

MEDIAN ALGEBRAS

BRIAN H. BOWDITCH

ABSTRACT. We give a self-contained account of the basic theory of median algebras. We explore notions of betweenness, convexity, walls, duality etc. In this context we include discussion of median metric spaces, cube complexes, spaces with measured walls, coarse median spaces, quasimedial graphs, and related structures.

CONTENTS

1. Introduction	5
1.1. Conventions	7
2. Distributive lattices	7
2.1. Lattices and medians	8
2.2. sublattices and Sperner families	10
2.3. Examples	11
3. Basic facts about median algebras	12
3.1. Definitions	12
3.2. Some fundamental constructs	13
3.3. Free median algebras	19
3.4. Examples	21
4. Intervals and betweenness	25
4.1. Description of median algebras in terms of intervals	25
4.2. The short distributive law	27
4.3. Isbell's condition	29
5. Free median algebras	30
5.1. Some general notions	30
5.2. Flows and superextensions	31
5.3. The structure of a superextension	33
5.4. Description of the first few cases	35
6. Expressions and identities	38
6.1. Formal definitions	38
6.2. Verification of identities	39
6.3. Explicit derivations of identities	41
7. Convex sets	46
7.1. Some definitions and basic facts	46

Date: First draft: 10th March 2022. Revised: 21st March 2024.

7.2.	Parallel sets and translations	48
7.3.	Gates	50
7.4.	Convex hulls	52
7.5.	Some final observations	56
8.	Walls and rank	56
8.1.	Walls and halfspaces	57
8.2.	Properties of rank and subalgebras	60
8.3.	Colourability	63
9.	Halfspaces and duality	65
9.1.	Prosets	65
9.2.	Duality	68
9.3.	The rank of a proset	71
9.4.	The discrete case	73
9.5.	The structure of the space of flows	75
9.6.	Formulation in terms of ideals	81
10.	Hypercubes	82
10.1.	Hypercubes and their faces	82
10.2.	Construction of subdivisions and realisations	83
10.3.	Finite cubes	84
11.	Discrete median algebras	87
11.1.	Some basic definitions	88
11.2.	Relation to cube complexes	90
11.3.	Some basic properties	92
11.4.	Links and local convexity	93
11.5.	Parallel edges and convex hulls	95
11.6.	Canonical paths	97
11.7.	Generating subalgebras	99
11.8.	Retracts	100
11.9.	More about free median algebras	103
11.10.	Subdivisions	106
11.11.	Infinite cubes	108
11.12.	The Roller boundary	111
11.13.	Event structures	117
12.	Topological median algebras	120
12.1.	Definition and examples	120
12.2.	Gates and local convexity	123
12.3.	Rank and dimension	125
12.4.	Totally ordered sets and connectedness	127
12.5.	Totally disconnected median algebras and duality	133
12.6.	A compactification procedure	134
12.7.	Rank-1 topological median algebras	136
13.	Median metric spaces	137
13.1.	Definition and examples	137

13.2.	Some basic properties	141
13.3.	Connectedness, geodesics and embeddings	145
13.4.	Connections with CAT(0) and injective metric spaces	148
14.	Submedian relations	148
14.1.	Characterisation of submedian relations	149
14.2.	Pretrees	153
15.	\mathbb{R} -trees	158
15.1.	Characterisations of \mathbb{R} -trees	158
15.2.	Boundaries of \mathbb{R} -trees	161
15.3.	Embedding in products of \mathbb{R} -trees	161
15.4.	Guirardel cores	163
16.	Median graphs	163
16.1.	Characterisations of median and modular graphs	164
16.2.	Equivalence with CAT(0) cube complexes	166
16.3.	Cubical structures	174
17.	Cube complexes	176
17.1.	Topologies, metrics, and statement of main results	176
17.2.	Proofs	179
17.3.	Subdivisions	185
17.4.	Subalgebras	186
18.	The CAT(0) property	187
18.1.	Definition and basic facts	187
18.2.	Hyperplanes and convexity in CAT(0) cube complexes	189
18.3.	Another description of the equivalence with median graphs	191
19.	Spaces with measured walls	193
19.1.	Definitions	194
19.2.	Duality with median metric spaces	195
19.3.	A construction of a median metric	201
19.4.	Examples	202
20.	Boolean functions and majority vote	205
20.1.	Definitions and connections with median algebras	205
20.2.	Clones	207
20.3.	Majority consensus algebras	209
21.	Group actions	214
21.1.	Fixed points	214
21.2.	Actions on cube complexes	218
21.3.	Some results regarding more general median metric spaces	223
22.	Gates	226
22.1.	General betweenness axioms and gates	226
22.2.	Applications to metric spaces and graphs	228
23.	Quasimedial graphs	229
23.1.	Definitions	230
23.2.	Examples	231

23.3.	Gates and prisms	234
23.4.	The prism complex	237
23.5.	Some consequences	244
23.6.	Existence and uniqueness of quasimedial triples	244
23.7.	Some further observations	246
24.	Coarse geometry	249
24.1.	Quasi-isometries and hyperbolicity	249
24.2.	Coarse median spaces	251
24.3.	An outline of some applications	253
24.4.	Coarse median algebras	255
25.	Injective metric spaces and helly graphs	255
25.1.	Injective metrics	256
25.2.	Helly graphs	257
26.	Notes	257
26.1.	Notes on Section 1 (Introduction)	257
26.2.	Notes on Section 2 (Distributive lattices)	258
26.3.	Notes on Section 3 (Basic facts about median algebras)	258
26.4.	Notes on Section 4 (Intervals and betweenness)	259
26.5.	Notes on Section 5 (Free median algebras)	260
26.6.	Notes on Section 6 (Expressions and identities)	260
26.7.	Notes on Section 7 (Convex sets)	262
26.8.	Notes on Section 8 (Walls and rank)	262
26.9.	Notes on Section 9 (Halfspaces and duality)	262
26.10.	Notes on Section 10 (Hypercubes)	263
26.11.	Notes on Section 11 (Discrete median algebras)	263
26.12.	Notes on Section 12 (Topological median algebras)	264
26.13.	Notes on Section 13 (Median metric spaces)	265
26.14.	Notes on Section 14 (Submedian relations)	266
26.15.	Notes on Section 15 (\mathbb{R} -trees)	266
26.16.	Notes on Section 16 (Median graphs)	267
26.17.	Notes on Section 17 (Cube complexes)	268
26.18.	Notes on Section 18 (The CAT(0) property)	268
26.19.	Notes on Section 19 (Spaces with measured walls)	269
26.20.	Notes on Section 20 (Boolean functions and majority vote)	270
26.21.	Notes on Section 21 (Group actions)	270
26.22.	Notes on Section 22 (Gates)	271
26.23.	Notes on Section 23 (Quasimedial graphs)	271
26.24.	Notes on Section 24 (Coarse geometry)	273
26.25.	Notes on Section 25 (Injective metric spaces and helly graphs)	275
	References	275

1. INTRODUCTION

A “median algebra” is a set, M , equipped with a symmetric ternary operation, $[(a, b, c) \mapsto abc] : M^3 \rightarrow M$, such that $aab = a$ and $(abd)cd = (acd)bd$ for all $a, b, c, d \in M$.

This definition is easy to state, but perhaps not particularly illuminating. It is difficult to use directly in this form. Nevertheless, these axioms (at least in one of several equivalent forms) have been proposed, or rediscovered, several times since the 1940s. They give rise to a very rich theory. Examples of such structures arise naturally in many different contexts, including algebra, geometry and computer science. They have found wide applications, in subjects as diverse as group theory and biology. Many of the earlier accounts tend to focus on universal algebra and lattice theory, and there is now a very extensive literature on the subject from this perspective. The geometric applications have come to the fore more recently. In this account, we aim to focus on geometry (though this may not be apparent from the first few sections). Most of what we do can be viewed as exposition, though it includes various results and arguments that I have not been able to find in the literature.

From a geometric viewpoint, the median, abc , of three points a, b, c can be thought of as a canonically determined point which lies “between” any two distinct points of the triple $\{a, b, c\}$.

A specific example which makes this explicit is a simplicial tree; that is, a connected graph with no embedded cycles. The vertex set, V , of such a tree has a natural structure of a median algebra. In this case, abc , is the unique vertex which lies on each of the three arcs of the tree pairwise connecting $a, b, c \in V$. It is readily verified that this satisfies the axioms given at the beginning.

For a more algebraic example, let X be any set, and let $\mathcal{P}(X)$ be its power set. Given any $A, B, C \in \mathcal{P}(X)$, let

$$ABC = (A \cap B) \cup (B \cap C) \cup (C \cap A) = (A \cup B) \cap (B \cup C) \cap (C \cup A).$$

Again, the axioms are readily verified.

These examples generalise in various directions. For example, simplicial trees generalise both to \mathbb{R} -trees and to CAT(0) cube complexes, both of which have played a major role in geometric group theory in recent years. In fact, we will see that “discrete” median algebras are essentially the same combinatorial structures as CAT(0) cube complexes, viewed in a slightly different way. Similarly, power sets generalise to distributive lattices, as well as to ternary boolean algebras. All of these have natural median algebra structures. These, and many other examples will be described in due course.

We note that one can give a number of equivalent definitions of a median algebra: see Theorems 3.2.2, 4.1.1 and 4.3.1.

A key notion is that of an “interval” in a median algebra. This is the set of points, x , lying between two other points, a and b : that is to say $abx = x$. It was shown by Sholander in the 1950s that one can equivalently define a median

algebra in terms of intervals. An interval has a natural intrinsic structure as a distributive lattice, which is where we will begin our account. Intervals naturally give rise to the notion of “convexity”, another central tool in the subject. Indeed “convex structures” can be studied in their own right, and median algebras are sometimes viewed in these terms. Convexity in turn gives rise to the notion of a “wall”: that is a partition of a median algebra into two non-empty convex subsets. From there, one can develop a theory of duality, analogous to that of Stone duality for boolean algebras. Another key fact is that a free median algebra on a finite set is finite. Hence the subalgebra of a median algebra generated by a finite subset is finite. This is exploited in many arguments.

We will give a general account of these notions the first ten sections. From there on, we mainly study particular cases, such as topological median algebras, median metric spaces, and discrete median algebras. In particular, discrete median algebras have other descriptions in terms of median graphs and CAT(0) cube complexes. These are major topics in their own right. Another particular class are “rank-1” median algebras. These can be alternatively formulated in terms of treelike structures. These overlap with median metric spaces in the theory of \mathbb{R} -trees, another much-used notion. We devote a chapter to looking at these from the perspective of median algebras.

The final chapters concern structures which are not in general median algebras, but which are related, or make use of the theory of median algebras. These include quasimedial graphs, coarse median spaces and injective metric spaces.

Most of the main text is focused on giving a logical development of the subject. This is generally intended to be self-contained, modulo fairly standard results from geometry, topology etc. The final chapters represent more of a survey. In order not to interrupt to flow too much, we have mostly limited the background material and references to particular facts which are needed for the development. More historical background, related discussion, and fuller references are given in the Notes to the different sections at the end (Section 26).

We have made an effort to keep different (sub)sections as logically independent as possible. For the most part, they should be individually readable with only some basic definitions and standard facts — any further results needed are explicitly referenced.

Some common notions used throughout include: betweenness relations and intervals, homomorphisms, adjacency, subalgebras, convex sets, halfspaces, walls, (hyper)cubes (all defined in Subsection 3.2), direct products (Subsection 3.1), median identities and the means of verifying them (Subsection 6.2), parallel sets, gates and convex hulls (Subsections 7.2, 7.3 and 7.4) and rank (Subsection 8.2). Some key facts are: Theorem 3.2.2 (the short distributive law), Proposition 3.3.3 (the subalgebra generated by a finite set is finite), Theorem 4.1.1 (formulation of median algebras in terms of intervals), Lemma 7.1.1 (the Helly Property for convex sets), Theorem 8.1.2 (any two disjoint convex sets are separated by a wall) and Lemma 8.2.1 (formulation of rank in terms of walls).

I thank Victor Chepoi, Elia Fioravanti, Anthony Genevois and Graham Niblo for their interest and comments on earlier drafts of this paper.

1.1. Conventions.

By default, our discussion is conducted in the context of Zermelo-Fraenkel set theory (ZF), though many of the formal arguments involving the median operation could be carried out in a much more restricted logical framework.

We will also by default be assuming the Axiom of Choice (ZFC), though for much of the discussion, this is not needed. We will generally make it clear when it is required. Where it is not needed, “finite” should be interpreted as being in bijective correspondence with a proper initial segment of the natural numbers.

We will use $\#A$ to denote the cardinality of a set (usually finite).

In the text and formulae, we often use “&” to mean “and”. (The more conventional symbol “ \wedge ” will be used to denote the meet operation in a lattice.) As usual, “ \neg ” denotes “not”.

We write “ $A := B$ ” or “ $B =: A$ ” to mean that A defined by equating it with B .

“LHS” and “RHS” are abbreviations for the “left-hand side” and “right-hand side” of an expression.

The median of a, b, c is generally denoted by abc . In a few places, where there might be some ambiguity (for example, with multiplication in a ring) we denote it instead by $\mu(a, b, c)$.

We use $\mathcal{P}(X)$ to denote the power set of a set X . If $X \neq \emptyset$, the “proper power set” is defined to be $\mathcal{P}_0(X) = \mathcal{P}(X) \setminus \{\emptyset, X\}$. By default, $\mathcal{P}(X)$ will be taken with its structure as a boolean algebra or median algebra, whereas $\mathcal{P}_0(X)$ will be considered with its structure as a “proset” (to be defined in Section 9).

We will use the term “hypercube” to mean a median algebra isomorphic to $\mathcal{P}(X) \cong \{0, 1\}^X$ for some (possibly infinite) set X . A “cube” is (isomorphic to) the subalgebra of finite subsets of X . Clearly, if X is finite, these notions coincide. In this case, we will refer to $[0, 1]^X$ as a “real cube”. Viewed as a polyhedron, its 1-skeleton will be called a “cubical graph”. Note that its 0-skeleton is a finite cube in the above sense. Cubes (or real cubes) will be referred to as “cells” when they are the building blocks of a polyhedral or CW-complex.

We will sometimes use the term “non-finite metric”. This is a metric, except that we are allowing it to take the value ∞ . In particular, it satisfies the triangle inequalities, with the obvious convention that $x + \infty = \infty + \infty = \infty$ for all $x \in [0, \infty)$. Such a non-finite metric induces a topology in the usual way.

2. DISTRIBUTIVE LATTICES

Distributive lattices arise as “intervals” in median algebras. They also provide a rich source of examples. A great deal has been written about median algebras from this perspective, so we will only touch on the subject here.

Here we will describe how the median can be defined in a distributive lattice, and give a brief account free distributive lattices. We give some examples of lattices at

the end of the section. (A simple one to keep in mind for the moment is that of a power set, equipped with the operations of union and intersection.)

2.1. Lattices and medians.

First, we recall some standard definitions.

Definition. A *semilattice* is a set, L , equipped with a binary operation, $[(x, y) \mapsto x \wedge y]$, satisfying $x \wedge x = x$, $x \wedge y = y \wedge x$ and $(x \wedge y) \wedge z = x \wedge (y \wedge z)$ for all $x, y, z \in L$.

Definition. A *lattice* is a set, L , equipped with two binary operations, \wedge and \vee , such that (L, \wedge) and (L, \vee) are both semilattices, and which satisfies the *absorption laws*, that is

$$x \wedge (x \vee y) = x \vee (x \wedge y) = x$$

for all $x, y \in L$.

Here \wedge and \vee are commonly referred to as the *meet* and *join* operations. In view of the associative rules, we can abbreviate $(x \wedge y) \wedge z$ to $x \wedge y \wedge z$ etc.

Given $x, y \in L$, we write $x \leq y$ to mean that $x \wedge y = x$. Note that in this case $x \vee y = (x \wedge y) \vee y = y$, and so we see that this is equivalent to saying that $x \vee y = y$.

If $x \leq y$ and $y \leq x$, then $x = x \wedge y = y$. Moreover, if $x \leq y \leq z$, then $x \wedge z = (x \wedge y) \wedge z = x \wedge (y \wedge z) = x \wedge y = x$, and so $x \leq z$. In other words, we see that \leq is a partial order on L . We write $x < y$ to mean $x \leq y$ and $x \neq y$. We call this a “strict” partial order.

Definition. A lattice, L , is *distributive* if we have

$$(x \wedge y) \vee z = (x \vee z) \wedge (y \vee z)$$

and

$$(x \vee y) \wedge z = (x \wedge z) \vee (y \wedge z)$$

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

A lattice is said to be *bounded* if there exist $a, b \in L$ such that for all $x \in L$, $a \wedge x = a$ and $b \vee x = b$. In this case, a and b are unique, and from the absorption laws, we see that $a \vee x = b \wedge x = x$ for all $x \in L$. We will refer to a, b respectively as the “minimum” and “maximum” elements. (In this context, they are sometimes referred to as “bottom” and “top” elements.) Although most of the distributive lattices we will be considering are bounded, we will not make that assumption in this section.

Let L be a distributive lattice.

The following is a key observation:

Lemma 2.1.1. *For all $x, y, z \in L$ we have*

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

Proof.

$$\begin{aligned}
 (x \wedge y) \vee (y \wedge z) \vee (z \wedge x) &= (x \wedge y) \vee (z \wedge (x \vee y)) \\
 &= ((x \wedge y) \vee z) \wedge ((x \wedge y) \vee (x \vee y)) \\
 &= (x \vee z) \wedge (y \vee z) \wedge (x \vee y).
 \end{aligned}$$

□

We will write

$$xyz := (x \wedge y) \vee (y \wedge z) \vee (z \wedge x).$$

Thus $[(x, y, z) \mapsto xyz]$ is a ternary operation on L . It is symmetric in x, y, z . Moreover, $xxxy = x$ for all $x, y \in L$.

For future reference, we also note:

Lemma 2.1.2. *If $x, y, z, w \in L$, then*

$$(xyw)zw = (x \wedge y \wedge z) \vee ((x \vee y \vee z) \wedge w).$$

Proof.

$$\begin{aligned}
 &(xyw)zw = \\
 &(((x \wedge y) \vee (y \wedge w) \vee (w \wedge x)) \vee z) \wedge (((x \wedge y) \vee (y \wedge w) \vee (w \wedge x)) \vee w) \wedge (z \vee w) \\
 &= ((x \wedge y) \vee (y \wedge w) \vee (w \wedge x) \vee (z \wedge w)) \wedge (z \vee w).
 \end{aligned}$$

Now:

$$\begin{aligned}
 (x \wedge y) \wedge (z \vee w) &= (x \wedge y \wedge z) \vee (x \wedge y \wedge w), \\
 (y \wedge w) \wedge (z \vee w) &= (y \wedge w \wedge z) \vee (y \wedge w) = y \wedge w, \\
 (w \wedge x) \wedge (z \vee w) &= (w \wedge x \wedge y) \vee (x \wedge w) = x \wedge w, \\
 (z \wedge w) \wedge (z \vee w) &= z \wedge w.
 \end{aligned}$$

Taking the join of these four elements and applying the distributive law, we get

$$\begin{aligned}
 (xyw)zw &= (x \wedge y \wedge z) \vee (x \wedge w) \vee (y \wedge w) \vee (z \wedge w) \\
 &= (x \wedge y \wedge z) \vee ((x \vee y \vee z) \wedge w).
 \end{aligned}$$

□

In particular, we see that $(xyw)zw$ is symmetric in x, y, z . We write

$$(xyz|w) := (xyw)zw.$$

We write $x.z.y$ to mean that $z = xzy$. We note:

Lemma 2.1.3. $x.z.y \Leftrightarrow x \wedge y \leq z \leq x \vee y$.

Proof. Suppose $x.z.y$. That is $(x \wedge y) \vee (y \wedge z) \vee (z \wedge x) = z$, so $z \wedge (x \vee y) = z$, that is $x \wedge y \leq z$. Similarly (swapping \wedge and \vee) we have $z \leq x \vee y$.

Conversely, suppose $x \wedge y \leq z \leq x \vee y$. Then,

$$xyz = (x \wedge y) \vee (y \wedge z) \vee (z \wedge x) = (x \wedge y) \vee (z \wedge (x \vee y)) = (x \wedge y) \vee z = z.$$

□

2.2. Sublattices and Sperner families.

Suppose that $A \subseteq L$. We write $\langle A \rangle$ for the sublattice generated by A . It is obtained by repeatedly applying meets and joins, starting with elements of A . Using the distributive laws, it is easily seen that any element, $x \in \langle A \rangle$ can be put into a “normal form”. Namely, we can write

$$x = y_1 \vee y_2 \vee \cdots \vee y_p,$$

where

$$y_i = a_{i1} \wedge a_{i2} \wedge \cdots \wedge a_{iq_i},$$

for $a_{ij} \in A$, where $p, q_i \in \mathbb{N}$. Moreover, we can assume that the y_i are all distinct, and that for any fixed i , the a_{ij} are all distinct. Therefore, if $\#A \leq n < \infty$, then there are at most 2^n possibilities for each y_i , hence at most 2^{2^n} possibilities for x . We deduce:

Lemma 2.2.1. *If $A \subseteq L$ with $\#A \leq n < \infty$, then $\#\langle A \rangle \leq 2^{2^n}$.*

We can elaborate on this a little. Suppose A is finite. Let $\mathcal{I}(A) = \mathcal{P}(A) \setminus \{\emptyset\}$ be the set of all non-empty subsets of A . Given $I \in \mathcal{I}(A)$, let $a_I = \bigwedge_{a \in I} a$. If $\mathcal{J} \subseteq \mathcal{I}(A)$ is non-empty, set $a(\mathcal{J}) = \bigvee_{I \in \mathcal{J}} a_I$. We can suppose that no element of \mathcal{J} is properly contained in another, since removing it would not change $a(\mathcal{J})$. Such a family is called a **Sperner family**. (In other words, it is an antichain with respect to inclusion.) Every element of $\langle A \rangle$ is thus of the form $a(\mathcal{J})$ for some Sperner family, \mathcal{J} .

We can use this idea to construct the free distributive on an abstract finite set X . We can define this formally as follows. Let $D(X)$ be the set of all Sperner families on X , thought of as a subset of the double power set, $\mathcal{P}(\mathcal{P}(X))$. Given $\mathcal{J} \in D(X)$, let $\mathcal{J}^\uparrow \subseteq \mathcal{P}(X)$ be the set of all $J \subseteq X$ such that $J \supseteq I$ for some $I \in \mathcal{J}$. The map $[\mathcal{J} \mapsto \mathcal{J}^\uparrow]$ gives a bijection from $D(X)$ to the set, $U(X) \subseteq \mathcal{P}(\mathcal{P}(X))$, of all subsets of $\mathcal{P}(X)$ which are closed under \supseteq . (We can recover \mathcal{J} from \mathcal{J}^\uparrow as the set of all \subseteq -minimal elements of \mathcal{J}^\uparrow .) Note that $U(X)$ is closed under union and intersection in $\mathcal{P}(\mathcal{P}(X))$, in other words, it is a sublattice. This induces a distributive lattice structure on $D(X)$ where $(\mathcal{J}_1 \wedge \mathcal{J}_2)^\uparrow = \mathcal{J}_1^\uparrow \cap \mathcal{J}_2^\uparrow$ and $(\mathcal{J}_1 \vee \mathcal{J}_2)^\uparrow = \mathcal{J}_1^\uparrow \cup \mathcal{J}_2^\uparrow$.

It is perhaps simpler to think of an element of $D(X)$ a formal expression of the form $a(\mathcal{J})$ for a Sperner family \mathcal{J} (of course, modulo permuting the entries of the expression). One can define the meet and join of such expressions in the obvious way, putting the result into normal form by formally applying the distributive laws. It is easy to see that $D(X)$ is indeed the free distributive lattice on X in the usual sense (i.e. any map of X into any distributive lattice extends uniquely to a lattice homomorphism).

If we have a map $X \rightarrow \{0, 1\}$, we can substitute the arguments of $a(\mathcal{J})$ accordingly, and evaluate in the lattice $\{0, 1\}$ to give an element of $\{0, 1\}$. (Here $0 \wedge 1 = 0$ and $0 \vee 1 = 1$: see Example (Ex2.2) below.)

Lemma 2.2.2. *Suppose $\mathcal{J}, \mathcal{J}'$ are Sperner families, and that $a(\mathcal{J})$ and $a(\mathcal{J}')$ evaluate to equal elements of $\{0, 1\}$ under any substitution $X \rightarrow \{0, 1\}$. Then $\mathcal{J} = \mathcal{J}'$.*

Proof. Let $I \in \mathcal{J}$. Set $a = 1$ for all $a \in I$, and $a = 0$ for all $a \in X \setminus I$. In this way, $a(\mathcal{J})$ evaluates to 1. So therefore must $a(\mathcal{J}')$. This means that there is some $I' \in \mathcal{J}'$ with $I' \subseteq I$. Conversely, there is some $I'' \in \mathcal{J}'$ with $I'' \subseteq I'$. Since \mathcal{J} is a Sperner family, we have $I = I' = I''$. This shows that $\mathcal{J} \subseteq \mathcal{J}'$. Conversely, $\mathcal{J}' \subseteq \mathcal{J}$, so $\mathcal{J} = \mathcal{J}'$ as required. \square

This gives a means of checking that two formal expressions involving \wedge and \vee , with arguments in X , represent the same element of the free distributive lattice: we check that they always evaluate to the same element of $\{0, 1\}$. (We will see an analogous principle for median algebras in Subsection 6.2.)

The following is generally referred to as ‘‘Sperner’s Lemma’’:

Lemma 2.2.3. *The cardinality of a Sperner family on a set with n elements is at most $\binom{n}{\lfloor n/2 \rfloor}$.*

This would allow us to improve on the above bounds somewhat. We will return to it in Section 11 (see Proposition 11.9.1).

Free distributive lattices are not essential to most of our discussion, but we will return to them briefly in Subsection 6.3.

2.3. Examples.

Here are a few examples of distributive lattices.

(Ex2.1): The empty set. Also, any one-point set.

(Ex2.2): The two-point set, $\{0, 1\}$, is a distributive lattice, where $0 \wedge 1 = 0$ and $0 \vee 1 = 1$.

Here xyz represents ‘‘majority vote’’, that is $000 = 001 = 0$ and $011 = 111 = 1$ etc.

If we think of x, y, z, w all casting votes and taking as outcome $(xyz|w)$, then the vote of w can only be overruled by the unanimous vote of x, y and z (as can be seen from the formula given by Lemma 2.1.2).

(Ex2.3): More generally, if (I, \leq) is a totally ordered set, then by defining $x \wedge y = \min(x, y)$ and $x \vee y = \max(x, y)$, (I, \wedge, \vee) becomes a distributive lattice. Here, xyz is the point ‘‘in the middle’’ of $\{x, y, z\}$.

One obvious example is the real line, \mathbb{R} , or any subset thereof.

We could also take the long line, and, if we want, adjoin a maximum and a minimum to it. We could also take any uncountable ordinal. Note that any successor ordinal is a bounded distributive lattice. These last serve as cautionary examples. The total orders of genuine interest to us here can all be embedded in \mathbb{R} .

(Ex2.4): The direct product of any family of distributive lattices is a distributive lattice, with meet and joins defined independently on each co-ordinate.

(Ex2.5): In particular, if X is any set, then the power, $\{0, 1\}^X$, is a distributive lattice. We can naturally identify $\{0, 1\}^X$ with the power set, $\mathcal{P}(X)$, so that $A \subseteq X$ gets identified with its characteristic function (i.e. taking the value 1 on elements of A). Under this identification, we have $A \wedge B = A \cap B$ and $A \vee B = A \cup B$. The median of three sets, A, B, C , consists of the set, ABC , of elements of X which lie in at least two of A, B, C . (This is not to be confused with the median defined by Example (Ex3.4) of Subsection 3.4.)

