msp

Algebraic & Geometric Topology 20 (2020) 1325–1370

Roller boundaries for median spaces and algebras

ELIA FIORAVANTI

We construct compactifications for median spaces with compact intervals, generalising Roller boundaries of CAT(0) cube complexes. Examples of median spaces with compact intervals include all finite-rank median spaces and all proper median spaces of infinite rank. Our methods also apply to general median algebras, where we recover the zero-completions of Bandelt and Meletiou (1993). Along the way, we prove various properties of halfspaces in finite-rank median spaces and a duality result for locally convex median spaces.

20F65; 20F67, 22F50, 51F99, 57M99

1 Introduction

The aim of this paper is to construct a compactification with good median properties for certain classes of median algebras and median spaces. Median algebras were originally introduced in order theory as a common generalisation of dendrites and lattices; they have been extensively studied in relation to semilattices (see eg Bandelt and Hedlíková [1], Isbell [29] and Sholander [37]) and, more recently, in more geometrical terms because of their connections to CAT(0) cube complexes and median spaces (for instance in Bowditch [3; 4] and Roller [34]).

A metric space X is said to be *median* if, for any three points x_1, x_2, x_3 of X, there exists a unique *median*, ie a unique point $m = m(x_1, x_2, x_3) \in X$ such that $d(x_i, x_j) = d(x_i, m) + d(m, x_j)$ for all $1 \leq i < j \leq 3$. In this case, we refer to

Published: 27 May 2020 DOI: 10.2140/agt.2020.20.1325

the induced map $m\colon X^3 \to X$ as the *median map*. The 0–skeleton of any CAT(0) cube complex becomes a median metric space if we endow it with the restriction of the intrinsic path metric of the 1–skeleton. More generally, every real tree and every ultralimit of median spaces is median. The latter include, in particular, all asymptotic cones of cube complexes.

Further examples of median spaces arise from Guirardel cores of pairs of actions on real trees — see Guirardel [24] — and from asymptotic cones of coarse median groups of finite rank — see Bowditch [3] and Zeidler [41]. Examples of the latter are provided by mapping class groups, cubulated groups and most irreducible 3–manifold groups. Finally, we remark that $L^1(X,\mu)$ is a median space for any measure space (X,μ). For additional literature on the subject, see for instance Bowditch [5], Chatterji, Druțu and Haglund [11], Nica [33], Verheul [39] and van de Vel [38].

The theory of median spaces has two essentially distinct flavours. On the one hand, the study of infinite-dimensional median spaces is strongly related to functional analysis. For instance, a locally compact, second countable group admits a (metrically) proper action on a median space if and only if it has the Haagerup property. Similarly, Kazhdan groups are precisely those that can act on median spaces only with bounded orbits; see Chatterji, Druțu and Haglund [11] and Cherix, Martin and Valette [14].

On the other hand, the study of finite-dimensional median spaces (*finite rank*[1] in our terminology) strongly resembles that of CAT(0) cube complexes, with the additional pathologies typical of real trees.

A key feature of cube complexes is that they come equipped with a collection of hyperplanes. These give each CAT(0) cube complex a canonical structure of *space with walls*; see Haglund and Paulin [28]. Conversely, every space with walls can be canonically embedded into a CAT(0) cube complex; see Chatterji and Niblo [13], Nica [32] and Sageev [36]. In fact, the relationship between spaces with walls and cube complexes can be viewed as a form of duality; see [32, Corollary 4.10]. A similar phenomenon arises for more general median spaces, as we now describe.

Spaces with measured walls were introduced by Cherix, Martin and Valette [14]; they provide useful characterisations for the Haagerup and Kazhdan properties; see Chatterji, Druțu and Haglund [11], Cherix, Martin and Valette [14] and De Cornulier, Tessera and Valette [16]. Each space with measured walls Z can be canonically embedded into

[1] For connected median spaces, the *rank* can be defined as the supremum of the topological dimensions of locally compact subsets; see Section 2.1.

a median space $\mathcal{M}(Z)$ called its *medianisation*. Conversely, to every median space X there corresponds a canonical set of walls, namely its *convex walls*; this induces a space-with-measured-walls structure on X such that $X \hookrightarrow \mathcal{M}(X)$; see [11].

We prove the following analogue of the duality result available for cube complexes [32, Corollary 4.10]. In particular, this applies to all complete, finite-rank median spaces:

Theorem A *For every complete, locally convex median space X, the inclusion $X \hookrightarrow \mathcal{M}(X)$ is a surjective isometry.*

As in cube complexes, walls split every median space into *halfspaces*. In general, the behaviour of halfspaces can be extremely complicated. For instance, if (X,μ) is a standard probability space with no atoms, every halfspace of $L^1(X,\mu)$ is dense (Example 2.24). This however does not happen in finite-rank median spaces:

Proposition B *In a complete, finite-rank median space every halfspace is either open or closed. If $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_k$ is a chain of halfspaces with $k > 2\cdot\mathrm{rank}(X)$, the closure $\overline{\mathfrak{h}}_k$ is contained in the interior of $\mathfrak{h}_1$.*

Many of the analogies between median spaces and CAT(0) cube complexes resemble those between real trees and simplicial trees. It is thus natural to wonder (see eg [11, Question 1.11]) whether a group acting on a median space with unbounded orbits (resp. properly) must have an action on a CAT(0) cube complex with unbounded orbits (resp. proper).

Both questions have a negative answer. Irreducible lattices in $O(4,1;\mathbb{R}) \times O(3,2;\mathbb{R})$, for instance, do not have property (T), but all their actions on CAT(0) cube complexes fix a point; see Cornulier [15, Theorem 6.14]. Furthermore, Baumslag–Solitar groups $\mathrm{BS}(m,n)$ with $m \neq n$ have the Haagerup property, but do not act properly on any CAT(0) cube complex; see Haglund [27].

We are however unable to answer the previous questions under the additional assumption that the median space be of *finite rank*. Even for groups acting on real trees, it is a delicate matter; see Minasyan [30]. See Chatterji and Druțu [10] for a discussion of similar problems.

In the present paper, we bring the analogies between finite-dimensional CAT(0) cube complexes and finite-rank median spaces one step further, by extending to median spaces the construction of the Roller compactification.

Roller boundaries of CAT(0) cube complexes are implicit in the work of Roller [34], although the definition that is most commonly used today probably first appeared in Brodzki, Campbell, Guentner, Niblo and Wright [7]. They have been profitably used to obtain various interesting results, for instance, without attempting to be exhaustive, in [7], Chatterji, Fernós and Iozzi [12], Fernós [19], Fernós, Lécureux and Mathéus [20] and Nevo and Sageev [31].

It is well known that Roller boundaries of cube complexes can be given the following two equivalent characterisations. We denote by $\mathscr{H}$ the set of halfspaces of the cube complex X and do not distinguish between X and its 0–skeleton.

(1) We can embed X into $2^{\mathscr{H}}$ by mapping each vertex v to the set

$$\sigma_v := \{\mathfrak{h} \in \mathscr{H} \mid v \in \mathfrak{h}\}.$$

The space $2^{\mathscr{H}}$ is compact with the product topology. Thus, the closure $\overline{X}$ of X inside $2^{\mathscr{H}}$ is compact; we refer to it as the *Roller compactification*. The Roller boundary is the set $\partial X := \overline{X} \setminus X$.

(2) An *ultrafilter* on $\mathscr{H}$ is a maximal subset $\sigma \subseteq \mathscr{H}$ such that any two halfspaces in σ intersect. The Roller compactification $\overline{X}$ coincides with the subset of $2^{\mathscr{H}}$ consisting of all ultrafilters. Boundary points correspond to ultrafilters that contain infinite descending chains of halfspaces.

It should be noted that a different definition of the Roller boundary appears in the work of Guralnik [25]; see Genevois [23] for a discussion of this alternative notion.

For a general median space X, we will give four equivalent definitions of the Roller compactification. We sketch them here to illustrate the issues that arise when leaving the discrete world of cube complexes.

(1) As in cube complexes, we denote the set of halfspaces by $\mathscr{H}$. In principle, one could try to define a compactification as we did above, namely by taking the closure of the image of $X \hookrightarrow 2^{\mathscr{H}}$. However, if X is not discrete, this results in a space that is too large and carries little geometrical meaning: the double dual $X^{\circ\circ}$; see [34]. See Remark 2.18 and Example 4.5 for the pathologies that may arise; here we simply remark that the inclusion $X \hookrightarrow X^{\circ\circ}$ needs not be continuous.

Instead, given $x, y \in X$, we denote by $I(x, y)$ the union of all medians $m(x, y, z)$ with $z \in X$; we refer to $I(x, y)$ as the *interval* between x and y. In many interesting cases, intervals are compact. We define the Roller compactification $\overline{X}$ as the closure of the image of the map

$$X \hookrightarrow \prod_{x,y \in X} I(x,y), \quad z \mapsto (m(x,y,z))_{x,y},$$

and we set $\partial X := \overline{X} \setminus X$. For CAT(0) cube complexes, this provides a new characterisation of the customary Roller boundary. A similar construction was considered by Ward [40] for dendrites.

(2) As in cube complexes, we could try to consider the set of ultrafilters on $\mathscr{H}$; this would however result again in the double dual $X^{\circ\circ}$. Instead, we endow $\mathscr{H}$ with a σ–algebra of subsets and a measure; we need a finer σ–algebra than the one in [11]; see Section 3. We then only consider measurable ultrafilters, identifying sets with null symmetric difference. This is an alternative description of the Roller compactification $\overline{X}$ (see Theorem 4.15), although no natural topology arises via this construction.

(3) For a general median algebra M, the *zero-completion* of M was introduced by Bandelt and Meletiou [2]. We will recover the same object from a more geometrical perspective in Section 4.1. When M is the median algebra arising from the median space X, the zero-completion of M is identified with the Roller compactification $\overline{X}$. This shows that $\overline{X}$ has itself a natural median-algebra structure.

(4) Finally, $\overline{X}$ can be naturally identified with the horofunction compactification of X. For 0–skeleta of CAT(0) cube complexes, this is an unpublished result of U Bader and D Guralnik; also see Caprace and Lécureux [8] and Fernós, Lécureux and Mathéus [20].

Each of the four definitions is better suited to the study of a particular aspect of $\overline{X}$. Their interplay yields:

Theorem C *Let X be a complete, locally convex median space with compact intervals. The Roller compactification $\overline{X}$ is a compact, locally convex, topological median algebra. The inclusion $X \hookrightarrow \overline{X}$ is a continuous morphism of median algebras, whose image is convex and dense. It is an embedding if, in addition, X is connected and locally compact.*

The class of locally convex median spaces with compact intervals encompasses both complete, finite-rank median spaces and (possibly infinite-dimensional) CAT(0) cube complexes. For the latter, we recover the usual Roller compactification. We also remark that every complete, connected, locally compact median space is proper[2] and thus has

[2] See Bowditch [5, Lemma 4.6] and Bridson and Haefliger [6, Proposition I.3.7].

compact intervals; examples of such spaces that are locally convex and infinite-rank appear eg in Chatterji and Druţu [10].

As for CAT(0) cube complexes, Roller compactifications of median spaces are endowed with an extended metric $d\colon \overline{X} \times \overline{X} \to [0, +\infty]$. Thus, they are partitioned into *components*, namely maximal subsets of points, any two of which are at finite distance. The space X always forms an entire component; moreover:

Theorem D *Let X be complete and finite rank. Every component of ∂X is a complete median space of strictly lower rank.*

In subsequent work, Theorem D will allow us to prove a number of results by induction on the rank. In [21], we use Roller boundaries to extend to finite-rank median spaces the machinery developed by Caprace and Sageev [9] and part of the theory of unidirectional boundary sets (UBS) of Hagen [26]. As a consequence, we obtain in [21] a version of the Tits alternative for groups acting freely on finite-rank median spaces.

In [22], we generalise to finite-rank median spaces the superrigidity result of Chatterji, Fernós and Iozzi [12]. As a consequence, if Γ is an irreducible lattice in a higher-rank semisimple Lie group, every action of Γ on a complete, connected, finite-rank median space must fix a point.

This is in sharp contrast to the behaviour of actions on infinite-rank median spaces. Indeed, Chatterji and Druţu showed in [10] that, for $k, n \geq 2$, all uniform lattices in $(\mathrm{PSL}_k \mathbb{R})^n$ admit proper, cocompact[3] actions on complete, connected, infinite-rank median spaces. Note that this phenomenon is specific to nondiscrete median spaces, as every cocompact cube complex is finite-dimensional.

Structure of the paper In Section 2 we give definitions and basic results. We study convexity, intervals and halfspaces; we prove Proposition B. In Section 3, we construct a σ–algebra and a measure on the set of halfspaces of a median space; we also prove Theorem A. In Section 4.1 we study zero-completions of median algebras; our perspective is different from the one in [2], but we show that our notions are equivalent. Section 4.2 is devoted to Roller compactifications of median spaces; we prove Theorem C there. Finally, we analyse components in Section 4.3 and prove Theorem D.

[3]To be precise, in [10], Chatterji and Druţu only prove that the action is *cobounded*, but it is possible to show that the target median space is proper. The authors have informed me that this stronger result (implying cocompactness) will appear in the next version of their preprint.

Acknowledgements The author warmly thanks Brian Bowditch, Pierre-Emmanuel Caprace, Indira Chatterji, Yves Cornulier, Talia Fernós and Mark Hagen for many helpful conversations and Anthony Genevois for his comments on an earlier version. The author expresses special gratitude to Cornelia Druţu for her excellent supervision and to Talia Fernós for her encouragement to pursue this project. We also thank the referee for many valuable comments.

This work was undertaken at the Mathematical Sciences Research Institute in Berkeley during the Fall 2016 programme "Geometric group theory", where the author was supported by the National Science Foundation under grant no. DMS-1440140 and by the GEAR Network. Part of this work was also carried out at the Isaac Newton Institute for Mathematical Sciences, Cambridge, during the programme "Non-positive curvature, group actions and cohomology" and was supported by EPSRC grant no. EP/K032208/1. The author was also supported by the Clarendon Fund and the Moussouris Scholarship of Merton College.

2 Preliminaries

2.1 Median algebras and median spaces

For definitions and various results on median algebras from a geometric perspective, we refer the reader to [11; 3; 5; 34]. In the following discussion we consider a median algebra M with median map m.

Given $x, y \in M$ the *interval* $I(x,y)$ is the set of points $z \in M$ satisfying $m(x,y,z) = z$. A finite or infinite sequence of points $x_k \in M$ is a *discrete geodesic* if $x_k \in I(x_m, x_n)$ whenever $m < k < n$. A subset $C \subseteq M$ is said to be *convex* if $m(x,y,z) \in C$ whenever $x, y \in C$ and $z \in M$; equivalently $I(x,y) \subseteq C$ for all $x, y \in C$. Any collection of pairwise-intersecting convex subsets has the finite intersection property; this fact is usually known as Helly's theorem; see eg [34, Theorem 2.2].

A *convex halfspace* is a nonempty convex subset $\mathfrak{h} \subseteq M$ whose complement $\mathfrak{h}^* := M \setminus \mathfrak{h}$ is also convex and nonempty. A *convex wall* is an unordered pair $\mathfrak{w} := \{\mathfrak{h}, \mathfrak{h}^*\}$, where $\mathfrak{h}$ is a convex halfspace; we will generally simply speak of *halfspaces* and *walls*, unless we need to avoid confusion with the notions in Section 2.2. The wall $\mathfrak{w}$ *separates* subsets A and B of M if $A \subseteq \mathfrak{h}$ and $B \subseteq \mathfrak{h}^*$ or vice versa. Any two disjoint convex subsets can be separated by a wall; see eg [34, Theorem 2.7].

We will denote the collections of all walls and all halfspaces of M by $\mathscr{W}(M)$ and $\mathscr{H}(M)$, respectively, or simply by $\mathscr{W}$ and $\mathscr{H}$ when the context is clear; there is a natural two-to-one projection $\pi\colon \mathscr{H} \to \mathscr{W}$. Given subsets $A, B \subseteq M$, we write $\mathscr{W}(A \mid B)$ for the set of walls separating them and set

$$\mathscr{H}(A \mid B) := \{\mathfrak{h} \in \mathscr{H} \mid B \subseteq \mathfrak{h} \text{ and } A \subseteq \mathfrak{h}^*\}.$$

We will simply write σ_A for $\mathscr{H}(\varnothing \mid A)$ and confuse the singleton $\{x\}$ with the point x. We will refer to sets of the form $\mathscr{W}(x \mid y)$ and $\mathscr{H}(x \mid y)$ as *wall-intervals* and *halfspace-intervals*, respectively.

A *pocset* $(\mathcal{P}, \preceq, *)$ consists of a poset $(\mathcal{P}, \preceq)$ equipped with an order-reversing involution $*$ such that every element $a \in \mathcal{P}$ is incomparable with a^*. Elements a and b of a pocset are *transverse* if any two elements of the set $\{a, a^*, b, b^*\}$ are incomparable. Considering the median algebra M, the triple $(\mathscr{H}, \subseteq, *)$ is a pocset. Halfspaces $\mathfrak{h}$ and $\mathfrak{k}$ are transverse if and only if all the intersections $\mathfrak{h} \cap \mathfrak{k}$, $\mathfrak{h} \cap \mathfrak{k}^*$, $\mathfrak{h}^* \cap \mathfrak{k}$ and $\mathfrak{h}^* \cap \mathfrak{k}^*$ are nonempty. We say that two walls are transverse if they correspond to transverse halfspaces.

A *partial filter* is a subset $\sigma \subseteq \mathcal{P}$ such that there do not exist $a, b \in \sigma$ with $a \preceq b^*$. If moreover for every $a \in \mathcal{P}$ we have either $a \in \sigma$ or $a^* \in \sigma$, we say that σ is an *ultrafilter*. A *filter* is a partial filter σ such that $b \in \sigma$ whenever $a \preceq b$ and $a \in \sigma$. Some care is needed when comparing our terminology to that of [12], as their notion of "partially defined ultrafilter" coincides with our notion of "filter".

Every partial filter $\sigma \subseteq \mathcal{P}$ is contained in a filter; the smallest such filter is the set of all $b \in \mathcal{P}$ such that $a \preceq b$ for some $a \in \sigma$. Every filter is contained in an ultrafilter; in fact, ultrafilters are precisely filters that are maximal under inclusion. Given any subset $A \subseteq M$, the set $\sigma_A \subseteq \mathscr{H}$ is a filter; it is an ultrafilter if and only if A consists of a single point.

A subset $\sigma \subseteq \mathcal{P}$ is said to be *inseparable* if, whenever $a \preceq b \preceq c$ and $a, c \in \sigma$, we also have $b \in \sigma$. Given a subset $\sigma \subseteq \mathscr{P}$, its *inseparable closure* is the smallest inseparable subset of $\mathcal{P}$ containing σ. It consists precisely of all those $b \in \mathcal{P}$ such that there exist $a, c \in \sigma$ with $a \preceq b \preceq c$. Note that the inseparable closure of a partial filter is again a partial filter; all filters are inseparable.

The set $\{-1, 1\}$ has a unique median-algebra structure. If $k \in \mathbb{N}$, a *k–hypercube* is the median algebra $\{-1, 1\}^k$ given by considering the median map of $\{-1, 1\}$ separately in all coordinates. The *rank* of the median algebra M is the maximal $k \in \mathbb{N}$ such that we can embed a k–hypercube into M; if M has at least two points, we have $\operatorname{rank}(M) \in [1, +\infty]$. The rank of M coincides with the maximal cardinality of a set of pairwise-transverse halfspaces; see [3, Proposition 6.2].

