Math. Proc. Camb. Phil. Soc. (2013), **154**, 85–95 © *Cambridge Philosophical Society* 2012
doi:10.1017/S0305004112000382
First published online 6 September 2012

Invariance of coarse median spaces under relative hyperbolicity

BY BRIAN H. BOWDITCH
Mathematics Institute, University of Warwick, Coventry, CV4 7AL.
e-mail: B.H.Bowditch@warwick.ac.uk

(*Received* 3 *January* 2012; *revised* 22 *May* 2012)

Abstract

We show that, for finitely generated groups, the property of admitting a coarse median structure is preserved under relative hyperbolicity.

1. *Introduction*

In [**Bo2**], we introduced the notion of a "coarse median group". This is a finitely generated group whose Cayley graph admits a "coarse median" as defined below. The existence of such a median can be thought of as a coarse non-positive curvature condition. Examples of such groups are hyperbolic groups, right-angled Artin groups, mapping class groups (see [**BeM**, **Bo2**]), and direct products of such groups. One can also define a notion of "rank" for such groups. For example, coarse median groups of rank 1 are precisely hyperbolic groups, and the "rank" of a mapping class group is the same as the maximal rank of a free abelian subgroup. Various applications of these notions are discussed in [**Bo2**] and [**Bo3**]. For example, the rank bounds the dimension of a quasi-isometrically embedded euclidean space; groups of finite rank have rapid decay, etc. It implies that the group is finitely presented, and has a quadratic Dehn function. We also note that the existence of a coarse median structure is quasi-isometry invariant.

The main result of this paper is:

THEOREM 1·1. *Suppose that the group* Γ *is hyperbolic relative to the finitely generated subgroups,* $\Gamma_1, \Gamma_2, \ldots, \Gamma_n$. *If each* Γ_i *is coarse median of rank at most* ν*, then so is* Γ.

Here $\nu \in \mathbb{N} \cup \{\infty\}$, and we will deem the statement "of rank at most ∞" to be vacuous. To accomodate the case of a hyperbolic group, when $n = 0$, or when all the Γ_i are finite, we should assume that $\nu \geqslant 1$.

The notion of relative hyperbolicity was defined in [**Gr**]. For other accounts, see [**F**, **Bo1**, **O**]. Note that Theorem 1·1 implies for example that geometrically finite kleinian groups and Sela's limit groups are coarse median.

Although we have expressed the result in terms of groups, it is more naturally a statement about geodesic metric spaces, which we will formulate as Theorem 2·1. In view of the fact that the existence of coarse median is quasi-isometry invariant, we can assume our space to be a connected graph with the combinatorial metric. To define the terms used in these theorems, we need the notion of a finite median algebra. For the purposes of this paper, we can define a finite median algebra in terms of cube complexes. Only very basic properties will be required here.

Let Υ be a finite CAT(0) cube complex (see, for example, [**BrH**]). Let Υ^0 and Υ^1 be the 0 and 1-skeletons of Υ and let ρ_Υ be the combinatorial path-metric on Υ^1. Given $x, y, z \in \Upsilon^0$, there is a unique $w \in \Upsilon^0$ which minimises $\rho_\Upsilon(w, x) + \rho_\Upsilon(w, y) + \rho_\Upsilon(w, z)$. This is the *median* of x, y, z, denoted $\mu_\Upsilon(x, y, z)$. A *finite median algebra* is a finite set, Π, with a ternary operation, μ_Π, such that there is a (necessarily unique) finite cube complex, Υ, and an identification of Π with Υ^0, such that $\mu_\Pi = \mu_\Upsilon$. One can equivalently express this in simple axiomatic terms, see for example, [**BaH**, **R**, **Bo2**]. We just note here that $\mu_\Pi(x, y, z) = \mu_\Pi(y, z, x) = \mu_\Pi(y, x, z)$ and $\mu_\Pi(x, x, y) = x$ for all $x, y, z \in \Pi$. We define the *rank* of Π to be the dimension of Υ. Note that the rank is 1 if and only if Υ is a simplicial tree.

Let (G, ρ) be a geodesic space, that is, a metric space in which every pair of points are connected by a geodesic. (In this paper, G will always be a connected graph, and ρ will be the combinatorial metric assigning each edge unit length.) A *coarse median* on G is a ternary operation satisfying:

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

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

and

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

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

and

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

for all $a \in A$.

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

We say that G has *rank at most* ν if in (C2) we can always choose Π to have rank at most ν.

We refer to (G, ρ, μ) as a *coarse median space (of rank at most* ν*)*, and to k, h as the *parameters* of (G, ρ, μ).

We note that the existence of a coarse median on a geodesic space is a quasi-isometry invariant. Moreover (after modifying μ up to bounded distance), we can assume that $\mu(a, b, c) = \mu(b, c, a) = \mu(b, a, c)$ and that $\mu(a, a, b) = a$ for all $a, b, c \in G$. We will always assume these properties to hold in this paper.

If G is a graph, then it enough to have μ defined on the vertex set, $V(G)$. We can assume that $\mu(V(G)^3) = V(G)$. Moreover, in this case, we can equivalently replace (C1) by the simpler statement:

(C1′): if $a, b, c, d \in V(G)$ with c, d adjacent, then

$$\rho(\mu(a, b, c), \mu(a, b, d)) \leqslant h_0$$

for some fixed $h_0 > 0$.

We recall the definition from [**Bo2**]:

https://doi.org/10.1017/S0305004112000382 Published online by Cambridge University Press

Definition. A *coarse median group (of rank at most ν)* is a finitely generated group whose Cayley graph with respect to a finite generating set admits a coarse median (of rank at most ν).

In view of quasi-isometry invariance, it does not matter which finite generating set we choose. Indeed we could take any locally finite graph on which the group acts properly discontinuously with finite quotient.

2. *Hyperbolic graphs*

In this section, we formulate a statement about graphs which implies Theorem 1·1.

