+1} = \langle H_n, h_{n+1} \rangle$ for some $h_{n+1} \in H$. Possibly replacing h_{n+1}, we can assume that the point $x_{n+1} := h_{n+1} p$ minimises the distance to p within the set $H_n h_{n+1} \cdot p$.

The claim provides an element $h \in H$ such that $d(p, hx_n) < d(p, x_n)$ occurs infinitely often. Since $h \in H$, there exists $N \geq 0$ such that $h \in H_n$ for all $n \geq N$. This contradicts the fact that x_{n+1} minimises the distance to p within $H_n \cdot x_{n+1}$ for $n \geq N$. □

Along with Propositions 4.1 and 4.11, the following is the only missing ingredient in the proof of Theorem 4.10. We refer the reader to the proof sketched in the introduction.

Lemma 4.12 *Let $G \curvearrowright X$ be a proper cocompact action on a* CAT(0) *cube complex. Consider a subgroup $H \leq G$. Suppose that there exists an H–invariant median subalgebra $M \subseteq X^{(0)}$ such that the action $H \curvearrowright M$ is cofinite. Then*:

(1) *H is finitely generated and undistorted in G.*

(2) *H admits a proper cocompact action on a* CAT(0) *cube complex.*

Proof Halfspaces and hyperplanes of the cube complex X, as usually defined, correspond exactly to halfspaces and hyperplanes of the median algebra $X^{(0)}$. As customary, we write $\mathscr{H}(X)$ and $\mathscr{W}(X)$ to mean $\mathscr{H}(X^{(0)})$ and $\mathscr{W}(X^{(0)})$. By Remark 2.2(1), we have a natural surjection $\mathrm{res}_M : \mathscr{H}_M(X) \to \mathscr{H}(M)$.

Since H acts cofinitely on the subalgebra M, it is an approximate median subalgebra of $(G, [\mu_X])$ and Proposition 4.11 implies that H is finitely generated. Thus, every H–orbit in X is coarsely connected and, since $H \curvearrowright M$ is cofinite, M is coarsely connected as well. It follows that there exists a uniform upper bound m to the cardinality of the fibres of the map res_M.

As in [89, Section 10; 31, Theorem 6.1], we can construct a CAT(0) cube complex $X(M)$ such that M is naturally isomorphic to the median algebra $X(M)^{(0)}$. Given $x, y \in M$, let us denote by $d(x, y)$ and $d_M(x, y)$ their distance in the 1–skeleta of X and $X(M)$, respectively.

By construction, $d_M(x, y)$ coincides with the number of walls of M separating x and y. It follows from the above discussion on res_M that

$$d_M(x, y) \le d(x, y) \le m \cdot d_M(x, y)$$

for all $x, y \in M$. Thus, the identification between $X(M)^{(0)}$ and $M \subseteq X^{(0)}$ gives a quasi-isometric embedding $X(M) \to X$ that is equivariant with respect to the inclusion $H \hookrightarrow G$.

The action $H \curvearrowright (M, d_M)$ is cofinite by assumption, and it follows from the above inequalities that it is also proper. This shows that the induced action $H \curvearrowright X(M)$ is proper and cocompact, proving part (2). The Milnor–Schwarz Lemma now implies that the inclusion $H \hookrightarrow G$ is a quasi-isometric embedding, which proves part (1). □

Proof of Theorem 4.10 For any vertex $p \in X$, the orbit $H \cdot p$ is an approximate median subalgebra of X. By Proposition 4.1, the subalgebra $M := \langle H \cdot p \rangle$ is at finite Hausdorff distance from $H \cdot p$. Since X is locally finite, it follows that the action $H \curvearrowright M$ is cofinite, hence Lemma 4.12 shows that H is finitely generated, undistorted and cocompactly cubulated. □

As discussed above, this completes the proof of Theorem B. The next example shows that, even for automorphisms of RAAGs, all of the claims in the statement of Theorem B can fail if the automorphism does not preserve the coarse median structure.

Example 4.13 Here is a recipe to construct automorphisms with unpleasant fixed subgroups. Consider a group G and a homomorphism $\alpha\colon G \to \mathbb{Z}$. These data define an automorphism $\psi \in \mathrm{Aut}(G \times \mathbb{Z})$ by the formula

$$\psi(g, n) := (g, n + \alpha(g)).$$

It is clear that $\mathrm{Fix}\, \psi = \ker \alpha \times \mathbb{Z}$.

Now, consider the situation where G is a right-angled Artin group $\mathcal{A}_\Gamma$ and $\alpha\colon \mathcal{A}_\Gamma \to \mathbb{Z}$ takes all standard generators to $+1$. The resulting automorphism $\psi \in \mathrm{Aut}(\mathcal{A}_\Gamma \times \mathbb{Z})$ is a product of finitely many twists (as defined in Section 3.4) and we have $\mathrm{Fix}\,\psi = BB_\Gamma \times \mathbb{Z}$, where BB_Γ denotes the Bestvina–Brady subgroup of $\mathcal{A}_\Gamma$; see [8].

We apply this construction to obtain examples where $\mathrm{Fix}\,\psi$ fails to have the three properties provided by Theorem B.

(1) The subgroup BB_Γ is finitely generated if and only if Γ is connected [80]. For instance, there exists $\psi \in \mathrm{Aut}(F_2 \times \mathbb{Z})$ such that $\mathrm{Fix}\,\psi$ is not finitely generated.

(2) If $\mathcal{A}_\Gamma$ is freely irreducible, directly irreducible and noncyclic, then BB_Γ is finitely generated and quadratically distorted [98, Theorem 1.1]. This gives examples where $\mathrm{Fix}\,\psi$ is finitely generated, but distorted.

(3) As shown in [8, Main Theorem], the finiteness properties and homological finiteness properties of BB_Γ are governed by the homology and homotopy groups of the flag simplicial complex L_Γ determined by Γ. The same is true of $\mathrm{Fix}\,\psi = BB_\Gamma \times \mathbb{Z}$; see [79; 22]. In particular, if L_Γ is not contractible, then $\mathrm{Fix}\,\psi$ is not of type F (hence not cocompactly cubulated, since there is no torsion). This can even be achieved while ensuring that $\mathrm{Fix}\,\varphi$ is undistorted: by [98, Theorem 1.1], it suffices to make sure that $\mathcal{A}_\Gamma$ splits as a product.

We emphasise that, by embedding $\mathcal{A}_\Gamma \times \mathbb{Z}$ as a parabolic subgroup of a larger RAAG and suitably extending the automorphism ψ, we can ensure that all these bad behaviours also occur for automorphisms of *irreducible* RAAGs.

We conclude this subsection by proving Corollaries G and H. All that is required is Proposition A, Proposition 4.1 and (part of) Lemma 4.12.

Proof of Corollary G Let $H \leq G$ be a finite-index subgroup with a proper cocompact action on a CAT(0) cube complex $H \curvearrowright X$. Replacing H with a finite-index subgroup, it is not restrictive to suppose that $H \lhd G$. By our assumption, the conjugation action $G \curvearrowright H$ preserves the coarse median structure $[\mu]$ induced on H by $H \curvearrowright X$.

It is well-known that the concept of induced representation can be generalised to actions on metric spaces; see eg [21, Section 2.1] or [9, Section 2.2]. In our context, this yields a proper action $G \curvearrowright X_1 \times \cdots \times X_n$, where n is the index of H in G and each X_i is isomorphic to X. Since H is normal, each factor is preserved by H and each action $H \curvearrowright X_i$ can be made equivariantly isomorphic to $H \curvearrowright X$ by twisting it by an automorphism of H corresponding to a conjugation by an element of G.

Since $[\mu]$ is preserved by the conjugation action $G \curvearrowright H$, it is the coarse median structure induced by all the cubulations $H \curvearrowright X_i$. This implies that, for every finite subset $A \subseteq X^{(0)}$, the orbit $H \cdot A$ is an approximate median subalgebra. Since H is normal, we can choose a finite subset $A \subseteq X^{(0)}$ so that $H \cdot A$ is G–invariant. Proposition 4.1 guarantees that the G–invariant median subalgebra $M := \langle H \cdot A \rangle$ is at finite Hausdorff distance from $H \cdot A$. Since each X_i is locally finite, this implies that the action $G \curvearrowright M$ is cofinite.

Since M is a discrete median algebra, the natural CAT(0) cube complex $X(M)$ with M as its 0–skeleton (as in [89, Section 10] or [31, Theorem 6.1]) gives the required cocompact cubulation of G. Here properness of the G–action on $X(M)$ can be checked as in the proof of Lemma 4.12, using the fact that M is coarsely connected to conclude that $X(M)$ and $\prod X_i$ induce bi-Lipschitz equivalent metrics on M. $\square$

Proof of Corollary H Let $G = \mathcal{A}_\Gamma$ or $G = \mathcal{W}_\Gamma$. Consider a finite subgroup $F \leq \operatorname{Out} G$ as in the statement. Let $\pi : \operatorname{Aut} G \to \operatorname{Out} G$ be the quotient projection. Our goal is to construct a proper, cocompact action on a CAT(0) cube complex $\pi^{-1}(F) \curvearrowright Y$. We can then take Q to be the quotient of Y by the finite-index normal subgroup $G \simeq \ker \pi \lhd \pi^{-1}(F)$.

Let $G \curvearrowright X$ be the standard action on the universal cover of the Salvetti/Davis complex. In both cases, Proposition A shows that F preserves the coarse median structure on G induced by this action. Thus, Corollary G provides the required proper cocompact action $\pi^{-1}(F) \curvearrowright Y$. $\square$

4.3 Staircases in cube complexes

In the rest of Section 4, our goal is to obtain a *quasiconvexity* criterion for median subalgebras of cube complexes, which will lead to the proof of Theorem C. Ultimately, we will restrict to universal covers of Davis/Salvetti complexes for right-angled groups and an important point will be that they do not admit infinite *staircases*.

In this subsection, we study staircases in general CAT(0) cube complexes.

Definition 4.14 Let M be a median algebra.

(1) A *length-n staircase* in M is the data of two chains of halfspaces $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_n$ and $\mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_n$ such that $\mathfrak{h}_i$ is transverse to $\mathfrak{k}_j$ for $j \leq i$, while $\mathfrak{k}_{i+1} \subsetneq \mathfrak{h}_i$.

(2) The *staircase length* of M is the supremum of $n \in \mathbb{N}$ such that M has a length-n staircase.

Figure 3 depicts part of a staircase of length ≥ 5.

When speaking of staircases in relation to a CAT(0) cube complex X, we always refer to the median algebra $M = X^{(0)}$. Note that the above notion of staircase seems to be a bit more general than the one in [63, page 51]: given hyperplanes bounding halfspaces as in Definition 4.14, there might not be a convex subcomplex of X with exactly these hyperplanes.

In view of the following discussion, it is convenient to introduce a notation for gate-projections to intervals. Given a median algebra M and points $x, y \in M$, we denote by $\pi_{xy} : M \to I(x, y)$ the map $\pi_{xy}(z) = m(x, y, z)$.

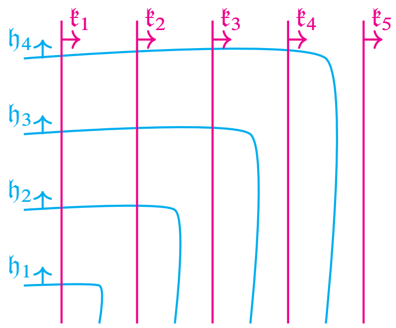

Figure 3

Lemma 4.15 *Let M be a median algebra of rank r and staircase length d. If there exist halfspaces $\mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_n$ and points $x, y \in \mathfrak{k}_1^*$ such that $\pi_{xy}(\mathfrak{k}_1) \supsetneq \cdots \supsetneq \pi_{xy}(\mathfrak{k}_n)$, then $n \le 2rd$.*

Proof The sets $C_i := \pi_{xy}(\mathfrak{k}_i)$ are convex, for instance by [50, Lemma 2.2(1)]. Since $C_{i+1} \subsetneq C_i$, there exist halfspaces $\mathfrak{h}_i \in \mathscr{H}(M)$ such that $\mathfrak{h}_i \in \mathscr{H}_{C_i}(M)$ and $C_{i+1} \subseteq \mathfrak{h}_i$.

Since both $\mathfrak{h}_i$ and $\mathfrak{h}_i^*$ intersect $C_i \subseteq I(x, y)$, we have $\mathfrak{h}_i \in \mathscr{H}(x|y) \sqcup \mathscr{H}(y|x)$ for all i. Possibly swapping x and y, we can assume that at least $n/2$ of the $\mathfrak{h}_i$ lie in $\mathscr{H}(x|y)$. By Dilworth's lemma, there exist $k \ge n/2r$ and indices $i_1 < \cdots < i_k$ such that $\mathfrak{h}_{i_1}, \ldots, \mathfrak{h}_{i_k}$ lie in $\mathscr{H}(x|y)$ and no two of them are transverse. Up to reindexing, we can assume that these are $\mathfrak{h}_1, \ldots, \mathfrak{h}_k$.

Since C_j is contained in $\mathfrak{h}_i$ if and only if $j > i$, we must have $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_k$. Note that $y \in \mathfrak{h}_i \cap \mathfrak{k}_j^*$ and $x \in \mathfrak{h}_i^* \cap \mathfrak{k}_j^*$ for all i, j. If $j \le i$, we have $\mathfrak{h}_i \in \mathscr{H}_{C_j}(X)$, hence $\mathfrak{h}_i \cap \mathfrak{k}_j$ and $\mathfrak{h}_i^* \cap \mathfrak{k}_j$ are both nonempty. This shows that $\mathfrak{h}_i$ and $\mathfrak{k}_j$ are transverse for $j \le i$, while the fact that $C_{i+1} \subseteq \mathfrak{h}_i$ implies that $\mathfrak{k}_{i+1} \subseteq \mathfrak{h}_i$. In conclusion, the $\mathfrak{h}_i$ and $\mathfrak{k}_j$ form a length-k staircase with $k \ge n/2r$. Since M has staircase length d, we have $n \le 2rk \le 2rd$. □

Lemma 4.16 *Let X be a* CAT(0) *cube complex of dimension r and staircase length d. Consider vertices $x, y \in X$ and $z \in I(x, y)$. Let $\alpha \subseteq I(x, z)$ be a (combinatorial) geodesic from x to z. Then the median subalgebra $M = X^{(0)} \cap I(x, y) \cap \pi_{xz}^{-1}(\alpha)$ has staircase length $\le d(1 + 2r^2)$.*

Proof Since $\pi_{xz}(y) = z$ and $x, z \in \alpha$, the three points x, y, z all lie in M. Since $M \subseteq I(x, y)$, every wall of M separates x and y. Recall that we use the notation $\mathscr{H}(X)$ and $\mathscr{W}(X)$ with the meaning of $\mathscr{H}(X^{(0)})$ and $\mathscr{W}(X^{(0)})$.

Claim 1 *If $\mathfrak{u}, \mathfrak{v} \in \mathscr{W}(M)$ separate x and z, then $\mathfrak{u}$ and $\mathfrak{v}$ are not transverse.*

Proof Pick halfspaces $\hat{\mathfrak{h}}, \hat{\mathfrak{k}} \in \mathscr{H}(X) \cap \mathscr{H}(x|z)$ such that $\mathfrak{h} := \hat{\mathfrak{h}} \cap M \in \mathscr{H}(M)$ is bounded by $\mathfrak{u}$ and $\mathfrak{k} := \hat{\mathfrak{k}} \cap M$ is bounded by $\mathfrak{v}$; this is possible by Remark 2.2(1). The intersections $\hat{\mathfrak{h}} \cap \alpha$ and $\hat{\mathfrak{k}} \cap \alpha$ are subsegments of α containing z. Without loss of generality, we have $\hat{\mathfrak{h}} \cap \alpha \subseteq \hat{\mathfrak{k}} \cap \alpha$. Then $\hat{\mathfrak{h}} \cap \hat{\mathfrak{k}}^* \cap \alpha = \varnothing$, hence $\varnothing = \hat{\mathfrak{h}} \cap \hat{\mathfrak{k}}^* \cap M = \mathfrak{h} \cap \mathfrak{k}^*$, proving the claim. ◁

Claim 2 *If $\widehat{\mathfrak{h}}, \widehat{\mathfrak{k}} \in \mathscr{H}(z|y)$ are halfspaces of X, then $\widehat{\mathfrak{h}}$ and $\widehat{\mathfrak{k}}$ are transverse if and only if $\widehat{\mathfrak{h}} \cap M$ and $\widehat{\mathfrak{k}} \cap M$ are transverse halfspaces of M.*

Proof Since $\pi_{xz}(I(z,y)) = \{z\}$, the vertex set of the interval $I(z,y) \subseteq X$ is entirely contained in M. Thus, $I(z,y)$ is a convex subset of both X and M. Remark 2.2(2) then shows that $\widehat{\mathfrak{h}}$ and $\widehat{\mathfrak{k}}$ are transverse if and only if $\widehat{\mathfrak{h}} \cap I(z,y)$ and $\widehat{\mathfrak{k}} \cap I(z,y)$ are transverse, if and only if $\widehat{\mathfrak{h}} \cap M$ and $\widehat{\mathfrak{k}} \cap M$ are transverse. ◁

Now, suppose that M contains a length-n staircase. Thus M has halfspaces $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_n$ and $\mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_n$ such that each $\mathfrak{h}_i$ is transverse to all $\mathfrak{k}_j$ with $j \leq i$, while $\mathfrak{k}_{i+1} \subseteq \mathfrak{h}_i$.

Since $\mathfrak{k}_n \subseteq \mathfrak{h}_{n-1} \subseteq \mathfrak{h}_1$, we have either $\{\mathfrak{h}_1, \mathfrak{k}_n\} \subseteq \mathscr{H}(x|y)$ or $\{\mathfrak{h}_1, \mathfrak{k}_n\} \subseteq \mathscr{H}(y|x)$. If we replace all $\mathfrak{h}_i$ and $\mathfrak{k}_j$ with $\mathfrak{k}^*_{n-i+1}$ and $\mathfrak{h}^*_{n-j+1}$, respectively, we obtain another length-n staircase. Thus, we can assume that $\{\mathfrak{h}_1, \mathfrak{k}_n\} \subseteq \mathscr{H}(x|y)$. It follows that all $\mathfrak{h}_i$ and $\mathfrak{k}_j$ lie in $\mathscr{H}(x|y)$.

Let $0 \leq a, b \leq n$ be the largest indices such that $z \in \mathfrak{h}_i$ and $z \in \mathfrak{k}_j$ hold for $1 \leq i \leq a$ and $1 \leq j \leq b$. Since $\mathfrak{h}_1$ and $\mathfrak{k}_1$ are transverse, Claim 1 shows that they cannot both lie in $\mathscr{H}(x|z)$. Thus $\min\{a,b\} = 0$. Since $\mathfrak{k}_{a+2} \subseteq \mathfrak{h}_{a+1}$, we have $z \notin \mathfrak{k}_{a+2}$, hence $b \leq a+1$. In conclusion, either $b = 0$, or $(a,b) = (0,1)$.

The halfspaces $\mathfrak{h}_i, \mathfrak{k}_j$ with $i, j > \max\{a,b\}$ all lie in $\mathscr{H}(z|y)$ and form a staircase of length $n - \max\{a,b\}$. By Remark 2.2(1) and Claim 2, this determines a staircase of halfspaces of X. Since X has staircase length d, we deduce that $n - \max\{a,b\} \leq d$.

If $b = 1$ and $a = 0$, we get $n \leq d+1$ and we are done. If instead $b = 0$, then $n \leq a + d$ and the proof is completed with the following claim:

Claim 3 *If $b = 0$, then $a \leq 2r^2 d$.*

Proof As a recap, M has halfspaces $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_a$ in $\mathscr{H}(x|z,y)$ and $\mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_a$ in $\mathscr{H}(x,z|y)$ forming a length-a staircase. By Remark 2.2(1), there exist halfspaces $\widehat{\mathfrak{h}}_i, \widehat{\mathfrak{k}}_j \in \mathscr{H}(X)$ such that $\mathfrak{h}_i = \widehat{\mathfrak{h}}_i \cap M$ and $\mathfrak{k}_i = \widehat{\mathfrak{k}}_i \cap M$.

By Dilworth's lemma, there exist $a' \geq a/r$ and indices $1 \leq j_1 < \cdots < j_{a'} \leq a$ such that no two among $\widehat{\mathfrak{k}}_{j_1}, \ldots, \widehat{\mathfrak{k}}_{j_{a'}}$ are transverse. Thus, up to reindexing, we can assume that $\widehat{\mathfrak{k}}_1 \supsetneq \cdots \supsetneq \widehat{\mathfrak{k}}_{a'}$.

Now, since the $\mathfrak{h}_i$ and $\mathfrak{k}_j$ form a staircase in M and $\widehat{\mathfrak{h}}_i \in \mathscr{H}(x|z)$, we have, for every $1 \leq j \leq a'$,

- $\varnothing = \mathfrak{h}^*_j \cap \mathfrak{k}_{j+1} = \widehat{\mathfrak{h}}^*_j \cap \widehat{\mathfrak{k}}_{j+1} \cap M$, hence $\pi_{xz}(\widehat{\mathfrak{k}}_{j+1}) \cap \widehat{\mathfrak{h}}^*_j \cap \alpha = \varnothing$;
- $\varnothing \neq \mathfrak{h}^*_j \cap \mathfrak{k}_j = \widehat{\mathfrak{h}}^*_j \cap \widehat{\mathfrak{k}}_j \cap M$, hence $\pi_{xz}(\widehat{\mathfrak{k}}_j) \cap \widehat{\mathfrak{h}}^*_j \cap \alpha \neq \varnothing$.

Note moreover that $x, z \in \widehat{\mathfrak{k}}^*_1$. If we had $a' > 2rd$, Lemma 4.15 would imply that there exists j with $\pi_{xz}(\widehat{\mathfrak{k}}_j) = \pi_{xz}(\widehat{\mathfrak{k}}_{j+1})$. However, $\pi_{xz}(\widehat{\mathfrak{k}}_j)$ intersects $\widehat{\mathfrak{h}}^*_j \cap \alpha$ while $\pi_{xz}(\widehat{\mathfrak{k}}_{j+1})$ does not.

We conclude that $a \leq ra' \leq 2r^2 d$, as required. ◁

As discussed before Claim 3, this proves the lemma. □

Recall that, if Γ is a finite simplicial graph, $\mathcal{X}_\Gamma$ and $\mathcal{Y}_\Gamma$ denote the universal covers, respectively, of the Salvetti complex for $\mathcal{A}_\Gamma$ and the Davis complex for $\mathcal{W}_\Gamma$.

Lemma 4.17 *The staircase length of $\mathcal{X}_\Gamma$ and $\mathcal{Y}_\Gamma$ is at most $\#\Gamma^{(0)}$.*

Proof We only run the proof for $\mathcal{X}_\Gamma$, since the argument for $\mathcal{Y}_\Gamma$ is identical. The important property, shared by both complexes, is that there is a map $\gamma\colon \mathscr{W}(\mathcal{X}_\Gamma) \to \Gamma^{(0)}$ such that, if $\mathfrak{u}, \mathfrak{v}$ are hyperplanes with intersecting carriers, then $\mathfrak{u}$ and $\mathfrak{v}$ are transverse if and only if $\gamma(\mathfrak{u})$ and $\gamma(\mathfrak{v})$ are joined by an edge of Γ. For simplicity, let us extend the map γ to $\mathscr{H}(\mathcal{X}_\Gamma)$, simply by composing it with the two-to-one map $\mathscr{H}(\mathcal{X}_\Gamma) \to \mathscr{W}(\mathcal{X}_\Gamma)$ pairing each halfspace with its hyperplane.

Consider halfspaces $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_n$ and $\mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_n$ such that $\mathfrak{h}_i$ is transverse to all $\mathfrak{k}_j$ with $j \le i$, while $\mathfrak{k}_{i+1} \subsetneq \mathfrak{h}_i$. We define the following subsets of $\Gamma^{(0)}$:

$$\Gamma_j := \gamma(\mathfrak{k}_1^*) \cup \gamma(\mathscr{W}(\mathfrak{k}_1^* | \mathfrak{k}_j)) \cup \{\gamma(\mathfrak{k}_j)\}.$$

It is clear that $\Gamma_j \subseteq \Gamma_{j+1}$ for all $j \ge 1$. The lemma is immediate from the following claim:

Claim *We have $\Gamma_j \subsetneq \Gamma_{j+1}$ for all $j \ge 1$.*

Suppose for the sake of contradiction that, for some $j \ge 1$, we have $\Gamma_{j+1} = \Gamma_j$.

Given $\mathfrak{j} \in \mathscr{H}(\mathfrak{h}_j^* | \mathfrak{k}_{j+1})$, we have $\mathfrak{j} \cap \mathfrak{k}_1 \supseteq \mathfrak{k}_{j+1} \neq \varnothing$. Moreover, $\mathfrak{j}^* \cap \mathfrak{k}_1 \neq \varnothing$ and $\mathfrak{j}^* \cap \mathfrak{k}_1^* \neq \varnothing$, since $\mathfrak{j}^*$ contains $\mathfrak{h}_j^*$, which is transverse to $\mathfrak{k}_1$. Thus, for each $\mathfrak{j} \in \mathscr{H}(\mathfrak{h}_j^* | \mathfrak{k}_{j+1})$, there are only two possibilities: either

(a) $\mathfrak{j} \cap \mathfrak{k}_1^* = \varnothing$, hence $\mathfrak{j} \subseteq \mathfrak{k}_1$ and $\mathfrak{j} \in \mathscr{H}(\mathfrak{k}_1^* | \mathfrak{k}_{j+1})$; or

(b) $\mathfrak{j}$ is transverse to $\mathfrak{k}_1$.

Note that no halfspace of type (a) can contain a halfspace of type (b). Moreover, each $\mathfrak{j}$ of type (b) is also transverse to $\mathfrak{k}_j$: we have $\mathfrak{j} \cap \mathfrak{k}_j \supseteq \mathfrak{k}_{j+1} \neq \varnothing$, $\mathfrak{j} \cap \mathfrak{k}_j^* \supseteq \mathfrak{j} \cap \mathfrak{k}_1^* \neq \varnothing$, $\mathfrak{j}^* \cap \mathfrak{k}_j \supseteq \mathfrak{h}_j^* \cap \mathfrak{k}_j \neq \varnothing$ and $\mathfrak{j}^* \cap \mathfrak{k}_j^* \supseteq \mathfrak{j}^* \cap \mathfrak{k}_1^* \neq \varnothing$. Thus, every $\mathfrak{j}$ of type (b) is transverse to the set $\mathscr{H}(\mathfrak{k}_1^* | \mathfrak{k}_j) \cup \{\mathfrak{k}_1^*, \mathfrak{k}_j\}$.

Now, consider a maximal chain of halfspaces $\mathfrak{j}_1 \supsetneq \cdots \supsetneq \mathfrak{j}_m$ in $\mathscr{H}(\mathfrak{h}_j^* | \mathfrak{k}_{j+1})$ with $m \ge 0$. We can enlarge this chain by adding $\mathfrak{j}_0 := \mathfrak{h}_j$ and $\mathfrak{j}_{m+1} = \mathfrak{k}_{j+1}$, which are, respectively, of type (b) and (a). Thus, there exists an index $0 \le k \le m$ such that $\mathfrak{j}_0, \dots, \mathfrak{j}_k$ are of type (b) and $\mathfrak{j}_{k+1}, \dots, \mathfrak{j}_{m+1}$ are of type (a). Since the chain is maximal, the set $\mathscr{W}(\mathfrak{j}_k^* | \mathfrak{j}_{k+1})$ is empty. Thus, since $\mathfrak{j}_k$ and $\mathfrak{j}_{k+1}$ are not transverse, the labels $\gamma(\mathfrak{j}_k)$ and $\gamma(\mathfrak{j}_{k+1})$ are not joined by an edge of Γ.

However, since $\Gamma_{j+1} = \Gamma_j$, we have

$$\gamma(\mathfrak{j}_{k+1}) \in \gamma(\mathscr{W}(\mathfrak{k}_1^* | \mathfrak{k}_{j+1})) \cup \{\gamma(\mathfrak{k}_1^*), \gamma(\mathfrak{k}_{j+1})\} = \gamma(\mathscr{W}(\mathfrak{k}_1^* | \mathfrak{k}_j)) \cup \{\gamma(\mathfrak{k}_1^*), \gamma(\mathfrak{k}_j)\},$$

while $\mathfrak{j}_k$ is transverse to $\mathscr{H}(\mathfrak{k}_1^* | \mathfrak{k}_j) \cup \{\mathfrak{k}_1^*, \mathfrak{k}_j\}$, a contradiction. This proves the claim and lemma. $\square$

4.4 A quasiconvexity criterion for median subalgebras

In this subsection, we provide a criterion (Proposition 4.25) for when a median subalgebra M of a CAT(0) cube complex X is quasiconvex. The subalgebra M will be required to satisfy two conditions, *edge-connectedness* and *weak quasiconvexity*, which we study separately in the next two subsections.

4.4.1 Edge-connected median subalgebras Let X be a CAT(0) cube complex.

Definition 4.18 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 4.19 If $A \subseteq X^{(0)}$ is edge-connected, then there do not exist distinct halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}_A(X)$ with $\mathfrak{h} \cap A = \mathfrak{k} \cap A$. Indeed, the intersections $\mathfrak{h} \cap \mathfrak{k}$ and $\mathfrak{h}^* \cap \mathfrak{k}^*$ would both be nonempty, so, possibly swapping $\mathfrak{h}$ and $\mathfrak{k}$, we would either have $\mathfrak{h} \subsetneq \mathfrak{k}$ or $\mathfrak{h}$ and $\mathfrak{k}$ would be transverse. However, since A is edge connected and intersects both $\mathfrak{h} \cap \mathfrak{k}$ and $\mathfrak{h}^* \cap \mathfrak{k}^*$, we must have $A \cap \mathfrak{h}^* \cap \mathfrak{k} \neq \varnothing$ if $\mathfrak{h} \subsetneq \mathfrak{k}$, and either $A \cap \mathfrak{h}^* \cap \mathfrak{k} \neq \varnothing$ or $A \cap \mathfrak{h} \cap \mathfrak{k}^* \neq \varnothing$ if $\mathfrak{h}$ and $\mathfrak{k}$ are transverse. This contradicts the fact that $\mathfrak{h} \cap A = \mathfrak{k} \cap A$.

Lemma 4.20 *For a median subalgebra $M \subseteq X^{(0)}$, the following are equivalent*:

(1) *M is edge-connected.*

(2) *For all $x, y \in M$, there exists a geodesic $\alpha \subseteq X$ joining x and y such that $\alpha \cap X^{(0)} \subseteq M$.*

(3) *The restriction map $\mathrm{res}_M : \mathscr{H}_M(X) \to \mathscr{H}(M)$ is injective.*

Proof The implication (2) $\Longrightarrow$ (1) is clear and the implication (1) $\Longrightarrow$ (3) follows from Remark 4.19. Let us show that (3) $\Longrightarrow$ (2).

Since M is a discrete median algebra, it is isomorphic to the 0–skeleton of a CAT(0) cube complex $X(M)$; see [31, Theorem 6.1] or [89, Section 10]. Given $x, y \in M$, let $\beta \subseteq X(M)$ be a geodesic joining x and y, and let $x_1 = x, x_2, \dots, x_n = y$ be the elements of $\beta \cap M$ as they appear along β. Since the restriction map $\mathrm{res}_M : \mathscr{H}_M(X) \to \mathscr{H}(M)$ is injective, there is only one hyperplane $\mathfrak{w}_i \in \mathscr{W}(X)$ separating x_i and x_{i+1}, that is, these two points are joined by an edge of X. If $i \neq j$, then $\mathfrak{w}_i \neq \mathfrak{w}_j$, or β would cross the corresponding wall of M twice. We conclude that there exists a geodesic $\alpha \subseteq X$ with $\alpha \cap M = \{x_1, \dots, x_n\}$. □

By the 3rd characterisation in Lemma 4.20, edge-connected subalgebras can be viewed as a middle ground between general median subalgebras and convex subcomplexes; cf part (2) of Remark 2.2.