Given a subset $C \subseteq M$, $x \in M$ and $y \in C$, we say that y is a *gate* for (x, C) if $y \in I(x, z)$ for every $z \in C$. We say that $C \subseteq M$ is *gate-convex* if a gate for (x, C) exists for every $x \in M$. If C is gate-convex, there is a unique gate for (x, C) for every $x \in M$; thus we can define a *gate-projection* $\pi_C \colon M \to C$. Gate-convex subsets are always convex, but the converse is not true in general; see Lemma 2.6 for an obstruction. The interval $I(x, y)$ is always gate-convex with gate-projection given by $\pi(z) = m(x, y, z)$; on the other hand, if $C \subseteq I(x, y)$ is gate-convex, we have $C = I(\pi_C(x), \pi_C(y))$.

Proposition 2.1 *A map $\phi \colon M \to M$ is a gate-projection to its image if and only if, for all $x, y, z \in M$, we have $\phi(m(x, y, z)) = m(\phi(x), \phi(y), z)$. In this case, we also have $\phi(m(x, y, z)) = m(\phi(x), \phi(y), \phi(z))$; in particular, gate-projections map intervals to intervals.*

Proof See [1, Proposition 5.1] and [29, 5.8]. Note that "retract" and "Čebyšev ideals" are alternative terminology for "gate-convex subset"; Isbell works in the more general context of isotropic media. □

Lemma 2.2 (1) *If $C_1 \subseteq M$ is convex and $C_2 \subseteq M$ is gate-convex, the projection $\pi_{C_2}(C_1)$ is convex. If moreover $C_1 \cap C_2 \neq \varnothing$, we have $\pi_{C_2}(C_1) = C_1 \cap C_2$.*

(2) *If $C_1, C_2 \subseteq M$ are gate-convex, the sets $\pi_{C_1}(C_2)$ and $\pi_{C_2}(C_1)$ are gate-convex with gate-projections $\pi_{C_1} \circ \pi_{C_2}$ and $\pi_{C_2} \circ \pi_{C_1}$, respectively.*

(3) *If $C_1, C_2 \subseteq M$ are gate-convex and $C_1 \cap C_2 \neq \varnothing$, then $C_1 \cap C_2$ is gate-convex with gate-projection $\pi_{C_1} \circ \pi_{C_2} = \pi_{C_2} \circ \pi_{C_1}$. In particular, if $C_2 \subseteq C_1$, we have $\pi_{C_2} = \pi_{C_2} \circ \pi_{C_1}$.*

(4) *If $C_1, C_2 \subseteq M$ are gate-convex, we have $\pi_{C_1} \circ \pi_{C_2} \circ \pi_{C_1} = \pi_{C_1} \circ \pi_{C_2}$.*

Proof For part (1), see [29, 1.8 and the corollary to 2.5]. Part (2) follows from Proposition 2.1 above and part (3) is an immediate consequence. Part (4) follows from part (3) and the observation that $\pi_{C_1} \circ \pi_{C_2}$ is the gate-projection to $\pi_{C_1}(C_2) \subseteq C_1$. □

Each convex subset $C \subseteq M$ is also a median subalgebra. In particular, we can consider the collection $\mathscr{H}(C)$ of all halfspaces of the median algebra C.

Proposition 2.3 *If $C \subseteq M$ is gate-convex, there is a one-to-one correspondence*

$$\{\mathfrak{h} \in \mathscr{H}(M) \mid \mathfrak{h} \cap C \neq \varnothing \text{ and } \mathfrak{h}^* \cap C \neq \varnothing\} \leftrightarrow \mathscr{H}(C)$$

given by

$$\mathfrak{h} \mapsto \mathfrak{h} \cap C \quad \text{and} \quad \pi_C^{-1}(\mathfrak{k}) \leftarrow\!\shortmid \mathfrak{k}.$$

Moreover, if $\mathfrak{h} \cap C, \mathfrak{k} \cap C \in \mathscr{H}(C)$, then $\mathfrak{h} \subseteq \mathfrak{k}$ if and only if $\mathfrak{h} \cap C \subseteq \mathfrak{k} \cap C$.

Proof Consider $\mathfrak{h} \in \mathscr{H}(M)$. If $\mathfrak{h} \cap C$ and $\mathfrak{h}^* \cap C$ are both nonempty, then they are halfspaces of C. All halfspaces of C arise this way: given a partition $C = C_1 \sqcup C_2$ where the C_i are both convex, we obtain a partition $M = \pi_C^{-1}(C_1) \sqcup \pi_C^{-1}(C_2)$ and Proposition 2.1 ensures that the $\pi_C^{-1}(C_i)$ are also convex. Finally, Lemma 2.2(1) implies that if $\mathfrak{h} \cap C \in \mathscr{H}(C)$, then $\pi_C(\mathfrak{h}) \subseteq \mathfrak{h} \cap C$; equivalently, $\mathfrak{h} \subseteq \pi_C^{-1}(\mathfrak{h} \cap C)$. Similarly $\mathfrak{h}^* \subseteq \pi_C^{-1}(\mathfrak{h}^* \cap C)$, and thus $\mathfrak{h} = \pi_C^{-1}(\mathfrak{h} \cap C)$. □

In view of Proposition 2.3, for all $x, y \in M$, we can make a canonical identification $\mathscr{W}(x \mid y) \simeq \mathscr{W}(I(x,y))$. We will also (slightly improperly) consider the sets $\mathscr{H}(C)$ as subsets of $\mathscr{H}(M)$ from now on. Given subsets $C_1, C_2 \subseteq M$, we say that (x_1, x_2) is a *pair of gates* for (C_1, C_2) if x_1 is a gate for (x_2, C_1) and x_2 is a gate for (x_1, C_2).

Lemma 2.4 *If $C_1, C_2 \subseteq M$ are gate-convex, for every $y_1 \in C_1$ and $y_2 \in C_2$ there exists a pair of gates (x_1, x_2) for (C_1, C_2) such that $y_1x_1x_2y_2$ is a discrete geodesic. Moreover, $\mathscr{H}(C_1 \mid C_2) = \mathscr{H}(x_1 \mid x_2)$.*

Proof Set $x_2 := \pi_{C_2}(y_1)$ and $x_1 := \pi_{C_1}(x_2)$; by Lemma 2.2(2), we have $\pi_{C_2}(x_1) = \pi_{C_2}\pi_{C_1}(x_2) = x_2$. Hence (x_1, x_2) is a pair of gates and the fact that $y_1x_1x_2y_2$ is a discrete geodesic follows from the gate property. Another consequence of x_1 and x_2 being gates is that the sets $\mathscr{H}(x_1 \mid x_2)$, $\mathscr{H}(x_1 \mid C_2)$ and $\mathscr{H}(C_1 \mid x_2)$ coincide; this yields the last part of the lemma. □

Lemma 2.5 *Let $\mathscr{K} \subseteq \mathscr{H}$ be a subset with $\mathscr{K} \cap \mathscr{H}(x \mid y) \neq \varnothing$ for every $x, y \in M$. The rank of M coincides with the maximal cardinality of a set of pairwise-transverse halfspaces in $\mathscr{K}$.*

Proof It suffices to prove that, if $\mathfrak{h}_1, \dots, \mathfrak{h}_{k-1}, \mathfrak{h} \in \mathscr{H}$ are pairwise transverse, there exists $\mathfrak{h}' \in \mathscr{K}$ such that $\mathfrak{h}_1, \dots, \mathfrak{h}_{k-1}, \mathfrak{h}'$ are pairwise transverse. Pick points

$$x \in \mathfrak{h}_1^* \cap \cdots \cap \mathfrak{h}_{k-1}^* \cap \mathfrak{h}^*, \quad y \in \mathfrak{h}_1 \cap \cdots \cap \mathfrak{h}_{k-1} \cap \mathfrak{h}^*,$$
$$u \in \mathfrak{h}_1^* \cap \cdots \cap \mathfrak{h}_{k-1}^* \cap \mathfrak{h}, \quad v \in \mathfrak{h}_1 \cap \cdots \cap \mathfrak{h}_{k-1} \cap \mathfrak{h};$$

these exist by Helly's theorem. The intervals $I := I(x, y)$ and $J := I(u, v)$ are disjoint since $I \subseteq \mathfrak{h}^*$ and $J \subseteq \mathfrak{h}$; thus there exists $\mathfrak{h}' \in \mathscr{H}(I \mid J) \cap \mathscr{K}$, by Lemma 2.4. It is immediate to check that $\mathfrak{h}_1, \dots, \mathfrak{h}_{k-1}, \mathfrak{h}'$ are pairwise transverse. □

A median algebra M is a *topological median algebra* if it is endowed with a Hausdorff topology so that the median map m is continuous. We speak of a *locally convex* median algebra if, in addition, every point has a basis of convex neighbourhoods. A topological median algebra is said to have *compact intervals* if, for every $x, y \in M$, the interval $I(x, y)$ is compact.

Lemma 2.6 *Let M be a topological median algebra M with compact intervals. A convex subset $C \subseteq M$ is gate-convex if and only if it is closed.*

Proof The fact that gate-convex subsets are closed holds in any topological median algebra. Indeed, if C is gate-convex with projection π and $x \notin C$, the points x and $\pi(x)$ are distinct. Setting $I := I(x, \pi(x))$, Lemma 2.2(1) gives $\pi_I(C) = I \cap C = \{\pi(x)\}$. Because the median map m is continuous by definition, the projection π_I is also continuous and $\pi_I^{-1}(I \setminus \{\pi(x)\})$ is an open neighbourhood of x disjoint from C. Hence C is closed.

Now suppose that M has compact intervals and that $C \subseteq M$ is closed and convex. Given $x \in M$, we consider the family $\mathscr{G} := \{I(x, y) \cap C \mid y \in C\}$; by Helly's theorem, any two elements of $\mathscr{G}$ intersect. Another application of Helly's theorem shows that $\mathscr{G}$ has the finite intersection property. By compactness, the intersection of all elements of $\mathscr{G}$ is nonempty. Any point in this intersection is a gate for (x, C); this proves that C is gate-convex. □

Lemma 2.7 *Let M be a compact topological median algebra.*

(1) *Projections to gate-convex sets are continuous.*

(2) *If $C_1, C_2 \subseteq M$ are convex and compact, the convex hull of $C_1 \cup C_2$ is compact.*

Proof Suppose that the projection π to a gate-convex subset $C \subseteq M$ is not continuous. There exist $y \in M$ and a net $(y_j)_{j \in J}$ converging to y such that $\pi(y_j)$ does not converge to $\pi(y)$. By compactness, there exists a subnet $(z_k)_{k \in K}$ such that $(\pi(z_k))_{z \in K}$ converges to a point $z \neq \pi(y)$; since C is closed by Lemma 2.6, we have $z \in C$. Thus, $\pi(z_k) = m(z_k, \pi(y), \pi(z_k))$ converges to $m(y, \pi(y), z) = \pi(y)$ for $k \in K$; this implies $z = \pi(y)$, a contradiction.

For part (2), the map $f\colon M \to M$ given by $f(x) := m(x, \pi_{C_1}(x), \pi_{C_2}(x))$ is continuous by part (1) and Lemma 2.6. The hull of $C_1 \cup C_2$ is precisely the fixed-point set of f; this easily follows from [34, Proposition 2.3] and the gate-property. We conclude that the hull is closed, hence compact. □

Given a metric space X and $x, y \in X$, we denote by $I(x, y)$ the *interval* between x and y, ie the set of points $z \in X$ such that $d(x, y) = d(x, z) + d(z, y)$. We say that X is a *median space* if, for all $x, y, z \in X$, the intersection $I(x, y) \cap I(y, z) \cap I(z, x)$ consists of a single point, which we denote by $m(x, y, z)$. This defines a median-algebra structure on X with the same notion of interval; in particular, we can define rank, convexity and gate-convexity for subsets of X.

A complete median space is geodesic if and only if it is connected; see [5, Lemma 4.6]. In this case, the interval $I(x, y)$ is simply the union of all geodesics with endpoints x and y.

If X is a geodesic median space, its rank coincides with the supremum of the topological dimensions of its locally compact subsets — even when either of the two quantities is infinite; see [3, Theorem 2.2 and Lemma 7.6] for one inequality and [5, Proposition 5.6] for the other. We prefer to speak of *rank*, rather than *dimension*, as the metric spaces that we will be interested in could well be disconnected, eg 0–skeleta of CAT(0) cube complexes.

Example 2.8 Let $(\Omega, \mathscr{B}, \mu)$ be a measure space.

(1) The space $L^1(\Omega, \mu)$ is median when endowed with the metric induced by its norm. The median map is determined by the property that $m(f, g, h)(x)$ is the middle value of $\{f(x), g(x), h(x)\}$ for almost every x and all $f, g, h \in L^1(\Omega, \mu)$.

(2) Given any $E \subseteq \Omega$, the collection $\mathcal{M}_E$ of all $F \subseteq \Omega$ such that $E \triangle F$ is measurable and finite-measure can be given the pseudometric

$$d(A, B) = \mu(A \triangle B).$$

This makes sense because $A \triangle B = (A \triangle E) \triangle (B \triangle E)$. Identifying sets at distance zero, the space $\mathcal{M}_E$ can be isometrically embedded into $L^1(\Omega, \mu)$ by mapping $F \mapsto \mathbb{1}_{F \triangle E}$ and it inherits a median metric. A point lies in the set $m(A, B, C)$ if and only if it lies in at least two of the sets $A, B, C \subseteq \Omega$; the interval $I(A, B)$ can be recognised as the collection of sets Z satisfying $A \cap B \subseteq Z \subseteq A \cup B$.

Let X be a median space throughout the rest of this section. The median map m is 1–Lipschitz if we endow X^3 with the ℓ^1 metric. If $C \subseteq X$ is convex and $x \in X$, a point $z \in C$ is a gate for (x, C) if and only if $d(x, C) = d(x, z)$; gate-projections are 1–Lipschitz. Gate-convex sets are closed and convex; the converse holds in complete median spaces. See [11] for further details and examples.

In complete median spaces we can complement Lemma 2.4 above.

Lemma 2.9 *If X is complete and $C_1, C_2 \subseteq X$ are closed and convex, the points $z_1 \in C_1$ and $z_2 \in C_2$ form a pair of gates for (C_1, C_2) if and only if*

$$d(z_1, z_2) = d(C_1, C_2).$$

In particular, disjoint closed convex sets always have positive distance.

Proof If $d(z_1, z_2) = d(C_1, C_2)$, it is immediate that (z_1, z_2) is a pair of gates. Conversely, given a pair of gates (z_1, z_2), we set $I := I(z_1, z_2)$; if $z_i' \in C_i$, we have $\pi_I(z_i') = z_i$ by Lemma 2.2(1) and the observation that $C_i \cap I = \{z_i\}$. Since π_I is 1–Lipschitz, we have $d(z_1, z_2) \leq d(z_1', z_2')$; hence $d(z_1, z_2) = d(C_1, C_2)$ by the arbitrariness of z_i'. □

Lemma 2.10 *If X has finite rank, it is locally convex.*

Proof Given $x \in X$, $\epsilon > 0$ and $y, z \in B(x, \epsilon)$, we have $I(y, z) \subseteq B(x, 2\epsilon)$. Thus X is "weakly locally convex" in the sense of [3] and we conclude by [3, Lemma 7.1]. □

The class of locally convex median spaces encompasses both finite-rank median spaces and infinite-dimensional CAT(0) cube complexes. An example of a median space that is not locally convex is provided by $L^1([0,1])$: as we shall see in Example 2.24, the convex hull of any nonempty open subset is the entire $L^1([0,1])$.

Lemma 2.11 *Suppose that X is complete and locally convex and let $\{C_i\}_{i \in I}$ be a collection of convex subsets of X with nonempty intersection K. The intersection of $\{\overline{C}_i\}_{i \in I}$ is $\overline{K}$.*

Proof We only need to prove that if $x \in \overline{C}_i$ for all $i \in I$, then $x \in \overline{K}$. Given $\epsilon > 0$, let $N_\epsilon \subseteq B(x, \epsilon)$ be a convex neighbourhood of x; denote by π_ϵ the gate-projection to $\overline{N}_\epsilon$. Since $C_i \cap N_\epsilon \neq \varnothing$ for all $i \in I$, Lemma 2.2(1) implies that $\pi_\epsilon(C_i) = C_i \cap \overline{N}_\epsilon$ and

$$\pi_\epsilon(K) \subseteq \bigcap_i \pi_\epsilon(C_i) \subseteq \left(\bigcap_i C_i\right) \cap \overline{N}_\epsilon = K \cap \overline{N}_\epsilon.$$

Hence K intersects $B(x, 2\epsilon)$ for all $\epsilon > 0$, and thus $x \in \overline{K}$, by the arbitrariness of ϵ. □

2.2 SMW's, PMP's and SMH's

Let X be a set. A *wall* is an unordered pair $\{\mathfrak{h}, \mathfrak{h}^*\}$ corresponding to any partition $X = \mathfrak{h} \sqcup \mathfrak{h}^*$. The wall *separates* subsets $A, B \subseteq X$ if $A \subseteq \mathfrak{h}$ and $B \subseteq \mathfrak{h}^*$ or vice versa. As in Section 2.1, we use the notation $\mathscr{W}(A \mid B)$ to refer to walls separating A and B.

Definition 2.12 We say that the 4–tuple $(X, \mathscr{W}, \mathscr{B}, \mu)$ is a *space with measured walls* (SMW) if $\mathscr{W}$ is a collection of walls of X and the measure μ, defined on the σ–algebra $\mathscr{B} \subseteq 2^{\mathscr{W}}$, satisfies $\mu(\mathscr{W}(x \mid y)) < +\infty$ for all $x, y \in X$.

If $(X, \mathscr{W}, \mathscr{B}, \mu)$ is a space with measured walls, the associated collection of *halfspaces* is the set $\mathscr{H}$ of those $\mathfrak{h} \subseteq X$ such that $\{\mathfrak{h}, \mathfrak{h}^*\} \in \mathscr{W}$. It is endowed with a two-to-one projection $\pi\colon \mathscr{H} \to \mathscr{W}$ given by $\pi(\mathfrak{h}) = \{\mathfrak{h}, \mathfrak{h}^*\}$. We can define a pseudometric on X by setting $\mathrm{pdist}_\mu(x, y) := \mu(\mathscr{W}(x \mid y))$. When this is a genuine metric, we speak of a *faithful* SMW.

Let $(X', \mathscr{W}', \mathscr{B}', \mu')$ be another SMW. Any map $f\colon X \to X'$ such that

$$\{f^{-1}(\mathfrak{h}), f^{-1}(\mathfrak{h}^*)\} \in \mathscr{W} \quad \text{whenever } \{\mathfrak{h}, \mathfrak{h}^*\} \in \mathscr{W}'$$

induces a map $f^*\colon \mathscr{W}' \to \mathscr{W}$. If f^* is measurable and $(f^*)_*\mu' = \mu$, we say that f is a *homomorphism* of spaces with measured walls. Note that this definition differs slightly from the one in [11]. If f is a homomorphism we have

$$\mu(\mathscr{W}(x \mid y)) = \mu'\big((f^*)^{-1}(\mathscr{W}(x \mid y))\big) = \mu'\big(\mathscr{W}(f(x) \mid f(y))\big),$$

so the map $f\colon X \to X'$ preserves the pseudometric. In particular, isomorphisms of SMW's are isometries for the pseudometrics.

