

CAT(0) cube complexes are determined by their boundary cross ratio

Jonas Beyrer, Elia Fioravanti, and Merlin Incerti-Medici

Abstract. We introduce a \mathbb{Z} -valued cross ratio on Roller boundaries of CAT(0) cube complexes. We motivate its relevance by showing that every cross-ratio preserving bijection of Roller boundaries uniquely extends to a cubical isomorphism. Our results are strikingly general and even apply to infinite dimensional, locally infinite cube complexes with trivial automorphism group.

Mathematics Subject Classification (2020). 20F65, 51F99, 53C24.

Keywords. CAT(0) cube complexes, cross ratios, Roller boundary

1. Introduction

Gromov boundaries of CAT(−1) spaces are naturally endowed with a notion of *cross ratio*. A classical example is provided by the standard projective cross ratio on $\partial_\infty \mathbb{H}^2 \simeq \mathbb{RP}^1$. In the present paper, we introduce a similar object on the *Roller boundary* of any CAT(0) cube complex X and show that this suffices to fully reconstruct the structure of X .

Our motivation essentially comes from two separate points of view. First, the theory of cube complexes has become fundamental within geometric group theory, proving extremely fruitful in relation to various questions stemming from low dimensional topology and group theory. The geometry of many interesting groups is encoded by a CAT(0) cube complex, from classical examples such as right-angled Artin and Coxeter groups, to more recent discoveries like hyperbolic 3-manifold or free-by-cyclic groups [1, 24, 21, 22], to pathological situations such as Thompson’s groups [16] and Higman’s group [27].

As a second perspective, boundary cross ratios provide a valuable tool in the study of negatively and non-positively curved spaces, often appearing in relation to strong rigidity results [30, 7, 8]. If X is a Gromov hyperbolic or CAT(0) space, its *quasi-isometry* type is fully determined by a cross ratio, respectively, on the Gromov or contracting boundary [32, 28, 13]. By contrast, it is an open question

whether these boundary cross ratios can be used to recover the *isometry* type of X (cf. [6, 2]).

When X is the universal cover of a closed, negatively-curved Riemannian manifold, the latter problem turns out to be essentially equivalent to the famous marked-length-spectrum rigidity conjecture (Problem 3.1 in [11]), which has been fully solved only in dimension two [30].

In the present paper, we aim to overcome some of the issues arising in general CAT(0) spaces by relying on the “combinatorial structure” available in cube complexes. This will allow us to define a simpler boundary cross ratio, which in fact only takes integer values. We will show that this cross ratio fully determines the CAT(0) cube complex X up to isometry and, in fact, even up to *cellular* isometries – “cubical isomorphisms” in our terminology.

To this end, we observe that every CAT(0) cube complex X is endowed with two natural metrics: the CAT(0) metric and the ℓ^1 (or *combinatorial*) metric. When X is finite dimensional, these are bi-Lipschitz equivalent. We will restrict our attention to the ℓ^1 metric – which we denote by d – as this enables us to better exploit the cellular structure of X .

It is then natural to consider the horoboundary of the metric space (X, d) , usually known as *Roller boundary* ∂X . This has by now become a standard tool in the study of cube complexes; see e.g. [10, 29, 14, 17, 18] for a (non-exhaustive) list of applications. We remark that, unlike Gromov and visual boundaries, Roller boundaries are always totally disconnected, as is reasonable to expect from objects associated to a cell complex.

By analogy with the CAT(−1) context, it is reasonable to define a *cross ratio*¹ in terms of *Gromov products*² in the metric space (X, d) . The result is a function $\text{cr}: \mathcal{A} \rightarrow \mathbb{Z} \cup \{\pm\infty\}$, defined on a subset $\mathcal{A} \subseteq (\partial X)^4$. We show that this cross ratio admits the alternative expression:

$$\text{cr}(x, y, z, w) = \#\mathcal{W}(x, z|y, w) - \#\mathcal{W}(x, w|y, z),$$

where $\mathcal{W}(x, z|y, w)$ denotes the set of hyperplanes of X that separate x and z from y and w . In particular, it is clear that cr is preserved by the diagonal action of $\text{Aut}(X)$ on $(\partial X)^4$ and independent of any choices in its definition.

Generalising a result of [5] for trees, we then prove the following.

Main Theorem. *Let X and Y be CAT(0) cube complexes with no extremal vertices and not isomorphic to \mathbb{R} . Every cross-ratio preserving bijection $f: \partial X \rightarrow \partial Y$ uniquely extends to a cubical isomorphism $F: X \rightarrow Y$.*

¹For us, a *cross ratio* is an \mathbb{R} -valued function defined on generic 4-tuples of boundary points and satisfying certain symmetries (see (i)–(iv) in Section 3). This should be compared to analogous notions in [31, 23, 25].

²See Definition III.H.1.19 in [9] or Section 3 below for a definition.

Note that the theorem does not require finite dimensional or locally finite cube complexes, nor any group action. The requirement that no vertex be extremal is necessary in order to prevent us from modifying a bounded portion of the cube complex without affecting the boundary; in the case of trees, this would amount to requiring that there are no leaves.

We introduce extremal vertices in Definition 2.2. Absence of extremal vertices can be viewed as an intermediate requirement between the geodesic extension property for the ℓ^1 and CAT(0) metrics.³ The Main Theorem holds more generally when every vertex satisfies Lemma 4.15.

In [4, 3], the Main Theorem is extended to cross-ratio preserving bijections between much smaller subsets of the Roller boundaries. This has applications to length-spectrum rigidity questions for actions on cube complexes. The price to pay is that stronger assumptions need to be imposed on X and Y .

We now briefly sketch the strategy of proof of the Main Theorem. To any three points $x, y, z \in \partial X$, we can associate two well-known objects: the *interval* $I(x, y) \subseteq X \cup \partial X$ and the *median* $m(x, y, z) \in I(x, y)$; see e.g. [29]. These should be interpreted, respectively, as the union of all infinite geodesics between x and y and as a barycentre for the triangle xyz . It is a natural attempt to define the map $F: X \rightarrow Y$ as $F(v) := m(f(x), f(y), f(z))$, assuming for simplicity that there exist points $x, y, z \in \partial X$ with $v = m(x, y, z)$.

However, even relying on the assumption that f preserves cross ratios, it is a priori unclear whether F is well-defined, i.e. independent of the choice of x, y, z . As an illustration of this, consider the two cube complexes X and Y pictured in in Figure 1. In both cases, the points x, y, z, z' lie in the Roller boundary and satisfy $\text{cr}(x, y, z, z') = 0$; the same holds if we permute the four points. In other words, cross ratios involving only x, y, z and z' cannot tell the two cases apart, even though we have $m(x, y, z) = m(x, y, z')$ in X and $m(x, y, z) \neq m(x, y, z')$ in Y .




Figure 1. The cube complex X , on the left, is a tree with a single branch point and four boundary points x, y, z and z' . Pictured on the right is a portion of $Y \simeq \mathbb{R}^3$; the points x, y, z and z' lie in the Roller boundary $\partial\mathbb{R}^3$.

³ For complete cube complexes, it is well-known that the CAT(0) metric satisfies the geodesic extension property if and only if X has no free faces (Proposition II.5.10 in [9]). Such spaces do not have extremal vertices (Remark 2.5) and every cube complex without extremal vertices has the geodesic extension property with respect to the ℓ^1 metric (Lemma 2.4).

We resolve the problem by only representing $v = m(x, y, z)$ with triples (x, y, z) where v disconnects the interval $I(x, y)$; in this case, we say that x and y are *opposite* with respect to v . Examining Figure 1, it is easy to see that x and y are opposite in X , but not in Y . It can be shown that most vertices v are of the form $v = m(x, y, z)$ for a triple (x, y, z) such that x and y are opposite (Lemma 4.15) and, moreover, such triples can be characterised in terms of cross ratios (Lemma 4.4).

We conclude the introduction by remarking that the Main Theorem does not generalise to cross-ratio preserving *embeddings* $\partial X \hookrightarrow \partial Y$. This stands in contrast with the behaviour of trees [5] and rank-one symmetric spaces [7]. A simple counterexample is provided by the cube complexes in Figure 1 and the map $\partial X \hookrightarrow \partial Y = \partial \mathbb{R}^3$ that pairs points of the same name.

For a counterexample involving *cocompact* spaces, the above can be adapted as follows. Let X be the 4-regular tree T_4 with each edge divided into three edges of length 1. Let Y be the 4-regular tree T_4 with each vertex blown up to a 3-cube as in Figure 1; thus, for every vertex of T_4 , there is a 3-cube in Y and, for every edge of T_4 , there is an edge of Y joining two cubes. It is not hard to check that the natural homeomorphism $\partial X \rightarrow \partial Y$ is cross-ratio preserving. Now, if we denote by S the universal cover of the Salvetti complex of $\mathbb{Z}^3 * \mathbb{Z}$, we can embed isometrically $Y \hookrightarrow S$ as a convex subcomplex. This gives rise to a cross-ratio preserving embedding $\partial X \simeq \partial Y \hookrightarrow \partial S$, which does not extend to an isometric embedding $X \hookrightarrow S$.