(Ex2.6): In the context of median algebras (to be defined in the next section), $\{0, 1\}^X$, is generally referred to as a “cube”. It naturally embeds in the “real cube” $[0, 1]^X$.

Example (Ex2.5) above is a boolean algebra.

A **boolean algebra** is a bounded distributive lattice, B , equipped with an involution $[x \mapsto x^*] : B \rightarrow B$, and minimum and maximum elements, 0 and $1 = 0^*$, respectively, such that $x \wedge x^* = 0$ and $x \vee x^* = 1$ for all $x \in B$. One can verify that $(x \wedge y)^* = x^* \vee y^*$ and $(x \vee y)^* = x^* \wedge y^*$ for all $x, y \in B$. (These identities are referred to as the “de Morgan Laws”.) Thus, for example, $\mathcal{P}(X)$ is a boolean algebra, where the involution is defined by taking complements in X . We will discuss boolean algebras further in Subsections 3.4 and 9.6.

3. BASIC FACTS ABOUT MEDIAN ALGEBRAS

The main aim of this section is to get the subject of median algebras “off the ground”. Most of the notions we introduce here will be revisited in more detail later. We define various basic notions: intervals, convexity, gates, etc. Two key facts shown here are that any median algebra embeds in a cube (Proposition 3.2.13) and that a subalgebra generated by a finite set is finite (Proposition 3.3.3). The latter entails a description of free median algebras (Proposition 3.3.1). We give some examples of median algebras at the end of the section.

3.1. Definitions.

We begin with a formal axiomatic definition, though this is not especially intuitive and difficult to apply directly. A more intuitive formulation will appear in Subsection 4.1. One can also formulate the notion in terms of cube complexes as we discuss Section 17.

Definition. A **median algebra** is a set, M , equipped with a symmetric ternary operation, $[(x, y, z) \mapsto xyz]$, such that for all $a, b, c, d \in M$ we have:

(M1): $aab = a$, and

(M2): $(abd)cd = (acd)bd$.

Recall that an operation is *symmetric* if it is invariant under permutation of its arguments. In other words, we are assuming that $abc = bac = bca$.

Writing $(abc|d) := (abd)cd$, Axiom (M2) tells us that this expression symmetric in a, b, c . (This is sometimes referred to as the *associative law* of a median algebra.)

For the ternary operation defined on a distributive lattice, this property is an immediate consequence of Lemma 2.1.2. Moreover, (M1) follows directly from the absorption rules. We therefore see:

Lemma 3.1.1. *A distributive lattice is a median algebra, where the median is defined by setting $abc := (a \vee b) \wedge (b \vee c) \wedge (c \vee a)$.*

This immediately gives us a rich source of examples. (In fact, any median algebra embeds into a distributive lattice — see Proposition 3.2.13 — though it is not necessarily natural to view it in this way.)

For more examples, note that any *direct product* of a family of median algebras is a median algebra with the median defined independently on each coordinate. (These may be infinite direct products.)

Another class of examples are trees of various sorts (for example, simplicial trees and \mathbb{R} -trees as we discuss later: see Subsection 14.2 and Section 15). Another important source are CAT(0) cube complexes. Median metric spaces, which we discuss in Section 13 can be viewed as a generalisation of both. Some further examples will be discussed in Subsection 3.4.

3.2. Some fundamental constructs.

Let us now proceed with the general theory.

Definition. A *subalgebra* of a median algebra is a subset $N \subseteq M$, which is closed under the ternary operation. We will write $N \leq M$.

Clearly N is intrinsically a median algebra.

Definition. The subalgebra *generated* by a subset $A \subseteq M$, is the smallest subalgebra containing A . We denote it by $\langle A \rangle$.

We see that $\langle A \rangle$ can be constructed by starting with elements of A , and repeatedly applying the median operation. In this way, each element of $\langle A \rangle$ can be written as a *median expression* with arguments in A . We will discuss such expressions more formally in Section 6. (In some cases one can put a bound on the complexity of such an expression — see for example Proposition 8.2.4 — though in general this is not possible.)

A map $\phi : M \rightarrow N$, between two median algebras M, N , is a *homomorphism* if it respects the ternary operation: that is $\phi(xyz) = \phi(x)\phi(y)\phi(z)$ for all $x, y, z \in M$. (Note this is equivalent to saying that its graph is a subalgebra of the direct product $M \times N$.) A homomorphism a *monomorphism* if it is injective, an *epimorphism* if it surjective, and an *isomorphism* if it is both.

Let M be a median algebra.

Suppose we fix some $a \in M$. We define a symmetric “meet” operation, \wedge , on M by setting $x \wedge y = axy$ for $x, y \in M$. Now $x \wedge x = x$ (by (M1)), and by (M2), we have that $(x \wedge y) \wedge z = (xya)za = (xyz|a)$ is symmetric in x, y, z . It follows that (M, \wedge) is a semilattice.

Given $a, b, c \in M$, we write $a.c.b$ to mean that $abc = c$. We say that c lies **between** a and b . We write

$$[a, b] := \{x \in M \mid a.x.b\}.$$

Definition. $[a, b]$ is the (median) **interval** from a to b .

We will write it as $[a, b]_M$, if there is any ambiguity regarding M . Note that if $N \leq M$ is a subalgebra, then $[a, b]_N = N \cap [a, b]_M$.

If $x \in M$, then $(abx)ab = (aax|b) = (aab)xb = axb = abx$. Thus, $a.abx.b$. This gives us an equivalent way of defining an interval as

$$[a, b] = \{abx \mid x \in M\}.$$

Given $x, y \in [a, b]$, we have already defined $x \wedge y := axy$. Note that $ab(axy) = (bxy|a) = ax(aby) = axy$, so $x \wedge y \in [a, b]$. Thus, $([a, b], \wedge)$ is a semilattice. We similarly define $a \vee b := bxy$, so that $([a, b], \vee)$ is also a semilattice. Moreover, $x \wedge (x \vee y) = ax(bxy) = (aby|x) = (abx)yx = xyx = x$. Similarly, $x \vee (x \wedge y) = x$. Therefore $([a, b], \wedge, \vee)$ is a lattice. It is also bounded: a and b are the minimum and maximum respectively. In fact, writing $x \leq y$ to mean $x \wedge y = x$, we see that $x \vee y = bxy = b(axy)y = (abx|y) = x(aby)y = xyx = y$. It follows that the statement $x \leq y$ is equivalent to $x \vee y = y$, or to $a.x.y$, or to $x.y.b$.

In general, we will use the notation $x_1.x_2.\dots.x_n$ to mean that $x_i.x_j.x_k$ holds whenever $i \leq j \leq k$. From the fact that \leq is a partial order, we have the general rule:

$$a.b.d \ \& \ b.c.d \Rightarrow a.b.c.d.$$

Indeed we can continue interpolating in this manner. For example, if in addition we have $b.e.c$, then $a.b.e.c.d$. We will use this principle many times in what follows, often without further comment. We refer to it as the (linear) **interpolation rule**.

Another important observation is what we will call the **median rule**, namely:

$$a.d.b \ \& \ b.d.c \ \& \ c.d.a \Rightarrow d = abc.$$

To see this, set $m = abc$. Now $mda = (abc)da = (abd)ca = dca = d$. Thus $m.d.a$. Similarly $m.d.b$. Now $d = dbm = (adm)bm = (abm)dm = mdm = m$. That is, $d = abc$ as claimed.

Putting this together with the fact that $abc \in [a, b]$ etc., we get:

Lemma 3.2.1. *If $a, b, c \in M$, then $[a, b] \cap [b, c] \cap [c, a] = \{abc\}$.*

Thus, the median operation is determined by the ternary betweenness relation.

As an immediate consequence we see that a map, ϕ , between median algebras is a homomorphism if and only if it satisfies $a.b.c \Rightarrow \phi a.\phi b.\phi c$.

To proceed, we would like to say that $[a, b]$ is a subalgebra and a distributive lattice (and lots more). This is indeed true, but we would be struggling to make much further progress applying axioms (M1) and (M2) directly. The following identity (sometimes referred to as the **short distributive law**) is much more practical.

Given any $a, b, c, d, e \in M$ consider the identity:

$$(M3): ab(cde) = (abc)(abd)e.$$

Note that, given (M1), this immediately implies (M2) (since $ab(ade) = (aba)(abd)e = ae(abd)$). In turns out that the converse is also true:

Theorem 3.2.2. *Any median algebra satisfies (M3) for all $a, b, c, d, e \in M$.*

In fact, this is surprisingly tricky to verify. The proof is not especially enlightening, and fits more naturally into the discussion in Section 4, so we postpone it until then. In the meantime, we could simply substitute (M3) for Axiom (M2) in the definition of a median algebra, and proceed on that basis. (In fact, the axioms of a median algebra are frequently given as (M1) and (M3).) We can now continue with our discussion of intervals.

Suppose $a, b, z \in M$ and $x, y \in [a, b]$. Then $ab(xyz) = (abx)(aby)z = xyz$, so $xyz \in [a, b]$. We immediately see:

Lemma 3.2.3. *Let $a, b \in M$ and $c, d \in [a, b]$. Then $[c, d] \subseteq [a, b]$.*

We also see:

Lemma 3.2.4. *$[a, b]$ is a subalgebra of M , and intrinsically a distributive lattice.*

Proof. The fact that it is a subalgebra is an immediate consequence of Lemma 3.2.3. We have already verified that it is a lattice. We need to check the distributive laws. Now $(x \wedge y) \vee z = (xya)bz = (bxz)(bzy)a = (x \vee z) \wedge (y \vee z)$. Similarly $(x \vee y) \wedge z = (x \wedge z) \vee (y \wedge z)$. \square

We also note that we can recover the median in $[a, b]$ from the meet and join operations.

Lemma 3.2.5. *If $x, y, z \in [a, b]$, then $xyz = (x \wedge y) \vee (y \wedge z) \vee (z \wedge x)$.*

Proof.

$$\begin{aligned} (x \wedge y) \vee (y \wedge z) \vee (z \wedge x) &= (x \wedge (y \vee z)) \vee (z \wedge y) = (ax(byz))(ayz)b \\ &= (a(ayz)b)(x(ayz)b)(byz) = (a(ayb)z)((ayz)bx)(byz) \\ &= (ayz)((ayz)bx)(byz) = (bx(byz)|(ayz)) = (b(byz)(ayz))x(ayz) \\ &= ((bba)yz)x(ayz) = (byz)(ayz)x = (bax)yz = xyz. \end{aligned}$$

\square

In particular, we see that a subalgebra of $[a, x]$ is the same as a sublattice (that is, closed under \wedge and \vee). Therefore, if $A \subseteq [a, b]$, then the median algebra, $\langle A \rangle$ generated by A also the sublattice generated by A . Thus, by Lemma 2.2.1, if A is finite, then so is $\langle A \rangle$. One of the main remaining aims of this section is to generalise this to all finite subsets of M : see Proposition 3.3.3 (as well as Subsection 6.3 for a constructive proof thereof).

Next, we note the following identity (sometimes referred to as the *long distributive law*):

Lemma 3.2.6. *If $a, b, x, y, z \in M$, then $(abx)(aby)(abz) = ab(xyz)$.*

Proof. Applying (M3) three times, and using $a.abx.b$, we get:

$$(abx)(aby)(abz) = ab((abx)yz) = (ab(abx))(aby)z = (abx)(aby)z = ab(xyz).$$

□

This tells us that the *gate map*, $\omega : M \rightarrow [a, b]$, defined by setting $\omega(x) = abx$, is a homomorphism. (This an instance of a more general notion of “gate map” defined in Subsection 7.3, and a yet more general notion defined in Section 22.)

We proceed with a few more general observations, needed shortly.

Lemma 3.2.7. *$a.b.c$ & $a.b.d$ & $c.e.d \Rightarrow a.b.e$.*

Proof. $abe = ab(cde) = (abc)(abd)e = bbe = b$.

□

Lemma 3.2.8.

$$a_1.c_1.b_1 \text{ \& } a_2.c_2.b_2 \text{ \& } c_1.c.c_2 \text{ \& } a_1.a.c \text{ \& } a_2.a.c \text{ \& } b_1.b.c \text{ \& } b_2.b.c \\ \Rightarrow a.c.b.$$

Proof. Let $m = abc$. We want to show that $m = c$. We have $a.m.c$ & $b.m.c$. Now $a_1.a.c$ & $a.m.c$ gives $a_1.m.c$ by the linear interpolation rule. Similarly we have $a_2.m.c$, $b_1.m.c$ and $b_2.m.c$. Now $a_1.m.c$ & $b_1.m.c$ & $a_1.c_1.b_1$ gives $c_1.m.c$ by Lemma 3.2.7. Similarly $c_2.m.c$. Now $c_1.c.c_2$, and so $c_1.m.c.m.c_2$. In particular, $m.c.m$ so $m = c$. □

We next give a brief discussion of convexity. This is a key notion in the subject, and will be the topic of Section 7. We restrict the present discussion to more immediate requirements.

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

Note that by Lemma 3.2.3, any interval $[a, b]$ is itself convex. Also any intersection of convex sets is convex, and any increasing union of convex sets is convex.

Given $A, B \subseteq M$, write

$$J(A, B) = \bigcup \{[a, b] \mid a \in A, b \in B\}.$$

This is often called the *join* of A and B (of course not be confused with the join operation in a lattice).

Lemma 3.2.9. *If $A, B \subseteq M$ are convex, so is $J(A, B)$.*

Proof. Let $c_1, c_2 \in J(A, B)$, and let $c \in [c_1, c_2]$. Now $c_i \in [a_i, b_i]$ for some $a_i \in A$ and $b_i \in B$. Let $a = a_1 a_2 c \in A$ and $b = b_1 b_2 c \in B$. By Lemma 3.2.8, $c \in [a, b] \subseteq J(A, B)$. \square

Definition. A *halfspace* is a subset, $H \subseteq M$, such that H and $M \setminus H$ are both non-empty and convex.

Proposition 3.2.10. *Let $B \subseteq M$ be convex, and let $a \in M \setminus B$. Then there is a halfspace, $H \subseteq M$, with $B \subseteq H$ and $a \notin H$.*

Proof. Consider the set, \mathcal{C} , of all convex sets containing B but not a . This is partially ordered by inclusion. As noted above, any chain in \mathcal{C} has an upper bound, namely its union. Therefore by Zorn's lemma, \mathcal{C} has a maximal element, H . We claim that H is a halfspace, i.e. $M \setminus H$ is convex.

Suppose, for contradiction, that $c_1, c_2 \in M \setminus H$ and $c \in [c_1, c_2] \cap H$. By Lemma 3.2.9, $J(H, \{c_i\})$ is convex, and so $a \in J(H, \{c_i\})$. In other words, $a \in [c_i, h_i]$ for some $h_i \in H$. Let $h = h_1 h_2 a \in H$. Now $c_i.a.h_i$ & $a.h.h_i$ gives $c_i.a.h$. Then $c_1.a.h$ & $c_2.a.h$ & $c_1.c.c_2$ gives $c.a.h$ by Lemma 3.2.7. But $c, h \in H$, so $a \in H$, contradicting $H \in \mathcal{C}$. \square

In other words, this shows that any convex set is an intersection of halfspaces. For the moment, we will only need Proposition 3.2.10 when B is a singleton. A stronger version of Proposition 3.2.10 will be given in Section 8: see Theorem 8.1.2.

(Note that Proposition 3.2.10 requires the Axiom of Choice, and so, a-priori, do some of the consequences given below. However, this can often be bypassed as we will mention later in this section.)

Definition. A *wall* of M is a (by default, unordered) partition of M into two non-empty convex sets, $M = A \sqcup B$.

Note that A, B are by definition halfspaces.

Definition. We say that a wall $\{A, B\}$ *separates* $a, b \in M$ if $(a \in A$ and $b \in B)$ or $(a \in B$ and $b \in A)$.

In these terms, an immediate consequence of Proposition 3.2.10 is:

Proposition 3.2.11. *Any two distinct points of M are separated by a wall.*

Definition. A *wall map* on M is an epimorphism $\phi : M \rightarrow \{0, 1\}$. We say that ϕ *separates* $a, b \in M$ if $\phi(a) \neq \phi(b)$.

Thus, if $\{A, B\}$ is a wall, then the map $\phi : M \rightarrow \{0, 1\}$ defined by $\phi(x) = 0$ if $x \in A$ and $\phi(x) = 1$ if $x \in B$, is a wall map. Conversely, a wall map gives rise to a wall. Of course, we could substitute $\{0, 1\}$ with any two-point median algebra.

There is a particular case where one can give a much more direct proof of Proposition 3.2.11, without recourse to the Axiom of Choice.

Definition. Two points, $a, b \in M$ are **adjacent** if $a \neq b$ and $[a, b] = \{a, b\}$.

In this case, the gate map $\omega : M \rightarrow \{a, b\}$, is a wall map separating a and b . In particular, we give direct proof of the following:

Lemma 3.2.12. *Any two distinct points in a finite median algebra are separated by a wall.*

Proof. Let Π be a finite median algebra, and let $a, b \in \Pi$ be distinct. Choose $c \in [a, b] \setminus \{a\}$ minimal with respect to the partial order \leq defined above. By Lemma 3.2.3, $[a, c] \subseteq [a, b]$. Also $[a, c] = \{a, c\}$, for if $d \in [a, c]$ then $c \wedge d \leq c$, and so by minimality, either $d = a$ or $d = c$. In other words a, c are adjacent in Π . Since $abc = c$, the gate map to $\{a, c\}$ separates a, b . \square

Note this argument only requires intervals to be finite, so it also applies to “discrete” median algebras to be discussed in Section 11.

We note that the above statements can be expressed in terms of embeddings into hypercubes.

Definition. A **hypercube** is a median algebra isomorphic to $\mathcal{P}(X) \cong \{0, 1\}^X$ for some set X . If X is finite, we will usually use the term **cube** for this notion. It is an n -**cube** if it has rank n for $n \in \mathbb{N}$.

(A notion of “infinite cube” will be defined in Subsection 11.11.)

Let $\mathcal{W}(M)$ be the set of all walls of M , which we can identify with the set of wall maps. Let $\Psi(M)$ be the hypercube $\{0, 1\}^{\mathcal{W}}$. Let $\iota : M \rightarrow \Psi(M)$ be the evaluation map: that is, $\iota(x)(\theta) = \theta(x)$. This is a median homomorphism (since all our maps here are homomorphisms). Moreover, by Proposition 3.2.11, ι is injective. Note also that if M is finite, so is $\Psi(M)$. We deduce:

Proposition 3.2.13. *Any median algebra embeds in a hypercube. Any finite median algebra embeds in a finite cube.*

In fact, for the latter statement, we only require that $\mathcal{W}(M)$ is finite. As an immediate corollary we have:

Corollary 3.2.14. *A median algebra with finitely many walls is finite.*

In fact, the argument shows that $\#M \leq 2^{\#\mathcal{W}(M)}$.

Remark. We note that this gives rise to another way of characterising a median algebra: it is a set equipped with a ternary “median” operation which embeds via a median monomorphism into $\mathcal{P}(X)$ for some set X . The median axioms are then consequences of the corresponding statements in $\mathcal{P}(X)$. Indeed this serves as a general means of verifying median identities. We return to this observation in Section 6.

This leads on to a discussion of free median algebras.

3.3. Free median algebras.

This is the usual notion of “free” with regard to an algebraic structure. A key fact is that the free median algebra on a finite set is finite. The description we give in this subsection makes use of Zorn’s Lemma, via Proposition 3.2.10. A different, constructive proof, avoiding the Axiom of Choice, is given in Subsection 6.3.

Definition. Let F be a median algebra, and let $A \subseteq F$. We say that F is **free** on A if any map, $\phi : A \rightarrow M$, into any median algebra, M , has a unique extension to a homomorphism, $\hat{\phi} : F \rightarrow M$ (that is, with $\hat{\phi}|_A = \phi$).

Note that F is generated by A . (Since the extension of the inclusion $A \hookrightarrow \langle A \rangle$, postcomposed with the inclusion $\langle A \rangle \hookrightarrow F$, extends the inclusion $A \hookrightarrow F$; and so, by uniqueness of extensions this must be the identity on F .)

Note also that if F is free on A and F' is free on A' , then any bijection from A to A' induces a unique isomorphism from F to F' . It therefore makes sense to talk about “the” free median algebra on a set X — assuming that it exists. It is well defined up to isomorphism, and only depends on the cardinality, $\#X$. We shall denote it by $F(X)$.

Regarding its existence, we will give an explicit construction of $F(X)$ as follows.

First, we make some general observations.

Suppose $E = E(\alpha_1, \dots, \alpha_p)$ is a formal expression which repeatedly applies a ternary relation in some formal alphabet which includes the letters $\alpha_1, \dots, \alpha_p$. The entries $\alpha_1, \dots, \alpha_p$ are called the “arguments” of the expression. (For example, $E(\alpha_1, \dots, \alpha_5)$ might be any expression like $(\alpha_1\alpha_2(\alpha_2\alpha_3\alpha_4))(\alpha_1\alpha_3\alpha_5)(\alpha_2\alpha_4\alpha_5)$, with arguments $\alpha_1, \dots, \alpha_5$.) We allow $E(\alpha) = \alpha$ as an expression with a single argument, α . (We will give a more formal treatment of expressions in Section 6.) Suppose we have a set X and any map $f : X \rightarrow M$ into a median algebra M . Given $x_1, \dots, x_p \in X$, we can take the expression $E(fx_1, \dots, fx_p)$ and evaluate it in M . Note that if $g : M \rightarrow N$ is a homomorphism to another median algebra, N , then we have $E(gfx_1, \dots, gfx_p) = gE(fx_1, \dots, fx_p)$, viewed as an identity in N .

Let X be any set. Let $\mathcal{P} = \{0, 1\}^X$ and let $\Psi = \Psi(X) = \{0, 1\}^{\mathcal{P}}$. We give Ψ the product median structure as a hypercube. Let $\iota : X \rightarrow \Psi$ be the evaluation map. That is, given any map $f : X \rightarrow \{0, 1\}$, we set $\iota(x)(f) = f(x)$. Clearly, ι is injective. We write $\pi_f : \Psi \rightarrow \{0, 1\}$ for the projection map to the f -coordinate. By the definition of the median on Ψ , this is a median homomorphism. Moreover, if $x \in X$, then $\pi_f(\iota x) = f(x)$.

Let $F = F(X) = \langle \iota X \rangle \leq \Psi$: that is, the median algebra generated by the image of ι in Ψ . Now any element $a \in F$ can be written as a median expression with arguments in ιX . In other words, we have $a = E(\iota x_1, \dots, \iota x_p)$, for some $x_1, \dots, x_p \in X$, interpreted as an identity in Ψ . If $f : X \rightarrow \{0, 1\}$, then $\pi_f(\iota x_i) = f(x_i)$ for all i . Since π_f is a homomorphism, we get $\pi_f(a) = \pi_f E(\iota x_1, \dots, \iota x_p) = E(\pi_f \iota x_1, \dots, \pi_f \iota x_p) = E(fx_1, \dots, fx_p)$ (evaluating in $\{0, 1\}$).

Now suppose that $\phi : X \rightarrow M$ is any map to any median algebra, M . Given any $a \in F$ and an expression E representing a , as above, we set $m = E(\phi x_1, \dots, \phi x_p) \in$

M . In fact, we claim that m is independent of the expression we chose to represent a in Ψ . For suppose that $a = E'(\iota y_1, \dots, \iota y_q)$ were another such expression with $y_1, \dots, y_q \in X$. Let $m' = E'(\phi y_1, \dots, \phi y_q)$. If $m \neq m'$, let $\theta : M \rightarrow \{0, 1\}$ be a wall map separating m and m' , as given by Proposition 3.2.11. Let $f = \theta \circ \phi : X \rightarrow \{0, 1\}$. Now $\theta(m) = \theta(E(\phi x_1, \dots, \phi x_p)) = E(\theta \phi x_1, \dots, \theta \phi x_p) = \pi_f(\iota x_1, \dots, \iota x_p) = \pi_f(a)$. Similarly, $\theta(m') = \pi_f(a)$, giving the contradiction that $\theta(m) = \theta(m')$. This proves the claim.

We can therefore unambiguously write $\hat{\phi}(a) = m$. This gives us a map $\hat{\phi} : F \rightarrow M$.

If $x \in X$, then we can represent ιx by the trivial expression $E(\iota x) = \iota x$, and so $\hat{\phi}x = \phi x$. This shows that $\hat{\phi}|_X = \phi$.

If $a_1, a_2, a_3 \in F$, we represent each a_i by an expression E_i in elements of ιX . The concatenation, $(E_1)(E_2)(E_3)$, is then an expression representing the median $a_1 a_2 a_3$ in Ψ . Replacing each of the arguments ιx by ϕx in this expression and evaluating in M , we see that $\hat{\phi}(a_1 a_2 a_3) = \hat{\phi}(a_1) \hat{\phi}(a_2) \hat{\phi}(a_3)$. This shows that $\hat{\phi}$ is a homomorphism.

Note also that the construction of $\hat{\phi}$ is completely determined if we want it to be a homomorphism into M extending ϕ .

After identifying X with $\iota X \subseteq F$, we have shown:

Proposition 3.3.1. *Given any set X , the median algebra $F(X) \subseteq \Psi(X)$ constructed above is free on X .*

This also shows that the free median algebra on any finite set is finite. In fact:

Proposition 3.3.2. *If $\#X \leq n < \infty$, then $\#F(X) \leq 2^{2^n}$.*

Proof. $\#\Psi(X) \leq 2^{2^n}$. □

As an immediate corollary we get:

Proposition 3.3.3. *Let M be a median algebra, and let $A \subseteq M$ with $\#A \leq n < \infty$. Then $\#\langle A \rangle \leq 2^{2^n}$.*

Proof. The inclusion $A \hookrightarrow M$ extends to a homomorphism $F(A) \rightarrow M$. Its image is $\langle A \rangle$. □

Another proof of Proposition 3.3.3 will be given in Section 8: see Proposition 8.2.5 and subsequent discussion.

The proofs we have presented of these statements are fairly standard. However they may make use of the Axiom of Choice (via Proposition 3.2.10) and are non-constructive. An alternative constructive approach will be described in Subsection 6.3. (See the discussion following Lemma 6.3.3.) In particular, once we have Proposition 3.2.2, we can reduce Proposition 3.3.1 to the finite case, by restricting everything to the subalgebra generated by ϕX in M .

We will give alternative descriptions of the free median algebra in Sections 5 and 6.

3.4. Examples.

We discuss a few examples of median algebras, in addition to the distributive lattices mentioned in Subsection 2.3, and free median algebras mentioned above. These are not essential to the logical development of the next few sections, though most of these examples will appear again later.

(Ex3.1): Simplicial trees.

These were mentioned briefly in the Introduction. A “simplicial tree”, T , is a connected graph with no non-trivial embedded circuits. We write V for its vertex set. Any two vertices, $a, b \in V$, are connected by a unique arc in T , and we write $I(a, b)$ for the set of vertices in this arc. If $a, b, c \in V$ then $I(a, b) \cap I(b, c) \cap I(c, a)$ consists of a single vertex, denoted abc . It is readily checked that the map $[(a, b, c) \mapsto abc]$ is symmetric and satisfies axioms (M1) and (M2) above. Moreover, $I(a, b)$ is precisely the median interval $[a, b]$.

Simplicial trees have various generalisations, for example to \mathbb{R} -trees discussed in Section 15, and to CAT(0) cube complexes discussed in Section 17.

(Ex3.2): Simplex graphs.

Let G be a graph with vertex set $V = V(G)$. A **clique** of G is a complete subgraph. By a “simplex” we mean the vertex set of a finite clique. Let $S = S(G) \subseteq \mathcal{P}(V)$ be the set of all simplices. We can view S as an abstract simplicial complex. (We have included the empty set as the “(-1)-cell”.) One readily checks that S is a subalgebra of $\mathcal{P}(V)$, hence intrinsically a median algebra.