Definition 2.13 A *pointed measured pocset* (PMP) is a 4–tuple $(\mathscr{P}, \mathscr{D}, \eta, \sigma)$, where $\mathscr{P}$ is a pocset, $\mathscr{D}$ is a σ–algebra of subsets of $\mathscr{P}$, the measure η is defined on $\mathscr{D}$ and $\sigma \subseteq \mathscr{P}$ is a (not necessarily measurable) ultrafilter.

In analogy to the terminology of [11], we say that an ultrafilter $\sigma' \subseteq \mathscr{P}$ is *admissible* if $\sigma' \triangle \sigma \in \mathscr{D}$ and $\eta(\sigma' \triangle \sigma) < +\infty$. We will denote by $\mathcal{M}(\mathscr{P}, \mathscr{D}, \eta, \sigma)$ (or simply $\mathcal{M}$) the set of admissible ultrafilters associated to the pointed measured pocset $(\mathscr{P}, \mathscr{D}, \eta, \sigma)$. As in Example 2.8, we can equip $\mathcal{M}$ with the median pseudometric $d(\sigma_1, \sigma_2) := \eta(\sigma_1 \triangle \sigma_2)$. We identify admissible ultrafilters at zero distance, so that $\mathcal{M}$ becomes a median space.

Two PMP's $(\mathscr{P},\mathscr{D},\eta,\sigma)$ and $(\mathscr{P}',\mathscr{D}',\eta',\sigma')$ are *isomorphic* if there exists an isomorphism of pocsets $f\colon \mathscr{P}\to\mathscr{P}'$ such that f and f^{-1} are measurable, $f_*\eta=\eta'$ and $\eta(f^{-1}(\sigma')\mathbin{\triangle}\sigma)<+\infty$. Any isomorphism of the two PMP's induces an isometry of the corresponding median spaces $\mathcal{M}$ and $\mathcal{M}'$.

Given a SMW $(X,\mathscr{W},\mathscr{B},\mu)$, we always obtain a PMP $(\mathscr{H},\pi^*\mathscr{B},\pi^*\mu,\sigma_x)$, where $\sigma_x\subseteq\mathscr{H}$ is the set of halfspaces containing x. Here $\pi^*\mathscr{B}$ denotes the σ–algebra $\{\pi^{-1}(E)\mid E\in\mathscr{B}\}$ and $\pi^*\mu(\pi^{-1}(E))=\mu(E)$. The choice of $x\in X$ does not affect the isomorphism type of the PMP.

We simply denote by $\mathcal{M}(X)$ the associated median space of admissible ultrafilters; unlike in [11; 10], for us this is a genuine metric space. We have a pseudo-distance-preserving map $X\to\mathcal{M}(X)$ given by $y\mapsto\sigma_y$. If $(X,\mathscr{W},\mathscr{B},\mu)$ is faithful, this is an isometric embedding. We have recovered the following:

Proposition 2.14 [11] *Any faithful space with measured walls can be isometrically embedded into a median space.*

The embedding is canonical in that every automorphism of X as a SMW extends to an isometry of $\mathcal{M}(X)$. Note however that the restriction of the metric to X does not have to be median, as X might not be a median subalgebra. The following is a partial converse to the previous proposition.

Theorem 2.15 [11] *Let Y be a median space, $\mathscr{W}$ its set of convex walls and $\mathscr{B}$ the σ–algebra generated by wall-intervals. There exists a measure μ on $\mathscr{W}$ such that $\mu(\mathscr{W}(x\mid y))=d(x,y)$ for all $x,y\in Y$. In particular, $(Y,\mathscr{W},\mathscr{B},\mu)$ is a faithful space with measured walls and we have isometric embeddings*

$$Y\hookrightarrow\mathcal{M}(Y)\hookrightarrow L^1(\mathscr{H},\pi^*\mu).$$

Summing up, we can associate to every faithful SMW $(X,\mathscr{W},\mathscr{B},\mu)$ a median space $\mathcal{M}(X)$ and to every median space a faithful SMW. One can wonder whether the compositions

$$\text{SMW}\rightsquigarrow\text{median space}\rightsquigarrow\text{SMW},$$
$$\text{median space}\rightsquigarrow\text{SMW}\rightsquigarrow\text{median space},$$

are the identity. While this has no hope of being true in the first case (we could have taken a set of nonconvex walls of a median space), we will show in Corollary 3.11 that $Y\simeq\mathcal{M}(Y)$ for all locally convex median spaces Y.

We conclude by introducing the following variation on the notion of SMW, which will be more useful to us in the following treatment.

Definition 2.16 We say that the 4–tuple $(X, \mathscr{H}, \mathscr{B}, \nu)$ is a *space with measured halfspaces* (SMH) if $\mathscr{H} \subseteq 2^X$ is a collection of subsets of X closed under taking complements, $\mathscr{B} \subseteq 2^{\mathscr{H}}$ is a σ–algebra and ν is a measure defined on $\mathscr{B}$ satisfying $\nu(\mathscr{H}(x \mid y)) = \nu(\mathscr{H}(y \mid x)) < +\infty$ for all $x, y \in X$.

If $\mathfrak{h} \in \mathscr{H}$, we set $\mathfrak{h}^* := X \setminus \mathfrak{h}$, and if $E \subseteq \mathscr{H}$, we define $E^* := \{\mathfrak{h}^* \mid \mathfrak{h} \in E\}$. We borrow the notation $\mathscr{H}(A \mid B)$ and σ_A from Section 2.1. Note that, if $(X, \mathscr{W}, \mathscr{B}, \mu)$ is a SMW with associated collection of halfspaces $\mathscr{H}$ and projection $\pi\colon \mathscr{H} \to \mathscr{W}$, the 4–tuple $(X, \mathscr{H}, \pi^*\mathscr{B}, \pi^*\mu)$ is *not* a space with measured halfspaces; indeed, $\mathscr{H}(x \mid y) \notin \pi^*\mathscr{B}$ if $x \neq y$.

A pseudometric on X and a notion of homomorphism of SMH's can be defined exactly as we did for SMW's. Given a space with measured halfspaces $(X, \mathscr{H}, \mathscr{B}, \nu)$ and a point $x \in X$, we obtain a PMP $(\mathscr{H}, \mathscr{B}, \nu, \sigma_x)$; we write $\mathcal{M}(X)$ instead of $\mathcal{M}(\mathscr{H}, \mathscr{B}, \nu, \sigma_x)$. The discussion in [11, Section 5] works identically if we replace the symbol $\mathscr{W}$ with $\mathscr{H}$ everywhere. In particular, we have:

Theorem 2.17 *Let Y be a median space, $\mathscr{H}$ the set of convex halfspaces and $\mathscr{B}$ the σ–algebra generated by halfspace-intervals. There exists a measure ν on $\mathscr{H}$ such that $\nu(\mathscr{H}(x \mid y)) = d(x, y)$ for all $x, y \in Y$. In particular, $(Y, \mathscr{H}, \mathscr{B}, \nu)$ is a faithful space with measured halfspaces and we have isometric embeddings*

$$Y \hookrightarrow \mathcal{M}(Y) \hookrightarrow L^1(\mathscr{H}, \nu).$$

Note that $*\colon \mathscr{H} \to \mathscr{H}$ is a measure-preserving involution. We will always denote the σ–algebras in Theorems 2.15 and 2.17 by the same symbol as there is no chance of confusion.

Remark 2.18 Embedding $Y \hookrightarrow \mathcal{M}(Y)$ as in Theorem 2.17, each point $y \in Y$ is represented by all measurable ultrafilters $\sigma \subseteq \mathscr{H}$ that have null symmetric difference with σ_y. Some of these ultrafilters can be "bad": the intersection of all halfspaces in σ can be empty, rather than $\{y\}$.

As an example, consider $Y = \mathbb{R}$ with its standard metric and $y = 0$. The natural ultrafilter σ_0 representing 0 consists of the halfspaces $(-\infty, a)$ for $a > 0$, $(-\infty, a]$ for $a \geq 0$, $(a, +\infty)$ for $a < 0$ and $[a, +\infty)$ for $a \leq 0$. However, $\overline{\sigma}_0 := \big(\sigma_0 \setminus \{(-\infty, 0]\}\big) \cup \{(0, +\infty)\}$ also is an admissible ultrafilter representing 0. Here, measurability and nullity of $\sigma_0 \triangle \overline{\sigma}_0$ follow from the observation that $\{(-\infty, 0]\}$ coincides with the intersection of the countable family of halfspace-intervals $\mathscr{H}(1/n \mid 0)$.

We remark that the measures in Theorems 2.15 and 2.17 can be extended to *complete* σ–algebras $\mathscr{B}_0 \supseteq \mathscr{B}$, ie the smallest σ–algebras with the property that any subset of a null set is measurable; see eg [35, Theorem 1.36].

2.3 Intervals and halfspaces

Let X be a median space throughout this section. It is not hard to use the arguments of [4] to prove the following generalisation of [7, Theorem 1.14]; still, we provide a proof below for the convenience of the reader.

Proposition 2.19 *Let X be complete of rank $r < +\infty$. For every $x, y \in X$, there exists an isometric embedding $I(x, y) \hookrightarrow \mathbb{R}^r$, where $\mathbb{R}^r$ is endowed with the ℓ^1 metric.*

In particular, we obtain the following useful fact (for connected median spaces, also see [4, Corollary 1.3]).

Corollary 2.20 *In a complete, finite-rank median space intervals are compact.*

To prove Proposition 2.19, observe that, if M is a median algebra and $\sigma_1, \sigma_2 \subseteq \mathscr{H}(M)$ are ultrafilters, antichains in the poset $\sigma_1 \setminus \sigma_2$ correspond to sets of pairwise-transverse halfspaces and thus have cardinality bounded above by $\operatorname{rank}(M)$. Hence, Dilworth's theorem [17] yields the following:

Lemma 2.21 *Let M be a median algebra with $\operatorname{rank}(M) = r < +\infty$, and let σ_1 and σ_2 be ultrafilters on $\mathscr{H}$. There is a decomposition $\sigma_1 \setminus \sigma_2 = \mathcal{C}_1 \sqcup \cdots \sqcup \mathcal{C}_k$, where $k \le r$ and each $\mathcal{C}_i$ is nonempty and totally ordered by inclusion.*

Note that, in general, we have no guarantee that the chains provided by the previous lemma are measurable.

Proof of Proposition 2.19 Assume that $X = I(x, y)$ for simplicity. We will produce maps $f_1, \dots, f_r \colon X \to [0, d(x, y)]$ such that for every $u, v \in X$ we have $d(u, v) = \sum |f_i(u) - f_i(v)|$.

We first do this under the assumption that X be finite. By Lemma 2.21, we can decompose $\mathscr{H}(x \mid y) = \mathcal{C}_1 \sqcup \cdots \sqcup \mathcal{C}_r$, where each $\mathcal{C}_i$ is a finite set that is totally ordered by inclusion. Let ν be the measure on $\mathscr{H}$ that is provided by Theorem 2.17. Since X is finite, the singletons of $\mathscr{H}$ are halfspace-intervals (see eg [5, Section 3]), hence measurable; thus, each $\mathcal{C}_i$ is measurable. For $z \in X$ and $1 \leq i \leq r$, we set $f_i(z) := \nu(\mathcal{C}_i \cap \mathscr{H}(x \mid z))$. It is straightforward to check that these define the required embedding.

In the general case, let $\mathfrak{M}$ be the set of finite subalgebras of X containing $\{x, y\}$. Every finite subset of X is contained in an element of $\mathfrak{M}$ by [3, Lemma 4.2]; in particular, $(\mathfrak{M}, \subseteq)$ is a directed set. Every $M \in \mathfrak{M}$ is a finite interval with endpoints x and y, so the previous discussion yields maps $f_1^M, \dots, f_r^M \colon M \to [0, d(x, y)]$ defining an isometric embedding $M \hookrightarrow \mathbb{R}^r$. We can extend each f_i^M to a function $\widetilde{f}_i^M \colon X \to [0, d(x, y)]$ that takes $X \setminus M$ to zero. This defines a net $P_i \colon \mathfrak{M} \to [0, d(x, y)]^X$ for each $1 \leq i \leq r$. The space $[0, d(x, y)]^X$ is compact by Tychonoff's theorem, and thus a subnet of P_i converges. Its limit is a function $f_i \colon X \to [0, d(x, y)]$ and it is immediate to check that $f_1, \dots, f_r$ yield the required embedding $X \hookrightarrow \mathbb{R}^r$. □

We now proceed to examine various properties of the halfspaces of X.

Proposition 2.22 *If X has finite rank and $\mathfrak{h} \in \mathscr{H}$, either $\partial\mathfrak{h} := \overline{\mathfrak{h}} \cap \overline{\mathfrak{h}^*}$ is empty or it is a closed, convex subset with* $\operatorname{rank}(\partial\mathfrak{h}) \leq \operatorname{rank}(X) - 1$.

Proof Follows from Lemma 2.10 above and [3, Lemma 7.5]. □

We remark that, in a median space, closures of convex sets are convex, but interiors of convex sets need not be; in particular, the closure of a halfspace needs not be a halfspace. For instance, consider a real tree T, a branch point $x \in T$ and a connected component $\mathfrak{h}$ of $T \setminus \{x\}$. This is a halfspace of T, as both $\mathfrak{h}$ and $\mathfrak{h}^*$ are convex. The interior of $\mathfrak{h}^*$, however, is not convex, as it coincides with $\mathfrak{h}^* \setminus \{x\}$. In particular, $\overline{\mathfrak{h}} = \mathfrak{h} \cup \{x\}$ is not a halfspace.

Nevertheless, we have the following:

Corollary 2.23 *In a complete, finite-rank median space, each halfspace is either open or closed* (*possibly both*).

Proof We proceed by induction on $\operatorname{rank}(X)$; if the rank is zero, $\mathscr{H} = \varnothing$ and there is nothing to prove. Now assume the result for all median spaces of rank at most $\operatorname{rank}(X) - 1$ and suppose $\mathfrak{h} \in \mathscr{H}(X)$ is neither open nor closed. Then $\partial\mathfrak{h}$ is nonempty and we have a partition $\partial\mathfrak{h} = (\partial\mathfrak{h} \cap \mathfrak{h}) \sqcup (\partial\mathfrak{h} \cap \mathfrak{h}^*)$. By Helly's theorem, the convex set $\partial\mathfrak{h}\cap\mathfrak{h}=\overline{\mathfrak{h}}\cap\overline{\mathfrak{h}^*}\cap\mathfrak{h}$ is nonempty; the same argument yields $\partial\mathfrak{h}\cap\mathfrak{h}^*\neq\varnothing$. The above partition of $\partial\mathfrak{h}$ must then arise from a halfspace of $\partial\mathfrak{h}$ and the inductive hypothesis guarantees that $\partial\mathfrak{h}\cap\mathfrak{h}$ is either open or closed. Note that $\mathfrak{h}=\pi_{\partial\mathfrak{h}}^{-1}(\partial\mathfrak{h}\cap\mathfrak{h})$, by Proposition 2.3. Since $\pi_{\partial\mathfrak{h}}$ is continuous, $\mathfrak{h}$ is either open or closed, a contradiction. □

The situation can be completely different in infinite-rank median spaces:

Example 2.24 The space $X=L^1([0,1])$ is complete, median and all its halfspaces are dense. In order to see the latter, let us write B_R for the R–ball around the origin and let us use the notation $\chi_\bullet$ for characteristic functions. Given a function $f\in L^1([0,1])$ with $\|f\|_1<2R$, there exists a measurable partition $[0,1]=P\sqcup Q$ such that $\|f\cdot\chi_P\|_1<R$ and $\|f\cdot\chi_Q\|_1<R$. By Example 2.8(1), the interval between $f\cdot\chi_P$ and $f\cdot\chi_Q$ contains f. This shows that B_{2R} is contained in the convex hull of B_R for every $R>0$; in particular, the hull of B_R is the entire X. Since any nonempty open subset of X can be translated so that it contains a ball around the origin, we conclude that the hull of any nonempty open set is the entire X. Thus, every proper convex subset of X must have empty interior and all halfspaces are dense.

Example 2.24′ Let us consider the compact topological median algebra $M=\{0,1\}^{\mathbb{N}}$ with the product topology; we identify M with the power set $2^{\mathbb{N}}$ of $\mathbb{N}$. There is a one-to-one correspondence between walls of M and ultrafilters[4] $\mathcal{U}\subseteq 2^{\mathbb{N}}$. Indeed, observe that, given $A,B\subseteq\mathbb{N}$, we have $I(A,B)=\{C\subseteq\mathbb{N}\mid A\cap B\subseteq C\subseteq A\cup B\}$. Since $M=I(\varnothing,\mathbb{N})$, every wall of M has a side $\mathfrak{h}\subseteq 2^{\mathbb{N}}$ containing $\mathbb{N}$ and a side $\mathfrak{h}^*$ containing $\varnothing$. Since $A\cap B\in I(A,B)$, the collection $\mathfrak{h}\subseteq 2^{\mathbb{N}}$ is closed under taking intersections. If $A\subseteq B$ and $A\in\mathfrak{h}$, we have $B\in I(A,\mathbb{N})\subseteq\mathfrak{h}$. Finally, for every $A\subseteq\mathbb{N}$, the collection $\mathfrak{h}$ contains exactly one among A and $\mathbb{N}\setminus A$, as $M=I(A,\mathbb{N}\setminus A)$. This shows that $\mathfrak{h}\subseteq 2^{\mathbb{N}}$ is an ultrafilter. Conversely, it is easy to check that every ultrafilter $\mathcal{U}\subseteq 2^{\mathbb{N}}$ is a halfspace of M.

Now, M has some obvious walls coming from its product structure. The associated halfspaces can be described explicitly by setting one coordinate of $\{0,1\}^{\mathbb{N}}$ to 0 or 1; note that, in a finite product $\{0,1\}^n$, all halfspaces would be of this form. In terms of the correspondence established above, these walls are exactly *principal* ultrafilters (see eg [18, Definition 10.15]). However, it is a well-known consequence of the axiom of

[4] Here we consider the classical set-theoretical notion of ultrafilter, not the one introduced in Section 2.1. See eg [18, Definition 10.12].

choice that there exist also *nonprincipal* ultrafilters. These will correspond to additional walls of M, which — unlike the previous ones — yield dense halfspaces of M.

Endowing the product $X = \prod\{0, 1/n^2\}$ with its ℓ^1 metric, we obtain a compact median metric space. As a median algebra, this is isomorphic to M above, which yields a natural correspondence between walls of X and walls of M. The space X is totally disconnected, but the walls of the *geodesic* median space $Y = \prod[0, 1/n^2]$ are not much better behaved: by [3, Lemma 6.5], every halfspace of X is of the form $\mathfrak{h} \cap Y$ for some $\mathfrak{h} \in \mathscr{H}(Y)$. We remark that — unlike $L^1([0,1])$ — M, X and Y are also locally convex.

Even in finite-rank median spaces, walls can display more complicated behaviours than hyperplanes in CAT(0) cube complexes. Consider for instance the rank-two median space pictured in Figure 1; it is obtained by glueing together three half-planes, each endowed with the ℓ^1 metric. The pictured halfspaces are closed and satisfy $\mathfrak{h} \subsetneq \mathfrak{k}$, but $d(\mathfrak{h}, \mathfrak{k}^*) = 0$; indeed, $\mathfrak{h}$ and $\mathfrak{k}$ share a portion of their frontier isometric to a ray.

Another pathology appears in the space I in Figure 2, which we view as a subset of $\mathbb{R}^2$ with the restriction of the ℓ^1 metric. The halfspaces $\mathfrak{h}$ and $\mathfrak{k}$ are open and $\mathfrak{h} \subsetneq \mathfrak{k}$, but $\overline{\mathfrak{h}} \not\subseteq \mathfrak{k}$.

