Groups Geom. Dyn. 10 (2016), 279–317
DOI 10.4171/GGD/350

Groups, Geometry, and Dynamics

Some properties of median metric spaces

Brian H. Bowditch

Abstract. We describe a number of properties of a median metric space. In particular, we show that a complete connected median metric space of finite rank admits a canonical bi-lipschitz equivalent CAT(0) metric. Metric spaces of this sort arise, up to bi-lipschitz equivalence, as asymptotic cones of certain classes of finitely generated groups, and the existence of such a structure has various consequences for the large scale geometry of the group.

Mathematics Subject Classification (2010). 20F65.

Keywords. Median metric space, median algebra, CAT(0), cube complex.

Contents

1. Introduction

In this paper, we discuss some basic properties of median metric spaces. One of the main objectives will be show that, under certain hypotheses, such a metric is bi-lipschitz equivalent to a CAT(0) metric. Median metrics have been studied by a number of authors; see for example, [20, 11] and the references therein. They arise, up to bi-lipschitz equivalence, as asymptotic cones of certain classes of groups, and the fact that these admit such a metric has various consequences for the structure of the group. We begin with some basic definitions (cf. [20, 11]), which will be elaborated upon in later sections.

Let (M, ρ) be a metric space. Given $a, b \in M$ let

$$I(a,b) = I_\rho(a,b) = \{x \in M \mid \rho(a,x) + \rho(x,b) = \rho(a,b)\}.$$

Definition. We say that (M, ρ) is a *median metric space* if, for all $a, b, c \in M$, $I(a,b) \cap I(b,c) \cap I(c,a)$ consists of exactly one element of M.

We will denote this element by $\mu(a,b,c)$. We refer to μ as the *median* induced by ρ. It turns out that (M, μ) is a median algebra. This follows by a result of Sholander [18] (see [11], and Section 2 here). For a general discussion of median algebras, see for example, [14, 2, 17]. Some further discussion relevant to this paper can be found in [6, 7]. Examples of median metric spaces include $\mathbb{R}^n$ in the l^1 metric, $\mathbb{R}$-trees, l^1 products of median metric spaces, and median subalgebras of such spaces.

We define the *rank* of a median algebra, M, to be the maximal n such that M contains a subalgebra isomorphic to the n-cube, $\{-1, 1\}^n$. If such cubes exist for all n, we say that the rank is infinite.

In the case where $\text{rank}(M) \leq n$, we will construct a new metric, $\sigma = \sigma_\rho$, canonically associated to ρ, and satisfying $\rho/\sqrt{n} \leq \sigma_\rho \leq \rho$. Among other things, we will show:

Theorem 1.1. *If (M, ρ) is a complete connected median metric space, then (M, σ_ρ) is a* CAT(0) *space.*

Given that the construction is canonical, any isometry of (M, ρ) is also an isometry of (M, σ_ρ).

The definition of a CAT(0) space will be given in Section 8. For a general discussion of such spaces, see for example, [10]. Connections with median algebras in a more combinatorial setting are described in [12].

I suspect that the assumption of completeness in Theorem 1.1 is unnecessary, and I will give a variation without this assumption (see Theorem 8.3).

We will prove a number of other results in this paper (mostly on the way to proving Theorem 1.1). For example, every complete connected median metric space is geodesic (Lemma 4.6). The convex hull of any cube of maximal (finite) rank in a median metric space is isometric to real l^1 cube (see Proposition 5.6). We also give a description of the geometry of a finite median subalgebra of a median metric space (see Section 6).

Median algebras arise in various contexts. We give some examples in Section 3. Our main motivation arises from geometric group theory, in particular, asymptotic cones of certain finitely generated groups.

An asymptotic cone of a finitely generated group is a complete metric space which captures much of the large-scale geometry of the group (see [19, 13]). It is known, for example, that any asymptotic cone of the mapping class group of a compact surface is bi-lipschitz equivalent to a median metric space [4]. More generally, the notion of a "coarse median group" was proposed in [6]. When such a group is "finitely colourable," any asymptotic cone admits a bi-lipschitz embedding into a finite product of $\mathbb{R}$-trees [7]. Its image is a connected median metric space. It follows that the asymptotic cone is bi-lipschitz equivalent to a CAT(0) space, though the CAT(0) metric might not be canonically determined by the metric on the asymptotic cone. (In fact, we can relax "finitely colourable" to "finite rank" for this particular statement [8].) In particular, it follows that the asymptotic cone is contractible. From this one can deduce that the group is FP_∞ and has polynomial isoperimetric functions in all dimensions, [16]. (For the mapping class groups, these follow by automaticity [15]. See also [3] for some more refined results.)

More generally, knowing that a space is (bi-lipschitz equivalent to) a CAT(0) space makes it easier to work with in several respects. For example, for various rigidity results, one needs information about local homology groups (Čech or singular), which require certain "straightening" constructions for continuous maps into the space, as well as the construction of homotopies between maps. Constructions of this sort are generally straightforward in a CAT(0) space, but can be technically quite complicated otherwise. These facts are exploited, for example, in [8], to give quasi-isometric rigidity results for the mapping class groups. These arguments can also be adapted to Teichmüller space [9].

We also note that median metric spaces are special cases of those discussed in [7], and so the results there also apply here. However, apart from some of the basic theory, these papers are largely independent.

We briefly outline the strategy for proving Theorem 1.1. We first consider the construction of the metric, σ_ρ. To motivate this, we note that the l^1 metric on $\mathbb{R}^n$ is a median metric, and that the euclidean metric is CAT(0). To get from the former to that latter, we can apply the Pythagorean formula to the l^1 distances in the coordinate directions. To do something similar in a general median metric space (M, ρ) we need to identify local "coordinate directions" in some sense. This will be based on the fact (Lemma 5.4) that any two points $a, b \in M$ are the diagonal of a unique maximal cube in M. Applying the Pythagorean formula to its edge-lengths in the metric ρ, we get a new "diagonal distance," $\omega(a, b)$, between these two points. To obtain the metric $\sigma_\rho(a, b)$, we consider any finite sequence of points in M, starting at a and ending at b, and sum the diagonal distances between consecutive points. We now take the infimum over all such sequences. This is the formula for σ_ρ given in Section 7. In Section 7, we check that this is indeed a metric bi-lipschitz equivalent to ρ.

We now need to verify (under appropriate connectedness and completeness assumptions) that (M, σ_ρ) is a CAT(0) space. For this, we show that (M, σ_ρ) is locally approximated by certain compact CAT(0) spaces. More precisely, if $A \subseteq M$ is finite, then A lies inside a compact CAT(0) space, $(\Upsilon, \sigma_\Upsilon)$, embedded in M, so that the metrics σ_ρ and σ_Υ agree on A to arbitrary precision (Lemma 7.8). From this, it is not hard to derive the relevant comparison statements (see Section 8).

In fact, the space $(\Upsilon, \sigma_\Upsilon)$ will be a CAT(0) cube complex, generalised so that the cells are rectilinear euclidean parallelepipeds (instead of just unit cubes). It is obtained by taking a sufficiently large finite subalgebra $\Pi \subseteq M$ containing A, and putting an appropriate metric on the realisation, $\Upsilon(\Pi)$, of Π. To get an embedding of $\Upsilon(\Pi)$ into M, we need to analyse more carefully the geometry of finite subalgebras of M. This is done in Section 6, using constructions from Section 3.

Our argument will entail some more general discussion of median algebras and median metric spaces – in particular in Sections 2 and 4. The geometry of cubes will be discussed in some detail in Section 5.

Acknowledgements. I thank Urs Lang and Hans Bandelt for their interest, and comments on earlier drafts of this paper.

2. Median metrics and median algebras

In this section, we give some basic definitions and review formulations of the notions of a median metric. Some of this can be found, expressed a little differently, in [20, 11]. We refer to [14, 2, 17, 1, 6] for further background.

First, we recall that a *median algebra* is a set, M, with a ternary operation

$$\mu\colon M^3 \longrightarrow M$$

satisfying the following for all $a, b, c, d, e \in M$:

(M1) $\mu(a,b,c) = \mu(b,c,a) = \mu(b,a,c)$,

(M2) $\mu(a,a,b) = a$, and

(M3) $\mu(a,b,\mu(c,d,e)) = \mu(\mu(a,b,c),\mu(a,b,d),e)$.

Given $a, b \in M$ we will write

$$[a,b] = [a,b]_M = \{x \in M \mid \mu(a,b,x) = x\}$$

for the median *interval* from a to b. One verifies that

(I1) for all $a \in M$, $[a,a] = \{a\}$;

(I2) for all $a, b \in M$, $[a,b] = [b,a]$;

(I3) for all $a, b \in M$, if $c \in [a,b]$, then $[a,c] \subseteq [a,b]$;

(I4) for all $a, b, c \in M$, there is a unique $m \in M$ such that

$$[a,b] \cap [b,c] \cap [c,a] = \{m\}.$$

In fact, in (I4), we have $m = \mu(a,b,c)$.

It turns out that (I1)–(I4) provide an alternative way of defining a median algebra, as follows. If we have a set M, and a map $[(a,b) \mapsto [a,b]]$ which assigns to any pair $a, b \in M$ a subset $[a,b] \subseteq M$ satisfying the above properties, then (M,μ) is a median algebra, where $\mu\colon M^3 \longrightarrow M$ is the ternary operation defined by setting

$$\mu(a,b,c) = m$$

in (I4). This follows from work of Sholander [18].

In fact, one only needs part of Sholander's paper for this. Since the logic might not be immediately apparent, a few comments are in order. To begin, we have included axiom (I2), so as not to worry about the order of a and b. (With judicious formulation of the remaining axioms, this can probably be circumvented.) We will ignore this issue henceforth. Note that Postulate Σ_1 of [18] is the conjunction of Postulates D, B_1 and F, which are respectively implied by our axioms (I4), (I1) and (I3). Now Paragraph (4.9) of [18] tells us that Σ_1 implies Postulate I of that

paper. From Paragraph (3.8) we now get that Σ_1 implies Postulates M and N, which are essentially the axioms of a median algebra (or a "median semilattice" in the terminology of that paper). This is observed in Paragraph (4.10) of [18]. In fact, for the purposes of the present paper, we could strengthen our axiom (I3) to say in addition that if $c \in [a,b]$ and $d \in [a,c]$ then $c \in [d,b]$. This property is formulated as (3.4) in [18], where its derivation from the original axioms is cited from a paper of Pitcher and Smiley. However, this assertion is almost immediate in the situation where we will want to apply it, that is in the case of a median metric. (Some further discussion of this can be found in [20].)

Note, in particular, that the median structure is completely determined by the set of intervals (though this fact can also be seen more directly).