Given a graph H, we will write $V(H)$ and $E(H)$ for the vertex and edge sets. (We will assume there are no loops or multiple edges.) We usually think of H as realised as a metric 1-complex with each edge of unit length. We write ρ_H for the induced combinatorial metric. (If H is not connected, this may take infnite values.)

A *retraction*, $\theta : G \to K$, is a surjective map to a graph, K, which sends each edge of G either to a vertex or to an edge of K. Let $E_0(G) \subseteq E(G)$ be the set of edges which get mapped to edges. Given $t \in V(K)$, write $G(t) \subseteq G$ for the subgraph, $\theta^{-1}(t)$. Note that

$$V(G) = \bigsqcup_{t\in V(K)} V(G(t))$$

and that

$$E(G) = E_0(G) \sqcup \bigsqcup_{t\in V(K)} E(G(t)).$$

We will assume that G is connected, and abbreviate $\rho = \rho_G$. We write $\rho_t = \rho_{G(t)}$ for the path metric induced on $G(t)$. Clearly, $\rho(a, b) \leqslant \rho_t(a, b)$ for all $a, b \in G(t)$. We will assume:

(G1): K is k-hyperbolic for some $k \geqslant 0$.

(G2): there is some function $F_1 : \mathbb{N} \to \mathbb{N}$ such that for all $t \in V(K)$ and for all $a, b \in G(t)$, we have $\rho_t(a, b) \leqslant F_1(\rho(a, b))$.

(G3): there is some $F_2 : \mathbb{N} \to \mathbb{N}$ with the following property. Suppose that $p \in \mathbb{N}$ and that $H \subseteq K$ is any 2-vertex connected subgraph with $|E(H)| \leqslant p$, then the ρ-diameter of $E_0(G) \bigcap \theta^{-1}(H)$ is at most $F_2(p)$.

Thus, (G2) is saying that the graphs $G(t)$ are uniformly uniformly embedded in G. Here the second "uniformly" refers to the standard notion of "uniform embedding" of one metric space in another, and the first "uniformly" means that the relevant parameters are independent of t. We can take (G1) and (G2) to retrospectively imply that G is connected (without needing to take this as hypothesis).

In (G3), $E_0(G) \bigcap \theta^{-1}(H)$ is the set of edges of $E(G)$ which map to edges of H. The term "2-vertex connected" means connected and without a global cut vertex. By this definition, a single edge is 2-vertex connected, so this implies that $E_0(G) \bigcap \theta^{-1}(e)$ has bounded ρ-diameter for all $e \in E(K)$. In fact, in (G3), it's enough to consider only those H which are circuits or single edges. This follows, for example, by noting that in a 2-vertex connected graph, any two distinct edges lie in a circuit.

The main result of this paper can now be stated as:

THEOREM 2·1. *Suppose that $\theta \to K$ is a retraction of graphs satisfying* (G1), (G2) *and* (G3). *Suppose that $G(t)$ is a coarse median space (of rank at most ν) for each $t \in V(K)$. Then G is a coarse median space of rank at most ν.*

In other words, we are assuming that $(G(t), \rho_t)$ admits a coarse median, μ_t, where the parameters, k, h are independent of t. We will construct a coarse median, μ, on (G, ρ) whose parameters depend only on those of $G(t)$ and the hypotheses, (G1)–(G3).

We relate the above to relatively hyperbolic groups via the following observation:

LEMMA 2·2. *Suppose that Γ is hyperbolic relative to the finitely generated subgroups, $\Gamma_1, \Gamma_2, \ldots, \Gamma_n$, where $n > 0$. Then there are connected graphs, G, K, and a retraction, $\theta : G \to K$ satisfying* (G1)–(G3) *above, together with Γ-actions on G and K such that θ is equivariant, G is locally finite, Γ acts freely on G, G/Γ is finite, and such that the vertex stabilisers of K are precisely the Γ-conjugates of $\Gamma_1, \ldots, \Gamma_n$.*

In other words, $\{\Gamma_1, \ldots, \Gamma_n\}$ is a Γ-conjugacy transversal of $\{\Gamma(t) \mid t \in V(K)\}$, where $\Gamma(t) = \{g \in \Gamma \mid gt = t\}$.

In fact, the above gives a characterisation of finitely generated relatively hyperbolic groups, though we only need one direction here. We can weaken the statement that Γ acts freely to say that edge stabilisers are finite. In what follows, we will asssume that each of the Γ_i is infinite. (Otherwise we would be in the case of a hyperbolic group, which is median of rank 1.)

There are several ways one can relate the above to the standard notion. For example, we recall the following notion from [**Bo1**].

We say that a connected graph, K is *fine* if every edge lies in only finitely many circuits of a given length. If there is a bound on this number in terms of the length, we say that K is *uniformly fine*. A group Γ is hyperbolic relative to $\Gamma_1, \ldots, \Gamma_n$ if and only if it acts on a fine hyperbolic graph with finite edge stablisers and finite quotient, and such that $\Gamma_1, \ldots, \Gamma_n$ is a conjugacy transversal of the set of vertex stabilisers, $\{\Gamma(t) \mid t \in V(K)\}$. In such a case, K is necessarily uniformly fine. We want to construct G and $\theta : G \to K$ satisfying (G1), (G2) and (G3). (It's not hard to see that conversely these conditions imply that K is uniformly fine, though we won't need that direction here.)

Suppose then that Γ acts on a fine hyperbolic graph K as above. Given $t \in V(K)$, let $G(t)$ be any Cayley graph of $\Gamma(t)$, and let $\hat{G} = \bigsqcup_{t \in V(K)} G(t)$. We can assume this to be equivariant with respect to a Γ-action on $\hat{G}$, so that for all $g \in \Gamma$, $G(gt) = gG(t)$. Thus, $g\Gamma(t)g^{-1}$ acts on $G(gt)$. One way to achieve this is to choose a finite generating set S_i for each Γ_i and let G_i be the (disconnected) "Cayley graph" of Γ with respect to S_i. (That is, $V(G_i) \equiv \Gamma$ and $g, h \in V(G_i)$ are adjacent if $g^{-1}h \in S_i$.) Now let $\hat{G} = \bigsqcup_{i=1}^{n} G_i$. This comes equipped with a Γ-action. By construction, the setwise stabiliser of each connected component of $\hat{G}$ is a Γ-conjugate of one of the Γ_i, and is therefore also the stabiliser of a unique vertex of K. This gives us a canonical, Γ-equivariant surjection, $\theta : \hat{G} \to V(K)$.

