

The simplicial boundary of a CAT(0) cube complex

MARK F HAGEN

For a CAT(0) cube complex \mathbf{X} , we define a simplicial flag complex $\partial_{\Delta}\mathbf{X}$, called the *simplicial boundary*, which is a natural setting for studying nonhyperbolic behavior of \mathbf{X} . We compare $\partial_{\Delta}\mathbf{X}$ to the Roller, visual and Tits boundaries of \mathbf{X} , give conditions under which the natural CAT(1) metric on $\partial_{\Delta}\mathbf{X}$ makes it isometric to the Tits boundary, and prove a more general statement relating the simplicial and Tits boundaries. The simplicial boundary $\partial_{\Delta}\mathbf{X}$ allows us to interpolate between studying geodesic rays in \mathbf{X} and the geometry of its *contact graph* $\Gamma\mathbf{X}$, which is known to be quasi-isometric to a tree, and we characterize essential cube complexes for which the contact graph is bounded. Using related techniques, we study divergence of combinatorial geodesics in \mathbf{X} using $\partial_{\Delta}\mathbf{X}$. Finally, we rephrase the rank-rigidity theorem of Caprace and Sageev in terms of group actions on $\Gamma\mathbf{X}$ and $\partial_{\Delta}\mathbf{X}$ and state characterizations of cubulated groups with linear divergence in terms of $\Gamma\mathbf{X}$ and $\partial_{\Delta}\mathbf{X}$.

05C25, 20F65, 57M99

Introduction

Since their introduction as a class of examples by Gromov in [20], CAT(0) cube complexes have become increasingly ubiquitous in group theory, and the class of groups acting on cube complexes is vast. The original examples come from Bass–Serre theory, since trees are 1–dimensional cube complexes, but, in the sense of Gromov’s density model, many more groups act on cube complexes than split (compare the results of Ollivier and Wise in [34] to those of Dahmani, Guirardel and Przytycki in [14]). Sageev provided a general means of obtaining actions on cube complexes by constructing, in [36], a G –cube complex from a *semisplitting* of the group G , which notion arose in work of Houghton and Scott [25; 37]. Many groups have been cubulated using Sageev’s construction and related techniques, including Coxeter groups (see Niblo and Reeves [32]), Artin groups of type FC (see Charney and Davis [8]), diagram groups, including Thompson’s group V (see Farley [16; 17]), small-cancellation groups (see Wise [39]), random groups at sufficiently low density in Gromov’s model [34] and groups with quasiconvex hierarchies, including one-relator groups with torsion, limit groups and fundamental groups of Haken hyperbolic 3–manifolds (see Wise [40]).

Right-angled Artin groups were shown to act on cube complexes by more direct means: the Salvetti complex of a right-angled Artin group was shown to be a nonpositively curved cube complex by Charney and Davis [7].

The structure of the cube complex \mathbf{X} is governed by the *hyperplanes*, which are convex subspaces, described independently by Sageev and Chepoi in [36; 10], that separate \mathbf{X} into exactly two complementary components. In this sense, CAT(0) cube complexes are “generalized trees”. Another salient property of a tree is that every geodesic triangle is a geodesic tripod. Generalizing this leads to the notion of a *median graph*, and in fact cube complexes and median graphs are essentially equivalent generalizations of trees: the 1–skeleton of \mathbf{X} is a median graph, and any median graph is the 1–skeleton of a uniquely determined cube complex.

It is natural to focus on how the hyperplanes interact. The *contact graph* $\Gamma\mathbf{X}$ of \mathbf{X} , introduced by the author in [21], is the intersection graph of the family of hyperplane carriers, and was shown in that paper to be quasi-isometric to a tree. The contact graph $\Gamma\mathbf{X}$ naturally generalizes the *crossing graph* $\Delta\mathbf{X}$ of \mathbf{X} , which is the intersection graph of the hyperplanes themselves, and is a spanning subgraph of $\Gamma\mathbf{X}$. Now, whenever \mathbf{X} decomposes as a nontrivial product, $\Gamma\mathbf{X}$ decomposes as a join and is thus a quasitree simply by virtue of being bounded. However, there are numerous infinite examples for which $\Gamma\mathbf{X}$ is bounded and \mathbf{X} does not have a cubical product structure; a natural example is the smallest subcomplex of the standard tiling of \mathbb{R}^2 by 2–cubes that contains $\{(x_1, x_2) \mid x_2 < x_1\}$, whose contact graph has diameter three. The motivating question in this paper is: *Under what geometric conditions on \mathbf{X} is $\Gamma\mathbf{X}$ (or $\Delta\mathbf{X}$) bounded, and what relevance does this have for groups acting on cube complexes?*

Studying this question led to the introduction of the *simplicial boundary* $\partial_\Delta\mathbf{X}$ of \mathbf{X} , a simplicial complex “at infinity” encoding much information about the nonhyperbolic behavior of $\partial_\Delta\mathbf{X}$. Our goal in this paper is to understand the simplicial boundary and to indicate some of its uses in cubical geometry.

We now summarize our results, which are related to this question, and briefly discuss our methods. Since the key objects in this paper are defined independently of the CAT(0) geometry, we mostly work in the 1–skeleton, taking advantage of the fact that it is a median graph (see the discussions in Bandelt and Chepoi [2; 10], Eppstein, Falmagne and Ovchinnikov [15], Imrich and Klavžar [27], Isbell [28], Nica [33], Roller [35] and van de Vel [38]). A large part of our use of the higher-dimensional cubes is in the application of disc diagram techniques, introduced by Casson and developed by Sageev in his thesis and by Wise in [40].

The simplicial boundary In view of [Theorem D](#) below, it is natural to attempt to detect unboundedness of $\Gamma\mathbf{X}$ by finding a geodesic ray in \mathbf{X} that contains points far from each hyperplane and does not lie in a cubical “flat sector”. The *simplicial boundary* of a CAT(0) cube complex \mathbf{X} , whose construction is this paper’s main innovation, is designed to keep track of infinite bounded subgraphs of $\Gamma\mathbf{X}$ corresponding to unbounded flat subcomplexes.

The set $\mathcal{W}(\gamma)$ of hyperplanes that cross the geodesic ray γ has several salient properties which are abstracted in the definition of a *unidirectional boundary set (UBS)* of hyperplanes in [Section 3.1](#). The simplices of $\partial_{\Delta}\mathbf{X}$ are defined to be equivalence classes of UBSs, where two UBSs are equivalent if their symmetric difference is finite. The definition of $\partial_{\Delta}\mathbf{X}$ is enabled by [Theorem 3.10](#), which says that each UBS decomposes as the disjoint union of *minimal* UBSs in an essentially unique way. Crucially, the subgraph of $\Delta\mathbf{X}$ generated by a UBS corresponding to an n -simplex of $\partial_{\Delta}\mathbf{X}$ is a spanning subgraph of an infinite complete $(n + 1)$ -partite graph. When $n \geq 1$, this subgraph is sufficiently connected to ensure that the corresponding subgraph of $\Gamma\mathbf{X}$ is bounded.

[Theorem 3.14](#) establishes basic properties of $\partial_{\Delta}\mathbf{X}$.

Theorem A *The simplicial boundary $\partial_{\Delta}\mathbf{X}$ is a flag complex, each simplex is contained in a maximal simplex, and, if \mathbf{X} is unbounded and hyperbolic, then $\partial_{\Delta}\mathbf{X}$ is a nonempty discrete set of 0-simplices.*

The last assertion of [Theorem A](#) follows from the relationship between UBSs and $\Gamma\mathbf{X}$ described above, together with the proof of [[21](#), [Theorem 7.6](#)], implies that $\partial_{\Delta}\mathbf{X}$ is totally disconnected when \mathbf{X} is unbounded and hyperbolic.

The definition of a UBS is motivated by sets of hyperplanes of the form $\mathcal{W}(\gamma)$, but the *eighth-flat* shown in [Figure 1](#) shows that not all UBSs are sets of hyperplanes dual to some ray. However, [Theorem 3.19](#) says that maximal simplices of $\partial_{\Delta}\mathbf{X}$ are *visible*, ie represented by UBSs determined by geodesic rays. The property of being *fully visible* – every simplex of $\partial_{\Delta}\mathbf{X}$ is visible – enables additional conclusions. In particular, by [Theorem 3.23](#), if \mathbf{X} is fully visible, then each n -simplex of $\partial_{\Delta}\mathbf{X}$ records the existence of an isometrically embedded copy of the standard tiling of $[0, \infty)^{n+1}$ by $(n + 1)$ -dimensional Euclidean unit cubes, and thus corresponds to a genuine infinite $(n + 1)$ -partite subgraph of $\Delta\mathbf{X}$.

This is one of several analogues between $\partial_{\Delta}\mathbf{X}$ – which is an invariant of the median graph $\mathbf{X}^{(1)}$ – and the Tits boundary of \mathbf{X} . Another similarity is given by [Theorem 3.30](#), which equates the existence of product decompositions of \mathbf{X} with simplicial join

decompositions of $\partial_{\Delta}\mathbf{X}$. While this is very similar to the corresponding theorem about spherical join decompositions of the Tits boundary of a CAT(0) space (see Bridson and Haefliger [5, Theorem II.9.24]), there are important differences in the hypotheses under which these results hold, reflecting the fact that \mathbf{X} generally contains many more combinatorial geodesics than CAT(0) geodesics.

However, $\partial_{\Delta}\mathbf{X}$ can be endowed with a CAT(1) metric by declaring each simplex to be a right-angled spherical simplex, and, if \mathbf{X} is fully visible and satisfies an additional technical condition, then this CAT(1) realization of $\partial_{\Delta}\mathbf{X}$ is isometric to the Tits boundary $\partial_T\mathbf{X}$, ie $\partial_{\Delta}\mathbf{X}$ is a triangulation of $\partial_T\mathbf{X}$. In Section 3.5, we give conditions ensuring that $\partial_{\Delta}\mathbf{X}$ is an isometric triangulation of the Tits boundary, explain the relationship between the two boundaries in general, and give examples showing that the two boundaries are in general slightly different.

Proposition 3.37 *Let \mathbf{X} be a fully visible CAT(0) cube complex for which $\partial_{\Delta}\mathbf{X}$ is defined. Let $(\partial_{\Delta}\mathbf{X}, \check{d})$ be the CAT(1) realization of $\partial_{\Delta}\mathbf{X}$. Then there is an isometric embedding $I: \partial_{\Delta}\mathbf{X} \rightarrow \partial_T\mathbf{X}$ that is a section of a surjective map $R: \partial_T\mathbf{X} \rightarrow \partial_{\Delta}\mathbf{X}$ such that*

- (1) *for each $z \in \partial_{\Delta}\mathbf{X}$, the preimage $R^{-1}(z)$ has Tits diameter less than $\frac{\pi}{2}$;*
- (2) *I is $\frac{\pi}{2}$ -quasisurjective, and has quasi-inverse R ;*
- (3) *I is an isometry if and only if every 0-simplex of $\partial_{\Delta}\mathbf{X}$ has bounded bending;*
- (4) *if $\gamma: \mathbb{R} \rightarrow \mathbf{X}$ is an axis of an isometry $g \in \text{Aut}(\mathbf{X})$, with endpoints $p_{\pm} \in \partial_T\mathbf{X}$, and $\mathcal{H}(\gamma)$ represents a pair of 0-simplices, then $p_{\pm} \in \text{im}(I)$.*

The CAT(0) cube complex \mathbf{X} has *bounded bending* if any two combinatorial geodesic rays representing the same 0-simplex of $\partial_{\Delta}\mathbf{X}$ fellow-travel. As was pointed out by Mike Carr, the bounded bending condition need not be satisfied even in the presence of a proper, cocompact group action on \mathbf{X} ; the universal cover of the Salvetti complex of a right-angled Artin group studied by Croke and Kleiner in [13] does not have bounded bending.

Boundedness of $\Gamma\mathbf{X}$ The motivating question is addressed in Section 4, where we prove the following theorem. The cube complex \mathbf{X} is *strongly locally finite* if it is locally finite and does not contain an infinite family of pairwise crossing hyperplanes.

Theorem C (Corollary 4.1, Theorem 4.3) *Let \mathbf{X} be a strongly locally finite, essential, one-ended CAT(0) cube complex. Then*

$$\text{diam}(\Gamma\mathbf{X}) \leq \text{diam}(\Delta\mathbf{X}) \leq \text{diam}((\partial_\Delta\mathbf{X})^1) \leq 2(\text{diam}(\Delta\mathbf{X}) - 1),$$

so that $\Delta\mathbf{X}$ has finite diameter if and only if the 1–skeleton of $\partial_\Delta\mathbf{X}$ is bounded, and, if $\partial_\Delta\mathbf{X}$ is bounded, then $\Gamma\mathbf{X}$ is bounded. If, in addition, \mathbf{X} has finite degree, then $\Gamma\mathbf{X}$ is bounded if and only if $\partial_\Delta\mathbf{X}$ has bounded 1–skeleton. Finally, if \mathbf{X} is a strongly locally finite CAT(0) cube complex and $\Gamma\mathbf{X}$ is bounded, then for each isolated 0–simplex v of $\partial_\Delta\mathbf{X}$, v lies in the image of $\partial_\Delta H$ for some hyperplane H .

If $\Delta\mathbf{X}$ is bounded, \mathbf{X} decomposes as an iterated pseudoproduct, a notion generalizing that of a product and providing a nexus between $\Delta\mathbf{X}$ and $\partial_\Delta\mathbf{X}$. If γ is a combinatorial geodesic ray in \mathbf{X} , then the projection $\text{Proj}_\mathbf{X} \gamma$ of γ to the contact graph is the embedded ray in $\Gamma\mathbf{X}$ that traverses the vertices corresponding to hyperplanes crossing γ , in the order that they cross γ . Projecting rays to $\Gamma\mathbf{X}$ allows one to detect unboundedness of $\Gamma\mathbf{X}$.

Theorem D (Theorems 2.2 and 2.4) *Let γ be a combinatorial geodesic ray, let $\text{Proj}_\mathbf{X} \gamma \subseteq \Gamma\mathbf{X}$ be its projection to the contact graph, and let $\Lambda(\gamma)$ be the full subgraph of $\Gamma\mathbf{X}$ generated by $\text{Proj}_\mathbf{X} \gamma$. Suppose there exists $B < \infty$ such that, if $K_{p,p} \subseteq \Lambda(\gamma)$, then $p \leq B$. Suppose, moreover, that for all $R \geq 0$, there exists T such that for all hyperplanes H , any subpath of γ that lies in $N_R(N(H))$ has length at most T . Then the inclusion $\text{Proj}_\mathbf{X} \gamma \hookrightarrow \Gamma\mathbf{X}$ is a quasi-isometric embedding. Moreover, if $\text{diam}_{\Gamma\mathbf{X}}(\text{Proj}_\mathbf{X} \gamma) < \infty$, then either γ lies in an isometrically embedded eighth-flat or γ lies in a uniform neighborhood of some hyperplane.*

Theorem D is used in conjunction with essentiality and local finiteness to show that rank-one combinatorial geodesic rays correspond to 0–simplices of the boundary, and is of interest in its own right. The remainder of the proof of **Theorem C** uses the duality between cube complexes and wallspaces, disc diagram arguments and basic facts about $\partial_\Delta\mathbf{X}$.

Combinatorial divergence Recently, Behrstock and Charney studied divergence in CAT(0) cube complexes associated to right-angled Artin groups. Motivated by their results, we examine, in **Section 6**, combinatorial divergence and divergence of geodesics in the strongly locally finite CAT(0) cube complex \mathbf{X} , without reference to any group action. For example, we prove the following.

Theorem 6.8 *Let \mathbf{X} be strongly locally finite. Then $\partial_\Delta\mathbf{X}^{(1)}$ is bounded if and only if \mathbf{X} has weakly uniformly linear divergence.*

Theorem 6.8 means that there exists a constant $A = A(\mathbf{X})$ such that for all combinatorial geodesic rays γ, γ' with common initial 0-cube, and for all $r \geq 0$, there exists a path P joining $\gamma(r), \gamma'(r)$ that avoids the ball of radius $r - 1$ about $\gamma(0)$ and has length bounded by $Ar + B$, where only B depends on γ, γ' . The situation is similar for divergence of combinatorial geodesics.

Corollary 6.9 *Let $\alpha: \mathbb{R} \rightarrow \mathbf{X}$ be a bi-infinite combinatorial geodesic representing the simplices v^\pm of $\partial_\Delta \mathbf{X}$. Then v^- and v^+ lie in different components of $\partial_\Delta \mathbf{X}$ if and only if the divergence of the geodesic α is superlinear.*

Simplicial boundary considerations can apparently be used to make more refined statements about divergence functions of cube complexes in terms of the simplicial structure of $\partial_\Delta \mathbf{X}$.

Rank-one isometries and divergence The slightly restricted notion of divergence used in stating the above results aligns with the standard definition (ie that introduced by Gersten in [19; 18]) when \mathbf{X} is combinatorially geodesically complete and admits a geometric group action; see the author [22, Lemma 7.3.5]. In Section 6.3, we observe the following.

Theorem G *Let G act properly, cocompactly and essentially on the combinatorially geodesically complete $CAT(0)$ cube complex \mathbf{X} . If $\Gamma \mathbf{X}$ is bounded, then G has linear divergence if and only if $\Gamma \mathbf{X}$ decomposes as the join of two infinite proper subgraphs (ie $\partial_\Delta \mathbf{X}$ decomposes as the simplicial join of two proper subcomplexes). Otherwise, G has at least quadratic divergence.*

Behrstock and Charney characterize right-angled Artin groups with linear divergence: they are exactly those for which the defining graph decomposes as a nontrivial join [3, Corollary 5.4]. It is evident that the defining graph Θ of the right-angled Artin group G_Θ decomposes as a join if and only if the contact graph (equivalently, simplicial boundary) of the universal cover of the Salvetti complex of G_Θ decomposes as a join. In this sense, either the contact graph or the simplicial boundary is the analogue of the defining graph of a right-angled Artin group needed to extend their characterization to cubulated groups that are not right-angled Artin groups.

A key ingredient in the proof of **Theorem G** is the *rank-rigidity theorem* of Caprace and Sageev [6, Theorem A, Theorem B], and in Section 5, we translate this result into the language of the contact graph and the simplicial boundary. Roughly speaking, the dichotomy established by rank-rigidity – if \mathbf{X} is not a product, then a group acting sufficiently nicely has a rank-one element – corresponds to the fact that $\Gamma \mathbf{X}$ is either bounded or unbounded. We also note, in **Theorem 5.2**, that whether an isometry is rank-one is detected by the contact graph.

Acknowledgements I am grateful to Jason Behrstock, Ruth Charney, Victor Chepoi, Michah Sageev and Daniel Wise for helpful discussions about cube complexes, boundaries and divergence. I enthusiastically thank the anonymous referee for making valuable comments and encouraging me to investigate in greater detail the material discussed in [Section 3.5](#). Finally, I thank Mike Carr for astute and helpful corrections and comments on that section.

1 Preliminaries

We refer the reader to Behrstock and Charney [3], Caprace and Sageev [6], Chatterji and Niblo [9], Chepoi [10], the author [22], Haglund [23] and Roller [35], Sageev [36] and Wise [40] for more comprehensive discussions of the background given in this section.

1.1 CAT(0) cube complexes and contact graphs

Throughout this paper, \mathbf{X} is a CAT(0) cube complex, ie a simply connected CW-complex whose cells are Euclidean cubes of the form $[-\frac{1}{2}, \frac{1}{2}]^d$ with $0 \leq d < \infty$, attached so that any two distinct cubes either have empty intersection or intersect in a common face. Also, the link of each 0-cube is a flag complex. The *dimension* of \mathbf{X} is the supremum of the set of dimensions of cubes of \mathbf{X} . The *degree* of $x \in \mathbf{X}^0$ is the number of vertices in its link, and the *degree* of \mathbf{X} is the supremum of the set of degrees of its 0-cubes. Note that the degree of \mathbf{X} is at least the dimension.

1.1.1 Hyperplanes For $d \geq 1$, a *midcube* of the cube $[-\frac{1}{2}, \frac{1}{2}]^d$ is a subspace obtained by restricting exactly one coordinate to 0. A *hyperplane* H of \mathbf{X} is a connected subspace such that, for each cube c , 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 such that $H \cap c \neq \emptyset$. The ambient CAT(0) cube complex is assumed to be countable and to contain at least two distinct 0-cubes, and hence at least one hyperplane.

In [36, Theorems 4.10, 4.11], Sageev proved the following for each hyperplane H : H is a CAT(0) cube complex with $\dim(H) < \dim(\mathbf{X})$; H is *2-sided* in the sense that $N(H) \cong H \times [-\frac{1}{2}, \frac{1}{2}]$; the complement $\mathbf{X} - H$ has exactly two nonempty components, H^+ and H^- . These are the *halfspaces* associated to H . These facts were established independently, from a slightly different viewpoint, by Chepoi in [10].

Subspaces $A, B \subset X$ lying in distinct halfspaces associated to H are *separated* by H . In particular, a 1-cube c whose endpoints are separated by H is *dual to* H and H is *dual to* c . The relation “dual to the same hyperplane” on the set of 1-cubes coincides

with the *Djoković–Winkler relation*, ie the transitive closure of the relation which contains (c, c') if the 1–cubes c and c' lie on opposite sides of a 4–cycle in \mathbf{X}^1 .

The distinct hyperplanes H and H' *contact* if no third hyperplane H'' separates H from H' . The hyperplanes H and H' contact if and only if $N(H) \cap N(H') \neq \emptyset$, and this can happen in one of two ways. First, H and H' can *cross*, ie $N(H) \cap N(H')$ contains a 2–cube s whose boundary 4–cycle $c_1 c_2 c'_1 c'_2$ has the property that c_1 and c'_1 are dual to H and c_2 and c'_2 to H' . Equivalently, H and H' cross if the four *quarter-spaces* $H^+ \cap (H')^+$, $H^+ \cap (H')^-$, $H^- \cap (H')^+$, $H^- \cap (H')^-$ are all nonempty. If H, H' contact and do not cross, then they *osculate*: there are 1–cubes c, c' , dual to H, H' respectively, such that c and c' have a common 0–cube and the path cc' does not lie on the boundary path of any 2–cube. We use the notation $H \perp H'$ to mean that H and H' cross and $H \perp \smile H'$ to mean that they contact.

If H_1, \dots, H_n are distinct pairwise crossing hyperplanes, then there is an n –cube $c \subset \bigcap_{i=1}^n N(H_i)$ containing a 1–cube dual to each H_i . We say that c is *dual* to the collection $\{H_i\}_{i=1}^n$. If $\{H_i\}$ is a maximal family of pairwise crossing hyperplanes, then c is unique. More generally, if H_1, \dots, H_n pairwise-contact, then there is at least one 0–cube c with an incident 1–cube dual to each H_i . Thus $\dim \mathbf{X}$ is the cardinality of a largest family of pairwise crossing hyperplanes, and the degree of \mathbf{X} is the cardinality of a largest family of pairwise-contacting hyperplanes.

1.1.2 Strong local finiteness The locally finite CAT(0) cube complex \mathbf{X} is *strongly locally finite* if every family of pairwise crossing hyperplanes is finite. Note that \mathbf{X} can be strongly locally finite with infinite dimension and degree. There are CAT(0) cube complexes that are locally finite but not strongly locally finite: these contain infinite families of pairwise crossing hyperplanes corresponding to infinite cubes “at infinity”. The cube complex dual to the wallspace shown by Hruska and Wise in [26, Figure 5] is of this type. We often assume that \mathbf{X} contains no infinite family of pairwise crossing hyperplanes without hypothesizing strong local finiteness.

1.1.3 Metrics, geodesics and hyperplane-equivalence Gromov showed that, if $\dim \mathbf{X} < \infty$, then by treating each cube as a Euclidean unit cube, one obtains a piecewise-Euclidean geodesic metric $\partial_{\mathbf{X}}: \mathbf{X} \times \mathbf{X} \rightarrow [0, \infty)^2$ such that $(\mathbf{X}, \partial_{\mathbf{X}})$ is a CAT(0) space [20]; more general results of this type were proved by Bridson [4]. Recently, Leary showed that this CAT(0) metric also exists when \mathbf{X} is infinite-dimensional; see Leary [30, Theorem 41].

We almost always use a different metric. A *combinatorial path* in \mathbf{X} is a continuous map $P: I \rightarrow \mathbf{X}^1$ of an interval I such that $P(i) \in \mathbf{X}^0$ for each $i \in I \cap \mathbb{Z}$ and P maps each $[n, n + 1]$ homeomorphically to a 1–cube. The *path metric* $d_{\mathbf{X}}: \mathbf{X}^1 \times \mathbf{X}^1 \rightarrow [0, \infty)$

is defined by letting $d_{\mathbf{X}}(x, y)$ be the length of a shortest combinatorial path joining x to y . If $x, y \in \mathbf{X}^0$ are 0-cubes, then $d_{\mathbf{X}}(x, y)$ counts the hyperplanes separating x from y .

Haglund showed that $d_{\mathbf{X}}$ extends to all of \mathbf{X} in such a way that $d_{\mathbf{X}}$ restricts to the ℓ^1 metric on each cube and to the path metric on \mathbf{X}^1 . For $x, y \in \mathbf{X}$, $d_{\mathbf{X}}(x, y)$ is the infimum of the lengths of paths joining x to y that are parallel to the 1-skeleton [23]. We usually work in the 1-skeleton and consider only combinatorial paths, and unless stated otherwise, we use this ‘‘cubical ℓ_1 ’’ metric $d_{\mathbf{X}}$, restricted to the 1-skeleton. It follows from [6, Lemma 2.2] that, if $\dim \mathbf{X} < \infty$, then $(\mathbf{X}, \partial_{\mathbf{X}})$ is quasi-isometric to $(\mathbf{X}, d_{\mathbf{X}})$. This fact is sometimes useful for applying facts about the CAT(0) metric in combinatorial contexts and vice versa.

We denote by $N_R(A)$ the smallest subcomplex containing the R -neighborhood of A^0 in \mathbf{X}^0 , where A is a subcomplex. The subcomplex $Y \subset \mathbf{X}$ is isometrically embedded if and only if $Y^1 \hookrightarrow \mathbf{X}^1$ is an isometric embedding. Given any connected subcomplex Y , the hyperplane H crosses Y if and only if $Y \cap H^+$ and $Y \cap H^-$ are nonempty, and Y is isometrically embedded if and only if $Y \cap H$ is connected for each hyperplane H that crosses Y . Thus $P: I \rightarrow \mathbf{X}^1$ is a geodesic if and only if $P(I)$ contains at most one 1-cube dual to each hyperplane. In this case, we also denote the isometric subcomplex $P(I)$ by P . If I is finite, then P is a (combinatorial) geodesic segment. If $I = \mathbb{R}$, then P is a (bi-infinite) geodesic. Otherwise, P is a geodesic ray.

The connected full subcomplex Y is convex if for each pair of hyperplanes H, H' that cross Y , if $H \perp \smile H'$ then $N(H) \cap N(H') \cap Y \neq \emptyset$ and, if in addition $H \perp H'$, then $H \cap Y$ and $H' \cap Y$ cross. (Y is full if it contains every cube c of \mathbf{X} whose 1-skeleton appears in Y . Convex subcomplex in this paper is understood to denote a full subcomplex with the preceding property. Convex subcomplexes are therefore CAT(0).) This can be proved using the fact that H and $N(H)$ are convex in \mathbf{X} , which was proved by Sageev [36], and which is true for either metric since the connected subcomplex Y is CAT(0)-convex if and only if it is combinatorially convex [23]. From this characterization of convexity, or from the median graph viewpoint, it follows that \mathbf{X} has the Helly property: if Y_1, \dots, Y_n are pairwise-intersecting convex subcomplexes of \mathbf{X} , then $\bigcap_{i=1}^n Y_n \neq \emptyset$.

Henceforth, \mathcal{W} will denote the set of all hyperplanes in \mathbf{X} , and \mathcal{W}^{\pm} the set of all halfspaces. For any isometrically embedded subcomplex Y , the set of hyperplanes that cross Y is $\mathcal{W}(Y)$. Occasionally, we shall use the notation $\mathcal{W}(\alpha)$, where α is a CAT(0) geodesic path in \mathbf{X} , to denote the set of hyperplanes crossing α .

Definition 1.1 (Hyperplane-equivalent, consuming, almost equivalent) The isometric subcomplex Y_1 consumes the isometric subcomplex Y_2 if $\mathcal{W}(Y_2) \subseteq \mathcal{W}(Y_1)$. If Y_1

consumes Y_2 and Y_2 consumes Y_1 , then Y_1 and Y_2 are *hyperplane-equivalent*. If $|\mathcal{W}(Y_1) \Delta \mathcal{W}(Y_2)| < \infty$, then Y_1 and Y_2 are *almost equivalent*.

Almost-equivalence is an equivalence relation, and we shall mostly be concerned with almost equivalence of combinatorial geodesic rays.

1.1.4 Medians Chepoi proved that the class of graphs that are 1–skeleta of CAT(0) cube complexes is precisely the class of *median graphs* [10]. By viewing hyperplanes as Djoković–Winkler classes of 1–cubes, he also established the properties of hyperplanes described above. A synopsis of the correspondence between these viewpoints appears in [12, Section 3] by Chepoi and the author.

Definition 1.2 (Median graph) The graph M is a *median graph* if for all triples $x, y, z \in M$ of distinct vertices, there exists a unique vertex $m = m(x, y, z)$, the *median* of x, y, z , such that $d(x, y) = d(x, m) + d(m, y)$, $d(x, z) = d(x, m) + d(m, z)$ and $d(y, z) = d(y, m) + d(m, y)$.

A detailed discussion of the relationship between median algebras, median graphs and cube complexes appears in [35], and the construction of a median graph from a *wallspace* is lucidly discussed in [33]. It is crucial to observe that the invariants of \mathbf{X} discussed in this paper – the crossing graph, the contact graph, and the simplicial boundary – are invariants of \mathcal{W} and hence of the median graph \mathbf{X}^1 only.

1.1.5 Essentiality and compact indecomposability Following [36], we define the hyperplane H to be *essential* if each of H^+ and H^- contains points arbitrarily far from H . If every hyperplane of \mathbf{X} is essential, then \mathbf{X} is *essential*. We say \mathbf{X} is *compactly decomposable* if there exists a compact, convex subcomplex K such that $\mathbf{X} - K$ is disconnected. Otherwise, \mathbf{X} is *compactly indecomposable*. If \mathbf{X} is essential and locally finite, then \mathbf{X} is compactly indecomposable if and only if \mathbf{X}^1 is one-ended.

1.1.6 Cubulating wallspace Sageev showed that a G –cube complex can be constructed from an appropriately chosen family of *codimension*–1 subgroups of a group G ; see [36, Theorem 3.1]. In the more general context of a *wallspace* provided by Haglund and Paulin [24], one can construct a CAT(0) cube complex in essentially the same way. We summarize the approach taken in [9].