Lemma 4.21 *If $A \subseteq X^{(0)}$ is an edge-connected subset, then $\langle A \rangle$ is an edge-connected subalgebra.*

Proof Suppose for the sake of contradiction that $\langle A \rangle$ is not edge-connected. Then there exist distinct halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}_{\langle A \rangle}(X)$ with $\mathfrak{h} \cap \langle A \rangle = \mathfrak{k} \cap \langle A \rangle$ by Lemma 4.20. Note that $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}_A(X)$, and $\mathfrak{h}^* \cap \mathfrak{k} \cap A = \varnothing$ and $\mathfrak{h} \cap \mathfrak{k}^* \cap A = \varnothing$. In particular, $\mathfrak{h} \cap A = \mathfrak{k} \cap A$, which violates Remark 4.19. □

Lemma 4.22 *Let $M \subseteq X^{(0)}$ be an edge-connected median subalgebra. Let $C \subseteq X$ be a convex subcomplex with gate-projection $\pi: X \to C$. Then:*

(1) *$\pi(M)$ is an edge-connected subalgebra of $C^{(0)}$.*

(2) *If $N \subseteq \pi(M)$ is an edge-connected subalgebra, then $M \cap \pi^{-1}(N)$ is edge-connected as well.*

Proof If vertices $x, y \in X$ are joined by an edge, then either $\pi(x)$ and $\pi(y)$ are joined by an edge or they are equal. Thus, part (1) is immediate from definitions.

Let us address part (2). Consider two points $x, y \in M \cap \pi^{-1}(N)$. Since N is edge-connected, there exists a geodesic $\alpha \subseteq C$ joining $\pi(x)$ and $\pi(y)$ with $\alpha \cap C^{(0)} \subseteq N$; see Lemma 4.20. It suffices to show that $M \cap \pi^{-1}(\alpha)$ is edge-connected.

In fact, since $\pi^{-1}(v) \cap M \neq \varnothing$ for every vertex $v \in \alpha$, it suffices to show that $M \cap \pi^{-1}(e)$ is edge-connected for every edge $e \subseteq \alpha$. In other words, we can suppose that $\pi(x)$ and $\pi(y)$ are joined by an edge $e \subseteq C$. Since M is edge-connected, there exists a geodesic $\beta \subseteq X$ joining x and y with $\beta \cap X^{(0)} \subseteq M$. Since π is a median morphism, the projection $\pi(\beta)$ is the image of a geodesic from $\pi(x)$ to $\pi(y)$, ie $\pi(\beta) = e$. Thus $\beta \cap X^{(0)} \subseteq M \cap \pi^{-1}(e)$, concluding the proof. □

4.4.2 Weakly quasiconvex median subalgebras Let X be a CAT(0) cube complex.

Definition 4.23 A subset $A \subseteq X^{(0)}$ is *weakly quasiconvex* if there exists a function $\eta: \mathbb{N} \to \mathbb{N}$ such that, for all $a, b, p \in X^{(0)}$ with $\mathscr{W}(p|a)$ transverse to $\mathscr{W}(p|b)$, we have

$$d(p, A) \leq \eta\big(\max\{d(a, A), d(b, A)\}\big).$$

Remark 4.24 (1) If $A \subseteq X^{(0)}$ is quasiconvex in the sense of Definition 2.30, then A is weakly quasiconvex. Indeed, suppose that $\mathcal{J}(A) \subseteq \mathcal{N}_R(A)$ and set $D = \max\{d(a, A), d(b, A)$. If $\mathscr{W}(p|a)$ and $\mathscr{W}(p|b)$ are transverse, then $p \in I(a, b)$. Thus, $p \in \mathcal{J}(\mathcal{N}_D(A))$ and Lemma 2.10 yields $d(p, A) \leq 2D + R =: \eta(D)$.

(2) If $A, B \subseteq X^{(0)}$ have finite Hausdorff distance, then A is weakly quasiconvex if and only if B is. This is straightforward, observing that η can always taken to be weakly increasing.

The following is the main result of this subsection.

Proposition 4.25 *If X has finite dimension and finite staircase length, then every edge-connected, weakly quasiconvex median subalgebra $M \subseteq X^{(0)}$ is quasiconvex.*

Proposition 4.25 fails for cube complexes of infinite staircase length, as the next example shows.

Example 4.26 Consider the standard structure of cube complex on $\mathbb{R}^2$. Let α be the geodesic line through all points (n,n) and $(n+1,n)$ with $n \in \mathbb{Z}$. Let $X \subseteq \mathbb{R}^2$ be the subcomplex that lies above α, including α itself. Note that X is a 2–dimensional CAT(0) cube complex of infinite staircase length, and $\alpha \subseteq X$ is an edge-connected median subalgebra that is not quasiconvex. It is not hard to see that α is weakly quasiconvex with $\eta(t) = 2t$.

The next lemma essentially proves the 2–dimensional case of Proposition 4.25.

Lemma 4.27 *Suppose that* $\dim X = 2$ *and that* X *has staircase length* d. *Let* $M \subseteq X^{(0)}$ *be an edge-connected median subalgebra. Consider* $x, y \in M$ *and* $z \in X^{(0)} \cap I(x,y)$. *Then there exist* $0 \le k \le d$ *and vertices* $z_0, z_1, z_2, \dots, z_k \in I(x,y)$ *and* $w_1, \dots, w_k \in I(x,y)$ *such that*

- $z_0 = z$, *while* $z_k \in M$ *and* $w_1, \dots, w_k \in M$;
- *the sets* $\mathscr{W}(z_i|w_{i+1}) \subseteq \mathscr{W}(X)$ *and* $\mathscr{W}(z_i|z_{i+1}) \subseteq \mathscr{W}(X)$ *are transverse for all* $0 \le i \le k-1$.

Proof If $z \in M$, we can simply take $k = 0$. If $z \notin M$, we begin with the following observation:

Claim *We can assume that there exist transverse hyperplanes* $\mathfrak{u} \in \mathscr{W}(x,z|y)$ *and* $\mathfrak{v} \in \mathscr{W}(y,z|x)$ *such that* x, z *lie in the carrier of* $\mathfrak{u}$ *and* y, z *lie in the carrier of* $\mathfrak{v}$.

Proof Up to replacing x and y with other points in the interval $I(x,y)$, we can assume that there do not exist points $x', y' \in I(x,y)$ with $z \in I(x',y')$, except for $\{x',y'\} = \{x,y\}$.

Since M is edge-connected, there exists a point $x' \in M \cap I(x,y)$ such that x and x' are separated by a single hyperplane $\mathfrak{u} \in \mathscr{W}(X)$. By the above assumption on x and y, we must have $z \notin I(x',y)$, hence $\varnothing \neq \mathscr{W}(z|x',y) = \mathscr{W}(z,x|x',y) \subseteq \{\mathfrak{u}\}$. It follows that $\mathscr{W}(z,x|x',y) = \{\mathfrak{u}\}$.

Observing that $\mathscr{W}(z|\mathfrak{u}) \subseteq \mathscr{W}(z|x',y) = \mathscr{W}(z,x|x',y) = \{\mathfrak{u}\}$, we conclude that $\mathscr{W}(z|\mathfrak{u})$ is empty. This shows that the carrier of $\mathfrak{u}$ contains z, while it is clear that it also contains x. The existence of $\mathfrak{v}$ is obtained similarly. Finally, since $\mathfrak{v} \in \mathscr{W}(y,z|x)$ and $\mathfrak{v} \neq \mathfrak{u}$, we must have $\mathfrak{v} \in \mathscr{W}(y,z|x,x')$. Recalling that $\mathfrak{u} \in \mathscr{W}(z,x|x',y)$, this shows that $\mathfrak{u}$ and $\mathfrak{v}$ are transverse. ◁

Now, the sets $\mathscr{H}(z|x)$ and $\mathscr{H}(z|y)$ are transverse, respectively, to $\mathfrak{u}$ and $\mathfrak{v}$. Since $\dim X = 2$, the set $\mathscr{H}(z|x)$ is a descending chain $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_m$, and $\mathscr{H}(z|y)$ is a descending chain $\mathfrak{k}_1 \supsetneq \cdots \supsetneq \mathfrak{k}_n$. Note that $\mathfrak{k}_1$ and $\mathfrak{h}_1$ are bounded, respectively, by $\mathfrak{u}$ and $\mathfrak{v}$, as depicted in Figure 4.

Since $\mathfrak{h}_1$ and $\mathfrak{k}_1$ are transverse, there exists a function $\tau\colon \{1,\dots,m\} \to \{1,\dots,n\}$ such that $\mathfrak{h}_i$ is transverse to $\mathfrak{k}_j$ if and only if $1 \le j \le \tau(i)$. Note that $\tau(1) = n$ and that τ is weakly decreasing.

Let $1 \le i_1 < \cdots < i_{k-1} < m$ be all indices i with $\tau(i+1) < \tau(i)$. Also define $i_k := m$ and set $\tau_s := \tau(i_s)$ for simplicity. Since the halfspaces $\mathfrak{h}^*_{i_k}, \dots, \mathfrak{h}^*_{i_1}$ and $\mathfrak{k}_{\tau_k}, \dots, \mathfrak{k}_{\tau_1}$ form a length-k staircase, while X has staircase length d, we must have $k \le d$.

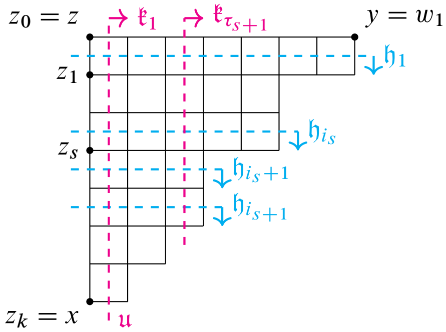

Figure 4

Set $z_0 = z$ and $w_1 = y$. For $1 \le s \le k$, let $z_s \in I(x,y)$ be the point with $\mathscr{H}(z|z_s) = \{\mathfrak{h}_1, \dots, \mathfrak{h}_{i_s}\}$. In particular, $z_k = x \in M$. Since M is edge-connected, there exist points

$$w_{s+1} \in M \cap \mathfrak{h}_{i_s} \cap \mathfrak{h}^*_{i_s+1} \cap \mathfrak{k}^*_{\tau_{s+1}+1}.$$

Observing that $\mathscr{H}(z_s|w_{s+1}) \subseteq \{\mathfrak{k}_1, \dots, \mathfrak{k}_{\tau_{s+1}}\}$ is transverse to $\mathscr{H}(z_s|z_{s+1}) = \{\mathfrak{h}_{i_s+1}, \dots, \mathfrak{h}_{i_{s+1}}\}$, this completes the proof of the lemma. □

The next lemma allows us to reduce the proof of Proposition 4.25 to the 2–dimensional case.

Lemma 4.28 *Let X have dimension r and staircase length d. Let $M \subseteq X^{(0)}$ be an edge-connected median subalgebra. For all points $x, y \in M$ and $z \in X^{(0)} \cap I(x,y)$, there exists a median subalgebra $N \subseteq X^{(0)} \cap I(x,y)$ with the following properties*:

- *$x, y, z \in N$ and* $\operatorname{rk} N \le 2$.
- *N has staircase length $\le d(1+2r^2)^2$.*
- *N and $N \cap M$ are edge-connected.*

Proof Let $\pi_{xz} \colon X \to I(x,z)$ be the gate-projection and note that $\pi_{xz}(y) = z$. By Lemma 4.22(1), the projection $\pi_{xz}(M)$ is an edge-connected median subalgebra containing x and z. Thus there exists a (combinatorial) geodesic $\alpha \subseteq I(x,z)$ joining x and z with $\alpha \cap X^{(0)} \subseteq \pi_{xz}(M)$.

By Lemma 4.22(2), the median subalgebras $N' := \pi_{xz}^{-1}(\alpha) \cap I(x,y) \cap X^{(0)}$ and $M \cap N'$ are edge-connected. Lemma 4.16 shows that N' has staircase length $\le d(1+2r^2)$, while it is clear that $\operatorname{rk} N' \le \dim X = r$.

Note that $x, y, z \in N'$. Since $\pi_{xz}(I(z,y)) = \{z\}$, the entire interval $I(z,y) \cap X^{(0)}$ is contained in N'. Consider the projection $\pi_{zy} \colon X \to I(z,y)$. Since $M \cap N'$ is edge-connected, Lemma 4.22 again shows that the projection $\pi_{zy}(M \cap N')$ is edge-connected, and we can join y and z by a geodesic β with $\beta \cap X^{(0)} \subseteq \pi_{zy}(M \cap N')$. Repeating the above argument, we see that $N := N' \cap \pi_{yz}^{-1}(\beta)$ has staircase length $\le d(1+2r^2)^2$, that N and $N \cap M$ are edge-connected, and that $x, y, z \in N$ (recall that N' is a finite median algebra, so it is naturally identified with the 0–skeleton of a CAT(0) cube complex and we can run the above argument in this cube complex).

We are left to show that $\operatorname{rk} N \le 2$. Since $x, y \in N \subseteq I(x, y)$, every wall of N either separates x from y, z, or it separates x, z from y. If two walls of N separate x and z, then they are not transverse; cf Claim 1 during the proof of Lemma 4.16. The same is true of walls separating z and y. This implies that $\operatorname{rk} N \le 2$, concluding the proof. □

Proof of Proposition 4.25 Let X have dimension r and staircase length d. Let M be an edge-connected, weakly quasiconvex subalgebra. We will show that $d_{\mathrm{Haus}}(I(x, y), M \cap I(x, y))$ remains uniformly bounded as x and y vary in M, which implies that M is quasiconvex.

Consider $x, y \in M$ and $z \in X^{(0)} \cap I(x, y)$. By Lemma 4.28, the points x, y, z lie in a median subalgebra $N \subseteq X^{(0)} \cap I(x, y)$ such that N and $N \cap M$ are edge-connected, $\operatorname{rk} N \le 2$, and N has staircase length $\le d(1 + 2r^2)^2$.

Viewing N as the vertex set of a finite CAT(0) cube complex and applying Lemma 4.27 to $M \cap N$, there exist points $z_0 = z, z_1, \dots, z_{k-1} \in N$ and $z_k, w_1, \dots, w_k \in N \cap M$ with $k \le d(1 + 2r^2)^2$, such that each wall of N separating z_i and z_{i+1} is transverse to every wall of N separating z_i and w_{i+1}. The same is true of hyperplanes of X separating these points.

Since M is weakly quasiconvex, it admits a function η as in Definition 4.23. Without loss of generality, we can take η to be weakly increasing. Then, since $d(w_i, M) = 0$, we have

$$\begin{aligned} d(z, M) \le \max\{\eta(d(z_1, M)), \eta(0)\} &\le \max\{\eta^2(d(z_2, M)), \eta^2(0), \eta(0)\} \\ &\le \dots \le \max\{\eta^k(0), \dots, \eta^2(0), \eta(0)\}. \end{aligned}$$

The last constant only depends on d, r and η, so this proves that M is quasiconvex. □

4.5 Fixed subgroups in right-angled groups

In this subsection, we combine the results of the previous two subsections to prove Theorem C.

Let Γ be a finite simplicial graph. Our focus will be on the right-angled Artin group $\mathcal{A} = \mathcal{A}_\Gamma$ and the universal cover of its Salvetti complex $\mathcal{X} = \mathcal{X}_\Gamma$. Throughout, we will identify $\mathcal{A} \cong \mathcal{X}^{(0)}$.

However, all results and proofs in this subsection (except for Remark 4.29) immediately extend to right-angled Coxeter groups $\mathcal{W} = \mathcal{W}_\Gamma$ and Davis complexes $\mathcal{Y}_\Gamma$, without requiring any adaptations. We suggest that the reader keep track of this as they make their way through the results, in view of Corollary 4.35 below. The relevant properties shared by RAAGs and RACGs are:

- The Cayley graph of $\mathcal{A}/\mathcal{W}$ associated to the standard generators (vertices of Γ) is the 1–skeleton of a CAT(0) cube complex (the universal cover of the Salvetti/Davis complex) of finite staircase length (Lemma 4.17).

- Hyperplanes are labelled by vertices of Γ and labels of transverse hyperplanes are joined by an edge of Γ.
- Elementary automorphisms of $\mathcal{A}$ and $\mathcal{W}$ (as defined in Section 3.4) have the same form with respect to standard generators.

We are interested in the subgroups $U_0(\mathcal{A}) \leq U(\mathcal{A})$ and $\mathrm{Aut}_0\,\mathcal{W} \leq \mathrm{Aut}\,\mathcal{W}$ generated by inversions, folds and partial conjugations, as defined at the end of Section 3.4.

Given a subset $\Delta \subseteq \Gamma^{(0)}$, it is convenient to introduce the notation

$$\Delta^{\perp} = \bigcap_{v \in \Delta} \mathrm{lk}\, v.$$

Remark 4.29 It is not hard to observe that a subgroup of $\mathcal{A}$ is an intersection of stabilisers of hyperplanes of $\mathcal{X}$ if and only if it is conjugate to a subgroup of the form $\mathcal{A}_{\Delta^{\perp}}$ for some $\Delta \subseteq \Gamma$.

Although we will not be using this remark in the present paper, we find it interesting in relation to Lemma 4.30 below: elements of $U_0(\mathcal{A})$ permute hyperplane-stabilisers while preserving labels.

Statements similar to the next lemma have been widely used in the literature, eg in [24, Proposition 3.2; 27, Proposition 3.2; 28, Section 3]). Compared to these references, we get a slightly stronger result because here we are only concerned with untwisted automorphisms.

Lemma 4.30 *For every $\varphi \in U_0(\mathcal{A})$ and $\Delta \subseteq \Gamma$, the subgroups $\mathcal{A}_{\Delta^{\perp}}$ and $\varphi(\mathcal{A}_{\Delta^{\perp}})$ are conjugate.*

Proof It suffices to prove the lemma for elementary generators. It is clear that it holds for inversions, so we are left to consider folds and partial conjugations.

If $\tau_{v,w}$ is a fold, then $\tau_{v,w}(\mathcal{A}_{\Delta^{\perp}}) = \mathcal{A}_{\Delta^{\perp}}$. This is immediate if $v \notin \Delta^{\perp}$. If instead $v \in \Delta^{\perp}$, we have $\Delta \subseteq \mathrm{lk}\, v \subseteq \mathrm{lk}\, w$, hence $w \in \Delta^{\perp}$.

If $\kappa_{w,C}$ is a partial conjugation, then $\kappa_{w,C}(\mathcal{A}_{\Delta^{\perp}})$ is either $\mathcal{A}_{\Delta^{\perp}}$ or $w^{-1}\mathcal{A}_{\Delta^{\perp}}w$. This is clear if $\Delta^{\perp}$ intersects at most one connected component of $\Gamma \setminus \mathrm{st}\, w$. Suppose instead that $\Delta^{\perp}$ intersects two distinct components of $\Gamma \setminus \mathrm{st}\, w$. Then, for every $a \in \Delta$, the fact that $\Delta^{\perp} \subseteq \mathrm{lk}\, a$ implies that $a \in \mathrm{lk}\, w$. Thus, $w \in \Delta^{\perp}$ and $\kappa_{w,C}(\mathcal{A}_{\Delta^{\perp}}) = \mathcal{A}_{\Delta^{\perp}}$ in this case. □

Corollary 4.31 *For every $\varphi \in U_0(\mathcal{A})$ and $g \in \mathcal{A}$, we have $\Gamma(\varphi(g))^{\perp} = \Gamma(g)^{\perp}$.*

Proof It suffices to show that $\Gamma(\varphi(g))^{\perp} \supseteq \Gamma(g)^{\perp}$ for all $\varphi \in U_0(\mathcal{A})$ and $g \in \mathcal{A}$. Note that g has a conjugate in $\mathcal{A}_{\Gamma(g)} \leq \mathcal{A}_{\Gamma(g)^{\perp\perp}}$. Thus, Lemma 4.30 implies that a conjugate of $\varphi(g)$ lies in $\mathcal{A}_{\Gamma(g)^{\perp\perp}}$. This shows that $\Gamma(\varphi(g)) \subseteq \Gamma(g)^{\perp\perp}$, hence $\Gamma(\varphi(g))^{\perp} \supseteq \Gamma(g)^{\perp\perp\perp} = \Gamma(g)^{\perp}$, as required. □

For the next results, recall that we are identifying elements of $\mathcal{A}$ and vertices of $\mathcal{X}$.

Lemma 4.32 *For every $\varphi \in U_0(\mathcal{A})$, there exists a constant $K(\varphi)$ with the following property. For all $x, y \in \mathcal{A}$, at most $K(\varphi)$ among the hyperplanes in $\mathscr{W}(\varphi(x)|\varphi(y))$ have label outside $\gamma(\mathscr{W}(x|y))^{\perp\perp}$.*

Proof It suffices to show that, for every $g \in \mathcal{A}$, at most $K(\varphi)$ among the hyperplanes in $\mathscr{W}(1|\varphi(g))$ have label outside $\gamma(\mathscr{W}(1|g))^{\perp\perp}$.

Since Γ has only finitely many subsets, Lemma 4.30 shows that there exists a constant $K'(\varphi)$ with the following property. For every $\Delta \subseteq \Gamma$ there exists $x_\Delta \in \mathcal{A}$ with $\varphi(\mathcal{A}_{\Delta^\perp}) = x_\Delta \mathcal{A}_{\Delta^\perp} x_\Delta^{-1}$ and $|x_\Delta| \leq K'(\varphi)$. Here $|\cdot|$ denotes word length with respect to the standard generators.

Now, consider $g \in \mathcal{A}$ and set $\Delta(g) := \gamma(\mathscr{W}(1|g))^\perp$. Then $g \in \mathcal{A}_{\Delta(g)^\perp}$ and the above observation shows that all but $2|x_{\Delta(g)}|$ hyperplanes in $\mathscr{W}(1|\varphi(g))$ have label in $\Delta(g)^\perp$. Taking $K(\varphi) := 2K'(\varphi)$, this concludes the proof. □

Proposition 4.33 *If $\varphi \in U_0(\mathcal{A})$, the subgroup* $\operatorname{Fix}\varphi$ *is a weakly quasiconvex subset of $\mathcal{X}^{(0)} \cong \mathcal{A}$.*

Proof Consider vertices $a, b, p \in \mathcal{X}$ with $\mathscr{W}(p|a)$ transverse to $\mathscr{W}(p|b)$. Set

$$D := \max\{d(a, \operatorname{Fix}\varphi), d(b, \operatorname{Fix}\varphi)\}.$$

Let $K = K(\varphi)$ be as in Lemma 4.32, let ζ_1, ζ_2 be the functions provided by Lemma 2.34 (without loss of generality, strictly increasing), and let C be a constant such that

$$\varphi(m(x,y,z)) \approx_C m(\varphi(x), \varphi(y), \varphi(z)) \quad \text{for all } x, y, z \in \mathcal{X}.$$

Let us write a', b', p' for $\varphi(a), \varphi(b), \varphi(p)$. Since $\mathscr{W}(p|a)$ and $\mathscr{W}(p|b)$ are transverse, we have $p \in I(a,b)$, so $\mathscr{W}(p|a,b) = \varnothing$. Observing that $m(a', b', p') \approx_C \varphi(m(a,b,p)) = p'$, we also have $\#\mathscr{W}(p'|a',b') \leq C$. Finally, by the first inequality in Lemma 2.34, we have $a' \approx_{D'} a$ and $b' \approx_{D'} b$, where $D' := \zeta_1^{-1}(D)$.

Putting together these inequalities, we obtain

$$\begin{aligned}
\#\mathscr{W}(p|p') &= \#\mathscr{W}(p|a',b',p') + \#\mathscr{W}(p,a'|b',p') + \#\mathscr{W}(p,b'|a',p') + \#\mathscr{W}(p,a',b'|p') \\
&\leq \#\mathscr{W}(p|a,b) + 2D' + \#\mathscr{W}(p,a'|b,p') + D' + \#\mathscr{W}(p,b'|a,p') + D' + \#\mathscr{W}(a',b'|p') \\
&\leq \#\mathscr{W}(p,a'|b,p') + \#\mathscr{W}(p,b'|a,p') + C + 4D'.
\end{aligned}$$

By Lemma 4.32, at most K elements of $\mathscr{W}(a'|p')$ have label in $\gamma(\mathscr{W}(a|p))^\perp$. Since $\mathscr{W}(p|a)$ and $\mathscr{W}(p|b)$ are transverse, we deduce that $\#\mathscr{W}(p,a'|b,p') \leq K$ and, similarly, $\#\mathscr{W}(p,b'|a,p') \leq K$. We conclude that

$$d(p, \varphi(p)) = \#\mathscr{W}(p|p') \leq 2K + C + 4D'.$$

Lemma 2.34 gives $d(p, \operatorname{Fix}\varphi) \leq \zeta_2(2K + C + 4 \cdot \zeta_1^{-1}(D))$, as required by Definition 4.23. □

Corollary 4.34 *For every $\varphi \in U_0(\mathcal{A})$, the subgroup* $\operatorname{Fix}\varphi$ *is convex-cocompact in $\mathcal{A} \curvearrowright \mathcal{X}$.*

Proof Set $H := \mathrm{Fix}\,\varphi$. By Theorem B, H is finitely generated, so there exists $R \geq 0$ such that $\mathcal{N}_R(H)$ is edge-connected, viewed as a subset of $\mathcal{X}$. By Lemma 4.21, the median subalgebra $M := \langle \mathcal{N}_R(H) \rangle$ is edge-connected. Since H is an approximate median subalgebra by Lemma 2.35, Proposition 4.1 shows that M is at finite Hausdorff distance from H. Since H is weakly quasiconvex by Proposition 4.33, so is M.

Finally, $\mathcal{X}$ has finite staircase length by Lemma 4.17. We have shown that $M \subseteq \mathcal{X}^{(0)}$ is edge-connected and weakly quasiconvex, so Proposition 4.25 implies that M is quasiconvex. By Lemma 2.10, $\mathrm{Hull}\, M$ is at finite Hausdorff distance from M, which is at finite Hausdorff distance from H. This implies that H acts cocompactly on the convex subcomplex $\mathrm{Hull}\, M \subseteq \mathcal{X}$. $\square$

The discussion in this subsection immediately extends to right-angled Coxeter groups $\mathcal{W}$ and the finite-index subgroup $\mathrm{Aut}_0\, \mathcal{W} \leq \mathrm{Aut}\, \mathcal{W}$ generated by folds and partial conjugations.

Corollary 4.35 *For every $\varphi \in \mathrm{Aut}_0\, \mathcal{W}$, the subgroup $\mathrm{Fix}\,\varphi$ is convex-cocompact in $\mathcal{W} \curvearrowright \mathcal{Y}$, where $\mathcal{Y}$ is the universal cover of the Davis complex.*

Recalling Lemma 3.2 and Remark 2.31, the previous two corollaries prove Theorem C.

5 Invariant splittings of RAAGs

This section only contains the proofs of Proposition D and Corollary E, which are independent from all other results mentioned in the introduction.

Let Γ be a finite simplicial graph and let $\mathcal{A} = \mathcal{A}_\Gamma$ be the corresponding right-angled Artin group. All results and proofs in this section immediately extend to the right-angled Coxeter group $\mathcal{W}_\Gamma$ and automorphisms in $\mathrm{Aut}_0\, \mathcal{W}_\Gamma$. We encourage the reader to verify this as they go through the material, emphasising that only Lemmas 5.3 and 5.4 and Corollary 5.10 require any kind of attention, as all other results in this section are purely about the finite graph Γ.

The following is Proposition D from the introduction.

Proposition 5.1 *Let $\mathcal{A}$ be directly irreducible, freely irreducible and noncyclic. Then there exists an amalgamated product splitting $\mathcal{A} = \mathcal{A}_+ *_{\mathcal{A}_0} \mathcal{A}_-$, with $\mathcal{A}_\pm$ and $\mathcal{A}_0$ parabolic subgroups of $\mathcal{A}$, such that the corresponding Bass–Serre tree $\mathcal{A} \curvearrowright T$ is $U_0(\mathcal{A})$–invariant. That is, for every $\varphi \in U_0(\mathcal{A})$, there exists an isometry $f : T \to T$ satisfying $f \circ g = \varphi(g) \circ f$ for all $g \in \mathcal{A}$.*

Proposition 5.1 follows from Corollary 5.4 and Proposition 5.5 below. The latter will be proved right after Lemma 5.9.

Given a partition $\Gamma^{(0)} = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$, we write $\mathcal{A}_+ := \mathcal{A}_{\Lambda \sqcup \Lambda^+}$ and $\mathcal{A}_- := \mathcal{A}_{\Lambda \sqcup \Lambda^-}$ for simplicity. If $\Lambda^{\pm}$ are nonempty and $d(\Lambda^+, \Lambda^-) \geq 2$ (where d denotes the graph metric on Γ), then the partition corresponds to a splitting as amalgamated product,

$$\mathcal{A} = \mathcal{A}_+ *_{\mathcal{A}_\Lambda} \mathcal{A}_-.$$

We denote by $\mathcal{A} \curvearrowright T_\Lambda$ the Bass–Serre tree of this splitting. This will not cause any ambiguity related to possible different choices of the sets $\Lambda^{\pm}$ in the following discussion.

We are interested in partitions of $\Gamma^{(0)}$ that satisfy a certain list of properties.

Definition 5.2 A partition $\Gamma^{(0)} = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$ into three nonempty subsets is *good* if:

(i) $d(\Lambda^+, \Lambda^-) \geq 2$, where d is the graph metric on Γ.

(ii) For every $\epsilon \in \{\pm\}$ and $w \in \Lambda^\epsilon$, there does not exist $v \in \Lambda \sqcup \Lambda^{-\epsilon}$ with $\operatorname{lk} v \subseteq \operatorname{lk} w \cup \Lambda^\epsilon$.

(iii) For every $\epsilon \in \{\pm\}$ and $w \in \Lambda^\epsilon$, the subgraph of Γ spanned by $(\Lambda \sqcup \Lambda^{-\epsilon}) \setminus \operatorname{st} w$ is connected.

We will simply write $\Gamma = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$, rather than $\Gamma^{(0)} = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$.

The motivation for Definition 5.2 comes from the next lemma and the subsequent corollary. Definition 5.2 actually contains slightly stronger requirements than what is strictly necessary to the two results: this will facilitate the inductive construction of good partitions of graphs Γ.

Lemma 5.3 *Let $\Gamma = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$ be a good partition. For every $\psi \in U_0(\mathcal{A})$, there exists $\varphi \in U_0(\mathcal{A})$ representing the same outer automorphism and simultaneously satisfying $\varphi(\mathcal{A}_+) = \mathcal{A}_+$ and $\varphi(\mathcal{A}_-) = \mathcal{A}_-$ (hence also $\varphi(\mathcal{A}_\Lambda) = \mathcal{A}_\Lambda$).*

Proof Inversions preserve $\mathcal{A}^+$ and $\mathcal{A}^-$. Given vertices $v, w \in \Gamma$ with $\operatorname{lk} v \subseteq \operatorname{lk} w$, condition (ii) implies that either $w \in \Lambda$, or $\{v, w\} \subseteq \Lambda^+$, or $\{v, w\} \subseteq \Lambda^-$. Thus, folds also preserve $\mathcal{A}^+$ and $\mathcal{A}^-$.

We are left to prove the lemma in the case when ψ is a partial conjugation $\kappa_{w,C}$. If $w \in \Lambda$, it is clear that $\kappa_{w,C}$ preserves $\mathcal{A}^+$ and $\mathcal{A}^-$. Thus, let us assume without loss of generality that $w \in \Lambda^+$. By condition (iii), the set $\Lambda \cup \Lambda^-$ intersects a unique connected component $K \subseteq \Gamma \setminus \operatorname{st} w$.

