

*Pacific
Journal of
Mathematics*

COARSE MEDIAN SPACES AND GROUPS

BRIAN H. BOWDITCH

Volume 261 No. 1

January 2013

COARSE MEDIAN SPACES AND GROUPS

BRIAN H. BOWDITCH

We introduce the notion of a coarse median on a metric space. This satisfies the axioms of a median algebra up to bounded distance. The existence of such a median on a geodesic space is quasi-isometry invariant, and so it applies to finitely generated groups via their Cayley graphs. We show that asymptotic cones of such spaces are topological median algebras. We define a notion of rank for a coarse median and show that this bounds the dimension of a quasi-isometrically embedded euclidean plane in the space. Using the centroid construction of Behrstock and Minsky, we show that the mapping class group has this property, and recover the rank theorem of Behrstock and Minsky and of Hamenstädt. We explore various other properties of such spaces, and develop some of the background material regarding median algebras.

1. Introduction

In this paper we introduce the notion of a “coarse median” on a metric space. The existence of such a structure can be viewed as a kind of coarse nonpositive curvature condition. It can also be applied to finitely generated groups. Many naturally occurring spaces and groups admit such structures. Simple examples include Gromov hyperbolic spaces and $CAT(0)$ cube complexes. It is also preserved under quasi-isometry, relative hyperbolicity and direct products. Moreover (using the construction of [Behrstock and Minsky 2011]), the mapping class group of a surface admits such a structure. One might conjecture that it applies to a much broader class of spaces that are in some sense nonpositively curved, such as $CAT(0)$ spaces. Much of this work is inspired by the results in [Behrstock and Minsky 2008; 2011; Bestvina et al. 2010; Behrstock et al. 2012; 2011; Chatterji et al. 2010]. It seems a natural general setting in which to view some of this work.

A “median algebra” is a set with a ternary operation satisfying certain conditions (see for example [Isbell 1980; Bandelt and Hedlíková 1983; Roller 1998; Chepoi 2000]). As we will see, for many purposes, one can reduce the discussion to a finite subalgebra. Any finite median algebra is canonically the vertex set of a $CAT(0)$

MSC2010: 20F65.

Keywords: median algebra, cube complex, rank, mapping class group .

cube complex, with the median defined in the usual way. One way to say this is that the median of three points is the unique point which minimises the sum of the distances in the 1-skeleton to these three points. For a fuller discussion, see Sections 2, 4 and 5.

We will also define a “coarse median” as a ternary operation on a metric space. We usually assume this to be a “geodesic space”, that is, every pair of points can be connected by a geodesic. The coarse median operation is assumed to satisfy the same conditions as a median algebra up to bounded distance. We can define the “rank” of such a space (which corresponds to the dimension of a CAT(0) complex). We show that the asymptotic cone [van den Dries and Wilkie 1984; Gromov 1993] of such a space is a topological median algebra. It has a “separation dimension” which is at most the rank, when this is finite. We remark that coarse median spaces of rank 1 are the same as Gromov hyperbolic spaces. In such a case, the asymptotic cone is an \mathbb{R} -tree.

The existence of a coarse median on a geodesic space is a quasi-isometry invariant, so we can apply this to finitely generated groups via their Cayley graphs. We can thus define a “coarse median group”. For example, a hyperbolic group is a coarse median group of rank 1, and a free abelian group is a coarse median group where “rank” agrees with the standard notion. More substantially we show that the mapping class group of a surface has a coarse median structure whose rank equals the maximal rank of a free abelian subgroup. The median we use for this is the centroid constructed in [Behrstock and Minsky 2011]. In particular, the asymptotic cone has at most (in fact precisely) this dimension, thereby giving another proof the rank theorem of [Behrstock and Minsky 2008; Hamenstädt 2005].

Another class of examples arise from relatively hyperbolic groups. We show in [Bowditch 2011b] that a group that is hyperbolic relative to a collection of coarse median groups (of rank at most ν) is itself coarse median (of rank at most ν). Examples of such are geometrically finite kleinian groups (of dimension ν) and Sela’s limit groups.

It is natural to ask what other classes of spaces or groups admit such a structure. For example, it is conceivable that every CAT(0) space does, where the rank might be bounded by the dimension. More modestly one could ask this for higher rank symmetric spaces. The only immediately evident constraint is that such a space should satisfy a quadratic isoperimetric inequality.

In [Bowditch 2011a], we show that a metric median algebra of the type that arises as an asymptotic cone of a finite rank coarse median space admits a bilipschitz embedding into a finite product of \mathbb{R} -trees. One consequence is that coarse median groups have rapid decay. In fact, their proof of rapid decay of the mapping class groups was the main motivation for introducing centroids in [Behrstock and Minsky 2011].

2. Statement of results

We begin by recalling the notion of a “median algebra”. This is a set equipped with a ternary “median” operation satisfying certain axioms. Discussion of these can be found in [Isbell 1980; Bandelt and Hedlíková 1983; Roller 1998; Chepoi 2000]. We will give a more detailed account in Sections 4 to 6. For the moment, we use more intuitive formulations of the definitions. A finite median algebra is essentially an equivalent structure to a finite cube complex. Recall that a (finite) cube complex is a connected metric complex built out of unit euclidean cubes. It is CAT(0) if it is simply connected and the link of every cube is a flag complex. See [Bridson and Haefliger 1999] for a general discussion. Note that a 1-dimensional CAT(0) cube complex is a simplicial tree.

Suppose M is a set, and $\mu : M^3 \rightarrow M$ is a ternary operation. Given $a, b \in M$, write $[a, b] = \{e \in M \mid \mu(a, b, e) = e\}$. This is the *interval* from a to b .

If $M = V(\Pi)$ is the vertex set of a finite cube complex, Π , we can define $[a, b]_\Pi$ to be the set of points of M which lie in some geodesic from a to b in the 1-skeleton of Π . One can show that there is a unique point, $\mu_\Pi(a, b, c)$, lying in $[a, b]_\Pi \cap [b, c]_\Pi \cap [c, a]_\Pi$. (In fact, it is the unique point which minimises the sum of the distances in the 1-skeleton to a , b and c .)

For the purposes of this section, we can define a “finite median algebra” to be a set M with a ternary operation: $\mu : M^3 \rightarrow M$ such that M admits a bijection to the vertex set, $V(\Pi)$, of some finite CAT(0) cube complex, Π , such that $\mu = \mu_\Pi$. (This is equivalent to the standard definition.) Given $a, b \in M$, write $[a, b] = \{e \in M \mid \mu(a, b, e) = e\}$. This is the *interval* from a to b . Under the bijection with $V(\Pi)$ it can be seen to agree with $[a, b]_\Pi$. Note that $\mu(a, b, c) = \mu(b, a, c) = \mu(b, c, a)$ and $\mu(a, a, b) = a$ for all $a, b, c \in M$. In fact, the complex Π is determined up to isomorphism by (M, μ) , so we can define the “rank” of M to be the dimension of Π . For more details, see Section 4.

In general, we say that a set, M , equipped with a ternary operation, μ , is a “median algebra”, if every finite subset $A \subseteq M$ is contained in another finite subset, $B \subseteq M$, which is closed under μ and such that (B, μ) is a finite median algebra. Note that, defining intervals in the same way, we again have $[a, b] \cap [b, c] \cap [c, a] = \{\mu(a, b, c)\}$ for all $a, b, c \in M$. We say that M has “rank at most ν ” if every finite subalgebra has rank at most ν . It has “rank ν ” if it has rank at most ν but not at most $\nu - 1$.

A median algebra of rank 1 is a treelike structure which has been studied under a variety of different names. They appear in [Sholander 1952] and as “tree algebras” in [Bandelt and Hedlíková 1983]. They have also been called “median pretrees”.

We introduce the following notion of a “coarse median space”. Let (Λ, ρ) be a metric space and $\mu : \Lambda^3 \rightarrow \Lambda$ be a ternary operation. We say that μ is a “coarse median” if it satisfies the following:

(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 $a \in A$.

Using (C1) and (C2) we can deduce that, if $a, b, c \in \Lambda$, then $\mu(a, b, c)$, $\mu(b, a, c)$ and $\mu(b, c, a)$ are a bounded distance apart, and that $\rho(\mu(a, a, b), a)$ is bounded. (These facts follow from the corresponding identities in the median algebra (Π, μ_Π) ; see the discussion in Section 8.) Thus, there is no essential loss in assuming μ to be invariant under permutation of a, b, c and assuming that $\mu(a, a, b) = a$.

If (Λ, ρ) is a geodesic space, then we can replace (C1) by a condition to the effect that if $\rho(c, d)$ is less than some fixed positive constant (for example, 1, for a graph) then $\rho(\mu(a, b, c), \mu(a, b, d))$ is bounded. It then follows for any a, b, c, d that $\rho(\mu(a, b, c), \mu(a, b, d))$ is, in fact, linearly bounded above in terms of $\rho(c, d)$.

Definition. We refer to μ as a *coarse median* on (Λ, ρ) if it satisfies (C1) and (C2) above. We refer to (Λ, ρ, μ) as a *coarse median space*.

If, in the above definition, we can strengthen (C2) to insist that Π has rank most v (independently of p), then we say that μ is a *coarse median of rank at most v* , and that (Λ, ρ, μ) is a *coarse median space of rank at most v* .

We refer to the multiplicative constant k and the additive constants, $h(p)$, featuring in the definitions as the *parameters* of the coarse median space.

Recall that a metric space is a “geodesic space” (or “length space”) if every pair of points are connected by a geodesic (that is, a path whose length equals the distance between its endpoints). In this context, coarse median spaces of rank 1 are precisely Gromov hyperbolic spaces (as defined in [Gromov 1987]).

Theorem 2.1. *Let (Λ, ρ) be a geodesic space. Then (Λ, ρ) is Gromov hyperbolic if and only if it admits a structure as a coarse median space of rank 1.*

In the above one can determine the parameters explicitly in terms of the hyperbolicity constant. The converse we offer here will be nonconstructive and based on the fact that any asymptotic cone is an \mathbb{R} -tree. (It is possible to give a constructive argument and explicit constants, but we will not pursue that matter here.)

By a *topological median algebra* we mean a topological space, M , equipped with a median, μ , which is continuous as a map from M^3 to M . Such structures are considered, for example, in [Bandelt and van de Vel 1989]. We will refer to a “metric median algebra” when the topology is induced by some particular metric.

We define a notion of “local convexity” in Section 7. For a finite-rank algebra this is equivalent to saying that an interval connecting two points close together is arbitrarily small. We will also define a notion of “separation dimension” of a topological space. This is analogous to (though weaker than) the standard notion of “inductive dimension”. The latter is equivalent to covering dimension [Hurewicz and Wallman 1941; Engelking 1995]. Every locally compact subspace of a space of separation dimension at most ν has covering dimension at most ν . In particular, such a space does not admit any continuous injective map of $\mathbb{R}^{\nu+1}$. We show:

Theorem 2.2. *A locally convex topological median algebra of rank at most ν has separation dimension at most ν .*

This notion of dimension is weaker than the standard notions of topological dimension referred to. For example, there is a totally disconnected space of positive covering dimension [Erdős 1940], but this has separation dimension 0. (I thank Klaas Hart for providing me with this reference.) Nevertheless, we see that every locally compact subspace of such a space has (covering) dimension at most ν . For the mapping class group, this follows from [Behrstock and Minsky 2008].

Topological median algebras arise as ultralimits of coarse median algebras. We will recall the basic definitions in Section 9. Suppose that $((\Lambda_i, \rho_i, \mu_i))_{i \in \mathcal{I}}$ is sequence of coarse median spaces, where the additive constants featuring in (C1) and (C2) tend to zero and where the multiplicative constant, k , featuring in (C1) remains constant. Let $e_i \in \Lambda_i$ be a sequence of basepoints. Given a nonprincipal ultrafilter on \mathcal{I} , we can pass to an ultralimit $(\Lambda_\infty, \rho_\infty, \mu_\infty)$, which is a topological median algebra. (In fact, $(\Lambda_\infty, \rho_\infty)$ is a complete metric space.)

Theorem 2.3. *If the $(\Lambda_i, \rho_i, \mu_i)$ all have rank at most ν (with respect to these constants) then $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is a locally convex topological median algebra of rank at most ν .*

Suppose we fix a coarse median space, (Λ, ρ, μ) , of rank at most ν . We take any sequence $(t_i)_i$ of positive real numbers tending to 0, rescale the metric $\Lambda_i = \Lambda$, $\rho_i = t_i \rho$ and $\mu_i = \mu$. Fixing a base point $e \in \Lambda$, and an ultrafilter, we then get an “asymptotic” cone, $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ as above. From this, we can deduce:

Corollary 2.4. *If (Λ, ρ) is a geodesic space admitting a coarse median of rank at most ν , then (Λ, ρ) admits no quasi-isometric embedding of $\mathbb{R}^{\nu+1}$ (with the euclidean metric).*

If it did, then an asymptotic cone would contain a bilipschitz copy of $\mathbb{R}^{\nu+1}$. But this contradicts a combination of Theorems 2.2 and 2.3.

The existence, or otherwise, of a coarse median (or rank at most ν) on a geodesic space is easily seen to be quasi-isometry invariant (Lemma 8.1). This justifies the following:

Definition. We say that a finitely generated group Γ is *coarse median (of rank at most ν)*, if its Cayley graph admits a coarse median (of rank at most ν).

Thus, in view of Theorem 2.1 “coarse median of rank 1” is the same as “hyperbolic”. We observed in the Introduction that \mathbb{Z}^ν is coarse median of rank ν . We also note (Corollary 8.3) that a coarse median group has (at worst) a quadratic Dehn function.

Note that we do not assume that the median is equivariant, though in the examples we describe, it can be assumed to be equivariant up to bounded distance.

One of the main motivations is to study mapping class groups. Let Σ be a compact orientable surface of genus g and with p holes. Let $\text{Map}(\Sigma)$ be its mapping class group. Set $\xi(\Sigma) = 3g - 3 + p$ for the *complexity* of Σ . We assume that $\xi(\Sigma) > 1$, in which case, $\xi(\Sigma)$ is exactly the maximal rank of any free abelian subgroup of $\text{Map}(\Sigma)$. Making use of ideas in [Behrstock and Minsky 2011], we show:

Theorem 2.5. $\text{Map}(\Sigma)$ is a coarse median group of rank at most $\xi(\Sigma)$.

We therefore recover the fact that the mapping class group has quadratic Dehn function [Mosher 1995]. Also, applying Corollary 2.4 we recover the result of [Behrstock and Minsky 2008; Hamenstädt 2005]:

Theorem 2.6. *There is no quasi-isometric embedding of $\mathbb{R}^{\xi(\Sigma)+1}$ into $\text{Map}(\Sigma)$.*

One can show that some (in fact any) free abelian subgroup of $\text{Map}(\Sigma)$ of rank $\xi(\Sigma)$ is necessarily quasi-isometrically embedded [Farb et al. 2001]. In other words, the rank of $\text{Map}(\Sigma)$ is exactly the maximal rank of a free abelian subgroup.

In Section 12, we briefly discuss a strengthening of rank to the notion of “colourability”. We show that the mapping class group has this property.

As mentioned in the Introduction, it is shown in [Bowditch 2011a] that an asymptotic cone that arises in this way admits a bilipschitz embedding into a finite product of \mathbb{R} -trees. From this, one can deduce the rapid decay of coarse median groups. For the mapping class group, such an embedding was constructed in [Behrstock et al. 2011] and rapid decay was shown directly using medians in [Behrstock and Minsky 2011].

3. Hyperbolic spaces

In this section, we briefly describe the rank-1 case which corresponds to Gromov hyperbolicity [Gromov 1987]. This case will be used again in Sections 10 and 11.

We suppose throughout this section that (Λ, ρ) is a geodesic space.

Let us suppose first that (Λ, ρ) is K -hyperbolic for some $K \geq 0$. This means that any geodesic triangle (α, β, γ) in Λ has a K -centre, that is, some point d , with $\rho(d, \alpha) \leq K$, $\rho(d, \beta) \leq K$ and $\rho(d, \gamma) \leq K$. If $a, b, c \in \Lambda$ we take a K -centre, d , of any geodesic triangle with vertices at a, b, c , and set $\mu(a, b, c) = d$. (We can assume this to be invariant under permutation of a, b, c .) This is well defined up to bounded distance. We claim:

Lemma 3.1. *(Λ, ρ, μ) is a rank-1 coarse median space whose parameters depend only on K .*

Lemma 3.1 can be viewed as an expression of the “treelike” nature of hyperbolicity. It is a simple consequence of the following standard fact which can be found in Section 6.2 of [Gromov 1987, p. 157]. A more detailed statement and proof is given as Proposition 6.7 of [Bowditch 2006a]. It will be formulated here as Lemma 3.2, and will be used again in Section 10 (see Lemma 10.3).

Before giving the statement, we give a few definitions. Suppose that $\tau \subseteq \Lambda$ is a simplicial tree in Λ (by which we mean a subset homeomorphic to a finite simplicial tree). Given $x, y \in \tau$, we write $[x, y]_\tau$ for the arc in τ with endpoints at x and y . We write $\rho_\tau(x, y)$ for the length of $[x, y]_\tau$, which we will always assume to be finite. (Thus, ρ_τ is the induced path-metric on τ .) Clearly, $\rho(x, y) \leq \rho_\tau(x, y)$.

Definition. Given $t \geq 0$, we say that τ is t -taut if $\rho_\tau(x, y) \leq \rho(x, y) + t$ for all $x, y \in \tau$.

