LARGE FACING TUPLES AND A STRENGTHENED SECTOR LEMMA

MARK HAGEN

ABSTRACT. We prove a strengthened sector lemma for irreducible, finite-dimensional, locally finite, essential, cocompact CAT(0) cube complexes under the additional hypothesis that the complex is *hyperplane-essential*; we prove that every quarterspace contains a halfspace. In aid of this, we present simplified proofs of known results about loxodromic isometries of the contact graph, avoiding the use of disc diagrams.

This paper has an expository element; in particular, we collect results about cube complexes proved by combining Ramsey's theorem and Dilworth's theorem. We illustrate the use of these tricks with a discussion of the Tits alternative for cubical groups, and ask some questions about "quantifying" statements related to rank-rigidity and the Tits alternative.

CONTENTS

1. INTRODUCTION

In various guises, CAT(0) cube complexes appear throughout mathematics. They appear in discrete mathematics as *median graphs* (the equivalence of median graphs and CAT(0) cube complexes is due to Chepoi [Che00]) and many other equivalent combinatorial structures: discrete median algebras [Ava61, Rol98], event structures [NPW81, BC93], etc. (see Bandelt-Chepoi [BC08] for a survey). After being introduced into group theory by Gromov as a source of examples [Gro87], CAT(0) cube complexes were understood by Sageev to provide the correct generalisation of trees needed to formulate a "high-dimensional Bass-Serre theory" [Sag95, Sag97].

The theory of groups acting on CAT(0) cube complexes has since proved extremely useful. Nonpositively-curved cube complexes provide the setting for Wise's cubical small-cancellation theory [Wis20], and the sub-class of *special* cube complexes defined by Haglund-Wise [HW08] provides a class of groups with many separable subgroups. These ideas were crucial to the resolution of several conjectures about 3–manifolds, notably the virtual Haken and virtual fibering conjectures [AGM13].

CAT(0) cube complexes are extremely organised spaces in which one has many tools far beyond CAT(0) geometry, largely because of the median structure, the *hyperplanes*, and the *(combinatorially) convex subcomplexes*. This has strong coarse-geometric consequences, e.g. finite

asymptotic dimension [Wri12] and various results related to quasi-isometric rigidity, e.g. [Hua17, HK18]. The nice geometric features of CAT(0) cube complexes have also led to the study of non-cubical spaces that can be "approximated" by cube complexes in one way or another, as in coarse median spaces [Bow13, Bow18, Bow19] and hierarchically hyperbolic spaces [BHS17b].

The purpose of this paper is threefold. First, we prove a statement, Proposition 1, about hyperplane-essential actions on CAT(0) cube complexes, needed elsewhere in the literature. In [CS11, Lemma 5.2], Caprace and Sageev show that, given an $\mathrm{Aut}(X)$–*essential*, irreducible cube complex X on which $\mathrm{Aut}(X)$ acts without a global fixed point at infinity, and given crossing hyperplanes v, h, there exist disjoint hyperplanes a, b such that a and b are separated by both v and h. This is crucial for their proof of rank-rigidity for CAT(0) cube complexes.

Simple examples also show that their lemma is sharp. Under a stronger hypothesis, *hyperplane-essentiality*, we get more (albeit *using* rank-rigidity):

Proposition 1. *Let X be an irreducible, locally finite, essential, hyperplane-essential CAT(0) cube complex such that $\mathrm{Aut}(X)$ acts cocompactly. Let h, v be distinct hyperplanes, let h^+, v^+ be halfspaces associated to h, v respectively, and suppose that $h \cap v \neq \emptyset$. Then $h^+ \cap v^+$ contains a hyperplane, and therefore contains a halfspace.*

A similar statement appears in [NS13]. The exact statement is [BF19, Proposition 2.11]. In the latter, the proof is attributed to still-in-progress work of the present author and Wilton [HW20]. So (with our collaborator's blessing), we extracted the proposition and its proof from [HW20] so that an account of Proposition 1 is readily available. The proof is in Section 5.

This seemed important to do because the results of [BF19] are significant and use the above proposition. In [BF19], conditions are given under which a cubulation of a group G is determined up to equivariant cubical isomorphism by function that assigns to each element of G its ℓ_1 translation length. In [BF19], the restrictions on the cube complex include the hypothesis that it has no free faces, and the same result holds when G is hyperbolic under the weaker hyperplane-essentiality hypothesis [BF18].

An essential action on a CAT(0) cube complex is the higher-dimensional version of a minimal action on a simplicial tree. Hyperplane-essentiality requires, in addition, that hyperplane-stabilisers act essentially on their hyperplanes. It always holds in the 1–dimensional case — hyperplanes in trees are points. A simple example is that the action of $\mathbb{Z}$ on the tiling of $\mathbb{R}$ by 1–cubes is hyperplane-essential, but the action of $\mathbb{Z}$ on the CAT(0) cube complex obtained by gluing squares corner-to-corner in the obvious way is not. See Figure 1
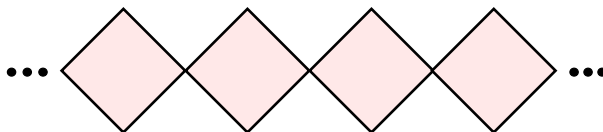

FIGURE 1. The $\mathbb{Z}$–action on this 2–dimensional CAT(0) cube complex by translations is essential but not hyperplane-essential. Indeed, the hyperplanes are nontrivial CAT(0) cube complexes — line segments — with trivial stabilisers, so the action is not hyperplane-essential. But every halfspace contains points in any fixed $\mathbb{Z}$–orbit arbitrarily far from its bounding hyperplane, so the action is essential.

The combination of essentiality and hyperplane-essentiality of a cocompact CAT(0) cube complex X can be viewed as a weak version of X having no free faces. The no free-faces property implies essentiality of hyperplanes of all codimensions.

Hyperplane-essentiality can often be arranged by modifying the cube complex [HT19], so it is a natural hypothesis to impose. Many motivating examples of actions on CAT(0) cube

complexes, like the cubulations of hyperbolic 3–manifold groups constructed by Bergeron-Wise using work of Kahn-Markovic [BW12, KM12], are hyperplane-essential [BF18]. In addition to its necessity for length spectrum rigidity, hyperplane-essentiality has recently proved useful in other contexts, e.g. [FH19]. Proposition 1 further illustrates the utility of the notion.

The proof of Proposition 1 in this paper relies (in a fairly soft way) on understanding which isometries of a locally finite CAT(0) cube complex X act loxodromically on the *contact graph*, the 1–skeleton of the nerve of the covering of X by hyperplane carriers. Such isometries were characterised in [Hag13]: a hyperbolic isometry g fails to be loxodromic on the contact graph when g is either not rank-one, or some axis of g fellow-travels a hyperplane. The proof in [Hag13] relies on *disc diagrams* in CAT(0) cube complexes, which were introduced by Casson in unpublished notes, and developed in work of Sageev [Sag95] and Wise [Wis20].

This brings us to the second purpose of this paper, which is expository. The study of rank-one/contracting isometries of CAT(0) spaces/cube complexes is a subject of current interest, see e.g. [CS15, QRT19, CH17]. So, in this paper we give a simpler proof of the preceding characterisation of contact graph loxodromics, avoiding the use of disc diagrams. This is Theorem 4.1.

One of the ingredients in the proof of Theorem 4.1 is a description of the convex hull of a combinatorial geodesic axis of an isometry (Lemma 4.8). Convex hulls of geodesics are examples of cube complexes in which no three hyperplanes *face*, i.e. for any three disjoint hyperplanes, one of them separates the other two. A naive question is: given a finite set of hyperplanes with no such "facing triple", is there a geodesic that crosses all hyperplanes in the set? Conversely, given a finite set of hyperplanes that cross a ball of a fixed size, can we find large subsets in which any three hyperplanes face?

The answer to the first question is no — the set of hyperplanes in the CAT(0) cube complex formed by arranging five squares cyclically around a common vertex (so that the cube complex is homeomorphic to a disc) is a counterexample. However, when the dimension of the cube complex is bounded, there is a geodesic that crosses *a definite proportion* of the hyperplanes; this is Corollary 3.4 below. A quantitative version of the second question also has a positive answer; see Corollary 3.9, which says that, if X is a D–dimensional CAT(0) cube complex that does not isometrically embed in the standard tiling of $\mathbb{E}^L$ by L–cubes for any L, then for all N, there exists R such that the set of hyperplanes intersecting an R–ball in X contains a facing N–tuple of hyperplanes.

These two results are proved by combining Ramsey's theorem and Dilworth's theorem. The idea to apply Dilworth's theorem to a collection of halfspaces appears throughout the literature (see e.g. [AOS12, BCG⁺09, Fio17]). The idea to apply Ramsey's theorem to a collection of hyperplanes appears in [CS11, CC19] and likely elsewhere. We are not aware of a reference where the two are applied in concert in this way, so decided to make matters explicit here.

The main statement on the above topic is Proposition 3.3, which says that if X is a D–dimensional CAT(0) cube complex, and $N \in \mathbb{N}$, and $\mathcal{W}$ is a finite set of hyperplanes with no $(N+1)$–tuple of facing hyperplanes, then there is a subset of $\mathcal{W}$ of size at least $|\mathcal{W}|/K$ such that we can choose one halfspace for each hyperplane in the subset in such a way that the associated halfspaces are totally ordered by inclusion. Here, K is a constant depending on D and N (which can be made explicit using Ramsey numbers).

The motivation for Corollary 3.4 is a question from Abdul Zalloum; the statement seems to be useful in current work of Murray-Qing-Zalloum on *sublinearly contracting boundaries*. The purpose of Corollary 3.9 is to illustrate the idea of Proposition 3.3, which we do by giving a simple proof of the Tits alternative for cubulated groups, in Proposition 3.10.

The proof of Proposition 3.10 is different from that of the more general statement due to Sageev-Wise [SW05], and more closely resembles the proof in [CS11]. In both cases, the main point is to find a facing 4–tuple of hyperplanes, divide this into two pairs, and apply the Double Skewering Lemma to find two hyperbolic isometries which, by ping-pong, generate a free group.

What we find intriguing is that a facing 4–tuple, if it exists (i.e. if the group in question is not virtually abelian), must be seen in a ball in the cube complex of quantifiable radius, because we found it using Corollary 3.9. In Section 6, we pose some questions aimed at effectivising various statements about actions on cube complexes. This is the third purpose of this paper.

2. Preliminaries

There are several introductions to CAT(0) cube complexes, emphasising different features; see e.g. [Che00, Hag08, Rol98, Sag14, Wis12]. We follow [FH19, Section 2].

Definition 2.1 (CAT(0) cube complex)**.** A *cube* is a Euclidean unit cube $[-\frac{1}{2}, \frac{1}{2}]^n$ for some $n \geq 0$, and a *face* of a cube c is a subcomplex obtained by restricting some of the coordinates to $\pm\frac{1}{2}$. A *midcube* of c is a subspace obtained by restricting exactly one coordinate to 0.

A *CAT(0) cube complex* is a simply connected CW complex X whose cells are cubes, where the attaching maps restrict to isometries on faces, and the following holds. For each 0–cube $v \in X$, and each collection $e_1, \ldots, e_k$ of 1–cubes incident to v, if for all $i \neq j$ the 1–cubes e_i, e_j span a 2–cube, then $e_1, \ldots, e_k$ span a unique k–cube. The *dimension* $\dim X$ is the supremum of the dimensions of the cubes of X.

We use the terms "vertex"/"0–cube", and the terms "edge"/"1–cube", interchangeably when talking about cube complexes.

By [Gro87, Bri91, Lea13], a CAT(0) cube complex X supports a CAT(0) metric d_2 in which each (Euclidean) cube is convex. We will instead work mainly with the path metric d obtained by equipping each cube with the ℓ_1 metric; see e.g [Mie14]. In fact, we mostly work with the restriction of d to $X^{(0)}$, which is isometric to the metric obtained by restricting to $X^{(0)}$ the usual graph metric on $X^{(1)}$.

We need the following language about paths in X. A *CAT(0) geodesic* is a geodesic for the metric d_2. Given $L \in \mathbb{Z}_{\geq 0}$, a *combinatorial path* $\gamma : [0, L] \to X$ is a continuous map sending $[0, L] \cap \mathbb{Z}$ to $X^{(0)}$ and sending each $[i, i+1]$ isometrically to a 1–cube. The combinatorial path γ is a *combinatorial geodesic* if $|i - j| = \mathsf{d}(\gamma(i), \gamma(j))$ for $0 \leq i \leq j \leq L$.

2.1. Hyperplanes and halfspaces.

The key features of CAT(0) cube complexes are their *hyperplanes* and *halfspaces*. There are different viewpoints on CAT(0) cube complexes, one emphasising hyperplanes and one emphasising halfspaces. The ubiquity of CAT(0) cube complexes comes from the fact that they are "geometric realisations" of very simple combinatorial data; if one finds the hyperplane viewpoint more natural, one thinks of CAT(0) cube complexes as dual to *wallspaces*, as explained in [Nic04, CN05, HW14]; if one favours halfspaces, one can think of CAT(0) cube complexes as dual to *pocsets*, as explained in [Sag14]. The viewpoints are equivalent, and also equivalent to *discrete median algebras* [Rol98] and *median graphs* [Che00]. All of these notions are intimately related, and there are different situations making each viewpoint optimal.

