

Original citation:

Bowditch, B. H.. (2014) Embedding median algebras in products of trees. *Geometriae Dedicata*, 170 (1). pp. 157-176.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/79498>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"The final publication is available at Springer via <http://dx.doi.org/10.1007/s10711-013-9874-x>."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

EMBEDDING MEDIAN ALGEBRAS IN PRODUCTS OF TREES

BRIAN H. BOWDITCH

ABSTRACT. We show that a metric median algebra satisfying certain conditions admits a bilipschitz embedding into a finite product of \mathbb{R} -trees. This gives rise to a characterisation of closed connected subalgebras of finite products of complete \mathbb{R} -trees up to bilipschitz equivalence. Spaces of this sort arise as asymptotic cones of coarse median spaces. This applies to a large class of finitely generated groups, via their Cayley graphs. We show that such groups satisfy the rapid decay property. We also recover the result of Behrstock, Druţu and Sapir, that the asymptotic cone of the mapping class group embeds in a finite product of \mathbb{R} -trees.

2010 Mathematics Subject Classification : 20F65

Keywords: median algebra, \mathbb{R} -tree, embedding, rapid decay

1. INTRODUCTION

The main aim of this paper is to show that a metric median algebra satisfying certain hypotheses admits a bilipschitz embedding into a finite product of \mathbb{R} -trees. Median algebras of this sort occur as asymptotic cones of coarse median spaces as defined in [Bo1]. The theory also applies to finitely generated groups via their Cayley graphs. Many natural groups have such a structure, notably the mapping class groups of surfaces, see [BehM, Bo1]. Other groups include right-angled Artin groups, groups hyperbolic relative to other such groups [Bo2], and direct products of such groups. Conjecturally it might apply to all CAT(0) groups. For mapping class groups, we recover the bilipschitz embedding theorem of [BehDS]. (This can be thought of as a kind of “asymptotic” version of the coarse embedding theorem given in [BesBF].) The results will be set in a broader context in Section 8.

We note that median structures, from a related but somewhat different perspective, are discussed in [ChaDH]. Also an analogous, though different, result for topological median algebras can be found in [BaV]. (There it is assumed that intervals are compact. We do not take that as an assumption here.) The issue of embedding median algebras in

Date: First Draft: November 2011. Revised: May 2013.

products of trees is also studied in the combinatorial setting of cube complexes, see for example [CheH], and the references therein.

We briefly recall the formal definitions (see [I, BaH, R, Che]). A *median algebra* is a set, Φ , equipped with a ternary operation, $\mu : \Phi^3 \rightarrow \Phi$, such that for all $a, b, c, d, e \in \Phi$, we have:

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

$$(M2): \mu(a, a, b) = a, \text{ and}$$

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

A *subalgebra* is a subset of Φ closed under this operation.

(In practice will not use the formal definition above. It is sufficient to note that any finite subset of Φ lies in a finite subalgebra, and that a finite median algebra can be naturally identified as the vertex set of a CAT(0) complex, as discussed in Section 2. Indeed these two statements combined could serve as an equivalent definition.)

A *homomorphism* is a map between median algebras respecting the ternary operations. Given $a, b \in \Phi$, the *interval* $[a, b]$ is defined by $[a, b] = \{c \in \Phi \mid \mu(a, b, c) = c\}$. A subset, $H \subseteq \Phi$, is *convex* if $[a, b] \subseteq H$ for all $a, b \in H$. Any convex subset is a subalgebra.

A *oriented wall*, W , is an ordered partition, $(H^-(W), H^+(W))$ of Φ into two non-empty convex subsets. A *wall* (or *unoriented wall*) is an unordered such partition, $\{H^-(W), H^+(W)\}$. We write $\mathcal{W}(\Phi)$ for the set of all walls. We say that two walls, W, W' , *cross*, and write $W \pitchfork W'$, if each of the sets $H^-(W) \cap H^-(W')$, $H^-(W) \cap H^+(W')$, $H^+(W) \cap H^-(W')$ and $H^+(W) \cap H^+(W')$ is non-empty.

Definition. We say that Φ had *rank at most* ν if there is no collection of $\nu + 1$ pairwise crossing walls of Φ .

Various equivalent formulations of this are described in [Bo1].

Definition. We say that Φ is ν -*colourable* if there is a map, $\chi : \mathcal{W}(\Phi) \rightarrow \{1, 2, \dots, \nu\}$, such that $\chi(W) \neq \chi(W')$ whenever $W \pitchfork W'$.

Clearly ν -colourable implies rank at most ν . There is a converse for intervals (see Lemma 2.3), though not in general. Also, any subalgebra of a ν -colourable algebra is also ν -colourable. The same statement applies to rank (see [Bo1]).

(We remark that in the context of a CAT(0) cube complex, the notion of a wall is essentially equivalent to that of a ‘‘hyperplane’’ — a totally geodesic codimension-1 subspace which cuts in half each cube that crosses the wall.)

Definition. By a *metric median algebra* we mean a triple (Φ, ρ, μ) , where (Φ, μ) is a median algebra, (Φ, ρ) is a metric space, and $\mu : \Phi^3 \rightarrow \Phi$ is continuous in the induced topology.

For the purposes of this paper, the following will serve as a definition of an \mathbb{R} -tree. It is one of many equivalent definitions.

Definition. An \mathbb{R} -tree is a connected metric space, (T, τ) , such that for all $a, b, c, d \in T$, we have:

$$(T1): \tau(a, b) + \tau(c, d) \leq \max\{\tau(a, c) + \tau(b, d), \tau(a, d) + \tau(b, c)\}.$$

An \mathbb{R} -tree admits a unique median, μ_T , in such a way that for all $a, b \in T$, the median interval $[a, b]$ is the unique topological arc from a to b in T . Moreover, $[a, b]$ is isometric to the real interval $[0, \tau(a, b)]$. In fact, (T, τ, μ_T) is a rank-1 median algebra. Note that it is clear from the definition that the metric completion of an \mathbb{R} -tree is an \mathbb{R} -tree.

Given a collection, $\Phi_1, \Phi_2, \dots, \Phi_\nu$, of ν metric median algebras, the direct product, $\prod_{i=1}^\nu \Phi_i$, with the l_1 -metric is also a metric median algebra (defining the median co-ordinatewise). There is a natural identification of $\mathcal{W}(\prod_i \Phi_i) \equiv \bigsqcup_i \mathcal{W}(\Phi_i)$. If Φ_i is ν_i -colourable, then $\prod_i \Phi_i$ is $\sum_i \nu_i$ -colourable. (A similar statement applies to rank.) In particular, we see that a direct product of ν \mathbb{R} -trees is a ν -colourable metric median algebra.

We can now state the main result. Let (Φ, ρ, μ) be a metric median algebra. We suppose:

(L1): $(\exists \nu \in \mathbb{N})$ such that (Φ, μ) is ν -colourable.

(L2): $(\exists k > 0)$ such that for all $a, b, c, d \in \Phi$, $\rho(\mu(a, b, c), \mu(a, b, d)) \leq k\rho(c, d)$.

(L3): $(\exists l \geq 1)$ such that for all $a, b \in \Phi$ there is an l -lipschitz path, $\alpha : [0, \rho(a, b)] \rightarrow \Phi$ with $\alpha(0) = a$ and $\alpha(\rho(a, b)) = b$.

Note that (L3) is equivalent to saying that two points, a, b , are connected by a rectifiable path of length at most $l\rho(a, b)$. In our main applications we have $l = 1$, that is, Φ will be a geodesic space.

Theorem 1.1. *Suppose that (Φ, ρ, μ) is a metric median algebra satisfying (L1), (L2) and (L3). Then there is an injective lipschitz median homomorphism, $f : \Phi \rightarrow \prod_{i=1}^\nu T_i$, of Φ into a product of ν \mathbb{R} -trees, T_i , with the l_1 -metric. Moreover, the inverse map $f^{-1} : f(\Phi) \rightarrow \Phi$ is also lipschitz, and both lipschitz constants depend only on the constants, ν , k and l , featuring in the hypotheses.*

Note that, conversely, hypotheses (L1) and (L2) follow from the conclusion of Theorem 1.1. Moreover, if $f(\Phi)$ is closed and connected, then it is a geodesic space in the l_1 -metric, and so it also satisfies (L3). (This is discussed further in Section 7.) If Φ is complete (as in the case of principal interest to us), then certainly, $f(\Phi)$ is closed. Moreover, there is no loss in assuming each of the \mathbb{R} -trees to be complete (by taking their completions). Thus, up to bilipschitz equivalence, we have a characterisation of closed connected subalgebras of a finite product of complete \mathbb{R} -trees.

We also note that it is an immediate consequence of Theorem 1.1 that Φ is bilipschitz equivalent to a median metric space in the sense described in [ChaDH]. For further discussion, see Section 7.

Corollary 1.2. *Let Φ be a metric median algebra of rank at most ν and satisfying (L2) and (L3). If $a, b \in \Phi$, then $[a, b]$ admits a bilipschitz embedding as a median algebra into a product of ν real intervals. The bilipschitz constants depend only on ν , k and l .*

Any interval in a metric median algebra is (topologically) closed, and so:

Corollary 1.3. *Let Φ be a complete metric median algebra of finite rank satisfying (L2) and (L3). Then any interval in Φ is compact.*

Another consequence, of course, is that the Hausdorff dimension of any interval is at most the rank. (Although there might be other compact subsets of infinite Hausdorff dimension, even in an \mathbb{R} -tree.)

In fact, our strategy will be to prove a variation of Corollary 1.2 first, and use this information to help us construct our embedding in Theorem 1.1 (see the proof of Lemma 6.2). We state this variation as follows. Here Δ will eventually play the role of an interval in Φ . Both the hypotheses and conclusion are weaker than those of Corollary 1.2.

Let Δ be a metric median algebra satisfying:

(I1): Δ has rank at most $\nu < \infty$.

(I2): There is a continuous path $\beta : [-1, 1] \rightarrow \Delta$ with $\Delta = [\beta(-1), \beta(1)]$.

Proposition 1.4. *Let (Δ, ρ, μ) be a metric median algebra satisfying (I1), and (I2). Then there is a continuous injective median homomorphism, $f : \Delta \rightarrow \Omega$ into the ν -cube $\Omega = [-1, 1]^\nu$.*

Note that the hypotheses make no reference to the metric. In fact, all we really require for this is that Δ be a Hausdorff topological space, and that the median is continuous.

Note also that, without loss of generality, we can suppose that $\phi(\beta(-1))$ and $\phi(\beta(1))$ are opposite corners of Ω . (At least with some mild additional hypotheses, one could deduce that the inverse map is also continuous, but since we do not need this, we will not pursue that matter here.)