Acknowledgements. We thank Cornelia Druțu, Ruth Charney, Tobias Hartnick and Viktor Schroeder for helpful conversations. We also express our gratitude to the organisers of the following conferences, where part of this work was carried out: *Third GEAR (Junior) Retreat*, Stanford, 2017; *Moduli Spaces*, Ventotene, 2017; *Young Geometric Group Theory VII*, Les Diablerets, 2018; *Topological and Homological Methods in Group Theory*, Bielefeld, 2018; *3-Manifolds and Geometric Group Theory*, Luminy, 2018; *Representation varieties and geometric structures in low dimensions*, Warwick, 2018. EF also thanks Viktor Schroeder for a visit at UZH.

Jonas Beyrer and Merlin Incerti-Medici acknowledge support by the Swiss National Science Foundation under Grant 200020\175567. Elia Fioravanti was supported by the Clarendon Fund and the Merton Moussouris Scholarship.

2. Preliminaries

2.1. CAT(0) cube complexes. For an introduction to CAT(0) cube complexes, we refer the reader to [34]. In this subsection, we only recall some of the relevant terminology.

Let X be a simply connected cube complex satisfying Gromov’s no- Δ -condition (see 4.2.C in [20] and Chapter II.5 in [9]). The Euclidean metrics on its cubes fit together to yield a CAT(0) metric on X . In the present paper, however, we prefer to endow X with its *combinatorial metric*. More precisely, if $v, w \in X$ are vertices, $d(v, w)$ is defined as the minimal length of a path connecting v and w within the 1-skeleton of X . When X is finite dimensional, the CAT(0) and combinatorial metrics are bi-Lipschitz equivalent and complete.

Unless specified otherwise, *all points $v \in X$ are implicitly understood to be vertices*; we do not distinguish between X and its 0-skeleton. Throughout the paper, the letter d denotes the combinatorial metric on X . All geodesics are meant with respect to the combinatorial metric d ; in particular, they are sequences of edges. We will nevertheless refer to X by the more familiar expression “CAT(0) cube complex.”

For every vertex $v \in X$, we define a graph $\text{lk}(v)$. Its vertices are the edges of X incident to v ; vertices of $\text{lk}(v)$ are joined by an edge if and only if the corresponding edges of X span a square. We refer to $\text{lk}(v)$ as the *link* of v .

Let $\mathcal{W}(X)$ and $\mathcal{H}(X)$ be, respectively, the sets of hyperplanes and halfspaces of X . We simply write \mathcal{W} and \mathcal{H} when there is no need to specify the cube complex. We denote by \mathfrak{h}^* the complement of the halfspace \mathfrak{h} .

Two distinct hyperplanes are *transverse* if they cross. If $e \subseteq X$ is an edge, we write $\mathfrak{w}(e)$ for the hyperplane dual to e . We say that a hyperplane \mathfrak{w} is *adjacent* to a point $v \in X$ if $\mathfrak{w} = \mathfrak{w}(e)$ for an edge e incident to v .

We will generally confuse geodesics and their images as subsets of X . If $\gamma \subseteq X$ is an (oriented) geodesic, we denote by $\gamma(0)$ its initial vertex and by $\gamma(n)$ its n -th vertex. We refer to bi-infinite geodesics simply as *lines*.

Every geodesic $\gamma \subseteq X$ can be viewed as a collection of edges; distinct edges $e, e' \subseteq \gamma$ must yield distinct hyperplanes $\mathfrak{w}(e)$ and $\mathfrak{w}(e')$. We write $\mathcal{W}(\gamma)$ for the collection of hyperplanes crossed by (the edges of) γ . If two geodesics γ and γ' share an endpoint $v \in X$, their union $\gamma \cup \gamma'$ is again a geodesic if and only if $\mathcal{W}(\gamma) \cap \mathcal{W}(\gamma') = \emptyset$.

Lemma 2.1. *Given $v \in X$ and rays $r_1, r_2 \subseteq X$ based at v , let $\mathcal{W}_i \subseteq \mathcal{W}(r_i)$ denote the subset of hyperplanes adjacent to v . The union $r_1 \cup r_2$ is a line if and only if $\mathcal{W}_1 \cap \mathcal{W}_2 = \emptyset$.*

Proof. If $r_1 \cup r_2$ is not a geodesic, there exists $\mathfrak{w} \in \mathcal{W}(r_1) \cap \mathcal{W}(r_2)$. If \mathfrak{w} is not adjacent to v , let u be a hyperplane closest to v among those that separate v from \mathfrak{w} ; otherwise, let us set $u = \mathfrak{w}$. If there existed a hyperplane u' separating v and u , we would have $d(v, u') < d(v, u)$ and u' would separate v and \mathfrak{w} ; this would contradict our choice of u . We conclude that u is adjacent to v and, since u must lie in both $\mathcal{W}(r_i)$, we have $u \in \mathcal{W}_1 \cap \mathcal{W}_2$. □

A subset $\sigma \subseteq \mathcal{H}$ is an *ultrafilter* if it satisfies the following two conditions:

- (1) given any two halfspaces $h, k \in \sigma$, we have $h \cap k \neq \emptyset$;
- (2) for any hyperplane $w \in \mathcal{W}$, a side of w lies in σ .

We say that σ is a *DCC ultrafilter* if, moreover, every descending chain of halfspaces in σ is finite. For every vertex $v \in X$, we denote by $\sigma_v \subseteq \mathcal{H}$ the set of halfspaces containing v . This is a DCC ultrafilter.

Consider now the map $\iota: X \rightarrow 2^{\mathcal{H}}$ taking each vertex v to the set σ_v . The image $\iota(X)$ coincides with the collection of all DCC ultrafilters. Endowing $2^{\mathcal{H}}$ with the product topology, we can consider the closure $\iota(\bar{X})$, which is precisely the set of all ultrafilters. Equipped with the subspace topology, this is a compact Hausdorff space known as the *Roller compactification* of X ; we denote it by \bar{X} . The *Roller boundary* is $\partial X := \bar{X} \setminus X$.

We prefer to imagine ∂X as a set of points at infinity, rather than a set of ultrafilters. We will therefore write $x \in \partial X$ for points in the Roller boundary and employ the notation $\sigma_x \subseteq \mathcal{H}$ to refer to the ultrafilter representing x .

Although this will not be needed in the present paper, it is interesting to observe that the Roller boundary ∂X is naturally homeomorphic to the horofunction boundary of the metric space (X, d) . This is an unpublished result of U. Bader and D. Guralnik; see [12] or [18] for a proof. Note that, on the other hand, the horofunction boundary with respect to the CAT(0) metric would simply coincide with the visual boundary of X .

Given a (combinatorial) ray $r \subseteq X$ and a hyperplane $w \in \mathcal{W}$, there exists a unique side h of w such that $r \setminus h$ is bounded. The collection of all such halfspaces forms an ultrafilter and we denote by $r^+ \in \partial X$ the corresponding point; we refer to r^+ as the *endpoint at infinity* of r .

Fixing a basepoint $v \in X$, every point of ∂X is of the form r^+ for a ray r based at v . This yields a bijection between points of ∂X and rays based at v , where we need to identify the rays r_1 and r_2 whenever $\mathcal{W}(r_1) = \mathcal{W}(r_2)$. See Proposition A.2 in [19] for details.

Note that, given $v \in X$ and $h \in \mathcal{H}$, we have $v \in h$ if and only if $h \in \sigma_v$. We thus extend the halfspace $h \subseteq X$ to a subset $\bar{h} \subseteq \bar{X}$ by declaring that a point $x \in \partial X$ lies in \bar{h} if and only if $h \in \sigma_x$. In particular, \bar{h} and \bar{h}^* provide a partition of \bar{X} with $\bar{h} \cap X = h$ and $\bar{h}^* \cap X = h^*$. For ease of notation, we will generally omit the overline symbol and will not distinguish between a halfspace $h \subseteq X$ and its extension $\bar{h} \subseteq \bar{X}$.

Given subsets $A, B \subseteq \bar{X}$, we employ the notation:

$$\begin{aligned} \mathcal{H}(A \mid B) &= \{h \in \mathcal{H} \mid B \subseteq h, A \subseteq h^*\}, \\ \mathcal{W}(A \mid B) &= \{w \in \mathcal{W} \mid \text{a side of } w \text{ lies in } \mathcal{H}(A \mid B)\}. \end{aligned}$$

It is immediate from the definitions that $\mathcal{W}(x \mid y) \neq \emptyset$ if and only if $x, y \in \bar{X}$ are distinct. If $u, v \in X$ are vertices, we have $d(u, v) = \#\mathcal{W}(u \mid v)$.

Given $x, y \in \bar{X}$, the *interval* between x and y is the set

$$I(x, y) = \{z \in \bar{X} \mid \mathcal{W}(z \mid x, y) = \emptyset\}.$$

We always have $I(x, x) = \{x\}$. In general, $I(x, y) \cap X$ coincides with the union of all (possibly infinite) geodesics with endpoints x and y . In particular, if $u, v, w \in X$ are vertices, we have $w \in I(u, v)$ if and only if $d(u, v) = d(u, w) + d(w, v)$.