Definition 2.2 (Hyperplane). Let X be a CAT(0) cube complex. A *hyperplane* is a connected subspace $h \subset X$ such that for each cube c of X, the intersection $h \cap c$ is either empty or a midcube of c. The *carrier* $\mathcal{N}(h)$ is the union of all (closed) cubes intersecting h.

Each midcube in X is contained in exactly one hyperplane.

By e.g. [Sag95], if h is a hyperplane, then h is again a CAT(0) cube complex whose cubes are midcubes of the form $h \cap c$, where c is a cube of X intersecting h. Moreover, $\mathcal{N}(h)$ is a CAT(0) cube complex isomorphic to $h \times [-\frac{1}{2}, \frac{1}{2}]$. A 1–cube e with $h \cap e \neq \emptyset$ is *dual to* h. The 0–skeleton of h (regarded as a CAT(0) cube complex) is the set of midpoints of 1–cubes dual to h. The hyperplanes of h are subspaces of the form $a \cap h$, as a varies over the hyperplanes of X that intersect h.

For each hyperplane h, the complement $X - h$ has exactly two components, called *halfspaces associated* to h. We usually denote these $\overleftarrow{h}$ and $\overrightarrow{h}$. If $\overleftarrow{h}$ is a halfspace, there is a unique hyperplane h such that $\overleftarrow{h}$ is a component of $X - h$, and we also say h is the hyperplane *associated* to $\overleftarrow{h}$.

Hyperplanes are not subcomplexes of X, but the following construction is often convenient: let X' be the CAT(0) cube complex obtained from X by subdividing each n–cube into 2^n n–cubes in the obvious way, by declaring the barycentre of each cube to be a 0–cube. Then the hyperplanes of X are now subcomplexes; they are no longer hyperplanes in the new cube complex X', which has two "parallel" copies of each of the original hyperplanes.

Definition 2.3 (Separation, crossing, parallelism). Given a hyperplane h and $A, B \subset X$, we say that h *separates* A, B if there are distinct halfspaces $\overleftarrow{h}, \overrightarrow{h}$ associated to h with $A \subset \overleftarrow{h}$ and $B \subset \overrightarrow{h}$. We are usually interested in the case where A, B are vertices, hyperplanes, or convex subcomplexes (see below).

If $A \subset X$ and the hyperplane h separates two points of A, we say that h *crosses* A.

The hyperplanes h, v cross if and only if h, v are distinct and have nonempty intersection. Equivalently, each halfspace associated to h intersects each halfspace associated to v.

The subcomplexes A, B are *parallel* if any hyperplane h crosses A if and only if h crosses B. Taking A, B to be 1–cubes gives an equivalence relation on 1–cubes in which 1–cubes are equivalent if they are dual to the same hyperplane (i.e. parallel). So, there is a bijection between hyperplanes and parallelism classes of 1–cubes.

Crucially, if γ is an embedded combinatorial path in X, then γ is a geodesic if and only if γ contains at most one 1–cube dual to each hyperplane. In particular, if x, y are vertices, then $\mathsf{d}(x, y)$ is the number of hyperplanes separating x from y.

Remark 2.4 (Helly property for hyperplanes). Let $h_1, \ldots, h_n$ be hyperplanes that pairwise cross. Then $\bigcap_{i=1}^n h_i \neq \emptyset$, i.e. there is a (not necessarily unique) n–cube whose barycentre is contained in each h_i. So $\dim X$ is the maximum possible cardinality of a set of pairwise crossing hyperplanes.

The same *Helly property* — any finite collection of pairwise intersecting subsets in a given class has nonempty total intersection — also holds for the class of convex subcomplexes, discussed presently.

2.2. **Medians, convexity, geodesics, and gates.** In [Che00], Chepoi established a correspondence between *median graphs* and CAT(0) cube complexes. Let Γ be a connected graph, with graph metric ρ. The *interval* $[a, b]$ between vertices a, b is the set of vertices c with $\rho(a, b) = \rho(a, c) + \rho(b, c)$. The graph Γ is *median* if for all $a, b, c \in \Gamma^{(0)}$, there is a unique vertex $\mu(a, b, c)$ with $\mu(a, b, c) \in [a, b] \cap [b, c] \cap [a, c]$. Chepoi's theorem says that the 1–skeleton of any CAT(0) cube complex is a median graph, and each median graph is the 1–skeleton of a uniquely determined CAT(0) cube complex.

Given a CAT(0) cube complex X, we let $\mu : (X^{(0)})^3 \to X^{(0)}$ be the median operator described above. There is a way to extend μ over the whole of X, but we will not need it here.

The subcomplex Y of X is *full* if Y contains every cube of X whose 0–skeleton appears in Y. The full subcomplex Y is *convex* if for all vertices $x, y \in Y$ and $z \in X$, we have $\mu(x, y, z) \in Y$. If x, y are vertices, then the median interval $[x, y]$ is convex and consists of the union of combinatorial geodesics from x to y. If Y is a convex subcomplex, then $[x, y] \subset Y$ for all $x, y \in Y^{(0)}$, and conversely.

If $\overleftarrow{h}$ is a halfspace, then the smallest subcomplex containing $\overleftarrow{h}$ is convex. If h is a hyperplane, then the carrier $\mathcal{N}(h)$ is convex.

Given a subspace A of X, the *convex hull* of A is defined as follows. First, let A' be the intersection of all halfspaces containing A. Then the convex hull is the union of all cubes contained in A'. Convex hulls are convex; the convex hull of a pair of vertices x, y is exactly the median interval $[x, y]$.

The median viewpoint enables a very useful construction, the *gate map*. Let $Y \subset X$ be a convex subcomplex. Then there is a map $\mathfrak{g} = \mathfrak{g}_Y : X \to Y$ with the following properties (see e.g. [BHS17a, Section 3]:

- $\mathfrak{g}$ is 1–lipschitz (for both d and d_2);
- if $x \in X^{(0)}$ and h is a hyperplane, then h separates x from $\mathfrak{g}(x)$ if and only if h separates x from Y;
- $\mathsf{d}(x, \mathfrak{g}(x)) = \mathsf{d}(x, Y)$, and $\mathfrak{g}(x)$ is the unique closest vertex of Y to x.

If Y, Z are convex subcomplexes, then $\mathfrak{g}_Y(Z)$ and $\mathfrak{g}_Z(Y)$ are isomorphic CAT(0) cube complexes, and a hyperplane h crosses $\mathfrak{g}_Y(Z)$ if and only if h crosses both Y and Z. Moreover, there is a convex subcomplex $\mathfrak{g}_Y(Z) \times I$ of X, where the hyperplanes crossing I are precisely those that separate Y from Z. In fact, $\mathfrak{g}_Y(Z) \times I$ is the convex hull of $\mathfrak{g}_Z(Y) \cup \mathfrak{g}_Y(Z)$. We will use this in the proof of Theorem 4.1, when we note that $\mathrm{diam}(\mathfrak{g}_Y(Z)) = \mathrm{diam}(\mathfrak{g}_Z(Y))$. (See e.g. [BHS17a, Lemma 2.6].)

The following lemma is standard, follows easily from the above itemised properties (specifically the second) and we will use it later. It appears in various places in the literature; see e.g. [Gen20, Proposition 2.6].

Lemma 2.5. *Let X be a CAT(0) cube complex and Y a convex subcomplex. Let $\mathfrak{g} : X \to Y$ be the gate map. Let $x, y \in X$ be 0–cubes. Then for all hyperplanes h, we have that h separates $\mathfrak{g}(x), \mathfrak{g}(y)$ if and only if h both intersects Y and separates x, y.*

Although a hyperplane h is not a subcomplex, we saw above that h becomes a subcomplex in the cubical subdivision X', and moreover it is convex. Accordingly, we also have a gate map $\mathfrak{g}_h : X \to h$. We will also use this in Theorem 4.1 and Proposition 1. More on gate maps to hyperplanes can be found in [FH19, Section 2.1.6].

Remark 2.6. If $Y \subset X$ is a subcomplex, then convexity of Y in the above sense follows from convexity of Y in the CAT(0) metric d_2 [Hag07]. However, this only works for subcomplexes. For general subspaces, convexity in the above sense implies CAT(0) convexity, but the converse might not hold: consider a diagonal line in the standard tiling of $\mathbb{E}^2$ by 2–cubes.

2.3. Isometries and skewering.

In this subsection, we mostly follow [Hag07] and [CS11].

By $\mathrm{Aut}(X)$, we mean the group of cubical automorphisms of the CAT(0) cube complex X. Automorphisms are isometries with respect to d and d_2, although (X, d_2) may have isometries that are not cubical (this is studied in [Bre17]).

The action $G \to \mathrm{Aut}(X)$ of the group G is *proper* if cube stabilisers are finite, and *metrically proper* if for all $x_0 \in X^{(0)}$ and all $R \geq 0$, the set of $g \in G$ with $\mathsf{d}(x_0, gx_0) \leq R$ is finite. If X is locally finite, then any proper action is metrically proper.

The action is *cocompact* if there is a compact subcomplex $K \subset X$ with $G \cdot K = X$.

The following is well-known and widely-used; see e.g. [FH19, Lemma 2.3] for a proof:

Lemma 2.7 (Hereditary cocompactness)**.** *Let X be a CAT(0) cube complex on which the group G acts cocompactly. Then for all hyperplanes h of X, the action of $\mathrm{Stab}_G(h)$ on h is cocompact.*

The cube complex X is *essential* if each halfspace $\overleftarrow{h}$, $\overrightarrow{h}$ contains points arbitrarily far from the associated hyperplane h. The action of G on X is *essential* if those points can all be chosen in a fixed G–orbit. When $\mathrm{Aut}(X)$ acts on X cocompactly, X is essential if and only if the action of $\mathrm{Aut}(X)$ is essential.

The cube complex X is *hyperplane-essential* if each hyperplane h of X is an essential CAT(0) cube complex, and $G \to \mathrm{Aut}(X)$ is a *hyperplane-essential* action if, for each hyperplane h, the action of $\mathrm{Stab}_G(h)$ on h is essential.

Given $g \in \mathrm{Aut}(X)$, we say that g is *combinatorially hyperbolic* if there is a combinatorial geodesic $\gamma : \mathbb{R} \to X$ and a positive integer ℓ such that $g\gamma(t) = \gamma(t + \ell)$ for all $t \in \mathbb{R}$, i.e. g acts on γ as a nontrivial translation. Such a γ is a *combinatorial axis* for g.

If g is combinatorially hyperbolic, then the set of hyperplanes h such that h intersects an axis of g is independent of the choice of axis. If h is a hyperplane, then g *skewers* h if $g\overleftarrow{h} \subset \overleftarrow{h}$, where $\overleftarrow{h}$ is one of the halfspaces associated to h.

Here is an exercise: if the convex hull A of the axis of g is finite-dimensional, then any hyperplane crossing A is skewered by some power of g.

A theorem of Haglund [Hag07] asserts that, under a mild assumption on X that can always be arranged by replacing X by its cubical subdivision, any $g \in \mathrm{Aut}(X)$ is either combinatorially hyperbolic or fixes a vertex. Even without subdividing, any g has a positive power that is either combinatorially hyperbolic or fixes a vertex, provided X is finite-dimensional.

Remark 2.8. Higher-dimensional versions of Haglund's combinatorial semisimplicity theorem, and related results, have been obtained by Woodhouse, Genevois, and Woodhouse-Wise [Gen19b, Woo17, WW17].

For the CAT(0) metric, one can deduce the following (see e.g. [BH99, Exercise II.6.6(2)]):

Lemma 2.9. *Let X be a finite-dimensional CAT(0) cube complex and let $g \in \mathrm{Aut}(X)$. Then either g fixes a point in X, or g is a hyperbolic isometry of the CAT(0) space (X, d_2).*

The finite dimensional hypothesis is necessary: there are infinite dimensional CAT(0) cube complexes with isometries that are combinatorially hyperbolic but CAT(0) parabolic [AKWW13].

Definition 2.10 (Rank one)**.** If $g \in \mathrm{Aut}(X)$ is hyperbolic for the CAT(0) metric, we say that g is *not rank one* if some CAT(0) geodesic axis for g lies in an isometrically embedded Euclidean half-flat $[0, \infty) \times \mathbb{R}$, and g is *rank one* otherwise.

The Double Skewering Lemma of Caprace-Sageev [CS11, p. 4] is a vital tool for identifying hyperbolic isometries of CAT(0) cube complexes, given an ambient essential action.