With regard to applications, we note that if Φ arises as the asymptotic cone of a coarse median algebra of rank at most ν , in the sense described in [Bo1], then Φ is complete, and automatically satisfies (L1), (L2) and (L3) (with $l = 1$). This holds for a large class of finitely generated groups, as mentioned earlier. This ties in with earlier work, as we describe in Section 8. One consequence of the embedding theorem is the rapid decay of a coarse median group, as we discuss in Section 9.

It is natural to ask the general question of when a median algebra admits a ν -colouring. For example, in the context of CAT(0) cube complexes it was asked independently by Chepoi and Sageev whether a bound on the cardinality of links implies that ν is finite (or equivalently bounded for finite complexes). In [CheH], this is shown to be true in dimension 2, but false in general in dimension 5. It remains open in dimensions 3 and 4.

2. GENERAL OBSERVATIONS

In this section, we make some general observations, and collect together a few facts we will need later.

To begin, we note the following canonical correspondence between finite median algebras and finite CAT(0) cube complexes. Let Υ be a finite CAT(0) cube complex, with vertex set $V(\Upsilon)$. We write Υ^1 for its 1-skeleton. (A graph arising in this way is sometimes termed a “median graph”, cf. [Che].) Now $V(\Upsilon) = V(\Upsilon^1)$ has a unique structure as a median algebra with the property that, for all $a, b \in V(\Upsilon)$, the median interval $[a, b] \subseteq V(\Upsilon)$ is precisely the set of vertices which lie on some geodesic from a to b in Υ^1 . In other words, the median $\mu(a, b, c)$ of $a, b, c \in V(\Upsilon)$ is the unique vertex which lies on some geodesic between each pair of distinct points of $\{a, b, c\}$. Conversely, if Π is any finite median algebra, one can construct a cube complex Υ with $V(\Upsilon) \equiv \Pi$, so that the median defined by Υ agrees with the original on Π . Moreover, Υ is completely determined (up to isomorphism fixing Π).

We write $E(\Upsilon)$ for the set of edges of Υ . Given a wall $W \in \mathcal{W}(\Pi)$, let $E_W \subseteq E(\Upsilon)$ be the set of edges with one endpoint in $H^-(W)$ and one endpoint in $H^+(W)$. One can check that $\{E_W \mid W \in \mathcal{W}(\Pi)\}$ is a

partition of $E(\Upsilon)$. We note that Π is 1-colourable if and only if it has rank at most 1. In this case, it is a simplicial tree.

Suppose that $\chi : \mathcal{W}(\Pi) \rightarrow \{1, \dots, \nu\}$ is a ν -colouring. Given a colour, i , define a relation \sim_i , on Π by setting $a \sim_i b$, if and only if no wall in $\chi^{-1}(i)$ separates a from b . This is an equivalence relation, and we set $\Pi_i = \Pi / \sim_i$. Since any two points of Π are separated by a wall of Π , the canonical homomorphism $\Pi \rightarrow \prod_i \Pi_i$ is injective. Moreover, it extends uniquely to an embedding of cube complexes, $\Upsilon \rightarrow \prod_i \Upsilon_i$, where Υ_i is the simplicial tree with $V(\Upsilon_i) = \Pi_i$. Thus, a colouring is essentially the same structure as an embedding into a finite product of trees.

Moving on to general median algebras, the following was proven in [Bo1]:

Proposition 2.1. *A median algebra is ν -colourable if and only if every finite subalgebra is.*

In fact, it is the latter statement that we will use here. Given this, Theorem 1.1 gives another, though much more indirect proof of the “if” part of the above statement.

The following observation will be used in Section 4.

Lemma 2.2. *Suppose that M is a median algebra of finite rank, with $|M| \geq 2$. Then there exist distinct points, $a, b \in M$, such that $[a, b]$ has (intrinsic) rank 1.*

For the proof, we recall the following definition from [Bo1]. A *hypercube* in Φ , is a subalgebra, Q , isomorphic to $\{-1, 1\}^\nu$ for some $\nu \in \mathbb{N}$, (the “dimension” of Q). Fixing some such isomorphism, we write a_0 for the point with all co-ordinates -1 , and a_i for the point with i th co-ordinate 1 and all other co-ordinates -1 . Thus, for each $i = 1, \dots, \nu$, $\{a_0, a_i\}$ is a side of the hypercube Q . We recall that the rank of a median algebra (defined here as the maximal cardinality of a set of pairwise crossing walls) is also the maximal dimension of a hypercube in the median algebra [Bo1].

Proof of Lemma 2.2. Let $Q \subseteq M$ be a hypercube of maximal dimension, n . With the above notation, we claim that $[a_0, a_1]$ has rank 1. To see this, for each i we choose any wall $W_i \in \mathcal{W}(M)$ separating a_0 from a_i . Thus, W_i will, in fact, separate all points of Q with i th co-ordinate -1 from all points with i th co-ordinate 1 . In particular, the W_i pairwise cross. If $[a_0, a_1]$ has rank at least 2, then we could find two walls, W_1 and W'_1 , separating a_0 and a_1 , which cross in M . Thus, $W_1, W'_1, W_2, \dots, W_n$ all pairwise cross in M , contradicting the maximality of the dimension n . \square

Suppose that (Φ, ρ, μ) is any metric median algebra. Suppose that $a, b \in \Phi$. Then the interval $[a, b] \subseteq \Phi$ is closed. There is a continuous retraction, $\pi : \Phi \rightarrow [a, b]$ defined by $\pi(x) = \mu(a, b, x)$. If Φ satisfies (L2), then π is k -lipschitz. Note that the properties (L1) and (L3) also pass to $[a, b]$. For (L1), we apply Proposition 2.1, and for (L3) we postcompose with π . (Note that the constants k and ν remain unchanged, but that l may increase to kl .) We can now focus on intrinsic median algebras of this form.

Let Δ be a metric median algebra, with points $e^-, e^+ \in \Delta$ such that $\Delta = [e^-, e^+]$. We can orient any wall, $W \in \mathcal{W}(\Delta)$, so that $e^- \in H^-(W)$ and $e^+ \in H^+(W)$. Given $W, W' \in \mathcal{W}(\Delta)$, we write $W \prec W'$ to mean that $H^-(W) \subseteq H^-(W')$, or equivalently, $H^+(W') \subseteq H^+(W)$. This is a partial order on $\mathcal{W}(\Delta)$. In fact, given any $W, W' \in \mathcal{W}(\Delta)$, exactly one of $W = W'$, $W \prec W'$, $W' \prec W$ or $W \pitchfork W'$ holds. It follows that the rank of Δ is exactly the maximal cardinality of any antichain in $(\mathcal{W}(\Delta), \prec)$. Dilworth's lemma [Di] now tells us that we can partition $\mathcal{W}(\Delta)$ into ν disjoint chains (cf. [BroCGNW]). This defines a ν -colouring of Δ . We deduce:

Lemma 2.3. *If Δ is a median algebra with $\Delta = [e^-, e^+]$ for some $e^-, e^+ \in \Delta$, then Δ is ν -colourable if and only if it has rank at most ν .*

Given $a, b \in \Delta$, we write $a \wedge b = \mu(e^-, a, b)$ and $a \vee b = \mu(e^+, a, b)$. One can verify that (Δ, \wedge, \vee) is a distributive lattice. (That is, for all $a, b, c \in \Delta$, we have $a \wedge b = b \wedge a$, $a \wedge (a \vee b) = a$ and $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$, and similarly, swapping \wedge and \vee .) Moreover, one can recover the median from the lattice structure as

$$\mu(a, b, c) = (a \wedge b) \vee (b \wedge c) \vee (c \wedge a) = (a \vee b) \wedge (b \vee c) \wedge (c \vee a).$$

We write $a \leq b$ to mean $a = a \wedge b$, or equivalently, $b = b \vee a$. This is a partial order on Δ , with $a \wedge b$ and $a \vee b$ respectively the greatest lower bound and least upper bound with respect to this order. Note that $a \leq b$ is equivalent to $a \in [e^-, b]$ or to $b \in [e^+, a]$.

Suppose that $c, d \in \Delta$ and $[c, d]$ has rank 1. Then we see that $\{c, d\} = \{c \wedge d, c \vee d\}$, so either $c \leq d$ or $d \leq c$. Thus, Lemma 2.2, applied intrinsically to any interval $[a, b]$ in Δ , gives:

Lemma 2.4. *If Δ is a finite rank median algebra, with $\Delta = [e^-, e^+]$. Suppose $a, b \in \Delta$ are distinct. Then there exist $c, d \in [a, b]$ with $c < d$, and with $[c, d]$ of rank 1.*

We finish this section by introducing the following notation regarding the real ν -cube, $\Omega = [-1, 1]^\nu$. We write $\omega^\pm = (\pm 1, \dots, \pm 1)$ for the opposite corners of Ω . Thus, Ω is a median algebra with $\Omega = [\omega^-, \omega^+]$.

As before, we write $x \wedge y = \mu(\omega^-, x, y)$ and $x \vee y = \mu(\omega^+, x, y)$. We write x_i for the i th co-ordinate of x . Thus, $(x \wedge y)_i = \min(x_i, y_i)$ and $(x \vee y)_i = \max(x_i, y_i)$. Note that the partial order, \leq , defined as before is equivalent to saying $x_i \leq y_i$ for all i .

3. TOTALLY ORDERED SETS

We make some observations regarding ordered sets used in proving Proposition 1.4. In this section, we consider total orders, and we describe two distinct ways of obtaining real intervals from such orders (see Lemmas 3.1 and 3.2).

Let $(O, <)$ be a totally ordered set. Given $x, y, z \in O$, we say that z lies *between* x and y if either $x \leq z \leq y$ or $y \leq z \leq x$. There is a unique median structure on O such that for all $x, y \in O$, the median interval $[x, y]$ is exactly the set of points between x and y . (With this convention, we always have $[x, y] = [y, x]$.) We write $(x, y) = [x, y] \setminus \{x, y\}$ and $[x, y) = [x, y] \setminus \{y\}$. We say that an interval is *non-trivial* if it contains at least two elements. As before, we say that a subset, $H \subseteq O$, is *convex* if $[x, y] \subseteq H$ for all $x, y \in H$.

A subset of O is *dense* if it meets every non-trivial interval in O . By a *rational* subset of O we mean a subset which is order isomorphic to the rational numbers. To simplify the exposition, we can assume that O has a minimum, e^- , and a maximum e^+ . (Such points can be adjoined, if necessary).

Here is one result that allows us to construct real intervals.