Lemma 3.2. *There is some function $h_0 : \mathbb{N} \rightarrow [0, \infty)$ such that if (Λ, ρ) is K -hyperbolic and $A \subseteq \Lambda$ with $|A| \leq p$, then there is a $(Kh_0(p))$ -taut simplicial tree, $\tau \subseteq \Lambda$, with $A \subseteq \tau$.*

Proof. This is essentially due to Gromov. It is a simple consequence of Proposition 6.7 of [Bowditch 2006a]. The conclusion there was stated a little differently, namely that $\rho_\tau(a, b) \leq \rho(a, b) + Kh_0(p)$ for all $a, b \in A$. To recover the statement above, first note that we can assume that every extreme (degree-1) vertex of τ is contained in A . (Otherwise, replace τ by the minimal subtree containing A .) Now, given any $x, y \in \tau$, it follows that there exist $a, b \in A$ such that $x, y \in [a, b]_\tau$. The statement that $\rho_\tau(x, y) \leq \rho(x, y) + Kh_0(p)$ is now a simple consequence of the same statement for a, b , using the triangle inequalities. \square

Note that if τ is a t -taut tree in Λ , and $x, y \in \tau$, then $[x, y]_\tau$ lies a Hausdorff distance at most s from any geodesic in Λ from x to y , where s depends only on t and K . This proven explicitly in [Bowditch 2006a], but is also an immediate consequence of the standard fact that quasigeodesics in a hyperbolic space fellow travel geodesics (where the distance bound depends only on the parameters of the quasigeodesic and the hyperbolicity constant). From this, one can easily deduce that

if $x, y, z \in \tau$, then $\mu(x, y, z)$ lies a bounded distance from the τ -median, $\mu_\tau(x, y, z)$, where the bound again depends only on t and K . In the situation described by Lemma 3.2, it therefore depends only on p and K .

We can now deduce Lemma 3.1.

Suppose that Λ is K -hyperbolic and that $A \subseteq \Lambda$ with $|A| \leq p$. Let $\tau \subseteq \Lambda$ be the tree given by Lemma 3.2. Let Π be the vertex set of τ , and let $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ be the inclusions. Property (C2) is now an immediate consequence of Lemma 3.2 and the subsequent discussion.

Finally, for (C1), it is well known (and also a consequence of Lemma 3.2) that if $a, b, c, d \in \Lambda$, then $\rho(\mu(a, b, c), \mu(a, b, d))$ is linearly bounded above in terms of $\rho(c, d)$. In fact, it is sufficient to note that if we move any one of the points a, b, c a bounded distance, say r , then the median thus defined moves a bounded distance depending only on K and r .

This proves Lemma 3.1, that is, one direction of Theorem 2.1.

For the converse, it is possible to give a constructive argument which gives an explicit constants. However, here we note that it is a consequence of the following statement proven in [Gromov 1993].

Theorem 3.3. *Let (Λ, ρ) be a geodesic space, and suppose that every asymptotic cone of (Λ, ρ) is an \mathbb{R} -tree, then (Λ, ρ) is Gromov hyperbolic.*

The notion of an asymptotic cone is due to Van den Dries and Wilkie [1984] and elaborated on in [Gromov 1993] (see Section 9 here). We will see (Theorem 2.2 and Lemma 9.6) that any asymptotic cone of a rank-1 median algebra is an \mathbb{R} -tree. From this we deduce the converse to Lemma 3.1. This then proves Theorem 2.1.

4. General median algebras

In this section we discuss some of the general theory regarding median algebras. We will elaborate on particular cases in Sections 5–7. We first describe some general terms, and then, in turn, finite, infinite and topological median algebras. Some of the basic material can be found elsewhere, though the references are somewhat scattered, and often pursued from quite different perspectives. Some general references are [Isbell 1980; Bandelt and Hedlíková 1983; Roller 1998; Chepoi 2000].

We begin with the standard formal definition, which is somewhat unintuitive. In practice, all one needs to know is that every finite subset of a median algebra is contained in a finite subalgebra (Lemma 4.2) which can be identified as the vertex set of a CAT(0) cube complex. (In fact, this could serve as an equivalent definition.)

Let M be a set. A *median* on M is a ternary operation, $\mu : M^3 \rightarrow M$, such that, for all $a, b, c, d, e \in M$,

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

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

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

The axioms are usually given in the above form, though, in fact, (M3) can be replaced by a condition on sets of four points [Kolibiar and Marcisová 1974; Bandelt and Hedlíková 1983].

We refer to (M, μ) as a *median algebra*.

Given $a, b \in M$ the *interval* $[a, b]$ between a and b is defined by $[a, b] = \{c \in M \mid \mu(a, b, c) = c\}$. Clearly $[a, a] = \{a\}$ and $[a, b] = [b, a]$. One can also verify that $[a, b] \cap [b, c] \cap [c, a] = \{\mu(a, b, c)\}$.

Definition. A (*median*) *subalgebra* of M is a subset closed under μ .

Given $A \subseteq M$, we write $\langle A \rangle$ for the subalgebra generated by A , that is, the smallest subalgebra containing A .

Definition. A subset $C \subseteq M$ is *convex* if $[a, b] \subseteq C$ for all $a, b \in C$.

Any convex subset is a subalgebra, but not necessarily conversely. One can check that any interval in M is convex.

Definition. A (*median*) *homomorphism* between median algebras is map which respects medians.

Note that a direct product of median algebras is a median algebra. Also the two-point set, $I = \{-1, 1\}$ has a unique structure as a median algebra. Given any set, X , the direct product, I^X , is naturally a median algebra.

Definition. A *hypercube* is a median algebra isomorphic to I^X for some set X . If $|X| = \nu < \infty$, we refer to it as a ν -*hypercube*. A *square* is a 2-hypercube.

If $Y \subseteq X$, then there is a natural projection epimorphism from I^X to I^Y . If $a \in I^{X \setminus Y}$, then $F = I^Y \times \{a\}$ is a convex hypercube in I^X , which we refer to as a *face* of I^X . There is a natural projection $\phi_F : I^X \rightarrow F$.

Let M be a median algebra.

Definition. A *directed wall*, W , is a pair, $(H^-(W), H^+(W))$, where $H^-(W)$ and $H^+(W)$ form a partition of M into two nonempty convex subsets. We refer to the unordered pair, $\{H^-(W), H^+(W)\}$, as an *undirected wall*, or simply a *wall*.

We write $\mathcal{W} = \mathcal{W}(M)$ for the set of all (undirected) walls in M .

Note that a directed wall, W , is equivalent to an epimorphism $\phi : M \rightarrow I$, where $H^\pm(W) = \phi^{-1}(\pm 1)$. We say that W , *separates* two subsets, $A, B \subseteq M$, if $A \subseteq H^-(W)$ and $B \subseteq H^+(W)$, or vice versa. We write $A|_W B$ to mean that A, B are separated by the wall W . We write $(A|B)$ or $(A|B)_M$ to mean that there is some $W \in \mathcal{W}$ such that $A|_W B$.

The following gets the whole subject going:

Lemma 4.1. *Any two distinct points of M are separated by a wall.*

A proof can be found in [Bandelt and Hedlíková 1983]. In fact, it can be reduced to the case of finite median algebras (cf. Lemma 6.1 here).

We note that Lemma 4.1 is equivalent to asserting that M can be embedded in a hypercube. Indeed, Lemma 4.1 tells us that the natural homomorphism from M to $I^{\mathcal{W}}$ (after arbitrarily assigning a direction to each wall) is injective.

Let S be any finite set. The *free median algebra*, $M(S)$, on S can be constructed as follows. First note that we can embed S in a hypercube Q such that the coordinate projections to I are precisely the set of all functions from S to I . Thus, Q has dimension $2^{|S|}$. Now let $M(S)$ be the subalgebra of Q generated by S . Note that S naturally embeds in $M(S)$. It has the property that any function of S to any median algebra, M , extends uniquely, to $M(S)$. Indeed, this property determines $M(S)$ uniquely up to isomorphism fixing S .

Little seems to be known about the general structure of free median algebras, though some discussion can be found in [Roller 1998]. Here we just note that $|M(S)| < 2^{2^{|S|}}$.

Suppose that M is a median algebra, and $A \subseteq M$ with $|A| \leq p$. The inclusion of A in M extends uniquely to a homomorphism of the free median algebra, $M(A)$, into M . Its image is a subalgebra of M containing A . (In fact it is precisely the subalgebra, $\langle A \rangle$, generated by A .) Thus, $|\langle A \rangle| \leq |M(A)| \leq 2^{2^p}$. We have therefore shown:

Lemma 4.2. *Suppose that $A \subseteq M$ with $|A| \leq p < \infty$, then $|\langle A \rangle| < 2^{2^p}$.*

Given $A \subseteq M$, write $G(A) = \{\mu(a, b, c) \mid a, b, c \in A\}$. Define $G^i(A)$ inductively by $G^0(A) = A$ and $G^i(A) = G(G^{i-1}(A))$. From the above, it follows that $\langle A \rangle = G^q(A)$ where $q = 2^{2^p}$.

5. Finite median algebras

We observed in Section 2 that the vertex set of a finite CAT(0) cube complex has a median algebra structure. (See, for example, [Bridson and Haefliger 1999], for a discussion of CAT(0) spaces.)

Conversely, suppose that M is a finite median algebra.

Definition. A *cube* in M is a convex subset isomorphic to a hypercube. If it has dimension $\nu < \infty$, then we refer to it as a ν -*cube*.

The set of all cubes in M gives M the structure of the vertex set, $V(\Upsilon)$, of a finite cube complex Υ . One way to view this is to embed M in the hypercube, $I^{\mathcal{W}}$, where \mathcal{W} is the set of walls of M . The complex Υ is then the full subcomplex of $I^{\mathcal{W}}$ with vertex set M . One can verify that Υ is simply connected, and that the link of every cube is a flag complex. Thus, Υ is CAT(0). Moreover, the median

structure induced by Υ (as described in Section 2) agrees with the original. We can look at this as follows. Given $a, b \in M$, let $\mathcal{W}(a, b) \subseteq \mathcal{W}$ be the set of walls separating a and b . We write $\rho_\Upsilon(a, b) = |\mathcal{W}(a, b)|$. Then, ρ_Υ is the same as the combinatorial metric on $M = V(\Upsilon)$ induced from the 1-skeleton of Υ . In fact, if α is any shortest path in the 1-skeleton from a to b , then the edges of α are in bijective correspondence with the elements of $\mathcal{W}(a, b)$ — the endpoints of each edge are separated by a unique element of $\mathcal{W}(a, b)$.

In other words, we see that $\Upsilon = \Upsilon(M)$ is canonically determined by M . We can define the “rank” of M as the dimension of $\Upsilon(M)$. Since this description is only applicable to finite median algebras, we describe some equivalent formulations below.

Let $W \in \mathcal{W}$. It’s sometimes helpful to view W geometrically as a closed totally geodesic codimension-1 subset, $\Upsilon^0(W)$, of Υ . This slices in half every cube of Υ which meets both $H^-(W)$ and $H^+(W)$. Geometrically this is closed and convex and has itself a natural structure of a cube complex (one dimension down). There is a natural nearest point retraction of Υ to $\Upsilon^0(W)$, which induces a median epimorphism. We will describe this more combinatorially later.

Suppose $W, W' \in \mathcal{W}$. There is a natural homomorphism, $\phi : M \rightarrow W \times W'$, to the square $W \times W'$.

Definition. We say that W and W' *cross* if ϕ is surjective.

In other words, each of the four sets $H^-(W) \cap H^-(W')$, $H^-(W) \cap H^+(W')$, $H^+(W) \cap H^-(W')$ and $H^+(W) \cap H^+(W')$ is nonempty. (It is also equivalent to saying that $\Upsilon^0(W) \cap \Upsilon^0(W') \neq \emptyset$.)

Lemma 5.1. *Suppose that P is a finite-dimensional hypercube, and that $A \subseteq P$ is a median subalgebra such that $\phi_F(A) = F$ for the projection ϕ_F to each square face, F . Then $A = P$.*

Proof. Suppose that $F \subseteq P$ is a square face. First note that if $A \cap F$ contains two opposite corners, a, b of F , then $F \subseteq A$. (Since, if $c \in F$, then $c = \mu(a, b, d)$ for any $d \in \phi_F^{-1}(c)$, and by assumption, $A \cap \phi_F^{-1}(c) \neq \emptyset$.) Now we proceed by induction on the dimension $v \geq 2$. Let $Q \subseteq P$ be any $(v - 1)$ -face. Applying the inductive hypothesis to $\phi_Q(A) \subseteq Q$, we see that $\phi_Q(A) = Q$. Now, by the (diagonal) observation above, we see easily that there must be some $a \in Q$ with $\phi_Q^{-1}(a) \subseteq A$. Again using the same observation, we see that if $b \in Q$ is adjacent to a (i.e., $\{a, b\}$ is a 1-face) then $\phi^{-1}(b) \subseteq Q$. Proceeding outwards from a , we eventually see that this holds for all elements of Q , and so $A = P$ as required. \square

One immediate consequence of this is the following. Suppose that $\mathcal{W}_0 \subseteq \mathcal{W}$ is a collection of pairwise crossing walls. Then the natural homomorphism, $M \rightarrow \prod \mathcal{W}_0 \cong I^{\mathcal{W}_0}$, is surjective. In other words, the sets $\bigcap_{W \in \mathcal{W}_0} H^{\epsilon(W)}(W)$ are nonempty for all functions $\epsilon : \mathcal{W}_0 \rightarrow I$.

(In terms of CAT(0) complexes, this can be interpreted as the statement that if the subspaces $\Upsilon^0(W)$ pairwise intersect, then $\bigcap_{W \in \mathcal{W}_0} \Upsilon^0(W) \neq \emptyset$.)

Suppose now that $\phi : M \rightarrow Q$ is an epimorphism of M to a hypercube, Q . (This corresponds to a collection of pairwise intersecting walls as above.) We say that a ν -cube, P , of M is *transverse* to ϕ if $\phi(P) = Q$, that is, $\phi|P$ is an isomorphism. Let $\mathcal{F} = \mathcal{F}(\phi)$ be the set of such faces, and write $F(\phi) = \bigcup \mathcal{F}(\phi)$. It's not hard to see that $F(\phi)$ is convex in M , and is isomorphic to the product $\mathcal{F}(\phi) \times Q$, where $\phi|F(\phi)$ is projection to the second factor, and where each $\{a\} \times Q$ is a transverse face. Note that the sets $\mathcal{F}(\phi) \times \{b\} \subseteq F(\phi)$ are all convex in $F(\phi)$ and so also in M . (Reinterpreting in terms of CAT(0) cube complexes, this corresponds to saying that the “walls” all intersect in a codimension ν subspace, which intrinsically has the structure of a cube complex naturally isomorphic to $\mathcal{F}(\phi)$.)

Proposition 5.2. *If $\phi : M \rightarrow Q$ is an epimorphism to a hypercube, then $\mathcal{F}(\phi) \neq \emptyset$.*

Proof. One can prove this by induction on the dimension, ν , of Q .

If $\nu = 1$, we have a single wall W . We can choose $a \in H^-(W)$ and $b \in H^+(W)$ so as to minimise $|\mathcal{W}(a, b)|$. In this case, one can verify that $\mathcal{W}(a, b) = \{W\}$, and so $\{a, b\}$ is a transverse face.

If $\nu > 1$, write $Q = P \times I$, and let $\psi : W \rightarrow I$ be the composition of ϕ with projection of Q to I . Given $a \in P$, note that $M(a) = \phi^{-1}(\{a\} \times I)$ is a convex subset of M . Now $\psi|M(a)$ is an epimorphism, so (by the case $\nu = 1$), $\mathcal{F}(\psi|M(a)) \neq \emptyset$. But $\mathcal{F}(\psi)$ is the disjoint union of the sets $\mathcal{F}(\psi|M(a))$ as a ranges over P . The natural epimorphism from $\mathcal{F}(\psi)$ to P is therefore surjective, so by induction, there must be a transverse $(\nu - 1)$ -face, say R , to this epimorphism. We see that $\bigcup R$ is now a transverse ν -cube to the original map ϕ . \square

Proposition 5.3. *Let M be a finite median algebra. The following are equivalent.*

- (1) *There is a ν -hypercube embedded in M .*
- (2) *There is an epimorphism of M to a ν -hypercube.*
- (3) *There is a set of ν pairwise crossing walls in M .*
- (4) *There is a ν -cube embedded in M .*

Proof. (1) implies (3): Let $Q \subseteq M$ be a ν -hypercube. If $\{a, b\}$ is any 1-face of Q , then any wall of M separating a and b will also separate the $(\nu - 1)$ -faces of Q containing a and b . In this way, we get a collection, \mathcal{W}_0 , of ν pairwise intersecting walls—one for each factor of Q .

(3) implies (2): As observed above, using Lemma 5.1, the map from M to the product, $\prod \mathcal{W}_0$ is surjective.

(2) implies (4): By Proposition 5.2.

(4) implies (1): Trivial. \square

Definition. We say that M has *rank at least ν* if any (hence all) the conditions of Proposition 5.3 are satisfied. We say that M has *rank ν* if it has rank at least ν but not at least $\nu + 1$.

Note that the cubes of M correspond exactly to the cubical cells of the complex $\Upsilon(M)$, so in view of (4), the definition is equivalent to that given earlier in Section 2.

Lemma 5.4. *Suppose that $A, B \subseteq M$ are disjoint nonempty convex subsets. Then there is a wall separating A and B .*

Proof. Choose $a \in A$ and $b \in B$ so as to minimise $|\mathcal{W}(a, b)|$. One can check that any $W \in \mathcal{W}$ will separate A and B . \square

In the case where $A = \{a\}$, there is unique $b \in B$ which minimises $\mathcal{W}(a, b)$. We write $\text{proj}_B(a) = b$. If $a \in B$, then we set $\text{proj}_B(a) = a$. This gives us a “nearest point” projection map $\text{proj}_B : M \rightarrow B$ to any nonempty convex subset, B , of M .

Now suppose $W \in \mathcal{W}$. We write $\mathcal{F}(W)$ for the set of transverse 1-faces. Note that $F(W) = \bigcup \mathcal{F}(W) \cong \mathcal{F}(W) \times I$. In particular, it follows that $\text{rank}(\mathcal{F}(W)) \leq \text{rank}(M) - 1$. Write $S^\pm = P \times \{\pm 1\} \subseteq H^\pm(W)$. If $a \in H^\pm(W)$, then $\text{proj}_{H^\mp(W)}(a) \in S^\mp(W)$. We set ψ_W to be the unique element of $\mathcal{F}(W)$ containing $\text{proj}_{H^\mp(W)}(a)$. This gives a map $\psi_W : M \rightarrow \mathcal{F}(W)$ which one can verify is a median epimorphism. (Geometrically, this corresponds to the nearest point projection of Υ the totally geodesic subspace $\Upsilon^0(W)$.)