Now let $E' \subseteq E(K)$ be a finite Γ-transversal of edges. For each e, we add an edge, $f(e)$, from a vertex of $G(t)$ to a vertex of $G(u)$, where $t, u \in V(K)$ are the endpoints of e. We now extend this Γ-equivariantly to give us a connected graph $G \supseteq \hat{G}$, and a Γ-equivariant extension $\theta : G \to K$.

Property (G1) is given, and (G2) is easily verified. For (G3), note that in a fine graph there are only finitely many 2-vertex connected graphs of any given cardinality containing any given edge. (Note that any two distinct edges of a 2-vertex connected graph are contained in a circuit.) In our situation, we see that there are only finitely many Γ-orbits of 2-vertex connected graphs of any given cardinality. We also note that if $e \in E(K)$ then $E_0(G) \cap \theta^{-1}(e)$ is the $(\Gamma(t) \cap \Gamma(u))$-orbit of a single edge, where $t, u \in V(K)$ are the endpoints of e. In particular, this is finite. Property (G3) now follows easily.

We will prove Theorem 2·1 in the remainder of this paper. We first make some preliminary observations. First, there is no loss in assuming:

(G4): $\theta : E_0(G) \to E(K)$ is bijective.

To see this, we select one edge from the preimage of each $e \in E(K)$ and delete the rest. It follows from the fact that such a preimage has bounded diameter that the inclusion of the resulting graph into the original is a quasi-isometry. (Here we are using (G3) applied to a single edge of K, as well as (G2).) Moreover, the existence of a coarse median on a space is quasi-isometry invariant. (Note that this process does not need be carried out in an equivariant fashion.)

We introduce the following notation. We will write $\hat{\rho}$ for the (possibly infinite) path metric on $\hat{G}$. In other words, $\hat{\rho}(x, y) = \rho_t(x, y)$ if there is some $t \in V(K)$ with $x, y \in G(t)$, and $\hat{\rho}(x, y) = \infty$ otherwise.

Given $e \in E(K)$, we write $\tilde{e} \in E_0(G)$ for its preimage under θ. Suppose that α is a non-trivial path in K. We write $\epsilon(\alpha)$ for the initial edge of α, and $\tilde{\epsilon}(\alpha)$ for its preimage in $E_0(G)$. We write $q(\alpha) = \tilde{\epsilon}(\alpha) \cap G(t) \in V(G(t))$, where t is the initial vertex of α.

If μ is any ternary operation on a set, we refer to a subset closed under μ as a *subalgebra* (without making any assumptions on μ). We refer to a map respecting ternary operations as a *homomorphism*. We define *epimorphism* and *isomorphism* in the obvious way.

3. *Trees of spaces*

We first prove Theorem 2·1 in the case where $K = T$ is a finite simplicial tree. Given $t, u \in V(T)$, write $[t, u]$ for the unique arc from t to u. Then $V(T)$ has the structure of a median algebra, where the median, μ_T is defined by $[t, u] \cap [u, v] \cap [v, t] = \{\mu_T(t, u, v)\}$. In other words, $\mu_T(u, v, w)$ is the centre of the "tripod" spanned by t, u, v. We also note that if $M \subseteq V(T)$ is any subalgebra, then we can identify M as the vertex set, $V(T_M)$, of another tree T_M obtained from T as follows. First, take the tree, T', spanned by M (i.e. the smallest subtree containing M). Then remove, from T', each degree-2 vertex of T' that is not in M and coalesce the incident edges to give us T_M. The median μ_{T_M} agrees with μ_T on M.

Suppose now that G is a connected graph with a map, $\theta : G \to T$, satisfying (G4) above. In this case, properties (G1), (G2) and (G3) are automatic. In particular, ρ_t agrees with ρ on each $G(t)$.

If $t \in V(T)$, then there is a well defined nearest point retraction, $\phi_t : G \to G(t)$. In fact, if $a \in V(G(t))$, then $\phi_t(a) = a$, and if $a \in V(G) \setminus V(G(t))$, then $\phi_t(a) = q(\alpha)$, where $\alpha = [t, \theta(a)]$.

Now given $a, b, c \in V(G)$, let $\mu(a, b, c) = \mu_t(\phi_t a, \phi_t b, \phi_t c)$, where $t = \mu_T(\theta a, \theta b, \theta c)$. Then $\mu : V(G)^3 \to V(G)$. By construction, we have $\theta\mu(a, b, c) = \mu_T(\theta a, \theta b, \theta c)$ for all $a, b, c \in V(G)$. (In other words, θ is a homomorphism.)

For future reference, we note that if $M \subseteq V(T)$ is a subalgebra, we can define a retraction of graphs, $\theta_M : G_M \to T_M$, as follows. We take the span T' of M in T as above. This gives us $\theta : \theta^{-1}(T') \to T'$. We now collapse to a point each graph $G(u)$ for $u \in V(T') \setminus M$. Each such vertex u has degree 2 in T', so we can coalesce the two edges in $E_0(G) \cap \theta^{-1}(T')$

https://doi.org/10.1017/S0305004112000382 Published online by Cambridge University Press

meeting $G(u)$. This gives us our graph G_M, with a natural map, $\theta_M : G_M \to T_M$. This satisfies (G4). Moreover, the nearest point retraction $\phi_t : G_M \to G(t)$ defined intrinsically to G_M agrees with the map induced from G. In particular, if $a, b, c \in G_M$, then the median $\mu(a, b, c)$ lies in G_M, and agrees with that defined intrinsically to G_M.