Lemma 3.1. *Let $(O, <)$ be a totally ordered set with a countable dense subset. Then there is a non-decreasing map, $\theta : O \rightarrow I$, where I is either $\{0\}$ or $[-1, 1]$, such that $\theta(O)$ is dense in I , and such that for all $x, y \in O$, we get $\theta(x) \neq \theta(y)$ if and only if x and y are separated by a rational subset of O .*

In this context, a wall $W = \{H^-(W), H^+(W)\}$, with $e^\pm \in H^\pm(W)$, is generally termed a *cut*. Note that $H^-(W)$ and $H^+(W)$ are respectively initial and final segments of O . A cut, W , is called a *jump* if $H^-(W)$ has a maximum and $H^+(W)$ has a minimum. A cut, W , is called a *gap* if $H^-(W)$ has no maximum and $H^+(W)$ has no minimum.

A standard completion process allows us to get rid of gaps. Note that $O \cup \mathcal{W}(O)$ admits a natural total order extending $<$ by setting $x < W$ if $x \in H^-(W)$, $W < x$ if $x \in H^+(W)$ and $W < W'$ if $W \prec W'$ (i.e. $H^-(W) \subseteq H^-(W')$). One verifies that $O \cup \mathcal{W}(O)$ has no gaps and that O is dense in $O \cup \mathcal{W}(O)$.

Here is a process for getting rid of jumps. Given $x, y \in O$, we say that a subset *lies between* x and y if it is contained in $[x, y]$. We define

a relation, \sim , on O , by writing $x \not\sim y$ if there is a rational subset between x and y . We check that \sim is an equivalence relation, and that all equivalence classes are convex. The quotient O/\sim is therefore equipped with a total order, and one verifies that it has no jumps.

If one performs the above constructions in turn (in either order) we arrive at a total order \hat{O} with a natural map $O \rightarrow \hat{O}$. Moreover, one checks that \hat{O} has no gaps or jumps. If O has a countable dense subset, this maps to a countable dense subset of \hat{O} ; and it follows from the well known characterisation of the order type of the reals (see for example [HY]) that \hat{O} is order isomorphic to a real interval. Lemma 3.1 now follows.

Here is another way of arriving at a real interval.

Lemma 3.2. *Suppose that $O = [e^-, e^+]$ is a total order with $e^- \neq e^+$. Suppose that O admits a hausdorff topology such that $[x, y]$ is closed for all $x, y \in O$. Suppose that there is a continuous map, $\beta : [-1, 1] \rightarrow O$, with $\beta(-1) = e^-$ and $\beta(1) = e^+$. Then there is an order-preserving homeomorphism from $[-1, 1]$ to O (necessarily taking -1 to e^- and 1 to e^+).*

Proof. If $x \in (e^-, e^+)$, then $[e^-, x]$ and $[e^+, x]$ are closed, and $[e^-, x] \cap [e^+, x] = \{x\}$ and $[e^-, x] \cup [e^+, x] = O$. In other words, x is a cut point separating e^+ from e^- . It follows that $x \in \beta([-1, 1])$, and so β is surjective. As the continuous image of a real interval, O is a metrisable continuum (in fact, locally connected). Any non-trivial metrisable continuum with at most two non-cut points is homeomorphic to a closed real interval [HY]. We therefore have a homeomorphism, $\gamma : [-1, 1] \rightarrow O$, with $\gamma(-1) = e^-$ and $\gamma(1) = e^+$. Note that, if $x, y \in [e^-, e^+)$, then $x \leq y$ if and only if y separates x from e^+ . It follows that γ is order-preserving. \square

Suppose now that $\Delta = [e^-, e^+]$ is a median algebra. In Section 2, we defined the order, $<$, on Δ . If Δ has rank 1, then this is a total order, and we are in the set-up of this section. In the following lemma, we say “metric”, though we only really need “hausdorff topological”. This hypothesis ensures that intervals are closed.

Lemma 3.3. *Suppose that $\Delta = [e^-, e^+]$ is a metric median algebra of rank 1, and suppose there is a continuous map, $\beta : [-1, 1] \rightarrow \Delta$ with $\beta(-1) = e^-$ and $\beta(1) = e^+$. Then there is a homeomorphism, $\gamma : [-1, 1] \rightarrow \Delta$, which is a median isomorphism.*

Proof. We apply Lemma 3.2 to the order described on Δ , and note that the median is determined by this order. \square

We will also need the following elementary observation:

Lemma 3.4. *Suppose that $\chi : [-1, 1] \rightarrow \{1, \dots, \nu\}$ is any function into a finite set. Then there exist $x < y$ in $[-1, 1]$ and $i \in \{1, \dots, \nu\}$ such that $[x, y] \cap \chi^{-1}(i)$ is dense in $[x, y]$.*

4. PROOF OF PROPOSITION 1.4

Let (Δ, ρ, μ) be a metric median algebra satisfying (I1) and (I2). In (I2), let $e^\pm = \beta(\pm 1)$, so that $\Delta = [e^-, e^+]$. In Section 2, we noted that Δ has the structure of a distributive lattice, (Δ, \wedge, \vee) , with induced partial order $<$.

First, note that any interval in Δ will satisfy the same hypotheses. In particular, for (I2) we note:

Lemma 4.1. *If $a, b \in \Delta$, then there is a continuous path, $\gamma : [-1, 1] \rightarrow [a, b]$ with $\gamma(-1) = a$ and $\gamma(1) = b$.*

Proof. Note that the paths $[t \mapsto a \wedge \beta(t)]$ and $[t \mapsto b \wedge \beta(t)]$ respectively connect e^- to a and to e^- to b . Concatenating them, we get a path, δ , from a to b in Δ . We now set $\gamma(t) = \mu(a, b, \delta(t))$. \square

Let \mathcal{H} be the set of closed convex subsets, H , of Δ such that $e^- \in H$, $e^+ \notin H$, and such that $\Delta \setminus H$ is convex. Thus, $W_H = \{H, \Delta \setminus H\}$ is a wall of Δ . We already have a partial order, \prec , on \mathcal{H} , as defined in Section 3. We define another partial order, \ll , on \mathcal{H} as follows. If $H, H' \in \mathcal{H}$, we write $H \ll H'$ to mean that $H \subseteq \text{int } H'$. Clearly $H \ll H'$ implies $H \prec H'$. Also, if $a \in \Delta$, and $H \in \mathcal{H}$, we write $a \ll H$ to mean $a \in \text{int}(H)$ and $H \ll a$ to mean that $a \notin H$. Given $a, b \in \Delta$, we will write $a \ll H \ll b$ to mean $a \ll H$ and $H \ll b$. (Note, however, that this does not necessarily imply $a \leq b$.)

Given $a \in \Delta$, let $\mathcal{H}(a) = \{H \in \mathcal{H} \mid a \notin H\}$. Note that if $H \in \mathcal{H}(a)$ and $H' \preceq H$, then $H' \in \mathcal{H}(a)$. Given $a, b \in \Delta$, write $\mathcal{H}(a, b)$ for the symmetric difference, $\mathcal{H}(a, b) = \mathcal{H}(a) \Delta \mathcal{H}(b)$. Note that if $c, d \in [a, b]$, then $\mathcal{H}(c, d) \subseteq \mathcal{H}(a, b)$.

Suppose that $c < d$ in Δ and that $[c, d]$ has rank 1. By Lemmas 4.1 and 3.3, $[c, d]$ is order isomorphic to the real interval, $[-1, 1]$. Given any $x \in [c, d)$ let $H(x) = \{z \in \Delta \mid \mu(c, d, z) \in [c, x]\}$. We check easily that $H(x) \in \mathcal{H}(c, d)$. (In fact, every element of $\mathcal{H}(c, d)$ has this form.) Note that, if $x, y \in [c, d)$, then $x < y$ if and only if $H(x) \ll H(y)$.

Now, given any distinct, $a, b \in \Delta$, Lemma 2.4 gives us $c, d \in [a, b]$ with $c < d$ and with $[c, d]$ of rank 1. Putting the above facts together, we conclude:

Lemma 4.2. *If $a, b \in \Delta$ are distinct, then $(\mathcal{H}(a, b), \ll)$ contains a subset order isomorphic to the closed real interval $[-1, 1]$.*

We now introduce colouring. Note that, by Lemma 2.3, Δ admits a ν -colouring, $\chi : \mathcal{W}(\Delta) \rightarrow \{1, \dots, \nu\}$. Given $H \in \mathcal{H}$, write $\chi(H) = \chi(W_H)$. Given $i \in \{1, \dots, \nu\}$, write $\mathcal{H}_i = \mathcal{H} \cap \chi^{-1}(i)$. We also write $\mathcal{H}_i(a) = \mathcal{H}_i \cap \mathcal{H}(a)$ and $\mathcal{H}_i(a, b) = \mathcal{H}_i \cap \mathcal{H}(a, b)$.

Note that (\mathcal{H}_i, \prec) is a total order, but that (\mathcal{H}_i, \ll) might not be. However, by Zorn's lemma, we can find a maximal subset, $\mathcal{I}_i \subseteq \mathcal{H}_i$, such that (\mathcal{I}_i, \ll) is totally ordered. Note that, on \mathcal{I}_i , the total orders \prec and \ll agree (so we will not need to specify the order). We write $\mathcal{I}_i(a) = \mathcal{I}_i \cap \mathcal{H}(a)$ and $\mathcal{I}_i(a, b) = \mathcal{I}_i \cap \mathcal{H}(a, b)$.

Lemma 4.3. *If $a, b \in \Delta$ are distinct, then there is some $i \in \{1, \dots, \nu\}$ such that $\mathcal{I}_i(a, b)$ contains a rational subset.*

Proof. By Lemma 4.2, there is a subset $\mathcal{H}' \subseteq \mathcal{H}(a, b)$ such that (\mathcal{H}', \ll) is order isomorphic to the real interval $[-1, 1]$. Applying Lemma 3.4, we can suppose that, for some i , $\mathcal{H}' \cap \mathcal{H}_i$ is dense in \mathcal{H}' , in particular, it contains a rational subset, say \mathcal{H}'_i . From the maximality of \mathcal{I}_i , between any two distinct elements of \mathcal{H}'_i , there must be some element of \mathcal{I}_i . Choosing one such element for every such pair in \mathcal{H}'_i , we get a countable subset of \mathcal{I}_i , which is easily seen to be rational. \square

Lemma 4.4. *For each $i \in \{1, \dots, \nu\}$, \mathcal{I}_i has a countable dense subset.*

Proof. Let $\beta : [-1, 1] \rightarrow \Delta$ be as in the hypothesis (I2). Given $H \in \mathcal{H}$, let $t(H) = \max(\beta^{-1}(H))$. Note that if $H \ll H'$, then $t(H) < t(H')$. Since $\mathcal{I}_i \subseteq \mathcal{H}$ is totally ordered, the map $t : \mathcal{I}_i \rightarrow [-1, 1]$ is strictly order preserving, and the statement follows. \square

Let $\theta_i : \mathcal{I}_i \rightarrow I$ be a non-decreasing map as given by Lemma 3.1. In order to simplify the exposition, we will assume that I is always equal to $[-1, 1]$. (We can simply discard those i for which it is $\{0\}$. They play no role in the discussion anyway.) In particular, if $H, H' \in \mathcal{I}_i$ are separated by a rational subset, then $\theta_i(H) \neq \theta_i(H')$.