If $K \neq C$, then $\kappa_{w,C}$ is the identity on $\mathcal{A}^-$, so $\mathcal{A}^{\pm}$ are both preserved. If $K = C$, then $\kappa_{w,C}$ represents the same outer automorphism as $\kappa_{w^{-1},K_1} \cdots \kappa_{w^{-1},K_k}$, where $K_1, \dots, K_k$ are the connected components of $\Gamma \setminus \operatorname{st} w$ other than K. Again, the latter is the identity on $\mathcal{A}^-$, so $\mathcal{A}^{\pm}$ are preserved. □

This shows that T_Λ is invariant under twisting by elements of $U_0(\mathcal{A})$:

Corollary 5.4 *Let $\Gamma = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$ be a good partition. For every $\varphi \in U_0(\mathcal{A})$, there exists an automorphism $f : T_\Lambda \to T_\Lambda$ satisfying $f \circ g = \varphi(g) \circ f$ for all $g \in \mathcal{A}$.*

Proof If φ is inner, we can take f to coincide with an element of $\mathcal{A}$. If $\varphi(\mathcal{A}_+) = \mathcal{A}_+$ and $\varphi(\mathcal{A}_-) = \mathcal{A}_-$, the statement is also clear, since the Bass–Serre tree can be defined in terms of cosets of $\mathcal{A}^{\pm}$. By Lemma 5.3, every element of $U_0(\mathcal{A})$ is a product of two automorphisms of these two types. □

Our next goal is to show that good partitions (almost) always exist. We say that Γ is *irreducible* if it does not split as a nontrivial join (equivalently, the opposite graph Γ^o is connected).

Proposition 5.5 *If Γ is connected, irreducible and not a singleton, then Γ admits a good partition.*

Proposition 5.5 and Corollary 5.4 immediately imply Proposition 5.1, as well as the analogous result for right-angled Coxeter groups.

Before proving Proposition 5.5, we need to obtain a few lemmas.

Lemma 5.6 *If Γ is connected and* $\operatorname{diam}\Gamma^{(0)} \geq 3$, *there exists a good partition of Γ.*

Proof Let $x, y \in \Gamma$ be arbitrary vertices with $d(x, y) \geq 3$. Let C_y be the connected component of $\Gamma \setminus \operatorname{st} x$ that contains y. Similarly, let C_x be the connected component of $\Gamma \setminus \operatorname{st} y$ that contains x.

Since $d(x, y) \geq 3$, we have $\operatorname{st} x \cap \operatorname{st} y = \varnothing$, hence $\operatorname{st} y \subseteq C_y$ and $\operatorname{st} x \subseteq C_x$. Since Γ is connected, $\Gamma \setminus C_x$ is also connected. Note that $\operatorname{st} x$ and $\Gamma \setminus C_x$ are disjoint and $y \in \Gamma \setminus C_x$. This implies that $\Gamma \setminus C_x \subseteq C_y$. In conclusion, $\Gamma = C_x \cup C_y$.

Note that, if $z \in \Gamma^{(0)}$ and $\operatorname{lk} z \cap C_y = \varnothing$, we cannot have $z \in C_y$. Indeed, this would imply that $C_y = \{z\}$ and $\operatorname{lk} z \subseteq \operatorname{st} x$. Since $y \in C_y$, we would then have $y = z$ and $\operatorname{lk} y \subseteq \operatorname{st} x$, contradicting the fact that Γ is connected and $d(x, y) \geq 3$.

Thus, we can define

$$\begin{aligned}\Lambda^+ &:= \{z \in \Gamma^{(0)} \mid \operatorname{st} z \cap C_y = \varnothing\} = \{z \in \Gamma^{(0)} \mid \operatorname{lk} z \cap C_y = \varnothing\},\\ \Lambda^- &:= \{z \in \Gamma^{(0)} \mid \operatorname{st} z \cap C_x = \varnothing\} = \{z \in \Gamma^{(0)} \mid \operatorname{lk} z \cap C_x = \varnothing\},\\ \Lambda &:= \Gamma^{(0)} \setminus (\Lambda^+ \sqcup \Lambda^-).\end{aligned}$$

Note that $x \in \Lambda^+$ and $y \in \Lambda^-$. If $z \in \Lambda^+$ and $w \in \Lambda^-$, we have $\operatorname{st} z \cap \operatorname{st} w = \varnothing$, since $\Gamma = C_x \cup C_y$. This shows that $d(\Lambda^+, \Lambda^-) \geq 3$. Since Γ is connected, we also conclude that $\Lambda \neq \varnothing$. We are left to verify conditions (ii) and (iii) of Definition 5.2.

If $v \in \Lambda$, then $\operatorname{lk} v$ intersects both C_x and C_y. Since C_y is disjoint from $\operatorname{lk} w \cup \Lambda^+$ for every $w \in \Lambda^+$ (and similarly with C_x and Λ^-), this implies condition (ii) when $v \in \Lambda$. On the other hand, the case with $v \in \Lambda^{-\epsilon}$ is immediate from the fact that $d(\Lambda^+, \Lambda^-) \geq 3$ and Γ is connected.

Finally, let us check condition (iii). Without loss of generality, we can suppose that $w \in \Lambda^+$. Note that C_y is connected, contained in $(\Lambda \sqcup \Lambda^-) \setminus \operatorname{st} w$, and it intersects the link of every point of Λ. Moreover, since $\Lambda^- \cap C_x = \varnothing$ and $\Gamma = C_x \cup C_y$, we have $\Lambda^- \subseteq C_y$. This shows that $(\Lambda \sqcup \Lambda^-) \setminus \operatorname{st} w$ is connected, concluding the proof. □

When the previous lemma cannot be applied, we will construct a good partition of Γ inductively, extending good partitions on subgraphs. We now prove a sequence of three lemmas aimed precisely at this, after which we will give the argument for Proposition 5.5.

For $x \in \Gamma^{(0)}$, let $\Gamma \setminus x$ be the graph obtained by removing x and all open edges incident to x.

Lemma 5.7 *Let $\Gamma \setminus x = \Delta^+ \sqcup \Delta \sqcup \Delta^-$ be a good partition. Then one of the following happens*:

1. *There exist $w \in \Delta^+$ and $z \in \Delta^-$ with $\operatorname{lk} x \subseteq \operatorname{lk} z \cap \operatorname{lk} w$.*
2. *The partition of Γ with $\Lambda^+ = \Delta^+ \sqcup \{x\}$, $\Lambda = \Delta$, $\Lambda^- = \Delta^-$ is good.*
3. *The partition of Γ with $\Lambda^+ = \Delta^+$, $\Lambda = \Delta \sqcup \{x\}$, $\Lambda^- = \Delta^-$ is good.*
4. *The partition of Γ with $\Lambda^+ = \Delta^+$, $\Lambda = \Delta$, $\Lambda^- = \Delta^- \sqcup \{x\}$ is good.*

Proof We begin with the following observation:

Claim *If there exists $w \in \Delta^+$ such that $\operatorname{lk} x \subseteq \operatorname{lk} w \cup \Delta^+$, we are either in case* (1) *or in case* (2).

Proof We assume that we are not in case (1) and show that the partition of Γ in case (2) is good. We need to verify conditions (i)–(iii) from Definition 5.2.

Since $d(\Delta^+, \Delta^-) \geq 2$ (both in $\Gamma \setminus x$ and in Γ), the set Δ^- is disjoint from $\operatorname{lk} w \cup \Delta^+$. Since $\operatorname{lk} x \subseteq \operatorname{lk} w \cup \Delta^+$, it follows that $\Delta^- \cap \operatorname{st} x = \varnothing$, hence $d(\Lambda^+, \Lambda^-) \geq 2$. This proves condition (i).

If condition (ii) fails, there exist $u \in \Lambda^\epsilon$ and $v \in \Lambda \sqcup \Lambda^{-\epsilon}$ with $\operatorname{lk} v \subseteq \operatorname{lk} u \cup \Lambda^\epsilon$. Since the partition of $\Gamma \setminus x$ is good, we must have either $v = x$ or $u = x$. If $v = x$, then $u \in \Delta^-$ and

$$\operatorname{lk} x \subseteq (\operatorname{lk} u \cup \Delta^-) \cap (\operatorname{lk} w \cup \Delta^+) = \operatorname{lk} u \cap \operatorname{lk} w,$$

which lands us in case (1). If instead $u = x$, we have $v \in \Delta \sqcup \Delta^-$ with

$$\operatorname{lk} v \subseteq \operatorname{lk} x \cup \Lambda^+ \subseteq \operatorname{lk} w \cup \Delta^+ \cup \{x\}.$$

This violates condition (ii) for the partition of $\Gamma \setminus x$.

Finally, suppose that condition (iii) fails. Thus, there exists $u \in \Lambda^\epsilon$ such that $(\Lambda \sqcup \Lambda^{-\epsilon}) \setminus \operatorname{st} u$ is disconnected. Since the partition of $\Gamma \setminus x$ is good, this can happen only in two ways: either $u = x$, or $u \in \Lambda^-$ and x is isolated in $(\Lambda \sqcup \Lambda^+) \setminus \operatorname{st} u$. In the latter case, we have $\operatorname{lk} x \subseteq \operatorname{lk} u \cup \Delta^-$, which again leads to case (1).

Suppose instead that $u = x$ and let us show that $(\Lambda \sqcup \Lambda^-) \setminus \operatorname{st} x = (\Delta \sqcup \Delta^-) \setminus \operatorname{lk} x$ is connected. Since $\operatorname{lk} x \subseteq \operatorname{lk} w \cup \Delta^+$, the set $(\Delta \sqcup \Delta^-) \setminus \operatorname{lk} x$ contains $(\Delta \sqcup \Delta^-) \setminus \operatorname{lk} w$. The latter is connected, as the partition of $\Gamma \setminus x$ satisfies condition (iii). Since condition (ii) is satisfied, every point of $(\Delta \sqcup \Delta^-) \cap \operatorname{lk} w = \Delta \cap \operatorname{lk} w$ is joined by an edge to a point of $\Gamma \setminus (\operatorname{lk} w \cup \Lambda^+) = (\Delta \sqcup \Delta^-) \setminus \operatorname{lk} w$. Thus, the star of every point of $(\Delta \sqcup \Delta^-) \setminus \operatorname{lk} x$ intersects the connected set $(\Delta \sqcup \Delta^-) \setminus \operatorname{lk} w$, proving that $(\Delta \sqcup \Delta^-) \setminus \operatorname{lk} x$ is connected. This completes the proof of the claim. $\triangleleft$

By the claim, if there exist either $w \in \Delta^+$ with $\mathrm{lk}\, x \subseteq \mathrm{lk}\, w \cup \Delta^+$ or $z \in \Delta^-$ with $\mathrm{lk}\, x \subseteq \mathrm{lk}\, z \cup \Delta^-$, then we are in cases (1), (2) or (4). In order to conclude the proof of the lemma, let us suppose that neither of the two inclusions is satisfied. We will show that the partition in case (3) is good.

Condition (i) is clear. Condition (ii) is immediate from the corresponding condition for $\Gamma \setminus x$ and our assumption that $\mathrm{lk}\, x$ be not contained in any subsets as in the previous paragraph.

Suppose that condition (iii) fails. Then there exists $u \in \Lambda^\epsilon$ such that $(\Lambda \sqcup \Lambda^{-\epsilon}) \setminus \mathrm{st}\, u$ is disconnected. Without loss of generality, we have $u \in \Lambda^+$. Since the partition of $\Gamma \setminus x$ satisfies condition (iii), the point x must be isolated in $(\Lambda \sqcup \Lambda^-) \setminus \mathrm{st}\, u$. Hence $\mathrm{lk}\, x \subseteq \mathrm{lk}\, u \cup \Delta^+$, again violating our assumption. □

Lemma 5.8 *Let Γ be an irreducible graph, and let $x \in \Gamma$ be a vertex such that there does not exist $y \in \Gamma^{(0)} \setminus \{x\}$ with $\mathrm{lk}\, x \subseteq \mathrm{lk}\, y$. Suppose that $\Gamma \setminus x$ is reducible. Then the partition of Γ given by $\Lambda^+ = \{x\}$, $\Lambda = \mathrm{lk}\, x$, $\Lambda^- = \Gamma \setminus \mathrm{st}\, x$ is good.*

Proof Write $\Gamma \setminus x$ as a join of nonempty subgraphs Γ_1 and Γ_2. Since Γ is irreducible, there exist points $a_1 \in \Gamma_1 \setminus \mathrm{lk}\, x$ and $a_2 \in \Gamma_2 \setminus \mathrm{lk}\, x$. Condition (i) is clear.

In order to verify condition (ii), we need to exclude the existence of $w \in \Lambda^\epsilon$ and $v \in \Lambda \sqcup \Lambda^{-\epsilon}$ with $\mathrm{lk}\, v \subseteq \mathrm{lk}\, w \cup \Lambda^\epsilon$. If $\epsilon = -$ and $v \in \Lambda$, then x lies in $\mathrm{lk}\, v$, but not in $\mathrm{lk}\, w \cup \Lambda^-$. If $\epsilon = -$ and $v = x$, then $\mathrm{lk}\, x$ is disjoint from Λ^-, and it cannot be contained in the link of any point of $\Gamma \setminus x$ by our hypotheses. If $\epsilon = +$, then $\mathrm{lk}\, w \cup \Lambda^\epsilon = \mathrm{st}\, x$, which cannot contain the link of any point of $\Gamma \setminus x$, as it does not contain a_1 and a_2.

Finally, let us show that, for every $w \in \Lambda^\epsilon$, the set $(\Lambda \sqcup \Lambda^{-\epsilon}) \setminus \mathrm{st}\, w$ is connected. If $\epsilon = +$, this amounts to showing that $\Gamma \setminus \mathrm{st}\, x$ is connected. This is immediate, since every point of $\Gamma \setminus x$ is joined by an edge to either a_1 or a_2, and these two points are themselves joined by an edge. If instead $\epsilon = -$, we need to show that $\mathrm{st}\, x \setminus \mathrm{st}\, w$ is connected for every $w \in \Gamma \setminus \mathrm{st}\, x$. This is also clear since this set is a cone over x. □

Consider the equivalence relation on $\Gamma^{(0)}$ where $v \sim w$ if and only if $\mathrm{lk}\, v = \mathrm{lk}\, w$. We define a graph $\overline{\Gamma}$ with a vertex for every $\sim$–equivalence class $[v] \subseteq \Gamma$ and an edge joining $[v]$ and $[w]$ exactly when v and w are joined by an edge (this is independent of the chosen representatives).

It is clear that $\overline{\Gamma}$ is again a simplicial graph, with at most as many vertices as Γ. We denote by $r\colon \Gamma \to \overline{\Gamma}$ the natural morphism of graphs.

Lemma 5.9 (1) *Γ is irreducible if and only if $\overline{\Gamma}$ is irreducible.*

(2) *If Γ has at least one edge, then Γ is connected if and only if $\overline{\Gamma}$ is connected.*

(3) *If $\overline{\Gamma} = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$ is a good partition, then so is $\Gamma = r^{-1}(\Lambda^+) \sqcup r^{-1}(\Lambda) \sqcup r^{-1}(\Lambda^-)$.*

Proof Parts (1) and (2) are straightforward, so we only prove part (3).

Consider a good partition $\overline{\Gamma} = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$. It is clear that the partition of Γ satisfies condition (i), while condition (ii) follows from the observation that $\mathrm{lk}\, r(x) = r(\mathrm{lk}\, x)$ for every $x \in \Gamma$.

Finally, we verify condition (iii). Given $w \in r^{-1}(\Lambda^{\epsilon})$, observe that r maps the subgraph
$$(r^{-1}(\Lambda) \sqcup r^{-1}(\Lambda^{-\epsilon})) \setminus \operatorname{st} w$$
onto the connected graph $(\Lambda \sqcup \Lambda^{-\epsilon}) \setminus \operatorname{st} r(w)$. As in part (2), this shows that $(r^{-1}(\Lambda) \sqcup r^{-1}(\Lambda^{-\epsilon})) \setminus \operatorname{st} w$ is connected, possibly except the case when $(\Lambda \sqcup \Lambda^{-\epsilon}) \setminus \operatorname{st} r(w)$ is a singleton. The latter is ruled out by the fact that the partition of $\overline{\Gamma}$ satisfies condition (ii). □

Proof of Proposition 5.5 We proceed by induction on the number of vertices of Γ. Since no graph with at most 3 vertices satisfies the hypotheses of the proposition, the base step is trivially satisfied. For the inductive step, we consider a connected irreducible graph Γ with at least 4 vertices, and assume that the proposition is satisfied by all graphs with fewer vertices than Γ.

If $\operatorname{diam} \Gamma^{(0)} \geq 3$, we can simply appeal to Lemma 5.6. If the graph $\overline{\Gamma}$ defined above has fewer vertices than Γ, then we can use the inductive hypothesis and Lemma 5.9. Thus, we can assume that $\Gamma = \overline{\Gamma}$ and $\operatorname{diam} \Gamma^{(0)} = 2$.

Pick a vertex $x \in \Gamma$ whose link is maximal under inclusion. Since $\Gamma = \overline{\Gamma}$, there does not exist $y \in \Gamma^{(0)} \setminus \{x\}$ with $\operatorname{lk} x = \operatorname{lk} y$. If $\Gamma \setminus x$ is reducible, Lemma 5.8 then shows that Γ admits a good partition. If $\Gamma \setminus x$ were disconnected, then the fact that $\operatorname{diam} \Gamma^{(0)} = 2$ would imply that $\operatorname{lk} x = \Gamma \setminus x$, contradicting the assumption that Γ is irreducible.

In conclusion, $\Gamma \setminus x$ is connected, irreducible, not a singleton, and it has fewer vertices than Γ. We conclude by applying the inductive hypothesis and Lemma 5.7 (case (1) of the latter is ruled out by our choice of x). □

The previous results prove Proposition 5.1. The following is Corollary E from the introduction:

Corollary 5.10 *Consider $\varphi \in U_0(\mathcal{A})$.*

1. *If $\mathcal{A}$ splits as a direct product $\mathcal{A}_1 \times \mathcal{A}_2$, then $\varphi(\mathcal{A}_i) = \mathcal{A}_i$ and $\operatorname{Fix} \varphi = \operatorname{Fix} \varphi|_{\mathcal{A}_1} \times \operatorname{Fix} \varphi|_{\mathcal{A}_2}$.*
2. *If $\mathcal{A}$ is directly irreducible, then the subgroup $\operatorname{Fix} \varphi \leq \mathcal{A}$ splits as a (possibly trivial) finite graph of groups with vertex and edge groups of the form $\operatorname{Fix} \varphi|_P$, for proper parabolic subgroups $P \leq \mathcal{A}$ with $\varphi(P) = P$ and $\varphi|_P \in U_0(P)$.*

Proof For simplicity, set $H := \operatorname{Fix} \varphi$. We distinguish three cases.

Case 1 ($\mathcal{A}$ is not directly irreducible) Let us write $\mathcal{A} = A \times \mathcal{A}_1 \times \cdots \times \mathcal{A}_m$, where A is a free abelian group and $\mathcal{A}_i$ are directly irreducible (noncyclic) right-angled Artin groups. This corresponds to a splitting of Γ as a join of a complete subgraph and irreducible subgraphs $\Gamma_1, \dots, \Gamma_m$.

Since $\varphi \in U_0(\mathcal{A})$, we have $\varphi(\mathcal{A}_k) = \mathcal{A}_k$ and $\varphi|_{\mathcal{A}_k} \in U_0(\mathcal{A}_k)$ for every $1 \leq k \leq m$, and $\varphi|_A$ is a product of inversions. Indeed, this is clear for inversions, folds and partial conjugations.

Thus $H = A' \times H_1 \times \cdots \times H_m$, where $H_i = \operatorname{Fix}(\varphi|_{\mathcal{A}_i})$ and A' is a standard direct factor of A. This proves part (1) of the corollary.

Case 2 ($\mathcal{A}$ is not freely irreducible) Write $\mathcal{A} = F * \mathcal{A}_1 * \dots * \mathcal{A}_m$, where F is a free group and $\mathcal{A}_i$ are freely irreducible (noncyclic) right-angled Artin groups of lower complexity. Since H is finitely generated by Theorem B, Kurosh's theorem guarantees that H decomposes as a free product $H = L * H_1 * \dots * H_n$, where L is a finitely generated free group and each H_i is a finitely generated subgroup of some $g_i \mathcal{A}_{k_i} g_i^{-1}$ with $g_i \in \mathcal{A}$ and $1 \le k_i \le m$.

By Grushko's theorem, the subgroup $\varphi(\mathcal{A}_k)$ is conjugate to some $\mathcal{A}_{k'}$ for every $1 \le k \le m$. Since φ fixes the nontrivial subgroup $H_i \le g_i \mathcal{A}_{k_i} g_i^{-1}$ pointwise, we must have $\varphi(g_i \mathcal{A}_{k_i} g_i^{-1}) = g_i \mathcal{A}_{k_i} g_i^{-1}$ for $1 \le i \le n$.

Consider the automorphism $\psi_i \in U_0(\mathcal{A})$ defined by $\psi_i(x) = g_i^{-1}\varphi(g_i x g_i^{-1}) g_i$. Note that $\psi_i(\mathcal{A}_{k_i}) = \mathcal{A}_{k_i}$ and $\operatorname{Fix} \psi_i|_{\mathcal{A}_{k_i}} = g_i^{-1} H_i g_i$. By Lemma 3.29, we have $\psi_i|_{\mathcal{A}_{k_i}} \in U_0(\mathcal{A}_{k_i})$. This proves part (2) of the corollary in the freely reducible case.

Case 3 ($\mathcal{A}$ is freely and directly irreducible) We can assume that $\mathcal{A} \not\simeq \mathbb{Z}$. By Proposition 5.5, Γ admits a good partition $\Gamma = \Lambda^+ \sqcup \Lambda \sqcup \Lambda^-$. By Corollary 5.4, there exists $f \in \operatorname{Aut} T_\Lambda$ satisfying $f \circ g = \varphi(g) \circ f$ for all $g \in \mathcal{A}$.

If H is elliptic in T_Λ, we have $H \le V$, where V is the $\mathcal{A}$–stabiliser of some vertex of T_Λ. The existence of the automorphism $f \in \operatorname{Aut} T_\Lambda$ guarantees that all subgroups $\varphi^n(V)$ with $n \in \mathbb{Z}$ are $\mathcal{A}$–stabilisers of vertices of T_Λ; in particular, they are all conjugate to either $\mathcal{A}_+$ or $\mathcal{A}_-$. We conclude that H is contained in the $\langle\varphi\rangle$–invariant parabolic subgroup $P := \bigcap_{n \in \mathbb{Z}} \varphi^n(V)$. Thus, we have $H = \operatorname{Fix} \varphi|_P$ and, by Lemma 3.29, $\varphi|_P \in U_0(P)$. This proves the corollary in this case, with H splitting as a trivial graph of groups.

Suppose instead that H is not elliptic in T_Λ and denote by $T_H \subseteq T_\Lambda$ the H–minimal subtree. Since H is finitely generated, the action $H \curvearrowright T_H$ is cocompact and gives a splitting of H as a (nontrivial) finite graph of groups. We are left to understand vertex-stabilisers of the action $H \curvearrowright T_H$.

As f normalises H in $\operatorname{Aut} T_\Lambda$, we have $f(T_H) = T_H$. It is convenient to distinguish two subcases.

Case 3a (f is elliptic in T_Λ) Since f commutes with every element of H, the tree T_H is fixed pointwise by f. For every $v \in T_H$, its $\mathcal{A}$–stabiliser $\mathcal{A}_v$ satisfies $\varphi(\mathcal{A}_v) = \mathcal{A}_v$ and is conjugate to either $\mathcal{A}_+$ or $\mathcal{A}_-$. By Lemma 3.29, we have $\varphi|_{\mathcal{A}_v} \in U_0(\mathcal{A}_v)$, proving the corollary in this case.

Case 3b (f is loxodromic in T_Λ) Let $\alpha \subseteq T_\Lambda$ be the axis of f. Since f commutes with every element of H, the geodesic α must be H–invariant and every nonloxodromic element of H fixes α pointwise. Note that T_H cannot be a singleton, or f would be elliptic. Thus, $T_H = \alpha$ and H contains a shortest loxodromic element $h \in H$. Moreover, $H = H_0 \rtimes \langle h \rangle$, where H_0 is the kernel of the action $H \curvearrowright \alpha$.

Let $Q \le \mathcal{A}$ be the intersection of the $\mathcal{A}$–stabilisers of the vertices of α. Being an intersection of parabolic subgroups, Q is itself a (possibly trivial) parabolic subgroup of $\mathcal{A}$. Since $f(\alpha) = \alpha$, we have $\varphi(Q) = Q$ and $H_0 = \operatorname{Fix} \varphi|_Q$. Lemma 3.29 guarantees that $\varphi|_Q \in U_0(Q)$. Thus, the HNN splitting $H = H_0 \rtimes \langle h \rangle$ is as required by the corollary. □

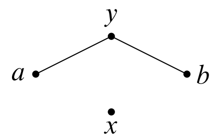

Figure 5

Remark 5.11 In Case 3b of the proof of Corollary 5.10, we can actually say more on the structure of $H = \mathrm{Fix}\,\varphi$. Specifically, $H = H_0 \times \langle h \rangle$ and h can be taken to be label-irreducible.

Indeed, since $h\alpha = \alpha$, the element h lies in the normaliser of Q in $\mathcal{A}$, which is a subgroup of the form $Q \times Q'$ (since Q is parabolic in $\mathcal{A}$). If $h = h_1 \cdots h_k$ is the decomposition of h into label-irreducible components, every h_i lies in either Q or Q'. Since φ is coarse-median preserving and fixes h, it must permute the h_i; Corollary 4.31 then shows that $\varphi(h_i) = h_i$ for every i. Thus, all the label-irreducible components of h that lie in Q actually lie in H_0. Up to replacing h, we can assume that all h_i lie in Q'; in particular, h lies in Q', hence it commutes with H_0. Since $H = \mathrm{Fix}\,\varphi$ is generated by H_0 and h, we must then have $k = 1$, ie h is label-irreducible.

In relation to Theorem C, it is natural to wonder if the proof of Corollary 5.10 can be used to give an alternative, inductive argument showing that $\mathrm{Fix}\,\varphi$ is convex-cocompact in $\mathcal{A}$ for every $\varphi \in U_0(\mathcal{A})$. In light of Remark 5.11, the only problematic situation is the one in Case 3a.

Unfortunately, cubical convex-cocompactness does not seem to be well-behaved with respect to graph-of-groups constructions, as the next example shows.

Example 5.12 Let Γ be the graph in Figure 5. Consider the subgroup $H = \langle ayx^{-1}, xby \rangle \leq \mathcal{A}_\Gamma$. We have an amalgamated product splitting $\mathcal{A}_\Gamma = \langle a, x, y \rangle *_{\langle x,y \rangle} \langle b, x, y \rangle$, which induces a splitting $H = \langle ayx^{-1} \rangle * \langle xby \rangle \simeq F_2$. The subgroups $\langle ayx^{-1} \rangle$ and $\langle xby \rangle$ are convex-cocompact, as they are each generated by a single label-irreducible element.

However, H is not convex-cocompact in $\mathcal{A}$: the element aby^2 lies in H, but no power of its label-irreducible components ab and y^2 does (which, for instance, violates Lemma 3.16).

6 Projectively invariant metrics on finite-rank median algebras

In this section, we initiate the lengthy proof of Theorem F, which will be completed in Section 7. Our main goal here is to formulate a criterion, for a group U and a subgroup $G \leq U$, guaranteeing that a U–action on a finite-rank median algebra admits a G–invariant compatible pseudometric for which U acts by homotheties (Corollary 6.23). An important tool will be the Lefschetz fixed point theorem for compact ANRs.

Throughout the section, M denotes a fixed median algebra of finite rank r.

6.1 Multibridges

The *bridge* of two gate-convex sets was first studied in [4; 30] for CAT(0) cube complexes and in [49, Section 2.2] for general median algebras. We will need an extension of this concept to arbitrary finite collections of gate-convex subsets: *multibridges*.

We briefly motivate why. As a recurring setup in the rest of the paper (especially in Sections 6.2.3 and 7.4), we will often find ourselves studying a group $G \leq \operatorname{Aut} M$ with a finite generating subset $S \subseteq G$ and a G–invariant compatible pseudometric η on M. It will be important to understand which points of M are moved as little as possible by all elements of S, ie which points realise the quantity $\overline{\tau}_S^{\eta}$ from Section 2.1. It turns out that the set of such points does not depend much on the specific pseudometric η, and can instead be characterised purely in terms of the median-algebra structure on M, using the notion of multibridge (Propositions 6.9 and 6.11).

Let $C_1, \dots, C_k \subseteq M$ be gate-convex subsets, with gate-projections $\pi_i \colon M \to C_i$. Let $\mathcal{H} \subseteq \mathscr{H}(M)$ be the set of halfspaces that contain at least one C_i and intersect each C_i. Then we have a partition

$$\mathscr{H}(M) = (\mathcal{H} \sqcup \mathcal{H}^*) \sqcup \Big(\bigcap_{1 \leq i \leq k} \mathscr{H}_{C_i}(M) \Big) \sqcup \Big(\bigcup_{1 \leq i, j \leq k} \mathscr{H}(C_i | C_j) \Big).$$

If $i \neq j$, the sets $\mathscr{H}_{C_i}(M) \cap \mathscr{H}_{C_j}(M)$ and $\mathscr{H}(C_i|C_j)$ are transverse. Thus, every halfspace in the second set of the above partition of $\mathscr{H}(M)$ is transverse to every halfspace in the third set.

Lemma 6.1 *The intersection of all halfspaces in $\mathcal{H}$ is a nonempty convex subset of M.*

Proof We will prove this by appealing to Lemma 2.4(1). It is clear that the elements of $\mathcal{H}$ intersect pairwise. Let us show that, for every chain $\mathscr{C} \subseteq \mathcal{H}$, the set $\mathfrak{k} := \bigcap \mathscr{C}$ is again an element of $\mathcal{H}$.

Note that there exist $1 \leq i_0 \leq k$ and a cofinal subset $\mathscr{C}' \subseteq \mathscr{C}$ consisting of halfspaces containing C_{i_0}. Thus, $C_{i_0} \subseteq \mathfrak{k}$ and $\mathfrak{k}$ is nonempty. Since $\mathfrak{k}$ is the intersection of a chain of halfspaces, both $\mathfrak{k}$ and $\mathfrak{k}^*$ are convex. It follows that $\mathfrak{k}$ is a halfspace of M.

For every $\mathfrak{h} \in \mathscr{C} \subseteq \mathcal{H}$, the fact that $\mathfrak{h}$ intersects each C_i implies that $\pi_i(\mathfrak{h}) = \mathfrak{h} \cap C_i$; see for instance [50, Lemma 2.2(1)]. Recalling that $\mathfrak{k} = \bigcap \mathscr{C}$, we deduce that $\pi_i(\mathfrak{k}) \subseteq \mathfrak{k} \cap C_i$ for $1 \leq i \leq k$, hence $\mathfrak{k}$ intersects all C_i. Since we have already seen that $C_{i_0} \subseteq \mathfrak{k}$, we conclude that $\mathfrak{k} \in \mathcal{H}$, as required. □

Definition 6.2 The intersection $\mathcal{B} = \mathcal{B}(C_1, \dots, C_k) \subseteq M$ of all halfspaces in $\mathcal{H}$ is the *multibridge* of the gate-convex sets $C_1, \dots, C_k$.

For every $\mathfrak{k} \in \mathscr{H}(M) \setminus \mathcal{H}^*$, the set $\mathcal{H} \sqcup \{\mathfrak{k}\}$ is again pairwise-intersecting. Hence, Lemma 2.4(1) yields

$$\mathscr{H}_{\mathcal{B}}(M) = \mathscr{H}(M) \setminus (\mathcal{H} \sqcup \mathcal{H}^*) = \Big(\bigcap \mathscr{H}_{C_i}(M) \Big) \sqcup \Big(\bigcup \mathscr{H}(C_i | C_j) \Big).$$

We have already observed that the two sets in this partition are transverse. By Remark 2.2(2) and Lemma 2.6, we obtain a natural product splitting

$$\mathcal{B} = \mathcal{B}_{/\!/} \times \mathcal{B}_{\perp}, \quad \text{where } \mathscr{H}_{\mathcal{B}_{/\!/}}(M) = \bigcap \mathscr{H}_{C_i}(M) \quad \text{and} \quad \mathscr{H}_{\mathcal{B}_{\perp}}(M) = \bigcup \mathscr{H}(C_i|C_j).$$

We can view $\mathcal{B}_{/\!/}$ and $\mathcal{B}_{\perp}$ as subsets of M by identifying them with any fibre of the splitting of $\mathcal{B}$.

