Groups Geom. Dyn. 10 (2016), 649–707
DOI 10.4171/GGD/360

Groups, Geometry, and Dynamics

Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries

Jason Behrstock[1] and Mark F. Hagen[2]

Abstract. Let G be a group acting geometrically on a CAT(0) cube complex $\mathbf{X}$. We prove first that G is hyperbolic relative to the collection $\mathbb{P}$ of subgroups if and only if the simplicial boundary $\partial_\triangle \mathbf{X}$ is the disjoint union of a nonempty discrete set, together with a pairwise-disjoint collection of subcomplexes corresponding, in the appropriate sense, to elements of $\mathbb{P}$. As a special case of this result is a new proof, in the cubical case, of a Theorem of Hruska and Kleiner regarding Tits boundaries of relatively hyperbolic CAT(0) spaces. Second, we relate the existence of cut-points in asymptotic cones of a cube complex $\mathbf{X}$ to boundedness of the 1-skeleton of $\partial_\triangle \mathbf{X}$. We deduce characterizations of thickness and strong algebraic thickness of a group G acting properly and cocompactly on the CAT(0) cube complex $\mathbf{X}$ in terms of the structure of, and nature of the G-action on, $\partial_\triangle \mathbf{X}$. Finally, we construct, for each $n \geq 0, k \geq 2$, infinitely many quasi-isometry types of group G such that G is strongly algebraically thick of order n, has polynomial divergence of order $n + 1$, and acts properly and cocompactly on a k-dimensional CAT(0) cube complex.

Mathematics Subject Classification (2010). 20F65, 20F69, 20F67.

Keywords. Thick metric space, thick group, cubulated group, cube complex, relatively hyperbolic group, boundary, simplicial boundary.

[1] Supported as an Alfred P. Sloan Fellow and by NSF grant #DMS-1006219.
[2] Partially supported by the NSF under Grant Number NSF 1045119.

Contents

Introduction

In this paper, we study the mutually exclusive properties of relative hyperbolicity and thickness of groups, in the context of groups acting properly and cocompactly on CAT(0) cube complexes. Since being first introduced as an interesting family of CAT(0) spaces by Gromov in [21], the class of CAT(0) cube complexes has been recognized as being sufficiently rich to warrant a theory encompassing more than just CAT(0) geometry. Applications of this theory range from their use by Charney and Davis to resolve the $K(\pi, 1)$ problem for hyperplane complements [11] to the recent resolution of the virtual fibering conjecture [44] and virtual Haken conjectures [1], among many others. In the setting of cubical complexes, we study the properties of relative hyperbolicity and thickness, the latter of which is a powerful obstruction to the former. We show that despite these two properties being antithetical, surprisingly, they admit similar characterizations using the boundary of the space. This is the only setting in which such a close relationship between these two properties is known.

A CAT(0) cube complex has a highly organized combinatorial structure that yields an associated space, the *simplicial boundary*, which encodes much of the large-scale structure of the cube complex. Our results show that, to a large extent, both relative hyperbolicity and thickness of a group G acting geometrically on a cube complex $\mathbf{X}$ correspond to simple properties of the simplicial boundary of $\mathbf{X}$ and the natural action of G on the simplicial boundary of $\mathbf{X}$.

The definition of a *thick metric space* and the attendant quasi-isometry invariant, *order of thickness*, were introduced by Behrstock, Druțu, and Mosher [4] as an obstruction to relative hyperbolicity and as a tool to study geometric commonalities between several classes of groups, notably mapping class groups of surfaces, outer automorphism groups of finitely generated free groups, and $\mathrm{SL}_n(\mathbb{Z})$. We review thick metric spaces in detail in Section 1.1.

The order of thickness of M, defined below, is intimately related to the divergence function of M. The relevant notion of divergence of a metric space originates in work by Gromov [22] and Gersten [19, 20], and, roughly speaking, estimates how far one must travel in M from a point a to a point b, avoiding a specified ball centered at a third point c. Divergence can be studied via asymptotic cones of M. In particular, Druţu, Mozes, and Sapir proved that, if M is quasi-isometric to a finitely generated group, then M has linear divergence if and only if it is wide [16]. Furthermore, the first author and Druţu proved in [3] that the divergence of M is bounded above by a polynomial of order $n + 1$ when M is a metric space that is strongly thick of order n.

The order of thickness of a metric space M is defined inductively. First, M is [strongly] thick of order 0 if M is *unconstricted* [*wide*], which means that some [any] asymptotic cone of M has no cut-point. M is [*strongly*] *thick of order at most* $n \geq 1$ if there is a collection of quasiconvex *thickly connecting* subspaces $\{S_i\}$ that coarsely cover M, with the additional property that each S_i is [strongly] thick of order at most $(n - 1)$. Being *thickly connected* means that for any $p, q \in M$, there is a sequence $S_{i_1}, \dots, S_{i_k}$ with $p \in S_{i_1}, q \in S_{i_k}$ and $\operatorname{diam}(S_{i_j} \cap S_{i_{j+1}}) = \infty$ for all j. An important variation on this notion occurs when M is quasi-isometric to a finitely generated group, and the sets S_i are cosets of a finite collection of quasi-convex subgroups, each of which is (*strongly*) *algebraically thick of order* $n - 1$. In this case, M is (*strongly*) *algebraically thick of order* n. *Algebraically thick of order* 0 means unconstricted, and *strongly algebraically thick of order* 0 means wide.

CAT(0) cube complexes are a generalization of trees in two fundamental ways. First, the class of graphs that are 1-skeleta of CAT(0) cube complexes is precisely the class of median graphs, of which trees are a special case, as was established independently by Chepoi [12] and by Roller [41]. Second, CAT(0) cube complexes contain large collections of convex subspaces with exactly two complementary components. These convex subspaces are the *hyperplanes*; in the 1-dimensional case, hyperplanes are midpoints of edges. A detailed discussion of basic properties of CAT(0) cube complexes occurs in Section 1.3.

Just as cube complexes generalize trees, the theory of groups acting on trees generalizes, yielding a theory of groups acting on cube complexes. See Sageev [42] and later developments in work of Chatterji and Niblo [14], Haglund and Paulin [28], Hruska and Wise [29], and Nica [36]. The class of groups known to be *cubulated* – i.e., to admit a metrically proper action by isometries on a CAT(0) cube complex – is ever-growing and contains many Coxeter groups [37], right-angled Artin groups [10], Artin groups of finite type [11], groups satisfying sufficiently strong small-cancellation conditions [45], random groups at sufficiently low density in Gromov's model [38], appropriately-chosen subgroups of fundamental groups of nonpositively-curved graph manifolds [33, 40], certain graphs of cubulated groups [30], and many others.

The connection between relative hyperbolicity and thickness for cube complexes results from the fact that these two properties of a group acting geometrically on a CAT(0) cube complex can both be detected by examining the action of the group on the simplicial boundary of the cube complex. The *simplicial boundary* $\partial_{\triangle}\mathbf{X}$ was introduced by Hagen [26] as a combinatorial analogue of the Tits boundary of $\mathbf{X}$. The simplicial boundary is a simplicial complex that is an invariant of the median graph $\mathbf{X}^{(1)}$, obtained by taking the 1-skeleton of $\mathbf{X}$, or, equivalently, of the hyperplanes and how they interact. In the event of a proper, cocompact action on $\mathbf{X}$, the two boundaries are quasi-isometric in a strong sense discussed in Section 6 [26, Section 3.5]. Simplices of $\partial_{\triangle}\mathbf{X}$ are represented by set of hyperplanes in $\mathbf{X}$ modeled on the set of hyperplanes separating some basepoint from a collection of points at infinity, and since an isometric action of a group G on $\mathbf{X}$ preserves the set of hyperplanes, such an action induces an action of G on $\partial_{\triangle}\mathbf{X}$ by simplicial automorphisms. A more discussion of the simplicial boundary is provided in Section 2.

Relative hyperbolicity. Relatively hyperbolic cubulated groups form a rich family. For instance, by recent work of Wise [44], if M is a finite-volume cusped hyperbolic 3-manifold with a geometrically finite incompressible surface, then M has a finite cover $\widehat{M}$ such that $\pi_1\widehat{M}$ is the fundamental group of a compact nonpositively-curved cube complex. The simplicial boundary of the universal cover of such a cube complex is described by Theorem 3.1 below.

Given a group G acting properly and cocompactly on a CAT(0) space Y, it is natural to search for characterizations of hyperbolicity of G relative to a collection of subgroups. A result of Hruska and Kleiner achieves this in the special case in which each peripheral subgroup is free abelian; they prove that G is hyperbolic relative to a collection of free abelian subgroups if and only if the Tits boundary $\partial_T Y$ decomposes as the union of an infinite set of isolated points and an infinite collection of spheres, which are boundaries of flats in Y corresponding to the peripheral subgroups [27]. The following two results generalize Hruska's and Kleiner's result in the cubical setting, by removing any assumptions on the peripheral subgroups. Just as Hruska's and Kleiner's result shows that the property of being hyperbolic relative to free abelian subgroups corresponds to the existence of a simple geometric description of the Tits boundary, the following theorems relate relative hyperbolicity of cubulated groups to the existence of a simple decomposition of the simplicial (and therefore Tits) boundary into pieces with simpler structure.

Theorem 3.1 *Let $(G, \mathbb{P})$ be a relatively hyperbolic structure and let G act properly and cocompactly on the* CAT(0) *cube complex $\mathbf{X}$. Then $\partial_{\triangle}\mathbf{X}$ consists of an infinite collection of isolated 0-simplices, together with a pairwise-disjoint collection $\{g\partial_{\triangle}\mathbf{Y}_P : P \in \mathbb{P}, g \in G\}$ of subcomplexes, with each $\mathbf{Y}_P$ the convex hull of a P-orbit in $\mathbf{X}$.*

When each $P \in \mathbb{P}$ is isomorphic to $\mathbb{Z}^{n_P}$ for some $n_P \geq 2$, the complex $\partial_{\triangle}\mathbf{Y}_P$ is isomorphic to the $(n-1)$-dimensional hyperoctahedron, and thus homeomorphic to $\mathbb{S}^{n-1}$; see Corollary 3.5. Conversely, the following shows that relative hyperbolicity can be identified by examining the action on the simplicial boundary:

Theorem 3.7. *Let G act properly and cocompactly on the* CAT(0) *cube complex* $\mathbf{X}$. *Let $\{\mathbf{S}_i\}_i$ be a G-invariant collection of pairwise-disjoint subcomplexes of $\partial_{\triangle}\mathbf{X}$, such that $\partial_{\triangle}\mathbf{X}$ consists of $\bigsqcup_i \mathbf{S}_i$ together with a G-invariant collection of isolated* 0*-simplices. Suppose each* $\mathrm{Stab}_G(\mathbf{S}_i)$ *acts with a quasiconvex orbit on* $\mathbf{X}$ *and has infinite index in G, and that $\mathbf{S}_i$ contains all limit simplices for the action of* $\mathrm{Stab}_G(\mathbf{S}_i)$. *Then G is hyperbolic relative to a collection of subgroups, each of which is commensurable with some* $\mathrm{Stab}_G(\mathbf{S}_i)$.

Corollary 6.1 provides an analogue of Theorem 3.1 and Theorem 3.7 in terms of the Tits boundary. In particular, this provides a characterization of relative hyperbolicity of a group acting geometrically on a cube complex $\mathbf{X}$ in terms of the action of G on $\partial_T\mathbf{X}$.

Thickness. Important motivating examples of cocompactly cubulated groups are the right-angled Artin groups, see Charney and Davis [10]. In contrast to the fundamental groups of finite volume hyperbolic manifolds mentioned above, right-angled Artin groups are cocompactly cubulated groups which are not relatively hyperbolic; in fact, these groups are thick [4]. Behrstock and Charney showed that one-ended right-angled Artin groups that are thick of order 0 (and thus have linear divergence) are precisely those whose presentation graphs decompose as nontrivial joins [2]. Motivated by this result, Hagen generalized this to show that a cocompactly cubulated groups has linear divergence if and only if it acts geometrically on a CAT(0) cube complex whose simplicial boundary decomposes as a nontrivial simplicial join [26]. Otherwise, the simplicial boundary is disconnected and contains many isolated 0-simplices corresponding to endpoints of axes of rank-one isometries [15, Corollary B]. Accordingly, as Theorem 4.3 we record the fact that if a CAT(0) cube complex $\mathbf{X}$ admits a geometric action by a group G, then $\mathbf{X}$ and G are each thick of order 0 exactly when the simplicial boundary of $\mathbf{X}$ is connected.

For proper, cocompact CAT(0) cube complexes, the property of being thick of order 1 admits a succinct characterization in terms of the simplicial boundary. We summarize this by:

Theorem 5.13 (characterization of thickness). *Let G act properly and cocompactly by isometries on the fully visible* CAT(0) *cube complex* $\mathbf{X}$. *If G is algebraically thick of order* 1 *relative to a collection of quasiconvex wide subgroups,* *then* $\partial_{\triangle}\mathbf{X}$ *is disconnected and contains a positive-dimensional,* G*-invariant connected component.*

Conversely, if $\partial_{\triangle}\mathbf{X}$ *is disconnected, and has a positive-dimensional* G*-invariant component, then* $\mathbf{X}$ *is thick of order* 1 *relative to a set of wide, convex subcomplexes, and, in particular,* G *is thick of order* 1.

Moreover, we obtain the following complete description of the boundary of a cube complex admitting a geometric action by a group that is strongly algebraically thick of order 1. This description of algebraic thickness closely parallels that of relative hyperbolicity provided by Theorem 3.7.

Theorem 5.13 (description of the boundary). *Let* G *act properly and cocompactly on the* CAT(0) *cube complex* $\mathbf{X}$*. Then* G *is strongly algebraically thick of order* 1 *if and only if* $\partial_{\triangle}\mathbf{X}$ *is disconnected and has a positive-dimensional,* G*-invariant connected subcomplex* $\mathfrak{C} = \bigcup_{A\in\mathcal{A},g\in G} gA$*, where* $\mathcal{A}$ *is a finite collection of bounded subcomplexes such that*

(1) *each* $\mathrm{Stab}(A)$ *acts on* $\mathbf{X}$ *with a quasiconvex orbit*;

(2) *for each* $A\in\mathcal{A}$*,* $f^{-1}(A)$ *belongs to the limit set of* $\mathrm{Stab}(A)$;

(3) $f^{-1}(\mathfrak{C})$ *is contained in the limit set of* $\langle\{\mathrm{Stab}(A)\colon A\in\mathcal{A}\}\rangle$.

Remark. Here, $f\colon\partial_{\infty}\mathbf{X}\to\partial_{\triangle}\mathbf{X}$ is a surjection from the visual boundary to the simplicial boundary which sends each asymptotic class of CAT(0) geodesic rays to a point in the simplex of $\partial_{\triangle}\mathbf{X}$ represented by the set of hyperplanes crossing some ray in the given asymptotic class; see Section 5. *Full visibility* of $\mathbf{X}$ is a technical condition on $\partial_{\triangle}\mathbf{X}$ saying roughly that each infinite family of nested halfspaces in $\mathbf{X}$ determines a combinatorial geodesic ray.

Condition (3) is used to verify that $\langle\{\mathrm{Stab}(A) : A\in\mathcal{A}\}\rangle$ has finite index in G, as required by the definition of algebraic thickness. In contrast to the situation for many other examples of thick groups (see [4]), in the present case there does not appear to be natural choice of generators of these subgroups from which one can easily see that the collection of them generate a finite index subgroup of G.

From Theorem 5.13, an application of Corollary 4.17 of [3] immediately yields:

Corollary. Let G act properly and cocompactly on the CAT(0) cube complex $\mathbf{X}$, and suppose that $\partial_{\triangle}\mathbf{X}$ has a G-invariant connected proper subcomplex satisfying (1)–(3) of Theorem 5.13. Then G has quadratic divergence function.

Theorem 5.13 and the above corollary are, respectively, equivalent to very similar statements about the G-action on the Tits boundary of $\mathbf{X}$; see Corollary 6.2 below.

A key ingredient in the proof of Theorem 5.4 is Theorem 4.1, which relates the existence of cut-points in some asymptotic cone of a cube complex (not necessarily cocompact) to boundedness of the 1-skeleton of the simplicial boundary. The proof of this theorem occupies much of Section 4, and relies in part on the relationship between divergence and wideness discussed in [16] and the relationship between divergence and the simplicial boundary discussed in [26].

We show that there are many cocompactly cubulated groups that are thick of any given order. Indeed we show this is already true for the class of groups that act geometrically on CAT(0) square complexes.

Theorem 7.3 (abundance of cubulated groups that are thick of order n). *For all $n \geq 0$, there are infinitely many quasi-isometry types of cocompactly cubulated groups that are algebraically thick (and hence metrically thick) of order n and have polynomial divergence of order precisely $n+1$.*

Furthermore, for any $k \geq 2$, there are infinitely many quasi-isometry types of such groups with the additional condition that the groups act properly and cocompactly on k-dimensional CAT(0) *cube complexes.*

The nature of the construction and the latter part of the proof are modeled on the construction by Behrstock and Druțu [3] of CAT(0) groups which are thick of order n and with polynomial divergence of degree $n+1$. CAT(0) groups of arbitrary order of polynomial growth were also constructed recently by Macura [34], who considered iterated HNN extensions of $\mathbb{Z}^2$. Dani and Thomas recently posted a preprint in which they show that for every integer there exists a Coxeter group whose divergence is polynomial of that degree [18] – it would be interesting to know if those Coxeter groups are each thick and to compute their simplicial boundaries.

Acknowledgements. The authors thank Dani Wise for sharing his enthusiasm for cubulated groups! The authors thank the NSF and the NSERC for funds to cover travel to the other's home institution to work on this project. The second author wishes to thank Juliana Nalerio and Tim Nest for allowing the use of their couch while he visited the first author in New York. We are grateful to Tim Susse for carefully reading an earlier draft and offering several helpful corrections. We also thank Michah Sageev, Eric Swenson, and Dani Wise for discussions relating to an early (unfortunately, unsuccessful) approach we had to the results in Section 5; we remain optimistic those ideas will bear fruit in future work. We finally thank the referee for their careful reading and many helpful comments and corrections.

1. Preliminaries

The summary of thick metric spaces and groups given in Section 1.1 is based on the discussion in [4]. Section 1.3 provides a brief review of CAT(0) cube complexes, and Section 1.2 recalls some facts about divergence.

1.1. Thick spaces and groups

1.1.1. Asymptotic cones. Let (M, d) be a metric space and let $\omega \subset 2^{\mathbb{N}}$ be an ultrafilter on $\mathbb{N}$. Given a sequence $m = (m_n \in M)_{n\in\mathbb{N}}$ of *observation points* and a positive sequence $s = (s_n)_{n\in\mathbb{N}}$ with $s_n \xrightarrow{n} \infty$, the *asymptotic cone* $\mathbf{Cone}_\omega(M, m, s)$ is the ultralimit of the based metric spaces $\lim_\omega \left(M, m_n, \frac{d}{s_n}\right)$. More precisely, define a pseudometric $\mathbf{d}_\omega$ on $\prod_n M$ by letting

$$\mathbf{d}_\omega(y, z) = \lim_\omega \frac{d(y_n, z_n)}{s_n},$$

and consider the induced pseudometric on the component containing m, i.e.,

$$\widehat{M} = \left\{(y_n)_{n\in\mathbb{N}} \in \prod_n \left(M, \frac{d}{d_n}\right) : \mathbf{d}_\omega(y_n, m_n) < \infty\right\}.$$

Then $\mathbf{Cone}_\omega(M, m, s)$ is the associated quotient metric space, obtained from $\widehat{M}$ by identifying points y and z for which $\mathbf{d}_\omega(y, z) = 0$. A priori, $\mathbf{Cone}_\omega(M, m, s)$ depends on the observation point m, the sequence s, and the ultrafilter ω.

When M admits an isometric action by a group G such that some bounded subset of M meets every G-orbit, then $\mathbf{Cone}_\omega(M, m, s)$ is independent of the choice of observation point m, and it suffices to consider $\mathbf{Cone}_\omega(M, m, s)$, where, for some fixed basepoint m_o, the observation point $m_n = m_o$ for all $n \in \mathbb{N}$. In most of our applications, M comes equipped with a geometric group action, and thus the asymptotic cone is independent of the choice of observation point.

1.1.2. Unconstricted spaces and groups. A point $c \in M$ is a *cut-point* if $M - \{c\}$ has at least two connected components. By convention, c is a cut-point of the space $\{c\}$.

Definition 1.1 (unconstricted space, wide space). The metric space (M, d) is *unconstricted* if it satisfies each of the following:

(1) there exists $\kappa < \infty$ such that for all $m \in M$, there exists a quasi-isometric embedding $\gamma: \mathbb{R} \to M$ such that $d(m, \gamma) < \kappa$;

(2) there exists an ultrafilter ω and a sequence s such that for any sequence m of observation points in M, there is no cut-point in $\mathbf{Cone}_\omega(M, m, s)$.

If for all ultrafilters ω, all sequences m of observation points, and all scaling sequences d, there is no cut-point in $\mathbf{Cone}_\omega(M, m, s)$, then M is *wide*.

Remark 1.2 (unconstricted group, wide group). Let the infinite finitely-generated group G act properly and cocompactly by isometries on (M, d). It is easy to see that Definition 1.1.(1) holds for M. Moreover, since $\mathbf{Cone}_\omega(M, m, s)$ is independent of m, Definition 1.1.(2) is satisfied exactly when at least one asymptotic cone of M does not have a cut-point. In particular, letting M be a Cayley graph of G and d the associated word-metric yields the notion of an *unconstricted group* and of a *wide* group.

The inductive definition of a thick metric space requires the notion of a *uniformly unconstricted* family of spaces.

Definition 1.3 (uniformly unconstricted, uniformly wide). The collection $(M_n, d_n)_{n\in\mathbb{N}}$ of metric spaces is *uniformly unconstricted* if there exists an ultrafilter ω and a sequence $(s_n)_{n\in\mathbb{N}}$ of scaling constants such that, for all observation points $m = (m_n \in M_n)_{n\in\mathbb{N}}$, the ultralimit $\lim_\omega \left(M_n, m_n, \frac{d_n}{s_n}\right)$ has no cut-point. If this ultralimit lacks cut-points for all choices of ultrafilter, observation points, and scaling constants, then $(M_n)_{n\in\mathbb{N}}$ is *uniformly wide*.

Definition 1.4 (Thick space, strongly thick space). The space (M, d) is *thick of order* 0 if it is unconstricted, and *strongly thick of order* 0 if it is wide. Let $\mathcal{S}$ be a collection of subsets of M which are each (strongly) thick of order at most n. Then M is *τ-thick* (*(τ, η)-strongly thick*) *of order at most $n + 1$ with respect to $\mathcal{S}$* if there exists $\tau, \eta \geq 0$ such that each of the following holds:

(1) for each $m \in M$, there exists $S \in \mathcal{S}$ with $d(m, S) \leq \tau$;

(2) each $S \in \mathcal{S}$ is τ-quasiconvex in M, i.e., any two points in S can be connected by a (τ, τ)-quasigeodesic in $N_\tau(S)$;

(3) for all $S, S' \in \mathcal{S}$, there exists a sequence

$$S = S_0, S_1, \ldots, S_k = S', \quad \text{with } S_i \in \mathcal{S}$$

such that for all $0 \leq i < k$, the subspace $\mathcal{N}_\tau(S_i \cap S_{i+1})$ is of infinite diameter and τ-path-connected. Strong thickness requires a strengthening of this condition, namely that for any S, S' that both intersect $\mathcal{N}_{3\tau}(x)$ for some $x \in M$, the preceding sequence can always be chosen so that $k \leq \eta$ and $x \in \mathcal{N}_\eta(S_i)$ for $0 \leq i \leq k$.

Further, we say a family of metric spaces $\mathcal{M}$ is *uniformly thick* (*uniformly strongly thick*) *of order at most $n + 1$* if it satisfies:

(4) (a) there exists constants τ and η as above such that each $M \in \mathcal{M}$ is τ-thick ((τ, η)-strongly thick) of order at most $n+1$ with respect to a collection, $\mathcal{S}_M$, of subsets of M;

(b) $\bigcup_{M\in\mathcal{M}} \mathcal{S}_M$ is uniformly thick (uniformly strongly thick) of order at most n.

We typically drop the constants τ and η from the notation, as the precise constants are rarely of interest; is is usually important only that some constants exist.

If M is (τ, η)-(strongly) thick of order at most n and is not (τ', η')-(strongly) thick of order at most $n-1$ for any τ', η', then M is (*strongly*) *thick of order* n.

Following [4] and [3], we define *algebraic thickness* and *strong algebraic thickness* of a group as follows.

Definition 1.5 (algebraically thick). The finitely generated group G is *algebraically thick of order 0* if it is unconstricted. For $n \geq 1$, the group G is *algebraically thick of order at most* $n+1$ if there exists a finite collection $\mathcal{G}$ of finitely generated undistorted subgroups of G such that:

(1) there exists a finite index subgroup $G' \leq G$ generated by a finite subset of $\bigcup_{H \in \mathcal{G}} H$;

(2) each $H \in \mathcal{G}$ is algebraically thick of order at most n;

(3) for all $H, H' \in \mathcal{G}$, there exists a finite sequence $H = H_1, \ldots, H_m = H'$ such that each $H_i \in \mathcal{G}$ and $H_i \cap H_{i+1}$ is infinite for $1 \leq i \leq m-1$.

If G is algebraically thick of order at most $n+1$ and is not algebraically thick of order at most n, then G is *algebraically thick of order* $n+1$.

Definition 1.6 (strongly algebraically thick). The finitely generated group Γ is *strongly algebraically thick of order* 0 if it is wide. For $n \geq 1$, the group Γ is *strongly algebraically thick of order at most* $n+1$ if there exists a finite collection $\mathbb{G}$ of finitely generated undistorted subgroups of G such that:

(1) there exists a finite index subgroup $\Gamma' \leq \Gamma$ generated by a finite subset of $\bigcup_{H \in \mathbb{G}} H$;

(2) each $H \in \mathbb{G}$ is strongly algebraically thick of order at most n;

(3) for $H, H' \in \mathbb{G}$, there exists a sequence $H = H_0, \ldots, H_n = H'$ such that $H_i \in \mathbb{G}$ for each i, and $H_i \cap H_{i+1}$ is infinite and M-path-connected for $0 \leq i < n$;

(4) there exists $M \geq 0$ such that each $H \in \mathbb{G}$ is M-quasiconvex.