We will show in this section that μ is a coarse median on G. In fact, we can make a stronger assertion. Recall that $\hat{\rho}$ is the (possibly infinite) metric on $\hat{G} = \bigsqcup_{t \in V(\tau)} G(t) \subseteq G$.

LEMMA 3·1. *Let $\theta : G \to T$ be a tree of spaces satisfying* (G3) *and* (G4)*, and such that $(G(t), \rho_t)$ admits a coarse median, μ_t, with uniform parameters (independent of t). Let μ be the ternary operation defined as above. Then:*

(CT1)*: there is some $h_0 \geqslant 0$ such that if $a, b, c, d \in V(G)$ with c, d adjacent, then either $\hat{\rho}(\mu(a, b, c), \mu(a, b, d)) \leqslant h_0$ or c, d are the endpoints of an edge of $E_0(G)$ and $\mu(a, b, c) = c$ and $\mu(a, b, d) = d$; and*

(CT2)*: in* (C2) *we make the stronger statements that*

$$\hat{\rho}(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z)) \leqslant h(p)$$

and that $\hat{\rho}(a, \lambda\pi a) \leqslant h(p)$.

Note that (CT1) implies (C1′) which implies (C1), and that (CT2) implies (C2).

The statement of Lemma 3·1 is taken to imply that if each $(G(t), \mu_t)$ has rank at most $\nu > 0$, then so does (G, μ).

We first note:

LEMMA 3·2. *Lemma* 3·1 *holds if T consists of a single edge.*

Proof. Let $E(T) = \{e\}$ and $V(T) = \{t_1, t_2\}$. We write $G_i = G(t_i)$ and $\mu_i = \mu_{t_i}$. Let $q_i = G_i \cap \tilde{e}$. Thus G is obtained from $G_1 \sqcup G_2$ by connecting $q_1 \in G_1$ to $q_2 \in G_2$ by a single edge $\tilde{e}$. Thus, $V(G) = V(G_1) \sqcup V(G_2)$.

Suppose that $a, b, c \in V(G)$. By construction, if $a, b, c \in G_1$, then $\mu(a, b, c) = \mu_1(a, b, c)$ and if $a, b \in G_1$, $c \in G_2$ then $\mu(a, b, c) = \mu_1(a, b, q_1)$. All other cases arise by permuting a, b, c and/or swapping G_1 and G_2.

We claim that μ satisfies the conclusion of Lemma 3·1

For (CT1) suppose that $a, b, c, d \in V(G)$, with c, d adjacent. If $c, d \in G_1$, then (C1′) in G_1 tells us that $\hat{\rho}(\mu(a, b, c), \mu(a, b, d))$ is bounded. This holds similarly if $c, d \in G_2$. Thus, we can suppose that $c \in G_1$ and $d \in G_2$, so that $c = q_1$ and $d = q_2$. If $a, b \in G_1$, then $\mu(a, b, c) = \mu(a, b, q_1) = \mu(a, b, d)$, and similarly, if $a, b \in G_2$. If $a \in G_1$ and $b \in G_2$, then $\mu(a, b, c) = \mu(a, q_1, q_2) = \mu_1(a, q_1, q_2) = q_1 = c$ and $\mu(a, b, d) = \mu(b, q_1, q_2) = \mu_2(b, q_1, q_2) = q_2 = d$.

For (CT2), suppose that $A \subseteq V(G)$, with $|A| \leqslant p$. Let $A_i = A \cap V(G_i)$, so $A = A_1 \sqcup A_2$. Let $B_1 = A_1 \cup \{q_1\}$ and $B_2 = A_2 \cup \{q_2\}$. Let $\pi_i : B_i \to \Pi_i$ and $\lambda_i : \Pi_i \to V(G_i)$ be the maps given by (C2) for G_i. Let $v_i = \pi_i(q_i) \in \Pi_i$. Let $B = B_1 \cup B_2$.

Now $\Pi_i = V(\Upsilon_i)$, where Υ_i is a finite CAT(0) cube complex. Let Υ be the cube complex obtained from $\Upsilon_1 \sqcup \Upsilon_2$ by adding an edge from v_1 to v_2. This is also a CAT(0) cube complex whose dimension is the maximum of 1 and those of Υ_1 and Υ_2. Thus, $\Pi = V(\Upsilon)$ is a finite median algebra, with $\Pi = \Pi_1 \sqcup \Pi_2$, and with Π_i a subalgebra.

We define $\pi : A \to \Pi$ and $\lambda : \Pi \to G$ by combining the maps π_1, π_2 and λ_1, λ_2.

Note that if $a \in B$, then $\rho_i(a, \lambda_i\pi_i a) \leqslant h(p)$. In particular, if $a \in A_i$, then $\hat{\rho}(a, \lambda\pi a) \leqslant h(p)$. Also $\rho_i(q_i, \lambda v_i) \leqslant h(p)$.

https://doi.org/10.1017/S0305004112000382 Published online by Cambridge University Press

Suppose now that $x, y, z \in \Pi$. We want to bound $\hat{\rho}(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z))$. If $x, y, z \in \Pi_i$, then the result follows directly from the statement for G_i. Thus, without loss of generality, we can assume that $x, y \in \Pi_1$ and $z \in \Pi_2$ so that $\lambda x, \lambda y \in V(G_1)$, $\lambda z \in V(G_2)$. Now, by construction, $\mu_\Pi(x, y, z) = \mu_\Pi(x, y, v_1)$ and $\mu(\lambda x, \lambda y, \lambda z) = \mu_1(\lambda x, \lambda y, q_1)$. But by (C2) in G_1, $\rho_1(\mu(\lambda x, \lambda y, \lambda v_1), \mu(\lambda x, \lambda y, q_1))$ is bounded, and by (C1) in G_1, $\rho_1(\lambda\mu_\Pi(x, y, v_1), \mu(\lambda x, \lambda y, \lambda v_1))$ is bounded. Putting these together, we bound $\hat{\rho}(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z))$ as required.