A *wallspace* consists of an *underlying set* \mathcal{S} equipped with a set \mathbb{W} of *walls*. A *wall* $W \in \mathbb{W}$ is a pair $(\mathfrak{h}(W), \mathfrak{h}^*(W))$ of disjoint nonempty subsets of \mathcal{S} whose union is all of \mathcal{S} . The wall W *separates* $s, s' \in \mathcal{S}$ if $s \in \mathfrak{h}(W)$ and $s' \in \mathfrak{h}^*(W)$. The walls W, W' *cross* if each of the four possible intersections of their halfspaces is nonempty. We

require that any two elements of \mathcal{S} be separated by only finitely many walls. For $s, s' \in \mathcal{S}$, denote by $\#(s, s')$ the number of walls separating s and s' .

Let \mathbb{H} be the set of halfspaces, and let $\pi: \mathbb{H} \rightarrow \mathbb{W}$ be the map that sends $\mathfrak{h}(W)$ and $\mathfrak{h}^*(W)$ to W for each $W \in \mathbb{W}$. The *dual cube complex* $\mathbf{X}(\mathcal{S}, \mathbb{W}) = \mathbf{X}$ is formed as follows. A *consistent orientation* is a section $x: \mathbb{W} \rightarrow \mathbb{H}$ of π , ie a choice of exactly one halfspace for each wall, such that $x(W) \cap x'(W) \neq \emptyset$ for all walls W, W' . If W, W' cross, any section x of π satisfies $x(W) \cap x(W') \neq \emptyset$. The consistent orientation x is *canonical* if there exists $s \in \mathcal{S}$ such that $s \in x(W)$ for all $W \in \mathbb{W}$. Each $s \in \mathcal{S}$ yields a canonical consistent orientation x_s , by letting $x_s(W)$ be whichever of $\mathfrak{h}(W), \mathfrak{h}^*(W)$ contains s . Clearly, x_s and $x_{s'}$ differ exactly on the finite set of walls W that separate s from s' .

The 0–skeleton \mathbf{X}^0 consists of all consistent orientations x such that x differs from some (any) x_s on finitely many walls. The 0–cubes x, y are joined by a 1–cube if and only if $x(W) \neq y(W)$ for a single wall W . The resulting graph \mathbf{X}^1 is median and thus the 1–skeleton of a CAT(0) cube complex. There is a bijection $\beta: \mathbb{W} \rightarrow \mathcal{W}$, and the hyperplanes $\beta(W), \beta(W')$ cross if and only if the walls W, W' cross. Similarly, $\beta(W)$ and $\beta(W')$ are separated by a hyperplane $\beta(W'')$ if and only if W'' separates W from W' , and $x_s, x_{s'}$ are separated by $\beta(W)$ if and only if s, s' are separated by W .

Consider the wallspace in which \mathcal{S} is the 0–skeleton of a CAT(0) cube complex and \mathbb{W} is the set of bipartitions of the 0–skeleton induced by the hyperplanes. We view a 0–cube $x \in \mathbf{X}^0$ as consistent, canonical orientation of all walls. A section x of π that is consistent but not canonical is a 0–cube at infinity.

A corollary of this construction is the existence of *restriction quotients* (see eg [6]). Let $\mathcal{V} \subseteq \mathcal{W}$ be a set of hyperplanes of \mathbf{X} . Then there exists a cubical quotient map $\rho: \mathbf{X} \rightarrow \mathbf{X}(\mathcal{V})$, where $\mathbf{X}(\mathcal{V})$ is the cube complex dual to the wallspace whose underlying set is \mathbf{X}^0 and whose walls are induced by the set \mathcal{V} of hyperplanes. $\mathbf{X}(\mathcal{V})$ is the *restriction quotient* associated to \mathcal{V} .

If $H, H' \in \mathcal{V}$ contact in \mathbf{X} , then $\rho(H) \perp \sim \rho(H')$ in $\mathbf{X}(\mathcal{V})$. The converse holds for crossings but not for osculations. More precisely, if $\rho(H) \perp \rho(H')$, then $H \perp H'$. However, if $H, H' \in \mathcal{V}$, and every hyperplane U separating H from H' in \mathbf{X} belongs to $\mathcal{W} - \mathcal{V}$, then $\rho(H) \perp \sim \rho(H')$.

1.1.7 Contact graphs and crossing graphs The *crossing graph* $\Delta \mathbf{X}$ of \mathbf{X} is a graph whose vertex set is the set \mathcal{W} of hyperplanes of \mathbf{X} , and we use the same notation for a vertex as for the corresponding hyperplane. The vertices H and H' are adjacent if and only if $H \perp H'$. Equivalently, $\Delta \mathbf{X}$ is the intersection graph of the family of hyperplanes. The crossing graph $\Delta \mathbf{X}$ appears as the *transversality* graph in work of

Niblo and of Roller, and an equivalent construction appears in work of Bandelt, Dress, van de Vel and others.

The *contact graph* $\Gamma\mathbf{X}$ of \mathbf{X} , introduced in [21], has vertex set \mathcal{W} , but now H, H' are adjacent if and only if $H \perp \smile H'$. In particular, $\Delta\mathbf{X}$ is a spanning subgraph of $\Gamma\mathbf{X}$. Equivalently, $\Gamma\mathbf{X}$ is the intersection graph of the family of hyperplane-carriers in \mathbf{X} . We denote by $H \perp \smile H'$ the closed edge of Γ corresponding to the contacting of H and H' .

Observe that for each n -clique in $\Delta\mathbf{X}$, there is an n -cube in \mathbf{X} and for each n -clique in $\Gamma\mathbf{X}$, there is a 0 -cube in \mathbf{X} of degree at least n . Hence the dimension of \mathbf{X} is the clique number of $\Delta\mathbf{X}$, and the degree of \mathbf{X} is the clique number of $\Gamma\mathbf{X}$.

Let $\gamma: I \rightarrow \mathbf{X}$ be a combinatorial geodesic path, and for each i , let H_i be the hyperplane dual to the 1 -cube $\gamma([i, i + 1])$. Then $H_i \perp \smile H_{i+1}$ for each i , and we obtain a path $\text{Proj}_{\mathbf{X}}(\gamma)$ in $\Gamma\mathbf{X}$, called the *projection of γ to $\Gamma\mathbf{X}$* , that is the concatenation of these edges. Since γ passes through each hyperplane at most once, $\text{Proj}_{\mathbf{X}}(\gamma)$ is an embedded path in $\Gamma\mathbf{X}$, and spans the subgraph $\Lambda(\gamma)$ generated by $\mathcal{W}(\gamma)$. This projection shows that, unlike the crossing graph, $\Gamma\mathbf{X}$ is necessarily connected.

Detailed discussions of the geometric and combinatorial properties of $\Gamma\mathbf{X}$ can be found in [21; 22; 12]. Some variants and uses of contact graphs are discussed by Chepoi in [11]. The main geometric property of contact graphs is [21, Theorem 4.1], which says that there exist constants $M \geq 1, C \geq 0$, independent of \mathbf{X} , and a tree $\mathcal{T} = \mathcal{T}(\mathbf{X})$, such that $\Gamma\mathbf{X}$ is (M, C) -quasi-isometric to \mathcal{T} .

We need the following two propositions, which can easily be proved by cubulating wallspaces. The first is essentially given in [6, Section 2], albeit not in contact graph terms, and characterizes cubical products. For graphs A_1, A_2 , we denote their join by $A_1 \star A_2$.

Proposition 1.3 *The following are equivalent for a CAT(0) cube complex \mathbf{X} .*

- (1) *There exist nonempty convex subcomplexes $\mathbf{X}_1, \mathbf{X}_2 \subset \mathbf{X}$ such that $\mathbf{X} \cong \mathbf{X}_1 \times \mathbf{X}_2$.*
- (2) *There exist nonempty disjoint subgraphs $A_1, A_2 \subset \Delta\mathbf{X}$ such that $\Delta\mathbf{X} \cong A_1 \star A_2$.*
- (3) *There exist nonempty disjoint subgraphs $A'_1, A'_2 \subset \Gamma\mathbf{X}$ such that $\Gamma\mathbf{X} \cong A'_1 \star A'_2$, and for all $H_1 \in A'_1, H_2 \in A'_2$, the edge $H_1 \perp \smile H_2$ is a crossing-edge.*

The second proposition requires the following notion, which also features in Section 3.

Definition 1.4 (Inseparable set) The set $\mathcal{U} \subseteq \mathcal{W}$ is *inseparable* if for all $U, U' \in \mathcal{U}$, every $H \in \mathcal{W}$ that separates U from U' belongs to \mathcal{U} .

Given a set \mathcal{U} of hyperplanes, let $\Lambda(\mathcal{U})$ denote the full subgraph of $\Gamma\mathbf{X}$ induced by \mathcal{U} . Given an isometrically embedded subcomplex $Y \subseteq \mathbf{X}$, let $\Lambda(Y) = \Lambda(\mathcal{W}(Y))$. Note that $\Lambda(Y)$ is the contact graph of the convex hull of Y , and $\Lambda(Y) \cap \Delta\mathbf{X}$ is its crossing graph. Note also that the subcomplex Y is compact if and only if $\mathcal{W}(Y)$ is finite, and that the convex hull of a compact subcomplex is again compact. The following is proved in [22] by constructing the cube complex dual to the wallspace $(\mathbf{X}^0, \mathcal{U})$ and showing that this complex embeds convexly in \mathbf{X} .

Proposition 1.5 *Let \mathcal{U} be a finite inseparable set of hyperplanes in the locally finite CAT(0) cube complex \mathbf{X} . Then there exists a compact, convex subcomplex $Y \subset \mathbf{X}$ such that $\mathcal{W}(Y) = \mathcal{U}$.*

There are other versions of Proposition 1.5 in which, under additional assumptions, \mathcal{U} is allowed to be infinite, but it is in general much more problematic to obtain a subcomplex of \mathbf{X} from a subgraph of $\Gamma\mathbf{X}$ than it is to obtain a quotient.

1.2 Actions on CAT(0) cube complexes

If \mathbf{X} is locally finite, then \mathbf{X}^1 is a proper metric space. We say that the group G acts on \mathbf{X} to mean that G acts by cubical automorphisms and thus by isometries (with respect to both $d_{\mathbf{X}}$ and $\partial_{\mathbf{X}}$). We shall always assume that G acts nontrivially, ie without a global fixed point. G acts properly if, for any $x \in \mathbf{X}^0$ and any infinite sequence $(g_n \in G)_{n \geq 0}$, we have that $d_{\mathbf{X}}(x, g_n x) \rightarrow \infty$ as $n \rightarrow \infty$. Equivalently, G acts properly if the stabilizer G_c of any cube c is finite, provided \mathbf{X} is locally finite.

Let G act on \mathbf{X} . For each $W \in \mathcal{W}$, let G_W denote the stabilizer of W and note that $G_W = \text{Stab}_{\mathbf{X}}(N(W))$. For all $g \in G$, and for all $W \in \mathcal{W}$, it is evident that gW is again a hyperplane, and $gW \perp \smile gW'$ (respectively, $gW \perp gW'$) if and only if $W \perp \smile W'$ (respectively, $W \perp W'$). Hence G acts by isometries on $\Gamma\mathbf{X}$, and the stabilizer of a vertex coincides with the stabilizer of the corresponding hyperplane.

1.2.1 Essential actions If G acts on \mathbf{X} , then the hyperplane H is G -essential if for some $x \in \mathbf{X}^0$, each halfspace H^{\pm} contains points of the orbit Gx arbitrarily far from H . Note that, if H is G -essential, then H is essential. If each hyperplane is G -essential, then G acts essentially on \mathbf{X} . If G acts essentially on \mathbf{X} , then \mathbf{X} is essential. The G -essential core $\mathbf{Y} \subseteq \mathbf{X}$ is a (possibly empty) convex G -invariant subcomplex on which G acts essentially. The essential core theorem [6, Proposition 3.5] shows that if G acts properly on the locally finite cube complex \mathbf{X} , either with finitely many orbits of hyperplanes or without a fixed point on the visual boundary, then the G -essential core is unbounded.

1.2.2 Semisimplicity and rank-one isometries Let G act on \mathbf{X} . The element $g \in G$ is *elliptic* if g stabilizes a cube c and *combinatorially hyperbolic* if there exists a combinatorial geodesic $\gamma: \mathbb{R} \rightarrow \mathbf{X}^1$ and some $\tau > 0$ such that $g\gamma(t) = \gamma(t + \tau)$ for all $t \in \mathbb{R}$. The path γ is a *combinatorial axis* for g , and τ is the *translation length* of g . The following result of Haglund [23, Theorem 1.4], is fundamental.

Theorem 1.6 (Semisimplicity) *Let G act on the CAT(0) cube complex \mathbf{X} . Then each $g \in G$ is either combinatorially hyperbolic, stabilizes a cube, or stabilizes a hyperplane H and exchanges the halfspaces associated to H .*

By passing to the first cubical subdivision of \mathbf{X} , Haglund shows that one can guarantee that $g \in G$ is either elliptic or combinatorially hyperbolic. If \mathbf{X} is locally finite and finite-dimensional, and $g \in \text{Aut}(\mathbf{X})$ has a combinatorial geodesic axis $\alpha: \mathbb{R} \rightarrow \mathbf{X}$, then g has a CAT(0) geodesic axis $\beta: \mathbb{R} \rightarrow \mathbf{X}$ that fellow-travels with α (with respect to either metric). The converse is not quite true, but it is easy to deduce the following from Haglund’s classification.

Proposition 1.7 *Let $g \in \text{Aut}(\mathbf{X})$. If g is CAT(0) hyperbolic and \mathbf{X} contains no infinite family of pairwise crossing hyperplanes, then g^k is combinatorially hyperbolic for some $k \in \mathbb{Z}$.*

The hyperbolic isometry g of \mathbf{X} is *rank-one* if some (and hence any) CAT(0) geodesic axis for g does not bound an isometrically embedded Euclidean half-plane in \mathbf{X} .

1.3 Divergence and divergence of geodesics

The notion of the divergence function of a geodesic space was introduced by Gersten in [19; 18]. The present exposition roughly follows that in [3]. Let (M, d) be a geodesic space, and let $\rho(k) = Ak - B$ with $A \in (0, 1), B \geq 0$. Given points $a, b, c \in M$, let $k = d(c, \{a, b\})$ and let $\overline{\text{div}}(a, b, c; \rho) = \inf_P |P|$, where P varies over all paths in $M - B_{\rho(k)}(c)$ that join a to b , where $B_{\rho(k)}(c)$ is the open ball of radius $\rho(k)$ centered at c . By convention, $\overline{\text{div}}(a, b, c; \rho)$ takes the value $+\infty$ if no such path exists. The *divergence of M* (with respect to ρ) is the function

$$\widetilde{\text{div}}(M, \rho)(r) = \sup\{\overline{\text{div}}(a, b, c; \rho) \mid a, b, c \in M, d(a, b) \leq r\}.$$

Given a function $f: [0, \infty) \rightarrow [0, \infty]$, we say that M has *divergence at most f* if there exists ρ such that $\widetilde{\text{div}}(M, \rho)(r) \leq f(r)$ for all $r \geq 0$. If $\gamma: \mathbb{R} \rightarrow \mathbf{X}$ is a bi-infinite geodesic, then the *divergence of γ* is $\overline{\text{div}}(\gamma(r), \gamma(-r), \gamma(0); \rho)$ for each $r \geq 0$. If M has divergence at most f , then each geodesic has divergence at most f , and if some geodesic has super- f divergence, then M has super- f divergence.

2 Projections of rays to the contact graph

We study subgraphs of ΓX induced by projections of geodesic rays. See [40; 21] for a discussion of disc diagrams as they are used in the proof of Theorem 2.2.

2.1 Combinatorial fractional flats and rank-one rays

A (*combinatorial*) *flat* is a cube complex isomorphic to the standard Euclidean tiling of \mathbb{R}^2 by 2-cubes. A *nondiagonal half-flat* is a cube complex isomorphic to the standard tiling of $\mathbb{R} \times [0, \infty)$ by 2-cubes. A *nondiagonal quarter-flat* is a cube complex isomorphic to the standard tiling of $[0, \infty)^2$ by 2-cubes.

Let $f: [0, \infty) \rightarrow [0, \infty)$ be an unbounded, nondecreasing function. The f -*sector* is the subspace $\{(x, y) \in [0, \infty)^2 \mid 0 \leq y \leq f(x)\}$. An *eighth-flat* is a CAT(0) cube complex isomorphic to the smallest subcomplex of the standard tiling of $[0, \infty)^2$ containing a given f -sector. The set of hyperplanes of an eighth-flat \mathbf{E} can be partitioned into two inseparable sets, \mathcal{H} and \mathcal{V} , such that each $H \in \mathcal{H}$ crosses all but finitely many $V \in \mathcal{V}$. The unique combinatorial geodesic ray γ in \mathbf{E} that crosses every hyperplane is the *top bounding ray*, and the unique ray γ' corresponding to the x -axis in $[0, \infty)^2$ is the *bottom bounding ray*; see Figure 1. A *diagonal quarter-flat* is a CAT(0) cube complex obtained from a pair $\mathbf{E}_1, \mathbf{E}_2$ of eighth-flats by identifying the bottom bounding ray of \mathbf{E}_1 with an infinite subray of the bottom bounding ray of \mathbf{E}_2 , using a cubical isometry. Let σ be a bi-infinite combinatorial geodesic in the standard tiling of \mathbb{R}^2 that crosses each hyperplane. The closure of a component of $\mathbb{R}^2 - \sigma$ is a CAT(0) cube complex called a *diagonal half-flat*. A cube complex \mathbf{F} of one of the above types is a *fractional flat*.

Remark 2.1 A nondiagonal fractional flat \mathbf{F} is the product of two 1-dimensional complexes and thus $\text{diam}(\Gamma \mathbf{F}) = \text{diam}(\Delta \mathbf{F}) = 2$. For \mathbf{F} a diagonal half-flat, this value is three. If \mathbf{F} is an eighth-flat or diagonal quarter-flat, then $\Gamma \mathbf{F}$ has finite diameter, and \mathbf{F} has a cofinite subcomplex whose crossing graph has diameter three.

The combinatorial geodesic ray $\gamma: [0, \infty) \rightarrow \mathbf{X}$ *bounds a fractional flat* if there exists a cubical isometric embedding $\mathbf{E} \rightarrow \mathbf{X}$ of an eighth-flat whose image contains γ . This includes any case in which γ lies in an isometric fractional flat. Otherwise, γ is (combinatorially) *rank-one*.

2.2 Rays of bounded projection

Theorem 2.2 gives the only possible obstructions to a combinatorial geodesic ray γ projecting to a quasigeodesic ray in the contact graph. By instead asking only that γ




Figure 1: The arrowed rays at left bound an eighth-flat contained in the larger eighth-flat. The arrowed ray at right is the top bounding ray of an eighth-flat in the standard tiling of \mathbb{R}^3 .

project to an unbounded ray, we reach the conclusion that either γ bounds an eighth-flat or γ lies uniformly close to a single hyperplane, which is [Theorem 2.4](#).

Theorem 2.2 (Projection trichotomy I) *Let \mathbf{X} be a CAT(0) cube complex, and let $\gamma: \mathbb{R}_+ \rightarrow \mathbf{X}$ be a combinatorial geodesic ray. Let $\text{Proj}_{\mathbf{X}}(\gamma)$ be the projection of γ to the contact graph, and let $\Lambda(\gamma)$ be the full subgraph of $\Gamma\mathbf{X}$ generated by $\text{Proj}_{\mathbf{X}}(\gamma)$. Then one of the following holds.*

- (1) *The inclusions $\text{Proj}_{\mathbf{X}}(\gamma), \Lambda(\gamma) \hookrightarrow \Gamma\mathbf{X}$ are quasi-isometric embeddings.*
- (2) *For all $p \geq 0$, we have $K_{p,p} \subset \Lambda(\gamma)$.*
- (3) *There exists $R \geq 0$ such that, for all $t \geq 0$, there exists a hyperplane H such that $\gamma \cap N_R(H)$ contains a connected subpath of length at least t .*

Proof For $i \geq 0$, let H_i be the hyperplane dual to $\gamma([i, i + 1])$. Note that, for all i, j , we have $d_{\text{Proj}_{\mathbf{X}}(\gamma)}(H_i, H_j) = |i - j|$. Moreover, for all i , the hyperplanes H_i and H_{i+1} contact, so that $d_{\Gamma\mathbf{X}}(H_i, H_j) \leq |i - j|$. Suppose that for all $0 < \lambda^{-1} \leq 1$, there exist $i, j \geq 0$ such that

$$N = d_{\Gamma\mathbf{X}}(H_i, H_j) < \lambda^{-1}|i - j| - \lambda^{-1}.$$

Let $H_i = U_0 \perp \cup U_1 \perp \cup \dots \perp U_N = H_j$ be a shortest path in $\Gamma\mathbf{X}$ joining H_i to H_j . We shall verify that, for some $R \leq 2B$, where B is the constant from [Lemma 2.3](#) below, there exists a hyperplane H and a subpath $\gamma'' \subset \gamma$ such that $\gamma'' \subset N_R(N(H))$ and $|\gamma''| \geq \lambda^{-1}$. Since B is independent of λ , and λ can be chosen arbitrarily large, this implies that (3) holds. To this end, let γ' be the subpath of γ between, but not including, the 1-cubes dual to H_i and H_j . Note that $|\gamma'| = |i - j| - 1$.

The disc diagram D For $0 \leq m \leq N$, let $A_m \rightarrow N(U_m)$ be a combinatorial geodesic path, and let these be chosen in such a way that A_m ends on the initial 0-cube of A_{m+1} for each $0 \leq m \leq N - 1$, and the concatenation $A = A_0 A_1 \dots A_N$ has the same

endpoints as γ' . Let $D \rightarrow \mathbf{X}$ be a disc diagram bounded by $A(\gamma')^{-1}$. We choose D subject to the following minimality constraints. First, D has minimal area among all disc diagrams with boundary path $A(\gamma')^{-1}$. Second, the paths A_m are chosen, subject to the constraint that each $A_m \rightarrow N(U_m)$, in such a way that the area of D is as small as possible and, among such minimal-area D , the length of A is as short as possible. Finally, the path $U_0 \perp \smile \cdots \perp \smile U_N$ is chosen among all $\Gamma\mathbf{X}$ -geodesics joining H_i to H_j so that the resulting D has area as small as possible.

Let K be a dual curve in D . As illustrated in Figure 2, K must travel from A to γ' . Indeed, K cannot have two ends on γ' , since γ' is a geodesic segment. Likewise, if K travels from A_m to A_m , for some m , then A_m is not a geodesic, a contradiction. If K travels from A_m to A_{m+1} , then we can employ hexagon moves (see eg [40]) and remove spurs, to reduce the area of D , and then the length of A , without affecting our path in $\Gamma\mathbf{X}$. See the discussion of diagrams with *fixed carriers* in [21, Section 2]. Hence, if K has two ends on A , then K travels from A_m to A_k , with $|m - k| \geq 2$. Let U be the hyperplane to which K maps. If $|m - k| > 2$, then

$$H_i = U_0 \perp \smile U_1 \perp \smile \cdots \perp \smile U_m \perp U \perp U_k \perp \smile \cdots \perp \smile U_N = H_j$$

is a path in $\Gamma\mathbf{X}$ of length less than N , a contradiction. If $m = k + 2$, then replacing U_{k+1} by U results in path from H_i to H_j whose minimal-area diagram is a proper subdiagram of D . Indeed, as in Figure 2, we can replace the part of A between and including the 1-cubes dual to K by a path in the carrier of K in D .




Figure 2: A heuristic picture of D , showing the hyperplane-carriers to which its boundary path maps. At left are impossible dual curves: the single-arrrowed dual curves are illegal, and the crossing pair is illegal. At right are the dual curves L_1 and L_M and their carriers.

Next, note that if K, K' are dual curves, both of which emanate from the path A_m , then K and K' do not cross, for otherwise they intersect in a 2-cube s of D mapping to a 2-cube \bar{s} of $N(U_m)$, and we can employ hexagon moves to pass to a lower-area diagram.

Thus $|A| = |i - j| - 1$, and hence there exists m with

$$M = |A_m| \geq \frac{|i - j| - 1}{N} > \lambda^{-1}.$$

Let E be the subdiagram of D bounded by A_m , the carriers of the dual curves L_1, L_M dual to the initial and terminal 1-cubes of A_m , and the subpath γ'' of γ' between and including the 1-cubes dual to L_1 and L_M . Then every dual curve in E travels from A_m to γ'' – there are M of these – from L_1 to L_M , or from γ'' to L_1 to L_M ; see Figure 3. Indeed, no dual curve K can travel from A_m and cross L_M (or L_1), for otherwise K and L_M (L_1) would cross in D , which is impossible since both emanate from A_m .




Figure 3: The diagram E has three types of dual curve.

Suppose that there exists a constant $B < \infty$ such that, for all $i', j' \geq 0$, the hyperplanes $H_{i'}$ and $H_{j'}$ cannot cross if $|i' - j'| > B$. Hence $M \leq |\gamma''| \leq M + 2B$: the first inequality follows from the fact that each dual curve emanating from A_m ends on γ'' , while the second follows from the fact that at most $2B$ dual curves emanating from γ'' can cross L_1 or L_M .

Let $y \in \gamma''$ be a 0-cube, and let K be a dual curve in E emanating from one of the 1-cubes of γ'' containing y . If K ends on A_m , then there is a path in E of length at most $|K|$ joining y to a 0-cube of A_m . But every dual curve L in E that crosses K has no end on A_m , since distinct dual curves emanating from A_m do not cross. Thus L crosses L_1 or L_M , so that $|K| \leq |L_1| + |L_M|$. Therefore, $d_X(y, U_m) \leq |L_1| + |L_M|$. If K ends on L_1 (or L_M), then y is at distance at most B from L_1 (say), and thus $d_X(y, U_m) \leq B + |L_1|$.

But every dual curve crossing L_1 crosses γ'' , and there are at most B of these, since L_1 maps to a hyperplane crossing $\gamma'' \subset \gamma$. Hence $|L_1| + |L_M| \leq 2B$. Therefore, γ'' is a path of length at least $M > \lambda^{-1}$ lying in the $2B$ -neighborhood of $N(U_m)$. Since B

is independent of λ , we have reached conclusion (3), or B does not exist and, by Lemma 2.3, either (2) or (3) holds. We conclude that if (2) and (3) fail, then there exists $\lambda \in [1, \infty)$ such that

$$|i - j| \geq d_{\Gamma\mathbf{X}}(H_i, H_j) \geq \lambda^{-1}(|i - j| - 1) = \lambda^{-1}d_{\text{Proj}_{\mathbf{X}}(\gamma)}(H_i, H_j) - \lambda^{-1}$$

for all $i, j \geq 0$, and hence $\text{Proj}_{\mathbf{X}}(\gamma) \hookrightarrow \Gamma\mathbf{X}$ is a quasi-isometric embedding. Since $\Lambda(\gamma)$ is a subgraph of $\Gamma\mathbf{X}$ spanned by $\text{Proj}_{\mathbf{X}}(\gamma)$, the inclusion $\Lambda(\gamma) \hookrightarrow \Gamma\mathbf{X}$ is also a quasi-isometric embedding. \square

Lemma 2.3 *Let \mathbf{X} and γ be as in Theorem 2.2. Suppose that conclusions (2) and (3) of Theorem 2.2 do not hold. Then there exists $B < \infty$ such that, if $H_i \perp \smile H_j$, then $|i - j| \leq B$.*

Proof Since Theorem 2.2(2) does not hold, there exists $B' < \infty$ such that, if the subgraph $\Lambda(\gamma)$ of $\Gamma\mathbf{X}$ generated by $\mathcal{W}(\gamma)$ contains $K_{p,p}$, then $p \leq B'$. Since statement (3) does not hold, for each $r \geq 0$, there exists $B''(r) \in [r, \infty)$ such that, if $\gamma' \subset \gamma$ is a connected subpath lying in $N_r(N(H))$, for any hyperplane H , then $|\gamma'| \leq B''(r)$.

Suppose that $H_i \perp \smile H_j$ (hence $|i - j| \geq 1$). Let $\gamma' \subset \gamma$ be the subpath of length $|i - j| - 1$ lying between the 1-cube dual to H_i and that dual to H_j , and joining $a_i \in N(H_i) \cap \gamma$ to $a_j \in N(H_j) \cap \gamma$. Let x be a closest 0-cube of $N(H_i) \cap N(H_j)$ to γ' . If $a_i = a_j$, there is nothing to prove. If $x = a_i$, then by convexity of $N(H_i)$, the path γ' lies in $N(H_i)$ and thus $N(H_i)$ contains a subpath of γ of length at least $|i - j| - 1$. Thus $|i - j| - 1 \leq B''(0)$.

Hence suppose that x, a_i, a_j are all distinct and let m be their median. Since the interval between a_i and x lies in $N(H_i)$, by convexity, $m \in N(H_i)$. Similarly $m \in N(H_j)$. Any hyperplane separating m from γ' separates m from a_i and a_j and hence separates x from a_i and a_j . Thus $m \in N(H_i) \cap N(H_j)$ is at least as close to γ' as is x , and therefore $m = x$, so that x lies on a geodesic segment joining a_i to a_j . It follows that every hyperplane H_k , with $i < k < j$ (say), separates exactly one of a_i and a_j from x . Moreover, every hyperplane separating x from a_i or a_j separates a_i from a_j , since x lies on a geodesic from a_i to a_j . Let $A_i \rightarrow N(H_i)$ be a geodesic segment joining a_i to x , and let $A_j \rightarrow N(H_j)$ be a geodesic segment joining x to a_j , chosen so that $A_i A_j$ is a geodesic segment.

Let e be the midpoint of γ' . Let γ_i be the subpath of γ' joining a_i to e , and let γ_j be the subpath joining e to a_j . Let h_i be the number of hyperplanes crossing γ_i and separating a_i and x from e and let h_j be the number of hyperplanes crossing γ_j and separating a_j and x from e . If H separates e from a_j and x , then H separates a_j

and x from a_i . Similarly, if H' separates e from a_i and x , then H' separates x and a_i from a_j . Hence $H \perp H'$ for all such H, H' . Thus $\Lambda(\gamma)$ contains K_{h_i, h_j} . Thus $\min(h_i, h_j) \leq B'$. As shown in Figure 4, γ_i lies in the h_i -neighborhood of $N(H_i)$, and hence $|\gamma_i| \leq B''(h_i)$. For the same reason, $|\gamma_j| \leq B''(h_j)$. But $|\gamma_i| = |\gamma_j|$, since e is the midpoint of γ' . Thus $|i - j| = |\gamma| + 1 \leq \max(2B''(B') + 1, B''(0)) = B$. \square




Figure 4: Either γ' contains a long path lying in the carrier of a hyperplane or a family of dual curves corresponding to a large biclique in $\Delta\mathbf{X}$.

