QUASIFLATS IN COARSE MEDIAN SPACES

BRIAN H. BOWDITCH

Abstract. We describe the geometry of a top-dimensional quasiflat (that is, a quasi-isometrically embedded copy of euclidean space of maximal possible dimension) in a coarse median space of finite rank. We show that such a quasiflat is a bounded Hausdorff distance from a finite union of subsets, each of which has a simple structure. In particular, each of these subsets is the image of a direct product of real intervals or rays under a quasi-isometric embedding which preserves medians up to bounded distance. As one consequence, we recover the result of Behrstock, Hagen and Sisto, regarding quasiflats in asymphoric hierarchically hyperbolic spaces. In the case of a median metric space, one can strengthen the conclusion. In particular, we recover the result of Huang regarding quasiflats in CAT(0) cube complexes, namely that a top-dimensional quasiflat is a finite Hausdorff distance from a finite union of orthants. In fact, we obtain some strengthening of both these results in the form of uniformity of various parameters involved. Examples of coarse median spaces include the mapping class groups of surfaces, as well as Teichmüller space in either the Teichmüller or the Weil-Petersson metric.

1. Introduction

In this paper, we study quasiflats in coarse median spaces. A "quasiflat" is a quasi-isometrically embedded copy of euclidean space, $\mathbb{R}^\nu$, which we will assume here to have maximal possible dimension, ν. A "quasi-isometric embedding" is a map which preserves distances to within linear bounds. The notion of a "coarse median space" was introduced in [Bo1], and applies to various naturally occurring spaces, such as the mapping class groups and Teichmüller space. Quasiflats have proven to be a useful tool in understanding the large-scale geometry of such spaces.

The main result here (Theorem 1.1) tells us that a quasiflat in a coarse median space, Λ, lies a bounded distance from a bounded finite number of "coarse panels", that is, quasi-isometrically embedded rectilinear subsets of $\mathbb{R}^\nu$, which preserve the median up to bounded distance. Moreover, the panels fit together in a natural way to give a "panel complex" also quasi-isometrically embedded in Λ. The constants of the conclusion are all "uniform" in the sense that they depend only on those of the hypotheses. In addition, one can also say that the quasiflat is a finite Hausdorff distance from a bounded number of coarse orthants, though the distance bound may no longer be uniform.

Date: First draft: 10th June 2019. Revised: 15th December 2022.

If the coarse median space is a connected median metric space, one can strengthen the conclusion (Theorem 1.3). In particular, this applies to CAT(0) cube complexes (Corollary 1.4), and we recover the result of Huang [H] regarding quasiflats in such spaces.

All "hierarchically hyperbolic spaces" in the sense of [BeHS1] are coarse median. Therefore, from Theorem 1.2, we recover the main result of [BeHS3] regarding quasiflats in asymphoric hierarchically hyperbolic spaces. To make sense of this statement, one needs to assume in addition that the coarse median space comes equipped with a family of projection maps to hyperbolic spaces, satisfying certain conditions which we discuss in Sections 13 and 14.

Related results have been obtained earlier in other contexts. Indeed many of the ideas have their origins in [KlLe] and [KaKL]. We will also make use of some ideas from [BeHS3], notably in Section 10 here, though overall, our argument is rather different. Some analogous recent results about quasiflats in CAT(0) spaces can be found in [KlLa].

We give more precise statements of these results as follows. Implicit in the statements are various constants or "parameters". In general, the parameters of the conclusion depend on those of the hypotheses, in a manner we will specify.

To begin, we recall the notion of a coarse median space. Briefly, it is a geodesic metric space, (Λ, ρ), equipped with a ternary "median" operation, $\mu : \Lambda^3 \longrightarrow \Lambda$, satisfying certain conditions. The key point is that on finite subsets, it behaves up to bounded distance like the standard median operation on the vertex set of a finite CAT(0) cube complex (or equivalently a finite median algebra). It is said to have "(coarse median) rank" at most ν if we can always take the cube complex to have dimension at most ν. See Section 7 for a formal definition. It is shown in [Bo1] that if $\mathbb{R}^n$ quasi-isometrically embeds in Λ, then $n \leq \nu$.

One can equip $\mathbb{R}^n$ with a natural median operation, namely the direct product of the standard median in each of the $\mathbb{R}$-factors: the median in $\mathbb{R}$ is the usual betweenness operation. In this context, it is natural to equip $\mathbb{R}^n$ with the l^1 metric (so that it is a median metric space, as discussed in Section 2). By a "panel" we mean a subset of $\mathbb{R}^n$ which is a direct product of non-trivial closed connected subsets of $\mathbb{R}$ in each of the factors. (It is generally convenient to assume that no such factor is all of $\mathbb{R}$.) We equip it with the induced median operation and l^1 metric. An ***orthant*** is a panel isometric (and median isomorphic) to $[0, \infty)^n$. We speak of "n-panels" and "n-orthants" when we want to specify the dimension. We refer to the lengths of the interval factors as the ***side-lengths*** of the panels. Therefore, all side-lengths of an orthant are ∞.

Definition. A ***quasimorphism*** of a panel, P, is a map $\phi : P \longrightarrow \Lambda$, which respects medians up to bounded distance. It is a ***strong quasimorphism*** if it is also a quasi-isometric embedding. (Note that we do not assume such maps to be continuous.) A ***coarse panel*** in Λ is the image, $\phi(P)$, of a panel under a strong quasimorphism, ϕ.

One can define a ***face*** of a panel in the obvious way: choose a subset of the factors, and replace each such factor by a point in its boundary in $\mathbb{R}$.

Definition. A ***panel complex*** is a complete geodesic space which is a finite union of subsets, called ***cells***, where each cell is isometric to a panel in the induced path metric, and where any two cells intersect, if at all, in a common face.

One can think of this as a cell complex, so that distinct cells have disjoint interiors. After subdividing, there is no loss in supposing that the induced metric on each cell is already a path metric.

Definition. An ***orthant complex*** is a panel complex with exactly one 0-cell.

In other words, all the cells are orthants, and we can think of the complex as a cone with vertex at the 0-cell.

The main result can be stated as follows. (We will give a more detailed discussion in Section 11.)

Theorem 1.1. *Let Λ be a coarse median space of rank $\nu < \infty$, and let $f : \mathbb{R}^\nu \longrightarrow \Lambda$ be a quasi-isometric embedding. Then there is a panel complex, Ω, which is a finite union, $\Omega = \bigcup_{i=1}^p P_i$, of ν-panels P_i, and a quasi-isometric embedding $\phi : \Omega \longrightarrow \Lambda$ with $\phi(\Omega)$ a bounded Hausdorff distance from $f(\mathbb{R}^\nu)$. Also, for each i, $\phi|P_i$ is a strong quasimorphism (so that $\phi(P_i)$ is a coarse ν-panel). The parameters of the coarse panels (i.e. the quasimorphism and quasi-isometry constants), the number of panels, and the Hausdorff distance bound, each depend only on the parameters of Λ and the quasi-isometry constants of f.*

Moreover, if $L \geq 0$ then there is a subcollection, say $P_1, \ldots, P_q$, of the ν-cells, P_i, of Ω, all of whose side-lengths are at least L, such that $f(\mathbb{R}^\nu)$ is a bounded Hausdorff distance from $\bigcup_{i=1}^q \phi(P_i)$, where the distance bound now depends on L, as well as the other parameters.

If Λ is proper (that is, complete and locally compact) then we can arrange that the parameters of the panels (the quasi-isometry and quasimorphism constants of the maps $\phi|P_i$) as well as the quasi-isometry constant of ϕ in the first paragraph depend only on the parameters of Λ (i.e. independently of f).

In the second paragraph of the theorem, we could take L to be larger than any of the finite side-lengths in any of the panels, in which case, each P_i for $i \leq q$ is an orthant. As an immediate consequence (taking L to be big enough) we see that a quasiflat in Λ is a finite Hausdorff distance from the union of a bounded number of coarse ν-orthants in Λ. Of course, we lose control on the Hausdorff distance in this case: we can only say that it is finite.

In fact, an elaboration of the argument gives the following:

Theorem 1.2. *Let Λ be a coarse median space of rank $\nu < \infty$, and let $f : \mathbb{R}^\nu \longrightarrow \Lambda$ be quasi-isometric embedding. Then there is an orthant complex, Ω, bilipschitz equivalent to $\mathbb{R}^\nu$, and consisting of finite union $\Omega = \bigcup_{i=1}^q P_i$ of orthants,*

P_i, together with a quasi-isometic embedding, $\phi : \Omega \longrightarrow \Lambda$, with $\phi(\Omega)$ a finite Hausdorff distance from $f(\mathbb{R}^\nu)$. Also, for each i, $\phi|P_i$ is a strong quasimorphism (so that each $\phi(P_i)$ is a coarse ν-orthant), where the parameters of the strong quasimorphism depend only on those of Λ. The number, q, of orthants depends only on the parameters of Λ and the quasi-isometry constants of f. (There is no uniform control on the quasi-isometry parameters of ϕ, nor the Hausdorff distance.)

The additional information here is the bilipschtitz equivalence of Ω and $\mathbb{R}^\nu$. In particular, they are homeomorphic. (This follows from Theorem 1.1 using Lemma 6.4.)

In the above results, our arguments will give a computable bound on p and q and on the strong quasimorphism parameters of the panels. However they do not give a means of computing the Hausdorff distance bound, since this relies on a limiting argument.

Examples of coarse median spaces are connected median metric spaces. (See Section 2 for a definition.) In this case, we can strengthen the conclusions so that each panel is isometrically embedded into Λ by a median homomorphism. We can summarise this as follows:

Theorem 1.3. *Suppose that Λ is a connected median metric space of rank ν, and $f : \mathbb{R}^\nu \longrightarrow \Lambda$ is a quasi-isometric embedding. Then in the conclusions of Theorems 1.1 and 1.2, we could take each $\Phi_i = \phi(P_i)$ to be a closed convex subset intrinsically isometric to a ν-panel. (Though we cannot assume ϕ to be continuous on Ω.)*

The notion of "convexity" is defined in Section 2. Theorem 1.3 will be proven in Section 12.

We should note that the "rank" of a median metric space may be greater than its rank as a coarse median space, so the hypotheses of Theorem 1.3 are stronger in that regard.

Any CAT(0) cube complex has the structure of a median metric space, on replacing the usual l^2 metric on each cube with the l^1 metric (see Section 3). In the conclusion of Theorem 1.3 in this case, we can assume that each of the panels, Φ_i, is a subcomplex. (Just move them a bounded distance so that their corners are 0-cells of the complex.) As an immediate consequence, we deduce:

Corollary 1.4. *Let Λ be a CAT(0) cube complex of dimension ν, and let $f : \mathbb{R}^\nu \longrightarrow \Lambda$ be a quasi-isometric embedding. Then $f(\mathbb{R}^\nu)$ is a bounded Hausdorff distance from $\bigcup_{i=1}^p \Phi_i$, where each Φ_i is a convex subcomplex isometric to a direct product of ν connected subsets of $\mathbb{R}$. It is also a finite Hausdorff distance from $\bigcup_{i=1}^q \Phi_i$, for $q \leq p$, where Φ_i is an orthant (a direct product of rays) for each $i \leq q$. Here, p as well as the Hausdorff distance bound in the first statement are bounded in terms of ν and the parameters of f.*

In this context "convexity" of a subcomplex in the median sense is equivalent to the usual geometric notion of convexity in the CAT(0) sense. (Though for an arbitrary subset of Λ, CAT(0) convexity is weaker.)

Again, in Theorem 1.3 and Corollary 1.4, the orthants fit together to form an orthant complex, Ω, homeomorphic to $\mathbb{R}^\nu$. The inclusions of the Φ_i into Λ combine to give a quasi-isometry of Ω into Λ. (Though again, we cannot in general assume this to be continuous.)

The second statement of Corollary 1.4 recovers the result of Huang [H], though our arguments are quite different. (This result of [H] is used in the proof of the main result of [BeHS3], though we derive it here independently.)

Clearly an example of a quasiflat would be a strong quasimorphism of $\mathbb{R}^\nu$ into Λ. Indeed, in some cases, all quasiflats arise in this way up to quasi-isometric reparameterisation. (For example, in the case where quasiflats are "isolated", and indeed in the case of $\mathbb{R}^\nu$ itself.) However, this is certainly not true in general. A quasiflat may "fold up", but only in a way that is strongly controlled.

As a simple example, let $T = \gamma_1 \cup \gamma_2 \cup \gamma_3$ be the tripod consisting of three rays, $\gamma_i \cong [0, \infty)$, joined at the basepoint, 0. This has the structure of a CAT(0) cube complex (subdividing the rays in each factor into unit intervals). Let $F = \bigcup_{i \neq j} (\gamma_i \times \gamma_j) \subseteq T \times T$. Then F is a quasiflat, and is the union of the 6 quadrants $\gamma_i \times \gamma_j$ for $i \neq j$.

Coarse median spaces often come equipped with a family of "projection maps" to hyperbolic spaces. Under certain assumptions, discussed in Section 13, one can say more about the structure of coarse panels. In particular, this applies to hierarchically hyperbolic spaces. In Section 14, we will illustrate this with an account of what this means for the mapping class groups and Teichmüller space in either the Teichmüller or Weil-Petersson metric (see Theorems 14.1 and 14.2). At the end of Section 14 we briefly describe how this relates to the main result of [BeHS3].

We remark that the results of [H] and [BeHS3] as stated allow the constants of the conclusion to depend on the quasiflat and the ambient space. However, as noted in [H], a volume-growth argument gives a uniform and computable bound on the number of orthants in the context of cube complexes. It seems quite plausible that by an elaboration of the argument there, one could obtain a version of Corollary 1.4, with computable bounds on the number of panels as well as Hausdorff distance. However, it would require some work, and I have not checked this. The proofs in [BeHS3] rely on a number of limiting arguments, so this does not give effective control of the constants. One can show the number of orthants is bounded in terms of the quasi-isometry constants of the quasiflat, as they note in their Corollary 4.16. However as stated, this bound may also depend on the ambient space. It is unclear whether their arguments could lead to a uniform bound, in the sense that it should only depend on the parameters of the hypotheses.

In the course of the proof, we give some general results which may have some independent interest.

For example, it is well known that a (finite) CAT(0) complex admits a natural structure as a median algebra (see Section 3). We show (Proposition 3.6) that we can recover a canonical structure as a CAT(0) cube complex from its structure as a median algebra — its vertex set is a finite subalgebra uniquely minimal with this property.

Lemma 6.1 tells us that a quasi-isometric embedding of $\mathbb{R}^\nu$ into a finite panel complex of dimension ν is a bounded distance from a quasi-isometry to a subcomplex. This can retrospectively be viewed as a consequence of Corollary 1.4, though its proof is more explicit, and does not involve any limiting argument. In particular, all the constants arising are explicitly computable.

The "coarsification" procedures described in Sections 7 to 9 have some potential for wider application. The general principle behind this is that many results about median algebras have analogues for coarse median spaces, where statements are interpreted as holding up to bounded distance.

In outline, this paper is structured as follows. In Section 2, we describe various properties of median algebras and median metric spaces. Many of these are standard, but we include some new statements. In Section 3, we describe CAT(0) cube complexes from the point of view of median algebra, and introduce panel complexes. Section 4 discusses cubulated sets. The main result (Lemma 4.2) tells us that a uniformly cubulated set is a subcomplex of a panel complex. Section 5 gives some regularity statements about maximal dimensional copies of euclidean spaces in a median metric space. The main result of Section 6 is Lemma 6.1 mentioned above. We introduce coarse median spaces in Section 7, and in Sections 8 and 9, we study coarse intervals and coarse cubes in such spaces. In Section 10, we describe asymptotic cones. The main result is Lemma 10.5 which is a key ingredient of the proof of Theorem 1.1 in Section 11. Section 12 strengthens these results in the case of a median metric space, and we deduce Theorem 1.3 and Corollary 1.4. In Sections 13 and 14, we give some elaboration on Theorem 1.1 for particular classes of coarse median spaces, such as the mapping class groups and Teichmüller space, as well as hierarchically hyperbolic spaces. In particular, we show how these results imply the main result of [BeHS3]. Finally, in Section 15, we give a variant of the Borsuk-Ulam Theorem which is used in the proof of Lemma 6.1.

I thank Jingyin Huang for his interest and comments on an earlier draft of this paper.

2. Median algebras

We describe some general facts about median algebras. For more background, see [BaH, I, Ve, Ro, Bo3, Bo9] and references therein. We begin with some fairly standard definitions and constructions. The literature is somewhat dispersed, and the references we give here are not necessarily the original ones.

A ***median algebra*** is a set, M, equipped with a ternary operation, $\mu : M^3 \longrightarrow M$, symmetric in the arguments and such that $\mu(a,a,b) = a$ and $\mu(a,b,\mu(c,d,e)) = \mu(\mu(a,b,c),\mu(a,b,d),e)$ for all $a,b,c,d,e \in M$. (This can be more intuitively thought of in terms of cube complexes as we discuss in Section 3.) A ***subalgebra*** is a subset closed under μ. A ***(median) homomorphism*** between two median algebras is a map which respects the median operation.

Note that a two-point set, $\{0,1\}$, admits a unique median algebra structure. An n-***cube*** is a median algebra isomorphic to $\{0,1\}^n$ with the product structure. The ***rank*** of a median algebra, M, is the maximal $\nu \in \mathbb{N}$ such that M contains (a subalgebra isomorphic to) a ν-cube. (For all the median algebras we consider in this paper, this will be finite.)

Given $A \subseteq M$, write $\langle A \rangle$ for the median algebra generated by A: that is the smallest subalgebra containing A. Equivalently, it the set of elements which can be written as a finite median expression in elements of A. The following is well known (see Lemma 4.2 of [Bo1] for one account).

Lemma 2.1. *If $|A| \leq n < \infty$, then $|\langle A \rangle| \leq 2^{2^n}$.*

In particular, if A is finite, so is $\langle A \rangle$. Moreover, this places a bound on the complexity of an expression we need to represent an element of $\langle A \rangle$.

If $\phi : M \longrightarrow M'$ is a median homomorphism between median algebras and $A \subseteq M$, then $\phi(\langle A \rangle) = \langle \phi A \rangle$ (since images and preimages of subalgebras are subalgebras).

Given $a,b \in A$, let $[a,b] = [a,b]_M = \{\mu(a,b,x) \mid x \in M\} = \{x \in M \mid \mu(a,b,x) = x\}$, for the ***interval*** from a to b. We say that a subset, $C \subseteq M$, is ***convex*** if $[a,b] \subseteq C$ for all $a,b \in C$. Given $A \subseteq M$, write $\operatorname{hull}(A)$ for the ***convex hull*** of A: the smallest convex set containing A. Clearly $\langle A \rangle \subseteq \operatorname{hull}(A)$. If $a,b \in M$, then $\operatorname{hull}(\{a,b\}) = [a,b]$.

The following, often referred to as the "Helly property", is a key fact about convex sets (see for example Section 22 of [Ro]).

Lemma 2.2. *Let $C_1, \dots, C_n$ be a non-empty finite family of pairwise intersecting convex subsets of M. Then $\bigcap_{i=1}^n C_i \neq \varnothing$.*

If C is convex and $x \in M$, we say that a point $y \in C$ is a ***gate*** for x if $[x,y] \cap C = \{y\}$. This is equivalent to saying that $y \in [x,c]$ for all $c \in C$. If a gate exists, then it is unique. A map $\omega : M \longrightarrow C$ is a ***gate map*** if $\omega(x)$ is a gate for x for all $x \in M$. We say that C is ***gated*** if a gate map exists. We note the following (see for example, Section 5 of [BaH], or Section 6 of [Bo5]).

Lemma 2.3. *A gate map $\omega : M \longrightarrow C$ to a gated convex set, C, is median epimorphism.*

In particular, it sends cubes to cubes (possibly of lower dimension). It also sends intervals to intervals. (For suppose $x,y \in M$. Let $a = \omega x$, $b = \omega y$, and suppose that $c \in [a,b]$. Let $z = \mu(x,y,c) \in [x,y]$. We can check that $c = \mu(a,b,z)$. Therefore $c = \omega c = \mu(\omega a, \omega b, \omega z) = \mu(a,b,\omega z) = \omega z$ since $\omega z \in [\omega x, \omega y] = [a,b]$.) If $a,b \in M$, then we have a gate map, $\omega_{a,b} : M \longrightarrow [a,b]$ defined by $\omega_{a,b}(x) = \mu(a,b,x)$. In other words, intervals are always gated.

Given $x,y \in [a,b]$, let $x \wedge y = \mu(a,x,y)$ and $x \vee y = \mu(b,x,y)$. With this structure of ***meet*** and ***join*** operations, $[a,b]$ is a distributive lattice (see [BaH]). We write $x \leq y$ to mean that $x = x \wedge y$. This is equivalent to saying that $y = x \vee y$, or to saying that $[a,x] \subseteq [a,y]$, or that $[b,y] \subseteq [b,x]$. Note that $\leq$ is a partial order on $[a,b]$.

Definition. If $\leq$ is a total order, we will say that a,b (or $[a,b]$) is ***straight***.

Note that this is equivalent to saying that $[a,b]$ has rank 1 as an intrinsic median algebra. (This definition is not standard terminology.)

Given $a,b,a',b' \in M$, we say that the pairs a,b and a',b' are ***parallel*** if $[a,b'] = [a',b]$. (If a,b,a',b' are all distinct, this is equivalent to saying that a,b,b',a' is a 2-cube.) In this case, the map $\omega_{a',b'}|[a,b]$ is a median isomorphism from $[a,b]$ to $[a',b']$, with inverse $\omega_{a,b}|[a',b']$. One also checks that parallelism is transitive, and that the gate maps between parallel intervals commute with each other. In particular, any intrinsic property (such as straightness) is preserved by parallelism. (See for example Section 2.4 of [Bo4] for more discussion.)

A ***wall*** in M is an unordered partition, $M = A \sqcup B$, of M into two disjoint non-empty convex sets. Write $\mathcal{W}(M)$ for the set of walls. We say that two walls, $W = \{A,B\}$ and $W' = \{A',B'\}$ ***cross*** if each of the sets $A \cap A'$, $A \cap B'$, $B \cap A'$ and $B \cap B'$ is non-empty. We note (Proposition 6.2 of [Bo1]):

Lemma 2.4. rank(M) *is the maximal cardinality of any set of pairwise crossing walls of* M*.*

We say that $W = \{A,B\}$ ***separates*** a from b if $\{a,b\}$ meets both A and B. We write $\mathcal{W}(a,b) \subseteq \mathcal{W}(M)$ for the set of walls separating a from b. Note that if $a,b,c \in M$, then $\mathcal{W}(a,b) \subseteq \mathcal{W}(a,c) \cup \mathcal{W}(c,b)$. More generally, a wall $W = \{A,B\}$ ***separates*** two non-empty subsets, C,D, if, up to swapping A,B, we have $C \subseteq A$ and $D \subseteq B$. The following, sometimes known as the "Kakutani separation property", is due to Nieminen [Nie] (see also Section 2.8 of [Ro] for a proof).

Proposition 2.5. *Any two disjoint non-empty convex subsets of* M *can be separated by a wall.*

As a special case, it follows that if $\mathcal{W}(a,b) = \varnothing$, then $a = b$.

If $M \longrightarrow M'$ is an epimorphism, we have a natural injective map $\mathcal{W}(M') \longrightarrow \mathcal{W}(M)$ obtained by taking preimages. From the above, we see that rank$(M') \leq$ rank(M). Also, note that if $M'' \subseteq M$ is a subalgebra, then rank$(M'') \leq$ rank(M).

Let $C \subseteq M$ be a gated convex subset. Given a wall, $W = \{A,B\} \in \mathcal{W}(C)$, there is a unique wall $\hat{W} = \{\hat{A},\hat{B}\} \in \mathcal{W}(M)$ such that $A = \hat{A} \cap C$ and $B = \hat{B} \cap C$, namely set $\hat{A} = \omega^{-1}A$ and $\hat{B} = \omega^{-1}B$, where $\omega : M \longrightarrow C$ is the gate map. The map $[W \mapsto \hat{W}] : \mathcal{W}(C) \longrightarrow \mathcal{W}(M)$ is injective. In this way, we can identify $\mathcal{W}(C)$ with a subset of $\mathcal{W}(M)$.

By an ***oriented wall*** we mean an ordered pair, $W = (W^-, W^+)$, such that $\{W^-, W^+\}$ is an (unoriented) wall. Given $X, Y \in \mathcal{W}(M)$, we write $X \pitchfork Y$ to mean that the unoriented walls cross in M, i.e. $X^\epsilon \cap Y^{\epsilon'} \neq \varnothing$, for all choices $\epsilon, \epsilon' \in \{+, -\}$. If $X_1, \ldots, X_n$ are pairwise crossing walls, then $\bigcap_{i=1}^n X_i^{\epsilon_i} \neq \varnothing$ for all choices of $\epsilon_i \in \{+, -\}$. This follows by the Helly property (Lemma 2.2 above).

From this discussion one can deduce the following (see Proposition 5.2 of [Bo1]).

Lemma 2.6. *If $\omega : M \longrightarrow Q$ is an epimorphism to an n-cube, Q, then there is an n-cube, $Q' \subseteq M$, such that $\omega|Q'$ is an isomorphism to Q.*

A ***topological median algebra*** is a hausdorff topological space, M, equipped with a median algebra structure such that the map $\mu : M^3 \longrightarrow M$ is continuous. We say that M is ***interval-compact*** if $[a, b]$ is compact for all $a, b \in M$. If M interval-compact and $C \subseteq M$ is closed and convex, then C is gated.

Topological median algebras often arise as median metric spaces. We recall the basic definitions.

Let (M, ρ) be a metric space. (We will generally use $\rho = \rho_M$ to denote a metric on a set, M.) Given $a, b \in M$, let $[a, b]_\rho = \{x \in M \mid \rho(a, x) + \rho(x, b) = \rho(a, b)\}$.

Definition. We say that M is a ***median metric space*** if for all $a, b, c \in M$, there is some $d \in M$ such that $[a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho = \{d\}$.

We write $\mu(a, b, c) = d$. The following follows from work of Sholander [S] (see [Bo3] for some elaboration):

Lemma 2.7. *With this ternary operation, (M, μ) is a topological median algebra. Moreover, and that $[a, b]_\rho = [a, b]$ is equal to the median interval between a and b.*

Also, the map $\omega_{a,b} : M \longrightarrow [a, b]$ is 1-lipschitz. Note that any subalgebra of M is a median metric space in the induced metric.

The completion of any median metric space is a median metric space of the same rank (see Corollary 2.16 of [ChaDH], or Lemma 4.3 of [Bo3]). A complete median metric space of finite rank is interval-compact (see for example, Corollary 5.2 of [Bo6]). As a simple consequence one can deduce:

Lemma 2.8. *A closed convex subset of a finite-rank complete median metric space is gated.*

In fact, if $C \subseteq M$ is non-empty closed and convex, then the gate, $\omega(x)$ of x in C is the unique nearest point to x. (Here it is easy to see that such exists: choose any $c \in C$ and let $\omega(x)$ be a nearest point to x in the compact set $C \cap [x, c]$. One checks that this is the unique nearest point in C.) The following is also easily verified (see for example, Section 2.4 of [Bo4]):

Lemma 2.9. *Let C be gated with gate map $\omega : M \longrightarrow C$. Suppose $D \subseteq M$ is convex and $D \cap C \neq \varnothing$, then $\omega(D) = C \cap D$.*

Also if $C' \subseteq C$ is closed and convex, then the gate map to C' factors through that to C. We also note that a connected complete median metric space is a geodesic space (see for example Lemma 4.6 of [Bo3]).

Lemma 2.10. *Let M be a median metric space of rank $\nu < \infty$. If $A \subseteq M$, then* $\operatorname{diam}(\operatorname{hull}(A)) \leq 3^\nu \operatorname{diam}(A)$.

Proof. If $a, b \in M$, the $\operatorname{diam}([a,b]) = \rho(a,b)$. It follows that $\operatorname{diam}(L(A)) \leq 3 \operatorname{diam}(A)$, where $L(A) = \bigcup_{a,b \in A} [a,b]$. Now $\operatorname{hull}(A)$ is the result of iterating L at most ν times (see Lemma 5.5 of [Bo1]). □

Of course, one could improve on this bound, though it is not clear what the optimal statement would be. We will only need here the fact that the convex hull of a bounded set is bounded. The closure of a convex set is convex, and so we see:

Lemma 2.11. *Any median metric space of finite rank admits an exhaustion by closed bounded convex subsets.*

The following discussion will be used in Section 7.

Let $\phi : M \longrightarrow M'$ be an epimorphism between two median algebras. It is easily checked that the image or preimage of any convex set is convex.

If $A, B \subseteq M$ are convex, and $A \cap B \neq \varnothing$, then $\phi(A \cap B) = \phi A \cap \phi B$. (For if $x \in \phi A \cap \phi B$, then the sets $A, B, \phi^{-1}(x)$ are all convex and pairwise intersect, and therefore have non-empty intersection, by Lemma 2.2.)