One can think of this geometrically as follows. For each element, $\beta \in S$, we take a real $(\#\beta)$ -cube, and glue these cubes together so that if one simplex is contained in a larger one then the corresponding cube is a face of the larger cube. This gives us a connected cube complex, with a central vertex (corresponding to the empty clique) whose link is the realisation of the simplicial complex. The 1-skeleton of the cube complex is the adjacency graph of the median algebra S . This is an example of a CCAT(0) cube complex: see Section 16.

This example will arise again later as a “star” in a discrete median algebras in Subsection 11.7. They also feature in relation to the Roller boundary (Subsection 11.12).

(Ex3.3): Linear median algebras.

We say that a median algebra, L , is **linear** if it arises from a total order on L , as described in Example (Ex2.3) of Subsection 2.3. (Recall that $a.b.c$ is equivalent to the disjunction $(a \leq b \leq c$ or $c \leq b \leq a)$.) Any subset of L is a subalgebra. Moreover, it is easily seen that a median algebra is linear if and only if every finite subalgebra is linear. Here is another description.

Suppose L is a median algebra such that $abc \in \{a, b, c\}$ for all $a, b, c \in M$. In other words, every subset of M is a subalgebra. Then M is either linear or a 2-cube (that is, isomorphic to $\{0, 1\}^2$).

To see this, we can suppose that L is finite. Choose $a, b \in L$ so as to maximise $\#[a, b]$. Then $L = [a, b]$, and so L is a distributive lattice with minimum a and maximum b . Suppose that the induced partial order on L not a total order. Choose $c, d \in L$ such that neither $c \leq d$ nor $d \leq c$ holds. Then $c \wedge d = a$ and $c \vee d = b$. In other words, $Q := \{a, c, b, d\}$ is a 2-cube.

We want to show that $L = Q$. To see this, suppose $e \in L \setminus Q$. We claim that we cannot have $a.e.c$: that is $e \leq c$. For if $b.e.d$, then $d \leq e \leq c$ giving a contradiction. If $b.d.e$, then $a = acd = ac(bde) = (acb)(acd)e = cae = e$ giving a contradiction. Similarly (swapping a with c and b with d) we get a contradiction to $d.b.e$. This shows that, as claimed, we cannot have $e \leq c$. Similarly (swapping the roles of a and b) we cannot have $c \leq e$. Therefore $c.a.e$ holds. Similarly (swapping the roles of c and d), we get $d.a.e$. In other words, we have $c \wedge e = d \wedge e = a = c \wedge d$, and so $cde = a \notin \{c, d, e\}$, giving a contradiction. This shows that $L = Q$, proving the claim.

(We remark that this argument could be simplified by considering the ‘‘adjacency graph’’, $\Gamma(\{a, b, c, d, e\})$, as defined in Subsection 5.1.)

(Ex3.4): The algebra of non-empty convex subsets of a median algebra.

Let M be a median algebra. Given subsets $A, B, C \subseteq M$, write $ABC = \{abc \mid a \in A, b \in B, c \in C\}$. Clearly we have $(ABD)CD = (ACD)BD$ for any subsets A, B, C, D . In other words, Axiom (M2) holds for this ternary operation. If $B \neq \emptyset$, then $A \subseteq AAB$. If A is convex, then $AAB \subseteq A$. One can also check that if A, B, C are convex, then so is ABC (see Lemma 7.1.2). This therefore defines a median algebra structure on the set, $\mathcal{K}(M)$, of all non-empty convex subsets of M . This of course is different from the median on an arbitrary power set, as defined in (Ex2.5) of Subsection 2.3. Note that the map $[x \mapsto \{x\}] : M \rightarrow \mathcal{K}(M)$ is a monomorphism to this median algebra. We will return to this example in Subsection 7.1. It makes a few further appearances later.

(Ex3.5): Boolean algebras.

Boolean algebras were briefly mentioned at the end of Subsection 2.3. They can be viewed as median algebras with some additional structure, which we go on to describe.

To begin, recall that a boolean algebra is essentially the same structure as a boolean ring. A ring, B , is **boolean** if $a^2 = a$ for all $a \in B$. This implies that B is commutative, and that $1 + 1 = 0$. (To be precise, we will assume that $0 \neq 1$, though there is no particular reason to rule out $B = \{0\}$ at this point.) To obtain a boolean algebra, we can define the meet and join respectively by $a \wedge b = ab$ and $a \vee b = a + b + ab$, and set $a^* = 1 + a$. Conversely, we can recover the ring structure as $ab = a \wedge b$ and $a + b = (a \wedge b^*) \vee (b \wedge a^*)$.

In terms of the ring structure one can describe the median as $\mu(a, b, c) = ab + bc + ca$. (We temporarily denote the median as $\mu(a, b, c)$ to avoid confusion with multiplication in the ring.) In this case, we also have an involution, $[a \mapsto a^*]$, defined by $a^* = 1 + a$. One easily checks that the median satisfies $\mu(a, b, c)^* = \mu(a^*, b^*, c^*)$, and $\mu(a^*, a, b) = b$ for all $a, b, c \in B$.

We note that one can define a new boolean ring structure on B as follows. Fix some $p \in B$. Given $a, b \in B$, set $a \oplus b = p + a + b$, and $a.b = pa + pb + ab$. Then $(B, \oplus, .)$ is a boolean ring with zero p , and one p^* . The median and involution remain unchanged. Note that the map $[a \mapsto a + p]$ gives us a ring isomorphism from the original structure to the new one.

The property of convexity can be described in terms of the ring structure as follows. Note that if $C \subseteq B$ is convex, then so is the translate, $C + p$, for any $p \in B$. (We have noted that $[a \mapsto a + p]$ preserves the median operation.) So let us suppose that $0 \in C$. If $x \in C$ and $a \in B$, then $ax = \mu(0, a, x) \in C$. If $x, y \in C$, then $x + y + xy = \mu(x, y, x + y) \in C$, so $x + y = \mu(0, x + y, x + y + xy) \in C$. This shows that C is an ideal of B . Conversely, suppose C is an ideal. If $x, y \in C$ and $a \in B$, then $\mu(a, x, y) = ax + bx + xy \in C$, so C is convex. In other words, the convex subsets of B are precisely the translates of ideals.

Some further discussion of boolean algebras is given in Subsection 9.6.

(Ex3.6): Ternary boolean algebras.

In this context, we briefly mention the notion of a “ternary boolean algebra” as defined by Grau. This can be thought of as equivalent to a boolean algebra, but without specifying a preferred 0 or 1.

Suppose that M is a median algebra, equipped with an involution denoted $[a \mapsto a^*] : M \rightarrow M$, which is an automorphism of the median structure (here denoted $\mu(., ., .)$). We also suppose that $\mu(a^*, a, b) = b$ for all $a, b \in M$. (In other words $M = [a, a^*]$ for all $a \in M$.) We also suppose that $a^* \neq a$ (otherwise $M = \{a\}$). We refer to such a structure as a **ternary boolean algebra**.

Now choose any $p \in M$, and redefine $0 = p$, $1 = 0^*$. We write $a \wedge b = \mu(0, a, b)$ and $a \vee b = \mu(1, a, b)$. Since $M = [0, 1]$, Lemma 3.2.4 tells us that (M, \wedge, \vee) is a distributive lattice. Moreover, one readily checks that $a \wedge 1 = a \vee 0 = a$, $a \wedge a^* = 0$ and $a \vee a^* = 1$ for all $a \in M$. Therefore, $(M, \wedge, \vee, 0, 1, *)$ is a boolean algebra. From this, one sees from the de Morgan laws of a boolean algebra that $\mu(a, b, c)^* = \mu(a^*, b^*, c^*)$.

Conversely, any boolean algebra gives rise such a ternary boolean algebra structure. In other words, one can axiomatise a boolean algebra in terms of the median and involution together with a preferred element designated as 0. As we observed in Example (Ex3.3), we can modify the operations on any boolean algebra, so that any given element turns into the zero.

We will mention ternary boolean algebras again in Subsection 20.2.

(Ex3.7): Quotients.

Suppose that \sim is an equivalence relation on a median algebra, M , such that $abx \sim aby$ for all $a, b, x, y \in M$ with $x \sim y$. We can form the quotient, M/\sim , by setting $[a][b][c] = [abc]$, where $[.]$ denotes the \sim -class. It is easily seen that M/\sim is a median algebra with this structure.

We give a particular example, which will feature in the discussion of Roller boundaries in Subsection 11.12.

Given $a, b \in M$, write $a \sim b$ to mean that $\#[a, b] < \infty$. Note that, by Corollary 3.2.14, this is equivalent to saying that only finitely many walls separate a from b . From this, and Proposition 3.2.11, it is easily checked that \sim has the above properties.

A more explicit argument is based on the following two observations. First, $\#[a, b] \leq (\#[a, c])(\#[b, c])$ for all $a, b, c \in M$. This is because the map $[x \mapsto (acx, bcx)] : [a, b] \rightarrow [a, c] \times [b, c]$ is injective. (For suppose $x, y \in [a, b]$ with $acx = acy$ and $bcx = bcy$. Then $y = cyy = cy(aby) = by(acy) = by(acx) = a(byc)(byx) = a(bxc)(bxy) = bx(acy) = bx(acx) = cx(abc) = cxx = x$.) Second, $\#[abc, abd] \leq \#[c, d]$ for all $a, b, c, d \in M$. This is because the map $[x \mapsto abx] : [c, d] \rightarrow [abc, abd]$ is surjective. (If $y \in [abc, abd]$ then $y = (abc)(abd)y = ab(cdy)$.)

The requisite properties of \sim now follow.

Note that we can proceed to do the same thing with M/\sim . Taking preimages in M , this gives rise to another relation on M , which includes \sim . Indeed, we can iterate by transfinite induction: taking the union of relations at each limit ordinal. The process eventually terminates on a relation, \approx , such that every non-trivial interval in M/\approx is infinite — equivalently, M/\approx has no adjacent pairs.

(Ex3.8): \mathbb{R}^n .

We equip \mathbb{R}^n with the product median structure on each of its factors, where \mathbb{R} is given the usual median of betweenness. Of course, \mathbb{R}^n is also a vector space. We write e_1, \dots, e_n for the standard unit basis vectors, and $\underline{0}$ for the zero vector. It is natural to ask when a (vector) linear map is also a median homomorphism, or when a vector subspace is also a subalgebra. It is not hard to give a complete answer.

First, consider a linear map, $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$. Let $\lambda_i = \phi(e_i)$. Suppose ϕ is also a median epimorphism. We claim that there is precisely one i for which $\lambda_i \neq 0$. For suppose $\lambda_i, \lambda_j \neq 0$ for $i \neq j$. Now $\underline{0} \in [e_i, e_j] \cap [e_i, -e_j]$, so $0 \in [\lambda_i, \lambda_j] \cap [\lambda_i, -\lambda_j]$, giving a contradiction. This shows that, up to linear automorphism of one of the factors of \mathbb{R}^n , ϕ is projection to one of the coordinates.

We can apply this to any linear map $\mathbb{R}^n \rightarrow \mathbb{R}^m$, treating each of the factors of \mathbb{R}^m independently. If $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear automorphism, it follows that ϕ consists of permuting the coordinates, composed with a linear automorphism of the factors. In other words, it is given by a matrix with precisely one non-zero entry in each row and each column.

An example of a vector subalgebra of \mathbb{R}^n is the diagonal, $D_n := \{(x, x, \dots, x) \mid x \in \mathbb{R}\} \subseteq \mathbb{R}^n$. In fact, up to linear median automorphism of \mathbb{R}^n (of the type just described) every vector subalgebra is a direct product of such diagonals.

To see this, let $M \subseteq \mathbb{R}^n$ be an m -dimensional vector subspace of \mathbb{R}^n . Up to permuting the factors of \mathbb{R}^n , M is the graph of a linear map $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^{n-m}$. If M is subalgebra, then ϕ is a median homomorphism. The projection to each factor of \mathbb{R}^{n-m} is either identically zero, or else the projection to a coordinate of \mathbb{R}^m followed by a linear automorphism of \mathbb{R} . It now follows that, up to permutation of the factors of \mathbb{R}^n , and linear automorphism of the factors, M has the form $D_{p_1} \times D_{p_2} \times \dots \times D_{p_m} \times \{\underline{0}\}$, where $p_i \geq 1$.

Note that if M is median convex, then each $p_i = 1$, so M has the form $\mathbb{R}^m \times \{\underline{0}\}$ up to linear automorphism of \mathbb{R}^n (though this fact can be verified more directly).

Some other naturally occurring median algebras arise from spaces with measured walls (Section 19), as asymptotic cones of various spaces (Subsection 24.3), as Guirardel cores (Subsection 15.4), and from quasimedial graphs (Section 23).

4. INTERVALS AND BETWEENNESS

In this section we explore further the notion of betweenness and describe an equivalent formulation of a median algebra in terms of intervals (Theorem 4.1.1). We give a proof of Theorem 3.2.2, and describe another characterisation of median algebras due to Isbell (Theorem 4.3.1). Some more general notions of betweenness will be discussed in Section 22.

4.1. Description of median algebras in terms of intervals.

Let M be a set, and suppose that to each pair of elements, $a, b \in M$, we have associated a subset $I(a, b) \subseteq M$. Consider the following conditions:

- (I1): $(\forall a \in M) I(a, a) = \{a\}$,
- (I2): $(\forall a, b \in M) I(a, b) = I(b, a)$,
- (I3): $(\forall a, b \in M)(\forall c \in I(a, b)) I(a, c) \subseteq I(a, b)$,
- (I4): $(\forall a, b, c \in M) \#(I(a, b) \cap I(b, c) \cap I(c, a)) = 1$.

The following is due to Sholander ([Sh]).

Theorem 4.1.1.

- (1) If M is a median algebra, then M satisfies (I1)–(I4), where $I(a, b) = [a, b]$. In this case, $I(a, b) \cap I(b, c) \cap I(c, a) = \{abc\}$.
- (2) Suppose that M is a set with a family of subsets, $\{I(a, b)\}_{a, b \in M}$ satisfying (I1)–(I4), then M is a median algebra where the median is the unique point of $I(a, b) \cap I(b, c) \cap I(c, a)$. Moreover, $I(a, b) = [a, b]$.

Here we using the original definition of a median algebra, so in the following discussion we only use the symmetry of the median and axioms (M1) and (M2).

We set about the proof of Part (1). Most of that work is already done.

Let M be a median algebra, and set $I(a, b) = [a, b]$.

Now (I1) is an immediate consequence of (M1), and (I2) follows from the symmetry of the median. Property (I3) asserts that:

$$a.c.b \ \& \ a.d.c \Rightarrow a.d.b.$$

This follows since $adb = a(adc)b = (bcd|a) = a(abc)d = acd = d$. (It is a particular case of the linear interpolation rule of a median algebra, mentioned in Subsection 3.2.)

Finally, (I4) is an immediate consequence of the median rule (see Lemma 3.2.1). This proves Part (1) of Theorem 4.1.1.

The converse, Part (2), is more involved.

Let M be set, and let $\{I(a, b)\}_{a, b \in M}$ be a family of intervals satisfying (I1)–(I4). Given $a, b, c \in M$, define $abc \in M$ by $I(a, b) \cap I(b, c) \cap I(c, a) = \{abc\}$. We want to verify (M1) and (M2).

Given $x, y, z \in M$, write $x.y.z$ to mean that $y \in I(x, z)$. Given $x_1, x_2, x_3, x_4, x_5 \in M$, write $x_1.x_2.x_3.x_4$ or $x_1.x_2.x_3.x_4.x_5$ to mean that $x_i.x_j.x_k$ holds for all $i < j < k$.

First note that by (I1), $a.b.a$ implies $a = b$. This shows that (M1) holds.

Note that, by (I1) and (I4), we have $\{a\} \cap I(a, b) = I(a, a) \cap I(a, b) \neq \emptyset$, so $a \in I(a, b)$. Thus, if $x \in I(a, b)$, then $x \in I(a, b) \cap I(a, x) \cap I(b, x)$, so $x = abx$. This shows that $a.x.b \Leftrightarrow x = abx$.

If $x \in M$ and $I(a, x) \cap I(b, x) = \{x\}$, then we must have $I(a, x) \cap I(b, x) \cap I(a, b) = \{x\}$, so $a.x.b$. Conversely, if $a.x.b$ holds then $I(a, x) \cap I(b, x) = \{x\}$.

Now (I3) says that $a.c.b \ \& \ a.x.c \Rightarrow a.x.b$. From $a.x.c$, we have $I(x, c) \subseteq I(a, c)$, so $I(x, c) \cap I(c, b) \subseteq I(a, c) \cap I(c, b) = \{c\}$ so $x.c.b$. This shows that $a.c.b \ \& \ a.x.c \Rightarrow a.x.c.b$.

We can apply this interpolation repeatedly. For example, if $c.y.b$ also holds then we have $a.x.c.y.b$. Similarly, $a.x.y.b \ \& \ x.z.y \Rightarrow a.x.z.y.c$.

Now suppose that $a, b, c \in M$ and $e \in I(a, b)$, that is $a.e.b$. Let $m = abc$. We first claim that $c.m.e$ holds.

To see this, set $p = cae$, $q = bae$ and $r = cpq$. We have:

$$\begin{aligned} c.r.p \ \& \ c.p.a &\Rightarrow c.r.p.a \Rightarrow c.r.a, \\ c.r.q \ \& \ c.q.b &\Rightarrow c.r.q.b \Rightarrow c.r.b, \\ a.e.b \ \& \ a.p.e \ \& \ b.q.e &\Rightarrow a.p.e.q.b \Rightarrow a.p.q.b, \\ a.p.q.b \ \& \ p.r.q &\Rightarrow a.p.r.q.b \Rightarrow a.r.b. \end{aligned}$$

By (I4), we have

$$a.r.b \ \& \ c.r.a \ \& \ c.r.b \Rightarrow r = abc.$$

Therefore $r = m$, so $c.r.p$ gives $c.m.p$. Now,

$$c.p.e \ \& \ c.m.p \Rightarrow c.m.p.e \Rightarrow c.m.e.$$

This proves the claim.

Now set $n = mae$. Given $c.m.e$, we have $I(m, e) \subseteq I(c, e)$; and given $a.m.c$, we have $I(m, a) \subseteq I(c, a)$, both from (I3). Therefore

$$n \in I(m, e) \cap I(m, a) \cap I(a, e) \subseteq I(c, e) \cap I(c, a) \cap I(a, e) = \{ace\},$$

so $n = ace$.

In summary, we have shown that if $a, b, c \in M$, and $e \in I(a, b)$, then $(abc)ae = ace$.

Now let $a, b, c, d \in M$ be arbitrary. Setting $e = abd$, we see that $a(abc)(abd) = ac(abd)$. The left-hand side is symmetric in c, d , so swapping c and d , we get $ac(abd) = ad(abc)$. In other words, we have verified (M2).

This proves Theorem 4.1.1.

Remark. We see retrospectively that if $x, y \in I(a, b)$, then $I(x, y) \subseteq I(a, b)$. The argument above could be considerably simplified if we knew this at the outset. However I can see no very direct way of verifying this.

4.2. The short distributive law.

Our argument of the previous subsection is a variation on one in Sholander's paper [Sh]. (He bypasses direct discussion of axiom (M2).) In fact, Sholander continues in a similar vein to deduce (M3). This still involves some amount of work. Instead of reproducing that proof (we can do no better) we show how one can deduce (M3) from (M2), directly applying the axioms. The argument below is due to Knuth, Veroff and McCune [VeroM] as we discuss in the Notes to this section.

Before giving the proof proper, we introduce the following condition which we will say more about in Subsection 4.3

Consider the following property of four points, a, b, c, d , in M :

$$(M4): a(abc)(dbc) = abc.$$

This can easily be derived from (M1) and (M3): $a(abc)(dbc) = (aad)bc = abc$. However we want a derivation from (M1) and (M2). Such a derivation is provided in [VeroM] as follows:

Lemma 4.2.1. $(M1) \ \& \ (M2) \Rightarrow (M4)$.

Proof.

$$\begin{aligned} a(abc)(dbc) &= a(ab(abc))(dbc) = (b(abc)(dbc)|a) \\ &= ((dbc)ba)(abc)a = (adc|b)(abc)a = ((abc)bd)(abc)a \\ &= (adb|(abc)) = (ab(abc))d(abc) = (abc)d(abc) = abc. \end{aligned}$$

□

So let us move on to the proof of Theorem 3.2.2.

We need to verify the identity

$$(abc)xy = (axy)(bxy)c$$

for all $a, b, c, x, y \in M$.

To simplify notation, we will use the following abbreviations. We will write: $A = axy$, $B = bxy$, $C = cxy$, $O = abc$ and $H = Oxy = (abc)xy$.

We therefore want to show that $H = ABc$. By the median rule, this is equivalent to asserting that $HAc = HBc = HAB = H$. (Recall that the median rule was verified in Subsection 3.2 directly from (M1) and (M2).)

We prove this in a series of steps. In what follows, a, b, c, x, y, o are arbitrary elements of M .

$$(E1): (oxy)(axy)c = (oxy)(axy)((oxy)ac).$$

Proof: Note that

$$axy = a(axy)(oxy) \quad [\text{by (M4): } b \rightarrow x, c \rightarrow y, d \rightarrow o].$$

So the LHS is:

$$(oxy)(a(axy)(oxy))c = (ac(axy)|(oxy)) = (ac(oxy))(axy)(oxy)$$

which is the RHS.

$$(E2): (axy)(bxy)(cxy) = a(bxy)(cxy).$$

Proof: We want to show that $ABC = aBC$.

Now

$$\begin{aligned} A &= axy = a(axy)(bxy) \quad [\text{by (M4): } b \rightarrow x, c \rightarrow y, d \rightarrow b] \\ &= (axy)a(bxy) \\ &= (a(axy)(cxy))a(bxy) \quad [\text{by (M4): } b \rightarrow x, c \rightarrow y, d \rightarrow c] \\ &= (aAC)aB. \end{aligned}$$

Therefore,

$$\begin{aligned} ABC &= ((aAC)aB)BC = ((aAC)Ca)|B) \\ &= (CaB)(aAC)B = B(BCa)(ACa) \\ &= BCa \quad [\text{by (M4): } a \rightarrow B, b \rightarrow C, c \rightarrow a, d \rightarrow A] \\ &= aBC. \end{aligned}$$

$$(E3): ((abc)xy)(axy)c = (abc)xy.$$

Proof: We want to show that

$$HAc = H.$$

To begin,

$$\begin{aligned} HAc &= (Oxy)(axy)c = (Oxy)(axy)((Oxy)ac) \quad [\text{by (E1): } o \rightarrow O] \\ &= HA(Hac) =: (*). \end{aligned}$$

Now

$$\begin{aligned} HAc &= H(HAc)(bac) \quad [\text{by (M4): } a \rightarrow H, b \rightarrow a, d \rightarrow b] \\ &= (HAc)HO. \end{aligned}$$

Therefore,

$$(*) = HA((HAc)HO) = ((HAc)AO|H) = (AOH)(HAc)H =: (**).$$

Now

$$\begin{aligned} AOH &= OHA = O(Oxy)(axy) \\ &= Oxy \quad [\text{by (M4): } a \rightarrow O, b \rightarrow x, c \rightarrow y, d \rightarrow a] \\ &= H. \end{aligned}$$

So

$$(**) = H(HAc)H = H.$$

This shows that $HAc = H$ as required. This proves (E3).

We can now complete the verification of the identity.

We have $HAc = H$, and $HBc = H$ [by (E3): $a \leftrightarrow b$].

Also

$$\begin{aligned} HAB &= (Oxy)(axy)(bxy) = a(Oxy)(bxy) \quad [\text{by (E2): } c \rightarrow O] \\ &= aHB = HBa = H \quad [\text{by (E3): } a \rightarrow b, b \rightarrow c, c \rightarrow a]. \end{aligned}$$

We have shown

$$HAc = HBc = HAB = H,$$

so by the median rule,

$$H = ABC.$$

This completes the proof of Theorem 3.2.2.

4.3. Isbell's condition.

Although we will have no further use for it in our discussion, we show that Property (M4) of Subsection 4.2 provides another “four-point” characterisation of a median algebra, which is due to Isbell [Is2]. We have already observed that any median algebra satisfies (M4).

Conversely, suppose that we have a symmetric ternary operation satisfying (M1) and (M4). Note that $ab(abc) = a(abb)(abc) = abc$. We can therefore give two equivalent definitions of an “interval” in the usual way, namely: $[a, b] = \{abc \mid c \in M\} = \{c \in M \mid abc = c\}$. Given (M1), we see that (M4) can be equivalently be expressed as:

(F): If $a, b, c \in M$ and $d \in [a, b]$, then $abc \in [c, d]$.

(Note that $cd(abc) = c(abd)(abc) = abc$.)

We claim:

Theorem 4.3.1. *Let M be a set equipped with a symmetric ternary operation $[(x, y, z) \mapsto xyz]$. Then M is a median algebra if and only if it satisfies (M1) and (M4).*

Proof. We have already observed that the “only if” direction holds. We therefore suppose that M satisfies (M1) and (M4). We proceed to verify the hypotheses (I1)–(I4) of Theorem 4.1.1.

Now (I1) and (I2) are immediate, so we proceed to (I3). Let $a, b \in M$ and $c \in [a, b]$. We claim that $[a, c] \subseteq [a, b]$. To see this, let $d \in [a, c]$ and let $p = abd$. Since $c \in [a, b]$, we have $p \in [c, d]$ by (F) above. Therefore, $p = pad = (pcd)a(acd) = acd = d$, and so $d = p \in [a, b]$, proving (I3).

To prove (I4) we first make the following two observations.

(1): Let $a, b, c \in M$ and $m = abc$. Then $[a, m] \cap [b, c] = \{m\}$.

Certainly, $m \in [a, m] \cap [b, c]$. Suppose $p \in [a, m] \cap [b, c]$, that is $p = pam = pbc$. Then $p = pam = (pbc)a(abc) = abc = m$ as required.

(2): If $a, b, c, t \in M$ with $[a, t] \cap [b, c] = \{t\}$, then $t = abc$.

Since $t \in [b, c]$ we have $abc \in [a, t]$ by (F). Therefore $abc \in [a, t] \cap [b, c] = \{t\}$, so $abc = t$.

To verify (I4), let $x, y, z \in M$, and let $m = xyz \in [x, y] \cap [y, z] \cap [z, x]$. Suppose $p \in [x, y] \cap [y, z] \cap [z, x]$. Since $p \in [x, z]$ we have $[x, p] \subseteq [x, z]$ by (I3). Applying (1) above, we have $[y, m] \cap [x, p] \subseteq [y, m] \cap [x, z] = \{m\}$. Moreover, since $p \in [y, z]$, we have $m = xyz \in [x, p]$ by (F). Therefore, $[y, m] \cap [x, p] = \{m\}$, and so by (2) we get $m = yxp = p$. We have shown that $[x, y] \cap [y, z] \cap [z, x] = \{xyz\}$.

This verifies (I4), and so by Theorem 4.1.1 (2), M is a median algebra where the median of x, y, z is xyz . \square

From here, we could proceed to verify (M2) and (M3) as before (though the argument can be shortened a little since we already have a few of the necessary ingredients). Of course, this is all rather involved. Again it would be nice to have a short direct derivation of these identities.

5. FREE MEDIAN ALGEBRAS

In this section, we will give a more concrete description of the free median algebra on a finite set. This will be in terms of “flows” on the proper power set. It is given here as Theorem 5.2.3, though we will postpone the proof of that result until Subsection 9.2. We give an explicit account of the free median algebra on a set with at most five elements. An understanding of the free median algebra on four elements will be helpful in later discussions (see, for example, Subsection 13.2). First, we introduce some general notions that will find further uses later.