Definition. The *convex hull*, $\text{hull}(A)$, of a subset $A \subseteq M$ is the smallest convex subset of M containing A .

One can verify that $a \notin A$ if and only if there is a wall of M separating a from A . We also note that if $a, b \in M$, then $\text{hull}\{a, b\} = [a, b]$.

Definition. If $A \subseteq M$, the *join*, $J(A)$, of A is defined by $J(A) = \bigcup_{a, b \in A} [a, b]$.

We define $J^i(A)$ iteratively by $J^0(A) = A$, and $J^i(A) = J(J^{i-1}(A))$. Clearly this must stabilise for some $p \in \mathbb{N}$, and we see that $\text{hull}(A) = J^p(A)$. In fact:

Lemma 5.5. *If $\text{rank}(M) \leq \nu$, and $A \subseteq M$, then $\text{hull}(A) = J^\nu(A)$.*

Proof. Clearly, $J^\nu(A) \subseteq \text{hull}(A)$. Suppose that $a \in \text{hull}(A) \setminus J^\nu(A)$. Choose $b \in A$ so as to minimise $|\mathcal{W}(a, b)|$. Choose $W \in \mathcal{W}(a, b)$ so that $a \in S^-(W)$ and $b \in H^+(W)$ (for example, corresponding to the first edge in the 1-skeleton of Υ in a shortest path from a to b). Since $a \in \text{hull}(A)$, A must meet both $H^-(W)$ and $H^+(W)$. Let $\psi_W : M \rightarrow \mathcal{F}(W)$ be the projection defined above. We see that $\bigcup \psi_W(A) \subseteq J(A)$. Now one can check (since ψ_W is an epimorphism) that $\psi_W(\text{hull}(A)) = \text{hull}(\psi_W(A))$. Now $\text{rank } \mathcal{F}(W) \leq \text{rank } M - 1 \leq \nu - 1$, so inductively, we have $\text{hull}(\psi_W(A)) = J_W^{\nu-1}(\psi_W(A))$ (where J_W denotes join in $\mathcal{F}(W)$). But $\bigcup J_W(\psi_W(A)) = J(\bigcup \psi_W(A))$, and so $\bigcup \text{hull}(\psi_W(A)) \subseteq \bigcup J_W^{\nu-1}(\psi_W(A)) = J^{\nu-1}(\bigcup \psi_W(A)) \subseteq J^{\nu-1}(J(A)) = J^\nu(A)$. Thus, $\bigcup \psi_W(\text{hull}(A)) \subseteq J^\nu(A)$. But

$a \in \text{hull}(A)$, and since $a \in S^-(W)$, we have $a \in \psi_W(a) \subseteq \psi_W(\text{hull}(A)) \subseteq J^v(A)$. Technically, this is a contradiction. In any case, we deduce that $\text{hull}(A) \subseteq J^v(A)$ as required. \square

This is all we need from Section 5 up until Section 9. We conclude this section with some observations relevant to the discussion of the mapping class group in Section 10.

Suppose that $N \subseteq M$ is a subalgebra of M . We write hull_N and J_N for the intrinsic hulls and joins in N . For future reference, we note that the following does not make any use of finiteness.

Lemma 5.6. *Suppose $A \subseteq N$, then $\text{hull}_N(A) = N \cap \text{hull}_M(A)$.*

Proof. Since $\text{hull}(A) = \bigcup_{i=0}^{\infty} J^i(A)$ and $\text{hull}_N(A) = \bigcup_{i=0}^{\infty} J_N^i(A)$, it is enough to show that $J_N^q(A) = N \cap J^q(A)$ for any q . Clearly $J_N^q(A) \subseteq J^q(A)$. Conversely, suppose that $a \in N \cap J^q(A)$. Then $a \in [b_0, b_1]$ where $b_0, b_1 \in J^{q-1}(A)$. (Here, $[\cdot, \cdot]$ denotes an interval in M .) Similarly, $b_0 \in [b_{00}, b_{01}]$, $b_1 \in [b_{10}, b_{11}]$, where $b_{00}, b_{01}, b_{10}, b_{11} \in J^{q-2}(A)$. Continuing in this way, we get points $b_w \in J^{q-j}$, where w is a word of length j in $\{0, 1\}$, so that $b_w \in [b_{w0}, b_{w1}]$. Let $B_j \subseteq J^{q-j}(A)$ be the set of such b_w . We terminate with a set $B_q \subseteq A$.

We now work backwards, to give us points $c_w \in \text{hull}_N(A)$, as follows. If w has length q , we set $c_w = b_w \in A$. If w has length less than q , we set $c_w = \mu(a, c_{w0}, c_{w1}) \in [c_{w0}, c_{w1}]_N$. By reverse induction, we end up with a point $c = \mu(a, c_0, c_1)$. We claim that $c = a$.

For suppose not. Then there is a wall $W \in \mathcal{W}(M)$ of M , with $a \in H^+(W)$ and $c \in H^-(W)$. Since $a \notin H^+(W)$, we cannot have $B_q \subseteq H^-(W)$. Thus, without loss of generality, we have $c_{0^q} = b_{0^q} \in H^+(W)$, where 0^j is the word consisting of j 0s. Working backwards, we see that $c_{0^j} \in H^+(W)$ for all j . Finally, when $j = 0$, we arrive at the contradiction that $c \in H^+(W)$.

This shows that $a = c \in \text{hull}_N(A)$. \square

Recall the notation $(A|B)_M$ to mean that subsets $A, B \subseteq M$ are separated by a wall in M . Note that, in view of Lemma 5.4 this is equivalent to saying that $\text{hull}(A) \cap \text{hull}(B) \neq \emptyset$. In fact, we note that:

Lemma 5.7. *Suppose $\text{hull}(A) \cap \text{hull}(B) \neq \emptyset$, then $\text{hull}(A) \cap \text{hull}(B) \cap \langle A \cup B \rangle \neq \emptyset$.*

Proof. Let $P(A) = \text{hull}(A) \cap \langle A \cup B \rangle$ and $P(B) = \text{hull}(B) \cap \langle A \cup B \rangle$. Suppose that $P(A) \cap P(B) = \emptyset$. Choose $a \in P(A)$ and $b \in P(B)$ so as to minimise $\rho(a, b) = |\mathcal{W}(a, b)|$. Choose any $W \in \mathcal{W}(a, b)$ with $a \in H^-(W)$ and $b \in H^+(W)$. Since $\text{hull}(A) \cap \text{hull}(B) = \emptyset$, we cannot have both $A \subseteq H^-(W)$ and $B \subseteq H^+(W)$, so without loss of generality, we can find $c \in B \cap H^-(W)$. Let $d = \mu(a, b, c)$. Since $d \in [a, b]$ we have $\rho(a, d) < \rho(a, b)$. But $d \in P(B)$, so we contradict the minimality of $\rho(a, b)$. \square

Lemma 5.8. *Let $N \subseteq M$ be a subalgebra of a finite median algebra M . If $A, B \subseteq N$, then $(A|B)_N$ if and only if $(A|B)_M$.*

Proof. Clearly $(A|B)_M$ implies $(A|B)_N$, so suppose that $(A|B)_M$ fails. By Lemma 5.7, $\text{hull}_M(A) \cap \text{hull}_M(B) \cap N \neq \emptyset$, so by Lemma 5.6, $\text{hull}_N(A) \cap \text{hull}_N(B) \neq \emptyset$, so $(A|B)_N$ fails. \square

If M, N are median algebras, then there are natural inclusions of $\mathcal{W}(M)$ and $\mathcal{W}(N)$ into $\mathcal{W}(M \times N)$ — by taking inverse images under the co-ordinate projections. In fact, under this identification, we have:

Lemma 5.9. $\mathcal{W}(M \times N) = \mathcal{W}(M) \sqcup \mathcal{W}(N)$.

Proof. This is best seen using the geometric description in terms of CAT(0) complexes. \square

This result extends to finite (and indeed infinite) direct products.

6. Infinite median algebras

We now drop the assumption that M be finite. Let \mathcal{M} be the set of all finite median subalgebras of M , which we view as a directed set under inclusion. By Proposition 5.2, \mathcal{M} is cofinal in the directed set of all finite subsets of M .

The definition of *convex*, *wall*, *crossing* etc. remain unchanged from Section 5. However, we don't have such an immediate geometrical interpretation in terms of complexes. (If M is discrete, that is, all intervals are finite, then it is again the vertex set of a CAT(0) cube complex. However, we are not assuming discreteness here.) Let \mathcal{W} be the set of walls. The following was proven in [Nieminen 1978].

Lemma 6.1. *If $A, B \subseteq M$ are disjoint convex subsets, then there is some wall, $W \in \mathcal{W}$, separating A from B .*

Proof. For finite median algebras, this was Lemma 5.4. For the general case, we use a compactness argument.

We identify the power set, \mathcal{P} , of M with the Tychonoff cube, $\{-1, 1\}^M$, of all functions from M to $\{-1, 1\}$. Here, a function, f , is identified with $f^{-1}(1)$. In particular, \mathcal{P} is compact in this topology.

Suppose that $C \subseteq \mathcal{M}$. Let $\mathcal{S}(C) \subseteq \mathcal{P}$ be the set of subsets, $C \subseteq \mathcal{P}$ with the property that $C \cap H$ and $C \setminus H$ are both convex in C and such that $C \cap H \subseteq A$ and $C \cap H \cap B = \emptyset$. In other words, $(C \cap H, C \setminus H)$ is an intrinsic wall in A which separates $C \cap A$ from $C \cap B$. By Lemma 5.4, $\mathcal{S}(C) \neq \emptyset$. Moreover, $\mathcal{S}(C)$ is closed in \mathcal{P} .

Note that if $C \subseteq D$, then $\mathcal{S}(D) \subseteq \mathcal{S}(C)$. Since \mathcal{M} is cofinal in the set of all finite subsets, it follows that $\{\mathcal{S}(C) \mid C \in \mathcal{M}\}$ has the finite intersection property. By compactness, $\bigcap_{C \in \mathcal{M}} \mathcal{S}(C) \neq \emptyset$. Let $H \in \bigcap_{C \in \mathcal{M}} \mathcal{S}(C)$.

If $a \in A$ and $b \in B$, then there is some $C \in \mathcal{M}$ with $a, b \in C$. Since $C \cap A \subseteq H$, we have $a \in H$, and since $C \cap H \cap B = \emptyset$, we have $b \notin H$. This shows that $B \subseteq H$ and $B \cap H = \emptyset$.

Also, H , and $M \setminus H$ are both convex. Suppose, for example, that $c, d \in H$, and $e \in [c, d]$ (the interval in M). Choose $C \in \mathcal{M}$ with $c, d, e \in C$. Now $[c, d] \cap A$ is an interval in C . Also, $c, d \in A \cap H$, which is convex in C . Thus, $e \in C \cap H \subseteq H$. This shows that H is convex. Similarly $M \setminus H$ is convex.

We have shown that $\{H, M \setminus H\}$ is a wall in M separating A and B . \square

In particular, any pair of distinct points of M are separated by a wall. (This shows how Lemma 4.1 can be reduced to the finite case.)

Proposition 6.2. *Let M be a median algebra. The following are equivalent.*

- (1) *There is a ν -hypercube embedded in M .*
- (2) *There is an epimorphism of M to a ν -hypercube.*
- (3) *There is a set of ν pairwise crossing walls in M .*

Proof. (1) implies (3): As in Proposition 5.3, this time using Lemma 6.1.

(3) implies (2): As in Proposition 5.3.

(2) implies (1): Let $\phi : M \rightarrow Q$ be an epimorphism to a ν -hypercube. There is some $A \in \mathcal{M}$ with $\phi(A) = Q$. By Proposition 5.2, A contains a ν -cube. This gives us a ν -hypercube in M . \square

Definition. We say that the rank of M is at least ν if any (hence all) the conditions of Proposition 6.2 hold. We say that it has rank ν if it has rank at least ν but not at least $\nu + 1$. We write $\text{rank}(M) \in \mathbb{N} \cup \{\infty\}$ for the rank of M .

Clearly the above agrees with the definition already given in the finite case. Also, using Lemma 4.2 and Proposition 6.2, we see that it is consistent with the descriptions of median algebras and rank as given in Section 2.

Let $A \subseteq M$. We define $\text{hull}(A)$, $J(A)$ and $J^i(A)$ in the same way as before. This time, $\text{hull}(A) = \bigcup_{i=1}^{\infty} J^i(A)$.

If $B \subseteq M$ is a finite median algebra, we write J_B for the intrinsic join in A , that is, $J_B(A) = B \cap J(A)$ for $A \subseteq B$. Note also that, by Lemma 5.6, $B \cap \text{hull}(A \cap B)$ is the intrinsic convex hull of $A \cap B$ in B .

Lemma 6.3. *If $A \subseteq M$, then $\text{hull}(A)$ is the union of the sets $B \cap \text{hull}(A \cap B)$ as B ranges over \mathcal{M} .*

Proof. Note that $\text{hull}(A) = \bigcup_{i=1}^{\infty} J^i(A)$. We prove inductively on i that $J^i(A) = \bigcup_{B \in \mathcal{M}} (J_B^i(A \cap B))$. First note that $J^0(A) = A = J_B^0(A)$ for any $B \in \mathcal{M}$ containing A . Suppose that $a \in J^i(A)$. Then $a \in [b, c]$ where $b, c \in J^{i-1}(A)$. By the inductive hypothesis, $b \in J_B^{i-1}(A \cap B)$ and $c \in J_C^{i-1}(A \cap C)$ for $B, C \in \mathcal{M}$. Now let $D \in \mathcal{M}$

with $\{a\} \cup B \cup C \subseteq D$. We see that $b, c \in J_D^{i-1}(A \cap D)$, so $a \in J_D(J_D^{i-1}(A \cap D)) = J_D^i(A \cap D)$. This proves the inductive statement. Now note that if $B \in \mathcal{M}$ then $J_B^i(A \cap B) \subseteq B \cap \text{hull}(A \cap B)$, proving the result. \square

Lemma 6.4. *Suppose that M has rank at most v . Then for any $A \subseteq M$, we have $\text{hull}(A) = J^v(A)$.*

Proof. If $a \in \text{hull}(A)$, then by Lemma 6.3, $a \in B \cap \text{hull}(A \cap B)$ for some $B \in \mathcal{M}$. But $B \cap \text{hull}(A \cap B)$ is the intrinsic convex hull of $A \cap B$ in B . (Indeed, in the proof of Lemma 6.3, we saw directly that $a \in \bigcup_{i=0}^{\infty} J_B^i(A \cap B)$.) Thus, by Lemma 5.5, we see that $a \in J_B^v(A \cap B) \subseteq J^v(A)$ as required. \square

Finally we note the following generalisation of Lemma 5.8 to arbitrary median algebras.

Lemma 6.5. *Let $N \subseteq M$ be a subalgebra of the median algebra M . If $A, B \subseteq N$, then $(A|B)_N$ if and only if $(A|B)_M$.*

Proof. First note that, by Lemma 5.6, for any $A \subseteq N$, we have $\text{hull}_N(A) = N \cap \text{hull}_M(A)$ (this did not make use of finiteness). We are therefore claiming that $N \cap \text{hull}_M(A) \cap \text{hull}_M(B) = \emptyset$ implies $\text{hull}_M(A) \cap \text{hull}_M(B) = \emptyset$. This was shown by Lemma 5.8, when M was finite. In the general case, suppose, for contradiction that there is some $c \in \text{hull}_M(A) \cap \text{hull}_M(B)$. It follows that $c \in \text{hull}_{\Pi}(A \cap \Pi) \cap \text{hull}_{\Pi}(B \cap \Pi)$ for some finite subalgebra, Π , of M . Now, $N \cap \Pi$ is a subalgebra of Π , and so, from the finite case, we have $N \cap \text{hull}_{\Pi}(A \cap \Pi) \cap \text{hull}_{\Pi}(A \cap \Pi) \neq \emptyset$. But this is contained in $N \cap \text{hull}_M(A) \cap \text{hull}_M(B)$, so we get a contradiction. \square

7. Topological median algebras

In this section we define the terms relevant to Theorem 2.2, and give a proof.

By a *topological median algebra* we mean a hausdorff topological space, M , together with a continuous ternary operation, $\mu : M^3 \rightarrow M$ such that (M, μ) is a median algebra.

Definition. We say that M is *locally convex* if every point has a base of convex neighbourhoods.

Put another way, if $a \in M$ and $U \ni a$ is open, then there is another open set $V \ni a$ with $\text{hull}(V) \subseteq U$.

Definition. We say that M is *weakly locally convex* if, given any $a \in M$, and any open $U \ni a$, there is an open set $V \ni a$ such that $[b, c] \subseteq U$ for all $b, c \in V$.

In other words, $J(V) \subseteq U$.

Lemma 7.1. *If M has finite rank and is weakly locally convex, then it is locally convex.*

Proof. Let $a \in U$, where $U \subseteq M$ is open. We inductively construct open sets U_i with $J^i(U_i) \subseteq U$. By Lemma 6.4 if $\nu = \text{rank}(M)$, then $\text{hull}(U_\nu) = J^\nu(U_\nu) \subseteq U$, so we can set $V = U_\nu$. \square

Given a set $C \subseteq M$, we write \bar{C} for its topological closure. The following is an elementary observation:

Lemma 7.2. *If C is convex, then so is \bar{C} .*

Suppose $W \in \mathcal{W}$. By Lemma 7.2, the closures, $\bar{H}^-(W)$ and $\bar{H}^+(W)$ are both convex. We write $L(W) = \bar{H}^-(W) \cap \bar{H}^+(W)$. It follows that $L(W)$ is also convex. Let $O^\pm(W) = M \setminus \bar{H}^\mp(W)$. Note that $O^\pm(W)$ is contained in the interior of $H^\pm(W)$.

Definition. We say that W *strongly separates* two points $a, b \in M$ if $a \in O^-(W)$ and $b \in O^+(W)$, or vice versa.