If $\operatorname{rank}(M') = \operatorname{rank}(M) = \nu < \infty$, and $A, B \subseteq M$ are convex with $\phi A = \phi B = M'$, then $A \cap B \neq \varnothing$. (For if not, by Proposition 2.5, there is a wall, W_0, of M separating A and B. Let $W'_1, \ldots, W'_\nu$ be pairwise crossing walls in M' and let $W_1, \ldots, W_\nu$ be the corresponding walls in M. Now the walls $W_0, W_1, \ldots, W_\nu$, pairwise cross in M. By Lemma 2.4, this gives contradiction that $\operatorname{rank}(M) \geq \nu + 1$.) In particular, from the previous paragraph, it follows that $\phi(A \cap B) = M'$.

Now suppose that M and M' are both topological median algebras, and that ϕ is continuous.

Lemma 2.12. *If ϕ is proper, there is a minimal closed convex subset $C \subseteq M$ such that $\phi C = M'$.*

Proof. This follows by Zorn's Lemma as follows. Let $\mathcal{C}$ be a chain of closed convex subsets C satisfying $\phi C = M'$, ordered by inclusion. Let $C_0 = \bigcap \mathcal{C}$. This is also closed and convex, and $\phi(C_0) = M'$. (For if $x \in M$, the family $\{C \cap \phi^{-1}(x) \mid C \in \mathcal{C}\}$ is a chain of closed non-empty subsets of the compact set $\phi^{-1}(x)$, and so has non-empty intersection.) □

Note that if $\operatorname{rank}(M) = \operatorname{rank}(M') = \nu < \infty$, then such a minimal C is unique. (For if $D \subseteq M$ were another such, then by the earlier discussion, $\phi(C \cap D) = M'$.)

Definition. A topological median algebra, M, of rank ν is ***thick*** if whenever we have $M = A \cup B$ with $A, B \subseteq M$ closed and convex and $\operatorname{rank}(A \cap B) \leq \nu - 2$, then either $A = M$ or $B = M$.

Lemma 2.13. *Let M, M' be topological median algebras, and let $\phi : M \longrightarrow M'$ be a continuous proper median epimorphism. Suppose that $\operatorname{rank}(M) = \operatorname{rank}(M') = \nu < \infty$. Suppose that M is connected and that M' is thick. Then there is a closed convex subset, $C \subseteq M$, which is intrinsically thick, and such that $\phi C = M'$.*

Proof. Let $C \subseteq M$ be the minimal closed convex subset with $\phi C = M'$, as given by Lemma 2.12. Note that $\operatorname{rank}(C) = \nu$. We claim that C is thick. For suppose that $C = A \cup B$, with $A, B \subseteq C$ non-empty closed and convex and with $\operatorname{rank}(A \cap B) \leq \nu - 2$. Since M is connected, so is C, and so $A \cap B \neq \varnothing$. Therefore $\phi(A \cap B) = \phi A \cap \phi B$. But $\operatorname{rank}(\phi(A \cap B)) \leq \operatorname{rank}(A \cap B) \leq \nu - 2$. Since M' is thick, we can assume that $\phi A = M'$. By minimality of C, it follows that $C = A$. This shows that C is thick. □

In fact, the proof gives a slightly stronger statement:

Lemma 2.14. *In the conclusion of Lemma 2.13, we can suppose that whenever $C = A \cup B$, with A, B closed convex subsets, then* $\operatorname{rank}(\phi(A \cap B)) \geq \nu - 1$.

Finally, we note the following, proven in [Bo3].

Theorem 2.15. *Let (M, ρ) be a complete connected median metric space of rank ν. Then M admits a (canonical) CAT(0) metric σ with $\sigma \leq \rho \leq \sigma\sqrt{\nu}$.*

As an example, if $M = \mathbb{R}^\nu$ and ρ is the l^1 metric, then σ is the l^2 metric.

3. Cube complexes

In this section, we describe the notion of a "CCAT(0) panel complex". Any subcomplex of such will be a "panel complex" as defined in Section 2. (Not every panel complex arises in this way, though all those that we consider in this paper do so.)

We begin with some more standard notions. For simplicity, we will deal here mostly with median metric spaces, though it can be seen that many of the ideas apply to more general median algebras.

A ***cube complex*** is a CW complex built out of real (euclidean) cubes attached along their faces. The cube complexes we consider in this section will be finite (i.e. have finitely many cells). We view a cube complex primarily as a combinatorial object. Note that the link of any cell is naturally a (finite) simplicial complex.

Definition. A ***CCAT(0) cube complex*** is a simply connected cube complex, Δ, such that the link of every cell is a flag simplicial complex.

(Here "flag" means that every complete graph in the 1-skeleton lies in a single cell.)

We can equip Δ with the "standard" euclidean (or l^2) metric, where each cell is given the structure of a unit euclidean cube, and we take the induced path metric. With this metric, Δ will be CAT(0) in the usual geometric sense. The term "CCAT(0)" means "combinatorially CAT(0)". We adopt this terminology since we will mostly be dealing with other metrics which are not CAT(0). In particular, we could instead take the "standard" l^1 metric, ρ. Here each cell is isometric to the l^1 direct product of unit real intervals. In this case, (Δ, ρ) is a median metric space, and in particular, a topological median algebra. Moreover, the vertex set, Π will be a subalgebra. In fact, every finite median algebra arises canonically in this way [Che].

One way of describing this is as follows. Let Π be a finite median algebra, and let $\mathcal{W} = \mathcal{W}(\Pi)$ be the set of walls of Π. To each wall, $W \in \mathcal{W}(\Pi)$, we associate a compact real interval, $I_W \subseteq \mathbb{R}$, with boundary $\partial I_W = \{\partial^- I_W, \partial^+ I_W\}$. We equip I_W with its standard median structure, so that ∂I_W is a two-point median algebra, and $Q(\Pi) := \prod_{W \in \mathcal{W}} \partial I_W$ is a $|\mathcal{W}|$-cube. We write $\Delta(Q(\Pi)) := \prod_{W \in \mathcal{W}} I_W$ for the real cube, with "corners" $Q(\Pi) \subseteq \Delta(Q(\Pi))$. Now Π embeds naturally into $Q(\Pi)$, and hence into $\Delta(Q(\Pi))$, by taking the I_W-coordinate of x to be $\partial^\pm I_W$ if $x \in W^\pm$. We identify Π with its image under this embedding, and we let $\Delta(\Pi)$ be the full subcomplex with vertex set Π. We get:

Lemma 3.1. *$\Delta(\Pi)$ is a CCAT(0) cube complex, and a subalgebra of $Q(\Pi)$.*

In particular, $\Delta(\Pi)$ has a natural structure as a median algebra. One can also check that $\mathrm{rank}(\Pi) = \dim(\Delta(\Pi))$ (see Proposition 5.3 of [Bo1]). We refer to $\Delta(\Pi)$ as the ***realisation*** of Π.

A ***cell*** of Π is a convex subset isomorphic to a cube. Given a cell $Q \subseteq \Pi$, write $\Delta(Q)$ for its convex hull in $\Delta(\Pi)$. These are precisely the cells of $\Delta(\Pi)$ in its structure as a CW complex.

The following is shown in [Bo9]:

Lemma 3.2. *Let $C \subseteq \Pi$ be a closed convex subset. Then C is gated, and median isomorphic to $\Delta(\omega C)$, where $\omega : \Delta(\Pi) \longrightarrow C$ is the gate map.*

We only really require this result for finite CAT(0) complexes. It can seen by embedding $\Delta(\Pi)$ in $\Delta(Q(\Pi))$ and taking the convex hull, H, of C in $Q(\Pi)$. Then $C = H \cap \Delta(\Pi)$ (see for example, Lemma 5.6 of [Bo1]). Any convex subset of $\Delta(Q(\Pi))$ is a direct product of intervals. Combining the gate map on each factor, we obtain the gate map $\omega : \Delta(Q(\Pi)) \longrightarrow H$, with $H \cong \Delta(\omega(Q))$. Intersecting with $\Delta(\Pi)$ this restricts to the gate map to C.

In fact, each cell of C is the intersection of C with a cell of $\Delta(\Pi)$.

Lemma 3.3. *Any epimorphism $\Pi \longrightarrow \Pi'$ between finite median algebras extends to an epimorphism of $\Delta(\Pi)$ to $\Delta(\Pi')$.*

Proof. Recall that there is a natural injective map, $\mathcal{W}(\Pi') \longrightarrow \mathcal{W}(\Pi)$ as described in Section 2. This gives us an epimorphism $\Delta(Q(\Pi)) \longrightarrow \Delta(Q(\Pi'))$, which collapses those factors which correspond to walls not in the image of $\mathcal{W}(\Pi')$. This restricts to an epimorphism $\Delta(\Pi') \longrightarrow \Delta(\Pi)$. □

The ***standard*** euclidean metric on $\Delta(\Pi)$ is the path-metric induced from the metric on $\Delta(Q(\Pi))$ as a unit real cube. If instead, we take the l^1 metric, this is already a path-metric, and we get the standard l^1 metric on $\Delta(\Pi)$. In this latter structure, $\Delta(\Pi)$ is a median metric space, inducing the median algebra structure. Note that if $a, b \in \Pi$, then $\rho(a,b) = |\mathcal{W}(a,b)|$, that is the number of walls of Π separating a from b.

We also note that associated to each wall, $W \in \mathcal{W}$, we have a ***hyperplane***, $\Delta_W \subseteq \Delta(\Pi)$. This can be defined by taking the set of points of $\Delta(Q(\Pi))$ whose coordinate in I_W is the midpoint of that interval, and then intersecting this set with $\Delta(\Pi)$. This hyperplane cuts $\Delta(\Pi)$ into two pieces, $\Delta^\pm_W$, with $\Delta_W = \Delta^-_W \cap \Delta^+_W$ and $\Delta(\Pi) = \Delta^-_W \cup \Delta^+_W$, and with $\Pi \cap \Delta^\pm_W = W^\pm$. All these sets are convex, both in the median structure, and in the CAT(0) metric structure.

More generally, suppose that (Π, ρ) is a finite median metric space. Given $W \in \mathcal{W}$, choose some 1-cell, $\{c, d\} \in E(W)$, and set $w(W) = \rho(c, d)$. Since all 1-cells of $E(W)$ are parallel, $w(W)$ does not depend on our choice. We refer to $w(W)$ as the ***width*** of W. (In the above combinatorial setting, the widths were all deemed to be 1.) In fact, of $a, b \in \Pi$, then $\rho(a,b) = \sum_{W \in \mathcal{W}(a,b)} w(W)$.

We can similarly give $\Delta(\Pi)$ the l^1 metric so that $\Pi \subseteq \Delta(\Pi)$ is an isometric embedding. One way to describe this is to put the l^1 product metric on the real cube $\Delta(Q(\Pi))$ so that the side corresponding to I_W has length $w(W)$. We then take the induced metric on $\Delta(\Pi)$. This is automatically a geodesic metric, and $\Delta(\Pi)$ is again a median metric space.

In summary, we have:

Lemma 3.4. *Let Π be a finite median algebra. Given any function $w : \mathcal{W}(\Pi) \longrightarrow (0, \infty)$, there is a natural median metric on the realisation, $\Delta(\Pi)$, inducing the median structure, and such that the width of any wall W of Π is precisely $w(W)$.*

In particular, if we start with a finite median metric space Π, then we can canonically extend this to a median metric on $\Delta(\Pi)$.

(We could alternatively take the corresponding euclidean metric, in which case we get a CAT(0) metric on $\Delta(\Pi)$ after taking the induced path-metric. The passage from the l^1 metric to the l^2 metric on $\Delta(\Pi)$ is another illustration of Theorem 2.15.)

The following will be used in the proof of Lemma 4.2.

Lemma 3.5. *Let Π be a finite median algebra, and let $\Delta(\Pi)$ be its realisation as a cube complex, with the induced median algebra structure. If $\Psi \subseteq \Delta(\Pi)$ is a connected subset such that $\Pi = \langle \Psi \cap \Pi \rangle$ in Π, then $\Delta(\Pi) = \langle \Psi \rangle$.*

Proof. Suppose $\{a, b\}$ is a 1-cell of Π. The gate map $[x \mapsto \mu(a, b, x)]$ is continuous. The image of Ψ under this map contains a, b, hence all of $[a, b]$. It follows that $[a, b] \subseteq \langle \Psi \rangle$. This shows that $\langle \Psi \rangle$ contains the 1-skeleton of $\Delta(\Pi)$. Now any real cube is generated as a median algebra by its 1-skeleton, and it follows easily that $\langle \Psi \rangle = \Delta(\Pi)$ as claimed. □

The fact that, as a median algebra, Δ is median isomorphic to $\Delta(\Pi)$ does not in itself determine the subalgebra Π uniquely. In view of this, we make the following definition:

Definition. A subalgebra, Π, of Δ, is a ***0-skeleton*** if there is a median isomorphism of $\Delta(\Pi)$ to Δ, and which is the identity on Π.

We first note that if $\Pi' \leq \Delta(\Pi)$ is any subalgebra with $\Pi \subseteq \Pi'$ then Π' is also a 0-skeleton. We can think of $\Delta(\Pi')$ as a subdivision of $\Delta(\Pi)$ as cube complex. One way to construct the isomorphism more formally is as follows.

Recall that there is a natural surjection $f : \mathcal{W}(\Pi') \longrightarrow \mathcal{W}(\Pi)$. If $W \in \mathcal{W}$, we can write $f^{-1}(W) = \{W_1, W_2, \ldots, W_n\}$, with $W_1^- \subseteq W_2^- \subseteq \cdots \subseteq W_n^-$. We subdivide the real interval I_W into subintervals, $I_W = I_1 \cup \cdots \cup I_n$, and map I_W into $D(W) := \prod_{i=1}^n I_{W_i}$, by a median monomorphism sending I_i to a 1-cell, J_i, of $D(W)$. More specifically, J_i is the 1-cell whose I_{W_j} coordinate is equal to $prt^+ I_{W_j}$ for $j < i$ and equal to $prt^- I_{W_j}$ for $j > i$. This gives us a path, J_W, in the 1-skeleton of $D(W)$ connecting a pair of antipodal points. Now $\Delta(Q(\Pi')) \equiv \prod_{W \in \mathcal{W}(\Pi)} D(W)$. By taking the direct product of the above paths, we get a monomorphism of $\Delta(Q(\Pi))$ into $\Delta(Q(\Pi'))$ which we can take to be the identity on Π'. The image of $\Delta(\Pi)$ is precisely $\Delta(\Pi')$. The inverse map gives the required isomorphism of $\Delta(\Pi')$ to $\Delta(\Pi)$.

Note that in the above construction, if $i < j$, then $\Pi \cap W_i^+ \cap W_j^- = \varnothing$, and that every wall of Π' that crosses W_i also crosses W_j.

There are plenty of such 0-skeleta. (Take $\langle A \cup \Pi \rangle$ for any finite subset, $A \subseteq \Delta(\Pi)$.) However, we will show:

Proposition 3.6. *There is a unique minimal 0-skeleton,* Π_Δ, *of* Δ.

That is to say, a subalgebra, Π', of Δ is a 0-skeleton if and only if $\Pi_\Delta \subseteq \Pi'$.

In fact, one can describe Π_Δ as the subalgebra generated by the set of "extreme" points of Δ. This construction will be used in the proof of Lemma 4.2. We now set about the proof.

Let Π be a finite median algebra, and let $\mathcal{W} = \mathcal{W}(\Pi)$. Given $X, Y \in \mathcal{W}$, we write $X \mid Y$ to mean that $X \neq Y$ and $X \not\pitchfork Y$. In this case, we can orient X, Y so that $X^+ \cap Y^+ = \varnothing$. We write $\Pi_{XY} = X^- \cap Y^-$ for the convex subset between X and Y. We can think of X and Y as being oriented away from Π_{XY}. Note that $\Pi_{XY}^C = \Pi \setminus \Pi_{XY} = X^+ \cup Y^+$ is a subalgebra, being the union of two convex subsets.

If $W \in \mathcal{W}$, let $\mathcal{W}_W^0 = \{X \in \mathcal{W} \mid X \pitchfork W\}$. Note that $E(W)$ is a finite median algebra with $\mathcal{W}(E(W)) \equiv \mathcal{W}_W^0$. Also, $X, Y \in \mathcal{W}_W^0$ cross in Π then they also cross in $E(W)$.

Given $X, Y \in \mathcal{W}$, write $X \parallel Y$ to mean that ($X = Y$ or $X \mid Y$) and $(\forall Z \in \mathcal{W})(Z \pitchfork X \Leftrightarrow Z \pitchfork Y)$. We say that X, Y are ***parallel***. (Intuitively, this means that the hyperplanes corresponding to X and Y bound a product region in Δ.) One checks easily that $\parallel$ is an equivalence relation on $\mathcal{W}$. Let Π_E be the intersection of all Π_{XY}^C for $X, Y \in \mathcal{W}$ with $X \parallel Y$ (with the convention that $\Pi_{XX}^C = \Pi$). Note that Π_E is a subalgebra of Π. One can naturally identify $\mathcal{W}(\Pi_E)$ with $\mathcal{W}/\parallel$ (via the natural projection map $\mathcal{W} \longrightarrow \mathcal{W}(\Pi_E)$: two walls are parallel in Π if and only if they are identified in Π_E). Note that $(\Pi_E)_E = \Pi_E$.

A wall $W \in \mathcal{W}(\Pi)$ is a ***neighbouring*** wall of an element $x \in \Pi$, if there is some $y \in \Pi$ with $\mathcal{W}(x, y) = \{W\}$ (so that $\{x, y\}$ is a 1-cell of Π). We say that $x \in \Pi$ is ***extreme*** if all neighbouring walls cross. This is equivalent to saying that there do not exist $X, Y \in \mathcal{W}$ with $X \mid Y$ and $x \in \Pi_{XY}$. (This is also equivalent to saying that x does not lie in the interior of any geodesic segment in the CAT(0) metric on $\Delta(\Pi)$.) We write $\mathrm{ext}(\Pi) \subseteq \Pi$ for the set of extreme points. Note that $\mathrm{ext}(\Pi) \subseteq \Pi_E$. In fact we claim:

Proposition 3.7. $\Pi_E = \langle \mathrm{ext}(\Pi) \rangle$.

For the proof, we will make use of the following general discussion. Let $W \in \mathcal{W}(\Pi)$. We fix an orientation on W. Recall that $\mathcal{W}_W^0 = \{X \in \mathcal{W}(\Pi) \mid X \pitchfork W\}$. Let $\mathcal{W}_W^-$ be the set of $X \in \mathcal{W} \setminus \{W\}$ for which we can write $X = \{X^-, X^+\}$ such that $X^+ \cap W^+ = \varnothing$. (This is equivalent to $X^+ \subseteq W^-$, or to $W^+ \subseteq X^-$.) We similarly define $\mathcal{W}_W^+$. We have $\mathcal{W}_W^- \cap \mathcal{W}_W^+ = \varnothing$. Moreover $X|W$ is equivalent to $X \in \mathcal{W}_W^- \cup \mathcal{W}_W^+$. We therefore have a partition of $\mathcal{W}(\Pi)$ as $\mathcal{W}(\Pi) = \{W\} \sqcup \mathcal{W}_W^0 \sqcup \mathcal{W}_W^- \sqcup \mathcal{W}_W^+$.

We write $\mathcal{W}_W^{0\pm} = \mathcal{W}_W^0 \cup \mathcal{W}_W^\pm$. Note that (using the gate map to $W^\pm$), we can identify $\mathcal{W}_W^{0\pm}$ with $\mathcal{W}(W^\pm)$. We also note that if $X, Y \in \mathcal{W}_W^{0\pm}$, then X, Y cross as walls in Π if and only if they cross as walls in $W^\pm$.

Lemma 3.8. *Suppose that* $|\Pi| \geq 2$ *and that* $\Pi' \subseteq \Pi$ *is a subalgebra of* Π *satisfying:*
(W1): if $X_1, \ldots, X_n \in \mathcal{W}(\Pi)$ *pairwise cross then* $\Pi' \cap \bigcap_{i=1}^n X_i^{\epsilon_i} \neq \varnothing$ *for all choices of* $\epsilon_i \in \{+, -\}$*, and*
(W2): if $X \mid Y$*, then* $\Pi' \cap \Pi_{XY} \neq \varnothing$*.*
Then $\Pi' = \Pi$*.*

Proof. The proof is by induction on $|\Pi|$. Since $|\Pi| \geq 2$, $\mathcal{W}(\Pi) \neq \varnothing$, and we choose any wall $W \in \mathcal{W}(\Pi)$. Now $\mathcal{W}_W^{0\pm} \equiv \mathcal{W}(W^\pm)$. We will verify that the subalgebra $\Pi' \cap W^\pm$ of $W^\pm$ intrinsically satisfies (W1) and (W2). Since $|W^\pm| < |\Pi|$, it then follows that $\Pi' \cap W^\pm = W^\pm$. In other words, $W^\pm \subseteq \Pi'$. We then get $\Pi = W^- \cup W^+ = \Pi'$ as required.

Allowing ourselves to swap $+$ and $-$, it therefore suffices to prove that W^- satisfies (W1) and (W2). We first note the following.

Claim: Let $X \in \mathcal{W}_W^{0-}$, and $\epsilon \in \{-,+\}$. Then $\Pi' \cap W^- \cap X^\epsilon \neq \varnothing$.

To see this, note first that if $X \pitchfork W$ then the statement follows immediately from (W1) (for Π). So we suppose $X|W$. In other words $X^+ \cap W^+ = \varnothing$. There are two cases. If $\epsilon = -$, then $\Pi' \cap W^- \cap X^- = \Pi' \cap \Pi_{WX} \neq \varnothing$ by (W2). If $\epsilon = +$, then $\Pi' \cap X^+ \neq \varnothing$ by (W1). But $X^+ \subseteq W^-$, and so $\Pi' \cap W^- \cap X^+ \neq \varnothing$. This proves the Claim.

We can now verify (W1) for $\Pi' \cap W^-$ in W^-. Let $X_1, \ldots, X_n \in \mathcal{W}_W^{0-}$ with $X_i \pitchfork X_j$ for all i, j, and let $\epsilon_i \in \{-,+\}$. We want to show that $\Pi' \cap W^- \cap \bigcap_{i=1}^n X_i^{\epsilon_i} \neq \varnothing$.

To see this, we proceed by induction on n. Note that the case $n = 1$ is the above Claim, so we suppose that $n \geq 2$. Now $\Pi' \cap W^- \cap \bigcap_{i=1}^{n-1} X_i^{\epsilon_i} \neq \varnothing$ (by the inductive hypothesis), $\Pi' \cap W^- \cap X_n^{\epsilon_n} \neq \varnothing$ (by the Claim) and $\Pi' \cap \bigcap_{i=1}^n X_i^{\epsilon_i} \neq \varnothing$ (by (W1)). Therefore (taking the median of three points respectively in each of the above sets) we see that $\Pi' \cap W^- \cap \bigcap_{i=1}^n X_i^{\epsilon_i} \neq \varnothing$. Property (W1) for W^- now follows by induction.

We can also deduce (W2) for $\Pi' \cap W^-$ in W^- as follows.

Let $X, Y \in \mathcal{W}_W^{0-}$ with $X|Y$. (Recall that this statement can be equivalently interpreted in Π or W^-.) In this paragraph, we orient X, Y so that $X^+ \cap Y^+ = \varnothing$, and $\Pi_{XY} = X^- \cap Y^-$. Now $\Pi' \cap X^- \cap Y^- \neq \varnothing$ (by (W2) in Π) and $\Pi' \cap W^- \cap X^- \neq \varnothing$ and $\Pi' \cap W^- \cap Y^- \neq \varnothing$ (both by the Claim). It follows by Lemma 2.2 that $(\Pi' \cap W^-) \cap \Pi_{XY} = \Pi' \cap W^- \cap X^- \cap Y^- \neq \varnothing$. This proves (W2) for W^-.

The lemma now follows by induction on $|\Pi|$ as described in the first paragraph. $\square$

Corollary 3.9. *Suppose that* $\Pi' \subseteq \Pi$ *is a subalgebra satisfying:*
(B1): $\mathrm{ext}(\Pi) \subseteq \Pi'$, *and*
(B2): if $X, Y \in \mathcal{W}(\Pi')$, *with* $X \neq Y$ *and* $X \parallel Y$, *then* $\Pi' \cap \Pi_{XY} \neq \varnothing$.
Then $\Pi' = \Pi$.

Proof. The statement is trivial for $|\Pi| \leq 1$, so we assume that $|\Pi| \geq 2$. It is easily seen that $\mathrm{ext}(\Pi) \cap \bigcap_{i=1}^n X_i^{\epsilon_i} \neq \varnothing$ for any set of pairwise crossing walls, $X_1, \ldots, X_n \in \mathcal{W}(\Pi)$, and all choices of $\epsilon_i \in \{+,-\}$. (Note that the walls X_i all cross some n-cell $Q \subseteq \Pi$, so that $Q \cap \bigcap_{i=1}^n X_i^{\epsilon_i}$ consists of a single point x. Now choose a point of $\bigcap_{i=1}^n X_i^{\epsilon_i}$ a maximal distance from x in the l^1 metric on Π.) Thus, Π' satisfies (W1) of Lemma 3.8.

For (W2), suppose that $X, Y \in \mathcal{W}(\Pi)$, with $X \mid Y$. If $X \parallel Y$, then $\Pi' \cap \Pi_{XY} \neq \varnothing$ by (B2). If not, then, up to swapping X, Y, there is some $Z \in \mathcal{W}(\Pi)$ with $Z \mid X$ and $Z \pitchfork Y$. Now $Z^+ \subseteq X^-$, so $Z^+ \cap Y^- \subseteq X^- \cap Y^- = \Pi_{XY}$. By the first observation, $\mathrm{ext}(\Pi) \cap Z^+ \cap Y^- \neq \varnothing$, so $\mathrm{ext}(\Pi) \cap \Pi_{XY} \neq \varnothing$. By (B1) $\Pi' \cap \Pi_{XY} \supseteq \mathrm{ext}(\Pi) \cap \Pi_{XY} \neq \varnothing$. Thus Π' also satisfies (W1) of Lemma 3.8, and so $\Pi' = \Pi$. $\square$

Proof of Proposition 3.7. Let $\Pi' = \langle \mathrm{ext}(\Pi) \rangle \subseteq \Pi_E \subseteq \Pi$. Note that $\mathrm{ext}(\Pi_E) = \mathrm{ext}(\Pi)$, so (B1) is satisfied for Π' in Π_E. Also, since $(\Pi_E)_E = \Pi_E$, no two walls of Π_E are parallel, and so (B2) is vacuously satisfied for Π' in Π_E. Therefore, by Corollary 3.9 we have $\Pi' = \Pi_E$ as required. $\square$

Given any median algebra, M, we can define an ***extreme point*** to be a point $x \in M$ such that whenever $[a, b]$ is a straight interval in M with $x \in [a, b]$ then $x \in \{a, b\}$. We write $\mathrm{ext}(M)$ for the set of extreme points. It is easily checked that this agrees with the previous definition when M is finite.

Proof of Proposition 3.6. Let $\Delta = \Delta(\Pi)$ for Π finite. Let $\mathrm{ext}(\Delta)$ be the set of extreme points of Δ, and let $\Pi_\Delta = \langle \mathrm{ext}(\Delta) \rangle$. We claim that Π_Δ is the minimal 0-skeleton.

It is easily seen that $\mathrm{ext}(\Delta) \subseteq \Pi$, and that $\mathrm{ext}(\Delta) = \mathrm{ext}(\Pi)$. Therefore $\Pi_\Delta \subseteq \Pi$, and by Proposition 3.7 we have $\Pi_\Delta = \Pi_E$. Now if we have two parallel walls, $X \parallel Y$, of Π, we can remove the elements of Π between them (namely Π_{XY}) without changing the isomorphism type of $\Delta(\Pi)$. Doing this a finite number of times (or simultaneously) we end up with $\Delta(\Pi_E)$. This shows that $\Pi_\Delta = \Pi_E$ is a 0-skeleton of Δ.

In fact, it is the minimal such, since the initial choice, Π, of 0-skeleton was arbritrary. $\square$

Next, recall the definition of a "thick" topological median algebra given in Section 2.

Let $\Delta = \Delta(\Pi)$ be the realisation of some finite median algebra, Π, of rank $\nu < \infty$.

Lemma 3.10. *Δ is thick if and only if whenever we have $\Delta = C \cup D$ with $C, D \subseteq \Delta$ convex subcomplexes and $\mathrm{rank}(C \cap D) \leq \nu - 2$, then either $C = \Delta$ or $D = \Delta$.*

In other words, we can weaken the original formulation of thickness to insist that C, D are both subcomplexes.

Proof. We only need to check the "if" direction. Let $A, B \subseteq \Delta$ be as in the original definition of thickness. Write $\Pi_A = A \cap \Pi$, $\Pi_B = B \cap \Pi$, $\Delta_A = \Delta(\Pi_A)$ and $\Delta_B = \Delta(\Pi_B)$. Thus, Δ_A, Δ_B are both subcomplexes of Δ, containing A, B respectively. There are two cases.