Lemma 6.3 *The sets $\mathcal{B}$, $\mathcal{B}_{/\!/}$ and $\mathcal{B}_{\perp}$ are gate-convex in M.*

Proof Since each C_i is gate-convex, Lemma 2.4(2) shows that, for every chain $\mathscr{C} \subseteq \bigcap \mathscr{H}_{C_i}(M)$, either $\bigcap \mathscr{C}$ is empty in M, or $\bigcap \mathscr{C} \in \bigcap \mathscr{H}_{C_i}(M)$. Hence $\mathcal{B}_{/\!/}$ is gate-convex in M.

If $\mathscr{C} \subseteq \bigcup \mathscr{H}(C_i|C_j)$ is a chain, a cofinal subset of $\mathscr{C}$ is contained in a single $\mathscr{H}(C_i|C_j)$. Hence $\bigcap \mathscr{C} \in \mathscr{H}(C_i|C_j)$. Invoking again Lemma 2.4(2), this shows that $\mathcal{B}_{\perp}$ is gate-convex.

Every chain in $\mathscr{H}_{\mathcal{B}}(M)$ has a cofinal subset contained in either $\bigcap \mathscr{H}_{C_i}(M)$ or $\bigcup \mathscr{H}(C_i|C_j)$. One last application of Lemma 2.4(2) shows that $\mathcal{B}$ is gate-convex. □

Corollary 6.4 *If $C_1, \dots, C_k \subseteq M$ are gate-convex subsets, their multibridge $\mathcal{B} = \mathcal{B}(C_1, \dots, C_k)$ is a gate-convex subset of M enjoying the following properties*:

1. *$\mathcal{B}$ splits as a product $\mathcal{B}_{/\!/} \times \mathcal{B}_{\perp}$ with $\mathscr{H}_{\mathcal{B}_{/\!/}}(M) = \bigcap \mathscr{H}_{C_i}(M)$ and $\mathscr{H}_{\mathcal{B}_{\perp}}(M) = \bigcup \mathscr{H}(C_i|C_j)$.*
2. *Each fibre $\{*\} \times \mathcal{B}_{\perp}$ intersects all of the C_i.*

Proof The only statement that has not already been proved is part (2). If it were false, there would exist an index i and $\mathfrak{h} \in \mathscr{H}(M)$ such that $C_i \subseteq \mathfrak{h}$ and $\{*\} \times \mathcal{B}_{\perp} \subseteq \mathfrak{h}^*$. Since $C_i \subseteq \mathfrak{h}$, we have $\mathfrak{h} \notin \mathscr{H}_{\mathcal{B}_{/\!/}}(M)$, so $\mathcal{B} \subseteq \mathfrak{h}^*$. Hence $\mathfrak{h}^* \in \mathcal{H}$, contradicting the fact that $C_i \subseteq \mathfrak{h}$. □

Recall the notation $\mathcal{PD}(M)$ and $\mathcal{D}(M)$ for compatible (pseudo)metrics, as in Section 2.3.

Remark 6.5 If $\eta \in \mathcal{PD}(M)$ and $x, y \in \mathcal{B}$ lie in the same fibre $\mathcal{B}_{/\!/} \times \{*\}$, then $\eta(x, C_i) = \eta(y, C_i)$ for all $1 \le i \le k$. Indeed, since $\mathscr{H}(x|y) \subseteq \mathscr{H}_{\mathcal{B}_{/\!/}}(M) = \bigcap \mathscr{H}_{C_i}(M)$, we have $\mathscr{W}(x|C_i) = \mathscr{W}(y|C_i)$ and it follows (eg by Remark 2.9) that $\eta(x, \pi_i(x)) = \eta(y, \pi_i(y))$ for every $\eta \in \mathcal{PD}(M)$.

Remark 6.6 If $\eta \in \mathcal{PD}(M)$, then $\eta(x, \mathcal{B}) \le r \cdot \max_i \eta(x, C_i)$ for every $x \in M$.

In order to see this, let $\mathfrak{h}_1, \dots, \mathfrak{h}_k$ be the minimal elements of $\mathscr{H}(x|\mathcal{B})$. Since the $\mathfrak{h}_i$ are pairwise transverse and $\operatorname{rk} M = r$, we have $k \le r$. Each $\mathfrak{h}_i$ must lie in $\mathcal{H}$, hence there exists an index j_i such that $C_{j_i} \subseteq \mathfrak{h}_i$. It follows that:

$$\mathscr{H}(x|\mathcal{B}) \subseteq \bigcup \mathscr{H}(x|\mathfrak{h}_i) \subseteq \bigcup \mathscr{H}(x|C_{j_i}).$$

Hence $\eta(x, \mathcal{B}) \le k \cdot \max_i \eta(x, C_i) \le r \cdot \max_i \eta(x, C_i)$.

Remark 6.7 If $\delta \in \mathcal{D}(M)$ and (M,δ) is complete, then $\mathcal{B}_\perp$ is compact in (M,δ).

In order to prove this, let $x_{i,j} \in C_i$ and $x_{j,i} \in C_j$ be a pair of gates for all distinct $1 \le i, j \le k$. Let K be the convex hull of the finite set $F = \{x_{i,j} \mid 1 \le i, j \le k\}$. Recall that $K = \mathcal{J}^r(F)$ by Remark 2.5, so it follows from [50, Corollary 2.20] that K is compact.

We have $K \cap \mathcal{B} \neq \varnothing$. Otherwise, the set $\mathscr{H}(K|\mathcal{B})$ would be nonempty and contained in $\mathcal{H}$. However, each element of $\mathcal{H}$ contains some C_i, so it cannot be disjoint from K.

Finally, observing that $\mathscr{H}_K(M)$ contains the set

$$\bigcup \mathscr{H}(x_{i,j}|x_{j,i}) = \bigcup \mathscr{H}(C_i|C_j) = \mathscr{H}_{\mathcal{B}_\perp}(M),$$

we deduce that $K \cap \mathcal{B}$ must contain a fibre $\{*\} \times \mathcal{B}_\perp$. Since $\mathcal{B}_\perp$ is gate-convex, it must be a closed subset of K, hence it is compact too.

Now, let $S \subseteq \operatorname{Aut} M$ be a finite set of automorphisms acting nontransversely and stably without inversions. By Theorem 2.16(1), the reduced cores $\bar{\mathcal{C}}(s)$ of $s \in S$ are all gate-convex. Let $\mathcal{B}(S)$ be their multibridge.

Definition 6.8 We refer to $\mathcal{B}(S)$ as the *multibridge* of the finite set $S \subseteq \operatorname{Aut} M$.

Recalling the notation introduced in Section 2.1, we have:

Proposition 6.9 *Let $S \subseteq \operatorname{Aut} M$ be a finite set of automorphisms acting nontransversely and stably without inversions. The multibridge $\mathcal{B}(S)$ is gate-convex and, for all $\eta \in \mathcal{PD}(M)^{\langle S\rangle}$:*

1. *We have $\tau_S^\eta(\pi_{\mathcal{B}}(x)) \le \tau_S^\eta(x)$ for all $x \in M$, where $\pi_{\mathcal{B}}\colon M \to \mathcal{B}(S)$ is the gate-projection.*
2. *$\tau_S^\eta(\cdot)$ is constant on each fibre $\mathcal{B}_{/\!/}(S) \times \{*\}$.*
3. *If $\delta \in \mathcal{D}(M)^{\langle S\rangle}$ and (M,δ) is complete, then there exists $z \in \mathcal{B}(S)$ with $\tau_S^\delta(x) = \bar{\tau}_S^\delta$.*

Proof Since the multibridge $\mathcal{B}(S)$ intersects each $\bar{\mathcal{C}}(s)$, we have $\mathscr{H}(\pi_{\mathcal{B}}(x)|\bar{\mathcal{C}}(s)) \subseteq \mathscr{H}(x|\bar{\mathcal{C}}(s))$ for all $x \in M$. Hence $\eta(\pi_{\mathcal{B}}(x), \bar{\mathcal{C}}(s)) \le \eta(x, \bar{\mathcal{C}}(s))$. Theorem 2.16(2) now implies that $\tau_S^\eta(\pi_{\mathcal{B}}(x)) \le \tau_S^\eta(x)$, proving part (1). By Remark 6.5, if $x, y \in \mathcal{B}(S)$ lie in the same fibre $\mathcal{B}_{/\!/}(S) \times \{*\}$, then $\eta(x, \bar{\mathcal{C}}(s)) = \eta(y, \bar{\mathcal{C}}(s))$. This proves part (2). Finally, part (3) follows from the previous two parts and Remark 6.7. □

Example 6.10 Let $G = \langle a, b\rangle$ be the free group over two generators. Let T be the standard Cayley graph of G, with all edges of length 1. Let (X,δ) be the (incomplete) median space obtained by removing from T all midpoints of edges. Then, taking $S = \{a, bab^{-1}\} \subseteq G \subseteq \operatorname{Isom} X$, there is no point $x \in X$ with $\tau_S^\delta(x) = \bar{\tau}_S^\delta = 2$.

Our interest in multibridges is due to the following result, which helps us understand the behaviour on M of the functions $\tau_S^\eta(\cdot)$ for $\eta \in \mathcal{PD}(M)^{\langle S\rangle}$.

Proposition 6.11 *Let $S \subseteq \operatorname{Aut} M$ be a finite set of automorphisms acting nontransversely and stably without inversions. Recall that $r = \operatorname{rk} M$. Then, the following hold for every $\eta \in \mathcal{PD}(M)^{\langle S\rangle}$:*

(1) *If $s_1, s_2 \in S$, then $\eta(\bar{\mathcal{C}}(s_1), \bar{\mathcal{C}}(s_2)) \le \bar{\tau}^\eta_S$.*

(2) *If $s \in S$ and $x \in \mathcal{B}(S)$, then $\eta(x, \bar{\mathcal{C}}(s)) \le r\bar{\tau}^\eta_S$.*

(3) *If $x \in \mathcal{B}(S)$, then $\tau^\eta_S(x) \le (2r+1)\bar{\tau}^\eta_S$.*

(4) *The η–diameter of each fibre $\{*\} \times \mathcal{B}_\perp(S)$ is at most $r^2\bar{\tau}^\eta_S$.*

(5) *If $x \in M$, then $\eta(x, \mathcal{B}(S)) \le \frac{1}{2} r\tau^\eta_S(x)$.*

(6) *For any $x \in M$ and any fibre $P = \mathcal{B}_{/\!/}(S) \times \{*\}$, we have $\eta(x, P) \le 2r^2\tau^\eta_S(x)$.*

Proof We begin with part (1). For every $x \in M$, we have

$$\mathscr{W}(\bar{\mathcal{C}}(s_1)|\bar{\mathcal{C}}(s_2)) = \mathscr{W}(x, \bar{\mathcal{C}}(s_1)|\bar{\mathcal{C}}(s_2)) \sqcup \mathscr{W}(\bar{\mathcal{C}}(s_1)|\bar{\mathcal{C}}(s_2), x) \subseteq \mathscr{W}(x|\bar{\mathcal{C}}(s_1)) \sqcup \mathscr{W}(x|\bar{\mathcal{C}}(s_2)).$$

Along with Theorem 2.16(2), this implies that

$$\tfrac{1}{2}\eta(\bar{\mathcal{C}}(s_1), \bar{\mathcal{C}}(s_2)) \le \max\{\eta(x, \bar{\mathcal{C}}(s_1)), \eta(x, \bar{\mathcal{C}}(s_2))\} \le \tfrac{1}{2}\max\{\eta(x, s_1x), \eta(x, s_2x)\} \le \tfrac{1}{2}\tau^\eta_S(x).$$

Part (1) follows by taking an infimum over $x \in M$.

Let us prove part (2). If $x \in \mathcal{B}(S)$ and $s \in S$, Corollary 6.4(2) implies that $\mathscr{H}(x|\bar{\mathcal{C}}(s))$ is contained in the union of the sets $\mathscr{H}(\bar{\mathcal{C}}(t)|\bar{\mathcal{C}}(s))$ with $t \in S \setminus \{s\}$. The maximal halfspaces in $\mathscr{H}(x|\bar{\mathcal{C}}(s))$ are pairwise transverse, so there are at most r of them. Hence, there exist $t_1, \dots, t_r \in S$ such that $\Omega := \bigcup_i \mathscr{H}(\bar{\mathcal{C}}(t_i)|\bar{\mathcal{C}}(s))$ contains every maximal element of $\mathscr{H}(x|\bar{\mathcal{C}}(s))$. In particular, $\mathscr{H}(x|\bar{\mathcal{C}}(s)) \subseteq \Omega$ and part (1) yields $\eta(x, \bar{\mathcal{C}}(s)) \le r\bar{\tau}^\eta_S$.

Part (3) of the proposition now follows from Theorem 2.16(2):

$$\tau^\eta_S(x) = \max_{s\in S}[\ell(s,\eta) + 2\eta(x, \bar{\mathcal{C}}(s))] \le \max_{s\in S}[\bar{\tau}^\eta_S + 2r\bar{\tau}^\eta_S] = (2r+1)\bar{\tau}^\eta_S.$$

Regarding part (4), consider two points x, y lying in the same fibre $\{*\} \times \mathcal{B}_\perp(S)$. Let $\mathfrak{h}_1, \dots, \mathfrak{h}_k$ be the minimal elements of $\mathscr{H}(x|y)$. Since $\operatorname{rk} M = r$, we have $k \le r$. By definition of $\mathcal{B}_\perp(S)$, there exist elements $s_i \in S$ with $\bar{\mathcal{C}}(s_i) \subseteq \mathfrak{h}_i$. Thus,

$$\mathscr{H}(x|y) \subseteq \bigcup \mathscr{H}(x|\mathfrak{h}_i) \subseteq \bigcup \mathscr{H}(x|\bar{\mathcal{C}}(s_i)).$$

Using part (2) of the proposition, it follows that $\eta(x,y) \le k \cdot \max_s \eta(x, \bar{\mathcal{C}}(s)) \le kr\bar{\tau}^\eta_S \le r^2\bar{\tau}^\eta_S$.

Finally, part (5) is a consequence of Remark 6.6 and the fact, due to Theorem 2.16(2), that $\tau^\eta_S(x) \ge 2\eta(x, \bar{\mathcal{C}}(s))$ for every $s \in S$. Part (6) is obtained by combining parts (4) and (5):

$$\eta(x, P) \le \eta(x, \mathcal{B}(S)) + r^2\bar{\tau}^\eta_S \le \tfrac{r}{2}\tau^\eta_S(x) + r^2\bar{\tau}^\eta_S \le 2r^2\tau^\eta_S(x). \qquad \square$$

6.2 Promoting median automorphisms to homotheties

Recall that M is a median algebra of finite rank r. In this subsection, we consider subgroups $G \lhd U \leq \operatorname{Aut} M$, with the goal of constructing G–invariant compatible pseudometrics $\eta \in \mathcal{PD}^G(M)$ with respect to which U acts by homotheties and G is nonelliptic. In general, this will only be possible after passing to a subalgebra of M. The final result in this direction is Corollary 6.23.

Our main technical tools are the notion of multibridge (exploited in Lemma 6.22) and the Lefschetz fixed point theorem applied to projectivisations of certain cones $\mathcal{C}$ in the topological vector space $\mathcal{PD}^G(M)$ (Proposition 6.17). Some extra work is required in order to ensure that our cones $\mathcal{C}$ have compact projectivisation and that they only contain pseudometrics η for which G acts nonelliptically (ie $\overline{\tau}^{\eta}_S > 0$ for some/any generating set $S \subseteq G$).

6.2.1 Preliminaries on normed spaces and ARs

Definition 6.12 Let V be a real vector space.

1. A *cone* is a convex subset $\mathcal{C} \subseteq V$ that is closed under multiplication by scalars in $[0, +\infty)$.
2. A *positive cone* is a cone $\mathcal{C} \subseteq V$ for which $\mathcal{C} \setminus \{0\}$ is convex. Equivalently, $\mathcal{C} \cap (-\mathcal{C}) = \{0\}$.
3. The *projectivisation* $\mathbb{P}(\mathcal{C})$ of a cone $\mathcal{C}$ is the quotient of $\mathcal{C} \setminus \{0\}$ obtained by identifying points that differ by multiplication by a scalar.

Given a countable probability space (Ω, σ) and a function $f : \Omega \to \mathbb{R}$, recall that

$$\|f\|_1 = \sum_{\omega \in \Omega} |f(\omega)|\sigma(\omega) \quad \text{and} \quad \|f\|_\infty = \sup_{\omega \in \Omega} |f(\omega)|.$$

We denote by $\ell^1(\Omega, \sigma)$ and $\ell^\infty(\Omega)$ the spaces of functions where $\|\cdot\|_1$ and $\|\cdot\|_\infty$ are finite, respectively.

The next result collects a few simple observations that will be useful later in this subsection. In particular, part (3) will be our compactness criterion for projectivised cones: we only need to ensure that $\|\cdot\|_1$ and $\|\cdot\|_\infty$ are bi-Lipschitz equivalent on the cone. This is one of the reasons we are forced to work with both norms $\|\cdot\|_1$ and $\|\cdot\|_\infty$.

Lemma 6.13 *Let (Ω, σ) be a countable set with a fully supported probability measure.*

1. *We have $\ell^\infty(\Omega) \subseteq \ell^1(\Omega, \sigma)$ and $\|\cdot\|_1 \leq \|\cdot\|_\infty$.*
2. *The topology of $(\ell^1(\Omega, \sigma), \|\cdot\|_1)$ is finer than the topology of pointwise convergence on Ω. The converse holds on those subsets of $\ell^1(\Omega, \sigma)$ where $\|\cdot\|_\infty$ is bounded.*
3. *Let $\mathcal{C} \subseteq \ell^1(\Omega, \sigma)$ be a positive cone that is closed in the topology of $\|\cdot\|_1$. Suppose that there exists $c > 0$ such that $\|f\|_\infty \leq c \cdot \|f\|_1$ for all $f \in \mathcal{C}$. Then $\mathbb{P}(\mathcal{C})$ is compact with respect to the quotient topology induced by $\|\cdot\|_1$.*

Proof Part (1) is clear. The two halves of part (2) follow respectively from the inequalities

$$|f(\omega)|\sigma(\{\omega\}) \le \|f\|_1 \quad \text{and} \quad \|f\|_1 \le \sum_{x\in F} |f(x)|\sigma(\{x\}) + \|f\|_\infty \cdot \sigma(\Omega \setminus F),$$

which hold for all $f \in \ell^1(\Omega,\sigma)$, all $\omega \in \Omega$ and every finite subset $F \subseteq \Omega$.

Finally, let us prove part (3). If S is the unit sphere in $\ell^1(\Omega,\sigma)$, then $\mathbb{P}(\mathcal{C})$ is homeomorphic to $\mathcal{C}\cap S$. Since the latter is metrisable, it suffices to show that every sequence $(f_k)_k \subseteq \mathcal{C}\cap S$ has a converging subsequence. Since $\|f_k\|_\infty \le c\cdot\|f_k\|_1 = c$, the sequence $(f_k(\omega))_k$ takes values in the compact interval $[-c,c]$ for all $\omega\in\Omega$. Since Ω is countable, a diagonal argument allows us to replace $(f_k)_k$ with a subsequence that converges pointwise to a function $f:\Omega\to[-c,c]$. Thus, part (2) shows that $\|f_k - f\|_1 \to 0$. Since $\mathcal{C}$ is closed in $\ell^1(\Omega,\sigma)$, we have $f\in\mathcal{C}\cap S$, as required. $\square$

Definition 6.14 A metrisable topological space X is an *absolute retract* (*AR*) if it enjoys the following property. For every metrisable topological space Y and every closed subset $A\subseteq Y$ homeomorphic to X, there exists a continuous retraction $Y\to A$.

The following summarises the key properties of ARs that we will need.

Theorem 6.15 (1) *Let X be a compact AR. Then every continuous map $f: X\to X$ has a fixed point.*

(2) *Let $(E,\|\cdot\|)$ be a normed space. If $\mathcal{C}\subseteq E$ is any positive cone, then $\mathbb{P}(\mathcal{C})$ is an AR* (*with the quotient of the norm topology of E*).

Proof Part (1) is a consequence of the Lefschetz fixed point theorem for compact ANRs [76; 77]. See for instance Theorem III.7.4 and Section I.6 in [70] for a clear statement.

If S is the unit sphere in the normed space E, then $\mathbb{P}(\mathcal{C})$ is homeomorphic to $\mathcal{C}\cap S$. Recall that every convex subset of a normed space is an AR; see for example [45, Corollary 4.2] or Corollary II.14.2 and Theorem III.3.1 in [70]. Every retract of an AR is again an AR; see [70, Proposition 7.7]. Thus, part (2) is immediate from the observation that $\mathcal{C}\cap S$ is a retract of the convex set $\mathcal{C}\setminus\{0\}$. $\square$

6.2.2 Finding a projectively invariant metric Let M be a *countable*, finite-rank median algebra. Consider a finite set $S\subseteq \operatorname{Aut} M$ and let $G\le \operatorname{Aut} M$ be the subgroup that it generates. Let $\alpha\in\operatorname{Aut} M$ be an element that normalises G.

Consider the locally convex real vector space $\mathcal{E}(M)=\mathbb{R}^{M\times M}$, endowed with the topology of pointwise convergence on $M\times M$. We have a continuous linear action $\operatorname{Aut} M \curvearrowright \mathcal{E}(M)$ given by

$$(\psi\cdot f)(x,y) = f(\psi^{-1}(x),\psi^{-1}(y)) \quad \text{for all } \psi\in\operatorname{Aut} M, \text{ all } f\in\mathcal{E}(M) \text{ and all } x,y\in M.$$

Remark 6.16 The sets $\mathcal{PD}(M)$ and $\mathcal{PD}^G(M)$ (introduced in Section 2.3) are closed positive cones in $\mathcal{E}(M)$. In addition, $\mathcal{PD}(M)$ is $(\operatorname{Aut} M)$–invariant and $\mathcal{PD}^G(M)$ is $\langle\alpha\rangle$–invariant.

Although $\mathcal{D}(M)\cup\{0\}$ also is a positive cone, it is only closed when M is a single point.

Given a function $\mathfrak{c}\colon M\times M\to(0,+\infty)$, consider the (not necessarily convex) subset

$$\mathcal{PD}^G_{\mathfrak{c}}(M):=\{\eta\in\mathcal{PD}^G(M)\mid \eta(x,y)\le\mathfrak{c}(x,y)\cdot\overline{\tau}^{\eta}_S \text{ for all } x,y\in M\}.$$

As we shall see, this serves two purposes: on the one hand all closed cones in $\mathcal{PD}^G_{\mathfrak{c}}(M)$ have compact projectivisation; on the other, they only contain pseudometrics with $\overline{\tau}^{\eta}_S>0$ (except for $\eta=0$).

Our main aim in this subsection is to prove the following result:

Proposition 6.17 *Suppose that, for some $\mathfrak{c}\colon M\times M\to(0,+\infty)$, there exists a nontrivial $\langle\alpha\rangle$–invariant cone $\mathcal{C}\subseteq\mathcal{PD}^G_{\mathfrak{c}}(M)$ that is closed in $\mathcal{E}(M)$ with respect to the topology of pointwise convergence. Then there exists $\eta\in\mathcal{C}\setminus\{0\}$ such that $\overline{\tau}^{\eta}_S>0$ and $\alpha\cdot\eta=\lambda\eta$ for some $\lambda>0$.*

In order to prove the proposition, let us fix a probability measure σ on M with full support. Given a function $\mathfrak{c}\colon M\times M\to(0,+\infty)$, for $f\in\mathcal{E}(M)$ we define

$$\|f\|^{\mathfrak{c}}_1:=\sum_{x,y\in M}\frac{|f(x,y)|}{\mathfrak{c}(x,y)}\sigma(x)\sigma(y)\quad\text{and}\quad \|f\|^{\mathfrak{c}}_\infty:=\sup_{x,y\in M}\frac{|f(x,y)|}{\mathfrak{c}(x,y)}.$$

Note that $\|f\|^{\mathfrak{c}}_1$ is a norm on the subspace $\mathcal{E}^1_{\mathfrak{c}}(M)\subseteq\mathcal{E}(M)$ where it is finite. (The same is true of $\|f\|^{\mathfrak{c}}_\infty$, but this will not be relevant to us.)

Remark 6.18 Rescaling functions $f\in\mathcal{E}(M)$ by $\mathfrak{c}$, we map $\mathcal{E}^1_{\mathfrak{c}}(M)$ onto $\ell^1(M\times M,\sigma\otimes\sigma)$ linearly isometrically, while taking $\|f\|^{\mathfrak{c}}_\infty$ to $\|f\|_\infty$. Thus, we can apply Lemma 6.13 in this context.

Lemma 6.19 *Consider a function $\mathfrak{c}\colon M\times M\to(0,+\infty)$.*

(1) *The subset $\mathcal{PD}^G_{\mathfrak{c}}(M)\subseteq\mathcal{E}(M)$ is closed under pointwise convergence.*

(2) *There exists a constant $c>0$ (depending on $\mathfrak{c}$ and σ) such that, for every $\eta\in\mathcal{PD}^G_{\mathfrak{c}}(M)$,*

$$\|\eta\|^{\mathfrak{c}}_1\le\|\eta\|^{\mathfrak{c}}_\infty\le\overline{\tau}^{\eta}_S\le c\cdot\|\eta\|^{\mathfrak{c}}_1.$$

Proof We begin with part (1). First, observe that the function $\eta\mapsto\overline{\tau}^{\eta}_S$ is upper-semicontinuous. Indeed, if $\eta_n\in\mathcal{PD}^G(M)$ converge pointwise to some $\eta\in\mathcal{PD}^G(M)$, then, for every $x\in M$,

$$\max_{s\in S}\eta(x,sx)=\lim_{n\to+\infty}\max_{s\in S}\eta_n(x,sx)\ge\limsup_{n\to+\infty}\overline{\tau}^{\eta_n}_S.$$

Hence $\overline{\tau}^{\eta}_S\ge\limsup\overline{\tau}^{\eta_n}_S$, which proves upper-semicontinuity. Now, if $\eta_n\in\mathcal{PD}^G_{\mathfrak{c}}(M)$, then

$$\eta(x,y)=\lim_{n\to+\infty}\eta_n(x,y)\le\limsup_{n\to+\infty}\mathfrak{c}(x,y)\cdot\overline{\tau}^{\eta_n}_S\le\mathfrak{c}(x,y)\cdot\overline{\tau}^{\eta}_S$$

for all $x,y\in M$. Along with Remark 6.16, this yields $\eta\in\mathcal{PD}^G_{\mathfrak{c}}(M)$, proving part (1).

Regarding part (2), the first inequality is in Lemma 6.13(1) and the second is immediate from the fact that $\eta \in \mathcal{PD}^G_{\mathfrak{c}}(M)$. In order to prove the third one, choose any point $x_0 \in M$. Then

$$\overline{\tau}^{\eta}_S = \inf_{x\in M} \max_{s\in S} \eta(x,sx) \le \max_{s\in S} \eta(x_0,sx_0) \le \|\eta\|^{\mathfrak{c}}_1 \cdot \max_{s\in S} \frac{\mathfrak{c}(x_0,sx_0)}{\sigma(\{x_0\})\sigma(\{sx_0\})}.$$

The constant appearing on the rightmost side is positive and well-defined, since $\mathfrak{c}$ takes positive values and σ has full support. This concludes the proof. □

Proof of Proposition 6.17 We want to apply the Lefschetz fixed point theorem to $\alpha\colon \mathbb{P}(\mathcal{C}) \to \mathbb{P}(\mathcal{C})$.

Since $\mathcal{C} \subseteq \mathcal{PD}^G(M)$, the cone $\mathcal{C}$ is actually a positive cone. By Lemma 6.19(2), the set $\mathcal{C}$ is contained in $\mathcal{E}^1_{\mathfrak{c}}(M)$. Thus, Theorem 6.15(2) shows that the projectivisation $\mathbb{P}(\mathcal{C})$, endowed with the quotient topology induced by $\|\cdot\|^{\mathfrak{c}}_1$, is an AR.

Since $\mathcal{C} \subseteq \mathcal{E}^1_{\mathfrak{c}}(M)$ is closed in the topology of pointwise convergence, the first half of Lemma 6.13(2) guarantees that $\mathcal{C}$ is also closed in the topology of $\|\cdot\|^{\mathfrak{c}}_1$. Thus, by Lemmas 6.19(2) and 6.13(3), the projectivisation $\mathbb{P}(\mathcal{C})$ is compact.

We are left to show that the action $\langle\alpha\rangle \curvearrowright \mathcal{C}$ is continuous with respect to the topology of $\|\cdot\|^{\mathfrak{c}}_1$. Note that, by Lemma 6.19(2), α takes $\|\cdot\|^{\mathfrak{c}}_1$–bounded subsets of $\mathcal{C} \subseteq \mathcal{PD}^G_{\mathfrak{c}}(M)$ to $\|\cdot\|^{\mathfrak{c}}_1$–bounded subsets of $\mathcal{C}$:

$$\|\alpha\cdot\eta\|^{\mathfrak{c}}_1 \le \overline{\tau}^{\alpha\cdot\eta}_S = \inf_{x\in M} \max_{s\in S} \eta(\alpha^{-1}x,\alpha^{-1}sx) = \overline{\tau}^{\eta}_{\alpha^{-1}S\alpha} \le |\alpha^{-1}S\alpha|_S \cdot \overline{\tau}^{\eta}_S \le c|\alpha^{-1}S\alpha|_S \cdot \|\eta\|^{\mathfrak{c}}_1.$$

Since the topology given by $\|\cdot\|^{\mathfrak{c}}_1$ is metrisable, it suffices to show that $\alpha\colon\mathcal{C}\to\mathcal{C}$ is sequentially continuous. Let $\eta_n \in \mathcal{C}$ be a sequence that $\|\cdot\|^{\mathfrak{c}}_1$–converges to $\eta \in \mathcal{C}$. By Lemma 6.13(2), η_n converges to η pointwise. Since the action $\operatorname{Aut} M \curvearrowright \mathcal{E}(M)$ is continuous, the sequence $\alpha\cdot\eta_n$ converges to $\alpha\cdot\eta$ pointwise. Note that the set $\{\eta_n\}_{n\ge0}\cup\{\eta\}$ is $\|\cdot\|^{\mathfrak{c}}_1$–bounded and, by the above observation, so must be $\{\alpha\cdot\eta_n\}_{n\ge0}\cup\{\alpha\eta\}$. By Lemma 6.19(2), this set is also $\|\cdot\|^{\mathfrak{c}}_\infty$–bounded, so Lemma 6.13(2) shows that $\alpha\cdot\eta_n$ $\|\cdot\|^{\mathfrak{c}}_1$–converges to $\alpha\cdot\eta$, as required.

In conclusion, α induces a homeomorphism of the compact AR $\mathbb{P}(\mathcal{C})$. Theorem 6.15(1) yields an $\langle\alpha\rangle$–fixed point $[\eta] \in \mathbb{P}(\mathcal{C})$. The fact that $\overline{\tau}^{\eta}_S > 0$ is clear since $\eta \in \mathcal{PD}^G_{\mathfrak{c}}(M)\setminus\{0\}$. □

In fact, Proposition 6.17 can be easily generalised to extensions of G by abelian groups.