For any three points $x, y, z \in \bar{X}$, there exists a unique $m(x, y, z) \in \bar{X}$ that lies in all three intervals $I(x, y)$, $I(y, z)$ and $I(z, x)$. We refer to $m(x, y, z)$ as the *median* of x, y and z and remark that it is represented by the ultrafilter

$$(\sigma_x \cap \sigma_y) \cup (\sigma_y \cap \sigma_z) \cup (\sigma_z \cap \sigma_x).$$

If $v_1, v_2, v_3 \in X$, the median $m = m(v_1, v_2, v_3)$ is the only vertex satisfying $d(v_i, v_j) = d(v_i, m) + d(m, v_j)$ for all $1 \leq i < j \leq 3$. The operator m determines a continuous map $m: \bar{X}^3 \rightarrow \bar{X}$ that endows \bar{X} with a structure of *median algebra*. See e.g. [33] for a definition of the latter notion.

Given $x, y, z \in \bar{X}$, the median $m = m(x, y, z)$ is the only point of $I(x, y)$ with the property that $m \in I(z, w)$ for every $w \in I(x, y)$. In particular, m is the unique point of $I(x, y)$ that is closest to z . For this reason, we also refer to $m(x, y, z)$ as the *gate-projection* of z to $I(x, y)$.

2.2. Extremal vertices and straight geodesics. Let X be a CAT(0) cube complex. We introduce the following two notions.

Definition 2.2. A vertex $v \in X$ is *extremal* if there exists an edge $e \subseteq X$ incident to v such that any other edge incident to v spans a square with e .

Equivalently, $\text{lk}(v)$ is a cone over one of its vertices. This happens if and only if a neighbourhood of v splits as $[0, 1) \times N$ for a subcomplex $N \subseteq X$.

Definition 2.3. A geodesic $\gamma \subseteq X$ is *straight* if no two hyperplanes in $\mathcal{W}(\gamma)$ are transverse.

Our interest in cube complexes with no extremal vertices is motivated by the following straightforward observation (proof omitted).

Lemma 2.4. *If X has no extremal vertices, every edge can be extended to a straight (bi-infinite) line.*

We say that X is *complete* if there is no infinite ascending chain of cubes (cf. [26]). In particular, finite dimensional cube complexes are always complete. A *free face* in X is a non-maximal cube $c \subseteq X$ that is contained in a unique maximal cube.

Remark 2.5. If X is complete and it has no free faces, then X has no extremal vertices. Indeed, consider a vertex $v \in X$ contained in an edge e . Since X is complete, there exists a maximal cube $c_1 \subseteq X$ containing e . Let $c \subseteq c_1$ be the face such that $v \in c$ and $c_1 \simeq c \times e$. Since c is not a free face, there exists a maximal cube $c_2 \subseteq X$ with $c_1 \cap c_2 = c$. Let $e' \subseteq c_2$ be an edge such that $v \in e'$ and $e' \not\subseteq c$. The edges e and e' do not span a square or e' and c_1 would span a cube properly containing c_1 . Hence v is not an extremal vertex.

Nevertheless, the reader will realise that CAT(0) cube complexes with no extremal vertices are much more common than cube complexes with no free faces. For instance, the (universal cover of) the Davis complex [15] associated to a right-angled Coxeter group G often has free faces,⁴ but it only has extremal vertices when $G \simeq \mathbb{Z}/2\mathbb{Z} \times H$ for a parabolic subgroup H .

3. Cross ratios on cube complexes

Let X be a CAT(0) cube complex with *combinatorial metric* d . Given a base vertex $v \in X$, the *Gromov product* of $x, y \in \bar{X}$ is given by

$$(x \cdot y)_v := \#\mathcal{W}(v \mid x, y) = d(v, m(v, x, y)) \in \mathbb{N} \cup \{+\infty\}.$$

Note that $(x \cdot y)_v = +\infty$ if and only if $m(v, x, y) \in \partial X$. Whenever $x, y \in X$, the above quantity coincides with the usual Gromov product:

$$(x \cdot y)_v = \frac{1}{2} \cdot [d(v, x) + d(v, y) - d(x, y)].$$

The following simple observation can be found as Lemma 2.3 in [4].

Lemma 3.1. *Consider $x, y, z \in \bar{X}$ and $v \in X$.*

- (1) *We have $m(x, y, z) \in X$ if and only if each of the three intervals $I(x, y)$, $I(y, z)$, $I(z, x)$ intersects X .*
- (2) *We have $(x \cdot y)_v < +\infty$ if and only if $I(x, y)$ intersects X .*

Fixing $v \in X$, we consider the subset $\mathcal{A} \subseteq (\bar{X})^4$ of 4-tuples (x, y, z, w) such that at most one of the three values $(x \cdot y)_v + (z \cdot w)_v$, $(x \cdot z)_v + (y \cdot w)_v$ and $(x \cdot w)_v + (y \cdot z)_v$ is infinite. By part (2) of Lemma 3.1, the set \mathcal{A} does not depend on the choice of v . The map $\text{cr}_v: \mathcal{A} \rightarrow \mathbb{Z} \cup \{\pm\infty\}$ defined by

$$\text{cr}_v(x, y, z, w) = (x \cdot z)_v + (y \cdot w)_v - (x \cdot w)_v - (y \cdot z)_v$$

satisfies the following identities for all 4-tuples (x, y, z, w) , (x, y, z, t) , and (x, y, t, w) in \mathcal{A} :

⁴ More precisely, this happens if and only if the defining flag complex has free faces.

- (i) $\text{cr}_v(x, y, z, w) = -\text{cr}_v(y, x, z, w)$;
- (ii) $\text{cr}_v(x, y, z, w) = \text{cr}_v(z, w, x, y)$;
- (iii) $\text{cr}_v(x, y, z, w) = \text{cr}_v(x, y, z, t) + \text{cr}_v(x, y, t, w)$;
- (iv) $\text{cr}_v(x, y, z, w) + \text{cr}_v(y, z, x, w) + \text{cr}_v(z, x, y, w) = 0$.

The next result shows that cr_v is moreover basepoint-independent.

Proposition 3.2. *For every $v \in X$ and every $(x, y, z, w) \in \mathcal{A}$, we have*

$$\text{cr}_v(x, y, z, w) = \#\mathcal{W}(x, z \mid y, w) - \#\mathcal{W}(x, w \mid y, z).$$

Proof. We show that every hyperplane $\mathfrak{w} \in \mathcal{W}$ gives the same contribution to both sides of the equality. Note that

$$\text{cr}_v(x, y, z, w) = \#\mathcal{W}(v \mid x, z) + \#\mathcal{W}(v \mid y, w) - \#\mathcal{W}(v \mid x, w) - \#\mathcal{W}(v \mid y, z).$$

Every $\mathfrak{w} \in \mathcal{W}(x, z \mid y, w)$ contributes to either $\mathcal{W}(v \mid x, z)$ or $\mathcal{W}(v \mid y, w)$ by $+1$, without affecting $\mathcal{W}(v \mid x, w)$ and $\mathcal{W}(v \mid y, z)$. Similarly, every hyperplane $\mathfrak{w} \in \mathcal{W}(x, w \mid y, z)$ decreases $-\mathcal{W}(v \mid x, w) - \mathcal{W}(v \mid y, z)$ by 1 and leaves $\mathcal{W}(v \mid x, z)$ and $\mathcal{W}(v \mid y, w)$ invariant. Thus, it suffices to check that hyperplanes \mathfrak{w} such that $\mathfrak{w} \notin \mathcal{W}(x, z \mid y, w) \sqcup \mathcal{W}(x, w \mid y, z)$ do not affect $\text{cr}_v(x, y, z, w)$.

This is clear if all four points x, y, z and w lie on the same side of \mathfrak{w} , or if $\mathfrak{w} \in \mathcal{W}(x, y \mid z, w)$. The remaining case is when exactly three of the four points lie on one side of \mathfrak{w} . Performing a sequence of moves $(x \leftrightarrow y, z \leftrightarrow w)$ and $(x \leftrightarrow z, y \leftrightarrow w)$, which leave $\text{cr}_v(x, y, z, w)$ invariant, we reduce to the case when $\mathfrak{w} \in \mathcal{W}(x \mid y, z, w)$. If v is not on the same side of \mathfrak{w} as x , the hyperplane \mathfrak{w} does not contribute to any summand of $\text{cr}_v(x, y, z, w)$. Otherwise $\mathfrak{w} \in \mathcal{W}(x, v \mid y, z, w)$; in this case the only contributions to $\text{cr}_v(x, y, z, w)$ arise from $\mathcal{W}(v \mid y, w)$ and $\mathcal{W}(v \mid y, z)$ and they cancel each other. \square

We remark that the right-hand side of the equality in Proposition 3.2 is in general defined on a set strictly larger than \mathcal{A} .

Corollary 3.3. *The map $\text{cr}_v: \mathcal{A} \rightarrow \mathbb{Z} \cup \{\pm\infty\}$ is independent of the choice of v . All automorphisms of X preserve cr_v .*

Definition 3.4. We will write $\text{cr}: \mathcal{A} \rightarrow \mathbb{Z} \cup \{\pm\infty\}$ from now on and refer to it as the *cross ratio* on \bar{X} (or ∂X).