Suppose $a \in \Delta$. Now $\mathcal{I}_i(a)$ is an initial segment of \mathcal{I}_i . We set $\phi_i(a) = \sup \theta_i(\mathcal{I}_i(a)) \in [-1, 1]$. (If $\mathcal{I}_i(a) = \emptyset$, we set $\phi_i(a) = -1$.) Note that $\phi_i(e^-) = -1$ and $\phi_i(e^+) = 1$. If $a \leq b$, then $\mathcal{I}_i(a) \subseteq \mathcal{I}_i(b)$, so $\phi_i(a) \leq \phi_i(b)$. Moreover, $\phi_i(a) < \phi_i(b)$ if and only if $\mathcal{I}_i(b) \setminus \mathcal{I}_i(a)$ contains a rational subset. In particular, there will be $H, H' \in \mathcal{I}_i$, with $a \ll H \ll H' \ll b$.

Now let $\Omega = [-1, 1]^\nu$. We define $\phi : \Delta \rightarrow \Omega$ by setting $\phi(a) = (\phi_1(a), \dots, \phi_\nu(a))$. Note that $\phi(e^-) = \omega^-$ and $\phi(e^+) = \omega^+$.

Lemma 4.5. *ϕ is a median homomorphism.*

Proof. Note that a median structure is determined by the set of intervals. Therefore, it's enough to show that if $a, b \in \Delta$ and $c \in [a, b]$, then

$\phi(c) \in [\phi(a), \phi(b)]$. In other words, $\phi_i(c) \in [\phi_i(a), \phi_i(b)]$ for all i . If not, then without loss in generality, we have $\phi_i(a) \leq \phi_i(b) < \phi_i(c)$. It follows that there is some $H \in \mathcal{I}_i$ with $a, b \in H$, but with $c \notin H$. But H is convex, so we get the contradiction $c \in [a, b] \subseteq H$. (We get a similar contradiction if $\phi_i(c) < \phi_i(a) \leq \phi_i(b)$.) \square

Lemma 4.6. *ϕ is injective.*

Proof. Suppose $a, b \in \Delta$ with $a \neq b$. By Lemma 4.3, there is some i such that $\mathcal{I}_i(a, b)$ contains a rational subset. Now $\mathcal{I}_i(a, b) = \mathcal{I}_i(a) \Delta \mathcal{I}_i(b)$, so we can assume that $\mathcal{I}_i(b) \setminus \mathcal{I}_i(a)$ contains a rational subset, and this implies that $\phi_i(a) < \phi_i(b)$. \square

Lemma 4.7. *ϕ is continuous.*

Proof. Let $(a^n)_n$ be a sequence in Δ converging to some $a \in \Delta$. We claim that $\phi(a^n) \rightarrow \phi(a)$ in Ω . If not, then passing to a subsequence, we would have $\phi(a^n) \rightarrow x$ for some $x \in \Omega \setminus \{\phi(a)\}$. In particular, there is some i with $x_i \neq \phi_i(a)$.

Suppose $\phi_i(a) < x_i$. Since $\phi_i(a^n) \rightarrow x_i$, and since $\theta_i(\mathcal{I}_i)$ is dense in $[-1, 1]$, we have $H, H' \in \mathcal{I}_i$ with $a \ll H \ll H' \ll a^n$ for all sufficiently large n . Thus $a \in H$ and $a^n \in \Delta \setminus \text{int}(H')$. But these are disjoint closed subsets of Δ , so we could not have $a^n \rightarrow a$.

If $x_i < \phi_i(a)$, we arrive at a similar contradiction. \square

Lemmas 4.5, 4.6 and 4.7 together now imply Proposition 1.4.

5. LATTICES IN THE CUBE

In this section we study further sublattices of $\Omega = [-1, 1]^\nu$. This is used in Section 6 for the proof of Theorem 1.1.

We recall from Section 2 that we define an order $<$ on Ω , by writing $x < y$ to mean that $x \neq y$ and $x_i \leq y_i$ for all i . Let $\omega^\pm = (\pm 1, \dots, \pm 1)$. By a *lattice* in Ω we mean a median subalgebra containing ω^+ and ω^- . (In particular, it is closed under \wedge and \vee .) Given $x, y \in \Omega$ write $N(x, y) = |\{i \mid x_i \neq y_i\}|$ for the number of coordinates in which x and y differ.

Suppose that (Π, ρ, μ) is a finite metric median algebra satisfying (L2). Recall that $\Pi = V(\Upsilon)$ where Υ is a finite CAT(0) cube complex. Given an edge $e \in E(\Upsilon)$ we set $\lambda(e) = \rho(x, y)$ where $x, y \in \Pi$ are the endpoints of e . Given $W \in \mathcal{W}$, recall that $E_W \subseteq E(\Upsilon)$ is the (non-empty) set of edges crossing W . We write $\lambda(W) = \min\{\lambda(e) \mid e \in E_W\}$. (We can intuitively think of $\lambda(W)$ as measuring the “width” of the wall W .)

Lemma 5.1. *If $e \in E_W$, then $\lambda(e) \leq k\lambda(W)$.*

Proof. If $e, e' \in E_W$ and $x, x' \in H^-(W)$ and $y, y' \in H^+(W)$ are their respective endpoints then x, x', y, y' forms a 2-hypercube in Π . In particular, $x' \in [x, y']$ and $y' \in [y, x']$ and so, by (L2), $\rho(x', y') \leq k\rho(x, y)$. \square

Given $x, y \in \Pi$, let $\mathcal{W}(x, y) \subseteq \mathcal{W}(\Pi)$ be the set of walls separating x from y . Let $\lambda(x, y) = \sum_{W \in \mathcal{W}(x, y)} \lambda(W)$. Note that any combinatorially shortest path from x to y in the 1-skeleton of Υ crosses each wall of $\mathcal{W}(x, y)$ exactly once, i.e. the sequence of edges in such a path has exactly one representative from E_W for each $W \in \mathcal{W}(x, y)$. Note also that, if $x, y, z \in \Pi$, then $\mathcal{W}(x, y) \subseteq \mathcal{W}(x, z) \cup \mathcal{W}(z, y)$. It follows that λ is a metric on Π . An immediate consequence of Lemma 5.1 is:

Lemma 5.2. *For any $x, y \in \Pi$, $\rho(x, y) \leq k\lambda(x, y)$.*

Now let $F \subseteq \Omega$ be a sublattice, and suppose that ρ is a metric on F satisfying (L2), and with the inclusion of (F, ρ) into Ω continuous. (A-priori, ρ might induce a finer topology than the subspace topology.) Suppose that we have a 1-lipschitz path, $\beta : [0, T] \rightarrow (F, \rho)$, with $\beta(0) = \omega^-$ and $\beta(T) = \omega^+$. Note that β is also continuous as a map into Ω . Given $z \in I$, and $i \in \{1, \dots, \nu\}$, let $P_i(z) = \{x \in F \mid x_i = z\}$. Note that β must cross $P_i(z)$ and so $P_i(z) \neq \emptyset$. Let $t_i(z) = \max \beta^{-1}(P_i(z))$. If $z, w \in [-1, 1]$ with $z \leq w$, then $t_i(z) \leq t_i(w)$. We write $s_i(z, w) = t_i(w) - t_i(z) \geq 0$. Given any $x, y \in F$ with $x \leq y$ we write $S(x, y) = \sum_{i=1}^{\nu} s_i(x_i, y_i)$.

The following will give an upper bound on ρ .

Lemma 5.3. *Suppose that $x, y \in F$, with $x \leq y$. Then $\rho(x, y) \leq 3^{\nu}kS(x, y)$.*

Proof. We will prove the following stronger statement by induction on p . If $x, y \in F$ with $x \leq y$ with $N(x, y) = p$, then $\rho(x, y) \leq 3^p kS(x, y)$. (Note that $p = |\{i \mid x_i < y_i\}|$.) If $p = 0$, then $x = y$, so this is trivial. We can therefore suppose, without loss of generality, that $x_1 < y_1$. Let $a = \beta(t_1(x_1))$ and $b = \beta(t_1(y_1))$, and let $c = \mu(x, y, a)$ and $d = \mu(x, y, b)$. Note that $x, a, c \in P_1(t_1(x_1))$ and that $y, b, d \in P_1(t_1(y_1))$. Now, since β is 1-lipschitz, $\rho(a, b) \leq t_1(y_1) - t_1(x_1) = s_1(x_1, y_1)$, and so by (L2), $\rho(c, d) \leq ks_1(x_1, y_1)$.

Note that $c \in [x, y]$, $x \leq c$ and $x_1 = c_1$. Thus, $N(x, c) < p$. Also $x_i \leq c_i$ for all i , so $t_i(x_i) \leq t_i(c_i) \leq t_i(y_i)$. In particular, $s_i(x_i, c_i) \leq$

$s_i(x_i, y_i)$ for all i . By the inductive hypothesis, we have:

$$\begin{aligned} \rho(x, c) &\leq 3^{p-1}kS(x, c) \\ &= 3^{p-1}k \sum_{i=1}^{\nu} s_i(x_i, c_i) \\ &\leq 3^{p-1}k \sum_{i=1}^{\nu} s_i(x_i, y_i) \\ &= 3^{p-1}kS(x, y). \end{aligned}$$

Similarly,

$$\rho(y, d) \leq 3^{p-1}kS(x, y).$$

Thus

$$\begin{aligned} \rho(x, y) &\leq \rho(c, d) + \rho(x, c) + \rho(y, d) \\ &\leq ks_1(x_1, y_1) + 2 \cdot 3^{p-1}kS(x, y) \\ &\leq (1 + 2 \cdot 3^{p-1})kS(x, y) \\ &\leq 3^p kS(x, y). \end{aligned}$$

The lemma now follows by induction, given that $p \leq \nu$. □

Now suppose that $\Pi \subseteq F$ is any finite subalgebra with $\omega^-, \omega^+ \in \Pi$. Let ρ be the metric induced on Π from F . Note that there is another metric, λ , on Π defined as above.

Lemma 5.4. $\lambda(\omega^-, \omega^+) \leq 3^\nu \nu kT$.