Theorem 2.4 (Projection trichotomy II) *Let \mathbf{X} and γ be as in Theorem 2.2, and assume that there is no infinite collection of pairwise crossing hyperplanes. Suppose that $\text{Proj}_{\mathbf{X}}(\gamma)$ is bounded in $\Gamma\mathbf{X}$. Then either $\gamma \subset N_R(H)$ for some $R \geq 0$ and some hyperplane H , or γ bounds an eighth-flat in \mathbf{X} .*

Proof Let H_i be the hyperplane dual to $\gamma([i, i + 1])$. Assume that there exists an infinite increasing sequence $0 \leq n_0 < n_1 < \dots$ so that for all $j \geq 1$, the hyperplane H_{n_j} separates $H_{n_{j-1}}$ from $H_{n_{j+1}}$. Since H_{n_i} separates H_{n_0} from H_{n_j} for $0 < i < j$, any path in $\Gamma\mathbf{X}$ from H_0 to H_{n_j} contains a hyperplane V that crosses H_{n_i} . Moreover, if U crosses H_{n_0} and H_{n_j} , then U crosses H_{n_i} . Hence, if (1) does not hold, it follows that there exists $K \geq 0$ such that for all $k > K$, some hyperplane V_k crosses H_{n_i} for all $i \geq k$. Let $\mathcal{H} = \{H_{n_i}\}$ and let \mathcal{V} be the set of all hyperplanes V that cross all but finitely many H_{n_i} .

Verifying (3) Suppose that $|\mathcal{V} \cap \mathcal{W}(\gamma)| < \infty$. By deleting a finite initial segment from γ , we may assume that no hyperplane in \mathcal{V} crosses γ , and some $V \in \mathcal{V}$ crosses H_n for all $n \geq 0$. Indeed, if V crosses every H_{n_i} and does not cross some H_n , then $H_n \in \mathcal{V} \cap \mathcal{W}(\gamma)$, and we can truncate the initial part of γ to remove the 1-cube dual to H_n .

For each $n \geq 0$, if a hyperplane U separates $\gamma(n)$ from $N(V)$, then U either separates all of γ from $N(V)$ or U crosses γ . If U crosses $\gamma([0, n])$, then U cannot cross

$\gamma([n, \infty))$, and hence U crosses H_m for $m \geq n$. In particular, $U \in \mathcal{V} \cap \mathcal{W}(\gamma)$, a contradiction. Thus U crosses $\gamma([n, \infty))$, or U separates γ from $N(V)$. Hence U is one of the finitely many hyperplanes separating $N(V)$ from $\gamma(0)$, and thus we have that $d_{\mathbf{X}}(\gamma(n), N(V)) \leq d_{\mathbf{X}}(\gamma(0), N(V))$ for all $n \geq 0$, and (3) holds.

Building an eighth-flat If for some n the ray $\gamma([n, \infty))$ bounds an eighth-flat, then so does γ . Also, we have reduced to the case in which $\mathcal{V} \cap \mathcal{W}(\gamma)$ is infinite. Hence, by truncating, assume that $H_0 \in \mathcal{V}$. For each $n \geq 0$, let $x_n = \gamma(n)$, and let g_n be the gate of x_n in $N(H_0)$, ie the unique closest 0-cube of the convex set $N(H_0)^0$ to x_n . Define a function f by $f(n) = d_{\mathbf{X}}(x_n, g_n)$. Now, any hyperplane separating x_n from x_m either crosses H_0 and separates g_m from g_n , or separates exactly one of x_m or x_n from $N(H_0)$, since $N(H_0)$ is convex. Thus f is a nondecreasing function. If f is bounded, then (3) holds.

For each $n \geq 0$, any hyperplane H separating g_n from g_{n+1} crosses H_0 , and hence does not separate x_m or x_n from H_0 . Thus H separates x_{n+1} from x_n , so there is a unique combinatorial geodesic ray $\sigma \rightarrow N(H_0)$ such that $\sigma(n) = g_n$ for each $n \geq 0$. (σ is the closest point projection of γ to the convex subcomplex $N(H)$). A discussion of gates in median graphs appears in [38].)

For each $n \geq 0$, let ω_n be a combinatorial geodesic joining x_n to g_n . Observe that, if $g_{n+1} = g_n$, then $\omega_{n+1} = \gamma([n, n+1])\omega_n$. If $g_n \neq g_{n+1}$, then $|\omega_n| = |\omega_{n+1}|$, and ω_n and ω_{n+1} lie on opposite sides of an isometrically embedded “strip” $S_n \cong \omega_n \times [-\frac{1}{2}, \frac{1}{2}]$ in \mathbf{X} . Indeed, the set of hyperplanes separating x_n from g_n coincides with that separating x_{n+1} from g_{n+1} in this case. From the definition, we see that $\mathbf{E} = \bigcup_n S_n$ is an embedded eighth-flat in \mathbf{X} whose bottom bounding ray is σ and whose top bounding ray is γ .

For each $x \in \mathbf{E}^0$, choose a smallest n so that x lies on ω_n . Let $\alpha_1(x) = n$, and let $\alpha_2(x) \leq f(n)$ be the distance from x to g_n in \mathbf{E} . Then for each $x, y \in \mathbf{X}^0$,

$$d_{\mathbf{E}}(x, y) = |\alpha_1(x) - \alpha_1(y)| + |\alpha_2(x) - \alpha_2(y)| \geq d_{\mathbf{X}}(x, y).$$

Let $n = \alpha_1(x), m = \alpha_1(y)$. Since x lies on a geodesic from x_n to g_n , g_n is the gate of x in $N(H_0)$. Likewise, g_m is the gate of y in $N(H_0)$. Hence the distance in \mathbf{X} between the gates in $N(H_0)$ of x and y is $|m - n|$. Since S_m, S_n are isometrically embedded, the distance from y to g_m in \mathbf{X} is $\alpha_2(y) = m$. Thus the number of hyperplanes that separate exactly one of x, y from $N(H_0)$ in \mathbf{X} is $|\alpha_2(x) - \alpha_2(y)|$, so that \mathbf{E} is isometrically embedded.

Existence of $\{H_{n_i}\}$ To conclude, we must show that the claimed family $\{H_{n_i}\}_{i \geq 0}$ of pairwise noncrossing hyperplanes exists. Were this not the case, then there would exist N such that $H_m \perp H_n$ for all $m, n \geq N$, a contradiction. \square

A combinatorially hyperbolic isometry of \mathbf{X} is *combinatorially rank-one* if some (and hence every) combinatorial geodesic axis does not lie in a combinatorial half-flat. The following simple observation relates this definition to rank-one rays.

Proposition 2.5 *Let \mathbf{X} be locally finite and let $g \in \text{Aut}(\mathbf{X})$ be a combinatorially hyperbolic isometry. Then g is combinatorially rank-one if and only if for some (and hence for every) combinatorial geodesic axis γ for g , and for all $m \in \mathbb{Z}$, the ray $\gamma([m, \infty))$ is rank-one.*

Proof Some subray of γ bounds an eighth-flat. Translating this eighth-flat by elements of $\langle g \rangle$ yields an increasing union of isometric eighth-flats, so that γ lies in a half-flat. \square

3 The simplicial boundary

We now turn to the definition and essential properties of the *simplicial boundary* $\partial_{\Delta} \mathbf{X}$ of \mathbf{X} .

3.1 Unidirectional boundary sets

We assume in this section that \mathbf{X} contains no infinite family of pairwise crossing hyperplanes, but not that \mathbf{X} is finite-dimensional or locally finite.

The distinct hyperplanes H_1, H_2, H_3 form a *facing triple* if for each $i \in \{1, 2, 3\}$, there exists a halfspace $\mathfrak{h}(H_i)$ that contains H_j and H_k , for $j, k \in \{1, 2, 3\} - \{i\}$. Equivalently, no set of three halfspaces, one associated to each of the H_i , is totally ordered by inclusion.

Definition 3.1 (Nested, seminested) The set \mathcal{U} of hyperplanes is *nested* if there exists consistent orientation χ of \mathcal{W} , corresponding to a 0-cube or a 0-cube at infinity, such that $\{\chi(U) \mid U \in \mathcal{U}\}$ is totally ordered by inclusion. Note that the hyperplanes in a nested set are pairwise noncrossing. \mathcal{U} is *seminested* if there exists a consistent orientation χ of \mathcal{W} such that $\{\chi(U) \mid U \in \mathcal{U}\}$ is partially ordered by inclusion, and, for all $U, U' \in \mathcal{U}$ such that $\chi(U) \not\subseteq \chi(U')$, either $\chi(U') \subseteq \chi(U)$ or $U \perp U'$.

Given a consistent orientation χ of \mathcal{W} , define a partial ordering $<$ on \mathcal{U} by $U < U'$ if and only if $\chi(U') \subseteq \chi(U)$. Then \mathcal{U} is nested if for some χ , $<$ is a total ordering, and \mathcal{U} is seminested if for some χ , any two distinct hyperplanes in \mathcal{U} cross or are $<$ -comparable. Nested sets are seminested, and a seminested set cannot contain a facing triple.

Definition 3.2 (Unidirectional) Let \mathcal{U} be a set of hyperplanes. Then \mathcal{U} is *unidirectional* if for each $U \in \mathcal{U}$, at least one halfspace associated to U contains only finitely many elements of \mathcal{U} .

Lemma 3.3 A set $\mathcal{U} \subset \mathcal{W}$ that is unidirectional, infinite, inseparable and without a facing triple is seminested.

Proof \mathcal{U} must contain an infinite set of pairwise noncrossing hyperplanes. Since \mathcal{U} is unidirectional, there exists a hyperplane $U_0 \in \mathcal{U}'$ such that the halfspace U_0^- does not contain any element of \mathcal{U} . Since \mathcal{U} does not contain a facing triple, \mathcal{U} contains a set $\mathcal{U}' = \{U_i\}_{i \geq 0}$ such that, for all $i \geq 1$, the hyperplanes U_{i-1} and U_{i+1} are separated by U_i . For each $i \geq 0$, let $\chi(U_i)$ be the halfspace associated to U_i that contains U_{i+1} . If some $U \in \mathcal{W}$ separates some U_i from some U_j , with $j > i$, let $\chi(U)$ be the halfspace containing U_j . Since \mathcal{U} is inseparable, such a U belongs to \mathcal{U} .

More generally, let \mathcal{A} be the set of all hyperplanes $H \in \mathcal{W}$ with $H \subset \mathbf{X} - \chi(U_0)$. Let \mathcal{B} be the set of all hyperplanes $H' \in \mathcal{W}$ with $H' \subset \chi(U_0)$. Let \mathcal{C} be the set of hyperplanes $H'' \in \mathcal{W}$ that cross U_0 . For each $H \in \mathcal{A}$, let $\chi(H)$ be the halfspace associated to H that contains U_0 , ie orient H toward U_0 . If $H \in \mathcal{U} \cap \mathcal{B}$, let $\chi(H)$ be the halfspace that does not contain U_0 . If $H \in \mathcal{B} - \mathcal{U}$, then for all $U \in \mathcal{U}$, either $U \perp H$ or U and U_0 lie in the same halfspace associated to H . In either event, let $\chi(H)$ be the halfspace containing U_0 . In other words, for all hyperplanes H that do not cross U_0 , χ orients H away from U_0 exactly when $H \notin \mathcal{U}$. If H does not cross U , and $H' \neq U$ does not cross U , then $\chi(H) \cap \chi(H') \neq \emptyset$.

If $H \perp U$, then any orientation of H is consistent with $\chi|_{\mathcal{A}}$, since χ orients each $A \in \mathcal{A}$ toward $H \cap U$. Similarly, any orientation of H is consistent with $\chi|_{\mathcal{B} - \mathcal{U}}$. If $H \subset (\mathbf{X} - \chi(U_i))$ for some $i \geq 1$, let $\chi(H)$ be the halfspace containing U_i . This is consistent with $\chi|_{\mathcal{B} \cap \mathcal{U}}$. Moreover, if H' is another hyperplane that crosses H but fails to cross some U_j , then $\chi(H') \cap \chi(H)$ contains $U_{\max(i,j)}$, so χ is consistent across all such H, H' . The remaining hyperplanes are those H that cross U_i for all $i \geq 0$. Choose a fixed 0-cube $x_0 \in N(U_0)$ and let χ orient all of these H toward x_0 . If H is such a hyperplane, and $H' \perp U$ and $H' \not\perp U_i$, then $N(H) \cap N(U_i)$ lies in $\chi(H')$ and contains a 0-cube in $\chi(H)$. Thus χ is a consistent orientation of all hyperplanes of \mathbf{X} . Since $\bigcap_{i \geq 0} \chi(U_i) = \emptyset$, χ is a 0-cube at infinity. Finally, for $U, U' \in \mathcal{U}$, let $U < U'$ if and only if $\chi(U') \subseteq \chi(U)$. This is clearly a partial ordering, and every $<$ -chain in \mathcal{U} has a minimum since \mathcal{U} is unidirectional.

Let \mathcal{C} be the set of hyperplanes C such that $C \perp U_i$ for all $i \geq 0$. Recall that $\chi(C)$ is the halfspace containing x_0 , and that the orientation χ is consistent on all of \mathcal{W} for arbitrarily chosen $x_0 \in N(H_0)$. Observe that since \mathcal{U} is inseparable and contains no

facing triple, \prec fails to have the desired comparability property on \mathcal{U} if and only if there exist $C, C' \in \mathcal{U} \cap \mathcal{C}$ that do not cross, with $x_0 \in N(C) \cap N(C')$. Hence the comparability requirement holds for χ except on the set $\mathcal{U} \cap \mathcal{W}(\chi(U_0) \cap \chi(C))$. Replacing χ by the 0-cube χ_1 at infinity that coincides with χ on all hyperplanes except C preserves consistency, but χ_1 yields a partial order satisfying the comparability requirement on a subset of \mathcal{U} that properly contains $\mathcal{U} \cap \mathcal{W}(\chi(U_0) \cap \chi(C))$. We can thus write \mathcal{U} as an increasing union of subsets \mathcal{U}^n such that, for each n , there is a 0-cube χ_n at infinity such that for all $U, U' \in \mathcal{U}^n$, either $U \perp U'$ or $\chi(U) \subseteq \chi(U')$ or $\chi(U') \subseteq \chi(U)$. Furthermore, the restriction of χ_m to \mathcal{U}^n coincides with χ_n whenever $m \geq n$. We thus define a 0-cube η at infinity by $\eta(H) = \chi_n(H)$ for some (and hence every) n for which $H \in (\mathcal{W} - \mathcal{U}) \cup \mathcal{U}_n$. By construction, η is consistent and the associated partial order on \mathcal{U} is as required. \square

We now abstract the essential properties of the set \mathcal{U} in [Lemma 3.3](#).

Definition 3.4 (Unidirectional boundary set, almost equivalent) A *unidirectional boundary set (UBS)* \mathcal{U} is an infinite, inseparable, unidirectional set of hyperplanes containing no facing triple. The UBSs \mathcal{U}' and \mathcal{U} are *almost equivalent* if their symmetric difference is finite. The UBS \mathcal{U} is *minimal* if any UBS $\mathcal{U}' \subseteq \mathcal{U}$ is almost equivalent to \mathcal{U} . Almost-equivalence of UBSs is an equivalence relation.

Example 3.5 If $\gamma: [0, \infty) \rightarrow \mathbf{X}$ is a geodesic ray, then $\mathcal{W}(\gamma)$ is a UBS. If γ' is another ray, then γ and γ' are almost equivalent if and only if $\mathcal{W}(\gamma)$ and $\mathcal{W}(\gamma')$ are almost equivalent UBSs. An associated 0-cube at infinity orients every $H \in \mathcal{W} - \mathcal{W}(\gamma)$ toward $\gamma(0)$, and orients each $H \in \mathcal{W}(\gamma)$ away from $\gamma(0)$.

Definition 3.6 (Inseparable closure) Let $\mathcal{U} \subset \mathcal{W}$. The *inseparable closure* $\overline{\mathcal{U}}$ of \mathcal{U} is the intersection of all inseparable sets \mathcal{V} such that $\mathcal{U} \subseteq \mathcal{V}$. Note that $\overline{\mathcal{U}}$ consists of \mathcal{U} , together with the set of hyperplanes V for which there exist $U, U' \in \mathcal{U}$ such that V separates U and U' .

Lemma 3.7 If \mathbf{X} contains no infinite family of pairwise crossing hyperplanes and \mathcal{W} contains a UBS, then \mathcal{W} contains a minimal UBS. Moreover, every essential hyperplane is contained in a UBS.

Remark 3.8 The hypotheses are satisfied if \mathbf{X} is unbounded and strongly locally finite, or, more generally, if \mathbf{X} has no infinite family of pairwise crossing hyperplanes and contains at least one deep halfspace.

Proof of Lemma 3.7 By hypothesis, there exists an infinite set $\mathcal{U}' = \{U_i\}_{i \geq 0}$ of hyperplanes such that, for all $i \geq 1$, the hyperplane U_i separates U_{i-1} from U_{i+1} . Let \mathcal{U} be the inseparable closure of \mathcal{U}' . Clearly \mathcal{U} is infinite and inseparable.

\mathcal{U} contains no facing triple Since \mathcal{U}' does not contain a facing triple, any facing triple in \mathcal{U} is of the form U_i, U_j, V_1 or U_i, V_1, V_2 or V_1, V_2, V_3 , where $V_1, V_2, V_3 \in \mathcal{U} - \mathcal{U}'$. If U_i, U_j, V_1 is a facing triple, then V_1 separates some U_k from U_i , and hence from U_j . Thus U_i, U_j, U_k is a facing triple, a contradiction. If U_i, V_1, V_2 is a facing triple, then V_1 separates some U_j from U_i , and thus U_i, U_j, V_2 form a facing triple, contradicting the previous statement. By like reasoning, V_1, V_2, V_3 cannot form a facing triple.

\mathcal{U} is unidirectional Let $U \in \mathcal{U}$ have the property that both associated halfspaces contain infinitely many elements of \mathcal{U} . Since there exist U_i, U_j separated by U , with $i < j$, and only finitely many hyperplanes separate U_i from U_j , either the halfspace U_i^- not containing U contains infinitely many hyperplanes, or infinitely many elements of \mathcal{U} cross U_i and do not cross U . Now, no element of \mathcal{U} crosses U_0 , by the definition of the inseparable closure. Hence infinitely many hyperplanes $V \in \mathcal{U}$ cross U_i but do not cross U or U_0 . Either U, U_0, V form a facing triple, or V is among the finitely many hyperplanes separating U from U_0 . This is a contradiction, and it follows that U is unidirectional.

\mathcal{U} contains a minimal UBS Let $\mathcal{S} \subseteq \mathcal{U}$ be a UBS. If \mathcal{S} does not contain U_i , then \mathcal{S} cannot contain any $U \in \mathcal{U}$ that separates U_i from U_0 . Thus there exists some minimal $I < \infty$ such that \mathcal{S} contains U_i for $i \geq I$. If $U \in \mathcal{U}$ does not separate U_I from U_0 and does not cross U_i , then inseparability of \mathcal{S} requires that $U \in \mathcal{S}$. Hence every element of $\mathcal{U} - \mathcal{S}$ either separates U_0 from U_I or crosses U_I and separates U_0 from U_i for some $i \geq I$.

Indeed, if $U \in \mathcal{U}$, then U cannot cross all but finitely many of the U_i , for otherwise $\mathcal{U} - \{U\}$ would be a smaller inseparable set containing \mathcal{U}' . Let \mathcal{C}_1 be the set of hyperplanes in \mathcal{U} that cross U_i . If \mathcal{C}_1 is finite, then $\mathcal{U} - \mathcal{S}$ has cardinality $|\mathcal{C}_1| + d_X(N(U_0), N(U_i))$ and so \mathcal{S} is almost equivalent to \mathcal{U} . If \mathcal{C}_1 is infinite, then let $\mathcal{U}_1 = \mathcal{U} - \mathcal{S}$. Note that \mathcal{C}_1 is unidirectional and contains no facing triple, since those properties are inherited by subsets. If $C, C' \in \mathcal{C}_1$ are separated by some hyperplane W , then W crosses U_I and crosses each hyperplane U_j crossed by both C and C' . Moreover, since W does not cross C or C' , W does not cross U_j for j sufficiently large. Hence $W \in \mathcal{C}_1$, which is therefore a UBS. Likewise, \mathcal{U}_1 is infinite, unidirectional and contains no facing triple. By the definition of \mathcal{C}_1 , \mathcal{U}_1 is also inseparable. Hence \mathcal{U}_1 is a UBS. Moreover, every element of \mathcal{U}_1 crosses all but finitely many elements of \mathcal{C}_1 .

Hence, if \mathcal{U} is not minimal, then \mathcal{U} contains two disjoint UBSs, \mathcal{U}_1 and \mathcal{C}_1 , such that every hyperplane in \mathcal{U}_1 crosses all but finitely many hyperplanes in \mathcal{C}_1 . If \mathcal{U}_1 is not minimal, then by the same reasoning, \mathcal{U}_1 is almost equivalent to a disjoint union $\mathcal{U}_2 \sqcup \mathcal{C}_2$ of UBSs such that each $U \in \mathcal{U}_2$ crosses all but finitely many $C \in \mathcal{C}_2$, and each

of U and C crosses all but finitely many $C' \in \mathcal{C}_1$. Apply the same argument to \mathcal{U}_2 , and continue in this way. Since \mathbf{X} does not contain an infinite set of pairwise crossing hyperplanes, we must, after finitely many applications of this argument, find a minimal UBS contained in \mathcal{U} .

Essential hyperplanes If $U \in \mathcal{W}$ is essential, then by definition there exists an infinite nested set of hyperplanes contained in a halfspace associated to U , and we argue as above. □

Remark 3.9 The proof of Lemma 3.7 shows that if \mathbf{X} contains no infinite collection of pairwise crossing hyperplanes, and $\mathcal{U} \subset \mathcal{W}$ is a UBS, then \mathcal{U} contains a minimal UBS. Indeed, \mathcal{U} must contain a unidirectional nested set of hyperplanes, whose inseparable closure is contained in \mathcal{U} .

The following theorem allows us to define the simplicial boundary of \mathbf{X} .

Theorem 3.10 Suppose that every collection of pairwise crossing hyperplanes in \mathbf{X} is finite. Let v be an almost equivalence class of UBSs. Then v has a representative of the form $\mathcal{V} = \bigsqcup_{i=1}^k \mathcal{U}_i$, where $k < \infty$, each \mathcal{U}_i is minimal, and for all $1 \leq i < j \leq k$, if $H \in \mathcal{U}_j$, then H crosses all but finitely many elements of \mathcal{U}_i .

Moreover, if $\mathcal{V}' = \bigsqcup_{i=1}^{k'} \mathcal{U}'_i$ is almost equivalent to \mathcal{V} , and each \mathcal{U}'_i is minimal, then $k = k'$ and, up to reordering, \mathcal{U}_i and \mathcal{U}'_i are almost equivalent for all i .

Proof If \mathcal{V} is a minimal UBS, then by definition any UBS $\mathcal{U} \subset \mathcal{V}$ is almost equivalent to \mathcal{V} , and the desired decomposition into minimal UBSs exists and is unique in the required sense. We now proceed inductively. Let \mathcal{U}_1 be a minimal UBS contained in \mathcal{V} . Let $\mathcal{V}_1 = \mathcal{V} - \mathcal{U}_1$. If \mathcal{V}_1 is finite, then \mathcal{V} is almost equivalent to \mathcal{U}_1 , and \mathcal{V} is already minimal.

Let $V \in \mathcal{V}_1$ be a hyperplane. If V crosses infinitely many hyperplanes of \mathcal{U}_1 , then V crosses all but finitely many hyperplanes of \mathcal{U}_1 since, for all $U \in \mathcal{U}_1$, all but finitely many hyperplanes of \mathcal{U}_1 lie in a common halfspace associated to U . Let $\mathcal{V}'_1 \subset \mathcal{V}_1$ be the set of hyperplanes in $\mathcal{V} - \mathcal{U}_1$ that cross only finitely many elements of \mathcal{U}_1 . If \mathcal{V}'_1 is infinite, then since \mathbf{X} is strongly locally finite, \mathcal{V}'_1 contains an infinite nested set $\{V_i\}_{i \geq 0}$ with initial hyperplane V_0 such that for all but finitely many $U \in \mathcal{U}_1$, V_0 separates U from each V_i with $i \geq 1$. This implies that $\mathcal{U}_1 \cup \mathcal{V}'_1$ contains a nested set with no initial hyperplane, contradicting the fact that \mathcal{V} is a UBS. Thus \mathcal{V}'_1 is finite and can be assumed to be empty without affecting the almost equivalence class of \mathcal{V} . Hence each $V \in \mathcal{V}_1$ crosses all but finitely many elements of \mathcal{U}_1 , and thus \mathcal{V}_1 contains an infinite inseparable set. \mathcal{V}_1 thus contains a minimal UBS \mathcal{U}_2 , by Lemma 3.7. Moreover, each $U \in \mathcal{U}_2$ crosses all but finitely many of the hyperplanes in \mathcal{U}_1 .

By induction, we obtain a (a priori, infinite) maximal index set $K \subseteq \mathbb{N}$ and a subset $\bigsqcup_{k \in K} \mathcal{U}_k \subseteq \mathcal{V}$, where each \mathcal{U}_k is a minimal UBS and for all $k < n$, each $U \in \mathcal{U}_n$ crosses all but finitely many elements of \mathcal{U}_k . We can thus choose, for each $k \in K$, $U_k \in \mathcal{U}_k$ such that the hyperplanes U_k, U_j cross whenever $j \neq k$. Hence $|K| < \infty$ since \mathbf{X} is strongly locally finite. (If $\dim \mathbf{X} < \infty$, then $|K| \leq \dim(\mathbf{X})$).

Consider the set $\mathcal{F} = \mathcal{V} - \bigsqcup_{k \in K} \mathcal{U}_k \subseteq \mathcal{V}$. If \mathcal{F} is infinite, \mathcal{F} contains a UBS, contradicting maximality of K . Hence \mathcal{F} is finite, and we may remove \mathcal{F} from \mathcal{V} without affecting the almost equivalence class v of \mathcal{V} .

Finally, suppose $\mathcal{V}' = \bigcup_{k' \in K'} \mathcal{U}'_{k'}$ is another such decomposition into minimal inseparable sets, with \mathcal{V}' almost equivalent to \mathcal{V} . Suppose there exist k'_1 and $k'_2 \in K'$ and $k \in K$ such that \mathcal{U}_k has infinite intersection with $\mathcal{U}'_{k'_i}$ for $i = 1, 2$. By inseparability, for $i \in \{1, 2\}$, we have that $\mathcal{U}_k \cap \mathcal{U}'_{k'_i}$ contains all but finitely many elements of \mathcal{U}_k and $\mathcal{U}'_{k'_i}$ (see below). Without affecting the almost equivalence class, we can remove finitely many hyperplanes from the symmetric difference of these sets to conclude that $\mathcal{U}_k = \mathcal{U}'_{k'_1} = \mathcal{U}'_{k'_2}$, and hence that $k_1 = k_2$. Thus we have \mathcal{U}_k is almost equivalent to at most one of the sets $\mathcal{U}'_{k'}$.

Suppose \mathcal{U}_k is not almost equivalent to $\mathcal{U}'_{k'}$ for any $k' \in K'$. Then \mathcal{U}_k has finite intersection with each $\mathcal{U}'_{k'}$. But \mathcal{U}_k is infinite and contained in the finite union $\bigcup_{k' \in K'} \mathcal{U}'_{k'}$, a contradiction. Thus \mathcal{U}_k is almost equivalent to at least, and thus exactly, one of the $\mathcal{U}'_{k'}$.

We used the fact that, if $\mathcal{U}_1, \mathcal{U}_2$ are minimal UBSs, then $\mathcal{U}_1 \cap \mathcal{U}_2$ is infinite if and only if \mathcal{U}_1 and \mathcal{U}_2 are almost equivalent. Indeed, if \mathcal{U}_1 is almost equivalent to \mathcal{U}_2 , then, since each is infinite and their symmetric difference finite, their intersection must be infinite. Conversely, if $\mathcal{U}_3 = \mathcal{U}_1 \cap \mathcal{U}_2$ is infinite, then it contains a UBS \mathcal{U}'_3 . But \mathcal{U}'_3 is almost equivalent to \mathcal{U}_1 . Likewise, \mathcal{U}'_3 is almost equivalent to \mathcal{U}_2 , whence $\mathcal{U}_1 \Delta \mathcal{U}_2$ is finite. □

3.2 Simplices at infinity

The unidirectional boundary set \mathcal{V} is k -dimensional if the decomposition of \mathcal{V} into minimal UBSs given by Theorem 3.10 has k factors. Theorem 3.10 implies that, if v is an almost equivalence class of UBSs, then there is a unique integer k such that every UBS representing v is k -dimensional. A 0-simplex at infinity is an almost equivalence class of minimal UBSs. More generally, a k -simplex at infinity is an almost equivalence class of UBSs whose representatives are $(k + 1)$ -dimensional. If u, v are simplices at infinity, let $u \leq v$ if and only if there exist representatives \mathcal{U} of u and \mathcal{V} of v such that $\mathcal{U} \subseteq \mathcal{V}$. By Theorem 3.10, \leq is a partial ordering on the set of almost equivalence classes.

Definition 3.11 (Simplicial boundary) Suppose that each collection of pairwise crossing hyperplanes in \mathbf{X} is finite. The *simplicial boundary* $\partial_{\Delta}\mathbf{X}$ of \mathbf{X} is the geometric realization of the abstract simplicial complex whose set of simplices is the set of simplices at infinity, in which u is a face of v if and only if $u \leq v$.

A 0-simplex $v \in \partial_{\Delta}\mathbf{X}$ is *isolated* if v is not contained in a 1-simplex. The graph $(\partial_{\Delta}\mathbf{X})^1$ is equipped with the standard path metric, and two 0-simplices are at distance $+\infty$ if they lie in distinct components of $\partial_{\Delta}\mathbf{X}$. If $(\partial_{\Delta}\mathbf{X})^1$ has finite diameter in this metric, we say $\partial_{\Delta}\mathbf{X}$ is *bounded*. Otherwise, $\partial_{\Delta}\mathbf{X}$ is *unbounded*. Define $\text{diam}(\partial_{\Delta}\mathbf{X})$ to be the diameter of $(\partial_{\Delta}\mathbf{X})^1$ with the standard path metric.

Example 3.12 The simplicial boundary of an eighth-flat or a nondiagonal quarter-flat is a single 1-simplex, that of a diagonal quarter-flat or nondiagonal half-flat is a subdivided interval of length two, that of a diagonal half-flat is a subdivided interval of length three, and that of a flat is a 4-cycle. The simplicial boundary of a tree with more than one end is totally disconnected.