Proof of Lemma 3·1. First, we prove a slightly weaker version in that we allow the function in (CT2) to depend on $n = |E(T)|$ as well as on the parameters of $G(t)$ and (G1)–(G3). That is, we have a bound $h_n(p)$, where $h_n : \mathbb{N} \to \mathbb{N}$.

Let $t_1 \in V(T)$ be an extreme (degree-1) vertex. Let $e \in E(T)$ be the incident edge, let $t_2 \in V(T)$ be the adjacent vertex, and let $T_0 \subseteq T$ be the subtree $T_0 = T \setminus e$. Let $\theta_0 : G \to e$ be defined by sending $\tilde{e}$ to e, $G(t_1)$ to t_1 and $\theta^{-1}(T_0)$ to t_2. Thus, $\theta_0 : G \to e$ is a tree of spaces of the sort described by Lemma 3·1, and so it satisfies (CT1) and (CT2). Also, by induction, we can assume these also hold for $\theta : \theta^{-1}(T_0) \to T_0$. Putting these together now gives us (CT1) and (CT2) for $\theta : G \to T$, though the constants of (CT2) may have increased, giving us our dependence on n.

To remove dependence on n, we make the following observation. Suppose $A \subseteq V(G)$ with $|A| \leqslant p$. Let $M \subseteq V(T)$ be the median algebra generated by $\theta(A)$. Then $|M| \leqslant 2p-2$. Let $\theta_M : G_M \to T_M$ be the corresponding tree of spaces. Now apply (CT2) to $A \subseteq V(G_M)$. This gives us $\pi : A \to \Pi$ and $\lambda : \Pi \to V(G_M)$ satisfying (CT2) with the bound $h_{2p-2}(p)$. But now the definitions of μ and $\hat{\rho}$, intrinsic to G_M, agree those obtained by restricting the definitions in G. Thus, (CT2) follows in G where we set $h(p) = h_{2p-2}(p)$.

We will use the idea of the last paragraph of the proof again in Section 4. One could give a proof of Lemma 3·1 without using induction, by constructing Π as the vertex set of a tree of cube complexes, though this seems more complicated to write out formally.

4. *Hyperbolic spaces*

Let (K, ρ_K) be a k-hyperbolic graph (see **[Gr]**, **[GhH]**). We write $\text{hd}(P, Q)$ for the Hausdorff distance between $P, Q \subseteq K$.

Definition. Given $l \geqslant 0$, we say that a path, α, in K is *l-taut* if $\text{length}(\alpha) \leqslant \rho(u, v) + l$, where u, v are the endpoints of α.

Note that any subpath of an l-taut path is l-taut, and that a 0-taut path is the same as a geodesic.

LEMMA 4·1. *Given $l, s \geqslant 0$, there is some $r_1 = r_1(l, s, k)$ with the following property. Suppose that α, α', are l-taut paths with endpoints u, v and u', v' respectively, and that $\rho_K(u, u') \leqslant s$ and $\rho(v, v') \leqslant s$. Then* $\text{hd}(\alpha, \alpha') \leqslant r_1$.

Proof. Note that taut paths are quasigeodesic, so the lemma is a simple consequence of the "fellow travelling" property of quasigeodesics in a hyperbolic space.

Definition. An *l-taut tree* is a simplicial tree, T, embedded in K such that each arc in T is l-taut in K.

LEMMA 4·2. *There is a function, $l_0 : \mathbb{N} \to \mathbb{N}$ such that if $B \subseteq K$ with $|B| \leqslant p < \infty$, then there is an l-taut tree, T, embedded in K, with $l = kl_0(p)$.*

https://doi.org/10.1017/S0305004112000382 Published online by Cambridge University Press

Proof. This is just a rephrasing of a standard fact due to Gromov [**Gr**].

We view T as a subgraph of K, so $V(T) = T \cap V(K)$. (It may have lots of degree-2 vertices.)

Note that there is no loss in assuming that T is spanned by B (i.e. is the minimal subtree containing B).

Definition. A *tripod* is a tree $\tau \subseteq K$ consisting of three arcs, $\alpha_1, \alpha_2, \alpha_3$, each starting at a single vertex, t, in $V(K)$.

We refer to $t = t(\tau)$ as the *centre* of the tripod, and to the other endpoints, u_1, u_2, u_3 of the arcs $\alpha_1, \alpha_2, \alpha_3$ as its *feet*. We assume that these are also vertices of K. (We allow the α_i to be trivial. Note, however, that if the u_i are distinct, then at most one of the α_i can be trivial.)

Writing $l_3 = kl_0(3)$, we see that any three points are feet of some l_3-taut tripod in K.

LEMMA 4·3. *Given $l \geqslant 0$, there is a constant, $r_2 = r_2(l, k) \geqslant 0$ with the following property. Suppose that $t, u \in V(K)$ are distinct, and that α, α' are l-taut arcs connecting t to u, with edges, e and e' incident on t. Then either $e = e'$, or else there is a (possibly empty) arc δ in K and initial segments, β, β' of α, α', respectively, such that $\gamma = \beta \cup \delta \cup \beta'$ is an (embedded) circuit in K of length at most r_2.*

Proof. Suppose that $e \neq e'$. By Lemma 4·1, $\mathrm{hd}(\alpha, \alpha') \leqslant r_1 = r_1(l, k)$. Let $v \in V(\alpha)$ be the first vertex of α also contained in $V(\alpha')$. If $\rho_K(t, v) \leqslant r_1$, we set $\delta = \emptyset$ and set β, β' to be the respective initial segments ending at v. Note that these have length at most $r_1 + l$, so $\mathrm{length}(\gamma) \leqslant 2(r_1 + l)$.

Now suppose that $\rho_K(t, v) > r_1$. Let $w_0 \in V(\alpha)$ be the first vertex of α with $\rho_K(t, w) = r_1 + 1$. Let $w' \in V(\alpha')$ be the nearest vertex of α' to w_0. Let δ_0 be any geodesic in K from w_0 to w'. Let $w \in V(\alpha) \cap V(\delta_0)$ be the last vertex of $V(\alpha)$ along δ_0, and let $\delta \subseteq \delta_0$ be the segment of δ from w to w'. Let β, β' be the respective intitial segments of α, α' ending at w and w', and let $\gamma = \beta \cup \delta \cup \beta'$. Then $\mathrm{length}(\gamma) \leqslant 2(r_1 + 1 + l) + r_1 = 3r_1 + 2l + 2$, so we set $r_2 = 3r_1 + 2l + 2$.