5.1. Some general notions.

Let M be a median algebra. Recall that $a, b \in M$ are **adjacent** if $[a, b] = \{a, b\}$. Let $\Gamma = \Gamma(M)$ be the graph with vertex set $V(\Gamma) = M$, and with two elements connected by an edge if they are adjacent. We refer to $\Gamma(M)$ as the **adjacency graph** associated to M .

Lemma 5.1.1. *If $a, b \in M$ and $[a, b]$ is finite, then a, b are connected by a path in $\Gamma(M)$.*

Proof. In fact, we see that any maximal chain, $a = x_0 < x_1 < \dots < x_n = b$ gives us a path in $\Gamma(M)$. \square

In particular, it follows that if M is finite, then $\Gamma(M)$ is connected. In fact, one can see in this case that $a.c.b$ holds if and only if c lies in some shortest path from a to b in $\Gamma(M)$. (One can check that such shortest paths are precisely the maximal chains in $[a, b]$: see Subsection 11.2.) From this one can read off the median in M from the graph $\Gamma(M)$. We remark that all we really need in the above is that intervals in M are finite. We will explore this in more detail in Sections 11 and 16.

Let M be any median algebra. Given elements, $a_1, \dots, a_n, p \in M$, write

$$(a_1 a_2 \dots a_n | p) = a_1 \wedge a_2 \wedge \dots \wedge a_n \in M,$$

where \wedge is the meet operation, as defined in Subsection 3.2 by setting $a \wedge b = abp$. Note that $(abc|p) = (abp)cp$, so this is consistent with the notation defined in Subsection 2.1. It is natural to set $(a|p) = a$ for all $a \in M$.

The element $(a_1 \dots a_n | p)$ does not change under permuting or repeating the elements a_i . Given any non-empty finite subset $A \subseteq M$, we can therefore write $(A|p) = (a_1 \dots a_n | p)$, where $A = \{a_1, \dots, a_n\}$.

In the case where $M = \{0, 1\}$, this again has an interpretation in terms of voting. If the jurors $A \cup \{p\}$ each cast a vote, 0 or 1, with the outcome deemed to be $(A|p)$, then the vote of p can only be overruled by the unanimous vote of A .

If we write $m = (A|p)$, then clearly we have $p.m.a$ for all $a \in A$. In fact, $p.m.x$ holds for all $x \in \langle A \rangle$. This follows by iterating the median operation and noting that $pm(xyz) = (pm.x)(pm.y)(pm.z)$ for all $x, y, z \in M$. (This has a geometrical interpretation in terms of convex hulls, as we will see in Subsection 7.4.)

5.2. Flows and superextensions.

Now let X be any set, and let $F(X)$ be the free median algebra on X . Given any $p \in X$, write $m(p) = ((A \setminus \{p\})|p)$.

Lemma 5.2.1. *For each $p \in X$, $m(p)$ is the unique point of $F(X)$ adjacent to p .*

Proof. Since $F(X) = \langle X \rangle$, we saw above that $p.m(p).x$ holds for all $x \in F(X)$. It remains to check that $p \neq m(p)$. To this end, define a map $\phi : X \rightarrow \{0, 1\}$ by $\phi(p) = 1$ and $\phi|(X \setminus \{p\}) \equiv 0$. This extends to a homomorphism $\hat{\phi} : F(X) \rightarrow \{0, 1\}$. Now $\hat{\phi}(m(p)) = (0 \dots 0 | 1) = 0$ in $\{0, 1\}$, and so $m(p) \neq p$. \square

Let $\Psi = \Psi(X) = \mathcal{P}(\mathcal{P}_0(X))$, where \mathcal{P} denotes power set, and \mathcal{P}_0 denotes the proper power set (removing the empty set and the whole set). This is essentially the same as the definition we gave in Subsection 3.3 after identifying a set with its characteristic function. (In Subsection 3.3, we used $\mathcal{P}(X)$ instead of $\mathcal{P}_0(X)$, but this makes no essential difference, as we note below.) In these terms, the evaluation map $\iota : X \rightarrow \Psi$ is given by setting $\iota(x) = \{A \subseteq X \mid x \in A\}$. This is the *principal ultrafilter* at x .

Given three families, $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \Psi$, recall that the median, \mathcal{ABC} , is defined by saying that a set $A \in \mathcal{P}_0(X)$ lies in \mathcal{ABC} if and only A lies in at least two of the families, $\mathcal{A}, \mathcal{B}, \mathcal{C}$.

Definition. We say that $\mathcal{A} \in \Psi$ is a *flow* if it satisfies the following three conditions:

(P1): if $A \in \mathcal{A}$, $B \in \mathcal{P}_0(X)$ and $A \subseteq B$ then $B \in \mathcal{A}$,

(P2): if $A, B \in \mathcal{A}$, then $A \cap B \neq \emptyset$, and

(P3): for all $A \in \mathcal{P}_0(X)$, then either $A \in \mathcal{A}$ or $A^* \in \mathcal{A}$.

Here $A^* := X \setminus A$. Note that by (P2) exactly one of A and A^* lies in \mathcal{A} . (Indeed, given (P1), we could remove (P2) and replace (P3) with this stronger statement.) The terminology of “flow” will be explained in Subsection 9.1, where such things will be defined more generally.

Remark. We could, instead, have taken $\Psi = \mathcal{P}(\mathcal{P}(X))$, with the same definition of “flow”. This gives rise to an essentially equivalent notion. We would just have to add (or remove) the whole set, X , to \mathcal{A} . (This is the definition of Ψ we gave in Section 3.3.) We have disallowed it here, because it fits better with our account of more general flows in Section 9. In any case, we can identify $\mathcal{P}(\mathcal{P}_0(X))$ as a subalgebra $\mathcal{P}(\mathcal{P}(X))$, under the map $[\mathcal{A} \mapsto \mathcal{A} \cup \{X\}]$.

Definition. The *superextension* of X is the set of all flows. We denote it $\Phi(X) \subseteq \Psi(X)$.

It is easy to check that $\Phi(X)$ is a subalgebra of $\Psi(X)$. Moreover $\iota(X) \subseteq \Phi(X)$.

Remark. Again, it is common elsewhere to include X itself in the superextension, but this makes no essential difference to the discussion. In this way, $\Phi(X)$ can also be viewed as subalgebra of $\mathcal{P}(\mathcal{P}(X))$, as in the previous remark.

In Subsection 3.3, we constructed $F(X)$ as $F(X) = \langle \iota(X) \rangle \subseteq \Psi(X)$. From the above we see:

Lemma 5.2.2. *For any set X , $F(X) \subseteq \Phi(X)$.*

In fact:

Theorem 5.2.3. *If X is finite, then $F(X) = \Phi(X)$.*

In other words for finite sets, free median algebras are the same as superextensions. Theorem 5.2.3 will follow from Lemma 9.2.3 and we postpone the proof until then. Another more explicit construction is given in Section 20.2, and a more general statement is given as Proposition 9.4.2.

The remainder of this section will mostly refer to $\Phi(X)$.

Given $\mathcal{A} \in \Phi(X)$, write $\mathcal{M}(\mathcal{A})$ for the set of minimal elements of \mathcal{A} (that is, with respect to inclusion). Note that we can reconstruct \mathcal{A} as the set of all subsets containing some element of $\mathcal{M}(\mathcal{A})$. In practice, it is often easiest to describe a particular element of $\Phi(X)$ by specifying $\mathcal{M}(\mathcal{A})$.

Note that $\mathcal{M}(A)$ is a Sperner family, as defined in Subsection 2.2. Moreover, $\mathcal{A} = (\mathcal{M}(A))^\uparrow$, as defined there. We see that the map $[\mathcal{A} \mapsto \mathcal{M}(\mathcal{A})]$ gives us a median monomorphism of $\Phi(X)$ into the free distributive lattice, $D(X)$. (The fact that it is a homomorphism follows directly from the definitions of the median in $\Phi(X)$ and in a distributive lattice, and from the lattice structure defined on $D(X)$.)

Another way to think of this embedding would be to write any given element of $F(X)$ as a median expression with arguments in X , then formally replace it by an expression involving \wedge and \vee by repeatedly applying the formula $xyz := (x \wedge y) \vee (y \wedge z) \vee (z \wedge x)$, and finally put the resulting expression into normal form. It is easily seen that this agrees with the above. (Another way to see that this operation is well defined, would be to note that any two median expressions representing the same element of $\Phi(X)$ would always evaluate to the same value in $\{0, 1\}$, and so the same applies to the two lattice expressions thus obtained.) This principle can be used to give explicit derivations of median identities, as discussed in Subsection 6.3.

5.3. The structure of a superextension.

Note that the halfspaces of $\Phi(X)$ correspond to proper subsets of X , and walls correspond to bipartitions (that is, partitions into two non-empty subsets).

Theorem 5.2.3 is equivalent to asserting that $\Phi(X)$ is generated by $\iota(X)$: that is the set of principal ultrafilters on X . The discussion in Subsection 5.4 will give explicit verifications of this when $\#X \leq 5$. We will postpone the general proof for the moment. In fact, we will eventually see three different proofs. One, in terms of a duality principle, is discussed in Section 9 (see Lemmas 9.2.2 and 9.2.3). It is also an immediate consequence of a more general result about “spaces with walls”, namely Proposition 9.4.2. A third, more explicit proof, in terms of boolean functions is given in Subsection 20.1: see Proposition 20.1.2.

Given $\mathcal{A} \in \Phi(X)$, write $\mathcal{M}(\mathcal{A})$ for the set of minimal elements of \mathcal{A} (that is, with respect to inclusion). Note that we can reconstruct \mathcal{A} as the set of all subsets containing some element of $\mathcal{M}(\mathcal{A})$. In practice, it is often easiest to describe a particular element of $\Phi(X)$ by specifying $\mathcal{M}(\mathcal{A})$.

To simplify notation, it will be convenient to identify X with $\iota X \subseteq \Phi(X)$. In these terms, if $a \in X$, then $\mathcal{M}(a) = \{\{a\}\}$. Given any $a, b, c \in X$, then

$\mathcal{M}(abc) = \{\{a, b\}, \{b, c\}, \{c, a\}\}$. More generally, if $a_1, \dots, a_n, p \in X$, are distinct, then

$$\mathcal{M}(\{a_1, \dots, a_n\}|p) = \{\{p, a_1\}, \{p, a_2\}, \dots, \{p, a_n\}, \{a_1, a_2, \dots, a_n\}\}.$$

(Note that this has an interpretation in terms of voting which we mentioned in Subsection 5.1.)

We can describe intervals and adjacency in $\Phi(X)$ as follows.

Let $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \Phi(X)$. By Lemma 2.1.3, $\mathcal{A}\mathcal{B}\mathcal{C}$ holds if and only if $\mathcal{A} \cap \mathcal{B} \subseteq \mathcal{C} \subseteq \mathcal{A} \cup \mathcal{B}$. But $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are all $*$ -transversals and so $\mathcal{A} \cap \mathcal{B} \subseteq \mathcal{C} \Leftrightarrow \mathcal{C} \subseteq \mathcal{A} \cup \mathcal{B}$. Therefore:

Lemma 5.3.1. *If $\mathcal{A}, \mathcal{B} \in \Phi(X)$, then*

$$[\mathcal{A}, \mathcal{B}] = \{\mathcal{C} \in \Phi(X) \mid \mathcal{A} \cap \mathcal{B} \subseteq \mathcal{C}\} = \{\mathcal{C} \in \Phi(X) \mid \mathcal{C} \subseteq \mathcal{A} \cup \mathcal{B}\}.$$

Let $\mathcal{A} \in \Phi(X)$, and let $A \in \mathcal{M}(\mathcal{A})$. Let $\mathcal{B} = (\mathcal{A} \setminus \{A\}) \cup \{A^*\}$. It is easily seen that $\mathcal{B} \in \Phi(X)$, and that $A^* \in \mathcal{M}(\mathcal{B})$. We say that \mathcal{B} is obtained from \mathcal{A} by **flipping** A . Conversely, \mathcal{A} is obtained from \mathcal{B} by flipping A^* . We refer to such an operation as a **flip**. Note that \mathcal{A}, \mathcal{B} differ by a flip if and only if $\#(\mathcal{A} \triangle \mathcal{B}) = 2$. In fact:

Lemma 5.3.2. *$\mathcal{A}, \mathcal{B} \in \Phi(X)$ are adjacent if and only if they differ by a flip.*

Proof. Suppose that \mathcal{B} is obtained by flipping $A \in \mathcal{M}(\mathcal{A})$. Then $\mathcal{A} \cap \mathcal{B} = \mathcal{A} \setminus \{A\} = \mathcal{B} \setminus \{A^*\}$. If $\mathcal{C} \in [\mathcal{A}, \mathcal{B}]$, then by Lemma 5.3.1, we have $\mathcal{A} \cap \mathcal{B} \subseteq \mathcal{C}$, and so either $\mathcal{A} \subseteq \mathcal{C}$ or $\mathcal{B} \subseteq \mathcal{C}$. Since these are all $*$ -transversals, this implies either $\mathcal{A} = \mathcal{C}$ or $\mathcal{B} = \mathcal{C}$. Thus \mathcal{A}, \mathcal{B} are adjacent.

Conversely, suppose that \mathcal{A}, \mathcal{B} are adjacent. Let A be a minimal element of $\mathcal{A} \setminus \mathcal{B}$. Now A is also minimal in \mathcal{A} . (For if $B \in \mathcal{A}$ were strictly contained in A , then $B \in \mathcal{B}$, so we get the contradiction that $A \in \mathcal{B}$.) Let \mathcal{C} be obtained by flipping A in \mathcal{A} . Then $\mathcal{A} \cap \mathcal{B} \subseteq \mathcal{C}$, so $\mathcal{C} \in [\mathcal{A}, \mathcal{B}]$. Since $\mathcal{C} \neq \mathcal{A}$, we get $\mathcal{C} = \mathcal{B}$. \square

In particular, we see that there is a natural bijection between $\mathcal{M}(\mathcal{A})$ and the set elements of $\Phi(X)$ adjacent to \mathcal{A} .

This will allow us to construct the adjacency graph, $\Gamma(\Phi(X))$, for small X . (Note that the above implies that the combinatorial distance between \mathcal{A} and \mathcal{B} in $\Gamma(\Phi(X))$ is at most $\frac{1}{2}\#(\mathcal{A} \triangle \mathcal{B})$. In fact, these are equal, as we discuss in Subsection 11.9.) The structure of $\Phi(X)$ is a little different depending on whether $\#X$ is odd or even.

Suppose $\#X = 2N - 1$, where N is a positive integer. In this case, $\Phi(X)$ has a **central element**, namely the set of all subsets of X of size at least N . In terms of “voting” this can be thought of as “majority vote”: that is when we have $2N - 1$ jurors all casting votes 0 or 1, and the outcome is determined by a simple majority. It is also not hard to see that the central element is the unique element of $\Phi(X)$ invariant under all permutations of X .

Suppose $\#X = 2N$, where N is a positive integer. In this case, we have a central cube of $\Phi(X)$ described as follows. First, let \mathcal{Y} be the set of subsets of X of size at least $N + 1$. Let \mathcal{W}_0 be the set of equal bipartitions of X : that is partitions

into two subsets of size N . Note that $\#\mathcal{W}_0 = \nu := \frac{1}{2} \binom{2N}{N} = \binom{2N-1}{N-1}$. We can view each element of \mathcal{W}_0 as a 2-element median algebra, and equip $Q := \prod \mathcal{W}_0$ with the product median structure. As such, Q is a ν -cube. We can view each element of Q as choosing some element from each equal bipartition so as to give us a family, \mathcal{C} , of subsets of X each of size N . We define a map $\phi : Q \rightarrow \Phi(X)$ by setting $\phi(\mathcal{C}) = \mathcal{C} \cup \mathcal{Y}$. This is clearly a median monomorphism, and its image, $\phi(Q)$, is a ν -cube in $\Phi(X)$. In fact, $\phi(Q)$ is convex in $\Phi(X)$. (For example, take any $\mathcal{C} \in Q$, and let $\mathcal{D} \in Q$ be the antipodal vertex: that is, from each bipartition we make the opposite choice of element. From Lemma 5.3.1, it is easily seen that $\phi(Q) = [\mathcal{C}, \mathcal{D}]$, and so is convex.) We refer to $\phi(Q)$ as the **central cube** of $\Phi(X)$.

Remark. While we are on the subject, we note that the operation of “majority vote” can be applied in any median algebra, M . One way to describe this is as follows.

Let $n = 2m + 1$ for $m \in \mathbb{N}$, and let $\underline{x} := \{x_1, \dots, x_n\} \subseteq M$. Let $\Phi = \Phi(\{1, \dots, n\})$, and let $\phi : \Phi \rightarrow M$ be the homomorphism extending the map $[i \mapsto x_i]$. Let $\mu(\underline{x})$ be the image of the central element of Φ under ϕ . Thus, $\mu : M^n \rightarrow M$ is a symmetric n -ary operation on M . We refer to it as the **standard majority vote**. Note that one can define $\mu(\underline{x})$ explicitly by applying any formula for majority vote. Thus, for $n = 1, 3, 5$ respectively, we have $\mu(a) = a$, $\mu(a, b, c) = abc$, and $\mu(a, b, c, d, e) = a(bcd)((bce)de)$ (as discussed in the next subsection).

Another way to describe $\mu(\underline{x})$ is as follows. Let \mathcal{H} be the set of halfspaces of M , and let $\mathcal{H}(\underline{x}) = \{H \in \mathcal{H} \mid \#\{i \mid x_i \in H\} \geq m + 1\}$; in other words, the set of halfspaces which contain a majority of the points x_i . Then $\bigcap \mathcal{H}(\underline{x}) = \{\mu(\underline{x})\}$. (The equivalence with the previous description can be seen from the fact that it is invariant under all permutations of the x_i , or directly in terms of the description of the central element as majority vote.)

This has a simple interpretation for a simplicial tree, T . If $x_1, \dots, x_n \in V(T)$, then $\mu(\underline{x})$ is the unique vertex such that any branch of T at $\mu(\underline{x})$ contains fewer than half of the x_i . Similarly, if M is a total ordered set, then $\mu(\underline{x})$ is the “median” value of \underline{x} in the commonly used sense of the word.

The notion of majority vote will appear at several places in our discussion, and is described in some detail in Subsection 20.3.

5.4. Description of the first few cases.

We now give an explicit description of $\Phi(X)$ when $\#X \leq 5$. It is natural to think of $\Gamma(\Phi(X))$ as the 1-skeleton of a cube complex, Δ . (One can imagine the cubes as euclidean cubes of unit side-length, meeting along common faces.) Note that the link of any vertex is a simplicial complex: the simplices corresponding to the cubes which contain that vertex. It turns out that in our case, such a simplicial complex will always be a “flag complex”: any clique (that is complete subgraph) of its 1-skeleton lies in a simplex. We will discuss such matters in greater detail and generality in Sections 11 and 17. The situation starts out very simply, but the

complexity grows very rapidly.

If $\#X \leq 2$, then $\Phi(X) = X$ (after identifying X and ιX).

If $X = \{a, b, c\}$ then we have $\Phi(X) = \{a, b, c, abc\}$, and we see that $\Gamma(\Phi(X))$ is a tripod with feet a, b, c and with central vertex abc .

Now let $X = \{a, b, c, d\}$. Suppose $\mathcal{A} \in \Phi(X)$. If \mathcal{A} contains a singleton, say $\{a\}$, then \mathcal{A} is just a (or more precisely ιa). If \mathcal{A} has no singleton, then it must contain a 2-element set. A convenient way to describe it is as follows. Let $G = G(\mathcal{A})$ be the graph with vertex set X and with $x, y \in X$ connected by an edge if $\{x, y\} \in \mathcal{A}$ (so $\{x, y\} \in \mathcal{M}(\mathcal{A})$). Now any two edges of G meet. Moreover, any pair of vertices are either connected by an edge or else the complementary pair of vertices are. From this we see easily that either G is a triangle, with vertices, a, b, c , say, together with an isolated vertex, d ; or else it is a tripod, with feet a, b, c , say, connected to a central vertex d . In the former case, $\mathcal{A} = abc$, and in the latter case, $\mathcal{A} = (abc|d)$ (after identifying a with ιa). Writing $m(d) = (abc|d)$ as above, and $h(d) = abc$, we see that $\Phi(X)$ has precisely 12 elements, namely, $a, b, c, d, m(a), m(b), m(c), m(d), v(a), v(b), v(c)$ and $v(d)$. We also see that $\Phi(X) = \langle a, b, c, d \rangle$, so that $F(X) = \Phi(X)$ verifying Theorem 5.2.3 in this case.

To complete the description of $\Phi(X)$, we note that the 8 elements, $m(a), m(b), m(c), m(d), v(a), v(b), v(c), v(d)$, form the central 3-cube of $\Phi(X)$. This is because the minimal set of each of these elements is obtained by selecting one element from each of the three equal bipartitions of X . Moreover, $m(a)$ is adjacent each of $v(b), v(c), v(d)$, etc. Now $m(a)$ is obtained by flipping $\{a\}$ in $a \in \Phi(X)$, so we see by Lemma 5.3.2 that a is adjacent to $m(a)$. In summary, we see that $\Gamma(\Phi(X))$ is the 1-skeleton of a cube complex consisting of one central 3-cube together with free edges attached at alternating vertices of this cube.

Now suppose $X = \{a, b, c, d, e\}$. Let $\mathcal{A} \in \Phi(X)$. Again, if \mathcal{A} contains a singleton then it lies in X . So suppose not. We again construct a graph, $G = G(\mathcal{A})$, with vertex set X and x, y adjacent if $\{x, y\} \in \mathcal{A}$. In other words, the edge set consists of the 2-element sets in $\mathcal{M}(\mathcal{A})$. The remaining elements of $\mathcal{M}(\mathcal{A})$ are the 3-element subsets of X which neither contain nor are disjoint from any edge of $G(\mathcal{A})$. In this way, $G(\mathcal{A})$ determines $\mathcal{M}(\mathcal{A})$ and hence also \mathcal{A} . As in the previous case, any two edges of $G(\mathcal{A})$ must be adjacent in $G(\mathcal{A})$.

If G has no edges, then \mathcal{A} is the central vertex of $\Phi(X)$: it consists of all subsets of X with at least 3 elements. One can give various expressions for \mathcal{A} . For example, it is given by $a(bcd)((bce)de)$ or by $(ab(cde))(cd(abe))e$. This example (“majority vote”) will be discussed further in Subsections 6.2 and 20.1.

If G has at least one edge, then it is easily seen that it is either a triangle, or else consists of a vertex with one, two, three, or four edges attached (together with the remaining isolated vertices). In other words up to permuting X , the edge-set

is one of: $\{\{a, b\}, \{b, c\}, \{c, a\}\}$, $\{\{a, e\}\}$, $\{\{a, e\}, \{b, e\}\}$, $\{\{a, e\}, \{b, e\}, \{c, e\}\}$ or $\{\{a, e\}, \{b, e\}, \{c, e\}, \{d, e\}\}$. One can verify that these correspond respectively to the elements, abc , $ae(bcd)$, $(abe)(ace)(bde)$, $(abc|e)$ and $(abcd|e)$. We see explicitly that $\Phi(X)$ is generated by a, b, c, d, e , again verifying Theorem 5.2.3 in this case.

Summing the numbers of each type of element, we see that in total there are $5 + 1 + \binom{5}{3} + \binom{5}{2} + 5\binom{4}{2} + 5 \cdot 4 + 5 = 81$. In other words, for $\#X = 5$, we have $\#\Phi(X) = 81$.

To describe the structure of $\Gamma(\Phi(X))$ we can think of $\Gamma(\Phi(X))$ as the 1-skeleton of a cube complex Δ . This consists of five 4-cubes and ten 3-cubes, all meeting at the central vertex, together with a free edge attached to the antipodal vertex of each 4-cube. Here is an explicit account of how these fit together.

First recall that the Petersen graph, P , is a graph with 10 vertices and 15 edges. (It can be described by taking 1-skeleton of the dodecahedron and quotienting by the antipodal map.) We are really interested in the complement graph, P^* — this has the same vertex set but complementary edge set. Note that P^* has 5 maximal cliques on 4 vertices, and 10 maximal cliques on 3 vertices. We can construct a simplicial (flag) complex, Σ , with 1-skeleton P^* by gluing in 3- and 2-simplices to these maximal cliques. Now to construct Δ , we take our 5 4-cubes and 10 3-cubes, and glue them together so that they all meet at a common vertex, and with Σ as the link of that vertex (cf. Example (Ex3.2) of Subsection 3.4). Finally we attach a free edge to the vertex of each 4-cube antipodal the central vertex.

We relate this back to $\Phi(X)$ as follows. Recall that the Petersen graph, P , can also be described as the graph whose vertex set is the set of 2-element subsets of X , and with adjacency given by disjointness. (This is the so-called “Kneser graph”, $KG_{5,2}$.) Thus, in P^* , two such subsets are adjacent if they intersect. The central element, \mathcal{A} , of $\Phi(X)$ corresponds to the central vertex of Δ . By Lemma 5.3.2, any adjacent element of $\Phi(X)$ is obtained by flipping some 3-element subset, $A \in \mathcal{A}$. Writing $A^* = \{a, b\} \subseteq X$, we see that gives us the element of $\Phi(X)$ with edge set $\{\{a, b\}\}$ in $\Gamma(\mathcal{A})$. In this way, we get a bijective correspondence between the vertices of P^* and the elements of $\Phi(X)$ adjacent to the central element. Now an edge of P^* has incident vertices of the form $\{a, b\}$ and $\{a, c\}$, where $a, b, c \in X$ are distinct. We have a third element of $\Phi(X)$ given by the edge set $\{\{a, b\}, \{a, c\}\}$ in the above description. Together with the central element, these form a 2-cube in $\Phi(X)$. In other words, the edges of P^* correspond to the 2-cubes of $\Phi(X)$ which contain the central element. We can now fill in the 4-cubes, which have antipodal vertices of the form $\{\{a, b\}, \{a, c\}, \{a, d\}, \{a, e\}\}$, and the 3-cubes which have antipodal vertices of the form $\{\{a, b\}, \{b, c\}, \{c, a\}\}$. Note that the 4-cubes also contain the elements of the form $\{\{a, b\}, \{a, c\}, \{a, d\}\}$. Finally, we attach each element $a \equiv \{\{a\}\}$ to $\{\{a, b\}, \{a, c\}, \{a, d\}, \{a, e\}\}$ by a free edge. This gives us our cube complex Δ , with 1-skeleton, $\Gamma(\Phi(X))$.

Beyond this point, $\Phi(X)$ becomes too complicated to give a simple explicit description: when $\#X = 6$, $\Phi(X)$ has a central 10-cube and a total of 2646 elements.

Nevertheless, one can make a number of further general observations. We say a bit more about this in Subsection 11.9, when we have some more facts about cube complexes at our disposal.

6. EXPRESSIONS AND IDENTITIES

In this section, we give a more formal treatment of median expressions. We go on to describe a simple procedure for verifying identities, which is usually much simpler than messing with the axioms directly. At the end of the section, we explain how such identities can be systematically proven directly from the axioms, though the procedure for carrying it out is somewhat involved.

6.1. Formal definitions.

Let X be any set, which we think of the **alphabet**. A (symmetric ternary) **expression**, E , in X consists of the following data. We have a finite tree, $T = T_E$, with a preferred **root** vertex, $r = r_E \in V(T)$, a set of **leaves** $L = L_E \subseteq V(T)$, and a map $\lambda = \lambda_E : L \rightarrow X$, such that r has valence 3, each element of L has valence 1, and every other vertex has valence 4. In addition, we allow for a **trivial** expression where $L_E = V(T_E) = \{r_E\}$. We think of the map λ as “labelling” the leaves of the tree. We view an expression as being defined up to isomorphism of the tree respecting the root and labelling.