Suppose $\Pi \cap A \cap B = \varnothing$. Then $W := \{\Pi_A, \Pi_B\}$ is a wall of Π_B. Let F be the closure of $\Delta \setminus (\Delta_A \cup \Delta_B)$. Then F is a subcomplex of Δ, isomorphic to $\Delta(E(W)) \times [0, 1]$, and so $\mathrm{rank}(F \cap \Delta_A) = \mathrm{rank}(F \cap \Delta_B) = \mathrm{rank}(\Delta(E(W)))$. Now $A \cap B \subseteq F$ separates Δ_A and Δ_B, so the projection of $A \cap B$ to $F \cap \Delta_A$ is surjective. It follows that $\mathrm{rank}(E(W)) \leq \mathrm{rank}(A \cap B) \leq \nu - 2$. Now $\Delta = \Delta_A \cup \Delta_B \cup F$. Applying the hypothesis on Δ to $\Delta_A, \Delta_B \cup F$, we get $\Delta \subseteq \Delta_A \subseteq A$ or $\Delta = \Delta_B \cup F$. In the latter case, applying the hypothesis to Δ_B, F, we get $\Delta = \Delta_B \subseteq B$, or $\Delta = F$, the last giving the contradiction that $\mathrm{rank}(\Delta) \leq \nu - 1$.

Suppose $\Pi \cap A \cap B \neq \varnothing$. We claim that $\Delta = \Delta_A \cup \Delta_B$. To see this, let Q be a cell of Π. We claim that $Q \subseteq \Pi_A$ or $Q \subseteq \Pi_B$. If not, by Helly property

(Lemma 2.2) applied to Q, Π_A, Π_B, we have $Q \cap \Pi_A \cap \Pi_B \neq \varnothing$. Now $Q \cap \Pi_A$ and $Q \cap \Pi_B$ are faces of the cube Q, and it follows that either $Q \subseteq \Pi_A$ or $Q \subseteq \Pi_B$ (formally a contradiction). Thus, $\Delta(Q) \subseteq \Delta_A \cup \Delta_B$, proving the claim. Moreover, $\Delta_A \cap \Delta_B \subseteq A \cap B$, and so we get either $\Delta = \Delta_A$ or $\Delta = \Delta_B$. □

Definition. We say that a finite median algebra, Π, of rank $\nu < \infty$ is ***fat*** if $\Delta(\Pi)$ is thick.

Of course, this can be viewed as a purely combinatorial property of Π. Note that it imples that $\text{rank}(E(W)) = \nu - 1$ for all $W \in \mathcal{W}(\Pi)$.

Suppose that $\phi : \Pi \longrightarrow \Pi'$ is an epimorphism of median algebras. This extends to an epimorphsim $\phi : \Delta(\Pi) \longrightarrow \Delta(\Pi')$ (by Lemma 3.3). In the proof of Lemma 2.1 we can restrict to subcomplexes, and so there is a minimal convex complex $C \subseteq \Delta(\Pi)$ with $\phi(C) = \Delta(\Pi')$. If $\text{rank}(\Pi) = \text{rank}(\Pi') = \nu < \infty$, then C is unique. In view of Lemma 3.10, we can apply the argument of Lemma 2.13, again restricted to subcomplexes. Writing $C = \Delta(\Pi'')$ where Π'' is a convex subset of Π, this shows:

Lemma 3.11. *Let $\phi : \Pi \longrightarrow \Pi'$ be an epimorphism of finite median algebras of rank $\nu < \infty$. Suppose that Π' is fat. Then there is a convex subset $\Pi'' \subseteq \Pi$ with is intrinsically fat and with $\phi\Pi'' = \Pi'$.*

In fact, we can say a bit more. Let $W \in \mathcal{W}(\Pi'')$. Then by Lemma 2.14, $\text{rank}(\phi(\Delta_W^- \cap \Delta_W^+)) \geq \nu - 1$. This means that we can find walls, $W_1, W_2, \ldots, W_{\nu-1} \in \mathcal{W}(\Pi')$ such that the corresponding walls $W_1', W_2', \ldots, W_{\nu-1}' \in \mathcal{W}(\Pi'')$ all pairwise cross, and all cross W. In particular, we can find a ν-cube, $Q \subseteq \Pi''$, with $\mathcal{W}(Q) = \{W_1', W_2', \ldots, W_{\nu-1}', W\}$. Note that $\phi|Q$ is either injective, or collapses the wall W.

(Of course, one could also give a combinatorial proof of the above statements.)

CCAT(0) complexes arise as subalgebras of median metric spaces in the following way. Let M be a median metric space and let $\Pi \subseteq M$ be a finite subalgebra. We say that a wall, $W \in \mathcal{W}(\Pi)$, is ***straight*** in M if c, d (or $[c, d]_M$) is straight in M for some, hence every, 1-cell, $\{c, d\} \in E(W)$. We say that Π is ***straight*** if each wall is straight. If $Q \subseteq M$ is a straight n-cube, then hull(Q) is an l^1 direct product of compact real intervals. In fact, if $a, b \in Q$ are opposite corners, then hull$(Q) = [a, b]_M$. If $e_1, \ldots, e_n$ are the adjacent corners of Q to a, then the map $\prod_{i=1}^n [a, e_i]_M \longrightarrow [a, b]_M$ which sends $(x_1, \ldots, x_n)$ to $x_1 \vee \cdots \vee x_n$ is an isomorphism to $[a, b]_M$. (Recall, from Section 2, that $[a, b]$ has the structure of a distributive lattice with initial point a.) Its inverse is given by sending $x \in [a, b]_M$ to the point whose ith coordinate is $\mu(a, e_i, x)$. (A coarse version of this construction will be used in Section 9.)

More generally, if $\Pi \subseteq M$ is straight, we write $\Upsilon(\Pi) = \Upsilon(\Pi, M)$ for the union of the sets hull(Q) as Q ranges of all cells of Π. The following is Lemma 6.1 of [Bo3]:

Lemma 3.12. *$\Upsilon(\Pi, M)$ is a subalgebra of M.*

Suppose now that M is connected. Then any straight interval in M is isometric to a compact real interval.

It follows that $\Upsilon(\Pi, M)$ is intrinsically isometric to $\Delta(\Pi)$ as defined above, where the width of each wall is determined by the induced metric on Π. (A more general construction, where Π is not assumed to be straight, is given in [Bo3]. This will only be used here in the proof of Lemma 5.1, where we elaborate on an argument in that paper.) Note that if Π is fat, then it is straight in M (since every 1-cell lies in a ν-cube).

This leads on to a more general notion.

Definition. An ***open CCAT(0) complex*** is a median algebra isomorphic to an open convex subset, $\Omega \subseteq \Delta = \Delta(\Pi)$, where Π is a finite median algebra.

We generally view Ω as coming equipped with a decomposition into cells, namely the intersection of Ω with the cells of Δ.

The closure, $\hat{\Omega} \subseteq \Delta$, is a CCAT(0) complex, and it is not hard to see that it is determined up to isomorphism by the median structure on Δ. In fact, if $\Delta_0 \subseteq \Omega$ is a compact convex subset containing all the 0-cells of Ω, then Δ_0 is isomorphic to $\hat{\Omega}$. Indeed, the gate map from $\hat{\Omega}$ to Δ_0 sends the 0-cells of $\hat{\Omega}$ isomorphically to the 0-cells of Δ_0. Note that Ω has a compact exhaustion by such convex sets.

We refer to the elements of $\mathcal{W}(\Pi)$ as ***walls*** of Ω, and denote the set of such walls by $\mathcal{W}(\Omega)$.

Let Ω be an open CCAT(0) complex. We can equip it with a complete median metric, inducing the median on Ω, as follows. Note that the walls of Ω fall into three classes, depending on whether the sides of Ω which cross it are compact intervals, half-open intervals or open intervals. If the last class is non-empty, then Ω is a product. For simplicity of exposition we will assume it to be empty. (This can be achieved by subdividing any such interval into two half-open intervals, so as to give rise to two parallel walls.) We then write $\mathcal{W}(\Omega) = \mathcal{W}_0(\Omega) \sqcup \mathcal{W}_\infty(\Omega)$, where $\mathcal{W}_0(\Omega)$ is the set of walls crossed by compact sides. If $w : \mathcal{W}_0(\Omega) \longrightarrow (0, \infty)$ is any map, then we can equip Ω with a complete median metric so that each side crossing $W \in \mathcal{W}_0(\Omega)$ is isometric to the real interval $[0, w(W)]$ and each side crossing $W \in \mathcal{W}_\infty(\Omega)$ is isometric to the ray $[0, \infty)$. This can be achieved, similarly as in the compact case: we can view Ω as a subcomplex of the product $\prod_{W \in \mathcal{W}_0(\Omega)} [0, w(W)] \times [0, \infty)^{\mathcal{W}_\infty(\Omega)} \subseteq \mathbb{R}^{\mathcal{W}(\Omega)}$ with the induced l^1 metric.

Definition. A ***CCAT(0) panel complex*** is an open CCAT(0) complex (as defined above) equipped with a compatible complete median metric.

By "compatible" we mean that it induces the given median on the CCAT(0) complex. We will refer to its cells as "panels" — a more general definition of this term will given in Section 4.

We can replace the l^1 metric on each panel with the l^2 metric with the same side-lengths, and take the induced path metric, σ. If Ω has dimension at most

ν, then $\sigma \leq \rho \leq \sigma\sqrt{\nu}$. The metric σ is CAT(0). This is another illustration of Theorem 2.15 (though it is directly verifiable in this case).

4. Cubulated sets

Cubulated sets provide a means of formulating some regularity results discussed in Section 5.

We begin by recalling a general definition from Section 1.

Definition. An n-***panel***, P, is (a metric space isometric to) an l^1 product, $P = \prod_{i=1}^n I_i$, where each $I_i \subseteq \mathbb{R}$ is a non-trivial closed connected subset.

In other words, each I_i is an interval, a ray, or all of $\mathbb{R}$. After subdividing, it is generally convenient to rule out the last case. If each I_i is a ray, we refer to P as an n-***orthant***.

We write $\mathrm{corner}(P) = \prod_{i=1}^n \partial I_i$, where ∂I_i denotes the boundary of $I_i \subseteq \mathbb{R}$. We refer to elements of $\mathrm{corner}(P)$ as ***corners*** of P. Note that they are determined by the median structure of P. An orthant has exactly one corner.

If P is a compact panel, and $Q = \mathrm{corner}(P)$, then Q is an n-cube, and $P = \mathrm{hull}(Q)$. Conversely, if $Q \subseteq M$ is a straight n-cube, then $P = \mathrm{hull}(Q)$ is a compact panel. Note the P is median isomorphic to $[0,1]^n$.

Let M be a complete connected median metric space of rank ν.

Definition. As subset of M is ***cubulated*** if it is a locally finite union of compact panels.

(This is slightly more general than the definition given in [Bo5] in that we are not assuming that the panels have dimension ν. This distinction will not matter in any essential way.) Clearly a cubulated set is locally compact, and closed in M.

A cubulated set can be identified with a subcomplex of a CCAT(0) complex in M as we now describe.

Recall that if $\Pi \subseteq M$ is a finite straight subalgebra, then $\Upsilon(\Pi)$ is a subalgebra isometric to $\Delta(\Pi)$ with the appropriate edge-lengths. Clearly this is a compact cubulated set.

Note that a cubulated set is closed and locally compact. If it is compact, then it is a union of a finite number of panels. Conversely, we have:

Lemma 4.1. *Let Ψ be a connected finite union of compact panels in M. Then there is a finite straight subalgebra, $\Pi \subseteq M$, such that $\Upsilon(\Pi) = \langle \Psi \rangle$. Moreover, Ψ is a subcomplex of $\Upsilon(\Pi) \subseteq M$. Also, $|\Pi|$ is bounded above in terms of ν and the number of panels of Ψ.*

Note that we can assume that $\Pi = \langle \Psi \cap \Pi \rangle$. After subdivision, we can also assume that Ψ is a full subcomplex of $\Upsilon(\Pi)$ with vertex set $\Psi \cap \Pi$.

Proof. Let $A \subseteq \Psi$ be the set of corners of all the panels of Ψ, and let $\Pi = \langle A \rangle$. We first note that Π is straight.

To see this, suppose that $W \in \mathcal{W}(\Pi)$. There is some panel P of Ψ which is crossed by W. If $\{c, d\} \in E(W)$ in Π, then $[c, d]_M$, is parallel to a subset of a side of P, hence has rank 1. This shows that Π is straight as claimed.

Now each panel of Ψ lies in $\Upsilon(\Pi)$, where it is subdivided into a collection of subpanels. See Lemma 3.3 of [Bo5] for more details. By construction, $A \subseteq \Pi \cap \Psi$, so $\Pi = \langle \Pi \cap \Psi \rangle$. Now $\Upsilon(\Pi)$ is isomorphic to the realisation, $\Delta(\Pi)$, of Π, so applying Lemma 3.5 of the present paper intrinsically to $\Upsilon(\Pi)$, we see that $\Upsilon(\Pi) = \langle \Psi \rangle$.

Note that $|A| \leq 2^\nu k$, and so $|\Pi| \leq 2^{2^{2^\nu k}}$, where k is the number of panels of Ψ. □

In view of Lemma 4.1, we will refer to a connected finite union of panels as a ***panel complex***.

We also note that if $\Psi \subseteq M$ is a simply connected cubulated set that is also a subalgebra then it is necessarily a CCAT(0) complex — it is easily verified that it satisfies the link condition.

Definition. A subset, $\Psi \subseteq M$ is ***uniformly cubulated*** if there is some $k \in \mathbb{N}$ such that if $A \subseteq \Psi$ is any bounded subset, there is subset $\Psi' \subseteq \Psi$, which is the union of at most k compact panels of M, and with $A \subseteq \Psi'$.

Note that there is no loss in assuming that the union of these panels is connected (since we can take the subset to be connected, and apply Lemma 4.1).

The following is used in the proof of Lemma 5.2.

Lemma 4.2. *Suppose that $\Psi \subseteq M$ is connected and uniformly cubulated. Let $\Upsilon = \langle \Psi \rangle \subseteq M$. Then $\Upsilon \subseteq M$ is a CCAT(0) panel complex, with Ψ as a subcomplex.*

(Note that we cannot assume a-priori that Υ is closed or locally compact, though this will retrospectively be a consequence of the conclusion.)

Proof. By definition, can write Ψ as an increasing union, $\Psi = \bigcup_{i=0}^{\infty} \Psi_i$, where each Ψ_i is a connected union of a bounded number of compact panels. By Lemma 4.1 and subsequent observations, Ψ_i is a subcomplex of $\Upsilon_i = \Upsilon(\Pi_i) = \langle \Psi_i \rangle$, where $\Pi_i \subseteq M$ is a finite straight subalgebra of M, with $|\Pi_i|$ bounded. (Recall that $\Upsilon(\Pi_i)$ can be identified with the realisation, $\Delta(\Pi_i)$, of Π_i as a cube complex.)

In fact, we can assume that there is a fixed element $p \in \Pi_i$ for all i. (Take any $p \in \Psi_0$ and replace Π_i by $\langle \Pi_i \cup \{p\} \rangle$. This has the effect of subdividing Υ_i. Moreover, $\langle \Pi_i \cup \{p\} \rangle$ remains bounded.)

Note that, after passing to a subsequence, we can suppose that each Π_i is isomorphic to a fixed median algebra Π_0. Let $f_i : \Pi_0 \longrightarrow \Pi_i$ be an isomorphism. Again, after passing to a subsequence, we can suppose that for each $x \in \Pi_0$, $f_i(x)$ is either bounded or else eventually leaves every bounded subset of M. Let $P \subseteq \Pi_0$ be the set of $x \in \Pi_0$ for which $f_i(x)$ is bounded. (This is a convex subset of Π_0.) Let $B \subseteq M$ be a bounded set with $f_i(P) \subseteq B$ for all i. By construction, if $B' \subseteq M$ is any other bounded set, then $B' \cap \Pi_i \subseteq B$ for all sufficiently large i.

We claim that if $C \subseteq M$ is any closed bounded convex set with $B \subseteq C$, then $\Upsilon \cap C = \langle \Psi \cap C \rangle$. The inclusion $\langle \Psi \cap C \rangle \subseteq \Upsilon \cap C$ is clear (since $\Upsilon \cap C$ is a subalgebra of M). We want to show that $\Upsilon \cap C \subseteq \langle \Psi \cap C \rangle$.

By Lemma 2.8, C is gated in M. Let $\omega : M \longrightarrow C$ be the gate map. This is an epimorphism by Lemma 2.3. Let $x \in \Upsilon \cap C$. Then $x \in \langle A \rangle$ for some finite subset $A \subseteq \Psi$. Let $a \in A$, and let $r = \rho(p, a)$. (Here ρ is the metric on M, and so restricts to an intrinsic median metric on the subalgebra Υ.) Now $\Pi_i \cap N(p; r) \subseteq B \subseteq C$ for all sufficiently large i. Choose some such i with $a \in \Psi_i$. Now Ψ_i is a subcomplex of $\Upsilon_i = \Upsilon(\Pi_i)$. By construction of $\Upsilon(\Pi_i)$, $a \in \Upsilon(Q)$ for some cell $Q \subseteq \Pi_i$, with $\Upsilon(Q) \subseteq \Psi_i$. Note that $\Upsilon(Q)$ is convex in M. Let b be the nearest point of $\Upsilon(Q)$ to p. This is the image of p under the gate map $\Upsilon(\Pi_i) \longrightarrow \Upsilon(Q)$. Since $p \in \Pi_i$ and $\Upsilon(\Pi_i) \cong \Delta(\Pi_i)$ it follows that $b \in Q$. Now $\rho(p, b) \leq \rho(p, a) \leq r$, so $b \in \Pi_i \cap N(p; r) \subseteq C$. In particular, this shows that $\Upsilon(Q) \cap C \neq \varnothing$. Since $\Upsilon(Q)$ is convex, by Lemma 2.9, we have $\omega(\Upsilon(Q)) = \Upsilon(Q) \cap C$. In particular, $\omega a \in \Upsilon(Q) \cap C \subseteq \Psi_i \cap C \subseteq \Psi \cap C$. This holds for all $a \in A$, and so $\omega(\langle A \rangle) = \langle \omega A \rangle \subseteq \langle \Psi \cap C \rangle$. Now $x \in C$, so $x = \omega x \in \omega(\langle A \rangle) \subseteq \langle \Psi \cap C \rangle$. We have shown that $\Upsilon \cap C = \langle \Psi \cap C \rangle$ as claimed.

Now $\Psi \cap C$ is compact, so we have $\Psi \cap C \subseteq \Psi_i$ for some (indeed all sufficiently large) i. Therefore $\Upsilon \cap C = \langle \Psi \cap C \rangle \subseteq \langle \Psi_i \rangle = \Upsilon_i = \Upsilon(\Pi_i)$. Now $\Upsilon \cap C = \Upsilon_i \cap C$ is an intrinsically closed subset of Υ_i. Therefore, $\Upsilon \cap C$ is itself a panel complex, $\Upsilon \cap C = \Upsilon(\omega' \Pi_i)$, where $\omega' : \Upsilon_i \longrightarrow \Upsilon \cap C$ is the gate map. (This follows by Lemma 3.2, on identifying Υ_i with the compact CCAT(0) complex, $\Delta(\Pi_i)$.) Let $\Pi_{\Upsilon \cap C}$ be the unique minimal subalgebra of $\Delta(\Pi_{\Upsilon \cap C}) \cong \Upsilon(\Pi_{\Upsilon \cap C})$ such that $\Upsilon \cap C = \Upsilon(\Pi_{\Upsilon \cap C})$ as given by Proposition 3.6. Note that $\Pi_{\Upsilon \cap C} \subseteq \omega' \Pi_i$. In particular, it follows that $|\Pi_{\Upsilon \cap C}|$ is bounded.

Now let $(C_i)_{i=0}^{\infty}$ be an exhaustion of M by finite-diameter closed convex sets all containing B. (Such exists, by Lemma 2.11.) Let $\Upsilon'_i = \Upsilon \cap C_i$ and $\Pi'_i = \Pi_{\Upsilon'_i}$, so that $\Upsilon'_i = \Upsilon(\Pi'_i)$ and $|\Pi'_i|$ is bounded. Note that $\Upsilon = \bigcup_{i=0}^{\infty} \Upsilon'_i$ is an exhaustion of Υ by compact subsets. We have gate maps $\omega_i : \Upsilon \longrightarrow \Upsilon'_i$. Note that if $i \geq j$, then $\Upsilon'_j = \Upsilon(\omega_j \Pi'_i)$ and so $\Pi'_j \subseteq \omega_j(\Pi'_i)$. (Again this follows by Lemma 3.2, on identifying Υ'_i with $\Delta(\Pi'_i)$.) In particular, $|\Pi'_j| \leq |\Pi'_i|$. Since $|\Pi'_i|$ is bounded, this number must eventually stabilise so that for all j and $i \geq j$, $\omega_j | \Pi'_i$ is an isomorphism from Π'_i to Π'_j. Moreover, $\omega_j = \omega_j \circ \omega_i$. Therefore, we have a fixed finite median algebra, Π, and an isomorphism $\phi_i : \Pi \longrightarrow \Pi'_i$ for all (sufficiently large) i. Moreover, if $i \geq j$, then $\omega_j \circ \phi_i = \phi_j$.

In summary, we have monomorphisms $\phi_i : \Pi \longrightarrow M$ with $\Upsilon \cap C_i = \Upsilon(\phi_i \Pi)$. Moreover, if $i > j$, then $\phi_j \circ \phi_i^{-1}$ is the gate map to $\Upsilon(\phi_j \Pi)$ restricted to $\phi_i \Pi$. Also, $\Psi \cap \Upsilon(\phi_i \Pi)$ is a subcomplex of $\Upsilon(\phi_i \Pi)$.

Let Π_B be the set of $x \in \Pi$ such that $\phi_i(x)$ is bounded, hence eventually constant, equal to $\phi(x)$ say. Then Π_B is a convex subset of Π, and $\phi : \Pi_B \longrightarrow$

$\Upsilon \subseteq M$ is a monomorphism. After subdivision, we can assume for simplicity that $\Pi_B \neq \varnothing$. We can identify $\Delta_B = \Delta(\Pi_B)$ as a subcomplex of $\Delta = \Delta(\Pi)$.

Let $\Delta_I \subseteq \Delta$ be the union of all closed cells of Δ which meet Δ_B. This is a convex subcomplex of Δ (possibly all of Δ). Let $\Pi_I = \Pi \cap \Delta_I$, so that $\Delta_I = \Delta(\Pi_I)$. For any i, let $\Upsilon''_i = \phi_i(\Pi_I) \subseteq \Pi'_i$, and let $\Upsilon''_i = \Upsilon(\Pi''_i) = \phi_i(\Delta_I)$. If $i \geq j$, then $\Upsilon''_j = \omega_j(\Upsilon''_i) \subseteq \Upsilon''_i$. Also $\Upsilon = \bigcup_{i=0}^{\infty} \Upsilon''_i$. (For suppose Q is any cell of Π with $Q \cap \Pi_B = \varnothing$. Fix any $p \in \Pi_B$. We have $\rho(p, \phi_i(\Delta(Q))) = \rho(p, \phi_i(Q)) \to \infty$. In other words, the corresponding cells of Υ'_i eventually leave every bounded set.)

Now let $\Delta_\infty \subseteq \Delta$ be the union of all open cells of Δ whose closures meet Δ_B. This is an open convex subset of Δ whose closure is Δ_I. As described at the end of Section 2, we can equip it with a complete median metric inducing the original metric on Δ_B. We can now find an exhaustion of Δ_∞ by compact convex subsets Δ_i, all containing Δ_B, such that Δ_i is isometric to Υ''_i. Moreover, these isometries commute with the respective gate maps to Δ_j and Υ''_j in Δ and Υ respectively. Taking the union of these isometries, we get an isometry from Δ_∞ to Υ. This shows that Υ is a CCAT(0) panel complex as claimed. □

5. Cubulation of top-dimensional submanifolds

The results of this section are analogues of regularity statements in [KlLe, KaKL].

Let M be a complete connected median metric space of rank ν. It is shown in [Bo5] that any closed subset of M homeomorphic to $\mathbb{R}^\nu$ is cubulated. If we add the stronger condition that the homeomorphism is bilipschitz, then we can achieve uniformity by a slight elaboration of the argument. For the statement below, it does not matter whether we equip $\mathbb{R}^\nu$ with the euclidean or the l^1 metric, as these are bilipschitz equivalent. For convenience, we use the euclidean metric for the proof.

Lemma 5.1. *Let $f : \mathbb{R}^\nu \longrightarrow M$ be an embedding of $\mathbb{R}^\nu$ into M which is κ-bilipschitz to its range. If $K \subseteq \mathbb{R}^\nu$ is compact, there is a compact panel complex, $\Psi \subseteq M$ with $f(K) \subseteq \Psi \subseteq f(\mathbb{R}^\nu)$, where the number of cells of Ψ is bounded above in terms of κ and ν.*

Proof. The argument follows that of Proposition 4.3 of [Bo5]. We can assume that $K \subseteq B_2 \subseteq B_1 \subseteq B_0$, where each B_i is a topological ball. In fact, we can take $B_2 = N(p; r)$, $B_1 = N(p; 2r)$ and $B_0 = N(p; 3r)$ for some $r > 0$. Since the statement is invariant under rescaling the domain and range, we may as well take $r = 1$.

We now triangulate $\partial B_0 \cong S^{\nu-1}$, and let $A \subseteq \partial B_0$ be the 0-skeleton. Let $\Pi = \langle f(A) \rangle \subseteq M$. As shown in [Bo3] (see also [Bo5]) there is a compact subset, $\Upsilon = \Upsilon(\Pi) \subseteq M$ containing Π, which in the induced metric is a CAT(0) cube complex, though the cells need not be straight. (This is a more general version of the construction described by Lemma 3.12 here.) If we take the triangulation to be sufficiently fine, then it was shown that there is a straight subcomplex, Ψ of Υ, with $B_2 \subseteq \Psi \subseteq B_0$. This suffices to show that $f(\mathbb{R}^\nu)$ is cubulated.

For the present statement we need in addition to bound $|A|$. This will in turn bound $|\Pi|$ and hence the number of cells of Υ as required.

To this end, let σ be the CAT(0) metric on M arising from Theorem 2.15 here, so that f is also $\sigma\sqrt{\nu}$-bilipschitz to (M, σ). To apply the argument of [Bo5], we need to construct a homotopy from $f|\partial B_0$ into Υ which does not meet $f(B_1)$. This can be achieved by coning along geodesic segments in (M, σ), while keeping A fixed. Provided the triangulation is fine enough, depending only on ν and κ, the images of all the cells will have diameter less than $1/\kappa\sqrt{\nu}$. Since $B_0 = N(B_1, 1)$, the image does not meet $f(B_1)$.

The other requirement of the construction of [Bo5] is that any cell of Υ which meets $f(B_2)$ should lie inside $f(B_1)$. It is enough that each cell has diameter less than 1, which can again be achieved by taking the triangulation sufficiently fine, depending only on ν and κ. □

In other words, we have shown that $f(\mathbb{R}^\nu)$ is uniformly cubulated.

Lemma 5.2. *Suppose that $f : \mathbb{R}^\nu \longrightarrow M$ is a κ-bilipschitz embedding. Let $\Upsilon = \langle f(\mathbb{R}^\nu)\rangle$. Then Υ is a thick CCAT(0) panel complex with $f(\mathbb{R}^\nu)$ as a subcomplex. Moreover, the number of cells of Υ is bounded above in terms of ν and κ.*

Proof. Let $\Psi = f(\mathbb{R}^\nu)$. By Lemma 5.1, Ψ is uniformly cubulated, so by Lemma 4.2, Υ is an CCAT(0) panel complex, and Ψ is a subcomplex. To see that Υ is thick, suppose that $\Upsilon = A \cup B$, and $A, B \subseteq \Upsilon$ is closed and convex in Υ. If $\dim(A \cap B) \leq \nu - 2$, then either $f(\mathbb{R}^\nu) \subseteq A$ or $f(\mathbb{R}^\nu) \subseteq B$, since otherwise $f^{-1}(A \cap B)$ would separate $\mathbb{R}^\nu$, giving a contradiction. Since $\Upsilon = \langle f(\mathbb{R}^\nu)\rangle$, it follows that $\Upsilon = A$ or $\Upsilon = B$. In other words, Υ is thick. □

We will also need the following in Section 11.

Lemma 5.3. *Let $\Upsilon \subseteq M$ be a CCAT(0) panel complex, and suppose that $F \subseteq M$ is closed subset homeomorphic to $\mathbb{R}^\nu$, and with $F \subseteq N(\Upsilon; t)$ for some $t \geq 0$. Then $F \subseteq \Upsilon$.*

Proof. Given that F is cubulated, this can be proven by a similar argument to that of Lemma 3.3 of [BeHS3].

Alternatively, we can again adapt the proof of Proposition 4.3 of [Bo5] as we now describe. For this we again use the CAT(0) metric, σ, on M, as given by Theorem 2.15 here. This induces the CAT(0) metric on Υ.