For the rest of this section, we will assume that M is locally convex.

Lemma 7.3. *Any two distinct points of M are strongly separated by a wall.*

Proof. Let $a, b \in M$ be distinct. Let $A \ni a$ and $B \ni b$ be disjoint convex neighbourhoods. By Lemma 6.1, there is a wall $W \in \mathcal{W}$ with $A \subseteq H^-(W)$ and $B \subseteq H^+(W)$. It now follows that $a \in O^-(W)$ and $b \in O^+(W)$. \square

Lemma 7.4. *Suppose that $Q \subseteq M$ is a finite dimensional hypercube, and that $\{P^-, P^+\}$ is an intrinsic wall of Q (i.e., a partition of Q into two codimension-1 faces). Then there is a wall $W \in \mathcal{W}$ with $P^- \subseteq O^-(W)$ and $P^+ \subseteq O^+(W)$.*

Proof. Choose $a \in P^-$ and $b \in P^+$ so that $\{a, b\}$ is a 1-face of Q . Let $W \in \mathcal{W}$ be a wall as given by Lemma 7.3. Suppose $c \in P^-$. Then $a \in [b, c]$. Since $\bar{H}^+(W)$ is convex, if $c \in \bar{H}^+(W)$, we would arrive at the contradiction that $a \in \bar{H}^+(W)$. It follows that $c \in O^-(W)$. Thus $P^- \subseteq O^-(W)$. Similarly, $P^+ \subseteq O^+(W)$. \square

Lemma 7.5. *If $\text{rank}(M) \leq \nu$ and $W \in \mathcal{W}$, then $\text{rank}(L(W)) \leq \nu - 1$.*

Proof. Suppose, for contradiction, that $Q \subseteq L(W)$ is a ν -hypercube. Let $a: I^\nu \rightarrow Q$ be an isomorphism. Given $\epsilon \in I^\nu$, we write $\epsilon_i \in I = \{-1, +1\}$ for the i -th coordinate. For each $i \in \{1, \dots, \nu\}$, we can partition Q as $P_i^- \sqcup P_i^+$, where P_i^- and P_i^+ correspond to $\epsilon_i = -1$ and $\epsilon_i = +1$. By Lemma 7.4, there is a wall, $W_i \in \mathcal{W}$ with $P_i^- \subseteq O^-(W_i)$ and $P_i^+ \subseteq O^+(W_i)$. Given $\epsilon \in I^\nu$, let $O(\epsilon) = \bigcap_{i=1}^\nu O^{\epsilon_i}$. Thus $O(\epsilon)$ is an open subset of M containing $a(\epsilon)$. Now $a(\epsilon) \in L(W) = \bar{H}^-(W) \cap \bar{H}^+(W)$. Thus, there are points, $a^\pm(\epsilon) \in O(\epsilon) \cap H^\pm(W)$. In particular, $a^\pm(\epsilon) \in H^{\epsilon_i}(W_i)$ for all i . It now follows that the walls, $W_1, W_2, \dots, W_\nu, W$, all pairwise intersect. We derive the contradiction that $\text{rank}(M) \geq \nu + 1$. \square

We also note that $L(W)$ is intrinsically a locally convex median algebra.

We now move on to our definition of “separation dimension”. (One can find related ideas in [Behrstock and Minsky 2008].)

Let \mathcal{D} be a collection of (homeomorphism classes) of hausdorff topological spaces. Let Θ be a hausdorff topological space. We say two points $x, y \in \Theta$ are \mathcal{D} -separated if there are closed sets, $X, Y \subseteq M$ with $x \notin Y, y \notin X, X \cup Y = M$ and $X \cap Y \in \mathcal{D}$.

Define $\mathcal{D}(n)$ inductively as follows. Set $\mathcal{D}(-1) = \{\emptyset\}$. We say $\Theta \in \mathcal{D}(n+1)$ if any two distinct points of Θ are $\mathcal{D}(n)$ -separated.

Definition. A space is has *separation dimension* n if it lies in $\mathcal{D}(n) \setminus \mathcal{D}(n-1)$.

Note that a space has separation dimension 0 if and only if it is nonempty and totally disconnected (in contrast to covering dimension [Erdős 1940]).

Suppose that $\Theta \in \mathcal{D}(n)$ and that $\Phi \subseteq \Theta$. Then $\Phi \in \mathcal{D}(n)$. This can be seen by induction on n as follows. Suppose that $x, y \in \Phi$ with $x \neq y$. There are closed sets $X, Y \subseteq \Theta$, with $x \notin Y, y \notin X, X \cup Y = \Theta$ and $X \cap Y \in \mathcal{D}(n-1)$. Inductively, $X \cap Y \cap \Phi \in \mathcal{D}(n-1)$. But $X \cap \Phi$ and $Y \cap \Phi$ are closed in $\Phi, x \notin Y \cap \Phi, y \notin X \cap \Phi$, and $(X \cap \Phi) \cup (Y \cap \Phi) = \Phi$, so x and y are $\mathcal{D}(n-1)$ -separated in Φ .

We claim that if $x, y \in \Theta \in \mathcal{D}(n)$, then there are open sets, $U \ni x$ and $V \ni y$ with $\bar{U} \cup \bar{V} = \Theta$ and $\bar{U} \cap \bar{V} \in \mathcal{D}(n-1)$. To see this, let X, Y be as in the definition of $\mathcal{D}(n)$. Let $U = \Theta \setminus Y$ and $V = \Theta \setminus X$. Now $U \subseteq X$, so $\bar{U} \subseteq X$. Thus, $\Theta \setminus X \subseteq \Theta \setminus \bar{U} = V$. Similarly, $\Theta \setminus \bar{V} \subseteq U$. In particular, $x \in U$ and $y \in V$. Also $\bar{U} \cup \bar{V} = \Theta$. We similarly have $\bar{V} \subseteq Y$, and so $\bar{U} \cap \bar{V} \subseteq X \cap Y \in \mathcal{D}(n-1)$. Thus, by the preceding paragraph, we have $\bar{U} \cap \bar{V} \in \mathcal{D}(n-1)$, thereby proving the claim.

Conversely, if U, V are as above, then \bar{U} and \bar{V} are as in the inductive definition of $\mathcal{D}(n)$. This therefore gives rise to an equivalent formulation of separation dimension.

Finally, putting together Lemmas 7.3 and 7.5, we see by induction on n that if $\text{rank}(M) \leq n$, then M has separation dimension at most n , thereby proving Theorem 2.2.

The usual notion of inductive dimension is similar—replacing separation of points with separation of disjoint closed sets. These notions are equivalent for locally compact spaces (see for example Section III(6) of [Hurewicz and Wallman 1941]). In particular, we note:

Lemma 7.6. *If Θ is a hausdorff topological space of separation dimension at most v , then every locally compact subset has (covering) dimension at most v .*

In particular, such a space does not admit any continuous injective map of \mathbb{R}^{v+1} .

We note that the conclusion of Lemma 7.6 suggests another notion of dimension for a topological space, namely the maximal dimension of a locally compact subspace. Indeed this was the notion that was used in [Behrstock and Minsky 2008].

8. Coarse median spaces

We establish some basic facts about coarse median spaces. We show that such a space satisfies certain quadratic isoperimetric inequality (Proposition 8.2).

Let (Λ, ρ) be a geodesic space. (A path-metric space would be sufficient.) Suppose that $\mu : \Lambda^3 \rightarrow \Lambda$ is (for the moment) any ternary operation on Λ .

Definition. If (Π, μ_Π) is a median algebra then a h -quasimorphism of Π into Λ is a map $\lambda : \Pi \rightarrow \Lambda$ satisfying

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

for all $x, y, z \in \Pi$.

Definition. We say that (Λ, ρ, μ) is a coarse median space if it satisfies:

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

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

(C2) There is a function $h : \mathbb{N} \rightarrow [0, \infty)$ such that $1 \leq |A| \leq p < \infty$, then there is a finite median algebra and a $h(p)$ -quasimorphism, $\lambda : \Pi \rightarrow \Lambda$ such that for all $a \in A$, $\rho(a, \lambda\pi a) \leq h(p)$.

We therefore have one multiplicative constant, k , and a sequence, $h(p)$, of additive constants. We can assume that $h(p)$ is increasing in p .

In (C2), we note that we can always assume that $\Pi = \langle \pi(A) \rangle$, so by Lemma 4.2, $|\Pi| \leq 2^{2^p}$. In particular, we can take Π to be finite. Our definition therefore agrees with that given in Section 2.

Remark. Note that in defining a coarse median space, there would be no loss in taking $\Pi = M(A)$ to be the free median algebra on A (since this will admit an epimorphism to any such Π). Also in (C2), there would be no loss in assuming that $\lambda\pi a = a$ for all $a \in A$. However, when we define a “coarse median space of rank ν ” below, we can no longer assume these things.

Definition. If we can always take Π to have rank at most ν , then we say that (Λ, ρ, μ) has *rank* at most ν .

Here, of course, the function h is fixed independently of ν .

Lemma 8.1. *Suppose that (Λ, ρ) and (Λ', ρ') are quasi-isometric geodesic spaces. Then (Λ, ρ) admits a coarse median (of rank ν) if and only if (Λ', ρ') does.*

Proof. Let $f : \Lambda \rightarrow \Lambda'$ and $g : \Lambda' \rightarrow \Lambda$ be quasi-inverse quasi-isometries. (That is, $f \circ g$ and $g \circ f$ are each a bounded distance from the respective identity maps) We define μ' on Λ' by setting $\mu'(a, b, c) = f\mu(ga, gb, gc)$. \square

Definition. A finitely generated group Γ is *coarse median* (of rank ν) if and only if its Cayley graph with respect to any finite generating set admits a coarse median.

Any two such Cayley graphs are quasi-isometric, so this is well defined by Lemma 8.1.

Returning to Λ , suppose $a, b, c \in \Gamma$. Let $A = \{a, b, c\}$, and let $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ be as in (C2). From the second part of (C2), we see that $\rho(a, \lambda\pi a)$, $\rho(b, \lambda\pi b)$ and $\rho(c, \lambda\pi c)$ are all bounded above by $h(3)$. Applying (C1), it follows that

$$\rho(\mu(a, b, c), \mu(\lambda\pi a, \lambda\pi b, \lambda\pi c)) \leq 3kh(3) + h(0).$$

Also from the first part of (C2),

$$\rho(\mu(\lambda\pi a, \lambda\pi b, \lambda\pi c), \lambda\mu_{\Pi}(\pi a, \pi b, \pi c)) \leq h(3),$$

and so

$$\rho(\mu(a, b, c), \lambda\mu_{\Pi}(\pi a, \pi b, \pi c)) \leq (3k + 1)h(3) + h(0).$$

The same holds for any permutation of a, b, c , and since μ_{Π} is invariant under such permutation, we deduce

$$\begin{aligned} \rho(\mu(a, b, c), \mu(b, c, a)) &\leq (6k + 2)h(3) + 2h(0), \\ \rho(\mu(a, b, c), \mu(b, a, c)) &\leq (6k + 2)h(3) + 2h(0). \end{aligned}$$

Since $\mu_{\Pi}(\pi a, \pi a, \pi b) = \pi a$, a similar argument gives

$$\rho(\mu(a, a, b), a) \leq (3k + 2)h(3) + h(0).$$

In view of this, there is no essential loss in assuming (M1) and (M2), namely, $\mu(a, b, c) = \mu(b, c, a) = \mu(b, a, c)$ and $\mu(a, a, b) = a$. We have already implicitly used this in Section 3.

Given this, we note that (C1) could be replaced by the assumption that

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

is uniformly bounded above in terms of $\rho(c, d)$. Given that (Λ, ρ) is a geodesic space, it is easy to see that such a bound can always be taken to be linear.

Next, we discuss the quadratic isoperimetric inequality. Suppose, $l, L > 0$.

Definition. An l -cycle is a cyclically ordered sequence of points, $a_0, a_1, \dots, a_p = a_0$ in Λ , with $\rho(a_i, a_{i+1}) \leq l$ for all i .

Definition. An L -disc consists of a triangulation of the disc, together with a map $b : V \rightarrow \Lambda$ of the vertex set, V , into Λ such that $\rho(b(x), b(y)) \leq L$ whenever $x, y \in V$ are adjacent in the 1-skeleton.

Definition. We say that b spans an l -cycle, $(a_i)_i$ if we can label the vertices on the boundary as x_i such that x_{i+1} is adjacent to x_i and with $a_i = b(x_i)$ for all i .

Proposition 8.2. *Suppose that Λ is a coarse median space. Given any $l > 0$, there is some $L > 0$, depending only on l and the parameters such that for any $p \in \mathbb{N}$, any l -cycle of length at most p bounds an L -disc with at most p^2 2-simplices.*

In fact, all we require of μ is (M1) and (M2) and the statement that

$$\rho(\mu(a, b, c), \mu(a, b, d)) \leq L/2$$

whenever $a, b, c, d \in \Lambda$ with $\rho(c, d) \leq l$.

To see this, we construct a triangulation of the disc as follows. Let

$$V = \{\{0\}\} \cup \{\{i, j\} \mid 1 \leq i, j \leq p-1\}.$$

We define the edge set by deeming $\{i, j\}$ to be adjacent to $\{i+1, j\}$ and to $\{i+1, j+1\}$ for all $1 \leq i, j \leq p-2$, and deeming $\{0\}$ to be adjacent to $\{1, i\}$ and to $\{p-1, i\}$ for all $1 \leq i \leq p-1$. Note that $\{i, i\} = \{i\}$, so $\{i\}$ is adjacent to $\{i+1\}$ for all $0 \leq i \leq p-2$, and $\{p-1\}$ is adjacent to $\{0\}$. Filling in every 3-cycle with a 2-simplex, we can see that this defines a triangulation of the disc whose boundary is the circuit with vertices $(\{i\})_i$. In total, it has $\frac{1}{2}(p^2 - p + 2)$ vertices, $\frac{p}{2}(3p - 5)$ edges and $p^2 - 2p$ triangles.

(We can realise this in the euclidean plane, \mathbb{R}^2 , as follows. We make the identification $V \subseteq \mathbb{Z}^2 \subseteq \mathbb{R}^2$, by identifying $\{i, j\}$ with the ordered pair, (i, j) , for $1 \leq j \leq i \leq p-1$, and identifying $\{0\}$ with $(p, 0)$. We can triangulate the convex hull of $\{(1, 1), (p-1, p-1), (p-1, 1)\}$ by cutting along straight lines with slope 0, 1, and ∞ through the integer lattice points. We then connect $(p, 0)$ by a geodesic segment to each of the points $(i, 1)$ and $(p-1, i)$ for $1 \leq i \leq p-1$. This gives us a triangulation of the convex hull, Δ , of $\{(1, 1), (p-1, p-1), (p, 0)\}$, with vertices $V \cong \mathbb{Z}^2 \cap \Delta$. Note that $V \cap \partial \Delta \cong \{(p, 0)\} \cup \{(i, i) \mid 1 \leq i \leq p-1\} \cong \{\{i\} \mid 0 \leq i \leq p-1\}$.)

Now suppose that $a_0, a_1, \dots, a_p = a_0$ is an l -cycle in Λ . Define $b : V \rightarrow \Lambda$ by $b(\{i, j\}) = \mu(a_0, a_i, a_j)$ thus, $b(\{i\}) = a_i$ for all i . Now, if $\{i', j'\}$ is adjacent to $\{i, j\}$, then $|i - i'| \leq 1$ and $|j - j'| \leq 1$, and so $\rho(b(\{i, j\}), b(\{i', j'\})) \leq 2(L/2) = L$.

This proves Proposition 8.2.

Note that, if Λ is the Cayley graph of a finitely generated group, then this implies that Γ is finitely presented, and that the Dehn function for any finite presentation is at most quadratic. In other words:

Corollary 8.3. *Any coarse median group is finitely presented, and has Dehn function that is at most quadratic.*

The following observations will be needed in the next section.

Lemma 8.4. *Suppose that Π is a finite median algebra generated by $B \subseteq \Pi$, with $|B| \leq p$. Suppose that $\lambda : \Pi \rightarrow \Lambda$ is a h -quasimorphism. then $\text{diam}(\lambda \Pi) \leq K_0(\text{diam}(\lambda B) + h(0) + h(p))$, where the constant, K_0 , depends only on k (the multiplicative constant of (C1)) and p .*

Proof. Given $C \subseteq \Pi$, let $G(C) = \{\mu(x, y, z) \mid x, y, z \in C\}$. Let $G^i(C)$ be the i -th iterate of G . Set $q = 2^{2^p}$. By Lemma 4.2, $|\Pi| \leq q$, so $\Pi = C^q(B)$.

Now suppose $x, y, z \in \Pi$ and set $w = \mu_\Pi(x, y, z)$. Now $\mu_\Pi(x, x, y) = x$, and so $\rho(\mu_\Pi(\lambda x, \lambda x, \lambda y), \lambda x) \leq h$. Also

$$\rho(\mu(\lambda x, \lambda y, \lambda z), \mu(\lambda x, \lambda x, \lambda y)) \leq k\rho(x, y) + h(0)$$

and $\rho(\lambda w, \mu(\lambda x, \lambda y, \lambda z)) \leq h$. Thus, $\rho(\lambda x, \lambda w) \leq k\rho(x, y) + h(0) + 2h$. It follows that if $C \subseteq \Pi$, then $\text{diam}(\lambda G(C)) \leq k \text{diam}(\lambda C) + h(0) + 2h$.

Now iterating this q times, starting with $B \subseteq \Pi$, we obtain $\text{diam}(\lambda \Pi) \leq K_0(\text{diam}(\lambda B) + h(0) + h)$ where $K_0 = k^q$. \square

Lemma 8.5. *Suppose that $A \subseteq \Lambda$ with $1 \leq |A| \leq p < \infty$ and that $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ are as in (C2), with $\Pi = \langle \pi A \rangle$. Then*

$$\text{diam}(\lambda \Pi) \leq K(\text{diam}(A) + h(0) + h(p)),$$

where K depends only on k and p .