Remark 3.13 If \mathcal{W} contains an infinite set \mathcal{U} of pairwise crossing hyperplanes, then the preceding method of defining simplices is unavailable. Indeed, \mathcal{U} is a UBS. However, let $\{U_n\}_{n=1}^{\infty}$ be a countably infinite subset of \mathcal{U} , and for each $m \geq 2$, note that $\{U_{mn}\}_{n \geq 1}$ is a UBS that is not almost equivalent to \mathcal{U} . In particular, $\dots\{U_{2^k n}\}_n \subset \{U_{2^{k-1} n}\}_n \dots \subset \{U_n\}_n \subset \mathcal{U}$ is an infinite chain of UBSs, no two of which are almost equivalent. Thus \mathcal{U} contains no minimal UBS and it is not clear how to associate a simplex with the almost equivalence class of \mathcal{U} in a way that is compatible with the assignment of simplices made available by [Theorem 3.10](#). It would be undesirable to ignore UBSs of this type, however, because one must then accept that there are geodesic rays that do not represent simplices of $\partial_{\Delta}\mathbf{X}$. This is why we disallow such \mathbf{X} when defining the simplicial boundary. Hence $\partial_{\Delta}\mathbf{X}$ is defined as long as every family of pairwise crossing hyperplanes is finite, and nonempty as long as some halfspace is deep.

Theorem 3.14 (Basic properties of $\partial_{\Delta}\mathbf{X}$) If $\partial_{\Delta}\mathbf{X}$ is defined, then

- (1) $\partial_{\Delta}\mathbf{X} = \emptyset$ if and only if for all hyperplanes H , each halfspace $\mathfrak{h}(H)$, $\mathfrak{h}^*(H)$ is contained in a uniform neighborhood of H , ie $\mathfrak{h}(H)$ is shallow;
- (2) each simplex of $\partial_{\Delta}\mathbf{X}$ is contained in a finite-dimensional maximal simplex;
- (3) $\partial_{\Delta}\mathbf{X}$ is a flag complex;
- (4) if \mathbf{X} is locally finite and is compactly decomposable into at least two unbounded components, then $\partial_{\Delta}\mathbf{X}$ is disconnected;
- (5) if \mathbf{X} is unbounded and hyperbolic, then $\partial_{\Delta}\mathbf{X}$ is a totally disconnected set.

Proof If each halfspace is shallow, then \mathcal{W} cannot contain a UBS, and thus $\partial_\Delta \mathbf{X}$ is empty. If $\partial_\Delta \mathbf{X}$ is empty, then \mathcal{W} does not contain a UBS. Lemma 3.7 implies that for each halfspace $\mathfrak{h}(H)$, every hyperplane $V \subset \mathfrak{h}(H)$ lies at finite Hausdorff distance from H , so that $\mathfrak{h}(H)$ is shallow. This proves assertion (1). To prove assertion (4), apply (1) and Theorem 3.15 below.

Assertion (2) follows from Theorem 3.10. Indeed, the proof of Theorem 3.10 shows that if $\partial_\Delta \mathbf{X}$ contains an infinite, increasing union of simplices, then \mathbf{X} contains an infinite family of pairwise crossing hyperplanes. In particular, if $\dim(\mathbf{X}) = D$, then $\dim(\partial_\Delta \mathbf{X}) \leq D - 1$. Assertion (3) then follows from (2) by induction on the size of a clique in $(\partial_\Delta \mathbf{X})^1$, using the definition of a UBS. Finally, if $\partial_\Delta \mathbf{X}$ contains a 1-simplex, then $\Delta \mathbf{X}$ contains $K_{p,p}$ for p arbitrarily large, by Theorem 3.10. Such an \mathbf{X} is not hyperbolic, by the proof of [21, Theorem 7.6]. \square

The simplicial boundary behaves in the same way with respect to convex subcomplexes as the visual boundary of a CAT(0) space does with respect to convex subspaces.

Theorem 3.15 *Let $\mathbf{Y} \subseteq \mathbf{X}$ be a convex subcomplex. Then $\partial_\Delta \mathbf{Y}$ is a subcomplex of $\partial_\Delta \mathbf{X}$. Moreover, if $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$, where \mathbf{X}_1 and \mathbf{X}_2 are convex subcomplexes, then $\partial_\Delta \mathbf{X} \cong A_1 \cup A_2$, where $A_1 \cong \partial_\Delta \mathbf{X}_1$, $A_2 \cong \partial_\Delta \mathbf{X}_2$ and $A_1 \cap A_2 \cong \partial_\Delta (\mathbf{X}_1 \cap \mathbf{X}_2)$.*

Proof Let u be a simplex of $\partial_\Delta \mathbf{Y}$. Let \mathcal{U} be a UBS of hyperplanes of \mathbf{Y} representing u . Each $U \in \mathcal{U}$ is of the form $U' \cap \mathbf{Y}$, where U' is a hyperplane of \mathbf{X} crossing \mathbf{Y} . Now, by convexity, $U' \cap \mathbf{Y}$ and $U'' \cap \mathbf{Y}$ contact in \mathbf{Y} if and only if U' and U'' contact (in the same way) in \mathbf{X} . Moreover, $U''' \cap \mathbf{Y}$ separates $U' \cap \mathbf{Y}$ and $U'' \cap \mathbf{Y}$ if and only if U''' separates U' and U'' in \mathbf{X} . Hence $\mathcal{U}' = \{U' \in \mathcal{W} \mid U' \cap \mathbf{Y} \in \mathcal{U}\}$ is a UBS in \mathbf{X} , representing a simplex u' of $\partial_\Delta \mathbf{X}$. Since elements of \mathcal{U}' that cross in \mathbf{X} cross in \mathbf{Y} , Theorem 3.10 implies that $\dim(u') = \dim(u)$.

The assignment $u \mapsto u'$ therefore yields a simplicial map $\partial_\Delta \mathbf{Y} \rightarrow \partial_\Delta \mathbf{X}$ since the intersection of two simplices in $\partial_\Delta \mathbf{Y}$ yields a simplex of $\partial_\Delta \mathbf{X}$ in the same way as above. This map is injective, since UBSs in \mathbf{Y} that correspond to almost equivalent UBSs in \mathbf{X} are almost equivalent.

In particular, if $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$ is the union of two convex subcomplexes, let $A_1 \cong \partial_\Delta \mathbf{X}_1$, $A_2 \cong \partial_\Delta \mathbf{X}_2$. Then every hyperplane crosses \mathbf{X}_1 or \mathbf{X}_2 , and hence $\partial_\Delta \mathbf{X} \cong A_1 \cup A_2$. Now, $\mathbf{X}_1 \cap \mathbf{X}_2$ is convex, being the intersection of convex subcomplexes, and thus $\partial_\Delta \mathbf{X}$ contains $\partial_\Delta (\mathbf{X}_1 \cap \mathbf{X}_2) = B$. If $u \subset A_1 \cap A_2$ is a simplex, then it is represented by a UBS, all of whose elements cross $\mathbf{X}_1 \cap \mathbf{X}_2$. Hence $A_1 \cap A_2 \subseteq B$, and the other inclusion is similar. \square

3.3 Visible simplices, visible pairs and combinatorial geodesic completeness

Recall that if $\gamma: [0, \infty) \rightarrow \mathbf{X}$ is a combinatorial geodesic ray, then $\mathcal{W}(\gamma)$ is a UBS. The converse is not true, as illustrated by [Example 3.17](#).

Definition 3.16 (Visible simplex, visible pair, fully visible, optical space) The simplex $v \subseteq \partial_{\Delta} \mathbf{X}$ is *visible* if there exists a combinatorial geodesic ray γ such that $\mathcal{W}(\gamma)$ represents v . The pair $u, v \subset \partial_{\Delta} \mathbf{X}$ of visible simplices is a *visible pair* if there exists a combinatorial geodesic $\gamma: \mathbb{R} \rightarrow \infty$ such that $\mathcal{W}(\gamma((-\infty, 0]))$ represents u and $\mathcal{W}(\gamma([0, \infty))$ represents v . \mathbf{X} is *fully visible* if each simplex of $\partial_{\Delta} \mathbf{X}$ is visible. \mathbf{X} is an *optical space* if, whenever the pair u, v of 0–simplices at infinity is invisible, $u = v$.

Example 3.17 (An invisible 0–simplex at infinity) In an eighth-flat, the set of “horizontal hyperplanes” is a UBS representing an invisible 0–simplex at infinity: every geodesic ray crosses infinitely many vertical hyperplanes.

The following lemma is a simple application of Sageev’s construction.

Lemma 3.18 *Let \mathcal{U} be a nonempty, finite, inseparable set of hyperplanes containing no facing triple. Then there exists a geodesic segment $P \rightarrow \mathbf{X}$ such that $\mathcal{W}(P) = \mathcal{U}$.*

Every simplex of $\partial_{\Delta} \mathbf{X}$ is contained in a visible simplex.

Theorem 3.19 *Let \mathbf{X} be a CAT(0) cube complex with no infinite family of pairwise crossing hyperplanes. Then every maximal simplex of $\partial_{\Delta} \mathbf{X}$ is visible. In particular, each isolated 0–simplex is visible.*

Proof Let \mathcal{V} be a UBS, seminested by a 0–cube f at infinity, inducing a partial ordering $<$ on \mathcal{V} such that any two elements of \mathcal{V} cross or are comparable, no infinite $<$ –chain has a maximum, and each infinite $<$ –chain has a minimum.

A sequence of segments By definition, there exists at least one $<$ –minimum $V_0 \in \mathcal{V}$. Fix V_0 and choose $\mathcal{V}_0 \subset \mathcal{V}_1 \subset \dots \subset \mathcal{V}_r \subset \dots$ such that $|\mathcal{V}_r| = r + 1$, each \mathcal{V}_r is inseparable, and $\bigcup_{r \geq 0} \mathcal{V}_r = \mathcal{V}$. Also, make this choice so that the hyperplane $V_0 \in \mathcal{V}_0$ is $<$ –minimal in \mathcal{V}_0 and has the property that every element of $\bigcup_r \mathcal{V}_r$ either lies in $f(V_0)$ or crosses V_0 . Note that if V_0 is maximal and minimal, then every hyperplane crosses V_0 , and we can choose $\mathcal{V}_0 = \{V_0\}$.

By [Lemma 3.18](#), for each $r \geq 0$, there exists a 0–cube $c_r \in N(V_0)$ and a geodesic segment $P_r \rightarrow \mathbf{X}$ emanating from c_r such that $\mathcal{W}(P_r) = \mathcal{V}_r$.

Let Ω be the graph whose vertex set is $\bigcup_{r \geq 0} \mathcal{R}_r$, where \mathcal{R}_r is the set of geodesic segments P_r whose 1-cubes are dual to exactly the hyperplanes in \mathcal{V}_r . Let an edge join $P_r \in \mathcal{R}_r$ to $P_{r+1} \in \mathcal{R}_{r+1}$ if and only if $P_r \subset P_{r+1}$. The graph Ω is infinite and locally finite. Indeed, local finiteness follows from the absence of infinite cubes in \mathbf{X} . If P_r is a geodesic segment crossing \mathcal{V}_r and extending to P_{r+1} , then let c' be the 1-cube of $P_{r+1} - P_r$. If V is dual to the terminal 1-cube of P_r adjacent to c' , and $V_1, \dots, V_n \in \mathcal{V}_{r+1} - \mathcal{V}_r$ are dual to 1-cubes c_1, \dots, c_n extending P_r , then since \mathcal{V} contains no facing triple, the set V, V_1, \dots, V_n contains a collection of n pairwise crossing hyperplanes, and thus n is bounded above by the dimension of the maximal cube containing c' . Hence each P_r extends in finitely many ways in each direction, and thus Ω is an infinite, locally finite graph.

Taking a limit to obtain a ray If Ω has an infinite connected component, then by König's lemma, there is an infinite path $\gamma \rightarrow \mathbf{X}$ that crosses every hyperplane of $\bigcup_{r \geq 0} \mathcal{V}_r = \mathcal{V}$ exactly once. Hence γ is a geodesic ray whose initial 1-cube is dual to the given hyperplane V_0 . By construction, any hyperplane crossing γ belongs to some (and hence all but finite many) \mathcal{V}_r , whence \mathcal{V} represents a visible simplex.

When the limit does not exist If every component of Ω is finite, then for each n and each $\phi \in \mathcal{R}_n$, there exists a maximal $k(\phi) \geq n$ such that ϕ extends to an element of $\mathcal{R}_{k(\phi)}$. Suppose there is a hyperplane $U(\phi)$ separating ϕ from infinitely many $\psi \in \bigcup_{m > k(\phi)} \mathcal{R}_m$. Since it separates the initial point of ϕ from the initial point of infinitely many such segments ψ , and ϕ and each such ψ emanate from $N(V_0)$, the hyperplane $U(\phi)$ crosses V_0 . Moreover, since $U(\phi)$ separates ϕ from each $\psi \in \bigcup_{m \geq k(\phi)} \mathcal{R}_m$, the hyperplane $U(\phi) \notin \mathcal{V}$, since each hyperplane of \mathcal{V} is dual to a 1-cube of infinitely many of the segments ψ .

Hence we can choose a subsequence $(i_j)_{j \geq 1}$ and, for each $j \geq 1$, a segment $\phi_j \in \mathcal{R}_{i_j}$ and a hyperplane $U_j \notin \mathcal{V}$ that separates ϕ_j from ϕ_k for all $k > j$, and thus crosses V_0 ; see the left side of [Figure 5](#).

Note that if $V \in \mathcal{V}_{i_j}$, then V crosses ϕ_k for all $k \geq j$, and hence V crosses U_k for all $k \geq j$. On the other hand, if some hyperplane U separates U_i from U_j with $i < j$, then U must cross every hyperplane crossed by both U_i and U_j . Hence there is an infinite, inseparable, partially ordered set \mathcal{U} , containing all of the U_j , such that, for all but finitely many $V \in \mathcal{V}$, V crosses all but finitely many $U \in \mathcal{U}$.

Finding a larger simplex The set $\mathcal{V}' = \mathcal{V} \cup \mathcal{U}$ is a UBS containing \mathcal{V} and representing a simplex v' of $\partial_\Delta \mathbf{X}$ that properly contains the simplex v represented by \mathcal{V} . Hence each maximal simplex of \mathbf{X} is visible. By [Theorem 3.14](#), the simplex v is contained in a maximal simplex and hence \mathcal{V} extends to a UBS almost equivalent to $\mathcal{W}(\gamma)$ for some geodesic ray γ .




Figure 5: The situation arising in the absence of a ray associated to a simplex at infinity

Proof that $U(\phi)$ exists Suppose $\phi \in \mathcal{R}_n$ does not extend to any $\psi \in \mathcal{R}_m$ for $m > k(n)$. Either infinitely many $\psi \in \bigcup_{m>k(n)} \mathcal{R}_m$ are separated from ϕ by a hyperplane $U \notin \mathcal{V}$, or we are in the following situation for all but finitely many such ψ . Let $m > k(n) \geq n + 1$, so that $\psi = \phi'\alpha$, where ϕ' is a length- n geodesic segment crossing \mathcal{V}_n and α is a geodesic segment crossing $\mathcal{V}_m - \mathcal{V}_n$. Then $\phi(n)$ and $\phi'(n)$ are separated by a hyperplane U that crosses each of the hyperplanes in $\mathcal{V}_m - \mathcal{V}_n$, for otherwise ϕ would extend to a segment $\phi\alpha \in \mathcal{R}_m$. If $U \notin \mathcal{V}$, then we are done, since U must then separate ϕ from ψ . Otherwise, U crosses ϕ or ϕ' and thus $U \in \mathcal{V}_n$. Hence $\phi \cap \phi' \neq \emptyset$, so that some subsegment of ϕ extends to a segment in \mathcal{R}_m . If this holds for arbitrarily large values of m , with ϕ fixed, then König's lemma yields a ray whose initial 0-cube lies on ϕ and whose set of dual hyperplanes is \mathcal{V} . Thus we may suppose that there exists $U \notin \mathcal{V}$ that separates ϕ and ψ . □

Theorems 3.14 and 3.19 imply that a hyperbolic cube complex is fully visible. More generally, we have the following.

Corollary 3.20 *Let γ be a rank-one combinatorial geodesic ray representing a simplex v of $\partial_\Delta \mathbf{X}$. Then v is an isolated 0-simplex.*

Proof Suppose that $\mathcal{W}(\gamma) = \mathcal{V}_1 \sqcup \mathcal{V}_2$, where \mathcal{V}_1 and \mathcal{V}_2 are disjoint UBSs. By Theorem 3.10, each element of \mathcal{V}_1 crosses all but finitely many elements of \mathcal{V}_2 . By the proof of Theorem 2.4, γ bounds an eighth-flat, a contradiction. □

Lemma 3.21 (Equating basepoints) *Let $\gamma_1, \gamma_2: [0, \infty) \rightarrow \mathbf{X}$ be combinatorial geodesic rays. Then there exist combinatorial geodesic rays $\gamma'_1, \gamma'_2: [0, \infty) \rightarrow \mathbf{X}$ such that $\gamma'_1(0) = \gamma'_2(0)$ and γ'_i is almost equivalent to γ_i for $i \in \{1, 2\}$.*

Proof Since only finitely many hyperplanes separate $\gamma_1(0)$ from $\gamma_2(0)$, there exists $T \geq 0$ such that for all $t \geq T$, the hyperplane dual to $\gamma_1([t, t+1])$ does not separate $\gamma_1(0)$ from $\gamma_2(0)$. Choose a segment P joining $\gamma_2(0)$ to $\gamma_1(T)$. The concatenation $P\gamma_1([T, \infty)) = \gamma'_1$ is almost equivalent to γ_1 , and $\gamma'_1(0) = \gamma_2(0)$. Any hyperplane H crossing $\gamma([T, \infty))$ and P separates $\gamma_2(0)$ from $\gamma_1(T)$. Since H crosses γ_1 , H cannot separate $\gamma_1(0)$ from $\gamma_1(T)$ and hence separates $\gamma_1(0)$ from $\gamma_2(0)$, but this contradicts the choice of T . Thus γ'_1 has the desired properties, and we take $\gamma'_2 = \gamma_2$; see Figure 6.




Figure 6: Equating basepoints

This completes the proof. □

Lemma 3.22 (Folding Lemma) *Let $\gamma_1, \gamma_2: [0, \infty) \rightarrow \mathbf{X}$ be geodesic rays with common initial 0-cube. Suppose that some hyperplane H crosses γ_1 and γ_2 , so that H is dual to $\gamma_1([s, s+1])$ and $\gamma_2([t, t+1])$ with $0 \leq s \leq t$. Then there exist geodesic rays $\gamma'_1, \gamma'_2: [0, \infty) \rightarrow \mathbf{X}$ such that $\gamma_1([s+1, \infty)) \subset \gamma'_1, \gamma_2([s+1, \infty)) \subset \gamma'_2$, and γ'_1, γ'_2 coincide in an initial segment P whose initial 0-cube is $\gamma_1(0)$ and whose terminal 1-cube is dual to H .*

Proof If $a_1 = \gamma_1(s+1) = \gamma_2(t+1) = a_2$, then replace $\gamma_1([0, s+1])$ and $\gamma_2([0, t+1])$ by a geodesic segment P joining a_1 to $\gamma_1(0) = \gamma_2(0) = b$. If not, then let m be the median of b, a_1, a_2 . Since H separates both a_1 and a_2 from b , H must separate m from b . Since $a_1, a_2 \in N(H)$ and that subcomplex is convex, $m \in N(H)$. Let P be a geodesic path joining b to m , so that the terminal 1-cube of P is dual to H . Choose geodesic segments $P_1, P_2 \rightarrow N(H)$ joining m to a_1 and to a_2 respectively. Then replace $\gamma_1([0, s+t])$ by PP_1 and $\gamma_2([0, t+1])$ by PP_2 ; see Figure 7. □

Theorem 3.23 *Let $v \subset \mathbf{X}$ be an n -simplex of $\partial_\Delta \mathbf{X}$, each of whose 0-simplices is visible. Then there exists an isometric embedding $\mathbf{R} \rightarrow \mathbf{X}$, where \mathbf{R} is the standard tiling of $[0, \infty)^{n+1}$ by $(n+1)$ -cubes.*




Figure 7: Folding

Proof Let $v = [v_0, v_1, \dots, v_n]$ and let $\gamma_0, \gamma_1, \dots, \gamma_n$ be combinatorial geodesic rays, with γ_i representing v_i . By finitely many applications of Lemma 3.21, we can assume that the γ_i all emanate from a common 0-cube x_0 . Since $v_i \cap v_j = \emptyset$ for $i \neq j$, we have $|\mathcal{W}(\gamma_i) \cap \mathcal{W}(\gamma_j)| < \infty$ for all $i \neq j$, so that by finitely many applications of the folding lemma, followed by the removal of a finite common initial segment of the γ_i , we can assume that $\mathcal{W}(\gamma_i) \cap \mathcal{W}(\gamma_j) = \emptyset$ when $i \neq j$.

For any $i \neq j$, let H_k^i be dual to the k -th 1-cube of γ_i and let H_m^j be dual to the m -th 1-cube of γ_j . If H_k^i does not cross H_m^j , then, since v_j is a 0-simplex, there are infinitely many hyperplanes H crossing γ_j and separated from H_k^i by H_m^j . But every hyperplane crossing γ_i crosses all but finitely many of the hyperplanes crossing γ_j , and thus $H_m^j \perp H_k^i$ for all $m, k \geq 1$ and all $i \neq j$.

By induction, $\mathbf{R}' \cong \prod_{i=0}^{n-1} \gamma_i$ is an isometrically embedded subcomplex of \mathbf{X} (the base case amounts to the existence of γ_0 , which is assumed). Indeed, every hyperplane H either crosses exactly one of \mathbf{R}' or γ_n , or H crosses neither of those subcomplexes. Each hyperplane crossing \mathbf{R}' crosses each hyperplane crossing γ_n . Hence $\mathbf{R} \cong \mathbf{R}' \times \gamma_n$ isometrically embeds in \mathbf{X} . □

3.3.1 Visible pairs A notable difference between the combinatorial metric and the CAT(0) metric is that distinct 0-simplices of the same simplex of $\partial_\Delta \mathbf{X}$ are joined by a bi-infinite combinatorial geodesic in \mathbf{X} , but the corresponding points on the visual boundary (see Section 3.5) are not joined by a CAT(0) geodesic. More precisely, we have the following.

Theorem 3.24 *Let \mathbf{X} be fully visible. Then \mathbf{X} is an optical space.*

Proof Let v, w be visible simplices of $\partial_\Delta \mathbf{X}$. We shall show that either u, v is a visible pair, or $u \cap v \neq \emptyset$. In the latter case, if u, v are 0-simplices, then $u = v$, and the claim follows.

Let γ_u and γ_v be combinatorial geodesic rays representing u and v respectively. By Lemma 3.21, we can take γ_u and γ_v to have the same initial 0-cube $\gamma_u(0) = \gamma_v(0)$. Let $\mathcal{C} = \mathcal{W}(\gamma_u) \cap \mathcal{W}(\gamma_v)$. Then \mathcal{C} contains no facing triple, being the intersection of UBSs. If $C_1, C_2 \in \mathcal{C}$ are separated by H , then H crosses γ_u and γ_v , and hence belongs to \mathcal{C} . Therefore, if \mathcal{C} is infinite, then it is a UBS representing a simplex $c \subset u \cap v$, and we are done. Hence suppose that \mathcal{C} is finite. If $|\mathcal{C}| = 0$, then define $\alpha: \mathbb{R} \rightarrow \infty$ by $\alpha(t) = \gamma_u(t)$ for $t \geq 0$ and $\gamma_v(t)$ otherwise, yielding a bi-infinite geodesic showing that u, v is a visible pair. If $|\mathcal{C}| = n > 0$, then choose $C \in \mathcal{C}$. By the folding lemma, we can assume that there exists $T \geq 0$ such that $\gamma_u(t) = \gamma_v(t)$ for $t \leq T$ and that C is dual to the 1-cube $\gamma_u([T, T + 1])$. Since \mathcal{C} is finite, we can choose C so that $\mathcal{W}(\gamma_u([T + 1, \infty))) \cap \mathcal{W}(\gamma_v([T + 1, \infty))) = \emptyset$. Now simply remove the path $\gamma_u([0, T])$ from $\gamma_u \cup \gamma_v$ and define α as before. \square

Remark 3.25 If γ_v, γ_u are combinatorial geodesic rays representing simplices u, v of \mathbf{X} , and $u \subset v$, then by Lemma 3.21, we can choose γ_u and γ_v so that we have $\mathcal{W}(\gamma_u) \subset \mathcal{W}(\gamma_v)$.

3.3.2 Combinatorial geodesic completeness The notion of combinatorial geodesic completeness is needed in stating Theorem 6.10. In that context, it coincides with CAT(0) geodesic completeness (see [22, Chapter 7] for a discussion of how combinatorial geodesic completeness is used in a mostly combinatorial proof of that theorem). However, in general, the notions differ, and their relationship warrants discussion.

Definition 3.26 (Combinatorial geodesic completeness) Suppose \mathbf{X} has no infinite family of pairwise crossing hyperplanes. Then \mathbf{X} is *combinatorially geodesically complete* if the following holds for every maximal set $\{W_1, \dots, W_n\}$ of pairwise crossing hyperplanes. For each $1 \leq i \leq n$, let $\chi(W_i)$ be a halfspace associated to W_i . Then $\bigcap_{i=1}^n \chi(W_i)$ contains 0-cubes arbitrarily far from $\bigcap_{i=1}^n N(W_i)$, for any choice of χ .

Remark 3.27 (Combinatorial geodesic completeness vs. essentiality and visibility) The cube complex $\mathbb{R} \times [-\frac{1}{2}, \frac{1}{2}]$ shows that combinatorial geodesic completeness does not imply essentiality. A diagonal half-flat shows that essentiality does not imply combinatorial geodesic completeness, though the two notions are equivalent for trees. Gluing a 2-cube to $\mathbb{R} \times [-\frac{1}{2}, \frac{1}{2}]$ along any proper face yields a an optical space that is not combinatorially geodesically complete. Let \mathbf{E} be an eight-flat. At each 0-cube e of \mathbf{E} that is contained in a single 2-cube, attach a ray by identifying its initial point with e to form a cube complex that is combinatorially geodesically complete but not fully visible.

Remark 3.28 (Combinatorial vs. CAT(0) geodesic completeness) A CAT(0) space Y is *geodesically complete* if each geodesic path P extends to a bi-infinite geodesic $\gamma: \mathbb{R} \rightarrow Y$.

If \mathbf{X} is CAT(0) geodesically complete, then \mathbf{X} is combinatorially geodesically complete, since each diagonal in a maximal cube is a CAT(0) geodesic path. However, there exists a combinatorially geodesically complete cube complex \mathbf{C} that is not CAT(0) geodesically complete: let $\mathbf{C}_0 \cong [0, \infty)^2$ be a nondiagonal quarter-flat, and let $\mathbf{C}_1, \mathbf{C}_2$ be copies of $[0, \infty) \times [0, 1]$. \mathbf{C} is formed by identifying $\{0\} \times [0, 1] \subset \mathbf{C}_1$ with $\{0\} \times [0, 1] \subset \mathbf{C}_0$ and $\{0\} \times [0, 1] \subset \mathbf{C}_2$ with $[0, 1] \times \{0\}$ in the obvious cubical way, as shown in Figure 8. For $a \geq 1$, the CAT(0) geodesic line segment joining $(0, a)$ to $(a, 0)$ does not extend to a bi-infinite CAT(0) geodesic. \mathbf{C} is, however, combinatorially geodesically complete.




Figure 8: The CAT(0) geodesic segment P does not extend to a geodesic line. In any of the possible lines extending P , there exists a point (v_1 or v_2) closer to u in \mathbf{C} than in the putative geodesic. However, each 2-cube s determines four quarter-spaces, each of which contains 0-cubes arbitrarily far from s .

The choice of nomenclature comes from the following characterization of combinatorial geodesic completeness.

Lemma 3.29 *If $\mathcal{W} \neq \emptyset$, then \mathbf{X} is combinatorially geodesically complete if and only if, for each combinatorial geodesic segment or ray $P \rightarrow \mathbf{X}$, there exists a bi-infinite combinatorial geodesic $\gamma: \mathbb{R} \rightarrow \mathbf{X}$ such that $P \subset \gamma$.*

Proof Suppose that every combinatorial geodesic segment in \mathbf{X} extends to a geodesic. Let W_1, \dots, W_n be a finite set of pairwise crossing hyperplanes and let χ choose a halfspace associated to each W_i . Let A_0 be a combinatorial geodesic segment in the unique maximal cube of $\bigcap_{i=1}^n N(W_i)$, so that the initial 0-cube of A_0 lies in $\bigcap_i (\mathbf{X} - \chi(W_i))$ and the terminal 0-cube lies in $\bigcap_{i=1}^n \chi(W_i)$. By assumption, there

is a combinatorial geodesic $\gamma: \mathbb{R} \rightarrow \infty$ that contains A_0 and hence crosses each W_i . Parameterize γ so that $\gamma(0)$ is the terminal 0-cube of A_0 and $\gamma(-n)$ is the initial 0-cube of A_0 . Then $\gamma([0, \infty))$ contains 0-cubes arbitrarily far from $\gamma(0)$ and hence from the cube $\bigcap_i N(W_i)$. Moreover, $\gamma([0, \infty))$ cannot cross any W_i , since γ is a geodesic, and thus \mathbf{X} is combinatorially geodesically complete.

Now suppose that \mathbf{X} is combinatorially geodesically complete. Let $P: [0, n] \rightarrow \infty$ be a finite combinatorial geodesic path. Suppose first that $n = 0$. Let c be a maximal cube containing $P = P(0)$. Now, $\bigcap_{W \in \mathcal{W}(c)} P(0)(W)$ is a maximal finite intersection of halfspaces, and so contains 0-cubes arbitrarily far from c and hence from $P(0)$, and the same holds for $\bigcap_{W \in \mathcal{W}(c)} (\mathbf{X} - P(0)(W))$. Thus, for any $N \geq 0$, P is contained in a geodesic segment γ_N such that each of the above two intersections 0-cubes of γ_N at distance N from $P(0)$. Applying the same argument to the paths consisting of the endpoints of γ_N shows that γ_{N+1} may be chosen with $\gamma_N \subsetneq \gamma_{N+1}$, whence P extends to a bi-infinite combinatorial geodesic.