Note that if $e \neq e'$, then $t \in \gamma \setminus \delta$, so e, e' are edges of γ.

LEMMA 4·4. *Given $l, s \geqslant 0$, there is some $r_3 = r_3(l, s, k)$ with the following property. Suppose that $t, t', u \in V(K)$ with $t \neq t'$, and $\rho_K(t, t') \leqslant s$. Let ζ be any geodesic from t to t'. Suppose that α, α' are l-taut paths which connect t and t' respectively to u. Then there is a (possibly empty) arc δ in K with $\delta \cap \zeta = \emptyset$, and initial segments, β, β' of α, α' respectively, such that $\gamma = \beta \cup \delta \cup \beta'$ is an arc from t to t' of length at most r_3.*

Proof. This follows by a similar argument to Lemma 4·3. Note that $\mathrm{hd}(\alpha, \alpha') \leqslant r_1 = r_1(l, s, k)$. We let v be the first vertex of $V(\alpha) \cap V(\alpha')$ along α, as before. This time, we split into two cases depending on whether or not $\rho_K(t, v) \leqslant r_1 + s$, and proceed as before to give us a path $\gamma = \beta \cup \delta \cup \beta'$. This time, we set $r_3 = 2(r_1 + s + 1 + l) + r_1 = 3r_1 + 2s + 2l + 2$.

Now suppose that $\theta : G \to K$ satisfies (G1)–(G4) as defined in Section 2. We recall the notations, $\tilde{e}$, $\tilde{\alpha}$, $q(\alpha)$ from there. In what follows, the various constants, or functions, we refer to will be implicitly assumed to depend on the parameters of the hypotheses (G1)–(G3).

LEMMA 4·5. *Given $l \geqslant 0$, there is some $r_4 = r_4(l)$ with the following property. Suppose that α, α' are l-taut arcs in K with the same endpoints $t, u \in V(K)$, where $t \neq u$. Then $\rho_t(q(\alpha), q(\alpha')) \leqslant r_4$.*

Proof. Let e, e' be the incident edges. If $e = e'$, then $q(\alpha) = q(\alpha')$, so we assume $e \neq e'$. Let γ be the circuit given by Lemma 4·3. Now, $e, e' \in E(\gamma)$, so by (G3), $\rho(\tilde{e}, \tilde{e}') \leqslant F_2(r_2(k, l))$. By definition, $q(\alpha) = \tilde{e} \cap G(t)$ and $q(\alpha') = \tilde{e}' \cap G(t)$, so by (G2) we have $\rho_t(q(\alpha), q(\alpha')) \leqslant F_1(F_2(r_2(l, k)))$.

Suppose that $a_1, a_2, a_3 \in V(G)$. Let $u_i = \theta(a_i) \in V(K)$. Suppose that $\tau = \alpha_1 \cup \alpha_2 \cup \alpha_3$ is a tripod with feet at u_1, u_2, u_3. Let $t = t(\tau)$ be the centre of τ. If $u_i = t$, set $q_i = a_i$, otherwise set $q_i = q(\alpha_i)$. Thus, $q_1, q_2, q_3 \in G(t)$. Let $\mu(a_1, a_2, a_3; \tau) = \mu_t(q_1, q_2, q_3) \in G(t)$.

In the rest of this section, we will use the abbreviation "**a**" to denote (a_1, a_2, a_3) etc. Thus, for example, we can rewrite the above as $\mu(\mathbf{a}; \tau) = \mu_t(\mathbf{q})$.

LEMMA 4·6. *There is some $r_5 = r_5(l)$ with the following property. Suppose that $a_1, a_2, a_3 \in V(G)$, and suppose that τ, τ' are each l-taut spanning tripods with feet $\theta(a_1), \theta(a_2), \theta(a_3)$. Then $\rho(\mu(\mathbf{a}; \tau), \mu(\mathbf{a}; \tau')) \leqslant r_5$.*

We will split the proof into two cases. The first gives a slightly stronger statement in the case where $t(\tau) = t(\tau')$.

LEMMA 4·7. *There is some $r_6 = r_6(l)$ with the following property. Suppose that* **a**, *τ, τ' are as in Lemma* 4·6, *and that $t = t(\tau) = t(\tau')$. Then $\rho_t(\mu(\mathbf{a}; \tau), \mu(\mathbf{a}; \tau')) \leqslant r_6$.*

Proof. Let $q_i, q_i' \in G(t)$ be as in the definitions of $\mu(\mathbf{a}; \tau)$ and $\mu(\mathbf{a}; \tau')$ respectively. By Lemma 4·5, we see that $\rho_t(q_i, q_i') \leqslant r_4$. Thus, by (C1) in $G(t)$, we see that $\rho_t(\mu(\mathbf{a}; \tau), \mu(\mathbf{a}; \tau'))$ is bounded.

For the case where $t(\tau) \neq t(\tau')$, we will need the following two general lemmas about tripods in K. Note that in this case, the feet, $u_i = \theta(a_i)$ must all be distinct.

Suppose that $\tau = \alpha_1 \cup \alpha_2 \cup \alpha_3$ and that $\tau' = \alpha_1' \cup \alpha_2' \cup \alpha_3'$ are l-taut tripods each with feet at $u_1, u_2, u_3 \in V(G)$. Let $t = t(\tau)$ and $t' = t(\tau')$.