Let $f : \mathbb{R}^\nu \longrightarrow F$ be a homeomorphism. Let $a \in \mathbb{R}^\nu$ and let $r_0 > r_1 > 0$ be sufficiently large as described below. Let $B_i = N(a; r_i)$. Thus $a \in B_1 \subseteq B_0 \subseteq \mathbb{R}^\nu$. (We will not need the set "B_2" of [Bo5] for the current argument.)

We triangulate ∂B_0 so that the f-image of each simplex has diameter bounded (by 1, say). For each vertex, x, of the triangulation, we choose some $\phi x \in F$, with $\sigma(\phi x, \phi y)$ bounded. We now extend this to a map $\partial B_0 \longrightarrow \Upsilon$ using the fact that Υ is CAT(0), to cone off using geodesic segments. In this way, the ϕ-image of each simplex is again bounded. Since Υ is contractible, this extends to a continuous map, $\phi : B_0 \longrightarrow \Upsilon$. Now construct a homotopy from $f|\partial B_0$ to $\phi|\partial B_0$ in M by linear interpolation along geodesics segments. The trajectories of this homotopy all have bounded length. Provided that $r_1 - r_0$ is chosen big enough, this homotopy will not meet $f(B_1)$. By the same homology argument as in [Bo5], we see that $f(B_1) \subseteq \phi(B_0) \subseteq \Upsilon$. In particular, $f(a) \in \Upsilon$. Since a was arbitrary, it follows that $F = f(\mathbb{R}^\nu) \subseteq \Upsilon$ as required. □

6. Quasiflats in finite panel complexes

Let Ψ be a CCAT(0) panel complex of dimension ν, as defined in Section 4. By hypothesis, this has finitely many cells. The main aim of this section is to show that a quasiflat in Ψ is a bounded Hausdorff distance from a union of ν-cells of Ψ. This will be used in the proof of Theorem 1.1 in Section 11. (It could retrospectively be deduced from Corollary 1.4, though that would not give us computable constants.) The statement is easily seen to be equivalent for the median (l^1) metric, and for the CAT(0) (l^2) metric. For convenience in the proof, we will use the euclidean metric on $\mathbb{R}^\nu$ and the CAT(0) metric on Ψ.

We begin with some definitions. By a ***subcomplex*** of a CCAT(0) panel complex, we mean a connected finite union of cells of Ω, with the induced l^1 path metric. (Note that each cell is an isometrically embedded panel.) If each cell is contained in a ν-cell, we refer to it as a ν-***subcomplex***. In particular, a subcomplex is a panel complex as defined in Section 2.

Lemma 6.1. *Let (Ψ, ρ) be a finite CCAT(0) panel complex of dimension ν. Let $f : \mathbb{R}^\nu \longrightarrow \Psi$ be a quasi-isometric embedding. Then there is a subcomplex, $\Omega \subseteq \Psi$, which is a union of ν-cells of Ψ, and a quasi-isometry $f' : \mathbb{R}^\nu \longrightarrow \Omega$, which is a bounded distance from f in Ψ. Here the distance bound and the parameters of quasi-isometry of f' depend only on those of f and the number of cells of Ψ.*

We first observe that we can assume f to be continuous — after moving it a bounded distance. (This is standard construction. Take any triangulation of $\mathbb{R}^\nu$ with all simplices of bounded diameter. Map the 0-skeleton via f, then extend inductively over the i-skeleta by coning over vertices, using the fact that Ψ with the l^2 metric is CAT(0).)

The proof will make use of the fact that a zero-degree map from the n-sphere to itself identifies some pair of antipodal points. (See Theorem 15.1 for an account of this.)

The idea behind the proof is as follows.

We want to push f into a subcomplex of Ψ, where its image will be cobounded. We can focus on ν-cells of Ψ. This is because if a cell is not contained in any ν-cell, then f can only penetrate it a bounded distance — otherwise it would have to "fold over" and identify a pair of distant points. We can therefore push f off any such cell. If P is a ν-cell of Ψ, there is a dichotomy. Either P lies in the image of f, or else again f can only penetrate a bounded distance. This is because if f is not surjective, it must map with degree 0 to P, and so fold over as before. (This is where we use Theorem 15.1.) In the latter case, we again push f off P. We end up with a map which is surjective to subcomplex. The pushing operations move f a bounded distance, so it remains a quasi-isometric embedding.

However, for the above argument to work, we would need that the side-lengths of all the panels (equivalently, the widths of all the walls) are large in relation to the quasi-isometry constants of f. We cannot assume this a-priori. However, if we consider the set of widths of all walls as a subset of $[0, \infty]$, it must have a large gap somewhere within a bounded distance of 0. We now collapse down all the "narrow" walls, whose widths lie below this threshold. We then apply the above argument. We finally need to reinstate the narrow walls. At this last step, f might not remain surjective to a subcomplex, but its image will remain cobounded there.

The proof will begin by describing various constants which will determine which of the walls are deemed to be narrow.

Proof. We begin with a simple observation. Let P be a panel. Given $t > 0$, let $P_t = N(\partial P; t)$ (in the euclidean metric, σ). If each side of P has length at least $3t$, then we can find a continuous retraction, $g : P_t \longrightarrow \partial P$, such that $\sigma(x, gx) \leq \tau t$ for all $x \in P_t$, where τ depends only on the dimension. (We could take $\tau = \sqrt{\nu}$.)

Next we define a series of constants, whose significance will become apparent later. We first let λ, κ be the constants of quasi-isometry of f (which we assume to be already continuous). In other words, $\sigma(fx, fy) \leq \lambda\sigma(x, y) + k$ and $\sigma(x, y) \leq \lambda\sigma(fx, fy) + k$ for all $x, y \in \mathbb{R}^\nu$. Let $\mathcal{W}$ be the set of walls of Ψ, and write $\omega = |\mathcal{W}|$.

Given any $t \geq 0$, set $k_0 = k + 2\omega t$, $t_0 = (\lambda + 2)k_0$, $k_1 = k_0 + 2\tau t_0$, $t_1 = (\lambda + 2)k_1$, $k_2 = k_1 + 2\tau t_1$, $t_2 = (\lambda + 2)k_2$, ... , $k_{\nu-1} = k_{\nu-2} + 2\tau t_{\nu-2}$, $t_{\nu-1} = (\lambda + 2)k_{\nu-1}$. Let $G(t) = 3t_{\nu-1}$. Clearly the k_i and t_i are both increasing in i. Write G^i for the ith iterate of G. (Note that this all depends on the initial choice of t.) Now let $T_0 = G^\omega(0)$. Note that T_0 only depends on ν, ω, λ, k.

There is some $T \in (0, T_0)$ and a subset $\mathcal{W}_0 \subseteq \mathcal{W}$ such that $w(W) \leq T$ for all $W \in \mathcal{W}_0$, and $w(W) \geq G(T)$ for all $W \in \mathcal{W} \setminus \mathcal{W}_0$. (To see this, arrange the values of $w(W)$ for $W \in \mathcal{W}$, in order as $w_1 \leq w_2 \leq \cdots \leq w_\omega$. Note that $w_\omega = \infty$ (since Ψ unbounded). Set $w_0 = 0$, and let i be maximal such that $w_i \leq G^i(0)$. Let $T = G^i(0)$, so that $w_{i+1} > G^{i+1}(0) = G(T)$. Let $\mathcal{W}_0$ be the set of walls corresponding to $w_1, \ldots, w_i$.)

In the above definition of constants, we now set $t = T$, so that $k_0 = k + 2\omega T$, etc. Note that all of the constants t_i and k_i are bounded in terms of ν, ω, λ, k.

We now collapse the walls $\mathcal{W}_0$ to give us a new complex, Ψ_0, with CAT(0) metric σ_0, and a retraction, $\psi : \Psi \longrightarrow \Psi_0$ satisfying $\sigma_0(\psi x, \psi y) \leq \sigma(x, y) \leq \sigma_0(\psi x, \psi y) + 2\omega T$, for all $x, y \in \Psi$. The walls of Ψ_0 are naturally identified with $\mathcal{W} \setminus \mathcal{W}_0$. In particular, they all have width at least $G(T) = 3t_{\nu-1}$. Let

$f_0 = \psi \circ f : \mathbb{R}^\nu \longrightarrow \Psi_0$. This is a continuous (λ, k_0)-quasi-isometric embedding, where $k_0 = k + 2\omega T$.

Our main goal now is to move f_0 a bounded distance so that it maps surjectively onto a connected subcomplex of Ψ_0. We will proceed by backward induction on dimension, starting with ν-cells. We will finally go back to Ψ and use this to modify f.

In what follows, we will use the notation $F \sim_r F'$ to mean that $\sigma_0(F(x), F'(x)) \leq r$ for all $x \in X$, where F, F' are functions from some set X into Ψ_0.

Let P be a ν-cell of Ψ_0. Since $f_0 : \mathbb{R}^\nu \longrightarrow \Psi_0$ is proper, it has a well defined degree, d_P, to P. One way to describe this is as follows. Let $S = \mathbb{R}^\nu \cup \{\infty\}$ and $S_P = \operatorname{int}(P) \cup \{\infty\}$ be the one-point compactifications of $\mathbb{R}^\nu$ and $\operatorname{int}(P)$ respectively. These are both topological ν-spheres. We have a continuous map, $f_P : S \longrightarrow S_P$, which agrees with f_0 on $f_0^{-1}(\operatorname{int}(P))$ and sends everything else to ∞. Let $d_P \in \mathbb{Z}$ be the degree of this map. Note that if $d_P \neq 0$, then f_P is surjective, so $P \subseteq f_0(\mathbb{R}^\nu)$.

Suppose that $p \in \operatorname{int}(P)$ and $0 < r < \sigma(p, \partial P)$. Then $B = N(p; r)$ is an embedded euclidean ball in $\operatorname{int}(P)$. Let $S_B = B/\partial B$ and $\infty = \partial B/\partial B \in S_B$. (In other words, this is the one-point compactification of $\operatorname{int}(B)$.) There is a deformation retraction of P onto B. This induces a homotopy equivalence from S_P to S_B. Postcomposing f_P by this homotopy equivalence, we get a map $f_B : S \longrightarrow S_B$, which agrees with f on $f^{-1}(\operatorname{int}(B))$ and sends everything else to ∞. Note that its degree is equal to d_P.

Suppose that $d_P = 0$. We claim that $f_0(\mathbb{R}^\nu) \subseteq P_{t_0}$. In other words, if $a \in \mathbb{R}^\nu$, then $\sigma(f_0(a), \partial P) \leq t_0$. To see this, let $p = f_0(a)$ and suppose that $\sigma(p, \partial P) > t_0 = (\lambda + 1)k_0$. Let $B_0 = N(a; k_0) \subseteq \mathbb{R}^\nu$, and let $B = N(p; t_0) \subseteq \operatorname{int} P$. Since f_0 is (λ, k_0)-quasi-isometric, we see that $f_0(B_0) \subseteq \operatorname{int}(B)$.

We now use stereographic projection from $\mathbb{R}^\nu$ to the round ν-sphere to put a spherical metric on S so that B_0 is identified with a hemisphere of S centred on a. Note that any two antipodal points of S are greater than k_0 apart in the euclidean metric, σ_0.

Now f_B has degree 0. So by Theorem 15.1, there are antipodal points, $x, y \in S$, with $f_B(x) = f_B(y)$. We can assume that $x \in B_0$, so $f_0(x) \in \operatorname{int}(B)$, so $x, y \neq \infty$, so $f_0(x) = f_0(y) = f_B(x) = f_B(y)$. This now gives the contradiction that $\sigma_0(x, y) \leq k_0$. This proves the claim that $f_0(\mathbb{R}^\nu) \subseteq P_{t_0}$.

It now follows that if P is any ν-cell of Ψ, then either $P \subseteq f_0(\mathbb{R}^\nu)$ or $f_0(\mathbb{R}^\nu) \subseteq P_{t_0}$. Note that $t_0 \leq t_{\nu-1} = \frac{1}{3}G(T)$. Therefore, in the latter case, we can postcompose f_0 with the retraction $g : P_{t_0} \longrightarrow \partial P$, which moves every point a distance at most τt_0.

We now perform this retraction simultaneously for each ν-cell, P, which is not contained in $f_0(\mathbb{R}^\nu)$, and remove the interior of P. This gives us a subcomplex, $\Psi_1 \subseteq \Psi_0$, and a continuous map, $f_1 : \mathbb{R}^\nu \longrightarrow \Psi_1$, with $f_1 \sim_{\tau t_0} f_0$. Note that each

ν-cell of Ψ_1 is contained in $f_1(\mathbb{R}^\nu)$. Also f_1 is a (λ, k_1)-quasi-isometric embedding, where $k_1 = k_0 + 2\tau t_0$.

Now let P be a $(\nu-1)$-cell of Ψ_1, not contained in any ν-cell of Ψ_1. We claim that $f_1(\mathbb{R}^\nu) \subseteq P_{t_1}$, where $t_1 = (\lambda+1)k_1$. We argue similarly as before. Suppose $a \in \mathbb{R}^\nu$ with $\sigma(f_1(a), \partial P) > t_1$. Let $p = f_1(a)$, $B = N(p; t_1) \subseteq \mathrm{int}(P)$, and $B_0 = N(a; k_1) \subseteq \mathbb{R}^\nu$. Since f_1 is (λ, k_1)-quasi-isometric, we have $f_1(B_0) \subseteq \mathrm{int}(B)$. Let $S_B = B/\partial B$. This is a topological $(\nu-1)$-sphere. We have a map $f_B : S \longrightarrow S_B$. We put a round metric on S as before, so that B_0 gets identified with a hemisphere centred at a. Now S is a ν-sphere, so by the (standard) Borsuk-Ulam Theorem, there are antipodal points, $x, y \in S$, with $f_B(x) = f_B(y)$, and we derive a contradiction $k_1 < \sigma_0(x,y) \leq k_1$ as before. This proves the claim.

Now $t_1 \leq t_{\nu-1} < \frac{1}{3}G(T)$, so we can again retract P_{t_1} onto ∂P moving points a distance at most τt_1. We perform this for all such $(\nu-1)$-cells, P, and then remove their interiors to arrive at a subcomplex, $\Psi_2 \subseteq \Psi_1$, and a map $f_2 : \mathbb{R}^\nu \longrightarrow \Psi_2$, with $f_2 \sim_{\tau t_2} f_1$. Note that f_2 is a (λ, k_2)-quasi-isometric embedding, where $k_2 = k_1 + 2\tau t_1$. Moreover Ψ_2 has the property that each $(\nu-1)$-cell lies in a ν-cell and hence in $f_2(\mathbb{R}^\nu)$.

We now proceed to $(\nu-2)$-cells, to give us $\Psi_3 \subseteq \Psi_2$, and $f_3 : \mathbb{R}^\nu \longrightarrow \Psi_3$ with $f_3 \sim_{\tau t_2} f_2$. Each $(\nu-2)$-lies in a $(\nu-1)$-cell hence in a ν-cell hence in $f_3(\mathbb{R}^\nu)$.

We eventually arrive at $\Psi_{\nu-1} \subseteq \Psi_0$, and $f_{\nu-1} : \mathbb{R}^\nu \longrightarrow \Psi_{\nu-1}$ with $\Psi_{\nu-1} = f_{\nu-1}(\mathbb{R}^\nu)$. Note that $f_{\nu-1} \sim_{r_0} f_0$, where $r_0 = \tau(t_0 + t_1 + \cdots + t_{\nu-1})$.

We now let $\Xi \subseteq \Psi$ be the union of all ν-cells of Ψ which map to ν-cells of $\Psi_{\nu-1}$ under the map $\psi : \Psi \longrightarrow \Psi_0$ defined earlier. (Thus $\Xi_{\nu-1} = \psi(\Psi)$.) Let $f' : \mathbb{R}^\nu \longrightarrow \Xi$ be any map such that $\psi \circ f' = f_{\nu-1}$. Now $f' \sim_{r_0+2T} f$, and $r_0 + 2T \leq r_0 + 2T_0$, which depends only on ν, ω, λ, k. Since $f_{\nu-1}$ is surjective, $\Xi = N(f'(\mathbb{R}^\nu), 2\omega T)$, where $2\omega T \leq 2\omega T_0$ is again bounded in terms of ν, ω, λ, k.

Of course, Ξ need not be connected. To rectify this, we set $\Omega = \psi^{-1}\Psi_{\nu-1}$. Then Ω is a (connected) subcomplex of Ψ containing Ξ, and with $\Omega \subseteq N(\Xi, 2\omega T)$. $\square$

Note that if, initially, $f(\mathbb{R}^\nu)$ already lies in some subcomplex of Ψ, then by construction, Ξ will be a subset of this subcomplex.

One can also elaborate on the proof as follows. Suppose that $L \geq 0$ is some constant. In the proof on Lemma 6.1, we could take $T_0 = G^\omega(L)$ (instead of $G^\omega(0)$). We then set w_i to be the maximal i such that $w_i \leq G^i(L)$. Then set $T = G^i(L)$ so that $w_{i+1} > G^{i+1}(L) = G(T)$. We then proceed as before. In this case, we have ensured that $w(W) > L$ for each wall W of Ψ_0. In the penultimate paragraph, all the panels of Ξ will have all side-lengths greater than L. Of course now, the quasi-isometry constants and distance bound will depend also on L.

The argument above does not directly relate Ξ and Ω. But by running the argument twice, we arrive at the following addendum to Lemma 6.1:

Lemma 6.2. *Let $L \geq 0$. In the conclusion to Lemma 6.1, there is a subset, Ξ, of Ω such that Ω lies in an s-neighbourhood of Ξ in Ψ, and such that all side-lengths* *of all panels of Ξ are greater than L. Here s depends on L, the parameters of f and the number of panels of Ψ.*

Proof. First apply Lemma 6.1 to give us $f' : \mathbb{R}^\nu \longrightarrow \Omega \subseteq \Psi$. Now run the argument again with f' replacing f, and modified as above. We now take Ξ as in the penultimate paragraph of the proof. $\square$

By choosing L bigger than any of the finite side-lengths of Ω, we get the following immediate corollary to Lemma 6.2:

Corollary 6.3. *In the conclusion to Lemma 6.1, there is a subset, Ξ of Ψ consisting of a union of ν-cells which are all orthants such that Ψ lies in a finite metric neighbourhood of Ξ.*

Let $\hat{\Psi}$ be the quotient complex of Ψ obtained by collapsing all walls of finite width. This is a CCAT(0) orthant complex. Let $\hat{\Xi} \subseteq \hat{\Psi}$ be the image of the set Ξ as given by Corollary 6.3. Postcomposing f' with projection to $\hat{\Psi}$, we get a quasi-isometry of $\mathbb{R}^\nu$ to the orthant complex $\hat{\Xi}$. We claim that Ξ is homeomorphic to $\mathbb{R}^\nu$:

Lemma 6.4. *An orthant complex quasi-isometric to $\mathbb{R}^\nu$ is bilipschitz equivalent to $\mathbb{R}^\nu$.*

Proof. This is a simple application of asymptotic cones, which we discuss in Section 10. We refer to that section for definitions.

Let $\theta : \mathbb{R}^\nu \longrightarrow \Omega$ be a quasi-isometry to an orthant complex, Ω. Passing to asymptotic cones with a fixed basepoint, we get a bilipschitz homeomorphism, $\theta^\infty : (\mathbb{R}^\nu)^\infty \longrightarrow \Omega^\infty$. Note that the isometry classes of $\mathbb{R}^\nu$ and Ω are both invariant under rescaling, and so they are both are isometric to their respective asymptotic cones. We therefore get a bilipschitz homeomorphism from $\mathbb{R}^\nu$ to Ω as required. $\square$

The above raises the following question:

Question. Is any self-quasi-isometry of $\mathbb{R}^\nu$ a bounded distance from a bilipschitz map?

In the above one could in addition ask for the constants of the conclusion to depend only on those of the quasi-isometry. (One could also weaken it to ask only for a homeomorphism.) One can show that this holds if $\nu \leq 2$, though I know of no result that would imply this for any $\nu \geq 3$.

One consequence of an affirmative answer would be that we could strengthen the conclusion Theorem 1.1 to say that f factors up to bounded distance though a bilipschitz embedding of $\mathbb{R}^\nu$ into Ω.

7. Coarse median spaces

We first introduce some general conventions regarding coarse geometry.

Let (Λ, ρ) be a metric space. Given subsets, $A, B \subseteq \Lambda$ and $r \geq 0$, we write $A \subseteq_r B$ to mean $A \subseteq N(B; r)$. We write $A \sim_r B$ to mean $A \subseteq_r B$ and $B \subseteq_r A$, in other words, $\mathrm{hd}(A, B) \leq r$, where "hd" denotes Hausdorff distance. If $x \in \Lambda$, we write $x \in_r A$ to mean $x \in N(A; r)$. If $x, y \in \Lambda$, we write $x \sim_r y$ to mean $\rho(x, y) \leq r$, in other words, $\{x\} \sim_r \{y\}$. If $f, g : X \longrightarrow \Lambda$ are functions from some set X, we write $f \sim_r g$ to mean that $f(x) \sim_r g(x)$ for all $x \in X$.

Note that $A \subseteq_r B \subseteq_r C$ implies $A \subseteq_{2r} C$ etc. We will sometimes omit reference to a particular constant, and just write $\subseteq_*$, $\sim_*$ or $\in_*$. The interpretation is that "$*$" substitutes for some constant whose value depends only on the constants introduced into the discussion. We can then behave as though the relations $\subseteq_*$ and $\sim_*$ were transitive. One can get explicit bounds for the value of $*$ at any particular point in the argument, by counting the number of times we have applied transitivity etc., but we won't usually keep track of this.

Next we recall the notion of a coarse median space.

Let (Λ, ρ) be a geodesic space, and let $\mu : \Lambda^3 \longrightarrow \Lambda$ be a ternary operation. If (Π, μ_Π) is a median algebra (or indeed, any set equipped with a ternary operation) we say that a map, $\lambda : \Pi \longrightarrow \Lambda$ is an h-***quasimorphism*** if for all $x, y, z \in \Pi$, we have $\lambda\mu_\Pi(x, y, z) \sim_h \mu(\lambda x, \lambda y, \lambda z)$.

We say that Λ is a ***coarse median space*** (of ***(coarse median) rank*** at most ν) if the following hold:

(C1): There is some $k \geq 0$ and $h_0 \geq 0$ such that for all $x, y, z, x', y', z' \in \Lambda$, we have $\rho(\mu(x, y, z), \mu(x', y', z')) \leq k(\rho(x, x') + \rho(y, y') + \rho(z, z')) + h_0$.

(C2): There is a function $h : \mathbb{N} \longrightarrow [0, \infty)$ such that if $n \in \mathbb{N}$ and $A \subseteq \Lambda$ with $|A| \leq n$ then there is a finite median algebra Π (of rank at most ν), a map $\pi : A \longrightarrow \Pi$, and an $h(n)$-quasimorphism, $\lambda : \Pi \longrightarrow \Lambda$, such that $\rho(a, \lambda\pi a) \leq h(n)$ for all $a \in A$.

After modifying μ up to bounded distance, we can also assume that μ is symmetric in the arguments, and that $\mu(x, x, y) = x$ for all $x, y \in \Lambda$.

We will usually use a variant of (C2) in this paper. We can assume in (C2), that $\lambda\pi a = a$ for all $a \in A$. (For this, one can take Π and adjoining A as a disjoint set of vertices, attaching each $a \in A$ by an edge to $\pi(a) \in \Pi$. One then replaces Π by this new complex, and replace π by the inclusion map of A.) We can also assume that λ is injective (moving points a small distance), and then take $\Pi \subseteq \Lambda$, replacing λ by the inclusion map. In this way, we end up with $A \subseteq \Pi \subseteq \Lambda$. While this construction is a bit artificial and not strictly necessary, it will allow us to simplify notation, suppressing mention of the maps π and λ.

Definition. An h-***quasisubalgebra*** of Λ is a finite subset, $\Pi \subseteq \Lambda$, equipped with the structure of a median algebra, such that the inclusion of Π into Λ is an h-quasimorphism.

Therefore, in the previous paragraph, $\Pi \subseteq \Lambda$ is an $h(n)$-quasisubalgebra.

One consequence is that any tautological identity in a median algebra holds up to bounded distance in a coarse median space. More precisely, suppose we have a formula equating two expressions involving the ternary relation μ, which holds for any values of the arguments in any median algebra. (In fact, it is sufficient to check it for the two-point median algebra, $\{0,1\}$.) Then if we substitute arguments in Λ instead, the two sides of the formula give points in Λ which are a bounded distance apart. Here the bound depends only on the complexity of the expressions and the parameters of Λ (that is k, h_0 and h). A similar statement holds for conditional identities. In other words, suppose some finite set of median relations tautologically imply another relation in any median algebra. If the initial relations hold up to bounded distance for some assignment of arguments in Λ, then the resulting relation also holds up to bounded distance in Λ with these arguments. For more formal statements of these principles, see [Z1, Bo5, NibWZ1].

Suppose that $\Pi \subseteq \Lambda$ is a (by definition, finite) h-quasisubalgebra. Given a wall, $W \in \mathcal{W}(\Pi)$, write $w(W,\Lambda) = \min\{\rho(c,d)\}$ as c,d ranges over all edges of Π which cross W. Since any two such c,d are parallel, it follows from Property (C1), that $\rho(c',d')$ agrees with $w(W,\Lambda)$ up to linear bounds, for any such edge c',d'. (See Section 9 for further discussion of this.) It is convenient to assume that $w(W,\Lambda) > 0$ (which is possible since we are really only interested in its value up to linear bounds). As discussed in Section 3 can now put a median metric on Π so that the width of each wall W is exactly $w(W,\Lambda)$. It now follows (again using Property (C1)), that the inclusion of Π into Λ is quasi-isometric, where the constants only depend on h and the parameters of Λ.

Given a map $f : C \longrightarrow \Lambda$ from any set, C, to Λ, we say that f is t-***separating*** if $\rho(fa, fb) > t$ for all distinct $a,b \in C$. A subset of Λ is t-***separated*** if its inclusion into Λ is t-separating.

Suppose again that $\Pi \subseteq \Lambda$ is an h-quasisubalgebra. Let $w(\Pi,\Lambda) = \min\{w(W,\Lambda) \mid W \in \mathcal{W}(\Pi)\}$. Clearly, if Π is t-separated, then $w(\Pi,\Lambda) \geq t$. Conversely, there is a linear function $[t \mapsto w(t)]$, such that if $w(\Pi,\Lambda) \geq w(t)$, then Π is t-separated.

Definition. An h-***quasicube of dimension*** n is an h-quasimorphism, $\phi : Q \longrightarrow \Lambda$, from an n-cube, Q, into Λ.

We will also refer to $Q \subseteq \Lambda$ as an h-***quasicube***, where $Q \subseteq \Lambda$ is an h-quasisubalgebra isomorphic to an n-cube.

If Λ is coarse median of rank at most ν, then for any $h \geq 0$ there is some $t \geq 0$ such that any t-separated h-quasicube has dimension at most ν. (Although we won't be needing it here, there is also a converse to this statement which characterises the rank of a coarse median space, see [Bo5, NibWZ1].)

We will need the following slightly technical statement, which will allow us to improve on a quasimorphism constant at the cost of enlarging the domain. For this statement, Λ could be any metric space with a ternary operation, $\mu = \mu_\Lambda$.

Lemma 7.1. *There are functions, $t, r : \mathbb{N} \longrightarrow [0, \infty)$ with the following property. Suppose that $\Pi_1 \subseteq \Pi_2 \subseteq \Lambda$ are both finite median algebras with respective medians μ_1 and μ_2. (We do not assume the inclusion of Π_1 into Π_2 to be a homomorphism.) Suppose that $|\Pi_1| \leq n$ and that $\Pi_2 = \langle \Pi_1 \rangle$. Suppose that for some $k \geq 0$, the inclusions of Π_1 and Π_2 into Λ are both k-quasimorphisms (to (Λ, μ_Λ)), and that Π_1 is $(kt(n))$-separated. Then there is a median epimorphism, $\omega : \Pi_2 \longrightarrow \Pi_1$, such that $\omega|\Pi_1$ is the identity and $\rho(x, \omega x) \leq kr(n)$ for all $x \in \Pi_2$.*

Proof. Write μ_i for the median on Π_i. Since $\Pi_2 = \langle \Pi_1 \rangle$, every element of Π_2 can be written as an expression involving the median, μ_2, with arguments in Π_1. Moreover, we can take the expression to have bounded complexity, in terms of n. (This follows from Lemma 2.1 as observed there.)

Let $x \in \Pi_2$. Write x as such an expression. Now replace μ_2 by μ_1. This gives us a median expression in Π_1. Evaluate this in Π_1 to give us some $x' \in \Pi_1$.

Since the inclusions of Π_1 and Π_2 are both k-quasimorphisms, and the expression has bounded complexity, it follows that $\rho(x, x')$ is bounded by some fixed multiple of k. (For example, at the first step, suppose that $x = \mu_2(a, b, c)$ with $a, b, c \in \Pi_1$. The construction gives us $x' = \mu_1(a, b, c)$. Note that $x = \mu_2(a, b, c) \sim_k \mu_\Lambda(a, b, c) \sim_k \mu_1(a, b, c) = x'$, so that $\rho(x, x') \leq 2k$. The general case follows by a similar principle, effectively iterating this procedure.)