A *subalgebra* of M is a subset closed under μ. Any finite subset of M is contained in a finite subalgebra. (This follows from the fact that the free median algebra on a finite set is finite.) A subset $C \subseteq M$ is *convex* if $[a,b] \subseteq C$ for all $a,b \in C$. The *convex hull,*

$$\mathrm{hull}(A) = \mathrm{hull}_M(A),$$

of a subset $A \subseteq M$ is the intersection of all convex subsets of M containing A. Clearly $\mathrm{hull}(A)$ is convex. A *wall* in M is an (unordered) partition of M into two disjoint non-empty convex subsets. A *homomorphism* between median algebras is a map respecting medians. Note also that a direct product of median algebras is itself a median algebra.

Given $a,b \in M$, we can define a projection map, $[x \mapsto \mu(a,b,x)]$, from M to $[a,b]$. One can verify that this is a median homomorphism. Moreover, if $c,d \in [a,b]$, then projection to $[a,b]$ followed by projection to $[c,d] \subseteq [a,b]$ agrees with projection directly to $[c,d]$.

The two-point set, $\{-1,1\}$, admits a unique median structure. By an *n-cube* in M, we mean a subalgebra isomorphic to $\{-1,1\}^n$. (Throughout this paper, the term "cube," unless otherwise qualified, will be used in this particular median sense.) The *rank* of M is the maximal n such that M contains an n-cube. The rank is deemed infinite if it contains cubes of all finite dimensions. We will refer to 2-cube as a *square*. (This is termed a "rectangle" in [11].) The following is a trivial, though useful observation:

Lemma 2.1. *A subset of $Q \subseteq M$ a square if and only if has exactly four elements, and we can cyclically order them* $\bmod 4$ *as $Q = \{a_1,a_2,a_3,a_4\}$ so that $a_i \in [a_{i-1},a_{i+1}]$ for all i.*

We will generally simply say that "a_1, a_2, a_3, a_4 is a square". We refer to the pairs a_i, a_{i+1} as its *sides*, and to the pairs a_i, a_{i+2} as the *diagonals*. We say that two ordered pairs, a, b and c, d are *parallel* if either a, b, d, c is a square, or else, $a = c$ and $b = d$. One can check that the parallel relation is transitive, so one can speak of "parallel classes".

More generally, if Q is a cube, we can define an *edge* of Q as a two-element subset of Q, which is intrinsincally convex in the subalgebra Q. A *diagonal* of Q is a two-element subset, $\{a, b\} \subseteq Q$, with $Q \subseteq [a, b]$. (Both these notions coincide with the obvious geometrical interpretations.) We remark that if a', b' is another diagonal of Q, then $[a, b] = [a', b']$, since clearly each of these sets in included in the other. In fact, this set is precisely $\mathrm{hull}_M(Q)$.

Now suppose that (M, ρ) is a metric space, with metric ρ. Given $a, b, c \in M$, write

$$\langle a, b\rangle_c = \frac{1}{2}(\rho(a, c) + \rho(b, c) - \rho(a, b)) \in [0, \infty)$$

for the "Gromov product" of a, b based at c. Thus,

$$I(a, b) = \{x \in M \mid \langle a, b\rangle_x = 0\}.$$

We write

$$S(a, b, c) = \frac{1}{2}(\rho(a, b) + \rho(b, c) + \rho(c, a)) = \langle a, b\rangle_c + \langle b, c\rangle_a + \langle c, a\rangle_b.$$

If $a, b, c, d \in M$, we write

$$T(a, b, c; d) = \rho(d, a) + \rho(d, b) + \rho(d, c).$$

Clearly $T(a, b, c; d) \geq S(a, b, c)$.

It is easily verified that the following are equivalent for $a, b, c, m \in M$:

(C1) $m \in I(a, b) \cap I(b, c) \cap I(c, a)$;

(C2) $T(a, b, c; m) = S(a, b, c)$;

(C3) $\rho(a, m) = \langle b, c\rangle_a$, $\rho(b, m) = \langle c, a\rangle_b$, and $\rho(c, m) = \langle a, b\rangle_c$.

Definition. We say that (M, ρ) is a *median metric space*, if, for all $a, b, c \in M$, there is exactly one $m \in M$ such that any (hence all) of the conditions (C1), (C2) or (C3) above hold.

In this case, we see easily that the maps $[(a,b) \mapsto I(a,b)]$ satisfy conditions (I1)–(I4) above, and so (M,μ) has the structure of a median algebra on setting $\mu(a,b,c) = m$.

We will refer to a finite sequence $a_0, a_1, \ldots, a_p$ as a *monotone sequence* if $\rho(a_0,a_p) = \sum_{i=1}^{p} \rho(a_{i-1},a_i)$. Note that this is equivalent to the median condition that $a_i \in [a_0, a_{i+1}]$ for all i. (In [I1], this is referred to as a "geodesic sequence." We use the term "monotone" since we will eventually be dealing with more than one metric.)

As in [I1], we note that squares are rectangular:

Lemma 2.2. *Suppose that (M,ρ) is a median metric space, and that a_1,a_2,a_3,a_4 is a square. Then*

$$\rho(a_1,a_2) = \rho(a_3,a_4) \quad \textit{and} \quad \rho(a_2,a_3) = \rho(a_4,a_1).$$

Proof. Let $t_i = \rho(a_i,a_{i+1})$. Then $a_{i+1} \in [a_i,a_{i+2}]$ and so $\rho(a_i,a_{i+2}) = t_i + t_{i+1}$, and we get $t_1 + t_2 = t_3 + t_4$ and $t_2 + t_3 = t_4 + t_1$. It follows that $t_1 = t_3$ and $t_2 = t_4$. □

Note that the diagonal lengths are also equal: $\rho(a_1,a_3) = \rho(a_2,a_4)$.

Lemma 2.2 can, of course, be expressed by saying that if a,b and c,d are parallel pairs in M, then $\rho(a,b) = \rho(c,d)$.

One can reformulate the discussion of metrics starting instead with median algebras. Suppose that (M,μ) is a median algebra admitting a metric ρ with the property that $\langle a,b\rangle_c = 0$ whenever $c \in [a,b]$; that is, $[a,b] \subseteq I(a,b)$. It follows that $[a,b] = I(a,b)$. For suppose $x \in I(a,b)$, and let $m = \mu(a,b,x)$. Now $\langle a,b\rangle_m = \langle a,x\rangle_m = \langle b,x\rangle_m = 0$, so $\rho(x,m) = \langle a,b\rangle_x = 0$, and so $x = m$. It follows that $x \in [a,b]$ as claimed. In other words, we see that a median metric space is the same thing as a median algebra (M,μ) with a metric ρ satisfying $\rho(a,b) = \rho(a,c) + \rho(c,b)$ whenever $a,b,c \in M$ with $c \in [a,b]$.

3. Examples of median metric spaces

In this section, we give some examples of median metric spaces which will feature in later discussions. Some related constructions have been discussed elsewhere, or have analogues in a combinatorial setting. For a recent survey, see [1].

The eventual aim of this section will be to construct a space, Φ, starting with a finite median algebra, Π, together with a collection $(\Phi_W)_W$ of median algebras,

one for each wall, W, of Π. In the particular case where each of the Φ_W is a non-trivial compact real interval, we get the realisation of Π as a (real) cube complex. We begin with a general discussion of cube complexes.

Let Υ be a CAT(0) cube complex, thought of as the topological realisation of a combinatorial cell complex. It is known that the vertex set

$$\Pi = V(\Upsilon)$$

is a median algebra. We write

$$\mathcal{W} = \mathcal{W}(\Pi)$$

for the set of walls of Π. (Alternatively, we can think of this as the set of hyperplanes of Υ.) We can also think of an element of $\mathcal{W}$ as corresponding to a parallel class of edges of Π. (The set of edges which cross any given wall will constitute such a parallel class.) Suppose we have a function, $\lambda\colon \mathcal{W} \longrightarrow (0, \infty)$. This gives rise to a path-metric on the 1-skeleton, so that if $a, b \in \Pi = V(\Upsilon)$ then $\rho(a, b) = \sum_W \lambda(W)$ where W ranges over the set of walls of Π which separate a from b. This naturally extends to a path-metric, ρ, on all of Υ, where each n-cell is given the structure of a rectilinear parallelepided in $\mathbb{R}^n$ with the l^1 metric. Note that (Υ, ρ) is uniquely determined up to a cell-preserving isometry.

A more formal way to describe this is as follows. Let $Q(\Pi)$ be the cube consisting of the direct product $\prod \mathcal{W}$, where each $W \in \mathcal{W}$ is viewed formally as a two-point median algebra. (We will only really need to consider the case where Π is finite.) There is a natural embedding of Π into $Q(\Pi)$. In this way, Υ can be seen as the full subcomplex of $Q(\Pi)$ with vertex set Π. Given our map $\lambda\colon \mathcal{W} \longrightarrow (0, \infty)$, let

$$P = \prod_{W \in \mathcal{W}} [0, \lambda(W)].$$

This is a median metric space in the l^1 metric, and Υ is a subalgebra, and itself a median metric in the induced path-metric.

Note that we could also put a euclidean structure on Υ. For this, we start in the same way, putting the same metric on the 1-skeleton, but instead of taking the l^1 metric on each cell, we take the euclidean metric. (Or in the formulation of the previous paragraph, we take the path-metric induced from the euclidean metric on P.) This gives us a path-metric σ with $\sigma \leq \rho$. It is also easy to see that it is CAT(0). (Thus usual construction demands that we take all sides to be of unit length, but the same holds in this more general situation. The links are all CAT(1) spaces.) Note also that, if Π has rank n (or equivalently, Υ has dimension n) then $\rho \leq \sigma\sqrt{n}$.

Terminology. Before continuing, we clarify the following terminology used throughout this paper. As noted, a *cube*, Q, is a direct product of two-point median algebras. An *edge* of Q is a convex two-element subset (a factor of Q). We refer to the realisation, P, of Q as a product of non-trivial compact real intervals as a *real cube*, which we can again view as a median algebra. We refer to the elements of $Q \subseteq P$ as the *corners* of Q. Given any edge $\{a, b\} \subseteq Q$, we refer to $[a, b] \subseteq P$ as a *side* of P; that is, a 1-face of the real cube.

In fact, any finite median algebra, Π, can be identified with the vertex set of a finite CAT(0) cube complex, as follows. Let $\mathcal{W} = \mathcal{W}(\Pi)$ be the set of walls in W. There is a natural embedding of Π into $Q(\Pi)$, and we can think of $\Upsilon(\Pi)$ as the full subcomplex with vertex set Π. Note that Υ is unique up to isomorphism. We will denote it by $\Upsilon(\Pi)$. (Thus, $\Upsilon(\Pi)$ is a subcomplex of $\Upsilon(Q(\Pi))$.) Note that each face of Π corresponds to a cell of $\Upsilon(\Pi)$.

If Π is a finite median metric space, then we also have a function $\lambda\colon \mathcal{W} \longrightarrow (0, \infty)$ which assigns the distance between the vertices of any edge crossing a given wall. By Lemma 2.2, this is well defined. From this we see that any finite median metric space can be embedded in such a complex. We will return to this construction at the end of Section 7, see Lemma 7.6.