If Γ is strongly algebraically thick of order at most $n+1$, but is not strongly algebraically thick of order n, then Γ is strongly algebraically thick of order $n+1$.

Note that if Γ is strongly algebraically thick of order n, then Γ is algebraically thick of order at most n.

1.2. Divergence. The notion of the divergence function of a metric space goes back to Gromov [22] and Gersten [20, 19]; the present summary follows [3].

Definition 1.7 (divergence). Let (M,d) be a geodesic metric space and fix $\lambda \in (0,1), \mu \geq 0$. For $a,b,c \in M$, with $d(c,\{a,b\}) = r > 0$, let $\mathbf{div}_{\lambda,\mu}(a,b,c)$ to be the infimum of the set $\{|P|\}$, where P varies over all paths in M that join a to b and satisfy $d(P(t),c) \geq \lambda r - \mu$ for all t.

The *divergence*

$$\mathbf{Div}^M_{\lambda,\mu} \colon \mathbb{N} \longrightarrow \mathbb{R}^+$$

of M with respect to λ, μ is defined by

$$\mathbf{Div}^M_{\lambda,\mu}(n) = \sup\{\mathbf{div}_{\lambda,\mu}(a,b,c) \colon d(a,b) \leq n\}.$$

For any function $f \colon \mathbb{N} \to \mathbb{R}^+$, the space M has divergence at most f if for some λ, μ, and for all $n \in \mathbb{N}$, we have $\mathbf{Div}^M_{\lambda,\mu}(n) \leq f(n)$, and the notion of a space with divergence at least f is defined analogously. As usual, for functions f, g, we write $f \preceq g$ if for all n, we have $f(n) \leq Kg(Kn+K)+K$ for some constant K, and $f \asymp g$ if $f \preceq g$ and $g \preceq f$. For $d \geq 1$, the space M has divergence *of order at most* d if $\mathbf{Div}^M_{\lambda,\mu} \preceq p$ for some λ, μ, where p is a polynomial of degree d, and *order* d if it has divergence of order at most d but does not have divergence of order at most $d-1$.

There are several alternative notions of divergence discussed in [3, Section 3]. In the situations of interest in this paper, M admits a proper, cocompact group action and thus the various divergence functions coincide up to $\asymp$, by [16, Corollary 3.2]. Further, under the hypotheses of [16, Corollary 3.2], the $\asymp$-class of the divergence of M is a quasi-isometry invariant, in the following sense: if $q \colon M \to M'$ is a quasi-isometry, then for some $\lambda, \lambda' \in (0,1), \mu, \mu' \geq 0$, we have $\mathbf{Div}^M_{\lambda,\mu} \asymp \mathbf{Div}^M_{\lambda',\mu'}$, and in particular the divergence order of M (if it exists) is a quasi-isometry invariant. Hence the divergence of a finitely-generated group is well-defined, and it is sensible to speak of groups with linear, quadratic, exponential, etc. divergence.

In this paper, we study divergence of cocompactly cubulated groups by studying thickness of cube complexes. The relationship between the thickness order and the divergence order of M is not yet fully understood (see, e.g. [3, Question 1.2]). One useful result that is established is the following, which we will use in Section 7, in conjunction with lower bounds on divergence for some cocompactly cubulated groups, in order to provide lower bounds on the order of thickness.

Proposition 1.8 (Corollary 4.17 of [3]). *Let M be a geodesic metric space that is strongly thick of order at most n. Then*

$$\mathbf{Div}^M_{\lambda,\mu}(r) \preceq r^{n+1}$$

for all $\lambda \in \left(0, \frac{1}{54}\right), \mu \geq 0$.

1.3. CAT(0) cube complexes

1.3.1. Cube complexes and hyperplanes. A *cube complex* $\mathbf{X}$ is a CW-complex whose cells are Euclidean unit cubes of the form $[-\frac{1}{2}, \frac{1}{2}]^d$ for $0 \leq d < \infty$, attached in such a way that any two cubes (not necessarily distinct) of $\mathbf{X}$ with nonempty intersection intersect in a common face. The *dimension* $\dim \mathbf{X}$ is the supremum of the set of $d \geq 0$ for which $\mathbf{X}$ contains a d-cube.

$\mathbf{X}$ is *nonpositively-curved* if for each $x \in \mathbf{X}^{(0)}$, the link of x is a simplicial flag complex, and "CAT(0)" if it is nonpositively-curved and simply connected. As observed by Gromov in [21] and in full generality by Leary [32], the CAT(0) cube complex $\mathbf{X}$ is endowed with a CAT(0) geodesic metric, denoted $\ddot{d}$, obtained by regarding each cube as a Euclidean unit cube (see also the more general results of Bridson and Moussong on the existence of CAT(0) metrics for many polyhedral complexes [9, 35]). It is often convenient to view the 1-cubes as unit intervals and use the combinatorial metric $\dot{d}$ on the graph $\mathbf{X}^{(1)}$.

These two geometries essentially agree when $\dim \mathbf{X} < \infty$ in the sense that $(\mathbf{X}, \ddot{d})$ is quasi-isometric to $(\mathbf{X}^{(1)}, \dot{d})$. The metric $\dot{d}$ is determined by hyperplanes, as explained below, and these hyperplanes can be used to provide a nice characterization of isometric embeddedness and convexity of subcomplexes. Since we are concerned with finite-dimensional cube complexes, we use whichever metric is most convenient in a given situation.

For $d \geq 1$, the d-cube c has d *midcubes*, which are subspaces obtained by restricting exactly one coordinate to 0. A *hyperplane* H of the CAT(0) cube complex $\mathbf{X}$ is a connected subspace such that for each cube c of $\mathbf{X}$, either $H \cap c = \emptyset$, or $H \cap c$ is a midcube of c. The *carrier* $N(H)$ of H is the union of all closed cubes c for which $H \cap c \neq \emptyset$. Each hyperplane H is itself a CAT(0) cube complex of dimension at most $\dim \mathbf{X} - 1$, and $N(H)$ is a CAT(0) cube complex isomorphic to $H \times [-\frac{1}{2}, \frac{1}{2}]$. Furthermore, H and $N(H)$ are convex with respect to $\ddot{d}$, and $N(H)^{(1)}$ is convex in $\mathbf{X}^{(1)}$, with respect to $\dot{d}$ (see [12, 42]).

Crucially, Sageev showed in [42] that, for each hyperplane H of $\mathbf{X}$, the complement $\mathbf{X} - H$ has exactly two components, called *halfspaces* (*associated to* H) and denoted $\overleftarrow{H}, \overrightarrow{H}$. We denote by $\mathcal{H}$ the set of hyperplanes in $\mathbf{X}$ and by $\widehat{\mathcal{H}}$ the set of halfspaces. If $A, B \subset \mathbf{X}$, then $H \in \mathcal{H}$ *separates* A and B if $A \subset \overleftarrow{H}$ and $B \subset \overrightarrow{H}$ or vice versa.

For each 1-cube c of $\mathbf{X}$, there is a unique hyperplane H that separates the endpoints of c. H is the hyperplane *dual* to c, and c is a 1-cube *dual* to H. It can be shown that a path $P \to \mathbf{X}^{(1)}$ is a $\dot{d}$-geodesic if and only if P contains at most one 1-cube dual to each $H \in \mathcal{H}$. Hence, for $x, y \in \mathbf{X}^{(0)}$, the number of hyperplanes separating x from y is exactly $\dot{d}(x, y)$. Usefully, it is also true that a path $P \to \mathbf{X}$ is an $\ddot{d}$-geodesic only if for each $K \in \widehat{\mathcal{H}}$, the intersection $P \cap K$ is connected.

Distinct $H_1, H_2 \in \mathcal{H}$ *contact* if $N(H_1) \cap N(H_2) \neq \emptyset$ (equivalently, no third hyperplane separates H_1 from H_2). This can happen in one of two ways: if $H_1 \cap H_2 \neq \emptyset$, then H_1 and H_2 *cross*. Crossing is also characterized by the fact that $\overleftarrow{H_1} \cap \overleftarrow{H_2} \neq \emptyset$, $\overrightarrow{H_1} \cap \overleftarrow{H_2} \neq \emptyset$, $\overrightarrow{H_1} \cap \overrightarrow{H_2} \neq \emptyset$, $\overleftarrow{H_1} \cap \overrightarrow{H_2} \neq \emptyset$, and by the fact that $N(H_1) \cap N(H_2)$ contains a 2-cube whose 1-cubes are dual to H_1 or H_2. If H_1 and H_2 contact and do not cross, then they *osculate*.

More generally, if $A \subset \mathbf{X}$ is a connected subspace and $H \in \mathcal{H}$, then H *crosses* A if $A \cap \overrightarrow{H}$ and $A \cap \overleftarrow{H}$ are both nonempty. We denote by $\mathcal{H}(A)$ the set of hyperplanes crossing A. A connected full subcomplex $\mathbf{Y} \subseteq \mathbf{X}$ is *isometrically embedded* if the inclusion $\mathbf{Y}^{(1)} \to \mathbf{X}^{(1)}$ is an isometric embedding. Equivalently, $\bigcap_i H_i \cap \mathbf{Y}$ is connected for each $\{H_i\} \subset \mathcal{H}(\mathbf{Y})$. Similarly, $\mathbf{Y}$ is *convex* if, for any collection $H_1, \ldots, H_n \in \mathcal{H}(\mathbf{Y})$ of pairwise-crossing hyperplanes, $\mathbf{Y}$ contains an n-cube of $\bigcap_{i=1}^n N(H_i)$. This notion turns out to coincide with CAT(0)–convexity for subcomplexes [23]; it also equivalent to the requirement that $\mathbf{Y}^{(1)}$ be a convex subgraph of $\mathbf{X}^{(1)}$ and every cube of $\mathbf{X}$ whose 1-skeleton lies in $\mathbf{Y}$ itself lies in $\mathbf{Y}$.

1.3.2. Actions on cube complexes. By $\mathrm{Aut}(\mathbf{X})$, we mean the group of cubical automorphisms of the CAT(0) cube complex $\mathbf{X}$, and by an action of the group G on $\mathbf{X}$, we mean a homomorphism $G \to \mathrm{Aut}(\mathbf{X})$. Such an action is also an action by $\dot{d}$-isometries on $\mathbf{X}^{(1)}$ and by $\ddot{d}$-isometries on $\mathbf{X}$.

This action is *proper* if the stabilizer of any cube of $\mathbf{X}$ is finite, and *metrically proper* if for all infinite sequences $(g_n \in G)_{n \geq 0}$ of distinct elements, and for all $x \in \mathbf{X}$, we have $\ddot{d}(x, g_n x) \to \infty$ as $n \to \infty$. Generally, we are concerned with cocompact actions, and in this situation the notions of properness and metric properness coincide. A proper action of G on a CAT(0) cube complex is a *cubulation* of G, and if such an action exists, G is *cubulated*. If G acts geometrically on a CAT(0) cube complex, then G is *cocompactly cubulated*.

Each $g \in \mathrm{Aut}(\mathbf{X})$ acts as an isometry of both the CAT(0) space $(\mathbf{X}, \ddot{d})$ and the median graph $(\mathbf{X}^{(1)}, \dot{d})$. According to [23], either g fixes the barycenter of a cube of $\mathbf{X}$, or there exists a combinatorial geodesic $\gamma\colon \mathbb{R} \to \mathbf{X}$ and some $N = N(\dim \mathbf{X}), \tau > 0$ such that $g^N \gamma(t) = \gamma(t + \tau)$ for all $t \in \mathbb{R}$; such an element g^N is *combinatorially hyperbolic* and γ is a *combinatorial axis for* g^N. Likewise, if g does not fix a point of $\mathbf{X}$, then since isometries of CAT(0) spaces are semisimple, g acts by translations on a CAT(0) geodesic $\alpha\colon \mathbb{R} \to \mathbf{X}$, called an *axis* for g. If γ is combinatorially rank-one (equivalently, α is rank-one) for some combinatorial axis γ (CAT(0) axis α), then g is a *rank-one isometry*.

The hyperplane $H \in \mathcal{H}$ is a *leaf* if at least one of $\overleftarrow{H}, \overrightarrow{H}$ fails to contain a hyperplane. $\mathbf{X}$ is *essential* if it contains no leaves. If G acts on $\mathbf{X}$, then H is a G*-leaf* if there exists $r \geq 0$ such that, for $A \in \{\overleftarrow{H}, \overrightarrow{H}\}$ and for all $x \in \mathbf{X}$, $\ddot{d}(gx, H) \leq r$ for all $g \in G$ such that $gx \in A$. The action of G on $\mathbf{X}$ is *essential* if $\mathbf{X}$ contains no G-leaves. Usually, we will assume that G acts essentially on $\mathbf{X}$, abetted by [15, Proposition 3.5] and Lemma 2.16 below. The former says, in particular, that if G acts geometrically on $\mathbf{X}$, then there is a convex, G-cocompact subcomplex $\mathbf{Y} \subseteq \mathbf{X}$ on which G acts essentially. The latter says that the simplicial boundaries of $\mathbf{X}$ and $\mathbf{Y}$ coincide.

We will occasionally need some notion of quasiconvexity of subgroups. Since the groups under consideration are not in general hyperbolic, quasiconvexity of a subgroup depends on the choice of generating set. However, the groups in this section come equipped with specific geometric actions on metric spaces; accordingly, we use:

Definition 1.9 (quasiconvex). Let the group G act properly and cocompactly on the metric space M. The subgroup $H \leq G$ is *quasiconvex* if for some (and hence any) $m \in M$, the orbit Hm is a quasiconvex subspace of M.

This definition is not intrinsic either to G or to M, but rather depends on the particular action of G on M. Note, in particular, that this property implies that for any fixed word metric on G, there exist uniform constants such that any pair of point in H can be joined by a uniform quality quasigeodesic contained inside a uniform neighborhood of H. This latter, weaker property is the one considered in [3], and it holds for subgroups that are quasiconvex as defined above.

2. The simplicial boundary

The definition and basic properties of the simplicial boundary of a CAT(0) cube complex are discussed in [26], and we recall these here briefly, before establishing some simple facts about the simplicial boundary that will be necessary in subsequent sections.

2.1. Boundary sets. Let $\mathbf{X}$ be a CAT(0) cube complex and suppose that the set $\mathcal{H}$ of hyperplanes contains no infinite set of pairwise-crossing hyperplanes. This holds for all cube complexes in this paper, since they are finite-dimensional by virtue of cocompactness.

Definition 2.1 (closed under separation). $\mathcal{U} \subseteq \mathcal{H}$ is *closed under separation* if for all $H_1, H_2 \in \mathcal{U}$, if some hyperplane H_3 separates H_1 from H_2, then $H_3 \in \mathcal{U}$.

For example, if $A \subset \mathbf{X}$ is a connected subspace, then $\mathcal{H}(A)$ is closed under separation.

Definition 2.2 (unidirectional). $\mathcal{U} \subseteq \mathcal{H}$ is *unidirectional* if for each $H \in \mathcal{U}$, at most one of $\overleftarrow{H}$ or $\overrightarrow{H}$ contains infinitely many elements of $\mathcal{U}$.

The motivating example of a set that is *not* unidirectional is the set $\mathcal{H}(\gamma)$, where γ is a bi-infinite combinatorial geodesic in a CAT(0) cube complex $\mathbf{X}$ in which every set of pairwise-crossing hyperplanes is finite.

Definition 2.3 (facing triple). A *facing triple* $\{H_1, H_2, H_3\} \subseteq \mathcal{H}$ is a set of three distinct hyperplanes, any two of which are contained in a single halfspace associated to the third. Equivalently, $\{H_1, H_2, H_3\}$ is a facing triple if no three of the associated halfspaces are totally ordered by inclusion.

Definition 2.4 (boundary set, boundary set equivalence). $\mathcal{U} \subseteq \mathcal{H}$ is a *boundary set* if $\mathcal{U}$ is infinite, unidirectional, closed under separation, and contains no facing triple.

Let $\mathcal{U}_1, \mathcal{U}_2$ be boundary sets. Then $\mathcal{U}_1 \precsim \mathcal{U}_2$ if $|\mathcal{U}_1 - \mathcal{U}_1 \cap \mathcal{U}_2| < \infty$. If $\mathcal{U}_1 \precsim \mathcal{U}_2$ and $\mathcal{U}_2 \precsim \mathcal{U}_1$, i.e., if $|\mathcal{U}_1 \triangle \mathcal{U}_2| < \infty$, then $\mathcal{U}_1$ and $\mathcal{U}_2$ are *equivalent* boundary sets, denoted $\mathcal{U}_1 \sim \mathcal{U}_2$. The boundary set $\mathcal{U}$ is *minimal* if for each boundary set $\mathcal{U}'$ with $\mathcal{U}' \precsim \mathcal{U}$, we have $\mathcal{U}' \sim \mathcal{U}$.

The following lemma from [26] explains why we assume that sets of pairwise-crossing hyperplanes are finite:

Lemma 2.5. *Any boundary set in $\mathcal{H}$ contains a minimal boundary set.*

Indeed, an infinite set of pairwise-crossing hyperplanes is, by definition, a boundary set, but such a set is easily seen to fail to contain a minimal boundary set. Lemma 2.5 is needed to prove Proposition 2.6 (which is [26, Proposition 3.10]), and this statement is in turn required when defining the simplicial boundary.

Proposition 2.6. *Let $\mathcal{U}$ be a boundary set. Then there exists $k \leq \dim \mathbf{X}$ and pairwise-disjoint minimal boundary sets $\mathcal{U}_1, \dots, \mathcal{U}_k$ such that $\bigsqcup_{i=1}^{k} \mathcal{U}_i \sim \mathcal{U}$ and, for each $1 \leq i < j \leq k$ and each $U \in \mathcal{U}_j$, the set of $V \in \mathcal{U}_i$ such that $U \cap V = \emptyset$ is finite.*

Moreover, if $\mathcal{U}'_1, \dots \mathcal{U}'_{k'}$ are pairwise-disjoint minimal boundary sets such that $\bigsqcup_{i=1}^{k'} \mathcal{U}'_i \sim \mathcal{U}$, then $k = k'$ and, after relabeling, $\mathcal{U}_i \sim \mathcal{U}'_i$ for all i.

2.2. Simplices at infinity. The *dimension* of the boundary set $\mathcal{U}$ is equal to $k - 1$, where k is the number of minimal boundary sets in the decomposition of $\mathcal{U}$ given by Proposition 2.6. In particular, the minimal boundary sets are exactly those that have dimension 0, and the dimension of any boundary set is finite, by Proposition 2.6, since $\mathbf{X}$ has no infinite set of pairwise-crossing hyperplanes. Note also that if $\mathcal{U} \sim \mathcal{U}'$, then their dimensions coincide. Accordingly, for each $k \geq 0$, let $\mathfrak{S}(k)$ be the set of $\sim$-classes u such that some (and hence every) representative $\mathcal{U}$ of u is a k-dimensional boundary set.

Definition 2.7 (simplicial boundary). Let $\mathbf{X}$ be a CAT(0) cube complex with no infinite set of pairwise-crossing hyperplanes. The *simplicial boundary* $\partial_{\triangle}\mathbf{X}$ of $\mathbf{X}$ is the simplicial complex $\bigcup_{k\geq 0}\mathfrak{S}(k)$, for $k \geq 0$, with the simplex u (represented by a boundary set $\mathcal{U}$) a face of v (represented by $\mathcal{V}$) exactly when $\mathcal{U} \precsim \mathcal{V}$.

For example, it is easily verified that the simplicial boundary of an infinite tree is a discrete set, and that the simplicial boundary of the standard tiling of $\mathbb{E}^2$ by 2-cubes is a 4-cycle. In [26], it is shown that $\partial_{\triangle}\mathbf{X}$ is a flag complex, every simplex of $\partial_{\triangle}\mathbf{X}$ is contained in a finite-dimensional maximal simplex.

2.3. Visibility and cubical flats. The motivating example of a boundary set is the set $\mathcal{H}(\gamma)$ of hyperplanes that cross the (combinatorial or CAT(0)) geodesic ray γ, but there are boundary sets not of this type: see [26, Example 3.17]. Following this example, a simplex v is called *visible* if there exists a combinatorial geodesic ray γ such that $\mathcal{H}(\gamma)$ represents the $\sim$-class v. By [26, Theorem 3.19]), each maximal simplex is visible. In this paper, $\mathbf{X}$ is often assumed to be *fully visible*, meaning that each simplex is visible. We believe the following is plausible and would remove the need for to hypothesis fully visible from several results in this paper, but a proof of this result appears to be tricky.

Conjecture 2.8. *Let* $\mathbf{X}$ *be a locally finite* CAT(0) *cube complex for which some* $G \leq \mathrm{Aut}(\mathbf{X})$ *acts cocompactly. Then* $\mathbf{X}$ *is fully visible.*

We shall occasionally use the fact that full visibility is inherited by convex subcomplexes.

Definition 2.9 (flat, orthant, cubical flat). For $d \geq 0$, a *d-flat* in $\mathbf{X}$ is the image of an isometric embedding $\mathbb{E}^d \to (\mathbf{X}, \ddot{d})$. An *orthant* is the image of an isometric embedding $([0,\infty)^d, \mathbf{d}_{\mathbb{E}^d}) \to (\mathbf{X}, \ddot{d})$. A *cubical flat* is an isometrically embedded subcomplex $\mathbf{F} \subseteq \mathbf{X}$ that is isomorphic to the standard tiling of $\mathbb{E}^d$ by unit d-cubes for some $d \geq 0$. A cubical orthant is defined similarly, in terms of the standard tiling of $[0,\infty)^d$.

The simplicial boundary of a d-dimensional cubical orthant is easily seen to be a $(d-1)$-simplex, for $d \geq 1$. Similarly, one checks that the simplicial boundary of a d-dimensional cubical flat is isomorphic to the $(d-1)$-dimensional *spherical hyperoctahedron* $\mathbb{O}_d$. This simplicial complex is defined as follows: $\mathbb{O}_1$ consists of a pair of 0-simplices, and for $d \geq 1$, $\mathbb{O}_d$ is the simplicial join of $\mathbb{O}_0$ and $\mathbb{O}_{d-1}$. Under the hypothesis of full visibility, the presence of a d-simplex at infinity ensures the presence of an isometric cubical orthant; likewise, the presence of a hyperoctahedra in the boundary yields a flat.

Proposition 2.10 (Theorem 3.23 of [26]). *Let* $\mathbf{X}$ *be fully visible and let* $v \subseteq \partial_{\triangle}\mathbf{X}$ *be a simplex. Then there is a cubical orthant* $\mathbf{F} \subseteq \mathbf{X}$ *with* $\mathcal{H}(\mathbf{F})$ *representing* v.

It will be necessary to reach conclusions similar to that of Proposition 2.10, but in the CAT(0) setting.

Proposition 2.11 (simplices yield orthants). *Let* $\mathbf{X}$ *be fully visible, and let* $\mathcal{V}$ *be a boundary set of dimension* $d \geq 1$. *Then there exists a* $(d+1)$*-dimensional orthant* $\mathbf{O} \subseteq \mathbf{X}$ *such that* $\mathcal{H}(\mathbf{O}) \sim \mathcal{V}$.

Proof. By Proposition 2.10, there exists an isometric cubical orthant $\mathbf{C}$ in $\mathbf{X}$ with $\mathcal{H}(\mathbf{C}) \sim \mathcal{V}$. Let $v_1, \dots, v_{d+1}$ be the 0-simplices of v. For $1 \leq i \leq d+1$, there is a combinatorial geodesic ray γ_i such that the γ_i all have common basepoint, and $\mathcal{H}(\mathbf{C}) = \bigsqcup_i \mathcal{H}(\gamma_i)$, and for $i' \neq j$, every $V \in \mathcal{H}(\gamma_i)$ crosses every $H \in \mathcal{H}(\gamma_j)$. As is shown in [26], there exists, for each i, a CAT(0) geodesic ray α_i in $\mathbf{X}$ with $\alpha_i(0) = \gamma_i(0)$ and $\mathcal{H}(\gamma_i) = \mathcal{H}(\alpha_i)$. The preceding crossing property ensures that $\mathbf{X}$ contains $\prod_i \alpha_i$, which is the desired CAT(0) orthant. □

Definition 2.12 (maximal orthant). The orthant $\mathbf{O} \subseteq \mathbf{X}$ is *maximal* if for all orthants $\mathbf{O}'$ that coarsely contain $\mathbf{O}$, $\dim \mathbf{O}' = \dim \mathbf{O}$.

Proposition 2.13 (orthants yield simplices). *Let* $\mathbf{X}$ *be fully visible and let* $\mathbf{O} \subseteq \mathbf{X}$ *be a* d*-dimensional maximal orthant or cubical orthant. Then* $\mathcal{H}(\mathbf{O})$ *represents a* $(d-1)$*-simplex of* $\partial_\triangle \mathbf{X}$.

Proof. Let $\mathcal{V} = \mathcal{H}(\mathbf{O})$, and let $\mathcal{V} = \bigcup_{i=1}^{e} \mathcal{V}_i$ be a decomposition into minimal boundary sets such that, for all $i \neq j$, if $H \in \mathcal{V}_i$ and $V \in \mathcal{V}_j$, then H crosses V. Now, $e \geq d$ since $\mathbf{O}$ is a d-flat. On the other hand, the proof of Proposition 2.10 shows that $\mathbf{O}$ is contained in an e-dimensional cubical orthant, whence $d = e$. Thus $\mathcal{V}$ represents a $(d-1)$-simplex. □

Remark 2.14. The conclusion of Proposition 2.13 fails in the absence of maximality. This is roughly because, while an isometric embedding $\mathbf{Y} \to \mathbf{X}$ induces an embedding of simplicial boundaries, the image of $\partial_\triangle \mathbf{Y}$ may not be a subcomplex if $\mathbf{Y}$ is not convex. For example, consider the geodesic ray L in $\mathbb{E}^2$ beginning at $(0,0)$ and containing $(1,1)$. Let $\mathbf{X}$ be the standard tiling of $\mathbb{E}^2$ by 2-cubes, and let $\mathbf{Y}$ be a combinatorial geodesic ray whose 0-cubes are the points $(n,n), (n+1,n)$ $n \geq 0$. No two hyperplanes of $\mathbf{Y}$ cross in $\mathbf{Y}$, so that $\partial_\triangle \mathbf{Y}$ is a 0-simplex. But $\mathcal{H}(\mathbf{Y})$ determines a 1-simplex of $\partial_\triangle \mathbf{X}$.