Suppose we chose a different expression for x (again of bounded complexity) so as to give us instead $x'' \in \Pi_1$. Then $\rho(x', x'') \leq \rho(x, x') + \rho(x, x'')$ is also bounded by some multiple of k. Therefore, provided $t(n)$ is chosen large enough, we will necessarily have $\rho(x', x'') < kt(n)$. Since Π_1 is $(kt(n))$-separated we get $x' = x''$. In other words, x' is uniquely determined, and we set $\omega x = x'$.

Suppose that $a \in \Pi_2$. In this case, the expression $x = a$ (vacuously involving μ_2) gives us $\omega x = x' = a$. In other words, $\omega a = a$, so that $\omega : \Pi_1 \longrightarrow \Pi_1$ is the identity on Π_1 as required.

To see that ω is a homomorphism, let $x, y, z \in \Pi_2$, and set $m = \mu_2(x, y, z)$. Write each of x, y, z as an expression involving μ_2 with arguments in Π_1. Substituting into $\mu_1(x, y, z)$ we get a similar expression for m, again of bounded complexity. Now apply the same procedure as before: replace μ_2 by μ_1, and evaluate in Π_1. This gives us $m' \in \Pi_1$. In fact, we see that $m' = \mu_1(x', y', z')$, where by construction, $x' = \omega x$, $y' = \omega y$ and $z' = \omega z$. This time the complexity of the expression for m' might be larger than considered before, but only by a controlled amount. We therefore get that $\rho(m', \omega m)$ is bounded by some multiple of k. Again, if $t(n)$ is large enough, we get $\rho(m', \omega m) < kt(n)$, so $m' = \omega m$. In other words, $\omega \mu_2(x, y, z) = \mu_1(\omega x, \omega y, \omega z)$, so ω is a homomorphism as claimed. Since $\omega|\Pi_1$ is the identity, it is an epimorphism. $\square$

We now suppose again that Λ is a coarse median space of rank at most ν.

Lemma 7.2. *There are functions, $h, t, r : \mathbb{N} \longrightarrow [0, \infty)$ with the following property. Suppose that $\Pi_1 \subseteq \Lambda$ is a median algebra with $|\Pi_1| \leq n$, and whose inclusion into Λ is a $(\kappa t(n))$-separating κ-quasimorphism for some $\kappa \geq h(n)$. Then there is an $h(n)$-quasisubalgebra, $\Pi_2 \subseteq \Lambda$, of rank ν, with $|\Pi_2| \leq 2^{2^n}$, together with an epimorphism: $\omega : \Pi_2 \longrightarrow \Pi_1$, such that $\rho(x, \omega x) \leq \kappa r(n)$ for all $x \in \Pi_2$.*

Proof. We take h to be the function given by Property (C2), and t, r as given by Lemma 7.1. Applying (C2) with $A = \Pi_1$ we get an $h(n)$-quasisubalgebra, $\Pi_2 \subseteq \Lambda$, with $\Pi_1 \subseteq \Pi_2$ (not necessarily a subalgebra). Moreover, we can suppose that $\Pi_2 = \langle \Pi_1 \rangle$, so $|\Pi_2| \leq 2^{2^n}$. Lemma 7.1 then gives us an epimorphism $\omega : \Pi_2 \longrightarrow \Pi_1$ with $\rho(x, \omega x) \leq \kappa r(n)$ for all $x \in \Pi_2$ as required. □

Recall the notion of "fatness" introduced in Section 3.

Lemma 7.3. *In Lemma 7.2, if Π_1 is fat, we can take Π_2 to be fat.*

Proof. Lemma 7.2 gives us an epimorphism $\omega : \Pi_2 \longrightarrow \Pi_1$. By Lemma 3.11, there is a convex subset, $\Pi'' \subseteq \Pi_2$, with Π'' intrinsically fat, and with $\omega\Pi'' = \Pi_1$. We now replace Π_2 by Π'' and denote it again by Π_2. □

For future reference (in the proof of Lemma 12.2) we remark that if Λ is a connected median metric space, then we can just set $h \equiv 0$ in Lemma 7.2.

Further refinements of Lemmas 7.2 and 7.3 are given in Section 9.

8. Coarse intervals

In the next two sections, we discuss coarse intervals and cubes in coarse median spaces. Much of what we do can be viewed as taking facts about median algebras and observing that the corresponding statements hold up to bounded distance in a coarse median space, via the general principle discussed in Section 7. We will refer to this process as "coarsifying". Further discussion of these notions can be found in [NibWZ1, NibWZ2].

Let Λ be a coarse median space. Given $a, b \in \Lambda$, let $[a, b] = [a, b]_\Lambda = \{\mu(a, b, x) \mid x \in \Lambda\}$.

Definition. $[a, b]$ is the ***coarse interval*** from a to b.

The map $\omega = \omega_{a,b} : \Lambda \longrightarrow [a, b]$ defined by $\omega(x) = \mu(a, b, x)$ is a median quasimorphism. Since Λ is connected, it follows that any two points, $x, y \in [a, b]$ are connected an r_0-***path*** in $[a, b]$; that is, a sequence of points $x = x_0, x_1, \dots, x_n = y$, with $x_i \in [a, b]$ and $x_i \sim_{r_0} x_{i+1}$ for all i. Indeed we can arrange for n to be linearly bounded above in terms of $\rho(x, y)$.

Up to bounded distance, $[a, b]$ is closed under μ. In other words, given $x, y, z \in [a, b]$, there is some $w \in [a, b]$ with $w \sim_* \mu(x, y, z)$. When dealing with a single interval, in order to simplify notation, we will assume that we always have $\mu(x, y, z) \in [a, b]$. Up to bounded distance, this makes no difference.

Here is another way of describing coarse intervals. Given $r \geq 0$, let $[a,b]_r = \{x \in \Lambda \mid \mu(a,b,x) \sim_r x\}$.

Lemma 8.1.
(1) Given $r \geq 0$, $[a,b]_r \subseteq_{r'} [a,b]$, where r' depends only on r and the parameters of Λ.
(2) There are some r_0, r'_0 depending only on the parameters of Λ such that $[a,b] \subseteq_{r'_0} [a,b]_{r_0}$.

In particular, for all sufficiently large r, we have $[a,b] \sim_* [a,b]_r$.

Proof. This follows by coarsifying the fact that in a median algebra, these correspond to equivalent ways of defining an interval: namely, $\{\mu(a,b,x) \mid x \in \Lambda\} = \{x \in \Lambda \mid \mu(a,b,x) = x\}$. Note that this equivalence can be expressed in terms of median expressions of bounded complexity. □

One can also check similarly that if $a \sim_* a'$ and $b \sim_* b'$, then $[a,b] \sim_* [a',b']$.

Given $x, y \in [a,b]$, write $x \wedge y = \mu(a,x,y) \in [a,b]$ and $x \vee y = \mu(b,x,y) \in [a,b]$. Coarsifying the corresponding statements in a median algebra discussed in Section 2, we see that with these operations, $[a,b]$ satisfies the axioms of a distributive lattice up to bounded distance. In other words, $(x \wedge y) \wedge z \sim_* x \wedge (y \wedge z)$, $x \wedge (x \vee y) \sim_* x$ etc. We write $x <_r y$ to mean that $x \sim_r x \wedge y$. Thus, $<_*$ is a partial order on $[a,b]$ up to bounded distance. In particular, $x <_* y <_* z$ implies $x <_* z$. Also $x <_* y$ and $y <_* x$ together imply $x \sim_* y$. In fact, the following statements are all equivalent up to bounded distance: $x <_* y$, $y \sim_* x \vee y$, $x \in_* [a,y]$, $[a,x] \subseteq_* [a,y]$, $y \in_* [b,x]$, $[b,y] \subseteq_* [b,x]$. Of course, the constants may change as we move between these statements. Note that if $x <_* y <_* z$, then $\mu(x,y,z) \sim_* y$.

Definition. We say that $[a,b]$ (or a,b) is s-***straight*** if given any $x, y \in [a,b]$ we have $x <_s y$ or $y <_s x$.

In other words, $<_*$ is a total order up to bounded distance.

Lemma 8.2. *Suppose that $[a,b] \subseteq \Lambda$ is s-straight. Let $J = [0, \rho(a,b)] \subseteq \mathbb{R}$. There is a quasimorphism, $\beta : J \longrightarrow [a,b]$, with image cobounded in $[a,b]$, and with $\beta(0) = a$ and $\beta(\rho(a,b)) = b$. Moreover, β is a quasi-isometric embedding. All the constants of the conclusion depend only on s and the parameters of Λ.*

Proof. Recall that any two points of $[a,b]$ are connected by an r_0-path for some r_0 depending only on the parameters of Λ. In particular, we can find an r_0-path, $a = y_0, y_1, \ldots, y_m = b$ from a to b.

Let $t \geq 0$ be chosen sufficiently large, depending on s and the parameters of Λ, as described below. We can suppose that $\rho(a,b) > t$. Let $A \subseteq [a,b]$ be a maximal t-separated subset with $a, b \in A$. For $x, y \in A$ write $x < y$ to mean that $x \neq y$ and $x <_s y$. We claim that $<$ is a total order on A. To see this, note that by straightness, we have $x < y$, $y < x$ for all $x \neq y$. Moreover, if $x < y < z < x$, then we must have $x \sim_* y \sim_* z \sim_* x$, which is a contradiction if t is chosen sufficiently large. We similarly get a contradiction to $x < y < x$. It now also follows that $x < y < z$ implies $x < z$. This proves the claim.

Now suppose that we have points $x_i \in A$ with $a = x_0 < x_1 < \cdots < x_n = b$. We claim that for all i, $y_i <_s x_i$. This follows by induction. For suppose $y_i <_s x_i$. Since $[a,b]$ is s-straight, if $y_{i+1} <_s x_{i+1}$ does not hold, then $x_{i+1} <_s y_{i+1}$ and so $y_i <_s x_i <_s x_{i+1} <_s y_{i+1} \sim_{r_0} y_i$, so $x_i \sim_* x_{i+1}$, which again is a contradiction if t is large enough. In particular, $y_i \neq b$ for $i < n$. Since $y_m = b$, we must have $n \leq m$. Since m is fixed, this places a bound on the length of such a sequence $(x_i)_i$. Since $<$ is a total order on A, it follows that A is finite. We can therefore retrospectively assume that $\{x_0, x_1, \ldots, x_n\} = A$. In particular, if $y \in [a,b]$, we have $y \sim_t x_i$ for some i.

We next claim that $\rho(x_i, x_{i+1})$ is bounded. Note that there must be some j such that $y_j \sim_t x_k$ for $k \leq i$ and $y_{j+1} \sim_t x_l$ for some $l \geq i+1$. Now $x_k \sim_t y_j \sim_{r_0} y_{j+1} \sim_t x_l$, so $x_l \sim_* x_k <_* x_i <_* x_{i+1} <_* x_l$, and so $x_i \sim_* x_{i+1}$ as claimed.

Next, we claim that $|i - j|$ is linearly bounded above in terms of $\rho(x_i, x_j)$. We can suppose that $j = i + q$ for $q > 0$. We can connect x_i to x_{i+q} by an r_0-path $x_i = z_0, z_1, , \ldots, z_p = x_{i+q}$, where p is linearly bounded in terms of $\rho(x_i, x_{i+q})$. By a similar argument as before, we see that $z_l <_s x_{i+l}$ for all l. In particular, we get $q \leq p$ as claimed.

We have now shown that $(x_i)_i$ is a quasigeodesic sequence in Λ with cobounded image in $[a,b]$. Moreover, if $i < j < k$, then $\mu(x_i, x_j, x_k) \sim_* x_j$. We can now divide J into n equal intervals, and interpolate to give the required quasi-isometric quasimorphism of J into $[a,b]$. □

9. Coarse cubes

In this section, we describe the geometry of coarse hulls of quasicubes in a coarse median space. (See also, [NibWZ1, Bo8].) Let $Q \subseteq \Lambda$ be an h-quasicube of dimension n. Recall the definition $w(Q, \Lambda) = \min\{w(W, \Lambda) \mid W \in \mathcal{W}(Q)\}$. If $a, b \in Q$ is an (ith) face of Q, we refer to $[a,b] \subseteq \Lambda$ as an (ith) ***side*** of Q.

Definition. We say that an h-quasicube, $Q \subseteq \Lambda$, is s-***straight*** if all its sides are s-straight.

We claim that any sufficiently separated h-quasicube of dimension ν is coarsely straight. More precisely:

Lemma 9.1. *There are functions, $t_0, s : \mathbb{N} \longrightarrow [0, \infty)$, with the following property. Let Λ be a coarse median space of rank ν. There is some $h_0 \geq 0$ depending only on the parameters of Λ, such that the following holds. Suppose $h \geq h_0$. If $Q \subseteq \Lambda$ is an h-quasicube of dimension ν with $w(Q, \Lambda) \geq h t_0(\nu)$, then Q is $hs(\nu)$-straight.*

Proof. Let c, d be face of Q. Let $x, y \in [c,d] \subseteq \Lambda$. Let $\Pi \subseteq \Lambda$ be a quasisubalgebra with $Q \cup \{x, y\} \subseteq \Lambda$ (the parameters of which depend only on those of Λ). Let

$\Pi' \subseteq \Pi$ be the subalgebra of Π generated by Q. Using Lemma 7.1 applied to $Q \subseteq \Pi' \subseteq \Lambda$, if t is large enough (in relation to $|Q| = 2^\nu$), then there is an epimorphism, $\omega : \Pi' \longrightarrow Q$ such that $\omega(x) \sim_* x$ for all $x \in Q$. By Lemma 2.6, there is a ν-cube $Q' \subseteq \Pi'$ with $\omega|Q'$ an isomorphism to Q. Let $c', d' \in Q'$ be such that $\omega(c') = c$ and $\omega(d') = d$. Thus, c', d' is a face of Q', and $c \sim_* c'$ and $d \sim_* d'$. Note that $[c', d']_\Pi$ is straight in Π (since Q' is also a ν-cube in Π and $\mathrm{rank}(\Pi) = \nu$). Let $x' = \mu_\Pi(c', d', x)$ and $y' = \mu_\Pi(c', d', y)$. Then $x' \sim_* x$ and $y' \sim_* y$. Up to swapping x and y, we have $x' \leq y'$ in $[c', d']_\Pi$. In other words, $x' = \mu_\Pi(c', x', y')$. It follows that $x \sim_* x' \sim_* \mu_\Lambda(c', x', y') \sim_* \mu_\Lambda(c, x, y)$. In other words, $x <_s y$ in $[c, d] \subseteq \Lambda$, where s depends only on the parameters of Λ. Since $x, y \in [c, d]$ were arbitrary, $[c, d]$ is s-straight, where s is some fixed multiple, $hs(\nu)$, of h, as required. ∎

Definition. A subset, $C \subseteq \Lambda$, is r-***convex*** if $[a, b] \subseteq_r C$ for all $a, b \in C$. We say that C is ***coarsely convex*** if it is r-convex for some $r \geq 0$.

Note that any interval, $[a, b] \subseteq \Lambda$ is coarsely convex. (This follows by coarsifying the statement that any interval in a median algebra is convex.)

Let $Q \subseteq \Lambda$ be an h-quasicube of dimension $n \leq \nu$. Let a, b be opposite corners of Q. Then $Q = [a, b]_Q$, and so $Q \subseteq_* [a, b] = [a, b]_\Lambda$. Since intervals are coarsely convex, we see that if $a', b' \in Q$, then $[a', b'] \subseteq_* [a, b]$. In particular, if a', b' is another pair of opposite corners of Q, then $[a', b'] \sim_* [a, b]$. Thus, up to bounded distance, we have a well defined subset, $H(Q) \subseteq \Lambda$, with $H(Q) \sim_* [a, b]$. We can assume that $Q \subseteq H(Q)$. Thus $H(Q)$ can be thought of as the "coarse hull" of Q in the following sense.

Lemma 9.2. *$Q \subseteq H(Q)$, and $H(Q)$ is r-convex, where r depends only on h and the parameters of Λ. If $H' \subseteq \Lambda$ is r'-convex and $Q \subseteq H'$, then $H(Q) \subseteq_{r''} H'$, where r'' depends only on h, r' and the parameters of Λ.*

The conclusion of Lemma 9.2 characterises $H(Q)$ up to bounded distance. (It is an instance of a more general construction of coarse hulls in a coarse median space of finite rank [Bo8].)

In fact, $H(Q)$ is a direct product of its sides, in the following sense.

Let $\mathcal{W}(Q) = \{W_1, \ldots, W_n\}$, and let a, c_i be the face of Q containing a and crossing W_i. In this way, Q is a direct product, $\prod_{i=1}^n \{a, c_i\}$, of the two-point median algebras $\{a, c_i\}$. Let $l_i = w(W_i)$, as defined in Section 7. Up to linear bounds, l_i agrees with $\rho(a, c_i)$. Let $I_i = [a, c_i] \subseteq \Lambda$. We can suppose for convenience that I_i is closed under μ. (This is necessarily true up to bounded distance.) We write $P_\Lambda = \prod_i I_i$ for the direct product equipped with the product ternary relation.

Let $f_i = \omega_{a,e_i} : \Lambda \longrightarrow I_i$, that is, $f_i(x) = \mu(a, c_i, x)$. Combining these maps, we get a quasimorphism, $f : \Lambda \longrightarrow P_\Lambda$. Conversely, given $x = (x_1, \ldots, x_n) \in P_\Lambda$, let $g(x) = x_1 \vee x_2 \vee \cdots \vee x_n \in [a, b] = H(Q)$. (Recall from Section 8 that $[a, b]$ is a distributive lattice up to bounded distance.) This gives a map $g : P_\Lambda \longrightarrow H(Q)$ which is also a quasimorphism. In fact, $g \circ f \sim_* \omega_{a,b} : \Lambda \longrightarrow H(Q)$.

In particular, $g \circ f|H(Q)$ is the identity up to bounded distance. Both maps are coarsely lipschitz, and so g is a quasi-isometric embedding. The above observations all follow by coarsifying the corresponding statements in a median algebra as discussed in Section 2.

Now if Q is straight, by Lemma 8.2, there is a median quasimorphism, $\beta_i : J_i \longrightarrow I_i$, where $J_i = [0, l_i] \subseteq \mathbb{R}$. Note that $\beta_i(J_i)$ is cobounded in I_i. Let $P = \prod_i J_i \subseteq \mathbb{R}^n$ with the l^1 metric. Thus P is a compact panel. Combining the maps β_i, we get quasimorphism, $\beta : P \longrightarrow P_\Lambda \subseteq \Lambda$. Postcomposing with g above, we get a map, $\psi = g \circ \beta : P \longrightarrow H(Q)$. All the above maps are coarsely lipschitz quasimorphisms. In summary, we have shown:

Lemma 9.3. *Let Λ be a coarse median space. Suppose that $Q \subseteq \Lambda$ is an r-straight h-quasicube of dimension n. Let $P = \prod_{i=1}^n [0, w(W_i)] \subseteq \mathbb{R}^n$, where $\mathcal{W}(Q) = \{W_1, \ldots, W_n\}$. Then there is a quasimorphism $\psi : P \longrightarrow H(Q)$, sending the corners of P to the corresponding elements of Q, with $\psi(P)$ cobounded in $H(Q)$, and with ψ a quasi-isometric embedding to Λ. All the constants of the conclusion depend only on h, r, n, and the parameters of Λ.*

In fact, our argument shows that if $\phi : P \longrightarrow H(Q)$ is any coarsely lipschitz quasimporphism sending the corners of P to the corresponding elements of Q, then $\phi(P) \sim_* H(Q)$.

Note that by Lemma 9.1, if Λ has rank ν, then this result applies to any sufficiently separated quasicube of dimension ν in Λ. In summary, we have shown that the coarse hull of a sufficiently separated ν-cube in Λ is a coarse panel as defined in Section 2.

Given $a, b, a', b' \in \Lambda$, we say that a, b is k-***parallel*** to a', b' if they form opposite faces of a k-quasicube of dimension 2. In this case, the gate maps, $w_{a',b'}|[a, b] \longrightarrow [a', b']$ and $w_{a,b}|[a', b'] \longrightarrow [a, b]$ are quasi-inverse quasimorphisms. From this one sees easily that if $[a, b]$ is r-straight, then $[a', b']$ is r'-straight, where r' depends only on r, k and the parameters of Λ. Also, if a'', b'' is k-parallel to a', b', then it is k'-parallel to a, b, where k' depends only on k and the parameters of Λ. Moreover the gate maps between them commute up to bounded distance.

Suppose that $\Pi \subseteq \Lambda$ is an h-quasisubalgebra. If $W \in \mathcal{W}(\Pi)$ and c, d and c', d' are 1-cells of Π crossing W, then if c, d is s-straight, then c', d' is s'-straight. We will say that W is s-***straight*** if all 1-cells crossing it are s-straight.

Definition. We say that Π is s-***straight*** if all walls are s-straight.

If Λ has rank ν, and $\mathrm{rank}(\Pi) = \nu$ and $\mathrm{rank}(W) = \nu - 1$, then by Lemma 9.1 and the above discussion, we see that if $w(W, \Lambda)$ is sufficiently large, then W is straight. Recalling the definition of "fat" from Section 3, we see that a sufficiently separated fat quasisubalgebra is straight. We have the following variant.

Lemma 9.4. *In the conclusion of Lemma 7.2, we can suppose in addition that Π_2 is $s(n)$-straight where $s(n)$ depends only on n and the parameters of Λ.*

Proof. We suppose that $t(n)$ sufficiently large in relation to the constant $t_0(n)$ of Lemma 9.1 as described below. Let $W \in \mathcal{W}(\Pi_2)$. By Lemma 2.6, and subsequent discussion, we can find a ν-cell $Q \subseteq \Pi$, with $\mathcal{W}(Q) = \{W'_1, \ldots, W'_{\nu-1}, W\}$, where W'_i corresponds to a wall $W_i \in \mathcal{W}(\Pi_2)$. By choosing $t(n)$ sufficiently large, we can assume that $w(W'_i, \Lambda) \geq ht_0(n)$ for each $i = 1, \ldots, \nu - 1$.

If $w(W, \Lambda) \leq t_0(n)$, then it is certainly uniformly straight. Otherwise, we apply Lemma 9.1. □

Suppose that $\Pi \subseteq \Lambda$ is an h-quasisubalgebra. As described in Section 3 we can put a median metric on $\Delta(\Pi)$ such that the width of a wall $W \in \mathcal{W}(\Pi)$ is equal to $w(W, \Lambda)$. With this structure, $\Delta(\Pi)$ is a compact CCAT(0) panel complex. We denote it by $\Delta(\Pi, \Lambda)$.

Let $\Upsilon(\Pi, \Lambda)$ be the union of the sets $H(Q)$ as Q ranges over all cells of Π (cf. the construction in Section 3).

Lemma 9.5. *Suppose that $\Pi \subseteq \Lambda$ is an r-straight h-quasisubalgebra. Then there is a strong quasimorphism, $\phi : \Delta(\Pi, \Lambda) \longrightarrow \Lambda$, with $\phi(\Delta(\Pi, \Lambda)) \subseteq \Upsilon(\Pi, \Lambda)$ cobounded in $\Upsilon(\Pi, \Lambda)$. The constants of the conclusion only depend on r, h, $|\Pi|$ and the parameters of Λ.*

Proof. Given any $W \in \mathcal{W}(\Pi)$ choose any 1-cell, c, d, of Π crossing W. Let $\beta : [0, w(W, \Lambda)] \longrightarrow [c, d]_\Lambda$ be the quasimorphism given by Lemma 8.1. We now postcompose by the gate maps to the other 1-cells crossing W. We do this for each wall W. This gives a quasimorphism on each 1-cell of $\Delta(\Pi, \Lambda)$. As with Lemma 9.3, we can now extend this to a quasimorphism from each cell $\Delta(Q)$ of $\Delta(\Pi, \Lambda)$ to $H(Q) \subseteq \Lambda$. Assembling these gives us the required map from $\Delta(\Pi, \Lambda)$ to $\Upsilon(\Pi, \Lambda)$.

Now ϕ is coarsely lipschitz, since its restriction to each cell is. We want to show that it is a quasi-isometry.

One way to see this is to embed $\Delta(\Pi, \Lambda)$ into the l^1 product $\Xi = \prod_{W \in \mathcal{W}(\Pi)} [0, w(W, \Lambda)]$. By the construction of $\Delta(\Pi, \Lambda)$, the inclusion is isometric. We can define a coarsely lipschitz map $\Lambda \longrightarrow \Xi$, where the W-coordinate is obtained by postcomposing the gate map $[x \mapsto \mu(c, d, x)]$ with a quasi-isometry from $[c, d]_\Lambda$ to $[0, w(W, \Lambda)]$, where $\{c, d\}$ is any side of Π crossing W. Restricting to $\Upsilon(\Pi, \Lambda)$ gives us a coarsely lipschitz quasi-inverse to ϕ.

It remains to show that ϕ is a quasimorphism. In principle, one could do this by coarsifying the proof of Lemma 3.12 of the present paper, as it was laid out in Section 6 of [Bo3]. However, we will give a somewhat different argument below.

To simplify notation, we will write $x' = \phi x$ for $x \in \Delta$. We will use μ and μ' for the medians in Δ and Λ respectively.

Given $x, y, z \in \Delta$, let $m = \mu(x, y, z)$. We claim that $m' \sim_* \mu'(x', y', z')$. We will say that the "claim holds" for $x, y, z \in \Delta$ if $m' \sim_r \mu'(x', y', z')$, for some fixed $r \geq 0$, which will depend on the parameters of Π and Λ and on the stage of the argument. We won't bother to explicitly estimate this. Since $\phi|\Pi$ is a quasimorphism by hypothesis, the claim certainly holds for $x, y, z \in \Pi$.

We begin by supposing that x, y lie in a 1-cell, $[c, d] \subseteq \Delta$, where $c, d \in \Pi$, and that $z \in \Delta$.

If $\mu(c, d, z) \neq c, d$, then it follows directly from the construction that the claim holds for x, y, z.

Suppose instead that $z \in \Pi$. Then without loss of generality, $\mu(c, d, z) = d$. In this case, $\mu(x, y, z) = \mu(\mu(x, c, d), \mu(y, c, d), z) = \mu(x, y, \mu(c, d, z)) = \mu(x, y, d)$. On coarsifying, $\mu'(x', y', z') \sim_* \mu'(x', y', d') = (\mu(x, y, z))'$. In particular, the claim again holds for x, y, z.

Now suppose $z \in \Delta$ is arbitrary. Let $P \subseteq \Delta$ be a cell containing z. We write $P = [a, b]_\Delta$, where $a, b \in \Pi$ are opposite corners of the cell, $P \cap \Pi$, of Π. First note that if $\mu(x, y, z) \neq c, d$, then the claim holds, as already noted. If not, we can assume, without loss of generality that P maps to d under the gate map to $[c, d]$. In particular, $\mu(c, d, a) = d$ and by the earlier discussion, $\mu(x, y, a) = \mu(x, y, d)$ and $\mu'(x', y', a') \sim_* \mu'(x', y', d')$. Similarly, $\mu(x, y, b) = \mu(x, y, d)$ and $\mu'(x', y', b') \sim_* \mu'(x', y', d')$. It follows that $\mu(x, y, z) = \mu(x, y, \mu(a, b, z)) = \mu(\mu(x, y, a), \mu(x, y, b), z) = \mu(\mu(x, y, d), \mu(x, y, d), z) = \mu(x, y, d)$. Coarsifying, we similarly get $\mu'(x', y', z') \sim_* \mu'(x', y', d')$. Using the previous case, we now get $(\mu(x, y, z))' = (\mu(x, y, d))' \sim_* \mu'(x', y', d') \sim_* \mu'(x'y', z')$. This shows that the claim holds whenever x, y lie in a 1-cell of Δ, and z is any point of Δ.

Now suppose that x, y both lie in some cell P_0 of Δ, and $z \in \Delta$. Then $\mu(x, y, z) \in P_0$ and since $P_0' = \phi(P_0)$ is coarsely convex $\mu'(x', y', z')$ lies in (or a bounded distance from) P_0. To check that $(\mu(x, y, z))' \sim_* \mu'(x', y', z')$ it is therefore enough to check that $\mu(x, y, z)$ projects to $\mu'(\mu'(c', d', x'), \mu'(c', d', y'), z')$ for any 1-face, c, d of P_0. But this follows easily from the previous case. Therefore the claim holds when x, y both lie in some cell, and for any $z \in \Delta$.