These issues can easily be circumvented, at least in finite-rank spaces; this is the content of Proposition 2.26 below.

Lemma 2.25 *Let $X = I(x, y)$ be complete and of finite rank. If, for $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}$, we have $y \in \mathfrak{h} \subseteq \mathfrak{k}$ and $d(x, \mathfrak{h}) = d(x, \mathfrak{k})$, then $\overline{\mathfrak{h}} = \overline{\mathfrak{k}}$.*

Proof Observe that the gate-projections of x to $\overline{\mathfrak{k}}$ and $\overline{\mathfrak{h}}$ coincide by Lemma 2.2(3); we denote them by z. If $w \in \overline{\mathfrak{k}}$, the sequence $xzwy$ is a discrete geodesic; thus, $w \in \overline{\mathfrak{h}}$. This proves that $\overline{\mathfrak{k}} \subseteq \overline{\mathfrak{h}}$, while the other inclusion is obvious. □

The hypotheses of Lemma 2.25 do not imply that $\mathfrak{h} = \mathfrak{k}$ or $\overline{\mathfrak{h}} = \mathfrak{k}$; see for instance $\mathfrak{h} \subseteq \mathfrak{k}$ in Figure 2.

Proposition 2.26 *Let X be complete of rank $r < +\infty$ and let $\mathfrak{h}_1 \supsetneq \cdots \supsetneq \mathfrak{h}_k$ be a chain of halfspaces.*

1. *If $d(\mathfrak{h}_1^*, \mathfrak{h}_k) = 0$ and each $\mathfrak{h}_i$ is open, then $k \leq r$.*
2. *In general, if $d(\mathfrak{h}_1^*, \mathfrak{h}_k) = 0$ we have $k \leq 2r$.*
3. *If there exists $x \in \mathfrak{h}_1^*$ such that $d(x, \mathfrak{h}_1) = d(x, \mathfrak{h}_k)$, then $k \leq r + 1$.*

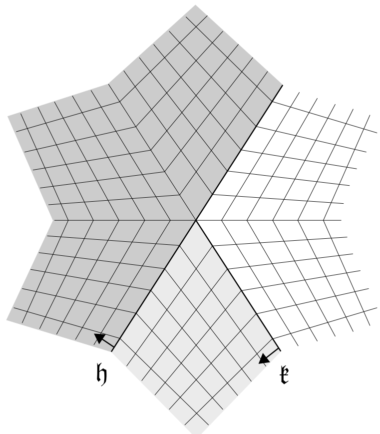

Figure 1

Proof If $d(\mathfrak{h}_1^*, \mathfrak{h}_k) = 0$, no $\mathfrak{h}_i$ can simultaneously be open and closed or we would have $\overline{\mathfrak{h}_1^*} \cap \overline{\mathfrak{h}_k} \subseteq \overline{\mathfrak{h}_i^*} \cap \overline{\mathfrak{h}_i} = \mathfrak{h}_i^* \cap \mathfrak{h}_i = \varnothing$ and $d(\mathfrak{h}_1^*, \mathfrak{h}_k) > 0$ by Lemma 2.9.

We prove part (1) by induction on r; the case $r = 0$ is trivial. If $r \geq 1$, observe that $C := \partial\mathfrak{h}_k$ is closed, convex and nonempty, since $\mathfrak{h}_k$ is not closed. If $i \leq k-1$, we have $\mathfrak{h}_i \supseteq \mathfrak{h}_k$ and $\mathfrak{h}_i \cap \mathfrak{h}_k^* \neq \varnothing$, hence $\mathfrak{h}_i \cap C \neq \varnothing$ by Helly's theorem. Similarly $\mathfrak{h}_i^* \cap C \neq \varnothing$, since $\mathfrak{h}_i^* \subseteq \mathfrak{h}_k^*$ and since by assumption we have $\mathfrak{h}_i^* \cap \overline{\mathfrak{h}_k} \supseteq \mathfrak{h}_1^* \cap \overline{\mathfrak{h}_k} = \overline{\mathfrak{h}_1^*} \cap \overline{\mathfrak{h}_k} \neq \varnothing$.

Proposition 2.3 implies that the $\mathfrak{h}_i \cap C$ are a chain of distinct halfspaces of C, for $i \leq k-1$, and by Proposition 2.22, the rank of C is at most $r-1$. By Helly's theorem and Lemma 2.9, we have $\overline{\mathfrak{h}_1^*} \cap \overline{\mathfrak{h}_{k-1}} \cap C \neq \varnothing$, and by Lemma 2.11, the sets $\overline{\mathfrak{h}_1^* \cap C}$ and $\overline{\mathfrak{h}_{k-1} \cap C}$ intersect. We conclude by applying the inductive hypothesis.

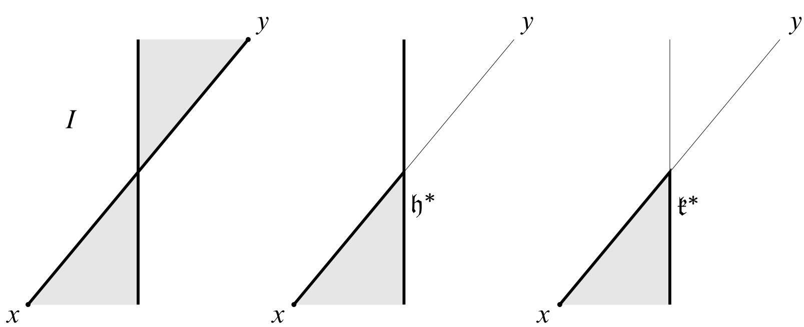

Figure 2

Part (2) is immediate from part (1) and Corollary 2.23, by splitting the chain into a subchain of open halfspaces and a subchain of halfspaces with open complement. To prove part (3), pick a point $y \in \mathfrak{h}_k$ and set $I := I(x,y)$. Note that $d(x,\mathfrak{h}_i) = d(x,\mathfrak{h}_i \cap I)$, since π_I is 1–Lipschitz, fixes x and maps $\mathfrak{h}_i$ onto $\mathfrak{h}_i \cap I$ by Lemma 2.2. An application of Lemma 2.25 yields $\overline{\mathfrak{h}_1 \cap I} = \overline{\mathfrak{h}_k \cap I}$; by Proposition 2.3, the $\mathfrak{h}_i \cap I$ are pairwise distinct so $\mathfrak{h}_i \cap I$ cannot be closed if $i \geq 2$. Corollary 2.23 and Proposition 2.3 imply that $\mathfrak{h}_i$ is open for $i \geq 2$. Moreover, $\overline{\mathfrak{h}_2^* \cap I} \cap \overline{\mathfrak{h}_k \cap I} = \overline{\mathfrak{h}_2^* \cap I} \cap \overline{\mathfrak{h}_2 \cap I} \neq \varnothing$, so $k-1 \leq r$ by part (1). □

Each of the bounds in Proposition 2.26 is sharp. An example with $r = 2$ is given by the median space in Figure 2. One can consider the chain $\mathfrak{h} \subseteq \mathfrak{k}$ for part (1) and $\mathfrak{h} \subseteq \mathfrak{k} \subseteq \overline{\mathfrak{k}}$ for part (3); for part (2) one needs to add $(\tau\mathfrak{h})^*$, where τ is the natural involution of I.

Corollary 2.27 *Let X be complete and finite rank. Every totally ordered subset $\mathcal{C} \subseteq \mathscr{H}$ has a countable subset $\mathcal{C}_0 \subseteq \mathcal{C}$ that is cofinal in $\mathcal{C}$.*

Proof Pick a point $x \in X$ and define $\delta_x\colon \mathcal{C} \to \mathbb{R}$ by $\delta_x(\mathfrak{h}) := d(x,\mathfrak{h}^*)$ if $x \in \mathfrak{h}$ and $\delta_x(\mathfrak{h}) := -d(x,\mathfrak{h})$ if $x \in \mathfrak{h}^*$. The map δ_x is monotone, and by Proposition 2.26(3), it has finite fibres. The image of δ_x is a separable metric space; let $A \subseteq \delta_x(\mathcal{C})$ be a countable dense subset. If $\delta_x(\mathcal{C})$ has a maximum or a minimum, add them to A and set $\mathcal{C}_0 := \delta_x^{-1}(A)$. It is immediate to check that this is cofinal in $\mathcal{C}$ (upwards and downwards). □

3 Measure theory on the pocset of halfspaces

3.1 A finer σ–algebra on the pocset of halfspaces

Let X be a median space. The σ–algebras introduced in Theorems 2.15 and 2.17 are often too restrictive to work comfortably with. As an example, the ultrafilters σ_x for $x \in X$ need not be measurable, for instance when working with nonseparable median spaces. The latter might look like pathological examples, but we remark that, on the contrary, an interesting class of finite-rank median spaces arises from asymptotic cones of coarse median groups and these are rarely separable.

Since the measures that were constructed in Theorems 2.15 and 2.17 originate from Carathéodory's construction, one might wish to consider instead the σ–algebras of additive sets for the outer measures μ^* and ν^*. Unfortunately, even in simple examples, one might obtain measure spaces that are not semifinite. We choose a different path.

We say that a subset $E \subseteq \mathscr{H}$ is *morally measurable* if $E \cap \mathscr{H}(x \mid y)$ lies in $\mathscr{B}_0$ for all $x, y \in X$; here $\mathscr{B}_0$ is the completion of $\mathscr{B}$; see the end of Section 2.2. Morally measurable sets form a σ–algebra $\widehat{\mathscr{B}} \supseteq \mathscr{B}$. For every $z \in X$, the ultrafilter $\sigma_z \subseteq \mathscr{H}$ is morally measurable since $\sigma_z \cap \mathscr{H}(x \mid y) = \mathscr{H}(x \mid m)$, where $m = m(x, y, z)$.

We say that a morally measurable subset $E \subseteq \mathscr{H}$ is *morally null* if $\nu(E \cap \mathscr{H}(x \mid y)) = 0$ for all $x, y \in X$. In particular, subsets of morally null sets are morally null. Given a morally measurable set E, we define

$$\widehat{\nu}(E) := \sup\left\{ \sum_{i \in I} \nu(E \cap \mathscr{H}(x_i \mid y_i)) \;\middle|\; \bigsqcup_{i \in I} \mathscr{H}(x_i \mid y_i) \subseteq \mathscr{H} \right\}.$$

We allow I to be of any cardinality, although restricting to finite or countable index sets would not affect the value of $\widehat{\nu}$. A morally measurable set E is morally null if and only if $\widehat{\nu}(E) = 0$.

We recall that a measure space (Ω, μ) is *semifinite* if every measurable subset $E \subseteq \Omega$ with $\mu(E) = +\infty$ contains a measurable subset $F \subseteq E$ with $0 < \mu(F) < +\infty$.

Proposition 3.1 *The triple* $(\mathscr{H}, \widehat{\mathscr{B}}, \widehat{\nu})$ *is a semifinite measure space.*

We prove Proposition 3.1 below; the nontrivial part is showing that $\widehat{\nu}$ is a measure. Since differences of halfspace-intervals are finite disjoint unions of halfspace-intervals (see [11]), the following is an easy observation.

Lemma 3.2 *Any finite (countable) union of halfspace-intervals is a finite (countable)* ***disjoint*** *union of halfspace-intervals.*

Proof of Proposition 3.1 It suffices to prove that $\widehat{\nu}$ is additive and σ–subadditive. Let $\{E_n\}_{n \geq 0}$ be a collection of pairwise-disjoint, morally measurable sets with union E. Given pairwise-disjoint $\mathscr{H}(x_1 \mid y_1), \dots, \mathscr{H}(x_k \mid y_k)$, we have

$$\sum_{i=1}^{k} \nu(E \cap \mathscr{H}(x_i \mid y_i)) = \sum_{n \geq 0} \sum_{i=1}^{k} \nu(E_n \cap \mathscr{H}(x_i \mid y_i)) \leq \sum_{n \geq 0} \widehat{\nu}(E_n),$$

and this proves the inequality $\widehat{\nu}(E) \leq \sum \widehat{\nu}(E_n)$. Now, if E and F are disjoint morally measurable sets and $\epsilon > 0$, we can find finitely many points $x_i, y_i \in X$ such that $\widehat{\nu}(E) - \epsilon \leq \sum \nu(E \cap \mathscr{H}(x_i \mid y_i))$ and, similarly, points $u_j, v_j \in X$ satisfying an analogous inequality for F. By Lemma 3.2, the union of all $\mathscr{H}(x_i \mid y_i)$ and $\mathscr{H}(u_j \mid v_j)$ can be decomposed as a finite disjoint union of halfspace-intervals $\mathscr{H}(w_k \mid z_k)$. Thus,

$$\begin{aligned}\widehat{\nu}(E)+\widehat{\nu}(F)-2\epsilon &\le \sum_i \nu(E\cap\mathscr{H}(x_i \mid y_i))+\sum_j \nu(F\cap\mathscr{H}(u_j \mid v_j))\\ &\le \sum_k \nu((E\sqcup F)\cap\mathscr{H}(w_k \mid z_k))\le \widehat{\nu}(E\sqcup F),\end{aligned}$$

and by the arbitrariness of ϵ, we obtain $\widehat{\nu}(E)+\widehat{\nu}(F)\le\widehat{\nu}(E\sqcup F)$. We have already shown subadditivity, so $\widehat{\nu}(E)+\widehat{\nu}(F)=\widehat{\nu}(E\sqcup F)$. □

Lemma 3.3 (properties of $\widehat{\nu}$) (1) *For every $E\in\widehat{\mathscr{B}}$ there exist pairwise-disjoint $\mathscr{H}(x_n \mid y_n)$ for $n\ge 0$ such that $\widehat{\nu}(E)=\sum\nu(E\cap\mathscr{H}(x_n \mid y_n))$.*

(2) *$\widehat{\nu}(E)\le\nu(E)$ for all $E\in\mathscr{B}$; in particular, $\widehat{\nu}\ll\nu$.*

(3) *$\widehat{\nu}(E)=\nu(E)$ if $E\in\mathscr{B}$ and $\nu(E)<+\infty$; in particular, if E is morally null, $\nu(E)$ is either 0 or $+\infty$.*

Proof Part (1) follows from Lemma 3.2 and part (2) is direct from definitions. The ring used to define the measure ν consists of unions of halfspace-intervals (see [11]); thus, if $\nu(E)<+\infty$, the set E is contained in a countable union of halfspace-intervals, hence in a disjoint one, by Lemma 3.2. Now part (3) is straightforward. □

3.2 Properties of the σ–algebra of morally measurable sets

Let X be a complete median space throughout this section.

Lemma 3.4 *Singletons in $\mathscr{H}$ are morally measurable and the following are equivalent for $\mathfrak{h}\in\mathscr{H}$:*

(1) *$\mathfrak{h}$ is an atom for $\widehat{\nu}$.*

(2) *$\mathfrak{h}$ is clopen.*

(3) $d(\mathfrak{h},\mathfrak{h}^*)>0$.

Proof If one side of the wall $\{\mathfrak{h},\mathfrak{h}^*\}$ is not closed, say $\mathfrak{h}$, we can find points $x_n\in\mathfrak{h}$ converging to some $y\in\mathfrak{h}^*$. Thus, $\{\mathfrak{h}\}$ lies in the intersection of the sets $\mathscr{H}(y \mid x_n)$ and it is morally null (hence morally measurable, since the σ–algebra $\mathscr{B}_0$ is complete). Otherwise, $\mathfrak{h}$ is clopen; by Lemma 2.9, this is equivalent to $d(\mathfrak{h},\mathfrak{h}^*)>0$. Lemma 2.4 provides a pair (x,y) of gates for $(\mathfrak{h},\mathfrak{h}^*)$, hence the set $\{\mathfrak{h}\}=\mathscr{H}(\mathfrak{h}^* \mid \mathfrak{h})=\mathscr{H}(y \mid x)$ lies in $\mathscr{B}$ and has positive measure. Conversely, if $\mathfrak{h}$ is an atom, it is easy to see that $d(u,v)\ge\widehat{\nu}(\{\mathfrak{h}\})$ for all $u\in\mathfrak{h}$ and $v\in\mathfrak{h}^*$. □

Lemma 3.5 *A complete median space X is connected if and only if the measure space $(\mathscr{H}(X), \widehat{\nu})$ has no atoms.*

Proof If X is connected, $\widehat{\nu}$ has no atoms by Lemma 3.4. Conversely, if $\widehat{\nu}$ has no atoms, we prove that, for every $x, y \in X$, there exists $z \in I(x, y)$ such that $d(x, z) = d(z, y)$. This implies that X is geodesic (see eg [6, Remark 1.4(1) in Chapter I.1]). Let $\mathscr{F}_x \subseteq I(x, y)$ be the subset of points z such that $d(x, z) \le d(z, y)$; we endow it with a poset structure by declaring that $z_1 \preceq z_2$ whenever $z_2 \in I(z_1, y)$. Chains in $\mathscr{F}_x$ correspond to Cauchy nets in $I(x, y)$ and these converge; thus, $\mathscr{F}_x$ is inductive and Zorn's lemma yields a maximal element $\bar{z}$. Exchanging the roles of x and y, the same construction provides a point $\overline{w}$. Suppose for the sake of contradiction that we have $d(x, \bar{z}) < \frac{1}{2} d(x, y) < d(x, \overline{w})$. By maximality of $\bar{z}$ and $\overline{w}$, the interval $I(\bar{z}, \overline{w})$ consists of the sole points $\bar{z}$ and $\overline{w}$. By Proposition 2.3, we conclude that $\mathscr{H}(\bar{z} \mid \overline{w})$ consists of a single halfspace, hence an atom, a contradiction. □

Given a point $x \in X$ and a convex subset $C \subseteq X$ we define their *adjacencies*:

$$\begin{aligned} \mathrm{Adj}_x &:= \{\mathfrak{h} \in \mathscr{H} \mid x \notin \mathfrak{h} \text{ and } x \in \overline{\mathfrak{h}}\}, \\ \mathrm{Adj}(C) &:= \{\mathfrak{h} \in \mathscr{H} \mid \mathfrak{h} \cap C = \varnothing \text{ and } \mathfrak{h} \cap \overline{C} \neq \varnothing\}. \end{aligned}$$

Note that in general $\mathrm{Adj}_x \neq \mathrm{Adj}(\{x\}) = \varnothing$.

Lemma 3.6 *If X is locally convex, adjacencies are morally null. In particular, σ_C and $\mathscr{H}(C)$ are morally measurable for every convex subset $C \subseteq X$.*

Proof Given $x, y \in X$, let K be the intersection of all halfspaces in $\mathrm{Adj}_x \cap \sigma_y$; by Lemma 2.11, we have $x \in \overline{K}$ and we can find points $x_n \in K$ converging to x. Thus $\mathrm{Adj}_x \cap \sigma_y$ is contained in the intersection of the sets $\mathscr{H}(x \mid x_n)$ and it is morally null. By the arbitrariness of y, we conclude that Adj_x is morally null. For every convex subset $C \subseteq X$ and points $u, v \in X$, we have

$$\mathrm{Adj}(C)^* \cap \mathscr{H}(u \mid v) \subseteq \mathrm{Adj}(C)^* \setminus \sigma_u \subseteq \mathscr{H}(\overline{u} \mid C),$$

where $\overline{u}$ is the gate-projection of u to $\overline{C}$; the last inclusion follows from Lemma 2.2(1). The set $\mathscr{H}(\overline{u} \mid C) \subseteq \mathrm{Adj}_{\overline{u}}$ is morally null, hence $\mathrm{Adj}(C)$ is morally null. Moreover, $\sigma_C \setminus \sigma_u = \mathscr{H}(u \mid C) = \mathscr{H}(u \mid \overline{u}) \sqcup \mathscr{H}(\overline{u} \mid C)$, where $\mathscr{H}(\overline{u} \mid C)$ is morally null; as a consequence, σ_C is morally measurable. The same holds for $\mathscr{H}(C) = \mathscr{H} \setminus (\sigma_C \sqcup \sigma_C^*)$. □

We will denote by $\mathscr{H}^0$ the collection of nowhere-dense halfspaces and set $\mathscr{H}^\times := \mathscr{H}^0 \cup (\mathscr{H}^0)^*$. If $\mathfrak{h} \in \mathscr{H} \setminus \mathscr{H}^\times$ we say that $\mathfrak{h}$ is *thick*.