Proposition 2.13 also requires full visibility. For example, if $\mathbf{X}$ is an *eighth-flat* (see [26, Example 3.17]), a maximal cubical orthant is 1-dimensional but the set of dual hyperplanes corresponds to a 1-simplex of $\partial_\triangle \mathbf{X}$.

The following proposition characterizes hyperbolic proper, cocompact CAT(0) cube complexes using $\partial_\triangle \mathbf{X}$. In the fully visible case, the proof is simplified slightly by Proposition 2.10 and Proposition 2.13.

Proposition 2.15. *Let the* CAT(0) *cube complex* $\mathbf{X}$ *admit a proper, cocompact group action.* $\partial_\triangle \mathbf{X}$ *is discrete if and only if* $\mathbf{X}$ (*and therefore* $\mathbf{X}^{(1)}$) *is hyperbolic.*

Proof. If $\partial_\triangle \mathbf{X}$ consists entirely of isolated 0-simplices, then $\mathbf{X}$ cannot contain an isometrically embedded flat of dimension $d \geq 2$: if $\mathbb{E}^d \cong \mathbf{F} \to (\mathbf{X}, \ddot{d})$ is such an isometric embedding, then the cubical convex hull of $\mathbf{F}$ contains a boundary set of positive dimension, resulting in a positive-dimensional simplex of $\partial_\triangle \mathbf{X}$. Hence, by the Flat Plane Theorem [5], $\mathbf{X}$ is hyperbolic. Conversely, if v is a d-simplex with $d \geq 2$, then the intersection graph of the set of hyperplanes contains arbitrarily large complete bipartite graphs $K_{n,n}$, by the definition of a boundary set, whence $\mathbf{X}$ is not hyperbolic [25]. □

2.4. Essential actions and the simplicial boundary. We will require the following lemma in Section 5.

Lemma 2.16. *Let the group* G *act properly and cocompactly on the* CAT(0) *cube complex* $\mathbf{X}$. *Let* $\mathbf{X}_1 \subseteq \mathbf{X}$ *be a convex,* G*-cocompact subcomplex on which* G *acts essentially. Then* $\partial_\triangle \mathbf{X} \cong \partial_\triangle \mathbf{X}_1$.

Proof. By [26, Theorem 3.15], the inclusion $\mathbf{X}_1 \hookrightarrow \mathbf{X}$ induces a simplicial embedding $\partial_\triangle \mathbf{X}_1 \to \partial_\triangle \mathbf{X}$. It suffices to show that this map is surjective. If not, there exists a 0-simplex v of $\partial_\triangle \mathbf{X}$ that does not belong to the image of $\partial_\triangle \mathbf{X}_1$. This means that v is represented by a minimal boundary set $\mathcal{V}$ such that, for all $V \in \mathcal{V}$, the intersection $V \cap \mathbf{X}_1 = \emptyset$. We thus have a sequence of hyperplanes $\{V_i \in \mathcal{V}\}_{i \geq 0}$ such that for all $i \geq 1$, we have $V_i \subset \overrightarrow{V}_{i-1}$ and $\mathbf{X}_1 \subset \overleftarrow{V}_{i-1}$. Now, by cocompactness, there exists $R < \infty$ such that every point of $\mathbf{X}$ is of the form gx, where $g \in G$ and x lies in the R-neighborhood of some fundamental domain $K \subset \mathbf{X}_1$ for the action of G on $\mathbf{X}_1$. For any $j \geq 0$, we can choose $gx \in \overrightarrow{V}_1$ to be separated from V_1, and hence from $\mathbf{X}_1$, by at least j of the hyperplanes V_i. This contradicts the fact that G stabilizes any regular neighborhood of $\mathbf{X}_1$. Thus the embedding $\partial_\triangle \mathbf{X}_1 \to \partial_\triangle \mathbf{X}$ is surjective. □

Lemma 2.16 will be used in conjunction with [15, Proposition 3.5] in the following way: if we wish to make a statement about $\partial_\triangle \mathbf{X}$, where $\mathbf{X}$ admits a proper, cocompact action, then there is no harm in passing to a convex, cocompact, essential subcomplex.

2.5. Limit simplices, limit sets, and the visual boundary. In this section, $\mathbf{X}$ is a CAT(0) cube complex admitting a proper, cocompact action by a group G. Let $\partial_\infty \mathbf{X}$ denote the visual boundary of $(\mathbf{X}, \ddot{d})$, endowed with the cone topology. For a geodesic ray $\gamma \subset \mathbf{X}$, we denote by $[\gamma]$ the point of $\partial_\infty \mathbf{X}$ represented by γ. It is shown in [26, Section 3] that, when $\mathbf{X}$ is fully visible, there is a surjection

$f : \partial_\infty \mathbf{X} \to \partial_\triangle \mathbf{X}$ such that, if γ is a CAT(0) geodesic and u is the simplex of $\partial_\triangle \mathbf{X}$ represented by $\mathcal{H}(\gamma)$, then $f([\gamma]) \in u$.

In the interest of an explicit, self-contained account, we now describe the map $f : \partial_\infty \mathbf{X} \to \partial_\triangle \mathbf{X}$ when $\mathbf{X}$ is a fully visible CAT(0) cube complex admitting a proper, cocompact action by some group G. Fix a base 0-cube x_o, and choose for each $[\gamma] \in \partial_\infty \mathbf{X}$ a CAT(0) geodesic ray γ representing $[\gamma]$, with $\gamma(0) = x_o$. Let $u_{[\gamma]}$ be the simplex of $\partial_\triangle \mathbf{X}$ represented by $\mathcal{H}(\gamma)$, which is easily seen to be a boundary set. Note that if γ' fellow-travels with γ, then $|\mathcal{H}(\gamma) \triangle \mathcal{H}(\gamma)| < \infty$, whence $u_{[\gamma]} = u_{[\gamma']}$. Hence $u_{[\gamma]}$ is well-defined. Moreover, every simplex u of $\partial_\triangle \mathbf{X}$ satisfies $u = u_{[\gamma]}$ for some $\gamma \in \partial_\infty \mathbf{X}$, by full visibility of $\mathbf{X}$.

If $[\gamma]$ has the property that $\mathcal{H}(\gamma)$ is a minimal boundary set, then $u_{[\gamma]}$ is a 0-simplex, and we let $f([\gamma]) = u_{[\gamma]}$.

Next, let γ be a combinatorial geodesic ray with $\gamma(0) = x_o$ and $\mathcal{H}(\gamma)$ a representative set for a d-simplex u of $\partial_\triangle \mathbf{X}$, with $d \geq 2$. By Proposition 2.11, there exists an isometrically embedded maximal flat orthant $Y \subset \mathbf{X}$ with $|\mathcal{H}(\gamma) - \mathcal{H}(\gamma) \cap \mathcal{H}(Y)| < \infty$, so that the cubical convex hull $\widehat{Y}$ has the property that the inclusion $\widehat{Y} \to \mathbf{X}$ induces the inclusion $\partial_\triangle Y \cong u \hookrightarrow \partial_\triangle \mathbf{X}$.

Choose a geodesic ray $\sigma \subset Y$ such that $\mathcal{H}(\sigma)$ and $\mathcal{H}(\gamma)$ have finite symmetric difference, and such that $\sigma(0)$ is the image of the origin under $[0, \infty)^D \cong Y \hookrightarrow \mathbf{X}$. Let $\gamma_0, \ldots, \gamma_D$, with $D \geq d$, be a collection of CAT(0) geodesic rays such that $\mathbf{Y} = \gamma_0 \times \ldots \times \gamma_D$, so that u is spanned by the 0-simplices $f([\gamma_0]), \ldots, f([\gamma_D])$. Then σ is determined by a unit vector $(\alpha_i)_{i=0}^D$, where α_i is the projection in $\mathbf{Y}$ of $\gamma(1)$ to γ_i. Let $f([\gamma]) = f([\sigma])$ be the point $\sum_{i=0}^D \alpha_i f([\gamma_i])$. Note that this is well-defined: if γ' fellow-travels with γ, then $|\mathcal{H}(\gamma') \triangle \mathcal{H}(\sigma)| < \infty$.

The map f is surjective, by construction, and has the additional property that if $\mathcal{H}(\gamma)$ represents a simplex $u \subset \partial_\triangle \mathbf{X}$, then $f([\gamma]) \in u$, and if $f([\gamma]) \in u$ for some simplex u, then $\mathcal{H}(\gamma)$ represents u or one of its faces. (A priori, for f to be injective requires that any two geodesic rays representing the same 0-simplex of $\partial_\triangle \mathbf{X}$ fellow-travel, and so there are in general many orthants that are coarsely inequivalent but represent the same simplex; each is coarsely equivalent to some orthant in the convex hull of any of them, however. This explains the failure of f to be injective; see [26, Proposition 3.37].)

Definition 2.17 (limit simplex, limit set). Let $H \leq G$. The simplex $a \subseteq \partial_\triangle \mathbf{X}$ is a *limit simplex* for the action of H on $\mathbf{X}$ (and on $\partial_\triangle \mathbf{X}$) if for some (and hence any) 0-cube $x \in \mathbf{X}$, there exists a sequence $(h_i \in H)$ such that the set of hyperplanes V such that V separates $h_i x$ from x for all but finitely many i is a boundary set representing a. The *limit complex* for H is the smallest subcomplex that contains every limit simplex.

A point $p \in \mathbf{X} \cup \partial_\infty \mathbf{X}$ is in the *limit set* of H if for some (and hence any) $x \in \mathbf{X}$, there exists $(h_i \in H)_{i \geq 0}$ such that $h_i x$ converges to p in the cone topology.

The following lemma relates limit sets (which live in the visual boundary) to limit complexes (which live in the simplicial boundary).

Lemma 2.18. *Let $H \leq G$ and let $\mathbf{X}$ be finite-dimensional, locally finite, and fully visible, and let $u \subset \partial_{\triangle}\mathbf{X}$ be a simplex. If $f^{-1}(u) \subset \partial_{\infty}\mathbf{X}$ is contained in the limit set of H, then u is contained in a limit simplex for H.*

Proof. Choose a (combinatorial or CAT(0)) geodesic ray γ such that $\mathcal{H}(\gamma)$ represents the simplex u; this is possible since $\mathbf{X}$ is fully visible. Since $f^{-1}(u)$ is contained in the limit set of H, there is a sequence $(h_i \in H)_{i\geq 0}$ such that $h_i x_o$ converges to $[\gamma] \in \partial_{\infty}\mathbf{X}$, where $x_o = \gamma(0)$. Hence there exists $K \geq 0$ such that for all sufficiently large i, there exists n_i such that $\ddot{d}(\gamma(n_i), p_i) \leq K$, where p_i is the projection of $h_i x_o$ onto the (CAT(0)-metric) sphere of radius n_i about x_o (and $\ddot{d}(h_i x_o, x_o) \geq n_i$).

Let $\mathcal{U}$ be the set of hyperplanes W such that W separates x_o from $h_i x_o$ for all but finitely many values of i. Write $\mathcal{U} = \mathcal{U}_1 \sqcup \mathcal{U}_2$, where $\mathcal{U}_1$ is the set of hyperplanes in $\mathcal{U}$ that separate p_i from x_o for all but finitely many i. Since p_i lies on the geodesic from $h_i x_o$ to x_o, we note that each $V \in \mathcal{U}_1$ separating p_i from x_o also separates $h_i x_o$ from x_o.

Observe that $|\mathcal{U}_1 - \mathcal{H}(\gamma)| \leq K$. Indeed, a hyperplane in $\mathcal{U}_1 - \mathcal{H}(\gamma)$ must separate $\gamma(n_i)$ from p_i for all sufficiently large i.

Conversely, suppose that $\mathcal{H}(\gamma) - \mathcal{U}_1$ is infinite. Each $V \in \mathcal{H}(\gamma) - \mathcal{U}_1$ fails to separate p_i from x_o for arbitrarily large values of i, while separating x_o from $\gamma(n_j)$ for all but finitely many j. Thus V separates p_i from $\gamma(n_i)$ for arbitrarily large values of i.

Suppose $V_1, V_2, \ldots$ are hyperplanes with this property, numbered according to the order in which one encounters them while traveling along γ. Let M be the Ramsey number $R(\dim\mathbf{X}+1, K+1)$. Then $\{V_1, \ldots, V_M\}$ contains either $\dim\mathbf{X}+1$ pairwise-crossing hyperplanes, which is impossible, or $K+1$ pairwise-disjoint hyperplanes. In the latter case, renumber so that $V_1, \ldots, V_{K+1}$ are pairwise-disjoint hyperplanes. Since γ is a geodesic, each V_j either separates p_i from $\gamma(n_i)$ for all sufficiently large i, or V_j separates p_i from $p_{i'}$ for infinitely many values of i, i'. However, if $V_j, V_{j'}$ are both hyperplanes of the latter type then, since they cannot cross, they separate $p_i, \gamma(n_i)$ for the same values of i. Hence there exists i such that $K+1$ hyperplanes separate p_i from $\gamma(n_i)$, which is impossible. Hence $|\mathcal{H}(\gamma) - \mathcal{U}| < \infty$.

Thus $|\mathcal{U}_1 \triangle \mathcal{H}(\gamma)| < \infty$, i.e. $\mathcal{U}_1$ and $\mathcal{H}(\gamma)$ represent the same simplex u of $\partial_{\triangle}\mathbf{X}$. Suppose that $V \in \mathcal{U}_1$ and $W \in \mathcal{U}_2$. Then there exists I such that for all $i \geq I$, the points x_o and $h_i x_o$ are separated by W, but there are infinitely many i such that W separates $h_i x_o$ from p_i. Hence, since V separates x_o from $h_i x_o$, all but finitely many such V cross W.

Hence, if $\mathcal{U}_2$ is finite, then u is a limit simplex for H. Otherwise, by local finiteness, $\mathcal{U}_2$ contains a boundary set $\mathcal{U}_2'$ representing a simplex v of $\partial_\triangle \mathbf{X}$ such that $u \star v$ is also a simplex of $\partial_\triangle \mathbf{X}$. By definition, $u \star v$ is a limit simplex for H. □

3. Relatively hyperbolic cubulated groups

Before studying cocompactly cubulated groups that are thick, we consider a natural class of such groups that are not thick, namely those that are relatively hyperbolic. We saw in Proposition 2.15 that if the infinite, finitely generated group G acts properly and cocompactly on the CAT(0) cube complex $\mathbf{X}$, then G is hyperbolic if and only if $\partial_\triangle \mathbf{X}$ is an infinite set of 0-simplices. It is natural to ask how this extends to relatively hyperbolic groups; in this section we shall provide a complete characterization of relatively hyperbolic cocompactly cubulated groups, in terms of the simplicial boundary.

Note that a subset of $\mathbf{X}^{(0)}$ is quasiconvex in $(\mathbf{X}, \ddot{d})$ if and only if it is quasiconvex in $(\mathbf{X}^{(1)}, \dot{d})$. Hence in what follows, we sometimes say that $A \subset \mathbf{X}$ is *quasiconvex in* $\mathbf{X}$ to mean that the set of 0-cubes of A is quasiconvex in $\mathbf{X}^{(0)}$.

3.1. The simplicial boundary of a relatively hyperbolic cube complex. Suppose that the group G acts properly and cocompactly on the CAT(0) cube complex $\mathbf{X}$, and is hyperbolic relative to a collection $\mathbb{P}$ of peripheral subgroups. Now, each $P \in \mathbb{P}$ is the stabilizer of a single vertex in an appropriately chosen fine hyperbolic graph for $(G, \mathbb{P})$ (see [8, 43]) and therefore acts on that graph with a quasiconvex orbit. (The latter condition is called *relative quasiconvexity* in [43].) By [43, Theorem 1.1], there exists a convex (and hence CAT(0)) P-invariant subcomplex $\mathbf{Y}_P \subseteq \mathbf{X}$. By [26, Theorem 3.15], the inclusion $\mathbf{Y}_P \to \mathbf{X}$ induces a simplicial embedding $\partial_\triangle \mathbf{Y}_P \to \partial_\triangle \mathbf{X}$. Now, if $\mathbf{Y}, \mathbf{Y}'$ are convex, P-cocompact subcomplexes, then each lies in a finite neighborhood of the other, and it follows that $\mathcal{H}(\mathbf{Y})$ and $\mathcal{H}(\mathbf{Y}')$ have finite symmetric difference, so that the images of $\partial_\triangle \mathbf{Y}$ and $\partial_\triangle \mathbf{Y}'$ in $\partial_\triangle \mathbf{X}$ coincide. We denote by $\mathcal{I}$ the set of isolated 0-simplices of $\partial_\triangle \mathbf{X}$.

Theorem 3.1. *Let G be hyperbolic relative to a collection $\mathbb{P}$ of peripheral subgroups, each of which has infinite index in G, and suppose that G acts properly and cocompactly on the* CAT(0) *cube complex* $\mathbf{X}$. *Then $\mathcal{I} \neq \emptyset$ and $\partial_\triangle \mathbf{X} \cong \mathcal{I} \cup \left(\bigsqcup_{P \in \mathbb{P}} \partial_\triangle \mathbf{Y}_P\right)$.*

Remark 3.2. Note that $\partial_\triangle \mathbf{Y}_P$ may be disconnected, and may contain simplices of $\mathcal{I}$.

Remark 3.3 (Metric relative hyperbolicity). Theorem 3.1 holds under more general conditions. Namely, if G acts properly and cocompactly on a CAT(0) cube complex $\mathbf{X}$ and there is a family $\{\mathbf{Y}_P\}$ of convex subcomplexes such that

$\mathbf{X} = \mathcal{N}_\tau(\bigcup_P \mathbf{Y}_P)$ for some $\tau \geq 0$, no distinct $\mathbf{Y}_P, \mathbf{Y}_Q$ have infinite coarse intersection, and the intersection graph of the τ-neighborhoods of the $\mathbf{Y}_P$ is fine and δ-hyperbolic for some $\delta \geq 0$, then $\partial_\triangle \mathbf{X}$ decomposes as in the conclusion of Theorem 3.1.

Remark 3.4 (limit simplices). If a is a limit simplex for the action of P on $\mathbf{X}$, then, fixing $y \in \mathbf{Y}$, we have a sequence $(p_j \in P)$ such that the set $\mathcal{A}$ of hyperplanes H that separates y from $p_j y$ for all but finitely many values of j represents a. Each such hyperplane separates two 0-cubes of the P-invariant subcomplex $\mathbf{Y}$, and thus crosses $\mathbf{Y}$. Hence $a \subseteq \partial_\triangle \mathbf{Y}$. Thus each $\partial_\triangle \mathbf{Y}_P$ contains every limit simplex for the action of P on $\mathbf{X}$. This verifies that each hypothesis in Theorem 3.7 below is necessary.

Proof of Theorem 3.1. That $\mathcal{I} \neq \emptyset$ follows from the rank-rigidity theorem [15, Corollary B] and the fact that the simplex represented by the boundary set consisting of hyperplanes that cross a sub-ray of an axis for a rank-one isometry is an isolated 0-simplex. Otherwise, $\mathbf{X}$ decomposes as the product of two unbounded subcomplexes and $\mathbb{P}$ consists of G itself.

We first show that, if $P, P' \in \mathbb{P}$ are distinct, then $\partial_\triangle \mathbf{Y}_P$ and $\partial_\triangle \mathbf{Y}_{P'}$ have disjoint images in $\partial_\triangle \mathbf{X}$. From this it follows that there is a simplicial embedding $\mathcal{I} \cup \left(\bigsqcup_{P \in \mathbb{P}} \partial_\triangle \mathbf{Y}_P\right) \hookrightarrow \partial_\triangle \mathbf{X}$.

Since $\mathbf{Y}_P \cap \mathbf{Y}_{P'}$ is the intersection of convex subcomplexes, it is convex and $P \cap P'$-cocompact, since $\mathbf{Y}_P$ and $\mathbf{Y}_{P'}$ are respectively P and P'-cocompact. Since $\mathbb{P}$ is almost-malnormal, $P \cap P'$ is finite, and $\mathbf{Y}_P \cap \mathbf{Y}_{P'}$ is therefore compact and, in particular, crossed by finitely many hyperplanes. The same is true of the intersection of any uniform neighborhoods of $\mathbf{Y}_P$ and $\mathbf{Y}_{P'}$. In particular, $\mathcal{H}(\mathbf{Y}_P) \cap \mathcal{H}(\mathbf{Y}_{P'})$ is finite, whence $\partial_\triangle \mathbf{Y}_P \cap \partial_\triangle \mathbf{Y}_{P'} = \emptyset$, as desired.

Consider a maximal simplex v of $\partial_\triangle \mathbf{X}$. If v is a 0-simplex, then it belongs to $\mathcal{I}$, so suppose that the dimension of v is positive. Let $\mathbf{O}$ be an orthant in $\mathbf{X}$ such that $\mathcal{H}(\mathbf{O})$ represents v. It suffices to verify that $\mathbf{O}$ is coarsely contained in some $\mathbf{Y}_P$, for it then follows that $v \subset \partial_\triangle \mathbf{Y}_P$ and the above embedding is surjective.

$\mathbf{O}$ is a maximal flat orthant, by maximality of v, and cannot have infinite coarse intersection with more than one $\mathbf{Y}_P$. Hence either $\mathbf{F}$ is coarsely contained in some $\mathbf{Y}_P$, or has finite intersection with each $\mathbf{Y}_P$. The latter case is impossible, since orthants are unconstricted, as shown in Section 4, and hence must lie near a peripheral subset by [17] and [4, Theorem 4.1, Remark 4.3]. Thus v belongs to a translate of some $\partial_\triangle \mathbf{Y}_P$, and the proof is complete. □

When the peripheral subgroups are virtually abelian, we obtain a cubical analogue of a result of Hruska and Kleiner [27, Theorem 1.2.1] which states that if X is a CAT(0) space admitting a proper, cocompact action by a group that is hyperbolic relative to maximal abelian subgroups, then the Tits boundary of X is isometric to the disjoint union of isolated points and spheres of various dimensions. This result of Hruska and Kleiner relates to the following:

Corollary 3.5. *Let G be hyperbolic relative to a collection $\mathbb{P}$ of virtually abelian subgroups of rank at least 2. Then for any* CAT(0) *cube complex* $\mathbf{X}$ *on which* G *acts properly and cocompactly,* $\partial_{\triangle}\mathbf{X}$ *is the disjoint union of a discrete set and a set of pairwise-disjoint spherical hyperoctahedra. If G is not virtually abelian, each of these sets is infinite.*

Proof. By Theorem 3.1, $\partial_{\triangle}\mathbf{X} \cong \mathcal{I} \sqcup \left(\bigsqcup_P \partial_{\triangle}\mathbf{Y}_P\right)$. The set of isolated 0-cubes, and the set of $\partial_{\triangle}\mathbf{Y}_P$, are obviously infinite if G is not virtually abelian. For each maximal virtually abelian subgroup P, we have $\partial_{\triangle}\mathbf{Y}_P \cong \mathbb{O}_d$, where $d \geq 2$ is the rank of P, by [24, Theorem A]. If $\partial_{\triangle}\mathbf{Y}_{P'}$ and $g\partial_{\triangle}\mathbf{Y}_P$ have nonempty intersection, containing a common simplex v, then $g\mathbf{Y}_P \cap \mathbf{Y}_{P'}$ is coarsely unbounded, since it is crossed by every hyperplane in a boundary set representing v. But then $gPg^{-1} \cap P'$ is infinite, contradicting almost-malnormality unless $gPg^{-1} = P'$. In the latter case, $g\partial_{\triangle}\mathbf{Y}_P = \partial_{\triangle}\mathbf{Y}_{P'}$. (If G is virtually abelian, then the above argument shows that $\partial_{\triangle}\mathbf{X}$ is a single hyperoctahedron.) □

Since each hyperoctahedron can be given a CAT(1) metric, in which simplices are spherical simplices with side length $\frac{\pi}{2}$, making it isometric to a sphere of the appropriate dimension (see Section 3 of [26]), Corollary 3.5 provides a new proof of the Hruska-Kleiner result in the CAT(0) cubical case.

3.2. Peripheral structures from collections of subcomplexes of $\partial_{\triangle}\mathbf{X}$. Conversely, one can recover a relatively hyperbolic structure on G from a decomposition of $\partial_{\triangle}\mathbf{X}$ like that in Theorem 3.1. Suppose G acts properly and cocompactly on the CAT(0) cube complex $\mathbf{X}$ and, as before, denote by $\mathcal{I}$ the set of isolated 0-simplices of $\partial_{\triangle}\mathbf{X}$.

Definition 3.6 (fine graph). The graph Λ is *fine* if for all $n \in \mathbb{N}$ and all edges e of Λ, there are finitely many n-cycles in Λ that contain e.

Theorem 3.7. *For some $k < \infty$, let $\mathbf{S}_1, \ldots, \mathbf{S}_k$ be subcomplexes of $\partial_{\triangle}\mathbf{X}$, with $P_i = \mathrm{Stab}(\mathbf{S}_i)$, and satisfying all of the following*:

(1) $\partial_{\triangle}\mathbf{X} = \mathcal{I}' \sqcup G\left(\bigsqcup_{i=1}^k \mathbf{S}_i\right)$*, where* $\mathcal{I}' \subseteq \mathcal{I}$;
(2) *for each i, the subcomplex $\mathbf{S}_i$ contains all limit simplices for the action of P_i on $\partial_{\triangle}\mathbf{X}$. Equivalently, when $\mathbf{X}$ is fully visible, each $f^{-1}(\mathbf{S}_i)$ contains the limit set of P_i*;
(3) *for all $1 \leq i \leq j \leq k$ and $g, h \in G$, we have $g\mathbf{S}_i \cap h\mathbf{S}_j = \emptyset$ unless $i = j$ and $gh^{-1} \in P_i$*;
(4) *either $k = 1$ and P_1 is a finite index subgroup of G, or each P_i has infinite index in G*;
(5) *each P_i is quasiconvex.*

Then G is hyperbolic relative to a collection $\{Q_i\}_{i=1}^k$ for which Q_i is commensurable with P_i for each $i \leq k$.

Proof. First, we assume that each $\mathbf{S}_i$ contains at least one positive-dimensional simplex, for otherwise the hypotheses are satisfied by a proper subset of $\{\mathbf{S}_i\}_{i=1}^k$. If the set of $\mathbf{S}_i$ is empty, then $\partial_\triangle \mathbf{X}$ consists entirely of isolated 0-simplices whence G is hyperbolic relative to $\{1\}$ by Proposition 2.15.