Suppose for some $n \geq 0$ that every geodesic segment of length at most n extends to a bi-infinite geodesic. Let $P: [0, n + 1] \rightarrow \mathbf{X}$ be a combinatorial geodesic segment, so that $P = P'c$, where c is the terminal 1-cube. By induction, there is a combinatorial geodesic γ containing P' . Let α be the subray of γ whose initial 0-cube is the initial 0-cube of c and which contains P' . Let χ be a (finite) maximal cube containing c and let χ be the terminal 0-cube of c . If α contains c , then we are done. Hence suppose that χ differs from every 0-cube of α on the hyperplane W dual to c . Since $\bigcap_{V \in \mathcal{W}(c)} \chi(V)$ contains 0-cubes arbitrarily far from χ , there is a geodesic ray $\beta \subset \bigcap_{V \in \mathcal{W}(c)} \chi(V)$ whose initial 0-cube is χ . Let $\gamma' = \alpha \cup c \cup \beta$. Any hyperplane U that meets γ in two distinct 1-cubes must cross α and $c \cup \beta$. Now, each element of $\mathcal{W}(c)$ separates α from β , and so U must cross each $V \in \mathcal{W}(C)$. Since $\mathcal{W}(c)$ is a maximal family of pairwise crossing hyperplanes, this is impossible. This argument also shows, mutatis mutandis, that every combinatorial geodesic ray extends to a geodesic. □

3.4 Products and joins

Theorem II.9.24 of [5] equates the existence of a spherical join decomposition of the Tits boundary of a CAT(0) space with the existence of a nontrivial product decomposition of that space. This generalizes a result of Schroeder on Hadamard manifolds (see Ballmann, Gromov and Schroeder [1, Appendix 4]). Here we prove an analogous result about simplicial join decompositions of $\partial_\Delta \mathbf{X}$. Remark 3.31 compares Theorem 3.30 to the corresponding CAT(0) result.

Given simplicial complexes A, B , we denote by $A \star B$ their (simplicial) join, ie the simplicial complex formed from $A \sqcup B$ by adding a simplex $[a_1, \dots, a_n, b_1, \dots, b_m]$ whenever $a_1, \dots, a_n \in A$ and $b_1, \dots, b_m \in B$ are pairwise distinct 0–simplices.

Theorem 3.30 *Let \mathbf{X} be an essential CAT(0) cube complex with no infinite collection of pairwise crossing hyperplanes. Consider the following two statements:*

- (1) *there exist unbounded convex subcomplexes \mathbf{X}_1 and \mathbf{X}_2 such that $\mathbf{X} \cong \mathbf{X}_1 \times \mathbf{X}_2$;*
- (2) *there exist disjoint, nonempty subcomplexes $A_1, A_2 \subset \partial_\Delta \mathbf{X}$ such that $\partial_\Delta \mathbf{X} \cong A_1 \star A_2$.*

Then (1) \Rightarrow (2). If, in addition, \mathbf{X} is fully visible, then (2) \Rightarrow (1).

Proof (1) \Rightarrow (2) Let \mathcal{W}_i be the set of hyperplanes of \mathbf{X}_i . Then $\mathcal{W} = \mathcal{W}'_1 \sqcup \mathcal{W}'_2$, where $\mathcal{W}'_1 = \{V \times \mathbf{X}_2 \mid V \in \mathcal{W}_1\}$ and, similarly, $\mathcal{W}'_2 = \{\mathbf{X}_1 \times H \mid H \in \mathcal{W}_2\}$, and for all $V \in \mathcal{W}'_1$ and $H \in \mathcal{W}'_2$, the hyperplanes V and H cross. This implies that any 0–simplex of $\partial_\Delta \mathbf{X}$ lies in the image of $\partial_\Delta \mathbf{X}_i$ for some $i \in \{1, 2\}$. From the above description of the crossings, the claimed join structure is immediate: every 0–simplex in the image of $\partial_\Delta \mathbf{X}_1$ is adjacent to every simplex in the image of $\partial_\Delta \mathbf{X}_2$. Essentiality ensures that the A_i are nonempty.

(2) \Rightarrow (1) This is proven by loose analogy with the proof of [5, Proposition II.9.25]. Fix a base 0–cube x . For each 0–simplex $u \in \partial_\Delta \mathbf{X}$, appealing to full visibility, choose a geodesic ray $\gamma: [0, \infty) \rightarrow \mathbf{X}$ such that $\gamma(0) = x$ and $\mathcal{W}(\gamma)$ represents the 0–simplex u . For $i \in \{1, 2\}$, let \mathbb{T}_i be the graph in \mathbf{X}^1 obtained by taking the union of all rays γ such that $\gamma(0) = x$ and γ represents a 0–simplex of A_i . Let $\mathbf{X}_i(x) = \mathbf{X}_i$ be the convex hull of \mathbb{T}_i .

$\mathbf{X}_i(x)$ is independent of basepoint in $\mathbf{X}_1(x)$ Let $x' \in \mathbf{X}_1(x)$ be a 0–cube, and construct the complex $\mathbf{X}_i(x')$ from rays γ' based at x' and representing 0–simplices of A_i , exactly as was done for \mathbf{X}_i . Since $\mathbf{X}_1(x)$ is convex, any hyperplane V that separates x from x' crosses $\mathbf{X}_1(x)$ and $\mathbf{X}_1(x')$. Indeed, if P is a geodesic segment joining x to x' , then $P \subset \mathbf{X}_1(x) \cap \mathbf{X}_1(x')$, since $x, x' \in \mathbf{X}_1(x) \cap \mathbf{X}_1(x')$.

If V is a hyperplane crossing $\mathbf{X}_1(x)$, then V crosses a geodesic ray γ emanating from x and representing a 0–simplex of A_1 . There is a geodesic ray γ' emanating from x' that is almost equivalent to γ . If V crosses γ and γ' , then V crosses $\mathbf{X}_1(x) \cap \mathbf{X}_1(x')$. If V separates x from x' , then V crosses $\mathbf{X}_1(x) \cap \mathbf{X}_1(x')$, as shown above. If not, then V separates γ' from an infinite subray α of γ . Hence V crosses all but finitely many of the hyperplanes that cross γ and γ' . Choose a ray γ_1 such that $\gamma_1(0) = x'$ and γ_1 is almost equivalent to γ , and γ_1 contains a 1–cube dual to V .

Hence V crosses $\mathbf{X}_1(x)$. Thus $\mathcal{W}(\mathbf{X}_1(x)) = \mathcal{W}(\mathbf{X}_1(x'))$, since $\mathbf{X}_1(x)$ is the smallest convex subcomplex containing all rays representing 0-simplices in A_1 and emanating from x' .

Verification that $\mathbf{X}_1 \cap \mathbf{X}_2 = \{x\}$ Let V be a hyperplane that crosses $\mathbf{X}_1 \cap \mathbf{X}_2$. Then $\mathbf{X}_1 - \mathbf{X}_2$ and $\mathbf{X}_2 - \mathbf{X}_1$ lie in distinct halfspaces associated to V . By essentiality, there exist minimal UBSs $\mathcal{V}_1, \mathcal{V}_2$ such that \mathcal{V}_i consists of hyperplanes that cross \mathbf{X}_i only, and V separates \mathcal{V}_1 from \mathcal{V}_2 . Let v_1, v_2 be the 0-simplices of A_1, A_2 , respectively, represented by \mathcal{V}_1 and \mathcal{V}_2 . Since \mathcal{V}_1 and \mathcal{V}_2 are separated by V , we have $d_{\partial_{\Delta}\mathbf{X}}(v_1, v_2) > 1$. Indeed, if v_1, v_2 are adjacent, then they are represented by UBS $\mathcal{V}'_1 \sqcup \mathcal{V}'_2$ such that each $V' \in \mathcal{V}'_1$ crosses all but finitely many $V'' \in \mathcal{V}'_2$, and hence crosses V , a contradiction. On the other hand, since $v_1 \in A_1$ and $v_2 \in A_2$ and $\partial_{\Delta}\mathbf{X} \cong A_1 \star A_2$, the simplices v_1 and v_2 are adjacent, ie $d_{\partial_{\Delta}\mathbf{X}}(v_1, v_2) = 1$, a contradiction. Hence no hyperplane crosses $\mathbf{X}_1 \cap \mathbf{X}_2$, which therefore consists entirely of the 0-cube x .

Defining orthogonal projections Define *orthogonal projections* $\rho_i: \mathbf{X}^{(0)} \rightarrow \mathbf{X}_i^{(0)}$ by letting $\rho_i(y)$ be the gate of y in \mathbf{X}_i . In particular, ρ_i is the identity on the 0-skeleton of \mathbf{X}_i .

Note that if $p, q \in \mathbf{X}$ are adjacent 0-cubes, then one of the following situations occurs. If $p, q \in \mathbf{X}_i$, then $\rho_i(p) = p, \rho_i(q) = q$ and in particular the projections are adjacent. If $p \in \mathbf{X}_i$ and $q \notin \mathbf{X}_i$, then $\rho_i(p) = \rho_i(q) = q$ since q is at distance at least 1 from \mathbf{X}_i and at distance exactly 1 from p . If $p, q \notin \mathbf{X}_i$, then either p is closer to \mathbf{X}_i than is q , in which case $\rho_i(p) = \rho_i(q)$, or they are at equal distance from \mathbf{X}_i , in which case p, q are separated by a single hyperplane V that does not separate either from \mathbf{X}_i , whence V crosses \mathbf{X}_i and is thus the unique hyperplane separating $\rho_i(p)$ from $\rho_i(q)$. Hence we obtain a map $\rho_i: \mathbf{X}^1 \rightarrow \mathbf{X}_i^1$ that sends 0-cubes to 0-cubes and sends each 1-cube c isometrically to a 1-cube $\rho_i(c)$ or collapses c to a 0-cube.

Verification that $\mathbf{X}_2(x) = \rho_1^{-1}(x)$ Let $p \in \mathbf{X}_2$ and let m be the unique median of x, p and $\rho_1(p)$. Then m lies on a geodesic joining x to $\rho_1(p)$, so that, by convexity of \mathbf{X}_1 , we have $m \in \mathbf{X}_1$. On the other hand, m lies on a geodesic joining x to p so that, by convexity of \mathbf{X}_2 , we have $m \in \mathbf{X}_2$. Therefore, $m \in \mathbf{X}_1 \cap \mathbf{X}_2$, whence $m = x$. Thus $x \in \mathbf{X}_1$ lies on a geodesic joining p to $\rho_1(p)$. Since $\rho_1(p)$ is the closest 0-cube of \mathbf{X}_1 to p , we have $\rho_1(p) = x$, so that $\mathbf{X}_2 \subseteq \rho_1^{-1}(x)$. On the other hand, since ρ_1 is the identity on \mathbf{X}_1 , if $p \in \rho_1^{-1}(x)$, then $p = x$ or $p \in \mathbf{X} - \mathbf{X}_1$. In the latter case, let V be a hyperplane separating p from \mathbf{X}_1 . Then by essentiality of \mathbf{X} , there exists a ray γ , containing a 1-cube dual to V , such that $\gamma(0) = x, \gamma(t) = p$ for some $t > 0$, and γ represents a simplex $u \in A_2$, from which it follows that $x \in \mathbf{X}_2$.

Conclusion Define a map $j: \mathbf{X}^{(0)} \rightarrow \mathbf{X}_1^{(0)} \times \mathbf{X}_2^{(0)}$ by $j(p) = (\rho_1(p), \rho_2(p))$, where $p \in \mathbf{X}$ is a 0-cube. For all 0-cubes $p, q \in \mathbf{X}$, we have

$$d_{\mathbf{X}_1 \times \mathbf{X}_2} = d_{\mathbf{X}_1}(\rho_1(p), \rho_1(q)) + d_{\mathbf{X}_2}(\rho_2(p), \rho_2(q));$$

the first term on the right counts hyperplanes in \mathbf{X} that cross \mathbf{X}_1 and separate p, q and the second term counts hyperplanes that cross \mathbf{X}_2 and separate p, q . Since each hyperplane in \mathbf{X} crosses exactly one of the \mathbf{X}_i , we see that j is an isometric embedding on 0-skeleta and thus extends to an isometric embedding $j: \mathbf{X}^1 \rightarrow \mathbf{X}_1^1 \times \mathbf{X}_2^1$. Since a CAT(0) cube complex is determined by its 1-skeleton, it suffices to show that j is surjective on 0-cubes.

Let $(p, q) \in \mathbf{X}_1 \times \mathbf{X}_2$ be a 0-cube. Let $\mathcal{V}_1(p)$ be the set of hyperplanes in \mathbf{X} separating p from x and let $\mathcal{V}_2(q)$ be the set of hyperplanes separating q from x . Suppose that some $V \in \mathcal{V}_2(q)$ fails to cross some $H \in \mathcal{V}_1(p)$. By essentiality, there are UBSs $\mathcal{V} \subset \mathcal{V}_2$ and $\mathcal{H} \subset \mathcal{V}_1$ that are separated by V and H and thus represent simplices $v \in A_2, h \in A_1$ that are nonadjacent in $\partial_\Delta \mathbf{X}$, a contradiction. Thus each hyperplane in $\mathcal{V}_1(p)$ crosses each hyperplane in $\mathcal{V}_2(q)$, so there is a 0-cube p' such that the set of hyperplanes separating p' from x is precisely $\mathcal{V}_1(p) \cup \mathcal{V}_2(q)$. But then $\rho_1(p') = p$ and $\rho_2(p') = q$, whence j is surjective. \square

Remark 3.31 To deduce [Theorem 3.30](#) from [[5](#), Theorem II.9.24], one would need to establish that each geodesic segment in $(\mathbf{X}, \partial \mathbf{X})$ extends to a geodesic line when \mathbf{X} is essential and fully visible and $\partial_\Delta \mathbf{X}$ has a nontrivial join decomposition. The example shown in [Figure 8](#) shows that the last fact would play an essential role in such a proof. More seriously, one would need the CAT(1) realization of $\partial_\Delta \mathbf{X}$ to be isometric to the Tits boundary (see [Section 3.5](#)), but there are essential, fully visible product cube complexes for which this is not the case.

3.5 Comparison to other boundaries

Throughout this section, $\partial_\infty \mathbf{X}$ denotes the visual boundary of $(\mathbf{X}, \partial \mathbf{X})$, endowed with the cone topology. The Tits boundary $\partial_T \mathbf{X}$ is obtained by endowing the visual boundary with the Tits metric. In this section, we discuss the relationship between $\partial_\Delta \mathbf{X}$, the visual boundary, the Tits boundary, and the *Roller boundary* which \mathbf{X} possesses as a CAT(0) cube complex.

The construction of $\partial_\Delta \mathbf{X}$ somewhat resembles that of the Roller boundary $\mathfrak{R}\mathbf{X}$ of \mathbf{X} (of which a compact description appears in Nevo and Sageev [[31](#), Section 2.2]). However, $\mathfrak{R}\mathbf{X}$ is necessarily compact, so that by [Theorem 3.14](#), even a tree suffices to distinguish the simplicial boundary from the Roller boundary. Likewise, $\partial_\Delta \mathbf{X}$ differs in general from $\partial_\infty \mathbf{X}$.

Topologically, $\partial_\infty \mathbf{X}$ is sensitive to perturbations of the CAT(0) metric, as shown in [13]: Croke and Kleiner demonstrated that the visual boundary of the CAT(0) cube complex associated to the right-angled Artin group $\langle a_1, a_2, a_3, a_4 \mid [a_i, a_{i+1}], 1 \leq i \leq 3 \rangle$ varies in response to changes in the angles at the corners of the 2-cubes. By contrast, $\partial_\Delta \mathbf{X}$ depends only on $\mathbf{X}^{(1)}$ – it does not “notice” the CAT(0) metric – and is thus invariant under such perturbations.

3.5.1 Simplices and the visual boundary Let $\sigma: [0, \infty) \rightarrow \mathbf{X}$ be a CAT(0) geodesic ray, and let $\mathcal{W}(\sigma)$ be the set of hyperplanes H with $H \cap \sigma = \emptyset$. Then $\mathcal{W}(\sigma)$ is a UBS. If the geodesic ray σ' lies at uniform Hausdorff distance from σ , then clearly $|\mathcal{W}(\sigma) \cap \mathcal{W}(\sigma')| < \infty$, so that the simplices of $\partial_\Delta \mathbf{X}$ represented by $\mathcal{W}(\sigma)$ and $\mathcal{W}(\sigma')$ coincide. Denote this simplex by $w(\sigma)$. Let $[\sigma]$ be the point of $\partial_\infty \mathbf{X}$ represented by σ and let $\mathfrak{S}\mathbf{X}$ be the set of simplices of $\partial_\Delta \mathbf{X}$. The earlier observation shows that the assignment $[\sigma] \mapsto w(\sigma)$ gives a well-defined map $\mathfrak{s}: \partial_\infty \mathbf{X} \rightarrow \mathfrak{S}\mathbf{X}$. The following simple lemma is crucial here and in the next section.

Lemma 3.32 (Representing 0-simplices with CAT(0) geodesic rays) *Let γ be a combinatorial geodesic ray such that $\mathcal{W}(\gamma)$ is minimal. Then there exists a combinatorial geodesic ray γ' such that $|\mathcal{W}(\gamma) \Delta \mathcal{W}(\gamma')| < \infty$, and a CAT(0) geodesic ray σ such that γ' is contained in the union of cubes c such that $c \cap \sigma \neq \emptyset$.*

Proof of Lemma 3.32 Let \mathbf{C} be the convex hull of γ , so that $\mathcal{W}(\mathbf{C}) = \mathcal{W}(\gamma)$. Since \mathbf{C} is CAT(0)-convex, it suffices to prove the claim for \mathbf{C} . Now, $\mathcal{W}(\gamma)$ is the inseparable closure of a set $\{H_i\}_{i=1}^\infty$ of pairwise noncrossing hyperplanes such that H_{i-1} and H_{i+1} are separated by H_i for all $i \geq 2$. Letting $x_o = \gamma(0)$, we see that for each $i \geq 1$, there exists a 0-cube x_i in \mathbf{C} that is separated from x_o by H_j for $j \leq i$, but not by H_{i+1} .

Now, x_o has finite degree. Indeed, each 1-cube of \mathbf{C} emanating from x_o is dual to a unique hyperplane, and these hyperplanes pairwise cross, since $\mathcal{W}(\mathbf{C}) = \mathcal{W}(\gamma)$. Since \mathbf{X} contains no infinite family of pairwise crossing hyperplanes, there are finitely many such 1-cubes. By induction, for each $i \geq 1$, there are finitely many 0-cubes x_i of the type described above. Hence there are finitely many CAT(0) geodesic segments σ_i that emanate from x_o and cross H_i but not H_{i+1} . By König’s lemma, there is a CAT(0) geodesic σ emanating from x_o that crosses each H_i . Since $\mathcal{W}(\gamma)$ is the inseparable closure of $\{H_i\}_i$, we have $\mathcal{W}(\sigma) = \mathcal{W}(\gamma)$. Let \mathbf{C}' be the union of all closed maximal cubes whose interiors have nonempty intersection with σ . Since each hyperplane crossing \mathbf{C}' crosses σ , the intersection of \mathbf{C}' with each hyperplane-carrier consists of a single maximal cube and is thus connected. Hence the inclusion $\mathbf{C}' \hookrightarrow \mathbf{C}$ is a (combinatorial) isometric embedding. We thus choose γ' to be any combinatorial geodesic ray in \mathbf{C}' emanating from x_o . □

Corollary 3.33 *If \mathbf{X} is fully visible, then \mathfrak{s} is surjective.*

Proof Let v be an n -simplex of $\partial_\Delta \mathbf{X}$. If $n = 0$, then by Lemma 3.32, there exists a CAT(0) geodesic ray σ such that $\mathcal{W}(\sigma)$ represents v , so that $\mathfrak{s}([\sigma]) = v$. Suppose $n \geq 1$, so that v is spanned by a set v_0, \dots, v_n of 0-simplices, each of which is visible. By Theorem 3.23, there exist combinatorial geodesic rays $\gamma_0, \dots, \gamma_n$ with common initial 0-cube such that for all $i \neq j$ and all $H \in \mathcal{W}(\gamma_i), H' \in (\gamma_j)$, the hyperplanes H and H' cross. Arguing as in Lemma 3.32 yields a CAT(0) geodesic ray σ with $\mathcal{W}(\sigma) = \bigcup_i \mathcal{W}(\gamma_i)$, whence $\mathfrak{s}([\sigma]) = v$. \square

Topologize $\mathfrak{S}\mathbf{X}$ by declaring the set $\mathfrak{W} \subset \mathfrak{S}\mathbf{X}$ to be open if $\mathfrak{s}^{-1}(\mathfrak{W}) \subset \partial_\infty \mathbf{X}$ is open in the cone topology. Hence $\mathfrak{S}\mathbf{X}$, with the given quotient topology, is compact.

Corollary 3.34 *If $\partial_\Delta \mathbf{X}$ is 0-dimensional, then $\partial_\Delta \mathbf{X}$ has a compact topology as a quotient of $\partial_\infty \mathbf{X}$. If $\mathbf{X}^{(1)}$ is hyperbolic, then the quotient $\partial_\infty \mathbf{X} \rightarrow \partial_\Delta \mathbf{X}$ is a homeomorphism (where $\partial_\Delta \mathbf{X}$ is given this compact topology).*

Proof Since each simplex of $\partial_\Delta \mathbf{X}$ is maximal, \mathbf{X} is fully visible by Theorem 3.19, so that the above quotient topology on $\mathfrak{S}\mathbf{X}$ exists. The identity $\mathfrak{S}\mathbf{X} \rightarrow \partial_\Delta \mathbf{X}$ endows $\partial_\Delta \mathbf{X}$ with the same topology. If \mathbf{X} is hyperbolic and the CAT(0) geodesic rays γ, γ' satisfy $w(\gamma) = w(\gamma') \in \partial_\Delta \mathbf{X}^{(0)}$, it is not hard to see that $[\gamma] = [\gamma']$, so that \mathfrak{s} is injective. (See Example 3.35 for a situation in which the quotient is not injective.) \square

3.5.2 The CAT(1) realization of $\partial_\Delta \mathbf{X}$ and $\partial_T \mathbf{X}$ For $n \geq 0$, let ζ_n be a right-angled spherical n -simplex whose 1-simplices have length $\pi/2$. Each n -simplex of $\partial_\Delta \mathbf{X}$ is metrized so as to be isometric to ζ_n . A *piecewise geodesic path* P in $\partial_\Delta \mathbf{X}$ is a path for which $P \cap v$ is a geodesic of v , for each simplex v . The length of P is $|P| = \sum_{v \in \mathfrak{S}\mathbf{X}} |P \cap v|_v$, where $|Q|_v$ denotes the length of the path Q in v (and $|\emptyset|_v = 0$). Given $x, y \in \partial_\Delta \mathbf{X}$, let $\ddot{d}(x, y) = \inf_P |P|$, where P varies over paths joining a to b . The simplicial boundary endowed with the metric $\ddot{d}: \partial_\Delta \mathbf{X} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is the *CAT(1) realization* of the simplicial boundary.

According to Theorem 3.14, $\partial_\Delta \mathbf{X}$ is a finite-dimensional flag complex, so the CAT(1) realization of $\partial_\Delta \mathbf{X}$ is indeed a CAT(1) space. The same is true of $\partial_T \mathbf{X}$ and it is natural to ask when $\partial_\Delta \mathbf{X}$ is actually an isometric triangulation of the Tits boundary.

Example 3.35 The following examples suggest necessary conditions for the CAT(1) realization of the simplicial boundary to be isometric to the Tits boundary.

- (1) Let \mathbf{X} be a diagonal half-flat. In this situation, $\partial_{\Delta}\mathbf{X}$ consists of a closed interval divided into three 1–simplices, so that $\text{diam}(\partial_{\Delta}\mathbf{X}, \ddot{d}) = \frac{3\pi}{2}$. However, $\text{diam}(\partial_T\mathbf{X})$ takes some value in $(\frac{\pi}{2}, \frac{3\pi}{2})$, depending on the constituent eighth-flats of \mathbf{X} . This shows that the existence of a proper, essential group action is insufficient to make $\partial_{\Delta}\mathbf{X} \cong \partial_T\mathbf{X}$.
- (2) The cube complex \mathbf{X} illustrated in Figure 9 has $\partial_{\Delta}\mathbf{X}$ a single 0–simplex since the set of all hyperplanes is a minimal UBS¹. The arrowed CAT(0) geodesic rays, however, determine points at positive distance in $\partial_T\mathbf{X}$.




Figure 9: The hyperplanes in this cube complex are compact, and the set of all hyperplanes is a minimal UBS. The single-headed combinatorial geodesic rays are hyperplane-equivalent, for example. The double-headed CAT(0) geodesic rays represent different points on the Tits boundary, however.

By Theorem 3.10, if the combinatorial geodesic ray γ represents a positive-dimensional simplex of $\partial_{\Delta}\mathbf{X}$, then $\mathcal{W}(\gamma)$ cannot have thin bicliques. Example 3.35(2) shows that it is possible for two almost equivalent minimal combinatorial geodesic rays to diverge, by virtue of the existence of arbitrarily large *finite* bicliques in the subgraph of $\Delta\mathbf{X}$ generated by the corresponding UBS.

Definition 3.36 (Bounded bending) The 0–simplex v of $\partial_{\Delta}\mathbf{X}$ has *bounded bending* if for all combinatorial geodesic rays γ, γ' such that $\mathcal{W}(\gamma)$ and $\mathcal{W}(\gamma')$ represent v , the rays γ, γ' lie at finite Hausdorff distance in $\mathbf{X}^{(1)}$.

In general, even in the presence of a proper, cocompact group action, there are 0–simplices of $\partial_{\Delta}\mathbf{X}$ that do not have bounded bending; in private communication, Mike Carr brought to my attention examples of such 0–simplices when \mathbf{X} is the universal cover of the Salvetti complex of the Croke–Kleiner group $\langle a, b, c, d \mid [a, b], [b, c], [c, d] \rangle$. The failure of bounded bending is the obstruction to the existence of an isometry between the CAT(1) realization of \mathbf{X} and the Tits boundary. However, we have the following.

¹The operative property of UBSs in this example is inseparability.

Proposition 3.37 *Let \mathbf{X} be a fully visible CAT(0) cube complex for which $\partial_\Delta \mathbf{X}$ is defined. Let $(\partial_\Delta \mathbf{X}, \ddot{d})$ be the CAT(1) realization of $\partial_\Delta \mathbf{X}$. Then there is an isometric embedding $I: \partial_\Delta \mathbf{X} \rightarrow \partial_T \mathbf{X}$ that is a section of a surjective map $R: \partial_T \mathbf{X} \rightarrow \partial_\Delta \mathbf{X}$ such that*

- (1) *for each $z \in \partial_\Delta \mathbf{X}$, the preimage $R^{-1}(z)$ has Tits diameter less than $\frac{\pi}{2}$;*
- (2) *I is $\frac{\pi}{2}$ -quasisurjective, and has quasi-inverse R ;*
- (3) *I is an isometry if and only if every 0-simplex of $\partial_\Delta \mathbf{X}$ has bounded bending;*
- (4) *if $\gamma: \mathbb{R} \rightarrow \mathbf{X}$ is an axis of an isometry $g \in \text{Aut}(\mathbf{X})$, with endpoints $p_\pm \in \partial_T \mathbf{X}$, and $\mathcal{H}(\gamma)$ represents a pair of 0-simplices, then $p_\pm \in \text{im}(I)$.*

Proof The plan of the proof is as follows. We first define an embedding I of the 0-skeleton of $\partial_\Delta \mathbf{X}$ in $\partial_T \mathbf{X}$. In the absence of bounded bending, this involves some arbitrary choices. We next show that for any admissible choice of I , we can extend I in a unique way to a map $\partial_\Delta \mathbf{X} \rightarrow \partial_T \mathbf{X}$ that is an isometric embedding on each simplex. We then verify that, if I, I' are maps constructed in this way, $\text{im}(I)$ is contained in the open $\frac{\pi}{2}$ -neighborhood of $\text{im}(I')$ and vice versa. From this, it follows that I is quasisurjective, and we also argue using this fact to show that I is an isometric embedding. At this point, it is easy to construct R and verify the remaining claims.

We use the following notation: if γ is a CAT(0) geodesic ray, then $[\gamma]$ is the corresponding point on the Tits boundary, $\mathcal{W}(\gamma)$ is the UBS consisting of hyperplanes crossing γ , and v_γ is the simplex of $\partial_\Delta \mathbf{X}$ represented by $\mathcal{W}(\gamma)$.

Defining I on the 0-skeleton Let $v \in \partial_\Delta \mathbf{X}$ be a 0-simplex. By full visibility of \mathbf{X} , and Lemma 3.32, there exists a CAT(0) geodesic ray γ with $v_\gamma = v$. Let $I(v) = [\gamma]$. Clearly, if $I(v) = I(v')$ for some 0-simplex v' , then since rays that are asymptotic are hyperplane-equivalent up to a finite set of hyperplanes, $v = v'$, ie $I: (\partial_\Delta \mathbf{X})^{(0)} \rightarrow \partial_T \mathbf{X}$ is injective.

Now, if v does not have bounded bending, then there is an arbitrary choice involved in defining I , since there may be a ray γ' such that $[\gamma'] \neq [\gamma]$ but $v_{\gamma'} = v$. However, in such a case, $\partial_T([\gamma], [\gamma']) < \frac{\pi}{2}$. Indeed, suppose to the contrary that $\partial_T([\gamma], [\gamma']) \geq \frac{\pi}{2}$. Without loss of generality, $\gamma(0) = \gamma'(0)$. If $\partial_T([\gamma], [\gamma']) \geq \pi$, then the concatenation of γ and γ' fellow-travels with a bi-infinite geodesic, so that $v_\gamma \neq v_{\gamma'}$, a contradiction. Otherwise, examining triangles determined by $\gamma(t)$ and $\gamma'(t)$ shows that $|\mathcal{W}(\gamma) \Delta \mathcal{W}(\gamma')|$ is infinite if γ and γ' meet at an angle at least $\frac{\pi}{2}$. Since this argument does not use minimality of $\mathcal{W}(\gamma)$, the same conclusion holds even if v_γ and $v_{\gamma'}$ are not necessarily 0-simplices: either $\partial_T([\gamma], [\gamma']) < \frac{\pi}{2}$ or $v_\gamma \neq v_{\gamma'}$. Conversely, it is easily checked that if the 0-simplex v has bounded bending, then any two rays representing v are asymptotic.

Extending I to $\partial_\Delta \mathbf{X}$ Let w be a d -simplex of $\partial_\Delta \mathbf{X}$, spanned by the 0-simplices v_0, \dots, v_d . For $0 \leq i \leq d$, let σ_i be the CAT(0) geodesic ray chosen above, representing $I(v_i)$. By induction on d , there exists $x \in \mathbf{X}$ and CAT(0) geodesic rays $\sigma'_0, \dots, \sigma'_{d-1}$ such that each σ'_i fellow-travels with σ_i , each σ'_i emanates from x , and \mathbf{X} contains an isometrically embedded subspace $E \cong \prod_{i=0}^{d-1} \sigma'_i$. In the base case, when $d = 0$, this is clear, since we can take $E = \sigma_0$.