Proof. Choose a combinatorially shortest path, $\omega^- = x^0, x^1, \dots, x^n = \omega^+$ in $\Pi = V(\Upsilon)$. Thus, for each j , x^{j-1} and x^j are adjacent in the 1-skeleton of Υ . In other words, they are the endpoints of some edge $e^j \in E_{W^j}$ for some $W^j \in \mathcal{W}(\Pi)$. This path must cross each wall of $\mathcal{W}(\Pi)$ exactly once, and $x^{j-1} \leq x^j$ for all j . Note that, for each i , we have $\sum_{j=1}^n s_i(x_i^{j-1}, x_i^j) = T$ (directly from the definition of s_i , noting that $t_i(x_i^0) = 0$ and $t_i(x_i^n) = T$). By Lemma 5.3, $\rho(x^{j-1}, x^j) \leq 3^\nu kS(x^{j-1}, x^j)$. Also, by the definition of $\lambda(e^j)$, we have $\lambda(W^j) \leq$

$\lambda(e^j) = \rho(x^{j-1}, x^j)$ and so:

$$\begin{aligned} \lambda(\omega^-, \omega^+) &= \sum_j \rho(x^{j-1}, x^j) \\ &\leq 3^\nu k \sum_j S(x^{j-1}, x^j) \\ &= 3^\nu k \sum_{j,i} s_i(x_i^{j-1}, x_i^j) \\ &= 3^\nu k \sum_i T \\ &= 3^\nu k(\nu T). \end{aligned}$$

Thus, $\lambda(\omega^-, \omega^+) \leq 3^\nu \nu k T$ as claimed. \square

6. PROOF OF THEOREM 1.1

We first describe, in general terms, how we aim to construct our embedding.

Suppose, for the moment, that Φ is any hausdorff topological space. Let σ be a continuous pseudometric on Φ . Define an equivalence relation, \sim , on Φ by $x \sim y$ if $\sigma(x, y) = 0$. The quotient Φ/\sim admits a metric, τ , given by $\tau([x], [y]) = \sigma(x, y)$, where $[.]$ denotes equivalence class. This is the *hausdorffification* of (Φ, σ) . Note that if Φ is connected, then so is Φ/\sim . If σ satisfies property (T1), defined in the Introduction, then so does τ .

Suppose now that (Φ, ρ) is a connected metric space. Suppose we have ν pseudometrics, $\sigma_1, \dots, \sigma_\nu$, on Φ , all satisfying (T1). Suppose there is some $K > 0$ such that, for all $x, y \in \Phi$, we have:

$$\frac{1}{K} \rho(x, y) \leq \sum_{i=1}^{\nu} \sigma_i(x, y) \leq K \rho(x, y).$$

In particular, each σ_i is continuous, so the hausdorffification, (T_i, τ_i) is an \mathbb{R} -tree. Let $f : \Phi \rightarrow \prod_i T_i$ be the natural map. Then, f is injective, and bilipschitz onto its image. In this way, Theorem 1.1, is reduced to finding a suitable set of pseudometrics σ_i .

We now suppose that Φ be a metric median algebra satisfying (L1), (L2) and (L3). Let \mathcal{M} be the set of all finite median algebras of Φ . This is a directed set under inclusion.

Suppose that $\Pi \in \mathcal{M}$. Now Π is ν -colourable. Let $\chi^\Pi : \mathcal{W}(\Pi) \rightarrow \{1, \dots, \nu\}$ be any ν -colouring. Given any $a, b \in \Pi$ write $\mathcal{W}_i^\Pi(a, b) = (\chi^\Pi)^{-1}(i) \cap \mathcal{W}^\Pi(a, b)$, where $\mathcal{W}^\Pi(a, b)$ is the set of walls of Π separating a and b in Π . With the notation of Section 5, we define $\lambda_i(a, b) =$

$\lambda_i^\Pi(a, b) = \sum\{\lambda(W) \mid W \in \mathcal{W}_i^\Pi(a, b)\}$. Thus, λ_i is a pseudometric on Π . Moreover, $\lambda(a, b) = \sum_{i=1}^\nu \lambda_i(a, b)$, where $\lambda = \lambda^\Pi$ is the metric defined in Section 5.

Lemma 6.1. *For each i , the pseudometric λ_i satisfies (T1).*

Proof. Let $\Pi_i = \Pi/\sim_i$ be the median algebra described in Section 2. We saw that $\Pi/\sim_i = V(\Upsilon_i)$, where Υ_i is a simplicial tree. The walls in $(\chi^\Pi)^{-1}(i)$ are in bijective correspondence with the edges, $E(\Upsilon_i)$. In particular, we can assign to each edge $e \in E(\Upsilon_i)$ a length equal to $\lambda(W)$, where $W \in \mathcal{W}(\Pi)$ is the corresponding wall. This defines a metric, τ_i , on Π_i , where we sum the lengths of the edges between any given pair of points. Since Υ_i is a tree, τ_i satisfies (T1). Now, if $a, b \in \Pi$, then directly from the definition of λ_i , we see that $\lambda_i(a, b) = \tau_i(a', b')$, where a', b' are the projections of a, b to Υ_i . The statement now follows. \square

Lemma 6.2. *If $a, b \in \Pi$, then $\lambda(a, b) \leq 3^\nu \nu k^2 l \rho(a, b)$.*

Proof. Let $\Delta = [a, b] \subseteq \Phi$. By (L3) there is an l -lipschitz map $\alpha : [0, \rho(a, b)] \rightarrow \Phi$ with $\alpha(0) = a$ and $\alpha(\rho(a, b)) = b$. Composing with the k -lipschitz projection $[x \mapsto \mu(a, b, x)]$, and reparameterising, we get a 1-lipschitz map $\beta : [0, T] \rightarrow \Delta$, with $T = kl\rho(a, b)$, $\beta(0) = a$ and $\beta(T) = b$. Now Δ satisfies the hypotheses of Proposition 1.4. Thus, there is a continuous embedding, $\phi : \Delta \rightarrow \Omega$. (The co-ordinates of this embedding need bear no relation to the ν -colouring of Π .) Note that $\phi(a) = \omega^-$ and $\phi(b) = \omega^+$. Also, $\phi \circ \beta$ gives us a 1-lipschitz map of $[0, T]$ into $\phi(\Delta)$ sending 0 to ω^- and T to ω^+ . Thus, Lemma 5.4 gives $\lambda(a, b) \leq 3^\nu \nu k T = 3^\nu \nu k^2 l \rho(a, b)$ and the result follows. \square

Let $K = 3^\nu \nu k^2 l$. By Lemma 5.2 we have $\rho(a, b) \leq k\lambda(a, b)$. In summary, we have $\lambda(a, b) = \sum_i \lambda_i(a, b)$, $\frac{1}{k}\rho(a, b) \leq \lambda(a, b) \leq K\rho(a, b)$, and each λ_i satisfies (T1).

Let P be the Tychonoff cube $\prod_{a, b \in \Phi} [0, K\rho(a, b)]$. Given $\Pi \in \mathcal{M}$ define $\sigma_i^\Pi \in P$ by $\sigma_i^\Pi(a, b) = \lambda_i^\Pi(a, b)$ if $a, b \in \Pi$, and $\sigma_i^\Pi(a, b) = 0$ otherwise. (The latter is just an arbitrary assignment which plays no essential role.)

Now, P is compact, and \mathcal{M} is a directed set under inclusion. Thus, there is a cofinal map, $[j \mapsto \Pi_j]$, from a directed indexing set, J , such that $\sigma_i^{\Pi_j}$ converges to some $\sigma_i \in P$. We abbreviate $\sigma_i^j = \sigma_i^{\Pi_j}$ and $\lambda_i^j = \lambda_i^{\Pi_j}$, etc. For all $a, b \in \Phi$, we have $\sigma_i(a, b) = \lim_j \sigma_i^j(a, b)$.

Any finite subset of Φ lies in some element of \mathcal{M} . Thus, for any $A \subseteq \Phi$ finite, there is some $j \in J$ with $A \subseteq \Pi_j$. In particular, $\sigma_i^j(a, b) =$

$\lambda_i^j(a, b)$ for all $a, b \in A$. Moreover, this holds for all $j' \geq j$ in J . Therefore, applying Lemma 6.1, we have:

Lemma 6.3. *σ_i is a pseudometric on Φ and satisfies (T1).*

Now set $\sigma^j(a, b) = \sum_i \sigma_i^j(a, b)$ and $\sigma(a, b) = \sum_i \sigma_i(a, b)$ for all $a, b \in \Phi$. Recall that $\lambda^j(a, b) = \sum_i \lambda_i^j(a, b)$. Thus, by Lemmas 5.2 and 6.2, for all sufficiently large j , we have $\frac{1}{k}\rho(a, b) \leq \sigma^j(a, b) \leq K\rho(a, b)$, so passing to the limit, we get $\frac{1}{k}\rho(a, b) \leq \sigma(a, b) \leq K\rho(a, b)$. In particular, $\sigma_i(a, b) \leq \sigma(a, b) \leq K\rho(a, b)$, so σ_i is continuous on (Φ, ρ) .

We have thus obtained a set of pseudometrics on Φ as described at the beginning of this section. This therefore gives us a bilipschitz embedding of Φ into a product of ν \mathbb{R} -trees. This proves Theorem 1.1.

7. MEDIAN METRIC SPACES

In this section, we briefly describe a partial converse to Theorem 1.1, and relate this to median metric spaces as described in [ChaDH].

Let (M, d) be any metric space. Given $x, y \in M$, write

$$I_d(x, y) = \{z \in M \mid d(x, z) + d(z, y) = d(x, y)\}.$$

Suppose that Ψ is a direct product of ν \mathbb{R} -trees, with the l_1 -metric, ρ_1 . Suppose that $\Phi \subseteq \Psi$ is a subalgebra. We denote by $[x, y]_\Phi$ and $[x, y]_\Psi$ the median intervals between x and y in Φ and Ψ respectively. Thus, $[x, y]_\Phi = \Phi \cap [x, y]_\Psi$.