In this proof, we use the metric $\dot{d}$ unless stated otherwise. Observe also that the hypotheses imply that each positive-dimensional component of $\partial_\triangle \mathbf{X}$ is contained in a single $g\mathbf{S}_i$.

Representing P_i in $\mathbf{X}$. Fix a 0-cube $x \in \mathbf{X}$. For $1 \leq i \leq k$, let $\mathbf{C}_i$ be the convex hull of the orbit $P_i x$. The subcomplex $\mathbf{C}_i$ is P_i-invariant because $\mathbf{C}_i$ is the largest subcomplex contained in the intersection of all halfspaces that contain $P_i x$, the set of which is obviously P_i-invariant. Thus $P_i \leq \mathrm{Stab}_G(\mathbf{C}_i)$. Each P_i is quasiconvex in G with respect to the action of G on $\mathbf{X}^{(1)}$. Hence the subcomplex $\mathbf{C}_i$ is contained in a uniform neighborhood of the orbit $P_i x$ and is therefore P_i-cocompact. Let $Q_i = \mathrm{Stab}_G(\mathbf{C}_i)$. Since $\mathbf{C}_i$ is contained in a finite neighbourhood of $P_i x$, the groups P_i and Q_i are commensurable.

Comparing $\partial_\triangle \mathbf{C}_i$, $\mathbf{S}_i$, and verifying almost-malnormality. The inclusion $\mathbf{C}_i \to \mathbf{X}$ induces an inclusion $\partial_\triangle \mathbf{C}_i \hookrightarrow \partial_\triangle \mathbf{X}$ whose image is a subcomplex. Now, suppose that $a \subseteq \partial_\triangle \mathbf{C}_i$ is a maximal, and therefore visible, simplex, and let $\gamma \to \mathbf{C}$ be a combinatorial geodesic ray such that $\mathcal{H}(\gamma)$ represents a. Since P_i acts cocompactly on $\mathbf{C}_i$, there exists a sequence $\{p_j \in P_i\}$ such that γ lies at finite Hausdorff distance from $\{p_j x\}$, and therefore that the set of hyperplanes H such that H separates x from $p_j x$ for all but finitely many values of j has finite symmetric difference with $\mathcal{H}(\gamma)$. Hence a is a limit simplex for the action of P_i on $\mathbf{X}$.

Under the hypothesis that each $\mathbf{S}_i$ contains every limit simplex for the action of its stabilizer P_i, this shows that $\partial_\triangle \mathbf{C}_i \subseteq \mathbf{S}_i$. Similarly, under the hypothesis that $f^{-1}(\mathbf{S}_i)$ contains the limit set for the action of P_i, this implies that $\partial_\triangle \mathbf{C}_i \subseteq \mathbf{S}_i$. Hence, if $g, h \in G$, then $g\partial_\triangle \mathbf{C}_i \cap h\partial_\triangle \mathbf{C}_j = \emptyset$ unless $i = j$ and $gh^{-1} \in P_i$. This implies that the set of hyperplanes crossing $g\mathbf{C}_i$ and $h\mathbf{C}_j$ is finite, whence, for any $R \geq 0$, the intersection of the R-neighborhood of $g\mathbf{C}_i$ with that of $h\mathbf{C}_j$ is compact.

Let $i, j \leq k$ and $h \in G$, and consider $P_i^h \cap P_j$. If this intersection is infinite, then $\mathbf{C}_j \cap (h\mathbf{C}_i)$ contain unbounded subsets at finite Hausdorff distance, a contradiction. Thus $\{P_i\}_{i=1}^k$ is an almost-malnormal collection, and the same is true of $\{Q_i\}$.

A Bowditch graph. For any $R \in \mathbb{N}$, and any convex subcomplex $Y \subset \mathbf{X}$, let $\mathfrak{K}_R(Y)$ be the following convex subcomplex containing Y with the property that every $x \in \mathfrak{K}_R(Y)$ satisfies $\dot{d}(x, Y) \leq R$. Let $t_R = \frac{R}{\dim \mathbf{X}}$ and let $\mathfrak{K}_R(Y)$ be the convex hull of the $\dot{d}$-neighborhood of Y of radius t_R. Then $Y \subseteq \mathfrak{K}_R(Y)$, the latter subcomplex is convex and contained in the uniform R-neighborhood of Y

as we now quickly show. Any geodesic joining $y \in \mathfrak{K}_R(Y)$ to a closest point of Y crosses a set of hyperplanes that cross the t_R-neighborhood of Y but do not cross Y. Further, this set of hyperplanes contains no facing triple, and each clique has cardinality at most $\dim \mathbf{X}$. Thus, there are at most $\dim \mathbf{X} t_r = R$ hyperplanes in the set, since otherwise we would have a contradiction as we would obtain a nested set of more than t_R hyperplanes separating y from Y and crossing $\mathcal{N}_{t_R}(Y)$.

Since G acts cocompactly, there exists $R < \infty$ such that $\bigcup_i G\mathfrak{K}_R(\mathbf{C}_i) = \mathbf{X}$. Fixing such an R, let Γ be the intersection graph of the collection of subspaces $\mathfrak{K}_R(\mathbf{C}_i)$ and all of their translates. More precisely, Γ has a vertex for each $\mathfrak{K}_R(g\mathbf{C}_i)$ and exactly one edge joining $\mathfrak{K}_R(g\mathbf{C}_i)$ to $\mathfrak{K}_R(h\mathbf{C}_j)$ if and only if $g\mathbf{C}_i \neq h\mathbf{C}_j$ and $\mathfrak{K}_R(g\mathbf{C}_i) \cap \mathfrak{K}_R(h\mathbf{C}_j) \neq \emptyset$.

Since $\mathbf{S}_i \cap \mathbf{S}_j = \emptyset$ for $i \neq j$, and $\mathbf{X}$ is locally finite, $\mathbf{C}_i \cap \mathbf{C}_j$ is compact, and in particular is crossed by finitely many hyperplanes. More strongly, the set of hyperplanes that crosses both $\mathbf{C}_i$ and $\mathbf{C}_j$ is finite, since otherwise $\mathcal{H}(\mathbf{C}_i) \cap \mathcal{H}(\mathbf{C}_j)$ would contain a boundary set. Hence finitely many hyperplanes cross $\mathfrak{K}_R(\mathbf{C}_i) \cap \mathfrak{K}_R(\mathbf{C}_j)$, and therefore there exists a compact convex subcomplex B such that for all $g, h \in G$, $1 \leq i, j \leq k$ there exists $a \in G$ such that $\mathfrak{K}_R(g\mathbf{C}_i) \cap \mathfrak{K}_R(h\mathbf{C}_j) \subset aB$.

By construction, G acts by isometries on Γ, in such a way that the set of vertex stabilizers is exactly the set of subgroups Q_i and their conjugates.

Edge-stabilizers. Almost-malnormality of $\{Q_i\}_i$ implies that the stabilizers of edges in Γ are finite.

Cofiniteness. There are finitely many G-orbits of edges in Γ. To see this, first observe that each P_i acts cocompactly on $\mathfrak{K}_R(\mathbf{C}_i)$. Also, there are clearly finitely many G-orbits of vertices in Γ: one for each $\mathbf{C}_i$ with $1 \leq i \leq k$.

For each vertex $\mathfrak{v}$ of Γ (corresponding to some translate of some $\mathfrak{K}_R(\mathbf{C}_i)$), let $E(\mathfrak{v}) = \{e_1, \dots, e_q\}$ be a set of edges of Γ incident to $\mathfrak{v}$, containing exactly one edge from each $\mathrm{Stab}_G(\mathfrak{v})$-orbit. This set is finite since $\mathrm{Stab}(\mathfrak{v})$ acts cocompactly on $\mathbf{C}_i$. Let $\{\mathfrak{v}_1, \dots, \mathfrak{v}_k\}$ contain exactly one vertex of Γ from each G-orbit. If $\mathfrak{v}$ is a vertex and e an incident edge, then $(\mathfrak{v}, e) = (g\mathfrak{v}_i, gpg^{-1}e_j)$, where $g^{-1}e_j \in E(\mathfrak{v}_i)$, and $g \in G$, and $p \in \mathrm{Stab}_G(\mathfrak{v})$. Thus $(\mathfrak{v}, e) = g(\mathfrak{v}_i, pg^{-1}e_j)$ is a translate of one of the finitely many pairs $(\mathfrak{v}_i, e_j)$. Hence there are finitely many G-orbits of edges in Γ.

Conclusion. Below we prove Γ is fine in Lemma 3.8 and hyperbolic in Lemma 3.9. Accordingly, the action of G on Γ satisfies all of the conditions of [8, Definition 2] and G is therefore hyperbolic relative to $\{Q_i\}_{i=1}^k$. □

Lemma 3.8. Γ *is fine.*

Proof. Since Γ contains no loops or bigons, every cycle has length at least 3.

3-CYCLES. Let $A_0 = \mathfrak{K}_R(g\mathbf{C}_i)$ and $A_1 = \mathfrak{K}_R(h\mathbf{C}_j)$ with $A_0 \cap A_1 \neq \emptyset$. Let e be the edge of Γ joining the vertices corresponding to A_0 and A_1. If A_2 is a subcomplex corresponding to some other vertex of Γ, and $A_0 \cap A_2 \neq \emptyset$ and $A_1 \cap A_2 \neq \emptyset$, then $A_0 \cap A_1 \cap A_2 \neq \emptyset$, since each A_i is convex and CAT(0) cube complexes have the Helly property. Now, $A_0 \cap A_1$ is compact, and thus contained in some translate aB of B. Hence, for each A_2 that intersects A_0 and A_1, the mutual intersection $A_0 \cap A_1 \cap A_2$ lies in aB. In particular, A_2 intersects aB. Hence, by cocompactness, there are only finitely many A_2 such that the vertices in Γ corresponding to A_0, A_1, A_2 form a 3-cycle.

4-CYCLES. As before, let $\{A_0, A_1\}$ be an edge of Γ. Let A_0', A_1' be vertices of Γ (we use the same notation for the corresponding subcomplexes of $\mathbf{X}$) such that $\{A_i, A_i'\}$ is an edge of Γ for $i \in \{0, 1\}$ and $\{A_0', A_1'\}$ is an edge of Γ.

Choose combinatorial geodesic paths $\rho_0, \rho_0', \rho_1, \rho_1'$ such that $\rho_i \to A_i$ and $\rho_i' \to A_i'$ for $i \in \{0, 1\}$ and $\rho_0\rho_1\rho_1'\rho_0'$ is a closed path in $\mathbf{X}$. Let $D \to \mathbf{X}$ be a disc diagram in $\mathbf{X}$ bounded by $\rho_0\rho_1\rho_1'\rho_0'$, as in Figure 1. Assume that D has minimal area among all diagrams with that boundary path, and, moreover, suppose that the ρ_i and ρ_i' are chosen among geodesic paths in the required A_i, A_i' in such a way that the resulting disc diagram D is as small as possible, in the following sense: $(\mathsf{Area}(D), |\partial_p D|)$ is as small as possible, where such pairs are taken in lexicographic order.

Suppose, for the moment, that $|\rho_i|, |\rho_i'| > 0$ for each i, so that D contains a dual curve emanating from each of the four named subpaths of its boundary path. If the dual curve K emanates from ρ_1, then K cannot end on ρ_1, since that path is a geodesic. Also, if K_1, K_2 are two dual curves emanating from ρ_1, then they cannot cross, for otherwise, by convexity of A_1, we could modify ρ_1 by finding a corner of a square of A_1 in the subdiagram bounded by A_1 and the arrowed path indicated in Figure 1, leading to a lower-area diagram. If K is a leftmost (or rightmost) dual curve emanating from ρ_1 and ending on ρ_1' (or ρ_0', if K is rightmost), then any dual curve emanating from the part of ρ_1 subtended by ρ_1' and K (respectively, ρ_0' and K) must cross K, and this is impossible. Hence K is dual to the terminal (respectively, initial) 1-cube c of ρ_1 and, by performing a series of *hexagon moves* (see [44, Section 2]), we find that ρ_1 and ρ_1' (respectively, ρ_1 and ρ_0') have a common 1-cube, namely c. We can thus remove c from ρ_1, ρ_1', resulting in a new diagram with the required properties, the same area as D, and strictly shorter boundary path. Since this is a contradiction, we conclude that every dual curve travels from ρ_1 to ρ_0' or from ρ_1' to ρ_0. Let $\mathcal{V}$ be the set of hyperplanes corresponding to dual curves of the former type, and $\mathcal{W}$ the set of hyperplanes corresponding to dual curves of the latter type. (Using this fact, the fact that geodesic segments cross each hyperplane at most once, and the fact that hyperplanes do not self-cross, it is easy to see that distinct dual curves in D map to distinct hyperplanes.)

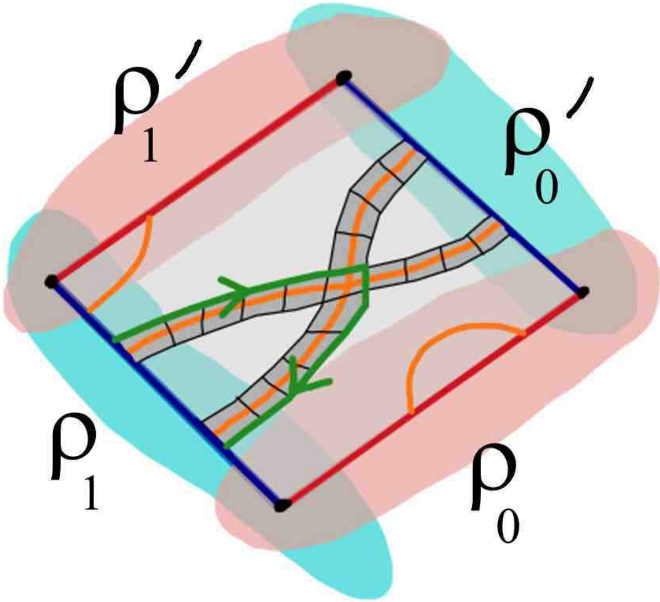

Figure 1. Some illegal dual curves, and an illegal crossing, in D.

This argument shows that $|\rho_1| = |\rho'_0|$ and $|\rho_0| = |\rho'_1|$. If $|\rho_1| = 0$, then A_0, A_1, A'_1 pairwise-intersect, and hence A'_1 is one of finitely many vertices of Γ that can be the third vertex in a 3-cycle containing the edge $\{A_0, A_1\}$. But A'_1, A'_0, A_0 form a 3-cycle in Γ, and thus there are only finitely many possible A'_0. In other words, if Γ contains infinitely many 4-cycles containing the edge $\{A_0, A_1\}$, then all but finitely many of these 4-cycles lead to disc diagrams with $|\rho_1| = 0$. An identical argument works for ρ_0, and hence $\mathcal{V}$ and $\mathcal{W}$ are nonempty for all but finitely many 4-cycles containing $\{A_0, A_1\}$.

Hence suppose that for all $m \geq 0$, there exist vertices $A'_0 = A'_0(m), A'_1 = A'_1(m)$ of Γ such that $A_0, A_1, A'_1, A'_0, A_0$ is a 4-cycle in Γ, and suppose that for all m, the sets $\mathcal{V}(m), \mathcal{W}(m)$ defined above are nonempty. Note that $\mathcal{V}(m) \subseteq \mathcal{H}(A_1) \cap \mathcal{H}(A'_0(m))$ and $\mathcal{W}(m) \subseteq \mathcal{H}(A_0) \cap \mathcal{H}(A'_1(m))$. Moreover, if $V \in \mathcal{V}(m)$ and $W \in \mathcal{W}(m)$, then V and W cross, since their corresponding dual curves in the associated disc diagram cross.

Next, we show that there exists $\xi < \infty$, depending only on R, such that $\max\{|\mathcal{V}(m)|, |\mathcal{W}(m)|\} \leq \xi$ for all m. $\mathcal{W}(m)$ is a set of hyperplanes H that cross both A_1 and $A'_0(m)$. If it were possible to choose $A'_0(m)$ in such a way as to make $H(A_1) \cap \mathcal{H}(A'_0(m))$ have arbitrarily large cardinality, then since $\mathrm{Stab}_G(A_1)$ acts cocompactly on A_1, there would exist some $A'_0(m)$ with $H(A_1) \cap \mathcal{H}(A'_0(m))$ infinite, contradicting the fact that distinct translates of the various $\mathbf{C}_i$ have disjoint simplicial boundaries.

By cocompactness of the action of $\mathrm{Stab}_G(A_0)$, we can assume that $\rho_0(m) \cap \rho_1(m)$ lies in a fixed compact set in A_0, of diameter $d < \infty$, and hence each $A_0'(m)$ and $A_1'(m)$ come within $d + \xi$ of $\rho_0(1) \cap \rho_1(1)$. There can only be finitely many such $A_0'(m)$ or $A_1'(m)$, and we conclude that each edge of Γ is contained in at most finitely many distinct 3-cycles or 4-cycles. (Alternatively, we see that $|\mathcal{V}(m)|$ and $|\mathcal{W}(m)|$ must both be unbounded as $m \to \infty$, and deduce that there exist infinite sets $\mathcal{V}_\infty \subset \mathcal{H}(A_1)$ and $\mathcal{W}_\infty \subset \mathcal{H}(A_0)$, with each $V \in \mathcal{V}$ crossing each $W \in \mathcal{W}$. Thus $\partial_\triangle \mathbf{X}$ contains a 1-simplex joining a 0-simplex of $\partial_\triangle A_1 = g\mathbf{S}_i$ to a 0-simplex of $\partial_\triangle A_0 = h\mathbf{S}_j$, and this is impossible.)

Large cycles. Let $p \geq 4$. Let A_0, A_1 be a pair of vertices of Γ connected by an edge. Let A_2, A_p be distinct vertices which are disjoint from A_0, A_1, and such that $\{A_1, A_2\}$ and $\{A_p, A_0\}$ are edges of Γ. Let σ be an embedded path of length at least 1 in Γ joining A_2 to A_p and not containing A_0 or A_1; for $2 \leq i \leq p$, let A_i denote the subcomplex corresponding to the $(i-1)^{th}$ vertex of σ. For each $0 \leq i \leq p$, let $\rho_i \to A_i$ be a combinatorial geodesic path such that $\rho_0 \dots \rho_p$ is a closed path in $\mathbf{X}$, bounding a disc diagram D that is minimal in the same sense as above (the details are identical to the 4-cycle case). Then every dual curve in D travels from some ρ_i to some ρ_j with $i \neq j$. For $0 \leq \ell \leq p$, let $\mathcal{V}_\ell$ be the set of distinct hyperplanes corresponding to dual curves emanating from ρ_ℓ. For each ℓ, there exists ℓ' such that $|\mathcal{V}_\ell \cap \mathcal{V}_{\ell'}| \geq \frac{|\mathcal{V}_\ell|}{p-2}$, since there are p possible destinations for each of the dual curves emanating from ρ_ℓ (minimality of D implies that such a dual curve cannot end on $\rho_{\ell\pm1}$). Now since $\mathcal{V}_\ell \subset \mathcal{H}(A_\ell)$ and $\mathcal{V}_{\ell'} \subset \mathcal{H}(A_{\ell'})$, we have $|\rho_\ell| \leq (p-2)\xi$ for all ℓ. As above, this implies that there are only finitely many paths ρ in Γ that combine with $\{A_0, A_1\}$ to make a $(p+1)$-cycle. Thus Γ is fine. □

Lemma 3.9. *There exists $\delta \in [0, \infty)$ such that Γ is δ-hyperbolic.*

Proof. We will verify that the G-cocompact graph Γ has thin triangles.

Superconvexity. The arguments supporting fineness work for any sufficiently large finite R. In particular, we first show that we can choose R large enough that $\mathfrak{K}_R(\mathbf{C}_i)$ is *superconvex* for $1 \leq i \leq k$, i.e., for any bi-infinite (combinatorial or CAT(0)) geodesic γ in $\mathbf{X}$, either $\gamma \subset \mathfrak{K}_R(\mathbf{C}_i)$, or $\gamma \cap \mathfrak{K}_r(\mathfrak{K}_R(\mathbf{C}_i))$ is bounded for all $r \geq 0$. By cocompactness, for all $r \geq 0$, there exists $m_r < \infty$ such that $\mathrm{diam}(\gamma \cap \mathfrak{K}_{R+r}(\mathbf{C}_i)) \leq m_r$ for any bi-infinite geodesic γ not contained in $\mathfrak{K}_R(\mathbf{C}_i)$.

To make this choice, suppose that for all $R \geq 0$, there exists a (CAT(0) or combinatorial) geodesic ray σ_R lying in $\mathfrak{K}_R(\mathbf{C}_i)$, with every point of σ_R at distance at least $R-1$ from $\mathbf{C}_i$. Applying cocompactness and a standard disc diagram argument shows that, in this situation, there is a boundary set $\mathcal{U} \subset \mathcal{H}(\mathbf{C}_i)$, representing a simplex u of $\mathbf{S}_i$, and a boundary set $\mathcal{V} \subset \mathcal{H} - \mathcal{H}(\mathbf{C}_i)$ representing a simplex v that is adjacent in $\partial_\triangle \mathbf{X}$ to u. But $v \not\subset \mathbf{S}_i$, since every simplex of $\mathbf{S}_i$

is represented by a boundary set consisting of hyperplanes crossing $\mathbf{C}_i$. Hence v lies in some $\mathbf{S}_j$ that differs from and intersects $\mathbf{S}_i$, a contradiction.

Non-peripheral rectangular discs. Convexity and superconvexity of $\mathfrak{K}_R(\mathbf{C}_i)$ together imply that any isometric flat $\mathbf{F} \subset \mathbf{X}$ lies entirely inside some $\mathfrak{K}_R(g\mathbf{C}_i)$. Cocompactness then implies that there exists N such that if $D \to \mathbf{X}$ is a combinatorial isometric embedding of the CAT(0) cube complex $[0,m]^2$, then either $m < N$ or the image of D is contained in exactly one $\mathfrak{K}_R(g\mathbf{C}_i)$.

Non-peripheral strips. N and R can be chosen so that if there exists a subspace $\mathfrak{K}_R(\mathbf{C})$ corresponding to a vertex of Γ and an isometrically embedded rectangle $S \cong [0,a] \times [0,b] \subset \mathbf{X}$ with $[0,a] \times \{0\} \subset \mathfrak{K}_R(\mathbf{C})$ and $a \geq N$, then $S \subset \mathfrak{K}_R(\mathbf{C})$. This follows from superconvexity of $\mathfrak{K}_R(\mathbf{C})$ and cocompactness of the action of its stabilizer.

Representing geodesics in Γ. Let $\gamma : [0,T] \to \Gamma$ be a geodesic segment. For $0 \leq i \leq T$, let $A_i = \gamma(i)$ be the i^{th} vertex. We also denote by A_i the corresponding subcomplex $\mathfrak{K}_R(\mathbf{C})$ of $\mathbf{X}$. A combinatorial piecewise-geodesic ρ is said to *represent* the geodesic γ in Γ if $\rho = \rho_0\rho_1 \dots \rho_{T-1}$, where ρ_i is a combinatorial geodesic of A_i for $0 \leq i \leq T-1$.

Properties of projection to Γ. The remainder of the proof requires establishing three claims. We note that there is a map $\mathbf{X} \to \Gamma$ sending each point to the vertex corresponding to the vertex corresponding to some $\mathfrak{K}_R(\mathbf{C})$ containing it. (There are many choices of such a map and we choose one arbitrarily. Although we don't use this fact, these maps are coarsely the same, since any point of $\mathbf{X}$ lies in a uniformly bounded number of subcomplexes $\mathfrak{K}_R(\mathbf{C})$.) Below, we discuss images of paths under this map. We note that these images need not be paths, but nevertheless are geometrically well-behaved in the following ways.

Claim 1. Let $\sigma'\sigma^{-1}$ be a geodesic bigon in $\mathbf{X}$. Then there exists δ' such that the image of σ is contained in the δ'-neighborhood of the image of σ' and vice versa.

Proof. Let $D \to \mathbf{X}$ be a minimal-area disc diagram with boundary path $\sigma'\sigma^{-1}$. Since σ, σ' are geodesics, every dual curve in D starts on σ and ends on σ'. Choose $x \in \sigma$ and $x' \in \sigma'$. Let $\mathcal{L}$ be the set of dual curves starting on σ to the left of x and ending on σ' to the right of x', and let $\mathcal{R}$ be the set of dual curves starting on σ to the right of x and ending on σ' to the left of x'. Then every dual curve in D separating x, x' belongs to one of these sets, whence

$$\dot{d}(x,x') \leq |\mathcal{L}| + |\mathcal{R}|.$$

If either of $\mathcal{L}$ or $\mathcal{R}$ has cardinality at most N, then x lies at distance at most N from σ and x' lies at distance at most ϵN from σ'. On the other hand, since each dual curve in $\mathcal{L}$ crosses each dual curve in $\mathcal{R}$, if $|\mathcal{L}|, |\mathcal{R}| \geq N$, then $\mathbf{X}$ contains an isometric flat rectangle F, each of whose sides has length at least N, containing x, x'. The rectangle F is contained in some subcomplex $\mathbf{C}$ corresponding to a vertex of Γ, whence the images of x, x' can be joined by a path of length 2 in Γ whose middle vertex is $\mathbf{C}$. Hence the image of σ is contained in the δ'-neighborhood of the image of σ', and vice versa, for δ' depending only on N. △

Claim 2. Let $\gamma\gamma'\gamma''$ be a geodesic triangle in $\mathbf{X}$. There exists δ such that the image of any of $\gamma, \gamma', \gamma''$ in Γ lies in the δ-neighborhood of the union of the images of the other two paths.

Proof. This follows from the fact that $\mathbf{X}^{(0)}$, endowed with the metric $\dot{d}$, is a median space, together with Claim 1. Indeed, let $\gamma\gamma'\gamma''$ be a geodesic triangle in $\mathbf{X}^{(1)}$. Then there is a combinatorial geodesic triangle $\alpha\alpha'\alpha''$ such that $\alpha\gamma^{-1}, \alpha'(\gamma')^{-1}, \alpha''(\gamma'')^{-1}$ are geodesic bigons and each of $\alpha, \alpha', \alpha''$ is contained in the union of the other two (each passes through the median of the three endpoints of $\gamma \cup \gamma' \cup \gamma''$). Hence, by Claim 1, the image of each of $\gamma, \gamma', \gamma''$ in Γ lies in the $\delta = 2\delta'$-neighborhood of the union of the other two. △

Claim 3. There exists $\mathfrak{L}$, independent of γ, such that a representative ρ can be chosen so that its image in Γ is contained in the $\mathfrak{L}$-neighborhood of the image of a geodesic σ of $\mathbf{X}$.