Corollary 6.20 *Let $U \le \operatorname{Aut} M$ be a countable subgroup such that $G \lhd U$, with abelian quotient U/G; let $p\colon U \to A$ be the quotient projection. Suppose that, for some $\mathfrak{c}$, there exists a nontrivial, U–invariant, closed cone $\mathcal{C} \subseteq \mathcal{PD}^G_{\mathfrak{c}}(M)$. Then there exists $\eta \in \mathcal{C}\setminus\{0\}$ with $\overline{\tau}^{\eta}_S > 0$ and a homomorphism $\lambda\colon A \to (\mathbb{R}_{>0},*)$ such that $u\cdot\eta = \lambda(p(u))\eta$ for all $u \in U$.*

Proof Let $\{a_i\}_{i\ge0}$ be a generating set for A. Consider the subgroups $A_n := \langle a_i \mid i < n\rangle$ and $U_n := p^{-1}(A_n)$; in particular, $A_0 = \{1\}$ and $U_0 = G$. We will show by induction on $n \ge 0$ that there exist nontrivial, U–invariant, closed cones $\mathcal{C}_n \subseteq \mathcal{PD}^G_{\mathfrak{c}}(M)$ and homomorphisms $\lambda_n \colon A_n \to (\mathbb{R}_{>0}, *)$ such that $u \cdot \eta = \lambda_n(p(u))\eta$ for all $\eta \in \mathcal{C}_n$ and $u \in U_n$. As base step, set $\mathcal{C}_0 := \mathcal{C}$.

Regarding the inductive step, suppose that we have constructed $\mathcal{C}_n$ and λ_n. By Proposition 6.17, there exists a point $[\eta_{n+1}] \in \mathbb{P}(\mathcal{C}_n)$ fixed by $p^{-1}(a_{n+1})$. In fact, since U_n acts trivially on $\mathbb{P}(\mathcal{C}_n)$, the entire group U_{n+1} fixes $[\eta_{n+1}]$ and there exists a homomorphism $\lambda_{n+1} \colon A_{n+1} \to (\mathbb{R}_{>0}, *)$ such that $u \cdot \eta_{n+1} = \lambda_{n+1}(p(u))\eta_{n+1}$ for all $u \in U_{n+1}$. We can then define $\mathcal{C}_{n+1}$ as the closed cone

$$\{\eta \in \mathcal{C}_n \mid u \cdot \eta = \lambda_{n+1}(p(u))\eta \text{ for all } u \in U_{n+1}\}.$$

Since $U \curvearrowright \mathcal{C}$ factors through the abelian group A, this cone is U–invariant, as required.

Finally, when A is not finitely generated, note that the intersection of the descending chain $\mathcal{C}_n$ is not just $\{0\}$. This is because, as we observed in the proof of Proposition 6.17, the sets $\mathbb{P}(\mathcal{C}_n)$ are compact. This concludes the proof. □

6.2.3 Universal uniform nonelementarity Let $G \curvearrowright M$ be an action by automorphisms on a median algebra of finite rank r. Consider the following strengthening of Definition 2.36 in the context of compatible metrics on median algebras.

Definition 6.21 The action $G \curvearrowright M$ is *universally uniformly nonelementary* (*WNE*) if there exists a constant $c > 0$ such that, for every $\eta \in \mathcal{PD}^G(M)$, the action $G \curvearrowright (M, \eta)$ is c–UNE.

This may seem an impossibly strong requirement to impose on $G \curvearrowright M$, but we will see in Corollary 7.24 that many actions arising from ultralimits of Salvetti complexes are WNE.

Lemma 6.22 *Let $G \leq \operatorname{Aut} M$ be generated by a finite set S of automorphisms acting nontransversely and stably without inversions. Let $G \lhd U \leq \operatorname{Aut} M$. Pick a point q in the multibridge $\mathcal{B}(S) \subseteq M$ and let $\mathfrak{M} \subseteq M$ be the median subalgebra generated by the orbit $U \cdot q$. Then*:

1. *There exists $\mathfrak{c}_1 \colon \mathfrak{M} \to (0, +\infty)$ such that $\tau^\eta_S(x) \leq \mathfrak{c}_1(x) \cdot \overline{\tau}^\eta_S$ for all $\eta \in \mathcal{PD}^G(M)$ and $x \in \mathfrak{M}$.*
2. *If $G \curvearrowright M$ is WNE, there exists $\mathfrak{c}_2 \colon \mathfrak{M} \times \mathfrak{M} \to (0, +\infty)$ such that $\eta(x, y) \leq \mathfrak{c}_2(x, y) \cdot \overline{\tau}^\eta_S$ for all $\eta \in \mathcal{PD}^G(M)$ and $x, y \in \mathfrak{M}$.*

Proof We only prove part (1), since part (2) then follows by setting $\mathfrak{c}_2(x, y) := c \cdot (\mathfrak{c}_1(x) + \mathfrak{c}_1(y))$, for a constant c as in Definition 6.21.

If part (1) holds for points $x, y, z \in \mathfrak{M}$, then it holds for their median $m(x, y, z)$. Indeed, we can take $\mathfrak{c}_1(m(x, y, z)) = \mathfrak{c}_1(x) + \mathfrak{c}_1(y) + \mathfrak{c}_1(z)$ and we have

$$\begin{aligned} \tau^\eta_S(m(x, y, z)) &= \max_{s \in S} \eta(m(x, y, z), m(sx, sy, sz)) \\ &\leq \max_{s \in S} [\eta(x, sx) + \eta(y, sy) + \eta(z, sz)] \\ &\leq \tau^\eta_S(x) + \tau^\eta_S(y) + \tau^\eta_S(z) \leq [\mathfrak{c}_1(x) + \mathfrak{c}_1(y) + \mathfrak{c}_1(z)] \cdot \overline{\tau}^\eta_S. \end{aligned}$$

Thus, it suffices to prove part (1) for $x \in U \cdot q$. Since $q \in \mathcal{B}(S)$, we have $uq \in \mathcal{B}(uSu^{-1})$ for all $u \in U$. Moreover, since U normalises G, the set uSu^{-1} is just another generating set of G. By Proposition 6.11(3), we have

$$\begin{aligned}\tau_S^{\eta}(uq) \leq |S|_{uSu^{-1}} \cdot \tau_{uSu^{-1}}^{\eta}(uq) &\leq |S|_{uSu^{-1}} \cdot (2r+1)\overline{\tau}_{uSu^{-1}}^{\eta} \\ &\leq |S|_{uSu^{-1}} \cdot (2r+1) \cdot |uSu^{-1}|_S \cdot \overline{\tau}_S^{\eta}.\end{aligned}$$

So we can take $\mathfrak{c}_1(uq) = (2r+1) \cdot |S|_{uSu^{-1}} \cdot |uSu^{-1}|_S$. This concludes the proof. □

Corollary 6.23 *Let $G \leq \operatorname{Aut} M$ be generated by a finite set S of automorphisms acting nontransversely and stably without inversions. Suppose that $G \curvearrowright M$ is WNE and that $\mathcal{D}^G(M) \neq \varnothing$. Consider a countable subgroup $U \leq \operatorname{Aut} M$ such that $G \lhd U$ and U/G is abelian. Then there exist a nonempty, countable, U–invariant, median subalgebra $\mathfrak{M} \subseteq M$, a pseudometric $\eta \in \mathcal{PD}^G(\mathfrak{M}) \setminus \{0\}$ with $\overline{\tau}_S^{\eta} > 0$, and a homomorphism $\lambda\colon U \to (\mathbb{R}_{>0}, *)$ (trivial on G) with $u \cdot \eta = \lambda(u)\eta$ for all $u \in U$.*

Proof Define the median subalgebra $\mathfrak{M} \subseteq M$ as in the statement of Lemma 6.22. Since $\mathfrak{M}$ is generated by a countable set, it is itself countable. The restriction map

$$\operatorname{res}_{\mathfrak{M}}\colon \mathcal{PD}(M) \to \mathcal{PD}(\mathfrak{M})$$

takes $\mathcal{PD}^G(M)$ into $\mathcal{PD}^G(\mathfrak{M})$ without decreasing the value of $\overline{\tau}_S^{\bullet}$. Thus, in the notation of Section 6.2.2, Lemma 6.22(2) yields

$$\operatorname{res}_{\mathfrak{M}}(\mathcal{PD}^G(M)) \subseteq \mathcal{PD}^G_{\mathfrak{c}_2}(\mathfrak{M}).$$

Choose $\delta \in \mathcal{D}^G(M)$ and let $\mathcal{C} \subseteq \mathcal{D}^G(M)$ be the smallest cone containing the U–orbit of δ. In other words, $\mathcal{C}$ is the convex hull of $U \cdot \delta$, saturated under multiplication by nonnegative scalars. Then $\operatorname{res}_{\mathfrak{M}}(\mathcal{C})$ is a nontrivial U–invariant cone contained in $\mathcal{PD}^G_{\mathfrak{c}_2}(\mathfrak{M})$.

Its closure $\overline{\operatorname{res}_{\mathfrak{M}}(\mathcal{C})} \subseteq \mathcal{E}(\mathfrak{M})$ in the topology of pointwise convergence is also a U–invariant cone. By Lemma 6.19(1), this is still contained in the set $\mathcal{PD}^G_{\mathfrak{c}_2}(\mathfrak{M})$. We can thus apply Corollary 6.20, obtaining $\eta \in \overline{\operatorname{res}_{\mathfrak{M}}(\mathcal{C})} \setminus \{0\}$ with $\overline{\tau}_S^{\eta} > 0$, and a homomorphism $\lambda\colon U \to (\mathbb{R}_{>0}, *)$ such that $u \cdot \eta = \lambda(u)\eta$ for all $u \in U$. □

7 Ultralimits and coarse-median preserving automorphisms

In this section we prove Theorem I (Corollary 7.23) and complete the proof of Theorem F (Theorem 7.25). Both results will follow quickly once we prove Theorem 7.21 in Section 7.4, which can be viewed as the main goal of this entire section.

This theorem claims that, in many cases, if $G \curvearrowright M$ is an action of a special group on a median algebra, η is a G–invariant compatible pseudometric and C is a large k–cube in M, then any subset of G that moves all points in C by a lot less than the "size" of C must commute with a copy of $\mathbb{Z}^k$ sitting inside G. This result holds, for instance, for co-special cubulations of G, ultralimits of these, and subalgebras thereof, with uniform constants that are independent of the specific choice of η.

The case $k = 1$ thus implies that all these actions are WNE (Definition 6.21) and that centreless special groups are UNE (Definition 2.36). The cases with $k > 1$ ensure that the actions on median spaces that we will construct for Theorem F are *moderate*, as defined in the introduction.

7.1 The Bestvina–Paulin construction

As sketched in the introduction, the first step in the proof of Theorem F will involve a standard Bestvina–Paulin construction, with some additional issues caused by the lack of hyperbolicity. In this subsection, we discuss the role played by UNE groups (Definition 2.36) in addressing these issues.

Consider a group G, a geodesic metric space (X, d), and a homomorphism $\rho\colon G \to \operatorname{Isom} X$ inducing a proper cocompact action $G \curvearrowright X$ (we simply write gx rather than $\rho(g) \cdot x$).

7.1.1 The classical Bestvina–Paulin construction Fix a finite generating set $S \subseteq G$ and let $|\cdot|_S$ be the induced word length on G. Denote by $\pi : \operatorname{Aut} G \to \operatorname{Out} G$ the quotient projection. Given $g, h \in G$, we write $\mathfrak{c}[g](h) := ghg^{-1}$.

Every group automorphism $\varphi\colon G \to G$ is bi-Lipschitz with respect to $|\cdot|_S$. By the Milnor–Schwarz lemma, φ induces a quasi-isometry $\widetilde{\varphi}\colon X \to X$ satisfying $\widetilde{\varphi} \circ \rho(g) = \rho(\varphi(g)) \circ \widetilde{\varphi}$ for all $g \in G$.

Consider a sequence $\varphi_n \in \operatorname{Aut} G$ and set $\rho_n := \rho \circ \varphi_n$ for all $n \geq 0$. Pick basepoints $p_n \in X$ with

$$\tau_S^{\rho_n}(p_n) - \overline{\tau}_S^{\rho_n} \leq 1.$$

We introduce the quantities $\epsilon_n := 1/\overline{\tau}_S^{\rho_n}$ to simplify the notation.

Assumption 7.1 In the rest of Section 7.1, we assume that no two elements of the sequence $\pi(\varphi_n) \in \operatorname{Out} G$ coincide. A classical argument due to Bestvina and Paulin (see eg [7] and [87, page 338]) then guarantees that $\epsilon_n \to 0$ for $n \to +\infty$.

Fix a nonprincipal ultrafilter ω and consider the ultralimit $(X_\omega, d_\omega, p_\omega) = \lim_\omega (X, \epsilon_n d, p_n)$. We have a homomorphism $\rho_\omega\colon G \to \operatorname{Isom} X_\omega$ obtained as ultralimit of the actions ρ_n, namely

$$\rho_\omega(g) \cdot (x_n) = (\rho_n(g) \cdot x_n) = (\varphi_n(g) x_n)$$

for all $g \in G$ and $(x_n) \in X_\omega$. This is well-defined since

$$\begin{aligned}
\lim_\omega \epsilon_n d(\varphi_n(g)x_n, p_n) &\leq \lim_\omega \epsilon_n [d(\varphi_n(g)x_n, \varphi_n(g)p_n) + d(\varphi_n(g)p_n, p_n)] \\
&\leq \lim_\omega \epsilon_n [d(x_n, p_n) + |g|_S \cdot \tau_S^{\rho_n}(p_n)] \\
&= d_\omega((x_n), p_\omega) + |g|_S < +\infty.
\end{aligned}$$

One easily checks that $\tau_S^{\rho_\omega}(p_\omega) = \overline{\tau}_S^{\rho_\omega} = 1$, so the action $G \curvearrowright X_\omega$ induced by ρ_ω does not have a global fixed point.

7.1.2 Automorphisms of UNE groups Suppose for a moment that we are in the special case where there exists $\varphi \in \operatorname{Aut} G$ such that $\varphi_n = \varphi^n$ for all $n \geq 0$ (thus $\rho_n = \rho \circ \varphi^n$). We want to show that φ induces a map $\Phi\colon X_\omega \to X_\omega$ with the property that $\Phi \circ \rho_\omega(g) = \rho_\omega(\varphi(g)) \circ \Phi$ for all $g \in G$. A natural attempt is setting $\Phi((x_n)) = (\widetilde{\varphi}(x_n))$ for all $(x_n) \in X_\omega$. However, for this to be well-defined we need $\lim_\omega \epsilon_n d(\widetilde{\varphi}(p_n), p_n) < +\infty$.

We are actually interested in the following more general setting.

Assumption 7.2 Let $N \leq \operatorname{Out} G$ be a subgroup with infinite centre $Z(N)$. Let $\varphi_n \in \operatorname{Aut} G$ be a sequence that is mapped by the projection $\pi : \operatorname{Aut} G \to \operatorname{Out} G$ to a sequence of pairwise distinct elements in $Z(N)$. Consider again $\rho_n = \rho \circ \varphi_n$ as above.

If $\psi \in \pi^{-1}(N)$, then $\pi(\psi)$ commutes with each $\pi(\varphi_n)$. For every $n \in \mathbb{Z}$, choose $g_{n,\psi} \in G$ with

$$\varphi_n \circ \psi = \mathfrak{c}[g_{n,\psi}] \circ \psi \circ \varphi_n.$$

We are about to prove that, if G is UNE, ψ induces a well-defined map $\zeta(\psi)\colon X_\omega \to X_\omega$ given by

$$\zeta(\psi)((x_n)) = (g_{n,\psi} \widetilde{\psi}(x_n)).$$

(Recall that $\widetilde{\psi} : X \to X$ is the quasi-isometry induced by ψ.) We essentially use the same argument as [88, pages 154–156], replacing hyperbolicity with the UNE condition.

The proof of this result is quite technical. On a first read, we suggest restricting to the situation where $N \simeq \mathbb{Z}$ and the automorphisms ψ and φ_n are all powers of a given automorphism, in which case the elements $g_{n,\psi}$ can all be taken to be the identity and our strategy boils down to what is described right before Assumption 7.2. This case is sufficient for Theorem F, though not for the more general Theorem 7.25 below.

Proposition 7.3 *Suppose that G is UNE. Let $N \leq \operatorname{Out} G$ and $\varphi_n \in \operatorname{Aut} G$ be as in Assumption 7.2. Then there exists a homomorphism $\zeta : \pi^{-1}(N) \to \operatorname{Homeo} X_\omega$ that extends ρ_ω, in the sense that $\zeta(\mathfrak{c}[g]) = \rho_\omega(g)$ for every $g \in G$. Every homeomorphism in the image of ζ is bi-Lipschitz.*

Proof Consider an element $\psi \in \pi^{-1}(N)$. Let $L \geq 1$ be a constant such that $\widetilde{\psi} : X \to X$ is an (L, L)–quasi-isometry and such that $\psi : G \to G$ is L–bi-Lipschitz with respect to $|\cdot|_S$.

Step 1 *The map $\zeta(\psi)$ described above is a well-defined bi-Lipschitz homeomorphism of X_ω.*

Since $\widetilde{\psi}$ is a quasi-isometry and $\epsilon_n \to 0$, it suffices to show that $\zeta(\psi)$ is a well-defined map, ie that $\lim_\omega \epsilon_n d(g_{n,\psi} \widetilde{\psi}(p_n), p_n)$ is finite.

We begin by observing that, since $\varphi_n \circ \psi = \mathfrak{c}[g_{n,\psi}] \circ \psi \circ \varphi_n$ and $\widetilde{\psi} \circ \rho(g) = \rho(\psi(g)) \circ \widetilde{\psi}$,

$$\begin{aligned}
\tau_S^{\rho_n}(g_{n,\psi}\widetilde{\psi}(p_n)) &= \max_{s\in S} d(\varphi_n(s)g_{n,\psi}\widetilde{\psi}(p_n), g_{n,\psi}\widetilde{\psi}(p_n)) = \max_{s\in S} d((\mathfrak{c}[g_{n,\psi}]^{-1}\varphi_n)(s)\widetilde{\psi}(p_n), \widetilde{\psi}(p_n)) \\
&= \max_{s\in S} d(\widetilde{\psi}((\psi^{-1}\mathfrak{c}[g_{n,\psi}]^{-1}\varphi_n)(s)p_n), \widetilde{\psi}(p_n)) = \max_{s\in S} d(\widetilde{\psi}(\varphi_n\psi^{-1}(s)p_n), \widetilde{\psi}(p_n)) \\
&\leq L\cdot \max_{s\in S} d(\varphi_n\psi^{-1}(s)p_n, p_n) + L = L\cdot\max_{s\in S} d(\rho_n(\psi^{-1}(s))\cdot p_n, p_n) + L \\
&\leq L\cdot \max_{s\in S} |\psi^{-1}(s)|_S \cdot \tau_S^{\rho_n}(p_n) + L \leq L^2\cdot\tau_S^{\rho_n}(p_n) + L.
\end{aligned}$$

Now, since G is UNE, there exists a constant $c>0$ such that, for every generating set $T\subseteq G$ and all $x,y\in X$, we have $d(x,y)\leq c\cdot(\tau_T^{\rho}(x)+\tau_T^{\rho}(y))$. For $T=\varphi_n(S)$, we obtain

$$\begin{aligned}
\lim_\omega \epsilon_n d(g_{n,\psi}\widetilde{\psi}(p_n), p_n) &\leq c\cdot \lim_\omega \epsilon_n(\tau^{\rho}_{\varphi_n(S)}(g_{n,\psi}\widetilde{\psi}(p_n)) + \tau^{\rho}_{\varphi_n(S)}(p_n)) \\
&= c\cdot\lim_\omega \epsilon_n(\tau_S^{\rho_n}(g_{n,\psi}\widetilde{\psi}(p_n)) + \tau_S^{\rho_n}(p_n)) \\
&\leq c(L^2+1)\cdot\lim_\omega \epsilon_n\tau_S^{\rho_n}(p_n) < +\infty.
\end{aligned}$$

Step 2 *The map ζ is a homomorphism.*

Since G is UNE, Example 2.38(3) shows that the centre $Z(G)\leq G$ is finite. Then, since G acts cocompactly on X, there exists a constant M such that $d(x,zx)\leq M$ for all $x\in X$ and $z\in Z(G)$. Given $\psi_1,\psi_2\in N$, we can take $\widetilde{\psi_1\psi_2}=\widetilde{\psi}_1\widetilde{\psi}_2$. Moreover,

$$\begin{aligned}
\mathfrak{c}[g_{n,\psi_1\psi_2}]\psi_1\psi_2\varphi_n = \varphi_n\psi_1\psi_2 &= \mathfrak{c}[g_{n,\psi_1}]\psi_1\varphi_n\psi_2 \\
&= \mathfrak{c}[g_{n,\psi_1}]\psi_1\mathfrak{c}[g_{n,\psi_2}]\psi_2\varphi_n \\
&= \mathfrak{c}[g_{n,\psi_1}]\mathfrak{c}[\psi_1(g_{n,\psi_2})]\psi_1\psi_2\varphi_n.
\end{aligned}$$

Hence $g_{n,\psi_1\psi_2}$ and $g_{n,\psi_1}\psi_1(g_{n,\psi_2})$ differ by multiplication by an element of $Z(G)$. It follows that, for every $x\in X$, we have $d(g_{n,\psi_1\psi_2}x, g_{n,\psi_1}\psi_1(g_{n,\psi_2})x)\leq M$. Thus, for every $(x_n)\in X_\omega$,

$$\begin{aligned}
\zeta(\psi_1\psi_2)((x_n)) &= (g_{n,\psi_1\psi_2}\widetilde{\psi_1\psi_2}(x_n)) = (g_{n,\psi_1}\psi_1(g_{n,\psi_2})\widetilde{\psi}_1(\widetilde{\psi}_2(x_n))) \\
&= (g_{n,\psi_1}\widetilde{\psi}_1(g_{n,\psi_2}\widetilde{\psi}_2(x_n))) = \zeta(\psi_1)((g_{n,\psi_2}\widetilde{\psi}_2(x_n))) \\
&= \zeta(\psi_1)\zeta(\psi_2)((x_n)).
\end{aligned}$$

Step 3 *We have $\zeta(\mathfrak{c}[g])=\rho_\omega(g)$ for all $g\in G$.*

Since $\mathfrak{c}[g]\colon G\to G$ is at bounded distance from left multiplication by g, the quasi-isometry $\widetilde{\mathfrak{c}[g]}$ is at bounded distance from $\rho(g)$. Moreover, observing that

$$\mathfrak{c}[\varphi_n(g)]\circ\varphi_n = \varphi_n\circ\mathfrak{c}[g] = \mathfrak{c}[g_{n,\mathfrak{c}[g]}]\circ\mathfrak{c}[g]\circ\varphi_n,$$

we deduce that $\mathfrak{c}[\varphi_n(g)g^{-1}]=\mathfrak{c}[g_{n,\mathfrak{c}[g]}]$, hence $g_{n,\mathfrak{c}[g]}\in Z(G)\varphi_n(g)g^{-1}$. Thus, for every $(x_n)\in X_\omega$,

$$\zeta(\mathfrak{c}[g])((x_n)) = (g_{n,\mathfrak{c}[g]}\widetilde{\mathfrak{c}[g]}(x_n)) = (g_{n,\mathfrak{c}[g]}gx_n) = (\varphi_n(g)g^{-1}gx_n) = (\varphi_n(g)x_n) = \rho_\omega(g)((x_n)).$$

This concludes the proof of the proposition. □

In the special case where there exists $\varphi \in \operatorname{Aut} G$ such that $\varphi_n = \varphi^n$ and $N = \langle \pi(\varphi) \rangle$, we have $\pi^{-1}(N) \simeq (G/Z(G)) \rtimes_\varphi \mathbb{Z}$ and we obtain:

Corollary 7.4 *Suppose that G is UNE and that $\pi(\varphi) \in \operatorname{Out} G$ has infinite order. Take $\varphi_n = \varphi^n$. Then the map $\Phi\colon X_\omega \to X_\omega$ given by $\Phi((x_n)) = (\widetilde{\varphi}(x_n))$ is a well-defined bi-Lipschitz homeomorphism of X_ω satisfying $\Phi \circ \rho_\omega(g) = \rho_\omega(\varphi(g)) \circ \Phi$ for all $g \in G$.*

7.1.3 Coarse-median preserving automorphisms of UNE groups Suppose now that X admits a coarse median μ of finite rank r. We can define a map $\mu_\omega\colon X_\omega^3 \to X_\omega$ by setting $\mu_\omega((x_n),(y_n),(z_n)) = (\mu(x_n, y_n, z_n))$. It was shown in [15, Section 9] that μ_ω is well-defined and the pair (X_ω, μ_ω) is a median algebra of rank $\leq r$.

If the coarse median structure $[\mu]$ is fixed by $G \curvearrowright X$, then the action $G \curvearrowright X_\omega$ is by automorphisms of the median algebra (X_ω, μ_ω). Moreover, if an automorphism $\psi \in \pi^{-1}(N) \leq \operatorname{Aut} G$ is such that $\widetilde{\psi}$ fixes $[\mu]$, then $\zeta(\psi) \in \operatorname{Aut}(X_\omega, \mu_\omega)$. Note that, although the metric d_ω on X_ω is G–invariant, it needs not be preserved by $\zeta(\psi)$.

Remark 7.5 If the space X is coarse median but not median, the metric d_ω may not be compatible with μ_ω (in the sense of Definition 2.8). However, it was shown by Zeidler [101, Proposition 3.3] that there always exists a metric $\delta \in \mathcal{D}^G(X_\omega, \mu_\omega)$ such that (X_ω, δ) is complete, geodesic, and bi-Lipschitz equivalent to (X_ω, d_ω). Theorem 2.14(2) and the fact that G does not fix a point in X_ω then imply that G acts on (X_ω, δ) with unbounded orbits (alternatively, one can appeal to [17]).

This is only tangentially relevant to us as we will only be interested in ultralimits of CAT(0) cube complexes in the forthcoming subsections.

Summing up the above discussion:

Corollary 7.6 *Let G be a UNE group. Let $N \leq \operatorname{Out} G$ be a subgroup with infinite centre. Let $(X, [\mu])$ be a geodesic coarse median space of finite rank r. Let $G \curvearrowright X$ be a proper cocompact action fixing the coarse median structure $[\mu]$. Suppose that the quasi-isometries of X induced by the elements of $\pi^{-1}(N)$ also preserve $[\mu]$.*

Then there exists a complete, geodesic median space X_ω of rank $\leq r$, and an action $\pi^{-1}(N) \curvearrowright X_\omega$ by bi-Lipschitz homeomorphisms that preserve the underlying median-algebra structure. The composition $G \to G/Z(G) \hookrightarrow \pi^{-1}(N) \curvearrowright X_\omega$ is an isometric G–action with unbounded orbits.

7.2 Equivariant embeddings in products of $\mathbb{R}$–trees

Let M be a median algebra and $G \curvearrowright M$ an action by median automorphisms. In the rest of Section 7, we will be interested in situations where M can be embedded G–equivariantly into a finite product of $\mathbb{R}$–trees. We reserve this subsection for a few general remarks on this setting.

Definition 7.7 An $\mathbb{R}$–*tree* is a geodesic, rank-1 median space.

This is equivalent to the usual definition of $\mathbb{R}$–trees as geodesic metric spaces where every geodesic triangle is a tripod. We stress that $\mathbb{R}$–trees are not required to be complete.

The next remark collects various simple observations for later use.

Remark 7.8 Consider isometric G–actions on $\mathbb{R}$–trees $T_1, \dots, T_k$. Equip $T_1 \times \cdots \times T_k$ with the diagonal G–action. Let $f = (f_i)\colon M \hookrightarrow \prod T_i$ be a G–equivariant, injective median morphism.

(1) The image $f(M)$ is a median subalgebra of $\prod_i T_i$. The set of halfspaces of the median algebra $\prod_i T_i$ is naturally identified with the disjoint union $\bigsqcup_i \mathscr{H}(T_i)$.

Every halfspace of T_i is either open or closed. Open halfspaces are precisely the single connected components of the sets $T_i \setminus \{p\}$, as p varies through all points of T_i (including when $T_i \setminus \{p\}$ is connected). Closed halfspaces are precisely the complements of open halfspaces.

If we let $\mathscr{H}_i \subseteq \mathscr{H}(M)$ be the set of halfspaces of the form $f_i^{-1}(\mathfrak{h})$ with $\mathfrak{h} \in \mathscr{H}(T_i)$, then the $\mathscr{H}_i$ cover $\mathscr{H}(M)$ by Remark 2.2(1). However, the $\mathscr{H}_i$ are usually not pairwise disjoint.

(2) Since the sets $\mathscr{H}_i$ are G–invariant and no two halfspaces in the same $\mathscr{H}_i$ are transverse, we see that each $g \in G$ must act nontransversely on M.

(3) Suppose that, for all i, all $x \in T_i$ and all $g \in G$, we have $g^2 x = x$ if and only if $gx = x$. Then the action $G \curvearrowright M$ has no wall inversions.

Indeed, suppose instead that there exists $\mathfrak{h} \in \mathscr{H}(M)$ such that $g\mathfrak{h} = \mathfrak{h}^*$. Pick i such that $\mathfrak{h} \in \mathscr{H}_i$, and choose $\mathfrak{k} \in \mathscr{H}(T_i)$ with $f_i^{-1}(\mathfrak{k}) = \mathfrak{h}$. Then $g\mathfrak{k} \cap \mathfrak{k}$ and $g\mathfrak{k}^* \cap \mathfrak{k}^*$ are disjoint from the $\langle g \rangle$–invariant median subalgebra $f_i(M)$. Note that we cannot have $g\mathfrak{k} \subseteq \mathfrak{k}$ or $g\mathfrak{k} \supseteq \mathfrak{k}$, so, without loss of generality, $g\mathfrak{k} \cap \mathfrak{k} = \varnothing$. It follows that $f_i(M) \subseteq \mathfrak{k} \cup g\mathfrak{k}$, hence g is elliptic and fixes a unique point p in the convex hull of $\mathfrak{k} \cup g\mathfrak{k}$. We conclude that $g^2\mathfrak{k} = \mathfrak{k}$, hence the points on the arc connecting p to $\mathfrak{k}$ are fixed by g^2, but not by g. This is a contradiction.

(4) Suppose that g acts on M stably without wall inversions. By Remark 2.18(2) and Theorem 2.14(3), a halfspace $\mathfrak{h} \in \mathscr{H}(M)$ lies in the set $\mathscr{H}_{\overline{\mathcal{C}}(g)}(M)$ if and only if either $\mathfrak{h} \subsetneq g\mathfrak{h}$, or $\mathfrak{h} \subsetneq g^{-1}\mathfrak{h}$, or $\mathfrak{h} = g\mathfrak{h}$.

It follows that, for every i, either g is loxodromic in T_i and $f_i(\overline{\mathcal{C}}(g, M))$ is contained in its axis, or g is elliptic in T_i and fixes $f_i(\overline{\mathcal{C}}(g, M))$ pointwise.

Now, let us fix a nonprincipal ultrafilter ω. Let the group G be generated by a finite subset S. Consider a sequence of actions by automorphism on median algebras $G \curvearrowright M_n$, along with metrics $\delta_n \in \mathcal{D}^G(M_n)$ and basepoints $p_n \in M_n$. Suppose moreover that

$$\max_{s \in S} \sup_n \delta_n(sp_n, p_n) < +\infty.$$

Define $(M_\omega, \delta_\omega, p_\omega) := \lim_\omega (M_n, \delta_n, p_n)$. The set M_ω becomes a median algebra if we endow it with the operator $m((x_n), (y_n), (z_n)) = (m(x_n, y_n, z_n))$. We have an action by median automorphisms $G \curvearrowright M_\omega$ given by $g(x_n) = (gx_n)$. Finally, note that $\delta_\omega \in \mathcal{D}^G(M_\omega)$, and that $(M_\omega, \delta_\omega)$ is a complete median space (every ultralimit of metric spaces is complete).

Given a sequence of subsets $A_n \subseteq M_n$, we will employ the notation

$$\lim_\omega A_n := \{(x_n) \in M_\omega \mid x_n \in A_n \text{ for } \omega\text{–all } n\} = \{(y_n) \in M_\omega \mid \lim_\omega \delta_n(y_n, A_n) = 0\}.$$

Note that $\lim_\omega A_n$ is a (possibly empty) *closed* subset of $(M_\omega, \delta_\omega)$ for *any* sequence of subsets $A_n \subseteq M_n$. It is also clear that $\lim_\omega A_n \subseteq M_\omega$ is convex as soon as $A_n \subseteq M_n$ is convex for ω–all n.