Lemma 2.11 (Double skewering)**.** *Let X be a finite-dimensional CAT(0) cube complex and let G act essentially on X. Suppose that one of the following holds:*

- *G acts with no fixed point in the visual boundary ∂X;*
- *X is locally finite and G acts cocompactly.*

Let h, v be disjoint hyperplanes, and let $\overleftarrow{h}$, $\overleftarrow{v}$ be halfspaces associated to h, v respectively, with $\overleftarrow{h} \subsetneq \overleftarrow{v}$. Then there exists a hyperbolic (in the combinatorial and CAT(0) metrics) element $g \in G$ such that $g\overleftarrow{v} \subsetneq \overleftarrow{h}$; in particular, h separates v from gv.

Proof. The "in particular" statement follows immediately from $g\overleftarrow{v} \subsetneq \overleftarrow{h}$. Second, it also follows immediately from $g\overleftarrow{v} \subsetneq \overleftarrow{h}$ that $\langle g \rangle$ has unbounded orbits in X. Hence, up to replacing g by

a positive power, g is combinatorially hyperbolic. Since X is finite-dimensional, the identity $(X, \mathsf{d}) \to (X, \mathsf{d}_2)$ is a G–equivariant quasi-isometry (by Lemma 3.6 below), so $\mathsf{d}_2(g^n x, x)$ grows linearly in n, for any $x \in X$, whence g is also hyperbolic in the CAT(0) metric.

So, it remains to find g with $g\overleftarrow{v} \subsetneq \overleftarrow{h}$. In the case where G acts without a fixed point at infinity, the desired statement appears on page 4 of [CS11].

Suppose G acts cocompactly on X and X is locally finite. In this case, Corollary 4.9 of [CS11] implies that $X = X_1 \times \cdots \times X_p \times Y$, where each X_i has compact hyperplanes, and the finite-index subgroup $G' \leq G$ preserving this decomposition does not fix a point in ∂Y.

If h, v are hyperplanes of the form $X_1 \times \cdots \times X_p \times \bar{h}, X_1 \times \cdots \times X_p \times \bar{v}$, where $\bar{h}, \bar{v}$ are hyperplanes of Y, then the claim follows from the version for actions without fixed points at infinity. So, it suffices to prove the claim in the case where hyperplanes of X are compact, X is locally finite, and G acts cocompactly and essentially; we leave this as an exercise. $\qquad\square$

2.4. The contact graph.

Let X be a CAT(0) cube complex and let $\mathcal{W}$ be the set of hyperplanes. Then $\{\mathcal{N}(h) : h \in \mathcal{W}\}$ is a covering of X, and we define the *contact graph* $\mathcal{C}X$ to be the (necessarily connected) 1–skeleton of the nerve of this covering, i.e. the intersection graph of the hyperplane carriers. This graph was initially defined in [Hag14]; it has a vertex for each hyperplane, with two vertices adjacent provided no third hyperplane separates the corresponding hyperplanes. By [Hag14, Theorem 4.1], $\mathcal{C}X$, equipped with its usual graph metric, is quasi-isometric to a tree (with constants independent of X). In particular, $\mathcal{C}X$ is hyperbolic.

Throughout the paper, if h is a hyperplane of X, we also use the letter h to mean the corresponding vertex of $\mathcal{C}X$.

We now define a coarse map $\pi : X \to 2^{\mathcal{C}X}$ by first defining π on vertices of X, and then on higher-dimensional open cubes.

Given $x \in X^{(0)}$, the set of hyperplanes h with $x \in \mathcal{N}(h)$ corresponds to a complete subgraph of $\mathcal{C}X$, which we denote $\pi(x)$. Hence $\pi : X^{(0)} \to 2^{\mathcal{C}X}$ is a coarse map. For each open edge e of X, we define $\pi(e)$ to be the vertex corresponding to the hyperplane dual to e. Hence we have a coarse map $\pi : X^{(1)} \to 2^{\mathcal{C}X}$. More generally, if c is an open n–cube, $n \geq 1$, let $\pi(c)$ be the complete subgraph with vertex set the hyperplanes intersecting c.

So, $\pi : X \to 2^{\mathcal{C}X}$ sends cubes to bounded sets, and, if f, c are open cubes with $\bar{f} \subset \bar{c}$, then $\pi(f) \subset \pi(c)$. If x is a vertex of $\bar{f}$, then $\pi(x) \cap \pi(f) \neq \emptyset$. (Here, $\bar{f}, \bar{c}$ denote the closures of f, c.) In particular, $\pi : X^{(1)} \to 2^{\mathcal{C}X}$ is coarsely lipschitz.

Since $\mathrm{Aut}(X)$ acts on the set of hyperplanes of X, preserving intersection and non-intersection of carriers, we get a homomorphism $\mathrm{Aut}(X) \to \mathrm{Aut}(\mathcal{C}X)$; the relationship between these two groups is studied in [Fio20].

With respect to this action of $\mathrm{Aut}(X)$ on $\mathcal{C}X$, the map π is equivariant, i.e. $\pi(gx) = g\pi(x)$ for all $x \in X, g \in \mathrm{Aut}(X)$.

We will chiefly be interested in when an element $g \in \mathrm{Aut}(X)$ acts on $\mathcal{C}X$ *loxodromically*, i.e. when the map $\mathbb{Z} \to \mathbb{Z}$ given by $n \mapsto \mathsf{d}_{\mathcal{C}X}(h, g^n h)$ is bounded above and below by increasing linear functions of n (the functions depend on h, but the existence of such functions for some hyperplane implies it for any other hyperplane).

Since π is equivariant and coarsely lipschitz, to prove that a hyperbolic isometry $g \in \mathrm{Aut}(X)$ is loxodromic, it suffices to prove the following: for any hyperplane h, there exists $C > 0$ such that $\mathsf{d}_{\mathcal{C}X}(h, g^n h) > Cn$ for all $n > 0$. A non-hyperbolic isometry can never be loxodromic on $\mathcal{C}X$, since it stabilises a clique $\pi(x)$, where $x \in X$ is a point fixed by g.

Finally, here is a simple observation that is used in the proof of Proposition 1:

Lemma 2.12. *Let X be a CAT(0) cube complex and let v, h be hyperplanes such that $\mathsf{d}_{\mathcal{C}X}(v, h) > 2$. Then $\mathfrak{g}_v(h)$ is a single point.*

Proof. Recall that, when v is regarded as a CAT(0) cube complex, $\mathfrak{g}_v(h)$ is a convex subcomplex of v whose hyperplanes have the form $a \cap v$, where a is a hyperplane of X that crosses both v and h. If $\mathsf{d}_{\mathcal{C}X}(v, h) > 2$, there are no such hyperplanes a, so $\mathfrak{g}_v(h)$ is a CAT(0) cube complex with no hyperplanes and hence no positive-dimensional cubes, i.e. $\mathfrak{g}_v(h)$ is a point. $\qquad\square$

3. Facing tuples and chains with Dilworth and Ramsey

Fix a CAT(0) cube complex X. Recall that $\dim X$ is equal to the supremum of the cardinalities of sets of pairwise intersecting hyperplanes.

Definition 3.1 (Facing tuple). Let $n \in \mathbb{N} \cup \{\infty\}$. A *facing n–tuple* is a set of hyperplanes $\{h_1, \ldots, h_n\}$ with the property that, for each h_i, we can choose an associated halfspace $\overleftarrow{h}_i$ such that $\overleftarrow{h}_i \cap \overleftarrow{h}_j = \emptyset$ for $i \neq j$. Equivalently, there do not exist i, j, k such that h_i separates h_j from h_k.

Definition 3.2 (Chain). A *chain* in X of length n is a set $\{h_1, \ldots, h_n\}$ of hyperplanes such that h_i separates h_{i-1} from h_{i+1} for $2 \leq i \leq n - 1$.

Let $h_1, \ldots, h_n$ be a chain. For $2 \leq i \leq n$, let $\overleftarrow{h}_i$ be the halfspace associated to h_i and containing h_1, and let $\overleftarrow{h}_1$ be the halfspace associated to h_1 not containing h_2. Then $\overleftarrow{h}_1 \subsetneq \cdots \subsetneq \overleftarrow{h}_n$. Conversely, let $h_1, \ldots, h_n$ be hyperplanes for which we can choose a halfspace $\overleftarrow{h}_i$ associated to each h_i in such a way that $\overleftarrow{h}_1 \subsetneq \cdots \subsetneq \overleftarrow{h}_n$. Then $\{h_1, \ldots, h_n\}$ is a chain.

The following is a useful trick:

Proposition 3.3 (Chains and facing tuples). *For all $D, N \in \mathbb{N}$, there exists $K(D, N) \geq 1$ such that the following holds. Let X be a D–dimensional CAT(0) cube complex, and let $\mathcal{W}$ be a finite set of hyperplanes in X such that any facing tuple in $\mathcal{W}$ has cardinality at most N. Then $\mathcal{W}$ contains a chain of cardinality at least $|\mathcal{W}|/K(D, N)$.*

Proof. Let $\widehat{\mathcal{W}} = \{\overleftarrow{h} : h \in \mathcal{W}\}$ be a set of halfspaces with exactly one associated to each hyperplane in $\mathcal{W}$. The set $\widehat{\mathcal{W}}$ is partially ordered by inclusion. For each $h \in \mathcal{W}$, recall the notation $\overrightarrow{h} = X - \overleftarrow{h}$.

Bounding $\subseteq$–antichains with Ramsey's theorem: Let $\mathcal{A} \subset \widehat{\mathcal{W}}$ be a set of halfspaces, no two of which are $\subseteq$–comparable. So, for all $\overleftarrow{h}, \overleftarrow{v} \in \mathcal{A}$, exactly one of the following holds:

(1) The hyperplanes h, v are distinct and $h \cap v \neq \emptyset$.

(2) The hyperplanes h, v are disjoint. So, one of the following holds: $\overleftarrow{v} \cap \overleftarrow{h} = \emptyset$, or $\overrightarrow{v} \cap \overrightarrow{h} = \emptyset$.

Suppose that $\overleftarrow{h}_1, \ldots, \overleftarrow{h}_n \in \mathcal{A}$ are distinct halfspaces with the property that item (1) holds for $\overleftarrow{h}_i, \overleftarrow{h}_j$ for all $i \neq j$. Then the set $\{h_1, \ldots, h_n\}$ contains n distinct, pairwise-intersecting hyperplanes. Hence $n \leq D$.

Suppose that $\overleftarrow{h}_1, \ldots, \overleftarrow{h}_n \in \mathcal{A}$ are distinct halfspaces with the property that item (2) holds for $\overleftarrow{h}_i, \overleftarrow{h}_j$ for all $i \neq j$.

For all $i, j, k \leq n$, the hyperplane h_i cannot separate h_j from h_k. Indeed, suppose that this is the case. Then, up to relabelling, $h_j \subset \overleftarrow{h}_i$. Since $\overleftarrow{h}_i$ and $\overleftarrow{h}_j$ are $\subseteq$–incomparable, $h_i \subseteq \overleftarrow{h}_j$, i.e. $\overrightarrow{h}_i \cap \overrightarrow{h}_j = \emptyset$. But then either $\overleftarrow{h}_k$ contains h_i, h_j and hence contains $\overleftarrow{h}_i$, or $\overleftarrow{h}_k$ is contained in $\overleftarrow{h}_j$. Either of these situations violates pairwise-incomparability of the elements of $\mathcal{A}$.

We have just shown that $\{h_1, \ldots, h_n\}$ form a facing tuple. Hence $n \leq N$.

Let G be the complete graph with vertex set $\mathcal{A}$. We colour an edge blue if the corresponding pair of halfspaces satisfy item (1) and red if the halfspaces satisfy item (2). This colours all of the edges. We have shown that blue cliques have at most D vertices, and red cliques have at

most N vertices. So, by **Ramsey's theorem** [Ram29], $|\mathcal{A}| \leq \mathrm{Ram}(D+1, N+1) - 1$, where $\mathrm{Ram}(\bullet, \bullet)$ denotes the Ramsey number.