Identities (i) and (ii) imply that $|\text{cr}(x, y, z, w)|$ is invariant under a subgroup of order 8 of $\text{Sym}(\{x, y, z, w\})$. Thus, we only need to record $24/8 = 3$ “meaningful” values for every subset $\{x, y, z, w\}$. These values are precisely the three cross ratios appearing in identity (iv), so they are not independent.

The purpose of Definition 3.5 below is precisely to record simultaneously all cross ratios obtained by permuting coordinates. We first introduce some notation.

Given $a_1, b_1, c_1, a_2, b_2, c_2 \in \mathbb{N} \cup \{+\infty\}$, we declare the triples (a_1, b_1, c_1) and (a_2, b_2, c_2) to be *equivalent* if there exists $n \geq 0$ such that

$$a_i = a_j + n, \quad b_i = b_j + n, \quad c_i = c_j + n,$$

where $\{i, j\} = \{1, 2\}$. The equivalence class of the triple (a, b, c) is denoted by $\llbracket a : b : c \rrbracket$. Note that $\llbracket +\infty : +\infty : +\infty \rrbracket$ is the only class consisting of a single triple; every other equivalence class has a unique representative with at least one zero entry. We also remark that all triples in a given equivalence class have the same infinite entries.

Definition 3.5. Given $x, y, z, w \in \bar{X}$ and $v \in X$, the *cross ratio triple*

$$\text{crt}_v(x, y, z, w)$$

is the equivalence class

$$\llbracket (x \cdot y)_v + (z \cdot w)_v : (x \cdot z)_v + (y \cdot w)_v : (x \cdot w)_v + (y \cdot z)_v \rrbracket.$$

Note that crt_v is always independent of the choice of v . This follows from Corollary 3.3 when $(x, y, z, w) \in \mathcal{A}$ and is clear otherwise. We are therefore allowed to simply write crt .

All entries of a cross ratio triple are nonnegative. The three cross ratios in identity (iv) above are recovered by taking the difference of two entries of the triple. We will employ asterisks $*$ when we do not want to specify a coordinate of $\text{crt}(x, y, z, w)$. For instance, we write $\text{crt}(x, y, z, w) = \llbracket * : 0 : 1 \rrbracket$ rather than $\text{crt}(x, y, z, w) = \llbracket a : 0 : 1 \rrbracket$ and $a \in \mathbb{N} \cup \{+\infty\}$.

Let now Y be another CAT(0) cube complex. We write $\mathcal{A}(X)$, rather than just \mathcal{A} , when it is necessary to specify the cube complex under consideration. The following makes the notion of “cross-ratio preserving” map more precise.

Definition 3.6. A map $f: \partial X \rightarrow \partial Y$ is *Möbius* if, for all $x, y, z, w \in \partial X$ with $(x, y, z, w) \in \mathcal{A}(X)$, we have $(f(x), f(y), f(z), f(w)) \in \mathcal{A}(Y)$ and

$$\text{cr}(f(x), f(y), f(z), f(w)) = \text{cr}(x, y, z, w).$$

The latter happens if and only if $\text{crt}(f(x), f(y), f(z), f(w)) = \text{crt}(x, y, z, w)$ for all $(x, y, z, w) \in \mathcal{A}(X)$.

We remark that a bijection $f: \partial X \rightarrow \partial Y$ is Möbius if and only if its inverse $f^{-1}: \partial Y \rightarrow \partial X$ is.

4. Möbius bijections between Roller boundaries

This section is devoted to the proof of the Main Theorem. Throughout it, let X and Y be CAT(0) cube complexes with no extremal vertices. We moreover consider a Möbius bijection $f: \partial X \rightarrow \partial Y$.

To avoid cumbersome formulas, we will employ the following notation for $x, y, z, w \in \partial X$ and $v \in Y$:

$$(x \cdot y)_v^f = (f(x) \cdot f(y))_v, \quad m^f(x, y, z) := m(f(x), f(y), f(z)),$$

$$\text{crt}^f(x, y, z, w) = \text{crt}(f(x), f(y), f(z), f(w)).$$

4.1. Opposite points. As described in the introduction, the following notion will be crucial to avoid the issues depicted in Figure 1.

Definition 4.1 (Definition 5.2 in [4]). Given $x, y, z \in \bar{X}$, we say that x and y are *opposite* with respect to z (written $x \text{ op}_z y$) if the median $m = m(x, y, z)$ lies in X and $I(x, y) = I(x, m) \cup I(m, y)$.

We will also write $x \text{ op}_z^f y$ with the meaning of $f(x) \text{ op}_{f(z)} f(y)$.

Remark 4.2. Let $x, y, z \in \bar{X}$ be points with $m = m(x, y, z) \in X$; denote by $\mathcal{W}_m \subseteq \mathcal{W}(X)$ the subset of hyperplanes adjacent to m . Lemma 5.1 in [4] shows that we have $x \text{ op}_z y$ if and only if no element of $\mathcal{W}(m \mid x) \cap \mathcal{W}_m$ is transverse to an element of $\mathcal{W}(m \mid y) \cap \mathcal{W}_m$.

Remark 4.3. Consider points $x, y, z \in \bar{X}$ with $x \text{ op}_z y$ and $m = m(x, y, z)$. Given any $w \in \partial X$, we either have $(x \cdot w)_m = 0$ or $(y \cdot w)_m = 0$. Indeed, the gate-projection $m(x, y, w)$ falls either in $I(y, m)$ or in $I(x, m)$.

Our goal is now to show that the property in Definition 4.1 is preserved by the Möbius bijection f . We will rely on the following analogue of Proposition 5.4 in [4].

Lemma 4.4. *Given points $x_1, x_2, y \in \partial X$ with $m(x_1, x_2, y) \in X$, the condition $x_1 \text{ op}_y x_2$ fails if and only if there exists a point $z \in \partial X$ such that $\text{crt}(x_1, x_2, y, z) = \llbracket a : b : c \rrbracket$ and $a < \min\{b, c\} < +\infty$.*

Proof. Setting $m = m(x_1, x_2, y)$, we have

$$\text{crt}_m(x_1, x_2, y, z) = \llbracket (y \cdot z)_m : (x_2 \cdot z)_m : (x_1 \cdot z)_m \rrbracket.$$

If $x_1 \text{ op}_y x_2$ and $z \in \partial X$, Remark 4.3 yields $\min\{(x_2 \cdot z)_m, (x_1 \cdot z)_m\} = 0$ and we cannot have $(y \cdot z)_m < 0$.

Consider instead the case when x_1 and x_2 are not opposite with respect to y . There exist transverse hyperplanes $\mathfrak{w}_i \in \mathcal{W}(m \mid x_i)$ adjacent to m . Denote by $m' \in X$ the vertex with $\mathcal{W}(m \mid m') = \{\mathfrak{w}_1, \mathfrak{w}_2\}$. By Lemma 2.4, there exists a straight ray r such that $r(1) = m'$ and $\mathcal{W}(r(0) \mid r(1)) = \{\mathfrak{w}_1\}$; we set $z = r^+$. Let γ and γ_2 be rays based at $r(0)$ satisfying $\gamma(1) = m, \gamma^+ = y$ and $\gamma_2^+ = x_2$. As $r(0)$,

y and x_2 are all on the same side of w_1 , Lemma 2.1 implies that the unions $\gamma \cup r$ and $\gamma_2 \cup r$ are lines. Hence $(y \cdot z)_m = 0$ and $(x_2 \cdot z)_m = 1$. Since $w_1 \in \mathcal{W}(m \mid x_1, z)$, we also have $(x_1 \cdot z)_m \geq 1$. We conclude that $\text{crt}(x_1, x_2, y, z) = \llbracket 0 : 1 : c \rrbracket$ with $c \geq 1$. \square

Proposition 4.5. *Given $x, y, z \in \partial X$, we have $x \text{ op}_z y$ if and only if $x \text{ op}_z^f y$.*

Proof. Fix a basepoint $v \in X$. Assume that $x \text{ op}_z y$; in particular, we have $m(x, y, z) \in X$. By Lemma 3.1, the latter is equivalent to the Gromov products $(x \cdot y)_v, (y \cdot z)_v$ and $(z \cdot x)_v$ being all finite. In other words, the 4-tuples (x, x, y, y) , (y, y, z, z) and (z, z, x, x) all lie in $\mathcal{A}(X)$. As f is Möbius, it takes these 4-tuples into $\mathcal{A}(Y)$ and we must have $m^f(x, y, z) \in Y$.

Now, if we did not have $x \text{ op}_z^f y$, Lemma 4.4 would yield $w \in \partial Y$ with $\text{crt}(f(x), f(y), f(z), w) = \llbracket a : b : c \rrbracket$ and $a < \min\{b, c\} < +\infty$. In particular $(f(x), f(y), f(z), w) \in \mathcal{A}(X)$ and hence $\text{crt}(x, y, z, f^{-1}(w)) = \llbracket a : b : c \rrbracket$, contradicting Lemma 4.4. Thus $x \text{ op}_z y \Rightarrow x \text{ op}_z^f y$ and the converse implication follows by considering $f^{-1}: B \rightarrow A$. \square

We can use triples of opposite points to obtain a well-defined map $X \rightarrow Y$. We now describe this procedure, culminating in Corollary 4.9 below.