Corollary 3.7 *If X is locally convex, the set $\mathscr{H}^{\times}$ is morally null.*

Proof This follows from Lemma 3.6 and the observation that, for every $x, y \in X$, we have $\mathscr{H}^0 \cap \mathscr{H}(x \mid y) \subseteq \mathrm{Adj}_y^*$. □

Proposition 3.8 *Let X be locally convex. If X is separable, the measure space $(\mathscr{H}, \widehat{\nu})$ is σ–finite. The converse holds if X has finite rank.*

Proof If $\{x_n\}_{n\geq 0}$ is a countable dense subset of X, then all halfspace-intervals $\mathscr{H}(x_n \mid x_m)$ have finite measure and their union contains $\mathscr{H} \setminus \mathscr{H}^{\times}$; thus $(\mathscr{H}, \widehat{\nu})$ is σ–finite by Corollary 3.7.

Conversely, if $(\mathscr{H}, \widehat{\nu})$ is σ–finite, Lemma 3.3(1) implies that there exists $\{x_n\}_{n\geq 0} \subseteq X$ such that the sets $\mathscr{H}(x_n \mid x_m)$ cover $\mathscr{H}$ up to a morally null set. By [3, Lemma 6.4] and Corollary 2.20 above, hulls of separable subsets of X are separable; thus, the hull C of $\{x_n\}_{n\geq 0}$ is separable. If there existed a point $z \notin \overline{C}$, the gate $\bar{z}$ for $(z, \overline{C})$ would produce a positive-measure set $\mathscr{H}(z \mid \bar{z})$ disjoint from the union of the $\mathscr{H}(x_n \mid x_m)$, a contradiction. Thus C is dense in X, and X is separable. □

Recall that a subset $\mathcal{C} \subseteq \mathscr{H}$ is said to be inseparable if it contains every halfspace $\mathfrak{j} \in \mathscr{H}$ such that there exist $\mathfrak{h}, \mathfrak{k} \in \mathcal{C}$ with $\mathfrak{h} \subseteq \mathfrak{j} \subseteq \mathfrak{k}$.

Lemma 3.9 *If X has finite rank, any inseparable subset $\mathcal{C} \subseteq \mathscr{H}$ is morally measurable. In particular, every filter $\sigma \subseteq \mathscr{H}$ is morally measurable.*

Proof It suffices to prove the lemma under the additional assumption that $\mathcal{C} \subseteq \mathscr{H}(x \mid y)$ for points $x, y \in X$. By Lemma 2.21 and Corollary 2.27, $\mathcal{C}$ is a countable union of subsets of the form $\mathscr{H}(\mathfrak{k}^* \mid \mathfrak{h})$ with $\mathfrak{h}, \mathfrak{k} \in \mathcal{C}$. Each of these is morally measurable by Lemma 3.6. □

We can extend the notion of admissibility in Section 2.2 and [11] as follows. We say that a partial filter $\sigma \subseteq \mathscr{H}$ is *tangible* if it is morally measurable and $\widehat{\nu}(\sigma \setminus \sigma_x) < +\infty$ for some (equivalently, all) $x \in X$. For a morally measurable ultrafilter σ, tangibility is equivalent to having $\widehat{\nu}(\sigma \triangle \sigma_x) < +\infty$, since $(\sigma_x \setminus \sigma) = (\sigma \setminus \sigma_x)^*$. For instance, all admissible ultrafilters σ on the PMP $(\mathscr{H}, \mathscr{B}, \nu, \sigma_x)$ are tangible. Indeed, they are morally measurable since $\sigma = \sigma_x \sqcup [(\sigma \triangle \sigma_x) \setminus \sigma_x]$, where σ_x is morally measurable and $\sigma \triangle \sigma_x \in \mathscr{B}$.

We denote by $\mathscr{M}(X)$ the set of tangible ultrafilters, identifying ultrafilters with $\widehat{\nu}$–null symmetric difference. The analogy with the notation $\mathcal{M}(X)$ of Section 2.2 is justified by Corollary 3.11 below; for now, we simply observe that there are isometric embeddings $X \hookrightarrow \mathcal{M}(X) \hookrightarrow \mathscr{M}(X)$ as a consequence of Lemma 3.3(2)–(3). The following is a key result.

Lemma 3.10 *Let X be locally convex. For every tangible filter $\sigma \subseteq \mathscr{H}$, there exists $x \in X$ such that $\widehat{\nu}(\sigma \setminus \sigma_x) = 0$.*

Proof If there exists $x_0 \in X$ such that $\widehat{\nu}(\sigma \cap \mathscr{H}(x_0 \mid y)) = 0$ for all $y \in X$, then $\widehat{\nu}(\sigma \setminus \sigma_{x_0}) = 0$. Indeed, given any $u, v \in X$, we can set $m := m(x_0, u, v)$, and $(\sigma \setminus \sigma_{x_0}) \cap \mathscr{H}(u \mid v) = \sigma \cap \mathscr{H}(m \mid v) \subseteq \sigma \cap \mathscr{H}(x_0 \mid v)$, the latter being morally null.

If instead $\widehat{\nu}(\sigma \cap \mathscr{H}(x_0 \mid y)) > 0$ for some $x_0, y \in X$, then there exists $z \in I(x_0, y)$ such that $z \neq x_0$ and $\mathscr{H}(x_0 \mid z) \subseteq \sigma$. Indeed, by Lemma 3.6 there exists $\mathfrak{h} \in \sigma \cap \mathscr{H}(x_0 \mid y)$ such that $d(x_0, \mathfrak{h}) > 0$. Letting z be the gate for $(x_0, \overline{\mathfrak{h}})$, if $\mathfrak{k} \in \mathscr{H}(x_0 \mid z)$, we have $\mathfrak{h} \subseteq \overline{\mathfrak{h}} \subseteq \mathfrak{k}$; hence, the fact that σ is a filter and $\mathfrak{h} \in \sigma$ implies that $\mathfrak{k} \in \sigma$.

Now, we construct a countable ordinal η and an injective net $(x_\alpha)_{\alpha \leq \eta}$ such that all the following are satisfied:

(1) $(\sigma \setminus \sigma_{x_{\alpha+1}}) \sqcup (\sigma_{x_{\alpha+1}} \setminus \sigma_{x_\alpha}) = (\sigma \setminus \sigma_{x_\alpha})$ for each $\alpha \leq \eta$.

(2) If α is a limit ordinal, $\bigcap_{\beta<\alpha} \sigma \setminus \sigma_{x_\beta} = \sigma \setminus \sigma_{x_\alpha}$ up to a morally null set.

(3) $\widehat{\nu}(\sigma \setminus \sigma_{x_\eta}) = 0$.

The net is constructed by transfinite induction, starting with an arbitrary choice of x_0. Suppose that x_β has been defined for all $\beta < \alpha$. If α is a limit ordinal, the inductive hypothesis implies that the disjoint union of the sets $\sigma_{x_{\beta+1}} \setminus \sigma_{x_\beta}$ for $\beta < \alpha$ is contained in $\sigma \setminus \sigma_{x_0}$, up to a morally null set. Hence $\sum_{\beta<\alpha} d(x_\beta, x_{\beta+1}) \leq \widehat{\nu}(\sigma \setminus \sigma_{x_0}) < +\infty$ and $(x_\beta)_{\beta<\alpha}$ is a Cauchy net. We define x_α to be its limit; we only need to check condition (2) and it follows from Lemma 3.6.

If $\alpha = \beta + 1$, we look at $\sigma \setminus \sigma_{x_\beta}$; if this is morally null, we stop and set $\eta = \beta$. Otherwise, we can find $y \in X$ so that $\widehat{\nu}(\sigma \cap \mathscr{H}(x_\beta \mid y)) > 0$ and we have already shown that there exists $z \in X$ with $\mathscr{H}(x_\beta \mid z) \subseteq \sigma$; we set $x_\alpha := z$. By construction $\sigma_{x_\alpha} \setminus \sigma_{x_\beta} \subseteq \sigma \setminus \sigma_{x_\beta}$, and since σ is a filter, we have $\sigma \cap (\sigma_{x_\beta} \setminus \sigma_{x_\alpha}) = \varnothing$; in particular, $\sigma \setminus \sigma_{x_\alpha} \subseteq \sigma \setminus \sigma_{x_\beta}$. From this we immediately get condition (1), ie $\sigma \setminus \sigma_{x_\beta} = (\sigma \setminus \sigma_{x_\alpha}) \sqcup (\sigma_{x_\alpha} \setminus \sigma_{x_\beta})$.

We conclude by remarking that, because $d(x_{\alpha+1}, x_\alpha) > 0$ for all ordinals α and $\sum d(x_\beta, x_{\beta+1}) \leq \widehat{\nu}(\sigma \setminus \sigma_{x_0}) < +\infty$, the process must terminate for some *countable* ordinal η. □

Corollary 3.11 (1) *If X is locally convex and $\sigma \subseteq \mathscr{H}$ is a tangible ultrafilter, there exists $x \in X$ such that $\widehat{\nu}(\sigma \triangle \sigma_x) = 0$.*

(2) *If X has finite rank and $\sigma \subseteq \mathscr{H}$ is a tangible filter, there exists a convex subset $C \subseteq X$ such that $\widehat{\nu}(\sigma \triangle \sigma_C) = 0$.*

Proof Let σ be a tangible filter and let $x \in X$ be such that $\widehat{\nu}(\sigma \setminus \sigma_x) = 0$, as provided by Lemma 3.10. If σ is an ultrafilter, we immediately obtain $\widehat{\nu}(\sigma \triangle \sigma_x) = 0$. Otherwise, suppose that X has finite rank and let C be the intersection of all $\mathfrak{h} \in \sigma \cap \sigma_x$; since $x \in C$, this is a nonempty convex subset and the filter σ_C contains $\sigma \cap \sigma_x$. In particular, we have $\widehat{\nu}(\sigma \setminus \sigma_C) = 0$.

Suppose $\mathfrak{k} \in \sigma_C \setminus \sigma$ and consider the gate-projection $\pi\colon X \to \overline{\mathfrak{k}^*}$. Observe that no $\mathfrak{h} \in \sigma \cap \sigma_x$ can be contained in $\mathfrak{k}$ or we would have $\mathfrak{k} \in \sigma$; hence every $\mathfrak{h} \in \sigma \cap \sigma_x$ intersects $\mathfrak{k}^*$. Lemma 2.2(1) then implies that, given any $y \in C$, the projection $\pi(y)$ still lies in C; in particular, $C \cap \overline{\mathfrak{k}^*} \neq \varnothing$. If $\mathfrak{k}$ were open, we would have $\overline{\mathfrak{k}^*} = \mathfrak{k}^*$ and the previous statement would contradict the fact that $C \subseteq \mathfrak{k}$; thus, $\mathfrak{k}$ is closed by Corollary 2.23.

If $u \in X$ and $\overline{u}$ is its gate-projection to $\overline{C}$, we have $\overline{u} \in \overline{\mathfrak{k}^*}$ for every $\mathfrak{k} \in \sigma_C \setminus (\sigma \cup \sigma_u)$, since $C \cap \overline{\mathfrak{k}^*} \neq \varnothing$. Since every such $\mathfrak{k}$ is closed, we have $\overline{u} \in \overline{\mathfrak{k}^*} \cap \mathfrak{k}$ and $\sigma_C \setminus (\sigma \cup \sigma_u) \subseteq \mathrm{Adj}^*_{\overline{u}}$. The arbitrariness of u and Lemma 3.6 imply that $\widehat{\nu}(\sigma_C \setminus \sigma) = 0$. □

An immediate consequence of Corollary 3.11(1) is:

Corollary 3.12 *For every complete, locally convex median space X, the isometric embedding $X \hookrightarrow \mathscr{M}(X)$ is surjective. In particular, the spaces X, $\mathcal{M}(X)$ and $\mathscr{M}(X)$ are isometric.*

An interesting consequence of Corollary 3.12 is the following.

Corollary 3.13 *Suppose that X is locally convex, with corresponding SMH–structure $(X, \mathscr{H}, \mathscr{B}, \nu)$. Then $\mathrm{Aut}(X, \mathscr{H}, \mathscr{B}, \nu) = \mathrm{Isom}(X)$.*

An analogous result holds for the SMW–structure.

4 Compactifying median spaces and algebras

4.1 The zero-completion of a median algebra

Let M be a median algebra. We denote by $\mathscr{I}(M)$ (or simply by $\mathscr{I}$) the poset of all intervals $I(x, y)$ with $x, y \in M$, ordered by inclusion; singletons are allowed. Note that the poset $(\mathscr{J}, \subseteq)$ is not a directed set. Still, whenever $I, I', I'' \in \mathscr{J}$ and $I \subseteq I' \subseteq I''$, we have $\pi_I|_{I''} = \pi_I|_{I'} \circ \pi_{I'}|_{I''}$ and it makes perfect sense to consider the inverse limit $\varprojlim_{I \in \mathscr{J}} I$ defined to be

$$\left\{ (x_I)_I \in \prod_{I \in \mathscr{J}} I \;\middle|\; \pi_{I \cap J}(x_I) = \pi_{I \cap J}(x_J) \text{ for all } I, J \in \mathscr{J} \text{ such that } I \cap J \neq \varnothing \right\}.$$

Note that $I \cap J \in \mathscr{J}$ whenever $I, J \in \mathscr{J}$; indeed, we showed in Section 2.1 that every gate-convex subset of I is itself an interval.

The product of all $I \in \mathscr{J}$ has a median-algebra structure given by considering median maps component by component; the inverse limit $\varprojlim I$ also inherits a median-algebra structure and we have a monomorphism

$$\iota\colon M \hookrightarrow \varprojlim I, \quad x \mapsto (\pi_I(x))_I.$$

Definition 4.1 We denote the median algebra $\varprojlim I$ simply by $\overline{M}$ and refer to it as the *zero-completion* of M.

This terminology comes from [2] where the same notion is defined from a different perspective; more on this in Remark 4.7.

We also have a monomorphism $i\colon M \hookrightarrow 2^{\mathscr{H}}$ given by mapping $x \mapsto \sigma_x$; the space $2^{\mathscr{H}}$ can be endowed with the product median-algebra structure and the product topology. Here the set $2 = \{0, 1\}$ is equipped with the discrete topology and its unique median-algebra structure.

Definition 4.2 The closure of $i(M)$ in $2^{\mathscr{H}}$ with the induced median-algebra structure will be denoted by $M^{\circ\circ}$ and we will refer to it as the *double dual* of M (compare [34]).

Lemma 4.3 (1) *The median algebra $M^{\circ\circ}$ coincides with the subset of $2^{\mathscr{H}}$ consisting of ultrafilters on $\mathscr{H}$.*

(2) *There is a monomorphism $j\colon \overline{M} \hookrightarrow M^{\circ\circ}$ such that $i = j \circ \iota$.*

Proof Ultrafilters on $\mathscr{H}$ form a closed subset of $2^{\mathscr{H}}$ with the product topology; thus every element of $M^{\circ\circ}$ is an ultrafilter. Every neighbourhood of an ultrafilter $\sigma \subseteq \mathscr{H}$ is of the form $\{A \subseteq \mathscr{H} \mid \mathfrak{h}_1, \dots, \mathfrak{h}_k \in A \text{ and } \mathfrak{h}_{k+1}^*, \dots, \mathfrak{h}_n^* \notin A\}$, where $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ lie in σ and hence intersect pairwise. Any such neighbourhood intersects $i(M)$ as it contains $i(x)$ for any $x \in \mathfrak{h}_1 \cap \dots \cap \mathfrak{h}_n$; the latter is nonempty by Helly's theorem. Hence, every ultrafilter on $\mathscr{H}$ lies in $M^{\circ\circ}$.

Regarding part (2), given $(x_I)_I \in \varprojlim I$ we will construct an ultrafilter $\sigma \subseteq \mathscr{H}$ such that $\sigma \cap \mathscr{H}(I) = \sigma_{x_I} \cap \mathscr{H}(I)$ for all $I \in \mathscr{J}$. This ultrafilter is unique and the corresponding map $j\colon \overline{M} \to M^{\circ\circ}$ is easily seen to be a monomorphism. It suffices to show that the sets

$$\Omega_1 := \{\mathfrak{h} \in \mathscr{H} \mid \text{there exists } I \text{ such that } x_I \in \mathfrak{h} \in \mathscr{H}(I)\},$$

$$\Omega_2 := \{\mathfrak{h} \in \mathscr{H} \mid x_I \in \mathfrak{h} \text{ for every } I \text{ such that } \mathfrak{h} \in \mathscr{H}(I)\},$$

coincide and set $\sigma := \Omega_1 = \Omega_2$. This is indeed an ultrafilter: Ω_1 contains at least one side of every wall of M, while any two halfspaces in Ω_2 intersect.

The inclusion $\Omega_2 \subseteq \Omega_1$ is immediate; proving the other amounts to showing that $x_J \in \mathfrak{h} \in \mathscr{H}(J)$ and $x_I \notin \mathfrak{h} \in \mathscr{H}(I)$ cannot happen at the same time. We argue by contradiction.

Observe that $m_I := \pi_I(x_J) \in \mathfrak{h}$ and $m_J := \pi_J(x_I) \in \mathfrak{h}^*$ by Lemma 2.2(1). Let $I' := I(x_I, m_I)$ and $J' := I(x_J, m_J)$; since $x_I \in I' \subseteq I$, we have $x_{I'} = \pi_{I'}(x_I) = x_I$ and similarly $x_{J'} = x_J$. Set $K := I(x_I, x_J)$. Since $I' \subseteq K$ and $J' \subseteq K$, we have $\pi_{I'}(x_K) = x_I$ and $\pi_{J'}(x_K) = x_J$. However, $\{\mathfrak{h}, \mathfrak{h}^*\} \in \mathscr{W}(I') \cap \mathscr{W}(J')$, so the previous equalities imply that $x_K \in \mathfrak{h}^*$ and $x_K \in \mathfrak{h}$, respectively, a contradiction. □

Corollary 4.4 (1) *The embedding $\iota\colon M \to \overline{M}$ has convex image. Thus, given $x, y \in M$, the notion of $I(x, y)$ is the same in M and $\overline{M}$.*

(2) *For every $J \in \mathscr{J}$, the projection $p_J\colon \overline{M} \to J$ that $\overline{M}$ inherits from $\prod I$ is precisely the gate-projection $\pi_J\colon \overline{M} \to J$.*

Proof Consider points $x, y \in M$ and set $J := I(x, y) \in \mathscr{J}$. Given $\underline{z} \in \overline{M}$, we write $\underline{z}_I$ instead of $p_I(\underline{z})$. Assuming $m(\iota(x), \iota(y), \underline{z}) = \underline{z}$, we will show that $\underline{z} = \iota(\underline{z}_J)$. This yields $\underline{z} \in \iota(M)$ and proves part (1).