Lemma 7.1. *For all $x, y \in \Phi$, $I_{\rho_1}(x, y) = [x, y]_\Phi$.*

Proof. Given that $[x, y]_\Phi = \Phi \cap [x, y]_\Psi$, we may as well assume that $\Phi = \Psi$. Now, $[x, y]_\Psi$ is isomorphic, as a metric median algebra, to a product of closed real intervals in the l_1 -metric. Now it is easily seen that \mathbb{R}^ν satisfies Lemma 7.1, and so, in particular, we see that $[x, y]_\Psi \subseteq I_{\rho_1}(x, y)$. For the reverse inclusion, write $\pi(w) = \mu_\Psi(x, y, w) \in [x, y]_\Psi$. If $w \notin [x, y]_\Psi$, then $\rho_1(x, \pi(w)) < \rho_1(x, w)$ and $\rho_1(y, \pi(w)) < \rho_1(y, w)$, and so $x \notin I_{\rho_1}(x, y)$. \square

Suppose now that Φ is connected. Since the projection map from Φ to $[x, y]_\Phi$ is continuous, we see that $[x, y]_\Phi$ is connected. In particular, given that $[x, y]_\Phi = I_{\rho_1}(x, y)$, we see that x, y have a midpoint, i.e. a point z with $\rho_1(x, z) = \rho_1(y, z) = \frac{1}{2}\rho_1(x, y)$. Let $T = \rho_1(x, y)$ and let $Q = \{Tp/2^n \mid n \in \mathbb{N}, p \in [0, 2^n] \cap \mathbb{N}\}$, in other words, the dyadic rationals in $[0, 1]$ rescaled by T . By iterating the midpoint construction, we get an isometric map, $\alpha : Q \rightarrow [x, y]_\Phi \subseteq \Phi$, with $\alpha(0) = x$ and $\alpha(T) = y$. If Φ is closed in Ψ , then $[x, y]_\Phi$ is compact (since it is a

closed subset of $[x, y]_\Psi$ which is product of closed intervals). Thus, α extends to an isometric map $\alpha : [0, T] \rightarrow \Phi$. We deduce:

Proposition 7.2. *If $\Phi \subseteq \Psi$ is closed and connected, then Ψ is a geodesic space.*

In other words, it satisfies (L3) with $l = 1$.

In [ChaDH], a metric space (M, d) is said to be a “median metric space” if, for all $x, y, z \in M$, $I_d(x, y) \cap I_d(y, z) \cap I_d(z, x)$ consists of a single point. Thus, Lemma 7.1 tells us that a subalgebra of Ψ is median metric space in the l_1 -metric. We see that a consequence of Theorem 1.1 is that a metric median algebra satisfying (L1), (L2) and (L3) is bilipschitz equivalent to a median metric space. If Φ is complete then the median metric space is a geodesic space. (I suspect this to be true without the assumption of completeness.)

8. APPLICATIONS

In this section, we put the results in context by noting some applications of the statements. These follow from the fact that metric median algebras of the type described arise as ultralimits of coarse median algebras.

We recall the following definition from [Bo1].

Let (Λ, ρ) be a metric space, and let $\mu : \Lambda^3 \rightarrow \Lambda$ be a ternary operation. We say that μ is a *coarse median*, and that (Λ, ρ, μ) is a *coarse median space*, if the following hold:

(C1): There are constants, $k, h(0)$, such that for all $a, b, c, a', b', c' \in \Lambda$ we have

$$\rho(\mu(a, b, c), \mu(a', b', c')) \leq k(\rho(a, a') + \rho(b, b') + \rho(c, c')) + h(0).$$

(C2): There is a function, $h : \mathbb{N} \rightarrow [0, \infty)$, with the following property. Suppose that $A \subseteq \Lambda$ with $1 \leq |A| \leq p < \infty$, then there is a finite median algebra, (Π, μ_Π) and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ such that for all $x, y, z \in \Pi$ we have:

$$\rho(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z)) \leq h(p)$$

and

$$\rho(a, \lambda\pi a) \leq h(p)$$

for all $a \in A$.

We refer to k and h as the *parameters* of (Λ, ρ, μ) .

We say that Λ has *rank at most ν* if we can always choose Π to have rank at most ν (i.e. is the vertex set of a CAT(0) cube complex of dimension at most ν). Similarly, we say that Λ is *ν -colourable* if we can always choose Π to be ν -colourable. (Implicit in these definitions are fixed parameters, k and h .)

Suppose that we have a sequence, $((\Lambda_n, \rho_n, \mu_n))_{n \in \mathbb{N}}$, of coarse median spaces with parameters k, h_n . We suppose that k is fixed and that for each p , $h_n(p) \rightarrow 0$ as $n \rightarrow \infty$. Fixing a non-principal ultrafilter on \mathbb{N} , and a basepoint in each Λ_n , we obtain an ultralimit, $(\Lambda_\infty, \rho_\infty, \mu_\infty)$, which is a metric median algebra [Bo1]. This automatically satisfies (L2). Moreover, any such limit metric, ρ_∞ , will be complete. We also note the following from [Bo1]:

Lemma 8.1. *If each Λ_n has rank at most ν then Λ_∞ has rank at most ν . If each Λ_n is ν -colourable, then Λ_∞ is ν -colourable.*

Here, “rank” and “colourability” for Λ_n refer to the coarse median structure of Λ_n , and are assumed to be with respect to parameters, k, h_n , as in the previous paragraph. For Λ_∞ , “rank” and “colourability” refer to its structure as a (metric) median algebra, as defined earlier in this paper.

If each Λ_n is a “geodesic space”, then so is Λ_∞ , that is, it satisfies (L3) with $l = 1$.

A particular case of this construction arises by taking a coarse median space, (Λ, ρ, μ) , and rescaling the metric by a sequence of positive numbers, $(t_n)_n$, with $t_n \rightarrow 0$. That is, we set $\Lambda_n = \Lambda$, $\rho_n = t_n \rho$, and $\mu_n = \mu$. In this case, the ultralimit, $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is referred to as an “asymptotic cone” (see [VW, G]). It depends on a choice of ultrafilter, the sequence $(t_n)_n$ and a sequence of basepoints in Λ . A particular case of interest is where Λ is a connected graph, and ρ is the combinatorial metric (assigning each edge unit length).

In [Bo1] we defined a “coarse median group” as a finitely generated group for which the Cayley graph with respect to some (or equivalently any) finite set admits a coarse median. This is quasi-isometry invariant. Moreover one can specify that it be “of rank ν ” or “ ν -colourable”. (These properties are also quasi-isometry invariant.)

Examples of such groups are right-angled Artin groups (such as free abelian groups), hyperbolic groups, or groups hyperbolic with respect to other coarse median groups [Bo2]. Also, following from work of [BehM], the mapping class groups are coarse median groups [Bo1]. One might conjecture that it applies to a much broader class, including CAT(0) groups.

More specifically, let Σ be a compact orientable surface, and let $\xi(\Sigma)$ be its *complexity*, that is $\xi(\Sigma) = 3g + p - 3$, where g is the genus, and p is the number of boundary component. We suppose $\xi(\Sigma) > 1$. Let $\text{Map}(\Sigma)$ be the mapping class group of Σ . The following was proven in [Bo1]:

Theorem 8.2. *If Σ is a compact surface, then $\text{Map}(\Sigma)$ is a coarse median group of rank $\xi(\Sigma)$. Moreover, it is $\nu(\Sigma)$ -colourable, where $\nu(\Sigma)$ depends only on the topological type of Σ .*

In fact, the constant $\nu(\Sigma)$ arises from the topological construction given in [BesBF] and can be explicitly bounded.

As a consequence, we recover the statement in [BehDS] that any asymptotic cone of $\text{Map}(\Sigma)$ admits a bilipschitz embedding in a finite product of \mathbb{R} -trees. Moreover, the median structure arising agrees with that from [BehM]. Also, we also note that every interval in the asymptotic cone is compact, and admits a bilipschitz embedding into $\mathbb{R}^{\xi(\Sigma)}$, again respecting the median structure.

Another application which we discuss in the next section is to the “rapid decay” property.

9. RAPID DECAY

The property of rapid decay features prominently in work towards the Novikov and Baum-Connes conjectures, one of the earliest applications being [CoM]. Further references can be found in [ChaR] and [DrS]. Indeed this was the motivation behind introducing medians in [BehM].

In this section, we show:

Theorem 9.1. *A coarse median group of finite rank has rapid decay.*

To this end we use the criterion given in [ChaR], described later. (We note that a related criterion is given in [DrS].) For us, this boils down to showing that “coarse intervals” have polynomial growth, in an appropriate sense (see Proposition 9.8 below).

Suppose, for the moment, that (Λ, ρ) is any geodesic space admitting a coarse median, μ . As observed in [Bo1], there is no essential loss in assuming that $\mu(a, a, b) = a$ and that $\mu(a, b, c) = \mu(b, c, a) = \mu(b, a, c)$ for all $a, b, c \in \Lambda$ (since this is always true up to bounded distance in Λ , and hence exactly in some space quasi-isometric to Λ). Also, in (C2) we can always assume that $\Pi = \langle \pi(A) \rangle$, so that $|\Pi| \leq 2^{2^p}$. If we don’t care about rank, we can take Π to be the free median algebra on A , and also arrange that $\lambda \pi a = a$ for all $a \in \Pi$. (This will be convenient

later.) We will refer to k and h in the definition, as the *parameters* of Λ .

Given $\kappa > 0$ and $a, b \in \Lambda$, write

$$[a, b]_\kappa = \{e \in \Lambda \mid \rho(e, \mu(a, b, e)) \leq \kappa\}.$$

Lemma 9.2. *There is some $\kappa \geq 0$, depending only on the parameters, such that for all $a, b, c \in \Lambda$, $\mu(a, b, c) \in [a, b]_\kappa$.*

Proof. Applying (C2), this is an easy consequence of the fact that in any median algebra, Π , we have $\mu_\Pi(a, b, \mu_\Pi(a, b, c)) = \mu_\Pi(a, b, c)$, and so the same relation holds up to bounded distance in a coarse median space. \square

In particular, it follows that for such κ , $[a, b]_\kappa \cap [b, c]_\kappa \cap [c, a]_\kappa \neq \emptyset$.

Lemma 9.3. *The diameter of $[a, b]_\kappa$ is bounded above by a linear function of $\rho(a, b)$ depending only on the parameters and κ .*

Proof. Suppose $e \in [a, b]_\kappa$, so that $\rho(e, \mu(a, b, e)) \leq \kappa$. Now $\rho(\mu(a, b, e), \mu(a, a, e)) \leq k\rho(a, b) + h(0)$. Also $\rho(a, \mu(a, a, e))$ is bounded above in terms of the parameters. This gives a bound of the required form on $\rho(a, e)$. \square

Given a finite subset, $A \subseteq \Lambda$, write $\text{sep}(A) = \min\{\rho(x, y) \mid x, y \in A, x \neq y\}$ for the ‘‘separation constant’’. Suppose that \mathcal{A} is a set of finite subsets of Λ which is closed under inclusion (that is, $B \subseteq A \in \mathcal{A}$ implies $B \in \mathcal{A}$). Let $f(r, t) = \max\{|A| \mid A \in \mathcal{A}, \text{sep}(A) \geq r, \text{diam}(A) \leq rt\}$. This is, $f(r, t)$ is non-decreasing in t . We will assume that $f(r, t)$ is always finite (as is the case for uniformly locally finite graphs which we use later).