Fix an integer $k \geq 1$. Suppose that each action $G \curvearrowright M_n$ is equipped with a G–equivariant, δ_n–isometric embedding $f_n = (f_n^i)\colon M_n \hookrightarrow \prod_i T_n^i$, where $\prod_i T_n^i$ is a product of k $\mathbb{R}$–trees endowed with an isometric, diagonal G–action as in Remark 7.8. (We have switched the index i from subscript to superscript to avoid confusion.)

It is straightforward to check that the ultralimits $\lim_\omega(T_n^i, f_n^i(p_n))$ yield isometric G–actions on $\mathbb{R}$–trees T_ω^i and a G–equivariant, δ_ω–isometric embedding $f_\omega = (f_\omega^i)\colon M_\omega \hookrightarrow \prod_i T_\omega^i$.

Lemma 7.9 *Consider the above setting. For every $g \in G$, we have*

(1) $\ell(g, T_\omega^i) = \lim_\omega \ell(g, T_n^i)$ *and* $\bar{\mathcal{C}}(g, T_\omega^i) = \lim_\omega \bar{\mathcal{C}}(g, T_n^i)$ *for all* $1 \leq i \leq k$.

If, in addition, (M_n, δ_n) is a geodesic space for ω–all n, then $(M_\omega, \delta_\omega)$ is geodesic and

(2) $\ell(g, \delta_\omega) = \lim_\omega \ell(g, \delta_n)$ *and* $\bar{\mathcal{C}}(g, M_\omega) = \lim_\omega \bar{\mathcal{C}}(g, M_n)$.

Proof We only prove part (2), since part (1) is a special case of it.

By Remarks 2.12 and 7.8(2), each $g \in G$ acts on M_ω stably without inversions and nontransversely; the same is true of the action on ω–all M_n. Theorem 2.16(2) shows that, for every $x = (x_n) \in M_\omega$, we have

$$\delta_\omega(x, gx) = \lim_\omega \delta_n(x_n, gx_n) = \lim_\omega [\ell(g, \delta_n) + 2\delta_n(x_n, \bar{\mathcal{C}}(g, M_n))] \geq \lim_\omega \ell(g, \delta_n).$$

Hence $\ell(g, \delta_\omega) \geq \lim_\omega \ell(g, \delta_n)$. By Theorem 2.16(1), the sets $\bar{\mathcal{C}}(g, M_n)$ are gate-convex. If y_n is the gate-projection of the basepoint $p_n \in M_n$ to $\bar{\mathcal{C}}(g, M_n)$, we have

$$\lim_\omega \delta_n(y_n, p_n) = \lim_\omega \delta_n(p_n, \bar{\mathcal{C}}(g, M_n)) \leq \lim_\omega \tfrac{1}{2}\delta_n(p_n, gp_n) < +\infty.$$

It follows that we have a well-defined point $y = (y_n) \in M_\omega$ and that $\delta_\omega(y, gy) = \lim_\omega \ell(g, \delta_n)$. This shows that $\ell(g, \delta_\omega) = \lim_\omega \ell(g, \delta_n)$.

Finally, since $\bar{\mathcal{C}}(g, M_\omega)$ is gate-convex, it is a closed subset of the complete median space $(M_\omega, \delta_\omega)$. Thus a point $x = (x_n) \in M_\omega$ lies in $\bar{\mathcal{C}}(g, M_\omega)$ if and only if $\delta_\omega(x, \bar{\mathcal{C}}(g, M_\omega)) = 0$, which happens if and only if $\delta_\omega(x, gx) = \ell(g, \delta_\omega)$ (again by Theorem 2.16). Equivalently, x lies in $\bar{\mathcal{C}}(g, M_\omega)$ if and only if $\lim_\omega \delta_n(x_n, \bar{\mathcal{C}}(g, M_n)) = 0$, ie if and only if $x \in \lim_\omega \bar{\mathcal{C}}(g, M_n)$. This concludes the proof. □

Lemma 7.10 *Consider again the above setting, with (M_n,δ_n) geodesic for ω–all n. Consider two elements $g,h\in G$ and $s\geq 1$.*

1. *Suppose that, for some $\mathfrak{w}\in\mathscr{W}(M_\omega)$, we have $\{\mathfrak{w},g^s\mathfrak{w}\}\subseteq\mathcal{W}_1(g,M_\omega)\cap\mathcal{W}_1(h,M_\omega)$. Then, for ω–all n, there exists $\mathfrak{w}_n\in\mathscr{W}(M_n)$ such that $\{\mathfrak{w}_n,g^s\mathfrak{w}_n\}\subseteq\mathcal{W}_1(g,M_n)\cap\mathcal{W}_1(h,M_n)$.*
2. *If there exist walls $\mathfrak{u},\mathfrak{v}\in\mathcal{W}_1(g,M_\omega)$ such that $\{\mathfrak{u},g^s\mathfrak{u}\}$ is transverse to $\{\mathfrak{v},g^s\mathfrak{v}\}$, then, for ω–all n, there exist $\mathfrak{u}_n,\mathfrak{v}_n\in\mathcal{W}_1(g,M_n)$ such that $\{\mathfrak{u}_n,g^s\mathfrak{u}_n\}$ is transverse to $\{\mathfrak{v}_n,g^s\mathfrak{v}_n\}$.*

Proof We begin with some general observations. We have already noted in Lemma 7.9 that (M_ω,δ_ω) is connected, hence g and h act stably without inversions. By parts (1) and (4) of Remark 7.8, each wall of M_ω arises from a wall of (at least) one of the trees T^i_ω. Moreover, each projection $f^i_\omega(\bar{\mathcal{C}}(g,M_\omega))$ is either fixed pointwise by g or it is a $\langle g\rangle$–invariant geodesic (and similarly for h).

We now prove part (1). By the above discussion, there exist an index i and $\mathfrak{v}\in\mathscr{W}(T^i_\omega)$ such that $\{\mathfrak{v},g^s\mathfrak{v}\}\subseteq\mathcal{W}_1(g,T^i_\omega)\cap\mathcal{W}_1(h,T^i_\omega)$. Thus, g and h are both loxodromic in T^i_ω, which implies that they are loxodromic in ω–all T^i_n. Let α_ω,α_n and β_ω,β_n be the axes in T^i_ω,T^i_n of g and h, respectively. By Lemma 7.9, we have $\alpha_\omega=\lim_\omega\alpha_n$ and $\beta_\omega=\lim_\omega\beta_n$. Since α_ω and β_ω both cross $\mathfrak{v}$ and $g^s\mathfrak{v}$, they must share a segment of length $\epsilon+s\cdot\ell(g,T^i_\omega)$ for some $\epsilon>0$.

If y and z are the endpoints of this segment, we can write $y=(y_n)=(y'_n)$ and $z=(z_n)=(z'_n)$ with $y_n,z_n\in\alpha_n$ and $y'_n,z'_n\in\beta_n$. Denoting by δ^i_n the metric of T^i_n, we have

$$\lim_\omega\delta^i_n(y_n,y'_n)=\lim_\omega\delta^i_n(z_n,z'_n)=0\quad\text{and}\quad\lim_\omega\delta^i_n(y_n,z_n)=\lim_\omega\delta^i_n(y'_n,z'_n)=\epsilon+s\cdot\lim_\omega\ell(g,T^i_n).$$

Hence α_n and β_n share a segment σ_n of length $>s\cdot\ell(g,T^i_n)$ for ω–all n. It follows that there exists a wall $\mathfrak{v}_n\in\mathscr{W}(T^i_n)$ such that σ_n crosses $\mathfrak{v}_n$ and $g^s\mathfrak{v}_n$. Hence $\{\mathfrak{v}_n,g^s\mathfrak{v}_n\}\subseteq\mathcal{W}_1(g,T^i_n)\cap\mathcal{W}_1(h,T^i_n)$, and it is clear that $\mathfrak{v}_n$ determines a wall $\mathfrak{w}_n$ of M with $\{\mathfrak{w}_n,g^s\mathfrak{w}_n\}\subseteq\mathcal{W}_1(g,M_n)\cap\mathcal{W}_1(h,M_n)$.

We now prove part (2). By Remark 7.8(4), $\mathfrak{u}$ and $\mathfrak{v}$ determine halfspaces $\mathfrak{h},\mathfrak{k}\in\mathscr{H}(M_\omega)$ satisfying $g\mathfrak{h}\subsetneq\mathfrak{h}$ and $g\mathfrak{k}\subsetneq\mathfrak{k}$. Since $\{\mathfrak{u},g^s\mathfrak{u}\}$ and $\{\mathfrak{v},g^s\mathfrak{v}\}$ are transverse, Helly's lemma implies that there exist points

$$x\in g^s\mathfrak{h}\cap g^s\mathfrak{k}\cap\bar{\mathcal{C}}(g,M_\omega),\qquad y\in g^s\mathfrak{h}\cap\mathfrak{k}^*\cap\bar{\mathcal{C}}(g,M_\omega),$$
$$z\in\mathfrak{h}^*\cap g^s\mathfrak{k}\cap\bar{\mathcal{C}}(g,M_\omega),\qquad w\in\mathfrak{h}^*\cap\mathfrak{k}^*\cap\bar{\mathcal{C}}(g,M_\omega).$$

Suppose that $\mathfrak{u}$ and $\mathfrak{v}$ arise from trees T^i_ω and T^j_ω, where g has axes α^i and α^j, respectively. Then the points $f^i_\omega(x),f^i_\omega(y),f^i_\omega(z),f^i_\omega(w)$ lie on α^i, and $\{f^i_\omega(x),f^i_\omega(y)\}$ is separated from $\{f^i_\omega(z),f^i_\omega(w)\}$ by a segment of length $>s\cdot\ell(g,T^i_\omega)$. Similarly, $\{f^j_\omega(x),f^j_\omega(z)\}$ and $\{f^j_\omega(y),f^j_\omega(w)\}$ are separated by a subsegment of α^j of length $>s\cdot\ell(g,T^j_\omega)$.

Writing $x=(x_n)$, $y=(y_n)$, $z=(z_n)$ and $w=(w_n)$, it follows that, for ω–all n, there exist walls $\mathfrak{u}'_n\in\mathcal{W}_1(g,T^i_n)$ and $\mathfrak{v}'_n\in\mathcal{W}_1(g,T^j_n)$ such that

$$\{\mathfrak{u}'_n,g^s\mathfrak{u}'_n\}\subseteq\mathscr{W}(f^i_n(x_n),f^i_n(y_n)|f^i_n(z_n),f^i_n(w_n)),$$
$$\{\mathfrak{v}'_n,g^s\mathfrak{v}'_n\}\subseteq\mathscr{W}(f^j_n(x_n),f^j_n(z_n)|f^j_n(y_n),f^j_n(w_n)).$$

Thus $\mathfrak{u}'_n,\mathfrak{v}'_n$ induce $\mathfrak{u}_n,\mathfrak{v}_n\in\mathcal{W}_1(g,M_n)$ with $\{\mathfrak{u}_n,g^s\mathfrak{u}_n\}$ transverse to $\{\mathfrak{v}_n,g^s\mathfrak{v}_n\}$; cf Lemma 4.4. □

7.3 Ultralimits of convex-cocompact actions on Salvettis

Let Γ be a finite simplicial graph, $\mathcal{A} = \mathcal{A}_\Gamma$ the associated right-angled Artin group, and $\mathcal{X} = \mathcal{X}_\Gamma$ the universal cover of its Salvetti complex. Denote by d the ℓ^1 metric on $\mathcal{X}$ and set $r = \dim \mathcal{X}$. Fix a nonprincipal ultrafilter ω.

When we speak *convex-cocompactness* in $\mathcal{A}$ from now on (Definition 3.1), this is always meant with respect to the standard action $\mathcal{A} \curvearrowright \mathcal{X}$. Note that a group G is isomorphic to a convex-cocompact subgroup of a right-angled Artin group if and only if G is the fundamental group of a compact special cube complex [68]. In particular, G must be torsionfree and finitely generated.

In the rest of Section 7 we make the following assumption.

Assumption 7.11 Let $G \leq \mathcal{A}$ be a convex-cocompact subgroup. Let $Y \subseteq \mathcal{X}$ be a G–invariant, convex subcomplex on which G acts with exactly q orbits of vertices. Let $[\mu]$ be the induced coarse median structure on G. Consider a sequence $\varphi_n \in \mathrm{Aut}(G, [\mu])$. Denote by $\rho\colon G \hookrightarrow \mathcal{A}$ the standard inclusion and set $\rho_n = \rho \circ \varphi_n$.

We say for simplicity that $g \in G$ is *label-irreducible* if $\rho(g)$ is a label-irreducible element of $\mathcal{A}$.

Remark 7.12 If $g \in G$ is label-irreducible, then Corollary 3.3 and Lemma 3.11(2) show that $\rho_n(g) \in \mathcal{A}$ is label-irreducible for all $n \geq 0$.

Let $S \subseteq G$ be a finite generating set. Choose basepoints $p_n \in Y_n$ with $\tau_S^{\rho_n}(p_n) = \overline{\tau}_S^{\rho_n}$ and define $\delta_n := d/\overline{\tau}_S^{\rho_n} \in \mathcal{D}^G(\mathcal{X})$. For ease of notation, let us write $G \curvearrowright \mathcal{X}_n$ and $G \curvearrowright Y_n$ for the actions of G on $\mathcal{X}$ and Y induced by the homomorphism ρ_n.

Recall that $\gamma\colon \mathscr{W}(\mathcal{X}) \to \Gamma^{(0)}$ is the map pairing each hyperplane with its label. For every $v \in \Gamma^{(0)}$, the hyperplanes in $\gamma^{-1}(v)$ are pairwise disjoint. Hence there is a natural simplicial tree $\mathcal{T}^v$ (usually locally infinite) that is dual to the collection $\gamma^{-1}(v)$. In the terminology of Section 2.5, the tree $\mathcal{T}^v$ is the restriction quotient of $\mathcal{X}$ associated to $\gamma^{-1}(v) \subseteq \mathscr{W}(\mathcal{X})$.

In particular, we have an $\mathcal{A}$–equivariant, surjective median morphism $\pi^v\colon \mathcal{X} \to \mathcal{T}^v$ taking cubes to edges or vertices, and an $\mathcal{A}$–equivariant, isometric median morphism $(\pi^v)\colon \mathcal{X} \hookrightarrow \prod_{v \in \Gamma} \mathcal{T}^v$.

Let $\mathcal{T}_n^v$ denote the tree $\mathcal{T}^v$ equipped with the twisted G–action induced by ρ_n and with its graph metric rescaled by $\overline{\tau}_S^{\rho_n}$. We obtain a G–equivariant, δ_n–isometric embedding $(\pi_n^v)\colon \mathcal{X}_n \hookrightarrow \prod_{v \in \Gamma} \mathcal{T}_n^v$.

Thus, our setting is a special case of the one in the second part of Section 7.2 (after Remark 7.8). If the automorphisms φ_n are pairwise distinct in $\mathrm{Out}\, G$, then we are also in a special case of Section 7.1, but we do not make this assumption for the moment.

As in Section 7.2, the sequence of actions $G \curvearrowright \mathcal{X}_n$ with metrics δ_n and basepoints p_n yields a limit action $G \curvearrowright \mathcal{X}_\omega$, along with a metric $\delta_\omega \in \mathcal{D}^G(\mathcal{X}_\omega)$, a basepoint $p_\omega \in \mathcal{X}_\omega$, and a G–equivariant, δ_ω–isometric embedding $(\pi^v_\omega)\colon \mathcal{X}_\omega \hookrightarrow \prod_{v\in\Gamma} \mathcal{T}^v_\omega$. The pair $(\mathcal{X}_\omega, \delta_\omega)$ is a complete, geodesic median space of rank $\le r$.

We now prove a sequence of fairly straightforward lemmas regarding the action of G on $\mathcal{X}_\omega$ and its median subalgebras. After that comes the most important part of this subsection, which is concerned with the notion of *cubical configurations* (Definition 7.17).

Lemma 7.13 *Consider label-irreducible elements $g, h \in G$.*

(1) *If there exist walls $\mathfrak{u}$ and $\mathfrak{w}$ with $\{\mathfrak{u}, \mathfrak{w}, h^{4r}\mathfrak{u}, g^{4r}\mathfrak{w}\} \subseteq \mathcal{W}_1(g, \mathcal{X}_\omega) \cap \mathcal{W}_1(h, \mathcal{X}_\omega)$, then $\langle g, h\rangle \simeq \mathbb{Z}$.*

(2) *There do not exist walls $\mathfrak{u}, \mathfrak{w} \in \mathcal{W}_1(g, \mathcal{X}_\omega)$ such that $\{\mathfrak{u}, g^{4r}\mathfrak{u}\}$ is transverse to $\{\mathfrak{w}, g^{4r}\mathfrak{w}\}$.*

Proof We begin with part (1). By Lemma 7.10(1), there exist hyperplanes $\mathfrak{u}_n, \mathfrak{w}_n \in \mathscr{W}(\mathcal{X}_n)$ for some n, such that $\{\mathfrak{u}_n, \mathfrak{w}_n, h^{4r}\mathfrak{u}_n, g^{4r}\mathfrak{w}_n\} \subseteq \mathcal{W}_1(g, \mathcal{X}_n) \cap \mathcal{W}_1(h, \mathcal{X}_n)$. Since $\rho_n(g)$ and $\rho_n(h)$ are label-irreducible by Remark 7.12, Lemma 3.13 guarantees that $\langle g, h\rangle \simeq \mathbb{Z}$.

Regarding part (2), if there existed such walls, Lemma 7.10(2) would yield hyperplanes $\mathfrak{u}_n, \mathfrak{w}_n \in \mathscr{W}(\mathcal{X}_n)$ for some n, such that the sets $\{\mathfrak{u}_n, g^{4r}\mathfrak{u}_n\}$ and $\{\mathfrak{w}_n, g^{4r}\mathfrak{w}_n\}$ were transverse and both contained in $\mathcal{W}_1(g, \mathcal{X}_n)$. This would violate Lemma 3.10, since $\rho_n(g)$ is label-irreducible. □

Lemma 7.14 *For every G–invariant median subalgebra $M \subseteq \mathcal{X}_\omega$, we have:*

(1) *The action $G \curvearrowright M$ has no wall inversions.*

(2) *Each element $g \in G$ is elliptic (resp. loxodromic) in M if and only if it is in $\mathcal{X}_\omega$.*

Proof Part (2) follows from part (1). Indeed, note that $\mathcal{H}_1(g, M) = \varnothing$ if and only if $\mathcal{H}_1(g, \mathcal{X}_\omega) = \varnothing$, for instance by Remark 2.18(3). Since the action $G \curvearrowright M$ has no inversions, Theorem 2.14(2) then shows that g is elliptic/loxodromic in M if and only if it is $\mathcal{X}_\omega$.

Regarding part (1), we will need the following observation:

Claim *Let an action $G \curvearrowright (T_\omega, d_\omega)$ be the ultralimit of a sequence of actions on $\mathbb{R}$–trees $G \curvearrowright (T_n, d_n)$. Suppose in addition that $g \in G$ is loxodromic in ω–all T_n. Then, for all $k \in \mathbb{Z} \setminus \{0\}$ and all $x \in T_\omega$, the point x is fixed by g^k if and only if it is fixed by g.*

Proof Let α_n be the axis of g in T_n and consider a point $y = (y_n) \in T_\omega$. Then

$$d_n(y_n, g^k y_n) = \ell(g^k, T_n) + 2d_n(y_n, \alpha_n) \ge \ell(g, T_n) + 2d_n(y_n, \alpha_n) = d_n(y_n, g y_n).$$

It follows that $d_\omega(y, g^k y) \ge d_\omega(y, gy)$ for all $k \in \mathbb{Z} \setminus \{0\}$, which proves the claim. ◁

Now, we will deduce that the action $G \curvearrowright M$ has no wall inversions from Remark 7.8(3). We need to show that, for every $v \in \Gamma$, every $x \in \mathcal{T}_\omega^v$ and every $g \in G$, we have $g^2 x = x$ if and only if $gx = x$. If $\rho_n(g)$ is loxodromic in $\mathcal{T}_n^v$ for ω–all n, this follows from the claim. If instead $\rho_n(g)$ is elliptic in $\mathcal{T}_n^v$ for ω–all n, then it follows from the observation that edge-stabilisers for the action $G \curvearrowright \mathcal{T}_n^v$ are closed under taking roots in G (since they are hyperplane-stabilisers for $G \curvearrowright \mathcal{X}_n$). $\square$

Lemma 7.15 *Consider $g \in G$ such that its label-irreducible components $g_1, \dots, g_k$ also lie in G (in general, they only lie in $\mathcal{A}$). Then, for every G–invariant median subalgebra $M \subseteq \mathcal{X}_\omega$:*

(1) *We have a partition $\mathcal{W}_1(g, M) = \mathcal{W}_1(g_1, M) \sqcup \dots \sqcup \mathcal{W}_1(g_k, M)$.*

(2) *Each wall in $\mathcal{W}_1(g_i, M)$ is preserved by each g_j with $j \neq i$.*

(3) *The sets $\mathcal{W}_1(g_1, M), \dots, \mathcal{W}_1(g_k, M)$ are transverse to each other.*

(4) *We have $\bar{\mathcal{C}}(g, M) = \bar{\mathcal{C}}(g_1, M) \cap \dots \cap \bar{\mathcal{C}}(g_k, M)$ and $\bar{\mathcal{C}}(g^m, M) = \bar{\mathcal{C}}(g, M)$ for all $m \geq 1$.*

(5) *For every $\eta \in \mathcal{PD}^G(M)$, we have $\ell(g, \eta) = \ell(g_1, \eta) + \dots + \ell(g_k, \eta)$.*

Proof Let us prove parts (1) and (2) first, except for disjointness of the sets $\mathcal{W}_1(g_i, M)$, which will follow from part (3). Note that it suffices to consider the case when $M = \mathcal{X}_\omega$. Indeed, by Remark 2.2, we have a surjection $\mathrm{res}_M : \mathscr{W}_M(\mathcal{X}_\omega) \to \mathscr{W}(M)$ and, by Remark 2.18(3), a wall $\mathfrak{w} \in \mathscr{W}_M(\mathcal{X}_\omega)$ lies in $\mathcal{W}_1(g, \mathcal{X}_\omega)$ if and only if $\mathrm{res}_M(\mathfrak{w})$ lies in $\mathcal{W}_1(g, M)$.

In fact, Remark 7.8(1) shows that it suffices to prove parts (1) and (2) "for the trees $\mathcal{T}_\omega^v$", ie prove that, for every $v \in \Gamma$, we have $\mathcal{W}_1(g, \mathcal{T}_\omega^v) = \mathcal{W}_1(g_1, \mathcal{T}_\omega^v) \cup \dots \cup \mathcal{W}_1(g_k, \mathcal{T}_\omega^v)$, and that g_j fixes the set $\mathcal{W}_1(g_i, \mathcal{T}_\omega^v)$ pointwise for $j \neq i$.

Note that distinct components g_i cannot be loxodromic in the same tree $\mathcal{T}_\omega^v$. Otherwise they would have the same axis, since they commute, and Lemma 7.13(1) would give a contradiction. Thus, at most one of the sets $\mathcal{W}_1(g_1, \mathcal{T}_\omega^v), \dots, \mathcal{W}_1(g_k, \mathcal{T}_\omega^v)$ can be nonempty, for each v.

Recalling that $g = g_1 \cdots g_k$ and that the g_i commute pairwise, we conclude that either $\mathcal{W}_1(g, \mathcal{T}_\omega^v)$ is empty, or it coincides with $\mathcal{W}_1(g_{i_v}, \mathcal{T}_\omega^v)$, where g_{i_v} is the only label-irreducible component that is loxodromic in $\mathcal{T}_\omega^v$. If $j \neq i_v$, then g_j is elliptic in $\mathcal{T}_\omega^v$ and, since it commutes with g_{i_v}, it must fix pointwise its entire axis. In particular, g_j preserves every wall in the set $\mathcal{W}_1(g_{i_v}, \mathcal{T}_\omega^v)$. This proves parts (1) and (2), except for disjointness of the sets $\mathcal{W}_1(g_i, M)$.

In order to prove part (3), note that part (2) shows that $\mathcal{W}_1(g_i, M) \subseteq \mathcal{W}_0(g_j, M)$ for $i \neq j$. By Lemma 7.14(1), the action $G \curvearrowright M$ has no wall inversions. Thus $\mathcal{W}_1(g_i, M)$ and $\mathcal{W}_1(g_j, M)$ are transverse by Theorem 2.14(3). In particular, $\mathcal{W}_1(g_i, M)$ and $\mathcal{W}_1(g_j, M)$ are disjoint, which completes the proof of part (1).

Regarding part (4), it suffices to prove the statements for $M = \mathcal{X}_\omega$. Indeed, G acts nontransversely on $\mathcal{X}_\omega$ and without inversions on M, so we have $\bar{\mathcal{C}}(g, M) = M \cap \bar{\mathcal{C}}(g, \mathcal{X}_\omega)$, for instance by [51, Proposition 3.40].

The same holds for the g_i. Now, Lemma 7.9(2) implies that $\bar{\mathcal{C}}(g,\mathcal{X}_\omega)$ coincides with $\bigcap_i \bar{\mathcal{C}}(g_i,\mathcal{X}_\omega)$ and $\bar{\mathcal{C}}(g^m,\mathcal{X}_\omega)$, since this is true for convex cores in all $\mathcal{X}_n$ (recall Lemma 3.7(2) and Remark 7.12).

Finally, we prove part (5). Parts (1) and (2) imply that a $\mathscr{B}$–measurable fundamental domain for the action $\langle g\rangle \curvearrowright \mathcal{H}_1(g,M)$ can be constructed as the disjoint union of $\mathscr{B}$–measurable fundamental domains for the actions $\langle g_i\rangle \curvearrowright \mathcal{H}_1(g_i,M)$. Since $G\curvearrowright M$ has no wall inversions, translation lengths coincide with a measure of these fundamental domains (Remark 2.19) and part (5) follows. □

Lemma 7.16 *Let $M\subseteq\mathcal{X}_\omega$ be a G–invariant median subalgebra with a pseudometric $\eta\in\mathcal{PD}^G(M)$. Consider an element $g\in G$ and a point $x\in M$.*

1. *For every $m\geq 1$, we have $\eta(x,gx)\leq\eta(x,g^m x)$.*
2. *If $h\in G$ is a label-irreducible component of g, then $\eta(x,hx)\leq\eta(x,gx)$.*

Proof Recall that the action $G\curvearrowright M$ is nontransverse by Remark 7.8(2), and without inversions by Lemma 7.14(1). Thus, Theorem 2.16(2) guarantees that $\eta(x,gx)=\ell(g,\eta)+2\eta(x,\bar{\mathcal{C}}(g,M))$.

Now, part (1) is obtained by observing that $\ell(g^m,\eta)=m\cdot\ell(g,\eta)$ and $\bar{\mathcal{C}}(g^m,M)=\bar{\mathcal{C}}(g,M)$, which follow from Remark 2.19 and Lemma 7.15(4), respectively. For part (2), it suffices to recall that $\bar{\mathcal{C}}(g,M)\subseteq\bar{\mathcal{C}}(h,M)$ and $\ell(h,\eta)\leq\ell(g,\eta)$, as shown in parts (4) and (5) of Lemma 7.15. □

We now introduce *cubical configurations*, which will be important in the proof of Theorem 7.21, hence in those of Theorems F and I. The idea is that large cubes in $\mathcal{X}_\omega$ that are moved very little by a subset $F\subseteq G$ will give rise to cubical configurations in $\mathcal{X}_\omega$ (Lemma 7.22).

After the definition, we will see how to transfer cubical configurations from $\mathcal{X}_\omega$ to $\mathcal{X}$ (Lemma 7.18) and how to use them to construct large abelian subgroups in the centraliser $Z_G(F)$ (Lemma 7.19).

Definition 7.17 Consider an action on a median algebra $G\curvearrowright M$ and a finite subset $F\subseteq G$. An (s,t,F)–*cubical configuration of width $m\geq 1$* in M is the datum of nonempty subsets $\mathcal{U}_1,\dots,\mathcal{U}_s\subseteq\mathscr{W}(M)$, walls $\mathfrak{v}_1,\dots,\mathfrak{v}_t\in\mathscr{W}(M)$ and a partition $F=F_0\sqcup\{g_1,\dots,g_t\}$ such that

1. the sets $\mathcal{U}_1,\dots,\mathcal{U}_s,\{\mathfrak{v}_1,g_1^m\mathfrak{v}_1\},\dots,\{\mathfrak{v}_t,g_t^m\mathfrak{v}_t\}$ are transverse to each other and their union is contained in $\mathcal{W}_0(f,M)$ for every $f\in F_0$,
2. for each $1\leq j\leq t$, we have $\{\mathfrak{v}_j,g_j^m\mathfrak{v}_j\}\subseteq\mathcal{W}_1(g_j,M)$, while $\mathcal{W}_0(g_j,M)$ contains $\mathcal{U}_1,\dots,\mathcal{U}_s$ and all sets $\{\mathfrak{v}_{j'},g_{j'}^m\mathfrak{v}_{j'}\}$ with $j'\neq j$.

We refer to $\mathcal{U}_1,\dots,\mathcal{U}_s$ as the *static sets* and to $g_1,\dots,g_t$ as the *skewering elements*.

The proof of the next result is quite similar in spirit to that of Lemma 7.9, but we repeat it for the reader's convenience, since it is a bit more technical.

We denote by $Y_\omega\subseteq\mathcal{X}_\omega$ the convex subset obtained as $\lim_\omega Y_n$. A subset $\mathscr{C}\subseteq\mathscr{W}(Y)$ is a *chain* if it is the set of hyperplanes associated to a set of halfspaces that is totally ordered by inclusion.

Lemma 7.18 *Suppose that the sequence φ_n is not ω–constant. Let $F \subseteq G$ be a finite subset of label-irreducible elements such that no two of them generate a cyclic subgroup. Suppose that Y_ω admits an (s,t,F)–cubical configuration of width $\geq 4r$ with skewering elements $g_1,\dots,g_t$.*

Then, for ω–all n, there exists a $(\sigma,\tau,\varphi_n(F))$–cubical configuration of width $\geq 4r$ in Y such that $\sigma+\tau=s+t$ and the $\varphi_n(g_i)$ are skewering elements (hence $\tau \geq t$ and $\sigma \leq s$). In addition, the static sets of this configuration can be taken to be arbitrarily long chains of hyperplanes.

Proof Let the cubical configuration in Y_ω consist of static sets $\mathcal{U}_1,\dots,\mathcal{U}_s$, walls $\mathfrak{v}_1,\dots,\mathfrak{v}_t$ and the partition $F=F_0 \sqcup \{g_1,\dots,g_t\}$. It suffices to assume that each $\mathcal{U}_i$ is a singleton $\{\mathfrak{u}_i\}$. Recall that the action $G \curvearrowright Y_\omega$ is nontransverse by Remark 7.8(2), and without inversions by Remark 2.12.

By Remark 7.8(1), there exist vertices $u_1,\dots,u_s,v_1,\dots,v_t \in \Gamma$ such that the walls $\mathfrak{u}_i$ and $\mathfrak{v}_j$ arise from walls $\overline{\mathfrak{u}}_i \in \mathscr{W}(\mathcal{T}_\omega^{u_i})$ and $\overline{\mathfrak{v}}_j \in \mathscr{W}(\mathcal{T}_\omega^{v_j})$, respectively. Note that each $\overline{\mathfrak{u}}_i$ is preserved by all elements of F, while $\overline{\mathfrak{v}}_j$ and $g_j^{4r}\overline{\mathfrak{v}}_j$ are preserved by $F \setminus \{g_j\}$ and cross the axis of g_j in $\mathcal{T}_\omega^{v_j}$.