We now proceed to a more general construction. We will need a formal description of "binary subdivision" as follows. To begin, let

$$Q = \{-1, 1\}^n$$

be the standard n-cube. Its realisation, $\Upsilon(Q)$, is a real n-cube. Let

$$F(Q) = \{-1, 0, 1\}^n.$$

Then $F(Q)$ is also a median algebra, containing Q as a subalgebra. We will think of $F(Q)$ as corresponding to the set of faces of $\Upsilon(Q)$. (It might also be thought of as the vertex set of the binary subdivision of $\Upsilon(Q)$.) Given $s, t \in F(Q)$ we will write $t \preceq s$ to mean that t corresponds to a face of s. (Formally, this means that s can be obtained from t be resetting some of the ± 1 coordinates equal to 0.) Given $s \in F(Q)$, let

$$Q(s) = \{t \in Q \mid t \preceq s\}.$$

In other words, $Q(s)$ is the face of Q corresponding to s. Note that if $r, s \in F(Q)$, then $r \preceq s$ if and only if $Q(r) \subseteq Q(s)$.

Now let Π be any finite median algebra, and let

$$Q(\Pi) = \prod \mathcal{W}$$

be as above. Given $W \in \mathcal{W}$, we write

$$W = \{H^+(W), H^-(W)\},$$

where the $\pm$ signs are assigned arbitrarily. In this way, the formal product, $Q(\Pi) = \prod \mathcal{W}$, can be identified with $\{-1, 1\}^{\mathcal{W}}$. Let $F(\Pi) \subseteq F(Q(\Pi))$ correspond to the set of faces of Π. Formally, we can set

$$F(\Pi) = \{s \in F(Q(\Pi)) \mid Q(s) \subseteq \Pi\}.$$

Note that if $s \in F(\Pi)$ and $t \in F(Q(\Pi))$ with $t \preceq s$, then $t \in F(\Pi)$.

We now proceed to the main construction of this section. It will be used in Section 6 – see, in particular, Lemma 6.2. The general idea is that, to each wall, W, of Π, we are prescribed a median algebra, Φ_W (which is intrinsically a median interval). From this data, we construct a space, Φ, which can be thought of as a kind of cell complex. The cells will be direct products of the Φ_W: one cell for each cell of $\Upsilon(\Pi)$, and glued together in the same combinatorial pattern. Indeed, if we take each Φ_W to be a non-trivial compact real interval, then we recover precisely $\Upsilon(\Pi)$.

Suppose, then that to each $W \in \mathcal{W}(\Pi)$, we have been assigned a median algebra, Φ_W, together with elements $p^-_W, p^+_W \in \Phi_W$ with $p^-_W \neq p^+_W$ and with $\Phi_W = [p^-_W, p^+_W]$. Let

$$P = \prod_{W \in \mathcal{W}} \Phi_W$$

be the product median algebra. We can identify $Q(\Pi) \equiv \{-1, 1\}^{\mathcal{W}}$ as a subalgebra $\prod_{W \in \mathcal{W}} \{p^-_W, p^+_W\} \subseteq P$. In this way, each element $s \in F(\Pi)$ gets canonically associated to a convex subset $P(s)$ of P, namely, $P(s) = \mathrm{hull}_P(Q(s))$. (This is a direct product of those Φ_W which correspond to the walls crossing $Q(s)$.) Let

$$\Phi = \bigcup_{s \in F(\Pi)} P(s).$$

We claim that Φ is a subalgebra of P.

To see this, define a map

$$h_W \colon \Phi_W \longrightarrow \{-1, 0, 1\}$$

for each $W \in \mathcal{W}$, by setting

$$h_W(p_W^{\pm}) = \pm 1$$

and

$$h_W(x) = 0 \quad \text{for all } x \in \Phi_W \setminus \{p_W^-, p_W^+\}.$$

(Note that h_W might not be a median homomorphism: for example, if $a, b \in \Phi_W$ with $\mu(a, b, p_W^+) = p_W^+$, then

$$\mu(h_W a, h_W b, h_W p_W^+) = \mu(0, 0, +1) = 0 \neq +1 = h_W(p_W).$$

This is however, essentially the only way it can fail.)

The maps h_W give rise to a map

$$h \colon P \longrightarrow F(Q(\Pi)) \equiv \{-1, 0, 1\}^{\mathcal{W}},$$

by taking h_W on each coordinate. By construction,

$$\Phi = h^{-1}(F(\Pi)).$$

Now (as we have observed) h need not be a median homomorphism. Nevertheless, if $a, b, c \in P$, then $h(\mu_P(a, b, c)) \preceq \mu_F(ha, hb, hc)$. Thus, if $a, b, c \in \Phi$, then $ha, hb, hc \in F(\Pi)$, so $\mu_F(ha, hb, hc) \in F(\Pi)$, since $F(\Pi)$ is a subalgebra of $F(Q(\Pi))$. Also $h(\mu_P(a, b, c)) \preceq \mu_F(ha, hb, hc)$, so $h(\mu_P(a, b, c)) \in F(\Pi)$, and we get $\mu_P(a, b, c) \in h^{-1}(F(\Pi)) = \Phi$ as required.

Note that if each Φ_W is a median metric space, then one can put a geodesic metric on P so that its restriction to each cell is the l^1 metric on the product. It is not hard to see that, in this structure, P is also median metric space, inducing the same median. (We omit details, since we will not formally need this fact.)

As an example of this construction, if each Φ_W is a non-trivial compact real interval, we recover the realisation,

$$\Phi = \Upsilon(\Pi)$$

of Π. In fact, if $\lambda \colon \mathcal{W}(\Pi) \longrightarrow (0, \infty)$ is any map (as above), we can set

$$\Phi_W = [0, \lambda(W)]$$

with $p_W^- = 0$ and $p_W^+ = \lambda(W)$. This gives $\Phi = \Upsilon(\Pi)$ naturally equipped with the l^1 metric ρ as described earlier.

4. Basic properties of median metric spaces

We discuss some of the basic properties of median metric spaces. We begin with the following general construction in a median algebra (M, μ).

Suppose that $a, b, c, d \in M$. Let

$$a' = \mu(b, c, d), \qquad b' = \mu(c, d, a),$$

$$c' = \mu(d, a, b), \qquad d' = \mu(a, b, c),$$

and

$$a'' = \mu(b', c', d'), \quad b'' = \mu(c', d', a'),$$

$$c'' = \mu(d', a', b') \quad d'' = \mu(a', b', c').$$

Now,

$$Q = \{a', b', c', d', a'', b'', c'', d''\}$$

is a (possibly degenerate) cube, with diagonals $\{a', a''\}$, $\{b', b''\}$, $\{c', c''\}$ and $\{d', d''\}$. Moreover, $Q \cup \{a, b, c, d\}$ is a subalgebra of M (Figure 1).

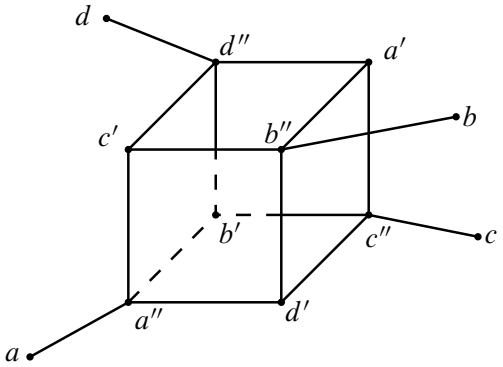

Figure 1. The median subalgebra generated by $\{a, b, c, d\}$

Of course, this can be verified directly from the axioms, but it is more naturally viewed as follows. The free median algebra, F, on a set of four elements, $\{a_0, b_0, c_0, d_0\}$, is the vertex set of a CAT(0) cube complex, consisting of a central 3-cube with four "free" sides attached to alternating corners. The points a_0, b_0, c_0, d_0 are identified with the terminal vertices of these free sides. We write a_0'' for the vertex of the cube adjacent to a_0, and a_0' for the vertex of the cube opposite a_0'' etc. Thus, the central 3-cube is

$$Q_0 = \{a_0', b_0', c_0', d_0', a_0'', b_0'', c_0'', d_0''\}.$$

The homomorphism from F to M sending a_0 to a etc. maps Q_0 to Q in the obvious way. In general, this map might not be injective. For example, Q_0 might collapse to a lower dimensional cube. (This is what we meant by a "possibly degenerate" cube.)

We now assume that (M, ρ) is a median metric space. The sides of the cube Q fall into three parallel classes, with all sides in a parallel class of equal length.

One immediate consequence is the following fact, proven in [11]:

Lemma 4.1. *If $a, b, c, d \in M$, then*

$$\rho(\mu(a,b,c), \mu(a,b,d)) \leq \rho(c,d).$$

Proof. In the above notation, c, c'', d'', d is a monotone sequence, and the pairs c'', d'' and d', c' are parallel. Therefore $\rho(c', d') = \rho(c'', d'') \leq \rho(c, d)$ as required. □

In particular, we see that if $c, d \in [a, b]$, then $\rho(c, d) \leq \rho(a, b)$.

We also note that this implies that $\mu\colon M^3 \longrightarrow M$ is continuous. In other words, (M, μ) is a topological median algebra.

Recalling the notation of Section 2, we also have:

Lemma 4.2. *If $a, b, c, d \in M$, then $\rho(d, \mu(a,b,c)) \leq T(a,b,c;d) - S(a,b,c)$.*

Proof. In the above notation, $d' = \mu(a,b,c)$. Let $A = \rho(a', d'')$, $B = \rho(b', d'')$ and $C = \rho(c', d'')$; that is, A, B, C are the three side-lengths of the central cube. Let $A_0 = \rho(a, a'')$, $B_0 = \rho(b, b'')$, $C_0 = \rho(c, c'')$ and $D_0 = \rho(d, d'')$; that is, the lengths of the four free sides (Figure 2)

Now,

$$\rho(a,b) = A_0 + B_0 + A + B$$

etc. and so

$$S(a,b,c) = A_0 + B_0 + C_0 + A + B + C.$$

Also,

$$\rho(d,a) = D_0 + A_0 + B + C$$

etc. and so

$$T(a,b,c;d) = 3D_0 + A_0 + B_0 + C_0 + 2(A + B + C)$$

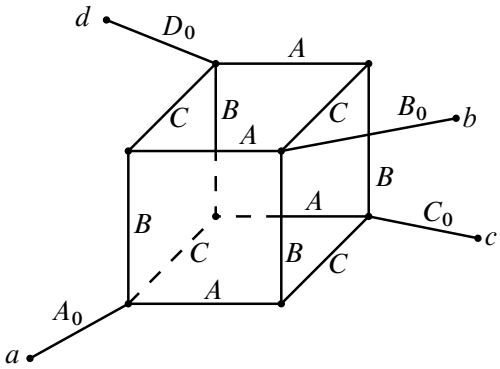

Figure 2. Distances in the metric ρ

Finally,

$$\begin{aligned}\rho(d,d') &= D_0 + A + B + C \\ &\leq 3D_0 + A + B + C \\ &= T(a,b,c;d) - S(a,b,c).\end{aligned}$$

□

We obtain the following, proven in [20] and [11].