Now, by removing a finite initial subpath from each σ'_i , and from σ_d , if necessary, we can assume that if H is a hyperplane crossing E , and V is a hyperplane crossing σ_d , then V and H cross (see the proof of [Theorem 3.23](#)). Therefore, for each hyperplane H crossing E , any hyperplane W that separates a point of σ_d from H must separate all of σ_d from H . Hence σ_d fellow-travels with a geodesic ray σ'_d that emanates from x and lies in a hyperplane crossing E . Arguing as in the proof of [Theorem 3.23](#), we find that the Euclidean orthant $\sigma_d \times E$ embeds isometrically in \mathbf{X} .

Hence \mathbf{X} contains a $(d + 1)$ -dimensional flat orthant $F \cong \prod_{i=0}^d \sigma'_i$, where each σ'_i is a CAT(0) geodesic ray representing $I(v_i)$ in the Tits boundary. Each point $z \in w$ is determined by barycentric coordinates $(\alpha_j(z))_{j=0}^d$, where $\alpha_j(z) \in [0, 1]$ and $\sum_{i=0}^d \alpha_j^2(z) = 1$. Let γ_z be the geodesic ray in F emanating from the origin x in the direction of the vector $(\alpha_i(z))_{i=0}^d$, and let $I(z)$ be the point of $\partial_T \mathbf{X}$ represented by γ_z (which is a geodesic of \mathbf{X} since F is isometrically embedded). This gives an isometric embedding $I: w \rightarrow \partial_T \mathbf{X}$ for each right-angled spherical simplex w of $\partial_\Delta \mathbf{X}$. Note that $I(w)$ is completely determined by $\{I(v_i)\}_{i=0}^d$.

(What we have actually proved is that the cubical convex hull of $\sigma'_0, \dots, \sigma'_d$ contains an orthant isometric to the product of those rays.)

I is quas surjective Let Λ be the set of all maps $\lambda: (\partial_\Delta \mathbf{X})^{(0)} \rightarrow \partial_T \mathbf{X}$ defined as in the first part of the proof. Note that for each $\lambda \in \Lambda$, there exists a unique isometric embedding $I_\lambda: \partial_\Delta \mathbf{X} \rightarrow \partial_T \mathbf{X}$ that extends the map λ and is constructed in the above manner.

Now, let $[\gamma] \in \partial_T \mathbf{X}$. If v_γ is a 0-simplex, we can choose $\lambda \in \Lambda$ such that $\lambda(v_\gamma) = [\gamma]$, and then extend to I_λ as above. Hence $[\gamma] \in \text{im}(I_\lambda)$ for some λ .

More generally, let $[\gamma] \in \partial_T \mathbf{X}$ and let v_γ be a d -simplex spanned by 0-simplices v_0, \dots, v_d . For some choice of λ , the ray γ is asymptotic to a ray in the orthant $\prod_i \lambda(v_i)$, and thus we can choose λ so that $\text{im}(I_\lambda)$.

We have shown that for each $[\gamma] \in \partial_T \mathbf{X}$, there exists $\lambda \in \Lambda$ such that $[\gamma] \in \text{im}(I_\lambda)$. On the other hand, let $z \in \partial_\Delta \mathbf{X}$ and let $\lambda, \lambda' \in \Lambda$. Then $I_\lambda(z)$ and $I_{\lambda'}(z)$ are represented by rays that are hyperplane-equivalent up to a finite set of hyperplanes; we have seen that this implies that $d_T(I_\lambda(x), I_{\lambda'}(x)) < \frac{\pi}{2}$. Combining these facts, fixing $\lambda \in \Lambda$

and letting $I = I_\lambda$, we see that for any $p \in \partial_T \mathbf{X}$, there exists $z \in \partial_\Delta \mathbf{X}$ such that $d_T(p, I(z)) < \frac{\pi}{2}$. Thus I is $\frac{\pi}{2}$ -quasisurjective.

I is an isometric embedding Let \mathfrak{M} be the set of all maximal simplices in $\partial_\Delta \mathbf{X}$, and for each $v \in \mathfrak{M}$, let \hat{v} be the open $\frac{\pi}{2}$ -neighbourhood of $I(v)$ in $\partial_T \mathbf{X}$, so that $\bigcup_{v \in \mathfrak{M}} \hat{v} = \partial_T \mathbf{X}$. Obviously, if $u \cap v \neq \emptyset$, then $\hat{u} \cap \hat{v} \neq \emptyset$. Conversely, using full visibility, let $u = v_{\gamma_u}$ and $v = v_{\gamma_v}$ for rays γ, γ' . Let $\hat{u} \cap \hat{v}$ contain a point $[\gamma]$. Then $\mathcal{W}(\gamma_u) \cap \mathcal{W}(\gamma)$ is infinite, for otherwise $[\gamma_u], [\gamma]$ are either the endpoints of a bi-infinite combinatorial geodesic, which is impossible if $d_T([\gamma_u], [\gamma]) < \frac{\pi}{2}$. Similarly, $\mathcal{W}(\gamma_v) \cap \mathcal{W}(\gamma)$ is infinite. Hence either $u \cap v \neq \emptyset$, or $u \cap v_\gamma$ and $v \cap v_\gamma$ are nonempty (or both). If $u \cap v = \emptyset$, then $\mathcal{W}(\gamma_v) \cap \mathcal{W}(\gamma)$ and $\mathcal{W}(\gamma_u) \cap \mathcal{W}(\gamma)$ are disjoint (up to finite subsets) UBSs in the UBS $\mathcal{W}(\gamma)$, whence u and v are contained in a common simplex. In short: if $\hat{u} \cap \hat{v} \neq \emptyset$, then either $u \cap v \neq \emptyset$, or u and v are contained in a common simplex.

Now, let $z, z' \in \partial_\Delta \mathbf{X}$. The above argument shows that $d_T(I(z), I(z')) < \infty$ if and only if $\check{d}(z, z') < \infty$, as follows. If $\check{d}(z, z')$ is infinite, ie if z, z' lie in distinct path-components of $\partial_\Delta \mathbf{X}$, then $I(z), I(z')$ cannot be joined by a sequence of subsets of the form \hat{v} , by the preceding argument, and thus $d_T(I(z), I(z')) = \infty$; the converse is almost identical. Hence it suffices to show that I is an isometric embedding on each path-component of $\partial_\Delta \mathbf{X}$, ie we can assume that $\check{d}(z, z') < \infty$.

For any $\epsilon > 0$, we can choose a path Q in $\partial_T \mathbf{X}$ such that $d_T(I(z), I(z')) < |Q| + \epsilon$. The path Q is a concatenation of paths joining points at finite Tits distance, and $|Q|$ is the sum of the lengths of these subpaths.

Since $\{\hat{v} \mid v \in \mathfrak{M}\}$ covers $\partial_T \mathbf{X}$, we can choose a shortest sequence $\hat{v}_0, \dots, \hat{v}_n$ such that $I(z) \in I(v_0), I(z') \in I(v_n)$, and $\hat{v}_i \cap \hat{v}_{i+1} \neq \emptyset$ for $0 \leq i \leq n-1$, and $Q \subset \cup_i \hat{v}_i$. If $n = 0$, z, z' lie in the same simplex, whence Q can be chosen in $\text{im}(I)$ since I is an isometric embedding on simplices.

For $n \geq 0$, by induction we can write $Q = Q'Q''$, where $Q' \subset \text{im}(I)$ and Q'' joins the terminal point z'' of Q' to z' , with $z'' \in I(v_{n-1})$. Now, if $v_{n-1} \cap v_n = \emptyset$, then there exists a simplex w containing v_{n-1} and v_n , contradicting minimality of n . Thus $I(v_n) \cap I(v_{n-1}) \neq \emptyset$, by the dichotomy established above. We shall show momentarily that, if u, v are simplices with nonempty intersection, then $I|_{u \cup v}$ is an isometric embedding. It follows that we can choose Q'' in the image of I , and thus $I(z)$ and $I(z')$ are joined by a path in $\text{im}(I)$ of length less than $d_T(I(z), I(z')) + \epsilon$. This completes the proof that I is an isometric embedding.

It remains to establish the supporting claim about $I_{u \cup v}$. Suppose that $a \in u, b \in v$ are interior points, with $I(a) = [\gamma_u]$ and $I(b) = [\gamma_v]$, where γ_u, γ_v are geodesic rays

emanating from a common point x . Since $u \cap v \neq \emptyset$, infinitely many hyperplanes cross both γ_u and γ_v , and since each of $I(u)$ and $I(v)$ is isometrically embedded in $\partial_T \mathbf{X}$, we have that

$$\ddot{d}(a, b) \leq \inf_{c \in u \cap v} [\ddot{d}(a, c) + \ddot{d}(c, b)] = \inf_{c \in u \cap v} [d_T(I(a), I(c)) + d_T(I(c), I(b))].$$

Now, if P is a path in $\partial_T \mathbf{X}$ from $I(a)$ to $I(b)$, then P contains a point of the form $I_\lambda(c)$ where $c \in u \cap v$ and $\lambda \in \Lambda$. Let σ be a ray emanating from x with $I(c) = [\sigma]$ and let ξ be a ray emanating from x with $I_\lambda(c) = [\xi]$. Let $\mathcal{U} = \mathcal{W}(\sigma) = \mathcal{W}(\xi)$ be the infinite subset of $\mathcal{W}(\gamma_a)$ and $\mathcal{W}(\gamma_b)$ representing the simplex $u \cap v$. For $H, H' \in \mathcal{U}$, write $H <_\sigma H'$ if a traveler leaving x along σ must cross H before H' , and define $<_{\gamma_a}, <_{\gamma_b}, <_\xi$ analogously. Each of these is clearly a partial ordering of \mathcal{U} . Now, by construction, γ_a and σ lie uniformly close to a common orthant, one of whose proper factors is an orthant containing σ , so that for all but finitely many $H, H' \in \mathcal{U}$, we have $H <_{\gamma_a} H'$ if and only if $H <_\sigma H'$. The same holds with γ_a replaced by γ_b . On the other hand, since σ and ξ are hyperplane-equivalent, either $[\sigma] = [\xi]$ or there are infinitely many $H \in \mathcal{U}$ and infinitely many $H' \in \mathcal{U}$ such that $H <_\sigma H'$ but $H' <_\xi H$. (This is illustrated in Figure 9.) In the latter case, this implies that $\mathfrak{d}(\gamma_a(t), \xi(t)) - \mathfrak{d}(\gamma_a(t), \sigma(t))$ is increasing, and the same is true with γ_a replaced by γ_b . Thus $|P| \geq d_T(I(a), I(c)) + d_T(I(c), I(b))$, and we are done.

The map R Given $[\gamma] \in \partial_T \mathbf{X}$, there exists a closest $z \in v_\gamma$ such that $d_T([\gamma], I(z)) < \frac{\pi}{2}$ and $I(z)$ is represented by a ray hyperplane-equivalent to γ . Let $R([\gamma]) = z$. For all $z \in \partial_\Delta \mathbf{X}$, we have $R \circ I(z) = z$. For all $[\gamma]$, $I \circ R([\gamma])$ lies in $I(v_\gamma)$, whence R is a quasi-inverse for I .

The bounded bending case If every 0-simplex has bounded bending, then there is a unique choice of $I(v)$ for each 0-simplex v . Since I is uniquely determined by its restriction to the 0-skeleton, I is uniquely determined, ie $|\Lambda| = 1$. Since each point in $\partial_T \mathbf{X}$ lies in the image of I_λ for some $\lambda \in \Lambda$, the map I is surjective, ie an isometry.

Conversely, if the 0-simplex v does not have bounded bending, then there exist hyperplane-equivalent geodesics γ, γ' that do not fellow-travel but whose sets of dual hyperplanes represent v . By construction, at most one of the corresponding points on the Tits boundary lies in $\text{im}(I)$, which is therefore not surjective.

Axes of isometries Let $g \in \text{Aut}(\mathbf{X})$ be a hyperbolic isometry and let γ be an axis for g , so that γ determines a pair v_-, v_+ of $\langle g \rangle$ -invariant 0-simplices of $\partial_\Delta \mathbf{X}$. Since γ is periodic and represents a pair of 0-simplices, the convex hull of γ is contained in a uniform neighborhood of γ , whence every ray that is hyperplane-equivalent to a subray of γ must fellow-travel with γ . Hence $I(v_\pm)$ are uniquely determined; they are the endpoints p_\pm of γ . □

A useful corollary of Proposition 3.37 is that I always induces a bijection from the set of components of $\partial_\Delta \mathbf{X}$ to the set of components of $\partial_T \mathbf{X}$. Also, if \mathbf{X} is hyperbolic, then each 0-simplex has bounded bending because $\Gamma \mathbf{X}$ has thin bicliques, whence $\partial_\Delta \mathbf{X}$ and $\partial_T \mathbf{X}$ coincide – both are discrete sets of equal cardinality.

It seems that the case in which every 0-simplex has bounded bending is in some sense rare, and we close with the following.

Question 3.38 Let G act properly, cocompactly and essentially on the CAT(0) cube complex \mathbf{X} . Under what additional conditions does every 0-simplex in $\partial_\Delta \mathbf{X}$ have bounded bending?

4 Bounded contact graphs

Although trees with a single point are easily classified, locally infinite quasipoints are considerably less so, and it is thus natural to ask for conditions on \mathbf{X} that make the quasitree $\Gamma \mathbf{X}$ bounded. For the moment, we study this question without using any group action on \mathbf{X} .

4.1 Applying the projection trichotomy

We begin with a simple corollary of Theorem 2.4.

Corollary 4.1 Let \mathbf{X} be a CAT(0) cube complex with no infinite collection of pairwise crossing hyperplanes. If $\text{diam}(\Gamma \mathbf{X}) < \infty$, then for every isolated 0-simplex $v \in \partial_\Delta \mathbf{X}$, there exists a hyperplane H such that v lies in the image of $\partial_\Delta N(H) \hookrightarrow \partial_\Delta \mathbf{X}$.

Proof By Corollary 3.20, a combinatorial geodesic ray γ representing an isolated 0-simplex v at infinity is rank-one. Hence, by Theorem 2.4, either γ projects to an unbounded path in $\Gamma \mathbf{X}$, or γ lies in a uniform neighborhood of a hyperplane H . In the latter case, it follows that all but finitely many hyperplanes crossing γ cross H . In this case, v lies in the image of $\partial_\Delta N(H)$. If $\partial_\Delta \mathbf{X}$ contains some isolated 0-simplex v , then Theorem 3.19 implies that v is visible, and the preceding argument shows that either $\Gamma \mathbf{X}$ has infinite diameter or v lies in the image of the simplicial boundary of a hyperplane. □

Example 4.2 (Spiral of quarter-flats) For $i \geq 1$, let \mathbf{F}_i be a nondiagonal quarter-flat, bounded by rays I_i and O_i . Let $\mathbf{X}_1 = \mathbf{F}_1$ and for $i \geq 2$, construct \mathbf{X}_i by attaching \mathbf{F}_i to \mathbf{X}_{i-1} by identifying the ray I_i with the subray $O'_{i-1} \subset O_{i-1}$ beginning at the

second 0–cube. \mathbf{X}_i is strongly locally finite and fully visible. For each $i < \infty$, the simplicial boundary $\partial_\Delta \mathbf{X}_i$ is a subdivided line segment of length i , and it is easily seen that $\text{diam}(\Delta \mathbf{X}_i) = i + 1$. The *spiral of quarter-flats* \mathbf{S} obtained by gluing two copies of $\bigcup_{i \geq 1} \mathbf{X}_i$ along I_0 has an unbounded contact graph, and $\partial_\Delta \mathbf{S}$ consists of \mathbb{R} together with a pair of isolated 0–simplices corresponding to a “spiraling” geodesic ray; see Figure 10.




Figure 10: An infinite spiral of quarter-flats. The arrowed geodesic spends finite time in each quarter-flat, and represents a pair of isolated 0–simplices in $\partial_\Delta \mathbf{S}$.

4.2 Relating $\Gamma \mathbf{X}$ and $\partial_\Delta \mathbf{X}$

The main theorem of this section gives conditions under which $\Gamma \mathbf{X}$ and $\Delta \mathbf{X}$ are bounded.

Theorem 4.3 *Let \mathbf{X} be a strongly locally finite, essential, compactly indecomposable CAT(0) cube complex with degree $D \in \mathbb{N} \cup \{\infty\}$. Then*

$$\text{diam}(\Gamma \mathbf{X}) \leq \text{diam}(\Delta \mathbf{X}) \leq \text{diam}(\partial_\Delta \mathbf{X}) \leq 2 \text{diam}(\Delta \mathbf{X}) - 2 \leq 2D \text{diam}(\Gamma \mathbf{X}) - 2.$$

Moreover,

- (1) if $\partial_\Delta \mathbf{X}$ is bounded, then $\text{diam}(\Gamma \mathbf{X}) < \infty$ and \mathbf{X} contains no rank-one geodesic ray;
- (2) \mathbf{X} decomposes as an iterated pseudoproduct if and only if $\Delta \mathbf{X}$ is bounded;
- (3) if \mathbf{X} decomposes as an iterated pseudoproduct, then $\partial_\Delta \mathbf{X}$ is bounded;
- (4) if $D < \infty$, then (2) holds with $\Delta \mathbf{X}$ replaced by $\Gamma \mathbf{X}$.

Proof Assertion (1) follows from Lemma 4.6 and Lemma 4.9. Assertion (2) is the content of Lemma 4.7. We get assertion (3) from Lemma 4.8, which establishes the bound $\text{diam}(\partial_{\Delta}\mathbf{X}) \leq 2 \text{diam}(\Delta\mathbf{X}) - 2$; the lower bounds on $\text{diam}(\partial_{\Delta}\mathbf{X})$ come from Lemma 4.9 and the fact that $\Delta\mathbf{X}$ is a spanning subgraph of $\Gamma\mathbf{X}$. Finally, Lemma 4.11 says that, if $D < \infty$, then $\Gamma\mathbf{X}$ is bounded if and only if $\Delta\mathbf{X}$ is bounded, which yields assertion (4). \square

Remark 4.4 Example 4.2 shows that the bounds on $\text{diam}(\partial_{\Delta}\mathbf{X})$ in Theorem 4.3 are attained. Also, if \mathbf{X} is the standard tiling of \mathbb{R}^2 by 2-cubes, then $\text{diam} \Delta\mathbf{X} = 2$ and $\text{diam}(\partial_{\Delta}\mathbf{X}) = 2$, since $\partial_{\Delta}\mathbf{X}$ is a 4-cycle. Looking ahead, one also sees that the combinatorial geodesic ray in \mathbf{S} that “spirals”, spending finite time in each quarter-flat, has quadratic divergence.

The assumption that \mathbf{X} is strongly locally finite is needed to give access to $\partial_{\Delta}\mathbf{X}$ and to guarantee that any infinite set of hyperplanes contains a UBS, which is needed to compare $\partial_{\Delta}\mathbf{X}$ to $\Delta\mathbf{X}$.

The following notion is used to relate $\Delta\mathbf{X}$ to $\partial_{\Delta}\mathbf{X}$.

Definition 4.5 (Pseudoproduct, iterated pseudoproduct) \mathbf{X} is the *pseudoproduct* of the cube complexes \mathbf{Q}_1 and \mathbf{Q}_2 , denoted by $\mathbf{X} \cong \mathbf{Q}_1 \boxtimes \mathbf{Q}_2$, if all of the following hold.

- (1) There is a convex embedding $a: \mathbf{Q}_1 \rightarrow \mathbf{X}$, and the image of \mathbf{Q}_1 is crossed by an infinite set \mathcal{V}_1 of hyperplanes. There is a restriction quotient $q_1: \mathbf{X} \rightarrow \mathbf{Q}_1$ such that a is a section of q_1 , and q_1 is obtained by restricting to the set \mathcal{V}_1 of hyperplanes.
- (2) There is a restriction quotient $q_2: \mathbf{X} \rightarrow \mathbf{Q}_2$ obtained by restricting to the set $\mathcal{V}_2 = \mathcal{W} - \mathcal{V}_1$ of hyperplanes. Moreover, \mathcal{V}_2 is infinite.
- (3) There is a cubical isometric embedding $e: \mathbf{X} \rightarrow \mathbf{Q}_1 \times \mathbf{Q}_2$ for which the convex hull of $e(\mathbf{X})$ is all of $\mathbf{Q}_1 \times \mathbf{Q}_2$. Also, if $\pi_1: \mathbf{Q}_1 \times \mathbf{Q}_2 \rightarrow \mathbf{Q}_1$ is the projection, then $\pi_1 \circ e \circ a$ is the identity on \mathbf{Q}_1 .
- (4) For each $H \in \mathcal{V}_2$, there are infinitely many $V \in \mathcal{V}_1$ such that $H \perp V$ in \mathbf{X} .

We say that \mathbf{X} is an *iterated pseudoproduct* if there exists a hyperplane $H_0 \subset \mathbf{X}$, an integer $k \geq 0$ called the *depth* and a sequence of cube complexes Q_1^i, Q_2^i , for $0 \leq i \leq k$, such that $\mathbf{X} \cong Q_1^0 \boxtimes Q_2^0$, for all $0 \leq i \leq k - 2$, and $Q_1^i \cong Q_1^{i+1} \boxtimes Q_2^{i+1}$, and finally $Q_1^k \cong H_0$, $Q_2^k \cong [-\frac{1}{2}, \frac{1}{2}]$ and $Q_1^{k-1} \cong Q_1^k \times Q_2^k$.

If $\mathbf{X} \cong \mathbf{Q}_1 \times \mathbf{Q}_2$, then $\mathbf{X} \cong \mathbf{Q}_1 \boxtimes \mathbf{Q}_2$. However, for example, an eighth-flat is a pseudoproduct of two infinite rays but is not a product of any two cube complexes since its crossing graph has diameter at least three.

Fix a base hyperplane H_0 . The *grade* $g(H)$ of the hyperplane H with respect to H_0 is the distance of H to H_0 in $\Delta\mathbf{X}$. In other words, H_0 has grade 0 and H has grade n if H crosses a grade- $(n-1)$ hyperplane but does not cross a grade- $(n-2)$ hyperplane.

4.2.1 Proof of Theorem 4.3 Theorem 4.3 is proved as the following sequence of lemmas. Throughout, \mathbf{X} is supposed to be strongly locally finite, leafless, and compactly indecomposable.

Lemma 4.6 *If $\partial_\Delta\mathbf{X}$ is bounded, then \mathbf{X} contains no rank-one combinatorial geodesic ray.*

Proof If $\partial_\Delta\mathbf{X}$ is bounded, then it is connected. Hence either $\partial_\Delta\mathbf{X}$ is a single 0-simplex, or $\partial_\Delta\mathbf{X}$ contains no isolated 0-simplex. In the latter case, \mathbf{X} contains no rank-one ray by Corollary 3.20. If $\partial_\Delta\mathbf{X}$ is a single 0-simplex v , then let \mathcal{V} be a minimal UBS representing v . But there is a hyperplane V that is initial in \mathcal{V} : every element of $\mathcal{V} - \{V\}$ lies in a single halfspace V^+ associated to V , or crosses V . But by essentiality, V^- contains a UBS representing a 0-simplex at infinity distinct from v , a contradiction. \square

Lemma 4.7 *$\Delta\mathbf{X}$ is bounded if and only if \mathbf{X} is an iterated pseudoproduct.*

Proof Choose $R \leq \text{diam } \Delta\mathbf{X} < \infty$ and a hyperplane H_0 such that, for all $H \in \mathcal{W}$, the grade $g(H)$ of H in $\Delta\mathbf{X}$, with respect to H_0 , is at most R . By essentiality, we can assume $R \geq 2$, for otherwise H_0 crosses each $H \neq H_0$, and thus H_0 is not essential. Moreover, choose R so that there exists H with $g(H) = R$.

Let \mathcal{V}_1 be the set of hyperplanes H with $0 \leq g(H) \leq R-1$, and let \mathcal{V}_2 be the set of grade- R hyperplanes. Since \mathbf{X} is compactly indecomposable, $N(H_0)$ cannot be compact, and thus there are infinitely many hyperplanes crossing H_0 . Hence \mathcal{V}_1 is infinite.

On the other hand, let H_R be a grade- R hyperplane. Then since $R \geq 2$, the hyperplanes H_R and H_0 do not cross. Since H_R is essential, there exists a hyperplane H' such that H_R separates H' from H_0 . Any path in $\Delta\mathbf{X}$ from H' to H_0 contains H_R or some hyperplane crossing H_R , and thus H' has grade at least R , whence, since every hyperplane has grade at most R , we have $H' \in \mathcal{V}_2$. Now H' is essential, and

thus separates some H'' from H_0 . Thus H'' has grade R , by the same argument. Hence \mathcal{V}_2 is infinite.

The quotients $\mathbf{X} \rightarrow \mathbf{Q}_i$ For $i \in \{1, 2\}$, let \mathbf{Q}_i be the cube complex dual to the wallspace $(\mathbf{X}^0, \mathcal{V}_i)$ whose walls are those induced by the hyperplanes in \mathcal{V}_i . This gives a cubical quotient $q_i: \mathbf{X} \rightarrow \mathbf{Q}_i$; this is just a restriction quotient that collapses the hyperplanes in $\mathcal{W} - \mathcal{V}_i$.

Let \mathbb{V}_i be the set of hyperplanes in \mathbf{Q}_i , so that there is a natural bijection $\mathcal{V}_i \rightarrow \mathbb{V}_i$ given by $V \mapsto q_i(V)$. Note that, if $V, V' \in \mathcal{V}_i$ are hyperplanes, then $V \perp V'$ if and only if $q_i(V) \perp q_i(V')$. Moreover, if V, V' osculate, then $q_i(V)$ and $q_i(V')$ osculate. For $V, V' \in \mathcal{V}_1$, the converse is true. If no hyperplane separates $q_1(V)$ from $q_1(V')$, then any hyperplane U separating V, V' must belong to \mathcal{V}_2 . But if $g(V), g(V') \leq R-1$ and U separates V, V' , then $g(U) \leq R-1$, a contradiction. In other words, \mathcal{V}_1 is inseparable, and hence $V, V' \in \mathcal{V}_1$ contact if and only if their images in \mathbf{Q}_1 do. However, there could exist noncontacting $V, V' \in \mathcal{V}_2$ such that $q_2(V)$ osculates with $q_2(V')$, if every hyperplane separating V, V' has grade at most $R-1$.

The embedding $e: \mathbf{X} \rightarrow \mathbf{Q}_1 \times \mathbf{Q}_2$ Fix a base 0-cube $x_o \in \mathbf{X}$. Recall that this is a section $x_o: \mathcal{W} \rightarrow \mathcal{W}^\pm$ of the map $\pi: \mathcal{W}^\pm \rightarrow \mathcal{W}$ that associates each halfspace of \mathbf{X} to the corresponding hyperplane. Moreover, x_o is automatically consistent and canonical. Fix a base 0-cube $(y_o^1, y_o^2) \in \mathbf{Q}_1 \times \mathbf{Q}_2$. Then y_o^1 is a consistent, canonical orientation of the hyperplanes $q_1(V) \in \mathbb{V}_1$ and y_o^2 a consistent canonical orientation of the hyperplanes in \mathbb{V}_2 . Define $y_o^i(q_i(V)) = q_i(x_o(V))$, ie the image in \mathbf{Q}_i of the halfspace associated to V that contains x_o . This is the halfspace of \mathbf{Q}_i containing $q_i(x_o)$. In other words, let $y_o^i = q_i(x_o)$.

We now define a map $\mathbf{X}^0 \ni x \mapsto \psi(x) = (\psi^1(x), \psi^2(x)) \in \mathbf{Q}_1 \times \mathbf{Q}_2$ as follows. Let $\psi^i(x)(q_i(V)) = q_i(x(V))$. Then, since x is consistent, for all $V, V' \in \mathcal{V}_i$, we have $x(V) \cap x(V') \neq \emptyset$ and hence $q_i(x(V)) \cap q_i(x(V')) \supseteq q_i(x(V) \cap x'(V)) \neq \emptyset$. Thus $\psi^i(x)$ is a consistent orientation of the hyperplanes of \mathbf{Q}_i . Since x differs from x_o on finitely many hyperplanes in \mathcal{V}_i , and $\psi^i(x)(q_i(V)) \neq \psi^i(x_o)(q_i(V))$ only if $x(V) \neq x_o(V)$, the 0-cube $\psi^i(x)$ is canonical. Hence $\psi(x)$ is a genuine 0-cube of $\mathbf{Q}_1 \times \mathbf{Q}_2$.

Now $\psi^i(x)(q_i(V))$ differs from $\psi^i(y)(q_i(V))$ when $q_i(x(V))$ and $q_i(y(V))$ are different halfspaces associated to $q_i(V)$. This happens if and only if $x(V) \neq y(V)$, so that $d_{\mathbf{Q}_1 \times \mathbf{Q}_2}(\psi(x), \psi(y))$ counts the hyperplanes in $\mathcal{V}_1 \sqcup \mathcal{V}_2 = \mathcal{W}$ on which x, y differ. Thus $\psi: \mathbf{X}^0 \rightarrow (\mathbf{Q}_1 \times \mathbf{Q}_2)^0$ is an isometric embedding, and hence the map ψ extends to an isometric embedding $e: \mathbf{X} \rightarrow \mathbf{Q}_1 \times \mathbf{Q}_2$, since each median graph uniquely determines a CAT(0) cube complex.

The embedding $a: \mathbf{Q}_1 \rightarrow \mathbf{X}$ Choose a base 0-cube $x_o \in N(H_0)$. Let \mathbf{Y}^0 be the set of 0-cubes y in \mathbf{X} such that every hyperplane separating y from x_o belongs to \mathcal{V}_1 . Let $\mathbf{Y} \subset \mathbf{X}$ be the convex hull of \mathbf{Y}^0 . Each 0-cube $y \in \mathbf{Y}$, viewed as a section of π , has the property that $y(W) = x_o(W)$ for $W \in \mathcal{V}_2$, since W cannot separate y from x_o . In other words, \mathbf{Y} is isomorphic to the restriction quotient obtained by restricting each 0-cube of \mathbf{X} to \mathcal{V}_1 ! More precisely, for each $q(x) \in \mathbf{Q}_1$, define $a(q(x)): \mathcal{W} \rightarrow \mathcal{W}^\pm$ by $a(q(x))(W) = x(W)$ for $W \in \mathcal{V}_1$ and, for all $W \in \mathcal{V}_2$, let $a(q(x))(W)$ be the halfspace containing H_0 (and thus x_o). This is well-defined since no $W \in \mathcal{W}_2$ crosses H_0 . This orientation is consistent, since x is consistent on \mathcal{W}_1 . If $a(q(x))$ orients $W \in \mathcal{W}_1$ inconsistently with $V \in \mathcal{W}_2$, then V separates W from H_0 , but we have seen this to be impossible. Finally, $a(q(x))$ agrees with x_o on all but finitely many hyperplanes, namely those hyperplanes in \mathcal{V}_1 on which x, x_o disagree. Hence $a(q(x))$ is canonical and belongs to \mathbf{Y} , and, if $x \in \mathbf{Y}$, then $a(q(x)) = x$. In the usual way, one checks that $q(x) \mapsto a(q(x))$ determines an isometric embedding $a: \mathbf{Q}_1 \rightarrow \mathbf{X}$. (In general, there is no isometric embedding $\mathbf{Q}_2 \rightarrow \mathbf{X}$.)