Proof. By Lemma 8.4, we have $\text{diam}(\lambda \Pi) \leq K_0(\text{diam}(\pi A) + h(0) + h(p))$. But if $a \in A$, then $\rho(a, \lambda \pi a) \leq h(p)$, so $\text{diam}(\lambda \pi A) \leq \text{diam}(A) + 2h(p)$, and the result follows. \square

9. Ultralimits

In this section we discuss ultralimits of coarse median spaces. When the ultralimit is obtained through a sequence of rescalings of a given space, we will refer to the resulting space as an ‘‘asymptotic cone’’. Asymptotic cones of groups and metric spaces were introduced by Van den Dries and Wilkie [1984] and elaborated upon by Gromov [1993]. They now play a major role in geometric group theory. We will show that the asymptotic cone of a coarse median space of rank at most ν is a locally convex topological median algebra of rank at most ν . (This was stated as Theorem 2.3.)

First, we give a general discussion. We fix an indexing set, \mathcal{J} , with a nonprincipal ultrafilter. Throughout this section, if $(t_i)_{i \in \mathcal{J}}$ is a sequence of real numbers, we will write $t_i \rightarrow t$ to mean that t_i tends to t with respect to this ultrafilter. We refer to a sequence as *bounded* if it is bounded with respect to the ultrafilter (i.e., there is some $K \geq 0$ so that the set of indices, $i \in \mathcal{J}$ for which $|t_i| \leq K$ lies in the ultrafilter). Note that any bounded sequence has a unique limit. We recall the following (e.g., [Gromov 1993]). Let $((\Lambda_i, \rho_i))_{i \in \mathcal{J}}$ be a collection of metric spaces indexed by \mathcal{J} . We will write $\mathbf{a} = (a_i)_i \in \prod_i \Lambda_i$ for a typical sequence of elements. We fix some *basepoint* $\mathbf{e} = (e_i)_i \in \prod_i \Lambda_i$. Let \mathcal{B} be the set of sequences \mathbf{a} in $\prod_i \Lambda_i$ such that $\rho_i(e_i, a_i)$ is bounded (in the above sense). Given $\mathbf{a}, \mathbf{b} \in \mathcal{B}$, write $\mathbf{a} \sim \mathbf{b}$ to mean that $\rho_i(a_i, b_i)$ is bounded. This is an equivalence relation, and we write $\Lambda_\infty = \mathcal{B}/\sim$.

Given $\mathbf{a} \in \mathcal{B}$, and $a \in \Lambda_\infty$, we write $a_i \rightarrow a$ to mean that a is the equivalence class of \mathbf{a} . Given $a, b \in \Lambda_\infty$, choose any $\mathbf{a}, \mathbf{b} \in \mathcal{B}$ with $a_i \rightarrow a$ and $b_i \rightarrow b$. Now $\rho_i(a_i, b_i)$ is bounded and we define $\rho_\infty(a, b)$ to be the limit of $\rho_i(a_i, b_i)$. One can easily check that this is well defined, and that ρ_∞ is a metric on Λ_∞ . With a bit more work, one can see that $(\Lambda_\infty, \rho_\infty)$ is complete.

Now suppose that $((\Lambda_i, \rho_i, \mu_i))_{i \in \mathcal{I}}$ is a sequence of coarse median spaces. We write k_i and h_i for the constants featuring in (C1) and (C2). We suppose:

(U1) k_i is bounded, and $h_i(p) \rightarrow 0$ for all $p \in \mathbb{N}$.

We may as well fix $k_i = k$.

Also, we will suppose that the spaces also satisfy properties (M1) and (M2) of a median algebra (that is, with no additive constant). As discussed earlier, there is no essential loss of generality in doing this.

Now suppose that $a, b, c \in \Lambda_\infty$. Choose $a_i \rightarrow a, b_i \rightarrow b$ and $c_i \rightarrow c$. Now

$$\rho_i(e_i, \mu_i(a_i, b_i, c_i)) \leq k(\rho_i(e_i, a_i) + \rho_i(e_i, b_i) + \rho_i(e_i, c_i)) + h_i(0),$$

so $\rho_i(e_i, \mu_i(a_i, b_i, c_i))$ is bounded. Moreover, if $a'_i \rightarrow a, b'_i \rightarrow b$ and $c'_i \rightarrow c$, is another such sequence, then

$$\rho_i(\mu_i(a_i, b_i, c_i), \mu_i(a'_i, b'_i, c'_i)) \leq k(\rho_i(a_i, a'_i) + \rho_i(b_i, b'_i) + \rho_i(c_i, c'_i)) + h_i(0),$$

so $\rho_i(\mu_i(a_i, b_i, c_i), \mu_i(a'_i, b'_i, c'_i)) \rightarrow 0$. It follows that the limit of $\mu_i(a_i, b_i, c_i)$ in Λ_∞ is well defined, and we write it as $\mu_\infty(a, b, c)$.

Now the metric ρ_∞ defines a topology in Λ_∞ . With respect to this topology, we claim:

Proposition 9.1. $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is a topological median algebra.

Proof. For this, we only need to consider a finite subset $A \subseteq \Lambda_\infty$. (In view of fact that the median axioms only require sets of four points we could restrict to the case where $|A| \leq 4$ here, and hence only require that $h_i(4) \rightarrow 0$. We will however need sets of arbitrary finite cardinality later, when we need to bound the rank.)

Let $A \subseteq \Lambda_\infty$ be finite, and set $p = |A|$. We define maps $f_i : A \rightarrow \Lambda_i$ by choosing a sequence $a_i \rightarrow a$ for all $a \in \Lambda_\infty$, and setting $f_i(a) = a_i$. We write $A_i = f_i(A) \subseteq \Lambda_i$. Thus $|A_i| \leq p$. Let $\pi_i : A_i \rightarrow \Pi_i$ and $\lambda_i : \Pi_i \rightarrow \Lambda_i$ be as in (C2). Thus λ_i is an $h_i(p)$ -quasimorphism, and we can assume that $\Pi_i = \langle \pi_i A_i \rangle$, so that $|\Pi_i| \leq 2^{2^p}$. There are only finitely many possibilities for the median algebra (Π_i, μ_{Π_i}) up to isomorphism, so we can assume that $\Pi_i = \Pi$ is fixed. We can now also assume that the compositions $\pi_i f_i : A \rightarrow \Pi$ are all equal to some fixed map $\pi : A \rightarrow \Pi$. Note again that $\Pi = \langle \pi A \rangle$.

Now $\text{diam}(A_i)$ is bounded. By Lemma 8.5,

$$\text{diam}(\lambda_i \Pi) \leq K(\text{diam}(A_i) + h_i(0) + h_i(p))$$

is also bounded. (Here K depends only on k and p and is therefore constant.) If $a \in A$, recall that $a_i = f_i(a) \rightarrow a$. Also $\rho_i(a_i, \lambda_i \pi_i a_i) \leq h_i(p) \rightarrow 0$, so $\lambda_i \pi_i a_i \rightarrow a$. Now if $x \in \Pi$, then $\rho_i(a_i, \lambda_i x)$ is bounded, by the above. So $\rho_i(e_i, \lambda_i x)$ is bounded, and so $\lambda_i x \rightarrow b$ for some $b \in \Lambda_\infty$. This gives us a well defined map $\lambda : \Pi \rightarrow \Lambda_\infty$, with $\lambda_i x \rightarrow \lambda x$.

Now $\Lambda_i : \Pi \rightarrow \Lambda_i$ is a $h_i(p)$ -quasimorphism where $h_i(p) \rightarrow 0$, so it follows that $\lambda : \Pi \rightarrow \Lambda$ is a homomorphism; that is, for all $x, y, z \in \Pi$, $\lambda \mu_\Pi(x, y, z) = \mu_\infty(\lambda x, \lambda y, \lambda z)$. Moreover, if $a \in A$, we have seen that $\lambda_i \pi a = \lambda_i \pi_i f_i a = \lambda_i \pi_i a_i \rightarrow a$. By definition of λ , we have $\lambda_i \pi a \rightarrow \lambda \pi a$, and so $\lambda \pi a = a$. Setting $B = \lambda \Pi$ we have $A \subseteq B$.

Now λ is a homomorphism, so it follows easily that B is closed under μ_∞ . Also, since Π is a median algebra, it follows easily that (B, μ_∞) is intrinsically a median algebra.

In summary, we have shown that any finite subset, $A \subseteq \Lambda_\infty$, is contained in another finite subset $B \subseteq \Lambda_\infty$ that is closed under μ_∞ and intrinsically a median algebra. It follows that $(\Lambda_\infty, \mu_\infty)$ is a median algebra. Note in particular, that $\mu_\infty(a, b, c)$ is invariant under permuting a, b, c .

Suppose that $a, b, c, d \in \Lambda_\infty$. Let $a_i \rightarrow a, b_i \rightarrow b, c_i \rightarrow c$ and $d_i \rightarrow d$. Then

$$\rho_i(\mu_i(a_i, b_i, c_i), \mu_i(a_i, b_i, d_i)) \leq k\rho_i(c_i, d_i) + h_i(0),$$

and so

$$\rho_\infty(\mu_\infty(a, b, c), \mu_\infty(a, b, d)) \leq k\rho_\infty(c, d).$$

We see that $\mu_\infty : \Lambda_\infty^3 \rightarrow \Lambda_\infty$ is continuous. In other words, $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is a topological median algebra. \square

In fact, we can say more. Suppose $a, b, c \in \Lambda_\infty$ with $c \in [a, b]$. Now $\rho_\infty(a, c) \leq \rho_\infty(\mu_\infty(a, a, c), \mu_\infty(a, b, c)) \leq k\rho_\infty(a, b)$. Therefore, $\text{diam}([a, b]) \leq k\rho_\infty(a, b)$. We deduce:

Lemma 9.2. $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is weakly locally convex.

Note that the conclusion of Lemma 9.2 is a consequence of the fact that

$$\rho_\infty(\mu_\infty(a, b, c), \mu_\infty(a, b, d)) \leq k\rho_\infty(c, d)$$

for all $a, b, c, d \in \Lambda_\infty$. This is a key property used in the embedding theorem in [Bowditch 2011a].

Suppose now that each $(\Lambda_i, \rho_i, \mu_i)$ is coarse median of rank at most ν . We now interpret property (U1) above to mean that the constants k_i and $h_i(p)$ of (C2) refer to median algebras Π_i of rank at most ν .

Following the proof of Proposition 9.1, we see that Π has rank at most ν . It follows that $B = \lambda \Pi$ also has rank at most ν (using, for example, condition (2) of Proposition 6.2. We deduce:

Proposition 9.3. *If the spaces $(\Lambda_i, \rho_i, \mu_i)$ all have rank at most v and satisfy (U1), then $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ has rank at most v .*

Putting together these results with Lemma 7.1, we deduce that $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is locally convex.

This proves Theorem 2.3.

Now suppose that (Λ, ρ, μ) is a coarse median space. Let $\mathcal{F} = \mathbb{N}$ with any nonprincipal ultrafilter. Let t_i be any sequence of positive numbers with $t_i \rightarrow 0$ (with respect to the ultrafilter is enough). Let $\Lambda_i = \Lambda$, $\rho_i = t_i \rho$ and $\mu_i = \mu$. Let $e \in \Lambda$, and set $e_i = e$ for all i to give us a fixed basepoint. The sequence $(\Lambda_i, \rho_i, \mu_i)$ satisfies the condition of Proposition 9.1, and so we get a topological median algebra $(\Lambda_\infty, \rho_\infty, \mu_\infty)$.

Definition. We refer to a topological median algebra arising in this way as an *asymptotic cone* of (Λ, ρ, μ) .

Thus $(\Lambda_\infty, \rho_\infty)$ is an asymptotic cone in the traditional sense. The following is an immediate consequence of the above:

Proposition 9.4. *If (Λ, ρ, μ) has rank at most v , then any asymptotic cone is locally convex and has rank at most v .*

We can now deduce Corollary 2.4 as explained in Section 2.

Finally, we note:

Lemma 9.5. *Any geodesic space which admits a structure as a rank-1 topological median algebra is an \mathbb{R} -tree.*

Proof. We see that any pair of distinct points are separated by a rank-0 subalgebra, in other words, a point. This implies that a geodesic connecting any pair of point must in fact be the unique arc connecting those points. In other words, any two points are connected by a unique arc which is isometric to a real interval. This is one of the standard definitions of an \mathbb{R} -tree. \square

Using Proposition 9.4, we deduce:

Lemma 9.6. *Let (Λ, ρ) be a geodesic space which admits a rank-1 coarse median. Then any asymptotic cone of (Λ, ρ) is an \mathbb{R} -tree.*

This now gives us what we need to complete the proof of Theorem 2.1 as explained in Section 3.

10. Projection maps

In this section, we explain how the existence of certain projection maps imply that a given ternary operation on a geodesic space is a coarse median. We first give the constructions in a formal manner. The main application we have in mind is to the mapping class group, as we explain in Section 11.

Let (Λ, ρ) be a geodesic space, and let $\mu : \Lambda^3 \rightarrow \Lambda$ be a ternary operation. Let \mathcal{X} be an indexing set, and suppose that to each $X \in \mathcal{X}$, we have associated a uniformly coarse median space $(\Theta(X), \sigma_X, \mu_X)$, together with a uniformly lipschitz quasimorphism, $\theta_X : \Lambda \rightarrow \Theta(X)$. Here, “uniform” means that the various parameters are independent of X . In particular, we are assuming that $\theta_X : (\Lambda, \rho) \rightarrow (\Theta(X), \sigma_X)$ is k_0 -lipschitz, and that $\theta_X : (\Lambda, \mu) \rightarrow (\Theta(X), \mu_X)$ is a h_0 -quasimorphism for fixed k_0 and h_0 .

We also assume:

(P1) For all l there is some l' such that if $a, b \in \Lambda$ satisfy $\sigma_X(\theta_X a, \theta_X b) \leq l$ for all $X \in \mathcal{X}$, then $\rho(a, b) \leq l'$.

Proposition 10.1. *A ternary operation μ satisfying the above is a coarse median on (Λ, ρ) . (In fact, we will see that the parameters of (Λ, ρ, μ) depend only on those arising in the hypotheses.)*

Before giving the proof, we note how the hypotheses arise in nature. In Section 11, Λ will be the “marking complex” of a compact surface Σ . This is quasi-isometric to the mapping class group of Σ . The map μ will be the “centroid” map defined in [Behrstock and Minsky 2011]. The set \mathcal{X} is the set of homotopy classes of essential subsurfaces of Σ . In this, we include annuli and Σ itself, but do not allow three-holed spheres. For a nonannular surface, the space $(\Theta(X), \sigma_X)$ will be the curve graph of X , which is hyperbolic by [Masur and Minsky 1999], and hence is coarse median of rank 1. If X is an annulus, then $(\Theta(X), \sigma_X)$ is a certain arc complex, which is quasi-isometric to the real line. In all cases, the maps $\theta_X : \Lambda \rightarrow \Theta(X)$ arises from the subsurface projection map described in [Masur and Minsky 2000]. The property (P1) can be shown using the distance formula used in the same reference. A consequence of Proposition 10.1, is that the mapping class group is coarse median. We recover the fact that it is finitely presented and has a quadratic Dehn function [Mosher 1995].

Proof of Proposition 10.1. We need to verify (C1) and (C2).

(C1) Let $a, b, c, a', b', c' \in \Lambda$, and write $e = \mu(a, b, c)$, $f = \mu(a', b', c')$. Let $X \in \mathcal{X}$. Write $t = \sigma(a, a') + \sigma(b, b') + \sigma(c, c')$. Since θ_X is a quasimorphism, we have

$$\begin{aligned} \sigma_X(\theta_X e, \mu_X(\theta_X a, \theta_X b, \theta_X c)) &\leq h_0, \\ \sigma_X(\theta_X f, \mu_X(\theta_X a', \theta_X b', \theta_X c')) &\leq h_0. \end{aligned}$$

Since μ_X satisfies (C1) and θ_X is k_0 -lipschitz, we have

$$\begin{aligned} \sigma_X(\mu_X(\theta_X a, \theta_X b, \theta_X c), \mu_X(\theta_X a', \theta_X b', \theta_X c')) \\ \leq k(\sigma_X(\theta_X a, \theta_X a') + \sigma_X(\theta_X b, \theta_X b') + \sigma_X(\theta_X c, \theta_X c')) \\ \leq k k_0(\rho(a, a') + \rho(b, b') + \rho(c, c')) = k k_0 t. \end{aligned}$$

Thus, $\sigma_X(\theta_X e, \theta_X f)$ is (linearly) bounded above in term of t .

Now since this holds uniformly for all $X \in \mathcal{X}$, it follows by (P1) that $\rho(e, f)$ is bounded above in terms of t . Since (Λ, ρ) is a geodesic space, this is sufficient to verify (C1) for μ (as observed in Section 8).

(C2): Let $A \subseteq \Lambda$, with $|A| \leq p < \infty$. Let $q = 2^{2^p}$. Let Π be the free median algebra on A , and write $\pi : A \rightarrow \Pi$ for the inclusion map. Note that $\Pi = \langle \pi A \rangle$, and recall from Section 4, that $\Pi = G^i$, where $G^i = G^i(\pi A)$ is defined by iterating the median operation, μ_Π .

We define $\lambda : \Pi \rightarrow \Lambda$ inductively as follows. Given $x \in G^0 = \pi A$, set $\lambda x = a$, where $x = \pi a$. Given $u \in G^{i+1} \setminus G^i$, choose any $x, y, z \in G^i$ with $u = \mu_\Pi(x, y, z)$ and set $\lambda u = \mu(\lambda x, \lambda y, \lambda z)$. By construction, we have $\lambda \pi a = a$ for all $a \in A$. We want to show that λ is a quasimorphism.

Let $X \in \mathcal{X}$. We have a quasimorphism $\theta_X : \Lambda \rightarrow \Theta(X)$. There is also a quasimorphism, $\omega_X : \Pi \rightarrow \Theta(X)$ such that $\omega_X \pi a = \theta_X a$ for all $a \in A$. (Certainly, such a quasimorphism exists from some median algebra to $\Theta(X)$, by (C2) applied to $\Theta(X)$. But since we have taken Π to be free on A , we can precompose this with a homomorphism from Π to the given median algebra which fixes A . Thus, we can take the domain to be Π .) By assumption, the additive constants depend only on the parameters and on p . In particular, they are independent of X .