Lemma 4.3. *The metric completion of a median metric space is a median metric space.*

Proof. Let (M,ρ) be a median metric space, and let $(\bar{M},\bar{\rho})$ be its completion. Let $a,b,c \in \bar{M}$, and choose sequences, a_i, b_i, c_i in M, with $a_i \to a$, $b_i \to b$ and $c_i \to c$. Let $m_i = \mu(a_i,b_i,c_i)$. By Lemma 4.1,

$$\rho(m_i,m_j) \leq \rho(a_i,a_j) + \rho(b_i,b_j) + \rho(c_i,c_j),$$

so $(m_i)_i$ is Cauchy, so m_i tends to some $m \in \bar{M}$. By continuity, we see that $m \in I_{\bar{\rho}}(a,b) \cap I_{\bar{\rho}}(b,c) \cap I_{\bar{\rho}}(c,a)$. Suppose that $d \in I_{\bar{\rho}}(a,b) \cap I_{\bar{\rho}}(b,c) \cap I_{\bar{\rho}}(c,a)$. Now

$$T(a_i,b_i,c_i;d) \to T(a,b,c;d) = S(a,b,c).$$

By Lemma 4.2,

$$\rho(d,m_i) \leq T(a_i,b_i,c_i;d) - S(a_i,b_i,c_i) \to 0,$$

so $\rho(d,m) = 0$, so $d = m$. □

Finally, we consider connectedness of a median metric space, M. If $J \subseteq \mathbb{R}$ is an interval, we say that a continuous path, $\gamma\colon J \longrightarrow M$ is *monotone* if $\gamma(v) \in [\gamma(t), \gamma(u)]$ whenever $t, u, v \in J$ with $t \leq v \leq u$. (This ties in with the notion of a "monotone sequence" defined in Section 2.) Note that if such a path exists, we can assume that it is injective, and we can reparametrise so that it is a ρ-geodesic, that is, $\rho(\gamma(t), \gamma(u)) = |t - u|$ for all $t, u \in J$. Conversely, any ρ-geodesic will be monotone. Recall that a *geodesic space* is a metric space in which any pair of points can be connected by a geodesic. We see that a metric median space is geodesic if and only if every pair of points can be connected by a monotone path.

The following is an easy consequence of the fact that the projection $[x \mapsto \mu(a, b, x)]$ from M to $[a, b]$ is 1-lipschitz for all $a, b \in M$, and the fact that $[c, d] \subseteq [a, b]$ for all $c, d \in [a, b]$ (that is, intervals are convex).

Lemma 4.4. *Let (M, ρ) be a median metric space. Then (M, ρ) is connected (respectively path connected; respectively geodesic) if any only if, for all $a, b \in M$, the interval $[a, b]$ is connected (respectively path connected; respectively geodesic).*

We suspect that these notions are all equivalent. We can certainly make the following observation:

Lemma 4.5. *Let (M, ρ) be a connected median metric space. Suppose that $a, b \in M$ and that $0 < t < \rho(a, b)$. Then there exists $c \in M$ with*

$$\rho(a, c) = t$$

and

$$\rho(b, c) = \rho(a, b) - t.$$

Proof. Suppose not. Let

$$U = \{x \in M \mid \rho(a, \mu(a, b, x)) < t\}$$

and

$$V = \{x \in M \mid \rho(b, \mu(a, b, x)) < \rho(a, b) - t\}.$$

Then U and V are open, $M = U \sqcup V$, $a \in U$ and $b \in V$, contradicting the fact that M is connected. □

In particular, setting $t = \frac{1}{2}\rho(a,b)$, we see that any pair of points of M must have a midpoint. A standard completion argument now shows:

Lemma 4.6. *Any complete connected median metric space is geodesic.*

In fact, as was pointed out to me by Hans Bandelt, a complete median metric space, M, is geodesic if and only if it has the "Menger property," which in this context means that $[a,b] \neq \{a,b\}$ for all distinct $a,b \in M$. Indeed, any complete metric space (M,ρ) is geodesic if and only if $I_\rho(a,b) \neq \{a,b\}$ for all distinct $a,b \in M$.

5. Cubes in median algebras

In the next two sections we describe some general median algebra constructions which we apply to median metric spaces in Sections 7 and 8. In the present section we make some observations about cubes in median algebras. Some related statements can found in, or derived from, the literature on distributive lattices. However, since much of this is presented in a form not so readily accessible to geometers, we give a self-contained account here.

Let (M,μ) be a median algebra, and let $a,b \in M$. In this section, we will adopt the convention that the interval denoted $[a,b]$ has a preferred "initial point," a, and "terminal point," b. Given $x,y \in [a,b]$, we write

$$x \wedge y = \mu(a,x,y)$$

and

$$x \vee y = \mu(b,x,y).$$

Then $([a,b],\wedge,\vee)$ is a distributive lattice. We write

$$x \leq y \quad \text{to mean that } x \wedge y = x,$$

or equivalently $x \vee y = y$. Then $\leq$ is a partial order on $[a,b]$, with minimum a and maximum b. We can recover the median on $[a,b]$ from the lattice structure as

$$\mu(x,y,z) = (x \wedge y) \vee (y \wedge z) \vee (z \wedge x) = (x \vee y) \wedge (y \vee z) \wedge (z \vee x).$$

In particular, any lattice homomorphism will be a median homomorphism.

Suppose that $e_1, e_2, \ldots, e_n \in [a,b]$. We write

$$I = \{1, 2, \ldots, n\}.$$

Given $J \subseteq I$ write

$$e_J = \bigvee_{i \in J} e_i,$$

with the convention that $e_\emptyset = a$. We say that $(e_i)_i$ *spans* $[a,b]$ if $e_I = b$. In this case, if $x \in [a,b]$, then $x = x \wedge b = \bigvee_{i \in I}(x \wedge e_i)$.

Let $P = \prod_{i \in I}[a, e_i]$ be the product median algebra. We define a map

$$\theta \colon [a,b] \longrightarrow P,$$

by setting

$$\theta(x) = (x \wedge e_1, x \wedge e_2, \ldots, x \wedge e_n),$$

and a map

$$\phi \colon P \longrightarrow [a,b]$$

by setting

$$\phi(x_1, x_2, \ldots, x_n) = \bigvee_{i \in I} x_i.$$

Now, by the above, if $(x_i)_i$ spans $[a,b]$, then $\phi \circ \theta$ is the identity on $[a,b]$. Note that P is itself intrinsically an interval, namely

$$P = [\theta(a), \theta(b)],$$

and hence also a distributive lattice. Indeed the lattice structure is the same as that induced from those on $[a, e_i]$ by defining $\wedge$ and $\vee$ coordinatewise. Note that for all $x, y \in [a,b]$, we have

$$\theta(x \wedge y) = \theta(x) \wedge \theta(y) \quad \text{and} \quad \theta(x \vee y) = \theta(x) \vee \theta(y).$$

It follows that θ is a monomorphism from $[a,b]$ into P.

We say that $(e_i)_i$ is *independent* if $e_i \wedge e_j = a$ whenever $i \neq j$. In this case, if $x \in [a, e_i]$ and $y \in [a, e_j]$, then

$$x \wedge y = (x \wedge e_i) \wedge (y \wedge e_j) = x \wedge y \wedge (e_i \wedge e_j) = a.$$

Now, if $\underline{x} = (x_1, \ldots, x_n) \in P$, then

$$\phi(\underline{x}) \wedge e_j = \Big(\bigvee_{i \in I} x_i \Big) \wedge e_j = x_j.$$

We now see that

$$\theta \circ \phi(\underline{x}) = \underline{x},$$

and so θ and ϕ are inverse homomorphisms. We have shown:

Lemma 5.1. *Suppose $a, b \in M$ and that $e_1, \ldots, e_n \in [a, b]$, where $(e_i)_i$ are independent and span $[a, b]$. Then $[a, b]$ is naturally isomorphic to the product median algebra $\prod_{i \in I} [a, e_i]$.*

In particular, this gives a monomorphism

$$\phi \colon \prod_{i \in I} \{a, e_i\} \longrightarrow [a, b],$$

with image

$$Q = \{e_J \mid J \subseteq I\}.$$

Thus, Q is an n-cube in $[a, b] \subseteq M$. Intrinsically, Q is the interval $[a, b] \cap Q$. In other words, a, b is a diagonal of the cube. Note that

$$e_J \wedge e_K = e_{J \cap K} \quad \text{and} \quad e_J \vee e_K = e_{J \cup K}$$

for $J, K \subseteq I$.

Suppose that $e_i = c \vee d$, with $c \wedge d = a$, for some $c, d \in [a, b] \setminus \{a\}$. Then a, c, e_i, d is a square, and we could replace e_i by c, d to obtain a larger independent spanning set for $[a, b]$. With this in mind, we say that $(e_i)_i$ is *maximal* if no $\{a, e_i\}$ is the diagonal of a square.

Definition. A *basis* for $[a, b]$ is a maximal independent spanning set for $[a, b]$.

Lemma 5.2. *Suppose that $e_1, \ldots, e_n$ is a basis for $[a, b]$, and that $e'_1, \ldots, e'_m$ is an independent spanning set. Then we can partition I into non-empty subsets as*

$$I = I_1 \sqcup \cdots \sqcup I_m$$

such that for all i we have

$$e'_i = e_{I_i} = \bigvee_{j \in I_i} e_j.$$

Proof. Given $j \in I$, we have

$$e_j = \bigvee_{i=1}^{m} (e_j \wedge e'_i).$$

In particular, there must be some $i \in \{1, \ldots, m\}$ with

$$e_j \wedge e'_i \neq a.$$

Note that

$$e_j = c \vee d \quad \text{and} \quad a = c \wedge d,$$

where

$$c = e_j \wedge e_i' \quad \text{and} \quad d = \bigvee_{k \neq i} (e_j \wedge e_k').$$

Since $c \neq a$, we must have

$$d = a$$

(otherwise, $\{a, c, e_j, d\}$ would be a square), and so

$$e_j \wedge e_k' = a \quad \text{for all } k \neq i.$$

It now also follows that

$$e_j = c.$$

In other words, for all $j \in I$, there is a unique $i(j) \in \{1, \ldots, m\}$ with

$$e_j \wedge e_{i(j)}' = e_j \quad \text{and} \quad e_j \wedge e_i' = a \quad \text{whenever } i \neq i(j).$$

Given $i \in \{1, \ldots, m\}$, let

$$I_i = \{j \in I \mid i(j) = i\}.$$

Clearly

$$I_i \cap I_k = \emptyset \quad \text{if } i \neq k.$$

Moreover, if $i \in I$, then since $e_j \wedge e_i' = a$ for all $j \notin I_i$, we have

$$e_i' = \bigvee_{j \in I} (e_j \wedge e_i') = \bigvee_{j \in I_i} (e_j \wedge e_i') = e_{I_i}.$$