Next, suppose that $x, y \in \Pi$ and $z \in \Delta$. Again, let $z \in P = [a, b]$ with $a, b \in \Pi$ as above. Let $m = \mu(x, y, z)$. So $m = \mu(x, y, \mu(a, b, z)) = \mu(\mu(x, y, a), \mu(x, y, b), z)$. Now $\mu(x, y, a)$ and $\mu(x, y, b)$ lie in the same cell of Δ, and so by the previous case, we get $m' \sim_* \mu'((\mu(x, y, a))', (\mu(x, y, b))', z')$. Since $x, y, a, b \in \Pi$, we get $\mu'((\mu(x, y, a))', (\mu(x, y, b))', z') \sim_* \mu'(\mu(x', y', a'), \mu'(x', y', b'), z') \sim_* \mu'(x', y', \mu'(a', b', z')) \sim_* \mu(x', y', (\mu(a, b, z)') \sim_* \mu'(x', y', z')$, using the fact that a, b lie in a cell of P for the penultimate step. This shows that the statement holds for $x, y \in \Pi$ and any $z \in \Delta$.

Now suppose $x \in \Pi$ and $y, z \in \Delta$. We repeat the above argument. The only difference is the step involving x, y, a, b, where this time $x, a, b \in \Pi$ and $y \in \Delta$. For this, we use the claim for the previous case. The claim now follows also for such x, y, z.

Finally, suppose $x, y, z \in \Delta$ are arbitrary. We repeat the argument a third time. On this occasion, $a, b \in \Pi$, and $x, y \in \Pi$, so again, we can use what we have already shown. We deduce that the claim holds for all $x, y, z \in \Delta$. $\square$

10. Asymptotic Cones

In this section we give some general results about asymptotic cones we will need in Section 11. Though expressed a bit differently, the idea behind many of the arguments here come from Section 3 of [BeHS3]. For simplicity of exposition, we will first prove the results for a given space, so that a-priori the resulting constants might depend on that space. We explain at the end what is needed to ensure uniformity of these constants.

We begin by recalling the general construction of an asymptotic cone [G, VaW].

Let $\mathcal{Z}$ be a countable indexing set, and let $((X_\zeta, \delta_\zeta))_{\zeta \in \mathcal{Z}}$ be a family of metric spaces indexed by $\mathcal{Z}$. A ***$\mathcal{Z}$-sequence*** is a family of points, $\underline{x} = (x_\zeta)_{\zeta \in \mathcal{Z}}$, indexed by $\mathcal{Z}$, such that $x_\zeta \in X_\zeta$ for all ζ. Let $\underline{X} = \prod_{\zeta \in \mathcal{Z}} X_\zeta$ be the set of $\mathcal{Z}$-sequences.

Fix a non-principal ultrafilter on $\mathcal{Z}$. A predicate in ζ holds ***almost always*** (or for ***almost all*** ζ) if the set of $\zeta \in \mathcal{Z}$ for which the predicate is true lies in the ultrafilter. Given $\underline{x}, \underline{y} \in \underline{X}$, write $\underline{x} \sim \underline{y}$ if $\delta_\zeta(x_\zeta, y_\zeta) \to 0$ (in the sense that, for all real $\epsilon > 0$, we have $\delta_\zeta(x_\zeta, y_\zeta) \leq \epsilon$ for almost all ζ). One can check that this is an equivalence relation on $\underline{X}$. Write $X^* = \underline{X}/\sim$. Given $\underline{x} \in \underline{X}$, we write $x_\zeta \to x \in X^*$ to mean that x is the equivalence class of $\underline{x}$ in X^*.

Suppose that $\underline{p} \in \underline{X}$. Let $\underline{X}(\underline{p})$ be the set of $\underline{x} \in \underline{X}$ such that $\delta_\zeta(x_\zeta, p_\zeta)$ is essentially bounded (in the sense that there is some $K > 0$ such that $\delta_\zeta(x_\zeta, p_\zeta) \leq K$ for almost all ζ). Clearly, if $\underline{y} \sim \underline{x} \in \underline{X}(\underline{p})$, then $\underline{y} \in \underline{X}(\underline{p})$. We can therefore write $X^\infty = X^\infty(\underline{p}) = \underline{X}(\underline{p})/\sim \subseteq X^*$. Given $x, y \in X^\infty$, choose $\underline{x}, \underline{y} \in \underline{X}(\underline{p})$ with $x_\zeta \to x$ and $y_\zeta \to y$. Then there is some $d \in \mathbb{R}$ such that $\delta_\zeta(x_\zeta, y_\zeta) \to d$ (in the sense that for all $\epsilon > 0$, $|\delta_\zeta(x_\zeta, y_\zeta) - d| \leq \epsilon$ for almost all ζ). This d is well defined independently of $\underline{x}, \underline{y}$, and we set $\rho^\infty(x, y) = d$. One can check that (X^∞, ρ^∞) is a complete metric space. If each X_ζ is a geodesic space, then so is X^∞.

Suppose now that (X, ρ) is a fixed metric space, and $(\tau_\zeta)_\zeta$ is a $\mathcal{Z}$-sequence in $[0, \infty) \subseteq \mathbb{R}$ with $\tau_\zeta \to \infty$ (in the sense that for all $K > 0$, $\tau_\zeta > K$ for almost all ζ). Let $X_\zeta = X$ and $\delta_\zeta = \rho/\tau_\zeta$. In this case, we refer to X^∞ as the ***asymptotic cone*** of X with ***observation points*** p_ζ, and ***scaling factors*** τ_ζ.

Suppose we have a $\mathcal{Z}$-sequence, $(A_\zeta)_\zeta$, of subsets of X. Let $A^\infty \subseteq X^\infty$ be the set of $a \in X^\infty$ such that for almost all ζ, there is some $a_\zeta \in A_\zeta$ such that $a_\zeta \to a$. Then A^∞ is a closed subset of X^∞. We write $A_\zeta \to \mathbf{A} \subseteq X^\infty$ to mean that $A^\infty = \mathbf{A}$. Note that if $B_\zeta \to \mathbf{B}$, then $A_\zeta \cup B_\zeta \to \mathbf{A} \cup \mathbf{B}$.

Suppose that $\phi_\zeta : X \longrightarrow Y$ is a $\mathcal{Z}$-sequence of uniformly coarsely lipschitz maps between fixed metric spaces X, Y. Given some scaling factors and some $p_\zeta \in X$, let X^∞ and Y^∞ be the asymptotic cones with observation points p_ζ and $\phi_\zeta(p_\zeta)$ respectively. We get a limiting lipschitz map, $\phi^\infty : X^\infty \longrightarrow Y^\infty$. If the ϕ_ζ are uniformly quasi-isometric embeddings, then ϕ^∞ is bilipschitz onto its range. If the ϕ_ζ are uniform quasi-isometries, then ϕ^∞ is a bilipschitz homeomorphism from X^∞ to Y^∞.

We now assume that X is a geodesic space. Let $\mathcal{A}$ be a collection of subsets of X, and let X^∞ be an asymptotic cone. (We fix once and for all the ultrafilter on $\mathcal{Z}$.) In what follows, all "sequences" are $\mathcal{Z}$-sequences, and "limits" are in the sense of ultrafilters.

Definition. We say that a subset, $\mathbf{A} \subseteq X^\infty$ is a ***strong limit*** of $\mathcal{A}$ if there are elements $A_\zeta \in \mathcal{A}$ with $A_\zeta \to \mathbf{A}$.
We say that $\mathbf{A}$ is a ***weak limit*** of $\mathcal{A}$ if given any bounded (i.e. finite-diameter) subset, $\mathbf{C} \subseteq \mathbf{A}$, there is a sequence of elements, $A_\zeta \in \mathcal{A}$, with $A_\zeta \to \mathbf{B}$, where $\mathbf{C} \subseteq \mathbf{B}$.

An equivalent way to say that $\mathbf{A}$ is a weak limit is to assert that for all $\lambda > 0$, there is a sequence $A_\zeta \in \mathcal{A}$ with $A_\zeta \cap N(p_\zeta; \lambda\tau_\zeta) \to \mathbf{D}$ where $\mathbf{A} \cap N(p;\lambda) \subseteq \mathbf{D}$ (where p_ζ are the observation points, and τ_ζ the scaling factors).

Let $\mathcal{E}, \mathcal{F}$ be two collections of subsets of X. Given any asymptotic cone, X^∞, of X we assume we have three subsets, $\mathcal{E}^\infty(X^\infty)$, $\mathcal{F}_0^\infty(X^\infty)$ and $\mathcal{F}_1^\infty(X^\infty)$, of the power set of X^∞, with $\mathcal{F}_0^\infty(X^\infty) \subseteq \mathcal{F}_1^\infty(X^\infty)$.

(In our eventual application in Section 11, we can think of $\mathcal{E}$ as a set of quasi-flats in a coarse median space, Λ, and $\mathcal{F}$ as a set of coarse panels in Λ. In the asymptotic cone, Λ^∞, $\mathcal{E}^\infty$ will be a family of bilipschitz embedded copies of $\mathbb{R}^\nu$, and $\mathcal{F}_0^\infty$ and $\mathcal{F}_1^\infty$ will be families of be CCAT(0) panel complexes. Here we make the assumption that these arise limits of $\mathcal{E}$ and $\mathcal{F}$ in an appropriate sense. We will eventually need to verify these hypotheses in our context (see Lemma 11.4). The conclusion of Lemma 10.5 tells us that elements of $\mathcal{E}$ lie close to elements of $\mathcal{F}$ over an arbitrary large distance. It is simplest just to think of a fixed coarse median space, Λ, though to achieve uniformity of parameters, one needs to consider simultaneously the family of all coarse median spaces with fixed parameters.)

We now introduce the hypotheses we will be using. Suppose:

(E1): For any asymptotic cone X^∞, the following hold.
(E1a): Every strong limit of $\mathcal{E}$ lies in $\mathcal{E}^\infty(X^\infty)$.
(E1b): Every strong limit of $\mathcal{F}$ lies in $\mathcal{F}_1^\infty(X^\infty)$.
(E1c): Every element of $\mathcal{F}_0^\infty(X^\infty)$ is a weak limit of $\mathcal{F}$.

(E2): For any asymptotic cone, X^∞, and any $\mathbf{E} \in \mathcal{E}^\infty(X^\infty)$ there is some $\mathbf{F} \in \mathcal{F}_0^\infty(X^\infty)$ with $\mathbf{E} \subseteq \mathbf{F}$.

The above are assumed to hold in any asymptotic cone of X, i.e. with any observation points or scaling factors (though we will work with a fixed ultrafilter on $\mathcal{Z}$).

Lemma 10.1. *Suppose $\mathcal{E}, \mathcal{F}$ satisfy (E1) and (E2). Then for all $\lambda > 0$, there is some $R_0 \geq 0$, such that for all $R \geq R_0$, $E \in \mathcal{E}$ and $a \in E$, there is some $F \in \mathcal{F}$ such that $E \cap N(a;\lambda R) \subseteq N(F;R)$.*

Proof. Suppose not. There is some $\lambda > 0$, and sequences $R_\zeta \to \infty$, $E_\zeta \in \mathcal{E}$, and $a_\zeta \in X$ such that for all $F \in \mathcal{F}$ there is some point of $E_\zeta \cap N(a_\zeta;\lambda R_\zeta)$ a distance more than R_ζ from F. We can suppose that $\lambda \geq 2$. Let X^∞ be the asymptotic cone with observation points a_ζ and scaling factors R_ζ. Let $E_\zeta \to \mathbf{E} \in \mathcal{E}^\infty(X^\infty)$. By (E2), there is some $\mathbf{F} \in \mathcal{F}_0^\infty(X^\infty)$ with $\mathbf{E} \subseteq \mathbf{F}$. By (E1c), there exist $F_\zeta \in \mathcal{F}$ with $F_\zeta \cap N(a_\zeta; \lambda R_\zeta) \to \mathbf{G}$ with $\mathbf{F} \cap N(a, \lambda) \subseteq \mathbf{G}$. We can find $x_\zeta \in E_\zeta \cap N(a_\zeta; \lambda R_\zeta)$ with $\rho(x_\zeta; F_\zeta) > R_\zeta$, and so $\delta_\zeta(a_\zeta; x_\zeta) \leq \lambda$ and $\delta_\zeta(x_\zeta; F_\zeta) > 1$ (where $\delta_\zeta = \rho/R_\zeta$).

Let $x_\zeta \to x \in X^\infty$. Then $x \in \mathbf{E}$ and $\rho^\infty(a, x) \leq \lambda$. Now $\mathbf{E} \subseteq \mathbf{F}$, so $x \in \mathbf{F} \cap N(a; \lambda) \subseteq \mathbf{G}$. In particular, there are $y_\zeta \in F_\zeta$ with $y_\zeta \to x$, so $\delta_\zeta(x_\zeta, F_\zeta) \leq \delta_\zeta(x_\zeta, y_\zeta) \to 0$, giving a contradiction. □

Suppose now:

(E3): Suppose X^∞ is any asymptotic cone and suppose that $\mathbf{E} \in \mathcal{E}^\infty(X^\infty)$, $\mathbf{F} \in \mathcal{F}_1^\infty(X^\infty)$, and that there is some $t \geq 0$ such that $\mathbf{E} \subseteq N(\mathbf{F}; t)$. Then $\mathbf{E} \subseteq \mathbf{F}$.

(There is no particular significance in the constant t. In fact, we could fix t once and for all to be any positive constant. It makes no difference to the statement.)

Henceforth, for a given asymptotic cone, X^∞, we will generally abbreviate $\mathcal{E}^\infty = \mathcal{E}^\infty(X^\infty)$ etc.

The idea behind the following lemma is that if $E \in \mathcal{E}$ lies close to an element $F \in \mathcal{F}$ in a certain neighbourhood of a point $a \in E$, then a itself must lie even closer to F.

Lemma 10.2. *Suppose that $\mathcal{E}, \mathcal{F}$ satisfy (E1a), (E1b) and(E3). Then there exist some $\lambda > 0$ and $k > 0$ such that if $E \in \mathcal{E}$, $F \in \mathcal{F}$ and $a \in E$, then*

$$\rho(a, F) \leq \frac{1}{2} \max(\{k\} \cup \{\rho(x, F) \mid x \in E \cap N(a; \lambda\rho(a, F))\}).$$

(For the factor of "$\frac{1}{2}$" here we could substitute our favourite number strictly between 0 and 1. The same proof will work, and the result would apply equally well.)

Proof. Suppose not. This means that there are sequences $\lambda_\zeta \to \infty$ and $k_\zeta \to \infty$, such that there exist $E_\zeta \in \mathcal{E}$, $F_\zeta \in \mathcal{F}$ and $a_\zeta \in E_\zeta$ with $\rho(a_\zeta, F_\zeta) \geq k_\zeta/2$ and with $\rho(a_\zeta, F_\zeta) \geq \frac{1}{2}\rho(x, F_\zeta)$ for all $x \in E_\zeta \cap N(a_\zeta; \lambda_\zeta \rho(a_\zeta, F_\zeta))$. Let $u_\zeta = \rho(a_\zeta, F_\zeta)$. Then $u_\zeta \to \infty$. (We can now forget k_ζ.) Thus, for each $x \in E_\zeta \cap N(a_\zeta; \lambda_\zeta u_\zeta)$ we have $\rho(x, F_\zeta) \leq 2u_\zeta$.

Let X^∞ be the asymptotic cone with observation points a_ζ and with scaling factors u_ζ. Let $E_\zeta \to \mathbf{E} \in \mathcal{E}^\infty$ and $F_\zeta \to \mathbf{F} \in \mathcal{F}_1^\infty$ and $a_\zeta \to a \in \mathbf{E}$. Suppose $x \in \mathbf{E}$. Let $x_\zeta \to x$ with $x_\zeta \in E_\zeta$. Now $\delta_\zeta(a_\zeta, x_\zeta) := \frac{1}{u_\zeta}\rho(a_\zeta, x_\zeta)$ is essentially bounded. Since $\lambda_\zeta \to \infty$, it follows that $\rho(a_\zeta, x_\zeta) \leq \lambda_\zeta u_\zeta$ for almost all ζ. As noted above, for such ζ and x_ζ, we have $\rho(x_\zeta, F_\zeta) \leq 2u_\zeta$, and so $\delta_\zeta(x_\zeta, F_\zeta) \leq 2$. We therefore see that $\rho^\infty(x, \mathbf{F}) \leq 2$. In other words, we have shown that $\mathbf{E} \subseteq N(\mathbf{F}; 2)$. By (E3), we now have $\mathbf{E} \subseteq \mathbf{F}$. But $\delta_\zeta(a_\zeta, F_\zeta) = u_\zeta/u_\zeta = 1$ for all ζ, and so $\rho^\infty(a, \mathbf{F}) = 1$. But $a \in \mathbf{E}$ giving a contradiction. □

The following applies Lemma 10.2 to give us a uniform bound on the distance of a to F, assuming again that E lies "close to" F over a certain range.

Lemma 10.3. *Suppose $\mathcal{E}, \mathcal{F}$ satisfy (E1a), (E1b) and (E3). There are some $\kappa > 0$ and $k \geq 0$ such that if $E \in \mathcal{E}$, $F \in \mathcal{F}$, $a \in E$ and $R \geq 0$ satisfy $E \cap N(a; \kappa R) \subseteq N(F; R)$, then we have $\rho(a, F) \leq k$.*

Proof. Let k and λ be as given by Lemma 10.2, and set $\kappa = 2\lambda$. Let $n = \lfloor \log_2(R/k) \rfloor$, so that $2^{n-1}k \leq R \leq 2^n k$. Given $i = 0, 1, \ldots, n$, let $R_i = R/2^i$, so that $R_0 = R$, $R_n \leq k$ and $R_i \geq k$ for all $i < n$. Let $\theta_i = \frac{1}{2^{n-1}} + \frac{1}{2^{n-2}} + \cdots + \frac{1}{2^i}$. Thus, $\theta_{n-1} \leq \theta_{n-2} \leq \cdots \leq \theta_0 \leq 2$, and $\theta_i - \theta_{i+1} = \frac{1}{2^i}$. Let $r_i = \lambda \theta_i R$, so $r_i - r_{i+1} = \lambda(\theta_i - \theta_{i+1})R = \lambda R/2^i = \lambda R_i$. Write $N_i = N(a; r_i)$, so that $N_0 = N(a; r_0) \subseteq N(a; 2\lambda R) = N(a; \kappa R)$. Also $N_i = N(N_{i+1}; \lambda R_i)$, since $r_i = r_{i+1} + \lambda R_i$.

Since $N_0 \subseteq N(a; \kappa R)$, our hypothesis on E tells us that $E \cap N_0 \subseteq N(F; R) = N(F; R_0)$. We claim inductively that for all $i = 0, \ldots, n$, we have $E \cap N_i \subseteq N(F; R_i)$. We have noted that this holds for $n = 0$. Assume that it holds for a given $i < n$. Note that $R_{i+1} = R_i/2$. We want to show that $E \cap N_{i+1} \subseteq N(F; R_i/2)$. Let $b \in E \cap N_{i+1}$. By Lemma 10.2, we have $\rho(b, F) \leq \frac{1}{2}(\max(\{k\} \cup \{\rho(x, F) \mid x \in E \cap N(b; \lambda \rho(b, F))\})$. Now $b \in E \cap N_{i+1} \subseteq E \cap N_i$, so by the inductive hypothesis, $\rho(b, F) \leq R_i$. Suppose $x \in E \cap N(b; \lambda \rho(b, F)) \subseteq N(b; \lambda R_i) \subseteq N(N_{i+1}; \lambda R_i) = N_i$. Then by the inductive hypothesis, $\rho(x, F) \leq R_i$. In other words, $\max\{\rho(x, F) \mid x \in E \cap N(b; \lambda \rho(b, F))\} \leq R_i$. Therefore $\rho(b, F) \leq \frac{1}{2} \max\{k, R_i\} = R_i/2$ since $k \leq R_i$, since $i < n$. In other words, $b \in N(F; R_i/2) = N(F; R_{i+1})$. We have shown that $E \cap N_{i+1} \subseteq N(F; R_{i+1})$ proving the inductive statement.

In particular, it follows that $E \cap N_n \subseteq N(F; R_n) \subseteq N(F; k)$ since $R_n \leq k$. But $a \in E \cap N_n$, and so $\rho(a, F) \leq k$ as required. □

Corollary 10.4. *Suppose $\mathcal{E}, \mathcal{F}$ satisfy (E1a), (E1b) and (E3). There exist $\kappa_0, k \geq 0$ such that if $E \in \mathcal{E}$, $F \in \mathcal{F}$, $a \in E$ and $R \geq 0$ satisfy $E \cap N(a; \kappa_0 R) \subseteq N(F; R)$, then $E \cap N(a; R) \subseteq N(F; k)$.*

Proof. Let κ be the constant of Lemma 10.3 and set $\kappa_0 = \kappa + 1$. If $b \in E \cap N(a; R)$, then $E \cap N(b; \kappa R) \subseteq N(a; \kappa_0 R) \subseteq N(F; R)$, so by Lemma 10.3, $\rho(b, F) \leq R$ as required. □

The following is the main result we are aiming for: any element of $E \in \mathcal{E}$ lies uniformly close to some element of $\mathcal{F}$ over an arbitrarily large range.

Lemma 10.5. *Suppose $\mathcal{E}, \mathcal{F}$ satisfy (E1), (E2) and (E3). Then there is some $k \geq 0$ such that for all $E \in \mathcal{E}$, $a \in E$ and $R \geq 0$, there is some $F \in \mathcal{F}$ with $E \cap N(a; R) \subseteq N(F; k)$.*

Proof. Let κ_0 be the constant given by Corollary 10.4. In Lemma 10.1, set $\lambda = \kappa_0$ and R_0 be the constant given by its conclusion. Let $E \in \mathcal{E}$ and $a \in E$. We can suppose that $R \geq R_0$. By Lemma 10.1, there is some $F \in \mathcal{F}$ with $E \cap N(a; \kappa_0 R) \subseteq N(F; R)$. By Corollary 10.4 we have $E \cap N(a; R) \subseteq N(F; k)$ as required. □

In the above results, the constants of the conclusion will depend on the particular space X, and the subsets $\mathcal{F}$. This would be sufficient if, for example, we were dealing with just one coarse median space. However, we can obtain stronger statements to the effect that the constants should depend only on the parameters of the hypotheses. For this we need to consider families of spaces (such as coarse median spaces of fixed parameters) simultaneously. This can be justified by the following general principle.

Suppose that $\mathcal{S}$ is a family of spaces, and that for each X in $\mathcal{S}$, we have families of subspaces, $\mathcal{E}(X)$ and $\mathcal{F}(X)$ of X. Suppose that (X_ζ, ρ_ζ) is a $\mathcal{Z}$-sequence of spaces in $\mathcal{S}$, that $p_\zeta \in X_\zeta$ is a $\mathcal{Z}$-sequence of observation points, and that τ_ζ is a $\mathcal{Z}$-sequence of scaling factors. We let $\delta_\zeta = \rho_\zeta/\tau_\zeta$ be the rescaled metric on X_ζ and pass to the asymptotic cone, X^∞, as defined at the beginning. (For example, in the proof of Lemma 10.1, if the conclusion fails with uniform constants, we would have a sequence $(X_\zeta)_\zeta$, of spaces and $E_\zeta \in \mathcal{E}(X_\zeta)$ with $N(a_\zeta; \lambda R_\zeta)$ interpreted in X_ζ etc.) Interpreting hypotheses (E1)–(E3) above to apply to all such asymptotic cones, the results go through as stated.

(We should note that the "family" $\mathcal{S}$ does not need to be a set. It can be determined by a predicate: for example, which asserts that any X in $\mathcal{S}$ is a coarse median space with fixed parameters. A $\mathcal{Z}$-sequence with values in $\mathcal{S}$ will still be a set by the Axiom of Replacement. In any any plausible application however, the elements of $\mathcal{S}$ would have cardinality at most some fixed cardinal, in which case, we could indeed take $\mathcal{S}$ to be a set.)

In fact, a simple trick to formally deal with this in terms of what we have already done (at least when $\mathcal{S}$ is a set) would be to let Y be the disjoint union of the spaces $X \in \mathcal{S}$. We can then set $\mathcal{E} = \bigsqcup_{X\in\mathcal{X}} \mathcal{E}(X)$ and $\mathcal{F} = \bigsqcup_{X\in\mathcal{X}} \mathcal{F}(X)$. We set $\rho(x, y) = \infty$ if x, y lie in different components. Note that Y is "geodesic" in the sense that if $\rho(x, y) < \infty$, then x, y are connected by a geodesic. This is sufficient for the above statements to go through.

In our application, we will take $\mathcal{S}$ to be the family of all coarse median spaces with fixed parameters, $\mathcal{E}(X)$ to be the space of all bounded-parameter quasiflats in $X \in \mathcal{S}$, and $\mathcal{F}(X)$ will be the class of all coarsely embedded cube complexes of fixed parameters. This will be explained in the next section: see Lemmas 11.4 and 11.5.

11. Proof of Theorems 1.1 and 1.3

We give a proof of Theorem 1.1. To this end, we describe some properties of an asymptotic cone of a coarse median space. In particular such is bilipschitz equivalent to a complete connected median metric space. We show that a limit of quasiflats lies inside CCAT(0) panel complex (Lemma 11.1). Conversely such a complex is a weak limit of coarse panels (Lemma 11.3). This enables us to verify the hypotheses of Lemma 10.5 (Lemma 11.4). As consequence, any quasiflat lies a bounded distance from a coarse panel over an arbitrarily large distance (Lemma

11.5). Taking larger and larger distances, using a diagonal sequence argument, and applying Lemma 6.1, we then arrive at our main result.

Let (Λ, ρ, μ) be a coarse median space of rank ν. Let τ_ζ be a $\mathcal{Z}$-sequence of scaling factors and let $p_\zeta \in \Lambda$ be a $\mathcal{Z}$-sequence of observation points. Let $\delta_\zeta = \rho/\tau_\zeta$, and let $(\Lambda^\infty, \rho^\infty)$ be the asymptotic cone. This is a geodesic metric space. We also get a limiting map, $\mu^\infty : (\Lambda^\infty)^3 \longrightarrow \Lambda^\infty$. With this operation, $(\Lambda^\infty, \mu^\infty)$ is a topological median algebra of rank ν. The metric ρ^∞ is uniformly bilipschitz equivalent to a median metric, here denoted ρ_0, which induces the same median, μ^∞ (see [Bo2, Bo5, Z2]).

Various observations follow easily from the general principle that additive constants disappear in the limit. For example, if a_ζ, b_ζ are $\mathcal{Z}$-sequences in Λ limiting on $a, b \in \Lambda^\infty$, then $[a_\zeta, b_\zeta] \to [a, b]$. Moreover, if the intervals $[a_\zeta, b_\zeta]$ are uniformly (coarsely) straight, then $[a, b]$ is straight.

Suppose that $(\Pi_\zeta, \delta_\zeta)$ is a $\mathcal{Z}$-sequence of finite median metric spaces. Suppose that $|\Pi_\zeta|$ is (essentially) bounded. Then for almost all ζ, there is a median isomorphism, $\hat{\Pi} \longrightarrow \Pi_\zeta$, from a fixed finite median algebra $\hat{\Pi}$. If we fix observation points, $p_\zeta \in \Pi_\zeta$, then the ultralimit, Π^∞, is also a finite median algebra.

If $\operatorname{diam}(\Pi_\zeta)$ is bounded, then the isomorphisms, $\hat{\Pi} \longrightarrow \Pi_\zeta$ limit on an epimorphism, $\hat{\Pi} \longrightarrow \Pi^\infty$. This can be described as follows. We can identify $\mathcal{W}(\hat{\Pi})$ with each $\mathcal{W}(\Pi_\zeta)$ via these isomorphisms. If $W \in \mathcal{W}(\hat{\Pi})$, then the widths, $w_\zeta(W)$, of W in Π_ζ converge to some $w^\infty(W) \in [0, \infty)$. Let $\mathcal{W}_0(\hat{\Pi}) = (w^\infty)^{-1}(0)$. Then Π^∞ is the result of collapsing the walls in $\mathcal{W}_0(\hat{\Pi})$. In this way, we can identify $\mathcal{W}(\Pi^\infty)$ with $\mathcal{W}(\hat{\Pi}) \setminus \mathcal{W}_0(\hat{\Pi})$. The width of $W \in \mathcal{W}(\Pi^\infty)$ is then just $w^\infty(W)$. It is also easily checked that $\Delta(\Pi_\zeta)$ converges to $\Delta(\Pi^\infty)$.

In general, Π^∞ will be a convex subset of the quotient of $\hat{\Pi}$. We write $\Delta_\zeta = \Delta(\Pi_\zeta)$ and let $\Delta_\zeta \to \Delta^\infty$. We claim that Δ^∞ is a CCAT(0) panel complex. In fact, if $B \subseteq \Delta^\infty$ is any bounded subset, we can find compact convex subsets, $B_\zeta \subseteq \Delta_\zeta$, with $B_\zeta \to B^\infty \supseteq B$. Now B_ζ has the form $B_\zeta = \Delta(\Pi'_\zeta)$ for a finite subalgebra, $\Pi'_\zeta \subseteq \Delta_\zeta$, with $|\Pi'_\zeta|$ bounded. Therefore $B^\infty = \Delta((\Pi')^\infty)$, where $\Pi'_\zeta \to (\Pi')^\infty$. We see that Δ^∞ has an exhaustion by convex subsets which are compact CCAT(0) panel complexes with a bounded number of cells. From this, it is easily seen that Δ^∞ is a CCAT(0) panel complex. (Alternatively, one can invoke Lemma 4.2 since Δ^∞ is uniformly cubulated — as a subset of itself.)