Proof. There are several steps.

STRATEGY. Suppose that γ has a representative ρ, so that $\rho = \rho_0\rho_1 \cdots \rho_{T-1}$ is a piecewise-geodesic with $\rho_j \to A_j$ for $0 \leq j \leq T-1$ that joins $x_0 \in A_0$ to $x_T \in A_{T-1} \cap A_T$. Let σ_0 be a geodesic joining x_0 to x_T. Let $D \to \mathbf{X}$ be a minimal-area disc diagram bounded by $\rho\sigma_0^{-1}$. Convexity of the A_j implies that no dual curve starts on ρ_j and ends on $\rho_{j\pm 1}$, for otherwise we could remove backtracks from the boundary path of D. Similarly, no two dual curves emanating from a common ρ_j can cross, for otherwise convexity of A_j would enable us to modify ρ_j, without changing its endpoints, to obtain a lower-area diagram.

If no dual curve in D has both ends on ρ, then ρ is a geodesic and the claim holds by setting $\sigma = \rho$. Hence, we suppose that K is a dual curve in D that is *outermost* in the sense that K is dual to two distinct 1-cubes on ρ, and the subpath of ρ subtended by these 1-cubes is not properly contained in a subpath subtended by two distinct 1-cubes dual to the same dual curve. If the image of K under the map $D \to \mathbf{X} \to \Gamma$ is at uniformly bounded Hausdorff distance from the image of ρ, then we can replace the part of ρ between and including the 1-cubes dual to K by a path in the carrier of K, yielding a new path ρ', whose image is at uniformly bounded Hausdorff distance from that of ρ, but which has strictly fewer pairs of

1-cubes dual to a common hyperplane. Finitely many repetitions of this procedure then yields the desired σ. Hence it suffices to find $\mathfrak{L}$ such that the $\mathfrak{L}$-neighborhood of the image of K in Γ contains the image of ρ.

THE SUBDIAGRAM D'. To this end, suppose that K starts on ρ_j and ends on $\rho_{j'}$, with $|j - j'| > 1$. Let P be a shortest path in $N(K) \subset D$ starting at $N(K) \cap \rho_j$ and ending at $N(K) \cap \rho_{j'}$, with P separated from the subtended part of ρ by K. Let ρ' be the subtended part of ρ, so that $\rho' = \rho'_j \rho_{j+1} \cdots \rho'_{j'}$, where $\rho'_j, \rho'_{j'}$ are respectively subpaths of $\rho_j, \rho_{j'}$. Let $D' \to \mathbf{X}$ be the subdiagram of D bounded by P and ρ'. As before, no dual curve travels from ρ'_j to ρ_{j+1}, or ρ_k to $\rho_{k\pm1}$ for $j-1 \leq k \leq j'+1$, or from $\rho_{j'-1}$ to $\rho'_{j'}$, and no two dual curves emanating from the same named subpath of P cross. Every dual curve emanating from P ends on ρ', since D has minimal area for its boundary path and therefore contains no bigon of dual curves (see e.g. [42, 44]). Note that the images of ρ'_j and $\rho'_{j'}$ in Γ are at distance at most 1 from the images of $A_j, A_{j'}$ and hence at distance at most 2 from the image of P.

THE DIAGRAMS D'_k. For $j+1 \leq k \leq j'-1$, we inductively define combinatorial paths a_k, b_k starting on ρ_k and ending on P as follows. Let a_{j+1} be a shortest path in D' joining a point of ρ_{j+1} to a point of P. Let b_{j+1} be of minimal length among all paths in D' joining a point of ρ_{j+2} to a point of P and not crossing a_{j+1} (these paths are allowed to coincide for some or all of their lengths). Given a_k joining ρ_k to P, let b_k be a minimal path joining ρ_{k+1} to P that does not cross a_k, and given b_k, let a_{k+1} be a minimal path joining ρ_{k+1} to P that does not cross b_k. See Figure 2.

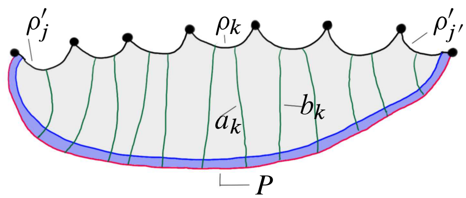

Figure 2. The diagram D'.

For each k, let P_k be the subpath of P between the endpoints of a_k and b_k. Let c_k be the subpath of ρ_k between the initial point of a_k and the terminal point of ρ_k, and let d_k be the part of ρ_{k+1} from the initial point of ρ_{k+1} to the initial point of b_k; these paths are shown in Figure 2. Consider the subdiagram D'_k bounded by a_k, P_k, b_k, d_k, and c_k. Every dual curve in D'_k emanating from P_k ends on c_k or d_k, and no two such dual curves cross. Indeed, if such a dual curve C ended on a_k (or b_k), then we could have chosen a_k (or b_k) to be shorter, as shown in Figure 3 at left. Similarly, no dual curve travels from a_k to c_k or b_k to d_k.

We conclude that D'_k is the union of two (possibly degenerate) flat rectangles, T_k, U_k and a subdiagram V_k shown at right in Figure 3. The subdiagram V_k is formed by the crossing of the dual curves emanating from P_k with the dual curves traveling from a_k to b_k. The rectangle T_k is formed from the dual curves traveling from a_k to d_k crossing those that emanate from c_k. The rectangle U_k is formed analogously. Now, if $|c_k| \geq N$, then the strip T_k actually lies in A_k and we could have chosen ρ_k to yield a lower-area diagram D. Hence $|c_k| < N$ and $|d_k| < N$. Thus $|P_k| < N$, and there is a path of length less than $2N$ joining $\rho_k \cap \rho_{k+1}$ to V_k. It follows that if, for any $\epsilon \geq 0$, at most ϵN dual curves travel from a_k to b_k, then $\dot{d}(\rho_k \cap \rho_{k+1}, P) \leq (2+\epsilon)N$. The images of ρ_k and ρ_{k+1} in Γ thus lie in the $[(2+\epsilon)N + 1]$-neighborhood of the image of P.

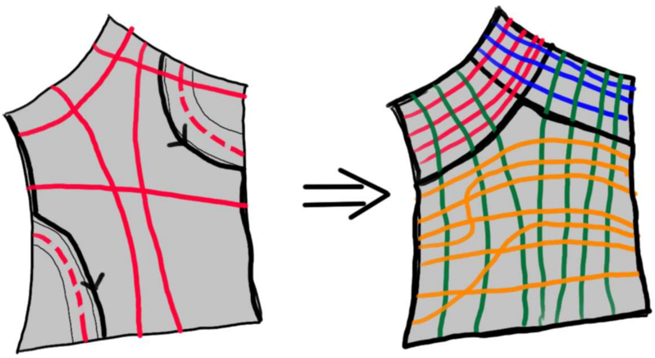

Figure 3. Left: the solid dual curves shown in D'_k are all possible. If either dotted dual curve occurs, then as shown, either a_k or b_k could be shortened (there are two other similar possibilities not shown). This leads to the conclusion at left: the rectangles T_k, U_k intersect in the smaller rectangle at the top of D'_k, each of whose sides has length less than N, and the remainder of the diagram is V_k.

The diagrams E_k. For each k, let Q_k be the subpath of P between the endpoints of b_k and a_{k+1} and let e_k be the subpath of ρ_{k+1} between the initial point of b_k and the initial point of a_{k+1}. The subdiagram E_k bounded by e_k, a_{k+1}, Q_k, and b_k has the property that all dual curves travel from b_k to a_{k+1} (by the minimality of those paths) or from Q_k to e_k. If there are at least N dual curves from b_k to a_{k+1}, then the convex hull of the image of E_k in $\mathbf{X}$ contains an $N \times N$ flat grid. Since this image of E_k contains an $N \times N$ flat grid, as we proved above in the paragraph on non-peripheral rectangular discs, we then have E_k contained in some A_i and thus the distance in Γ between the images of ρ_k and P is at most 3. Thus $|\rho_k| \leq 3N$ for all k for which the distance between some point in the image of ρ_k and the image of P is at least 4.

Choose k_1, k_2 with $j \leq k_1 \leq k_2 \leq j'$ such that for all $k \in \{k_1, \ldots, k_2\}$, the diagram D'_k has more than ϵN dual curves traveling from a_k to b_k, and for $k \in \{k_1, \ldots, k_2 - 1\}$, the diagram E_k has more than ϵN such dual curves, and the distance from some point of each ρ_k to P in Γ is at least 4, and such that $k_2 - k_1$ is as large as possible.

Define the subdiagram E of D' to consist of the union of the D'_k, for $k_1 \leq k \leq k_2$, together with E_k for $j \leq k \leq j'$. Let $\mathcal{V}$ be the set of *vertical* dual curves, i.e., those that have an end on some P_k or Q_k. By the above discussion, at most $3N$ vertical dual curves end on each ρ_k. Observe that there is a path of length $2\epsilon N + |\mathcal{V}|$ joining ρ_{k_1} to ρ_{k_2}, and hence a path of length at most $2\epsilon N + |\mathcal{V}| + 2$ in Γ joining A_{k_1} to A_{k_2}. Hence $|\mathcal{V}| \geq k_2 - k_1 - 2\epsilon N - 1$. If $k_2 - k_1 \leq 2(2\epsilon N + 1)$, then we have a uniform bound of $2(2\epsilon N + 1) + 3$ on the distance from any point of the image of any ρ_k to the image of P, for $k_1 \leq k \leq k_2$. Hence we can assume $|\mathcal{V}| \geq \frac{k_2 - k_1}{2}$. Since there is a bound of $3N$ on the number of vertical dual curves intersecting each ρ_k, there exists an integer $p = p(N) \geq 1$, independent of ϵ, such that any concatenation of p consecutive paths of the form ρ_k, with $k_2 \leq k \leq k_1$, crosses at least N vertical dual curves.

To conclude, consider a path $\rho_k \rho_{k+1} \cdots \rho_{k+p}$ with $k_1 \leq k \leq k + p \leq k_2$. This path crosses at least N vertical dual curves. There are at least $\epsilon N - pN$ dual curves in E that cross b_k and a_{k+p}, and thus cross each intervening vertical dual curve emanating from P, since for each k' at most N non-vertical dual curves leave the diagram through $c_{k'}$. Hence take $\epsilon = p + 1$. Then there are at least N horizontal dual curves, each of which crosses each of the at least N vertical dual curves, in the subdiagram between b_k, a_{k+p}, and the subtended parts of ρ and P. Hence there is an $N \times N$ flat grid whose convex hull intersects P and $\rho_k, \rho_{k+1}, \ldots, \rho_{k+p}$. Thus each such path projects to a subspace of the 3-neighborhood of the image of P in Γ. Either every ρ_k is contained in such a path, or $k_2 - k_1 \leq p$ and we can bound the distance from any ρ_k to P in Γ. △

Conclusion. Let $\gamma, \gamma', \gamma'' \to \Gamma$ be geodesics forming a triangle in Γ. Let ρ, ρ', ρ'' be combinatorial paths respectively representing $\gamma, \gamma', \gamma''$ as above, chosen so that $\rho\rho'\rho''$ is a closed path in $\mathbf{X}$. For each $p \in \rho$, there is some subspace $\mathbf{C}$ representing a vertex of γ and containing p. Hence the image of p in Γ lies at distance at most 1 from γ. Similarly, the image of ρ' [respectively ρ''] lies in the 1-neighborhood of the image of γ' [respectively γ'']. By Claim 3, there exist geodesics $\sigma, \sigma', \sigma''$, the $\mathfrak{L}$-neighborhoods of whose images in Γ respectively contain the images of ρ, ρ', ρ''. By Claim 2, the image of the geodesic triangle $\sigma\sigma'\sigma''$ has the property that the image of any side is contained in the δ-neighborhood of the image of the other two sides. Hence $\gamma\gamma'\gamma''$ is $\delta + (2\mathfrak{L} + 1)$-thin, whence Γ is $\delta + 2(\mathfrak{L} + 1)$-hyperbolic. □

In particular, when the $\mathbf{S}_i$ are hyperoctahedra of dimension at least 1 satisfying the hypotheses of Theorem 3.1, then we may conclude that G is hyperbolic relative to a finite collection of virtually abelian subgroups, as we now explain. First, consider the action of Q_i on $\mathbf{C}_i$. This action is proper and cocompact, and by Lemma 2.16 and [15, Proposition 3.5], we may assume that this action is essential. Now, $\mathbf{C}_i$ is fully visible because any invisible simplex is non-maximal and contained in a unique maximal simplex, by the proof of [26, Theorem 3.19], and no such simplices exist in a hyperoctahedron. By [26, Theorem 3.30], the decomposition $\mathbf{S}_i \cong \mathbb{O}_{d-1} \star \mathbb{O}_0$ corresponds to a decomposition $\mathbf{C}_i \cong X_{d-1} \times X_0$, where $\partial_{\triangle} X_0 \cong \mathbb{O}_0$ and $\partial_{\triangle} X_{d-1} \cong \mathbb{O}_{d-1}$. Since the boundary of X_0 is a single pair of points, and X_0 is cocompact, there exists a periodic geodesic γ such that X_0 lies in a finite neighborhood of γ. By induction on dimension, X_{d-1} contains a periodic flat $F \cong \mathbb{R}^{d-1}$ which coarsely contains all of X_{d-1}. Hence $\mathbf{C}_i$ is coarsely contained in a flat $F \times \gamma$ of dimension d that is stabilized by a finite-index subgroup of Q_i. Thus Q_i is virtually $\mathbb{Z}^d$, by Bieberbach's theorem.

Example 3.10. We conclude this section with some examples and non-examples of relatively hyperbolic cocompactly cubulated groups:

(1) Right-angled Artin groups. The results of [2] and [4] combine to show that one-ended right-angled Artin groups are never relatively hyperbolic since they are all either thick of order 0 (in the case the group is a direct product) or thick of order 1 and thus not relatively hyperbolic by [4, Corollary 7.9]. Theorem 3.1 above provides another proof of non-relative hyperbolicity for these groups, since the simplicial boundary of a one-ended right-angled Artin group, A, has only one positive-dimensional connected component.

(2) Hyperbolic relative to a right angled Artin group. Figure 4 shows a cubical subdivision of the Salvetti complex $\overline{\mathbf{C}}$ of

$$F_2 \times \mathbb{Z} \cong \langle a, b, t \mid [a,t], [b,t] \rangle$$

at left and a nonpositively-curved cube complex $\overline{\mathbf{Y}}$ at right that is a tiling by 2-cubes of a closed, orientable genus-3 surface. The fundamental group of $\overline{\mathbf{Y}}$ is presented by

$$\pi_1 \mathbf{Y} \cong \langle p_1, q_1, p_2, q_2, p_3, q_3 \mid [p_1, q_1][p_2, q_2][p_3, q_3] \rangle$$

and we form a compact nonpositively-curved cube complex $\overline{\mathbf{X}}$ by attaching a cylinder to $\overline{\mathbf{C}}$ and $\overline{\mathbf{Y}}$ as shown, so that $G = \pi_1 \overline{\mathbf{X}}$ is isomorphic to the following

$$(\pi_1 \mathbf{C} * \pi_1 \mathbf{Y}) / \langle\langle b = p_1 q_1^{-1} p_1^{-1} p_2 q_2 p_2^{-1} \rangle\rangle$$

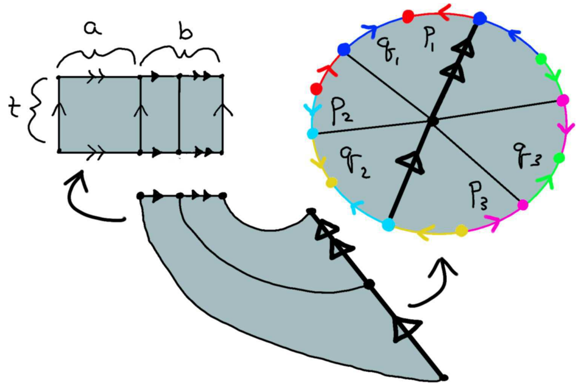

Figure 4

Since the attaching maps of the cylinder are locally convex circles, $\bar{\mathbf{C}}$ and $\bar{\mathbf{Y}}$ are locally convex in $\bar{\mathbf{X}}$, and hence the universal cover $\mathbf{C}$ is a convex, $P \cong \pi_1\bar{\mathbf{C}}$-cocompact subcomplex of the universal cover $\mathbf{X}$. Now, $\mathbf{S} = \partial_\triangle \mathbf{C}$ is isomorphic to the join of an infinite discrete set with a pair of 0-simplices, and $\mathbf{S} \subset \partial_\triangle \mathbf{X}$. Any two distinct translates of $\mathbf{C}$ intersect in a translate of a convex periodic geodesic lying in a translate of the universal cover $\mathbf{Y}$, which is a convex copy of $\mathbb{H}^2$ in $\mathbf{X}$. Hence, since cyclic subgroups of $\pi_1\bar{\mathbf{Y}}$ are malnormal, $\mathbf{S} \cap g\mathbf{S} = \emptyset$ for $g \notin P$. Now, every flat orthant in $\mathbf{X}$ lies in some translate of $\mathbf{C}$. Therefore, $\partial_\triangle \mathbf{X}$ is the union of translates of $\mathbf{S}$ together with a nonempty set of isolated points arising from translates of $\partial_\triangle \mathbf{Y}$, and Theorem 3.7 confirms that G is hyperbolic relative to P.

(3) Cusped hyperbolic 3-manifolds. There are many cusped, hyperbolic 3-manifolds $\widehat{M}$ for which $\pi_1\widehat{M}$ is the fundamental group of a compact nonpositively-curved cube complex. Such manifolds arise as finite covers of finite-volume cusped hyperbolic 3-manifolds that contain a geometrically finite incompressible surface [44, Theorem 14.29]. In this case, the cusp subgroups correspond to isolated 4-cycles in the simplicial boundary of the cocompact cubulation of $\pi_1 M$, the remainder of which consists of an infinite collection of isolated 0-simplices.

4. Unconstricted and wide cube complexes

We assume throughout this section that $\mathbf{X}$ is a locally finite, finite-dimensional CAT(0) cube complex.

$\mathbf{X}$ is *geodesically complete* if each CAT(0) geodesic segment is contained in a bi-infinite CAT(0) geodesic. If $\mathbf{X}$ is geodesically complete, then it is *combinatorially geodesically complete* in the sense that, for any maximal set $W_1, \ldots, W_n$

of pairwise-crossing hyperplanes, each of the 2^n maximal intersections of half-spaces associated to those hyperplanes contains 0-cubes arbitrarily far from the cube $\bigcap_{i=1}^{n} N(W_i)$. Equivalently, $\mathbf{X}$ is combinatorially geodesically complete if every combinatorial geodesic segment extends to a bi-infinite combinatorial geodesic, as is shown in [26]. If $\mathbf{X}$ is (combinatorially or CAT(0)) geodesically complete, then $\mathbf{X}$ satisfies the first requirement of the definition of an unconstricted space, since each point of $\mathbf{X}$ lies at distance 0 from a bi-infinite (combinatorial or CAT(0)) geodesic and hence lies uniformly close to a CAT(0) quasigeodesic.

Let ω be an ultrafilter, $(s_n)_{n\geq 1}$ a sequence of scaling constants, and $(x_n)_{n\geq 1}$ be a sequence of observation points in $\mathbf{X}$. Given a sequence $(y_n \in \mathbf{X})_{n\geq 1}$, we write $[y_n]$ to denote the associated point of $\mathbf{Cone}_\omega(\mathbf{X},(x_n),(s_n))$. Since $\mathbf{X}$ is finite-dimensional the CAT(0) metric and the path metric on $\mathbf{X}^{(1)}$ are quasi-isometric, and thus $\mathbf{Cone}_\omega(\mathbf{X},(x_n),(s_n))$ is bilipschitz homeomorphic to $\mathbf{Cone}_\omega(\mathbf{X}^{(1)},(x_n'),(s_n))$, where x_n' is a closest 0-cube to x_n. Where the ultrafilter, scaling constants, and observation points are understood, we denote this asymptotic cone by $\mathbf{X}_\omega$.

We say $\partial_\triangle \mathbf{X}$ is *bounded* if its 1-skeleton (with the usual graph metric) is finite diameter.

Theorem 4.1. *Let* $\mathbf{X}$ *be a locally finite, finite-dimensional* CAT(0) *cube complex such that* $|\partial_\triangle \mathbf{X}| > 1$. *If* $\partial_\triangle \mathbf{X}$ *is bounded then no asymptotic cone of* $\mathbf{X}$ *is separated by a finite closed ball, in the sense that in no asymptotic cone do there exist points* $\mathbf{a},\mathbf{b},\mathbf{x}$ *such that* $\mathbf{d}_\omega(\mathbf{x},\{\mathbf{a},\mathbf{b}\}) > 3$ *and every path from* $\mathbf{a}$ *to* $\mathbf{b}$ *passes through the 1-ball about* $\mathbf{x}$. *Under the additional hypotheses that every combinatorial geodesic segment can be extended to a ray: if* $\partial_\triangle \mathbf{X}$ *is bounded, then* $\mathbf{X}$ *is wide.*

Proof. Although $\mathbf{X}$ is not assumed to be fully visible, we always work with visible simplices, justified by the fact that maximal simplices are visible [26, Theorem 3.19].

Let α,β be combinatorial geodesics, representing simplices h_α, h_β of $\partial_\triangle \mathbf{X}$ respectively. Without loss of generality, α and β have a common initial point x_o. The *cubical divergence*, $\mathbf{div}(\alpha,\beta)(r)$, is the length of a shortest combinatorial path $P_r \to \mathbf{X}$ which joins $\alpha(r)$ to $\beta(r)$ and contains no 0-cube at distance less than r from x_o. Now, h_α and h_β lie in the same component of $\partial_\triangle \mathbf{X}$ if and only if $\mathbf{div}(\alpha,\beta)(r)$ is bounded above by a linear function of r, by [26, Theorem 6.8]. In this case, for all $r \geq 0$,

$$A_1 r + B_1 \leq \mathbf{div}(\alpha,\beta)(r) \leq A_2 r + B_2$$

where A_1, A_2 depend linearly on the distance between h_α and h_β in $\partial_\triangle \mathbf{X}^{(1)}$ and B_1, B_2 are constants depending on α and β. We first exhibit a cut-ball in an asymptotic cone when $\partial_\triangle \mathbf{X}$ is disconnected, and then do the same when $\partial_\triangle \mathbf{X}$ is connected but unbounded.

Disconnected $\partial_\triangle \mathbf{X}$ implies cut-ball. Suppose that h_α and h_β lie in distinct components of $\partial_\triangle \mathbf{X}$. Then, for each $M \geq 1$, there exists a smallest $r_M \geq M$ such that $\mathbf{div}(\alpha,\beta)(r_M) \geq Mr_M$. From the definition of r_M, it follows immediately that $\mathbf{div}(\alpha,\beta)(Kr_M) \geq (2-2K+M)r_M$ for any fixed $K \geq 1$.

Consider an asymptotic cone, $\mathbf{Cone}_\omega(\mathbf{X},\mathbf{x},(r_n))$, where the scaling constants are given by the (r_n) above, and the sequence of observation points is $\mathbf{x} = (x_o)$.

For each $n \geq 0$, let $a_n = \alpha(Kr_n)$, where $K \geq 3$ is some fixed integer, and likewise let $b_n = \beta(Kr_n)$. Then $\dot{d}(a_n, x_o)r_n^{-1} = K = \dot{d}(b_n, x_o)r_n^{-1}$, so that $\mathbf{a} = [(a_n)], \mathbf{b} = [(b_n)]$ define points of $\mathbf{Cone}_\omega(\mathbf{X},\mathbf{x},(r_n))$, and these points are each at distance K from $\mathbf{x}$.

By construction, any path P_n in X from a_n to b_n either has length at least $(2-2K+n)r_n$ or travels through the interior of the r_n-ball about x_o, i.e., through the closed (r_n-1)-ball. We see this as follows. By prepending the part of α joining $\alpha(r_n)$ to a_n, and appending the part of β joining b_n to $\beta(r_n)$, to P_n, we obtain a path P'_n of length $2(K-1)r_n + |P_n|$ joining $\alpha(r_n)$ to $\beta(r_n)$. Either P'_n travels through the interior of the forbidden r_n-ball or else, by our choice of r_n, $|P'_n| \geq nr_n$ and thus $|P_n| \geq (2-2K+n)r_n$ as claimed.

By construction, $\mathbf{d}_\omega(\mathbf{a},\mathbf{b}) \leq 2K$ and, as noted above, $\mathbf{d}_\omega(\mathbf{a},\mathbf{x}) = \mathbf{d}_\omega(\mathbf{b},\mathbf{x}) = K$. We shall show that the closed ball of radius 1 about $\mathbf{x}$ separates $\mathbf{a}$ from $\mathbf{b}$.

Let $\mathfrak{P}$ be a finite length path in $\mathbf{Cone}_\omega(\mathbf{X},\mathbf{x},(r_n))$ joining $\mathbf{a}$ to $\mathbf{b}$ and let P_n be a path in $\mathbf{X}$ joining a_n to b_n for which the ω-limit of these paths is $\mathfrak{P}$. Either P_n passes through the (r_n-1)-ball about x_o for ω-almost all n, or $|P_n| \geq (2-2K+n)r_n$ for ω-almost all n. Now, the latter case can't occur, since if it did then we would have $\lim_\omega |P_n|r_n^{-1} = \infty$, and thus $\mathfrak{P}$ has infinite length, contradicting our hypothesis. In the former case, by taking the ω-limit of these balls, we have that $\mathfrak{P}$ passes through a ball of radius $\lim_\omega \frac{r_n-1}{r_n} = 1$. Taking $K > 3$, the claim is proved.

Unbounded $\partial_\triangle \mathbf{X}$ implies cut-ball. By [26, Theorem 6.9], for each $n \geq 0$, we have $r_n \geq 0$ and combinatorial geodesic rays α_n, β_n emanating from x_o with $\mathbf{div}(\alpha_n,\beta_n)(r) \geq nr$ for all $r \geq r_n$. From this point the argument then finishes exactly as above.