The next three results also appear in [4] as Lemmas 5.21, 5.22 and Proposition 5.23. We include them here along with their proofs for the convenience of the reader, but also because the standing assumptions of [4] are much stronger than the current ones.

Given points $x_1, x_2, x, y_1, y_2, y \in \partial X$ with $x_1 \text{ op}_x x_2$ and $y_1 \text{ op}_y y_2$, we set

$$\begin{aligned} m_x &= m(x_1, x_2, x), & m_y &= m(y_1, y_2, y), \\ m'_x &= m^f(x_1, x_2, x), & m'_y &= m^f(y_1, y_2, y). \end{aligned}$$

Lemma 4.6. *Given $u \in \partial X$ with $(x_1, x_2, x, u) \in \mathcal{A}(X)$, we have*

$$(x_1 \cdot u)_{m_x} = (x_1 \cdot u)_{m'_x}^f, \quad (x_2 \cdot u)_{m_x} = (x_2 \cdot u)_{m'_x}^f, \quad (x \cdot u)_{m_x} = (x \cdot u)_{m'_x}^f.$$

Proof. Observe that

$$\begin{aligned} \text{crt}_{m_x}(x_1, x_2, x, u) &= \llbracket (x \cdot u)_{m_x} : (x_2 \cdot u)_{m_x} : (x_1 \cdot u)_{m_x} \rrbracket, \\ \text{crt}_{m'_x}^f(x_1, x_2, x, u) &= \llbracket (x \cdot u)_{m'_x}^f : (x_2 \cdot u)_{m'_x}^f : (x_1 \cdot u)_{m'_x}^f \rrbracket. \end{aligned}$$

Since $x_1 \text{ op}_x x_2$, Remark 4.3 shows that either $(x_1 \cdot u)_{m_x} = 0$ or $(x_2 \cdot u)_{m_x} = 0$. Since $x_1 \text{ op}_x^f x_2$ by Proposition 4.5, also one among $(x_1 \cdot u)_{m'_x}^f$ and $(x_2 \cdot u)_{m'_x}^f$ must vanish. The equality $\text{crt}(x_1, x_2, x, u) = \text{crt}^f(x_1, x_2, x, u)$ then implies that $(x_1 \cdot u)_{m_x} = (x_1 \cdot u)_{m'_x}^f$, $(x_2 \cdot u)_{m_x} = (x_2 \cdot u)_{m'_x}^f$ and $(x \cdot u)_{m_x} = (x \cdot u)_{m'_x}^f$. \square

Lemma 4.7. *Let $u, v \in \partial X$ be two points such that the 4-tuples (x_1, x_2, u, v) , (x_1, x_2, x, u) and (x_1, x_2, x, v) all lie in $\mathcal{A}(X)$. Then $(u \cdot v)_{m_x} = (u \cdot v)_{m'_x}^f$.*

Proof. We have $\text{crt}(x_1, x_2, u, v) = \text{crt}^f(x_1, x_2, u, v)$. Equating the cross ratio triples $\text{crt}_{m_x}(x_1, x_2, u, v)$ and $\text{crt}_{m'_x}^f(x_1, x_2, u, v)$, we obtain

$$\llbracket (u \cdot v)_{m_x} : b : c \rrbracket = \llbracket (u \cdot v)_{m'_x}^f : b' : c' \rrbracket,$$

where Lemma 4.6 yields $b = b'$ and $c = c'$. Hence $(u \cdot v)_{m_x} = (u \cdot v)_{m'_x}^f$. \square

Proposition 4.8. *We have*

$$d(m_x, m_y) = (y_1 \cdot y_2)_{m_x} + |(y_1 \cdot y)_{m_x} - (y_2 \cdot y)_{m_x}|.$$

Proof. Set $v = m(y_1, y_2, m_x)$. As v is the gate-projection of m_x to the interval $I(y_1, y_2)$, we have

$$d(m_x, m_y) = d(m_x, v) + d(v, m_y),$$

where $d(m_x, v) = (y_1 \cdot y_2)_{m_x}$. Up to exchanging y_1 and y_2 , we can assume that v lies within $I(m_y, y_2)$. Since no element of $\mathcal{W}(v \mid y_2) = \mathcal{W}(m_x, y_1 \mid y_2)$ separates m_x and y , it follows that the set $\mathcal{W}(m_x, y_1 \mid y_2, y)$ is empty. We conclude that $(y_2 \cdot y)_{m_x} = \#\mathcal{W}(m_x \mid y_1, y_2, y)$. On the other hand, observing that $\mathcal{W}(v \mid m_y) = \mathcal{W}(m_x, y_2 \mid y_1, y)$, we have

$$\mathcal{W}(m_x \mid y_1, y) = \mathcal{W}(m_x \mid y_1, y_2, y) \sqcup \mathcal{W}(v \mid m_y)$$

and $(y_1 \cdot y)_{m_x} = (y_2 \cdot y)_{m_x} + d(v, m_y)$. \square

Lemma 4.7 and Proposition 4.8 immediately yield the following.

Corollary 4.9. *Suppose that (x_1, x_2, u, v) and (x_1, x_2, x, u) lie in $\mathcal{A}(X)$ whenever u and v are distinct elements of the set $\{y_1, y_2, y\}$. Then, we have $d(m_x, m_y) = d(m'_x, m'_y)$. In particular, m'_x and m'_y coincide if and only if m_x and m_y do.*

4.2. Straight points. In order to ensure that the hypotheses of Corollary 4.9 are satisfied, we will consider the following class of boundary points.

Definition 4.10. A point $x \in \partial X$ is *straight* if there exists a straight ray $r \subseteq X$ with $r^+ = x$; equivalently, x is an endpoint of a straight line. We denote by $\partial_s X \subseteq \partial X$ the set of straight boundary points.

Observe that two points $x, y \in \partial X$ are endpoints of a straight line γ if and only if the interval $I(x, y) \cap X$ is isomorphic to \mathbb{R} . Indeed, $I(x, y) \cap X$ coincides with γ in this case. The following result characterises such situations in a similar way to Lemma 4.4.

Lemma 4.11. *Two points $x, y \in \partial X$ are endpoints of a straight line if and only if both the following are verified:*

- (1) $I(x, y) \cap X \neq \emptyset$;
- (2) *there do not exist points $z, w \in \partial X$ with $\text{crt}(x, y, z, w) = \llbracket a : b : c \rrbracket$ and $a < \min\{b, c\} < +\infty$.*

Proof. We begin by assuming that x and y are endpoints of a straight line γ . Condition (1) is clearly satisfied and we are now going to prove condition (2) for points $z, w \in \partial X$.

If $m(x, y, z) \in X$, then $x \text{ op}_z y$ and Lemma 4.4 shows that we cannot have $\text{crt}(x, y, z, w) = \llbracket a : b : c \rrbracket$ and $a < \min\{b, c\} < +\infty$. The same happens if $m(x, y, w) \in X$, as can be observed by simply swapping z and w . We are left to examine the situation where $m(x, y, z)$ and $m(x, y, w)$ lie in the boundary; note that both medians must then belong to the set $\{x, y\}$. We can assume that $m(x, y, z) \neq m(x, y, w)$ as otherwise the first coordinate of $\text{crt}(x, y, z, w)$ is infinite. If $m(x, y, z) = x$ and $m(x, y, w) = y$, then $\text{crt}(x, y, z, w) = \llbracket * : \infty : 0 \rrbracket$; otherwise, $\text{crt}(x, y, z, w) = \llbracket * : 0 : \infty \rrbracket$. In all cases $\text{crt}(x, y, z, w)$ is not of the form $\llbracket a : b : c \rrbracket$ with $a < \min\{b, c\} < +\infty$ and condition (2) is satisfied.

We now assume that $I(x, y) \cap X \neq \emptyset$, but x and y are not endpoints of a straight line. We will show that condition (2) fails. The intersection $I(x, y) \cap X$ cannot be one-dimensional or it would be isomorphic to \mathbb{R} . Hence $I(x, y) \cap X$ contains a square s ; denote by \mathfrak{w} and \mathfrak{w}' its hyperplanes. Let v_x and v_y be the vertices of s such that $\{\mathfrak{w}, \mathfrak{w}'\}$ is disjoint from $\mathcal{W}(x | v_x)$ and $\mathcal{W}(y | v_y)$. Lemma 2.4 shows that there exist straight rays r_x and r_y , based at v_x and v_y respectively, such that their first crossed hyperplane is \mathfrak{w} . We set $z = r_x^+$ and $w = r_y^+$.