LEMMA 4·8. *If τ, τ' are l-taut, then $\rho(t, t') \leqslant s_1$, where $s_1 = s_1(l, k)$ depends only on l and k.*

Proof. This is a simple consequence of hyperbolicity. Note that each of the paths $\alpha_i \cap \alpha_j$ and $\alpha_i' \cap \alpha_j'$ are l-taut, and therefore remains a bounded distance from any geodesic which connects the same endpoints. It follows that t and t' are each a bounded distance from the centre of any geodesic triangle in K with vertices at u_1, u_2, u_3.

We suppose that $t \neq t'$, so that the u_i are all distinct. Let ζ be any geodesic from t to t'. Let $\gamma_i = \beta_i \cup \delta_i \cup \beta_i'$ be the arc from t to t' given by Lemma 4·4 (with $\alpha = \alpha_i$, $\alpha' = \alpha_i'$ and $u = u_i$). Thus $\delta_i \cap \zeta = \emptyset$, and $\text{length}(\gamma_i) \leqslant r_7(l)$, where $r_7(l) = r_4(l, s_1(k, l), k)$. Let $L \subseteq K$ be the image of $\zeta \cup \gamma_1 \cup \gamma_2 \cup \gamma_3$ in K. Note that $|E(L)| \leqslant s_1(k, l) + 3r_7(l)$.

LEMMA 4·9. *L is 2-vertex connected.*

Proof. Suppose that $v \in L$ were a cut point of L. Since ζ and each γ_i is an arc, v must separate t from t' in L. Thus, $v \in \zeta \setminus \{t, t'\}$, and so $v \notin \delta_i$. For each i, we have $v \in \gamma_i \setminus \delta_i =$

https://doi.org/10.1017/S0305004112000382 Published online by Cambridge University Press

$\beta_i \cup \beta'_i \subseteq \alpha_i \cup \alpha'_i$. It follows that v must lie in at least two of the α_i or at least two of the α'_i. We respectively arrive at the contradictions $v = t$ or $v = t'$.

Now suppose that $\mathbf{a}, \tau, \tau'$ are as in the hypotheses of Lemma 4·6. Let $u_i = \theta(a_i)$, and let q_i, q'_i be as in the definitions of μ, so that $\mu(\mathbf{a}; \tau) = \mu_t(\mathbf{q})$ and $\mu(\mathbf{a}; \tau') = \mu_{t'}(\mathbf{q}')$. We suppose that $t \neq t'$.

LEMMA 4·10. *There is some $r_8 = r_8(l)$ such that $\rho(\mu(\mathbf{a}; \tau), \mu(\mathbf{a}; \tau')) \leqslant r_8(l)$.*

Proof. Let $L = \zeta \cup \gamma_1 \cup \gamma_2 \cup \gamma_3$ as above. By Lemma 4·9, L is 2-vertex connected. Note that if $\beta_i = \{t\}$, then since $\delta_i \cap \zeta = \emptyset$, we must have $\delta_i = \emptyset$, and so $t \in \beta'_i$. This can hold for at most one i. In other words, at most one of the β_i can be trivial, so we can suppose that β_1 and β_2 are non-trivial. Let e_1 and e_2 be the initial edges of β_1 and β_2. Since e_1 and $e_2 \in E(L)$, by (G3), we see that $\rho(\tilde{e}_1, \tilde{e}_2) \leqslant s_2$, where $s_2 = s_2(l) = F_2(s_1(k, l) + 3r_7(l))$. By definition, $q_1 = \tilde{e}_1 \cap G(t)$ and $q_2 = \tilde{e}_2 \cap G(t)$, so $\rho(q_1, q_2) \leqslant s_2 + 2$, and so, by (G2), $\rho_t(q_1, q_2) \leqslant F_1(s_2 + 2)$. By (C1) applied to μ_t, we get that $\rho_t(q_1, \mu_t(\mathbf{q})) \leqslant s_3$ and $\rho_t(q_2, \mu_t(\mathbf{q})) \leqslant s_3$, where $s_3 = s_3(l)$ depends only on l (and the parameters of the hypotheses).

Now, without loss of generality, we also have β'_1, β'_j non-trivial, where $j \in \{2, 3\}$. Thus, by a similar argument applied to τ', we get $\rho_{t'}(q'_1, \mu_{t'}(\mathbf{q}')) \leqslant s_3$. Moreover, $e_1, e'_1 \in E(L)$, where e'_1 is the initial edge of β'_1. Thus, we also get $\rho(q_1, q'_1) \leqslant s_2 + 2$. This therefore places a bound on $\rho(\mu_t(\mathbf{q}), \mu_{t'}(\mathbf{q}'))$ as required.

Lemmas 4·7 and 4·10 together give Lemma 4·6.

We can now define medians in G.

Given $a_1, a_2, a_3 \in V(G)$, choose τ to be any l_3-taut tripod with feet at $\theta(a_1), \theta(a_2), \theta(a_3)$. We set $\mu(a_1, a_2, a_3) = \mu(\mathbf{a}) = \mu(\mathbf{a}; \tau)$. Note that, by Lemma 4·6, this is well defined up to a bounded distance $r_9 = r_5(l_3)$, depending only on the parameters of the hypotheses.

Note that in the case where K is a tree, τ is unique. Moreover, in this case, $q_i = \phi_t(a_i)$, where $t = t(\tau)$. Thus, this definition agrees with that given for trees in Section 3.

Now suppose that $T \subseteq K$ is an embedded l-taut tree. Suppose that $M \subseteq V(\tau)$ is some median subalgebra of $V(T)$. Then we can identify $M = V(T_M)$ for the tree, T_M, described in Section 3. Moreover, we have $\theta : \theta^{-1}(T) \to T$, and $\theta : G_M \to T_M$, with $\hat{G}_M = \bigsqcup_{t \in M} G(t)$. Note that $\hat{\rho}_M$ agrees with $\hat{\rho}$ on $\hat{G}_M$.

As discussed in Section 3, we have a median, μ_M, defined on G_M, satisfying (CT1) and (CT2).