Any vertex $v \in V(T)$ determines a **branch**, S , of T — that is the subtree consisting of those points (including v) separated from the root by v . This gives us an expression, F , defined by setting $T_F = S$, $r_F = v$, $L_F = L_E \cap S$ and $\lambda_F = \lambda_E|_{L_F}$. We refer to S as a **subexpression** of E .

Given three expressions, A, B, C , we can form an expression, D , by taking the disjoint union $T_A \sqcup T_B \sqcup T_C$ and connecting r_A, r_B, r_C , to a new vertex, r_D , by three new edges. This gives us a tree, T_D , with root r_D , and with $L_D = L_A \sqcup L_B \sqcup L_C$. We write $D = ABC$. Note that any non-trivial expression can be written in this way.

Identifying an element $a \in X$, with the trivial expression labelled a , we see that we can represent any expression as string in the alphabet X together with parentheses “(” and “)”. Note that we are allowing permutations. Thus for example, $((abc)ad)(bce)$ and $(ceb)(da(bac))$ are the same expression, and $(acb)da$ is a subexpression thereof. (Putting a preferred linear order on the entries of an expression would be unnatural from our point of view.)

Let $\mathcal{E} = \mathcal{E}(X)$ be the set of all expressions in X . Let \sim be the smallest equivalence relation satisfying:

- (1): if $A, B \in \mathcal{E}$, then $AAB \sim A$,
- (2): if $A, B, C, D \in \mathcal{E}$ then $(ABD)CD \sim (ACD)BD$, and
- (3): if $A, B, C, D \in \mathcal{E}$ and $C \sim D$, then $ABC \sim ABD$.

Let $\Theta = \Theta(X) = \mathcal{E}/\sim$. We write $[A]$ for the \sim -class of A . By (3), we have a well defined ternary operation on Θ defined by $[ABC] = [A][B][C]$. By (1) and (2), this satisfies axioms (M1) and (M2), and so Θ is a median algebra. Define $\iota : X \rightarrow \Theta$ by $\iota(a) = [a]$.

Let $\phi : X \rightarrow M$ be any map into a median algebra, M . Given an expression $E \in \mathcal{E}$, we define $\phi(E)$ recursively by the rule $\phi(ABC) = \phi(A)\phi(B)\phi(C)$, and taking it to be given by the original map on trivial expressions. Thus $\phi(E)$ is the **evaluation** of E in M . Directly from the axioms, (M1) and (M2) in M , we see that if $E \sim F$, then $\phi(E) = \phi(F)$. We therefore get a map $\hat{\phi} : \Theta \rightarrow M$, defined by setting $\hat{\phi}([E]) = \phi(E)$. By construction, $\hat{\phi}$ is a median homomorphism satisfying $\hat{\phi}(\iota a) = \phi(a)$ for all $a \in X$. Note that ι is injective. (To see this, take $M = \{0, 1\}$. If $a, b \in X$ are distinct, let ϕ be any map with $\phi(a) \neq \phi(b)$, then $\hat{\phi}(\iota(a)) \neq \hat{\phi}(\iota(b))$.) After identifying X with $\iota X \subseteq \Theta$, we have shown the following:

Proposition 6.1.1. *Let X be any set, and let $\Theta \supseteq X$ be the set of \sim -classes of median expressions in X , with the median operation as defined above. Then Θ is the free median algebra on X .*

We note that one can equivalently define the relation \sim by saying that two expressions are equivalent if we can get between them by a finite sequence of operations which modify some branch of the tree by performing operations of type (1) or (2) above (i.e. interchanging AAB with A , or interchanging $(ABD)CD$ with $(ACD)BD$ in some subexpression). This can be thought of as a finite sequence of applications of axioms (M1) and (M2).

6.2. Verification of identities.

A **tautological** identity between two median expressions is one that holds however it is evaluated in any median algebra. By Proposition 6.1.1, this is the same as saying that the expressions are equivalent under the relation, \sim . For example, we have seen that the identity $ab(xyz) = (abx)(aby)(abz)$ is tautological. From the above, we see that any tautological identity can be explicitly verified by a repeated application of the axioms (M1) and (M2). We refer to such a sequence as a “derivation” of the identity.

However, there is a simpler way of checking identities in practice. We have seen that any two distinct points in any median algebra, M , are separated by a wall (Proposition 3.2.12). Equivalently, M can be embedded into a hypercube (Proposition 3.2.13). (Of course, we had to verify quite a few identities directly in order to get to that point, but that work is now done.) To check that an identity is tautological, it is therefore sufficient to verify that it holds in $\{0, 1\}$ for all possible assignments of the arguments. One can often argue more efficiently in the contrapositive.

To illustrate this, we will write $x\downarrow$ and $x\uparrow$ to mean that x takes the value 0 and 1 respectively. So for example, if $abc\uparrow$ holds, then up to permuting a, b, c we can assume that $a\uparrow$ and $b\uparrow$ hold.

As an example, we stated in Subsection 5.4 that $a(bcd)((bce)de)$ represents the majority vote of five jurors. One consequence is that it is symmetric under permuting a, b, c, d, e . In particular (transposing a and e), the identity

$$a(bcd)((bce)de) = e(bcd)((bca)da)$$

is tautological.

Suppose we had not noticed that this is majority vote, and wanted to verify this identity. We might proceed as follows.

Without loss of generality (swapping \downarrow and \uparrow), we can suppose that $a(bcd)((bce)de)\downarrow$ and $e(bcd)((bca)da)\uparrow$. Also (swapping a and e) we can suppose that $bcd\uparrow$. From the first expression, it follows that $a\downarrow$ and $(bce)de\downarrow$.

Now if $e\uparrow$, then the last statement implies that $d\downarrow$ and $bce\downarrow$. Given $e\uparrow$, we therefore get $b\downarrow, c\downarrow$, and so $bcd\downarrow$, contradicting an earlier statement.

Therefore, $e\downarrow$. Now $e(bcd)((bca)da)\uparrow$ gives $(bca)da\uparrow$. Given $a\downarrow$, this implies $d\uparrow$ and $bca\uparrow$. Again, given $a\downarrow$, we get $b\uparrow, c\uparrow$, so $bce\uparrow$. Thus, given $d\uparrow$, we get $(bce)de\uparrow$, again contradicting an earlier statement.

This shows that the identity is tautological as claimed.

Of course, there are lots of possible variations on the above logic, but they all lead quickly to a similar contradiction.

One can proceed to verify the other symmetries of the expression or its equivalence with $(ab(cde))(cd(abe))e$, for example.

A similar discussion applies to conditional identities. Let I_1, \dots, I_n, I be ternary identities. We say that $I_1 \& \dots \& I_n$ **tautologically implies** I if any median algebra in which $I_1 \& \dots \& I_n$ holds for some assignment of the arguments, then I also holds for that assignment of the arguments. (For example, we have seen that $a.x.y \& a.y.c$ tautologically implies $x.y.c$. Here, of course, $a.b.c$ is shorthand for the identity $abc = b$.) By a similar argument as for tautological identities, we see that a conditional identity can be verified by a repeated application of the identities I_1, \dots, I_n together with the axioms (M1) and (M2).

However, it is again sufficient to verify it in $\{0, 1\}$. To see this, suppose that the two sides of I evaluate to different elements, say p and q , in some median algebra, M , in which $I_1 \& \dots \& I_n$ holds. We postcompose by a wall map separating p and q . Now $I_1 \& \dots \& I_n$ still holds in $\{0, 1\}$, whereas I does not. This gives a contradiction.

We illustrate this with a few examples (which will be useful later).

$$(Ex6.1) \quad a.b.c \& b.c.d \& a.d.e \Rightarrow b.c.e$$

If not, we can suppose $c\uparrow, b\downarrow$ and $e\downarrow$. From $a.b.c$, we get $a\downarrow$. From $b.c.d$, we get $d\uparrow$. But from $a.d.e$, we get $d\downarrow$, giving a contradiction.

$$(Ex6.2) \quad x.z.y \& x.a.c \& y.b.c \& z.c.a \Rightarrow a.c.b$$

If not, we can suppose $c\uparrow, a\downarrow$ and $b\downarrow$. From $x.a.c$ and $y.b.c$, we get $x\downarrow$ and $y\downarrow$. Now $x.z.y$ gives $z\downarrow$ and so $z.c.a$ gives $c\downarrow$, a contradiction. (We could alternatively use

(Ex6.1) to deduce $a.c.y$, which together with $y.b.c$ gives $a.c.b.y$.)

$$(Ex6.3) \quad abc.y.def \Rightarrow y = (ady)(bey)(cfy)$$

If not, we can suppose $y\uparrow$, $ady\downarrow$ and $bey\downarrow$. Then $a\downarrow$, $d\downarrow$, $b\downarrow$ and $e\downarrow$, and so $abc\downarrow$ and $def\downarrow$, so $y\downarrow$, a contradiction.

We can also eliminate variables:

$$(Ex6.4) \quad aa'x = bb'y \Rightarrow b.aa'b.b'$$

If not, we can suppose $aa'b\uparrow$, $b\downarrow$ and $b'\downarrow$. Thus, $aa'x = bb'y\downarrow$, so either $a\downarrow$ or $a'\downarrow$, so $aa'b\downarrow$, giving a contradiction.

Of course, one can derive these statements directly from the axioms, or identities we have already established. Some of the above examples are straightforward. Others are more challenging. For example, we can derive (Ex6.4) above as:

$$\begin{aligned} (aa'b)bb' &= ((aa'b)bb')((aa'b)bb')y = (((aa'b)bb')(aa'b)b')(b(aa'b)b')y \\ &= (((aa'b)bb')by)(aa'b)b' = ((bb'y)(aa'b)b)(aa'b)b' \\ &= ((aa'x)(aa'b)b)(aa'b)b' = (aa'(xbb))(aa'b)b' \\ &= (aa'b)(aa'b)b' = aa'b. \end{aligned}$$

6.3. Explicit derivations of identities.

In summary, the above discussion gives a means of checking whether two expressions are tautologically equivalent. It also shows that a tautological identity can be derived by a finite sequence of applications of axioms (M1) and (M2). However, the argument made appeal to Zorn's lemma (via Proposition 3.2.10), and is thus non-constructive. It gives a proof, within the framework of ZFC set theory, of the formal statement that "there exists a derivation from the axioms of a median algebra". It does not exhibit such a derivation (nor even a proof in the first-order theory of median algebras). If we want an explicit recipe to do this, we need another approach.

In this subsection, we start again, avoiding any use of Lemma 3.2.12. In the process, explain how our constructions can be made explicit. The basic idea will be to reduce to the case of a free distributive lattice. The general idea was hinted at at the end of Subsection 5.2.

Let $\mathcal{E}(X)$ be the set of median expressions in a formal alphabet, X . For our purposes, there is no loss in assuming X to be finite. Given $E \in \mathcal{E}(X)$, and a map $\epsilon : X \rightarrow M$ to a median algebra, M , write $E(\epsilon) \in M$ for the result of evaluating E in M under this substitution. We say that $E, F \in \mathcal{E}(X)$ **coevaluate** in M if $E(\epsilon) = F(\epsilon)$ for all maps $\epsilon : X \rightarrow M$. Define two equivalence relations on $\mathcal{E}(X)$ by writing $E \simeq F$ and $E \approx F$ to mean respectively that E, F coevaluate in $I := \{0, 1\}$ and in any median algebra M . We also recall the relation, \sim , defined in Subsection 6.1. Clearly, $E \sim F \Rightarrow E \approx F \Rightarrow E \simeq F$.

We have seen, using Lemma 3.2.12, that $E \simeq F \Rightarrow E \approx F$; and the discussion of Subsection 6.1 tells us that $E \approx F \Rightarrow E \sim F$. Together this gives:

Proposition 6.3.1. $E \simeq F \Rightarrow E \sim F$.

We now aim to give a different, constructive, proof of this fact, and to relate it to the superextension, $\Phi(X)$.

Recall that $\Phi(X)$ is a subalgebra of $\mathcal{P}(\mathcal{P}(X)) \equiv I^{\mathcal{P}(X)}$ (see Subsection 5.2). Suppose $A \subseteq \mathcal{P}(X) \equiv I^X$. Thinking of A as a map $X \rightarrow I$ (namely its characteristic function), then given any $E \in \mathcal{E}(X)$, we can evaluate under this substitution to give us $E(A) \in I$. This gives us a map, $\alpha : \mathcal{E}(X) \rightarrow I^{\mathcal{P}(X)} \equiv \mathcal{P}(\mathcal{P}(X))$, and one readily checks that its image lies in $\Phi(X)$. (This can be verified directly. Alternatively, recall that $\Phi(X) \leq \mathcal{P}(\mathcal{P}(X))$ is a subalgebra, and that the image, $\alpha(x)$, of a one-letter expression, x , is the principal ultrafilter on $x \in X$, hence lies in $\Phi(X)$.) So, in fact, we get a map $\alpha : \mathcal{E}(X) \rightarrow \Phi(X)$. By construction, we see that $\alpha(E) = \alpha(F) \Leftrightarrow E \simeq F$. Moreover, if $P, Q, R \in \mathcal{E}(X)$, then $\alpha(PQR) = \alpha(P)\alpha(Q)\alpha(R)$. Also, Theorem 5.2.3 tells us that α is surjective to $\Phi(X)$ — a fact we will return to later.

Now let X^+ be the formal alphabet $X^+ := X \sqcup \{p, q\}$, obtained by adjoining two new symbols, p, q , not featuring in X . We can identify $\mathcal{E}(X)$ as a subset of $\mathcal{E}(X^+)$. We define a map, $\beta : \Phi(X) \rightarrow \mathcal{E}(X^+)$ as follows.

Recall that $\mathcal{P}(\mathcal{P}(X))$ has the structure of a bounded distributive lattice, where $\mathcal{A} \wedge \mathcal{B} = \mathcal{A} \cap \mathcal{B}$ and $\mathcal{A} \vee \mathcal{B} = \mathcal{A} \cup \mathcal{B}$ for $\mathcal{A}, \mathcal{B} \in \mathcal{P}(\mathcal{P}(X))$. Given $\mathcal{A} \in \mathcal{P}(\mathcal{P}(X))$, we define a formal expression in the binary operations, \wedge, \vee , namely $\theta(\mathcal{A}) := \bigvee_{A \in \mathcal{A}} \bigwedge_{x \in A} x$ (cf. the discussion of free distributive lattices in Subsection 2.2). To be precise, we should specify how we bracket the terms \bigvee and \bigwedge , but one can always specify some canonical way of doing this, after taking some fixed linear order on X . (The precise specification is not important, since we will see that we can freely apply the associative laws for \wedge and \vee .) Note that $\theta(\mathcal{A})$ and $\theta(\mathcal{M}(\mathcal{A}))$ coevaluate in I , where $\mathcal{M}(\mathcal{A}) \subseteq \mathcal{A}$ is the set of elements minimal with respect to inclusion. We now take the expression, $\theta(\mathcal{M}(\mathcal{A}))$, and first replace each occurrence of the argument $x \in X$ by the median expression (xpq) . We then recursively replace each subexpression $C \wedge D$ by CDp and each subexpression $C \vee D$ by CDq . We denote the resulting expression by $\beta(\mathcal{A})$. (As an illustration, suppose $X = \{a, b, c, d\}$, and $\mathcal{A} = \{\{a, b, c\}, \{a, b, d\}, \{a, d\}\}$. Then $\theta(\mathcal{A}) = (a \wedge b \wedge c) \vee (a \wedge b \wedge d) \vee (a \wedge d)$, $\theta(\mathcal{M}(\mathcal{A})) = (a \wedge b \wedge c) \vee (a \wedge d)$, and $\beta(\mathcal{A}) = (((apq)(bpq)p)(cpq)p)((apq)(dpq)p)q$.) This now gives us a map, β , defined on $\mathcal{P}(\mathcal{P}(X))$, which we restrict to give us a map, $\beta : \Phi(X) \rightarrow \mathcal{E}(X^+)$.

We have a relation, \sim , defined on $\mathcal{E}(X^+)$, which for clarity, we will denote by \sim^+ . (Given that $\Phi(X)$ embeds in $\Phi(X^+)$ it turns out that \sim^+ extends \sim on $\mathcal{E}(X)$, but we don't need to know that for the present argument.)

We claim:

Lemma 6.3.2. *If $E \in \mathcal{E}(X)$, then $\beta\alpha E \sim^+ Epq$.*

We will give a constructive proof of this shortly. First, we note a few consequences.

Given $E, P, Q \in \mathcal{E}(X)$, write $\gamma(E, P, Q)$ for the result of substituting each occurrence of p, q in $\beta\alpha E$ respectively by the expressions P, Q . Note that $E \simeq F \Rightarrow \gamma(E, P, Q) = \gamma(F, P, Q)$ (since $\alpha E = \alpha F$). Now the fact that $\beta\alpha E \sim^+ Epq$ also shows that in $\mathcal{E}(X)$, we have:

Lemma 6.3.3. $\gamma(E, P, Q) \sim EPQ$.

This is simply a matter of performing the above substitutions throughout the process of getting from $\beta\alpha E$ to Epq in $\mathcal{E}(X^+)$.

In particular, it follows that $E \simeq F \Rightarrow EPQ \sim FPQ$. This now constructively proves Proposition 6.3.1 since we have $E \sim EEF \sim FEF \sim F$.

We can also give a constructive proof of Proposition 3.3.3. In fact, we can explicitly extend any map, $\phi : X \rightarrow M$ to a map $\hat{\phi} : \Phi(X) \rightarrow M$, in any median algebra, M , thereby showing constructively that $\Phi(X)$ is indeed the free median algebra on X .

To this end, we first recall that $\alpha : \mathcal{E}(X) \rightarrow \Phi(X)$ is surjective. In fact, given any $\mathcal{A} \in \Phi(X)$, we can explicitly express \mathcal{A} as a median expression, E in X , so that $\mathcal{A} = \alpha(E)$. (One way of doing this is described in Subsection 20.1.) We can now set $\hat{\phi}(\mathcal{A}) = E(\phi)$, i.e. evaluating E in M under the substitution ϕ . Note that if $\alpha(F) = \alpha(E)$, then $E \simeq F$, so Proposition 6.3.1 tells us that $E \approx F$, so $E(\phi) = F(\phi)$. Thus, retrospectively, we could have chosen any such F with the same result.

Now if $x \in X$, where X is identified as a subset of $\Phi(X)$, then we can assume that we have chosen E to be just the one-letter expression x , and it follows that $\hat{\phi}$ does indeed extend ϕ . We finally need to check that $\hat{\phi}$ is a homomorphism. Suppose $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \Phi(X)$. We have chosen particular $P, Q, R \in \mathcal{E}(X)$ with $\mathcal{A} = \alpha(P)$, $\mathcal{B} = \alpha(Q)$ and $\mathcal{C} = \alpha(R)$. Now $\mathcal{A}\mathcal{B}\mathcal{C} = \alpha(PQR)$, and so by a previous observation, we have $\hat{\phi}(\mathcal{A}\mathcal{B}\mathcal{C}) = (PQR)(\phi) = P(\phi)Q(\phi)R(\phi) = \hat{\phi}(\mathcal{A})\hat{\phi}(\mathcal{B})\hat{\phi}(\mathcal{C})$ as required.

We now set about proving Lemma 6.3.2.

The idea is to reinterpret various lemmas we have already proven as statements in $\mathcal{E}(X^+)$ modulo the relation \sim^+ . This is justified by replacing equality by \sim^+ everywhere in their respective proofs. Thus, for example, we have a version of (M3), namely $AB(CDE) \sim^+ (ABC)(ABD)E$: replace $=$ by \sim^+ everywhere in the derivation presented in Subsection 4.2. (Of course, $P.Q.R$ should be interpreted as $PQR \sim^+ Q$.)

We proceed as follows. Given any $C \in \mathcal{E}(X^+)$, let $C' = Cpq \in \mathcal{E}(X^+)$. Let $X' = \{a' \mid a \in X\}$. Viewing X' , as a formal alphabet, we write $\mathcal{E}(X')$, for the set of expressions with arguments in X' . Viewing each element of X' instead as an expression in $\mathcal{E}(X^+)$, we can identify $\mathcal{E}(X')$ as a subset of $\mathcal{E}(X^+)$. As with the long distributive law (Lemma 3.2.6), we have $(CDE)' \sim^+ C'D'E'$ for all $C, D, E \in \mathcal{E}(X^+)$ (as seen by applying the principle described above). By iterating this, we

see that $C' \sim^+ \hat{C}$, where \hat{C} is the result of replacing x by x' for all arguments, x , in C .

We say that $C \in \mathcal{E}(X^+)$ lies “between” p and q if $Cpq \sim^+ C$ (and assume that we have found an explicit derivation of this fact). Note that if $C \in \mathcal{E}(X^+)$, then $C'pq = (Cpq)pq \sim^+ Cpq = C'$, so C' lies between p and q . We can now apply the discussion of intervals in a median algebra, as in Subsection 3.2. Given $C, D \in \mathcal{E}(X^+)$, write $C \wedge D = CDp$ and $C \vee D = CDq$. Note that if $C \sim^+ Cpq$ then $(C \wedge D)pq = (CDp)pq \sim^+ (Cqp)Dp \sim^+ CDp = C \wedge D$, and similarly for \vee . Therefore these relations (verifiably) preserve the property of lying between p and q . Now, as in Lemma 3.2.5, we have an explicit derivation showing that for $C, D, E \in \mathcal{E}(X^+)$ between p and q , we have $CDE \sim^+ (C \wedge D) \vee (D \wedge E) \vee (E \wedge C)$. Similarly, the axioms of a distributive lattice hold modulo \sim^+ . For example, we have $(C \wedge D) \wedge E \sim^+ C \wedge (D \wedge E)$ and $(C \wedge D) \vee E \sim^+ (C \vee E) \wedge (D \vee E)$ etc., for C, D, E between p and q .

Suppose now we start with some $E \in \mathcal{E}(X)$. We saw above that that $E' \sim^+ \hat{E} \in \mathcal{E}(X')$. We rewrite the latter in terms of the operations \wedge and \vee , as above, again with arguments in X' . Applying the distributive laws, we can put our expression into a normal form, namely $\bigvee_{B \in \mathcal{B}} \bigwedge_{x \in B} x'$, for some Sperner family, \mathcal{B} (as with the free distributive lattice, modulo \sim^+). We can also assume that the terms \bigvee and \bigwedge are bracketed in a standard way (as in the earlier discussion) since the associative laws correspond to applying (M2). This gives us an expression $\tilde{E} \sim^+ E'$.

We claim:

Lemma 6.3.4. $\tilde{E} = \beta\alpha E$.

First, we elaborate on some of our earlier definitions. Given a map $\phi : X \rightarrow I$, write $A_\phi := \{x \in X \mid \phi(x) = 1\}$. In other words, ϕ is the characteristic function of A_ϕ , and the map $[\phi \mapsto A_\phi] : X^I \rightarrow \mathcal{P}(X)$ gives us our identification of X^I with $\mathcal{P}(X)$.

Given an expression in \wedge, \vee and $\phi \in X^I$, we can evaluate the expression in the lattice, I . We can therefore speak of two such expressions “coevaluating” in I , and write \simeq for this relation, as with median expressions.

If $\mathcal{A} \in \mathcal{P}(\mathcal{P}(X))$ and $\phi \in X^I$, then $\theta(\mathcal{A})(\phi) = 1$ if and only if there is some $B \in \mathcal{A}$ with $B \subseteq A_\phi$. Thus, if $\mathcal{A} \in \Phi(X)$, then $\theta(\mathcal{A})(\phi) = 1 \Leftrightarrow A_\phi \in \mathcal{A}$. By definition of α , if $E \in \mathcal{E}(X)$, then $A_\phi \in \alpha E \Leftrightarrow E(\phi) = 1$, so $\theta(\alpha E)(\phi) = 1 \Leftrightarrow E(\phi) = 1$. In other words, $\theta(\alpha E)$ coevaluates with E in I .

Given $\mathcal{A} \in \mathcal{P}(\mathcal{P}(X))$, write $\mathcal{M}(\mathcal{A}) \subseteq \mathcal{A}$ for the set of elements of \mathcal{A} minimal with respect to inclusion. Thus, $\mathcal{M}(\mathcal{A})$ is a Sperner family. Moreover, $\theta(\mathcal{A}) \simeq \theta(\mathcal{M}(\mathcal{A}))$. Recall that, by Lemma 2.2.2, if $\mathcal{B}, \mathcal{C} \in \mathcal{P}(\mathcal{P}(X))$ are Sperner families with $\theta(\mathcal{B}) \simeq \theta(\mathcal{C})$, then $\mathcal{B} = \mathcal{C}$.

Now, given $E \in \mathcal{E}(X)$, we have noted that E coevaluates in I with $\theta(\alpha E) \simeq \theta(\mathcal{M}(\alpha E))$. Thus, if \mathcal{B} is a Sperner family such that $\theta(\mathcal{B})$ coevaluates with E , then $\mathcal{B} = \mathcal{M}(\mathcal{A})$.

Proof of Lemma 6.3.4. By construction, $\beta\alpha E$ has the form $\bigvee_{A \in \mathcal{M}(\mathcal{A})} \bigwedge_{x \in A} x'$, where $\mathcal{A} = \alpha(E) \in \Phi(X)$; and \tilde{E} has the form $\bigvee_{B \in \mathcal{B}} \bigwedge_{x \in B} x'$, for a Sperner family \mathcal{B} . We therefore need to check that $\mathcal{B} = \mathcal{M}(\mathcal{A})$. By the above discussion, it is enough to show that $\theta(\mathcal{B})$ coevaluates with E .

To see this, we extend any map $X \rightarrow I$ to a map $X^+ \rightarrow I$, by setting $p = 0$ and $q = 1$. In this way, \tilde{E} coevaluates with $\theta(\mathcal{B})$, and E coevaluates with E' . Now $\tilde{E} \sim^+ E'$, so certainly these expressions coevaluate in I (whatever values we assign to p, q). Therefore $\theta(\mathcal{B})$ coevaluates with E as required. \square

This now proves Lemma 6.3.2: we have $E p q = E' \sim^+ \tilde{E} = \beta\alpha E$.

We note that this gives an explicit way of deriving identities (or disproving them). Suppose $E, F \in \mathcal{E}(X)$, and we have checked that $E \simeq F$. Then following through on the above, we can explicitly derive the fact that $E \sim E E F \sim \gamma(E, E, F) = \gamma(F, E, F) \sim F E F \sim F$, directly applying the axioms (M1) and (M2).

We can similarly derive conditional identities as follows.

Note that any element of $\mathcal{E}(X)$ is \sim -equivalent to an expression of bounded length. We can choose a particular ‘‘canonical’’ representative for each \sim -class (e.g. minimal with respect to some measure of complexity). The above gives a systematic means of putting a median expression into its canonical form. We write $\mathcal{E}_0(X) \subseteq \mathcal{E}(X)$ for the set of such canonical forms.

We can apply this to median identities more generally. We can represent any median identity in the form $A \sim B$, for $A, B \in \mathcal{E}_0(X)$. (Note in particular there are only finitely many such identities.) By the above, the equivalence of such an identity with the original is explicitly derivable.

We can now give a procedure for explicitly deriving conditional median identities. Let I_1, \dots, I_n be a set of input identities. We write I_i as $P_i \sim Q_i$ for $P_i, Q_i \in \mathcal{E}_0(X)$. We define a sequence of equivalence relations, $(\equiv_n)_{n \in \mathbb{N}}$, on $\mathcal{E}_0(X)$ inductively as follows. Let \equiv_0 be the equivalence relation generated by the relations $P_i \equiv_0 Q_i$ for all i . Suppose we have defined \equiv_n . Given $A, B \in \mathcal{E}_0(X)$, say A is related to B if there exist $C, D, P, Q \in \mathcal{E}_0(X)$ with $A \sim C D P$, $B \sim C D Q$ and $P \equiv_n Q$. Let \equiv_{n+1} be the transitive closure of this relation. Note that if $A \equiv_n B$, then since $A \sim A B A$ and $B \sim A B B$, we have $A \equiv_{n+1} B$. Therefore the relations \equiv_n stabilise, after a bounded number of steps, on an equivalence relation, \equiv , on $\mathcal{E}_0(X)$. Since reduction to $\mathcal{E}_0(X)$, and the relation \sim are explicitly derivable, so is \equiv .