Applying Dilworth's theorem: Let $K(D, N) = \mathrm{Ram}(D+1, N+1) - 1$. We have shown that antichains in $\widehat{\mathcal{W}}$ have cardinality at most $K(D, N)$. So, by **Dilworth's theorem** [Dil50], we have a partition

$$\widehat{\mathcal{W}} = \bigsqcup_{i=1}^{C} \widehat{\mathcal{W}}_i,$$

where $C \leq K(D, N)$ and each $\widehat{\mathcal{W}}_i$ is a set of halfspaces that is totally ordered by inclusion. Let $\mathcal{W}_i$ be the set of hyperplanes h such that the halfspace associated to h and belonging to $\widehat{\mathcal{W}}$ appears in $\widehat{\mathcal{W}}_i$. Then $\mathcal{W}_i$ is a chain in the sense of Definition 3.2, and for some i, we have $|\mathcal{W}_i| \geq |\mathcal{W}|/K(D, N)$. $\qquad\square$

Here are some consequences. The first answers a question posed by Abdul Zalloum:

Corollary 3.4. *Let X be a CAT(0) cube complex of dimension $D < \infty$. Let $\mathcal{W}$ be a set of hyperplanes in X that does not contain a facing triple. Then X contains a (combinatorial or CAT(0)) geodesic γ such that γ intersects at least $|\mathcal{W}|/K(D, 2)$ of the hyperplanes in $\mathcal{W}$.*

Proof. Use Proposition 3.3, with $N = 2$, to find a chain $\mathcal{C} = \{h_1, \ldots, h_n\}$ of cardinality at least $|\mathcal{W}|/K(D, 2)$ in $\mathcal{W}$. Choose associated halfspaces $\overleftarrow{h}_1, \ldots \overleftarrow{h}_n$ with $\overleftarrow{h}_1 \subset \cdots \subset \overleftarrow{h}_n$. Choose vertices $x \in \overleftarrow{h}_1$ and $y \in \overrightarrow{h}_n$, and let γ be any geodesic from x to y. $\qquad\square$

In a similar vein, we can slightly strengthen Lemma 2.1 from [CS11]. This was pointed out by Elia Fioravanti and Abdul Zalloum.

Corollary 3.5. *Let X be a CAT(0) cube complex of dimension $D < \infty$. Let $k \in \mathbb{N}$. Let γ be a geodesic (in the combinatorial or CAT(0) metric) that crosses at least $D \cdot k$ hyperplanes. Then the set of hyperplanes crossing γ contains a chain of cardinality k.*

Proof. Let $\mathcal{W}$ be the set of hyperplanes intersecting γ and note that $\mathcal{W}$ cannot contain a facing triple. Applying Proposition 3.3 would yield the desired statement with Dk replaced by $K(D, 2)k$. However, one can do a bit better: for each $h \in \mathcal{W}$, let $\overleftarrow{h}$ be the halfspace associated to h and containing $\gamma(0)$. Given $h, v \in \mathcal{W}$, the halfspaces $\overleftarrow{h}, \overleftarrow{v}$ are $\subset$–incomparable only if h, v cross, so the claim follows by applying Dilworth's theorem to $\{\overleftarrow{h} : h \in \mathcal{W}\}$. $\qquad\square$

It is well known that, when X is finite-dimensional, the CAT(0) metric is quasi-isometric to the combinatorial metric (see e.g. [CS11, Lemma 2.2]). In the interest of self-containment, here we give a cosmetically different statement and proof [CS11] working with arbitrary points instead of vertices:

Lemma 3.6. *Let X be a CAT(0) cube complex of dimension $D < \infty$. Then there exist $\lambda_0 \geq 1, \lambda_1 \geq 0$, depending only on D, such that the following holds. Let $x, y \in X$ be arbitrary points, and let $\mathcal{W}(x, y)$ be the set of hyperplanes separating x from y. Then*

$$\frac{1}{\lambda_0} \mathsf{d}_2(x, y) - \lambda_1 \leq |\mathcal{W}(x, y)| \leq \lambda_0 \mathsf{d}_2(x, y) + \lambda_1.$$

Proof. Let γ be a CAT(0) geodesic joining x to y. Note that a hyperplane h intersects γ in a single point if and only if $h \in \mathcal{W}(x, y)$ (any other hyperplane either contains γ or is disjoint from γ). By Corollary 3.5, $\mathcal{W}(x, y)$ contains a chain $\mathcal{C}$ of hyperplanes with cardinality at least $\lfloor |\mathcal{W}(x, k)|/D \rfloor$. If $h, v \in \mathcal{C}$ are distinct, then $\mathsf{d}_2(h, v) \geq 1$, so $|\gamma| \geq |\mathcal{C}| - 1$. Hence

$$|\mathcal{W}(x, y)| \leq D\mathsf{d}_2(x, y) + 2D.$$

To prove the other inequality, let c_x, c_y be cubes containing x, y respectively. Then the set $\mathcal{W}(c_x, c_y)$ of hyperplanes separating c_x, c_y satisfies $|\mathcal{W}(c_x, c_y)| \leq |\mathcal{W}(x, y)|$. Now, fix a combinatorial geodesic $\alpha : [0, L] \to X$ from c_x to c_y having length $|\mathcal{W}(c_x, c_y)|$. Then

$$\mathsf{d}_2(x, y) \leq \mathsf{d}_2(x, \alpha(0)) + \mathsf{d}_2(y, \alpha(L)) + \mathsf{d}_2(\alpha(0), \alpha(L)).$$

Just because d_2 is a path-metric, and edges of X have d_2–length 1,

$$\mathsf{d}_2(\alpha(0), \alpha(L)) \leq \mathsf{d}(\alpha(0), \alpha(L)) = |\mathcal{W}(c_x, c_y)|.$$

Since $x, \alpha(0)$ lie in a common cube c_x, we have $\mathsf{d}_2(x, \alpha(0)), \mathsf{d}_2(y, \alpha(L)) \leq \sqrt{D}$. So

$$\mathsf{d}_2(x, y) - 2\sqrt{D} \leq |\mathcal{W}(x, y)|,$$

as required. $\qquad\square$

We can also produce large facing tuples, given a bound on chains.

Definition 3.7. Let X be a CAT(0) cube complex, let $x_0 \in X^{(0)}$, and let $B_R(x_0)$ be the set of vertices $y \in X$ with $\mathsf{d}(x_0, y) \leq R$. Then $\mathcal{H}_R$ denotes the set of hyperplanes h such that h crosses $B_R(x_0)$.

Observe:

Lemma 3.8. *Any chain in $\mathcal{H}_R$ has cardinality at most $2R$.*

Now we can use Proposition 3.3:

Corollary 3.9. *Let X be a CAT(0) cube complex with dimension $D < \infty$, let $x_0 \in X^{(0)}$, and for each R, let $\mathcal{H}_R$ be as in Definition 3.7. Let $N \in \mathbb{N}$. Then one of the following holds:*

- *There exists R_0 such that for all $R \geq R_0$, there is a facing $(N + 1)$–tuple in $\mathcal{H}_R$.*
- *(X, d) admits an isometric embedding in the standard tiling of $\mathbb{E}^L$ by L–cubes, where $L \leq K(D, N)$.*

In the second case, $|\mathcal{H}_R|$ grows at most linearly in R.

Proof. Suppose that the first conclusion fails. Then N bounds the cardinality of facing tuples in $\mathcal{H}_R$ for all $R \geq 0$. Let $\widehat{\mathcal{H}_R}$ be a set of halfspaces associated to hyperplanes in $\mathcal{H}_R$, with one halfspace per hyperplane. By the proof of Proposition 3.3, $\widehat{\mathcal{H}_R}$ can be partitioned into $L \leq K(D, N)$ sets $\widehat{\mathcal{H}_R}^1, \ldots, \widehat{\mathcal{H}_R}^L$, each totally ordered by inclusion.

Hence $\mathcal{H}_R$ can be partitioned into L chains $\mathcal{H}_R^1, \ldots, \mathcal{H}_R^L$. For $i \leq R$, let $f_i : X \to X_i$ be the *restriction quotient* (see [CS11, Section 2]) obtained by cubulating the wallspace $(X^{(0)}, \mathcal{H}_R^i)$. Since $\mathcal{H}_R^i$ is a finite chain, X_i is isomorphic to the tiling of a finite segment by 1–cubes.

Taking the product gives a map $f' : X \to \prod_i X_i$ which is an isometric embedding on $B_R(x_0)$. Since $\mathrm{diam}(X_i) \leq 2R$, by Lemma 3.8, we can isometrically embed X_i in the cubical tiling of $[-R, R]$, take the product over i of these embeddings, and compose with f' to get a (combinatorial) isometric embedding $f_R : B_R(x_0) \to [-R, R]^L$.

For each R, there are finitely many such embeddings, and for $R = 0$, the embedding is unique. Let G be the graph whose vertex set is the set of isometric embeddings $f_R : B_R(x_0) \to [-R, R]^L$ for $R \geq 0$. Join vertices f_R and f_{R+1} by an edge if f_{R+1} restricts to f_R on $B_R(x_0)$. Then G is a connected, locally finite graph. So, by König's lemma [Kő50], we obtain a combinatorial isometric embedding of X in $\mathbb{E}^L$, as required.

Finally, in this situation, Lemma 3.8 and Proposition 3.3 combine to prove that $|\mathcal{H}_R| \leq 2RK(D, N)$ for all R. $\qquad\square$

Using the above, we recover a Tits alternative for cubulated groups; see also [CS11, SW05].

Proposition 3.10. *Let the group G act essentially on the D–dimensional CAT(0) cube complex X. Suppose that the action of G is either cocompact and X is locally finite, or G has no global fixed point in ∂X. Then either G contains a nonabelian free group, or X admits a combinatorial isometric embedding in $\mathbb{E}^L$, where $L \leq K(D, 3)$.*

Proof. By Corollary 3.9, either the second conclusion holds, or there exists R such that $\mathcal{H}_R$ contains a facing 4–tuple a, b, c, d. Let $\overleftarrow{a}$ be the halfspace associated to a that is disjoint from b, c, d, and define $\overleftarrow{b}, \overleftarrow{c}, \overleftarrow{d}$ analogously. Let $\overrightarrow{a} = X - \overleftarrow{a}$, and define $\overrightarrow{b}, \overrightarrow{c}, \overrightarrow{d}$ analogously.

Apply the Double Skewering Lemma to find $g, h \in G$ such that $g\overrightarrow{b} \subsetneq \overleftarrow{a}$ and $h\overrightarrow{d} \subsetneq \overleftarrow{c}$.

Hence $g(X - \overleftarrow{b}) \subset \overleftarrow{a}$ and $g^{-1}(X - \overleftarrow{a}) \subset \overleftarrow{b}$. Similarly, $h(X - \overleftarrow{d}) \subset \overleftarrow{c}$ and $h^{-1}(X - \overleftarrow{c}) \subset \overleftarrow{b}$. So, applying the ping-pong lemma to the four disjoint sets $\overleftarrow{a}, \overleftarrow{b}, \overleftarrow{c}, \overleftarrow{d}$ and the elements g, h shows that $\langle g, h \rangle \cong F_2$. $\qquad\square$

Corollary 3.11. *Let the group G act properly and cocompactly on the D-dimensional CAT(0) cube complex X. Then G contains a nonabelian free group or G is virtually finite-rank abelian.*

Proof. Using [CS11, Proposition 3.5], we can assume that the action of G on X is essential. By Proposition 3.10, either G contains a nonabelian free group, or X isometrically embeds in $\mathbb{E}^L$ for some L. Observe that G is finitely generated, and the composition of an orbit map $G \to X$ with the embedding $X \to \mathbb{E}^L$ shows that G has polynomial growth. Thus G is virtually nilpotent [Gro81] and hence virtually abelian [BH99, Theorem II.7.8]. $\qquad\square$

Remark 3.12. The above proof is similar to the more general proof in [CS11], but focused on what we hope could be a more "quantitative" way to find the necessary facing 4–tuple in the cocompact case (they use the "Flipping lemma" to find a facing 4–tuple).

One can get the same conclusion without cocompactness, as long as there is a bound on the order of finite subgroups of G [SW05]. One can combine the above argument with a result of Caprace from [CFI16] to obtain the same conclusion as in [SW05], but the details are tangential to our goal here. Instead, we wish to highlight that, by understanding the growth of $\mathcal{H}_R$, and by cooking up an "effective" version of the Double-Skewering Lemma, one could hope to use Corollary 3.9 to prove effective versions of the Tits alternative. There is some discussion of this in the last section of this paper.

4. Loxodromic isometries of $\mathcal{C}X$

The main theorem in this section is Theorem 4.1. This restates results in [Hag13, Section 5]; it can also be assembled from more recent results of Genevois [Gen19a, Gen16, Gen19c], who has significantly extended the study of rank-one isometries of cube complexes and of contact graphs and related objects. We give a simplified proof using gates and not mentioning disc diagrams.

Theorem 4.1. *Let X be a finite-dimensional, locally finite CAT(0) cube complex. Let $g \in \mathrm{Aut}(X)$ be a combinatorially hyperbolic isometry of X. Then the following are equivalent:*

(1) *g acts on $\mathcal{C}X$ as a loxodromic isometry.*

(2) *$\langle g \rangle$ has an unbounded orbit in $\mathcal{C}X$.*

(3) *g acts on X as a rank-one isometry, and for all $n > 0$ and all hyperplanes h of X, we have $g^n h \neq h$.*

(4) *For all $x \in X^{(0)}$ there exists $R = R(g, x)$ such that for all hyperplanes h and all $n \in \mathbb{Z}$, we have $\mathsf{d}(\mathfrak{g}_h(x), \mathfrak{g}_h(g^n x)) \leq R$.*

Moreover, if $\langle g \rangle$ has a bounded orbit in $\mathcal{C}X$, then for any hyperplane h intersecting an axis of g, the orbit $\langle g \rangle \cdot h$ has diameter at most 3 in $\mathcal{C}X$.