If we had $\underline{z} \neq \iota(\underline{z}_J)$, there would exist $I \in \mathscr{J}$ such that $\underline{z}_I \neq p_I\iota(\underline{z}_J)$; note that $p_I\iota(\underline{z}_J) = \pi_I(\underline{z}_J)$, where π_I denotes the gate-projection $M \to I$. Let $\mathfrak{h} \in \mathscr{H}(M)$ be a halfspace lying in $\mathscr{H}(\pi_I(\underline{z}_J) \mid \underline{z}_I)$; by Proposition 2.3, we have $\mathfrak{h} \in \mathscr{H}(\underline{z}_J \mid \underline{z}_I)$. The equality $\Omega_1 = \Omega_2$ in the proof of Lemma 4.3 now shows that $\mathfrak{h} \notin \mathscr{H}(J)$. Since $\underline{z}_J \in \mathfrak{h}^*$, we have $J \subseteq \mathfrak{h}^*$; in particular, $\pi_I(x)$ and $\pi_I(y)$ lie in $\mathfrak{h}^*$. However, since $m(\iota(x), \iota(y), \underline{z}) = \underline{z}$, we have $m(\pi_I(x), \pi_I(y), \underline{z}_I) = p_I m(\iota(x), \iota(y), \underline{z}) = \underline{z}_I \in \mathfrak{h}$, a contradiction.

For part (2), note that $\pi_J(\underline{z}) = m(x, y, \underline{z})$ lies in $\iota(M)$ by part (1). Thus,

$$m(x, y, \underline{z}) = p_J m(x, y, \underline{z}) = m(p_J(x), p_J(y), p_J(\underline{z})) = m(x, y, \underline{z}_J) = \underline{z}_J. \quad \square$$

The median algebras $\overline{M}$ and $M^{\circ\circ}$ can coincide; for instance, this is the case for 0–skeleta of CAT(0) cube complexes. However, the following example shows that $\overline{M}$ and $M^{\circ\circ}$ differ in general.

Example 4.5 Consider the median algebras $N = \mathbb{N}$ and $M = \mathbb{N} \cup \{+\infty\}$; in both cases $m(x,y,z) = y$ if $x \leq y \leq z$. For every $k \in \mathbb{N}$, both M and N have a wall $\mathfrak{w}_k$ separating k and $k+1$; in M, there is an additional wall $\mathfrak{w}_\infty$ separating $\mathbb{N}$ and $+\infty$. Observe that $M = \overline{N} = N^{\circ\circ}$, and since $M = I(0, +\infty)$, we also have $M = \overline{M}$.

However, $M^{\circ\circ} = M \cup \{\infty_-\}$, where ∞_- is represented by the ultrafilter that picks the side containing $+\infty$ for every wall $\mathfrak{w}_k$ and the side containing $\mathbb{N}$ for the wall $\mathfrak{w}_\infty$. The point ∞_- is "bigger than any natural number" but still "smaller than $+\infty$"; also compare Remark 2.18.

Given $a \in M$, we say that a convex subset $C \subseteq M$ is *a–directed* if $a \in C$ and, for every $x, y \in C$, there exists $z \in C$ such that $x, y \in I(a,z)$.

Lemma 4.6 *Fix $a \in M$. There is a one-to-one correspondence between points of $\overline{M}$ and gate-convex, a–directed subsets $C \subseteq M$. Points $b \in M \subseteq \overline{M}$ correspond to intervals $I(a,b)$.*

Proof If $C \subseteq M$ is gate-convex and a–directed, the projection $C_I := \pi_I(C)$ is a gate-convex, $\pi_I(a)$–directed subset of I by Lemma 2.2(2). It follows that there exists a (unique) point $x_I \in C_I$ such that $C_I = I(\pi_I(a), x_I)$. By Proposition 2.1, we obtain a point $\xi_C := (x_I)_I \in \varprojlim I$.

Conversely, given $\xi \in \overline{M}$, we can consider the interval $I(a,\xi) \subseteq \overline{M}$ and set $C_\xi := I(a,\xi) \cap M$. Since M is convex in $\overline{M}$, the map $u \mapsto m(a,\xi,u)$ takes M into itself and it is a gate-projection $M \to C_\xi$ by Proposition 2.1. If $x, y \in C_\xi$, we have $x, y \in I(a,z)$ with $z := m(x,y,\xi) \in C_\xi$. Thus C_ξ is gate-convex and a–directed.

Observe that, setting $z := m(a,\xi,\pi_I(\xi))$, we have $\pi_I(z) = \pi_I(\xi)$ and $z \in I(a,\xi) \cap M$; hence, $\pi_I(I(a,\xi) \cap M) = I(\pi_I(a), \pi_I(\xi))$ for all $I \in \mathscr{I}$, ie $\xi = \xi_{C_\xi}$ for every $\xi \in \overline{M}$. We conclude by observing that $C_1 \neq C_2$ implies $\xi_{C_1} \neq \xi_{C_2}$; indeed, if $x \in C_1 \setminus C_2$ and $J := I(a,x)$, we have $\pi_J(\xi_{C_1}) = x$ and $\pi_J(\xi_{C_1}) \neq x$, since $\pi_J(C_1) = J$ and $\pi_J(C_2) = C_2 \cap J \neq J$. □

Remark 4.7 When zero-completions were originally defined in [2], points of $\overline{M}$ were compatible meet-semilattice operations; see [2, Theorem 1]. The terminology "zero-completion" is due to the fact that all semilattice operations have a zero in $\overline{M}$, while they need not have one in M. By Lemma 4.6 above and [1, Theorem 5.5], our definition of $\overline{M}$ yields the same object. Zero-completions do not seem to have been studied outside of [2]; we will develop their theory further in the next sections, especially in the case when M arises from a median space.

Lemma 4.8 *If $C \subseteq M$ is gate-convex, the zero-completion $\overline{C}$ canonically embeds into $\overline{M}$ as a gate-convex subset with $\overline{C} \cap M = C$.*

Proof If $\xi = (x_I)_I \in \varprojlim I$, we set $\pi(\xi) := (\pi_I \pi_C(x_I))_I \in \prod I$. This is an element of $\varprojlim I$ as, if $J \subseteq K$ are intervals of M, we have

$$\pi_J(\pi_K \pi_C(x_K)) = \pi_J \pi_C \pi_K(\xi) = \pi_J \pi_C \pi_J \pi_K(\xi) = \pi_J \pi_C \pi_J(\xi) = \pi_J \pi_C(x_J),$$

by Corollary 4.4 and Lemma 2.2(4). We obtain a map $\pi\colon \overline{M} \to \overline{M}$, which is a gate-projection onto its image, by Proposition 2.1. The image of π is the set of those $(x_I)_I$ with $\pi_I \pi_C(x_I) = x_I$, ie $x_I \in \pi_I(C)$, for all $I \in \mathscr{J}$. Restricting the index set to $\mathscr{J}(C)$, we obtain a morphism $f\colon \operatorname{im}\pi \to \overline{C}$. Embedding $\overline{M} \hookrightarrow M^{\circ\circ}$ as in Lemma 4.3, all points of $\operatorname{im}\pi$ are represented by ultrafilters containing σ_C. In terms of ultrafilters, f is the restriction of the map that takes each ultrafilter on $\mathscr{H}(M)$ to its intersection with the subset $\mathscr{H}(C) \subseteq \mathscr{H}(M)$; in particular, f is injective.

We now show that f is surjective; given $\xi \in \overline{C}$, let $\sigma \subseteq \mathscr{H}(C)$ be the ultrafilter representing ξ, as provided by Lemma 4.3. If $I \in \mathscr{J}(M)$, there exist points $u, v \in C$ such that $\pi_I(C) = I(\pi_I(u), \pi_I(v))$; we set $J := I(u, v)$. If ξ_J is the coordinate of ξ corresponding to $J \in \mathscr{J}(C)$, we set $x_I := \pi_I(\xi_J)$. Note that x_I is represented by the ultrafilter $(\sigma_C \cap \mathscr{H}(I)) \sqcup (\sigma \cap \mathscr{H}(I))$ on $\mathscr{H}(I)$. In particular, our definition of x_I does not depend on the choice of the points u and v, and $(x_I)_I$ satisfies the compatibility condition necessary to define a point $\eta \in \overline{M}$. It is clear that $\eta \in \operatorname{im}\pi$ and $f(\eta) = \xi$. Thus, we can make the identification $\overline{C} \simeq \operatorname{im}\pi$; the fact that $\overline{C} \cap M = C$ is a trivial observation. □

Given $\mathfrak{h} \in \mathscr{H}(M)$, there exists a unique halfspace $\widetilde{\mathfrak{h}} \in \mathscr{H}(\overline{M})$ satisfying $\widetilde{\mathfrak{h}} \cap M = \mathfrak{h}$. This can be observed by applying Proposition 2.3 to an interval $I(x, y)$ with $x \in \mathfrak{h}$ and $y \in \mathfrak{h}^*$.

The halfspace $\widetilde{\mathfrak{h}}$ can be further characterised by $\xi \in \widetilde{\mathfrak{h}} \iff \mathfrak{h} \in j(\xi)$, where j is as in Lemma 4.3. Indeed, Corollary 4.4 implies that, for every interval $I \subseteq M$, we have $j(\xi) \cap \mathscr{H}(I) = \sigma_{\pi_I(\xi)} \cap \mathscr{H}(I)$. If $\mathfrak{h} \in \mathscr{H}(I)$, Proposition 2.3 then yields

$$\xi \in \widetilde{\mathfrak{h}} \iff \pi_I(\xi) \in \mathfrak{h} \iff \mathfrak{h} \in j(\xi).$$

Note that not all halfspaces of $\overline{M}$ are of the form $\widetilde{\mathfrak{h}}$; indeed, by convexity of M, we have $\mathscr{W}(M \mid \xi) \neq \varnothing$ for every $\xi \in \overline{M} \setminus M$. An explicit example of such a wall was given in Example 4.5 for the median algebra N (in the notation of the example, we have $\mathfrak{w}_\infty \in \mathscr{W}(N \mid +\infty)$ and $+\infty \in \partial N$).

Nevertheless, we can identify the set of halfspaces of M with a subset $\mathscr{H}(M) \subseteq \mathscr{H}(\overline{M})$, and any two points of $\overline{M}$ are separated by an element of $\mathscr{H}(M)$. Indeed, if $\xi, \eta \in \overline{M}$ are distinct, then $\pi_I(\xi) \neq \pi_I(\eta)$ for some interval $I \subseteq M$, and for any $\mathfrak{h} \in \mathscr{H}(\pi_I(\xi) \mid \pi_I(\eta))$, we have $\widetilde{\mathfrak{h}} \in \mathscr{H}(\xi \mid \eta)$.

Lemma 4.9 $$\operatorname{rank}(\overline{M}) = \operatorname{rank}(M).$$

Proof This is immediate from the discussion above and Lemma 2.5. ◻

Lemma 4.10 *If there exists a topology on M for which it is a compact topological median algebra, then $M = \overline{M}$.*

Proof Given $\xi = (x_I)_I \in \varprojlim I$, consider the projections $\pi_I \colon \overline{M} \to I$ and set $C_I := \pi_I^{-1}(x_I) \subseteq \overline{M}$. These are convex sets by Proposition 2.1 and they pairwise intersect as they all contain ξ; moreover, they all intersect M, which is convex in $\overline{M}$. Helly's theorem implies that $\{C_I \cap M \mid I \in \mathscr{I}\}$ has the finite intersection property.

Now endow M with its compact topology. Each $C_I \cap M = (\pi_I|_M)^{-1}(x_I)$ is closed in M, hence compact. Thus, the intersection of all the $C_I \cap M$ is nonempty; any x in this intersection satisfies $\pi_I(x) = x_I$ for every interval I, ie $\xi = \iota(x)$. ◻

In the rest of the section, we suppose that M is a topological median algebra. Endowing the product of all intervals with the product topology, the zero-completion $\overline{M}$ inherits a topology. The inclusion $\iota \colon M \hookrightarrow \overline{M}$ is always continuous, but in general not a topological embedding (compare Proposition 4.20). Classical examples of this phenomenon are provided by locally infinite trees, since in that case $\overline{M}$ coincides with the usual Roller compactification.

For instance, if T is a geodesically complete (real or simplicial) tree and $x \in T$ is a point of infinite degree, $\iota(x)$ does not lie in the interior of $\iota(T)$; in particular, $\iota(T)$ is not open in $\overline{T}$. Another interesting case is when T is a bounded tree and $T \setminus \{x\}$ contains infinitely many connected components of diameter at least 1; here the map $\iota \colon T \to \overline{T}$ is a continuous bijection, but still not a homeomorphism.

The following is an easy observation.

Lemma 4.11 *If M is locally convex, $\overline{M}$ is as well.*

Lemma 4.12 *If M has compact intervals, $\overline{M}$ is compact and M is dense in $\overline{M}$.*

Proof Compactness is immediate from the observation that $\overline{M}$ is a closed subset of $\prod I$. We prove density by showing that for every $\xi \in \overline{M}$ and for every finite collection of intervals $I_1, \dots, I_k \in \mathscr{J}(M)$, there exists $x \in M$ such that $\pi_{I_i}(x) = \pi_{I_i}(\xi)$ for all i.

Let C be the convex hull of $I_1 \cup \dots \cup I_k$ in $\overline{M}$; it is compact by Lemma 2.7 and $C \subseteq M$ since M is convex in $\overline{M}$. Note that C is gate-convex in $\overline{M}$ by Lemma 2.6; let $\pi\colon \overline{M} \to C$ be the corresponding gate-projection. By Lemma 2.2(3), we can set $x := \pi(\xi) \in C \subseteq M$. □

If M has compact intervals and $C \subseteq M$ is gate-convex, Lemma 4.12 implies that the subset of $\overline{M}$ that we identified with the zero-completion $\overline{C}$ in Lemma 4.8 coincides with the closure of C in the topology of $\overline{M}$.

4.2 The Roller compactification of a median space

Throughout this section, let X be a complete, locally convex median space with compact intervals. This encompasses all finite-rank median spaces (see Corollary 2.20), all (possibly infinite-dimensional) CAT(0) cube complexes and all (possibly infinite rank) complete, connected, locally compact, locally convex median spaces (see [10] for examples).

Definition 4.13 In this context, we will refer to the zero-completion $\overline{X}$ as the *Roller compactification* and to $\partial X := \overline{X} \setminus X$ as the *Roller boundary*.

The renaming is justified by the fact that the median metric on X induces an additional structure on $\overline{X}$, which we shall study in this and the following section. There is a strong analogy with Roller boundaries of CAT(0) cube complexes, and indeed, if X is the 0–skeleton of a CAT(0) cube complex, our notion of Roller boundary coincides with the usual one. The following proposition sums up what we already know about $\overline{X}$; it roughly corresponds to the first and third definitions of $\overline{X}$ that we gave in the introduction.

Theorem 4.14 (1) *The Roller compactification $\overline{X}$ is a locally convex, compact, topological median algebra.*

(2) *The inclusion $\iota\colon X \hookrightarrow \overline{X}$ is a continuous morphism with convex, dense image.*

(3) *For every closed convex subset $C \subseteq X$, the closure of C in $\overline{X}$ is gate-convex and naturally identified with the Roller compactification of C.*

(4) *If X is separable, the topology of $\overline{X}$ is separable and metrisable.*

We remark that ∂X does not need to be closed; compare Proposition 4.20 below. In analogy with the space $\mathscr{M}(X)$ and Definition 4.2, we introduce

$$\overline{\mathscr{M}}(X) := \{\sigma \subseteq \mathscr{H} \mid \sigma \in \widehat{\mathscr{B}} \text{ and } \sigma \text{ is an ultrafilter}\}/\sim,$$

where $\sigma_1 \sim \sigma_2$ if $\widehat{\nu}(\sigma_1 \triangle \sigma_2) = 0$. We can give $\overline{\mathscr{M}}(X)$ a median-algebra structure by defining the median map as in Example 2.8(2). If X has finite rank, all ultrafilters are morally measurable by Lemma 3.9 and $\overline{\mathscr{M}}(X)$ is simply a quotient of $X^{\circ\circ}$. By Corollary 3.12, we have a monomorphism $X \simeq \mathscr{M}(X) \hookrightarrow \overline{\mathscr{M}}(X)$. The next result corresponds to the second definition of the Roller compactification from the introduction.

Theorem 4.15 *The map $j\colon \overline{X} \hookrightarrow X^{\circ\circ}$ introduced in Lemma 4.3 takes values in the set of morally measurable ultrafilters and it descends to an isomorphism $\bar{j}\colon \overline{X} \xrightarrow{\simeq} \overline{\mathscr{M}}(X)$ extending $X \hookrightarrow \overline{\mathscr{M}}(X)$.*

Proof If $\xi = (x_I)_I \in \varprojlim I$, we have $j(\xi) \cap \mathscr{H}(I) = \sigma_{x_I} \cap \mathscr{H}(I)$ for every interval I, hence $j(\xi)$ is morally measurable. We get a morphism $\bar{j}\colon \overline{X} \to \overline{\mathscr{M}}(X)$ extending $X \hookrightarrow \overline{\mathscr{M}}(X)$. If $\eta = (y_I)_I$ satisfies $\widehat{\nu}(j(\xi) \triangle j(\eta)) = 0$, we have $\widehat{\nu}((\sigma_{x_I} \triangle \sigma_{y_I}) \cap \mathscr{H}(I)) = 0$, ie $x_I = y_I$, for all $I \in \mathscr{J}$. Thus, $\bar{j}$ is injective.

If $\sigma \subseteq \mathscr{H}$ is a morally measurable ultrafilter, each $\sigma \cap \mathscr{H}(I)$ is a morally measurable ultrafilter on $\mathscr{H}(I)$ and it is tangible since $\widehat{\nu}(\mathscr{H}(I)) < +\infty$. Corollary 3.12 provides the existence of $z_I \in I$ with $\widehat{\nu}((\sigma_{z_I} \triangle \sigma) \cap \mathscr{H}(I)) = 0$. We obtain $\zeta := (z_I)_I \in \varprojlim I$ with $\widehat{\nu}(j(\zeta) \triangle \sigma) = 0$; hence, $\bar{j}$ is also surjective. □

In general, given $\xi \in \overline{X}$ and a morally measurable ultrafilter $\sigma \subseteq \mathscr{H}$ representing ξ, we could have $\mathfrak{h} \in \sigma$ even if $\xi \notin \widetilde{\mathfrak{h}}$; see eg Remark 2.18. However, we have already observed that this does not happen for $\sigma = j(\xi)$. Since $j(x) = \sigma_x$ for every $x \in X$, we will also denote $j(\xi)$ by σ_ξ from now on. This should be viewed as a canonical choice of an ultrafilter representing ξ.

Lemma 4.16 *A sequence $(\xi_n)_{n\geq 0}$ in $\overline{X}$ converges to $\xi \in \overline{X}$ if and only if*

$$\widehat{\nu}(\limsup(\sigma_\xi \triangle \sigma_{\xi_n})) = 0.$$

Proof Given $I \in \mathscr{J}$, Lemma 3.6 implies that $\pi_I(\xi_n) \to \pi_I(\xi)$ if and only if

$$0 = \widehat{\nu}(\limsup_{n\to+\infty}(\sigma_{\pi_I(\xi)} \triangle \sigma_{\pi_I(\xi_n)})) = \widehat{\nu}(\limsup_{n\to+\infty}(\sigma_\xi \triangle \sigma_{\xi_n}) \cap \mathscr{H}(I)).$$

Since $\xi_n \to \xi$ if and only if $\pi_I(\xi_n) \to \pi_I(\xi)$ for all $I \in \mathscr{J}$, convergence corresponds precisely to $\limsup(\sigma_\xi \triangle \sigma_{\xi_n})$ being morally null. □

We can endow $\overline{X} \simeq \overline{\mathscr{M}}(X)$ with an *extended metric*

$$d(\sigma_1, \sigma_2) := \tfrac{1}{2} \cdot \widehat{\nu}(\sigma_1 \triangle \sigma_2) \in [0, +\infty];$$

the restriction to $\mathscr{M}(X) \simeq X$ coincides with the usual metric on X.