Lemma 9.4. *If $r, t, u > 0$, we have $f(r, tu) \leq f(r, 2t)f(rt, u)$.*

Proof. Let $A \in \mathcal{A}$ with $\text{sep}(A) \geq r$ and $\text{diam}(A) \leq rtu$. Let $B \subseteq A$ be maximal such that $\text{sep}(B) \geq rt$. Since $\text{diam}(B) \leq rtu$, we have $|B| \leq f(rt, u)$. Given $a \in A$, there is some $b(a) \in B$ with $\rho(a, b(a)) \leq rt$. Given $b \in B$, let $A(b) = \{a \in A \mid b(a) = b\}$. Thus, $\text{diam}(A(b)) \leq 2rt$ and $\text{sep}(A(b)) \geq r$, so $|A(b)| \leq f(r, 2t)$. Now $A \subseteq \bigcup_{b \in B} A(b)$, so $|A| \leq f(r, 2t)f(rt, u)$. Since this applies to any such A , we have $f(t, tu) \leq f(r, 2t)f(rt, u)$ as required. \square

Lemma 9.5. *Suppose that $(\exists C, \nu > 0)(\forall t > 0)(\exists s > 0)(\forall r \geq s)(f(r, t) \leq Ct^\nu)$. Then $f(1, r) = o(r^{\nu+\epsilon})$ for all $\epsilon > 0$.*

Proof. Let $\delta > 0$. Choose $u > (2^\nu C)^{\frac{1}{\delta}}$, so that $C(2u)^\nu \leq u^{\nu+\delta}$. Let s be as in the hypotheses with $t = 2u$. Given $m \in \mathbb{N}$, we have:

$$\begin{aligned} f(1, su^m) &\leq f(1, 2s)f(s, u^m) \\ &\leq f(1, 2s)f(s, 2u)f(su, u^{m-1}) \\ &\leq f(1, 2s)f(s, 2u)f(su, 2u) \cdots f(su^{m-1}, 2u) \\ &\leq f(1, 2s)(C(2u)^\nu)^m \\ &\leq f(1, 2s)u^{(\nu+\delta)m} \end{aligned}$$

using induction and the fact that $f(su^{m-1}, u) \leq f(su^{m-1}, 2u)$. Setting $K = f(1, 2s)$, we have shown that $(\forall \delta > 0)(\exists K, s, u > 0)(\forall m \in \mathbb{N})(f(1, su^m) \leq Ku^{(\nu+\delta)m})$. From this it follows that $f(1, r) = o(r^{\nu+\epsilon})$ for all $\epsilon > 0$. This is perhaps most easily seen by taking logarithms. We have a non-decreasing function, $\log f(1, r)$, which, for all $\delta > 0$, is bounded by $(\nu + \delta) \log r$ on an infinite arithmetic sequence of values of $\log r$. It follows that $(\log f(1, r))/\log r$ is eventually less than $\nu + \epsilon$ for any $\epsilon > 0$. \square

The aim now is to verify the hypotheses of Lemma 9.5 in the case where \mathcal{A} is set of all finite subsets of intervals of the form $[a, b]_\kappa$. We begin with the following lemma, which effectively tells us that we can assume that a, b are not too far away from a given $A \in \mathcal{A}$.

Lemma 9.6. *Given any $p \in \mathbb{N}$ and $\kappa \geq 0$, together with the parameters, k, h , of Λ , there exist constants κ', ζ and ξ which depend only on p, κ, k, h , such that the following holds. Suppose $a, b \in \Lambda$ and $A \subseteq [a, b]_\kappa$ with $|A| \leq p$. Then there exist $c, d \in \Lambda$ with $A \subseteq [c, d]_{\kappa'}$ and with $\text{diam}(A \cup \{c, d\}) \leq \zeta \text{diam}(A) + \xi$.*

Proof. Write $D = \text{diam}(A)$ and set $h_0 = h(p+2) \geq h(0)$. Let $\pi : A \cup \{a, b\} \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ be the maps given by the hypothesis (C2). We can assume that $\Pi = \langle \pi(A \cup \{a, b\}) \rangle$, so that $|\Pi| \leq q = 2^{2^{p+2}}$. Moreover (since we make no reference to rank in the statement of the theorem) we can assume that $\lambda \pi e = e$ for all $e \in A$. (This will serve to simplify notation.)

The basic idea of the proof is that if a is a long way from A , then πa will be an extreme point of Π (identified as the vertex set of a cube complex), and we can replace a by λx , where x is the unique point of Π adjacent to πa . Similarly, if b is far from A , we can replace it by λy , where $y \in \Pi$ is adjacent to πb . We will need some more technical statements in order to set this up.

We first claim that there are constants ζ_0 and ξ_0 , depending only on κ, k, h_0, p such that $\text{diam}(\lambda(\Pi \setminus \{\pi a, \pi b\})) \leq \zeta_0 D + \xi_0$.

To see this, we note that Π is generated from $\pi(A \cup \{a, b\})$ by iterating the median operation at most q times. More precisely, set $M_1 = \pi(A \cup \{a, b\})$, and define M_m inductively by $M_{m+1} = \mu_\Pi(M_m^3) \subseteq \Pi$. Thus, $\Pi = M_q$. We set $N_m = M_m \setminus \{\pi a, \pi b\}$.

We claim, inductively, that there are constants, ζ_m, ξ_m , depending only on κ, k, h_0, m such that $\text{diam}(\lambda N_m) \leq \zeta_m D + \xi_m$. For $m = 1$, we simply note that $N_1 \subseteq \pi A$, so that $\lambda N_1 \subseteq \lambda \pi A$, and we set $\zeta_1 = 1$ and $\xi_1 = 0$.

Suppose that $u \in N_{m+1}$. Then, $u = \mu_\Pi(x, y, z)$ where $x, y, z \in M_m = N_m \cup \{\pi a, \pi b\}$. Since $u \notin \{\pi a, \pi b\}$, we can assume that $z \notin \{\pi a, \pi b\}$. We distinguish two cases.

First suppose that $\{x, y\} \neq \{\pi a, \pi b\}$. Then, we can assume that $y \in M_m$, and so the inductive hypothesis gives $\rho(\lambda y, \lambda z) \leq \zeta_m D + \xi_m$. Thus, by (C1), we have $\rho(\lambda z, \lambda u) \leq k\rho(\lambda y, \lambda z) + h_0 \leq k\zeta_m D + k\xi_m + h_0$.

Now suppose that $\{x, y\} = \{\pi a, \pi b\}$, so that $u = \mu_\Pi(\pi a, \pi b, z)$. Since $z \in M_m$ and $A \subseteq M_m$, there is some $e \in A$ with $\rho(e, \lambda z) \leq \zeta_m D + \xi_m$. By hypothesis, $e \in [a, b]_\kappa$, that is, $\mu(e, \mu(a, b, e)) \leq \kappa$. We also note that $\rho(\mu(a, b, e), \mu(a, b, \lambda z)) \leq k\rho(e, \lambda z) \leq k(\zeta_m D + \xi_m) + h_0$ and $\rho(\mu(a, b, \lambda z), \lambda u) = \rho(\mu(\lambda \pi a, \lambda \pi b, \lambda z), \lambda \mu(\pi a, \pi b, z)) \leq h_0$, so $\rho(e, \lambda u) \leq \kappa + k(\zeta_m D + \xi_m) + 2h_0$.

In both cases, we have bounded $\rho(\lambda u, \lambda N_m)$ linearly in terms of D . We conclude that $\text{diam}(\lambda N_m) \leq (\zeta_m D + \xi_m) + 2(\kappa + k(\zeta_m D + \xi_m) + 2h_0) \leq (1 + 2k)\xi_m D + (1 + 2k)\xi_m + 2\kappa + 4h_0 \leq \zeta_{m+1} D + \xi_{m+1}$, on setting $\zeta_{m+1} = (1 + 2k)\xi_m$ and $\xi_{m+1} = (1 + 2k)\xi_m + 2\kappa + 4h_0$. The first claim now follows by induction, noting that $\Pi \setminus \{\pi a, \pi b\} \subseteq N_q$, and setting $\zeta_0 = \zeta_q$ and $\xi_0 = \xi_q$.

Suppose $\pi a \in [y, z]_\Pi$, where $y, z \in \Pi \setminus \{\pi a, \pi b\}$. Now $a = \lambda \pi a = \lambda \mu_\Pi(\pi a, y, z)$, so $\rho(a, \mu(a, \lambda x, \lambda y)) \leq h_0$. Also $\rho(a, \lambda x) \leq k(\zeta_0 D + \xi_0) + 2h_0$, and $\rho(\lambda y, A) \leq \zeta_0 D + \xi_0$. Thus, $\rho(a, A) \leq L$, where $L = (1 + k)\zeta_0 D + (1 + k)\xi_0 + 2h_0$. Similarly, if $\pi b \in [y, z]_\Pi$, for $y, z \in \Pi \setminus \{\pi a, \pi b\}$, then $\rho(b, A) \leq L$.

Suppose that $[\pi a, \pi b]_\Pi = \{\pi a, \pi b\}$ has at most two points. Then, swapping a with b , if necessary, we can suppose that there is some $e \in A$ with $\pi b \in [\pi a, \pi e]$. Now $b = \lambda \pi b = \mu(\pi a, \pi b, \pi e)$, so $\rho(b, \mu(a, b, e)) \leq \rho(\lambda \mu_\Pi(\pi a, \pi b, \pi e), \mu(\lambda \pi a, \lambda \pi b, \lambda \pi e)) \leq h_0$. But $e \in [a, b]_\kappa$, so $\rho(a, \mu(a, b, e)) \leq \kappa$, thus, $\rho(b, a) \leq \kappa + h_0$. In particular, $\rho(b, a) \leq L$. By a similar argument, if $[\pi a, \pi b]$ has at most three points, there is some $x \in \Pi$ with $\rho(\lambda x, A) \leq L$.

Recall that $\Pi = V(\Upsilon)$, where Υ is a finite CAT(0) complex. We say that $w \in \Pi$ is *strictly between* $u, v \in \Pi$ if $w \in [u, v]_\Pi \setminus \{u, v\}$. We say that u is *terminal* if it is not strictly between any two other points

of Π . In other words, it is incident on just one edge of Υ . We write $w(x) \in \Pi$ for the adjacent vertex.

At least one of the following five cases must occur:

- (1) $\pi a, \pi b$ are both terminal. We set $x = w(\pi a)$ and $y = w(\pi b)$.
- (2) πa is terminal, and b lies between two points of $\Pi \setminus \{\pi a, \pi b\}$, We set $x = w(\pi a)$ and $y = \pi b$.
- (3) The same as (2), swapping a with b and x with y .
- (4) Both πa and πb lie between two points of $\Pi \setminus \{\pi a, \pi b\}$, and we set $x = \pi a$ and $y = \pi b$.
- (5) The median interval $[\pi a, \pi b]_\Pi$ has at most three points, and is connected to the rest of Υ by a single cut vertex, $z \in \Pi$.

In this case, we set $x = y = z$.