The lemmas supporting the proof of Theorem 4.1 can be found in Section 4.1 below. Recall that to prove that $g \in G$ acts on $\mathcal{C}X$ loxodromically, we need to prove that for some (hence any) hyperplane h, the quantity $\mathsf{d}_{\mathcal{C}X}(h, g^n h), n \in \mathbb{Z}$ is bounded below by a linear function of $|n|$.

Proof of Theorem 4.1. The implication $(1) \implies (2)$ is immediate.

The implication $(2) \implies (3)$**:** Suppose that $\langle g \rangle$ has an unbounded orbit in $\mathcal{C}X$. So, $g^n h \neq h$ for all hyperplanes h and all $n > 0$. It remains to check that g is rank-one. Since g does not have a power stabilising a hyperplane, no axis of g can lie in a neighbourhood of a hyperplane, since X is locally finite. Suppose that g is not rank-one and let A be a combinatorial geodesic axis for g. Lemma 4.8 below implies that for all $N > 0$, there are edges e_N, f_N of A such that $\mathsf{d}(e_N, f_N) > N$ and $\mathsf{d}_{\mathcal{C}X}(\pi(e_N), \pi(h_N)) = 1$, since by the lemma the hyperplanes h_N, v_N respectively dual to e_N, f_N intersect. Now, if a, b are arbitrary hyperplanes intersecting A, we can choose e_N, f_N so that the subpath of A between e_N, f_N contains the edges dual to a and b. Now, each of a, b must cross either h_N or v_N, so $\mathsf{d}_{\mathcal{C}X}(a, b) \leq 3$. Hence $\mathrm{diam}(\pi(A)) \leq 3$, i.e. $\langle g \rangle$ has a bounded orbit in $\mathcal{C}X$, a contradiction.

For **the implication** $(3) \implies (4)$, assume that g is rank-one and has no power stabilising a hyperplane. Fix a combinatorial geodesic axis A of g and let Y be its convex hull. By Lemma 4.8, there exists $R_0 < \infty$ such that $Y \subseteq \mathcal{N}_{R_0}(A)$. We will show that there exists R_1 such that $\mathrm{diam}(\mathfrak{g}_Y(h)) = \mathrm{diam}(\mathfrak{g}_h(Y)) \leq R_1$ for all hyperplanes h. From this, we get assertion (4) as follows: let $x \in X^{(0)}$ and let $y = \mathfrak{g}_Y(x)$. For any hyperplane h and any $n \in \mathbb{Z}$, we have

$$\mathsf{d}(\mathfrak{g}_h(x), \mathfrak{g}_h(g^n x)) \leq \mathsf{d}(\mathfrak{g}_h(y), \mathfrak{g}_h(g^n y)) + \mathsf{d}(\mathfrak{g}_h(x), \mathfrak{g}_h(y)) + \mathsf{d}(\mathfrak{g}_h(g^n x), \mathfrak{g}_h(g^n y)),$$

by the triangle inequality. Since $\mathfrak{g}_h$ is 1–lipschitz, $\mathsf{d}(\mathfrak{g}_h(x), \mathfrak{g}_h(y)) \leq \mathsf{d}(x, y) = \mathsf{d}(x, Y)$ and

$$\mathsf{d}(\mathfrak{g}_h(g^n x), \mathfrak{g}_h(g^n y)) \leq \mathsf{d}(g^n x, g^n y) = \mathsf{d}(x, y) = \mathsf{d}(x, Y).$$

Hence $\mathsf{d}(\mathfrak{g}_h(x), \mathfrak{g}_h(g^n x)) \leq R_1 + 2\mathsf{d}(x, Y)$, and (4) holds with $R(g, x) = R_1 + 2\mathsf{d}(x, Y)$. So, it remains to produce R_1.

Suppose that no such R_1 exists, so that for all $N > 0$, there exists a hyperplane h_N with $\mathrm{diam}(\mathfrak{g}_Y(h_N)) > N$. Fix a base 0–cube $a \in A$.

Now, if h is a hyperplane separating h_N from Y, we have $\mathrm{diam}(\mathfrak{g}_Y(h_N)) \leq \mathrm{diam}(\mathfrak{g}_Y(h))$, and we can replace h_N by h. Hence we can assume that no hyperplane separates h_N from Y. Thus $\mathcal{N}(h_N)$ contains a point x_N lying in Y. Note that $\mathsf{d}(x_N, A) \leq R_0$.

So, by translating by an appropriate power of g and enlarging R_0 by an amount depending on g, we can assume x_N satisfies $\mathsf{d}(a, x_N) \leq R_0$ for the base 0–cube $a \in A$ chosen above. Since X is locally finite, only finitely many hyperplanes intersect the R_0–ball about a.

So $\{h_N\}_{N>0}$ is finite, whence there exists a hyperplane h such that $\mathfrak{g}_Y(h)$ is unbounded and a is R_0–close to h. Now, since $Y \cap \mathcal{N}(h) \neq \emptyset$, we have $\mathfrak{g}_Y(\mathcal{N}(h)) = Y \cap \mathcal{N}(h)$. Thus $Y \cap \mathcal{N}(h)$ is unbounded. Now, $Y \subset \mathcal{N}_{R_0}(A)$. Hence A has a sub-ray A' lying in the R_0–neighbourhood of $\mathcal{N}(h)$. In other words, for all $n \geq 0$ (say), we have $\mathsf{d}(h, g^n a) \leq R_0$. Hence, for all n and all $0 \leq i \leq n$, we have $\mathsf{d}(g^i h, g^n a) \leq R_0$. Let K be the number of hyperplanes crossing the R_0–ball about a. Then for $n > K$, the list $h, gh, \cdots, g^n h$ must contain two identical elements, by the pigeonhole principle, so $h = g^i h$ for some $i \neq 0$, a contradiction. This completes the proof that $(3) \implies (4)$.

Now **we prove** $(4) \implies (1)$. Recall that we need to verify that for each hyperplane h, there exists $C > 0$ such that $\mathsf{d}_{\mathcal{C}X}(h, g^n h) \geq Cn$ for all $n > 0$.

Let h be a hyperplane and let $x \in \mathcal{N}(h)$. Fix $n > 0$ and consider the vertices x and $g^n x$. By Lemma 4.11 below, there exists a combinatorial geodesic $\gamma = \gamma_1 \cdots \gamma_k$ joining x to $g^n x$ and having the following properties:

- there is a sequence $h = h_1, \ldots, h_k = g^n h$ of hyperplanes such that $\mathcal{N}(h_i) \cap \mathcal{N}(h_{i+1}) \neq \emptyset$ for $1 \leq i \leq k - 1$;
- $h = h_1, \ldots, h_k = g^n h$ is a geodesic of $\mathcal{C}X$;
- the geodesic γ_i lies in $\mathcal{N}(h_i)$ for $1 \leq i \leq k$;
- γ_i has length at most $\mathsf{d}(\mathfrak{g}_{\mathcal{N}(h_i)}(x), \mathfrak{g}_{\mathcal{N}(h_i)}(g^n x))$.

By (4) and the fourth bullet point above, there exists $R = R(g, x) \geq 1$ such that $|\gamma_i| \leq R$ for all i. Hence $k \geq \frac{1}{R}\mathsf{d}(x, g^n x)$ for all n. On the other hand, since g is combinatorially hyperbolic, there exists $\tau \geq 1$ (depending on g) such that $\mathsf{d}(x, g^n x) \geq \tau n$ for all n. So, taking $C = \tau/2R$ completes the proof.

To conclude, we prove the **"moreover" clause**. Suppose that $\langle g \rangle$ has a bounded orbit in $\mathcal{C}X$, so that by the equivalence of (2) and (3), either $g^n h = h$ for some hyperplane h and some $n > 0$, or g is not rank-one (or both). We saw above that if the former does not hold, and g is not rank-one, then $\langle g \rangle \cdot v$ has diameter at most 3 in $\mathcal{C}X$ for any hyperplane v crossing a g–axis.

If the former holds, then since g^n stabilises h, the carrier $\mathcal{N}(h)$ contains an axis B of g^n. Any hyperplane v crossing B crosses h, whence $\mathsf{d}_{\mathcal{C}X}(v, g^j v) \leq 2$ for all $j \in \mathbb{Z}$. Any g–axis is also a g^n–axis, and any two axes of g^n cross the same hyperplanes. So any hyperplane v intersecting any g–axis must intersect B and thus satisfy $\mathsf{d}_{\mathcal{C}X}(v, g^j v) \leq 2$ for all j, as required. $\square$

Remark 4.2 (The local finiteness hypothesis and other proofs in the literature)**.** We can obtain a similar conclusion without the local finiteness hypothesis. Specifically, Lemmas 4.11 and 4.8 do not use that assumption: the former works for *arbitrary* CAT(0) cube complexes, and the latter uses only that X is finite-dimensional. Provided X is finite-dimensional, but with no local finiteness hypothesis, a variant of the above proof gives that (1),(2),(4) are all equivalent.

There are other ways to prove parts of Theorem 4.1 from results in the literature. For example, equivalence of (4) and (1) can easily be deduced from a result of Genevois [Gen19a, Proposition 4.2]. Meanwhile, it is straightforward to show (2) implies (4), and (1) implies (2) is obvious.

Remark 4.3 (Non-equivariant versions)**.** One can state a non-equivariant version of Theorem 4.1 about projecting geodesic rays to $\mathcal{C}X$; see Section 2 of [Hag13]. Here is a simple version. Let X be a CAT(0) cube complex (with no local finiteness or dimension assumption). Let $\gamma : [0, \infty) \to X$ be a combinatorial geodesic ray. Then $\pi \circ \gamma : [0, \infty) \to \mathcal{C}X$ is a quasigeodesic if and only if there exists R such that $\mathrm{diam}(\mathfrak{g}_{\mathcal{N}(h)}(\gamma)) \leq R$ for all hyperplanes h. Indeed, given such a bound, Lemma 4.11 implies that $\pi \circ \gamma$ uniformly fellow-travels a geodesic ray in $\mathcal{C}X$. On the other hand, if $\mathrm{diam}(\mathfrak{g}_{\mathcal{N}(h)}(\gamma))$ is unbounded, then γ has arbitrarily large subpaths projecting to stars in $\mathcal{C}X$, so $\pi \circ \gamma$ is not a (parameterised) quasigeodesic. Such a $\pi \circ \gamma$ still has image lying at finite Hausdorff distance from a geodesic ray or segment in $\mathcal{C}X$, again by Lemma 4.11.

Here are some corollaries. The first appears in [Hag13, Section 5], but we reproduce it since we will use it in Section 5.

Corollary 4.4 ($\mathcal{C}X$ loxodromics skewer specified pairs)**.** *Let X be a finite-dimensional CAT(0) cube complex. Let $G \to \mathrm{Aut}(X)$ be an essential group action. Assume that one of the following holds: G does not fix a point in ∂X, or X is locally finite and G acts cocompactly.*

Suppose that h, v are hyperplanes of X satisfying $\mathsf{d}_{\mathcal{C}X}(h, v) > 3$. Then there exists a hyperbolic isometry $g \in G$ such that g acts loxodromically on $\mathcal{C}X$ and v separates h from gh.

Proof. By the Double-Skewering Lemma, there exists a hyperbolic isometry $g \in \mathrm{Aut}(X)$ of X such that v separates h from gh. It follows that $\mathsf{d}_{\mathcal{C}X}(h, gh) > 3$, so g is loxodromic on $\mathcal{C}X$ by Theorem 4.1. $\square$

Hyperplanes h, v are *strongly separated* if they are disjoint and no hyperplane intersects both. This notion is due to Behrstock-Charney [BC12] and plays an important role in Caprace-Sageev's rank rigidity theorem [CS11]. For generalisations and applications of this property, see, for example, [CFI16, CM19, Gen16, Lev18, CS15].

Corollary 4.5 (Strong separation criterion)**.** *Let X be a finite-dimensional CAT(0) cube complex. Let $G \to \mathrm{Aut}(X)$ be an essential group action. Assume that one of the following holds: G does not fix a point in ∂X, or X is locally finite and G acts cocompactly.*

Suppose that h, v are strongly separated hyperplanes of X. Then there exists a hyperbolic isometry $g \in G$ such that g acts loxodromically on CX and v separates h from gh (i.e. g double-skewers the pair h, v).

Remark 4.6. As mentioned above, a result of Genevois implies that if v, h is a pair of strongly separated hyperplanes and g is an isometry of X double-skewering the pair v, h (i.e. $g\overleftarrow{h} \subsetneq \overleftarrow{v} \subsetneq \overleftarrow{h}$), then g is loxodromic on CX [Gen19a, Proposition 4.2]. Conversely, if g is loxodromic on CX, and h is a hyperplane crossing some (hence any) axis of g, then we can choose $n > 0$ such that $\mathsf{d}_{CX}(h, g^n h) > 2$, so $h, g^n h$ are strongly separated. Moreover, $g^{2n}h$ and h are separated by $g^n h$, since g translates along the axis (preserving the orientation), and thus we have a strongly-separated pair skewered by a power of g.