The convex hull of $e(\mathbf{X})$ Let $\mathbf{Z} \subseteq \mathbf{Q}_1 \times \mathbf{Q}_2$ be the convex hull of $e(\mathbf{X})$. Every hyperplane of $\mathbf{Q}_1 \times \mathbf{Q}_2$ is of the form $q_1(V) \times \mathbf{Q}_2$ or $\mathbf{Q}_1 \times q_2(H)$, where $V \in \mathcal{V}_1$ and $H \in \mathcal{V}_2$, and each hyperplane of the former type crosses each hyperplane of the latter type. It follows that every hyperplane of $\mathbf{Q}_1 \times \mathbf{Q}_2$ crosses $e(\mathbf{X})$, whence $\mathbf{Z} = \mathbf{Q}_1 \times \mathbf{Q}_2$. Indeed, \mathbf{Z} is either all of $\mathbf{Q}_1 \times \mathbf{Q}_2$, or \mathbf{Z} is the intersection of all halfspaces that contain $e(\mathbf{X})$. The latter is impossible since every hyperplane crosses $e(\mathbf{X})$.

The 2-cubes Let $H \in \mathcal{V}_2$ be a grade- R hyperplane. By essentiality, there exists an infinite set U_0, U_1, \dots of grade- R hyperplanes such that H separates U_i from H_0 , for all $i \geq 0$, and for all $i \geq 1$, the hyperplanes U_{i+1} and U_{i-1} are separated by U_i . Every U_i crosses a set $\{W_i^j\}_j$ of grade- $R-1$ hyperplanes, each of which must cross H . Let Ω be the full subgraph of $\Delta\mathbf{X}$ generated by $\{W_i^j\} \cup \{H\}$. Then every path in $\Delta\mathbf{X}$ from U_i to H_0 passes through H or through one of the W_i^j , since H separates U_i from H_0 . Thus $\Delta\mathbf{X} - \Omega$ has at least two nonempty components, namely the component containing U_0 and the component containing H_0 . Hence, by Lemma 4.10 and compact indecomposability, Ω is infinite, whence infinitely many grade- $(R-1)$ hyperplanes cross H in \mathbf{X} . We conclude that $\mathbf{X} \cong \mathbf{Q}_1 \boxtimes \mathbf{Q}_2$.

Induction \mathbf{Q}_1 has infinitely many hyperplanes and $\Delta\mathbf{Q}_1$ has finite diameter – grading from H_0 , every hyperplane has grade at most $R-1$. If $R=2$, then $\mathbf{Q}_1 \cong H_0 \times [-\frac{1}{2}, \frac{1}{2}]$. Otherwise, by essentiality, the set \mathcal{V}_2^1 of grade $(R-1)$ hyperplanes is infinite, and, letting \mathcal{V}_1^1 be the (infinite) set of hyperplanes of grade at most $R-2$, we proceed as before, to find that $\mathbf{Q}_1 \cong \mathbf{Q}_1^1 \boxtimes \mathbf{Q}_2^1$. Continuing in this way, since the highest grade decreases at each step, we find that as long as $R > 2$, we can continue

the pseudoproduct decompositions, using the above argument. This terminates with $Q_1^{R-1} \cong H_0 \times [-\frac{1}{2}, \frac{1}{2}] \cong N(H_0)$. Hence \mathbf{X} is an iterated pseudoproduct.

Bounding $\text{diam}(\Delta\mathbf{X})$ Conversely, suppose that \mathbf{X} is an iterated pseudoproduct, and let k be the constant from Definition 4.5. Then, for each hyperplane $H \neq H_0$ of \mathbf{X} , either there exists a maximal $m \leq k - 1$ such that H survives in the restriction quotient $\mathbf{X} \rightarrow \mathbf{Q}_1^0 \rightarrow \mathbf{Q}_1^1 \rightarrow \dots \rightarrow \mathbf{Q}_1^m$, or H crosses a hyperplane crossing \mathbf{Q}_1^0 . In the latter case, $Q_1^m \cong Q_1^{m+1} \boxtimes Q_{m+1}^2$, so that, by the definition of a pseudoproduct, H crosses a hyperplane that crosses Q_1^{m+1} . Hence H is at distance at most $k + 1$ from H_0 in $\Delta\mathbf{X}$. □

Lemma 4.8 *If \mathbf{X} is an iterated pseudoproduct, then $\text{diam}(\partial_\Delta\mathbf{X}) \leq 2(\text{diam}(\Delta\mathbf{X}) - 1)$.*

Proof Let $R < \infty$ be the maximal grade in $\Delta\mathbf{X}$ with respect to a hyperplane H_0 , and let $\mathbf{X} \cong \mathbf{Q}_1 \boxtimes \mathbf{Q}_2$ be the pseudoproduct decomposition from Lemma 4.7. Since $a(\mathbf{Q}_1) \subset \mathbf{X}$ is convex, we have a simplicial embedding $\partial_\Delta a: \partial_\Delta\mathbf{Q}_1 \rightarrow \partial_\Delta\mathbf{X}$, by Theorem 3.15. Let $A_1 = \partial_\Delta a(\partial_\Delta\mathbf{X})$ be its image. Hence, by Theorem 3.30, we have $\partial_\Delta(\mathbf{Q}_1 \times \mathbf{Q}_2) \cong A_1 \star \partial_\Delta\mathbf{Q}_2$, since $e \circ a: \mathbf{Q}_1 \rightarrow \mathbf{Q}_2$ embeds \mathbf{Q}_1 convexly in $\mathbf{Q}_1 \times \mathbf{Q}_2$ (with image $\mathbf{Q}_1 \times y_0^2$).

Let $A_2 \cong \partial_\Delta\mathbf{Q}_2$, so that $\partial_\Delta(\mathbf{Q}_1 \times \mathbf{Q}_2) = A_1 \star A_2$. In particular, $A_1 \cap A_2 = \emptyset$. Thus every simplex of $\partial_\Delta(\mathbf{Q}_1 \times \mathbf{Q}_2)$ is of the form $[u_0, u_1, \dots, u_p, v_0, \dots, v_q]$, where each u_i represented in $\partial_\Delta\mathbf{X}$ by a UBS $\mathcal{U}_i \subset \mathcal{V}_1$ and each v_j is represented by a UBS $q_2(\mathcal{W}_j)$, where $\mathcal{W}_j \subset \mathcal{V}_2$ is a UBS.

Let Q^2 be a UBS of hyperplanes in \mathbf{Q}_2 that represents the 0-simplex q . Let $\{W_i\}_{i \geq 0} \subseteq Q^2$ be a set of hyperplanes such that for all $i \geq 1$, the hyperplane W_i separates W_{i-1} from W_{i+1} . By the pseudoproduct decomposition, for each $i \geq 0$, there exist infinitely many hyperplanes H of grade at most $R - 1$ that cross W_i in \mathbf{X} , and thus cross W_j in \mathbf{X} , for all $j < i$. Hence there is a UBS Q^1 representing a 0-simplex $q^1 \in \partial_\Delta\mathbf{Q}_1$ such that $[q^1, q]$ is a 1-simplex of $\text{im}(\partial_\Delta\mathbf{X} \rightarrow \partial_\Delta(\mathbf{Q}_1 \times \mathbf{Q}_2))$. In other words, each 0-simplex of $\partial_\Delta\mathbf{Q}_2$ is adjacent in the image of $\partial_\Delta\mathbf{X}$ to a 0-simplex of $\partial_\Delta\mathbf{Q}_1$.

We now argue by induction on R . First suppose that $R \geq 3$ and, by induction, that there is a function $K: \mathbb{N} \rightarrow \mathbb{N}$ such that, for all 0-simplices $s, s' \in \partial_\Delta\mathbf{Q}_1$, there is a path in $\partial_\Delta\mathbf{X}$ of length at most $K(R - 1)$ joining s to s' . Then, if $q, q' \in \partial_\Delta\mathbf{Q}_2$ are 0-simplices, they are respectively at distance 1 from some 0-simplices $s, s' \in \partial_\Delta\mathbf{Q}_1$, by the preceding paragraph, and hence there is a path in $\partial_\Delta\mathbf{X}$ of length at most $K(R) = K(R - 1) + 2$ joining q, q' . Hence we have

$$\text{diam}(\partial_\Delta\mathbf{X}) \leq K(R) = K(2) + 2(R - 2).$$

Next, let $R = 2$ and let s^1, s^2 be 0-simplices of $\partial_\Delta \mathbf{Q}_1$, so that for $i \in \{1, 2\}$, the simplex s^i is represented by a UBS \mathcal{S}^i consisting only of grade-1 hyperplanes. Let $\{S_j^i\}_{j \geq 0} \subseteq \mathcal{S}^i$ be a maximal set of pairwise noncrossing hyperplanes. Either $[s^1, s^2]$ is a 1-simplex, and we are done, or we can choose \mathcal{S}^i so that $\mathcal{S}^1 \cap \mathcal{S}^2 = \emptyset$.

In the latter case, for all $j, k \geq 0$, there is a geodesic path $S_j^1 \perp H_0 \perp S_k^2$ in $\Delta \mathbf{X}$, where H_0 is the base hyperplane of the grading. Thus the set of hyperplanes H such that $H \perp S_j^1$ and $H \perp S_k^2$ is nonempty. On the other hand, the set of such H , together with the two element set $\{S_j^1, S_k^2\}$, generates a subgraph $\Lambda \subset \Delta \mathbf{X}$ such that $\Delta \mathbf{X} - \Lambda$ has at least two infinite components, since \mathbf{X} is essential. Since \mathbf{X} is compactly indecomposable, there are infinitely many such H , by Lemma 4.10, and hence there is a minimal UBS \mathcal{U} such that, for all $U \in \mathcal{U}$, we have a path $S_j^1 \perp U \perp S_k^2$ in $\Delta \mathbf{X}$. But then s^1 and s^2 are both adjacent to the 0-simplex u represented by \mathcal{U} . Thus $K(2) = 2$ and $\text{diam}(\partial_\Delta \mathbf{X}) \leq 2R - 2 \leq 2 \text{diam}(\Delta \mathbf{X}) - 2$, as required. \square

Lemma 4.9 *If \mathbf{X} is essential and strongly locally finite, then $\text{diam}(\Delta \mathbf{X}) \leq \text{diam}(\partial_\Delta \mathbf{X})$.*

Proof Choose $R \leq \text{diam}(\Delta \mathbf{X})$. Also, assume that we have chosen $R \geq 2$, for if we can not make this choice then essentiality is contradicted. Choose hyperplanes U_0 and V_0 such that $d_{\Delta \mathbf{X}}(U_0, V_0) = R \geq 2$. By essentiality, there exists a minimal UBS \mathcal{U} and a minimal UBS \mathcal{V} such that $U_0 \in \mathcal{U}$ and $V_0 \in \mathcal{V}$ are initial hyperplanes in the given UBSs: for all $U' \in \mathcal{U}$, U_0 separates U' from V_0 , and for all $V' \in \mathcal{V}$, V_0 separates V' from U_0 .

Let u, v be the 0-simplices at infinity represented by \mathcal{U} and \mathcal{V} respectively. If u, v lie in different components of $\partial_\Delta \mathbf{X}$, then $\text{diam}(\partial_\Delta \mathbf{X}) = \infty$ and we are done. Therefore, let u, s_1, \dots, s_k, v be a path in $(\partial_\Delta \mathbf{X})^1$ joining u to v . Let \mathcal{S}_i be a UBS representing the 0-simplex s_i . Then there is a path $\sigma = U \perp S_1 \perp \dots \perp S_k \perp V$ in $\Delta \mathbf{X}$, where $U \in \mathcal{U}, V \in \mathcal{V}$ and for each $1 \leq t \leq k$, we have $S_t \in \mathcal{S}_t$. Without loss of generality, U_0 and V_0 both separate U from V . Hence σ must either contain U_0 (respectively, V_0) or σ must contain a hyperplane S_t that crosses U_0 (respectively, a hyperplane S_r that crosses V_0). Hence, in the worst case, we have a path $U_0 \perp S_t \perp S_{t+1} \perp \dots \perp S_r \perp V_0$ of length $r - t + 2$ joining U_0 to V_0 . Thus we have $k + 1 \geq r - t + 2 \geq R$, whence $\text{diam}(\partial_\Delta \mathbf{X}) \geq \text{diam}(\Delta \mathbf{X})$. \square

Lemma 4.10 allows us to assume that no finite subgraph separates the crossing graph; this is how compact indecomposability manifests itself in $\Delta \mathbf{X}$.

Lemma 4.10 *Let \mathbf{X} be strongly locally finite and compactly indecomposable. Then for any finite subgraph Λ of $\Delta \mathbf{X}$, the complement $\Delta \mathbf{X} - \Lambda$ is connected.*

Proof that ΔX is connected If ΔX is disconnected, X is the wedge sum of two proper subcomplexes, and is thus not compactly indecomposable. Indeed, let H, H' be hyperplanes in different components of the crossing graph. If H'' separates H from H' , then H'' cannot lie in the same component of ΔX as both H and H' , so we may choose these hyperplanes so that $H \perp \sim H'$. Now $N(H) \cap N(H')$ consists of a single 0-cube x , since, if $N(H) \cap N(H')$ contained distinct 0-cubes x, y , each hyperplane separating x from y would cross H and H' .

Neither of H, H' is compact, since X is compactly indecomposable. Thus we can choose 0-cubes $y \in N(H)$ and $y' \in N(H')$ so that x, y, y' are pairwise distinct. Suppose that $P \rightarrow X$ is a path joining y to y' with $x \notin P$. Let $A \rightarrow N(H)$ be a shortest path joining x to y and let $B \rightarrow N(H')$ be a shortest path joining y' to x , and consider a minimal-area disc diagram $D \rightarrow X$ bounded by APB , as in Figure 11. Let K be the dual curve emanating from the initial 1-cube of A and let K' be the dual curve emanating from the terminal 1-cube of B . Let U be the hyperplane to which K maps, and let U' be the hyperplane to which K' maps. If $U = U'$, then $H \perp U \perp H'$, a contradiction. If K crosses K' , then $H \perp U \perp U' \perp H'$, another contradiction. By minimality of area, K and K' end on P . Let P' be the subpath of P between and including the 1-cubes c and c' dual to K and K' , as in Figure 11, so that $P' = cQc'$, where Q is a path such that $|Q| < |P|$ and $x \notin Q$. On the other hand, $x \in N(U) \cap N(U')$, and no hyperplane W crosses U and U' , since otherwise the path $H \perp U \perp W \perp U' \perp H'$ would join H and H' in ΔX . Thus we can replace H and H' by U and U' respectively, and replace P by Q , and obtain a lower-length counterexample. Eventually, we are in a situation where $|Q| = 0$ and $x \notin Q$. But then $Q \in N(U) \cap N(U')$ is separated from $x \in N(U) \cap N(U')$, and thus U and U' are crossed by a common hyperplane, a contradiction. Hence no such P can exist, ie every path from $N(H)$ to $N(H')$ passes through x . It is now obvious that x disconnects X .




Figure 11: If ΔX is disconnected, then X has a cut-0-cube.

Disconnecting ΔX in a restriction quotient Now suppose Λ is a finite graph that separates ΔX . Let $\mathcal{W}' \subset \mathcal{W}$ be the set of hyperplanes that do not correspond to vertices

of Λ . Let $\mathbf{X}' = \mathbf{X}(\mathcal{W}')$ be the restriction quotient corresponding to this subset, and let $q: \mathbf{X} \rightarrow \mathbf{X}'$ be the quotient map. If W, W' are hyperplanes of \mathbf{X} , then W, W' cross if and only if $q(W)$ crosses $q(W')$ in \mathbf{X}' . Therefore, if $\Delta\mathbf{X} - \Lambda$ is disconnected, it follows that $\Delta\mathbf{X}'$ is disconnected, and therefore that $\mathbf{X}' \cong \mathbf{Y}_1 \cup_{\{a\}} \mathbf{Y}_2$, where \mathbf{Y}_i is a nonempty subcomplex containing at least one hyperplane, and a is a single 0-cube. Let $K = q^{-1}(\mathbf{Y}_1) \cap q^{-1}(\mathbf{Y}_2) = q^{-1}(a)$. Since q is continuous, K separates \mathbf{X} into two nonempty subcomplexes. On the other hand, if $b, b' \in K$ are 0-cubes, then $q(b) = q(a) = q(b')$, so that $b(H) \neq b(H')$ if and only if H is a hyperplane corresponding to a vertex of Λ . Thus K is finite and disconnects \mathbf{X} , contradicting compact indecomposability. \square

Lemma 4.11 *Let \mathbf{X} be compactly indecomposable with $\deg D < \infty$. If $\text{diam}(\Gamma\mathbf{X}) < \infty$, then $\text{diam}(\Delta\mathbf{X}) < \infty$.*

Proof Since \mathbf{X} is compactly indecomposable, $\Delta\mathbf{X}$ is path-connected, by Lemma 4.10. If H and H' are osculating hyperplanes and $d_{\Delta\mathbf{X}}(H, H') = R$, then a simple disc diagram argument shows there is a shortest path $H = H_0 \perp H_1 \perp \cdots \perp H_{R-1} \perp H_R = H'$ in $\Delta\mathbf{X}$ with the property that $H_i \perp \smile H_j$ for all i, j , and hence $N(H) \cap N(H')$ contains a 0-cube of degree at least $R + 1$. Thus, if H, H' osculate, then $d_{\Delta\mathbf{X}}(H, H') \leq D - 1$. Thus $d_{\Delta\mathbf{X}}(H, H') \leq D$ whenever $H \perp \smile H'$, and it follows from the triangle inequality that $\text{diam}(\Delta\mathbf{X}) \leq (D - 1) \text{diam}(\Gamma\mathbf{X})$. \square

5 Rank-one isometries, $\Gamma\mathbf{X}$ and $\partial_{\Delta}\mathbf{X}$

In this section, we first use Theorems 2.2 and 2.4 to analyze the action of hyperbolic isometries of \mathbf{X} on $\Gamma\mathbf{X}$, and then reinterpret the rank-rigidity theorem of Caprace and Sageev [6] in terms of the contact graph and the simplicial boundary. Rank-one isometries of a \mathbf{X} whose axes diverge from every hyperplane are exactly those that act as hyperbolically on $\Gamma\mathbf{X}$.

Proposition 5.1 *Let G act on the CAT(0) cube complex \mathbf{X} and let $g \in G$. If g acts on $\Gamma\mathbf{X}$ with an unbounded orbit, then g is rank-one.*

Proof By Theorem 1.6, g has a combinatorial axis γ . If g is not rank-one, then for any hyperplane H crossing γ , we have $d_{\Gamma\mathbf{X}}(g^n H, H) \leq 3$ for all $n \in \mathbb{Z}$, since $g^n H$ and H cross a common isometric half-flat by Proposition 2.5. If U is any other hyperplane, then $d_{\Gamma\mathbf{X}}(g^n U, U) \leq 3 + 2d_{\Gamma\mathbf{X}}(H, U)$, and thus every g -orbit in $\Gamma\mathbf{X}$ is bounded. \square

Theorem 5.2 *Let G act on the strongly locally finite CAT(0) cube complex \mathbf{X} and let $g \in G$ be a combinatorially rank-one element. Suppose that for all $n > 0$, and for all hyperplanes H , $g^n \notin \text{Stab}(H)$. Then g has a quasigeodesic axis in $\Gamma\mathbf{X}$.*

Proof Apply Theorems 2.2 and 2.4 to a combinatorial geodesic axis for g . □

We now restate the rank-rigidity theorem [6, Theorem A, Theorem B] in terms of the contact graph and the simplicial boundary. Theorem 5.4 follows directly from rank-rigidity, since the contact graph of a nontrivial product is a nontrivial join, and a hyperbolic element obtained as in [6] by applying double-skewering to a pair of strongly separated (ie at distance at least three in $\Delta\mathbf{X}$) hyperplanes necessarily acts with an unbounded orbit on $\Gamma\mathbf{X}$. We first need the following.

Remark 5.3 A group G acting on \mathbf{X} acts by simplicial automorphisms on $\partial_\Delta\mathbf{X}$. Also, the set of hyperplanes crossing a CAT(0) geodesic ray is a UBS, so that $g \in G$ has a fixed point in $\partial_\infty\mathbf{X}$ if and only if g fixes the barycenter of a simplex of $\partial_\Delta\mathbf{X}$.

Theorem 5.4 (Rank-rigidity, contact graph form) *Let G act essentially on the finite-dimensional CAT(0) cube complex \mathbf{X} . Suppose in addition that either G acts properly and cocompactly, or G does not stabilize a simplex of $\partial_\Delta\mathbf{X}$. If $\Gamma\mathbf{X}$ is unbounded, then G contains a combinatorially rank-one element g , no power of which stabilizes a hyperplane. Otherwise, $\Gamma\mathbf{X}$ has diameter 2 and there exist unbounded convex subcomplexes $\mathbf{Q}_1, \mathbf{Q}_2 \subset \mathbf{X}$ such that $\mathbf{X} \cong \mathbf{Q}_1 \times \mathbf{Q}_2$.*

By Theorem 3.30, a product decomposition of \mathbf{X} with unbounded factors corresponds to a join decomposition of $\partial_\Delta\mathbf{X}$, while an axis for a rank-one isometry yields a pair of isolated 0-simplices of $\partial_\Delta\mathbf{X}$. Rank-rigidity therefore implies the following.

Theorem 5.5 (Rank-rigidity, simplicial boundary form) *Under the same hypotheses on G and \mathbf{X} as in Theorem 5.4, if $\partial_\Delta\mathbf{X}$ is connected, there exist unbounded convex subcomplexes $\mathbf{Q}_1, \mathbf{Q}_2 \subset \mathbf{X}$ such that $\mathbf{X} \cong \mathbf{Q}_1 \times \mathbf{Q}_2$, and $\partial_\Delta\mathbf{X}$ decomposes as a nontrivial simplicial join. Otherwise, G contains a rank-one element g , no power of which stabilizes a hyperplane.*

Problem 5.6 Without using [6, Proposition 5.1], show that if G acts on \mathbf{X} , satisfying the hypotheses of either form of the rank-rigidity theorem, and \mathbf{X} decomposes as an iterated pseudoproduct of depth 2, then $\mathbf{X} \cong \mathbf{Q}_1 \times \mathbf{Q}_2$ for unbounded subcomplexes $\mathbf{Q}_1, \mathbf{Q}_2$. A solution to this problem would yield an alternate proof of rank-rigidity (in the locally finite case) using Theorem 4.3 together with double-skewering; this is discussed in [22, Chapter 6].

6 Cubical divergence

We study combinatorial divergence of geodesic rays in an arbitrary strongly locally finite CAT(0) cube complex \mathbf{X} , and obtain characterizations of linear and superlinear divergence and divergence of geodesics in terms of $\partial_\Delta \mathbf{X}$. In Section 6.3, we state a criterion for linear divergence of cocompactly cubulated groups stated in terms of $\partial_\Delta \mathbf{X}$ or $\Gamma \mathbf{X}$.

6.1 Simplex-paths and combinatorial divergence

Choose a 0-cube $x \in \mathbf{X}$. Note that, by Lemma 3.21, each visible simplex of $\partial_\Delta \mathbf{X}$ is represented by the set $\mathcal{W}(\gamma)$, where $\gamma: [0, \infty) \rightarrow \mathbf{X}$ is a combinatorial geodesic ray with $\gamma(0) = x$. Given such a ray γ , let $v(\gamma)$ denote the simplex at infinity represented by the UBS $\mathcal{W}(\gamma)$.

Definition 6.1 (Simplex path, simplicial Tits distance) Let u, v be simplices of $\partial_\Delta \mathbf{X}$. A *simplex-path* is a sequence $u = u_0, u_1, \dots, u_n = v$ such that $u_i \cap u_{i+1} \neq \emptyset$ for $0 \leq i \leq n - 1$. The *length* of this simplex-path is n . The *simplicial Tits distance* $\eta(u, v)$ from u to v is the length of a shortest simplex-path whose initial simplex is u and whose terminal simplex is v . Note that $\eta(u, v) \leq 1$ if and only if $u \cap v \neq \emptyset$, and $\eta(u, v) \leq 2$ if u, v are contained in a common simplex.

The next lemma allows us to induct on the simplicial Tits distance between visible simplices without assuming that \mathbf{X} is fully visible.

Lemma 6.2 Let u_0, u_1, \dots, u_n be a shortest simplex-path joining u_0 to u_n . Then, for $1 \leq i \leq n$, the u_i can be chosen so that the simplex u_i is visible.

Proof Each simplex is contained in a visible simplex by Theorems 3.14 and 3.19. \square

Definition 6.3 (r -avoiding path, combinatorial divergence) With respect to $x \in \mathbf{X}^{(0)}$, for $r \geq 0$, the combinatorial path $P \rightarrow \mathbf{X}$ is r -avoiding if, for all 0-cubes $p \in P$, we have that $d_{\mathbf{X}}(p, x) \geq r$. Let $\gamma, \gamma': [0, \infty) \rightarrow \mathbf{X}$ be combinatorial geodesic rays with $\gamma(0) = \gamma'(0) = x$. For $r \geq 0$, let $\mathbf{div}(\gamma, \gamma')(r) = \inf_P |P|$, where P varies over all r -avoiding paths joining $\gamma(r)$ to $\gamma'(r)$. This is the *combinatorial divergence* of the rays γ, γ' .

We say that the rays γ, γ' diverge *linearly* if there exist $A > 0, B \in \mathbb{R}$ such that $\mathbf{div}(\gamma, \gamma')(r) \leq Ar + B$ for all $r \geq 0$. The bi-infinite combinatorial geodesic $\gamma: \mathbb{R} \rightarrow \mathbf{X}$ has *linear (combinatorial) divergence* if the rays $\gamma([0, \infty))$ and $\gamma((-\infty, 0])$ diverge

linearly. \mathbf{X} has *linear divergence of rays* if each pair of combinatorial geodesic rays in \mathbf{X} diverges linearly, and \mathbf{X} has *(uniformly) linear divergence* if there exist $A \geq 0, B \in \mathbb{R}$ such that $\mathbf{div}(\gamma, \gamma')(r) \leq Ar + B$ for all combinatorial geodesic rays γ, γ' emanating from x and for all $r \geq 0$.

From now on, all geodesic rays are understood to emanate from a fixed initial 0-cube x_0 . If $\gamma, \gamma', \gamma''$ are three such rays, then $\mathbf{div}(\gamma, \gamma')(r) \leq \mathbf{div}(\gamma, \gamma'')(r) + \mathbf{div}(\gamma', \gamma'')(r)$ for all $r \geq 0$.

6.2 Divergence of pairs of combinatorial geodesic rays

Let $\gamma, \gamma': [0, \infty) \rightarrow \mathbf{X}$ be combinatorial geodesic rays emanating from x_0 and let $v, v' \subset \partial_\Delta \mathbf{X}$ be the simplices at infinity represented by γ and γ' respectively.

Lemma 6.4 *If $v \subseteq v'$, then there exists $B \in \mathbb{R}$ such that $\mathbf{div}(\gamma, \gamma')(r) \leq 2r + B$ for all $r \geq 0$. If v and v' lie in a common simplex of $\partial_\Delta \mathbf{X}$, then there exists B such that the same inequality holds for all $r \geq 0$.*

Proof Suppose that $v \subseteq v'$. Then the set $\mathcal{H} = \mathcal{W}(\gamma) - \mathcal{W}(\gamma')$ of hyperplanes that cross γ but not γ' is finite. The set $\mathcal{H}' = \mathcal{W}(\gamma') - \mathcal{W}(\gamma)$ of hyperplanes that cross γ' but not γ is finite if and only if $v = v'$.

First, if $\mathcal{H} = \emptyset$, then for all $r \geq 0$, $\mathbf{div}(\gamma, \gamma')(r) = d_{\mathbf{X}}(\gamma(r), \gamma'(r)) \leq 2r$. To see this, first fix $r \geq 0$. Consider the set of hyperplanes crossing the geodesic segment $P_r = \gamma([0, r])$. Since $\mathcal{H} = \emptyset$, this set of hyperplanes is partitioned into two subsets: \mathcal{A}_r is the set of hyperplanes that cross P_r and $P'_r = \gamma'([0, r])$. The set of remaining hyperplanes, \mathcal{B}_r , all cross P_r and $\gamma' - P'_r$, since every hyperplane that crosses γ also crosses γ' .

The set of hyperplanes crossing P'_r consists of \mathcal{A}_r , together with a set \mathcal{C}_r of hyperplanes U such that each U crosses $\gamma - P_r$, or separates an infinite subray of γ from an infinite subray of γ' , and thus crosses all but finitely many elements of \mathcal{H} . Note that $|\mathcal{A}_r| + |\mathcal{B}_r| = |\mathcal{A}_r| + |\mathcal{C}_r| = r$. Note also that if $B \in \mathcal{B}_r$ and $C \in \mathcal{C}_r$, then $B \perp C$, so that the hyperplanes in \mathcal{B}_r are orientable independently of those in \mathcal{C}_r .

We now construct an r -avoiding geodesic segment Q_r joining $x'_r = \gamma'(r)$ to $x_r = \gamma(r)$, from which the claimed equality follows. Let $Q'_r(0) = x'_r$. Let B_0, \dots, B_b be the hyperplanes in \mathcal{B}_r , numbered so that $x'_r \in N(B_0)$ and $B_0 \perp \smile B_1 \perp \smile \dots \perp \smile B_b$. This assumption is justified, since any hyperplane separating B_i from B_{i+1} must cross P_r and thus belong to \mathcal{B}_r . Any hyperplane separating x'_r from B_0 is likewise in \mathcal{B}_r . For $0 \leq t \leq b$, let $Q'_r(t)$ orient all hyperplanes $U \in \mathcal{W} - \mathcal{B}_r$ toward the halfspace $x'_r(U)$, and




Figure 12: The sets $\mathcal{A}_r, \mathcal{B}_r, \mathcal{C}_r$ are shown at left. At right is the path Q_r .

let $Q'_r(t)(B_s) = x'_r(B_s)$ for $s > t$. For $s \leq t$, let $Q'_r(t)(B_s) = \mathbf{X} - x'_r(B_s)$. This creates a geodesic segment from x'_r to $Q'_r(t)$ that crosses exactly the set \mathcal{B}_r of hyperplanes. Note that each $Q'_r(t)$ is separated from x_o by $\mathcal{A}_r \cup \mathcal{C}_r$, and thus $d_{\mathbf{X}}(Q'_r(t), x_o) \geq r$ for all t . In the same manner, we flip the hyperplanes \mathcal{C}_r successively to yield a geodesic segment Q''_r from $Q'_r(b)$ to x_r , with Q''_r separated from x_o by $\mathcal{A}_r \cup \mathcal{B}_r$. The segment $Q_r = Q'_r Q''_r$ is the desired r -avoiding geodesic segment. This yields the desired bound on the divergence.