LEMMA 4·11. *Suppose that $T \subseteq K$ is an l-taut tree, and that $M \subseteq V(T)$ is a median subalgebra. Let μ_M be the median defined on G_M. If $a_1, a_2, a_3 \in V(G_M)$, then $\rho(\mu(\mathbf{a}), \mu_M(\mathbf{a})) \leqslant r_5(l)$.*

Proof. We can suppose that $l \geqslant l_3$. Let $u_i = \theta(a_i)$. Let $\tau \subseteq K$ be the tripod used in the definition of μ, that is, $\mu(\mathbf{a}) = \mu(\mathbf{a}; \tau)$. Let $\tau' \subseteq T$ be the tripod spanned by u_1, u_2, u_3. This is l-taut in K. By construction of μ_M we have $\mu_M(\mathbf{a}) = \mu(\mathbf{a}; \tau')$. Lemma 4·6 now tells us that $\rho(\mu(\mathbf{a}; \tau), \mu(\mathbf{a}; \tau')) \leqslant r_5(l)$ as required.

Proof of Theorem 2·1. We prove (C1′) and (C2).

(C1′) Let μ be the median defined on $V(G)$ as above. Suppose that $a, b, c, d \in V(G)$, with c, d adjacent. Let $t = \theta(a)$, $u = \theta(b)$, $v = \theta(c)$ and $w = \theta(d)$. Suppose first that $v = w$. Let $T \subseteq K$ be an l_3-taut tripod spanning $\{t, u, v\}$. Let $M \subseteq V(T)$ be the median

https://doi.org/10.1017/S0305004112000382 Published online by Cambridge University Press

algebra spanned by $\{t, u, v\}$ (so that $|M| \leqslant 4$), and let $\theta_M : G_M \to T_M$ be the corresponding tree of graphs as in Lemma 4·11. Thus, $\rho(\mu(a, b, c), \mu_M(a, b, c)) \leqslant r_9 = r_5(l_3)$ and $\rho(\mu(a, b, d), \mu_M(a, b, d)) \leqslant r_9$. Lemma 3·1 tells us that (CT1) holds in G_M, and so

$$\rho(\mu_M(a, b, c), \mu_M(a, b, d)) \leqslant \hat{\rho}(\mu_M(a, b, c), \mu_M(a, b, d)) \leqslant h_0.$$

Thus, $\rho(\mu(a, b, c), \mu(a, b, d)) \leqslant h_0 + 2r_9$.

The case where c, d are the endpoints of an edge, $\tilde{e} \in E_0(G)$ is similar. Let $e = \theta(\tilde{e}) \in E(K)$. This has endpoints $v, w \in V(K)$. We can easily construct an $(l_3 + 2)$-taut tripod $T \subseteq K$, with $t, u, v, w \in V(T)$ and with $e \in E(T)$. (Start with an l_3-taut tripod for $\{t, u, v\}$, and suppose that it does not already contain e. If it does not contain w, then add in e. If it does contain w, we can assume that w lies in the arc from t to u, and we can divert this to pass through e using a leg of the tripod. In this case, we end up with an arc from t to u containing v and w.) Let $M \subseteq V(T)$ be the median algebra spanned by $\{t, u, v, w\}$ (so that $|M| \leqslant 5$). Let $\theta_M : G_M \to T_M$ be the corresponding tree of graphs. By construction, c, d are also adjacent in G_M. We now proceed similarly as before, applying (CT1) to G_M.

(CT2): let $A \subseteq V(G)$, with $|A| \leqslant p$. Let $B = \theta(A) \subseteq V(K)$ and let $T \subseteq K$ be a $(kl_0(p))$-taut tree with $B \subseteq V(T)$, as given by Lemma 4·2. Let $M \subseteq V(T)$ be the median algebra generated by B. Let $\theta_M : G_M \to T_M$ be the associated tree of graphs. Let $\pi : A \to \Pi$ and $\lambda : \Pi \to V(G_M) \subseteq V(G)$ be the maps given by (CT2) for G_M as in Lemma 3·1 If $x, y, z \in \Pi$, then

$$\rho(\lambda\mu_\Pi(x, y, z), \mu_M(\lambda x, \lambda y, \lambda z)) \leqslant \hat{\rho}(\lambda\mu_\Pi(x, y, z), \mu_M(\lambda x, \lambda y, \lambda z)) \leqslant h(p).$$

By Lemma 4·11, $\rho(\mu(\lambda x, \lambda y, \lambda z), \mu_M(\lambda x, \lambda y, \lambda z)) \leqslant r_5(kl_0(p))$, and so $\rho(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z)) \leqslant h(p) + r_5(kl_0(p))$ which depends only on p and the parameters.

Finally, if $a \in A$, then $\rho(a, \lambda\pi a) \leqslant \hat{\rho}(a, \lambda\pi a) \leqslant h(p)$.

REFERENCES

[BaH] H.-J. BANDELT and J. HEDLIKOVA. Median algebras. *Discrete Math.* **45** (1983), 1–30.

[BeM] J. A. BEHRSTOCK and Y. N. MINSKY. Centroids and the rapid decay property in mapping class groups. Preprint (2008).

[Bo1] B. H. BOWDITCH. Relatively hyperbolic groups. To appear in *Internat. J. Algebra Comput.*

[Bo2] B. H. BOWDITCH. Coarse median spaces and groups. To appear in *Pacific J. Math.*

[Bo3] B. H. BOWDITCH. Embedding median algebras in products of trees. Preprint, Warwick (2011).

[BrH] M. BRIDSON and A. HAEFLIGER. Metric spaces of non-positive curvature. *Grundlehren der Math. Wiss.* No. 319 (Springer 1999).

[F] B. FARB. Relatively hyperbolic groups. *Geom. Funct. Anal.* **8** (1998), 810–840.

[GhH] E. GHYS and P. DE LA HARPE (eds.). Sur les groupes hyperboliques d'après Mikhael Gromov. *Prog. Math.* 83 (Birkhäuser 1990).

[Gr] M. GROMOV. Hyperbolic groups. In *Essays in Group Theory. Math. Sci. Res. Inst. Publ.* No. 8 (Springer 1987), 75–263.

[O] D. V. OSIN. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.* **179** (2006).

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