Proof of Corollary 4.5. Apply Double-Skewering to h, v to obtain a hyperbolic element g double-skewering the pair h, v; in particular, v separates h, gh (as required by the statement).

We will show that $\mathsf{d}_{CX}(h, g^4 h) > 3$. Since g double-skewers h and $g^4 h$, it will then follow from Theorem 4.1 (as in the proof of Corollary 4.4) that g is CX–loxodromic.

Suppose, for the sake of a contradiction, that $\mathsf{d}_{CX}(h, g^4 h) \le 3$. Consider the hyperplanes $h, gh, g^2 h, g^3 h, g^4 h$. Let $h = u_1, u_2, u_3, u_4 = g^4 h$ be the vertex sequence of a CX–path of length at most 3 from h to $g^4 h$. Then for some $j \le 4$ and some $i \le 3$, we have that u_j crosses $g^i h$ and $g^{i+1} h$, and hence crosses $g^i v$ and $g^i h$. But then $g^{-i} u_j$ crosses h and v, contradicting strong separation. $\qquad\square$

Corollary 4.7 (Rank-rigidity and the contact graph)**.** *Let X be a finite-dimensional, irreducible, locally finite, essential CAT(0) cube complex, and suppose X has at least one hyperplane (i.e. X is not a single vertex). Suppose that the action of the group G on X is cocompact. Then G contains an element g acting loxodromically on CX.*

Proof. Since X is locally finite and G acts cocompactly, Corollary 4.9 of [CS11] applies (X is unbounded since it contains at least one hyperplane and is essential; the action is *hereditarily essential*, in the language of [CS11], since X is locally finite and G acts cocompactly). Since X is irreducible, we therefore have one of the following:

(1) G does not fix a point in the visual boundary ∂X. In this case, Proposition 5.1 of [CS11] says that X contains a strongly separated pair of hyperplanes.
(2) All hyperplanes of X are compact. Since G acts cocompactly, there are finitely many orbits of hyperplanes. So we can choose $D \ge 0$ that exceeds the diameter of every hyperplane. Since X is unbounded, it contains hyperplanes h, v with $\mathsf{d}(h, v) > D$. Such a pair must be strongly separated.

In either case, X contains a strongly separated pair of hyperplanes, and we conclude by applying Corollary 4.5. $\qquad\square$

4.1. Supporting lemmas.
The following lemma is known (compare e.g. [Gen19c, Proposition 4.2]), but we prove it here for self-containment:

Lemma 4.8 (Convex hulls of axes)**.** *Let X be a finite-dimensional CAT(0) cube complex and let $g \in \mathrm{Aut}(X)$ be combinatorially hyperbolic. Let A be a combinatorial geodesic axis for g, and let Y be the cubical convex hull of A. Then:*

(1) *If no regular neighbourhood of $A^{(0)}$ contains $Y^{(0)}$, the element g is not rank-one.*
(2) *If some regular neighbourhood of $A^{(0)}$ contains $Y^{(0)}$, then either g is rank-one or A lies in a regular neighbourhood of a hyperplane.*

In particular, if g is not rank-one and A does not lie in a neighbourhood of a hyperplane, then for all $N \ge 0$, there exist hyperplanes h, v, dual to edges e_h, e_v of A, such that $\mathsf{d}(e_h, e_v) > N$ but $h \cap v \ne \emptyset$.

Before the proof of Lemma 4.8, we need two auxiliary lemmas:

Lemma 4.9. *Let X, g, A, Y be as in Lemma 4.8. For $r \geq 0$, let A_r be the subgeodesic from $A(-r)$ to $A(r)$ and let Y_r be the cubical convex hull of A_r. Then $Y = \bigcup_{r \geq 0} Y_r$.*

Proof. Since any convex subcomplex containing A must contain A_r for all r, we have $\bigcup_{r \geq 0} Y_r \subseteq Y$. To prove the reverse inclusion, let $y \in Y$ be a vertex. Choose $r \geq 0$ such that every hyperplane separating y from $A(0)$ separates $A(-r)$ from $A(r)$; this is possible since each hyperplane separating y from $A(0)$ crosses Y and hence A, and since there are finitely many such hyperplanes. Any halfspace containing A_r contains $A(-r), A(0), A(r)$, and hence the associated hyperplane does not separate $A(0)$ from y. Thus $y \in Y_r$. Hence $Y = \bigcup_{r \geq 0} Y_r$. $\qquad\square$

In the next lemma, by a *half-flat* in the CAT(0) cube complex X, we mean an isometric embedding $F : [0, \infty) \times \mathbb{R} \to X$ where X is given the CAT(0) metric and $[0, \infty) \times \mathbb{R}$ is given the Euclidean metric. We also use the notation F, and the term "half-flat", for its image in X. The bi-infinite CAT(0) geodesic β *bounds* the half-flat F if $F|_{\{0\} \times \mathbb{R}} = \beta$ (here we are also using the notation β both for the isometric embedding $\beta : \mathbb{R} \to X$ and for its image).

Lemma 4.10 (Hyperplanes and half-flats). *Let X be a finite-dimensional CAT(0) cube complex. Let β be a CAT(0) geodesic that bounds a half-flat F. Let h be a hyperplane of X such that $h \cap F \neq \emptyset$. Then one of the following holds:*
- *$F \subseteq h$;*
- *$F \cap h$ is a line in F parallel to β;*
- *$F \cap h$ is a ray in F whose initial point is on β.*

Proof. This is an exercise in CAT(0) geometry using CAT(0) convexity of hyperplanes and halfspaces and the product structure of hyperplane carriers; see e.g. [HP15, Remark 3.4]. $\qquad\square$

Proof of Lemma 4.8. We first prove assertion (1). Fix $y \in Y^{(0)}$. We claim that y lies on some combinatorial geodesic with endpoints on A. Indeed, by Lemma 4.9, we can choose $r \geq 0$ such that $y \in Y_r$, where Y_r is the convex hull of the subgeodesic of A from $A(-r)$ to $A(r)$. Since Y_r is, equivalently, the interval between $A(-r)$ and $A(r)$, we have that y lies on a geodesic from $A(-r)$ to $A(r)$.

Suppose that Y does not lie in a regular neighbourhood of A. Then for any $R > 0$, the above argument shows that there is a combinatorial geodesic that has endpoints on A but does not lie in $\mathcal{N}_R(A)$. Hence A is not a *Morse geodesic* (see e.g. [ACGH17, Definition 1.2] for the definition; the notion goes back in some form to [Mor24]).

Now, Y is a proper CAT(0) space, because it is the convex hull of a bi-infinite combinatorial geodesic.[1] Moreover, $\langle g \rangle$ acts on Y by isometries, with g acting hyperbolically. Finally, any CAT(0) geodesic axis in Y for g is not Morse. Hence, by [Sul14, Lemma 3.3] and [BF09, Theorem 5.4], g is not rank-one.

We now prove assertion (2) and the "in particular" statement.

For both, we suppose that g is not rank-one and A does not lie in any neighbourhood of any hyperplane. Let $\beta : \mathbb{R} \to Y$ be a CAT(0) geodesic axis of g, and let F be a half-flat bounded by β. (The half-flat F need not be unique or g–invariant, but we will not need this.)

Let h be a hyperplane intersecting F. By Lemma 4.10, $h \cap F$ is either all of F, a line parallel to β, or a ray starting on β. In either of the first two cases, β is contained in a neighbourhood of h, so, since A and β lie at finite Hausdorff distance, the same is true for A, a contradiction. Hence h intersects β in a point and intersects F in a ray.

---

[1] This follows from [BCG$^+$09, Theorem 1.14], and one can also prove it using Corollary 3.9 and the fact that Y contains no facing triple.

From this, we obtain assertion (2) as follows. We have seen that every hyperplane crossing F crosses β, and hence crosses A, since A and β cross the same hyperplanes. Let Z be the cubical convex hull of F. Any hyperplane crossing Z crosses F, and hence crosses Y.

Consider the gate map $\mathfrak{g}_Y : X \to Y$. We claim that there exists $N_0 \in \mathbb{N}$ such that for all $z \in Z^{(0)}$, we have $\mathsf{d}_2(z, \mathfrak{g}_Y(z)) \leq N_0$. Indeed, let w be a hyperplane separating z from $\mathfrak{g}_Y(z)$. Recall from Section 2 that such a w must separate z from Y. Since every hyperplane crossing Z crosses Y, we see that w must not cross Z, and thus separates all of Z from Y. There are finitely many such hyperplanes, so Lemma 3.6 provides the constant N_0.

Let $N \geq 0$ be given. Choose $f \in F$ such that $\mathsf{d}_2(f, \beta) > N$, which is possible since F is a half-flat bounded by β. Choose a 0–cube $f' \in Z$ such that $\mathsf{d}_2(f, f') \leq \sqrt{\dim X}/2$. So $\mathfrak{g}_Y(f') \in Y$ satisfies
$$\mathsf{d}_2(f, \mathfrak{g}_Y(f')) \leq \sqrt{\dim X}/2 + N_0.$$
Hence
$$\mathsf{d}_2(\mathfrak{g}_Y(f'), \beta) > N - \sqrt{\dim X}/2 - N_0.$$
Since β and A are at finite Hausdorff distance, and we could have chosen N arbitrarily large, Y contains 0–cubes arbitrarily far from A, as required. (We have used the CAT(0) metric here but the same conclusion applies in the ℓ_1 metric by Lemma 3.6.)

To prove the "in particular" statement, let h be a hyperplane intersecting F, and recall that $h \cap F$ is a ray starting on β. Let $\rho = h \cap F$ be such a ray, and let $t \in \mathbb{R}$ be such that $\rho \cap \beta = \beta(t)$. Since ρ is a CAT(0) geodesic ray in X, for all $n \in \mathbb{N}$ there exists a hyperplane h_n such that $\mathsf{d}(h_n \cap \rho, \beta(t)) > n$. Now, since h_n intersects F, it does so in a ray ν_n intersecting β at a point $\beta(t_n)$. Since any interval in β intersects finitely many hyperplanes (by, say, Lemma 3.6), the quantity $|t_n|$ is unbounded as $n \to \infty$. Hence, for any $r \geq 0$, there exist hyperplanes h, v such that $h \cap v \neq \emptyset$ but $\mathsf{d}_2(h \cap \beta, v \cap \beta) > r$. Since β and A fellow travel (in the CAT(0) and ℓ_1 metrics), and cross the same hyperplanes, we obtain the desired conclusion. $\qquad\square$

The next lemma is Lemma 3.1 in [BHS17a]. Here we give essentially the same proof, except using gates instead of disc diagrams.

Lemma 4.11 ("Hierarchy paths" in $\mathcal{C}X$). *Let X be a CAT(0) cube complex and let $x, y \in X^{(0)}$. Let h_x, h_y be hyperplanes such that $x \in \mathcal{N}(h_x)$ and $y \in \mathcal{N}(h_y)$. Then there exists a sequence $h_x = h_1, \ldots, h_k = h_y$ of hyperplanes such that:*

- *$\mathcal{N}(h_i) \cap \mathcal{N}(h_{i+1}) \neq \emptyset$ for $1 \leq i \leq k - 1$, and*
- *$h_1, \ldots, h_k$ is a $\mathcal{C}X$–geodesic from h_x to h_y;*
- *there is a path $\gamma = \gamma_1 \cdots \gamma_k$ from x to y, where each γ_i is a combinatorial geodesic in $\mathcal{N}(h_i)$;*
- *γ is a combinatorial geodesic;*
- *for each i, $|\gamma_i| \leq \mathsf{d}(\mathfrak{g}_{\mathcal{N}(h_i)}(x), \mathfrak{g}_{\mathcal{N}(h_i)}(y))$.*

Hence x and y are joined by a geodesic γ such that $\pi \circ \gamma$ is an unparameterised quasigeodesic in $\mathcal{C}X$ (with constants independent of X).

Remark 4.12. Lemma 4.11 says roughly that any two points in X can be joined by a geodesic that tracks a geodesic between their projections to the contact graph. This is reminiscent of "hierarchy paths" in the marking complex of a surface [MM00], with the curve graph playing the role of the contact graph. This similarity is part of the motivation for the notion of a *hierarchically hyperbolic space* [BHS17a].

Proof of Lemma 4.11. Let $h_x = h_1, \ldots, h_k = h_y$ be a sequence of hyperplanes satisfying the first two properties, which is possible just because $\mathcal{C}X$ is a connected graph.

Let $x = x_1 \in \mathcal{N}(h_1)$. For $2 \leq i \leq k$, suppose that $x_{i-1} \in \mathcal{N}(h_{i-1})$ has been chosen, and let $x_i = \mathfrak{g}_{\mathcal{N}(h_i)}(x_{i-1}) \in \mathcal{N}(h_i)$. Let $x_{k+1} = y$. For each i, let γ_i be a combinatorial geodesic in $\mathcal{N}(h_i)$ joining x_i to x_{i+1}. Let $\gamma = \gamma_1 \cdots \gamma_k$.