We claim that this precisely describes those identities which are consequences of the input identities. To see this, note that if $A, B, C, P, Q, R \in \mathcal{E}_0(X)$ with $A \equiv P$, $B \equiv Q$ and $C \equiv R$, then $A B C \equiv P Q R$. We therefore get a ternary operation induced on $\mathcal{E}_0(X)/\equiv$. With this structure, $\mathcal{E}_0(X)/\equiv$ is a median algebra where all the identities I_1, \dots, I_n are satisfied. Moreover, it is universal with this property. We see that if $A, B \in \mathcal{E}_0(X)$, then $A \equiv B$ if and only if the identity $A \sim B$ is tautologically implied by the input identities. Moreover, we see that this can be

explicitly derived from the axiom (M1) and (M2) together with the input identities.

From the above, we deduce:

Proposition 6.3.5. *There is a primitive recursive algorithm which takes as input a finite set of ternary identities I_1, \dots, I_n, I , and either proves or refutes the statement that I is tautologically implied by $I_1 \& \dots \& I_n$. In the former case, it provides an explicit derivation of this fact in terms of the axioms (M1) and (M2).*

Here “primitive recursive” means that it terminates in primitive recursive time as a function of complexity of the input. The latter might be measured as the sum of the lengths of the identities (though any sensible measure would give an equivalent statement). In fact the time of our procedure is bounded by some tower of exponentials.

In contrast, we remark that the general first-order theory of median algebras is undecidable (see the Notes to this section).

7. CONVEX SETS

Convexity is a central notion in the theory of median algebras. Indeed the subject is sometimes approached from the point of view of convex structures. We have briefly discussed convex sets in Section 3. Here we start again with a more systematic treatment.

A fundamental fact is the Helly Property (Lemma 7.1.1). We digress slightly to discuss the notions of parallel sets and translations. We say more about gate maps to convex sets, which were defined briefly in Subsection 3.2. We give a number of descriptions of the convex hull of a set, and show how this behaves when passing to subalgebras. We finish the section with a few additional observations that will be used later.

It will be noted that some definitions and discussion only make direct reference to the betweenness relation. Some of this can be set in a more general context, as we discuss in Section 22. Even if that is the case, the statements can often be made stronger, and the proofs simpler, in the case of median algebras. We will restrict attention to that case in this section.

7.1. Some definitions and basic facts.

Let M be a median algebra. From Section 3, we recall:

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

Clearly any convex set is a subalgebra. The intersection of any family of convex sets is convex. Also, an interval, $[a, b] \subseteq M$ is itself convex.

A useful property is the “Helly Property”:

Lemma 7.1.1. *Let C_1, \dots, C_n be a non-empty finite family of pairwise intersecting convex subsets of M . Then $\bigcap_{i=1}^n C_i \neq \emptyset$.*

Proof. For $n = 3$, choose points respectively in $C_1 \cap C_2$, $C_2 \cap C_3$ and $C_3 \cap C_1$. Their median lies in $C_1 \cap C_2 \cap C_3$.

We now proceed by induction. Suppose it holds for a given $n \geq 3$. Let C_1, C_2, \dots, C_{n+1} pairwise intersect. From the case $n = 3$, we see that the sets $C_1 \cap C_{n+1}, C_2 \cap C_{n+1}, \dots, C_n \cap C_{n+1}$ pairwise intersect, and we apply the inductive hypothesis. \square

We can describe an algebra of non-empty convex sets as follows.

Given subsets, $A, B, C \subseteq M$, write $ABC = \{abc \mid a \in A, b \in B, c \in C\}$ (see Example (Ex3.4) of Subsection 3.4).

Lemma 7.1.2. *If A, B, C are convex, then so is ABC .*

Proof. Let $a, b, c, a', b', c' \in M$. Let $x = abc$ and $x' = a'b'c'$, and suppose $y \in [x, x']$. In Example (Ex6.3) in Subsection 6.2, we showed that $y = (aa'y)(bb'y)(cc'y) \in ABC$. \square

Here is another way of interpreting Lemma 7.1.2.

Given sets, $A, B \subseteq M$, as in Subsection 3.2 we write $J(A, B) = \bigcup_{a \in A, b \in B} [a, b]$. Recall that Lemma 3.2.9 tells us that if A, B are convex, so is $J(A, B)$.

Lemma 7.1.3. *If $A, B, C \subseteq M$, then $ABC = J(A, B) \cap J(B, C) \cap J(C, A)$.*

Proof. The inclusion $ABC \subseteq J(A, B) \cap J(B, C) \cap J(C, A)$ is clear. For the reverse inclusion, let $x \in J(A, B) \cap J(B, C) \cap J(C, A)$. That is, $x = ab'x = bc'x = ca'x$ for some $a, a' \in A, b, b' \in B$, and $c, c' \in C$. We claim that $x = (aa'x)(bb'x)(cc'x)$. For if not, we can suppose $x \uparrow, aa'x \downarrow, bb'x \downarrow$, so $a \downarrow, b' \downarrow$, so $x \downarrow$, giving a contradiction. We see that $x \in ABC$ as required. \square

From Lemmas 3.2.9 and 7.1.3, we see again that if A, B, C are convex, then so is ABC .

Although we won't be using it, we note that the join operation is associative:

Lemma 7.1.4. *For all $A, B, C \subseteq M$, $J(J(A, B), C) = J(A, J(B, C))$.*

Proof. By symmetry, it is enough to verify the inclusion \subseteq . To this end, let $y \in [x, c]$ where $x \in J(A, B)$ and $c \in C$. Then $x \in [a, b]$ for some $a \in A$ and $b \in B$. Let $z = bcy \in J(B, C)$. We claim that $y \in [a, z]$. For if not, we can suppose $a \downarrow, z \downarrow$ and $y \uparrow$. From $z = bcy$, we get $b \downarrow$ and $c \downarrow$. From $a.x.b$, we get $x \downarrow$. From $x.y.c$, we get the contradiction $y \downarrow$. This shows that $y \in [a, z] \subseteq J(A, J(B, C))$ as required. \square

Let $\mathcal{K} = \mathcal{K}(M)$ be the median algebra of all non-empty convex subsets of M , as defined by Example (Ex3.4). As usual, we write $[A, B]_{\mathcal{K}} \subseteq \mathcal{K}$ for the median interval between the convex sets, A, B .

The following are easily checked for non-empty convex sets: $A \cap B \in [A, B]_{\mathcal{K}}$ (if $A \cap B \neq \emptyset$), $J(A, B) \in [A, B]_{\mathcal{K}}$, $\bigcup [A, B]_{\mathcal{K}} \subseteq J(A, B)$, and if $A \subseteq C \subseteq B$ then $C \in [A, B]_{\mathcal{K}}$.

(These observations will be useful when we come to discuss cube complexes, in Sections 10 and 11.)

We note that the preimage of any convex set under a homomorphism is convex. Also:

Lemma 7.1.5. *Let $\phi : M \rightarrow N$ be an epimorphism between two median algebras, M and N . Then $\phi([x, y]) = [\phi x, \phi y]$ for all $x, y \in M$.*

Proof. Let $\phi : M \rightarrow N$ be an epimorphism between two median algebras, M and N . Certainly $\phi([x, y]) \subseteq [\phi x, \phi y]$. For the reverse inclusion, and let $a \in [\phi x, \phi y] \subseteq N$. Then $a = \phi z$ for some $z \in M$. Let $w = xyz \in [x, y]$. Then $\phi w \in [\phi x, a] \cap [a, \phi y] = \{a\}$, so $\phi w = a$, so $a \in \phi([x, y])$ as required. \square

In particular, it follows that the image of a convex set under an epimorphism is convex.

7.2. Parallel sets and translations.

Before continuing with convex sets, we make a digression concerning parallel sets and translations.

Let M be a median algebra. Given $a, b, c, d \in M$, we say that a, b, c, d forms a **rectangle** if $a.b.c$ & $b.c.d$ & $c.d.a$ & $d.a.c$. It is readily checked that this is equivalent to saying that $[a, c] = [b, d]$.

Given pairs a, a' and b, b' in M , we write $aa' \parallel bb'$ to mean that a, a', b', b forms a rectangle. This is equivalent to saying that $ab \parallel a'b'$.

Definition. We say that the pairs a, a' and b, b' are **parallel** if $aa' \parallel bb'$.

Note that if $aa' \parallel aa''$, then $a' = a''$.

Lemma 7.2.1. *Parallelism is an equivalence relation on the set of pairs of M .*

Proof. Suppose $aa' \parallel bb' \parallel cc'$. We have $b.a.a'$ & $a.a'.b$ & $a'.b'.c$ & $c.b.b'$. By Example (Ex6.1) of Subsection 6.2, we have $c.a.a'$. By symmetry, swapping a with a' , and a with c , we also have $c.a'.a$, $a.c.c'$ and $a.c'.c$. In other words, $aa' \parallel cc'$. \square

Lemma 7.2.2. *Suppose $aa' \parallel bb'$ and $x \in M$. Let $d = abx$ and $d' = a'b'x$. Then $dd' \parallel aa'$.*

Proof. Note that $abd' = (aba')(abb')x = bax = d$. Similarly, $a'b'd = d'$. Now $b.a.a'$ & $b.d.a$ gives $b.d.a.a'$ so $d.a.a'$. Similarly, $d'.a'.a$. Also, $abd' = d$ gives $a.d.d'$. Similarly $a'.d'.d$. It follows that $dd' \parallel aa'$. \square

Lemma 7.2.3. *Suppose $a, a', b, b', c, c' \in M$ with $aa' \parallel bb' \parallel cc'$. Then $(abc)(a'b'c') \parallel aa'$.*

Proof. Let $d = abc$. Since $abd = d$, we see by Lemma 7.2.2 that $d(a'b'd) \parallel aa' \parallel bb'$. Similarly, $d(b'c'd) \parallel bb' \parallel cc'$ and $d(c'a'd) \parallel cc' \parallel aa'$. Thus $d(a'b'd) \parallel d(b'c'd) \parallel d(c'a'd)$. It follows that $a'b'd = b'c'd = c'a'd$. Writing d' for this element, we have $dd' \parallel aa'$. Moreover, we have $a.d'.b$ & $b.d'.c$ & $c.d'.a$ and so $d' = a'b'c'$. \square

Lemma 7.2.3 can be expressed by saying that any parallel class, viewed as a subset of $M \times M$, is a subalgebra of $M \times M$ with the product median structure.

Definition. Given two subsets, $A, B \subseteq M$, a **translation** from A to B is a bijection, $\tau : A \rightarrow B$, such that for all $a, b \in A$, $ab \parallel (\tau a)(\tau b)$. We say that A, B be are **parallel** if there is a translation between them. We write $A \parallel B$.

Clearly the inverse of a translation is a translation. By Lemma 7.2.1, it is also closed under composition. Therefore parallelism of sets is an equivalence relation. If $a \in A$ and $b \in B$, then $\tau a \in [a, b]$ (since $a(\tau^{-1}b) \parallel (\tau a)b$). It follows that the translation, $\tau : A \rightarrow B$ is unique. (For if τ' were another, we would have $a.\tau a.\tau'a$ and $a.\tau'a.\tau a$, so $\tau a = \tau'a$ for all $a \in A$.) It follows that any self-translation of A is the identity. If $A \parallel B \parallel C$, then translating from A to B then to C is the same as translating directly from A to C . If $A \parallel B$ and $A \cap B \neq \emptyset$, then $A = B$. (For let $b \in A \cap B$. If $a \in A$, then $a.\tau a.b$, and by the above, $\tau a.a.b$. Therefore $a = \tau a \in B$.) We also note from Lemma 7.2.3 that if $A \parallel B$ and A is a subalgebra of M , then so is B , and the translation between them is an isomorphism.

As an example, if $a, b \in M$ and $aa' \parallel bb'$, then $[a, a'] \parallel [b, b']$. The translations between them are the gate maps $[x \mapsto bb'x]$ and $[y \mapsto aa'y]$.

Let \mathcal{A} be a parallel class non-empty subsets of M . We fix some arbitrary representative, say $P \in \mathcal{A}$. For any $A \in \mathcal{A}$, we have a translation, $\tau_A : P \rightarrow A$. As observed above, this is a median isomorphism.

Let $p \in P$, and write $T_p = \{\tau_{AP} \mid A \in \mathcal{A}\}$. Thus T_p is a transversal of \mathcal{A} .

Lemma 7.2.4. T_p is convex in M .

Proof. Let $a \in [\tau_{AP}, \tau_{BP}]$ where $A, B \in \mathcal{A}$. Define a map, $\theta : P \rightarrow M$ by $\theta(x) = (\tau_A x)(\tau_B x)a$. Thus $\theta(p) = a$. If $y \in P$, then $(\tau_A x)(\tau_B x) \parallel (\tau_A y)(\tau_B y)$. Equivalently, we have $(\tau_A x)(\tau_A y) \parallel (\tau_B x)(\tau_B y)$, and so by Lemma 7.2.2 we have $(\theta x)(\theta y) \parallel (\tau_A x)(\tau_A y) \parallel xy$. It follows that $\theta : A \rightarrow \theta(P)$ is a translation, so $\theta(P) \parallel P$, so $\theta(P) \in \mathcal{A}$ and $\theta = \tau_{\theta(P)}$. Therefore $a = \theta(p) \in T_p$ as required. \square

Lemma 7.2.5. If $q \in P$, then $T_p \parallel T_q$. Moreover, the map $[\tau_{AP} \mapsto \tau_{AQ}] : T_p \rightarrow T_q$ is a translation.

Proof. If $A, B \in \mathcal{A}$, then $(\tau_{AP})(\tau_{AQ}) \parallel (\tau_{BP})(\tau_{BQ})$, so $(\tau_{AP})(\tau_{BP}) \parallel (\tau_{AQ})(\tau_{BQ})$. \square

This gives \mathcal{A} the structure of a median algebra, where the median is induced from transversal T_p . That is, $(\tau_{AP})(\tau_{BP})(\tau_{CP}) = \tau_{ABCp}$. By Lemmas 7.2.5 and 7.2.3, this is independent of the choice of p .

Now let \mathcal{T} be the entire parallel class of T_p in M . Thus, $T_q \in \mathcal{T}$ for all $q \in P$. Each $A \in \mathcal{A}$ is a subset of a transversal, A' , of \mathcal{T} (constructed as above). Let $\mathcal{A}' = \{A' \mid A \in \mathcal{A}\}$. This is again a parallel class. By Lemmas 7.2.4 and 7.2.5 applied to \mathcal{T} , we see that each element of \mathcal{A}' is convex. For this reason, there is no essential loss in restricting attention to parallel classes of convex sets. In this case, \mathcal{A} will be a subalgebra of $\mathcal{K}(M)$ (see Example (Ex3.4)).

Let \mathcal{A} be a parallel class of convex sets. Thus $\bigcup \mathcal{A} \subseteq M$ is a disjoint union of the elements of \mathcal{A} .

Lemma 7.2.6. $\bigcup \mathcal{A}$ is convex in M .

Proof. Let $x \in [\tau_A a, \tau_B b]$, where $a, b \in P$ and $A, B \in \mathcal{A}$. Let $c = (\tau_A a)(\tau_B b)x \in A$ and $d = (\tau_B b)(\tau_A a)x \in B$. Let $q = abx$. By Lemma 7.2.2, $qc \parallel aa'$, so $c = \tau_A q$. Similarly $d = \tau_B q$, so $c, d \in T_q$. We claim that $x \in [c, d]$. If not, we can suppose $x \uparrow, c \downarrow, d \downarrow$. Also (up to swapping a with b , and A with B), we can suppose $\tau_A a \downarrow$. Thus $c \downarrow$, giving a contradiction. By Lemma 7.2.4, T_q is convex, and so $x \in [c, d] \subseteq T_q \subseteq \bigcup \mathcal{A}$. \square

We can put the product median structure on $\mathcal{A} \times P$.

Lemma 7.2.7. The map $\mathcal{A} \times P$ given by $[(A, a) \mapsto \tau_A a]$ is a median isomorphism.

Proof. By construction, the map is bijective. We therefore want to show that for all $x, y, z \in P$ and $A, B, C \in \mathcal{A}$, we have $(\tau_A x)(\tau_B y)(\tau_C z) = \tau_{ABC}(xyz)$. If not, we can suppose that $\tau_{ABC}(xyz) \uparrow, \tau_A x \downarrow, \tau_B y \downarrow$. Now (from the definition of the median on \mathcal{A}) we have $\tau_{ABC}(xyz) = (\tau_A(xyz))(\tau_B(xyz))(\tau_C(xyz))$, so (up to swapping A and B) we can suppose $\tau_A(xyz) \uparrow$. Now (by Lemma 7.2.3 and subsequent discussion) $\tau_A(xyz) = (\tau_A x)(\tau_A y)(\tau_A z)$. Given $\tau_A x \downarrow$, we have $\tau_A y \uparrow$. But $\tau_A x \cdot \tau_A y \cdot \tau_B y$, and $\tau_A x \downarrow, \tau_B y \downarrow$, so $\tau_A y \downarrow$ giving a contradiction. \square

7.3. Gates.

We now move on to consider ‘‘gates’’ of convex sets. Much of this can be set in a more general context, and also applied to metric spaces, as we discuss in Section 22.

Let $A \subseteq M$ be convex.

Definition. Let $a \in M$. We say that $b \in A$ is a **gate** for a in A if $[a, b] \cap A = \{b\}$.

This is equivalent to saying that for all $c \in A$, $a.b.c$ holds (since $abc \in [a, b] \cap A$). It follows that if a gate for a exists, then it is unique. Note also that if $a \in A$, then a is a gate for a in A .

Definition. A convex set, A , is **gated** if each point of M has a gate in A .

In this case, given $a \in M$, we write $\omega a = \omega_A a$ for the gate of a in A . This gives us a map $\omega : M \rightarrow A$. It is called the **gate map** to A . Note that ω satisfies $\omega^2 = \omega$ and $a.\omega a.\omega b$ for all $a, b \in M$. (In fact, one can readily check for any such map, $\omega(M)$ is convex, and ω is the gate map to $\omega(M)$.)

Lemma 7.3.1. A gate map, $\omega : M \rightarrow A$, is a homomorphism.

Proof. From the median rule (see Lemma 3.2.1) it is enough to show that for all $x, y \in M$, we have $\omega([x, y]) \subseteq [\omega x, \omega y]$. Write $a = \omega x, b = \omega y$. Let $z \in [x, y]$ and write $c = \omega z$. Then $x.z.y$ & $x.a.c$ & $y.b.c$ & $z.c.a$. Example (Ex6.2) of Subsection 6.2 tells us that this implies $a.c.b$ as required. \square

(One also readily checks that if $\omega : M \rightarrow M$ is any homomorphism with $\omega^2 = \omega$ and ωM convex, then ω is gate map to ωM .)

Given a convex set, $B \subseteq M$, write $A_B = \omega_A B \subseteq A$.

Lemma 7.3.2. *A_B is convex in M .*

Proof. $\omega_A : B \rightarrow A$ is an epimorphism to A , so by Lemma 7.1.5, A_B is convex in A , hence also convex in M . \square

Now suppose that $B \subseteq M$ is also gated. We write $B_A = \omega_B A \subseteq B$.

Lemma 7.3.3. *$\omega_A \omega_B : M \rightarrow A_B$ is a gate map to A_B .*

Proof. Let $x \in M$ and $c \in A_B$. Write $b = \omega_B x \in B$, $a = \omega_A b \in A$, and choose $d \in B$ such that $c = \omega_A d$. We have $x.b.d$ & $b.a.c$ & $a.c.d$. We claim this implies $x.a.c$. If not, we can assume $a \uparrow$, $x \downarrow$, $c \downarrow$. Now $b.a.c$ gives $b \uparrow$, and $a.c.d$ gives $d \downarrow$, and then $x.b.d$ gives $b \downarrow$, which is a contradiction.

We have shown that $x.\omega_A \omega_B x.c$ holds for all $x \in M$ and $c \in A_B$, so $\omega_A \omega_B$ is a gate map. \square

Now $\omega_A \omega_B |_{A_B}$ is the identity on A_B . Similarly, $\omega_B \omega_A |_{B_A}$ is the identity on B_A . Therefore $\omega_B |_{A_B}$ and $\omega_A |_{B_A}$ are inverse isomorphisms between A_B and B_A . In fact:

Lemma 7.3.4. *$\omega_B |_{A_B}$ is a translation of A_B to B_A .*

Proof. Let $a, a' \in A_B$, and write $b = \omega_B a$ and $b' = \omega_B a'$. Since ω_B is a gate map, we have $a.b.b'$ and $a'.b'.b$. Also $a = \omega_A b$ and $a' = \omega_A b'$, and so $b.a.a'$ and $b'.a'.a$. In other words, $aa' \parallel bb'$. \square

This shows that $A_B \parallel B_A$. In particular, either $A_B = B_A$ or $A_B \cap B_A = \emptyset$. Since we always have $A \cap B \subseteq A_B \cap B_A$, it follows that if $A \cap B \neq \emptyset$, then $A \cap B = A_B = B_A$.

Note that as an immediate consequence, we have

Lemma 7.3.5. *If A, B are gated convex sets and $A \cap B \neq \emptyset$, then $A \cap B$ is gated and $\omega_A B = \omega_B A = A \cap B$. Moreover, $\omega_A \omega_B = \omega_B \omega_A : M \rightarrow A \cap B$ is the gate map to $A \cap B$.*

By the Helly Property (Lemma 7.1.1) we see that the intersection of any non-empty finite family of pairwise intersecting gated convex sets is gated convex.

Given convex sets, A, B , we say that $a \in A$ and $b \in B$ are **mutual gates** if b is a gate for a in B and a is a gate for b in A .

We note:

Lemma 7.3.6. *If $A, B \subseteq M$ are gated convex sets, then there are mutual gates $a \in A$ and $b \in B$.*

Proof. Choose any $a \in A_B$ and set $b = \omega_B a \in B_A$. \square

Lemma 7.3.7. *Any non-empty finite convex subset of a median algebra is gated.*

Proof. Let M be a median algebra, and $C \subseteq M$ be non-empty finite and convex. Let $a \in M$, and choose any $b \in A$. Thus, $[a, b]$ has a partial order, \leq , as defined in Subsection 3.2. Let $c \in C \cap [a, b]$ be minimal with respect to \leq . We claim that c is a gate for C . For suppose that $d \in C \cap [a, c]$. Let $m = c \wedge d = acd$. Then $m \in C \cap [a, b]$ and $m \leq c$. Therefore $m = c$. In other words, $a.c.d$. But $a.d.c$, and so $c = d$. This shows that $C \cap [a, c] = \{c\}$. \square

Of course, this applies to finite median algebras. We remark that the existence of mutual gates (Lemma 7.3.6) is simpler in this case. If A and B are convex, choose $a \in A$ and $b \in B$ so as to minimise $\#[a, b]$. Then a similar argument shows that a, b are mutual gates.

Recall that the set, $\mathcal{K} = \mathcal{K}(M)$, of all non-empty convex subsets of M is itself a median algebra (see Example (Ex3.4) of Subsection 3.4). We claim that gated convex sets form a subalgebra of $\mathcal{K}(M)$:

Lemma 7.3.8. *If A, B, C are gated, then so is ABC .*

Proof. Let $x \in M$, and let a, b, c be gates for A, B, C respectively. We claim that $d := abc$ is a gate for x in ABC . We need to check that $[x, d] \cap ABC = \{d\}$. So let $d' \in [x, d] \cap ABC$. Then $d' = a'b'c'$ for some $a' \in A, b' \in B$ and $c' \in C$. If $d \neq d'$, we can suppose that $d \uparrow$ and $d' \downarrow$. By $x.d'.d$, we get $x \downarrow$. Using $d = abc$ and $d' = a'b'c'$, we can suppose (up to permuting a, b, c) that $a \uparrow$ and $a' \downarrow$. But this contradicts $x.a.a'$. \square

For future reference (see Lemma 11.3.3) we also make the following observation.

Lemma 7.3.9. *Let $A, B \in \mathcal{K}$ and suppose that A is gated. Let $\omega_A : M \rightarrow A$ be the gate map, and write $C = \omega_A B \in \mathcal{K}$. Then $C \in [A, B]_{\mathcal{K}}$.*

Proof. In other words, we want to show that $ABC = C$.

To show $C \subseteq ABC$, choose any $b \in B$. If $c \in C$, then $c = cbc \in ABC$.

To show that $ABC \subseteq C$, let $a \in A, b \in B, c \in C$, and set $d = \omega_A b$ and $m = abc$. We have $a.d.b, c.m.a$ and $c.m.b$. By Lemma 3.2.7, this implies $c.m.d$. But $c, d \in C$, so since C is convex, $m \in C$. \square

We will see many more of examples of gated sets later (for example, in Sections 12 and 15).

7.4. Convex hulls.

We now move on to discuss joins and convex hulls.

Given a subset, $A \subseteq M$, write

$$J(A) = \bigcup_{a, b \in A} [a, b].$$

(Thus, $J(A) = J(A, A)$ as defined in Subsection 3.2. Clearly $J(A, B) \subseteq J(A \cup B)$. If A, B are convex, then $J(A, B) = J(A \cup B)$.)

Definition. $J(A)$ is the *join* of A .

Given $n \in \mathbb{N}$, we write $J^n(A)$ for the n th iterated join of A . In other words, $J^n(A)$ is defined by setting $J^0(A) = A$, and $J^{n+1}(A) = J(J^n(A))$.

Definition. The *convex hull*, $\text{hull}(A)$, of A is the intersection of all convex subsets of M containing A .

Thus, for example, $\text{hull}(\{a, b\}) = [a, b]$.

We also note that $\text{hull}(A) = \bigcup_{n=0}^{\infty} J^n(A)$. In general, one needs to take an infinite union. (For example, consider the convex hull of $\{0, 1\}^{\mathbb{N}}$ in $[0, 1]^{\mathbb{N}}$.) However, in certain cases, a finite union is sufficient. This is true, for example, if A is finite: as can be seen by inductively applying Lemma 3.2.9, or using Lemma 7.4.4 below. Another example is given by Proposition 8.2.3 in the next section.

From this description of convex hulls, we also see that if A, B are parallel, then so are $\text{hull}(A)$ and $\text{hull}(B)$.

We note the following elementary observation:

Lemma 7.4.1. $\text{hull}(A)$ is the union of all sets $\text{hull}(B)$ as B ranges over all finite subsets of A .

Proof. The union is convex, since if $x \in \text{hull}(B)$, $y \in \text{hull}(B')$, then $[x, y] \subseteq \text{hull}(B \cup B')$. \square

For future reference, we also note:

Lemma 7.4.2. If $A, B, C \subseteq M$, then $\text{hull}(ABC) = \text{hull}(A)\text{hull}(B)\text{hull}(C)$.

Proof. By Lemma 7.1.2, $\text{hull}(A)\text{hull}(B)\text{hull}(C)$ is convex, and clearly contains ABC , therefore the inclusion \subseteq holds.

For the inclusion \supseteq , we first claim that $J(A)BC \subseteq J(ABC)$. To this end, let $d \in J(A)$, $b \in B$, $c \in C$. Now $d = aa'd$ for some $a, a' \in A$. Thus, $dbc = (aa'd)bc = (abc)(a'bc)d \in [abc, a'bc] \subseteq J(ABC)$, as claimed. Iterating this we get $J^n(A)J^n(B)J^n(C) \subseteq J^{3n}(ABC) \subseteq \text{hull}(ABC)$. Thus, $\text{hull}(A)\text{hull}(B)\text{hull}(C) = \bigcup_{n=0}^{\infty} J^n(A)J^n(B)J^n(C) \subseteq \text{hull}(ABC)$ as required. \square

Note that one consequence is that the set of convex hulls of finite sets is a subalgebra of $\mathcal{K}(M)$.