Bounded $\partial_\triangle \mathbf{X}$ implies no cut-ball. First we show: if $\partial_\triangle \mathbf{X}$ is bounded and $|\partial_\triangle \mathbf{X}| > 1$, then the combinatorial metric on $\mathbf{X}^{(1)}$ has linear divergence function.

Let $a,b,c \in \mathbf{X}^{(1)}$, with $\dot{d}(a,b) \leq n$ and $\dot{d}(\{a,b\},c) = r > 0$. Choose $\delta \in \left(0,\frac{1}{2}\right)$ and $\kappa \geq 0$. Let μ be the median of a,b,c and let γ be a bi-infinite path with $\gamma(0) = \mu$ and $\gamma(-t_a) = a, \gamma(t_b) = b$ for $t_a, t_b \in [0,n]$ and with the property that both $\gamma|_{(-\infty,0]}$ and $\gamma|_{[0,\infty)}$ are geodesic rays. Here we have used the combinatorial geodesic-ray completeness hypothesis.

Since X is finite-dimensional and locally finite, the hypothesis of [26, Theorem 6.8] is satisfied, and thus, since $\partial_{\triangle}\mathbf{X}$ is bounded, the divergence of γ is bounded above by a linear function with uniform additive and multiplicative constants. Note that to use [26, Theorem 6.8] implicitly requires $|\partial_{\triangle}\mathbf{X}| \geq 2$, since a pair of distinct infinite geodesic rays is required in order to apply that theorem.

If $\dot{d}(\mu, c) > \delta r - \kappa$, then the subpath of γ joining a to b has length $t_a + t_b \leq n$ and avoids the $(\delta r - \kappa)$-ball about c. In this case we thus have that $\mathbf{div}_{\delta,\kappa}(a, b, c) = t_a + t_b \leq n$.

Hence, we restrict our attention to the alternate case where $\dot{d}(\mu, c) \leq \delta r - \kappa$. Let $T = 2\max\{t_a, t_b\}$. Note that since $\delta < \frac{1}{2}$ we have $\min\{t_a, t_b\} \geq \frac{r}{2}$. Since, as noted above, γ has linear divergence, there exists a path P connecting $\gamma(-T)$ to $\gamma(T)$ whose length is linear in T and which avoids the ball of radius T about $\gamma(0)$, i.e., for each $p \in P$ we have $\dot{d}(p, \mu) \geq T$. Since $\dot{d}(\mu, c) \leq \delta r - \kappa$, the triangle inequality implies that for each $p \in P$ we have $\dot{d}(p, c) \geq T - \delta r - \kappa \geq \delta r - \kappa$. Thus concatenating P with the subpaths of γ from $\gamma(-T)$ to $\gamma(-t_a)$ and from $\gamma(t_b)$ to $\gamma(T)$ (which are each of length at most n), we get a path P' connecting a to b, which is of linear length and which avoids the $(\delta r - \kappa)$-ball about c.

Hence, for any choices of a, b, c we have obtained that $\mathbf{div}_{\delta,\kappa}(a, b, c)$ is bounded above by a linear function with uniform constants, as desired.

The remainder of the argument is a routine application of linear divergence. For a fixed space $\mathbf{Cone}_{\omega}(\mathbf{X}, \mathbf{x}, (s_n))$, we want to show that for each closed ball $\mathcal{B}$ in $\mathbf{Cone}_{\omega}(\mathbf{X}, \mathbf{x}, (s_n))$ and distinct points $\mathbf{a}, \mathbf{b} \in \mathbf{Cone}_{\omega}(\mathbf{X}, \mathbf{x}, (s_n)) - \mathcal{B}$, there exists a path in $\mathbf{Cone}_{\omega}(\mathbf{X}, \mathbf{x}, (s_n)) - \mathcal{B}$ joining $\mathbf{a}$ to $\mathbf{b}$. To do this we fix sequences $(a_n), (b_n)$ representing $\mathbf{a}, \mathbf{b}$, respectively, and let (c_n) be a sequence representing $\mathbf{c}$, the center of the ball $\mathcal{B}$. Since the divergence of $\mathbf{X}$ is linear, following the proof of [16, Lemma 3.14] shows that no ball in $\mathbf{Cone}_{\omega}(\mathbf{X}, \mathbf{x}, (s_n))$ about $\mathbf{c}$ of radius less than δ can separate $\mathbf{a}$ from $\mathbf{b}$. Any ball of radius at least r about $\mathbf{c}$ contains an element of $\{\mathbf{a}, \mathbf{b}\}$ and hence cannot separate those points. □

The following corollary is a characterization of wide cube complexes in a slightly more general framework than we shall later apply. Cocompactness of the action of $\mathrm{Aut}(\mathbf{X})$ is needed to find a cut-point in an asymptotic cone given a cut-ball in some other asymptotic cone. For the converse, the failure to be wide implies that the simplicial boundary is unbounded, and this assumption is unnecessary. We have hypothesized finite-dimensionality so that $\mathbf{X}$ with the CAT(0) metric is quasi-isometric to $\mathbf{X}^{(1)}$, which is the natural setting for working with the simplicial boundary.

Corollary 4.2. *Let* $\mathbf{X}$ *be a locally finite, geodesically complete, finite-dimensional* CAT(0) *cube complex on which* $\mathrm{Aut}(\mathbf{X})$ *acts cocompactly. Then* $\mathbf{X}$ *is wide if and only if* $\partial_{\triangle}\mathbf{X}$ *is bounded.*

Proof. By geodesic completeness, every point of $\mathbf{X}$ lies in a bi-infinite geodesic. By Theorem 4.1, if $\partial_\triangle \mathbf{X}$ is unbounded then some asymptotic cone of $\mathbf{X}$ has a finite cut-ball. More precisely, there exists $\delta \geq 0$ and points $\mathbf{a}, \mathbf{b}, \mathbf{c}$ in some asymptotic cone, with $\mathbf{d}_\omega(\mathbf{c}, \{\mathbf{a}, \mathbf{b}\}) > 3\delta$, such that the closed δ-ball about $\mathbf{c}$ separates $\mathbf{a}$ from $\mathbf{b}$. By [16, Lemma 3.16], $\mathbf{X}$ is not wide.

Conversely, if $\mathbf{X}$ is not wide, then $\partial_\triangle \mathbf{X}$ is unbounded, by Theorem 4.1. □

In the event of a proper, cocompact, essential group action, that $\mathbf{X}$ is wide corresponds to $\partial_\triangle \mathbf{X}$ being connected can be seen without directly analyzing the asymptotic cones.

Theorem 4.3. *Let* $\mathbf{X}$ *be a* CAT(0) *cube complex on which the group* G *acts properly and cocompactly. Then* $\mathbf{X}$ *is wide if and only if* $\partial_\triangle \mathbf{X}$ *is connected.*

Hence, if G *is a cocompactly cubulated group, then* G *is wide if and only if* G *acts geometrically on a* CAT(0) *cube complex with connected simplicial boundary.*

Proof. Throughout the proof, by appealing to [15, Proposition 3.5] and Proposition 2.16, we assume that G acts essentially on $\mathbf{X}$.

The end-points of the axis stabilized by any rank-one element in G are isolated 0-simplicies in the the boundary. Thus, $\partial_\triangle \mathbf{X}$ is connected if and only if G does not contain any rank-one elements. By the rank-rigidity theorem [15, Theorem 6.3] G does not contain any rank-one elements if and only if there exists unbounded convex subcomplexes $\mathbf{X}_1, \mathbf{X}_2 \subset \mathbf{X}$ satisfying $\mathbf{X} = \mathbf{X}_1 \times \mathbf{X}_2$. If the space $\mathbf{X}$ is such a direct product then it has linear divergence; if it has a rank-one element then its divergence is superlinear. By [16, Proposition 1.1] a space linear divergence if and only if it is wide. □

5. Characterizing thickness of order 1

Throughout this section, $\mathbf{X}$ will denote a CAT(0) cube complex on which a group G acts properly and cocompactly. Let $\mathcal{I}$ denote the subcomplex of $\partial_\triangle \mathbf{X}$ consisting of all isolated 0-simplices. Since maximal simplices of $\partial_\triangle \mathbf{X}$ are visible, each $v \in \mathcal{I}$ is represented by a combinatorial geodesic ray that is rank-one in the sense of [26], and conversely, each rank-one geodesic ray represents an isolated 0-simplex of $\partial_\triangle \mathbf{X}$. In this section, we adopt the following notation: if $Y \subset \mathbf{X}$ is a subspace, we denote by $\widehat{Y}$ its cubical convex hull.

5.1. Simplicial boundaries of algebraically thick cube complexes. A *cubical flat sector* is a CAT(0) cube complex of the form $\mathbb{R}^p \times [0, \infty)^q$, with $p + q \geq 2$, tiled in the standard Euclidean fashion by unit $(p+q)$-cubes. The class of cubical flat sectors includes cubical orthants, half-flats, and flats of dimension at least 2.

Our first theorem describes simplicial boundaries of CAT(0) cube complexes admitting geometric actions by groups that are algebraically thick of order 1.

Theorem 5.1. *Let G act properly and cocompactly on a fully visible* CAT(0) *cube complex* $\mathbf{X}$*, and suppose that G is algebraically thick of order* 1 *relative to a collection $\mathbb{G}$ of quasiconvex wide subgroups. Then $\mathcal{I} \neq \emptyset$ and $\partial_\triangle \mathbf{X}$ has at least one G-invariant positive-dimensional component.*

Remark 5.2. Note that full visibility of $\mathbf{X}$ is hypothesized. This hypothesis can be removed if Conjecture 2.8 is true. The conclusion of the above theorem holds in slightly more generality, namely when $\mathbf{X}$ is thick relative to a G-invariant collection of convex subcomplexes with connected simplicial boundaries. We also note that in many examples the G-invariant component is in fact the unique positive-dimensional component.

Proof of Theorem 5.1*.* Throughout the proof, by appealing to [15, Proposition 3.5] and Proposition 2.16, we assume that G acts essentially on $\mathbf{X}$. We can assume that G is one-ended, since otherwise $\partial_\triangle \mathbf{X}$ is disconnected and G is not thick.

Necessarily, $\mathcal{I} \neq \emptyset$. To see this, note first that if $\mathcal{I} = \emptyset$, then G cannot contain a rank-one isometry of $\mathbf{X}$, since the set of hyperplanes crossing an axis for such an element would represent a pair of isolated 0-simplices of $\partial_\triangle \mathbf{X}$. In such a case, by rank-rigidity, $\mathbf{X}$ decomposes as the product of two convex subcomplexes, each of which has nonempty simplicial boundary, and therefore $\partial_\triangle \mathbf{X}$ decomposes as a nontrivial simplicial join. Hence we have, in particular, that $\partial_\triangle \mathbf{X}$ is bounded and hence connected; thus by Theorem 4.3 G is wide, i.e., strongly algebraically thick of order 0, a contradiction.

Representing $\mathbb{G}$ in $\mathbf{X}$. Fix a 0-cube $x_o \in \mathbf{X}$. For each $H \in \mathbb{G}$, the orbit Hx_o is quasiconvex; denote by S_H the convex hull of this orbit. By quasiconvexity, S_H is contained in a uniform neighbourhood of Hx_o, and therefore S_H is an H-cocompact CAT(0) cube complex. Let $\mathcal{S} = \{gS_H : g \in G, H \in \mathbb{G}\}$. Since G acts cocompactly on $\mathbf{X}$, the set $\mathcal{S}$ coarsely covers $\mathbf{X}$ (denote by τ a constant such that the τ-neighborhoods of the various S_H together cover $\mathbf{X}$). Now, each $H \in \mathbb{G}$ is wide, so since the property of being wide is quasi-isometry invariant, S_H is likewise wide. By Theorem 4.3, $\partial_\triangle S_H$ is a connected, positive-dimensional subcomplex of $\partial_\triangle \mathbf{X}$. Finally, since G is algebraically thick with respect to $\mathbb{G}$, $\mathbf{X}$ is thick with respect to $\mathcal{S}$, i.e., for all $S, S' \in \mathcal{S}$, there exists a sequence $S = S_0, S_1, \ldots, S_k = S'$ such that for $1 \leq i \leq k$, the intersection $\mathcal{N}_\tau(S_{i-1}) \cap \mathcal{N}_\tau(S_i)$ is τ-path-connected and coarsely unbounded. We now make a series of modifications to $\mathcal{S}$ to put it in a particularly nice form.

Thickness relative to flat sectors. As stated above, for each $S \in \mathcal{S}$, $\partial_{\triangle} S$ is connected. Let $\mathcal{F}_S$ be the set of all cubical flat orthants in S of dimension exceeding 1. If $F, F' \in \mathcal{F}_S$, then the simplices v_F, v'_F of $\partial_{\triangle} S$ are joined by a sequence $v_F = v_0, \ldots, v_n = v_{F'}$ of positive-dimensional simplices such that $v_{i-1} \cap v_i \neq \emptyset$ for $1 \leq i \leq n$. By full visibility of S – here we use the fact that full visibility is inherited by convex subcomplexes, by definition – there exists a sequence $F = F_0, \ldots, F_n = F'$ such that F_i and F_{i-1} are crossed by infinitely many common hyperplanes, for $1 \leq i \leq n$. Employing the *Flat Bridge Trick* (Lemma 5.3 below), we can assume that $\widehat{F}_i \cap \widehat{F}_{i-1}$ is path-connected and unbounded for all i.

For $S, S' \in \mathcal{S}$, suppose that $F \in \mathcal{F}_S$ and $F' \in \mathcal{F}_{S'}$ are flat orthants of dimension at least 2. Choose a sequence $S = S_0, \ldots, S_k = S'$ such that the intersection of consecutive terms is coarsely connected and unbounded, i.e., $\mathcal{H}(S_i) \cap \mathcal{H}(S_{i+1})$ is infinite for all i. Applying the flat Bridge Trick, we find a cubical flat sector F_i (containing a half-flat) such that the intersection of F_i with each of S_i and S_{i+1} contains a flat orthant. Hence F can be thickly connected to F' by a chain of flat orthants. Moreover, any new flat orthant added during an application of the Flat Bridge Trick is coarsely contained in a flat orthant belonging to some $S \in \mathcal{S}$, and can thus be thickly connected to any other such flat orthant by a sequence of flat orthants.

Conclusion. Thus $\mathbf{X}$ contains a collection $\mathcal{S}'$ of convex subcomplexes $\widehat{F}$, where each F is a flat orthant of dimension at least 2, such that $\mathbf{X} = \mathcal{N}_\tau(\bigcup_{\widehat{F} \in \mathcal{S}'} \widehat{F})$ and, for all $\widehat{F}, \widehat{F}' \in \mathcal{S}'$, there exists a sequence $F = F_0, F_1, \ldots, F_k = F'$ such that $\widehat{F}_i \in \mathcal{S}'$ for all i and $\widehat{F}_i \cap \widehat{F}_{i-1}$ is connected and unbounded for $1 \leq i \leq k$.

Now, for each $\widehat{F} \in \mathcal{S}'$, let $\chi_{\widehat{F}}$ be the image in $\partial_{\triangle}\mathbf{X}$ of the simplicial boundary of $\widehat{F}$. Each $\chi_{\widehat{F}}$ is a positive-dimensional simplex by Proposition 2.13. The above discussion shows that $\bigcup_{\widehat{F} \in \mathcal{S}} \chi_{\widehat{F}}$ is a connected subcomplex of $\partial_{\triangle}\mathbf{X}$. Finally, for each $\widehat{F} \in \mathcal{S}'$, either F is a flat orthant of some $S \in \mathcal{S}$, or F is a flat orthant such that, for some $S, S' \in \mathcal{S}$, each of the intersections $\widehat{F} \cap S$ and $\widehat{F} \cap S'$ is unbounded and path-connected. Since $\mathcal{S}$ is G-invariant – it is the set of G-translates of convex hulls of the various Hx_o – the set of all such orthants F, and hence $\bigcup_{\widehat{F} \in \mathcal{S}} \chi_{\widehat{F}}$, is likewise G-invariant. The component containing this subcomplex is thus G-invariant. □

Lemma 5.3 (the flat bridge trick). *Let* $\mathbf{X}$ *be finite-dimensional, locally finite a* CAT(0) *cube complex, and let* $S, S' \subseteq \mathbf{X}$ *be wide convex subcomplexes with connected simplicial boundaries, such that there exist flat sectors* $F \subseteq S$, $F' \subseteq S'$, *and* $\mathcal{H}(S) \cap \mathcal{H}(S')$ *is infinite. Then there exists a sequence* $F = F_0, F_1, \ldots, F_n = F'$ *of flat sectors such that for all* i, *the intersection* $\widehat{F}_i \cap \widehat{F}_{i+1}$ *is unbounded and path-connected.*

Proof. FLAT BRIDGES FOR PAIRS OF FLAT SECTORS. First, let F_i, F_{i+1} be flat sectors such that $\partial_{\triangle} \widehat{F}_j$ contains a simplex u_j, for $j \in \{i, i+1\}$, such that $u_i \cap u_{i+1} \neq \emptyset$. Then there exists for each i an infinite set of hyperplanes H such that H crosses both $\widehat{F}_i$ and $\widehat{F}_{i+1}$. Hence $\mathcal{H}(F_i) \cap \mathcal{H}(F_{i+1})$ contains a boundary set $\mathcal{V}$ representing a 0-simplex $v \in u_i \cap u_{i+1}$. Choose disjoint minimal boundary sets $\mathcal{W}_i \subset \mathcal{H}(F_i)$ and $\mathcal{W}_{i+1} \subset \mathcal{H}(F_{i+1})$. Then the smallest set of hyperplanes containing $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ that is closed under separation is of the form $\mathcal{H}(\alpha)$ for some bi-infinite geodesic α containing an infinite ray in F_i and an infinite ray in F_{i+1}.

Now, since $\mathcal{V} \cap \mathcal{W}_i = \mathcal{V} \cap \mathcal{W}_{i+1} = \emptyset$, for any geodesic ray β in F_i or F_{i+1} with initial point on α and $\mathcal{H}(\beta) \subseteq \mathcal{V}$, every hyperplane dual to a 1-cube of β crosses every hyperplane dual to a 1-cube of α, and thus there is an isometric embedding $\alpha \times \beta \to \mathbf{X}$. The half-flat $\alpha \times \beta$ has the property that its convex hull contains a flat orthant in F_i and a flat orthant in F_{i+1}, since β has the same set of dual hyperplanes as a ray in F_i and a ray in F_{i+1}. The half-flat $\alpha \times \beta$ is a flat sector $F_{i+\frac{1}{2}}$, and $\widehat{F}_{i+\frac{1}{2}}$ must have unbounded convex intersection with $\widehat{F}_i, \widehat{F}_{i+1}$.

FLAT BRIDGES FOR SUBCOMPLEXES WITH CONNECTED BOUNDARIES. The same argument can be applied to arbitrary wide convex subcomplexes S, S' with $\mathcal{H}(S) \cap \mathcal{H}(S')$ infinite. Indeed, there exist combinatorial geodesic rays γ, γ' in S, S' respectively, such that $\mathcal{H}(\gamma) = \mathcal{H}(\gamma')$. Since S, S' are wide, γ and γ' can be chosen to lie in flat sectors $F_i \subseteq S$, $F_{i+1} \subseteq S'$, and we argue as above. If $F \subseteq S$, $F' \subseteq S'$ are the given flat sectors, then since S has connected simplicial boundary, we can chain F to F_i and F_{i+1} to F' by thickly connecting sequences of convex hulls of flat sectors, and the proof is complete. □

The Flat Bridge Trick is also used in the next section.

5.2. Identifying thickness and algebraic thickness of order 1.

The goal of this section is to prove Theorem 5.4, which allows one to identify thickness of order 1, and algebraic thickness of order 1, of a group G acting geometrically on the cube complex $\mathbf{X}$ by examining the action of G on the simplicial boundary and on the visual boundary. For thickness, one only need examine the action on the simplicial boundary, while a convenient statement of hypotheses implying algebraic thickness also involves the action on the visual boundary.

In the following, $f: \partial_\infty \mathbf{X} \to \partial_{\triangle} \mathbf{X}$ denotes the surjection defined in Section 2.5.

Theorem 5.4. *Let G be a group which acts properly and cocompactly on a fully visible* CAT(0) *cube complex $\mathbf{X}$. If $\mathcal{I} \neq \emptyset$ and $\partial_{\triangle}\mathbf{X}$ has a positive-dimensional G-invariant connected subcomplex $\mathfrak{C}$, then G is thick of order* 1 *relative to a collection of wide subsets.*

Suppose, further, that there is a finite collection $\mathcal{A}$ of bounded subcomplexes of $\mathfrak{C}$ such that

(1) *the stabilizer H_A of A is quasiconvex for all $A \in \mathcal{A}$,*
(2) *for all $A \in \mathcal{A}$, the set $f^{-1}(A) \subset \partial_\infty \mathbf{X}$ is contained in the limit set of H_A,*
(3) *$\mathfrak{C} = \bigcup_{g \in G, A \in \mathcal{A}} gA$ and $f^{-1}(\mathfrak{C})$ is contained in the limit set of the subgroup of G generated by $\{H_A : A \in \mathcal{A}\}$.*

Then G is algebraically thick of order 1 relative to the collection $\{H_A : A \in \mathcal{A}\}$ of wide subgroups.

Remark 5.5. Note that $\partial_\triangle \mathbf{X}$ has a connected, positive-dimensional, G-invariant subcomplex if and only if $\partial_\triangle \mathbf{X}$ has a positive-dimensional, G-invariant component. Theorem 5.4 is stated in terms of connected subcomplexes, rather than components, since (3) is rarely satisfied if $\mathfrak{C}$ is required to be an entire component.

Theorem 5.4 could be stated in terms of the G-action on $\partial_\triangle \mathbf{X}$ alone, with each hypothesis about limit sets in $\partial_\infty \mathbf{X}$ replaced by the appropriate statement about limit subcomplexes in $\partial_\triangle \mathbf{X}$: the appropriate modification of condition (2) would require each A to lie in the limit subcomplex of H_A and that of (3) would require $\mathfrak{C}$ to lie in the limit subcomplex of $\langle\{H_A\}\rangle$.

Proof of Theorem 5.4. Suppose that $\partial_\triangle \mathbf{X} - \mathfrak{I}$ is nonempty and has a G-invariant positive-dimensional connected subcomplex $\mathfrak{C}$. Since $\dim(\mathbf{X}) < \infty$, there is no infinite family of pairwise-crossing hyperplanes and hence every simplex of $\partial_\triangle \mathbf{X}$ is contained in a finite-dimensional maximal simplex, by Theorem 3.14.(2) of [26]. Let v be a maximal positive dimensional simplex of $\mathfrak{C}$. From Proposition 2.11, it follows that there exists an isometrically embedded maximal flat orthant $F_v \cong [0,\infty)^n \subset \mathbf{X}$, for some $n \geq 2$, whose boundary is v. Hence the set $\mathcal{F}$ of *flat sectors* whose convex hulls represent positive-dimensional connected subcomplexes of $\mathfrak{C}$ is nonempty. Moreover, $\mathcal{F}$ is G-invariant, since $\mathfrak{C}$ is. To see this, note that for all $F \in \mathcal{F}$, the inclusion $g\widehat{F} \hookrightarrow \mathbf{X}$ induces an inclusion of simplicial boundaries whose image lies in $g\mathfrak{C} = \mathfrak{C}$. By definition, $gF \in \mathcal{F}$. Hence, by cocompactness, there exists $\tau \geq 0$ such that $\mathbf{X} = \bigcup_{F \in \mathcal{F}} N_\tau(F)$.

For each $F \in \mathcal{F}$, let $\widehat{F}$ be the convex hull of $N_\tau(F)$. Since $\widehat{F}$ is convex, it is a CAT(0) cube complex, and moreover, $\partial_\triangle \widehat{F}$ is bounded and positive-dimensional, being a connected subcomplex of the simplicial boundary of a cubical flat of dimension at least 2 and containing the boundary of a flat orthant. Thus $\widehat{F}$ is wide, by Theorem 4.1. We conclude that $\{\widehat{F} : F \in \mathcal{F}\}$ is a set of convex (and hence uniformly quasiconvex) wide subcomplexes that covers $\mathbf{X}$. By definition, each $\widehat{F}$ has the property that every $f \in \widehat{F}$ is contained in a bi-infinite combinatorial geodesic, and therefore each point in $\widehat{F}$ is uniformly close to a bi-infinite CAT(0) quasigeodesic. To conclude that $\{\widehat{F}\}$ is uniformly wide, it remains to check that no ultralimit of a sequence in $\{\widehat{F}\}$ has a cut-point; this is the content of Lemma 5.9 below.

Let $p, q \in \mathbf{X}$ be 0-cubes, and choose $F, F' \in \mathcal{F}$ so that $p \in N_k(F)$, $q \in N_k(F')$. By assumption, there exists a sequence $v_F = u_0, u_1, \ldots, u_n = v_{F'}$ of maximal simplices in $\mathfrak{C}$ such that $u_i \cap u_{i+1} \neq \emptyset$ for all i. For each i, let $\widehat{F}_i \in \mathcal{F}$ be the convex hull of a maximal flat sector containing an orthant representing u_i. If for each i, there exist geodesic rays $\gamma \subset \widehat{F}_i, \gamma' \subset \widehat{F}_{i+1}$ that fellow-travel at distance τ, then we have thickly connected p to q using convex hulls of flat orthants. (Note that the intersection of CAT(0) τ-neighborhoods of convex subcomplexes is convex.) Otherwise, for any pair F_i, F_{i+1} not containing such a pair of geodesic rays, we construct a third flat orthant $F_{i+\frac{1}{2}}$ whose convex hull has unbounded intersection with $\widehat{F}_i$ and $\widehat{F}_{i+1}$, using the Flat Bridge Trick. Adding the new $\widehat{F}_{i+\frac{1}{2}}$ for each i yields the desired thickly connecting sequence.

Thus far, we have shown that for any two points $x, y \in \mathbf{X}$, and any $\widehat{F}_0, \widehat{F}_n$ with $x \in \mathcal{N}_\tau(\widehat{F}_0)$ and $y \in \mathcal{N}_\tau(\widehat{F}_n)$, there exists a sequence $\widehat{F}_0, \ldots, \widehat{F}_n$ of subcomplexes, where each $\widehat{F}_i$ is the convex hull of a d-dimensional flat, where $d \geq 2$, such that $\mathcal{N}_\tau(\widehat{F}_i) \cap \mathcal{N}_\tau(\widehat{F}_{i+1})$ is unbounded and path-connected for each i. In so doing, we have verified that $\mathbf{X}$, and therefore G, is thick of order at most 1. Since $\mathcal{I} \neq \emptyset$, $\partial_\triangle \mathbf{X}$ is disconnected and hence G contains a rank-one isometry of $\mathbf{X}$, whence G is not unconstricted and is therefore thick of order exactly 1.