The *complexity* of the pair $((h_1, \ldots, h_k), (\gamma_1, \ldots, \gamma_k))$ of k–tuples is the tuple $(|\gamma_1|, \ldots, |\gamma_k|)$, taken in lexicographic order. Suppose that $\gamma_1, \ldots, \gamma_k$ have been chosen as above so as to minimise the complexity.

γ **is a geodesic:** We claim that γ is a combinatorial geodesic. Suppose to the contrary that some hyperplane h is dual to two distinct edges of γ, respectively lying in γ_i, γ_j for $1 \le i \le j \le k$. We cannot have $i = j$, since γ_i is a geodesic. We also cannot have $j > i + 2$, because h intersects $\mathcal{N}(h_i)$ and $\mathcal{N}(h_j)$, which would yield a path $h_1, \ldots, h_i, h, h_j, \ldots, h_k$ in $\mathcal{C}X$ from h_1 to h_k. This path has length less than $k - 1$, contradicting that $h_1, \ldots, h_k$ is a geodesic.

Hence $j = i + 1$ or $j = i + 2$. If $j = i + 1$, then h intersects $\mathcal{N}(h_i)$ and $\mathcal{N}(h_{i+1})$, and thus does not separate x_i from $\mathcal{N}(h_{i+1})$. Hence h does not separate x_i from $\mathfrak{g}_{\mathcal{N}(h_{i+1})}(x_i)$, which is a contradiction since h is dual to an edge of γ_i, and therefore separates the endpoints of γ_i. Thus $j \ne i + 1$.

If $j = i + 2$, then we can replace h_{i+1} by h to yield a lower-complexity pair. Indeed, we replace $h_1, \ldots, h_k$ by $h_1, \ldots, h_i, h, h_{i+2}, \ldots, h_k$, obtaining a new geodesic in $\mathcal{C}X$ from h_x to h_y. We replace x_{i+1} by $x'_{i+1} = \mathfrak{g}_{\mathcal{N}(h)}(x_i)$ and replace x_{i+2} by $\mathfrak{g}_{\mathcal{N}(h_{i+2})}(x'_{i+1})$. The geodesics $\gamma_1, \ldots, \gamma_{i-1}$ are unchanged, but γ_i is replaced by a geodesic of length

$$\mathsf{d}(x_i, \mathfrak{g}_{\mathcal{N}(h)}(x_i)) = \mathsf{d}(x_i, \mathcal{N}(h)) < \mathsf{d}(x_i, \mathcal{N}(h_{i+1})) = |\gamma_i|,$$

so we have reduced complexity. This contradicts our initial choice of pair of k–tuples, and we conclude that $j \ne i + 2$. Hence no hyperplane is dual to two distinct edges of γ, so γ is a geodesic.

Length of γ_i: Fix i and let h be a hyperplane crossing γ_i. Since γ is a geodesic, h separates x from y. On the other hand, h crosses the convex subcomplex $\mathcal{N}(h_i)$, so h must separate $\mathfrak{g}_{\mathcal{N}(h_i)}(x)$ from $\mathfrak{g}_{\mathcal{N}(h_i)}(y)$, by Lemma 2.5. Since $|\gamma_i|$ is the number of hyperplanes h crossing γ_i, we have $\mathsf{d}(\mathfrak{g}_{\mathcal{N}(h_i)}(x), \mathfrak{g}_{\mathcal{N}(h_i)}(y)) \ge |\gamma_i|$.

Unparameterised quasigeodesic: Let γ be a geodesic from x to y provided by the first part of the lemma, so that $\gamma = \gamma_1 \cdots \gamma_k$, where each γ_i lives in the carrier of a hyperplane h_i, and the sequence $h_1, \ldots, h_k$ is a $\mathcal{C}X$–geodesic from a point $h_1 \in \pi(x)$ to a point $h_k \in \pi(y)$. By construction, $\pi(\gamma_i)$ lies in the 1–neighbourhood in $\mathcal{C}X$ of h_i, so $\pi \circ \gamma$ lies at uniformly bounded Hausdorff distance in $\mathcal{C}X$ from some, and hence any, geodesic from $\pi(x)$ to $\pi(y)$, as required. $\qquad\square$

5. The sector lemma

Our goal is to prove Proposition 1. Fix a CAT(0) cube complex X satisfying the hypotheses: X is irreducible, locally finite, essential, hyperplane-essential, and $\mathrm{Aut}(X)$ acts cocompactly (hence X is finite-dimensional). We continue to use the convention that, if h is a hyperplane, then $\overleftarrow{h}, \overrightarrow{h}$ denote the associated halfspaces.

The main lemma is:

Lemma 5.1. *Let h, v be distinct hyperplanes of X such that $h \cap v \ne \emptyset$. Suppose that $\overleftarrow{h} \cap \overleftarrow{v}$ contains a hyperplane a such that $\mathsf{d}_{\mathcal{C}X}(a, h) > 2$. Then $\overleftarrow{h} \cap \overrightarrow{v}$ contains a hyperplane a' such that $\mathsf{d}_{\mathcal{C}X}(a', h) > 2$.*

Proof. Since $\mathsf{d}_{\mathcal{C}X}(a, h) > 2$, Lemma 2.12 implies that $\mathfrak{g}_h(a)$ is a single point, which we denote by p.

Now, h is a locally finite CAT(0) cube complex, and $\mathrm{Stab}_{\mathrm{Aut}(X)}(h)$ acts on h cocompactly by Lemma 2.7, since $\mathrm{Aut}(X)$ acts on X cocompactly. Since the action of $\mathrm{Aut}(X)$ on X is assumed to be hyperplane-essential, the action of $\mathrm{Stab}_{\mathrm{Aut}(X)}(h)$ on h is essential, so by Proposition 3.2 of [CS11], there exists $g \in \mathrm{Stab}_{\mathrm{Aut}(X)}(h)$ such that, regarded as a hyperplane of h, the intersection $h \cap v$ separates $g^{-1}(h \cap v)$ from $g(h \cap v)$, and $g\overrightarrow{v} \subsetneq \overrightarrow{v}$. By replacing g by g^2 if necessary, we can assume that g stabilises $\overleftarrow{h}$ and $\overrightarrow{h}$. By replacing g with a further positive power, we

have $gp \in \overrightarrow{v}$. In other words, $g\mathfrak{g}_h(a) = \mathfrak{g}_{gh}(ga) = \mathfrak{g}_h(ga)$ is contained in $\overrightarrow{v}$. Hence, since v crosses h, we have $ga \subset \overrightarrow{v}$.

Since g stabilises $\overleftarrow{h}$, we get $ga \subset \overleftarrow{h} \cap \overrightarrow{v}$, as required. Moreover, $\mathsf{d}_{\mathcal{C}X}(ga, h) = \mathsf{d}_{\mathcal{C}X}(ga, gh) = \mathsf{d}_{\mathcal{C}X}(a, h) > 2$, so taking $a' = ga$ completes the proof. $\qquad\square$

Now we can prove the proposition.

Proof of Proposition 1. The proof is summarised in Figure 2.

Since X is irreducible, Corollary 4.7 implies that $\mathcal{C}X$ is unbounded. Hence there exists a hyperplane a such that $\mathsf{d}_{\mathcal{C}X}(a, h) > 2$. Since $\mathsf{d}_{\mathcal{C}X}(h, v) = 1$, we have $\mathsf{d}_{\mathcal{C}X}(a, v) > 1$. So, up to relabelling halfspaces, we have $a \subset \overleftarrow{h} \cap \overleftarrow{v}$. Applying Lemma 5.1, we obtain a hyperplane a' such that $\mathsf{d}_{\mathcal{C}X}(h, a') > 2$ and $a' \subset \overleftarrow{h} \cap \overrightarrow{v}$.
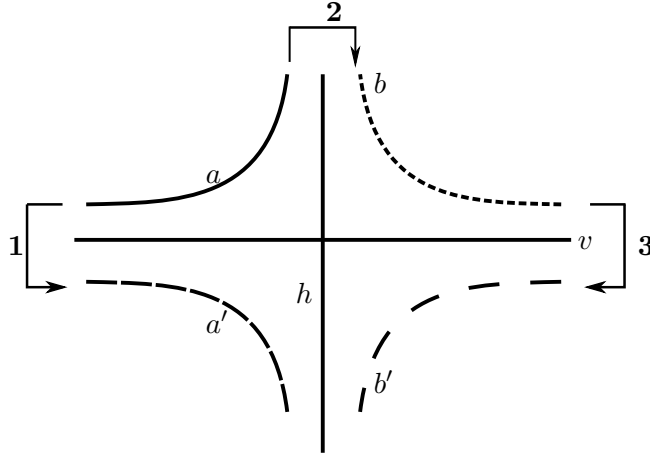

FIGURE 2. Proof of Proposition 1. Start with the hyperplane a. The arrow labelled **1** shows the first application of Lemma 5.1 to a. The arrow labelled **2** shows the application of Corollary 4.5 to the strongly separated pair a, h. The arrow labelled **3** shows the application of Lemma 5.1 to the hyperplane b. The hyperplane b provided by Corollary 4.5 could also have been in the bottom right sector, in which case we would use Lemma 5.1 to move b "up" into the top right.

Apply Corollary 4.5 to the strongly separated pair a, h to obtain a hyperplane $b \subset \overrightarrow{h}$ with $\mathsf{d}_{\mathcal{C}X}(b, h) > 2$ (the hyperplane b arises as a translate of a by some element acting loxodromically on $\mathcal{C}X$ and double-skewering a and h). Again, $\mathsf{d}_{\mathcal{C}X}(b, v) > 1$, so $b \subset \overleftarrow{v}$ or $b \subset \overrightarrow{v}$. Suppose the former holds, i.e. $b \subset \overrightarrow{h} \cap \overleftarrow{v}$. Applying Lemma 5.1 (with the roles of $\overleftarrow{h}$ and $\overrightarrow{h}$ switched), we obtain a hyperplane $b' \subset \overrightarrow{h} \cap \overrightarrow{v}$. By an identical argument, if $b \subset \overrightarrow{h} \cap \overrightarrow{v}$, we obtain a hyperplane $b' \subset \overrightarrow{h} \cap \overleftarrow{v}$. We have shown that each of the four intersections $\overleftarrow{h} \cap \overleftarrow{v}, \overleftarrow{h} \cap \overrightarrow{v}, \overrightarrow{h} \cap \overleftarrow{v}, \overrightarrow{h} \cap \overrightarrow{v}$ contains a hyperplane. One of these four intersections is $h^+ \cap v^+$, so we are done. $\qquad\square$

6. QUESTIONS

We close with some questions about using Proposition 3.3 to effectivise statements about actions on CAT(0) cube complexes.

Question 6.1 (Effective double skewering)**.** Let the group G act cocompactly and essentially on the finite-dimensional, locally finite CAT(0) cube complex X. Find an explicit estimate of the function $f : \mathbb{N} \to \mathbb{N}$ such that the following holds: let v, h be disjoint hyperplanes and let $L = \mathsf{d}(\mathcal{N}(v), \mathcal{N}(h))$. Then there exists $g \in G$ such that v separates h from gh and g has combinatorial translation length at most $f(L)$.

The function f should be allowed to depend on invariants of X like its dimension and the maximum degree of 0–cubes. It also seems reasonable (and necessary) to allow f to depend on the number of orbits of hyperplanes, or the diameter of a smallest compact convex subcomplex whose G–orbit covers X, or some other parameter of the action.

Next, recall that, in the proof of Corollary 3.9, when X is finite-dimensional, we found a constant K, depending on the dimension of X, such that for all $R \geq 0$, either the set $\mathcal{H}_R$ of hyperplanes crossing the R–ball about a fixed basepoint contains a facing 4–tuple, or $|\mathcal{H}_R| \leq KR$.

So, if one knew the growth rate of the function $R \mapsto |\mathcal{H}_R|$, and this growth rate was superlinear, one could compute a minimal R_0 such that $\mathcal{H}_{R_0}$ contains a facing 4–tuple. In conjunction with an answer to Question 6.1, the proof of Proposition 3.10 would then yield $g, h \in G$, each with translation length at most $f(2R_0)$, such that g, h freely generate a free subgroup of G. So:

Question 6.2 (Effective Tits alternative 1). Let X be a locally finite, finite-dimensional CAT(0) cube complex on which the group G acts cocompactly and essentially. Find a constant L_0 such that either $R \mapsto |\mathcal{H}_R|$ grows at most linearly, or G contains a free group generated by elements g, h whose combinatorial translation lengths are at most L_0. The constant L_0 should depend on specific parameters of the G–action in an explicit way, as in Question 6.1.

Remark 6.3. The proof of the Tits alternative in [CS11] also involves an application of the Double Skewering Lemma to two pairs of hyperplanes drawn from a facing 4–tuple. In that setting, the facing 4–tuple is found by applying the *Flipping Lemma* to a facing triple, and then concluding that, if X has no facing triple, then the G–action fixes a point at infinity. So, it would also be interesting to try to effectivise the Flipping Lemma (of which the Double Skewering Lemma is an easy consequence).