Here is another description of convex hulls.

Recall that, given any finite non-empty subset, $B \subseteq M$, and $x \in M$, we have defined $(B|x) = (b_1 \dots b_p|x)$, where $B = \{b_1, \dots, b_p\}$ (see Subsection 5.1). From the definition, it is clear that $(B|x) \in J^{p-1}(B)$. We also note the following properties. For any x , $(B|(B|x)) = (B|x)$. If $x \in B$, then $(B|x) = x$. Also:

Lemma 7.4.3. If $B \subseteq M$ is finite and non-empty, and $x, y, z \in M$, then $(B|xyz) = (B|x)(B|y)z = (B|x)(B|y)(B|z)$.

Proof. This can be conveniently be thought of in terms of voting. Recall that in a vote with outcome $(B|a)$, the vote of a can only be overruled by the unanimous vote of B . Therefore, if B is not unanimous, then in each of the three expressions

in the identity, the outcome is determined by the majority vote of $\{x, y, z\}$. (More formally, this shows that the expressions take the same value whenever they are evaluated in $\{0, 1\}$.) \square

Lemma 7.4.4. *Let $B \subseteq M$ be finite and non-empty. Then $\text{hull}(B) = \{x \in M \mid (B|x) = x\}$. Moreover, $[x \mapsto (B|x)] : M \rightarrow \text{hull}(B)$ is a gate map.*

Proof. Write $n = \#B$. Let $H = \{x \in M \mid (B|x) = x\}$. Clearly $H \subseteq J^{n-1}(B) \subseteq \text{hull}(B)$. Moreover, H is convex, for if $x, y \in H$ and $z \in [x, y]$, then $(B|xyz) = (B|x)(B|y)z = xyz$, so $xyz \in H$. Since $(B|a) = a$ for all $a \in B$, we have $B \subseteq H$. Therefore $\text{hull}(B) \subseteq H$, and so $H = \text{hull}(B)$.

To see that $(B|x)$ is a gate for x , let $y = (B|x)$ and suppose $z \in [x, y] \cap H$. Then $z = (B|z) = B(xyz) = (B|x)(B|z)y = yzy = y$. Thus, $[x, y] \cap H = \{y\}$. \square

Therefore $\text{hull}(B)$ is gated. Also, we have shown that $\text{hull}(B) = J^{n-1}(B)$, where $n = \#B$. (One can no doubt do a lot better, though it's not clear what the optimal result is.)

Now let $A \subseteq M$ be any subset. We note that $\text{hull}(A)$ is the union of all sets $\text{hull}(B)$ as B ranges over all finite subsets of A . As a consequence, we see:

Lemma 7.4.5. *$x \in \text{hull}(A)$ if and only if $x = (B|x)$ for some finite subset $B \subseteq A$.*

As a corollary, suppose $N \leq M$ is a subalgebra, and $A \subseteq N$. Write $\text{hull}_N(A)$ for the intrinsic convex hull of A in N .

Lemma 7.4.6. *If $A \subseteq N$, then $\text{hull}_N(A) = N \cap \text{hull}_M(A)$.*

Proof. The inclusion $\text{hull}_N(A) \subseteq N \cap \text{hull}_M(A)$ is clear. For the reverse inclusion, let $x \in N \cap \text{hull}_M(A)$. By Lemma 7.4.5, we have $x = (B|x)$ for some finite subset, $B \subseteq A \subseteq N$. Applying Lemma 7.4.5 intrinsically to N we get $x \in \text{hull}_N(A)$. \square

Here are a couple of further observations.

Lemma 7.4.7. *Let $A \subseteq M$, and let $C = \text{hull}_M(A)$. Suppose that $x \in M$ and that $c \in C$ is a gate for x in C . Then there is a finite subset, $B \subseteq A$, such that $c = (B|x)$.*

Proof. By Lemma 7.4.5, there is some finite $B \subseteq A$ such that $c = (B|c)$. We claim that $c = (B|x)$. To see this, let $d = (B|x)$. Since $d \in C$, we have $x.c.d$. Also, by Lemma 7.4.3, $xdc = x(B|x)(B|c) = (B|xc) = (B|x) = d$. In other words, $x.d.c$. Therefore $c = d$ as claimed. \square

Lemma 7.4.8. *Let $N \leq M$ be subalgebra, and suppose that $A \subseteq N$ is convex in N . Let $C = \text{hull}_M(A)$. Suppose that C is gated in M , and let $\omega_C : M \rightarrow C$ be the gate map. Then $\omega_C(N) \subseteq A$, and $\omega_C|_N$ is a gate map to A in N . (In particular, A is gated in N .)*

Proof. Let $x \in N$ and let $c = \omega_C x \in C$. By Lemma 7.4.7, there is some finite subset, $B \subseteq A \subseteq N$ with $c = (B|x)$. Since A is convex in N , we have $c \in A$. For the second statement, let $d \in [x, c]_N \subseteq [x, c]_M$. Since c is a gate for x in C inside M , we have $d = c$. It follows that c is also a gate for x in A . \square

To proceed we need the following variation on Lemma 7.4.3.

Lemma 7.4.9. *Let $x \in M$ and let $A, B \subseteq M$ be non-empty finite subsets. Let $C = \{(A|b) \mid b \in B\}$. Then $(A|(B|x)) = (C|x)$.*

For the proof we will write $X\uparrow$ to mean $(\forall x \in X)(x\uparrow)$. Thus $\neg X\uparrow$ means $(\exists x \in X)(x\downarrow)$.

Proof. We want to show that $\text{LHS}\uparrow$ precisely when $\text{RHS}\uparrow$.

Suppose $A\uparrow$. Then $\text{LHS}\uparrow$. Since $(A|b)\uparrow$ for all $b \in B$, we have $C\uparrow$, and so $\text{RHS}\uparrow$. Similarly, if $A\downarrow$, we have $\text{LHS}\downarrow$ and $\text{RHS}\downarrow$. We therefore assume that $\neg A\uparrow$ and $\neg A\downarrow$.

Now suppose $B\uparrow$. Then $(B|x)\uparrow$. Since $\neg A\downarrow$, we have $\text{LHS}\uparrow$. Also, for each $b \in B$, we have $(A|b)\uparrow$, so $C\uparrow$, so $\text{RHS}\uparrow$. Similarly, if $B\downarrow$, we have $\text{LHS}\downarrow$ and $\text{RHS}\downarrow$. We therefore assume in addition that $\neg B\uparrow$ and $\neg B\downarrow$.

Now, without loss of generality, $x\uparrow$. Since $\neg B\downarrow$, we have $(B|x)\uparrow$, so $\text{LHS}\uparrow$. Again, since $\neg B\downarrow$, there is some $b \in B$ with $b\uparrow$. Thus, $(A|b)\uparrow$, and so $\neg C\downarrow$. Therefore $\text{RHS}\uparrow$. \square

The following is a variation on Example (Ex6.4) of Subsection 6.2.

Lemma 7.4.10. *Suppose that $A, B \subseteq M$ are non-empty finite subsets, and that there is some $x \in M$ such that $(A|x) = (B|x)$. Then for all $b \in B$ we have $(B|(A|b)) = (A|b)$.*

Proof. Suppose not. We can assume that $(B|(A|b))\uparrow$ and $(A|b)\downarrow$, so $B\uparrow$. Now $(A|x) = (B|x)\uparrow$, so $\neg A\downarrow$. Since $(A|b)\downarrow$, we have $b\downarrow$. But $b \in B$, so $b\uparrow$, giving a contradiction. \square

Proposition 7.4.11. *Let $N \leq M$ be a subalgebra, and $A, B \subseteq N$ be any subsets of N . Then $\text{hull}_M(A) \cap \text{hull}_M(B) = \text{hull}_M(\text{hull}_N(A) \cap \text{hull}_N(B))$.*

Proof. Since $\text{hull}_N(A) \cap \text{hull}_N(B) \subseteq \text{hull}_M(A) \cap \text{hull}_M(B)$, and the latter set is convex in M , the inclusion \supseteq is clear. We therefore want to prove that $\text{hull}_M(A) \cap \text{hull}_M(B) \subseteq \text{hull}_M(\text{hull}_N(A) \cap \text{hull}_N(B))$.

First let us suppose that A, B are both finite. Let $x \in \text{hull}_M(A) \cap \text{hull}_M(B)$ (assuming such exists). Therefore, by Lemma 7.4.4, interpreted in M , we have $x = (A|x) = (B|x)$. By Lemma 7.4.4, interpreted in N , we have $(A|b) \in \text{hull}_N(A)$ for all $b \in B$. By Lemma 7.4.10, we have $(A|b) = (B|(A|b))$. By Lemma 7.4.4, interpreted in N again, we have $(A|b) \in \text{hull}_N(B)$. Let $C = \{(A|b) \mid b \in B\} \subseteq \text{hull}_N(A) \cap \text{hull}_N(B)$. By Lemma 7.4.9 and Lemma 7.4.4, interpreted in M , we have $x = (A|x) = (A|(B|x)) = (C|x) \in \text{hull}_M C \subseteq \text{hull}_M(\text{hull}_N(A) \cap \text{hull}_N(B))$, as required.

Now suppose that A, B are arbitrary. If $x \in \text{hull}_M(A) \cap \text{hull}_M(B)$, then by Lemma 7.4.1, interpreted in M , we have $x \in \text{hull}_M(A_0) \cap \text{hull}_M(B_0)$, for some finite non-empty subsets, $A_0 \subseteq A$ and $B_0 \subseteq B$. Thus, by the previous paragraph, $x \in \text{hull}_M(\text{hull}_N(A_0) \cap \text{hull}_N(B_0)) \subseteq \text{hull}_M(\text{hull}_N(A) \cap \text{hull}_N(B))$. \square

7.5. Some final observations.

We end this section with a couple of further observations which will be used later.

Lemma 7.5.1. *Let $M = M_1 \times \cdots \times M_n$ be a direct product of median algebras M_i . Let $C \subseteq M$ be convex. Then $C = C_1 \times \cdots \times C_n$, where each $C_i \subseteq M_i$ is convex.*

Proof. By induction, it is enough to prove the statement for $n = 2$. Let C_i be the image of C under the projection map to M_i . Clearly, $C \subseteq C_1 \times C_2$. For the reverse inclusion, let $x_1 \in C_1$ and $x_2 \in C_2$. Then there exist $y_1 \in M_1$ and $y_2 \in M_2$ with $(x_1, y_2), (y_1, x_2) \in C$. Now $(x_1, y_2) \cdot (x_1, x_2) \cdot (y_1, x_2)$, so $(x_1, x_2) \in C$. \square

Let M be any median algebra. We can identify the power set, $\mathcal{P}(M)$ with $\{0, 1\}^M$ via characteristic functions. In this way, $\mathcal{P}(M)$ can be given the product topology, and as such is compact, by Tychonoff's theorem. Given $x \in M$, let $\mathcal{S}(x) = \{A \in \mathcal{P}(M) \mid x \in A\}$. This is clopen in $\mathcal{P}(M)$. (Note that $\mathcal{S}(x)$ is the preimage of 1 under the projection to the x -coordinate of $\{0, 1\}^M$. Its complement, $\mathcal{S}(x)^*$, is the preimage of 0.) Given $a, b, c \in M$, let $\mathcal{T}(a, b, c) = \mathcal{S}(a)^* \cup \mathcal{S}(b)^* \cup \mathcal{S}(c)$. This is also clopen. The set of all convex subsets of M is thus equal to $\bigcap \{\mathcal{T}(a, b, c) \mid a, b, c \in M, a.c.b\}$. This is closed in $\mathcal{P}(M)$, and therefore inherits a compact topology. (Note however, that the ternary operation $[(A, B, C) \mapsto ABC]$ described above is not in general continuous in this topology.)

Now the involution $[A \mapsto A^*]$ is continuous on $\mathcal{P}(M)$. In particular, the set of subsets whose complements are convex is also closed in $\mathcal{P}(M)$. Intersecting this with the set of convex subsets, we see that $\mathcal{H}(M) \cup \{\emptyset, M\}$ is also closed. Here, $\mathcal{H}(M)$ is the set of halfspaces of M , to which need to adjoin \emptyset and M itself. Note that $\{H \in \mathcal{H}(M) \cup \{\emptyset, M\} \mid A \subseteq H\}$ is clopen in $\mathcal{H}(M)$ for any finite subset $A \subseteq M$.

We can also factor $\mathcal{P}(M)$ by the involution. The quotient space is again a compact totally disconnected space. The image of $\mathcal{H}^\infty(M) := \mathcal{H}(M) \cup \{\emptyset, M\}$ in this quotient is $\mathcal{W}^\infty(M) := \mathcal{W}(M) \sqcup \{\{\emptyset, M\}\}$; that is, the set of walls, $\mathcal{W}(M)$, to which we have adjoined the trivial partition of M , namely $\{\emptyset, M\}$. Thus, $\mathcal{W}^\infty(M)$ has a natural topology as a compact totally disconnected space. Finally note that if $x, y \in M$, then the set of walls $\mathcal{W}(x, y)$ separating x and y is clopen in $\mathcal{W}^\infty(M)$. (Its preimage in $\mathcal{H}^\infty(M)$ is $(\mathcal{S}(x) \cap \mathcal{S}(y)^*) \cup (\mathcal{S}(x)^* \cap \mathcal{S}(y))$.)

Topological median algebras will be discussed in Section 12. We will make use of the above observations regarding $\mathcal{P}(M)$ in Sections 8, 11 and 19.

8. WALLS AND RANK

We defined the notion of a wall in Subsection 3.2. This is another central notion in the subject, and we explore it in more detail here. We also define the notion of “rank” of a median algebra, and show that it can be equivalently described in terms of walls (Lemma 8.2.1). We define the notion of colourability, which is relevant to embedding theorems we consider later (for example, Proposition 15.3.1).

The special case of rank-1 median algebras will be discussed in Subsection 12.7, 14.2 and in Section 15.

8.1. Walls and halfspaces.

Let M be a median algebra. We begin by recalling two closely related definitions from Subsection 3.2.

Definition. A *wall*, $W = \{C, D\}$, of M is a partition of M into two disjoint non-empty convex subsets, $M = C \sqcup D$.

Definition. A *halfspace* of M is a non-empty convex subset, $H \subseteq M$, such that $M \setminus H$ is non-empty and convex.

We can “direct” a wall, by representing it as an ordered pair, (C, D) . In this way, a directed wall is essentially equivalent to a halfspace (where we take C to be our halfspace). By default however, walls are considered to be undirected.

We write $\mathcal{W} = \mathcal{W}(M)$ for the set of walls of M . Given $W \in \mathcal{W}$, we will often write $W = \{W^-, W^+\}$, though generally this is just for notational convenience, and does not imply a preferred direction on W .

Note that a directed wall is essentially equivalent to a wall map, also defined in Subsection 3.2; that is, an epimorphism $\phi : M \rightarrow \{0, 1\}$. The corresponding directed wall is $(\phi^{-1}(1), \phi^{-1}(0))$. We will mostly refer to walls from now on.

Definition. We say that a wall $W \in \mathcal{W}$ *separates* two subsets $A, B \subseteq M$ if either $(A \subseteq W^-$ and $B \subseteq W^+)$ or $(B \subseteq W^-$ and $A \subseteq W^+)$. We write $A|_W B$.

Although $A|_W B$ is equivalent to $B|_W A$, we will usually adopt the convention that when writing $A|_W B$ we will take $A \subseteq W^-$ and $B \subseteq W^+$. We write $A|B$ to mean that $A|_W B$ for some $W \in \mathcal{W}$. This clearly implies $A \cap B = \emptyset$.

If $a, b \in M$, we abbreviate $\{a\}|_W \{b\}$ to $a|_W b$, etc. By Proposition 3.2.11, and two distinct points are separated by a wall. In other words, $a|b$ is equivalent to $a \neq b$.

We will write $\mathcal{W}(a, b) = \{W \in \mathcal{W} \mid a|_W b\}$. We make the following observations, all of which are easily verified:

Lemma 8.1.1. *For $a, b, c, d \in M$, we have:*

- (1): $\mathcal{W}(a, b) = \mathcal{W}(b, a)$.
- (2): $\mathcal{W}(a, b) \subseteq \mathcal{W}(a, c) \cup \mathcal{W}(c, b)$.
- (3): $\mathcal{W}(a, c) \cap \mathcal{W}(c, b) = \emptyset \Leftrightarrow \mathcal{W}(a, b) = \mathcal{W}(a, c) \cup \mathcal{W}(c, b) \Leftrightarrow c \in [a, b]$.
- (4): $c, d \in [a, b] \Rightarrow \mathcal{W}(c, d) \subseteq \mathcal{W}(a, b)$.
- (5): $\mathcal{W}(a, b) = \emptyset \Leftrightarrow a = b$.
- (6): $\#\mathcal{W}(a, b) = 1 \Leftrightarrow a, b$ are adjacent.

For Part (6), recall that $a, b \in M$ are “adjacent” if $a \neq b$ and $[a, b] = \{a, b\}$. Now $\mathcal{W}(a, b) \neq \emptyset$. Suppose $W_1, W_2 \in \mathcal{W}(a, b)$ are distinct. We can suppose that $a \in W_1^- \cap W_2^-$, $b \in W_1^+ \cap W_2^+$, and that $W_1^+ \cap W_2^- \neq \emptyset$. Let $c \in W_1^+ \cap W_2^-$. Then

$abc \in W_1^+ \cap W_2^-$. In particular, $abc \notin \{a, b\}$, giving a contradiction. Therefore $\#\mathcal{W}(a, b) = 1$. In this case, if $\mathcal{W}(a, b) = \{W\}$, we will write $W(a, b) := W$.

We move on to the consideration of separation of convex sets. Suppose $A, B \subseteq M$ are disjoint convex subsets.

Suppose first that A and B are gated. Let $a \in A$ and $b \in B$ be mutual gates (as given by Lemma 7.3.6). Now (similarly as for adjacent points above) we see that any wall separating a from b will also separate A from B . Therefore (since $\mathcal{W}(a, b) \neq \emptyset$) we have $A|B$.

It follows immediately from Lemma 7.3.7 that any two disjoint convex sets, A, B , in a finite median algebra are separated by a wall. (As noted after Lemma 7.3.7, the existence of mutual gates is easier in this case.)

In fact, this result holds generally. The following is often referred to as the ‘‘Kakutani separation property’’.

Theorem 8.1.2. *Any two disjoint convex sets in any median algebra are separated by a wall.*

Proof. The argument is similar to that of Proposition 3.2.10.

Let $A, B \subseteq M$ be disjoint convex subsets of a median algebra M . Let \mathcal{C} be the set of convex subsets, C , of M with $B \subseteq C$ and $A \cap C = \emptyset$. We order \mathcal{C} by inclusion. By Zorn’s Lemma, \mathcal{C} has a maximal element H . We claim that H is a halfspace.

Suppose, to the contrary that $c_1, c_2 \in M \setminus H$ and $c \in [c_1, c_2] \cap H$. By Lemma 3.2.9, $J(H, \{c_i\})$ is convex, and so must intersect A . Let $a_i \in A \cap J(H, \{c_i\})$. Then $a_i \in [c_i, h_i]$ for some $h_i \in H$. Let $a = a_1 a_2 c$ and $h = h_1 h_2 a$.

We claim that $a \in [c, h]$. For suppose not. We can suppose $a \uparrow$ and $c \downarrow$, $h \downarrow$. Therefore $a_1 \uparrow$ and $a_2 \uparrow$. Also $h_1 \downarrow$ and $h_2 \downarrow$. Since $a_i \in [c_i, h_i]$, we must have $c_i \uparrow$. Now $c_1.c.c_2$, so $c \uparrow$ given a contradiction. This proves the claim.

But now, $c, h \in H$, so $a \in H$, contradicting $H \in \mathcal{C}$. \square

(Note that we have implicitly used Proposition 3.2.10 in arguing that $a \in [c, h]$, though one could, of course, give an explicit verification of this.)

As a consequence, we see that $A, B \subseteq M$ are arbitrary subsets, then $A|B$ holds if and only if $\text{hull}(A) \cap \text{hull}(B) = \emptyset$.

Suppose $N \leq M$ is a subalgebra. From the previous observation and Proposition 7.4.11, we see that for any subsets, $A, B \subseteq N$, $A|B$ holds in N if and only if it holds in M . Note that we can apply this to (the (halfspaces of) any wall in N . This shows that any such wall has the form $W_N := \{N \cap W^-, N \cap W^+\}$ for some wall W of M . If $M = \text{hull}_M(N)$, then we have a map $\mathcal{W}(M) \rightarrow \mathcal{W}(N)$ given by $[W \mapsto W_N]$ which is surjective by the above observation.

Suppose now that $M \rightarrow N$ is an epimorphism. In this case, we have a natural injective map $\mathcal{W}(N) \rightarrow \mathcal{W}(M)$ obtained by taking preimages of halfspaces.

Suppose $(M_i)_{i \in \mathcal{I}}$ is a family of median algebras indexed by some set \mathcal{I} . Let $M = \prod_{i \in \mathcal{I}} M_i$ be the direct product. Then we have a natural inclusion of the disjoint union, $\bigsqcup_{i \in \mathcal{I}} \mathcal{W}(M_i)$, into $\mathcal{W}(M)$ induced by the projection maps to the

factors, as in the previous paragraph. It is easily checked that this is injective. If \mathcal{I} is finite, then it is also surjective:

Lemma 8.1.3. *Let M_1, \dots, M_n be median algebras. Then the natural map $\mathcal{W}(M_1) \sqcup \dots \sqcup \mathcal{W}(M_n) \hookrightarrow \mathcal{W}(M_1 \times \dots \times M_n)$ described above is a bijection.*

Proof. By induction, it is enough to check this for $n = 2$. Let $M = M_1 \times M_2$ and suppose $W \in \mathcal{W}(M_1 \times M_2)$. Let H_i^\pm be the image of W^\pm under the projection map to M_i . By Lemma 7.5.1, $W^\pm = H_1^\pm \times H_2^\pm$. Now $M = W^- \sqcup W^+$ and it follows easily that, up to swapping the indices, we have $M_1 = H_1^- \sqcup H_1^+$ and $H_2^- = H_2^+ = M_2$. Thus, $W_\pm = H_1^\pm \times M_2$ and $\{H_1^-, H_1^+\}$ is a wall of M_1 . \square

In particular, if $Q = \{0, 1\}^{\mathcal{I}}$ is a hypercube, then for each $i \in \mathcal{I}$, we have a wall $W_i \in \mathcal{W}(Q)$ arising from the projection to the i -coordinate. If \mathcal{I} is finite, then $\mathcal{W}(Q) = \{W_i \mid i \in \mathcal{I}\}$.

Remark. This need not be true if \mathcal{I} is infinite. For example, let $Q_0 \subseteq Q$ be the set of elements for which all but finitely many coordinates are 0. Similarly define Q_1 . Then Q_0 and Q_1 are disjoint convex subsets, and therefore separated by a wall. But this wall cannot be any of the W_i . We return to this example in Subsection 11.11.

Definition. We say that a subset $A \subseteq M$ and a wall W **cross** if $A \cap W^- \neq \emptyset$ and $A \cap W^+ \neq \emptyset$. We write $A \pitchfork W$.

Definition. We say that two walls, W_1, W_2 **cross** if each of the four subsets $W_1^- \cap W_2^-, W_1^- \cap W_2^+, W_1^+ \cap W_2^-$ and $W_1^+ \cap W_2^+$ are all non-empty. We write $W_1 \pitchfork W_2$.

The following simple observation is worth noting:

Lemma 8.1.4. *Suppose $W, W' \in \mathcal{W}(M)$ and $C \subseteq M$ is convex. Suppose $C \pitchfork W$, $C \pitchfork W'$ and $W \pitchfork W'$. Then $W_C \pitchfork W'_C$ intrinsically in C .*

Proof. By the Helly Property (Lemma 7.1.1), the sets $C \cap W^\pm \cap (W')^\mp$ are all non-empty. \square

Suppose $W_1, \dots, W_n \in \mathcal{W}(M)$ pairwise cross; that is, $W_i \pitchfork W_j$ for all $i \neq j$. Given $\epsilon = \{\epsilon_1, \dots, \epsilon_n\} \in \{+, -\}^n$, we write $O(\epsilon) = \bigcap_{i=1}^n W_i^{\epsilon_i}$. We refer to $O(\epsilon)$ as an **orthant** of this family of walls. By the Helly Property (Lemma 7.1.1) we have $O(\epsilon) \neq \emptyset$ for all ϵ . In fact, if $C \subseteq M$ is convex and crosses each of the walls W_i , then (applying the Helly Property to the family $\{W_1, \dots, W_n, C\}$) we see that $C \cap O(\epsilon) \neq \emptyset$ for all ϵ . In particular, $\#C \geq 2^n$. Suppose that $\#C = 2^n$. Let $c(\epsilon)$ be the unique point of $C \cap O(\epsilon)$. Then the map $[\epsilon \mapsto c(\epsilon)]$ is an isomorphism from the cube $\{+, -\}^n$ to C . In other words, C is an n -cube.

Lemma 8.1.5. *Let W_1, \dots, W_n be pairwise crossing walls. Then there is an n -cube, $Q \subseteq M$ such that $\#(Q \cap O(\epsilon)) = 1$ for all $\epsilon \in \{+, -\}^n$.*

Proof. First observe that we can reduce to the case where M is finite. To see this, choose any $c(\epsilon) \in O(\epsilon)$ and let $\Pi = \langle \{c(\epsilon) \mid \epsilon \in \{+, -\}^n\} \rangle$. This is finite by Proposition 3.3.3. The walls W_i give us pairwise crossing walls, $(W_i)_\Pi$, in $\mathcal{W}(\Pi)$. Now replace M with Π .

So we assume that M is finite. Let $Q \subseteq M$ be a convex set with $Q \pitchfork W_i$ for all i and such that $\#Q$ is minimal with this property. We claim that Q is an n -cube.

If not, by the earlier discussion, we must have $\#(Q \cap O(\epsilon)) > 1$ for some ϵ . Let $a, b \in Q \cap O(\epsilon)$ be distinct, and let $W \in \mathcal{W}(a, b)$. Note that $W \neq W_i$ for any i .

Suppose first that there is some i such that $W \not\pitchfork W_i$. We can suppose that $W_i^- \cap W^+ = \emptyset$, so $W_i^- \subseteq W^-$. Let $C = Q \cap W^-$. Then C is convex. We claim that $C \pitchfork W_j$ for all j . If $j = i$ this follows directly from the construction. If $j \neq i$, note that $W_j^\pm \cap W^- \supseteq W_j^\pm \cap W_i^- \neq \emptyset$, $Q \cap W_j^\pm \neq \emptyset$ and $Q \cap W^- \neq \emptyset$, and so by the Helly Property (Lemma 7.1.1), we have $C \cap W_j^\pm = Q \cap W^- \cap W_j^\pm \neq \emptyset$. Thus $C \pitchfork W_j$ as claimed.

On the other hand, if $W \pitchfork W_i$ for all i then (since Q meets all the orthants of $\{W_1, \dots, W_n, W\}$) we again have $C \pitchfork W_i$ for all i , where $C = W^-$.

We could therefore replace Q by C , thereby contradicting the minimality of $\#Q$. This shows that Q is an n -cube as claimed. \square

8.2. Properties of rank and subalgebras.

We now make a key definition.

Definition. The *rank*, $\text{rank}(M)$ of a median algebra, M , the maximal $\nu \in \mathbb{N}$ such that M contains an ν -cube. We write $\text{rank}(M) = \infty$ if there is no upper bound.