Lemma 4.17 (1) *The median map $m\colon \overline{X}^3 \to \overline{X}$ is 1–Lipschitz with respect to the extended metric d.*

(2) *If $C \subseteq \overline{X}$ is closed and convex, the gate-projection $\pi\colon \overline{X} \to C$ is 1–Lipschitz with respect to the extended metric d.*

Proof Part (1) follows from part (2) applied to intervals. Denote by σ_C the set of $\mathfrak{h} \in \mathscr{H}$ such that $C \subseteq \widetilde{\mathfrak{h}}$ and set $\mathscr{H}(C) := \mathscr{H} \setminus (\sigma_C \sqcup \sigma_C^*)$. By Lemma 2.2(1), we have $\sigma_{\pi(\xi)} = (\sigma_\xi \cap \mathscr{H}(C)) \sqcup \sigma_C$ for all $\xi \in \overline{X}$. Thus, for all $\xi, \eta \in \overline{X}$,

$$2 \cdot d(\pi(\xi), \pi(\eta)) = \widehat{\nu}((\sigma_\xi \triangle \sigma_\eta) \cap \mathscr{H}(C)) \leq \widehat{\nu}(\sigma_\xi \triangle \sigma_\eta) = 2 \cdot d(\xi, \eta). \qquad \square$$

If $\xi \in \overline{X}$ and $(\xi_n)_{n\geq 0}$ is a sequence in $\overline{X}$ with $d(\xi_n, \xi) \to 0$, Lemma 4.17(2) shows that $\pi_I(\xi_n) \to \pi_I(\xi)$ for every interval $I \subseteq X$; hence $\xi_n \to \xi$ in $\overline{X}$. The converse does not hold: consider a sequence $(x_n)_{n\geq 0}$ in $X \subseteq \overline{X}$ that converges to a point $\xi \in \partial X$ (these exist by Lemma 4.12). Corollary 3.12 shows that $d(x_n, \xi) = +\infty$ for all $n \geq 0$, so $d(x_n, \xi)$ does not converge to 0.

The following notion is due to [25] for CAT(0) cube complexes.

Definition 4.18 A *component* of $\overline{X}$ is a $\approx$–equivalence class of morally measurable ultrafilters for the equivalence relation $\approx$ defined by

$$\sigma_1 \approx \sigma_2 \iff d(\sigma_1, \sigma_2) < +\infty.$$

Note that the subset $X \subseteq \overline{X}$ always forms a single component. Example 2.8(2) implies the following.

Proposition 4.19 *The restriction of the metric d to any component of $\overline{X}$ gives it a median-space structure. Each component is convex in $\overline{X}$.*

The study of components of ∂X will be the subject of Section 4.3.

Proposition 4.20 *If X is connected and locally compact, the inclusion $\iota\colon X \to \overline{X}$ is a topological embedding.*

Proof Given $x_0 \in X$, choose $0 < \delta < +\infty$ and let F be a finite, δ–dense subset of $\overline{B}_{2\delta}(x_0)$; this exists since X is proper by [6, Proposition 3.7 in Chapter I.3]. Let $\mathcal{U}$

be the set of points $\xi \in \overline{X}$ such that $d(\pi_I(\xi), x_0) < \delta$ for all intervals $I = I(x_0, y)$ with $y \in F$; this is a neighbourhood of x_0 in $\overline{X}$. If $\eta \in \overline{X}$ and $d(\eta, x_0) \geq 2\delta$, there exists $z \in I(x_0, \eta)$ with $d(x_0, z) = 2\delta$, since X is geodesic. Choosing $y \in F$ with $d(y, z) < \delta$ we have

$$d(m(\eta, x_0, y), x_0) > d(m(\eta, x_0, z), x_0) - \delta = d(z, x_0) - \delta = \delta.$$

Thus $\mathcal{U} \subseteq B_{2\delta}(x_0) \subseteq X$ and $x_0 \in \mathcal{U}$. Since x_0 and δ were arbitrary, this shows that the map ι is open. □

Note that connectedness cannot be dropped from the statement of Proposition 4.20. For instance, let T be the tree obtained by glueing countably many rays $\{r_n \mid n \in \mathbb{N}\}$ by their origins. Consider the closed subset $X \subseteq T$ obtained by removing the open interval $(0, n) \subseteq [0, +\infty)$ from the ray r_n for each $n \geq 0$. The median space X is proper, but the inclusion $\iota\colon X \to \overline{X}$ does not have open image; indeed, in $\overline{X}$ the endpoints at infinity of the rays r_n converge to their common origin.

We conclude this section by presenting one more characterisation of $\overline{X}$. Fixing $x_0 \in X$, we denote by $\mathcal{C}_{\mathrm{Lip}}(X)_{x_0}$ the set of 1–Lipschitz functions $X \to \mathbb{R}$ taking x_0 to 0; we endow this space with the topology of pointwise convergence, which is compact and coincides with the topology of uniform convergence on compact subsets. The map

$$B_{x_0}\colon X \hookrightarrow \mathcal{C}_{\mathrm{Lip}}(X)_{x_0}, \quad x \mapsto d(x, \cdot) - d(x, x_0),$$

is continuous and it is customary to refer to $\overline{B_{x_0}(X)}$ as the *horofunction compactification* (or *Busemann compactification*) of X; indeed, it does not depend on the basepoint x_0. The following is an extension of an unpublished result of Bader and Guralnik in the case of CAT(0) cube complexes (see eg the appendix to [8]).

Proposition 4.21 *The identity map of X extends to a homeomorphism between its Roller and Busemann compactifications.*

Proof Since $B_{x_0}(x)[z] = d(z, m(z, x, x_0)) - d(x_0, m(z, x, x_0))$, we can construct an extension of B_{x_0} taking values in the space $\mathbb{R}^X_{x_0}$ of functions $X \to \mathbb{R}$ taking x_0 to 0:

$$\widetilde{B}_{x_0}\colon \overline{X} \to \mathbb{R}^X_{x_0}, \quad \xi \mapsto d(\cdot, m(\cdot, \xi, x_0)) - d(x_0, m(\cdot, \xi, x_0)).$$

This is well defined because of the convexity of $X \subseteq \overline{X}$. If $\xi, \eta \in \overline{X}$ and $\xi \neq \eta$, there exists $\mathfrak{h} \in \mathscr{H}$ with $\widetilde{\mathfrak{h}} \in \mathscr{H}(\xi \mid \eta)$. Without loss of generality, we can assume that $x_0 \in \mathfrak{h}^*$.

Pick a point $x \in \mathfrak{h}$, and set $u := m(x_0, x, \eta)$ and $v := m(x_0, u, \xi)$; since $\mathfrak{h} \in \mathscr{H}(v \mid u)$, we have $u \neq v$. In particular,

$$\widetilde{B}_{x_0}(\xi)[u] = d(u,v) - d(x_0,v) > -d(x_0,v) > -d(x_0,u) = \widetilde{B}_{x_0}(\eta)[u].$$

This shows that $\widetilde{B}_{x_0}$ is injective; we now prove that it is continuous for the topology of pointwise convergence. Given $\xi \in \overline{X}$, $x \in X$ and $\epsilon > 0$, there exists a neighbourhood U of ξ such that, for every $\eta \in U$, the projections of ξ and η to $I(x_0, x)$ are at distance smaller than $\frac{1}{2}\epsilon$. In particular, for $\eta \in U$, the triangle inequality yields

$$|\widetilde{B}_{x_0}(\xi)[x] - \widetilde{B}_{x_0}(\eta)[x]| \leq 2 \cdot d(m(x,\xi,x_0), m(x,\eta,x_0)) < \epsilon.$$

Continuity of $\widetilde{B}_{x_0}$ and Lemma 4.12 imply that $\widetilde{B}_{x_0}(\overline{X})$ coincides with the Busemann compactification. Finally, since $\overline{X}$ is compact, $\widetilde{B}_{x_0}$ is a closed map, hence a homeomorphism. □

As a consequence of the proof of Proposition 4.21, we can define 1–Lipschitz Busemann functions for points in the Roller boundary.

Corollary 4.22 *For every $\xi \in \overline{X}$ and $x_0 \in X$, the function $X \to \mathbb{R}$ defined by*

$$z \mapsto d(z, m(z,\xi,x_0)) - d(x_0, m(z,\xi,x_0))$$

is 1–Lipschitz.

4.3 Components of the Roller boundary

Let X be a complete, locally convex median space with compact intervals. In this section, we study the structure of the median spaces arising as components of ∂X. Our first goal is to obtain the following.

Proposition 4.23 *Components of $\overline{X}$ are complete.*

To do so, we need to relate the extended metric on $\overline{X}$ to its restriction to the intervals of X.

Proposition 4.24 *For every $\xi, \eta \in \overline{X}$, we have*

$$d(\xi,\eta) = \sup_{I \in \mathscr{I}(X)} d(\pi_I(\xi), \pi_I(\eta)).$$

Proof The inequality $\geq$ follows from Lemma 4.17. Given $\epsilon > 0$, we will produce an interval $I \subseteq X$ with $d(\pi_I(\xi), \pi_I(\eta)) \geq d(\xi,\eta) - \epsilon$. By the definition of $\widehat{\nu}$, there exist points $x_1, \dots, x_n, y_1, \dots, y_n \subseteq X$ such that $\mathscr{H}(x_i \mid y_i)$ are pairwise-disjoint and

$$\sum_{k=1}^{n} \hat{\nu}((\sigma_\xi \setminus \sigma_\eta) \cap \mathscr{H}(x_i \mid y_i)) \geq d(\xi, \eta) - \epsilon.$$

Suppose that n is minimal among the integers for which such an inequality holds; we will show that $n = 1$, which will conclude the proof. Suppose for the sake of contradiction that $n \geq 2$, and set $u := m(\eta, x_1, x_2)$ and $v := m(\xi, y_1, y_2)$. Observe that

$$(\sigma_\xi \setminus \sigma_\eta) \cap (\mathscr{H}(x_1 \mid y_1) \sqcup \mathscr{H}(x_2 \mid y_2)) \subseteq (\sigma_\xi \setminus \sigma_\eta) \cap \mathscr{H}(u \mid v),$$

which, applying $\hat{\nu}$, violates the minimality of n. □

Proof of Proposition 4.23 Let $(\xi_n)_{n \geq 0}$ be a Cauchy sequence in a component of the Roller boundary. By Lemma 4.17, the sequence $\pi_I(\xi_n)$ is also Cauchy for every $I \in \mathscr{J}$ and it has a limit $\xi_I \in I$. These points define a point $\xi := (\xi_I)_I \in \varprojlim I$. By Proposition 4.24,

$$\begin{aligned} d(\xi, \xi_n) = \sup_{I \in \mathscr{J}} d(\xi_I, \pi_I(\xi_n)) &= \sup_{I \in \mathscr{J}} \lim_{m \to +\infty} d(\pi_I(\xi_m), \pi_I(\xi_n)) \\ &\leq \sup_{I \in \mathscr{J}} \limsup_{m \to +\infty} d(\xi_m, \xi_n) = \limsup_{m \to +\infty} d(\xi_m, \xi_n), \end{aligned}$$

and the latter converges to zero as n goes to infinity. □

Proposition 4.25 *Components of $\overline{X}$ have compact intervals.*

Proof Given points ξ and η in the component $Z \subseteq \overline{X}$, let π_n be the gate-projection to an interval $I_n \subseteq X$ with $\hat{\nu}((\sigma_\xi \triangle \sigma_\eta) \setminus \mathscr{H}(I_n))) \leq 1/n$; these exist for every $n \geq 1$, by Proposition 4.24. Since X has compact intervals, every sequence in $I(\xi, \eta)$ has a subsequence $(\zeta_k)_{k \geq 0}$ with the property that $(\pi_n(\zeta_k))_{k \geq 0}$ converges for every $n \geq 1$. For all $k, h, n \geq 1$, we have $d(\pi_n(\zeta_k), \pi_n(\zeta_h)) \geq d(\zeta_k, \zeta_h) - 1/n$; thus, $(\zeta_k)_{k \geq 0}$ is Cauchy, and compactness of $I(\xi, \eta)$ follows from Proposition 4.23. □

The following should better justify the terminology introduced in Definition 4.18.

Proposition 4.26 *If X is connected, each component $Z \subseteq \overline{X}$ is connected.*

Note however that Z is not a connected component of $\overline{X}$ as the latter is connected, being the closure of X.

Proof By Lemma 3.5 and Proposition 4.23, it suffices to prove that no halfspace of Z is an atom. Suppose for the sake of contradiction that there exists $\mathfrak{k} \in \mathscr{H}(Z)$

with $d(\mathfrak{k},\mathfrak{k}^*)>0$; let (ξ,η) be a pair of gates for $(\mathfrak{k},\mathfrak{k}^*)$, as provided by Lemma 2.4. Since Z is convex in $\overline{X}$, the interval between ξ and η in $\overline{X}$ consists of the sole points ξ and η. Let $\pi\colon \overline{X}\to I(\xi,\eta)=\{\xi,\eta\}$ be the corresponding gate-projection; since π is 1–Lipschitz and X is connected, we must have either $\pi(X)=\{\xi\}$ or $\pi(X)=\{\eta\}$. However, since $\xi\neq\eta$, there exists $\mathfrak{h}\in\mathscr{H}$ such that $\widetilde{\mathfrak{h}}\in\mathscr{H}(\xi\mid\eta)$; hence $\pi(\mathfrak{h}^*)=\{\xi\}$ and $\pi(\mathfrak{h})=\{\eta\}$, a contradiction. □

Observe that if Z is a component of the Roller boundary and $\mathfrak{h}\in\mathscr{H}$ is such that $\widetilde{\mathfrak{h}}\cap Z$ and $\widetilde{\mathfrak{h}}^*\cap Z$ are both nonempty, then they are halfspaces for the median-space structure of Z. The corresponding walls of Z are enough to separate points in Z. However, not all walls of the median space Z arise this way, essentially because of the fact that not every wall of the median algebra $\overline{X}$ arises from a wall of X. Still, *almost every* wall of Z comes from the above construction; see Proposition 4.29 below.

Lemma 4.27 *Let Z be a component of the Roller boundary. The sets*

$$\sigma_Z:=\{\mathfrak{h}\in\mathscr{H}\mid Z\subseteq\widetilde{\mathfrak{h}}\}\quad\text{and}\quad\mathscr{H}_Z:=\mathscr{H}\setminus(\sigma_Z\sqcup\sigma_Z^*)$$

are morally measurable.

Proof By Lemma 2.2(1), we have $\sigma_Z\cap\mathscr{H}(I)=\sigma_{\pi_I(Z)}\cap\mathscr{H}(I)$ for every interval $I\subseteq X$, and $\pi_I(Z)$ is convex. The statement now follows from Lemma 3.6. □

Lemma 4.27 allows us to define gate-projections to boundary components. Namely, if Z is a component of the Roller boundary, we have a morally measurable decomposition $\mathscr{H}=\mathscr{H}_Z\sqcup\sigma_Z\sqcup\sigma_Z^*$ and we can consider the map $\mathrm{res}_Z\colon 2^{\mathscr{H}}\to 2^{\mathscr{H}}$ that takes $E\subseteq\mathscr{H}$ to $(E\cap\mathscr{H}_Z)\sqcup\sigma_Z$. Using Lemma 4.27, it is immediate to observe that res_Z sends morally measurable ultrafilters to morally measurable ultrafilters and hence induces a map $\pi_Z\colon\overline{X}\to\overline{X}$.

Proposition 4.28 *The map π_Z is the gate-projection to the closure of Z in $\overline{X}$; this is canonically identified with the Roller compactification of Z and will be denoted unambiguously by $\overline{Z}\subseteq\overline{X}$.*

Proof By Proposition 2.1, the map π_Z is a gate-projection to some gate-convex set $C\subseteq\overline{X}$. To avoid confusion, we denote by W the closure of Z in $\overline{X}$. Given $\xi\in\overline{X}$, we have $\pi_Z(\xi)=\xi$ if and only if $\widehat{\nu}(\sigma_Z\setminus\sigma_\xi)=0$; hence, π_Z is the identity on Z, and by Lemma 2.7(1), it follows that $W\subseteq C$.

Suppose for the sake of contradiction that there exists $\xi \in C \setminus W$. By Lemma 2.6, W is gate-convex, so there exists a gate η for (ξ, W). Since $\xi \neq \eta$, we have $\pi_I(\xi) \neq \pi_I(\eta)$ for some interval $I \subseteq X$; every $\mathfrak{h} \in \mathscr{H}$ such that $\mathfrak{h} \in \mathscr{H}(\pi_I(\xi) \mid \pi_I(\eta))$ satisfies $\widetilde{\mathfrak{h}} \in \mathscr{H}(\xi \mid \eta) = \mathscr{H}(\xi \mid W)$, hence $\mathfrak{h} \in \sigma_Z \setminus \sigma_\xi$. This implies that $\widehat{\nu}(\sigma_Z \setminus \sigma_\xi) > 0$, contradicting the fact that $\xi \in C$.

We are left to identify W with the Roller compactification $\overline{Z}$. We have a continuous morphism $f \colon \overline{X} \to \overline{Z}$ mapping each $\xi \in \overline{X}$ to $(\pi_J(\xi))_{J \in \mathscr{J}(Z)}$. Since $f(Z) = Z$, Lemma 4.12 and Proposition 4.25 imply that f is surjective. Moreover, $f(\xi) = f(\eta)$ if and only if no $\mathfrak{h} \in \mathscr{H}_Z$ satisfies $\widetilde{\mathfrak{h}} \in \mathscr{H}(\xi \mid \eta)$. If ξ and η are distinct and lie in W, we have $0 < d(\xi, \eta) = \widehat{\nu}((\sigma_\xi \setminus \sigma_\eta) \cap \mathscr{H}_Z)$. Thus, the restriction of f to W is an isomorphism. □

In the rest of the section we will have to assume in addition that X has finite rank; the necessity of this will be discussed below.

Proposition 4.29 *Suppose X is a complete, finite-rank median space and let Z be a component of ∂X. Then:*

1. $\pi_Z(X) \subseteq Z$.
2. *Every thick halfspace of Z is of the form $\widetilde{\mathfrak{h}} \cap Z$ for a unique $\mathfrak{h} \in \mathscr{H}$.*
3. $\mathrm{rank}(Z) \leq \mathrm{rank}(X) - 1$.