Our proof of Proposition 1 also yields a facing 4–tuple of hyperplanes a, b, c, d that are pairwise strongly separated (because they are pairwise at large distance in $\mathcal{C}X$, by construction). Applying the Double Skewering Lemma as in the proof of Proposition 3.10 would then yield a free subgroup of G generated by two elements acting on $\mathcal{C}X$ loxodromically, in view of Corollary 4.5.

Question 6.4 (Effective rank rigidity). Is there an effective version of the Rank-Rigidity Theorem for cocompact actions on CAT(0) cube complexes? Specifically: let X be a finite-dimensional, essential, cocompact, locally finite CAT(0) cube complex, let $x_0 \in X^{(0)}$ be a base vertex. For $R \geq 0$, let $B_R(x_0)$ and $\mathcal{H}_R$ be defined as above. Let $vol(R)$ be the number of 0–cubes in $B_R(x_0)$ and let $Hvol(R) = |\mathcal{H}_R|$. Can one characterise, in terms of the growth rates of $vol(R)$ and $Hvol(R)$, when X splits as a nontrivial product? If X does not split as a nontrivial product, can one estimate the value R_0, depending on $vol(R)$ and $Hvol(R)$, such that $\mathcal{H}_{R_0}$ contains a strongly separated pair of hyperplanes?

Given a cocompact action of G on X, one could then combine this with an answer to Question 6.1 and Corollary 4.5 to produce a rank-one element of bounded translation length.

One can imagine a more complicated version of Question 6.4 about facing 4–tuples of strongly separated hyperplanes, and a free subgroup of G acting purely loxodromically on $\mathcal{C}X$, generated by two elements of bounded translation length on X. One can also imagine a version about the translation lengths on $\mathcal{C}X$, rather than on X. In fact:

Question 6.5. Let X be a finite-dimensional, locally finite, irreducible, essential CAT(0) cube complex with a group G acting cocompactly. Estimate the minimal translation length on $\mathcal{C}X$ of elements of G acting loxodromically.

We finally ask whether this is related to *uniform exponential growth* for cubulated groups. Here the question boils down to: let the finitely generated group G act properly and cocompactly

on the CAT(0) cube complex X. Find a constant λ such that either G is virtually abelian or, for any finite generating set of G, the λ–ball in the corresponding Cayley graph of G contains two elements that freely generate a free (semi)group. There is quite a strong result about this in the 2–dimensional case, due to Kar and Sageev [KS19], and this is an actively-studied question in higher dimensions.

References

[ACGH17]   Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction. *Doc. Math.*, 22:1193–1224, 2017.

[AGM13]   Ian Agol, Daniel Groves, and Jason Manning. The virtual Haken conjecture. *Doc. Math*, 18:1045–1087, 2013.

[AKWW13]   Yael Algom-Kfir, Bronislaw Wajnryb, and Pawel Witowicz. A parabolic action on a proper, CAT(0) cube complex. *J. Group Theory*, 16(6):965–984, 2013.

[AOS12]   Federico Ardila, Megan Owen, and Seth Sullivant. Geodesics in CAT(0) cubical complexes. *Adv. in Appl. Math.*, 48(1):142–163, 2012.

[Ava61]   Sherwin P. Avann. Median algebras. *Proceedings of the American Mathematical Society*, 12:407–414, 1961.

[BC93]   Jean-Pierre Barthélemy and Julien Constantin. Median graphs, parallelism and posets. *Discrete mathematics*, 111(1-3):49–63, 1993.

[BC08]   Hans-Jurgen Bandelt and Victor Chepoi. Metric graph theory and geometry: a survey. *Contemporary Mathematics*, 453:49–86, 2008.

[BC12]   Jason Behrstock and Ruth Charney. Divergence and quasimorphisms of right-angled Artin groups. *Math. Ann.*, 352(2):339–356, 2012.

[BCG$^+$09]   Jacek Brodzki, Sarah J. Campbell, Erik Guentner, Graham A. Niblo, and Nick J. Wright. Property A and CAT(0) cube complexes. *J. Funct. Anal.*, 256(5):1408–1431, 2009.

[BF09]   Mladen Bestvina and Koji Fujiwara. A characterization of higher rank symmetric spaces via bounded cohomology. *Geom. Funct. Anal.*, 19(1):11–40, 2009.

[BF18]   Jonas Beyrer and Elia Fioravanti. Cross ratios and cubulations of hyperbolic groups. *arXiv preprint arXiv:1810.08087*, 2018.

[BF19]   Jonas Beyrer and Elia Fioravanti. Cross ratios on CAT(0) cube complexes and marked length-spectrum rigidity. *arXiv preprint arXiv:1903.02447*, 2019.

[BH99]   Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BHS17a]   Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[BHS17b]   Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *arXiv preprint arXiv:1704.04271*, 2017.

[Bow13]   Brian H. Bowditch. Coarse median spaces and groups. *Pacific Journal of Mathematics*, 261(1):53–93, 2013.

[Bow18]   Brian H. Bowditch. Convex hulls in coarse median spaces. *Preprint*, 2018.

[Bow19]   Brian H. Bowditch. Quasiflats in coarse median spaces, 2019.

[Bre17]   Corey Bregman. Isometry groups of CAT(0) cube complexes. *arXiv preprint arXiv:1712.04805*, 2017.

[Bri91]   Martin R. Bridson. Geodesics and curvature in metric simplicial complexes. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 373–463. World Sci. Publ., River Edge, NJ, 1991.

[BW12]   Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[CC19]   Jérémie Chalopin and Victor Chepoi. 1-safe petri nets and special cube complexes: equivalence and applications. *ACM Transactions on Computational Logic (TOCL)*, 20(3):1–49, 2019.

[CFI16]   Indira Chatterji, Talia Fernós, and Alessandra Iozzi. The median class and superrigidity of actions on CAT(0) cube complexes. *J. Topol.*, 9(2):349–400, 2016. With an appendix by Pierre-Emmanuel Caprace.

[CH17]   Matthew Cordes and David Hume. Stability and the Morse boundary. *J. Lond. Math. Soc. (2)*, 95(3):963–988, 2017.

[Che00]   Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.

[CM19]   Indira Chatterji and Alexandre Martin. A note on the acylindrical hyperbolicity of groups acting on CAT(0) cube complexes. In *Beyond hyperbolicity*, volume 454 of *London Math. Soc. Lecture Note Ser.*, pages 160–178. Cambridge Univ. Press, Cambridge, 2019.

[CN05]   Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005.

[CS11]   Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21(4):851–891, 2011.

[CS15]   Ruth Charney and Harold Sultan. Contracting boundaries of CAT(0) spaces. *Journal of Topology*, 8(1):93–117, 2015.

[Dil50]   Robert P. Dilworth. A decomposition theorem for partially ordered sets. *Ann. of Math. (2)*, 51:161–166, 1950.

[FH19]   Elia Fioravanti and Mark Hagen. Deforming cubulations of hyperbolic groups. *arXiv preprint arXiv:1912.10999*, 2019.

[Fio17]   Elia Fioravanti. Roller boundaries for median spaces and algebras. *arXiv preprint arXiv:1708.01005*, 2017.

[Fio20]   Elia Fioravanti. Automorphisms of contact graphs of CAT(0) cube complexes. *arXiv preprint arXiv:2001.08493*, 2020.

[Gen16]   Anthony Genevois. Acylindrical action on the hyperplanes of a CAT(0) cube complex. *arXiv preprint arXiv:1610.08759*, 2016.

[Gen19a]   Anthony Genevois. Acylindrical hyperbolicity from actions on CAT(0) cube complexes: a few criteria. *New York J. Math*, 25:1214–1239, 2019.

[Gen19b]   Anthony Genevois. A cubical flat torus theorem and some of its applications. *arXiv preprint arXiv:1902.04883*, 2019.

[Gen19c]   Anthony Genevois. Rank-one isometries of CAT(0) cube complexes and their centralisers. *arXiv preprint arXiv:1905.00735*, 2019.

[Gen20]   Anthony Genevois. Contracting isometries of CAT(0) cube complexes and acylindrical hyperbolicity of diagram groups. *Algebraic & Geometric Topology*, 20(1):49–134, 2020.

[Gro81]   Mikhael Gromov. Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math.*, (53):53–73, 1981.

[Gro87]   Mikhail Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Hag07]   Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. *arXiv preprint arXiv:0705.3386*, 2007.

[Hag08]   Frédéric Haglund. Finite index subgroups of graph products. *Geom. Dedicata*, 135:167–209, 2008.

[Hag13]   Mark F. Hagen. The simplicial boundary of a CAT(0) cube complex. *Algebr. Geom. Topol.*, 13(3):1299–1367, 2013.

[Hag14]   Mark F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.

[HK18]   Jingyin Huang and Bruce Kleiner. Groups quasi-isometric to right-angled Artin groups. *Duke Mathematical Journal*, 167(3):537–602, 2018.

[HP15]   Mark F. Hagen and Piotr Przytycki. Cocompactly cubulated graph manifolds. *Israel J. Math.*, 207(1):377–394, 2015.

[HT19]   Mark F. Hagen and Nicholas W. M. Touikan. Panel collapse and its applications. *Groups Geom. Dyn.*, 13(4):1285–1334, 2019.

[Hua17]   Jingyin Huang. Top-dimensional quasiflats in cat (0) cube complexes. *Geometry & Topology*, 21(4):2281–2352, 2017.

[HW08]   Frédéric Haglund and Daniel T Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, 2008.

[HW14]   G. Christopher Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. *Compos. Math.*, 150(3):453–506, 2014.

[HW20]   Mark Hagen and Henry Wilton. Guirardel cores for cubical actions. *In preparation*, 2020.

[Kö50]   Dénes König. *Theorie der endlichen und unendlichen Graphen. Kombinatorische Topologie der Streckenkomplexe.* Chelsea Publishing Co., New York, N. Y., 1950.

[KM12]   Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

[KS19]   Aditi Kar and Michah Sageev. Uniform exponential growth for CAT(0) square complexes. *Algebr. Geom. Topol.*, 19(3):1229–1245, 2019.

[Lea13]   Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013.

[Lev18]   Ivan Levcovitz. Divergence of CAT(0) cube complexes and Coxeter groups. *Algebr. Geom. Topol.*, 18(3):1633–1673, 2018.

[Mie14]   Benjamin Miesch. Injective metrics on cube complexes. *arXiv preprint arXiv:1411.7234*, 2014.

[MM00]   Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[Mor24]   Harold Marston Morse. A fundamental class of geodesics on any closed surface of genus greater than one. *Trans. Amer. Math. Soc.*, 26(1):25–60, 1924.

[Nic04]   Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004.

[NPW81]   Mogens Nielsen, Gordon Plotkin, and Glynn Winskel. Petri nets, event structures and domains, part i. *Theoretical Computer Science*, 13(1):85–108, 1981.

[NS13]   Amos Nevo and Michah Sageev. The Poisson boundary of CAT(0) cube complex groups. *Groups Geom. Dyn.*, 7(3):653–695, 2013.

[QRT19]   Yulan Qing, Kasra Rafi, and Giulio Tiozzo. Sublinearly morse boundary I: Cat (0) spaces. *arXiv preprint arXiv:1909.02096*, 2019.

[Ram29]   Frank P. Ramsey. On a Problem of Formal Logic. *Proc. London Math. Soc. (2)*, 30(4):264–286, 1929.

[Rol98]   Martin Roller. Poc sets, median algebras and group actions. *arXiv preprint arXiv:1607.07747*, 1998.

[Sag95]   Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Sag97]   Michah Sageev. Codimension-1 subgroups and splittings of groups. *Journal of Algebra*, 189(2):377–389, 1997.

[Sag14]   Michah Sageev. CAT(0) cube complexes and groups. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 7–54. Amer. Math. Soc., Providence, RI, 2014.

[Sul14]   Harold Sultan. Hyperbolic quasi-geodesics in CAT(0) spaces. *Geom. Dedicata*, 169:209–224, 2014.

[SW05]   Michah Sageev and Daniel T. Wise. The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.*, 37(5):706–710, 2005.

[Wis12]   Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.

[Wis20]   Daniel T Wise. The structure of groups with a quasiconvex hierarchy. *Preprint*, 2020.

[Woo17]   Daniel J. Woodhouse. A generalized axis theorem for cube complexes. *Algebr. Geom. Topol.*, 17(5):2737–2751, 2017.

[Wri12]   Nick Wright. Finite asymptotic dimension for CAT(0) cube complexes. *Geometry & Topology*, 16(1):527–554, 2012.

[WW17]   Daniel T. Wise and Daniel J. Woodhouse. A cubical flat torus theorem and the bounded packing property. *Israel J. Math.*, 217(1):263–281, 2017.

School of Mathematics, University of Bristol, Bristol, UK

Email address: markfhagen@posteo.net