Returning to the coarse median space, (Λ, ρ, μ), let $(\Lambda^\infty, \rho^\infty, \mu^\infty)$ be the asymptotic cone, and let ρ_0 be a uniformly bilipschitz equivalent median metric on Λ^∞.

Let $(\Pi_\zeta)_\zeta$ be a $\mathcal{Z}$-sequence of finite median algebras, $\Pi_\zeta \subseteq \Lambda$, with $|\Pi_\zeta|$ bounded, whose inclusions into Λ are uniform quasimorphisms and which are uniformly straight. Let $\Upsilon_\zeta = \Upsilon(\Pi_\zeta) \subseteq \Lambda$ be the subset defined in Section 9. Let $\Upsilon_\zeta \longrightarrow \Upsilon^\infty \subseteq \Lambda^\infty$. Note that, since the Υ_ζ are uniform quasisubalgebras, Υ^∞ is a subalgebra of Λ^∞. In fact:

Lemma 11.1. Υ^∞ *is CCAT(0) a panel complex.*

(This allows for the possibility that $\Upsilon^\infty = \emptyset$.)

Proof. Let $\Delta_\zeta = \Delta(\Pi_\zeta, \Lambda)$ as defined in Section 9. By Lemma 9.5, there is a uniformly strong quasimorphism, $\phi_\zeta : \Delta_\zeta \longrightarrow \Lambda$, with $\phi_\zeta(\Delta_\zeta) \sim_* \Upsilon_\zeta$. (In other words, ϕ_ζ is a quasimorphism and a quasi-isometric embedding with the Hausdorff distance from $\phi_\zeta(\Delta_\zeta)$ to Υ_ζ bounded — all constants being independent of ζ.) We get a limiting map $\phi^\infty : \Delta^\infty \longrightarrow \Lambda^\infty$ which is bilipschitz onto its range and a median homomorphism. By the above discussion, we know that Δ^∞ is a panel complex. It follows that Υ^∞ is a panel complex in the induced metric, ρ_0. ☐

Note in particular, that if $\Pi_\zeta \subseteq N(p; t\tau_\zeta)$ for some fixed $t \geq 0$ (i.e. the Π_ζ are uniformly bounded in the rescaled metric, $\Delta_\zeta = \rho/\tau_\zeta$), then $\Pi_\zeta \to \Pi^\infty \subseteq \Lambda^\infty$, where Π^∞ is a finite subalgebra of Λ^∞. Moreover, $\Upsilon^\infty = \Upsilon(\Pi^\infty, \Lambda^\infty)$ is naturally isomorphic to Δ^∞.

We will also require a certain converse to these statements. We begin with:

Lemma 11.2. *Given $n \in \mathbb{N}$, there is some $h(n) \geq 0$ and $s(n) \geq 0$ with the following property. Suppose that $\Pi \subseteq \Lambda^\infty$ is a subalgebra with $|\Pi| \leq n < \infty$. Then there are median algebras, $\Pi_\zeta \subseteq \Lambda$, whose inclusions into Λ are $h(n)$-quasimorphisms, with $|\Pi_\zeta| \leq 2^{2^n}$, and with $\Pi_\zeta \to \Pi$. Moreover, if Π is fat, we can take all the Π_ζ to be fat and $s(n)$-straight in Λ. The functions h and s depend only on the parameters of Λ.*

Proof. Let κ be the separation constant of Π. Let $\Pi_\zeta^0 \subseteq \Lambda$ be subsets with a natural bijection to Π, such that $\Pi_\zeta^0 \to \Pi$. Let μ_ζ be the median induced on Π_ζ^0 via the median μ_Π on Π. Now $\mu_\zeta \to \mu_\Pi$, and so the inclusion of Π_ζ^0 into Λ is an h_ζ-quasimorphism, where $h_\zeta/\tau_\zeta \to 0$. Also Π_ζ^0 is t_ζ-separated, where $t_\zeta/\tau_\zeta \to \kappa > 0$. Let $t(n)$ be the constant given by Lemma 7.2. Now $t_\zeta > h_\zeta t(n)$ for almost all ζ. Therefore, by Lemma 7.2 (with $\Pi_1 = \Pi_\zeta^0$ and $\kappa = h_\zeta$) there is an $h(n)$-quasisubalgebra $\Pi_\zeta \subseteq \Lambda$, with $|\Pi_\zeta| \leq 2^{2^n}$, together with an epimorphism $\omega_\zeta : \Pi_\zeta \longrightarrow \Pi_\zeta^0$, such that $\rho(x, \omega_\zeta x) \leq h_\zeta r(n)$ for all $x \in \Pi_\zeta$ (namely $\Pi_\zeta = \Pi_2$ and $\omega_\zeta = \omega$ as given by Lemma 7.2). Now $h_\zeta r(n)/\tau_\zeta \to 0$. Also Π_ζ is almost always isomorphic to a fixed median algebra $\hat{\Pi}$, and so ω_ζ converges on an epimorphism $\omega^\infty : \hat{\Pi} \longrightarrow \Pi$. It follows that $\Pi_\zeta \to \Pi$, respecting the median operations.

Applying Lemma 7.3, we see that if Π is fat, we can take all of the Π_ζ to be fat.

Moreover, by the addendum to Lemma 7.2 given by Lemma 9.4, we also see that the Π_ζ can be assumed uniformly straight. ☐

Let $\Pi \subseteq \Lambda^\infty$ be a fat subalgebra, and let Π_ζ be as given by Lemma 11.2. If Π is fat, then so is Π_ζ. It follows that Π_ζ is uniformly straight in Λ, and we can define $\Upsilon_\zeta = \Upsilon(\Pi_\zeta, \Lambda) \subseteq \Lambda$ as above. Since $\Pi_\zeta \to \Pi$, we see that $\Upsilon_\zeta \to \Upsilon(\Pi, \Lambda^\infty) \subseteq \Lambda^\infty$.

Lemma 11.3. *Given $n \in \mathbb{N}$, there are constants, $h(n), s(n) \geq 0$, and $p(n) \in \mathbb{N}$, with the following property. Let $\Upsilon \subseteq \Lambda^\infty$ be a fat CCAT(0) panel complex with at* *most n panels. Then Υ is a weak limit of sets of the form $\Upsilon(\Pi, \Lambda)$, where $\Pi \subseteq \Lambda$ is a fat $s(n)$-straight $h(n)$-quasisubalgebra with $|\Pi| \leq p(n)$.*

Proof. Let $B \subseteq \Upsilon$ be any compact subset. Then there is some finite subalgebra, $\Pi_0 \subseteq \Upsilon$, with $B \subseteq \Upsilon(\Pi_0, \Lambda^\infty) \subseteq \Upsilon$. Moreover, we can take Π_0 to be fat. (In fact, we can take Π_0 so that $\Delta(\Pi_0)$ is isomorphic to the compact panel complex $\hat{\Upsilon}$ as defined in Section 3.)

Let $(\Pi_\zeta)_\zeta$ be a sequence of finite median algebras, $\Pi_\zeta \subseteq \Lambda$, as given by Lemma 11.2 (with $\Pi = \Pi_0$). From the observation above, it follows that $\Upsilon(\Pi_\zeta, \Lambda) \to \Upsilon(\Pi_0, \Lambda^\infty)$. Since $B \subseteq \Upsilon(\Pi_0, \Lambda^\infty)$, the statement follows. □

We are now in a position to apply the results of Section 10.

To this end, we fix constants $\lambda_0, h_0, s_0 \geq 0$ and $K_1, K_2 \in \mathbb{N}$, depending only on ν and the parameters of Λ and of the quasiflat, as determined below. (They will be chosen in the order $\lambda_0, K_2, K_1, h_0, s_0$.)

Let $\mathcal{E}$ be the set of all quasiflats in Λ with the given quasi-isometry constants. Let $\mathcal{F}$ be the set of subsets of Λ of the form $\Upsilon(\Pi, \Lambda)$, where $\Pi \subseteq \Lambda$ is a fat h_0-quasisubalgebra, s_0-straight in Λ, with $|\Pi| \leq K_1$.

Let Λ^∞ be any asymptotic cone of Λ. Let $\mathcal{E}^\infty$ be the set of images of $\mathbb{R}^\nu$ in Λ^∞ under λ_0-bilipschitz maps. Let $\mathcal{F}_1^\infty$ be the set of all CCAT(0) panel complexes in Λ^∞. Let $\mathcal{F}_0^\infty \subseteq \mathcal{F}_1^\infty$ consist of those CCAT(0) panel complexes which are fat and which have at most K_2 panels.

Lemma 11.4. *For suitable choice of constants, $\lambda_0, h_0, s_0, K_1, K_2$, the collections $\mathcal{E}, \mathcal{F}, \mathcal{E}^\infty, \mathcal{F}_0^\infty, \mathcal{F}_1^\infty$ defined above satisfy properties (E1)–(E3) of Section 10.*

Proof.
(E1a): If $f_\zeta : \mathbb{R}^\nu \longrightarrow \Lambda$ is a sequence of uniform quasi-isometric embeddings, then the limiting map, $f^\infty : \mathbb{R}^\nu \longrightarrow \Lambda^\infty$, is a uniformly bilipschitz to its range (with respect to the metric, ρ^∞, hence also with respect to ρ_0). Choosing λ_0 accordingly gives the condition on $\mathcal{E}^\infty$.
(E1b): Lemma 11.1 tells us that any strong limit of elements of $\mathcal{F}$ lies in $\mathcal{F}_1^\infty$.
(E2): By Lemma 5.2, any element of $\mathcal{E}^\infty$ is contained in a CCAT(0) panel complex, where the number of cells is bounded in terms of the bilipschitz constant λ_0 chosen above. We choose K_2 in the definition of $\mathcal{F}_0^\infty$ accordingly.
(E1c): Lemma 11.3 tells us that any element of $\mathcal{F}_0^\infty$ is a weak limit of elements of $\mathcal{F}$, where K_1 and h_0, s_0 are chosen according to the conclusion of Lemma 11.3 given the bound K_2.
(E3): This follows immediately from Lemma 5.3. □

Lemma 11.5. *Suppose that $f : \mathbb{R}^\nu \longrightarrow \Lambda$ is a quasi-isometric embedding. There is some $n \in \mathbb{N}$, $s, h, k \geq 0$, such that for all $a \in \Lambda$ and all $r \geq 0$, there is a finite fat s-straight median algebra, $\Pi \subseteq \Lambda$, with $|\Pi| \leq n$, whose inclusion into Λ is an h-quasimorphism, and with $f(\mathbb{R}^\nu) \cap N(a; r) \subseteq N(\Upsilon(\Pi, \Lambda); k)$. Here n, h, s, k depend only on the parameters of Λ and f. (In particular, they are independent of r.)*

Proof. Let $\mathcal{E}, \mathcal{F}, \mathcal{E}^\infty, \mathcal{F}_0^\infty, \mathcal{F}_1^\infty$, be as described above. By Lemma 11.4, these satisfy (E1)–(E3). Therefore, by Lemma 10.5, there is some $F \in \mathcal{F}$ with $f(\mathbb{R}^\nu) \cap N(a;R) \subseteq N(F;k)$ for some $k \geq 0$. By the definition of $\mathcal{F}$, F has the form $F = \Upsilon(\Pi, \Lambda)$, where Π is an s-straight h-quasisubalgebra with $|\Pi| \leq n$. The uniformity of the constants, n, h, k, follows by considering simultaneously all spaces and families of subsets of fixed parameters, as described by the discussion following Lemma 10.5. $\square$

Let $f : \mathbb{R}^\nu \longrightarrow \Lambda$ be a quasi-isometric embedding. Given $i \in \mathbb{N}$, Lemma 11.5 gives us a quasisubalgebra, Π_i, with $f(\mathbb{R}^\nu) \cap N(a;i) \subseteq N(\Upsilon_i;k)$, where $\Upsilon_i = \Upsilon(\Pi_i, \Lambda)$, and with $|\Pi_i|$ bounded. By Lemma 9.5, we have uniform strong quasimorphisms, $\phi_i : \Delta_i \longrightarrow \Lambda$, with $\Upsilon_i \sim_* \phi_i(\Delta_i)$, where $\Delta_i = \Delta(\Pi_i, \Lambda)$. In particular, $f(\mathbb{R}^\nu) \cap N(a;i) \subseteq N(\phi_i\Delta_i; k_0)$ for some fixed constant, k_0. (Note also that the strong quasimorphism constants of ϕ_i depend only on the parameters of Λ. The constant k_0 depends also on the quasi-isometry parameters of f.)

We can therefore find an exhaustion, $(N_i)_i$, of $\mathbb{R}^\nu$ by compact convex subsets, and maps $\theta_i : N_i \longrightarrow \Delta_i$ such that $\phi_i \circ \theta_i \sim_* f|N_i$ for all i. Note that these are all uniform quasi-isometric embeddings. Note that all the Δ_i have a bounded number of cells.

Up to bounded distance, and after passing to a subsequence, we can assume that $(\Delta_i)_i$ is an increasing sequence of convex subsets exhausting a fixed CCAT(0) panel complex, Ψ, and moreover if $i < j$, then $\theta_j|N_i \sim_* \theta_i$. (For example, up to bounded distance, we can assume that each Δ_i is, after subdivision, a finite CAT(0) cube complex with all side-lengths equal to some fixed positive constant. By a diagonal sequence argument, we can pass to a subsequence so that Δ_i is identified with a subset of Δ_j whenever $i \leq j$, and with the maps θ_i agreeing up to bounded distance, whenever they are defined. We now take Ψ to be the union of the Δ_i.) Again up to bounded distance, there is a map $\theta : \mathbb{R}^\nu \longrightarrow \Psi$ such that $\theta_i \sim_* \theta|N_i$ for all i.

Now θ is a quasi-isometric embedding (since all the θ_i are). Therefore, by Lemma 6.1, there is a subcomplex, $\Omega \subseteq \Psi$, consisting of a union of ν-panels, with $\theta(\mathbb{R}^\nu) \sim_* \Omega$. Up to bounded distance, we can assume that $\theta(\mathbb{R}^\nu) \subseteq \Omega$, and so $\theta : \mathbb{R}^\nu \longrightarrow \Omega$ is a quasi-isometry. Let $\psi : \Omega \longrightarrow \mathbb{R}^\nu$ be a quasi-inverse of this map, and let $\phi = f \circ \psi : \Omega \longrightarrow \Lambda$. This is a quasi-isometric embedding, and by construction, $\phi(\Omega) \sim_* f(\mathbb{R}^\nu)$.

It remains to check that ϕ is a uniform quasimorphsim on each ν-panel, P, of Ω. But on any compact subset of P, ϕ agrees up to bounded distance with ϕ_i for some (indeed all sufficiently large) i (since $\phi_i \sim_* \phi_i \circ \theta \circ \psi \sim_* f \circ \psi = \phi$ by construction). The claim follows since each ϕ_i is a uniform quasimorphism.

In summary, we have shown:

Lemma 11.6. *Let $f : \mathbb{R}^\nu \longrightarrow \Lambda$ be a quasi-isometric embedding. There is a panel complex, Ω and quasi-isometric embedding $\phi : \Omega \longrightarrow \Lambda$, which is a strong quasimorphism restricted to each panel, and such that* $\mathrm{hd}(f(\mathbb{R}^\nu), \phi(\Omega))$ *is bounded.*

The constants of the conclusion (namely the number of cells of Ω, the strong quasimorphism constants of ϕ, and the Hausdorff distance bound) depend only on ν and the parameters of Λ and f.

This is the first part of Theorem 1.1.

For the second part, let $P_1, \ldots, P_q$ be those panels which of side-lengths at least L. By Lemma 6.2, there is some (non-uniform) $r \geq 0$, such that $\Omega \subseteq N(\bigcup_{i=1}^q P_i; r)$. It follows that the Hausdorff distance from $f(\mathbb{R}^\nu)$ to $\phi(\Omega)$ hence also to $\bigcup_{i=1}^q \phi_i(P_i)$ is finite.

For the final part of Theorem 1.1, suppose that Λ is proper (i.e. complete and locally compact). We use a slightly different argument. By a diagonal sequence argument, we can assume that on any given compact set the maps $\phi_i : \Delta_i \longrightarrow \Lambda$ agree up to bounded distance, where the bound depends only on the parameters of Λ. Therefore, after passing to a subsequence, these converge up to bounded distance on a map $\phi : \Omega \longrightarrow \Lambda$. By construction, $\phi(\Lambda)$ is again a bounded Hausdorff distance from $f(\mathbb{R}^\nu)$. This latter bound will depend on the parameters of f. However, by this argument, the quasi-isometry constant of ϕ, as well as the coarse quasimorphism constants of the panels, depend only on the parameters of Λ.

This concludes the proof of Theorem 1.1.

Theorem 1.2 now follows using Corollary 6.3 and Lemma 6.4.

12. Quasiflats in median metric spaces

In this section, we explain how the conclusion of Theorem 1.1 can be strengthened in the case of a connected median metric space. In particular, this will prove Theorem 1.3. One immediate application of this is to quasiflats in CAT(0) cube complexes.

Let (M, ρ) be a connected complete median metric space of rank ν. (There is no loss in assuming completeness, since the completion of a median metric space is again a median metric space of the same rank.) Then M is a geodesic space. Indeed, it is a coarse median space: we can take the function, h, in (C2) to be identically 0. (Given $A \subseteq M$, we can just take $\Pi = \langle A \rangle$ with the induced median.)

Recall that a "panel" is a space isometric to a direct product of non-trivial real intervals with the l^1-metric. After subdividing, we can assume each factor to be either a non-trivial compact real interval, or a ray, $[0, \infty)$. If all factors are rays, we refer to a panel as an "orthant". We use the terms "n-panel" and "n-orthant" to mean that there are n factors. By a "panel" or "orthant in M", we mean a convex subset of M, which is intrinsically a panel.

Here is a restatement of Theorem 1.3.

Proposition 12.1. *Let M be a complete connected median metric space of rank ν. Let $f : \mathbb{R}^\nu \longrightarrow M$ be a quasi-isometric embedding. Then there are ν-panels, $\Phi_1, \ldots, \Phi_p$, in M such that* $\mathrm{hd}(f(\mathbb{R}^\nu), \bigcup_{i=1}^p \Phi_i) \leq s$*, where k and s depend only* *on the parameters of the hypotheses. If $\Phi_1, \ldots, \Phi_q$ are those panels which are orthants, then* $\mathrm{hd}(f(\mathbb{R}^\nu), \bigcup_{i=1}^q \Phi_i) < \infty$.

Note that "rank" here refers to the rank of M as a median algebra, as defined in Section 2. (See the Remark after Lemma 12.2.)

We remark that the bound on the number of panels (or orthants) is a computable function of ν and the quasi-isometry parameters. However, our argument does not give any explicit means of computing s.

Deducing Proposition 12.1 from the coarse version of the last section is essentially a matter of approximating quasicubes by genuine cubes.

Lemma 12.2. *Given $n \in \mathbb{N}$ and $k \geq 0$, there exist $r, t \geq 0$ with the following property. Suppose $R \subseteq M$ is a t-separated k-quasicube of dimension n. Then there is an n-cube, $Q \subseteq M$, with* $\mathrm{hd}(Q, R) \leq r$. *Indeed, there is a median isomorphism, $\theta : Q \longrightarrow R$ with $\rho(x, \theta x) \leq r$ for all $x \in Q$.*

Proof. Let $t = kt(n)$ as in Lemma 7.2, and set $\Pi_0 = R$. In this case, we can take the function h to be identically 0. (In the proof it arose from the function of property (C2) which can be taken to be 0 for a median metric space.) Lemma 7.2 therefore gives us a finite median subalgebra, $\Pi \subseteq M$, and a median epimorphism, $\omega : \Pi \longrightarrow R$. By Lemma 2.6, there is a cube $Q \subseteq \Pi$ such that $\omega|Q$ is an isomorphism to R. We set $\theta = \omega|Q$. □

Remark. This shows that if $\mathrm{rank}(M) \leq \nu$, then any $(\nu+1)$-cube is "degenerate", in the sense that there is an upper bound on the width of at least one of its walls. It follows that the coarse median rank of M is at most ν. (See [Bo5] or [NibWZ1].) It could be strictly less. For example, taking a direct product with an interval $[0, 1]$ increases the median-algebra rank by 1, but does not change the coarse median rank.)

Lemma 12.3. *Let P be a ν-panel, and let $\phi : P \longrightarrow M$ be a quasimorphism which is also a quasi-isometric embedding. We assume that all the side lengths of P are at least L, where L is a constant which depends only on the quasimorphism and quasi-isometry parameters. Then there is a ν-panel, $\Phi \subseteq M$, such that the Hausdorff distance between ϕP and Φ is finite and bounded above by another constant depending only on the quasimorphism and quasi-isometry parameters.*

We first deal with the case where P is compact. Let L be a sufficiently large constant as determined below, and let $\phi : P \longrightarrow M$ be an h-quasimorphism and a quasi-isometry. Let $\hat{Q} \subseteq P$ be the set of corners, and let $R = \phi\hat{Q} \subseteq M$. Then R is an h-quasicube. Provided L is large enough, we can suppose that R is t-separated for sufficiently large t so that Lemma 12.2 gives us a ν-cube, $Q \subseteq M$ with $Q \sim_* R$. Let $H(R)$ be the coarse hull of R, as discussed in Section 9. Then $\phi(P) \sim_* H(R)$. (See Lemma 9.3 and subsequent discussion.) We also have $\Phi = \mathrm{hull}(Q) \sim_* H(R)$. (Moving a quasicube a bounded distance moves its coarse hull a bounded distance. This can be seen from the fact that its coarse hull is the coarse interval between two opposite vertices.) We therefore get $\Phi \sim_* \phi(P)$. This proves Lemma 12.3 for compact panels.

We now deal with the general case. For this, we begin with the following observation.

Given a panel, P, and $r \geq 0$, let $C(P, r) \subseteq P$ be the set of points a distance at least r from its boundary, ∂P. If the side-lengths of P are all greater than $2r$, then this is also a panel.

Suppose $P_0 = \mathrm{hull}(Q_0)$ and $P_1 = \mathrm{hull}(Q_1)$ are n-panels. Let $\pi : M \longrightarrow P_1$ be the gate map (nearest point projection) to P_1. This is a median epimorphism. If $\pi|Q_0$ is injective, then $\pi(Q_0)$ is also an n-cube, and $\mathrm{hull}(\pi(Q_0)) = \pi(\mathrm{hull}(Q_0)) = P_0 \cap P_1$. Clearly this applies if, for some $r \geq 0$, we have $Q_0 \subseteq_r P_1$ and all the side-lengths of P_0 are greater than $2r$. In this case, $C(P_0, r) \subseteq P_1$.

Now suppose that, for sufficiently large L as determined below, P is a panel with all side-lengths at least L, and that $\phi : P \longrightarrow M$ is an h-quasimorphism. Write $P = \bigcup_{j=0}^{\infty} P_j$ as an increasing union of compact panels, $P_j = \mathrm{hull}(\hat{Q}_j)$, again with all side-lengths at least L. Let $R_j = \phi(\hat{Q}_j) \subseteq M$. This is an h-quasicube. As before, if L is large enough, we have have a cube $Q_j \subseteq M$, with $Q_j \sim_* R_j$. Let $\Phi_j = \mathrm{hull}(Q_i)$. Then $\Phi_j \sim_* H(R_j) \sim_* \phi(P_j)$.

Now if $j \leq k$, then $Q_i \subseteq_r \Phi_k$ for some uniform $r \geq 0$. We can suppose that all side-lengths of Φ_j are greater than $2r$. Setting $C_j = C(\Phi_j, r)$, we see $C_j \subseteq \Phi_k$.

In summary, for all k, we have $\bigcup_{j=0}^{k} C_j \subseteq \Phi_k$. Now $H_k = \mathrm{hull}(\bigcup_{i=0}^{k} C_j)$ is a ν-panel. Let Φ be the closure of $\bigcup_{k=0}^{\infty} H_k = \mathrm{hull}(\bigcup_{j=0}^{\infty} C_j)$. This is also a panel, and we have $\Phi \sim_* \phi(P)$.

This proves Lemma 12.3 in the general case.

Proof of Proposition 12.1. By Theorem 1.1, $f(\mathbb{R}^\nu)$ is a bounded distance from $\bigcup_{i=1}^{p} \phi_i(P_i)$, where each P_i is a panel, and each ϕ_i is a strong quasimorphism, in particular, an h-quasimorphism for some h depending only on the parameters of Λ and f. We choose $L \geq 0$, depending on h as above. By the second paragraph of Theorem 1.1, the same statement holds if we restrict to those panels P_i which have all side-lengths at least L. (This may increase the Hausdorff distance, but by an amount depending only on the original parameters.) From the choice of L, we see that each $\phi(P_i)$ is a bounded distance from a panel, Φ_i, in M.

The statement about orthants follows similarly. □

This proves Theorem 1.3.

In particular, this applies when M is a CAT(0) cube complex of dimension ν. In this case, we can say a bit more. In the conclusion of Lemma 12.3 we can assume that Φ is a subcomplex of M. This is achieved by insisting that the corners of all the panels involved in the proof are all 0-cells of M.

Corollary 1.4 is now follows directly from Theorem 1.3.

We remark that any median convex subset, C, of M is also convex in the CAT(0) (l^2) metric in the usual sense (that is, the geodesic connecting any two points of C is contained in C). One way to see this is to note that the gate map to C is 1-lipschitz in the l^2 metric. For finite complexes, this can be seen from the description following Lemma 3.2. One can reduce to the finite case, since we only really need to consider projections to intervals.

In fact, the converse also holds for subcomplexes. A subcomplex of M is convex (in either sense) if and only if it is connected, and locally convex. The latter is equivalent to saying that it intersects the link, L, of any cell in a full subcomplex K (i.e. any simplex of L with vertices in K lies entirely in K).

Further discussion of these matters can be found in [Bo9].

13. Projection maps

There is another context in which the conclusion of Theorem 1.1 can be strengthened. One can say more about the structure of coarse panels and orthants in the case where the coarse median space, Λ, comes equipped with a family of coarsely lipschitz maps to uniformly hyperbolic spaces. Indeed many coarse median spaces arise in this way. The original context was that of the mapping class groups pioneered in [MM]. Similar principles apply to Teichmüller space [MM, Ra, D]. More general contexts are discussed in [BeHS1] and in [Bo5], where various axioms for projection maps are listed. All the properties discussed in this section hold for spaces satisfying either of these sets of axioms, namely axioms (1)–(10) of Section 13 of [BeHS1], or axioms (A1)–(A10) of Section 7 of [Bo5]. In particular, we recover the result about quasiflats in asymphoric hierarchically hyperbolic spaces given in [BeHS3], as described at the end of Section 14.

First, suppose that $\mathcal{X}$ is an indexing set, and that to each $X \in \mathcal{X}$ we have associated a space $(\Theta(X), \sigma_X)$ which is k-hyperbolic for some fixed $k \geq 0$. Suppose also that we have a map, $\theta_X : \Lambda \longrightarrow \Theta(X)$, which is uniformly coarsely lipschitz. That is, there are constants, $k_1, k_2 \geq 0$ such that for all $X \in \mathcal{X}$ and all $x, y \in \Lambda$, we have $\sigma_X(\theta_X x, \theta_X y) \leq k_1 \rho(x, y) + k_2$. We will often abbreviate $\sigma_X(x, y) = \sigma_X(\theta_X x, \theta_X y)$. (In this way, σ_X can be viewed as a pseudometric on Λ.)

As noted earlier, a hyperbolic space is coarse median of rank 1, where the median, $\mu_X(x, y, z)$, is a centre of a geodesic triangle with vertices x, y, z. We assume that each θ_X is a uniform quasimorphism. That is, there is some fixed r_0 such that for all $X \in \mathcal{X}$ and all $x, y, z \in \Lambda$, we have $\rho(\theta_X \mu(x, y, z), \mu_X(\theta_X x, \theta_X y, \theta_X z)) \leq r_0$.

Definition. A ***monotone*** path, is a coarsely lipschitz quasimorphism from a connected subset, $I \subseteq \mathbb{R}$, to Λ.

Remark. In the cases of interest to us, one can always reparameterise a monotone path so that it becomes a quasi-isometric embedding (see, for example, Theorem 1.1 of [Bo5]). Moreover, after moving it a bounded distance, we can take it to be quasigeodesic in the traditional sense (that is, a rectifiable path, $\alpha : I \longrightarrow \Lambda$, such that if $t, u \in I$, then $\text{length}(\alpha([t, u]))$ is bounded above by a fixed linear function of $\rho(\beta(t), \beta(u))$). The constants involved depend only on those of Λ and the original monotone path. If the coarse median space is hyperbolic, then the converse is also true: any (unparameterised) quasigeodesic is monotone.