We now set $c = \lambda x$ and $d = \lambda y$. In all cases, $\rho(c, A) \leq L$ and $\rho(d, A) \leq L$, so $\text{diam}(A \cup \{c, d\}) \leq 2L + D = \zeta D + \xi$, where ζ and ξ depend only on κ, k, h_0, p .

Moreover, if $e \in A \subseteq \Pi \setminus \{\pi a, \pi b\}$, then by choice of x, y we have $\mu_\Pi(\pi a, \pi b, \pi e) = \mu_\Pi(x, y, \pi e)$. Now $\mu(c, d, e) = \mu(\lambda x, \lambda y, \lambda \pi e)$, and so we have $\rho(\mu(c, d, e), \lambda \mu(\pi a, \pi b, \pi e)) = \rho(\mu(\lambda x, \lambda y, \lambda \pi e), \lambda \mu(x, y, \pi e)) \leq h_0$. Also, $\rho(\mu(a, b, e), \lambda \mu_\Pi(\pi a, \pi b, \pi e)) \leq h_0$, and $\rho(e, \mu(a, b, e)) \leq \kappa$, so $\rho(e, \mu(c, d, e)) \leq \kappa'$, where $\kappa' = \kappa + 2h_0$.

We have shown that $A \subseteq [c, d]_{\kappa'}$ and $\text{diam}(A \cup \{c, d\}) \leq \zeta D + \xi$, as required. \square

We now suppose that Λ has rank at most ν .

Lemma 9.7. $(\exists C > 0)(\forall \kappa, t > 0)(\exists s > 0)(\forall r \geq s)$ if $a, b \in \Lambda$ and $A \subseteq [a, b]_\kappa$ with $\text{sep}(A) \geq r$ and with $\text{diam}(A) \leq rt$ then $|A| \leq Ct^\nu$.

Here C is an explicit function of ν and the parameters of Λ . So we are saying that uniformly separated subsets of intervals have cardinality eventually bounded above by a polynomial in the diameter. As we have stated it, the scale, s , at which this bound starts might depend on Λ . In fact, we shall see later that it depends only on κ, ν, t and the parameters of Λ .

Proof. We set $C = NK^2$, where N depends only on ν and where K is the bilipschitz constant arising from an application of Theorem 1.1 to the ultralimit Λ_∞ constructed below. We will see that K depends only on the parameters, k, h , of Λ . We will specify N and K later.

We now fix κ and assume that the conclusion fails. That is, there is a particular t for which the result fails. We set $p = [Ct^\nu + 1] > Ct^\nu$. In other words, there is a sequence, $(r_n)_n$ with $r_n \rightarrow \infty$, points $a_n, b_n \in \Lambda$ and subsets, $A_n \subseteq [a_n, b_n]_\kappa$, with $\text{sep}(A_n) \geq r_n$, $\text{diam}(A_n) \leq r_n t$ and $|A_n| = p$.

By Lemma 9.6, we find $c_n, d_n \in \Lambda$ so that $\text{diam}(A_n \cup \{c_n, d_n\}) \leq \zeta \text{diam}(A_n) + \xi \leq \zeta r_n t + \xi$, and with $A_n \subseteq [c_n, d_n]_{\kappa'}$.

We now rescale the metric ρ by a factor of $1/r_n$ to give new spaces Λ_n . With respect to the new metrics, we have $\text{sep}(A_n) \geq 1$, $\text{diam}(A_n) \leq t$ and $\text{diam}(A_n \cup \{c_n, d_n\}) \leq \zeta t + \xi/r_n$ which is bounded. We choose basepoints, $e_n \in A_n$, and pass to an ultralimit $(\Lambda_\infty, \rho_\infty, \mu_\infty)$. We have $A_n \rightarrow A_\infty$, $c_n \rightarrow c$ and $d_n \rightarrow d$. Now $A_\infty \subseteq [c, d]$, $\text{sep}(A_\infty) \geq 1$, $\text{diam}(A_\infty) \leq t$ and $|A_\infty| = p$.

By Corollary 1.2, we can find a K -bilipschitz embedding, ϕ , of $[c, d]$ into \mathbb{R}^ν with the l_1 -metric, where K depends only on k and h . (Here $l = 1$.) Let $B = \phi(A_\infty)$. We have $\text{sep}(B) \geq 1/K$, $\text{diam}(B) \leq Kt$, and $|B| = p$. Thus, $p \leq NK^2 t^\nu$, where N depends only on ν . But we originally set $C = NK^2$ and $p = [Ct^\nu + 1]$, so this is a contradiction. \square

We now let \mathcal{A} be the set of finite subsets, A , of Λ for which there exist $a, b \in \Lambda$ with $A \subseteq [a, b]_\kappa$. Let $f(r, t)$ be the function arising, as defined before Lemma 9.3. Reinterpreting Lemma 9.6, we have now verified the hypotheses of Lemma 9.5 and so we have $f(1, r) = o(t^{\nu+\epsilon})$ for all $\epsilon > 0$.

Suppose now that Λ is a connected uniformly locally finite graph (each vertex has bounded finite degree). In this case, $f(r, t)$ is always finite, and we conclude:

Proposition 9.8. *Suppose that Λ is a connected, uniformly locally finite graph with combinatorial metric, ρ , and that μ is a coarse median of rank at most ν on (Λ, ρ) . Then, given any $\kappa > 0$, there is a function, $P : \mathbb{N} \rightarrow \mathbb{N}$, with $P(r) = o(r^{\nu+\epsilon})$ for all $\epsilon > 0$, such that if $a, b \in \Lambda$ and $Q \subseteq [a, b]_\kappa \cap V(\Lambda)$, then $|Q| \leq P(\text{diam}(Q))$.*

Proof. We just set $P(r) = f(1, r)$. \square

Although we won't need it here, we remark that we can take P to depend only on the parameters, the constant κ , and the maximal degree of a vertex of Λ . To see this, note that we can allow the graph Λ to vary in the proof of Lemma 9.7, taking a sequence of putative counterexamples, and passing to an ultralimit to arrive at a contradiction. Nevertheless, the argument remains non-constructive. One could ask whether the Proposition 9.8 holds with P a polynomial of degree ν , and moreover, whether this is explicitly computable.

Now let Γ be a finitely generated group.

In [ChaR], the following sufficient condition was given for a finitely generated group, Γ , to have the rapid decay property.

Suppose that to each pair, $a, b \in \Gamma$, we have associated set, $C(a, b) = C(b, a) \subseteq \Gamma$. We assume this to be Γ -equivariant, i.e. $C(ga, gb) =$

$gC(a, b)$ for all $g \in \Gamma$. We also suppose:

(R1): For all $a, b, c \in \Gamma$, such that $C(a, b) \cap C(b, c) \cap C(c, a) \neq \emptyset$.

(R2): There is a polynomial, P , such that, for all $a, b \in \Gamma$, and all $r \in \mathbb{N}$, we have $|\{e \in C(a, b) \mid \rho(a, e) \leq r\}| \leq P(r)$.

(R3): There is a polynomial, R , such that for all $a, b \in \Gamma$, $\text{diam } C(a, b) \leq R(\rho(a, b))$.

Then Γ has rapid decay [ChaR].

Suppose now that Γ is a coarse median group of rank at most ν . Let (Λ, ρ) be any Cayley graph, so that $\Gamma = V(\Lambda)$. By definition, Λ is a coarse median space. We fix κ as in Lemma 9.2 and set $C(a, b) = [a, b]_\kappa \cap \Gamma$. Now (R1), (R2) and (R3) follow respectively from Lemma 9.2, Proposition 9.8 and Lemma 9.3. We deduce that Γ has rapid decay, proving Theorem 9.1.

REFERENCES

- [BaH] H.-J.Bandelt, J.Hedlikova, *Median algebras* : Discrete Math. **45** (1983) 1–30.
- [BaV] H.-J.Bandelt, M.van de Vel, *Embedding topological median algebras in products of dendrons* : Proc. London Math. Soc. **58** (1989) 439–453.
- [BehDS] J.A.Behrstock, C.Druţu, M.Sapir, *Median stuctures on asymptotic cones and homomorphisms into mapping class groups* : Proc. London Math. Soc. **3** (2011) 503–554.
- [BehM] J.A.Behrstock, Y.N.Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.
- [BesBF] M.Bestvina, K.Bromberg, K.Fujiwara, *The asymptotic dimension of the mapping class groups is finite* : preprint, 2010.
- [Bo1] B.H.Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.
- [Bo2] B.H.Bowditch, *Invariance of coarse median spaces under relative hyperbolicity* : Math. Proc. Camb. Phil. Soc. **154** (2013) 85–95.
- [BriH] M.Bridson, A.Haefliger, *Metric spaces of non-positive curvature* : Grundlehren der Math. Wiss. No. 319, Springer (1999).
- [BroCGNW] 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.
- [ChaDH] I.Chatterji, C.Druţu, F.Haglund, *Kazhdan and Haagerup properties from the median viewpoint* : Adv. Math. **225** (2010) 882–921.
- [ChaR] I.Chatterji, K.Ruane, *Some geometric groups with rapid decay* : Geom. Funct. Anal. **15** (2005) 311–339.
- [Che] V.Chepoi, *Graphs of some CAT(0) complexes* : Adv. in Appl. Math. **24** (2000) 125–179.
- [CheH] V.Chepoi, M.F.Hagen, *On embeddings of CAT(0) cube complexes into products of trees* : preprint, 2011.

- [CoM] A.Connes, H.Moscovici, *Cyclic cohomology, the Novikov conjecture and hyperbolic groups* : Topology **29** (1990) 345–388.
- [Di] R.P.Dilworth, *A decomposition theorem for partially ordered sets* : Ann. of Math. **51** (1950) 161–166.
- [DrS] C.Druţu, M.Sapir, *Relatively hyperbolic groups with the rapid decay property* : Int. Math. Res. Not. **19** (2005) 1181–1194
- [G] M.Gromov, *Asymptotic invariants of infinite groups* : “Geometric group theory, Vol. 2” London Math. Soc. Lecture Note Ser. No. 182, Cambridge Univ. Press (1993)
- [HY] G.J.Hocking, G.S.Young, *Topology* : Addison-Wesley, 1961.
- [I] J.R.Isbell, *Median algebras* : Trans. Amer. Math. Soc. **260** (1980) 319–362.
- [R] M.A.Roller, *Poc-sets, median algebras and group actions, an extended study of Dunwoody’s construction and Sageev’s theorem* : Habilitationsschrift, Regensburg, 1998.
- [VW] L.Van den Dries, A.J.Wilkie, *On Gromov’s theorem concerning groups of polynomial growth and elementary logic* : J. Algebra **89** (1984) 349–374.

MATHEMATICS INSTITUTE, UNIVERSITY OF WARWICK, COVENTRY, CV4 7AL,
GREAT BRITAIN