Claim *There exist nontrivial arcs $\alpha_i \subseteq \mathcal{T}_\omega^{u_i}$ and $\beta_j \subseteq \mathcal{T}_\omega^{v_j}$, with endpoints $\alpha_i^\pm$ and $\beta_j^\pm$, such that:*

(a) *Each α_i is fixed by F and each β_j is fixed by $F \setminus \{g_j\}$.*

(b) *β_j is contained in the axis of g_j in $\mathcal{T}_\omega^{v_j}$ and it has length $> 4r \cdot \ell(g_j,\mathcal{T}_\omega^{v_j})$.*

(c) *These arcs induce transverse sets of walls of Y_ω. More precisely, for every $(\epsilon,\zeta) \in \{\pm\}^s \times \{\pm\}^t$, there exists a point $x^{\epsilon,\zeta} \in Y_\omega$ such that, for all i and j, the nearest-point projection of $\pi_\omega^{u_i}(x^{\epsilon,\zeta})$ to α_i is $\alpha_i^{\epsilon_i}$, and the nearest-point projection of $\pi_\omega^{v_j}(x^{\epsilon,\zeta})$ to β_j is $\beta_j^{\zeta_j}$.*

Proof The walls $\mathfrak{u}_i,\mathfrak{v}_j \in \mathscr{W}(Y_\omega)$ correspond to halfspaces $\mathfrak{u}_i^\pm,\mathfrak{v}_j^\pm \in \mathscr{H}(Y_\omega)$, which we label so that, for each j, the halfspaces $\mathfrak{v}_j^-$ and $g_j^{4r}\mathfrak{v}_j^+$ are disjoint. Since the sets $\{\mathfrak{u}_i\}$ and $\{\mathfrak{v}_j, g_j^{4r}\mathfrak{v}_j\}$ are all transverse to each other, Helly's lemma allows us to find points $x^{\epsilon,\zeta} \in Y_\omega$ so that $x^{\epsilon,\zeta}$ lies in $\mathfrak{u}_i^{\epsilon_i}$ for all i and so that, for all j, it lies in $\mathfrak{v}_j^-$ if $\zeta_j=-$ and in $g_j^{4r}\mathfrak{v}_j^+$ if $\zeta_j=+$.

For each i, there exists a point $q_i \in \mathcal{T}_\omega^{u_i}$ such that one of the two halfspaces associated to $\overline{\mathfrak{u}}_i$ is a connected component κ_i of $\mathcal{T}_\omega^{u_i} \setminus \{q_i\}$; in particular, κ_i is open. Since G acts on Y_ω without inversions, F fixes q_i and leaves κ_i invariant. Since F is finite, it fixes nontrivial arc of $\mathcal{T}_\omega^{u_i}$ with one endpoint equal to q_i and the other lying in κ_i. We let α_i be this arc, possibly shrinking it a bit to ensure that the finitely many points $\pi_\omega^{u_i}(x^{\epsilon,\zeta})$ have the correct projections to α_i.

Finally, consider an index $1 \leq j \leq t$. The set of points $x^{\epsilon,\zeta}$ with $\zeta_j=-$ is contained in $\mathfrak{v}_j^-$, whereas the set of points $x^{\epsilon,\zeta}$ with $\zeta_j=+$ is contained in $g_j^{4r}\mathfrak{v}_j^+$. Note that either $\pi_\omega^{v_j}(\mathfrak{v}_j^-)$ or $\pi_\omega^{v_j}(g_j^{4r}\mathfrak{v}_j^+)$ is an *open* halfspace of $\mathcal{T}_\omega^{v_j}$. It follows that the set of points $\pi_\omega^{v_j}(x^{\epsilon,\zeta})$ with $\zeta_j=-$ is separated by the set of points $\pi_\omega^{v_j}(x^{\epsilon,\zeta})$ with $\zeta_j=+$ by an arc β_j that is contained in the axis of g_j in $\mathcal{T}_\omega^{v_j}$ and has length $> 4r \cdot \ell(g_j,\mathcal{T}_\omega^{v_j})$. Since $F \setminus \{g_j\}$ preserves the halfspaces $\mathfrak{v}_j^-$ and $g_j^{4r}\mathfrak{v}_j^+$, we can shrink β_j a bit to ensure that it is fixed by $F \setminus \{g_j\}$, while retaining length $> 4r \cdot \ell(g_j,\mathcal{T}_\omega^{v_j})$. This concludes the proof of the claim. $\triangleleft$

Now, it is straightforward to approximate, for ω–all n, the data provided by the claim by arcs $\alpha_i(n) \subseteq \mathcal{T}_n^{u_i}$, $\beta_j(n) \subseteq \mathcal{T}_n^{v_j}$ and points $x_n^{\epsilon,\zeta} \in Y_n$ satisfying analogous conditions.

Here we need to account for the fact that some elements of F that are elliptic in one of the trees $\mathcal{T}_\omega^\bullet$ might be loxodromic in the trees $\mathcal{T}_n^\bullet$, with translation lengths converging to zero. In any case, we can ensure that the following are satisfied:

- (a′) For all $f \in F$ and $1 \le i \le s$, either f fixes α_i pointwise, or α_i is contained in the axis of f in $\mathcal{T}_n^{u_i}$ and has length $> 4r \cdot \ell(f, \mathcal{T}_n^{u_i})$; the same holds for $f \in F \setminus \{g_j\}$ and β_j.
- (b′) $\beta_j(n)$ is contained in the axis of g_j in $\mathcal{T}_n^{v_j}$ and it has length $> 4r \cdot \ell(g_j, \mathcal{T}_n^{v_j})$.
- (c′) The nearest-point projection of $\pi_n^{u_i}(x_n^{\epsilon,\zeta})$ to $\alpha_i(n)$ is $\alpha_i^{\epsilon_i}(n)$, and the nearest-point projection of $\pi_n^{v_j}(x_n^{\epsilon,\zeta})$ to $\beta_j(n)$ is $\beta_j^{\zeta_j}(n)$.

Fix a value of n such that the above are satisfied. Condition (c′) implies that the subsets of $\mathscr{W}(Y_n)$ corresponding to the arcs $\alpha_i(n)$ and $\beta_j(n)$ are all transverse to each other. Hence:

- *Each $f \in F$ fixes all arcs $\alpha_i(n), \beta_j(n)$ except at most one.* Otherwise, conditions (a′) and (b′) would yield $\mathfrak{w}, \mathfrak{w}' \in \mathcal{W}_1(f, Y_n)$ such that $\{\mathfrak{w}, f^{4r}\mathfrak{w}\}$ and $\{\mathfrak{w}', f^{4r}\mathfrak{w}'\}$ are transverse. Along with Lemma 3.10, this would contradict label-irreducibility of $\rho_n(f)$ (Remark 7.12).
- *Each of the arcs $\alpha_i(n), \beta_j(n)$ is fixed by all elements of F except at most one.* Indeed, if neither of $f_1, f_2 \in F$ fixed a given arc, then the same conditions would yield $\mathfrak{w}_1, \mathfrak{w}_2$ such that $\{\mathfrak{w}_1, f_1^{4r}\mathfrak{w}_1, \mathfrak{w}_2, f_2^{4r}\mathfrak{w}_2\} \subseteq \mathcal{W}_1(f_1, Y_n) \cap \mathcal{W}_1(f_2, Y_n)$. Since $\rho_n(f_1)$ and $\rho_n(f_2)$ are label-irreducible, Lemma 3.13 would then imply that $\langle f_1, f_2\rangle \simeq \mathbb{Z}$, contradicting our assumptions.

In conclusion, up to reordering, there exists $0 \le \sigma \le s$ such that, for $1 \le i \le \sigma$, the arcs $\alpha_i(n)$ are fixed by the whole F, while, for $\sigma < i \le s$, there exists $f_i \in F$ such that f_i contains $\alpha_i(n)$ in its axis and $\alpha_i(n)$ is fixed by $F \setminus \{f_i\}$. We obtain a (σ, τ, F)–cubical configuration of width $\ge 4r$ in Y_n, where the static sets are given by the hyperplanes of Y_n originating from the arcs $\alpha_i(n)$ with $i \le \sigma$, while the skewering elements are $f_{\sigma+1}, \dots, f_s$ and $g_1, \dots, g_t$ (thus, $\tau = s + t - \sigma$).

This immediately translates into a $(\sigma, \tau, \varphi_n(F))$–cubical configuration of width $\ge 4r$ in Y. The fact that the static sets can be taken with arbitrarily large cardinality is also immediate, recalling that, since φ_n is not ω–constant, the scaling factors $\overline{\tau}_S^{\rho_n}$ diverge; cf Assumption 7.1 above. □

The next result only requires the material in Section 3.3 for its proof. However, it is best stated in terms of cubical configurations, as defined above.

Recall that q is the number of orbits of vertices for the action $G \curvearrowright Y$.

Lemma 7.19 *Let $F \subseteq G$ be a finite set of label-irreducible elements. Suppose that there is an (s, t, F)–cubical configuration of width $\ge 4r$ in Y, where all the static sets are chains, each containing $\ge q$ hyperplanes. Then the centraliser $Z_G(F)$ contains a copy of $\mathbb{Z}^k$ with $k = s + t$.*

Proof Let the cubical configuration consist of static sets $\mathcal{U}_1,\dots,\mathcal{U}_s$, hyperplanes $\mathfrak{v}_1,\dots,\mathfrak{v}_t$ and the partition $F = F_0 \sqcup \{g_1,\dots,g_t\}$.

Form a set $\mathcal{U}_i'$ by adding to $\mathcal{U}_i$ all hyperplanes of Y that separate hyperplanes of $\mathcal{U}_i$. This guarantees that there exist vertices $x_i, y_i \in Y$ such that $\mathscr{W}(x_i|y_i) = \mathcal{U}_i'$. The sets $\mathcal{U}_i'$ and $\{\mathfrak{v}_j, g_j^{4r}\mathfrak{v}_j\}$ are still all transverse to each other and the elements of F still fix each $\mathcal{U}_i'$ pointwise.

Since $\#\mathscr{W}(x_i|y_i) \geq \#\mathcal{U}_i \geq q$, any geodesic joining x_i to y_i must contain two points in the same G–orbit. Thus, there exist $z_i \in Y$ and $h_i \in G\setminus\{1\}$ with $\mathscr{W}(z_i|h_iz_i) \subseteq \mathcal{U}_i'$ for all $1 \leq i \leq s$.

Lemma 3.15 shows that the h_i commute pairwise. In addition, for every $f \in F$, we have $\mathscr{W}(z_i|h_iz_i) \subseteq \mathcal{U}_i' \subseteq \mathcal{W}_0(f)$, which is transverse to $\mathcal{W}_1(f)$. Since $\mathcal{W}_1(f)$ contains $\mathscr{W}(z|fz)$ for any $z \in \overline{\mathcal{C}}(f)$, another application of Lemma 3.15 guarantees that h_i and f commute. Finally, for $1 \leq j \leq t$, the hyperplanes $\mathfrak{v}_j$ and $g_j^{4r}\mathfrak{v}_j$ are preserved by all elements of $F\setminus\{g_j\}$. Since g_j is label-irreducible, Corollary 3.14(1) implies that g_j commutes with every element of F.

In conclusion, we have shown that the subgroup generated by $A := \{h_1,\dots,h_s,g_1,\dots,g_t\}$ is abelian and contained in $Z_G(F)$. We are left to show that A is a basis for $\langle A\rangle$.

Observe that $\mathfrak{v}_j$ is preserved by all elements of $A\setminus\{g_j\}$, but lies in $\mathcal{W}_1(g_j)$. Similarly, there exist hyperplanes $\mathfrak{u}_i \in \mathscr{W}(z_i|h_iz_i)$ that are preserved by $A\setminus\{h_i\}$, but lie in $\mathcal{W}_1(h_i)$. If a product $h_1^{m_1}\cdots h_s^{m_s}\cdot g_1^{n_1}\cdots g_t^{n_t}$ represents the identity, then it must preserve all hyperplanes $\mathfrak{u}_i$ and $\mathfrak{v}_j$, which implies that $m_i = 0$ and $n_j = 0$ for all i, j. This concludes the proof. □

7.4 Ultralimits of Salvettis and the WNE property

This subsection is devoted to the proof of Theorems F and I. We keep the exact same setting as the previous subsection:

Assumption 7.20 Let $G \leq \mathcal{A}$ be a convex-cocompact subgroup. Let $Y \subseteq \mathcal{X}$ be a G–invariant, convex subcomplex on which G acts with q orbits of vertices. Set $r = \dim \mathcal{X}$. Denote by d the ℓ^1 metric on $\mathcal{X}$ and Y. Let $[\mu]$ be the induced coarse median structure on G.

Consider a sequence $\varphi_n \in \mathrm{Aut}(G,[\mu])$. Denote by $\rho\colon G \hookrightarrow \mathcal{A}$ the standard inclusion and set $\rho_n = \rho\circ\varphi_n$. Fixing a nonprincipal ultrafilter ω, define $\mathcal{X}_\omega$, $\mathcal{X}_n$, Y_n and $Y_\omega = \lim_\omega Y_n$ as in Section 7.3.

The following result is the coronation of our efforts from Section 3.3 and the previous portion of Section 7. Its second part (with $k = 1$) proves Theorem I, while its first part is the last remaining ingredient in the proof of Theorem F (together with Corollary 6.23).

Theorem 7.21 *Let $F \subseteq G$ be a finite subset and suppose that* **one** *of the following holds.*

(1) *Let φ_n not be ω–constant. Let $M \subseteq Y_\omega$ be a G–invariant median subalgebra and consider $\eta \in \mathcal{PD}^G(M)$. There exists a k–cube $C \subseteq M$ such that, for any two distinct points $x, y \in C$,*

$$\eta(x,y) > 4r^2q\cdot[\tau_F^\eta(x)+\tau_F^\eta(y)].$$

(2) *There exists a (generalised) k–cube $C \subseteq Y^{(0)}$ such that, for any two distinct points $x, y \in C$,*

$$d(x,y) > \left(2r^2q + \tfrac{1}{2}rq \cdot \max\{4r, q\}\right) \cdot [\tau_F^d(x) + \tau_F^d(y)].$$

Then the centraliser $Z_G(F)$ contains a copy of $\mathbb{Z}^k$.

The theorem will follow quickly from Lemma 7.22 below, which constructs a cubical configuration in Y_ω (in case (1)) or directly in Y (in case (2)). Indeed, we can then use Lemma 7.18 to always obtain a cubical configuration in Y, and this yields the required copy of $\mathbb{Z}^k$ in $Z_G(F)$ by Lemma 7.19.

Lemma 7.22 *Consider the setting of Theorem 7.21. There exists an (s,t,F')–cubical configuration of width $\geq 4r$ in Y_ω (in case (1)) or in Y (in case (2)), where $s+t=k$ and $F' \subseteq G$ is a finite subset with $Z_G(F') = Z_G(F)$. All elements of F' are label-irreducible and no two of them generate a cyclic subgroup.*

In addition, in case (2), the static sets are chains of hyperplanes of cardinality $\geq q$.

Proof We prove the lemma simultaneously in the two cases of the theorem. In fact, in this proof it is irrelevant whether the φ_n are ω–constant or not, so we can view case (2) as a special instance of case (1) by taking $\varphi_n \equiv \mathrm{id}_G$, $Y_\omega = Y$, $M = Y^{(0)}$ and $\eta = d$.

Recall that, by Remark 7.8(2) and Lemma 7.14(1), the action $G \curvearrowright M$ is nontransverse and without inversions. We begin by constructing the finite subset $F' \subseteq G$.

Claim 1 *There exists $F' \subseteq G$ such that $Z_G(F') = Z_G(F)$, all elements of F' are label-irreducible and no two of them generate a cyclic subgroup. In addition, $\tau_{F'}^\eta(x) \leq q \cdot \tau_F^\eta(x)$ for all $x \in M$.*

Proof Let $a_1, \dots, a_N \in \mathcal{A}$ be a choice of generator for each maximal cyclic subgroup of $\mathcal{A}$ that contains a label-irreducible component of an element of F. Let $m_i \geq 1$ be the smallest integer such that $a_i^{m_i}$ lies in G; by Lemma 3.16, m_i is well-defined and, by Remark 3.17, we have $1 \leq m_i \leq q$. Define $F' := \{a_i^{m_i} \mid 1 \leq i \leq N\}$.

It is clear that every element of F' is label-irreducible and that any two elements of F' generate a noncyclic subgroup. Since all nontrivial powers of any given element of $\mathcal{A}$ have the same centraliser, Lemma 3.7(3) implies that $Z_G(F') = Z_G(F)$.

For every a_i, there exist $n \geq 1$ and $f \in F$ such that a_i^n is a label-irreducible component of f. Thus $a_i^{nm_i}$ is a label-irreducible component of f^{m_i}. Applying Lemma 7.16, it follows that

$$\eta(x, a_i^{m_i}x) \leq \eta(x, a_i^{nm_i}x) \leq \eta(x, f^{m_i}x) \leq m_i \cdot \eta(x, fx) \leq q \cdot \eta(x, fx) \leq q \cdot \tau_F^\eta(x).$$

Hence $\tau_{F'}^\eta(x) \leq q \cdot \tau_F^\eta(x)$ for all $x \in M$, as required. ◁

Now, consider the multibridge $\mathcal{B}(F') \subseteq M$ introduced in Definition 6.8. Pick any fibre $P = \mathcal{B}_{/\!/}(F') \times \{*\}$. Let $\pi_P : M \to P$ be the gate-projection.

Claim 2 *The set $C' := \pi_P(C)$ is again a k–cube and, for all distinct points $x', y' \in C'$, we have $\eta(x', y') \geq 4r^2 \cdot \overline{\tau}^{\eta}_{F'}$. Under the assumptions of case* (2)*, we further have $\eta(x', y') \geq rq \cdot \overline{\tau}^{\eta}_{F'}$.*

Proof If $x, y \in C$ are distinct, note that we have $\eta(x, y) > 4r^2 q \cdot [\tau^{\eta}_F(x) + \tau^{\eta}_F(y)]$ under the assumptions of both case (1) and case (2). Also recall that, by Proposition 6.11(6), we have $\eta(x, \pi_P(x)) \leq 2r^2 \cdot \tau^{\eta}_{F'}(x)$ for all $x \in M$. Combining these inequalities with Claim 1, we obtain

$$\begin{aligned}\eta(\pi_P(x), \pi_P(y)) &\geq \eta(x, y) - 2r^2 \cdot [\tau^{\eta}_{F'}(x) + \tau^{\eta}_{F'}(y)] \\ &> 4r^2 q \cdot [\tau^{\eta}_F(x) + \tau^{\eta}_F(y)] - 2r^2 \cdot [\tau^{\eta}_{F'}(x) + \tau^{\eta}_{F'}(y)] \\ &\geq 2r^2 \cdot [\tau^{\eta}_{F'}(x) + \tau^{\eta}_{F'}(y)] \geq 4r^2 \cdot \overline{\tau}^{\eta}_{F'}.\end{aligned}$$

In particular, distinct points of C project to distinct points of C', which guarantees that C' is again a k–cube. Moreover, $\eta(x', y') \geq 4r^2 \cdot \overline{\tau}^{\eta}_{F'}$ whenever x', y' are distinct points of C'.

Under the assumptions of case (2), we also have $\eta(x, y) > (2r^2 q + rq^2/2) \cdot [\tau^{\eta}_F(x) + \tau^{\eta}_F(y)]$ for all $x, y \in C$. Using this instead of $\eta(x, y) > 4r^2 q \cdot [\tau^{\eta}_F(x) + \tau^{\eta}_F(y)]$ in the above chain of inequalities, we obtain $\eta(\pi_P(x), \pi_P(y)) \geq rq \cdot \overline{\tau}^{\eta}_{F'}$, as required. ◁

Let $\{C'_{i,-}, C'_{i,+}\}$ be the k pairs of opposite codimension-1 faces of C'. Setting $\mathcal{H}_i := \mathscr{H}(C'_{i,-} | C'_{i,+})$, we obtain sets of halfspaces $\mathcal{H}_1, \dots, \mathcal{H}_k \subseteq \mathscr{H}(M)$ that are transverse to each other. If ν_η is the measure introduced in Remark 2.9, we have $\nu_\eta(\mathcal{H}_i) > 4r^2 \cdot \overline{\tau}^{\eta}_{F'}$ by Claim 2.

The set $\mathcal{H}_i$ can be partitioned into at most r measurable subsets such that no two halfspaces in the same subset are transverse; this follows from Corollary A.3, proved in the appendix (note that $\mathcal{D}(M) \neq \varnothing$ since $\mathcal{D}(\mathcal{X}_\omega) \neq \varnothing$, even though η is just a pseudometric). Define $\mathcal{H}'_i \subseteq \mathcal{H}_i$ as the subset with the largest measure among those in this partition. No two halfspaces in $\mathcal{H}'_i$ are transverse, and

$$\nu_\eta(\mathcal{H}'_i) \geq \frac{1}{r} \cdot \nu_\eta(\mathcal{H}_i) > 4r \cdot \overline{\tau}^{\eta}_{F'}.$$

Let $\mathcal{U}'_i \subseteq \mathscr{W}(M)$ be the set of walls associated to $\mathcal{H}'_i$. Recall that

$$\mathcal{U}'_i \subseteq \mathscr{W}_{C'}(M) \subseteq \mathscr{W}_P(M) \subseteq \bigcap_{f \in F'} \mathscr{W}_{\overline{\mathcal{C}}(f)}(M) = \bigcap_{f \in F'} \big(\mathcal{W}_1(f, M) \sqcup \mathcal{W}_0(f, M)\big).$$

Since the sets $\mathcal{W}_1(f, M)$ and $\mathcal{W}_0(f, M)$ are transverse, while no two walls in $\mathcal{U}'_i$ are transverse, we must have either $\mathcal{U}'_i \subseteq \mathcal{W}_1(f, M)$ or $\mathcal{U}'_i \subseteq \mathcal{W}_0(f, M)$ for every index i and element $f \in F'$. Consider the partitions $F' = F_i \sqcup F_i^{\perp}$ such that $\mathcal{U}'_i \subseteq \mathcal{W}_1(f, M)$ if $f \in F_i$ and $\mathcal{U}'_i \subseteq \mathcal{W}_0(f, M)$ if $f \in F_i^{\perp}$.

Claim 3 *We have $\#F_i \leq 1$ for all $1 \leq i \leq k$, and $F_i \cap F_j = \varnothing$ for $i \neq j$.*

Proof For every $f \in F'$, we have

$$\nu_\eta(\mathcal{H}_i') > 4r \cdot \overline{\tau}^{\eta}_{F'} \geq 4r \cdot \ell(f,\eta) = \ell(f^{4r},\eta).$$

If $\mathcal{U}_i' \subseteq \mathcal{W}_1(f,M)$, it follows that there exists a wall $\mathfrak{w} \in \mathcal{U}_i'$ such that $f^{4r}\mathfrak{w} \subseteq \mathcal{U}_i'$.

Thus, if $f,g \in F_i$, there exist walls $\mathfrak{w}$ and $\mathfrak{w}'$ such that $\{\mathfrak{w}, f^{4r}\mathfrak{w}, \mathfrak{w}', g^{4r}\mathfrak{w}'\} \subseteq \mathcal{W}_1(f,M) \cap \mathcal{W}_1(g,M)$. By Remarks 2.2(1) and 2.18(3), there is an analogous inclusion involving walls of $\mathcal{X}_\omega$, so Lemma 7.13(1) yields $\langle f,g\rangle \simeq \mathbb{Z}$. Since $f,g \in F'$, this means that $f=g$. Hence $\#F_i \leq 1$.

Suppose towards a contradiction that there exists $f \in F_i \cap F_j$. Then there are walls $\mathfrak{w}_i,\mathfrak{w}_j$ such that $\{\mathfrak{w}_i, f^{4r}\mathfrak{w}_i\} \subseteq \mathcal{U}_i'$ and $\{\mathfrak{w}_j, f^{4r}\mathfrak{w}_j\} \subseteq \mathcal{U}_j'$. In particular, the subsets $\{\mathfrak{w}_i, f^{4r}\mathfrak{w}_i\}$ and $\{\mathfrak{w}_j, f^{4r}\mathfrak{w}_j\}$ are transverse to each other and contained in $\mathcal{W}_1(f,M)$. Again, Remarks 2.2(1) and 2.18(3) give walls of $\mathcal{X}_\omega$ with the same properties, which contradicts Lemma 7.13(2). ◁

Up to permuting the sets $\mathcal{U}_i'$, we can assume that there exists an index $0 \leq s \leq k$ such that $F_i = \varnothing$ for $1 \leq i \leq s$, while $F_i = \{f_i\}$ for $s < i \leq k$ and pairwise-distinct elements $f_i \in F'$. This all follows from Claim 3. Each $f \in F'$ preserves every wall in $\mathcal{U}_i'$ except if $i > s$ and $f = f_i$. In addition, the proof of Claim 3 gives walls $\mathfrak{v}_i$ for $i > s$, such that $\{\mathfrak{v}_i, f_i^{4r}\mathfrak{v}_i\} \subseteq \mathcal{U}_i' \subseteq \mathcal{W}_1(f_i,M)$. Also recall that the sets $\mathcal{U}_i'$ are all transverse to each other.

This gives an $(s,k-s,F)$–cubical configuration of width $\geq 4r$ in M, where $\mathcal{U}_1',\dots,\mathcal{U}_s'$ are the static sets and $f_{s+1},\dots,f_k$ are the skewering elements. Since $M \subseteq Y_\omega$, it is straightforward to transfer this to a cubical configuration in Y_ω with the same parameters using Remarks 2.2(1) and 2.18(3). This completes the proof of the lemma in Case (1).

In Case (2), we are left to show that the static sets $\mathcal{U}_1',\dots,\mathcal{U}_s'$ contain at least q hyperplanes each. Recall that $M = Y^{(0)}$ and $\eta = d$, so ν_η is just the counting measure. By Claim 2, we have $\#\mathcal{H}_i' = \nu_\eta(\mathcal{H}_i') \geq (1/r)\nu_\eta(\mathcal{H}_i) \geq q \cdot \overline{\tau}^{\eta}_{F'} \geq q$, which concludes the proof since $\#\mathcal{H}_i' = \#\mathcal{U}_i'$. □

Combining Lemmas 7.18, 7.19 and 7.22, we can finally prove Theorem 7.21.

Proof of Theorem 7.21 Our goal is to construct an (s,t,F'')–cubical configuration of width $\geq 4r$ in Y, where $s+t=k$, the centraliser $Z_G(F'')$ is isomorphic to $Z_G(F)$, all elements of F'' are label-irreducible, and all static sets are chains of hyperplanes of cardinality $\geq q$. If we manage to do this, then Lemma 7.19 guarantees that $Z_G(F) \simeq Z_G(F'')$ contains the required copy of $\mathbb{Z}^k$.

In case (2) of the theorem, a cubical configuration with these properties is provided by Lemma 7.22. In case (1), we first apply Lemma 7.22 to obtain an (s,t,F')–cubical configuration of width $\geq 4r$ in Y_ω, where $Z_G(F') = Z_G(F)$. Then we obtain the required cubical configuration in Y from Lemma 7.18, with $F'' = \varphi_n(F')$ for some n; this is the only place where it is important that φ_n is not ω–constant. Since $Z_G(F'') = \varphi_n(Z_G(F')) \simeq Z_G(F)$, this proves the theorem. □

The following two corollaries collect the key takeaways from Theorem 7.21 that we will need in the rest of the paper.

Corollary 7.23 *Every special group with trivial centre is UNE (Definition 2.36).*

Proof Let G be a special group with trivial centre. Embed G as a convex-cocompact subgroup of a RAAG and apply Theorem 7.21(2), taking $k=1$ and letting F be an arbitrary generating set for G. This shows that the proper cocompact action $G \curvearrowright Y$ is UNE, hence G is a UNE group. □

Corollary 7.24 *Consider the setting of Assumption 7.20. Suppose that the φ_n are pairwise distinct.*

(1) *If $C \subseteq Y_\omega$ is a k–cube and $H \leq G$ fixes C pointwise, then $Z_G(H)$ contains a copy of $\mathbb{Z}^k$.*

(2) *Let G have trivial centre. Then, for every G–invariant median subalgebra $M \subseteq Y_\omega$, the action $G \curvearrowright M$ is WNE (in the sense of Definition 6.21).*

Proof By Remark 3.8, it suffices to prove part (1) under the additional assumption that H is finitely generated. So let us suppose that H is generated by a finite set F that fixes the k–cube C. We have observed in Section 7.3 that $\mathcal{D}(\mathcal{X}_\omega)^G \neq \varnothing$. Applying Theorem 7.21(1) to any choice of $\eta \in \mathcal{D}(Y_\omega)^G$, we obtain the required copy of $\mathbb{Z}^k$ inside $Z_G(F)=Z_G(H)$.

Part (2) also follows from Theorem 7.21(1), setting $k=1$ and letting F generate G. □

The following implies parts (1) and (2) of Theorem F as a special case; parts (3) and (4) are obtained below in Remark 7.27. Note that the essentiality requirement in Theorem 7.25(3) is equivalent to the minimality requirement in Theorem F(2), because of [51, Theorem C].

Recall that we denote by $\pi : \mathrm{Aut}\, G \to \mathrm{Out}\, G$ the quotient projection. If G has trivial centre and $A \leq \mathrm{Out}\, G$ is a subgroup, we have $G \lhd \pi^{-1}(A)$ and $\pi^{-1}(A)/G \simeq A$.

Theorem 7.25 *Let $G \leq \mathcal{A}$ be a convex-cocompact subgroup with trivial centre. Let $[\mu]$ be the induced coarse median structure on G. Let $A \leq \mathrm{Out}(G,[\mu])$ be an infinite abelian subgroup. Then there exists an action $\pi^{-1}(A) \curvearrowright X$ with the following properties:*

(1) *X is a geodesic median space X with $\mathrm{rk}\, X \leq r$.*

(2) *$\pi^{-1}(A) \curvearrowright X$ is an action by homotheties.*

(3) *The restriction $G \curvearrowright X$ is isometric, essential and with unbounded orbits.*

(4) *If $C \subseteq X$ is a k–cube and $H \leq G$ fixes C pointwise, then $Z_G(H)$ contains a copy of $\mathbb{Z}^k$.*

Proof Consider a sequence of pairwise distinct automorphisms $\varphi_n \in A$ and set $\rho_n = \rho \circ \varphi_n$. Choose a finite generating set $S \subseteq G$ and consider the action $G \curvearrowright Y_\omega$ as in Section 7.3.

Corollary 7.23 shows that G is UNE. Thus, denoting by $\mathrm{Aut}\, Y_\omega$ the group of automorphisms of the underlying median algebra, Proposition 7.3 yields a homomorphism $\zeta : \pi^{-1}(A) \to \mathrm{Aut}\, Y_\omega$ that extends the isometric action $G \curvearrowright Y_\omega$.

By Corollary 7.24(2), the action $G \curvearrowright Y_\omega$ is WNE. Thus, Corollary 6.23 yields a nonempty, countable, $\pi^{-1}(A)$–invariant, median subalgebra $\mathfrak{M} \subseteq Y_\omega$, and a pseudometric $\eta \in \mathcal{PD}^G(\mathfrak{M}) \setminus \{0\}$ for which $\overline{\tau}^{\eta}_S > 0$ and $\pi^{-1}(A) \curvearrowright (\mathfrak{M}, \eta)$ is homothetic.

Let $(\mathfrak{M}_\circ, \delta)$ be the quotient median space obtained by identifying points $x, y \in \mathfrak{M}$ with $\eta(x, y) = 0$. By Remark 2.1, we have $\operatorname{rk} \mathfrak{M}_\circ \leq \operatorname{rk} \mathfrak{M} \leq \operatorname{rk} X_\omega \leq r$. Since $\overline{\tau}^{\delta}_S = \overline{\tau}^{\eta}_S > 0$, the action $G \curvearrowright \mathfrak{M}_\circ$ does not have a global fixed point. Moreover, since the action $G \curvearrowright \mathfrak{M}$ has no wall inversions by Lemma 7.14(1), the action $G \curvearrowright \mathfrak{M}_\circ$ also has no inversions. Theorem 2.14(2) then guarantees that G acts on $\mathfrak{M}_\circ$ with unbounded orbits.