Finally note that

$$b = \bigvee_{i=1}^{m} e_i' = \bigvee_{i=1}^{m} e_{I_i} = e_{\bigcup_i I_i}.$$

It follows that $I = \bigcup_i I_i$, for if $k \in I \setminus \bigcup_i I_i$, we would get the contradiction that $e_k = b \wedge e_k = e_{\{k\} \cap \bigcup_i I_i} = e_\emptyset = a$. □

Corollary 5.3. *Any two bases for* $[a, b]$ *agree up to permutation.*

Proof. Let $(e_i)_{i=1}^n$ and $(e_i')_{i=1}^m$ be bases. By Lemma 5.2, we see that $m = n$, and that each I_i is a singleton. □

Suppose now that $a, b \in M$, and that $Q \subseteq [a, b]$ is a cube with $a, b \in Q$, that is a, b is a diagonal of Q. If $e_1, \ldots, e_n \in Q$ are the points adjacent to a in Q (that is the sets $\{a, e_i\}$ are the edges of Q containing a), then $(e_i)_i$ are an independent spanning set of $[a, b]$ in M. Conversely, we have already observed that if $e_1, \ldots, e_n$ are independent and span a, b, then Q is isomorphic to $\prod_{i=1}^n \{a, e_i\}$. We see that a, b is not the diagonal of a strictly larger cube. (In other words, if $Q \subseteq Q'$ where Q' is a cube containing a, b, then $Q = Q'$.) This is the same as saying that $(e_i)_i$ is a basis for $[a, b]$. As a consequence of Corollary 5.3, we have:

Lemma 5.4. *Given* $a, b \in M$*, there is at most one maximal cube in* M *with diagonal* a, b*.*

If M has finite rank, then such a cube must always exist – take a cube of maximal possible rank which has diagonal a, b.

Note also that Lemma 5.2 can be interpreted as saying that if Q is a maximal cube with diagonal a, b, then any other cube with diagonal a, b must be a subcube of Q.

If we start with a cube, Q, then $\mathrm{hull}_M(Q) = [a, b]$ where a, b is any diagonal of Q. We can associate to each wall $W \in \mathcal{W}(Q)$, a median algebra Φ_W, well defined up to isomorphism (and inclusion into M up to the relation of parallelism). In fact, we can let $\Phi_W = [a, e_i] = \mathrm{hull}\{a, e_i\}$, where $\{e_1, \ldots, e_n\}$ is the basis for $Q = [a, b] \cap Q$, and where $\{a, e_i\}$ is the edge crossing Q. From the above, we see:

Lemma 5.5. *The convex hull,* $\mathrm{hull}(Q)$*, of the cube* $Q \subseteq M$ *is naturally isomorphic to the product median algebra* $\prod_{W \in \mathcal{W}(Q)} \Phi_W$*.*

Although it is not needed for the proof of the main theorem, for applications elsewhere, we will show that the convex hulls of cubes of maximal dimension are isomorphic to real cubes.

First note that if $c, d \in [a, b]$, then $[c, d] \subseteq [a, b]$. Moreover, if $c = a$ then the order on $[a, d]$ is precisely the restriction of the partial order on $[a, b]$. We also note that an interval $[c, d] \subseteq M$ has (intrinsic) rank 1 if and only if the partial order on $[c, d]$ is a total order. In this case, we will say that the interval $[c, d]$ is *linear*.

Now if $(e_i)_i$ is an independent spanning set of $[a, b]$ and that $[a, e_i]$ is linear for each i. Clearly, this implies that $(e_i)_i$ is a basis.

In fact, suppose rank $M = n < \infty$, and $e_1, \ldots, e_n$ in an independent basis for $[a, b]$, then, $[a, e_i]$ is linear for each i. For if not, we can find a square $S \subseteq [a, e_i]$ for some i. Using Lemma 5.1, we see that $S \times \prod_{j \neq i} \{a, e_j\}$ is a cube of dimension $n + 1$, giving a contradiction.

Putting the above facts together we conclude:

Proposition 5.6. *Suppose that M is a connected metrisable topological median algebra and that $Q \subseteq M$ is a cube in M whose dimension equals* $\mathrm{rank}(M) < \infty$. *Then* $\mathrm{hull}(Q)$ *is isomorphic to a real n-cube.*

Proof. This follows from the topological characterisation of a real interval as a connected metrisable space with exactly two non-cut points. □

Note that if M is a connected median metric space, then this cube will be isometric to a compact real n-cube with an l^1 metric. (We have already noted that a median metric space is a topological median space, by Lemma 4.1.)

6. Finite subalgebras of median algebras

In this Section, we show how to associate, to any finite subalgebra, Π, of any median algebra, M, a larger subalgebra, $M(\Pi) \supseteq \Pi$, which has a kind of "cellular" structure, where the "cells" are product median spaces in bijective correspondence with the faces of the complex Π (or equivalently, the cells of $\Upsilon(\Pi)$). This construction will be used in the proof of Lemma 7.6.

First we make the following definition. Suppose that $A, B \subseteq M$. A *parallel map* between A and B is a bijection $\psi : A \longrightarrow B$ such that, for all $x, y \in A$, $\psi x, \psi y$ is parallel to x, y. Clearly its inverse is also parallel. It is not hard to see that, if A is a subalgebra, then so is B, and ψ is a median isomorphism between them (though this will be clear in the case of interest here). In particular, if a, b, d, c is a square in M, then the projection of M to $[a, b]$ defined by

$$\phi(x) = \mu(a, b, x)$$

restricted to $[c, d]$ is a parallel map from $[c, d]$ to $[a, b]$. Moreover, the projection of M to $[c, d]$ composed with $\psi|[c, d]$ is also equal to ψ.

Suppose that $\Pi \subseteq M$ is a finite subalgebra, and let $F(\Pi)$ be the set of faces as described in Section 3. Let

$$M(\Pi) = \bigcup_{s \in F(\Pi)} \text{hull}(Q(s)).$$

We claim:

Lemma 6.1. *$M(\Pi)$ is a subalgebra of M.*

In fact, we can describe the structure of $M(\Pi)$ in terms of the construction of Section 5. Given $W \in \mathcal{W}(\Pi)$, let $\{p_W^-, p_W^+\}$ be an edge of Π which crosses W. Let

$$\Phi_W = [p_W^-, p_W^+].$$

Note that Φ_W is well defined up to a parallel map in M, and that there is a well defined projection,

$$\theta_W : M \longrightarrow \Phi_W$$

(namely, $\theta_W(x) = \mu(p_W^-, p_W^+, x)$). Note that θ_W is an epimorphism. Let Φ be the subalgebra of $P = \prod_{W \in \mathcal{W}(\Pi)} \Phi_W$, as defined in Section 3.

We also get a natural map

$$\phi : \Phi \longrightarrow M$$

as follows. Note that if $s \in F(\Pi)$, then by Lemma 5.5, there is a natural isomorphism,

$$\phi_s : P(s) \longrightarrow \text{hull}(Q(s)).$$

In fact, this is naturally isomorphic to $\prod_{W \in \mathcal{W}_s} \Phi_W$, where $\mathcal{W}_s = \mathcal{W}(Q(s)) \subseteq \mathcal{W}$ is the set of walls crossing s. Note that if $t \preceq s$, then $P(t) \subseteq P(s)$ and $\phi_t = \phi_s | P(s)$. Assembling these, we get a map $\phi : \Phi \longrightarrow M$. By construction,

$$\phi(\Phi) = M(\Pi).$$

We claim:

Lemma 6.2. *$\phi : \Phi \longrightarrow M$ is a monomorphism.*

Clearly this implies Lemma 6.1.

Lemma 6.2 will follow from the following observation. Note that if $W \in \mathcal{W}(\Pi)$, we have a projection, $\psi_W \colon P \longrightarrow \Phi_W$, namely the median projection (defined as for M), or equivalently, simply the coordinate projection to the factor Φ_W of P. We claim:

Lemma 6.3. *If $W \in \mathcal{W}(\Pi)$, then*

$$\psi_W = \theta_W \circ \phi.$$

Proof. Let $x \in \Phi$, and set $s = h(x) \in F(\Pi)$, where $h \colon \Phi \longrightarrow F(\Pi)$ is the map defined in Section 3. We distinguish two cases.

Suppose $W \in \mathcal{W}_s$. In this case, the statement follows from the fact that $\phi_s \colon P(s) \longrightarrow \mathrm{hull}(Q(s))$ is an isomorphism.

Suppose that $W \notin \mathcal{W}_s$. In this case, without loss of generality, we have

$$Q(s) \subseteq H^-(W) \cap \Pi.$$

Let $\{p^-_W, p^+_W\}$ be an edge of Π crossing W with $p^\pm_W \in H^\pm(W)$. Since Π is a subalgebra of M, we see that $\theta_W(Q(s)) = \{p^-_W\}$ (since the projection of $Q(s)$ to $\{p^-_W, p^+_W\}$ in Π is $\{p^-_W\}$). Now $\phi(x) = \phi_s(x) \in \mathrm{hull}(Q(s))$, by the construction of ϕ, and so $\theta_W(\phi(x)) = p^-_W$. (This follows from the fact that θ_W is a monomorphism.) But, by the construction of $\Phi \subseteq P$, we also have $\psi_W(x) = p^-_W$, so the result follows. □

Proof of Lemma 6.2. First, to see that ϕ is injective, suppose that $a, b \in \Phi$ with $\phi a = \phi b$. By Lemma 6.3, we get $\psi_W a = \psi_W b$ for all $W \in \mathcal{W}$, so $a = b$ (since $(\psi_W)_W$ is the full set of coordinate projections to P).

To see that ϕ is a homomorphism, suppose that $a, b, c \in \Phi$ with $c \in [a, b]_\Phi$. Suppose, for contradiction, that $\phi c \notin [\phi a, \phi b]_M$. Let $W_0 \in \mathcal{W}(M)$ be a wall of M separating ϕc from $\mu_M(\phi a, \phi b, \phi c)$. Without loss of generality, we have

$$\phi a, \phi b \in H^-(W_0) \quad \text{and} \quad \phi c \in H^+(W_0).$$

Note that

$$Q(h(a)) \cap H^-(W_0) \neq \emptyset, \quad \text{and} \quad Q(h(c)) \cap H^+(W_0) \neq \emptyset.$$

In particular, $H^\pm(W) \neq \emptyset$, where $H^\pm(W) = H^\pm(W_0) \cap \Pi$, and so we deduce that $W = \{H^-(W), H^+(W)\}$ is a wall of Π. Let $\Phi_W = [p^-_W, p^+_W]$, with $p^\pm_W \in H^\pm(W)$, and with $\{p^-_W, p^+_W\}$ an edge of Π. (We have chosen some edge, $\{p^-_W, p^+_W\}$, of Π crossing W, so we have $p^-_W, p^+_W \in \Pi \cap \Phi_W \subseteq M$. The exact choice does not matter to this discussion.) See Figure 3.