Proof Given $x \in X$ and $\xi \in Z$, let $\{\mathfrak{h}_1, \dots, \mathfrak{h}_k\}$ be a maximal set of pairwise-transverse halfspaces in $(\sigma_\xi \setminus \sigma_x) \cap \mathscr{H}_Z$. We have

$$\begin{aligned} d(\xi, \pi_Z(x)) &= \widehat{\nu}((\sigma_\xi \setminus \sigma_x) \cap \mathscr{H}_Z) \\ &\leq \sum_{i=1}^{k} \widehat{\nu}(\mathscr{H}(x \mid \mathfrak{h}_i)) + \sum_{i=1}^{k} \widehat{\nu}(\mathscr{H}(\mathfrak{h}_i^* \mid \xi)) \\ &\leq \sum_{i=1}^{k} d(x, \mathfrak{h}_i) + \sum_{i=1}^{k} d(\xi, \widetilde{\mathfrak{h}}_i^* \cap Z) < +\infty. \end{aligned}$$

We now prove part (2). Given the partition $Z = \mathfrak{k} \sqcup \mathfrak{k}^*$ associated to a halfspace of Z, we obtain a partition of X into the convex subsets $\pi_Z^{-1}(\mathfrak{k})$ and $\pi_Z^{-1}(\mathfrak{k}^*)$. These are halfspaces of X, unless $\pi_Z(X) \subseteq \mathfrak{k}$ or $\pi_Z(X) \subseteq \mathfrak{k}^*$. We show that this cannot happen if $\mathfrak{k}$ is thick. Pick a point $\xi \in \mathfrak{k}^*$ with $d(\xi, \mathfrak{k}) > 0$; let η be the gate for $(\xi, \overline{\mathfrak{k}})$. Since $\xi \neq \eta$, there exists a halfspace $\mathfrak{h} \in \mathscr{H}$ with $\widetilde{\mathfrak{h}} \in \mathscr{H}(\xi \mid \eta) \subseteq \mathscr{H}(\xi \mid \mathfrak{k})$. Thus, $\pi_Z(x) \in \mathfrak{k}^*$ for every $x \in \mathfrak{h}^*$; in particular, $\pi_Z(X) \not\subseteq \mathfrak{k}$. A symmetric argument shows that $\pi_Z(X) \not\subseteq \mathfrak{k}^*$.

Finally, we prove part (3). Suppose for the sake of contradiction that $\operatorname{rank}(Z) \geq r = \operatorname{rank}(X)$. We have already observed that the halfspaces $\widetilde{\mathfrak{h}} \cap Z$ with $\mathfrak{h} \in \mathscr{H}_Z$ are enough to separate points of Z. Lemma 2.5 then shows that there exist $\mathfrak{h}_1, \dots, \mathfrak{h}_r \in \mathscr{H}$ such that $\mathfrak{k}_i := \widetilde{\mathfrak{h}_i} \cap Z \in \mathscr{H}(Z)$ are pairwise transverse; in particular, $\mathfrak{h}_1, \dots, \mathfrak{h}_r$ must be pairwise transverse. Pick $x \in \mathfrak{h}_1^* \cap \dots \cap \mathfrak{h}_r^*$ and $\xi \in \mathfrak{k}_1 \cap \dots \cap \mathfrak{k}_r$. Given two halfspaces $\mathfrak{h}, \mathfrak{k} \in \sigma_\xi \setminus \sigma_x \subseteq \mathscr{H}$, either $\mathfrak{h} \subseteq \mathfrak{k}$ or $\mathfrak{k} \subseteq \mathfrak{h}$ or they are transverse; observe that $\{\mathfrak{h}_1, \dots, \mathfrak{h}_r\} \subseteq \sigma_\xi \setminus \sigma_x$. If $\mathfrak{h} \in (\sigma_\xi \setminus \sigma_x) \cap \sigma_Z$, we cannot have $\mathfrak{h} \subseteq \mathfrak{h}_i$ for any i since σ_Z is a filter and $\mathfrak{h}_i \in \mathscr{H}_Z$. Moreover, since $(\sigma_\xi \setminus \sigma_x) \cap \mathscr{H}_Z$ has finite measure by part (1), the set $(\sigma_\xi \setminus \sigma_x) \cap \sigma_Z$ has infinite measure. Finally, the halfspaces $\mathfrak{h} \in (\sigma_\xi \setminus \sigma_x) \cap \sigma_Z$ such that $\mathfrak{h}_i \subseteq \mathfrak{h}$ for some i form a subset of finite measure, bounded above by the sum of the distances from x to each $\mathfrak{h}_i$; we conclude that there exists $\mathfrak{h} \in \sigma_\xi \setminus \sigma_x$ that is transverse to $\mathfrak{h}_1, \dots, \mathfrak{h}_r$, a contradiction. □

Proposition 4.29(1) can fail without the finite-rank assumption. Let X be the 0–skeleton of the CAT(0) cube complex whose vertex set is the restricted product $\{0,1\}^{(\mathbb{N})}$ and whose edges join sequences with exactly one differing coordinate. Hyperplanes are in one-to-one correspondence with natural numbers and the Roller compactification $\overline{X}$ can be identified with the unrestricted product $\{0,1\}^{\mathbb{N}}$. Let $\xi \in \overline{X}$ be the point whose coordinates are all 1; its component Z consists of sequences with only finitely many zeroes. It is immediate to observe that $\bar{Z} = \overline{X}$; in particular, π_Z is the identity on all of $\overline{X}$.

In general, even in finite rank, nonthick halfspaces of a component $Z \subseteq \partial X$ need not be of the form $\widetilde{\mathfrak{h}}$ for some $\mathfrak{h} \in \mathscr{H}$. Consider the median space X in Figure 3. It is an infinite descending staircase with steps of constant height and exponentially decreasing width; we consider X as a subset of $\mathbb{R}^2$ with the restriction of the ℓ^1 metric. It is a complete median space of rank two. Let $\xi \in \overline{X}$ be the point "at the bottom" of the staircase X, and let $Z \subseteq \overline{X}$ be its component of the Roller boundary. It is easy to notice that $\{\xi\}$ is a halfspace of Z, while, for every $\mathfrak{h} \in \mathscr{H}$, the set $\widetilde{\mathfrak{h}} \cap Z$ either does not contain ξ or contains a neighbourhood of ξ in Z.

Proposition 4.30 *Let X be a complete finite-rank median space with distinct components $Z_1, Z_2 \subseteq \partial X$ satisfying $\operatorname{rank}(Z_1) = \operatorname{rank}(Z_2) = k$. There exists a component $W \subseteq \overline{X}$ such that $\operatorname{rank}(W) \geq k+1$ and $W \cap I(\eta_1, \eta_2) \neq \varnothing$ for every $\eta_1 \in Z_1$ and $\eta_2 \in Z_2$.*

Proof As in the proof of Proposition 4.29, Lemma 2.5 yields pairwise-transverse halfspaces $\mathfrak{h}_1, \dots, \mathfrak{h}_k \in \mathscr{H}_{Z_1}$. Suppose that $\mathfrak{h}_i \in \mathscr{H}_{Z_2}$ if and only if $i \leq s$, for

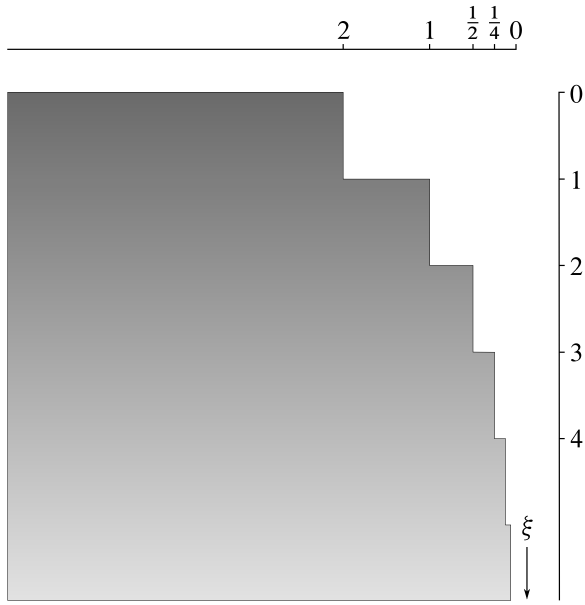

Figure 3

some $0 \le s \le k$. Similarly, let $\mathfrak{k}_1, \dots, \mathfrak{k}_k \in \mathscr{H}_{Z_2}$ be pairwise transverse, with $\mathfrak{k}_j \in \mathscr{H}_{Z_1}$ if and only if $j \le t$, for some $0 \le t \le k$.

Up to replacing some of these halfspaces with their complements, we can assume that $\mathfrak{h}_i \cap \mathfrak{k}_j \neq \varnothing$ and $\mathfrak{h}_i^* \cap \mathfrak{k}_j^* \neq \varnothing$ for all $1 \le i, j \le k$ and, in addition, $\mathfrak{k}_j^* \cap Z_1 \neq \varnothing$ and $\mathfrak{h}_i \cap Z_2 \neq \varnothing$. This can be achieved as follows. We start by ensuring that $\mathfrak{h}_i \in \sigma_{Z_2}$ and $\mathfrak{k}_j \in \sigma_{Z_1}^*$ if $i > s$ and $j > t$. If $i \le s$ and there exists $1 \le j \le k$ such that $\mathfrak{h}_i$ and $\mathfrak{k}_j$ are not transverse, we pick the side of the wall $\{\mathfrak{h}_i, \mathfrak{h}_i^*\}$ that intersects both $\mathfrak{k}_j$ and $\mathfrak{k}_j^*$; since the $\mathfrak{k}_h$ are pairwise transverse, they all determine the same side of $\{\mathfrak{h}_i, \mathfrak{h}_i^*\}$. Finally, we pick sides for the walls $\{\mathfrak{k}_j, \mathfrak{k}_j^*\}$ with $j \le t$ in a similar way. Now, Helly's theorem implies that there exist points

$$\xi_1 \in \widetilde{\mathfrak{h}}_1^* \cap \dots \cap \widetilde{\mathfrak{h}}_k^* \cap \widetilde{\mathfrak{k}}_1^* \cap \dots \cap \widetilde{\mathfrak{k}}_k^* \cap Z_1,$$
$$\xi_2 \in \widetilde{\mathfrak{h}}_1 \cap \dots \cap \widetilde{\mathfrak{h}}_k \cap \widetilde{\mathfrak{k}}_1 \cap \dots \cap \widetilde{\mathfrak{k}}_k \cap Z_2.$$

The set $\sigma_{\xi_2} \setminus \sigma_{\xi_1}$ has infinite measure, since Z_1 and Z_2 are distinct. On the other hand, the sets $\sigma_{\mathfrak{h}_i} \cap \sigma_{\xi_1}^*$, $\sigma_{\mathfrak{k}_j} \cap \sigma_{\mathfrak{h}_i^*}^*$ and $\sigma_{\xi_2} \cap \sigma_{\mathfrak{k}_j^*}^*$ all have finite measure; indeed, $d(\xi_1, \widetilde{\mathfrak{h}}_i \cap Z_1)$, $d(\mathfrak{h}_i^*, \mathfrak{k}_j)$ and $d(\xi_2, \widetilde{\mathfrak{k}}_j^* \cap Z_2)$ are finite. We conclude that there exists $\mathfrak{h} \in \sigma_{\xi_2} \setminus \sigma_{\xi_1}$ not lying in any of these sets; in particular, $\mathfrak{h}$ is either transverse to all the $\mathfrak{h}_i$ or it is transverse to all the $\mathfrak{k}_j$. Without loss of generality, let us assume that we are in the former case. By Helly's theorem, we can choose points $x_1 \in \mathfrak{h}_1^* \cap \cdots \cap \mathfrak{h}_k^* \cap \mathfrak{h}^*$ and $x_2 \in \mathfrak{h}_1 \cap \cdots \cap \mathfrak{h}_k \cap \mathfrak{h}$; we set $\xi_1' := m(\xi_1, \xi_2, x_1)$ and $\xi_2' := m(\xi_1, \xi_2, x_2)$. Observe that ξ_1' and ξ_2' belong to the interval $I(\xi_1, \xi_2)$, and by Lemma 4.17, we have that $d(\xi_1', \xi_2') \le d(x_1, x_2) < +\infty$. In particular, the points ξ_1' and ξ_2' lie in the same component of $\overline{X}$, which we denote by W. Because $\widetilde{\mathfrak{h}}_1, \dots, \widetilde{\mathfrak{h}}_k, \widetilde{\mathfrak{h}}$ all separate x_1 and x_2 and they intersect $I(\xi_1, \xi_2)$ nontrivially, they also separate ξ_1' and ξ_2'. Hence, $\mathfrak{h}_1, \dots, \mathfrak{h}_k, \mathfrak{h}$ all lie in $\mathscr{H}_W$, and $\operatorname{rank}(W) \ge k+1$.

Finally, $I(\eta_1, \eta_2)$ intersects W for all $\eta_i \in Z_i$. Indeed, projecting ξ_1' to $I(\eta_1, \eta_2)$ we only move it by a finite amount:

$$\begin{aligned} d(\xi_1', m(\eta_1, \eta_2, \xi_1')) &= d(m(\xi_1, \xi_2, \xi_1'), m(\eta_1, \eta_2, \xi_1')) \\ &\le d(\xi_1, \eta_1) + d(\xi_2, \eta_2) < +\infty. \end{aligned}$$

□

Corollary 4.31 *Every convex subset $C \subseteq \overline{X}$ intersects a unique component of $\overline{X}$ of maximal rank.*

Proof Suppose C intersects two distinct components Z_1 and Z_2 of $\overline{X}$ of maximal rank. Given $\eta_1 \in C \cap Z_1$ and $\eta_2 \in C \cap Z_2$, we have $I(\eta_1, \eta_2) \subseteq C$, and this interval intersects a component of $\overline{X}$ of strictly higher rank by Proposition 4.30, a contradiction.

□

References

[1] **H-J Bandelt**, **J Hedlíková**, *Median algebras*, Discrete Math. 45 (1983) 1–30 MR

[2] **H-J Bandelt**, **G C Meletiou**, *The zero-completion of a median algebra*, Czechoslovak Math. J. 43(118) (1993) 409–417 MR

[3] **B H Bowditch**, *Coarse median spaces and groups*, Pacific J. Math. 261 (2013) 53–93 MR

[4] **B H Bowditch**, *Embedding median algebras in products of trees*, Geom. Dedicata 170 (2014) 157–176 MR

[5] **B H Bowditch**, *Some properties of median metric spaces*, Groups Geom. Dyn. 10 (2016) 279–317 MR

[6] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999) MR

[7] **J Brodzki**, **S J Campbell**, **E Guentner**, **G A Niblo**, **N J Wright**, *Property A and* CAT(0) *cube complexes*, J. Funct. Anal. 256 (2009) 1408–1431 MR

[8] **P-E Caprace**, **J Lécureux**, *Combinatorial and group-theoretic compactifications of buildings*, Ann. Inst. Fourier (Grenoble) 61 (2011) 619–672 MR

[9] **P-E Caprace**, **M Sageev**, *Rank rigidity for* CAT(0) *cube complexes*, Geom. Funct. Anal. 21 (2011) 851–891 MR

[10] **I Chatterji**, **C Druţu**, *Median geometry for spaces with measured walls and for groups*, preprint (2017) arXiv

[11] **I Chatterji**, **C Druţu**, **F Haglund**, *Kazhdan and Haagerup properties from the median viewpoint*, Adv. Math. 225 (2010) 882–921 MR

[12] **I Chatterji**, **T Fernós**, **A Iozzi**, *The median class and superrigidity of actions on* CAT(0) *cube complexes*, J. Topol. 9 (2016) 349–400 MR

[13] **I Chatterji**, **G Niblo**, *From wall spaces to* CAT(0) *cube complexes*, Internat. J. Algebra Comput. 15 (2005) 875–885 MR

[14] **P-A Cherix**, **F Martin**, **A Valette**, *Spaces with measured walls, the Haagerup property and property* (T), Ergodic Theory Dynam. Systems 24 (2004) 1895–1908 MR

[15] **Y Cornulier**, *Irreducible lattices, invariant means, and commensurating actions*, Math. Z. 279 (2015) 1–26 MR

[16] **Y De Cornulier**, **R Tessera**, **A Valette**, *Isometric group actions on Banach spaces and representations vanishing at infinity*, Transform. Groups 13 (2008) 125–147 MR

[17] **R P Dilworth**, *A decomposition theorem for partially ordered sets*, Ann. of Math. 51 (1950) 161–166 MR

[18] **C Druţu**, **M Kapovich**, *Geometric group theory*, American Mathematical Society Colloquium Publications 63, Amer. Math. Soc., Providence, RI (2018) MR

[19] **T Fernós**, *The Furstenberg–Poisson boundary and* CAT(0) *cube complexes*, Ergodic Theory Dynam. Systems 38 (2018) 2180–2223 MR

[20] **T Fernós**, **J Lécureux**, **F Mathéus**, *Random walks and boundaries of* CAT(0) *cubical complexes*, Comment. Math. Helv. 93 (2018) 291–333 MR

[21] **E Fioravanti**, *The Tits alternative for finite rank median spaces*, Enseign. Math. 64 (2018) 89–126 MR

[22] **E Fioravanti**, *Superrigidity of actions on finite rank median spaces*, Adv. Math. 352 (2019) 1206–1252 MR

[23] **A Genevois**, *Contracting isometries of* CAT(0) *cube complexes and acylindrical hyperbolicity of diagram groups*, Algebr. Geom. Topol. 20 (2020) 49–134 MR

[24] **V Guirardel**, *Cœur et nombre d'intersection pour les actions de groupes sur les arbres*, Ann. Sci. École Norm. Sup. 38 (2005) 847–888 MR

[25] **D Guralnik**, *Coarse decompositions of boundaries for* CAT(0) *groups*, PhD thesis, Technion – Israel Institute of Technology (2005)

[26] **M F Hagen**, *The simplicial boundary of a* CAT(0) *cube complex*, Algebr. Geom. Topol. 13 (2013) 1299–1367 MR

[27] **F Haglund**, *Isometries of* CAT(0) *cube complexes are semi-simple*, preprint (2007) arXiv

[28] **F Haglund**, **F Paulin**, *Simplicité de groupes d'automorphismes d'espaces à courbure négative*, from "The Epstein birthday schrift" (I Rivin, C Rourke, C Series, editors), Geom. Topol. Monogr. 1, Geom. Topol. Publ., Coventry (1998) 181–248 MR

[29] **J R Isbell**, *Median algebra*, Trans. Amer. Math. Soc. 260 (1980) 319–362 MR

[30] **A Minasyan**, *New examples of groups acting on real trees*, J. Topol. 9 (2016) 192–214 MR

[31] **A Nevo**, **M Sageev**, *The Poisson boundary of* CAT(0) *cube complex groups*, Groups Geom. Dyn. 7 (2013) 653–695 MR

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

[33] **B Nica**, *Group actions on median spaces*, preprint (2008) arXiv

[34] **M A Roller**, *Poc sets, median algebras and group actions: an extended study of Dunwoody's construction and Sageev's theorem*, Habilitationsschrift, Universität Regensburg (1998) arXiv

[35] **W Rudin**, *Real and complex analysis*, 3rd edition, McGraw-Hill, New York (1987) MR

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

[37] **M Sholander**, *Trees, lattices, order, and betweenness*, Proc. Amer. Math. Soc. 3 (1952) 369–381 MR

[38] **M L J van de Vel**, *Theory of convex structures*, North-Holland Mathematical Library 50, North-Holland, Amsterdam (1993) MR

[39] **E R Verheul**, *Multimedians in metric and normed spaces*, CWI Tract 91, Centrum voor Wiskunde en Informatica, Amsterdam (1993) MR

[40] **L E Ward, Jr**, *On dendritic sets*, Duke Math. J. 25 (1958) 505–513 MR

[41] **R Zeidler**, *Coarse median structures and homomorphisms from Kazhdan groups*, Geom. Dedicata 180 (2016) 49–68 MR

Mathematical Institute, University of Oxford
Oxford, United Kingdom

fioravanti@mpim-bonn.mpg.de

Received: 25 August 2018 Revised: 24 May 2019