Obtaining algebraic thickness. Fix a base 0-cube $x_o \in \mathbf{X}$, and let $\mathbf{C}_A$ denote the cubical convex hull of the quasiconvex orbit $H_A x_o$, for each $A \in \mathcal{A}$. Then $\mathbf{C}_A$ is an H_A-cocompact subcomplex, by quasiconvexity. By passing to the H_A-essential core of $\mathbf{C}_A$, if necessary, we may assume that $\mathbf{C}_A$ is a CAT(0) cube complex on which H_A acts properly, cocompactly, and essentially.

Let $u \subseteq \partial_\triangle \mathbf{C}_A$ be a simplex represented by $\mathcal{H}(\gamma)$ for some geodesic ray γ emanating from x_o. Since H_A acts cocompactly on $\mathbf{C}_A$, γ is contained in the limit of a sequence of H_A-periodic geodesics, from which it is easily verified that u is a limit simplex of H_A. Conversely, if u is a limit simplex of H_A that is not contained in $\mathbf{C}_A$, then u is represented by $\mathcal{H}(\gamma)$ for some geodesic ray γ that contains points arbitrarily far from $\mathbf{C}_A$. There exists a sequence $(h_i \in H_A)$ such that $\mathcal{H}(\gamma)$ is the set of hyperplanes H such that H separates all but finitely many $h_i x_o$ from x_o. Since $\mathbf{C}_A$ is convex and γ contains points arbitrarily far from $\mathbf{C}_A$, infinitely many $H \in \mathcal{H}(\gamma)$ separate points of γ from $\mathbf{C}_A$, whence $(h_i x_o)$ contains points not in $\mathbf{C}_A$, contradicting the fact that the latter contains $H_A x_o$ by definition. Thus $\partial_\triangle \mathbf{C}_A$ coincides with the limit complex for H_A. Our hypothesis that A is contained in the limit complex for H_A implies that $A \subseteq \partial_\triangle \mathbf{C}_A$. ($A$ is contained in the limit complex for H_A since $f^{-1}(A)$ is contained in the limit set of H_A, by Lemma 2.18.)

Suppose H_A contains a rank-one isometry g of $\mathbf{C}_A$. We shall show that this contradicts the fact that A is bounded. Let $a \in A$ be a visible 0-simplex (this must exist because A contains a maximal simplex, at least one of whose 0-simplices must be visible, by the proof of [26, Theorem 3.19]). Then either the orbit $\langle g \rangle a$ is unbounded and contained in $A^{(1)}$, by Lemma 5.6 below, or g fixes a. The former possibility contradicts boundedness of A. Hence g fixes each $a \in A^{(0)}$. This contradicts the fact that g is rank-one unless A consists of a pair of 0-simplices represented by a combinatorial geodesic axis for g, which is impossible since A is connected. Hence H_A contains no rank-one elements.

By Corollary B of [15], $\mathbf{C}_A$ decomposes as a non-trivial product equal to the limit complex for H_A, and thus $\mathbf{C}_A$ has bounded simplicial boundary that contains A and is contained in $\mathfrak{C}$. We may thus assume that $\partial_{\triangle}\mathbf{C}_A = A$, by adding to A, if necessary, any simplices of $\mathfrak{C}$ that lie in $\partial_{\triangle}\mathbf{C}_A$ but not in A.

Since $A = \partial_{\triangle}\mathbf{C}_A$ is connected, $\mathbf{C}_A$, and therefore H_A, is wide by Theorem 4.3.

Verifying thick connectivity. Let $A, A' \in \mathcal{A}$ and let $H = H_A$, $H' = H_{A'}$. Suppose that $A \cap A' \neq \emptyset$. Then $\mathcal{H}(\mathbf{C}_A) \cap \mathcal{H}(\mathbf{C}_{A'})$ is infinite, and by cocompactness of the actions of H on $\mathbf{C}_A$ and H' on $\mathbf{C}_{A'}$, it follows that $H_A \cap H'_A$ is infinite (the same holds for conjugates of H, H': if the corresponding G-translates of A, A' have nonempty intersection, then the corresponding conjugates of H and H' have infinite intersection). Conversely, if $H \cap H'$ is infinite, then the intersection contains a hyperbolic isometry of $\mathbf{X}$, and thus each of $\mathbf{C}_A$ and $\mathbf{C}_{A'}$ contains a bi-infinite combinatorial geodesic such that these two geodesics are parallel, and hence represent the same simplices of $\mathfrak{C}$. Thus $A \cap A' \neq \emptyset$. Now, without loss of generality, $\bigcup_{A\in\mathcal{A}} A$ is a connected subcomplex of $\mathfrak{C}$, which can be achieved by choosing conjugacy class representatives of the various $A \in \mathcal{A}$ so that the corresponding subcomplex is connected, and replacing $\mathcal{A}$ by this collection of subgroups. Hence, for any $A, A' \in \mathcal{A}$, there exists a sequence $A = A_0, \dots, A_n = A'$ such that $A_i \in \mathcal{A}$ for all i and $A_i \cap A_{i+1} \neq \emptyset$ for $0 \leq i \leq n-1$. Hence $H_{A_i} \cap H_{A_{i+1}}$ is infinite for $0 \leq i \leq n-1$.

Verifying that $\bigcup_{A\in\mathcal{A}} H_A$ generates. To complete the proof of algebraic thickness of G, it suffices to show that $G' = \langle\{H_A : A \in \mathcal{A}\}\rangle$ has finite index in G, and, to this end, we will verify that there exists $R \geq 0$ such that $\mathbf{X} = \mathcal{N}_R(G'(\bigcup_{A\in\mathcal{A}} \mathbf{C}_A))$.

If the preceding equality does not hold, then for all $r \geq 0$, there exists $x_r \in \mathbf{X}$ such that $\ddot{d}(x_r, h\mathbf{C}_A) > r$ for all $A \in \mathcal{A}$ and all $h \in G'$. By cocompactness of the G-action on $\mathbf{X}$, we may choose $\{x_r\}_{r\geq 0}$ so that for some fixed $A \in \mathcal{A}$ and $g \in G - G'$, each $x_r \in g\mathbf{C}_A$ and x_r converges to a point $x_\infty \in f^{-1}(g\partial_{\triangle}\mathbf{C}_A) \subset \partial_\infty \mathbf{X}$. Thus $f(x_\infty) \in g\partial_{\triangle}\mathbf{C}_A$, but x_∞ fails, by construction, to be a limit point of G', a contradiction. Hence $\mathbf{X}$ is contained in a finite neighborhood of the union of G'-translates of the various $\mathbf{C}_A$, and the stabilizer of each $\mathbf{C}_A$ is a subgroup of G', whence G' generates a finite-index subgroup of G, as required. □

Lemma 5.6. *Let H act properly and cocompactly on the* CAT(0) *cube complex* $\mathbf{C}$, *and let $g \in H$ be a rank-one element. Then for any simplex v of $\partial_{\triangle}\mathbf{X}$ not stabilized by g, the orbit $\langle g\rangle v$ is unbounded in $(\partial_{\triangle}\mathbf{C})^{(1)}$.*

Proof. If $v \in \partial_{\triangle}\mathbf{C}$ is an isolated 0-simplex not fixed by g, then $\langle g\rangle v$ is disconnected and therefore unbounded. Hence it suffices to consider a visible 0-simplex v that is not fixed by g. Let α be a combinatorial geodesic axis for g, and let γ be a ray representing v and emanating from a 0-cube of α. Suppose that there exists $M < \infty$ such that $g^n v$ is joined to v by a path of length at most M in $\partial_{\triangle}\mathbf{C}^{(1)}$, for all $n \in \mathbb{Z}$.

Then, applying the Flat Bridge Trick, we find for each $n \in \mathbb{Z}$ some $m \leq 2M$ and a sequence $F_0, \ldots, F_m$ of flat sectors such that $\widehat{F}_i \cap \widehat{F}_{i+1}$ is unbounded and path-connected for all i and such that $v \in \partial_{\triangle}\widehat{F}_0$ and $g^n v \in \partial_{\triangle}\widehat{F}_m$; moreover it is no loss of generality to assume that F_0 is always the same flat sector, as can be seen by applying the Flat Bridge Trick between any pair of F_0s obtained as above. Moreover, these flat sectors can be chosen, again using the Flat Bridge Trick, so that $g^n F_0 = F_m$ and so that γ contains a sub-ray lying in F_0. Hence for all n, the distance from α to F_m is uniformly bounded. It follows that there exists $\eta \geq 0$ such that there are hyperplanes H, H' crossing the subpath α_n of α subtended by γ and $g^n(\gamma)$, satisfying $\dot{d}(N(H) \cap \alpha, N(H') \cap \alpha) \geq |\alpha_n| M^{-1} - \eta$ and H, H' crossing a common $\widehat{F}_i$.

Since α is a rank-one periodic geodesic, there exists $p < \infty$ such that if H, H' are hyperplanes that cross α, either $H \cap H' = \emptyset$ or the subpath of α between the 1-cubes dual to H, H' has length at most p (see [26, Section 2]).

Note that if H, H' are hyperplanes crossing α_n, and H, H' both cross $\widehat{F}_i$, and H, H' do not cross, then $\dot{d}(N(H), N(H')) \leq q$ for some q depending on g (but independent of H, H' and F_i). Indeed, analyzing a minimal-area disc diagram bounded by geodesics in $N(H), N(H')$, the subtended part of α, and a geodesic in a hyperplane of F_i crossing H, H' (as in [26, Section 2]) shows that if H, H' can be chosen arbitrarily far apart, then either there are hyperplanes V, V', that cross α arbitrarily far apart and cross each other, or there are arbitrarily large isometric flat discs of the form $[0, N]^2$ with one side on α. This contradicts that g is a rank-one isometry.

Now, there must exist hyperplanes H, H' with $\dot{d}(N(H), N(H')) \geq |\alpha_n| M^{-1}$ that both cross $\widehat{F}_i$ for some i. If $|\alpha_n| > M \max\{p, q\}$, then H, H' cannot cross, and cannot cross a common flat sector, a contradiction. □

Remark 5.7. Let G act properly, cocompactly and essentially on $\mathbf{X}$, and suppose that G is algebraically thick of order 1 with respect to a finite collection $\mathbb{G} = \{H_A : A \in \mathcal{A}\}$ of quasiconvex, wide subgroups, as in Theorem 5.1. For each $A \in \mathcal{A}$, let S_A be the H_A-cocompact convex subcomplex constructed in the proof of Theorem 5.1. That proof shows that $\bigcup_{A \in \mathcal{A}, g \in G} \partial_{\triangle} S_A = \mathfrak{C}$ is positive-dimensional, connected, and G-invariant. Hence $\partial_{\triangle}\mathbf{X}$ has a positive-dimensional G-invariant component, namely that containing $\mathfrak{C}$. Moreover, Theorem 4.3 implies that $\partial_{\triangle}\mathbf{X}$ is disconnected, since G is not wide.

Now, since H_A acts cocompactly on S_A for all $A \in \mathcal{A}$, and each H_A is wide, each $\partial_\triangle S_A$ is connected by Theorem 4.3. Each $f^{-1}(\partial_\triangle S_A)$ is contained in the limit set of H_A, by the general fact that bi-infinite geodesics in proper, cocompact spaces are limits of sequences of periodic geodesics. Likewise, since $\{H_A : A \in \mathcal{A}\}$ generates a finite-index subgroup $G' \leq G$, by algebraic thickness, G' acts cocompactly on $\mathbf{X}$, which is the coarse union of G'-translates of the various S_A, and thus $f^{-1}\big(\bigcup_{A\in\mathcal{A},g\in G'}\big) = f^{-1}(\mathfrak{C})$ is contained in the limit set of G'.

This discussion shows that, if G is algebraically thick of order 1 relative to a finite collection $\{H_A\}$ of quasiconvex, wide subgroups, then $\partial_\triangle \mathbf{X}$ has a G-invariant component $\mathfrak{C}$, and a finite collection $\mathcal{A}$ of connected subcomplexes, satisfying hypotheses (1)–(3) of Theorem 5.4. This conclusion is used in the proof of Theorem 5.13.

The following characterization of convex hulls of flat sectors is immediate from the definitions.

Lemma 5.8. *Let $\mathbf{X}$ be as in Theorem 5.4. For $n \geq 2$, let $\mathcal{A}_n$ be the class of* CAT(0) *cube complexes $A \subseteq \mathbf{X}$ such that*

(1) *A contains an isometrically embedded cubical flat sector F satisfying $2 \leq \dim F \leq n$;*
(2) *every hyperplane of A crosses F.*

Then the convex hull of each cubical flat sector F in $\mathbf{X}$ belongs to $\mathcal{A}_n$ for $n = \dim \mathbf{X}$.

Lemma 5.9. *$\mathcal{A}_n$ is uniformly wide. Equivalently, $\{A^{(1)} : A \in \mathcal{A}_n\}$ is uniformly wide.*

Proof. The two assertions are equivalent since the collection of elements in $\mathcal{A}_n$ have uniformly bounded dimension and are thus each quasi-isometric to their 1-skeleta, with uniform quasi-isometry constants (see, e.g., [15, Lemma 2.2]).

Let $(A_i)_{i\geq 0}$ be a sequence of cube complexes in $\mathcal{A}_n$, and denote by d_i the standard path-metric on $A_i^{(1)}$. Recall that $\mathcal{A}_n$ is uniformly wide if and only if for any sequence $(a_i \in A_i)_{i\geq 0}$, any positive sequence $(s_i)_{i\geq 0}$ with $\lim_i s_i = \infty$, and any ultrafilter ω, the ultralimit $\lim_\omega \big(A_i, a_i, \frac{d_i}{s_i}\big)$ has no cut-point. We will prove that exhibit a uniform linear bound on the divergences of the $A_i^{(1)}$, from which the result then follows from [16, Proposition 1.1] which relates divergence and wideness.

Let $a, b, c \in A_i^{(0)}$, with $d_i(a, b) = m$ and $d_i(\{a, b\}, c) = r$. Choose $\delta \in \big(0, \frac{1}{2}\big)$ and $\kappa \geq 0$. Let μ be the median of a, b, c and let γ be a bi-infinite geodesic with $\gamma(0) = \mu$ and $\gamma(-t_a) = a, \gamma(t_b) = b$ for $t_a, t_b \in (0, m)$.

If $d_i(\mu, c) > \delta r - \kappa$, then the subpath of γ joining a to b has length m and avoids the $(\delta r - \kappa)$-ball about c.

Otherwise, $d_i(\mu, c) \leq \delta r - \kappa$, so that for any $t \in \mathbb{R}$ we have $d_i(\gamma(t), c) \geq t - \delta r + \kappa$.

Let $T = 2\max\{t_a, t_b\}$. Since $\delta < \frac{1}{2}$ we have $T \geq \delta r - \kappa$.

In the proof of [26, Lemma 6.5], it is shown that there exists a combinatorial path P connecting $\gamma(-T)$ to $\gamma(T)$, with the property that each point of P lies at distance at least T from μ, having length at most $5T + B$, where B counts a certain set of hyperplanes separating a or b from μ, whence $B \leq 2T$. Since $d(\mu, c) \leq \delta r - \kappa$, the triangle inequality implies that for each $p \in P$ we have $d(p, c) \geq T - \delta r - \kappa \geq \delta r - \kappa$. Thus concatenating P with the subpaths of γ from $\gamma(-T)$ to $\gamma(-t_a)$ and from $\gamma(t_b)$ to $\gamma(T)$ (which are each of length at most m), we get a path P' connecting a to b, which is of linear length (at most $11T \leq 22m$) and which avoids the $(\delta r - \kappa)$-ball about c.

Hence, the divergence of a, b, c is at most $22m$. Since the constants for divergence do not depend on i, it follows immediately that the ultralimit of the sequence $A_i^{(0)}$ does not have any cut-points. □

5.3. Strong algebraic thickness. Obtaining quadratic divergence bounds using the result in [3] requires strong thickness of order 1, which we obtain here as a consequence of strong algebraic thickness. First, we note that little stands between the conclusion of Theorem 5.4 and the conclusion of strong algebraic thickness of a cocompactly cubulated group G:

Proposition 5.10. *Let G act properly and cocompactly on the* CAT(0) *cube complex* $\mathbf{X}$*, and suppose that G is algebraically thick of order $n \geq 1$ with respect to a finite collection $\{H_A : A \in \mathcal{A}\}$ of quasiconvex subgroups, each of which is strongly algebraically thick of order at most $n-1$. Then G is strongly algebraically thick of order n relative to $\{H_A\}$.*

Hence $\mathbf{X}$ *and G are strongly thick of order n and have polynomial divergence of order at most $n + 1$.*

Proof. By hypothesis, each H_A is strongly algebraically thick of order $n - 1$. Moreover, since each H_A acts with an M_A-quasiconvex orbit on $\mathbf{X}$, each H_A is $M = \max_{A \in \mathcal{A}}$-quasiconvex in G. In particular, each H_A acts properly and cocompactly on a convex subcomplex $\mathbf{C}_A$ of $\mathbf{X}$ that is contained in the tubular M-neighborhood of the orbit $H_A x_o$, where x_o is a fixed 0-cube. Now, if $H_A, H_{A'}$ are among the given finite collection, then by algebraic thickness, there exists a sequence $A = A_0, \ldots, A_n = A'$ such that for $0 \leq i < n$, the intersection $H_{A_i} \cap H_{A_{i+1}}$ is infinite. This implies that $H_{A_i} x_o \cap H_{A_{i+1}} x_o$ is infinite, whence $\mathbf{C}_{A_i} \cap \mathbf{C}_{A_{i+1}}$ is unbounded, and hence path-connected, since the intersection of convex subcomplexes of $\mathbf{X}$ is again convex. Thus any geodesic segment starting and ending in $H_{A_i} x_o \cap H_{A_{i+1}} x_o$ lies inside of the M-neighborhood of $H_{A_i} x_o \cap H_{A_{i+1}} x_o$, whence $H_{A_i} \cap H_{A_{i+1}}$ is M-path-connected. Finally, $\langle\{H_A\}\rangle$ has finite index in G since G is algebraically thick relative to $\{H_A\}$. Thus G is strongly algebraically thick of order at most n. Obviously, if G were strongly algebraically thick of order $k < n$, then G would be thick of order k, a contradiction. Hence G is strongly algebraically thick of order exactly n.

It is now readily verified that $\mathbf{X}$ is strongly thick of order n relative to the collection $\{g\mathbf{C}_A : g \in G, A \in \mathcal{A}\}$. Thus $\mathbf{X}$, and G, have polynomial divergence of order at most $n + 1$ by Corollary 4.17 of [3]. □

Corollary 5.11. *Let G act properly and cocompactly on the fully visible* CAT(0) *cube complex $\mathbf{X}$, and suppose that $\partial_\triangle \mathbf{X}$ contains isolated* 0*-simplices and a G-invariant connected subcomplex $\mathfrak{C} = \bigcup_{g\in G, A\in\mathcal{A}} gA$, with $\mathcal{A}$ and the collection $\{H_A = \mathrm{Stab}_G(A) \colon A \in \mathcal{A}\}$ as in Theorem 5.4. Then G has quadratic divergence function.*

Proof. By Theorem 5.4, G is algebraically thick of order 1 relative to $\{H_A\}$, and thus strongly algebraically thick by Proposition 5.10, from which it also follows that the divergence of G is at most quadratic. On the other hand, if the divergence is subquadratic, then it is linear [31, Proposition 3.3], which implies that $\partial_\triangle \mathbf{X}$ is connected, contradicting the fact that the set of isolated 0-simplices is nonempty. □

Example 5.12 (the Croke–Kleiner example). The following example confirms that $\mathbf{X}$ satisfies the conclusions of Theorem 5.1, Theorem 5.4, and Corollary 5.11 when $\mathbf{X}$ is the universal cover of the Salvetti complex of a right-angled Artin group; here we have chosen the Croke–Kleiner group [13]. The same reasoning applies to any one-ended right-angled Artin group that is not a product, and these are known to be thick of order 1 and have quadratic divergence; see [4] and [2].

Let $\mathbf{X}$ be the universal cover of the Salvetti complex of the right-angled Artin group

$$G \cong \langle a, b, c, d \mid [a,b], [b,c], [c,d] \rangle.$$

(This group is studied by Croke-Kleiner in [13].) $\mathbf{X}$ decomposes as a tree T of spaces: the vertex-spaces are the obvious periodic 2-dimensional cubical flats whose edges are labeled by generators, and the edge-spaces are bi-infinite combinatorial geodesics representing cosets of $\langle a\rangle$, $\langle b\rangle$, $\langle c\rangle$, or $\langle d\rangle$.

Each flat F corresponding to a vertex of T is convex in $\mathbf{X}$, so $\partial_\triangle F$ embeds as a subcomplex in $\partial_\triangle \mathbf{X}$. Each F is labeled by a pair $(x, y) \in \{a, b, c, d\}^2$ of distinct generators corresponding to the labels of the 1-cubes of the constituent squares of F. The x-labeled combinatorial geodesics in F represent a pair of 0-simplices in $\partial_\triangle F$, and the same is true of the y-labeled geodesics, and $\partial_\triangle F$ is a 4-cycle, being the join of the x-labeled 0-simplices and the y-labeled 0-simplices.

Now, fix a root of T and let F_0 be the corresponding flat; for concreteness, take F_0 to be a flat labeled (a, b). For each $n \geq 0$, let $\mathcal{S}_n$ be the set of flats that correspond to vertices of T at distance n from the vertex corresponding to F_0. Each F corresponding to a vertex in $\mathcal{S}_1$ is labeled (b, c), and for each such F, $\partial_\triangle \mathbf{X}$ contains a copy of $\partial_\triangle F$ attached to $\partial_\triangle F_0$ along the pair of b-labeled 0-simplices. If $F, F' \in \mathcal{S}_1$ are distinct, then the images of their c-labeled 0-simplices are distinct. By induction on n, one checks that the union of the images of all $\partial_\triangle F$ is connected; this union is clearly G-invariant.

Now, for each geodesic ray $\bar{\gamma}$ in T, there exists a rank-one geodesic ray γ in $\mathbf{X}$ such that γ has nonempty intersection with exactly those F that correspond to vertices of $\bar{\gamma}$. Conversely, each rank-one ray in $\mathbf{X}$ projects to a geodesic ray in T, and two rays projecting to the same ray in T represent the same 0-simplex of $\partial_\triangle \mathbf{X}$. Hence $\partial_\triangle \mathbf{X}$ contains exactly one isolated 0-simplex for each point of $\partial_\triangle T$, as shown in Figure 5.

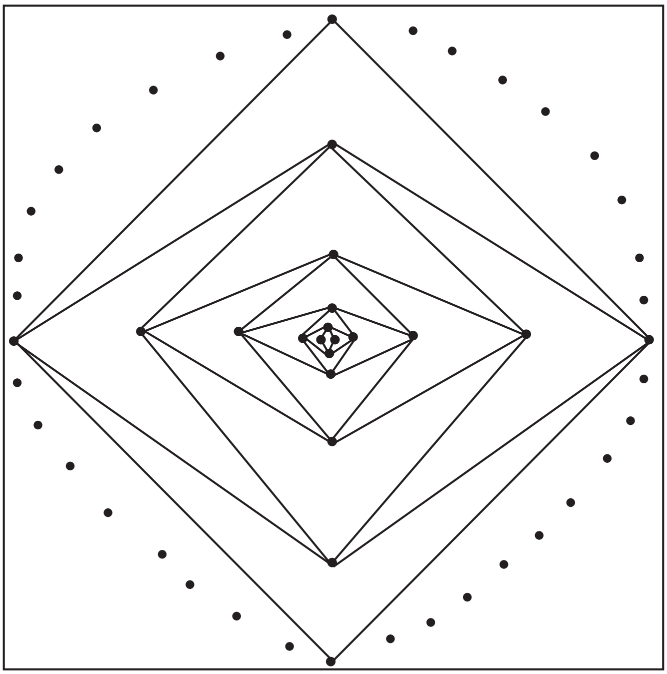

Figure 5. Part of the simplicial boundary of the universal cover of the Salvetti complex of the Croke-Kleiner group.

5.4. Necessary and sufficient conditions for thickness of order 1. The following is a culmination of the results of this section.

Theorem 5.13. *Let G act properly and cocompactly by isometries on the fully visible* CAT(0) *cube complex* $\mathbf{X}$*. If G is algebraically thick of order* 1 *relative to a collection of quasiconvex wide subgroups, then $\partial_\triangle \mathbf{X}$ is disconnected and contains a positive-dimensional, G-invariant connected component. Conversely, if $\partial_\triangle \mathbf{X}$ is disconnected, and has a positive-dimensional G-invariant component, then $\mathbf{X}$ is thick of order* 1 *relative to a collection of wide, convex subcomplexes, whence G is thick of order* 1.

Moreover, G is strongly algebraically thick of order 1 *if and only if $\partial_\triangle \mathbf{X}$ is disconnected and has a positive-dimensional, G-invariant connected subcomplex $\mathfrak{C} = \bigcup_{A \in \mathcal{A}, g \in G} gA$, where $\mathcal{A}$ is a finite collection of bounded subcomplexes such that*:

(1) *each* $\mathrm{Stab}(A)$ *acts on* $\mathbf{X}$ *with a quasiconvex orbit*;

(2) *for each $A \in \mathcal{A}$, $f^{-1}(A)$ belongs to the limit set of* $\mathrm{Stab}(A)$;

(3) *$f^{-1}(\mathfrak{C})$ is contained in the limit set of* $\langle\{\mathrm{Stab}(A) : A \in \mathcal{A}\}\rangle$.