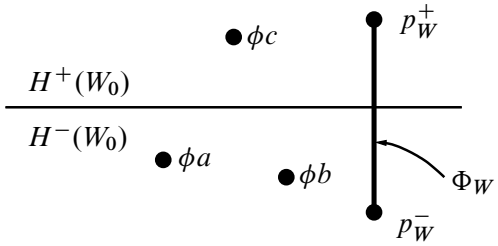

Figure 3. The wall W_0

Now,

$$\psi_W a = \theta_W \phi a = \mu(p_W^-, p_W^+, \phi a) \in [p_W^-, \phi a] \subseteq H^-(W_0),$$

so

$$\psi_W a \in H^-(W_0).$$

Similarly,

$$\psi_W b \in H^-(W_0) \quad \text{and} \quad \psi_W c \in H^+(W_0).$$

But ψ_W is just coordinate projection in P to an interval, hence a homomorphism, so

$$\psi_W c \in \psi_W([a, b]) \subseteq [\psi_W a, \psi_W b] \subseteq H^-(W_0),$$

giving a contradiction. □

7. Cubes in median metric spaces

In this section, we will describe how to define a metric associated to a median metric space of finite rank. In Section 8, we will show this to be CAT(0).

Let (M, ρ) be a median metric space. Suppose that $Q \subseteq M$ is an m-cube with edge-lengths $t_1, \ldots, t_m$ (that is, $t_i = \rho(a, e_i)$, where $\{a, e_1\} \ldots \{a, e_m\}$ are the edges containing a). Set

$$\omega(Q) = \sqrt{\sum_i t_i^2}.$$

Note that, if a, b is a diagonal of Q, then $\rho(a, b) = \sum_i t_i$, and so

$$\rho(a, b)/\sqrt{m} \leq \omega(Q) \leq \rho(a, b).$$

Suppose that $Q' \subseteq Q$ is a cube in Q containing a, b (that is, a subalgebra isomorphic to a cube). Its edge-lengths have the form $\sum_{j \in I_i} t_j$, where $I = \bigsqcup_i I_i$ is a partition of $I = \{1, \dots, m\}$. It follows that $\omega(Q') \geq \omega(Q)$.

Suppose that rank $M = n < \infty$. Given $a, b \in M$, write $Q(a, b) \subseteq M$ for the unique maximal cube in M with diagonal a, b. We write

$$\omega(a, b) = \omega(Q(a, b)).$$

Thus,

$$\rho(a, b)/\sqrt{n} \leq \omega(a, b) \leq \rho(a, b).$$

Moreover, if $Q' \subseteq M$ is any other cube with diagonal a, b, then Lemma 5.2 tells us that Q' is a subcube of Q, and so

$$\omega(Q') \geq \omega(Q).$$

Suppose $a, b, c, d \in M$, and Q is a cube with diagonal a, b. Then $\pi(Q)$ is a cube of diagonal πa, πb, where π is the projection map $[x \mapsto \mu(c, d, x)]$. Since π is 1-lipschitz, the edge-lengths of πQ are at most those of Q and so

$$\omega(\pi Q) \leq \omega(Q).$$

By taking Q to be maximal, and applying the above, we see that

$$\omega(\pi a, \pi b) \leq \omega(a, b).$$

We now define a metric $\sigma = \sigma_\rho$ on M as follows. If $\underline{a} = a_0, a_1, \dots, a_p$ is a sequence in M, we write

$$\omega(\underline{a}) = \sum_{i=1}^{p} \omega(a_{i-1}, a_i).$$

See Figure 4.

Given $a, b \in M$, let $\sigma(a, b) = \inf\{\omega(\underline{a})\}$ as $\underline{a}$ ranges over all sequences (of any finite length) with $a_0 = a$ and $a_p = b$. Clearly σ is (at least) a pseudometric. In fact, we have:

Lemma 7.1. *$\sigma(a, b) = \inf\{\omega(\underline{a})\}$ as $\underline{a}$ ranges over all monotone sequences from a to b.*

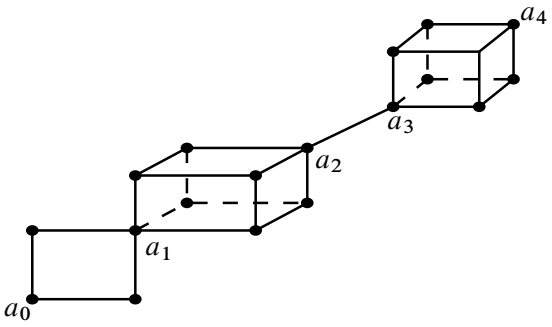

Figure 4. Defining the metric σ_ρ

Proof. Let $\underline{c} = c_0, \ldots, c_p$ be any sequence with $c_0 = a$ and $c_p = b$. First, project $\underline{c}$ to $[a, b]$ to give another sequence, $\underline{d}$, from a to b; that is, set

$$d_i = \mu(a, b, c_i).$$

From the above observation, we have $\omega(\underline{d}) \leq \omega(\underline{c})$. Now project to $d_1, \ldots, d_p = b$ to give a sequence $e_1, \ldots, e_p$ from d_1 to b. Project $e_2, \ldots, e_p$ to $[e_2, b]$ to give $f_2, \ldots, f_p$, etc. After p steps, we arrive at a monotone sequence,

$$\underline{a} = a, \quad d_1, \quad e_2, \quad f_3, \quad \ldots, \quad b,$$

with $\omega(\underline{a}) \leq \omega(\underline{c})$. □

If $\underline{a}$ is monotone, then

$$\rho(a, b) = \sum_{i=1}^{p} \rho(a_{i-1}, a_i),$$

and so it follows that

$$\rho(a, b)/\sqrt{n} \leq \sigma(a, b) \leq \rho(a, b).$$

In particular, it now follows that σ is a metric on M.

To motivate the following, note that if c is a point in a rectilinear euclidean parallelepiped with diagonal (opposite corners) a, b, then in the euclidean metric, we have

$$\sigma(a, c)^2 + \sigma(c, b)^2 \leq \sigma(a, b)^2$$

(since the euclidean angle acb is at least $\pi/2$). We can view Lemma 7.3 below as a generalisation of this fact. First, we note:

Lemma 7.2. *Let $Q \subseteq M$ be a cube with diagonal a, b, and let $c \in [a, b]$. Let Q^-, Q^+ be, respectively, the projections of Q to $[a, c]$ and $[c, b]$. Then*

$$\omega(Q^-)^2 + \omega(Q^+)^2 \leq \omega(Q)^2.$$

Proof. Let t_i be the ith side-length of Q. On projecting c to the interval of M corresponding to the ith side of Q, we can write

$$t_i = t_i^- + t_i^+,$$

where $t_i^\pm$ is either 0 or a side-length of $Q_i^\pm$. From this we get that

$$\omega(Q^-)^2 + \omega(Q^+)^2 = \sum_i (t_i^-)^2 + \sum_i (t_i^+)^2 \leq \sum_i t_i^2 = \omega(Q)^2. \qquad \square$$

Lemma 7.3. *Suppose $a, b \in M$ and $c \in [a, b]$, then*

$$\sigma_\rho(a, c)^2 + \sigma_\rho(b, c)^2 \leq \sigma_\rho(a, b)^2.$$

Proof. Let $\delta > 0$, and let

$$a = d_0, d_1, \ldots, d_p = b$$

be a monotone sequence with

$$\sum_i \omega(d_{i-1}, d_i) \leq \sigma_\rho(a, b) + \delta.$$

Let

$$Q_i = Q(d_{i-1}, d_i)$$

be the maximal cube with diagonal d_{i-1}, d_i. Let

$$\omega_i = \omega(Q_i).$$

Let

$$d_i^- = \mu(a, c, d_i) \quad \text{and} \quad d_i^+ = \mu(c, b, d_i).$$

Thus

$$a = d_0^-, d_1^-, \ldots, d_p^- = c \quad \text{and} \quad c = d_0^+, d_1^+, \ldots, d_p^+ = b$$

are monotone sequences. Let Q_i^- and Q_i^+ be the projections of Q_i to $[a, c]$ and $[c, b]$ respectively. Since projection is a homomorphism, these are also cubes with diagonals d_{i-1}^-, d_i^- and d_{i-1}^+, d_i^+ respectively (Figure 5).

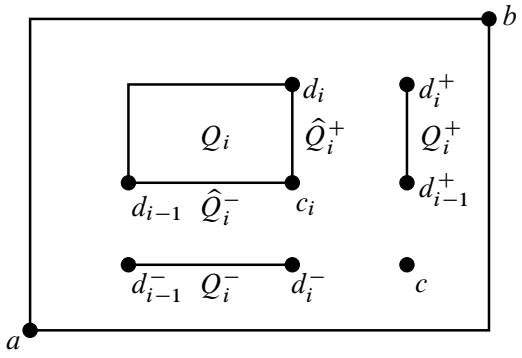

Figure 5. Cubes in the interval $[a, b]$

Let

$$\omega_i^{\pm} = \omega(Q_i^{\pm}).$$

Then

$$\sigma_\rho(a, c) \leq \sum_i \omega_i^- \quad \text{and} \quad \sigma_\rho(c, b) \leq \sum_i \omega_i^+.$$

Let

$$c_i = \mu(c, d_{i-1}, d_i).$$

Now d_{i-1}, c_i is parallel to d_{i-1}^-, d_i^-. (In the notation of Section 5 applied to the interval $[a, b]$, note that

$$d_{i-1} \leq d_i,$$

$$c_i = (c \vee d_{i-1}) \wedge d_i = (c \wedge d_i) \vee d_{i-1},$$

$$d_{i-1}^- = c \wedge d_{i-1}$$

and

$$d_i^- = c \wedge d_i.$$

Now

$$d_i^- \wedge d_{i-1} = (c \wedge d_i) \wedge d_{i-1} = c \wedge d_{i-1} = d_{i-1}^- \quad \text{and} \quad d_i^- \vee d_{i-1} = c_i,$$

and so also,

$$d_{i-1}^- \leq d_i \leq c_i \quad \text{and} \quad d_{i-1}^- \leq d_{i-1} \leq c_i.$$

It follows that $d_{i-1}, c_i, d_i^-, d_{i-1}^-$ is a square.) We see that d_{i-1}, c_i is the diagonal of a cube, $\widehat{Q}_i^-$ (not necessarily maximal) parallel to Q_i^-. Similarly, c_i, d_i is the diagonal of a cube, $\widehat{Q}_i^+$, parallel to Q_i^+. Note that $\widehat{Q}_i^-$ and $\widehat{Q}_i^+$ are the projections of Q_i respectively to $[d_{i-1}, c_i]$ and to $[c_i, d_i]$. By Lemma 7.2, we therefore see that,

$$(\omega_i^-)^2 + (\omega_i^+)^2 \leq \omega_i^2.$$

Now

$$\Big(\sum_i \omega_i^-\Big)^2 + \Big(\sum_i \omega_i^+\Big)^2 \leq \Big(\sum_i \omega_i\Big)^2.$$

(Note that $(\omega_i^-\omega_j^- + \omega_i^+\omega_j^+)^2 \leq ((\omega_i^-)^2 + (\omega_i^+)^2)((\omega_j^-)^2 + (\omega_j^+)^2) \leq \omega_i^2\omega_j^2$, so $\omega_i^-\omega_j^- + \omega_i^+\omega_j^+ \leq \omega_i\omega_j$, and the inequality follows on expanding both sides.) We get that $\sigma_\rho(a,c)^2 + \sigma_\rho(b,c)^2 \leq (\sigma_\rho(a,b) + \delta)^2$. Since this holds for all $\delta > 0$, the statement follows. □

In particular, we see that if $a, b \in M$ and $c \in [a,b] \setminus \{b\}$, then

$$\sigma_\rho(a,c) < \sigma_\rho(a,b).$$

Recall that, by definition of the median structure, $I_\rho(a,b) = [a,b]$.