Lemma 2.1 implies that $(z \cdot w)_{v_x} = 0$, $(x \cdot z)_{v_x} = 0$ and $(y \cdot w)_{v_x} = 1$. Moreover, $(y \cdot z)_{v_x} \geq 1$ as \mathfrak{w} separates v_x from y and z . We conclude that $\text{crt}(x, y, z, w) = \llbracket 0 : 1 : c \rrbracket$ with $c = (x \cdot w)_{v_x} + (y \cdot z)_{v_x} \geq 1$. \square

Proposition 4.12. (1) *We have $x \in \partial_s X$ if and only if $f(x) \in \partial_s Y$.*

- (2) *If $x, y \in \partial X$ are endpoints of a straight line, so are the points $f(x)$ and $f(y)$.*

Proof. As a boundary point is straight if and only if it is an endpoint of a straight line, part (1) follows from part (2). If $x, y \in \partial_s X$ are endpoints of a straight line $\gamma \subseteq X$, we have $I(x, y) \cap X \neq \emptyset$. By part (2) of Lemma 3.1, this is equivalent to the fact that (x, x, y, y) lies in $\mathcal{A}(X)$. We conclude that $I(f(x), f(y)) \cap Y \neq \emptyset$.

Now, if $f(x)$ and $f(y)$ were not endpoints of a straight line, Lemma 4.11 would yield points $z, w \in \partial Y$ with $\text{crt}(f(x), f(y), z, w) = \llbracket a : b : c \rrbracket$ and $a < \min\{b, c\} < +\infty$. However, $\text{crt}(x, y, f^{-1}(z), f^{-1}(w))$ would then have the same form, contradicting Lemma 4.11. \square

The next result is our main motivation for considering straight points.

Lemma 4.13. *Consider $x \in \partial_s X$ and a vertex $v \in X$. Given $y, z \in \partial_s X$ with $(y \cdot z)_v < +\infty$, at least one of the Gromov products $(x \cdot y)_v$, $(x \cdot z)_v$ is finite.*

Proof. Let r_x, r_y and r_z be straight rays representing x, y and z , respectively. As the symmetric difference $\mathcal{W}(r_x) \Delta \mathcal{W}(v | x)$ is contained in $\mathcal{W}(r_x(0) | v)$, the intersection $\mathcal{U}_x = \mathcal{W}(r_x) \cap \mathcal{W}(v | x)$ is cofinite in $\mathcal{W}(v | x)$ and does not contain transverse hyperplanes. The same holds for $\mathcal{U}_y = \mathcal{W}(r_y) \cap \mathcal{W}(v | y)$ and $\mathcal{U}_z = \mathcal{W}(r_z) \cap \mathcal{W}(v | z)$.

If we had $(x \cdot y)_v = (x \cdot z)_v = +\infty$, the set $\mathcal{W}(v | x)$ would have infinite intersection with both $\mathcal{W}(v | y)$ and $\mathcal{W}(v | z)$. In particular, both $\mathcal{U}_x \cap \mathcal{U}_y$ and $\mathcal{U}_x \cap \mathcal{U}_z$ would be infinite. As any hyperplane separating two elements of $\mathcal{W}(r_y)$ must lie in $\mathcal{W}(r_y)$, the intersections $\mathcal{U}_x \cap \mathcal{U}_y$ and $\mathcal{U}_x \cap \mathcal{U}_z$ would then be cofinite in \mathcal{U}_x . Hence $\mathcal{U}_y \cap \mathcal{U}_z$ would be infinite, contradicting the fact that $(y \cdot z)_v < +\infty$. \square

Corollary 4.14. *Given points $x_1, x_2, x, y_1, y_2, y \in \partial_s X$ with $x_1 \text{ op}_x x_2$ and $y_1 \text{ op}_y y_2$, we have*

$$d(m(x_1, x_2, x), m(y_1, y_2, y)) = d(m^f(x_1, x_2, x), m^f(y_1, y_2, y)).$$

Proof. We only need to verify the hypotheses of Corollary 4.9. To this end, let u and v be distinct elements of the set $\{y_1, y_2, y\}$ and fix a basepoint $p \in X$. As the Gromov products $(x_1 \cdot x_2)_p$, $(x_1 \cdot x)_p$ and $(x_2 \cdot x)_p$ are all finite, Lemma 4.13 shows that $(x_1, x_2, x, u) \in \mathcal{A}(X)$. If (x_1, x_2, u, v) does not lie in $\mathcal{A}(X)$, we can assume, up to permuting the points, that either $(x_1 \cdot u)_p = (x_2 \cdot u)_p = +\infty$ or $(x_1 \cdot u)_p = (x_1 \cdot v)_p = +\infty$. As $(x_1 \cdot x_2)_p$ and $(u \cdot v)_p$ are finite, both situations are ruled out by Lemma 4.13. \square

4.3. Skinny vertices. We now address the problem of which vertices $v \in X$ can be represented as median of a triple such as those in Corollary 4.14. Throughout this subsection, X is required to have at least two vertices.

We say that a vertex $v \in X$ is *skinny* if $\text{deg}(v) = 2$. Skinny vertices are exactly those that are “invisible from the boundary,” as we now describe.

Lemma 4.15. *For a vertex $v \in X$, the following are equivalent:*

- (1) v is not skinny;
- (2) there exist $x, y, z \in \partial_s X$ such that $m(x, y, z) = v$ and $x \text{ op}_z y$.

Proof. If there exist $x, y, z \in \partial X$ with $m(x, y, z) = v$, the rays from v to x, y and z must begin with three pairwise distinct edges. Hence $\text{deg}(v) \geq 3$, which shows (2) \implies (1).

Regarding the implication (1) \implies (2), we have $\deg(v) \geq 3$ as soon as v is not skinny. Indeed, $\deg(v) = 0$ can only happen if X is a single point and $\deg(v) = 1$ would violate the assumption that X has no extremal vertices. Let e_1, e_2 and e_3 be pairwise distinct edges incident to v ; as v is not extremal, we can assume that $\mathfrak{w}(e_1)$ and $\mathfrak{w}(e_2)$ are not transverse. Lemma 2.4 allows us to extend each e_i to a straight ray r_i . By Lemma 2.1, each union $r_i \cup r_j$ is a line if $i \neq j$. Setting $x = r_1^+, y = r_2^+$ and $z = r_3^+$, we thus have $m(x, y, z) = v$. As $r_1 \cup r_2$ is straight, we also have $x \operatorname{op}_z y$. \square

Denote by $\mathcal{V} \subseteq X$ the set of skinny vertices. Let $\mathcal{S} \subseteq X$ be the union of all edges intersecting \mathcal{V} . Let $\mathcal{F} \subseteq X$ be the full subcomplex with vertex set $X \setminus \mathcal{V}$. We call \mathcal{S} and \mathcal{F} the *skinny* and *fat parts* of X , respectively. We remark that $X = \mathcal{F} \cup \mathcal{S}$ and that every vertex in $\mathcal{S} \setminus \mathcal{F}$ is skinny. We will employ the notation $\mathcal{V}(X)$, $\mathcal{S}(X)$ and $\mathcal{F}(X)$ when it is necessary to specify the cube complex under consideration.

If $X \not\cong \mathbb{R}$, each connected component of \mathcal{S} is either a straight ray or a straight segment; we refer to these as *skinny rays* and *skinny segments*. Every skinny ray intersects \mathcal{F} at a single vertex; given $v \in \mathcal{F}$, we denote by $\mathcal{R}(v) \subseteq \partial_{\mathcal{S}} X$ the set of endpoints at infinity of skinny rays based at v .

Lemma 4.16. *A vertex $v \in X$ and a point $x \in \partial X$ are endpoints of a skinny ray if and only if there exist $y, z \in \partial_{\mathcal{S}} X$ with the following properties:*

- (1) $m(x, y, z) = v$ and $x \operatorname{op}_z y$;
- (2) for every $w \in \partial_{\mathcal{S}} X \setminus \{x\}$ we have $\operatorname{crt}(x, y, z, w) = \llbracket * : * : 0 \rrbracket$.

Proof. Suppose that v and x are endpoints of a skinny ray r . Lemma 2.4 allows us to extend r to a straight line γ ; let y be the endpoint of γ other than x . As $\deg(v) \geq 3$ by definition, we can construct a straight ray r' based at v and disjoint from γ . Setting $z = (r')^+$, we have $m(x, y, z) = v$ and $x \operatorname{op}_z y$. If $w \in \partial X$ and $w \neq x$, we must have $m(x, y, w) \notin I(x, v) \setminus \{v\}$; Remark 4.3 then shows that $m(x, y, w) \in I(v, y)$ and $(x \cdot w)_v = 0$. As $(y \cdot z)_v$ also vanishes, $\operatorname{crt}(x, y, z, w)$ is of the form $\llbracket * : * : 0 \rrbracket$ as required.