Note that the coordinate maps in any coarse panel in Λ will be monotone paths. More precisely, suppose that $P = \prod_{j=1}^{n} I_j \subseteq \mathbb{R}^n$ is a panel. Suppose that $\alpha_j : I_j \longrightarrow P$ is a map of the form $\alpha_i(t) = (x_1, , \ldots, x_{i-1}, t, x_{i+1}, \ldots, x_n)$ for some $i \in \{1, \ldots, n\}$ and $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n$ fixed. Then it follows directly from the definitions that if $\phi : P \longrightarrow \Lambda$ is a coarsely lipschitz quasimorphism, then $\beta_j = \phi \circ \alpha_j : I \longrightarrow \Lambda$ is a monotone path. (Note also that if ϕ is a quasi-isometric embedding, then so is β_j.)

In fact, we can take the β_i to be independent of the x_j, for $j \neq i$. This is based on the following discussion.

By an h-***quasisquare*** in Λ, we mean a cyclically ordered sequence of points, $a_1, a_2, a_3, a_4 \in \Lambda$ with $a_i \in_h [a_{i-1}, a_{i+1}]$ for all i. In other words, the inclusion of $\{a_1, a_2, a_3, a_4\}$ into Λ is a quasicube of dimension 2, where the pairs of opposite corners are a_1, a_3 and a_2, a_4.

If a_1, a_2, a_3, a_4 is a quasisquare in a hyperbolic space Θ, then either ($a_1 \sim_* a_4$ and $a_2 \sim_* a_3$) or ($a_1 \sim_* a_2$ and $a_3 \sim_* a_4$), (or both). In other words, a quasisquare in a hyperbolic space is degenerate: it collapses onto one of its 1-faces up to bounded distance (or both, in the case where the quasisquare has bounded diameter). This is a simple exercise in hyperbolic spaces.

Clearly, this generalises directly to cubes of any dimension. In summary:

Lemma 13.1. *Suppose Θ is hyperbolic, and $Q \subseteq \Theta$ be a quasicube. Then the inclusion of Q into Θ factors, up to bounded distance, through projection to one of its 1-faces. The bound depends only on the constants of hyperbolicity and quasimorphism.*

As an immediate consequence we get:

Lemma 13.2. *Let $P = \prod_{j=1}^{n} I_j$ be an n-panel, and $\phi : P \longrightarrow \Lambda$ be a coarsely lipschitz quasimorphism. Then for each $X \in \mathcal{X}$, there is some $j(X) \in \{1, \ldots, n\}$ and a monotone path $\gamma_X : I_{j(X)} \longrightarrow \Theta(X)$ such that $\theta_X \circ \phi$ agrees up to bounded distance with $\gamma_X \circ \pi_{j(X)}$, where $\pi_j : P \longrightarrow I_j$, is projection to the ith coordinate.*

Of course, it is possible that γ_X may have bounded image, in which case, the image of $\theta_X \circ \phi$ is also bounded. In this case, we could just leave $j(X)$ undefined.

We will choose a constant, K, sufficiently large as determined below, and let $\mathcal{X}_0$ be the set of $X \in \mathcal{X}$ for which the diameter of $\theta_X(\phi(P))$ is at most K. In particular, we can suppose that K is large enough so that $j(X)$ is defined for all $X \in \mathcal{X} \setminus \mathcal{X}_0$. For $i \in \{1, \ldots, n\}$, we write $\mathcal{X}_j = \{X \in \mathcal{X} \setminus \mathcal{X}_0 \mid j(X) = j\}$. Thus, $\mathcal{X} = \bigsqcup_{j=0}^{n} \mathcal{X}_j$.

To proceed we will assume the following "distance bound" hypothesis:

(P1): Given any $r \geq 0$, there is some $r' \geq 0$ such that if $x, y \in \Lambda$ and $\sigma_X(x, y) \leq r$ for all $X \in \mathcal{X}$, then $\rho(x, y) \leq r'$.

In this case if $\gamma : I \longrightarrow \Lambda$ then γ is monotone if and only if each of the maps $\theta_X \circ \gamma : I \longrightarrow \Theta(X)$ is uniformly monotone for all $X \in \mathcal{X}$.

The latter condition (or some slight variation of it) is frequently used to define a "hierarchy path" in the situation where we have a family of maps to hyperbolic spaces. For example, in the case of the mapping class groups, a path arising from the resolution of a "hierarchy" in the sense of [MM] is easily seen to be a hierarchy path in this sense.

(Under certain other assumptions, one can show that there is always a monotone, or hierarchy, path between any two points of Λ. This is true in any hierarchically hyperbolic space [BeHS2]. See also [Bo8], where this is shown under more general assumptions. This is not directly relevant here, since as we have observed, our coarse panels automatically give rise to monotone paths.)

Under another condition one can say more. Let us suppose that there is a relation, $\perp$, on $\mathcal{X}$ which is antireflexive and symmetric. We suppose:

(P2): Given h there is some $r_0 \geq 0$ such that if $a_1, a_2, a_3, a_4 \in \Lambda$ is an h-quasisquare and $X, Y \in \mathcal{X}$ with $\sigma_X(a_1, a_2) > r_0$ and $\sigma_Y(a_1, a_3) > r_0$, then $X \perp Y$.

The above is a consequence of other axioms of projection maps, in particular it follows from Axiom (P4) of [Bo1]. It therefore also follows from Properties (A1)–(A10) of [Bo5], or from the axioms of a hierarchically hyperbolic space.

We can take the constant K (used in defining the partition of $\mathcal{X}$) greater than r_0. It then follows that if $X \in \mathcal{X}_j$ and $Y \in \mathcal{X}_k$ with $j \neq k$, then $X \perp Y$. Note also that, in view of the assumption (P1) above, if we take the constant L of Theorem 1.1 sufficiently large, then $\mathcal{X}_j \neq \varnothing$ for all j.

Under some further assumptions of "orthogonality" and "nesting", one can show that there are pairwise $\perp$-related elements, $U_1, \ldots, U_n \in \mathcal{X}$, such that all of the elements of $\mathcal{X}_j$ are "nested" in U_j for each $j \in \{1, \ldots, n\}$. The projection of $\phi(P)$ to any $X \in \mathcal{X}_0$ will have bounded diameter, so we can disregard these. For example, the above holds under the axioms of "orthogonality" given in [BeHS1], or with the family of subsurfaces of a compact surface as we discuss in the next section.

14. Mapping class groups, Teichmüller space and hierarchically hyperbolic spaces

Before discussing hierarchically hyperbolic spaces, we first illustrate how the statements of Section 13 apply to the mapping class groups and Teichmüller space.

In these cases, the indexing set $\mathcal{X}$ consists of subsurfaces of a fixed surface, "orthogonal" means "disjoint" and "nested in" means "included in".

To be more precise, let Σ be a compact orientable surface of complexity, $\xi = \xi(\Sigma)$ (i.e. 3 times the genus plus the number of boundary components minus 3). Let $\Lambda = \mathbb{M}(\Sigma)$ be the marking graph. (There are a number of different but related formulations of this, see [MM] or [Bo5]. One could also take it to be the Cayley graph of the mapping class group with respect to any finite generating set. Since these are all quasi-isometric, any one would serve for our purposes here.) Let $\mathcal{X}$ be the collection of subsurfaces of Σ, in the sense of [MM]. (These are assumed to be connected, π_1-injective, neither discs nor three-holed spheres, not homotopic into the boundary of Σ, and to be defined up to homotopy.) If $X, Y \in \mathcal{X}$, we write $X \perp Y$ to mean that X, Y can be homotoped to be disjoint. (This was denoted $\wedge$ in [Bo1, Bo4, Bo5].) We will write $Y \preceq X$ to mean that Y can be homotoped into X but not into the boundary of X. (In the latter case, Y must be an annulus, and we would have $X \perp Y$.) If none of the conditions, $X = Y$, $X \perp Y$, $X \preceq Y$ or $Y \preceq X$ hold, we say that X, Y are ***transverse***. Given $X \in \mathcal{X}$, let $\Theta(X)$ be the intrinsic curve graph of X (appropriately defined for annuli) and let $\theta_X : \mathbb{M}(\Sigma) \longrightarrow \Theta(X)$ be the usual subsurface projection map. We also have maps, $\psi_X : \mathbb{M}(\Sigma) \longrightarrow \mathbb{M}(X)$, to the intrinsic marking graph of X. These maps satisfy axioms (A1)–(A10) of [Bo5]. In particular, $\mathbb{M}(\Sigma)$ is coarse median of rank ξ.

Remark. Note that we are only considering connected subsurfaces of Σ, as in [MM]. However, it [BeHS2, BeHS3] it is necessary to include disconnected subsurfaces in the indexing set in order to satisfy their "orthogonality axiom". This introduces some complications into the proceedings, but it does not directly affect the discussion here.

Now suppose $\phi : P \longrightarrow \mathbb{M}(\Sigma)$ is a coarse ξ-panel. Let $\mathcal{X} = \bigsqcup_{j=0}^{\xi} \mathcal{X}_j$ be the partition of $\mathcal{X}$ described in Section 13, where the constant, K, is chosen sufficiently large as described below. Thus, if $X \in \mathcal{X}_j$ for $j > 0$, we have a path $\gamma_X : I_j \longrightarrow \Theta(X)$ such that $\gamma_X \circ \pi_j$ agrees up to bounded distance with $\theta_X \circ \phi$. Moreover, if $Z \notin \mathcal{X}_j$, then $\theta_Z \circ \phi(P)$ has bounded diameter. In particular, this applies to all $Z \in \mathcal{X}_0$.

The elements of distinct $\mathcal{X}_j$ are disjoint. It follows that there are disjoint subsurfaces, $U_j \subseteq \Sigma$, such that $X \preceq U_j$ for all $X \in \mathcal{X}_j$. In fact, we can take U_j minimal containing all elements of $\mathcal{X}_j$. We note that each U_j must be an annulus, a one-holed torus, or a four-holed sphere. (One way to see this is to take a curve δ_j in each U_j, so that δ_j is a core curve if U_j is an annulus, or else an essential non-peripheral curve in U_j otherwise. The resulting curves are disjoint and non-homotopic in Σ. There can be at most ξ such curves, and this bound is attained precisely when each complementary region is a three-holed sphere. In particular, if U_j is non-annular, then $U_j \setminus \delta_j$ consists of either one, or two, three-holed spheres.)

Suppose that $Z \in \mathcal{X}$ does not lie in any U_j. Then it must be transverse to, or strictly contain, at least one U_j. We write $a_Z \in \Theta(Z)$ for the projection of U_j to $\Theta(Z)$. Since all the U_j is disjoint, a_Z is well defined up to bounded distance, regardless of which such U_j we choose. We claim that $\theta_Z(\phi(P))$ lies in a bounded neighbourhood of a_Z. To see this, first note that by the minimality of U_j, there must be some $X \in \mathcal{X}_j$ with Z transverse to, or containing, X. The projection of X to $\Theta(Z)$ lies a bounded distance from a_Z. Now suppose, for contradiction, that $\theta_Z(\phi(P))$ is a long way from this projection. Suppose first that Z is transverse to X. Let $a \in \phi(P)$. Then by Behrstock's Lemma [Be], $\theta_X a$ is a bounded distance from the projection of Z to $\Theta(X)$. (This is stated as (A9) in [Bo5].) Since this holds for all such a, we deduce that $\theta_X(\phi(P))$ has bounded diameter. This contravenes the definition of $\mathcal{X}_j$, assuming that we have chosen K large enough in relation to the constant of Behrstock's Lemma. Suppose on the other hand that X is contained in Z. Let $a, b \in \phi(P)$. This time, we can use the Bounded Geodesic Image theorem of Masur and Minsky [MM]. (This is stated as (A8) in [Bo5].) to show that $\theta_X a$ is a bounded distance from $\theta_X b$. Since this holds for all $a, b \in \phi(P)$, we again get a contradiction. This proves the claim.

Now let $j > 0$, and let $\alpha_j : I_j \longrightarrow P$ be a coordinate map. Let $\beta = \phi \circ \alpha_j : I_j \longrightarrow \mathbb{M}(\Sigma)$, and let $\gamma_j = \psi_{U_j} \circ \beta_j : I_j \longrightarrow \mathbb{M}(U_j)$. These are both monotone paths. Then if $t \in (t_1, \dots, t_\xi) \in P$, and $X \preceq U_j$, then $\theta_X(\phi(t))$ is a bounded distance from $\theta_X(\gamma_j(t_j))$. If Z does not lie in any U_j, then $\theta_Z(\phi(t))$ is a bounded distance from a_Z. Note that, by the distance bound (P1) of Section 13, these projections determine $\phi(t)$ up to bounded distance.

In all cases above, $\mathbb{M}(U_j)$ is hyperbolic. Therefore, saying that a path in $\mathbb{M}(U_j)$ is "monotone" is equivalent to saying that it is an unparameterised quasigeodesic. We can therefore take it to be a uniform quasigeodesic. Indeed, given that $\mathbb{M}(U_j)$ is a locally finite graph, we can take it be geodesic.

Conversely, given such a family of geodesics, $\gamma_1, \dots, \gamma_\xi$, we can reconstruct ϕ as we now describe.

To describe this, first note that there is a coarsely lipschitz embedding of $\mathbb{M}(U_1) \times \cdots \times \mathbb{M}(U_n)$, into $\mathbb{M}(\Sigma)$ such that postcomposition with ψ_{U_j} is just projection to the $\mathbb{M}(U_j)$ coordinate, and such that if Z does not lie in any U_j, then the image of the embedding projects under θ_Z into a bounded neighbourhood of $a_Z \in \Theta(Z)$ as defined above. Moreover, the embedding is a strong quasimorphism. For more detail, see the discussion of product regions in [Bo5].

Now given geodesics $\gamma_j : I_j \longrightarrow \mathbb{M}(U_j)$, we can combine them to give a quasimorphism, $\hat{\phi} = \gamma_1 \times \cdots \times \gamma_\xi : P \longrightarrow \mathbb{M}(\Sigma)$ via this embedding. Suppose the γ_j arise from a map ϕ as described above. Then by construction, we see that $\theta_X(\hat{\phi}(t))$ lies a bounded distance from $\theta_X(\phi(t))$ for all $X \in \mathcal{X}$. Therefore, by (P1) again, we see that $\hat{\phi}$ agrees with ϕ up to bounded distance. In other words, any coarse panel $\phi : P \longrightarrow \mathbb{M}(\Sigma)$ (for L sufficiently large) arises from this construction.

In summary, we have shown:

Theorem 14.1. *Let Σ be a compact orientable surface of complexity ξ, and let $\mathbb{M}(\Sigma)$ be the marking graph of Σ. Let $f : \mathbb{R}^\xi \longrightarrow \mathbb{M}(\Sigma)$ be a quasi-isometric embedding. Then there is some $p \in \mathbb{N}$, depending only on ξ and the quasi-isometry parameters of f, such that for each $i \in \{1, \ldots, p\}$, we have a collection of disjoint subsurfaces, $U_{i1}, \ldots, U_{i\xi}$, of Σ, together with a family of geodesics, $\gamma_{ij} : I_{ij} \longrightarrow \mathbb{M}(U_{ij})$, with the following properties. Let $P_i = \prod_{j=1}^{\xi} I_{ij}$. Then the product map $\phi_i = \gamma_{i1} \times \cdots \gamma_{i\xi} : P_i \longrightarrow \mathbb{M}(\Sigma)$, obtained by combining the paths γ_{ij} as described above, is a quasi-isometric embedding, with constants depending only on ξ. Moreover, the Hausdorff distance between $f(\mathbb{R}^\xi)$ and $\bigcup_{i=1}^{p} \phi_i(P_i)$ is bounded in terms of ξ and the quasi-isometry parameters of f.*

If we allow the Hausdorff distance between $f(\mathbb{R}^\nu)$ and $\bigcup_{i=1}^{p} \phi_i(P_i)$ to be finite, but not necessarily uniformly bounded, then we can take each I_{ij} to be $[0, \infty)$ (so that each γ_{ij} is a geodesic ray, and each P_i is a ξ-orthant). The orthants, P_i, can be assembled into an orthant complex, Ω, bilipschitz equivalent to $\mathbb{R}^\xi$; and the maps, ϕ_i, combine to give us a quasi-isometric embedding, $\phi : \Omega \longrightarrow \mathbb{M}(\Sigma)$.

Essentially the same discussion applies to Teichmüller space with the Teichmüller metric, where $\Theta(X)$ is appropriately modified when X is an annulus. This is again coarse median of rank ξ. However in this case, quasiflats only exist when Σ has genus at most 1, or is a closed surface of genus 2 [Bo4]. (Indeed, this result might be used to give another proof of this fact. A quasiflat gives rise an orthant complex bilipschitz equivalent to $\mathbb{R}^\xi$. This is the cone over an embedded homology $(\xi - 1)$-sphere in the curve complex of Σ. Such can only arise in the above cases. We will not give details here since the logic is similar to that of the original argument.)

Similarly, the Weil-Petersson metric (or equivalently the pants graph) is coarse median of rank $\lfloor (\xi + 1)/2 \rfloor$. In this case, $\mathcal{X}$ consists only of non-annular surfaces. (Alternatively, we could instead take $\Theta(X)$ to be a singleton whenever X is an annulus.)

Here the maps ψ_X can be defined in a number of equivalent ways, for example, via the respective combinatorial model spaces (see [Ra, D, Bro]) as used in [Bo4, Bo7].

In summary, we have:

Theorem 14.2. *The statement of Theorem 14.1 holds on replacing $\mathbb{M}(\Sigma)$ with Teichmüller space in either the Teichmüller or Weil-Petersson metric, together in the latter case, on replacing the dimension ξ by $\lfloor (\xi + 1)/2 \rfloor$.*

More generally, Theorems 14.1 and 14.2 hold for any space, with projection maps satisfying the axioms (A1)–(A10) as laid out in Section 7 of [Bo5]. (Though we may need to take the paths γ_{ij} to be uniform quasigeodesics, rather than geodesics.) We just replace the space, $\mathbb{M}(X)$, by the space $\mathcal{M}(X)$ defined there, and replace ξ by the rank, ν, of $\mathcal{M}(\Sigma)$. Everything we need, such as a version of Behrstock's Lemma and the Bounded Geodesic Image Theorem, is incorporated into the axioms.

Finally, we make a few comments to relate the above to the main theorem formulated in [BeHS3] for an asymphoric hierarchically hyperbolic space, $\mathcal{M}$. There the indexing set $\mathcal{X}$ (which the authors denote by "$\mathfrak{S}$") is assumed to come equipped with relations corresponding to $\preceq$ (there denoted "$\sqsubseteq$") and $\perp$. To each $X \in \mathcal{X}$ they associate a uniformly hyperbolic space, CX, together with a coarsely lipschitz map from $\mathcal{M}$ to CX. If $X \preceq Y$, there is assumed to be a map ρ^Y_X from CY to the power set of CX. (It is not made explicit when sets in the image of this map are assumed to be non-empty, but that does not directly affect the present discussion.) These, together with certain other maps, are required to satisfy a list of axioms; so that the space, $\mathcal{M}$, equipped with this structure is "(asymphorically) hierarchically hyperbolic". Given $U \in \mathcal{X}$, the authors define a space of "consistent tuples", denoted F_U. A consistent tuple is an element of the direct product of the spaces CX as X varies over those $X \in \mathcal{X}$ for which $X \preceq U$, and which satisfies certain "consistency" conditions as laid out in [BeHS2]. A path in F_U is said to be a "hierarchy path" if the projection to each factor, CX, is an unparameterised quasigeodesic.

One can interpret this in terms of the axioms (A1)–(A10) of [Bo5]. (There the indexing set was taken to consist of subsurfaces of Σ, rather than an abstract set with relations satisfying "orthogonality" axioms, but either is sufficient for the earlier discussion to apply.) In this context, the role of CX is played by $\Theta(X)$ (denoted $\mathcal{G}(X)$ in [Bo5]), and the role of F_U is played by $\mathcal{M}(U)$. The projection of $\mathcal{M}$ to $\Theta(X)$ corresponds to the projection of F_U to the factor CX. From the axioms of [Bo5], the median on $\mathcal{M}(U)$ is characterised by the fact that the projection maps are uniform quasimorphisms. In particular, it follows that a monotone path in $\mathcal{M}(U)$ corresponds to a hierarchy path in F_U. Therefore coarse orthants, as we have defined them, will be orthants in the sense defined in [BeHS3]. Any asymphoric hierarchically hyperbolic space of "rank ν" in their sense is coarse median of rank ν in our sense. From this, one sees that Theorem A of [BeHS3] follows from Theorem 1.2 here.

15. A variant of the Borsuk-Ulam theorem

The following strengthening of the Borsuk-Ulam Theorem is a straightforward consequence of the fact that a self-map of a sphere which commutes with the antipodal map has odd degree. The latter statement can be found in [Bre] (Theorem IV.20.6 thereof) though by a somewhat more involved argument. For convenience, we give a self-contained proof, based on an argument for a related result given in [W]. We reproduce this, with some simplification for this particular case, below. This result is used in the proof of Lemma 6.1.

Let $S^n = \{x \in \mathbb{R}^{n+1} \mid ||x|| = 1\}$ be the n-sphere. Let $\tau : S^n \longrightarrow S^n$ be the antipodal map: $\tau(x) = -x$. Let H_i be the ith reduced singular homology group of S^n with $\mathbb{Z}_2$ coefficients. So $H_n \cong \mathbb{Z}_2$ and $H_i = 0$ for $i \neq n$. We say that a continuous map $f : S^n \longrightarrow S^n$ has ***even degree*** if it maps H_n to 0. (This accords with the standard notion of "degree" in this context.)

Theorem 15.1. *If $f : S^n \longrightarrow S^n$ has even degree then there is some $x \in S^n$ with $f(\tau x) = f(x)$.*

Proof. Suppose not. Define $F : S^n \times [0, 1] \longrightarrow S^n$ by $F(x, t) = \frac{f(x)-tf(-x)}{||f(x)-tf(-x)||}$. Then $F(x, 0) = f(x)$ and $F(-x, 1) = -F(x, 1)$. Define $g : S^n \longrightarrow S^n$ by $g(x) = F(x, 1)$. Then F is a homotopy from f to g, so g also has even degree. Note that $g \circ \tau = \tau \circ g$.

Let C_i be the ith reduced singular chain complex (so $C_{-1} = \mathbb{Z}_2$). Let $\beta = g_* : C_i \longrightarrow C_i$, and $\theta = 1_* + \tau_* : C_i \longrightarrow C_i$ (i.e. $\theta(c) = c + \tau_*(c)$). Let $\partial : C_i \longrightarrow C_{i-1}$ be the boundary map. Then β, θ, ∂ all commute. Also, $\theta^2 = 0$. Note that if $c \in C_i$, for $i < n$, and $\partial c = 0$, then there is some $b \in C_{i+1}$ with $c = \partial b$.

By taking hemispheres, we can find elements $h_i \in C_i$ for $i = 0, \ldots, n$ such that $\partial h_i = \theta h_{i-1}$, such that $\partial h_0 = 1 \in \mathbb{Z}_2$, and such that θh_n generates H_n. (Note that $\partial\theta h_i = \theta^2 h_{i-1} = 0$.)

Let $k_i = \beta h_i$. So $\partial k_0 = 1$ and $\partial k_i = \theta k_{i-1}$. We claim that for all $i \leq n$, there exists $a_i \in C_i$ such that $\partial(h_i + k_i + \theta a_i) = 0$. We show this by induction on i. Set $a_0 = 0$. Suppose we have found a_i for $i < n$. There is some $a_{i+1} \in C_{i+1}$ with $h_i + k_i + \theta a_i = \partial a_{i+1}$. Now $\partial(h_{i+1} + k_{i+1} + \theta a_{i+1}) = \theta h_i + \theta k_i + \theta\partial a_{i+1} = 2\theta h_i + 2\theta k_i + \theta^2 a_i = 0$. The claim follows.

In particular, $\partial(h_n + k_n + \theta a_n) = 0$. Since H_n is generated by the class of θh_n, there is some $b \in C_{n+1}$ such that either $h_n + k_n + \theta a_n = \partial b$, or $h_n + k_n + \theta a_n + \theta h_n = \partial b$. Either way, $\theta h_n + \theta k_n = \theta\partial b = \partial\theta b$. Therefore, θk_n represents the generator, θh_n, of H_n. But since g has even degree, $\theta k_n = \beta\theta h_n$ must be trivial, giving a contradiction. ∎

References

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

[Be] J.A.Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space* : Geom. Topol. **10** (2006) 1523–1578.

[BeHS1] J.A.Behrstock, M.F.Hagen, A.Sisto, *Hierarchically hyperbolic spaces I: curve complexes for cubical groups* : Geom. Topol. **21** (2017) 1731–1804.

[BeHS2] J.A.Behrstock, M.F.Hagen, A.Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula* : Pacific J. Math. **299** (2019) 257–338.

[BeHS3] J.A.Behrstock, M.F.Hagen, A.Sisto, *Quasiflats in hierarchically hyperbolic spaces* : Duke Math. J. **170** (2021) 909–996.

[BeKMM] J.A.Behrstock, B.Kleiner, Y.Minsky, L.Mosher, *Geometry and rigidity of mapping class groups* : Geom. Topol. **16** (2012) 781–888.

[BeM] J.A.Behrstock, Y.N.Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.

[Bo1] B.H.Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.

[Bo2] B.H.Bowditch, *Embedding median algebras in products of trees* : Geom. Dedicata **170** (2014) 157–176.

[Bo3] B.H.Bowditch, *Some properties of median metric spaces* : Groups, Geom. Dyn. **10** (2016) 279–317.

[Bo4] B.H.Bowditch, *Large-scale rank and rigidity of the Teichmüller metric* : J. Topol. **9** (2016) 985–1020.

[Bo5] B.H.Bowditch, *Large-scale rigidity properties of the mapping class groups* : Pacific J. Math. **293** (2018) 1–73.

[Bo6] B.H.Bowditch, *Median and injective metric spaces* : Math. Proc. Camb. Phil. Soc. **168** (2020) 43–55.

[Bo7] B.H.Bowditch, *Large-scale rank and rigidity of the Weil-Petersson metric* : Groups Geom. Dyn. **14** (2020) 607–652

[Bo8] B.H.Bowditch, *Convex hulls in coarse median spaces* : preprint, Warwick 2018.

[Bo9] B.H.Bowditch, *Median algebras* : preprint, Warwick 2022.

[Bre] G.E.Bredon, *Topology and Geometry* : Graduate Texts in Mathematics, Vol. 139, Springer (1993).

[Bro] J.F.Brock, *The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores* : J. Amer. Math. Soc. **16** (2003) 495–535.

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

[Che] V.Chepoi, *Graphs of some CAT(0) complexes* : Adv. in Appl. Math. **24** (2000) 125–179.

[D] M.G.Durham, *The augmented marking complex of a surface* : J. Lond. Math. Soc. **94** (2016) 933–969.

[G] M.Gromov, *Asymptotic invariants of infinite groups* : “Geometric group theory, Vol. 2” London Math. Soc. Lecture Note Ser. No. 182, Cambridge Univ. Press (1993).

[H] J.Huang, *Top-dimensional quasiflats in CAT(0) cube complexes* : Geom. Topol. **21** (2017) 2281–2352.

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

[KaKL] M.Kapovich, B.Kleiner, B.Leeb, *Quasi-isometries and the de Rham decomposition* : Topology **37** (1998) 1193–1211.

[KlLa] B.Kleiner, U.Lang, *Higher rank hyperbolicity* : Invent. Math. **221** (2020) 597–664.

[KlLe] B.Kleiner, B.Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings* : Inst. Hautes Études Sci. Publ. Math. **86** (1997) 115–197.

[MM] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.

[NibWZ1] G.A.Niblo, N.Wright, J.Zhang, *A four point characterisation for coarse median spaces* : Groups Geom. Dyn. **13** (2019) 939–980.

[NibWZ2] G.A.Niblo, N.Wright, J.Zhang, *Coarse median algebras: the intrinsic geometry of coarse median spaces and their intervals* : Selecta Math. **27** (2021) No. 2, Paper No. 20, 50 pp.

[Nie] J.Nieminen, *The ideal structure of simple ternary algebras* : Colloq. Math. **40** (1978/79) 23–29.

[Ra] K.Rafi, *A combinatorial model for the Teichmüller metric* : Geom. Funct. Anal. **17** (2007) 936–959.

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

[S] M.Sholander, *Medians and betweenness* : Proc. Amer. Math. Soc. **5** (1954) 801–807.

[VaW] L.Van den Dries, A.J.Wilkie, *On Gromov's theorem concerning groups of polynomial growth and elementary logic* : J. Algebra **89** (1984) 349–374.

[Ve] E.R.Verheul, *Multimedians in metric and normed spaces* : CWI Tract, 91, Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam, 1993.

[W] J.W.Walker, *A homology version of the Borsuk-Ulam theorem* : Amer. Math. Monthly **90** (1983) 466–468.

[Z1] R.Zeidler, *Coarse median structures on groups* : Masterarbeit, University of Vienna, 2013.

[Z2] R.Zeidler, *Coarse median structures and homomorphisms from Kazhdan groups* : Geom. Dedicata **180** (2016) 49–68.

Mathematics Institute, University of Warwick, Coventry, CV4 7AL, Great Britain