Corollary 7.4. *For all $a, b \in M$, $I_{\sigma_\rho}(a,b) \subseteq I_\rho(a,b)$.*

Proof. Suppose $x \in M \setminus I_\rho(a,b)$. Let

$$c = \mu(a,b,x),$$

so $x \neq c$. Now

$$c \in [a,x] \setminus \{x\},$$

so

$$\sigma_\rho(a,c) < \sigma_\rho(a,x).$$

Similarly,

$$\sigma_\rho(b,c) < \sigma_\rho(b,x),$$

so

$$\sigma_\rho(a,b) \leq \sigma_\rho(a,c) + \sigma_\rho(c,b) < \sigma_\rho(a,x) + \sigma_\rho(x,b),$$

so

$$x \notin I_{\sigma_\rho}(a,b).$$

□

As an example of the above construction, if ρ is the l^1 metric on $\mathbb{R}^n$, then cubes are the vertex sets of rectilinear parallelepipeds, and we see that σ_ρ is the euclidean metric on $\mathbb{R}^n$. More generally, if (Υ, ρ) is a CAT(0) cube complex, and we are given $\lambda : \mathcal{W}(\Upsilon) \longrightarrow (0, \infty)$, we get a median metric on Υ as discussed in Section 2. In this case, σ_ρ coincides with the corresponding euclidean structure, σ, on Υ. (First note that one easily sees that σ_ρ and σ agree on each cell of Υ. From the fact that (Υ, σ) is geodesic, we can deduce that $\sigma_\rho \leq \sigma$, and directly from the definition of σ_ρ, we see that $\sigma \leq \sigma_\rho$.)

Now, as discussed in Section 2, any finite median metric space arises as the vertex set of such a complex. We therefore get:

Lemma 7.5. *Suppose that (Π, ρ) is a finite median metric space. Then we can canonically identify Π as the vertex set of a* CAT(0) *cube complex, $\Upsilon(\Pi)$, admitting a* CAT(0) *metric σ_Π. Moreover, $\Upsilon(\Pi)$ also canonically admits a median metric ρ_Π, which agrees with ρ on Π, and is such that $\sigma_\Pi = \sigma_{\rho_\Pi}$. In fact, any cell, P, of $\Upsilon(\Pi)$, of any dimension n, can be embedded into $\mathbb{R}^n$ as a rectilinear parallelepiped in such a way that ρ_Π and σ_Π on P respectively agree with the usual l^1 metric and the euclidean metric induced from $\mathbb{R}^n$.*

Now suppose that (M, ρ) is a median metric space. Given any finite subalgebra $\Pi \subseteq M$, the metric restricted to Π is an intrinsic median metric, and so we can construct $\Upsilon(\Pi)$, as in Lemma 7.5.

For the rest of this section, we will be assuming that (M, ρ) is a geodesic space. Note that Corollary 7.4 implies that any σ_ρ-geodesic in M can be reparameterised to give a ρ-geodesic.

Lemma 7.6. *If (M, ρ) is a geodesic median metric space, and $\Pi \subseteq M$ a finite subalgebra, then there is a median monomorphism*

$$f : \Upsilon(\Pi) \longrightarrow M$$

extending the inclusion of Π into M. Moreover, f is an isometric embedding as a map $(\Upsilon(\Pi), \rho_\Pi) \longrightarrow (M, \rho)$. Also, if M has finite rank, then f is 1-lipschitz as a map $(\Upsilon(\Pi), \rho_\Pi) \longrightarrow (M, \sigma_\rho)$.

Here, in general, f is not canonically defined. Note that, in the last clause, we need that M is finite rank just so that σ_ρ is defined.

Proof. We begin by describing how to construct f. If $Q \subseteq \Pi \subseteq M$ is a face of Π, we can choose a diagonal, a, b, and let $e_1, \ldots, e_n$ be the adjacent vertices to a in Q. By Lemma 5.1, $\text{hull}_M(Q)$ is isomorphic as a median algebra to the direct product $\prod_i [a, e_i]$. Now the cell, P_Q, of $\Upsilon(\Pi)$ is isometric to $\prod_i [0, r_i] \subseteq \mathbb{R}^n$ in the l^1 metric, where $r_i = \rho(a, e_i)$. Since M is a geodesic space, we can find a ρ-geodesic

$$\gamma_i : [0, r_i] \longrightarrow [a, e_i],$$

with

$$\gamma_i(0) = 0 \quad \text{and} \quad \gamma_i(r_i) = e_i.$$

Combining these, we get a distance-preserving map

$$f_Q : P_Q \longrightarrow \text{hull}_M(Q).$$

In fact, we can assume that all the paths of the form γ_i crossing any given wall of Π are parallel, and so the maps f_Q fit together to give a map

$$f : \Upsilon(\Pi) \longrightarrow M.$$

We need to check that this is a monomorphism. To do this, we start again with a more formal description of f in terms of the constructions of Sections 5 and 6.

By Lemma 6.2, there is an isomorphism,

$$\phi : \Phi \longrightarrow M(\Pi),$$

where Φ is the subalgebra of the product $P = \prod_{W \in \mathcal{W}} \Phi_W$, as described in Section 5, and where $M(\Pi)$ is the subalgebra of M described in Section 5. We can write $\Phi_W = [p_W^-, p_W^+]$, and set

$$\delta_W : [0, \lambda(W)] \longrightarrow \Phi_W$$

to be a ρ-geodesic from p_W^- to p_W^+. Doing this on every coordinate, we get a product map

$$\prod_{W \in \mathcal{W}} [0, \lambda(W)] \longrightarrow P.$$

Restricting to $\Upsilon(\Pi)$, we get a median homomorphism,

$$\delta : \Upsilon(\Pi) \longrightarrow \Phi.$$

Let

$$f = \phi \circ \delta \colon \Upsilon(\Pi) \longrightarrow M(\Pi).$$

From this description, it is clear that f is a median monomorphism. This in turn implies that f is an isometric embedding from $(\Upsilon(\Pi), \rho_\Pi)$ to (M, ρ).

Finally, to see that

$$f : (\Upsilon(\Pi), \sigma_\Pi) \longrightarrow (M, \sigma_\rho)$$

is 1-lipschitz, let $a, b \in \Upsilon(\Pi)$ and let

$$a = a_0, \quad a_1, \quad \ldots, \quad a_p = b$$

now be a sequence of points, in order, along the geodesic from a to b in $(\Upsilon(\Pi), \sigma_\Pi)$, and with a_{i-1}, a_i lying in some cell of $\Upsilon(\Pi)$ for all i. (This is also a monotone sequence.) Let Q_i be the maximal cube in $\Upsilon(\Pi)$ with diagonal a_{i-1}, a_i. Now $f(Q_i)$ is a cube in M with diagonal $f(a_{i-1})$, $f(a_i)$, so by the earlier observation, and the fact that f is a median monomorphism, we have

$$\omega(f(a_{i-1}), f(a_i)) \leq \omega(f(Q_i)) = \omega(Q_i) = \sigma_\Pi(a, b).$$

Thus,

$$\begin{aligned} \sigma_\rho(f(a), f(b)) &\leq \sum_i \omega(f(a_{i-1}), f(a_i)) \\ &\leq \sum_i \sigma_\Pi(a_{i-1}, a_i) \\ &= \sigma_\Pi(a, b). \end{aligned}$$

□

Lemma 7.7. *Suppose that M is a geodesic median metric space of finite rank. Suppose that $A \subseteq M$ is any finite subset, and $\delta > 0$. Then there is a finite subalgebra, $\Pi \subseteq M$ with $A \subseteq \Pi$, such that if $(\Upsilon(\Pi), \sigma_\Pi)$ is the* CAT(0) *cube complex with vertex set Π as given by Lemma 7.5, then*

$$\sigma_\Pi(a, b) \leq \sigma(a, b) + \delta$$

for all $a, b \in A$.

Proof. For each pair, $a, b \in A$, let

$$a = a_0, \quad a_1, \quad \ldots, \quad a_p = b$$

be a monotone sequence with

$$\sum_{i=1}^{p} \omega(a_{i-1}, a_i) \leq \sigma(a, b) + \delta.$$

Let $Q(a_{i-1}, a_i)$ be the maximal cube in M with diagonal a_{i-1}, a_i. Thus,

$$\omega(a_{i-1}, a_i) = \omega(Q(a_{i-1}, a_i)).$$

Let

$$B(a, b) = \bigcup_{i=1}^{p} Q(a_{i-1}, a_i),$$

and let

$$B = \bigcup_{a,b \in A} B(a, b).$$

Let $\Pi \subseteq M$ be a finite subalgebra containing B.

Now suppose that $a, b \in A$. Let $a = a_0, a_1, \ldots, a_p = b$ be as above. Now $Q(a_{i-1}, a_i) \subseteq \Pi$, and ρ_Π agrees with ρ on Π. Since $\sigma_\Pi = \sigma_{\rho_\Pi}$, we have

$$\sigma_\Pi(a_{i-1}, a_i) \leq \omega(Q(a_{i-1}, a_i)) = \omega(a_{i-1}, a_i),$$

and so

$$\sigma_\Pi(a, b) \leq \sum_{i=1}^{p} \sigma_\Pi(a_{i-1}, a_i) \leq \sum_{i=1}^{p} \omega(a_{i-1}, a_i) \leq \sigma(a, b) + \delta. \qquad \square$$

Putting Lemmas 7.5, 7.6 and 7.7 together, we have shown:

Lemma 7.8. *Suppose that (M, ρ) is a geodesic median metric space of finite rank. Given any finite $A \subseteq M$, and any $\delta > 0$, there is a compact* CAT(0) *space $(\Upsilon, \sigma_\Upsilon)$ with $A \subseteq \Upsilon$, and a 1-lipschitz map,*

$$f : (\Upsilon, \sigma_\Upsilon) \longrightarrow (M, \sigma_\rho)$$

extending the inclusion of A into M, such that for all $a, b \in A$, we have

$$\sigma_\Upsilon(a, b) \leq \sigma_\rho(a, b) + \delta.$$

8. CAT(0) spaces

Let (M, σ) be a geodesic metric space. Suppose $\delta \geq 0$. By a *δ-kite* in M we mean an ordered quadruple of points,

$$K = \{a, b, c, d\}$$

with

$$\rho(a, d) + \rho(d, b) \leq \rho(a, b) + \delta.$$

Given $\epsilon \geq 0$, an *ϵ-comparison* of K is a map

$$\xi \colon K \longrightarrow \mathbb{R}^2,$$

into the plane with euclidean metric σ_0 such that

$$|\sigma_0(\xi(x), \xi(y)) - \sigma(x, y)| \leq \epsilon,$$

whenever $x, y \in K$ and $\{x, y\} \neq \{c, d\}$. Note that $\xi(K)$ is then a $(\delta + 3\epsilon)$-kite in $\mathbb{R}^2$.