Conversely, suppose that y and z are given satisfying condition (1). If the intersection $I(x, v) \cap X$ is not a skinny ray, it contains a vertex $u \neq v$ with $\deg(u) \geq 3$. Let $\mathfrak{w}_x \in \mathcal{W}(u \mid x)$ and $\mathfrak{w}_v \in \mathcal{W}(u \mid v)$ be hyperplanes adjacent to u . Let e be an edge incident to u and not crossing \mathfrak{w}_x or \mathfrak{w}_v ; let γ be a straight ray extending e and set $w = \gamma^+$. We have $\mathfrak{w}_x \in \mathcal{W}(w \mid x)$ so $w \neq x$. Similarly, we have $\mathfrak{w}_v \in \mathcal{W}(w \mid v)$; this shows that $(x \cdot w)_v > 0$ and, along with $x \operatorname{op}_v y$, it guarantees that $m(x, y, w) \in I(x, v)$ and $(y \cdot w)_v = 0$. Thus condition (2) fails, as $\operatorname{crt}(x, y, z, w) = \llbracket * : 0 : c \rrbracket$ with $c = (x \cdot w)_v > 0$. \square

4.4. The isomorphism and its uniqueness. In this subsection, we complete the proof of the Main Theorem. We now require X and Y to be neither single points, nor isomorphic to \mathbb{R} .

Consider a non-skinny vertex $v \in X$. Lemma 4.15 provides three points $x_1, x_2, x \in \partial_s X$ with $m(x_1, x_2, x) = v$ and $x_1 \text{ op}_x x_2$. We define a map $F: \mathcal{F}(X) \rightarrow \mathcal{F}(Y)$ by setting $F(v) = m^f(x_1, x_2, x)$. Note that F is well-defined and distance-preserving by Corollary 4.14.

Applying the same construction to the inverse $f^{-1}: \partial Y \rightarrow \partial X$, we obtain $H: \mathcal{F}(Y) \rightarrow \mathcal{F}(X)$. By Propositions 4.5 and 4.12, the compositions $F \circ H$ and $H \circ F$ are the identity. We conclude that F is surjective and, in fact, an isometric bijection of fat parts.

Theorem 4.17. *The map F extends to a cubical isomorphism $F: X \rightarrow Y$.*

Proof. Observe that two vertices $v_1, v_2 \in \mathcal{F}(X)$ are endpoints of a skinny segment if and only if there does not exist any $v_3 \in \mathcal{F}(X) \setminus \{v_1, v_2\}$ satisfying $d(v_1, v_2) = d(v_1, v_3) + d(v_3, v_2)$. Thus v_1 and v_2 are endpoints of a skinny segment of length ℓ if and only if $F(v_1)$ and $F(v_2)$ are. We can therefore isometrically extend F over all skinny segments in X .

We are left to deal with skinny rays. We conclude by showing that $f(\mathcal{R}(v)) = \mathcal{R}(F(v))$ for all $v \in \mathcal{F}(X)$. It suffices to prove the inclusion $f(\mathcal{R}(v)) \subseteq \mathcal{R}(F(v))$ and then apply the same argument to f^{-1} .

Consider $x \in \mathcal{R}(v)$ and let $y, z \in \partial_s X$ be the points provided by Lemma 4.16. Since $x \in \partial_s X$, we have $m^f(x, y, z) = F(v)$ and, by Proposition 4.5, also $x \text{ op}_z^f y$. Given $w \in \partial_s Y \setminus \{f(x)\}$, the point $f^{-1}(w)$ lies in the set $\partial_s X \setminus \{x\}$ by Proposition 4.12. Lemma 4.16 shows that $\text{crt}(x, y, z, f^{-1}(w))$ is of the form $\llbracket * : * : 0 \rrbracket$ and, by Lemma 4.13, the 4-tuple $(x, y, z, f^{-1}(w))$ lies in $\mathcal{A}(X)$. Hence $\text{crt}(f(x), f(y), f(z), w) = \llbracket * : * : 0 \rrbracket$ for all $w \in \partial_s Y \setminus \{f(x)\}$. Lemma 4.16 finally implies that $f(x) \in \mathcal{R}(F(v))$. \square

Now, the isomorphism $F: X \rightarrow Y$ extends to an isomorphism of median algebras $\bar{F}: \bar{X} \rightarrow \bar{Y}$. We conclude the proof of the Main Theorem via:

Theorem 4.18. *The map F is the only cubical isomorphism with $\bar{F}|_{\partial X} = f$.*

Proof. The uniqueness of F is clear from our construction. We need to prove that $\bar{F}(x) = f(x)$ for every $x \in \partial X$. First, we suppose that $x \in \partial_s X$.

Let γ be a straight line with an endpoint at x ; denote by y the other endpoint of γ . We can assume that x is not endpoint of a skinny ray, as \bar{F} and f clearly coincide on those. Thus, there exist vertices $v_n \in \gamma$ with $\text{deg}(v_n) \geq 3$ and $v_n \rightarrow x$; we can moreover assume that $v_{n+1} \in I(v_n, x) \setminus \{v_n\}$.

Let e_n be an edge with $e_n \cap \gamma = \{v_n\}$. Extending e_n to a straight ray, we construct $z_n \in \partial_s X$ with $m(x, y, z_n) = v_n$ and $x \text{ op}_{z_n} y$. Note that $\text{crt}_{v_0}(x, y, z_0, z_n) = \llbracket * : 0 : c_n \rrbracket$, where $c_n = (z_n \cdot x)_{v_0} = d(v_0, v_n)$ is strictly increasing. By construction, $\bar{F}(x)$ is the limit of the sequence $F(v_n) = m^f(x, y, z_n)$. By Proposition 4.12, there exists a straight line $\gamma' \subseteq Y$ with endpoints $f(x)$ and $f(y)$. Now, the fact that $c_n \rightarrow +\infty$ implies that the points $m^f(x, y, z_n) \in \gamma'$ converge to $f(x)$. Hence $f(x) = \bar{F}(x)$.

We are left to handle points $x \in \partial X \setminus \partial_s X$. Suppose for the sake of contradiction that $\bar{F}(x)$ and $f(x)$ are separated by a hyperplane \mathfrak{w} . Consider an edge e crossing \mathfrak{w} and extend e to a straight line γ ; let u be the endpoint of e on the same side of \mathfrak{w} as $\bar{F}(x)$. Name y and z the endpoints of γ so that $\mathfrak{w} \in \mathcal{W}(\bar{F}(x), y \mid f(x), z)$. The situation is portrayed in Figure 2. We are going to construct a point $w \in \partial_s X$ such that $\min\{(\bar{F}(x) \cdot y)_u, (\bar{F}(x) \cdot w)_u\} < +\infty$ and $m(y, z, w) \in X$. We first show how to use w to conclude the proof.




Figure 2. The case when $x \in \partial X \setminus \partial_s X$.


First, observe that we have $(\bar{F}(x) \cdot z)_u = 0$ and $(y \cdot z)_u = 0$. Our choice of w also implies that $(y \cdot w)_u, (z \cdot w)_u$ and at least one among $(\bar{F}(x) \cdot y)_u$ and $(\bar{F}(x) \cdot w)_u$ are finite, so $(\bar{F}(x), y, z, w) \in \mathcal{A}(X)$. As \bar{F} and f coincide on the set $\{y, z, w\} \subseteq \partial_s X$, we have

$$\begin{aligned} \text{cr}(\bar{F}(x), w, y, z) &= \text{cr}(x, \bar{F}^{-1}(w), \bar{F}^{-1}(y), \bar{F}^{-1}(z)) \\ &= \text{cr}(x, f^{-1}(w), f^{-1}(y), f^{-1}(z)) \\ &= \text{cr}(f(x), w, y, z). \end{aligned}$$

On the other hand, observe that we have $(\bar{F}(x) \cdot y)_u \geq (f(x) \cdot y)_u = 0$ and $0 = (\bar{F}(x) \cdot z)_u < (f(x) \cdot z)_u$. Hence $\text{cr}(\bar{F}(x), w, y, z) > \text{cr}(f(x), w, y, z)$, a contradiction.

Regarding the construction of the point w , observe that γ contains at least two vertices v, v' of degree at least 3; this is because $\bar{F}(x)$ and $f(x)$ are not straight and project to different points of γ . We can assume that $v' \in I(v, z)$. Let e_y and e_z be the only edges at v that lie in $I(y, v)$ and $I(v, z)$, respectively; cf. Figure 3. Let moreover \mathcal{E} be the set of edges ϵ at v with $\mathfrak{w}(\epsilon) \in \mathcal{W}(v \mid \bar{F}(x))$. We distinguish three cases.

Case 1: either $(\bar{F}(x) \cdot y)_v < +\infty$ or $\#\mathcal{E} = 1$. It suffices to pick any edge ϵ incident to v and distinct from e_y and e_z ; extending ϵ to a straight ray we obtain $w \in \partial_s X$ with $m(y, z, w) = v$. If $(\bar{F}(x) \cdot y)_v < +\infty$, we are done. If $(\bar{F}(x) \cdot y)_v = +\infty$ and $\#\mathcal{E} = 1$, we have $\mathcal{E} = \{e_y\}$; hence $(\bar{F}(x) \cdot w)_v = 0$ by Lemma 2.1.

Figure 3. The general setup for the construction of the point w .

Case 2: $(\bar{F}(x) \cdot y)_v = +\infty$ and no edge in \mathcal{E} spans a square with e_z . Replacing v with v' , we end up again in the situation where $\#\mathcal{E} = 1$, which was handled in the previous case.