Now, if \mathcal{H} is nonempty, there exists a combinatorial geodesic ray γ'' with $\gamma''(0) = x_o$, such that $d_{\mathbf{X}}(\gamma(r), \gamma''(r)) \leq 2|\mathcal{H}|$ for each $r \geq 0$. Moreover, $\mathcal{W}(\gamma) = \mathcal{W}(\gamma'') \cup \mathcal{H}$, so that $\mathcal{W}(\gamma'') \subseteq \mathcal{W}(\gamma')$. Given such a ray γ'' , we see that

$$\mathbf{div}(\gamma, \gamma')(r) \leq \mathbf{div}(\gamma, \gamma'')(r) + \mathbf{div}(\gamma', \gamma'')(r) \leq 2|\mathcal{H}| + 2r,$$

and the first assertion of the lemma follows, with $B = 2|\mathcal{H}|$.

It thus suffices to produce the ray γ'' . Let $\mathcal{A}_r, \mathcal{B}_r, \mathcal{C}_r$ be as above and let \mathcal{H}_r be the set of $H \in \mathcal{H}$ that cross P_r . Since \mathcal{H} is finite, \mathcal{H}_r has uniformly bounded cardinality $h \leq |\mathcal{H}|$. Let $\mathcal{H} = \{H_0, \dots, H_h\}$, where for $i < j$, the hyperplane H_i separates the 1-cube $N(H_j) \cap \gamma$ from x_o . Let R be the subpath of γ joining x_o to the 0-cube of $N(H_h) \cap \gamma$ separated from x_o by H_h . Let α be the combinatorial geodesic subray of γ emanating from the terminal 0-cube of R . Then H_h crosses every hyperplane dual to α and $\alpha \rightarrow N(H_h)$. For $t \geq 0$, define $\alpha'(t)$ to be the 0-cube that orients each hyperplane U toward $\alpha(t)(U)$, except for $U = H_h$, in which case $\alpha'(t)(H_h) = \mathbf{X} - \alpha(t)(H_h)$. This orientation is consistent and canonical for each t , and the map $t \mapsto \alpha'(t)$ yields a geodesic ray $\alpha' \rightarrow N(H_h)$ emanating from the penultimate 0-cube of R . Concatenating α' with the first $|R| - 1$ 1-cubes of R yields a geodesic $\gamma''_h \rightarrow \mathbf{X}$ such that $\gamma''_h(t)$ is separated from $\gamma(t + 1)$ by at most one hyperplane, namely H_h , for each $t \geq 0$. The claim now follows by induction on $|\mathcal{H}|$. Note that by

our previous claim, for all $r \geq 0$, $\mathbf{div}(\gamma, \gamma'')(r) \leq 2|\mathcal{H}|$, since $\mathcal{W}(\gamma'') \subseteq \mathcal{W}(\gamma)$. On the other hand, $\mathcal{W}(\gamma'') \subseteq \mathcal{W}(\gamma')$ by construction.

Suppose that v, v' belong to a common simplex $v'' \subseteq \partial_{\Delta} \mathbf{X}$. If $v'' = v$ or $v'' = v'$, then the claim follows from the first assertion of the lemma. There are two remaining possibilities: first, we could have that v, v' are distinct 0-simplices belonging to a common 1-simplex v'' . Second, if $v \not\subseteq v'$ and $v' \not\subseteq v$, and v is not a 0-simplex, then there exists a simplex $u = v \cap v'$ that is properly contained in both v, v' .

Let $\mathcal{V}, \mathcal{V}'$ be as before. Then $\mathcal{V} = \mathcal{V}_0 \sqcup \mathcal{H}$ and $\mathcal{V}' = \mathcal{V}'_0 \sqcup \mathcal{H}'$, where $\mathcal{H}, \mathcal{H}'$ are finite, and $\mathcal{V}_0, \mathcal{V}'_0$ satisfy the following conditions:

- (1) if $U \in \mathcal{V}_0$, then either U crosses both γ and γ' , or U crosses every element of \mathcal{V}'_0 ;
- (2) if $U' \in \mathcal{V}'_0$, then either U' crosses both γ and γ' , or U' crosses every element of \mathcal{V}_0 .

In either case ($v \cap v' = u$ or $v \cap v' = \emptyset$), the proof of assertion (1) shows, mutatis mutandis, that $\mathbf{div}(\gamma, \gamma')(r) \leq d_{\mathbf{X}}(\gamma(r), \gamma'(r)) + 2(|\mathcal{H}| + |\mathcal{H}'|)$ for all $r \geq 0$. □

Lemma 6.4 enables us to bound the $\mathbf{div}(\gamma, \gamma')(r)$ using the simplicial Tits distance, as follows.

Lemma 6.5 *There exists $B \in \mathbb{R}$ such that $\mathbf{div}(\gamma, \gamma')(r) \leq 2[\eta(v, v') + 1]r + B$ for all $r \geq 0$.*

Proof If $\eta(v, v') = \infty$, then $A = \eta(v, v')$ suffices, so suppose $\eta(v, v') = N < \infty$. The claim holds for $N = 0$, with $A = 1$, by **Lemma 6.4**. More generally, if v and v' lie in a common simplex, then the claimed inequality holds, for some B and $A = 1$, by **Lemma 6.4**.

Suppose that the claim holds for some $N - 1 \geq 0$. Let γ, γ' satisfy $\eta(\gamma, \gamma') = N$. Then there exist simplices $u_0, u_N \subset \partial_{\Delta} \mathbf{X}$ such that $v = u_0, v' = u_N$, and there is a minimal simplex-path $P = u_0, u_1, \dots, u_N$ in $\partial_{\Delta} \mathbf{X}$ joining u_0 to u_N . Let γ'' be a geodesic ray emanating from x_o and representing $u_N \cap u_{N-1}$. By induction, $\mathbf{div}(\gamma, \gamma'')(r) \leq 2Ar + B_1$ for all r and some fixed A depending on N , since $\eta(v, u_N \cap u_{N-1}) = N - 1$. Since $u_N \cap u_{N-1}$ and v' lie in a common simplex, namely u_N , then by **Lemma 6.4**, we have that $\mathbf{div}(\gamma', \gamma'')(r) \leq 2r + B_1$, so that $\mathbf{div}(\gamma, \gamma')(r) \leq 2(A + 1)r + B_1 + B_2$. Choosing $A = 1$ suffices for $N = 0$, so we may take $A = N + 1$ and $B = B_1 + B_2$, and we may conclude that we have $\mathbf{div}(\gamma, \gamma')(r) \leq 2r(\eta(\gamma, \gamma') + 1) + B$. □

Let γ, γ' be geodesic rays with $\gamma(0) = \gamma'(0) = x_0$ and for each $r \geq 0$, let $x_r = \gamma(r)$ and $x'_r = \gamma'(r)$. Let S_r be a shortest r -avoiding path joining x_r to x'_r . Then S_r decomposes as a “minimal” concatenation of geodesic segments as follows. Let P_1^r be a subpath of S_r beginning at x_r that is a geodesic segment of \mathbf{X} and let P_1^r have maximal length among such paths. If P_1^r is properly contained in S_r , then let P_2^r be a maximal geodesic subpath of $S_r - \text{Int}(P_1^r)$ beginning at the terminal 0-cube of P_1^r . In this way, we can write $S_r = P_1^r P_2^r \cdots P_{c_r}^r$, where each P_i^r is a geodesic segment and the hyperplane dual to the initial 1-cube of P_i^r crosses a 1-cube of P_{i-1}^r . This is a *fan decomposition* of S_r , and, if $c_r = 1$ for some fan decomposition of S_r , for all $r \geq 0$, then the simplices v, v' of $\partial_\Delta \mathbf{X}$ represented by γ, γ' either have nonempty intersection or lie in a common simplex.

Indeed, let $C_r = \gamma([0, r])$ and let $C'_r = \gamma'([0, r])$. Suppose that S_r is a geodesic segment for each $r \geq 0$, and suppose that $v \cap v' = \emptyset$. Let $\mathcal{U} = \mathcal{W}(\gamma) \cap \mathcal{W}(\gamma')$. Now \mathcal{U} is finite, since $v \cap v' = \emptyset$ (see [Theorem 3.24](#)). Then every hyperplane separating $\gamma(r)$ from x_0 either belongs to \mathcal{U} or separates $\gamma'(r)$ from $\gamma(r)$. Analogously, every hyperplane separating $\gamma'(r)$ from x_0 separates $\gamma(r)$ from $\gamma'(r)$ or belongs to \mathcal{U} . Every hyperplane separating $\gamma(r)$ from $\gamma'(r)$ separates either $\gamma(r)$ or $\gamma'(r)$ from x_0 . Hence $|P_1^r| = \text{div}(\gamma, \gamma')(r) = 2r - 2|\mathcal{U}|$ for all $r \geq 0$. Applying [Lemma 3.22](#) to γ and γ' , and then truncating the common initial segment of the resulting rays, we can assume that $\mathcal{U} = \emptyset$ and $|P_1^r| = r$.

Now let $P_r = QQ'$, where Q is a geodesic segment of length r emanating from x_r and Q' is a geodesic segment of length r terminating at x'_r . Let $D \rightarrow \mathbf{X}$ be a minimal-area disc diagram bounded by the closed path $C_r QQ'(C'_r)^{-1}$. Without affecting the points x_r, x_0, x'_r or the sets $\mathcal{W}_r, \mathcal{W}'_r$, we can modify D and the paths C_r, C'_r so that no two dual curves in D emanating from C_r (and no two dual curves emanating from C'_r) cross.

Since $\partial_p D$ is the concatenation of three geodesic segments, and any two dual curves have a total of four ends on $\partial_p D$, distinct dual curves in D must map to distinct hyperplanes. Thus $D \rightarrow \mathbf{X}$ is an isometric embedding and, since $\mathcal{U} = \emptyset$, each dual curve travels from QQ' to C_r or from QQ' to C'_r . Since every 0-cube of QQ' is separated from x_0 by at least r dual curves, it follows that each dual curve emanating from C'_r crosses each dual curve emanating from C_r , and hence that v and v' lie in a common simplex. More generally, we have the following.

Lemma 6.6 *Suppose that there exists $K \in \mathbb{R}$ such that for all $r \geq 0$, we have $|S_r| \leq d_{\mathbf{X}}(x_r, x'_r) + 2K$. Then, if v and v' are 0-simplices, they lie in a common simplex. If the above inequality holds and v is not a 0-simplex, then $v \cap v' \neq \emptyset$.*

Proof As before, suppose that $v \cap v' \neq \emptyset$ and that $\mathcal{W}(\gamma) \cap \mathcal{W}(\gamma') = \emptyset$. To achieve the latter requires that we apply Lemma 3.22 and truncate the common initial segment of the resulting rays. This does not affect v or v' , and modifies γ and γ' in only finitely many 1-cubes.

In case $K = 0$, the claim follows from the previous discussion. The hypothesis says that for all $r \geq 0$, there are at most K hyperplanes that cross S_r in more than one 1-cube. Suppose that v, v' do not belong to any common simplex. Let H_r be the hyperplane dual to the r -th 1-cube of γ , and define H'_r analogously for γ' . There exists $R, N \geq 0$ such that, for all $r \geq R$, the hyperplane H_r does not cross H'_N . Also, we can choose R such that, for all $r \geq 0$, there are at most R hyperplanes U that cross H_r and H'_r .

Each 0-cube $c \in S_r$ is separated from x_o by at least r hyperplanes. Let Q_r, Q'_r be the initial, length- r segments of γ, γ' respectively, emanating from x_o . Let D_r be a minimal-area disc diagram bounded by $Q_r, S_r, (Q'_r)^{-1}$. Since γ and γ' have no common dual hyperplanes, every dual curve in D_r travels from Q_r or Q'_r to S_r . Without modifying $x_o, \gamma(r), \gamma'(r)$ or v, v' , we may change Q_r, Q'_r so that no two dual curves in D_r emanating from Q_r cross, and no two dual curves emanating from Q'_r cross; see Figure 13. Let u_r be the number of dual curves in D_r that start and end on S_r , so that $|S_r| = 2r + 2u_r$.




Figure 13: The diagram D_r showing that either $\eta(v, v') \leq 1$ or $u_r \rightarrow \infty$.

Let $c \in S_r$ be the 0-cube in $S_r \cap N(H'_N)$ that is separated from $\gamma'(r)$ by each H'_i with $N \leq i \leq r$, as shown in Figure 13. Then there is a path in the image of D_r joining c to x_o that crosses only N of the dual curves mapping to hyperplanes H'_i : simply travel along the carrier of the dual curve K_N mapping to H_N , until arriving at $\gamma'(N - 1)$. Then travel along γ' to x_o . Hence every dual curve in D_r separating c from x_o maps to a hyperplane that crosses H_0 . Thus $u_r \geq r - R - N$ for all $r \geq 0$, and thus $|S_r| - d_X(\gamma(r), \gamma'(r))$ is unbounded as $r \rightarrow \infty$, a contradiction. We conclude that v and v' either intersect or lie in a common simplex. \square

Lemma 6.7 Suppose that there exist $M, N \in \mathbb{R}$ such that, for all $r \geq 0$, the rays γ, γ' satisfy $\mathbf{div}(\gamma, \gamma')(r) \leq 2Nr + M$. Then $N \geq \eta(v, v')$.

Proof First fix $r \geq 0$, and let S_r be a shortest r -avoiding path joining $\gamma(r)$ to $\gamma'(r)$, so that $|S_r| \leq 2Nr + M$. Let $k = \lceil \frac{M}{N} \rceil$. Let $S_r = P_1^r \cdots P_N^r$ be a concatenation of combinatorial paths such that each P_i^r has length at most $2r + k$, and $||P_i^r| - |P_j^r|| \leq 1$ for all i, j . In other words, subdivide S_r into N paths, each of length as close as possible to $\frac{|S_r|}{N}$.

For $1 \leq i \leq N$, let $f_i(r)$ be the terminal 0-cube of P_i^r , let $f_0(r) = \gamma(r)$, $f_N(r) = \gamma'(r)$. Let $Q_i(r)$ be a geodesic segment joining $f_i(r)$ to x_o , for $0 \leq i \leq n$, so that $Q_0(r), Q_N(r)$ are hyperplane-equivalent to the initial length- r segments of γ, γ' respectively.

Let $t_i(r)$ be the number of hyperplanes crossing both $Q_i(r)$ and $Q_{i+1}(r)$. Let $p_i(r)$ be the number of hyperplanes dual to at least two 1-cubes of P_i^r . Then we have $2r - t_i(r) + 2p_i(r) \leq 2r + k$, so that $2t_i(r) \geq 2p_i(r) - k$. By König's lemma, there are combinatorial geodesic rays γ_i , for $1 \leq i \leq N$, such that $\gamma_i(r) = f_i(r)$ for arbitrarily large values of r . If $t_i(r)$ is unbounded as $r \rightarrow \infty$, then γ_i, γ_{i+1} have infinitely many common hyperplanes, and thus the associated simplices $v_i, v_{i+1} \subset \partial_\Delta \mathbf{X}$ satisfy $\eta(v_i, v_{i+1}) \leq 1$. If $t_i(r)$ is uniformly bounded as $r \rightarrow \infty$, then so is $p_i(r)$, and thus $P_i(r)$ fails to be a geodesic segment by a uniformly bounded number of 1-cubes, whence we can take $\eta(v_i, v_{i+1}) \leq 1$ by Lemma 6.6. Hence there is a simplex path of length at most N joining v to v' . Indeed, if $t_i(r)$ is uniformly bounded, then either $v_i \cap v_{i+1}$ is nonempty, or v_i and v_{i+1} belong to a common simplex. \square

6.2.1 Linear divergence of rays \mathbf{X} has weakly uniformly linear divergence if there exists $A \geq 0$ such that, for all $x_0 \in \mathbf{X}^0$ and for all combinatorial geodesic rays γ, γ' with $\gamma(0) = \gamma'(0) = x_0$, we have $\mathbf{div}(\gamma, \gamma')(r) \leq Ar + B$ for all $r \geq 0$, where B (but not A) depends on γ and γ' .

Theorem 6.8 Let \mathbf{X} be strongly locally finite. Then $\partial_\Delta \mathbf{X}$ is bounded if and only if \mathbf{X} has weakly uniformly linear divergence.

Proof Suppose there exists $A < \infty$ such that $\eta(v, v') \leq A$ for all simplices v, v' of $\partial_\Delta \mathbf{X}$. By Lemma 6.5, if γ, γ' are rays emanating from x_0 , then we have that $\mathbf{div}(\gamma, \gamma')(r) \leq 2(A + 1)r + B$ for all $r \geq 0$. Conversely, suppose that for some $A \in \mathbb{R}$, and for each γ, γ' , there exists $B \in \mathbb{R}$ such that the above inequality holds for all $r \geq 0$. Then by Lemma 6.6, we have $\eta(v, v') \leq A + 1$ for all visible simplices v, v' .

Therefore, it suffices to show that $\eta(v, v')$ is uniformly bounded as v, v' vary over the set of visible simplices if and only if the graph $(\partial_\Delta \mathbf{X})^1$ has finite diameter. First, suppose

there exists $D < \infty$ such that, for all visible simplices v, v' , we have $\eta(v, v') \leq D$. Let u, u' be 0-simplices. Then, by [Theorem 3.19](#), there exist visible simplices v, v' with $u \subseteq v$ and $u' \subseteq v'$. Let $v = v_0, v_1, \dots, v_d = v'$ be a simplex-path joining v to v' , where $d \leq D$. For $0 \leq i \leq d - 1$, let w_i be a 0-simplex of $v_i \cap v_{i+1}$. Then w_i is adjacent to w_{i+1} for $0 \leq i \leq d - 1$. Moreover, w_0 and u lie in v_0 and are thus adjacent; similarly, w_{d-1} and u' are adjacent, so that u and u' are at distance at most $d - 1 + 2 \leq D + 1$ in $(\partial_\Delta \mathbf{X})^1$. On the other hand, suppose that $\text{diam}(\partial_\Delta \mathbf{X}) \leq D < \infty$. Let v, v' be simplices. Let $u \subseteq v, u' \subseteq v'$ be 0-simplices. Let $P: [0, d] \rightarrow (\partial_\Delta \mathbf{X})^1$ be a combinatorial path, where $P(0) = u$ and $P(d) = u'$, and $d \leq D$. For each i , let $s_i = [P(i), P(i + 1)]$ be a 1-simplex of P . Then $v, s_0, s_1, \dots, s_{d-1}, v'$ is a simplex-path of length $d + 2$ joining v to v' , whence $\eta(v, v') \leq \text{diam}(\partial_\Delta \mathbf{X}) + 2$. \square

6.2.2 Divergence of combinatorial geodesics Recall that $\overline{\text{div}}(a, b, c; \rho)$ is the divergence of the points $a, b \in \mathbf{X}^1$ with respect to the standard path metric, the basepoint c and the linear function ρ . It follows from the proof of [Theorem 6.8](#) that the bi-infinite combinatorial geodesic $\gamma: \mathbb{R} \rightarrow \mathbf{X}$ has linear divergence if and only if the simplices of $\partial_\Delta \mathbf{X}$ represented by the rays $\gamma((-\infty, 0])$ and $\gamma([0, \infty))$ belong to the same component of $\partial_\Delta \mathbf{X}$. Note that if, for some function $f: \mathbb{R}^+ \rightarrow \mathbb{R}^+$, the divergence $\overline{\text{div}}(\alpha(r), \alpha(-r), \alpha(0); \rho)$ of the combinatorial geodesic $\alpha: \mathbb{R} \rightarrow \mathbf{X}$ exceeds $f(r)$, then the divergence $\overline{\text{div}}(\mathbf{X}, \rho)(r)$ of $(\mathbf{X}, d_{\mathbf{X}})$ exceeds $f(r)$. The next observation is a consequence of the proof of [Theorem 6.8](#).

Corollary 6.9 *Let \mathbf{X} be strongly locally finite. Let $\alpha: \mathbb{R} \rightarrow \mathbf{X}$ be a bi-infinite combinatorial geodesic. Let v^-, v^+ be the simplices of $\partial_\Delta \mathbf{X}$ represented by the rays $\gamma_- = \alpha((-\infty, 0])$ and $\gamma_+ = \alpha([0, \infty))$. The simplices v^- and v^+ lie in different components of $\partial_\Delta \mathbf{X}$ if and only if the divergence of the geodesic α is superlinear, with respect to the combinatorial metric $d_{\mathbf{X}}$. In particular, if either γ_+ or γ_- is rank-one, then α has superlinear divergence.*

Proof By the proof of [Theorem 6.8](#), $\overline{\text{div}}(\gamma_-, \gamma_+)(r)$ is at most linear if and only if v^-, v^+ lie in the same component of $\partial_\Delta \mathbf{X}$. Hence, if v^- and v^+ lie in the same component of $\partial_\Delta \mathbf{X}$, then for all $r \geq 0$, there is an r -avoiding path P_r of length at most $Ar + B$ joining $\gamma_-(r)$ to $\gamma_+(r)$. Thus $\overline{\text{div}}(\alpha(r), \alpha(-r), \alpha(0); Cr - D) \leq Ar + B$ for any $C < 1, D \geq 0$ and for all $r \geq 0$, ie α has linear divergence with respect to the combinatorial metric.

Conversely, suppose that there exist $0 < C < 1, D \geq 0$ such that for all $r \geq 0$, the preceding inequality holds. Then $\gamma_+((D + r)/C)$ and $\gamma_-((D + r)/C)$ are joined by an r -avoiding path $P_r \rightarrow \mathbf{X}$, for each $r \geq 0$, with $|P_r| \leq Ar + B$. Let Q_r be the subpath of γ joining $\gamma_+(r)$ to $\gamma_+((D + r)/C)$ and let Q'_r be the subpath of γ_-

joining $\gamma_-((D+r)/C)$ to $\gamma_-(r)$. Then $Q_r P_r Q'_r$ is an r -avoiding path joining $\gamma_+(r)$ to $\gamma_-(r)$, so that

$$\mathbf{div}(\gamma_+, \gamma_-)(r) \leq |Q_r| + |Q'_r| + Ar + B = (2(C^{-1} - 1) + A)r + B + 2DC^{-1},$$

and v, v' lie in the same component of $\partial_\Delta \mathbf{X}$. \square

6.3 Divergence of cubulated groups

Let the group G act properly, essentially and cocompactly on the combinatorially geodesically complete CAT(0) cube complex \mathbf{X} . Denote by $\widetilde{\mathbf{div}}(G, \rho)(r)$ the divergence of G with respect to the word-metric arising from some finite generating set and with respect to the linear function ρ . The following is an immediate consequence of rank-rigidity and a result of Kapovich and Leeb [29, Proposition 3.3], and the fact that nontrivial product decompositions of \mathbf{X} correspond to nontrivial join decompositions of $\partial_\Delta \mathbf{X}$ and $\Gamma \mathbf{X}$.

Theorem 6.10 *Let G act properly, cocompactly and essentially on the combinatorially geodesically complete CAT(0) cube complex \mathbf{X} . If $\Gamma \mathbf{X}$ is bounded, then G has linear divergence if and only if $\Gamma \mathbf{X}$ decomposes as the join of two infinite proper subgraphs (ie $\partial_\Delta \mathbf{X}$ decomposes as the simplicial join of two proper subcomplexes). Otherwise, G has at least quadratic divergence.*

It also appears as though the techniques used to prove [Theorem 6.8](#) can be brought to bear on the question of when the divergence of G is at most quadratic.

References

- [1] **W Ballmann, M Gromov, V Schroeder**, *Manifolds of nonpositive curvature*, Progress in Mathematics 61, Birkhäuser, Boston, MA (1985) [MR823981](#)
- [2] **H-J Bandelt, V Chepoi**, *Metric graph theory and geometry: a survey*, from: “Surveys on discrete and computational geometry”, (J E Goodman, J Pach, R Pollack, editors), Contemp. Math. 453, Amer. Math. Soc., Providence, RI (2008) 49–86 [MR2405677](#)
- [3] **J Behrstock, R Charney**, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. 352 (2012) 339–356 [MR2874959](#)
- [4] **M R Bridson**, *Geodesics and curvature in metric simplicial complexes*, from: “Group theory from a geometrical viewpoint”, (É Ghys, A Haefliger, A Verjovsky, editors), World Sci. Publ., River Edge, NJ (1991) 373–463 [MR1170372](#)
- [5] **M R Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, Grundle. Math. Wissen. 319, Springer, Berlin (1999) [MR1744486](#)

- [6] **P-E Caprace, M Sageev**, *Rank rigidity for $CAT(0)$ cube complexes*, *Geom. Funct. Anal.* 21 (2011) 851–891 [MR2827012](#)
- [7] **R Charney, M W Davis**, *Finite $K(\pi, 1)$ s for Artin groups*, from: “Prospects in topology”, (F Quinn, editor), *Ann. of Math. Stud.* 138, Princeton Univ. Press (1995) 110–124 [MR1368655](#)
- [8] **R Charney, M W Davis**, *The $K(\pi, 1)$ -problem for hyperplane complements associated to infinite reflection groups*, *J. Amer. Math. Soc.* 8 (1995) 597–627 [MR1303028](#)
- [9] **I Chatterji, G Niblo**, *From wall spaces to $CAT(0)$ cube complexes*, *Internat. J. Algebra Comput.* 15 (2005) 875–885 [MR2197811](#)
- [10] **V Chepoi**, *Graphs of some $CAT(0)$ complexes*, *Adv. in Appl. Math.* 24 (2000) 125–179 [MR1748966](#)
- [11] **V Chepoi**, *Nice labeling problem for event structures: a counterexample*, *SIAM J. Comput.* 41 (2012) 715–727 [MR2974750](#)
- [12] **V Chepoi, M Hagen**, *On embeddings of $CAT(0)$ cube complexes into products of trees* [arXiv:1107.0863v1](#)
- [13] **C B Croke, B Kleiner**, *Spaces with nonpositive curvature and their ideal boundaries*, *Topology* 39 (2000) 549–556 [MR1746908](#)
- [14] **F Dahmani, V Guirardel, P Przytycki**, *Random groups do not split*, *Math. Ann.* 349 (2011) 657–673 [MR2755002](#)
- [15] **D Epstein, J-C Falmagne, S Ovchinnikov**, *Media theory*, Springer, Berlin (2008) [MR2368792](#)
- [16] **DS Farley**, *Finiteness and $CAT(0)$ properties of diagram groups*, *Topology* 42 (2003) 1065–1082 [MR1978047](#)
- [17] **DS Farley**, *Actions of picture groups on $CAT(0)$ cubical complexes*, *Geom. Dedicata* 110 (2005) 221–242 [MR2136028](#)
- [18] **SM Gersten**, *Divergence in 3-manifold groups*, *Geom. Funct. Anal.* 4 (1994) 633–647 [MR1302334](#)
- [19] **SM Gersten**, *Quadratic divergence of geodesics in $CAT(0)$ spaces*, *Geom. Funct. Anal.* 4 (1994) 37–51 [MR1254309](#)
- [20] **M Gromov**, *Hyperbolic groups*, from: “Essays in group theory”, (S M Gersten, editor), *Math. Sci. Res. Inst. Publ.* 8, Springer, New York (1987) 75–263 [MR919829](#)
- [21] **M Hagen**, *Weak hyperbolicity of cube complexes and quasi-arboreal groups* [arXiv:1101.5191v5](#)
- [22] **M Hagen**, *Geometry and combinatorics of cube complexes*, Ph.D. thesis, McGill University (2012)
- [23] **F Haglund**, *Isometries of $CAT(0)$ cube complexes are semi-simple* [arXiv:0705.3386](#)

- [24] **F Haglund, F Paulin**, *Simplicité de groupes d'automorphismes d'espaces à courbure négative*, from: “The Epstein birthday schrift”, (I Rivin, C Rourke, C Series, editors), Geom. Topol. Monogr. 1 (1998) 181–248 [MR1668359](#)
- [25] **C H Houghton**, *Ends of locally compact groups and their coset spaces*, J. Austral. Math. Soc. 17 (1974) 274–284 [MR0357679](#)
- [26] **C Hruska, D T Wise**, *Finiteness properties of cubulated groups*, Preprint (2010)
- [27] **W Imrich, S Klavžar**, *Product graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York (2000) [MR1788124](#)
- [28] **J R Isbell**, *Median algebra*, Trans. Amer. Math. Soc. 260 (1980) 319–362 [MR574784](#)
- [29] **M Kapovich, B Leeb**, *3-manifold groups and nonpositive curvature*, Geom. Funct. Anal. 8 (1998) 841–852 [MR1650098](#)
- [30] **I Leary**, *A metric Kan–Thurston theorem*, Preprint (2010)
- [31] **A Nevo, M Sageev**, *The Poisson boundary of CAT(0) cube complex groups* [arXiv: 1105.1675v1](#)
- [32] **G A Niblo, L D Reeves**, *Coxeter groups act on CAT(0) cube complexes*, J. Group Theory 6 (2003) 399–413 [MR1983376](#)
- [33] **B Nica**, *Cubulating spaces with walls*, Algebr. Geom. Topol. 4 (2004) 297–309 [MR2059193](#)
- [34] **Y Ollivier, D T Wise**, *Cubulating random groups at density less than 1/6*, Trans. Amer. Math. Soc. 363 (2011) 4701–4733 [MR2806688](#)
- [35] **MA Roller**, *Poc Sets, Median Algebras and Group Actions* (1998) Available at <http://www.personal.soton.ac.uk/gan/Roller.pdf>
- [36] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995) 585–617 [MR1347406](#)
- [37] **P Scott**, *Ends of pairs of groups*, J. Pure Appl. Algebra 11 (1977) 179–198 [MR487104](#)
- [38] **M L J van de Vel**, *Theory of convex structures*, North-Holland Mathematical Library 50, North-Holland Publishing Co., Amsterdam (1993) [MR1234493](#)
- [39] **D T Wise**, *Cubulating small cancellation groups*, Geom. Funct. Anal. 14 (2004) 150–214 [MR2053602](#)
- [40] **D T Wise**, *The structure of groups with a quasiconvex hierarchy*, Preprint (2011)

Department of Mathematics, University of Michigan
530 Church Street, 1859 East Hall, Ann Arbor, MI 48109, USA

markfhagen@gmail.com

<http://www-personal.umich.edu/~mfhagen/index.html>

Received: 4 January 2012 Revised: 1 May 2012