In what follows, it will be convenient to adopt the following convention. Given points, x, y in a metric space (namely Λ or $\Theta(X)$), we will write $x \sim y$ to mean that, at any particular stage in the argument, the distance between x and y is bounded above by some explicit constant, depending only on the parameters and on p . The bound may increase as the argument proceeds, though we won't keep track of it explicitly here.

We first claim that $\theta_X \lambda x \sim \omega_X x$ for all $x \in \Pi$. We show this by induction on i , where $x \in G^{i+1} \setminus G^i$. Note first that if $x \in G^0$, then setting $x = \pi a$, we have $\theta_X \lambda x = \theta_X a = \omega_X \pi a = \omega_X x$.

Now suppose that $u \in G^{i+1} \setminus G^i$. Let $x, y, z \in G^i$ be the three points that were chosen in the definition of λ , so that $\lambda u = \mu(\lambda x, \lambda y, \lambda z)$. We now have

$$\begin{aligned} \theta_X \lambda u &= \theta_X \mu(\lambda x, \lambda y, \lambda z) \\ &\sim \mu_X(\theta_X \lambda x, \theta_X \lambda y, \theta_X \lambda z) \\ &\sim \mu_X(\omega_X x, \omega_X y, \omega_X z) \\ &\sim \omega_X u. \end{aligned}$$

(The above follow respectively from the fact that $\theta_X : \Lambda \rightarrow \Theta(X)$ is a quasimorphism; the inductive hypothesis; and the fact that $\omega_X : \Pi \rightarrow \Theta(X)$ is a quasimorphism.) This proves that $\theta_X \lambda x \sim \omega_X x$ for all $x \in \Pi = G^q$.

Now suppose that $x, y, z \in \Pi$ are any three points. We have

$$\begin{aligned} \theta_X \lambda \mu_\Pi(x, y, z) &\sim \omega_X \mu_\Pi(x, y, z) \\ &\sim \mu_X(\omega_X x, \omega_X y, \omega_X z) \\ &\sim \mu_X(\theta_X \lambda x, \theta_X \lambda y, \theta_X \lambda z) \\ &\sim \theta_X \mu(\lambda x, \lambda y, \lambda z). \end{aligned}$$

(These relations follow respectively from the claim already proven above; the fact that $\omega_X : \Pi \rightarrow \Theta(X)$ is a quasimorphism; the claim again, together with property (C1) applied to $(\Theta(X), \mu_X)$; and the fact that $\theta_X : \Lambda \rightarrow \Theta(X)$ is a quasimorphism.)

In other words, we have shown that

$$\theta_X \lambda \mu_\Pi(x, y, z) \sim \theta_X \mu(\lambda x, \lambda y, \lambda z)$$

for all $X \in \mathcal{X}$, and for all $x, y, z \in \Pi$. Applying (P1), we get

$$\lambda \mu_\Pi(x, y, z) \sim \mu(\lambda x, \lambda y, \lambda z).$$

Thus $\lambda : \Pi \rightarrow \Lambda$ is a quasimorphism. The constants depend only on p and the parameters inputted. This verifies (C2). \square

We have shown that (Λ, ρ, μ) is a coarse median space. With some additional hypotheses (justified for the mapping class group in Section 11), we can control the rank. For this we will assume the spaces $\Theta(X)$ to be uniformly hyperbolic. In this regard, we introduce the following notation.

Suppose that (Θ, σ) is k_0 -hyperbolic. Given $x, y, z, w \in \Theta$, we write

$$(x, y : z, w) = \frac{1}{2}(\max\{\sigma(x, z) + \sigma(y, w), \sigma(x, w) + \sigma(y, z)\} - (\sigma(x, y) + \sigma(z, w))).$$

Up to an additive constant, depending only on k_0 , this ‘‘crossratio’’ is equal to the distance between any geodesic from x to y and any geodesic from z to w . Note that $(x, y : z, w) \leq \sigma(x, z)$, and that $(x, x : y, y) = \sigma(x, y)$. Also, $(x, y : z, z)$ is the ‘‘Gromov product’’ of x and y with respect to z . Again, up to an additive constant, this equals the distance from z to any geodesic from x to y .

We now make the following additional hypotheses. We suppose that \mathcal{X} comes equipped with a symmetric relation, \wedge , with not $X \wedge X$ for all $X \in \mathcal{X}$. We suppose:

- (P2) There is some $k_0 \geq 0$ such that each $(\Theta(X), \sigma_X)$ is k_0 -hyperbolic.
- (P3) There is some $\nu \in \mathbb{N}$ such that if we have a subset $\mathcal{Y} \subseteq \mathcal{X}$ with $X \wedge Y$ for all distinct $X, Y \in \mathcal{Y}$, then $|\mathcal{Y}| \leq \nu$.
- (P4) There is some $l_0 \geq 0$ such that if $X, Y \in \mathcal{X}$ and there exist $a, b, c, d \in \Lambda$ with

$$(\theta_X a, \theta_X b : \theta_X c, \theta_X d) \geq l_0 \quad \text{and} \quad (\theta_Y a, \theta_Y c : \theta_Y b, \theta_Y d) \geq l_0,$$

then $X \wedge Y$.

In relation to the mapping class group, where Λ is the marking complex, these are interpreted as follows. The relation, \wedge , refers to disjointness of the subsurfaces in Σ . Thus, (P3) is a purely topological observation, where $\nu = \xi(\Sigma)$ as defined in Section 2. For (P2), we have already noted that curve complexes are hyperbolic [Masur and Minsky 1999]. Property (P4) follows from properties of subsurface projection as we discuss in Section 11.

Proposition 10.2. *Suppose that (Λ, ρ, μ) satisfies the above — in particular, conditions (P1)–(P4). Then (Λ, ρ, μ) is a coarse median space of rank at most ν .*

Here, ν is the constant featuring in (P3). As usual, the parameters outputted depend only on those of the hypotheses.

Before giving the proof, we need a general observation regarding hyperbolic spaces. Let (Θ, σ) be k_0 -hyperbolic. Let μ be the median as defined in Section 3. We know that (Θ, σ, μ) is coarse median of rank 1. In fact:

Lemma 10.3. *Given $k_0, l \geq 0$ and $p \in \mathbb{N}$, there is some $h \geq 0$ with the following property. Suppose that (Θ, σ) is k_0 -hyperbolic, and that μ is the median on (Θ, σ) . Suppose that A is any set with $|A| \leq p < \infty$ and that $\theta : A \rightarrow \Theta$ is any map. Then there is a rank-1 median algebra, Π , and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Theta$, satisfying:*

(L1) $\sigma(\theta a, \lambda \pi a) \leq h$ for all $a \in A$.

(L2) λ is an h -quasimorphism.

(L3) If $a, b, c, d \in A$ with $(\pi a, \pi b | \pi c, \pi d)_\Pi$, then $(\theta a, \theta b : \theta c, \theta d) \geq l$.

Here, of course, Π is just the vertex set, $V(\tau)$, of a simplicial tree, τ . As in Section 4, we use the notation $(x, y | z, w)_\Pi$ to mean that the sets $\{x, y\}$ and $\{z, w\}$ are separated by a wall in Π . Here, this is equivalent to saying that the arcs $[x, y]_\tau$ and $[z, w]_\tau$ are disjoint.

Proof. Let $\tau_0 \subseteq \Theta$ be the embedded tree arising from $\theta(A) \subseteq \Theta$, as given by Lemma 3.2. Thus, if $a, b \in A$, then $\sigma_{\tau_0}(\theta a, \theta b) \leq \sigma(\theta a, \theta b) + k_1$, where $k_1 = k_0 h_0(p)$. Let $t = l + 2k_1$ and let τ be the metric tree obtained from τ_0 by collapsing down each edge of length at most t . Let $\Pi = V(\tau)$. Given $x \in \Pi$, let $\tau(x) \subseteq \tau$ be the preimage of x under the collapsing map. Thus, $\tau(x)$ is a subtree of diameter at most $k_2 = pt$.

Now let $\pi : A \rightarrow \Pi$ be the postcomposition of θ with the collapsing map of τ_0 to τ , define $\lambda : \Pi \rightarrow \Theta$ by setting λx to be any vertex of $\tau(x)$.

If $a \in A$, then $\theta a, \lambda \pi a \in \tau(\theta a)$, so $\sigma(\theta a, \lambda \pi a) \leq k_2$. This gives (L1) provided $h \geq k_2$.

For (L2), suppose that $x, y, z \in \Pi$. By definition, $\lambda x \in \tau(x)$, $\lambda y \in \tau(y)$ and $\lambda z \in \tau(z)$. Let $w = \mu_\Pi(x, y, z)$. Let $w' = \mu_{\tau_0}(\lambda x, \lambda y, \lambda z) \in \tau_0 \subseteq \Theta$. Now $w', \lambda w \in \tau(w)$, and so $\sigma(w', \lambda w) \leq k_2$. Now, as in the proof of Lemma 3.1, the

median $\mu_{\Theta}(\lambda x, \lambda y, \lambda z)$ in Θ is a bounded distance from the median w' in τ_0 , where the bound depends only on p and k_0 . This gives a bound on $\sigma(\lambda w, \mu_{\Theta}(\lambda x, \lambda y, \lambda z))$ as required.

Finally, suppose that $a, b, c, d \in A$ with $(\pi a, \pi b | \pi c, \pi d)_{\Pi}$. It follows that $[\pi a, \pi b]_{\tau_0} \cap [\pi c, \pi d]_{\tau_0} = \emptyset$, and so the crossratio $(\theta a, \theta b : \theta c, \theta d)$ defined intrinsically to τ_0 must be at least t . But this agrees with the crossratio defined in Θ up to an additive constant $2k_1$. This proves property (L3). \square

Now let l_0 be the constant in property (P4). Suppose that $A \subseteq \Lambda$ with $|A| \leq p < \infty$. Let $X \in \mathcal{X}$. Property (P2) tells us that $(\Theta(X), \sigma_X)$ is k_0 -hyperbolic, where k_0 depends only on $\xi(\Sigma)$. Let μ_X be the median operation on $\Theta(X)$. Lemma 10.3 now gives us a rank-1 median algebra $\Pi(X)$, and maps $\pi_X : A \rightarrow \Pi(X)$ as well as a h -quasimorphism, $\lambda_X : \Pi(X) \rightarrow \Theta(X)$, such that if $a, b, c, d \in A$, with $(\pi_X a, \pi_X b | \pi_X c, \pi_X d)_{\Pi(X)}$ then $(\theta_X a, \theta_X b : \theta_X c, \theta_X d) > l_0$.

Now let $\Pi_0 = \prod_{X \in \mathcal{X}} \Pi(X)$, and let $\psi_X : \Pi_0 \rightarrow \Pi(X)$ be the projection map. We define $\pi : A \rightarrow \Pi_0$ so that $\psi_X \pi a = \pi_X a$ for all $a \in A$. Let $\Pi = \langle \pi A \rangle \subseteq \Pi_0$, be the subalgebra generated by πA . Note that Π is finite.

(We note that, a-priori, Π_0 might be infinite. In fact, in the application to the mapping class group, we will see that $\Pi(X)$ is trivial for all but finitely many X , so in fact, Π_0 , can be taken to be finite. We do not formally need that here.)

Recall that we can naturally identify the set of walls, ${}^{\circ}\mathcal{W}(\Pi_0)$, with $\bigsqcup_{X \in \mathcal{X}} {}^{\circ}\mathcal{W}(\Pi(X))$ via the projection maps, ψ_X . Also, by Lemma 6.5 any wall, W , in Π arises from a wall in Π_0 , and hence from a wall in $\Pi(X)$ for some $X \in \mathcal{X}$. (In fact, Lemma 5.8 will suffice in the case of the mapping class group, where Π_0 is finite.) We write $X(W)$ for some such X . (It might not be uniquely determined by W .) Note that the map $[W \mapsto X(W)]$ is injective.

Lemma 10.4. *Suppose that $W, W' \in {}^{\circ}\mathcal{W}(\Pi)$ cross. Then $X(W) \wedge X(W')$.*

Proof. Write $X = X(W)$ and $Y = X(W')$. Since W and W' cross, there is a natural epimorphism of Π to the square $W \times W'$. Since $\Pi = \langle \pi A \rangle$, the restriction to πA is also surjective (since any subset of $W \times W'$ is a subalgebra). In other words, we can find $a, b, c, d \in A$ satisfying $(\pi_X a, \pi_X b | \pi_X c, \pi_X d)_{\Pi(X)}$ and $(\pi_Y a, \pi_Y c | \pi_Y b, \pi_Y d)_{\Pi(Y)}$. Thus, by the construction of $\Pi(X)$ and $\Pi(Y)$, we have

$$(\theta_X a, \theta_X b : \theta_X c, \theta_X d) \geq l_0 \quad \text{and} \quad (\theta_Y a, \theta_Y c : \theta_Y b, \theta_Y d) \geq l_0.$$

By (P4) it now follows that $X \wedge Y$. \square

Corollary 10.5. *Π has rank at most ν .*

Proof. Suppose that ${}^{\circ}\mathcal{W} \subseteq {}^{\circ}\mathcal{W}(\Pi)$ is a set of pairwise crossing walls. By Lemma 10.4, we have $X(W) \wedge X(W')$ for all distinct $W, W' \in {}^{\circ}\mathcal{W}_0$. It now follows by (P3) that $|{}^{\circ}\mathcal{W}_0| \leq \nu$. \square

Proof of Proposition 10.2. We proceed as in the proof of Proposition 10.1. We already have (C1).

For (C2), we need that the rank of Π is at most ν . Instead of taking the free median algebra on A , we take Π as constructed above. In the verification of (C2) we only used the fact that $\Pi = \langle \pi A \rangle$, together with the existence of uniform quasimorphisms $\omega_X : \Pi \rightarrow \Theta(X)$ with $\theta_X a \sim \omega_X \pi a$ for all $a \in A$. (In the proof of Proposition 10.1, we had $\theta_X a = \omega_X \pi a$, but we only need that these agree up to bounded distance.)

This time, we have $\Pi = \langle \pi A \rangle$ by construction. The quasimorphism ω_X can now be defined as the composition $\omega_X = \lambda_X \psi_X$.

The proof now proceeds as before. \square

11. Surfaces

In this section we verify the hypotheses of Proposition 10.2 in the case where Λ is a connected locally finite graph on which the mapping class group, $\text{Map}(\Sigma)$, acts properly discontinuously with finite quotient. This shows that $\text{Map}(\Sigma)$ is a coarse median group of rank at most $\xi(\Sigma)$.

Here, Θ will be the curve graph $\mathcal{C} = \mathcal{C}(\Sigma)$, of Σ , \mathcal{X} will be the set of subsurfaces of Σ , and $\Theta(X) = \mathcal{C}(X)$ will be the curve graph defined intrinsically to $X \in \mathcal{X}$ (appropriately interpreted if X is an annulus). Briefly, Property (P1) is a consequence of the distance formula of [Masur and Minsky 2000] (see Lemma 11.5), Property (P2) is the hyperbolicity of the curve complex proven in [Masur and Minsky 1999], Property (P3) is an elementary topological observation (see Lemma 11.1) and Property (P4) follows from a result in [Behrstock 2006] which is reformulated here as Lemma 11.3 (see Lemma 11.7).

For the graph, Λ , we could use a Cayley graph with respect to a finite generating set, though we will find it more convenient to work with a “marking complex”; compare [Masur and Minsky 2000].

We now give more formal definitions. Let Σ be a compact orientable surface with (possibly empty) boundary $\partial\Sigma$. Let $\xi(\Sigma) = 3g + p - 3$, where g is the genus, and p the number of boundary components. We assume that $\xi(\Sigma) > 1$. Let $\mathcal{C}^0 = \mathcal{C}^0(\Sigma)$ be the set of homotopy classes of essential nonperipheral simple closed curves in Σ , referred to here simply as “curves”. Given $\alpha, \beta \in \mathcal{C}^0$, we write $\iota(\alpha, \beta)$ for their geometric intersection number, in other words, the minimal possible number of intersections taken over all representative curves in the respective homotopy classes. (We remark that given any finite set of curves in Σ , we can find realisations which simultaneously achieve these minima for all pairwise intersections — for example, take geodesic representatives with respect to any complete hyperbolic structure on the interior of Σ .) The *curve graph*, $\mathcal{C} = \mathcal{C}(\Sigma)$, is the graph with vertex set,

$V(\mathcal{C}) = \mathcal{C}^0$, where $\alpha, \beta \in \mathcal{C}^0$ are adjacent if $\iota(\alpha, \beta) = 0$. (This is the 1-skeleton of Harvey’s curve complex.) We write σ for the combinatorial metric on \mathcal{C} . It was shown in [Masur and Minsky 1999] that \mathcal{C} is hyperbolic. (A constructive proof can be found in [Bowditch 2006b].) It is not hard to see that $\sigma(\alpha, \beta)$ is bounded above in terms of $\iota(\alpha, \beta)$ (for example, $\sigma(\alpha, \beta) \leq \iota(\alpha, \beta) + 1$). We will write $\alpha \pitchfork \beta$ to mean that $\iota(\alpha, \beta) > 0$.