Case 3: $(\bar{F}(x) \cdot y)_v = +\infty$, $\#\mathcal{E} \geq 2$, and some $\epsilon \in \mathcal{E}$ spans a square with e_z . Since v is not an extremal vertex, there exists an edge ϵ' at v that does not span a square with ϵ . In particular, $\epsilon' \neq e_z$ and $\epsilon' \notin \mathcal{E}$, which also ensures that $\epsilon' \neq e_y$. We extend ϵ' to a straight ray r and set $w = r^+$. Lemma 2.1 implies that $m(y, z, w) = v$ and $(\bar{F}(x) \cdot w)_v = 0$. \square

References

- [1] N. Bergeron and D. T. Wise, A boundary criterion for cubulation. *Amer. J. Math.* **134** (2012), no. 3, 843–859. [Zbl 1279.20051](#) [MR 2931226](#)
- [2] J. Beyrer, Cross ratios on boundaries of symmetric spaces and Euclidean buildings. To appear in *Transform. Groups*. Preprint, 2017. [arXiv:1701.09096](#) [math.DG]
- [3] J. Beyrer and E. Fioravanti, Cross ratios and cubulations of hyperbolic groups. Preprint, 2018. [arXiv:1810.08087v3](#) [math.GT]
- [4] J. Beyrer and E. Fioravanti, Cross ratios on CAT(0) cube complexes and marked length-spectrum rigidity. Preprint, 2019. [arXiv:1903.02447v2](#) [math.GT]
- [5] J. Beyrer and V. Schroeder, Trees and ultrametric Möbius structures. *p-Adic Numbers Ultrametric Anal. Appl.* **9** (2017), no. 4, 247–256. [Zbl 1387.53022](#) [MR 3719682](#)
- [6] K. Biswas, On Moebius and conformal maps between boundaries of CAT(−1) spaces. *Ann. Inst. Fourier (Grenoble)* **65** (2015), no. 3, 1387–1422. [Zbl 1328.53051](#) [MR 3449184](#)
- [7] M. Bourdon, Sur le birapport au bord des CAT(−1)-espaces. *Inst. Hautes Études Sci. Publ. Math.* (1996), no. 83, 95–104. [Zbl 0883.53047](#) [MR 1423021](#)
- [8] M. Bourdon, Immeubles hyperboliques, dimension conforme et rigidité de Mostow. *Geom. Funct. Anal.* **7** (1997), no. 2, 245–268. [Zbl 876.53020](#) [MR 1445387](#)
- [9] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften, 319. Springer-Verlag, Berlin, 1999. [Zbl 0988.53001](#) [MR 1744486](#)

- [10] J. Brodzki, S. J. Campbell, E. Guentner, G. A. Niblo, and N. J. Wright, Property A and CAT(0) cube complexes. *J. Funct. Anal.* **256** (2009), no. 5, 1408–1431. [Zbl 1233.20036](#) [MR 2490224](#)
- [11] K. Burns and A. B. Katok, Manifolds with nonpositive curvature. *Ergodic Theory Dynam. Systems* **5** (1985), no. 2, 307–317. In collaboration with W. Ballman, M. Brin, P. Eberlein and R. Osserman. [MR 796758](#)
- [12] P.-E. Caprace and J. Lécureux, Combinatorial and group-theoretic compactifications of buildings. *Ann. Inst. Fourier (Grenoble)* **61** (2011), no. 2, 619–672. [Zbl 1266.51016](#) [MR 2895068](#)
- [13] R. Charney, M. Cordes, and D. Murray, Quasi-Möbius homeomorphisms of Morse boundaries. *Bull. Lond. Math. Soc.* **51** (2019), no. 3, 501–515. [Zbl 1450.20008](#) [MR 3964503](#)
- [14] I. Chatterji, T. Fernós, and A. Iozzi, The median class and superrigidity of actions on CAT(0) cube complexes. *J. Topol.* **9** (2016), no. 2, 349–400. With an appendix by P.-E. Caprace. [Zbl 1387.20032](#) [MR 3509968](#)
- [15] M. W. Davis, *The geometry and topology of Coxeter groups*. London Mathematical Society Monographs Series, 32. Princeton University Press, Princeton, N.J., 2008. [Zbl 1142.20020](#) [MR 2360474](#)
- [16] D. S. Farley, Finiteness and CAT(0) properties of diagram groups. *Topology* **42** (2003), no. 5, 1065–1082. [Zbl 1044.20023](#) [MR 1978047](#)
- [17] T. Fernós, The Furstenberg–Poisson boundary and CAT(0) cube complexes. *Ergodic Theory Dynam. Systems* **38** (2018), no. 6, 2180–2223. [Zbl 1400.37029](#) [MR 3833346](#)
- [18] T. Fernós, J. Lécureux, and F. Mathéus, Random walks and boundaries of CAT(0) cubical complexes. *Comment. Math. Helv.* **93** (2018), no. 2, 291–333. [Zbl 06897414](#) [MR 3811753](#)
- [19] A. Genevois, Contracting isometries of CAT(0) cube complexes and acylindrical hyperbolicity of diagram groups. *Algebr. Geom. Topol.* **20** (2020), no. 1, 49–134. [Zbl 07188017](#) [MR 4071367](#)
- [20] M. Gromov, Hyperbolic groups. In S. M. Gersten (ed.), *Essays in group theory*. Mathematical Sciences Research Institute Publications, 8. Springer-Verlag, New York, 1987, 75–263. [Zbl 0634.20015](#) [MR 919829](#)
- [21] M. F. Hagen and D. T. Wise, Cubulating hyperbolic free-by-cyclic groups: the general case. *Geom. Funct. Anal.* **25** (2015), no. 1, 134–179. [Zbl 1368.20050](#) [MR 3320891](#)
- [22] M. F. Hagen and D. T. Wise, Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *Duke Math. J.* **165** (2016), no. 9, 1753–1813. [Zbl 1398.20051](#) [MR 3513573](#)
- [23] U. Hamenstädt, Cocycles, Hausdorff measures and cross ratios. *Ergodic Theory Dynam. Systems* **17** (1997), no. 5, 1061–1081. [Zbl 0906.58035](#) [MR 1477033](#)
- [24] J. Kahn and V. Markovic, Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)* **175** (2012), no. 3, 1127–1190. [Zbl 1254.57014](#) [MR 2912704](#)

- [25] F. Labourie, Anosov flows, surface groups and curves in projective space. *Invent. Math.* **165** (2006), no. 1, 51–114. [Zbl 1103.32007](#) [MR 2221137](#)
- [26] I. J. Leary, A metric Kan–Thurston theorem. *J. Topol.* **6** (2013), no. 1, 251–284. [Zbl 1343.20044](#) [MR 3029427](#)
- [27] A. Martin, On the cubical geometry of Higman’s group. *Duke Math. J.* **166** (2017), no. 4, 707–738. [Zbl 1402.20054](#) [MR 3619304](#)
- [28] S. C. Mousley and J. Russell, Hierarchically hyperbolic groups are determined by their Morse boundaries. *Geom. Dedicata* **202** (2019), 45–67. [Zbl 07106666](#) [MR 4001807](#)
- [29] A. Nevo and M. Sageev, The Poisson boundary of CAT(0) cube complex groups. *Groups Geom. Dyn.* **7** (2013), no. 3, 653–695. [Zbl 1346.20084](#) [MR 3095714](#)
- [30] J.-P. Otal, Le spectre marqué des longueurs des surfaces à courbure négative. *Ann. of Math. (2)* **131** (1990), no. 1, 151–162. [Zbl 0699.58018](#) [MR 1038361](#)
- [31] J.-P. Otal, Sur la géométrie symplectique de l’espace des géodésiques d’une variété à courbure négative. *Rev. Mat. Iberoamericana* **8** (1992), no. 3, 441–456. [Zbl 0777.53042](#) [MR 1202417](#)
- [32] F. Paulin, Un groupe hyperbolique est déterminé par son bord. *J. London Math. Soc. (2)* **54** (1996), no. 1, 50–74. [Zbl 0854.20050](#) [MR 1395067](#)
- [33] M. A. Roller, Poc sets, median algebras and group actions. Habilitationsschrift. Universität Regensburg, Regensburg, 1998. [arXiv:1607.07747](#) [math.GN]
- [34] M. Sageev, CAT(0) cube complexes and groups. In M. Bestvina, M. Sageev, and K. Vogtmann (eds.), *Geometric group theory*. IAS/Park City Mathematics Series, 21. American Mathematical Society, Providence, R.I., and Institute for Advanced Study (IAS), Princeton, NJ, 2014, 7–54. [Zbl 1440.20015](#) [MR 3329724](#)

Received March 29, 2019

Jonas Beyrer, IHES, 35 route de Chartres, 91440 Bures-sur-Yvette, France

e-mail: beyrer@ihes.fr

Elia Fioravanti, Universität Bonn, Endenicher Allee 60, 53115 Bonn, Germany

e-mail: fioravan@math.uni-bonn.de

Merlin Incerti-Medici, IHES, 35 route de Chartres, 91440 Bures-sur-Yvette, France

e-mail: merlin.medici@gmail.com