Theorem 7.21(1) (applied to the pseudometric η on $\mathfrak{M}$) shows that $G \curvearrowright \mathfrak{M}_\circ$ satisfies part (4). Thus, we are only left to ensure that the median space is geodesic and the action essential.

In order to make our space geodesic, note that the homothetic $\pi^{-1}(A)$–action extends to the metric completion $\overline{\mathfrak{M}}_\circ$ of $\mathfrak{M}_\circ$. This is a median space of rank $\leq r$ by [29, Proposition 2.21] and [50, Lemma 2.5]. Note that $G \curvearrowright \overline{\mathfrak{M}}_\circ$ still satisfies part (4) because of Theorem 7.21(1). Now, "filling in cubes" as in [48, Corollary 2.16], the space $\overline{\mathfrak{M}}_\circ$ embeds into a complete, connected median space Z of the same rank. By [17, Lemma 4.6], the space Z is geodesic. The isometric G–action extends to Z and one can similarly check that so does the homothetic $\pi^{-1}(A)$–action.

Summing up, we have constructed an action $\pi^{-1}(A) \curvearrowright Z$ that satisfies conditions (1)–(4), possibly except essentiality of the G–action (in addition, Z is complete). By Theorem 2.14(4), there exists a $\pi^{-1}(A)$–invariant, nonempty, convex subset $K \subseteq Z$ and a $\pi^{-1}(A)$–invariant splitting $K = K_0 \times K_1$ such that the action $G \curvearrowright K_1$ is essential. We conclude by taking $X = K_1$. (Note that K is not closed in Z in general, so we may have lost completeness along the way.) $\square$

Remark 7.26 In Theorem 7.25, we cannot both require the space X to be complete and the action $G \curvearrowright X$ to be essential. There is a very good reason for this.

Consider the special case where G is hyperbolic. Then Y_ω is an $\mathbb{R}$–tree, which forces X to also be an $\mathbb{R}$–tree. Note that an isometric action on an $\mathbb{R}$–tree is essential if and only if it is minimal.

Let us show that, if G is a finitely generated group and $G \curvearrowright T$ is a minimal action on a complete $\mathbb{R}$–tree not isometric to $\mathbb{R}$, then no homothety $\Phi\colon T \to T$ with factor $\lambda \neq 1$ can normalise G.

If G is generated by $s_1, \dots, s_k$ and $x \in T$ is any point, the union of all segments $g[x, s_i x]$ with $g \in G$ is a G–invariant subtree. Since $G \curvearrowright T$ is minimal, T must be covered by the segments $g[x, s_i x]$. In particular, the action $G \curvearrowright T$ is cocompact. If Φ normalised G, then every orbit of $G \curvearrowright T$ would be dense; see eg [88, Proposition 3.10]. Since $T \not\simeq \mathbb{R}$, this implies that each segment $g[x, s_i x]$ is nowhere-dense. This violates Baire's theorem, since a complete metric space cannot be covered by countably many nowhere-dense subsets.

I learned this argument from [54, Example II.6].

The following proves parts (3) and (4) of Theorem F.

Remark 7.27 Consider the special case of Theorem 7.25 with $A=\mathbb{Z}$, generated by an outer automorphism $\phi \in \mathrm{Out}(G,[\mu])$. Picking a representative $\varphi \in \mathrm{Aut}(G,[\mu])$, we have $\pi^{-1}(A) = G \rtimes_\varphi \mathbb{Z}$. The theorem gives an isometric action $G \curvearrowright X$ and a homothety $H\colon X \to X$ of factor λ such that $H \circ g = \varphi(g) \circ H$ for all $g \in G$. We keep the notation of the proof of Theorem 7.25.

(1) Each $g \in \mathrm{Fix}\,\varphi$ is elliptic in X. Indeed, Lemma 7.9 shows that g is elliptic in $\mathcal{X}_\omega$, since $\ell(\varphi^n(g),\mathcal{X})$ does not diverge. Lemma 7.14(2) then implies that g is elliptic in $\mathfrak{M}$, and it is clear that a fixed point in $\mathfrak{M}$ will translate into a fixed point in X.

Recalling that $\mathrm{Fix}\,\varphi$ is finitely generated by Theorem B, Theorem 2.14(2) actually implies that $\mathrm{Fix}\,\varphi$ has a global fixed point $x_0 \in X$. This proves Theorem F(3).

(2) Fix a finite generating set $S \subseteq G$. Recall from Section 2.1, that we denote conjugacy length by $\|\cdot\|_S$. Let $\Lambda(\varphi)$ be the maximal exponential growth rate of the quantity $\|\varphi^n(g)\|_S^{1/n}$,

$$\Lambda(\varphi) := \sup_{g\in G} \limsup_{n\to+\infty} \|\varphi^n(g)\|_S^{1/n}.$$

Note that $\Lambda(\varphi)$ is independent of the generating set S. For every $g \in G$, we have

$$\lambda^n \ell(g,X) = \ell(H^n g H^{-n}, X) = \ell(\varphi^n(g), X) \le \|\varphi^n(g)\|_S \cdot \overline{\tau}_S^X,$$

where the last inequality follows from the identities in Section 2.1. Since there exist elements $g \in G$ with $\ell(g,X) > 0$, we deduce that $\lambda \le \Lambda(\varphi)$ and, similarly, $\lambda^{-1} \le \Lambda(\varphi^{-1})$.

If φ has *subexponential growth* (in the sense that $\Lambda(\varphi) = \Lambda(\varphi^{-1}) = 1$), then these inequalities force $\lambda = 1$. Hence the homothetic action $G \rtimes_\varphi \mathbb{Z} \curvearrowright X$ provided by Theorem 7.25 is actually isometric, which proves Theorem F(4).

Appendix Measurable partitions of halfspace-intervals

This appendix is devoted to the proof of Corollary A.3 below. This is needed in the proof of Theorem 7.21 in order to get the exact constant $4r^2q$, and could be avoided if we contented ourselves with the worse bound $4rq\cdot\#\Gamma^{(0)}$. However, Corollary A.3 is important in the general theory of median spaces and we think it is likely to prove useful elsewhere.

Let M be a median algebra. Given a subset $P \subseteq M \times M$, let us write $\mathcal{H}_P := \bigcup_{(x,y)\in P} \mathscr{H}(x|y)$.

Lemma A.1 *Every subset $P \subseteq [0,1]^n \times [0,1]^n$ contains a countable subset $\Delta \subseteq P$ with $\mathcal{H}_\Delta = \mathcal{H}_P$.*

Proof First, we prove the case $n=1$. We can assume that $x<y$ for every $(x,y)\in P$.

Let $\Omega(P) \subseteq [0,1]$ be the union of the closed arcs $I(x,y)$ with $(x,y)\in P$. Let $\mathcal{D}(P)$ be the set of points that lie in the interior of $\Omega(P)$, but not in the interior of any arc $I(x,y)$ with $(x,y)\in P$. Thus each point of $\Omega(P)$ lies either in the frontier of $\Omega(P)$, or in the interior of some $I(x,y)$, or in the set $\mathcal{D}(P)$, and these three possibilities are disjoint. There is a unique partition of $\Omega(P)$ into maximal segments J_i (closed, open, or half-open) such that

- the interior of J_i does not intersect $\mathcal{D}(P)$, and
- if J_i intersects the interior of $I(x,y)$ for some $(x,y)\in P$, then $I(x,y)\subseteq J_i$.

Observe that $\mathcal{H}_P=\bigsqcup_i \mathscr{H}_{J_i}([0,1])\cap\mathscr{H}(0|1)$.

It is classical to see that there exists a countable subset $\Delta\subseteq P$ with $\Omega(\Delta)=\Omega(P)$. Note that $\mathcal{D}(\Delta)$ is countable and it contains $\mathcal{D}(P)$. Adding to Δ countably many pairs in P, we can thus ensure that $\mathcal{D}(\Delta)=\mathcal{D}(P)$. Hence, P and Δ determine the same the segments J_i, and $\mathcal{H}_P=\mathcal{H}_\Delta$.

Now consider a general $n\geq 1$. Let $I_i\subseteq[0,1]^n$ be the segment where all coordinates but the i^{th} vanish. Let $\pi_i\colon[0,1]^n\to I_i$ be the coordinate projections. Setting $P_i:=(\pi_i\times\pi_i)(P)\subseteq[0,1]^n\times[0,1]^n$, we have $\mathcal{H}_P=\bigcup_i\mathcal{H}_{P_i}$. By the case $n=1$, there exist countable subsets $\Delta_i\subseteq P_i$ with $\mathcal{H}_{\Delta_i}=\mathcal{H}_{P_i}$. Choosing countable sets $\Delta_i'\subseteq P$ with $(\pi_i\times\pi_i)(\Delta_i')=\Delta_i$, we have $\mathcal{H}_{\Delta_i}\subseteq\mathcal{H}_{\Delta_i'}\subseteq\mathcal{H}_P$. Hence, taking $\Delta:=\bigcup_i\Delta_i'$, we obtain $\mathcal{H}_P=\mathcal{H}_\Delta$. □

Recall that $\mathscr{B}(M)$ is the σ–algebra generated by halfspace-intervals, as in Remark 2.9.

Lemma A.2 *Suppose that $M\subseteq[0,1]^n$ is a median subalgebra containing the points $\underline{0}=(0,\dots,0)$ and $\underline{1}=(1,\dots,1)$. Let $\pi_i\colon M\to[0,1]$ denote the coordinate projections. Then the induced maps $\pi_i^*\colon\mathscr{H}([0,1])\to\mathscr{H}(M)$ (as in Remark 2.1) map $\mathscr{B}$–measurable sets to $\mathscr{B}$–measurable sets.*

Proof Since π_i^* is injective, we have

$$\pi_i^*(\mathscr{H}([0,1])\setminus E)=\pi_i^*(\mathscr{H}(0|1))\cup\pi_i^*(\mathscr{H}(1|0))\setminus\pi_i^*(E)$$

for every $E\subseteq\mathscr{H}([0,1])$. Thus, it suffices to show that, for all $0\leq a<b\leq 1$, the set $\pi_i^*\mathscr{H}(a|b)$ is $\mathscr{B}$–measurable.

Let a' and b' be, respectively, the infimum and the maximum of $\pi_i(M)\cap[a,b]$. Pick sequences of elements $a'\leq a_{n+1}<a_n<b_n<b_{n+1}\leq b'$ so that $a_n,b_n\in\pi_i(M)$ and $a_n\to a'$, $b_n\to b'$. These sequences can be empty if $\pi_i(M)\cap[a,b]=\varnothing$, or consist of single elements if $a',b'\in\pi_i(M)$. Then

$$\pi_i^*\mathscr{H}(a|b)=\bigcup\pi_i^*\mathscr{H}(a_n|b_n)\cup\{\pi_i^{-1}((a,1])\}\cup\{\pi_i^{-1}([b,1])\}.$$

Observing that singletons are $\mathscr{B}$–measurable, it suffices to show that, for every $x,y\in M$, the set $\pi_i^*\mathscr{H}(\pi_i(x)|\pi_i(y))$ is $\mathscr{B}$–measurable.

This means that it actually suffices to prove that the sets $\pi_i^*\mathscr{H}(0|1)$ are $\mathscr{B}$–measurable. We will achieve this by showing that each set $\mathscr{H}(M)\setminus\pi_i^*\mathscr{H}(0|1)$ is a countable union of halfspace-intervals.

Note that $\mathfrak{h}\in\mathscr{H}(M)$ lies in $\pi_i^*\mathscr{H}([0,1])$ if and only if the projections $\pi_i(\mathfrak{h})$ and $\pi_i(\mathfrak{h}^*)$ are disjoint. Thus, $\mathfrak{h}$ lies in $\mathscr{H}(M)\setminus\pi_i^*\mathscr{H}(0|1)$ if and only if there exist $x,y\in M$ such that $\mathfrak{h}\in\mathscr{H}(x|y)$ and $\pi_i(x)\geq\pi_i(y)$. This gives a subset $P\subseteq M\times M$ with $\mathscr{H}(M)\setminus\pi_i^*\mathscr{H}(0|1)=\mathcal{H}_P$.

In view of Lemma A.1 and Remark 2.2(1), there exists a countable subset $\Delta \subseteq P$ with $\mathcal{H}_\Delta = \mathcal{H}_P$. This concludes the proof. $\square$

The following would be an immediate consequence of Dilworth's lemma, were it not for the measurability requirement.

Corollary A.3 *Let X be a median space of finite rank r. For all $x, y \in X$, there exists a $\mathscr{B}$–measurable partition $\mathscr{H}(x|y) = \mathcal{H}_1 \sqcup \cdots \sqcup \mathcal{H}_r$ such that no two halfspaces in the same $\mathcal{H}_i$ are transverse.*

Proof Taking the metric completion of X and applying [50, Proposition 2.19], we obtain an isometric embedding $\iota\colon I(x,y) \hookrightarrow \mathbb{R}^r$. The image of ι is contained in a product $J_1 \times \cdots \times J_r$ of compact intervals $J_i \subseteq \mathbb{R}$, which is isomorphic to the median algebra $[0,1]^r$. Let $\pi_i : M \to J_i$ be the composition of ι with the projection to J_i, and set $\mathcal{H}_i' := \mathscr{H}(x|y) \cap \pi_i^*(\mathscr{H}(J_i))$. We have $\mathscr{H}(x|y) = \mathcal{H}_1' \cup \cdots \cup \mathcal{H}_r'$, no two halfspaces in the same $\mathcal{H}_i'$ are transverse, and each $\mathcal{H}_i'$ is $\mathscr{B}$–measurable by Lemma A.2. We conclude by taking $\mathcal{H}_i := \mathcal{H}_i' \setminus (\mathcal{H}_1' \cup \cdots \cup \mathcal{H}_{i-1}')$. $\square$

References

[1] **H-J Bandelt**, **J Hedlíková**, *Median algebras*, Discrete Math. 45 (1983) 1–30 MR Zbl

[2] **B Beeker**, **N Lazarovich**, *Resolutions of* CAT(0) *cube complexes and accessibility properties*, Algebr. Geom. Topol. 16 (2016) 2045–2065 MR Zbl

[3] **B Beeker**, **N Lazarovich**, *Stallings' folds for cube complexes*, Israel J. Math. 227 (2018) 331–363 MR Zbl

[4] **J Behrstock**, **R Charney**, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. 352 (2012) 339–356 MR Zbl

[5] **J Behrstock**, **M F Hagen**, **A Sisto**, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. 21 (2017) 1731–1804 MR Zbl

[6] **J Behrstock**, **M Hagen**, **A Sisto**, *Hierarchically hyperbolic spaces, II: Combination theorems and the distance formula*, Pacific J. Math. 299 (2019) 257–338 MR Zbl

[7] **M Bestvina**, *Degenerations of the hyperbolic space*, Duke Math. J. 56 (1988) 143–161 MR Zbl

[8] **M Bestvina**, **N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129 (1997) 445–470 MR Zbl

[9] **M Bestvina**, **K Bromberg**, **K Fujiwara**, *Proper actions on finite products of quasi-trees*, Ann. H. Lebesgue 4 (2021) 685–709 MR Zbl

[10] **M Bestvina**, **M Feighn**, *Stable actions of groups on real trees*, Invent. Math. 121 (1995) 287–321 MR Zbl

[11] **M Bestvina**, **M Feighn**, **M Handel**, *The Tits alternative for* $\mathrm{Out}(F_n)$*, I: Dynamics of exponentially-growing automorphisms*, Ann. of Math. 151 (2000) 517–623 MR Zbl

[12] **M Bestvina**, **M Handel**, *Train tracks and automorphisms of free groups*, Ann. of Math. 135 (1992) 1–51 MR Zbl

[13] **J Beyrer**, **E Fioravanti**, *Cross-ratios on* CAT(0) *cube complexes and marked length-spectrum rigidity*, J. Lond. Math. Soc. 104 (2021) 1973–2015 MR Zbl

[14] **J Beyrer**, **E Fioravanti**, *Cross ratios and cubulations of hyperbolic groups*, Math. Ann. 384 (2022) 1547–1592 MR Zbl

[15] **B H Bowditch**, *Coarse median spaces and groups*, Pacific J. Math. 261 (2013) 53–93 MR Zbl

[16] **B H Bowditch**, *Embedding median algebras in products of trees*, Geom. Dedicata 170 (2014) 157–176 MR Zbl

[17] **B H Bowditch**, *Some properties of median metric spaces*, Groups Geom. Dyn. 10 (2016) 279–317 MR Zbl

[18] **B H Bowditch**, *Convex hulls in coarse median spaces*, preprint (2018) Available at `https://homepages.warwick.ac.uk/~masgak/papers/hulls-cms.pdf`

[19] **B H Bowditch**, *Large-scale rigidity properties of the mapping class groups*, Pacific J. Math. 293 (2018) 1–73 MR Zbl

[20] **C Bregman**, **R Charney**, **K Vogtmann**, *Outer space for RAAGs*, Duke Math. J. 172 (2023) 1033–1108 MR Zbl

[21] **M R Bridson**, *The rhombic dodecahedron and semisimple actions of* $\mathrm{Aut}(F_n)$ *on* $\mathrm{CAT}(0)$ *spaces*, Fund. Math. 214 (2011) 13–25 MR Zbl

[22] **K-U Bux**, **C Gonzalez**, *The Bestvina–Brady construction revisited: geometric computation of* Σ*–invariants for right-angled Artin groups*, J. London Math. Soc. 60 (1999) 793–801 MR Zbl

[23] **P-E Caprace**, **M Sageev**, *Rank rigidity for* $\mathrm{CAT}(0)$ *cube complexes*, Geom. Funct. Anal. 21 (2011) 851–891 MR Zbl

[24] **R Charney**, **J Crisp**, **K Vogtmann**, *Automorphisms of* 2*–dimensional right-angled Artin groups*, Geom. Topol. 11 (2007) 2227–2264 MR Zbl

[25] **R Charney**, **N Stambaugh**, **K Vogtmann**, *Outer space for untwisted automorphisms of right-angled Artin groups*, Geom. Topol. 21 (2017) 1131–1178 MR Zbl

[26] **R Charney**, **H Sultan**, *Contracting boundaries of* $\mathrm{CAT}(0)$ *spaces*, J. Topol. 8 (2015) 93–117 MR Zbl

[27] **R Charney**, **K Vogtmann**, *Finiteness properties of automorphism groups of right-angled Artin groups*, Bull. Lond. Math. Soc. 41 (2009) 94–102 MR Zbl

[28] **R Charney**, **K Vogtmann**, *Subgroups and quotients of automorphism groups of RAAGs*, from "Low-dimensional and symplectic topology" (M Usher, editor), Proc. Sympos. Pure Math. 82, Amer. Math. Soc., Providence, RI (2011) 9–27 MR Zbl

[29] **I Chatterji**, **C Druţu**, **F Haglund**, *Kazhdan and Haagerup properties from the median viewpoint*, Adv. Math. 225 (2010) 882–921 MR Zbl

[30] **I Chatterji**, **T Fernós**, **A Iozzi**, *The median class and superrigidity of actions on* $\mathrm{CAT}(0)$ *cube complexes*, J. Topol. 9 (2016) 349–400 MR Zbl

[31] **V Chepoi**, *Graphs of some* $\mathrm{CAT}(0)$ *complexes*, Adv. in Appl. Math. 24 (2000) 125–179 MR Zbl

[32] **M M Cohen**, **M Lustig**, *On the dynamics and the fixed subgroup of a free group automorphism*, Invent. Math. 96 (1989) 613–638 MR Zbl

[33] **D Cooper**, *Automorphisms of free groups have finitely generated fixed point sets*, J. Algebra 111 (1987) 453–456 MR Zbl

[34] **M Cordes**, **M G Durham**, *Boundary convex cocompactness and stability of subgroups of finitely generated groups*, Int. Math. Res. Not. 2019 (2019) 1699–1724 MR Zbl

[35] **L J Corredor**, **M A Gutierrez**, *A generating set for the automorphism group of a graph product of abelian groups*, Internat. J. Algebra Comput. 22 (2012) art. id. 1250003 MR Zbl

[36] **M Culler**, *Finite groups of outer automorphisms of a free group*, from "Contributions to group theory" (K I Appel, J G Ratcliffe, P E Schupp, editors), Contemp. Math. 33, Amer. Math. Soc., Providence, RI (1984) 197–207 MR Zbl

[37] **M Culler**, **J W Morgan**, *Group actions on $\mathbb{R}$–trees*, Proc. London Math. Soc. 55 (1987) 571–604 MR Zbl

[38] **P Dani**, **I Levcovitz**, *Subgroups of right-angled Coxeter groups via Stallings-like techniques*, J. Comb. Algebra 5 (2021) 237–295 MR Zbl

[39] **M W Davis**, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series 32, Princeton Univ. Press (2008) MR Zbl

[40] **M W Davis**, **T Januszkiewicz**, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra 153 (2000) 229–235 MR Zbl

[41] **M B Day**, *Peak reduction and finite presentations for automorphism groups of right-angled Artin groups*, Geom. Topol. 13 (2009) 817–855 MR Zbl

[42] **M B Day**, **A W Sale**, **R D Wade**, *Calculating the virtual cohomological dimension of the automorphism group of a RAAG*, Bull. Lond. Math. Soc. 53 (2021) 259–273 MR Zbl

[43] **M B Day**, **R D Wade**, *Relative automorphism groups of right-angled Artin groups*, J. Topol. 12 (2019) 759–798 MR Zbl

[44] **C Druţu**, **M Kapovich**, *Geometric group theory*, American Mathematical Society Colloquium Publications 63, Amer. Math. Soc., Providence, RI (2018) MR Zbl

[45] **J Dugundji**, *An extension of Tietze's theorem*, Pacific J. Math. 1 (1951) 353–367 MR Zbl

[46] **J L Dyer**, **G P Scott**, *Periodic automorphisms of free groups*, Comm. Algebra 3 (1975) 195–201 MR Zbl

[47] **T Fernós**, *The Furstenberg–Poisson boundary and* CAT(0) *cube complexes*, Ergodic Theory Dynam. Systems 38 (2018) 2180–2223 MR Zbl

[48] **E Fioravanti**, *The Tits alternative for finite rank median spaces*, Enseign. Math. 64 (2018) 89–126 MR Zbl

[49] **E Fioravanti**, *Superrigidity of actions on finite rank median spaces*, Adv. Math. 352 (2019) 1206–1252 MR Zbl

[50] **E Fioravanti**, *Roller boundaries for median spaces and algebras*, Algebr. Geom. Topol. 20 (2020) 1325–1370 MR Zbl

[51] **E Fioravanti**, *Convex cores for actions on finite-rank median algebras*, preprint (2021) arXiv 2101.03060v1

[52] **E Fioravanti**, *On automorphisms and splittings of special groups*, Compos. Math. 159 (2023) 232–305 MR Zbl

[53] **D Gaboriau**, **A Jaeger**, **G Levitt**, **M Lustig**, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. 93 (1998) 425–452 MR Zbl

[54] **D Gaboriau**, **G Levitt**, *The rank of actions on $\mathbb{R}$–trees*, Ann. Sci. École Norm. Sup. 28 (1995) 549–570 MR Zbl

[55] **D Gaboriau**, **G Levitt**, **M Lustig**, *A dendrological proof of the Scott conjecture for automorphisms of free groups*, Proc. Edinburgh Math. Soc. 41 (1998) 325–332 MR Zbl

[56] **S M Gersten**, *On fixed points of automorphisms of finitely generated free groups*, Bull. Amer. Math. Soc. 8 (1983) 451–454 MR Zbl Addendum in 49 (1984) 340–342

[57] **S M Gersten**, *Fixed points of automorphisms of free groups*, Adv. in Math. 64 (1987) 51–85 MR Zbl

[58] **R Z Goldstein**, **E C Turner**, *Automorphisms of free groups and their fixed points*, Invent. Math. 78 (1984) 1–12 MR Zbl

[59] **D Groves**, **M Hull**, *Abelian splittings of right-angled Artin groups*, from "Hyperbolic geometry and geometric group theory" (K Fujiwara, S Kojima, K Ohshika, editors), Adv. Stud. Pure Math. 73, Math. Soc. Japan, Tokyo (2017) 159–165 MR Zbl

[60] **V Guirardel**, *Cœur et nombre d'intersection pour les actions de groupes sur les arbres*, Ann. Sci. École Norm. Sup. 38 (2005) 847–888 MR Zbl

[61] **V Guirardel**, **A Sale**, *Vastness properties of automorphism groups of RAAGs*, J. Topol. 11 (2018) 30–64 MR Zbl

[62] **M F Hagen**, *Cocompactly cubulated crystallographic groups*, J. Lond. Math. Soc. 90 (2014) 140–166 MR Zbl

[63] **M F Hagen**, **T Susse**, *On hierarchical hyperbolicity of cubical groups*, Israel J. Math. 236 (2020) 45–89 MR Zbl

[64] **M F Hagen**, **N W M Touikan**, *Panel collapse and its applications*, Groups Geom. Dyn. 13 (2019) 1285–1334 MR Zbl

[65] **M F Hagen**, **H Wilton**, *Guirardel cores for cube complexes*, in preparation

[66] **F Haglund**, *Finite index subgroups of graph products*, Geom. Dedicata 135 (2008) 167–209 MR Zbl

[67] **F Haglund**, *Isometries of CAT(0) cube complexes are semi-simple*, Ann. Math. Qué. 47 (2023) 249–261 MR

[68] **F Haglund**, **D T Wise**, *Special cube complexes*, Geom. Funct. Anal. 17 (2008) 1551–1620 MR Zbl

[69] **S Hensel**, **D Kielak**, *Nielsen realisation for untwisted automorphisms of right-angled Artin groups*, Proc. Lond. Math. Soc. 117 (2018) 901–950 MR Zbl

[70] **S-t Hu**, *Theory of retracts*, Wayne State Univ. Press, Detroit, MI (1965) MR Zbl

[71] **J Huang**, **B Kleiner**, *Groups quasi-isometric to right-angled Artin groups*, Duke Math. J. 167 (2018) 537–602 MR Zbl

[72] **W Jaco**, **P B Shalen**, *Surface homeomorphisms and periodicity*, Topology 16 (1977) 347–367 MR Zbl

[73] **T Koberda**, **J Mangahas**, **S J Taylor**, *The geometry of purely loxodromic subgroups of right-angled Artin groups*, Trans. Amer. Math. Soc. 369 (2017) 8179–8208 MR Zbl

[74] **M Kolibiar**, **T Marcisová**, *On a question of J Hashimoto*, Mat. Časopis Sloven. Akad. Vied 24 (1974) 179–185 MR Zbl

[75] **M R Laurence**, *A generating set for the automorphism group of a graph group*, J. London Math. Soc. 52 (1995) 318–334 MR Zbl

[76] **S Lefschetz**, *On locally connected and related sets*, Ann. of Math. 35 (1934) 118–129 MR Zbl

[77] **S Lefschetz**, *On locally connected and related sets (second paper)*, Duke Math. J. 2 (1936) 435–442 MR Zbl

[78] **G Levitt**, *Automorphisms of hyperbolic groups and graphs of groups*, Geom. Dedicata 114 (2005) 49–70 MR Zbl

[79] **J Meier**, **H Meinert**, **L VanWyk**, *Higher generation subgroup sets and the Σ–invariants of graph groups*, Comment. Math. Helv. 73 (1998) 22–44 MR Zbl

[80] **J Meier**, **L VanWyk**, *The Bieri–Neumann–Strebel invariants for graph groups*, Proc. London Math. Soc. 71 (1995) 263–280 MR Zbl

[81] **A Minasyan**, **D Osin**, *Fixed subgroups of automorphisms of relatively hyperbolic groups*, Q. J. Math. 63 (2012) 695–712 MR Zbl

[82] **W D Neumann**, *The fixed group of an automorphism of a word hyperbolic group is rational*, Invent. Math. 110 (1992) 147–150 MR Zbl

[83] **G A Niblo**, **N Wright**, **J Zhang**, *A four point characterisation for coarse median spaces*, Groups Geom. Dyn. 13 (2019) 939–980 MR Zbl

[84] **G A Niblo**, **N Wright**, **J Zhang**, *Coarse median algebras: the intrinsic geometry of coarse median spaces and their intervals*, Selecta Math. 27 (2021) art. id. 20 MR Zbl

[85] **F Paulin**, *Topologie de Gromov équivariante, structures hyperboliques et arbres réels*, Invent. Math. 94 (1988) 53–80 MR Zbl

[86] **F Paulin**, *Points fixes des automorphismes de groupe hyperbolique*, Ann. Inst. Fourier (Grenoble) 39 (1989) 651–662 MR Zbl

[87] **F Paulin**, *Outer automorphisms of hyperbolic groups and small actions on $\mathbb{R}$–trees*, from "Arboreal group theory" (R C Alperin, editor), Math. Sci. Res. Inst. Publ. 19, Springer (1991) 331–343 MR Zbl

[88] **F Paulin**, *Sur les automorphismes extérieurs des groupes hyperboliques*, Ann. Sci. École Norm. Sup. 30 (1997) 147–167 MR Zbl

[89] **M Roller**, *Poc sets, median algebras and group actions*, preprint (2016) arXiv 1607.07747

[90] **M Sageev**, *CAT(0) cube complexes and groups*, from "Geometric group theory" (M Bestvina, M Sageev, K Vogtmann, editors), IAS/Park City Math. Ser. 21, Amer. Math. Soc., Providence, RI (2014) 7–54 MR Zbl

[91] **M Sageev**, **D T Wise**, *Cores for quasiconvex actions*, Proc. Amer. Math. Soc. 143 (2015) 2731–2741 MR Zbl

[92] **A Sale**, **T Susse**, *Outer automorphism groups of right-angled Coxeter groups are either large or virtually abelian*, Trans. Amer. Math. Soc. 372 (2019) 7785–7803 MR Zbl

[93] **Z Sela**, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups, II*, Geom. Funct. Anal. 7 (1997) 561–593 MR Zbl

[94] **H Servatius**, *Automorphisms of graph groups*, J. Algebra 126 (1989) 34–60 MR Zbl

[95] **J R Stallings**, *Topology of finite graphs*, Invent. Math. 71 (1983) 551–565 MR Zbl

[96] **J R Stallings**, *Foldings of G–trees*, from "Arboreal group theory" (R C Alperin, editor), Math. Sci. Res. Inst. Publ. 19, Springer (1991) 355–368 MR Zbl

[97] **W P Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. 19 (1988) 417–431 MR Zbl

[98] **H C Tran**, *Geometric embedding properties of Bestvina–Brady subgroups*, Algebr. Geom. Topol. 17 (2017) 2499–2510 MR Zbl

[99] **M L J van de Vel**, *Theory of convex structures*, North-Holland Mathematical Library 50, North-Holland, Amsterdam (1993) MR Zbl

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

[101] **R Zeidler**, *Coarse median structures and homomorphisms from Kazhdan groups*, Geom. Dedicata 180 (2016) 49–68 MR Zbl

Universität Bonn
Bonn, Germany
elia.fioravanti@kit.edu

Proposed: Mladen Bestvina
Seconded: David Fisher, Urs Lang

Received: 25 February 2021
Revised: 10 March 2022

GEOMETRY & TOPOLOGY

msp.org/gt

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2024 is US $805/year for the electronic version, and $1135/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Geometry & Topology

Volume 28 Issue 1 (pages 1–496) 2024

Homological invariants of codimension 2 contact submanifolds 1

Laurent Côté and François-Simon Fauteux-Chapleau

The desingularization of the theta divisor of a cubic threefold as a moduli space 127

Arend Bayer, Sjoerd Viktor Beentjes, Soheyla Feyzbakhsh, Georg Hein, Diletta Martinelli, Fatemeh Rezaee and Benjamin Schmidt

Coarse-median preserving automorphisms 161

Elia Fioravanti

Classification results for expanding and shrinking gradient Kähler–Ricci solitons 267

Ronan J Conlon, Alix Deruelle and Song Sun

Embedding calculus and smooth structures 353

Ben Knudsen and Alexander Kupers

Stable maps to Looijenga pairs 393

Pierrick Bousseau, Andrea Brini and Michel van Garrel