Given $a \subseteq \mathcal{C}^0$, we write $\iota(a) = \max\{\iota(\alpha, \beta) \mid \alpha, \beta \in a\}$ for the *self-intersection* of a . If $\iota(a) < \infty$ then a is finite. (In fact, $\sum\{\iota(\alpha, \beta) \mid \alpha, \beta \in a\}$ is bounded above in terms of $\iota(a)$ and $\xi(\Sigma)$.) We say that a *fills* Σ if, for all $\gamma \in \mathcal{C}^0$, there is some $\alpha \in a$ with $\alpha \pitchfork \gamma$. Given $p \in \mathbb{N}$, we write $L(p)$ for the set of subsets $a \subseteq \mathcal{C}^0$ with $\iota(a) \leq p$ and which fill Σ . Given $p, q \in \mathbb{N}$ we write $\Lambda(p, q)$ for the graph with vertex set $L(p)$ where $a, b \in L(p)$ are deemed to be adjacent if $\iota(a \cup b) \leq q$. Thus, $\Lambda(p, q)$ is locally finite, and $\text{Map}(\Sigma)$ acts on $\Lambda(p, q)$ with finite quotient. For a “marking complex”, we could take any connected $\text{Map}(\Sigma)$ -invariant subgraph of $\Lambda(p, q)$ for some p, q (which might be allowed to depend on $\xi(\Sigma)$). The notion is quite robust, so it doesn’t much matter exactly what construction we use. For definiteness, we can set Λ to be the marking complex used in [Masur and Minsky 2000]. In this case, $\Lambda \subseteq \Lambda(4, 4)$. (We could also use $\Lambda(p, q)$ itself for sufficiently large p, q .)

We define a map $\chi : \Lambda \rightarrow \mathcal{C}$, which chooses some element $\chi(a) \in a$ from each $a \in V(\Lambda)$. Note that this is uniformly lipschitz with respect to the metrics ρ and σ on $V(\Lambda)$ and $V(\mathcal{C}) = \mathcal{C}^0$. (We can extend to a map $\Lambda \rightarrow \mathcal{C}$, by first collapsing each of Λ to an incident vertex.)

We now move on to consider subsurfaces.

Definition. By a *subsurface realised in* Σ we mean a compact connected subsurface $X \subseteq \Sigma$ such that each boundary component of X is either a component of $\partial\Sigma$, or else an essential nonperipheral simple closed curve in $\Sigma \setminus \partial\Sigma$, and such that X is not homeomorphic to a three-holed sphere.

Note that we are allowing Σ itself as a subsurface, as well as nonperipheral annuli.

Definition. A *subsurface* is a free homotopy class of realised subsurfaces.

We will sometimes abuse notation and use the same symbol for a subsurface and some realisation of it in Σ .

We write $\mathcal{X} = \mathcal{X}(\Sigma)$ for the set of subsurfaces of Σ . We write $\mathcal{X} = \mathcal{X}_A \sqcup \mathcal{X}_N$ where \mathcal{X}_A and \mathcal{X}_N are respectively the sets of annular and nonannular subsurfaces. Note that there is a natural bijective correspondence between \mathcal{X}_A and the set of curves, \mathcal{C}^0 . (We will, however, treat them as distinct from the point of view of the notation introduced below.)

Suppose $X \in \mathcal{X}_N$. We have $0 < \xi(X) \leq \xi(\Sigma)$, and write $\mathcal{C}^0(X)$, $\mathcal{C}(X)$, $\Lambda(X)$ respectively for \mathcal{C}^0 , \mathcal{C} , Λ defined intrinsically to X . (In the exceptional cases where $\xi(X) = 1$, $\mathcal{C}(X)$ is defined by deeming two curves to be adjacent if they have minimal possible intersection for that surface, that is, 1 for a one-holed torus, and 2 for a four-holed sphere. In both cases this gives us a Farey graph.) Note that we can identify $\mathcal{C}^0(X)$ as a subset of \mathcal{C}^0 . We write σ_X and ρ_X for the combinatorial metrics on $\mathcal{C}(X)$ and $\Lambda(X)$. Let $\mathcal{C}^0(\Sigma, X)$ and $\mathcal{C}^0(\Sigma, \partial X)$ be the subsets of \mathcal{C}^0 consisting of curves of Σ homotopic into X or ∂X respectively. In this way, $\mathcal{C}^0(\Sigma, X) = \mathcal{C}^0(X) \sqcup \mathcal{C}^0(\Sigma, \partial X)$.

If $X \in \mathcal{X}_A$, the set $\mathcal{C}(X)$ is defined as an arc complex in the cover of Σ corresponding to X , as in [Masur and Minsky 2000]. This is quasi-isometric to the real line. We set $\Lambda(X) = \mathcal{C}(X)$.

Given $X, Y \in \mathcal{X}$, we distinguish five mutually exclusive possibilities denoted as follows:

- (1) $X = Y$.
- (2) $X \prec Y$: $X \neq Y$, and X can be homotoped into Y but not into ∂Y .
- (3) $Y \prec X$: $Y \neq X$, and Y can be homotoped into X but not into ∂X .
- (4) $X \wedge Y$: $X \neq Y$ and X, Y can be homotoped to be disjoint.
- (5) $X \pitchfork Y$: none of the above.

In (2)–(4) one can find realisations of X, Y in Σ such that $X \subseteq Y$, $Y \subseteq X$, $X \cap Y = \emptyset$, respectively. (Note that $X \wedge Y$ covers the case where X is an annulus homotopic to a boundary component of Y , or vice versa.) We can think of (5) as saying that the surfaces “overlap”.

Lemma 11.1. *Suppose $\mathcal{Y} \subseteq \mathcal{X}$ satisfies $X \wedge Y$ for all distinct $X, Y \in \mathcal{Y}$. Then $|\mathcal{Y}| \leq \xi(\Sigma)$.*

Proof. For each $Y \in \mathcal{Y}$, choose an essential curve, α_Y in Y which is nonperipheral if $Y \in \mathcal{X}_N$ and the core curve if $Y \in \mathcal{X}_A$. The curves α_Y are all pairwise nonhomotopic in Σ , so there can be at most $\xi(\Sigma)$ of them. \square

Next we consider subsurface projections. These were defined in [Masur and Minsky 2000].

Let $X \in \mathcal{X}$. If $\alpha \in \mathcal{C}^0$, write $\alpha \pitchfork X$ to mean that either $\alpha \in \mathcal{C}^0(X)$ or $\alpha \pitchfork \gamma$ for some $\gamma \subseteq \partial X$. In other words, α cannot be homotoped to be disjoint from X . (This is consistent with the notation above if we identify α with an annular neighbourhood.) In this case, we write $\theta_X \alpha$ for a projection of α in $\mathcal{C}(X)$, as defined in [Masur and Minsky 2000]. There is some ambiguity in the definition, but it is well defined up to bounded distance. In fact, if $X \in \mathcal{X}_N$, we can take $\theta_X \alpha \in \mathcal{C}^0(X)$, and this case, it is well defined up to bounded intersection. Moreover, if $\alpha, \beta \pitchfork X$,

then $\iota(\theta_X\alpha, \theta_X\beta)$ is bounded above in terms of $\iota(\alpha, \beta)$. Note that if a fills Σ , then at least one $\alpha \in a$ must satisfy $a \pitchfork X$. The resulting curve, $\theta_X\alpha \in \mathcal{C}^0(X)$, is well defined up to bounded intersection number in X , where the bound depends only on $\iota(a)$. This gives rise to a map $\theta_X : \Lambda \rightarrow \mathcal{C}(X)$, well defined up to bounded distance. Moreover, θ_X is uniformly lipschitz with respect to the metrics ρ and σ_X .

Suppose that $a \in L(p)$, for $p \geq 4$. Let $a_X \subseteq a$ be the set of curves, $\alpha \in a$, with $\alpha \pitchfork X$. This must be nonempty. Note that $\{\theta_X\alpha \mid \alpha \in a_X\}$ has bounded self-intersection. Moreover, if p is large enough it's not hard to see that this set must fill X . Given these observations, we see that we have also a map $\phi_X : \Lambda \rightarrow \Lambda(X)$, well defined up to bounded distance, and uniformly lipschitz with respect to the metrics ρ and ρ_X . (Namely, set $\phi_X(\alpha) = \theta_X\alpha$ for some $\alpha \in a_X$.) Moreover, writing $\chi_X : \Lambda(X) \rightarrow \mathcal{C}(X)$, for the map χ defined intrinsically to X , we see that the map θ_X agrees up to bounded distance with the composition $\chi_X\phi_X$.

Suppose that $X, Y \in \mathcal{X}$ with $X \pitchfork Y$ or $Y \prec X$. We define a point $\theta_X Y \in \mathcal{C}(X)$ as follows. If $Y \in \mathcal{X}_A$, we set $\theta_X Y = \theta_X\alpha$, where $\alpha \in \mathcal{C}^0$ is the curve homotopic to Y . If $Y \in \mathcal{X}_N$, we choose any $\alpha \in \mathcal{C}^0(\Sigma, \partial X)$ with $\alpha \pitchfork Y$ and set $\theta_X Y = \theta_X\alpha$. Note that this is well defined up to bounded distance.

We list a few properties of subsurface projections.

First note that if $X \prec Y$, we have a subsurface projection, θ_{XY} defined intrinsically to Y . In other words, we can replace Σ by Y in the earlier discussion, and work intrinsically with Y . (Note that $Y \in \mathcal{X}_N$.)

Lemma 11.2. *If $\alpha \in \mathcal{C}^0(\Sigma)$ with $\alpha \pitchfork X$, then $\alpha \pitchfork Y$, and $\sigma_X(\theta_X\alpha, \theta_{XY}\theta_Y\alpha)$ is bounded in terms of $\xi(\Sigma)$.*

Proof. This is an easy consequence of the construction in [Masur and Minsky 2000]. \square

In fact, using that same construction, we see that the intersection number between $\theta_X\alpha$ and $\theta_{XY}\theta_Y\alpha$ is also bounded. In view of this, we can henceforth drop the suffix “ Y ”, and write θ_{XY} as θ_X .

Lemma 11.3. *There is some constant l_1 , depending only on $\xi(\Sigma)$, with the following property. Suppose that $X, Y \in \mathcal{X}$ with $X \pitchfork Y$, and that $a \in V(\Lambda)$. Then $\min\{\sigma_X(\theta_X a, \theta_X Y), \sigma_Y(\theta_Y a, \theta_Y X)\} \leq l_1$.*

Proof. This is an immediate consequence of the result in [Behrstock 2006]; see also [Mangahas 2010]. This was stated for curves, namely that if $\alpha \in \mathcal{C}^0$ with $\alpha \pitchfork X$ and $\alpha \pitchfork Y$, then $\min\{\sigma_X(\theta_X\alpha, \theta_X Y), \sigma_Y(\theta_Y\alpha, \theta_Y X)\}$ is bounded above in terms on $\xi(\Sigma)$. To relate this to our statement, it is a simple exercise to find such a curve, $\alpha \in \mathcal{C}^0$, with $\iota(a \cup \{\alpha\})$ bounded in terms of $\xi(\Sigma)$. Thus, $\sigma_X(\theta_X\alpha, \theta_X a)$ and $\sigma_Y(\theta_Y\alpha, \theta_Y a)$ are bounded. \square

Lemma 11.4. *There is some l_2 , depending only on $\xi(\Sigma)$ with the following property. Suppose $X, Y \in \mathcal{X}$ with $Y < X$, and suppose that $a, b \in \Lambda$ with $(\theta_X a, \theta_X b : \theta_X Y, \theta_X Y) \geq l_2$. Then $\sigma_Y(\theta_Y a, \theta_Y b) \leq l_2$.*

Proof. Choosing $\alpha \in a$ and $\beta \in b$ with $\alpha \curvearrowright Y$ and $\beta \curvearrowright Y$, we will also have $\alpha \curvearrowright X$ and $\beta \curvearrowright X$. We can therefore interpret the lemma as a statement about curves rather than markings (perhaps with a different constant). Also, in view of Lemma 11.2, we may as well assume that $X = \Sigma$, so that $\alpha = \theta_\Sigma \alpha$ and $\beta = \theta_\Sigma \beta$, and we set $\gamma = \theta_\Sigma Y \in \mathcal{C}^0(\Sigma, Y)$. Now $\mathcal{C}^0(\Sigma, Y)$ has diameter at most 2 in \mathcal{C} . Thus, if the Gromov product $(\alpha, \beta | \gamma, \gamma)$ is sufficiently large in relation to the hyperbolicity constant of \mathcal{C} , then any geodesic from α to β in \mathcal{C} will miss $\mathcal{C}^0(\Sigma, Y)$. By the bounded geodesic image theorem of Masur and Minsky [2000], it then follows that $\sigma_Y(\theta_Y \alpha, \theta_Y \beta)$ and hence $\sigma_Y(\theta_Y a, \theta_Y b)$ is bounded as required. \square

The following two lemmas are both consequences of the distance formula in [Masur and Minsky 2000] (though can also be seen more directly). The first of these implies (P1).

Lemma 11.5. *Given any $l \geq 0$, there is some $l' \geq 0$, depending only on l and $\xi(\Sigma)$ with the following property. Suppose that $a, b \in \Lambda$ and that $\sigma_X(\theta_X a, \theta_X b) \leq l$ for all $X \in \mathcal{X}$, then $\rho(a, b) \leq l'$.*

Lemma 11.6. *There is some l_3 depending only on $\xi(\Sigma)$ such that if $a, b \in \Lambda$, then $\{X \in \mathcal{X} \mid \sigma_X(\theta_X a, \theta_X b) \geq l_3\}$ is finite.*

We can now verify property (P4) of Proposition 10.1.

Lemma 11.7. *There is some $l_0 \geq 0$, depending only on $\xi(\Sigma)$ such that if $X, Y \in \mathcal{X}$ and there exist $a, b, c, d \in \Lambda$ with*

$$(\theta_X a, \theta_X b : \theta_X c, \theta_X d) \geq l_0 \quad \text{and} \quad (\theta_Y a, \theta_Y c : \theta_Y b, \theta_Y d) \geq l_0,$$

then $X \wedge Y$.

Proof. Since $\mathcal{C}(X)$ and $\mathcal{C}(Y)$ are hyperbolic, we must have $X \neq Y$, provided that l_0 is large enough in relation to the hyperbolicity constant. We will also assume that $l_0 \geq 2 \max\{l_1, l_2\}$ (the constants of Lemmas 11.3 and 11.4). If not $X \wedge Y$, then either $X \curvearrowright Y$ or, without loss of generality, $Y < X$.

Note that the hypotheses on a, b, c, d remain unchanged if we simultaneously swap a with b and c with d . Since $(\theta_X a, \theta_X b : \theta_X c, \theta_X d) \geq l_0 > 2 \max\{l_1, l_2\}$, we can assume that $(\theta_X a, \theta_X b : \theta_X Y, \theta_X Y) \geq \max\{l_1, l_2\}$. In particular, this implies that $\sigma_X(\theta_X a, \theta_X Y) > l_1$ and $\sigma_X(\theta_X b, \theta_X Y) > l_1$. Now, if $X \curvearrowright Y$, then Lemma 11.3 tells us that $\sigma_Y(\theta_Y a, \theta_Y X) \leq l_1$ and $\sigma_Y(\theta_Y b, \theta_Y X) \leq l_1$, so that $\sigma_Y(\theta_Y a, \theta_Y b) \leq 2l_1$, giving the contradiction that $(\theta_Y a, \theta_Y c : \theta_Y b, \theta_Y d) \leq l_1$. If $Y < X$, then by Lemma 11.4, we have $\sigma_Y(\theta_Y a, \theta_Y b) \leq l_2$ again giving a contradiction. \square

We have now verified each of the hypotheses of Proposition 10.2 for the mapping class group, where $\nu = \xi(\Sigma)$. This proves Theorem 2.5.

12. Colourability

In this section we briefly describe the notion of colourability for median algebras and coarse median spaces. In general, this is a strengthening of the rank condition. This property is used in [Bowditch 2011a] to give embeddings of median algebras into products of trees.

Let M be a median algebra.

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

Clearly this implies that the rank of M is at most ν . The converse does not hold in general, but it does for intervals (see Lemma 12.4).

Proposition 12.1. *A median algebra is ν -colourable if and only if every finite subalgebra is.*

(In fact, it is the latter condition that is applied in practice, so in principle one could bypass this discussion by defining colourability in that way.)

Lemma 12.2. *Any subalgebra of a ν -colourable median algebra is ν -colourable.*

Proof. Let N be a subalgebra of a ν -colourable median algebra, M . Let $\nu : M \rightarrow \{1, \dots, \nu\}$ be a ν -colouring. If $W \in \mathcal{W}(N)$, then by Lemma 6.1, there is a wall in M separating $H^-(W) \subseteq N$ from $H^+(W) \subseteq N$. Let W_M be any such wall. We write $\chi(W) = \chi(W_M)$. Now if $W, W' \in \mathcal{W}(N)$ cross in N , then certainly W_M and W'_M cross in M , and so $\chi(W) \neq \chi(W')$. Thus, $\chi : \mathcal{W}(N) \rightarrow \{1, \dots, \nu\}$ is a ν -colouring of N . \square

Lemma 12.3. *If every finite subalgebra of a median algebra M is ν -colourable then M is ν -colourable.*

Proof. We first note that it's enough to show that for any finite subset, $\mathcal{W}_0 \subseteq \mathcal{W}(M)$, we can find a map $\chi : \mathcal{W}_0 \rightarrow \{1, \dots, \nu\}$ such that $\chi(W) \neq \chi(W')$ whenever $W, W' \in \mathcal{W}_0$ with $W \pitchfork W'$. To deduce Lemma 12.3 from this, we recall the standard compactness result from graph theory, namely that a graph is vertex ν -colourable if and only if every finite subgraph is. Here we construct a graph, \mathcal{G} , with vertex set $\mathcal{W}(M)$, where $W, W' \in \mathcal{W}(M)$ are deemed adjacent if and only if $W \pitchfork W'$. Thus, colouring M is equivalent to vertex-colouring the graph \mathcal{G} . Our claim therefore says that every full subgraph of \mathcal{G} is ν -colourable.