Proof. The first assertion is the content of Theorem 5.1. Remark 5.7 shows that $\mathfrak{C}$ satisfies (1) – (3). The converse is Theorem 5.4, with the equivalence of strong algebraic thickness of order 1 is equivalent to algebraic thickness of order 1 relative to quasiconvex wide subgroups being established by Proposition 5.10. □

6. Characterizations of thickness and relative hyperbolicity via the Tits boundary

When regarding $\mathbf{X}$ as a combinatorial object, it is natural to use the simplicial boundary; as a CAT(0) space, $\mathbf{X}$ also has a Tits boundary $\partial_T \mathbf{X}$. By viewing each simplex of $\partial_\triangle \mathbf{X}$ as a right-angled spherical simplex whose 1-simplices have length $\frac{\pi}{2}$, one realizes $\partial_\triangle \mathbf{X}$ as a piecewise-spherical CAT(1) space. Proposition 3.37 of [26] asserts that, when $\mathbf{X}$ is fully visible, there is an isometric embedding $I\colon \partial_\triangle \mathbf{X} \to \partial_T \mathbf{X}$ such that $\partial_T \mathbf{X} \subseteq \mathcal{N}_{\frac{\pi}{2}}(\mathrm{im}\, I)$. (The map I is an isometric embedding with respect to the piecewise-spherical CAT(1) metric on $\partial_\triangle \mathbf{X}$.) This map sends each 0-simplex v – which, by full visibility, is represented by some CAT(0) geodesic ray γ – to the point of $\partial_\triangle \mathbf{X}$ represented by γ. It follows that I is G-equivariant, and induces a bijection from the set of components (respectively, the set of isolated 0-simplices) of $\partial_\triangle \mathbf{X}$ to the set of components (respectively, the set of isolated points) of the Tits boundary.

Moreover, I is a section of a surjective map $R: \partial_T \mathbf{X} \to \partial_\triangle \mathbf{X}$ such that the R-preimage of any point is connected, has diameter at most $\frac{\pi}{2}$, and consists of points represented by rays that represent the same simplex in $\partial_\triangle \mathbf{X}$. Furthermore, in the cocompact case, if the simplicial boundary contains infinitely many isolated points, then so does the Tits boundary.

Corollary 6.1. *Let the group G act geometrically on the fully visible* CAT(0) *cube complex* $\mathbf{X}$.

Suppose that G is hyperbolic relative to a collection $\mathbb{P}$ of peripheral subgroups. Then $\partial_T \mathbf{X}$ consists of a nonempty set of disjoint closed balls of radius less than $\frac{\pi}{2}$, together with a collection $\{g\mathbf{T}_P : P \in \mathbb{P}, g \in G\}$ of subspaces such that $\mathrm{Stab}(\mathbf{T}_P) = P$ *for all $P \in \mathbb{P}$ and $g\mathbf{T}_P \cap h\mathbf{T}_{P'} = \emptyset$ unless $P = P'$ and $gh^{-1} \in P$.*

Conversely, suppose that the set of isolated points of $\partial_T \mathbf{X}$ is nonempty, and that there is a pairwise-disjoint, G-finite collection $G(\{\mathbf{S}_i\}_{i=1}^k)$ of subspaces of $\partial_T \mathbf{X}$ such that each $P_i = \mathrm{Stab}_G(\mathbf{S}_i)$ is quasiconvex and of infinite index in G, each $\mathbf{S}_i$ contains the limit set for the action of P_i on $\partial_\infty \mathbf{X}$, and every point of $\partial_T \mathbf{X}$ lies in some $g\mathbf{S}_i$ or in some isolated ball of radius less than $\frac{\pi}{2}$. Then G is hyperbolic relative to $\{P_i\}_{i=1}^k$.

Proof. If G is relatively hyperbolic, then each $\mathbf{T}_P = R^{-1}(\mathbf{S}_P)$, where $\mathbf{S}_P$ is one of the subcomplexes arising from Theorem 3.1. It is easily verified that the resulting family of subspaces has the desired properties. Every other point in $\partial_T \mathbf{X}$ lies in $R^{-1}(p)$ for some isolated 0-simplex p. Any two points in the preimage of the same isolated point correspond to rays that are almost-equivalent and thus represent points at Tits distance strictly less than $\frac{\pi}{2}$.

Conversely, suppose that

$$\partial_T \mathbf{X} = \mathbf{B} \cup \Big(\bigsqcup_{g \in G, P \in \mathbb{P}} g\mathbf{T}_P \Big),$$

where $\mathbf{B}$ is the disjoint union of the isolated balls. Then for each g, P, let $g\mathbf{S}_P = R(g\mathbf{T}_P) = gR(\mathbf{T}_P)$. This is a P^g-invariant subcomplex, and any two of these subcomplexes are disjoint. For each $b \in \mathbf{B}$, $R(b)$ must be an isolated 0-simplex, and it follows from Theorem 3.7 that G is hyperbolic relative to $\mathbb{P}$. □

Corollary 6.2. *Let G act properly and cocompactly on the fully visible* CAT(0) *cube complex* $\mathbf{X}$. *If G is algebraically thick of order* 1, *then $\partial_T \mathbf{X}$ has a proper G-invariant connected component.*

Conversely, if $\partial_T \mathbf{X}$ has this feature, then G is thick of order 1 *relative to a collection of wide subsets. Suppose, in addition, that $\partial_T \mathbf{X}$ has a connected G-invariant subspace $\mathfrak{C} = \bigcup_{g \in G, A \in \mathcal{A}}$, where $\mathcal{A}$ is a finite set of connected subspaces satisfying*:

(1) *for all* $A \in \mathcal{A}$*, the stabilizer* H_A *of* A *is quasiconvex*;

(2) *for all* $A \in \mathcal{A}$*, the limit set of* H_A *(in the cone topology on* $\partial_\infty \mathbf{X}$*) contains* A;

(3) *the limit set of* $\langle \{H_A : A \in \mathcal{A}\} \rangle$ *contains* $\mathfrak{C}$.

Then G *is strongly algebraically thick of order* 1 *relative to a collection of quasiconvex, wide subgroups, and* G *has polynomial divergence function of order exactly* 2.

Proof. If G is algebraically thick of order 1, then $\partial_\triangle \mathbf{X}$ has a G-invariant connected subspace $\mathfrak{C}'$ that is properly contained in $\partial_\triangle \mathbf{X}$, by Theorem 5.1. Let $\mathfrak{C} = R^{-1}(\mathfrak{C}')$. The definition of R implies that $\mathfrak{C}$ is connected: each simplex has connected R-preimage. Also, $\mathfrak{C}$ does not contain all of $\partial_T \mathbf{X}$ since R is surjective and distance-nonincreasing, and $\partial_\triangle \mathbf{X}$ has more than one component.

Conversely, if $\mathfrak{C}$ is a G-invariant connected subspace of $\partial_T \mathbf{X}$, then $R(\mathfrak{C})$ is a G-invariant connected subspace of $\partial_\triangle \mathbf{X}$, whence $\mathbf{X}$ is thick by Theorem 5.4. It is easily verified that $\{R(A) : A \in \mathcal{A}\}$ satisfies the hypotheses of Theorem 5.4, from which strong algebraic thickness of order 1 follows. □

7. Cubulated groups with arbitrary order of thickness

The goal of this section is to produce cocompactly cubulated groups of any order of thickness; in fact, the groups we produce will be strongly algebraically thick of the desired order.

Notation 7.1. For $n \geq 1$, we will let $\mathbb{G}_n$ denote the class of groups such that each $G \in \mathbb{G}_n$ acts properly and cocompactly on a CAT(0) cube complex, is strongly algebraically thick of order at most n, and has polynomial divergence of order $n + 1$.

Note that $\mathbb{G}_n$ does not contain any groups of dimension 1, since a 1-dimensional CAT(0) cube complex is a tree, and hence such a group could not have polynomial divergence.

Lemma 7.2. *For each dimension* $k > 1$*, the class* $\mathbb{G}_1$ *has an infinite subclass of pairwise non-quasi-isometric groups of geometric dimension* k.

Proof. Let Γ be a connected graph with at least two vertices that does not decompose as a nontrivial join. The universal cover $\mathbf{X}_\Gamma$ of the Salvetti complex of the associated right-angled Artin group $G(\Gamma)$ is a combinatorially geodesically complete CAT(0) cube complex on which $G(\Gamma)$ acts properly, cocompactly, and essentially.

According to [2], the right-angled Artin group $G(\Gamma)$ is algebraically thick of order 1 and has quadratic divergence, since Γ is not a nontrivial join.

A connected graph Γ is *atomic* if it has no leaves, if its girth is at least 5, and no vertex-star is separating. It is shown in [7] that, if Γ_1, Γ_2 are atomic graphs, then $G(\Gamma_1)$ and $G(\Gamma_2)$ are quasi-isometric if and only if $\Gamma_1 \cong \Gamma_2$. Since there are obviously infinitely many isomorphism types of finite atomic graphs, it follows that $\mathbb{G}_1$ contains infinitely many pairwise non-quasi-isometric groups each of dimension 2.

For each $k > 2$, the irreducible k-tree groups constructed in [6] provide an infinite family of k-dimensional right-angled Artin groups which are all algebraically thick of order 1. Further, it was shown in [6] that this family contains infinitely many pairwise non-quasi-isometric groups. □

Theorem 7.3. *For each dimension $k > 1$ and each $n \geq 1$, the class $\mathbb{G}_n$ contains an infinite class of pairwise non-quasi-isometric groups of geometric dimension k.*

Proof. The claim holds when $n = 1$ by Lemma 7.2. For $n \geq 1$, by induction there exists a group $G_n \in \mathbb{G}_n$ acting freely, cocompactly, and essentially on a k-dimensional CAT(0) cube complex $\mathbf{X}_n$ that is algebraically thick of order n and has divergence of order $n + 1$.

Construction of G_{n+1} and $\mathbf{X}_{n+1}$. By [15, Corollary B], there exists $g \in G_n$ acting on $\mathbf{X}_n$ as a rank-one isometry. Let $\gamma \subset \mathbf{X}_n$ be a CAT(0) geodesic axis for g. By induction, we can choose g so that γ has divergence of order at least $n + 1$. Since g is rank-one, the cubical convex hull K_n of γ lies in a finite neighborhood of γ. Hence the stabilizer $C_n \leq G$ of K_n contains $\langle \gamma \rangle$ as a finite-index subgroup.

Let $G_{n+1} = G_n *_{C_n} G_n$, and denote by T_n the associated Bass-Serre tree. The space $\mathbf{X}_{n+1}$ is defined to be the total space of the tree of spaces whose underlying tree is T_n, whose vertex-spaces are copies of $\mathbf{X}_n$ and whose edge-spaces are copies of K_n corresponding to cosets of C_n. The attaching maps are inclusions. Since $\mathbf{X}_{n+1}$ is obtained by gluing CAT(0) cube complexes along convex subcomplexes, it is nonpositively curved and therefore a CAT(0) cube complex, by virtue of being simply connected. There is an obvious free, cocompact, essential action of G_{n+1} on $\mathbf{X}_{n+1}$, where the vertex-stabilizers are conjugate to G_n and the edge-stabilizers are conjugate to C_n.

We remark that collapsing each edge-space $K_n \times [-1, 1]$ to K_n within $\mathbf{X}_{n+1}$ yields a new G_{n+1}-cocompact CAT(0) cube complex $\mathbf{X}'_{n+1}$ with $\dim \mathbf{X}'_{n+1} = \dim \mathbf{X}_n$. Although we work in $\mathbf{X}_{n+1}$ for convenience, this observation shows, by induction on n, that G_{n+1} can always be chosen to act properly and cocompactly on a CAT(0) cube complex of dimension $\dim \mathbf{X}_1$, where $\mathbf{X}_1$ corresponds to some $G_1 \in \mathbb{G}_1$. To prove that $\mathbb{G}_n$ contains infinitely many quasi-isometry types of k-dimensionally cocompactly cubulated groups, one needs only to add to the induction hypothesis that $\dim \mathbf{X}_n = k$ and note that Lemma 7.2 has already accounted for the base case.

An upper bound on order of thickness. By Lemma 7.4 below, $h\mathbf{X}_n$ is a convex subcomplex of $\mathbf{X}_{n+1}$ for each $h \in G$, and $h\mathbf{X}_n$ is thick of order n.

By construction, $\mathbf{X}_{n+1}$ is contained in the 1-neighborhood of $G_{n+1}\mathbf{X}_n$. Therefore, for any $x, y \in \mathbf{X}_{n+1}$, there exist $h_0, h_m \in G$ such that $\dot{d}(x, h_0\mathbf{X}_n) \leq 1$ and $\dot{d}(y, h_m\mathbf{X}_n) \leq 1$. Let $h_0\mathbf{X}_n, h_1\mathbf{X}_n, \ldots, h_m\mathbf{X}_n$ be the sequence of vertex-spaces corresponding to the sequence of vertices in the projection to T_n of a geodesic in $\mathbf{X}_{n+1}$ joining x to y. By construction $h_i\mathbf{X}_n \cap h_{i+1}\mathbf{X}_n$ is a translate of K_n for $0 \leq i \leq m-1$. Since K_n is unbounded, the set $\{h\mathbf{X}_n : h \in G_{n+1}\}$ is thickly connecting, whence $\mathbf{X}_{n+1}$, and therefore G_{n+1}, is thick of order at most $n+1$. Since each translate of $\mathbf{X}_n$ is stabilized by a conjugate of one of the two vertex groups in the splitting $G_{n+1} \cong G_n *_{C_n} G_n$, and K_n has infinite stabilizer, we see that G_{n+1} is algebraically thick of order at most $n+1$.

A lower bound on divergence. By Lemma 7.5, C_n is a malnormal subgroup of G_{n+1} and the action of G_{n+1} on T_n is acylindrical by Lemma 7.7. The proof of [3, Proposition 5.2] can now be repeated almost verbatim to show that for any $g' \in G_{n+1}$ acting axially on T_n, any geodesic axis for g' in $\mathbf{X}_{n+1}$ has divergence of order at least $n+2$. The only difference is that the "separating geodesics" discussed in [3] are replaced here by tubular neighborhoods of γ that contain K_n and therefore separate $\mathbf{X}_{n+1}$.

Infinitely many quasi-isometry types. Denote by A and B the copies of G_n that are vertex groups of the splitting $G_{n+1} \cong G_n *_{C_n} G_n$, so that $\{A, B\}$ is a set of subgroups showing that G_{n+1} has order of algebraic thickness at most $n+1$. Let $G'_{n+1} \in \mathbb{G}_{n+1}$ and define $A', B' \leq \mathbb{G}_{n+1}$ analogously (so that A' and B' are both isomorphic to some $G'_n \in \mathbb{G}_n$). If $q : G_{n+1} \to G'_{n+1}$ is a quasi-isometry, then $q(A)$ and $q(B)$ are respectively coarsely equal to A and B (or B and A), as in the construction in [3, Section 5], because of quasi-isometry invariance of the splitting over $\mathbb{Z}$, which follows from [39, Theorem 7.1]. Hence G_n and G'_n are quasi-isometric, and therefore the set of quasi-isometry types represented in $\mathbb{G}_{n+1}$ has cardinality at least that of the set of quasi-isometry types represented in $\mathbb{G}_1$, and the latter is infinite by Lemma 7.2. □

Lemma 7.4. *$\mathbf{X}_n$ and K_n are convex subcomplexes of $\mathbf{X}_{n+1}$.*

Proof. $\mathbf{X}_{n+1}$ is the union of copies of $\mathbf{X}_n$ and copies of $K_n \times [-1, 1]$. We denote by K_n the subspace $K_n \times \{-1\}$ of $\mathbf{X}_n$.

Since $\mathbf{X}_{n+1}$ is CAT(0), it is sufficient to verify that $\mathbf{X}_n$ and K_n are locally convex. Suppose to the contrary that s is a 2-cube whose boundary path is a 4-cycle $abcd$ with $ab \subset K_n$. If $s \subset \mathbf{X}_n$, then cd is a combinatorial geodesic segment in $\mathbf{X}_n$ starting and ending on K_n, whence $s \subset K_n$ since K_n is convex in $\mathbf{X}_n$. Otherwise, s lies in the copy of $K_n \times [-1, 1]$ projecting to the edge of T_n corresponding to K_n. The unique possibility in this case is that $s \subset K_n$. Hence K_n is convex.

A 2-cube with two consecutive boundary 1-cubes in $\mathbf{X}_n$ has two consecutive boundary 1-cubes in some $\mathrm{Stab}(\mathbf{X}_n)$-translate of K_n, and must therefore lie in $K_n \subset \mathbf{X}_n$. Thus $\mathbf{X}_n$ is convex. □

Lemma 7.5. *C_n is a malnormal subgroup of G.*

Proof. If C_n fails to be malnormal, then there exists $h \in G_{n+1} - C_n$ and nonzero integers r, s such that $g^r = hg^s h^{-1}$. Since C is quasi-isometrically embedded in G_{n+1}, we must have $|r| = |s|$, so that without loss of generality, $r = s$ and h is a hyperbolic isometry of $\mathbf{X}_{n+1}$. There is thus a $\langle g^r, h\rangle$-invariant flat in $\mathbf{X}_{n+1}$ coarsely containing γ, and this contradicts the fact that g is a rank-one isometry of $\mathbf{X}_n$ and that $\mathbf{X}_{n+1}$ is a tree of spaces where the vertex spaces are copies of $\mathbf{X}_n$ and the edge spaces are each contained in finite neighborhoods of copies of the axis of g. □

Definition 7.6 (acylindrical). The isometric action of the group G on the graph Y is *acylindrical* if for some $\ell > 0$, there exists $M < \infty$ such that $|\mathrm{Stab}(x) \cap \mathrm{Stab}(y)| \leq M$ whenever x and y are at distance at least ℓ in Y.

Lemma 7.7. *The action of G_{n+1} on T_n is acylindrical.*

Proof. Let x, y be vertices corresponding to $h_x G_n$ and $h_y G_n$, with $d_{T_n}(x, y) = 2$. Let z be the midpoint of the unique geodesic joining x to y, and denote by $h_z G_n$ the corresponding coset. If $k \in G_n^{h_x} \cap G_n^{h_y}$, then k stabilizes two distinct edges in T_n, corresponding to distinct translates of K_n in $h_z \mathbf{X}_n$. Lemma 7.5 implies that $k = 1$.

If $d_{T_n}(x, y) > 2$, then since geodesics in trees are unique, there exist x', y' between x, y such that $d_{T_n}(x', y') = 2$ and every element of G_{n+1} stabilizing x and y must also stabilize x' and y'. □

References

[1] I. Agol, The virtual haken conjecture. With an appendix by I. Agol, D. Groves, and J. Manning. *Doc. Math.* **18** (2013), 1045–1087. Zbl 1286.57019 MR 3104553

[2] J. Behrstock and R. Charney, Divergence and quasimorphisms of right-angled Artin groups. *Math. Ann.* **352** (2012), no. 2, 339–356. Zbl 1251.20036 MR 2874959

[3] J. Behrstock and C. Druțu, Divergence, thick groups, and short conjugators. *Illinois J. Math.* **58** (2014), no. 4, 939–980. Zbl 06510955 MR 3421592

[4] J. Behrstock, C. Druțu, and L. Mosher, Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Math. Ann.* **344** (2009), no. 3, 543–595. Zbl 1220.20037 MR 2501302

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

[6] J. Behrstock, T. Januszkiewicz, and W. Neumann, Quasi-isometric classification of high dimensional right angled artin groups. *Groups Geom. Dyn.* **4** (2010), no. 4, 681–692. Zbl 1226.20033 MR 2727658

[7] M. Bestvina, B. Kleiner, and M. Sageev, The asymptotic geometry of right-angled Artin groups. I. *Geom. Topol.* **12** (2008), no. 3, 1653–1699. Zbl 1203.20038 MR 2421136

[8] B. H. Bowditch, Relatively hyperbolic groups. *Internat. J. Algebra Comput.* **22** (2012), no. 3, article id. 1250016, 66 pp. Zbl 1259.20052 MR 2922380

[9] M. R. Bridson, Geodesics and curvature in metric simplicial complexes. In E. Ghys, A. Haefliger, and A. Verjovsky (eds.), *Group theory from a geometrical viewpoint.* (Trieste, 1990). World Scientific, Singapore, 1991, 373-463. Zbl 0844.53034 MR 1170372

[10] R. Charney and M. W. Davis, Finite $K(\pi,1)$s for Artin groups. In F. Quinn (ed.), *Prospects in topology.* (Princeton, N.J., 1994) Annals of Mathematics Studies, 138. Princeton University Press, Princeton, NJ, 1995, 110–124. Zbl 0930.55006 MR 1368655

[11] R. Charney and M W. Davis, The $K(\pi,1)$-problem for hyperplane complements associated to infinite reflection groups. *J. Amer. Math. Soc.* **8** (1995), no. 3, 597–627. Zbl 0833.51006 MR 1303028

[12] V. Chepoi, Graphs of some CAT(0) complexes. *Adv. in Appl. Math.* **24** (2000), no. 2, 125–179. Zbl 1019.57001 MR 1748966

[13] C. Croke and B. Kleiner, Spaces with nonpositive curvature and their ideal boundaries. *Topology* **39** (2000), no. 3, 549–556. Zbl 0959.53014 MR 1746908

[14] I. Chatterji and G. Niblo, From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.* **15** (2005), no. 5-6, 875–885. Zbl 1107.20027 MR 2197811

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

[16] C. Druțu, S. Mozes, and M. Sapir, Divergence in lattices in semisimple Lie groups and graphs of groups. *Trans. Amer. Math. Soc.* **362** (2010), no. 5, 2451–2505. Zbl 1260.20065 MR 2584607

[17] C. Druţu and M. Sapir, Tree-graded spaces and asymptotic cones of groups. With an appendix by D. Osin and M. Sapir. *Topology* **44** (2005), no. 5, 959–1058. Zbl 1101.20025 MR 2153979

[18] P. Dani and A. Thomas, Divergence in right-angled coxeter groups. *Trans. Amer. Math. Soc.* **367** (2015), no. 5, 3549–3577. Zbl 06429021 MR 3314816

[19] S. M. Gersten, Divergence in 3-manifold groups. *Geom. Funct. Anal.* **4** (1994), no. 6, 633–647. Zbl 0841.57023 MR 1302334

[20] S. M. Gersten, Quadratic divergence of geodesics in CAT(0) spaces. *Geom. Funct. Anal.* **4** (1994), no. 1, 37–51. Zbl 0809.53054 MR 1254309

[21] M. Gromov, Hyperbolic groups. In S. M. Gersten (ed.), *Essays in group theory*. Mathematical Sciences Research Institute Publications, 8. Springer-Verlag, New York etc., 1987, 75–263. Zbl 0634.20015 MR 0919829

[22] M. Gromov, Asymptotic invariants of infinite groups. In G. Niblo and M. Roller (eds.), *Geometric group theory*. Vol. 2. (Sussex, 1991) London Mathematical Society Lecture Note Series, 182. Cambridge University Press, Cambridge, 1993, 1–295. Zbl 0841.20039 MR 1253544

[23] F. Haglund, Isometries of CAT(0) cube complexes are semi-simple. Preprint 2007. arXiv:0705.3386 [math.GR]

[24] M. F. Hagen, Cocompactly cubulated crystallographic groups. *J. Lond. Math. Soc.* (2) **90** (2014), no. 1, 140–166. Zbl 06327931 MR 3245140

[25] M. F. Hagen, Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.* **7** (2014), no. 2, 385–418. Zbl 06366501 MR 3217625

[26] M. F. Hagen, The simplicial boundary of a CAT(0) cube complex. *Algebr. Geom. Topol.* **13** (2013), no. 3, 1299–1367. Zbl 1267.05302 MR 3071129

[27] G. Ch. Hruska and B. Kleiner, Hadamard spaces with isolated flats. With an appendix by the authors and M. Hindawi. *Geom. Topol.* **9** (2005), 1501–1538. Zbl 1087.20034 MR 2175151

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

[29] Ch. Hruska and D. T. Wise, Finiteness properties of cubulated groups. *Compos. Math.* **150** (2014), no. 3, 453–506. Zbl 3187627 MR 06293660

[30] T. Hsu and D. T. Wise, Cubulating malnormal amalgams. *Invent. Math.* **199** (2015), no. 2, 293–331. Zbl 1320.20039 MR 3302116

[31] M. Kapovich and B. Leeb, 3-manifold groups and nonpositive curvature. *Geom. Funct. Anal.* **8** (1998), no. 5, 841–852. Zbl 0915.57009 MR 1650098

[32] I. Leary, *J. Topol.* **6** (2013), no. 1, 251–284. Zbl 06272201 MR 3029427

[33] Y. Liu, Virtual cubulation of nonpositively curved graph manifolds. *J. Topol.* **6** (2013), no. 4, 793–822. Zbl 1286.57002 MR 3145140

[34] N. Macura, CAT(0) spaces with polynomial divergence of geodesics. *Geom. Dedicata* **163** (2013), 361–378. Zbl 1329.57002 MR 3032700

[35] G. Moussong, *Hyperbolic Coxeter groups*. Ph.D. thesis, Ohio State University, Columbus, OH, 1987.

[36] B. Nica, Cubulating spaces with walls. *Algebr. Geom. Topol.* **4** (2004), 297–309. Zbl 1131.20030 MR 2059193

[37] G. A. Niblo and L. D. Reeves, Coxeter groups act on CAT(0) cube complexes. *J. Group Theory* **6** (2003), no. 3, 399–413. Zbl 1068.20040 MR 1983376

[38] Y. Ollivier and D. T. Wise, Cubulating random groups at density $< \frac{1}{6}$. *Trans. Amer. Math. Soc.* **363** (2011), no. 9, 4701–4733. Zbl 1277.20048 MR 2806688

[39] P. Papasoglu, Quasi-isometry invariance of group splittings. *Ann. of Math.* (2) **161** (2005), no. 2, 759–830. Zbl 1129.20027 MR 2153400

[40] P. Przytycki and D. T. Wise, Graph manifolds with boundary are virtually special. *J. Topol.* **7** (2014), no. 2, 419–435. Zbl 1331.57023 MR 3217626

[41] M. A. Roller, Poc-sets, median algebras and group actions. An extended study of Dunwoody's construction and Sageev's theorem. Preprint 1998. http://www.personal.soton.ac.uk/gan/Roller.pdf

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

[43] M. Sageev and D. T. Wise, Cores for quasiconvex actions. *Proc. Amer. Math. Soc.* **143** (2015), no. 7, 2731–2741. Zbl 06428953 MR 3336599

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

[45] D. T. Wise, Cubulating small cancellation groups. *Geom. Funct. Anal.* **14** (2004), no. 1, 150–214. Zbl 1071.20038 MR 2053602

Received May 7, 2013

Jason Behrstock, Lehman College and The Graduate Center,
City University of New York, 365 Fifth Avenue, New York, NY 10016, USA

e-mail: jason.behrstock@lehman.cuny.edu

Mark F. Hagen, DPMMS, University of Cambridge, Wilberforce Road,
Cambridge CB3 0WB, UK

e-mail: markfhagen@gmail.com