If K is a 0-kite, then we can always find a 0-comparison, ξ, of K, and $\xi(K)$ is then a 0-kite. Moreover, the image is uniquely determined up to isometry of $\mathbb{R}^2$. In this case, we can speak of "the" 0-comparison to $\mathbb{R}^2$.

Since we do not know a-priori that our space is geodesic, we begin with the following:

Definition. A metric space, (M, σ), is *weakly* CAT(0) if given any 0-kite, $a, b, c, d \in M$, and if ξ is the 0-comparison of a, b, c, d in $\mathbb{R}^2$, then

$$\sigma(c, d) \leq \sigma_0(\xi(c), \xi(d)).$$

The following is now the standard definition of a CAT(0) space:

Definition. A CAT(0) space is a weakly CAT(0) geodesic metric space.

The following assertion about comparisons is a simple exercise in euclidean geometry:

Lemma 8.1. *Given $\eta, r > 0$, there is some $\delta \geq 0$ with the property that if (M, σ) is any weakly* CAT(0) *space, $K = \{a, b, c, d\} \subseteq M$ is any δ-kite in M of diameter at most r, and $\xi \colon K \longrightarrow \mathbb{R}^2$ is any δ-comparison of K, then*

$$\sigma(c, d) \leq \sigma_0(\xi(c), \xi(d)) + \eta.$$

We can now prove the following:

Lemma 8.2. *If (M, ρ) is a median metric space of finite rank, then (M, σ_ρ) is weakly* CAT(0).

Proof. Suppose that $K = \{a, b, c, d\} \subseteq M$ is a 0-kite (in the metric σ_ρ). Let

$$\xi\colon K \longrightarrow \mathbb{R}^2$$

be a 0-comparison. Let $r = \operatorname{diam} K$. Given any $\eta > 0$, let δ be as given by Lemma 8.1. Let $f : (\Upsilon, \sigma_\Upsilon) \longrightarrow (M, \sigma_\rho)$ be as given by Lemma 7.8 with $A = K$ and δ as given. In particular, for each $x, y \in K$, we have

$$\sigma_\rho(x, y) \leq \sigma_\Upsilon(x, y) + \delta.$$

Therefore K is a δ-kite in $(\Upsilon, \sigma_\Upsilon)$ and ξ is a δ-comparison with respect to the metric σ_Υ. Since $(\Upsilon, \sigma_\Upsilon)$ is CAT(0), it follows from Lemma 8.1 that

$$\sigma_\Upsilon(c, d) \leq \sigma_0(\xi(c), \xi(d)) + \eta.$$

It follows that

$$\sigma_\rho(c, d) \leq \sigma_0(\xi(c), \xi(d)) + \eta.$$

Since this holds for all $\eta > 0$, and ξ is fixed, we have

$$\sigma_\rho(c, d) \leq \sigma_0(\xi(c), \xi(d)).$$

Thus, by definition, (M, σ_ρ) is weakly CAT(0) as claimed. □

If M is complete, then it is enough to assume that it is connected:

Theorem 8.3. *Suppose that (M, ρ) is a median metric space of finite rank, and that (M, σ_ρ) is geodesic. Then (M, σ_ρ) is* CAT(0).

Proof. Note that using Corollary 6.2, any σ_ρ-geodesic can be reparameterised as a ρ-geodesic, and so (M, ρ) is also a geodesic space. Thus, by Lemma 8.2, (M, σ_ρ) is weakly CAT(0), hence CAT(0). □

Lemma 8.4. *If (M, ρ) is a complete connected median metric space of finite rank, then (M, σ_ρ) is geodesic.*

Proof. Note that, by Lemma 4.6, we already know that (M, ρ) is geodesic. Since the metrics ρ and σ_ρ are bi-lipschitz equivalent, (M, σ_ρ) is also complete. Therefore, it is enough to prove the existence of midpoints in (M, σ_ρ).

Let $a, b \in M$. Given $\delta > 0$, we first claim that there is some $c \in M$ with

$$\sigma_\rho(a, c), \sigma_\rho(b, c) \leq \frac{1}{2}(\sigma_\rho(a, b) + \delta).$$

To this end, let $f : (\Upsilon, \sigma_\Upsilon) \longrightarrow (M, \sigma_\rho)$ be the map given by Lemma 7.8, with $A = \{a, b\}$. Let x be a midpoint of a, b in $(\Upsilon, \sigma_\Upsilon)$, and let $c = f(x)$. Now,

$$\sigma_\rho(a, c) \leq \sigma_\Upsilon(a, x) = \frac{1}{2}\sigma_\Upsilon(a, b) \leq \frac{1}{2}(\sigma_\rho(a, b) + \delta).$$

Similarly,

$$\sigma_\rho(b, c) \leq \frac{1}{2}(\sigma_\rho(a, b) + \delta)$$

as claimed.

Suppose d is another such point. We can view $K = \{a, b, c, d\}$ as a δ-kite in M. We have a δ-comparison, $\xi : K \longrightarrow \mathbb{R}^2$, such that $\xi(c) = \xi(d)$ is the midpoint of $\xi(a), \xi(b)$. Now, by Lemma 8.2, (M, σ_ρ) is weakly CAT(0). Therefore, if $\eta > 0$, then by choosing $\delta > 0$ sufficiently small depending on η, r, Lemma 8.1 tells us that $\sigma_\rho(c, d) \leq \eta$.

In this way we obtain a sequence of points $(c_i)_i$ in M with

$$\sigma_\rho(a, c_i) \to \frac{1}{2}\sigma_\rho(a, b)$$

and

$$\sigma_\rho(b, c_i) \to \frac{1}{2}\sigma_\rho(a, b),$$

and with $(c_i)_i$ Cauchy. Thus c_i converges to a midpoint of a, b in (M, σ_ρ) as required. □

Now we already know by Lemma 8.2 that, under these assumptions, (M, σ_ρ) is weakly CAT(0), hence it is CAT(0). This proves Theorem 1.1.

Finally, it is worth observing that if C is a closed convex subset of (M, ρ), then C is also a complete connected median metric space in the restricted metric, ρ. It is a simple consequence of the construction that C is a totally geodesic subset of (M, σ), and that the metric σ restricted to C is the same as the metric σ_ρ, obtained intrinsically from ρ in C.

References

[1] H.-J. Bandelt and V. Chepoi, Metric graph theory and geometry: a survey. In J. E. Goodman, J. Pach and R. Pollack (eds.), *Surveys on discrete and computational geometry.* (Snowbird, UT, 2006.) Contemporary Mathematics, 453. American Mathematical Society, Providence, R.I., 2008, 49–86. Zbl 1169.05015 Zbl 1134.52001 (collection) MR 2405677 MR 2407791 (collection)

[2] H.-J. Bandelt and J. Hedlikova, Median algebras. *Discrete Math.* **45** (1983), no. 1, 1–30. Zbl 0506.06005 MR 0700848

[3] J. A. Behrstock and C. Druțu, Higher dimensional filling and divergence for mapping class groups. Preprint 2013. arXiv:1305.2994 [math.GT]

[4] J. A. Behrstock, C. Druțu, and M.Sapir, Median structures on asymptotic cones and homomorphisms into mapping class groups. *Proc. Lond. Math. Soc.* (3) **102** (2011), no. 3, 503–554. Zbl 2783135 MR 1270.20044

[5] J. A. Behrstock and Y. N. Minsky, Centroids and the rapid decay property in mapping class groups. *J. Lond. Math. Soc.* (2) **84** (2011), no. 3, 765–784. Zbl 1269.20032 MR 2855801

[6] B. H. Bowditch, Coarse median spaces and groups. *Pacific J. Math.* **261** (2013), no. 1, 53–93. Zbl 1283.20048 MR 3037559

[7] B. H. Bowditch, Embedding median algebras in products of trees. *Geom. Dedicata* **170** (2014), 157–176. Zbl 06299144 MR 3199482

[8] B. H. Bowditch, Large-scale rigidity properties of the mapping class groups. Preprint 2015. http://homepages.warwick.ac.uk/~masgak/papers/teich-rigid.pdf

[9] B. H. Bowditch, *Large-scale rank and rigidity of the Teichmüller metric.* Preprint 2015. http://www.warwick.ac.uk/~masgak/papers/mappingclass.pdf

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

[11] I. Chatterji, C. Druțu, and F. Haglund, Kazhdan and Haagerup properties from the median viewpoint. *Adv. Math.* **225** (2010), no. 2, 882–921. Zbl 1271.20053 MR 2671183

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

[13] M. Gromov, Asymptotic invariants of infinite groups. In G. A. Niblo and M. A. Roller (eds.), *Geometric group theory.* Vol. 2. (Brighton, 1991.) London Mathematical Society Lecture Note Series, 182. Cambridge University Press, Cambridge, 1993, 1–295. Zbl 0841.20039 MR 1253543

[14] J. R. Isbell, Median algebra. *Trans. Amer. Math. Soc.* **260** (1980), no. 2, 319–362. Zbl 0446.06007 MR 0574784

[15] L. Mosher, Mapping class groups are automatic. *Ann. of Math.* (2) **142** (1995), no. 2, 303–384. Zbl 0867.57004 MR 1343324

[16] T. R. Riley, Higher connectedness of asymptotic cones. *Topology* **42** (2003), no. 6, 1289–1352. Zbl 1038.20031 MR 1981358

[17] M. A. Roller, *Poc-sets, median algebras and group actions, an extended study of Dunwoody's construction and Sageev's theorem.* Habilitationschrift, Regensberg, 1998.

[18] M. Sholander, Medians and betweenness. *Proc. Amer. Math. Soc.* **5**, (1954), 801–807. Zbl 0056.26101 MR 0064749

[19] L. Van den Dries and A. J. Wilkie, Gromov's theorem on groups of polynomial growth and elementary logic. *J. Algebra* **89** (1984), no. 2, 349–374. Zbl 0552.20017 MR 0751150

[20] E. R. Verheul, *Multimedians in metric and normed spaces.* CWI Tract, 91. Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam, 1993. Zbl 0790.46008 MR 1244813

Received March 4, 2014

Brian H. Bowditch, Mathematics Institute, University of Warwick,
Coventry CV4 7AL, UK

home page: http://homepages.warwick.ac.uk/~masgak/