Let $\mathcal{W}_0 \subseteq \mathcal{W}(M)$ be finite. Given any pair, $W, W' \in \mathcal{W}_0$ with $W \pitchfork W'$, choose any $a \in H^-(W) \cap H^-(W')$, $b \in H^+(W) \cap H^-(W')$, $c \in H^-(W) \cap H^+(W')$ and $d \in H^+(W) \cap H^+(W')$. Let A be the union of all such $\{a, b, c, d\}$ as (W, W')

ranges over all such pairs. Let Π be a finite median algebra of M containing A . By hypothesis, there is a ν -colouring, $\chi : \mathcal{W}(\Pi) \rightarrow \{1, \dots, \nu\}$. Now each $W \in \mathcal{W}_0$ determines a wall, $\hat{W} = \{H^-(W) \cap \Pi, H^+(W) \cap \Pi\}$ in $\mathcal{W}(\Pi)$. Clearly, if W, W' cross in M , then \hat{W}, \hat{W}' cross in Π , and so we can set $\chi(W) = \chi(\hat{W})$ for any such W to prove the claim. \square

Lemmas 12.2 and 12.3 now give Proposition 12.1

Suppose Δ is a metric median algebra with points $a, b \in \Delta$ such that $\Delta = [a, b]$. We can orient any wall, $W \in \mathcal{W}(\Delta)$, so that $a \in H^-(W)$ and $b \in H^+(W)$. Given $W, W' \in \mathcal{W}(\Delta)$, we write $W \leq W'$ 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 $W = W', W < W', W' < W$ or $W \pitchfork W'$ holds. It follows that the rank of Δ is exactly the maximal cardinality of any antichain in $(\mathcal{W}(\Delta), <)$. Dilworth's lemma [Dilworth 1950] now tells us that we can partition $\mathcal{W}(\Delta)$ into ν disjoint chains (compare [Brodzki et al. 2009]). This defines a ν -colouring of Δ . We deduce:

Lemma 12.4. *Let M be a median algebra of rank at most ν . If $a, b \in M$, then the interval $[a, b]$ is intrinsically ν -colourable as a median algebra.*

The definition for coarse median spaces is now a simple variation on that for rank:

Definition. A coarse median space is ν -colourable, if in (C2), we can always take the finite median algebra Π to be ν -colourable.

Suppose now that $(\Lambda_i, \rho_i, \mu_i)$ is a directed set of coarse median space as in Theorem 2.3 (where the additive constants tend to 0, and the multiplicative constants are bounded with respect to the ultrafilter). Let $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ be the ultralimit constructed as in Proposition 9.1

Proposition 12.5. *If each of the $(\Lambda_i, \rho_i, \mu_i)$ is ν -colourable (for the given parameters) then $(\Lambda_\infty, \rho_\infty, \mu_\infty)$ is ν -colourable (as a median algebra).*

Proof. Substituting colourability for rank in the proof of Theorem 2.3 in Section 9, exactly the same argument shows that every finite subalgebra of Λ_∞ is ν -colourable. We now apply Lemma 12.3. \square

Again the notion is quasi-isometry invariant, so we can apply it to finitely generated groups via their Cayley graphs. We note:

Theorem 12.6. *The mapping class group $\text{Map}(\Sigma)$ is ν -colourable for some $\nu = \nu(\Sigma)$.*

In fact, we can get an explicit bound on $\nu(\Sigma)$ from the statement in [Bestvina et al. 2010] which gives us a map: $\chi : \mathcal{X} \rightarrow \{1, \dots, \nu(\Sigma)\}$ such that if $\chi(X) = \chi(Y)$, then $X \pitchfork Y$.

The proof of Lemma 11.2 now only requires a slight modification of that of Theorem 2.5. Recall that the median algebra Π used for Property (C2) was constructed using projection maps, before the statement of Lemma 10.4. We now need to check that this is $\nu(\Sigma)$ -colourable — a slight modification of Corollary 10.5. For this we need a variation of Property (P3), namely:

(P3') If $X, Y \in \mathcal{X}$ with $X \wedge Y$, then $\chi(X) \neq \chi(Y)$.

In the present situation, Property (P3') is an immediate consequence of the definition of the relation \pitchfork in Section 11, and the construction of [Bestvina et al. 2010] mentioned above.

Now let Π be the median algebra defined before Lemma 10.4. We define a map $\chi : \mathcal{W}(\Pi) \rightarrow \{1, \dots, \nu(\Sigma)\}$ by setting $\chi(W) = \chi(X(W))$. We claim that this is a $\nu(\Sigma)$ -colouring of Π . To see this, suppose that $W, W' \in \mathcal{W}(\Pi)$ cross. Lemma 10.4 then tells us that $X(W) \wedge X(W')$ and so, by (P3'), $\chi(W) \neq \chi(W')$, as required. We can thus replace Corollary 10.5 by the statement that Π is $\nu(\Sigma)$ -colourable, and so Theorem 12.6 follows.

As a consequence, from [Bowditch 2011a] we recover the result of Behrstock, Druţu and Sapir [Behrstock et al. 2011] that any asymptotic cone of $\text{Map}(\Sigma)$ admits a bilipschitz embedding in a finite product of \mathbb{R} -trees. Moreover, using Lemma 12.3, any interval in the asymptotic cone is compact, and admits a bilipschitz embedding in $\mathbb{R}^{\xi(\Sigma)}$. From this one can recover the fact that $\text{Map}(\Sigma)$ has rapid decay [Behrstock and Minsky 2011].

References

- [Bandelt and Hedlíková 1983] H.-J. Bandelt and J. Hedlíková, “Median algebras”, *Discrete Math.* **45**:1 (1983), 1–30. MR 84h:06015 Zbl 0506.06005
- [Bandelt and van de Vel 1989] H.-J. Bandelt and M. van de Vel, “Embedding topological median algebras in products of dendrons”, *Proc. London Math. Soc.* (3) **58**:3 (1989), 439–453. MR 90j:52001 Zbl 0682.05031
- [Behrstock 2006] J. A. Behrstock, “Asymptotic geometry of the mapping class group and Teichmüller space”, *Geom. Topol.* **10** (2006), 1523–1578. MR 2008f:20108 Zbl 1145.57016
- [Behrstock and Minsky 2008] J. A. Behrstock and Y. N. Minsky, “Dimension and rank for mapping class groups”, *Ann. of Math.* (2) **167**:3 (2008), 1055–1077. MR 2009d:57031 Zbl 05578711
- [Behrstock and Minsky 2011] J. A. Behrstock and Y. N. Minsky, “Centroids and the rapid decay property in mapping class groups”, *J. London Math. Soc.* (2) **84**:3 (2011), 765–784. MR 2855801 Zbl 05987716 arXiv 0810.1969
- [Behrstock et al. 2011] J. A. Behrstock, C. Druţu, and M. Sapir, “Median structures on asymptotic cones and homomorphisms into mapping class groups”, *Proc. London Math. Soc.* (3) **102**:3 (2011), 503–554. MR 2012c:20110 Zbl 05869985
- [Behrstock et al. 2012] J. A. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, “Geometry and rigidity of mapping class groups”, *Geom. Topol.* **16**:2 (2012), 781–888. Zbl 06035994 arXiv 0801.2006

- [Bestvina et al. 2010] M. Bestvina, K. Bromberg, and K. Fujiwara, “The asymptotic dimension of the mapping class groups is finite”, preprint, 2010. Last revised in 2012 as v3. arXiv 1006.1939v1
- [Bowditch 2006a] B. H. Bowditch, *A course on geometric group theory*, MSJ Memoirs **16**, Mathematical Society of Japan, Tokyo, 2006. MR 2007e:20085 Zbl 1103.20037
- [Bowditch 2006b] B. H. Bowditch, “Intersection numbers and the hyperbolicity of the curve complex”, *J. Reine Angew. Math.* **598** (2006), 105–129. MR 2009b:57034 Zbl 1119.32006
- [Bowditch 2011a] B. H. Bowditch, “Embedding median algebras in products of trees”, preprint, 2011, Available at <http://wrap.warwick.ac.uk/id/eprint/48818>.
- [Bowditch 2011b] B. H. Bowditch, “Invariance of coarse median spaces under relative hyperbolicity”, *Math. Proc. Camb. Phil. Soc.* (2011). To appear.
- [Bridson and Haefliger 1999] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften **319**, Springer, Berlin, 1999. MR 2000k:53038 Zbl 0988.53001
- [Brodzki et al. 2009] J. Brodzki, S. J. Campbell, E. Guentner, G. A. Niblo, and N. Wright, “Property A and CAT(0) cube complexes”, *J. Funct. Anal.* **256**:5 (2009), 1408–1431. MR 2010i:20044 Zbl 1233.20036
- [Chatterji et al. 2010] I. Chatterji, C. Druţu, and F. Haglund, “Kazhdan and Haagerup properties from the median viewpoint”, *Adv. Math.* **225**:2 (2010), 882–921. MR 2011g:20059 Zbl 05777258
- [Chepoi 2000] V. Chepoi, “Graphs of some CAT(0) complexes”, *Adv. in Appl. Math.* **24**:2 (2000), 125–179. MR 2001a:57004 Zbl 1019.57001
- [Dilworth 1950] R. P. Dilworth, “A decomposition theorem for partially ordered sets”, *Ann. of Math.* (2) **51** (1950), 161–166. MR 11,309f Zbl 0038.02003
- [van den Dries and Wilkie 1984] L. van den Dries and A. J. Wilkie, “Gromov’s theorem on groups of polynomial growth and elementary logic”, *J. Algebra* **89**:2 (1984), 349–374. MR 85k:20101 Zbl 0552.20017
- [Engelking 1995] R. Engelking, *Theory of dimensions finite and infinite*, Sigma Series in Pure Mathematics **10**, Heldermann, Lemgo, 1995. MR 97j:54033 Zbl 0872.54002
- [Erdős 1940] P. Erdős, “The dimension of the rational points in Hilbert space”, *Ann. of Math.* (2) **41** (1940), 734–736. MR 2,178a Zbl 0025.18701
- [Farb et al. 2001] B. Farb, A. Lubotzky, and Y. Minsky, “Rank-1 phenomena for mapping class groups”, *Duke Math. J.* **106**:3 (2001), 581–597. MR 2001k:20076 Zbl 1025.20023
- [Gromov 1987] M. Gromov, “Hyperbolic groups”, pp. 75–263 in *Essays in group theory* (Berkeley, CA, 1985), edited by S. M. Gersten, Math. Sci. Res. Inst. Publ. **8**, Springer, New York, 1987. MR 89e:20070 Zbl 0634.20015
- [Gromov 1993] M. Gromov, “Asymptotic invariants of infinite groups”, pp. 1–295 in *Geometric group theory*, 2 (Sussex, 1991), edited by J. W. S. Cassels et al., London Math. Soc. Lecture Note Ser. **182**, Cambridge University Press, Cambridge, 1993. MR 95m:20041 Zbl 0841.20039
- [Hamenstädt 2005] U. Hamenstädt, “Geometry of the mapping class groups, III: Quasi-isometric rigidity”, preprint, 2005. Last revised in 2007 as v2. arXiv math/0512429v1
- [Hurewicz and Wallman 1941] W. Hurewicz and H. Wallman, *Dimension theory*, Princeton Mathematical Series **4**, Princeton University Press, Princeton, NJ, 1941. MR 3,312b Zbl 0060.39808
- [Isbell 1980] J. R. Isbell, “Median algebra”, *Trans. Amer. Math. Soc.* **260**:2 (1980), 319–362. MR 81i:06006 Zbl 0446.06007
- [Kolibiar and Marcisová 1974] M. Kolibiar and T. Marcisová, “On a question of J. Hashimoto”, *Mat. Časopis Sloven. Akad. Vied* **24** (1974), 179–185. MR 50 #4427 Zbl 0285.06008

- [Mangahas 2010] J. Mangahas, “Uniform uniform exponential growth of subgroups of the mapping class group”, *Geom. Funct. Anal.* **19**:5 (2010), 1468–1480. MR 2011d:57002 Zbl 1207.57005
- [Masur and Minsky 1999] H. A. Masur and Y. N. Minsky, “Geometry of the complex of curves, I: Hyperbolicity”, *Invent. Math.* **138**:1 (1999), 103–149. MR 2000i:57027 Zbl 0941.32012
- [Masur and Minsky 2000] H. A. Masur and Y. N. Minsky, “Geometry of the complex of curves, II: Hierarchical structure”, *Geom. Funct. Anal.* **10**:4 (2000), 902–974. MR 2001k:57020 Zbl 0972.32011
- [Mosher 1995] L. Mosher, “Mapping class groups are automatic”, *Ann. of Math. (2)* **142**:2 (1995), 303–384. MR 96e:57002 Zbl 0867.57004
- [Nieminen 1978] J. Nieminen, “The ideal structure of simple ternary algebras”, *Colloq. Math.* **40**:1 (1978), 23–29. MR 80c:20095 Zbl 0415.06002
- [Roller 1998] M. A. Roller, *Poc-sets, median algebras and group actions: an extended study of Dunwoody’s construction and Sageev’s theorem*, Habilitationsschrift, Regensburg, 1998, Available at <http://www.personal.soton.ac.uk/gan/Roller.pdf>.
- [Sholander 1952] M. Sholander, “Trees, lattices, order, and betweenness”, *Proc. Amer. Math. Soc.* **3** (1952), 369–381. MR 14,9b

Received November 16, 2011. Revised July 24, 2012.

BRIAN H. BOWDITCH
MATHEMATICS INSTITUTE
UNIVERSITY OF WARWICK
COVENTRY, CV4 7AL
UNITED KINGDOM

B.H.Bowditch@warwick.ac.uk
<http://www.warwick.ac.uk/~masgak>

PACIFIC JOURNAL OF MATHEMATICS

msp.org/pjm

Founded in 1951 by E. F. Beckenbach (1906–1982) and F. Wolf (1904–1989)

EDITORS

V. S. Varadarajan (Managing Editor)
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
pacific@math.ucla.edu

Paul Balmer
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
balmer@math.ucla.edu

Don Blasius
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
blasius@math.ucla.edu

Vyjayanthi Chari
Department of Mathematics
University of California
Riverside, CA 92521-0135
chari@math.ucr.edu

Daryl Cooper
Department of Mathematics
University of California
Santa Barbara, CA 93106-3080
cooper@math.ucsb.edu

Robert Finn
Department of Mathematics
Stanford University
Stanford, CA 94305-2125
finn@math.stanford.edu

Kefeng Liu
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
liu@math.ucla.edu

Jiang-Hua Lu
Department of Mathematics
The University of Hong Kong
Pokfulam Rd., Hong Kong
jhlu@maths.hku.hk

Sorin Popa
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
popa@math.ucla.edu

Jie Qing
Department of Mathematics
University of California
Santa Cruz, CA 95064
qing@cats.ucsc.edu

Paul Yang
Department of Mathematics
Princeton University
Princeton NJ 08544-1000
yang@math.princeton.edu

PRODUCTION

Silvio Levy, Scientific Editor, production@msp.org

SUPPORTING INSTITUTIONS

ACADEMIA SINICA, TAIPEI
CALIFORNIA INST. OF TECHNOLOGY
INST. DE MATEMÁTICA PURA E APLICADA
KEIO UNIVERSITY
MATH. SCIENCES RESEARCH INSTITUTE
NEW MEXICO STATE UNIV.
OREGON STATE UNIV.

STANFORD UNIVERSITY
UNIV. OF BRITISH COLUMBIA
UNIV. OF CALIFORNIA, BERKELEY
UNIV. OF CALIFORNIA, DAVIS
UNIV. OF CALIFORNIA, LOS ANGELES
UNIV. OF CALIFORNIA, RIVERSIDE
UNIV. OF CALIFORNIA, SAN DIEGO
UNIV. OF CALIF., SANTA BARBARA

UNIV. OF CALIF., SANTA CRUZ
UNIV. OF MONTANA
UNIV. OF OREGON
UNIV. OF SOUTHERN CALIFORNIA
UNIV. OF UTAH
UNIV. OF WASHINGTON
WASHINGTON STATE UNIVERSITY

These supporting institutions contribute to the cost of publication of this Journal, but they are not owners or publishers and have no responsibility for its contents or policies.


See inside back cover or msp.org/pjm for submission instructions.

The subscription price for 2013 is US \$400/year for the electronic version, and \$485/year for print and electronic. Subscriptions, requests for back issues and changes of subscribers address should be sent to Pacific Journal of Mathematics, P.O. Box 4163, Berkeley, CA 94704-0163, U.S.A. The Pacific Journal of Mathematics is indexed by Mathematical Reviews, Zentralblatt MATH, PASCAL CNRS Index, Referativnyi Zhurnal, Current Mathematical Publications and the Science Citation Index.

The Pacific Journal of Mathematics (ISSN 0030-8730) at the University of California, c/o Department of Mathematics, 798 Evans Hall #3840, Berkeley, CA 94720-3840, is published monthly except July and August. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices. POSTMASTER: send address changes to Pacific Journal of Mathematics, P.O. Box 4163, Berkeley, CA 94704-0163.

PJM peer review and production are managed by EditFLOW[®] from Mathematical Sciences Publishers.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2013 Mathematical Sciences Publishers

PACIFIC JOURNAL OF MATHEMATICS

Volume 261 No. 1 January 2013

Hierarchies and compatibility on Courant algebroids	1
PAULO ANTUNES, CAMILLE LAURENT-GENGOUX and JOANA M. NUNES DA COSTA	
A new characterization of complete linear Weingarten hypersurfaces in real space forms	33
CÍCERO P. AQUINO, HENRIQUE F. DE LIMA and MARCO A. L. VELÁSQUEZ	
Calogero–Moser versus Kazhdan–Lusztig cells	45
CÉDRIC BONNAFÉ and RAPHAËL ROUQUIER	
Coarse median spaces and groups	53
BRIAN H. BOWDITCH	
Geometrization of continuous characters of \mathbb{Z}_p^\times	95
CLIFTON CUNNINGHAM and MASOUD KAMGARPOUR	
A note on Lagrangian cobordisms between Legendrian submanifolds of \mathbb{R}^{2n+1}	101
ROMAN GOLOVKO	
On slope genera of knotted tori in 4-space	117
YI LIU, YI NI, HONGBIN SUN and SHICHENG WANG	
Formal groups of elliptic curves with potential good supersingular reduction	145
ÁLVARO LOZANO-ROBLEDO	
Codimension-one foliations calibrated by nondegenerate closed 2-forms	165
DAVID MARTÍNEZ TORRES	
The trace of Frobenius of elliptic curves and the p -adic gamma function	219
DERMOT MCCARTHY	
(DN) - (Ω) -type conditions for Fréchet operator spaces	237
KRZYSZTOF PISZCZEK	



0030-8730(201301)261:1;1-D