Note that any 2-element subset is a subalgebra, and so if M has rank 0, then it is a single point. (In line with dimension theory conventions, it may be natural to view the empty set as having rank -1 , though that won't concern us much.)

Lemma 8.2.1. *The rank of M is the maximal ν such that there is a set, W_1, \dots, W_ν , of pairwise crossing walls in M .*

Proof. The fact that $\nu \leq \text{rank}(M)$ follows immediately from Lemma 8.1.5. To see that $\text{rank}(M) \leq \nu$, let $Q \subseteq M$ be a ν -cube. Then $\#\mathcal{W}(Q) = \nu$. Moreover, as observed above, each wall of Q has the form $W_Q = \{Q \cap W^-, Q \cap W^+\}$ for some $W \in \mathcal{W}$. These walls pairwise cross in M . \square

We make a few simple observations.

If $N \leq M$ is a subalgebra, then $\text{rank}(N) \leq \text{rank}(M)$. If $M \rightarrow M'$ is an epimorphism, then $\text{rank}(M) \geq \text{rank}(M')$ (since we have seen that there is a natural injective map $\mathcal{W}(M') \rightarrow \mathcal{W}(M)$, and this respects crossing).

Lemma 8.2.2. *If M_1, M_2 are median algebras then $\text{rank}(M_1 \times M_2) = \text{rank}(M_1) + \text{rank}(M_2)$.*

Proof. We have seen (Lemma 8.1.3) that there is a natural identification of $\mathcal{W}(M_1 \times M_2)$ with $\mathcal{W}(M_1) \sqcup \mathcal{W}(M_2)$. Moreover, any wall of $M_1 \times M_2$ arising from M_1 crosses any wall arising from M_2 . \square

We mention the following result:

Proposition 8.2.3. *Suppose $\text{rank}(M) \leq \nu$ and $A \subseteq M$. Then $\text{hull}(A) = J^\nu(A)$.*

(Recall that $J^n(A)$ is the n th iterated join of A , as defined in Subsection 7.4.)

We just note, for the moment, that this reduces to the case where M is finite. To see this, let $x \in \text{hull}(A)$. By Lemma 7.4.1, $x \in \text{hull}(B)$ for some finite $B \subseteq A$. Let $\Pi = \langle B \cup \{x\} \rangle$. Now $\text{rank}(\Pi) \leq \text{rank}(M) \leq \nu$, so assuming the finite case, we have $x \in J^\nu(B)$ in Π , so $x \in J^\nu(B) \subseteq J^\nu(A)$ in M . The reverse inclusion is clear.

Although we have what we need to prove the finite case now, we will postpone it until Subsection 11.5 (see Proposition 11.5.4), since it fits more naturally with the general discussion there.

One can doubtless do better than the ν th iterate in this result. But it is not clear what the optimal statement would be.

We also mention a similar result regarding subalgebras. Let M be a median algebra. Given $A \subseteq M$, let $T(A) = \{abc \mid a, b, c \in A\}$. Define $T^n(A)$ inductively by $T^0(A) = A$ and $T^{n+1}(A) = T(T^n(A))$. One sees easily that $\langle A \rangle = \bigcup_{n=0}^{\infty} T^n(A)$.

In fact, we have the following result:

Proposition 8.2.4. *If $\text{rank}(M) \leq \nu$, then $\langle A \rangle = T^{2\nu}(A)$.*

The statement with 2ν replaced by a different explicit function of ν can be found in [Fi5]. Again is not clear what the optimal statement should be.

We will again postpone the proof to Subsection 11.7 (see Lemma 11.7.3 and subsequent proof). We just observe here that one can reduce to the case where A is finite and where $M = \langle A \rangle$. (Note that if $a \in \langle A \rangle$, there is some finite $C \subseteq A$ such that $a \in \langle C \rangle$, and we can replace A by C and M by $\langle C \rangle$.)

While we are on the subject, we note that following gives another description of the subalgebra generated by a subset of M .

Let \mathcal{H} be the set of halfspaces of M . Given $p \in M$, let $\mathcal{H}(p) = \{H \in \mathcal{H} \mid p \in H\}$.

Proposition 8.2.5. *Let $A \subseteq M$ and $p \in M$. Then $p \in \langle A \rangle$ if and only if $H_1 \cap H_2 \cap A \neq \emptyset$ for all $H_1, H_2 \in \mathcal{H}(p)$.*

Put another way, $\langle A \rangle$ is the intersection of all sets of the form $H \cup H'$ containing A , where H, H' are halfspaces of M .

Proof. For the “if” direction, suppose $H_1, H_2 \in \mathcal{H}(p)$ and $H_1 \cap H_2 \cap A = \emptyset$. Now $M \setminus (H_1 \cap H_2)$ is a subalgebra (being the union of two convex sets, namely $(M \setminus H_1) \cup (M \setminus H_2)$) and it contains A . Therefore $\langle A \rangle \subseteq M \setminus (H_1 \cap H_2)$, and so $p \notin \langle A \rangle$.

For the converse, let us assume first that A is finite, and that $H_1 \cap H_2 \cap A \neq \emptyset$ for all $H_1, H_2 \in \mathcal{H}(p)$. Let $\mathcal{C} = \mathcal{C}(p) = \{\text{hull}(H \cap A) \mid H \in \mathcal{H}(p)\}$. By hypothesis, the elements of \mathcal{C} pairwise intersect, so by the Helly Property (Lemma 7.1.1), $\bigcap \mathcal{C} \neq \emptyset$. In fact, we claim that $\bigcap \mathcal{C} = \{p\}$. For if $q \in \bigcap \mathcal{C} \setminus \{p\}$, then by Proposition 3.2.11, there is some $H \in \mathcal{H}(p)$ with $q \notin H$. But also $q \in \text{hull}(H \cap A) \subseteq H$, giving a contradiction.

Now let $\mathcal{C}' = \{C \cap \langle A \rangle \mid C \in \mathcal{C}\}$. The elements of \mathcal{C}' are all convex as subsets of $\langle A \rangle$, and again pairwise intersect. Therefore, by the Helly Property in $\langle A \rangle$, we have $\bigcap \mathcal{C}' \neq \emptyset$. But $\bigcap \mathcal{C}' \subseteq \langle A \rangle$ and $\bigcap \mathcal{C}' \subseteq \bigcap \mathcal{C} = \{p\}$, and so $p \in \langle A \rangle$ as required.

For the general case, recall from Subsection 7.5 that $\mathcal{H} \sqcup \{\emptyset, M\}$ admits a compact topology such that $\mathcal{H}(a)$ is clopen for all $a \in M$. Given any finite $B \subseteq M$, let $K(B) = \{(H_1, H_2) \in \mathcal{H}^2 \mid H_1 \cap H_2 \cap B = \emptyset, p \in H_1 \cap H_2\}$. This is clopen in \mathcal{H}^2 with the product topology. (To see this, write $\pi_i : \mathcal{H}^2 \rightarrow \mathcal{H}$ for projection to the i th coordinate. Then $K(B) = \mathcal{H}(p)^2 \cap \bigcap_{b \in B} ((\pi_1^{-1}(M \setminus \mathcal{H}(b)) \cup (\pi_2^{-1}(M \setminus \mathcal{H}(b))))$, and each of these sets are clopen in \mathcal{H}^2 .)

Now suppose that $p \notin \langle A \rangle$. By the finite case above, we have $K(B) \neq \emptyset$, for all finite $B \subseteq A$. Thus, the family of all such $K(B)$ has the finite intersection property in the compact space, \mathcal{H}^2 , so there is some (H_1, H_2) in the intersection. We now get $H_1 \cap H_2 \cap A = \emptyset$ as required. \square

Note that, in the proof, we showed that $\bigcap \mathcal{C}(p) = \{p\}$, and so the map $[p \mapsto \mathcal{C}(p)]$ is injective. There are at most $2^{\#A}$ possibilities for the sets $H \cap A$, and hence at most $2^{2^{\#A}}$ possibilities for $\mathcal{C}(p)$. This shows that $\#\langle A \rangle \leq 2^{2^{\#A}}$, thereby giving another (non-constructive) proof of Proposition 3.3.3.

Many of the median algebras arising in geometry have finite rank. For example, finite dimensional CAT(0) cube complexes (see Sections 16 and 17) as well as those arising as certain asymptotic cones (see Section 24).

However, many do not. An obvious example is an infinite direct product of non-trivial median algebras. The median algebras arising from spaces of measured walls (Section 19) typically have infinite rank.

Here is an example of a median algebra in which all non-trivial convex subsets have infinite rank.

Example. Let M be any median algebra with $\#M \geq 2$. Let $M^{\mathbb{Z}}$ be the set of bi-infinite sequences, $\underline{x} = (x_i)_{i \in \mathbb{Z}}$, with the product median structure. Given $n \in \mathbb{N} \setminus \{0\}$, let $P_n \subseteq M^{\mathbb{Z}}$ consist of all n -periodic sequences (i.e. $x_{i+n} = x_i$ for all i). This is a subalgebra of $M^{\mathbb{Z}}$, isomorphic to M^n . If $p, q, r \in \mathbb{N} \setminus \{0\}$, then $P_p P_q P_r \subseteq P_{pqr}$, and so $P := \bigcup_{n=1}^{\infty} P_n$ is also a subalgebra of $M^{\mathbb{Z}}$.

Suppose $\underline{x}, \underline{y} \in P$. We have that $\underline{x}, \underline{y} \in P_p$ for some p . If $\underline{x} \neq \underline{y}$, then there is some $k \in \mathbb{Z}$ such that $x_i \neq y_i$ for all $i \in \mathcal{I} := \{pm + k \mid m \in \mathbb{Z}\}$. Let $I = [x_i, y_i]_M$ for some $i \in \mathcal{I}$. Given any $n \in \mathbb{N} \setminus \{0\}$, $[\underline{x}, \underline{y}]_P$ contains a copy of I^n , namely, those $\underline{z} \in P_{pn}$ for which $z_i \in [x_i, y_i]$ for all $i \in \mathcal{I}$ and $z_i = x_i$ for all $i \in \mathbb{Z} \setminus \mathcal{I}$. In

particular, $[\underline{x}, \underline{y}]_P$ contains an n -cube for all $n \in \mathbb{N}$.

8.3. Colourability.

A stronger notion related to rank is that of colourability. Certain naturally occurring spaces are finitely colourable (for example most of those mentioned in Section 24 — see the Notes to that section). There are a number of results to the effect that a finitely colourable space can be embedded in a finite product of trees — for example Proposition 15.3.1.

Here is the definition:

Definition. A median algebra, M , is ν -**colourable** if there is a map $\chi : \mathcal{W}(M) \rightarrow \{1, \dots, \nu\}$, such that if $W, W' \in \mathcal{W}(M)$ with $W \pitchfork W'$, then $\chi(W) \neq \chi(W')$.

We write $\mathcal{W}_i(M) = \chi^{-1}(i)$. We think of these walls as labelled with the “colour”, i .

Remark. We are making no assumption at all regarding the regularity of the function χ . Indeed, this would make no sense for a general median algebra. In certain cases, it may be possible to arrange for χ to have some degree of regularity.

From Lemma 8.2.1, we see that if M is ν -colourable, then $\text{rank}(M) \leq \nu$.

Remark. It is natural to ask conversely when $\text{rank}(M) \leq \nu$ implies that that M is finitely-colourable. This is true if $\nu \leq 2$ (the case $\nu = 1$ is elementary) but fails in general for $\nu \geq 5$. These facts were proven in [ChepH]. (The authors refer to CAT(0) cube complexes rather than median algebras, but as we shall see, the statements are equivalent.) It remains an open question as to whether the implication is true when $n = 3$ or 4.

If N is a subalgebra of M and M is ν -colourable, then so is N . (To see this, note that any wall of N has the form W_N for some wall W of M . We choose any such W , and set the colour of W_N to be $\chi(W)$.)

If $M \rightarrow M'$ is an epimorphism, and M' is colourable, then so is M . (Precompose the colouring of M with the natural map $\mathcal{W}(M') \rightarrow \mathcal{W}(M)$.)

Typically, what one uses in practice is that any finite subalgebra of M is ν -colourable. Given the Axiom of Choice, this is in fact equivalent. For the proof, we use the fact (from Proposition 3.3.3) that finite subalgebras form a directed set under inclusion. We will make use of this principle again later.

For the proof, if $\Pi \leq M$ and $W, W' \in \mathcal{W}(M)$, we write $W \pitchfork_{\Pi} W'$ to mean that $W \pitchfork \Pi$, $W' \pitchfork \Pi$ and $W_{\Pi} \pitchfork W'_{\Pi}$ in Π .

Lemma 8.3.1. *A median algebra is ν -colourable if and only if every finite subalgebra is.*

Proof. We have already noted that any subalgebra of ν -colourable median algebra is ν -colourable.

For the converse, suppose that every finite subalgebra of M is ν -colourable. Let $\Psi = \{1, \dots, \nu\}^{\mathcal{W}(M)}$. We give Ψ the product topology, which is compact by Tychonoff's Theorem. Let \mathcal{A} be the set of finite subalgebras of M . Given $\Pi \in \mathcal{A}$, let Ψ_Π be the set of maps $\chi : \mathcal{W}(M) \rightarrow \{1, \dots, \nu\}$ such that if $W, W' \in \mathcal{W}(M)$ with $W \pitchfork_\Pi W'$ then $\chi(W) \neq \chi(W')$. Thus, Ψ_Π is a closed subset of Ψ . Note that if $W \pitchfork \Pi$, then $W_\Pi \in \mathcal{W}(\Pi)$, and so any ν -colouring, χ_Π , of Π gives rise to an element, $\chi \in \Psi_\Pi$ by setting $\chi(W) = \chi_\Pi(W_\Pi)$ for all $W \in \mathcal{W}(M)$ with $W \pitchfork \Pi$, and setting $\chi(W)$ arbitrarily otherwise.

We claim that the family $\{\Psi_\Pi\}_{\Pi \in \mathcal{A}}$ has the finite intersection property. To see this, suppose that $\Pi_1, \dots, \Pi_\nu \in \mathcal{A}$. Let $\Pi = \langle \Pi_1 \cup \dots \cup \Pi_\nu \rangle \in \mathcal{A}$. Now any ν -colouring of Π gives us some element of Ψ_Π as described above. But if $W \in \mathcal{W}(M)$ crosses some Π_i , then it crosses Π , and so $\Psi_\Pi \subseteq \Psi_{\Pi_i}$ for all i . Therefore $\Psi_{\Pi_1} \cap \dots \cap \Psi_{\Pi_\nu} \neq \emptyset$ as claimed. Since Ψ is compact, $\bigcap_{\Pi \in \mathcal{A}} \Psi_\Pi \neq \emptyset$. Moreover, $W \pitchfork_{\Pi_i} W'$ implies $W \pitchfork_\Pi W'$.

Now choose any $\chi \in \bigcap_{\Pi \in \mathcal{A}} \Psi_\Pi$. We claim that χ is a ν -colouring of M . For if $W, W' \in M$ with $W \pitchfork W'$, then certainly there is a finite subalgebra $\Pi \leq M$ with $W \pitchfork_\Pi W'$. (See for example, Lemma 8.1.5.) Since $\chi \in \Psi_\Pi$, it follows that $\chi(W) \neq \chi(W')$ as required. \square

We now consider intervals in a median algebra. (The following few results can be interpreted intrinsically to an interval, and are really about bounded distributive lattices.)

Let $a, b \in M$. We have $\mathcal{W}(a, b) = \mathcal{W}([a, b])$. Given $W \in \mathcal{W}(a, b)$, we will use the convention that $a \in W^-$ and $b \in W^+$. We write $W_1 \leq W_2$ to mean that $W_1^- \cap W_2^+ = \emptyset$. This is equivalent to saying $W_1^- \subseteq W_2^-$ or to $W_2^+ \subseteq W_1^+$. We write $W_1 < W_2$ to mean $W_1 \leq W_2$ and $W_1 \neq W_2$. We therefore have the following tetrachotomy: if $W_1, W_2 \in \mathcal{W}(a, b)$, then exactly one of the relations $W_1 = W_2$, $W_1 < W_2$, $W_2 < W_1$ or $W_1 \pitchfork W_2$ holds. We note that \leq is a partial order on $\mathcal{W}(a, b)$. An antichain in $\mathcal{W}(a, b)$ is then a set of pairwise crossing walls. Thus, any antichain has cardinality at most $\text{rank}([a, b])$.

We recall Dilworth's Lemma [Dil] which says that if we have a partially ordered set such that every antichain has cardinality at most n , then the set decomposes as disjoint union of at most n chains.

Lemma 8.3.2. *Suppose that $a, b \in M$, and that $\text{rank}([a, b]) \leq \nu$. Then $[a, b]$ is ν -colourable.*

Proof. By Dilworth's Lemma, we can write $\mathcal{W}(a, b) = \mathcal{W}_1 \sqcup \dots \sqcup \mathcal{W}_\nu$, where each \mathcal{W}_i is a chain. This means that if $W, W' \in \mathcal{W}_i$, then either $W \leq W'$ or $W' \leq W$. This therefore gives a ν -colouring of $[a, b]$. \square

Recall that we also have a partial order on $[a, b]$ (where $x \leq y$ is equivalent to $a.x.y$ or to $x.y.b$). Note that $x \leq y$ is equivalent to saying $\mathcal{W}(a, x) \subseteq \mathcal{W}(a, y)$.

Given $n \in \mathbb{N}$, let $I_n = \{1, \dots, n\}$, thought of as a finite rank-1 median algebra (with the standard notion of betweenness).

Lemma 8.3.3. *Suppose $a, b \in M$, $\text{rank}([a, b]) \leq \nu$ and $\#\mathcal{W}(a, b) = p < \infty$. Then there is a monomorphism, $[a, b] \hookrightarrow \prod_{i=1}^{\nu} I_{p_i}$, where $p_i \in \mathbb{N}$ and $\sum_{i=1}^{\nu} p_i = p$.*

Proof. By Lemma 8.3.2, $[a, b]$ admits a ν -colouring: $\mathcal{W}(a, b) = \mathcal{W}_1 \sqcup \dots \sqcup \mathcal{W}_{\nu}$. Let $p_i = \#\mathcal{W}_i$. Given $i \in \{1, \dots, \nu\}$ and $x \in [a, b]$, let $\phi_i(x) = \#(\mathcal{W}_i \cap \mathcal{W}(a, x))$. If $x \leq y$, then $\mathcal{W}(a, x) \subseteq \mathcal{W}(a, y)$, and so $\phi_i(x) \leq \phi_i(y)$. Since the order on a bounded distributive lattice determines the median structure, it follows that $\phi_i : [a, b] \rightarrow I_{p_i}$ is a homomorphism. Now let $\phi(x) = (\phi_1(x), \dots, \phi_{\nu}(x))$, so that $\phi : [a, b] \rightarrow \prod_{i=1}^{\nu} I_{p_i}$ is also a homomorphism. Moreover, ϕ is injective. For if $x \neq y$, choose some $W \in \mathcal{W}(x, y) \subseteq \mathcal{W}(a, b)$. Up to swapping x and y , we can suppose that $W \in \mathcal{W}(a, y) \setminus \mathcal{W}(a, x)$. Suppose $W \in \mathcal{W}_i$. Since no two walls in \mathcal{W}_i cross, $\mathcal{W}_i \cap \mathcal{W}(a, x) \setminus \mathcal{W}(a, y) = \emptyset$, and so $\phi_i(x) < \phi_i(y)$. Thus $\phi(x) \neq \phi(y)$. \square

Note that it immediately follows that $[a, b]$ is finite. In fact, we knew this already. We can identify $\mathcal{W}(a, b)$ as the set of walls of $[a, b]$ viewed intrinsically as a median algebra. It follows that $\#[a, b] \leq 2^p$ (see Corollary 3.2.14). Moreover, this does not a-priori require $[a, b]$ to have finite rank.

9. HALFSACES AND DUALITY

In this section we study the structure of halfspaces in a median algebra. This gives rise to a duality between median algebras and certain structures which we call “prosets”. A proset is almost the same thing as what Roller calls a “poc set” in [R] (except we are not allowing for the “zero” element of a poc set).

We give some examples of prosets. In particular, Example (Ex9.5) provides a useful way of visualising such a structure in terms of arrangements of lines (though it does not give a completely general picture). The duality alluded to is expressed in terms of “flows” on a proset. The flows on a proper power set described in Subsection 5.1 are an important examples. The duality is precise in the finite case: the statements are given by Propositions 9.2.2 and 9.2.7. This will complete the proof of Theorem 5.2.3. In fact, a more general statement, due to Nica, is given here as Proposition 9.4.2. (There also are generalisations of the duality results to the infinite case, see for example, Theorem 12.4.2.) In Subsection 9.5 we say more about the structure of the space of flows, which will be relevant to the discussion of Roller boundaries in Subsection 11.12. At the end of the section, we return to a brief discussion of boolean algebras.

9.1. Prosets.

Here is our definition of a proset:

Definition. A **proset** is a set, Ω , equipped with an involution $[a \mapsto a^*]$, and a partial order, \leq , which satisfies:

- (P1): For all $a, b \in \Omega$, $a \leq b \Rightarrow b^* \leq a^*$, and
- (P2): There is no element $a \in \Omega$ for which $a \leq a^*$.

We write $a < b$ to mean that $a \neq b$ and $a \leq b$. Note that, given (P1), axiom (P2) is equivalent to saying that for any $a, b \in \Omega$, at most one of the six relations, $a = b$, $a = b^*$, $a < b$, $b < a$, $a^* < b$, $a < b^*$, holds.

We write $a \pitchfork b$, to mean than none of these six relations hold. Note that $a \pitchfork b \Leftrightarrow b \pitchfork a \Leftrightarrow a \pitchfork b^*$. If this holds, we say that a and b **cross**.

Remark. If we assume that exactly one of these six relations holds for any $a, b \in \Omega$ (that is, we never have $a \pitchfork b$) then we arrive at the definition of a “protree” as defined by Dunwoody. This corresponds to the “rank-1” situation. (We will give the definition of “rank” in this context in Subsection 9.3.)

An intuitive way of thinking about a proset is as follows. We imagine some ambient space. We think of each element of the proset as separating the space into two pieces, and pointing towards one of them. The elements a and a^* separate it into the same two pieces, but point in opposite directions. One can think of the relation $a \pitchfork b$ to mean that a, b cross, so that together they separate the space into four pieces. If $a \not\pitchfork b$, then they separate the space into just three pieces. The statement $a^* < b$ (equivalently $b^* < a$) can be interpreted to mean that a, b point towards each other. Similarly the statement $a < b^*$ (equivalently $b < a^*$) means that they a, b point towards each other. This idea can be made explicit various contexts. (See for example (Ex9.5) below.)

Note that any $*$ -invariant subset of Ω is also a proset: a **subproset**.

Definition. A **flow** on a proset is a subset $R \subseteq \Omega$ such that:

(F1): For all $a \in \Omega$, exactly one of $a \in R$ or $a^* \in R$ holds, and

(F2): If $a \in R$, $b \in \Omega$ and $a \leq b$, then $b \in R$.

Thus, (F1) says that R is a $*$ -transversal, and (given (F1)) (F2) is equivalent to saying that we never have $a \leq b^*$, for $a, b \in R$. Note that a flow restricted to a subproset is also a flow on that subproset.

Let $\mathcal{F}(\Omega)$ be the set of all flows on Ω . We can view $\mathcal{F}(\Omega)$ as a subset of the power set, $\mathcal{P}(\Omega)$, with the standard median. (Recall that the median, RST , of $R, S, T \subseteq \Omega$, is the set of elements of Ω which lie in at least two of the sets, R, S, T .) It is readily checked that $\mathcal{F}(\Omega)$ is a subalgebra.

By Lemma 2.1.3, the median interval, $[R, S] \subseteq \mathcal{P}(\Omega)$, between two subsets $R, S \subseteq \Omega$ is given by $[R, S] = \{T \subseteq \Omega \mid R \cap S \subseteq T \subseteq R \cup S\}$. (Note that this can equivalently be written as $R \triangle S = (R \triangle T) \sqcup (S \triangle T)$, where “ \triangle ” denotes symmetric difference, and “ \sqcup ” implies disjoint union. If R, S, T are all $*$ -transversals, then the statements $R \cap S \subseteq T$ and $T \subseteq R \cup S$ are equivalent. In this case, $R \triangle S$, $R \triangle T$ and $S \triangle T$ are all subprosets of Ω . In particular, this applies to any interval in $\mathcal{F}(\Omega)$.) Viewing a flow as selecting an element from $\{a, a^*\}$, for all $a \in \Omega$, we can express this by saying that T agrees with R and S wherever R and S agree.

Here are some examples:

(Ex9.1): Let X be any set, and let $\mathcal{P}(X)$ be its power set. Let $\mathcal{P}_0(X) = \mathcal{P}(X) \setminus \{\emptyset, X\}$ be the “proper power set” of X . Given $A \in \mathcal{P}_0(X)$, let $A^* = X \setminus A$, and write $A \leq B$ to mean $A \subseteq B$. Then $\mathcal{P}_0(X)$ is a proset. The notion of a flow on $\mathcal{P}_0(X)$ coincides precisely with that defined in Subsection 5.1. (Note that $A \cap B = \emptyset$ is equivalent to $A \subseteq B^*$, which we are ruling out in each of the respective formulations.) Thus, $\mathcal{F}(\mathcal{P}_0(X))$ is precisely the superextension, $\Phi(X)$, of X , defined in Subsection 5.1. (We remark that, with the same definitions, $\mathcal{P}(X)$, is a poc set, as defined in [R], with $0 = \emptyset$ and $0^* = X$.)

(Ex9.2): Suppose Ω_1, Ω_2 are prosets. Let $\Omega = \Omega_1 \sqcup \Omega_2$, with the relation \leq defined as the union of these relations on Ω_1 and Ω_2 : (so that $a \uparrow b$ for all $a \in \Omega_1$ and $b \in \Omega_2$). With this structure, Ω is a proset. The flows on Ω are of the form $R_1 \sqcup R_2$, where R_i is a flow in Ω_i . Identifying $\mathcal{P}(\Omega) \equiv \mathcal{P}(\Omega_1) \times \mathcal{P}(\Omega_2)$, we see that $\mathcal{F}(\Omega)$ is median isomorphic to the direct product $\mathcal{F}(\Omega_1) \times \mathcal{F}(\Omega_2)$. This construction clearly generalises to arbitrary disjoint unions.

(Ex9.3): Let M be any median algebra. Then the set of halfspaces, $\mathcal{H}(M)$, of M is a subproset of $\mathcal{P}(M)$. We can think of a flow on $\mathcal{H}(M)$ as choosing a preferred direction on each wall $W \in \mathcal{W}(M)$, where we “choose” the halfspace W^- to lie in the flow. We think of the wall as directed away from W^+ and towards W^- . Property (P2) tells us that no two walls are directed away from each other. This gives the intuitive sense of a “flow”. We note that if $R \in \mathcal{F}(\mathcal{H}(M))$, then by definition, the elements of R pairwise intersect, and so by the Helly Property (Lemma 7.1.1) any finite subset of R has non-empty intersection.

(Ex9.4): Let M be the vertex set of a simplicial tree. Then M has the structure of a rank-1 median algebra with the obvious betweenness relation (Ex3.1). We can naturally identify $\mathcal{W}(M)$ with the edge set of M . A directed wall is then a directed edge. A flow is then an assignment of directions which has no source. Such a flow either has exactly one sink in M , or else converges on an ideal point in the “boundary” of M . We will explain this further in Subsection 11.12 (see Example (Ex11.2)).

(Ex9.5): A geometric example, which might help to visualise some of the constructions, comes from families of lines in the plane, $D \cong \mathbb{R}^2$. By a “line” in D we mean a properly embedded copy of the real line. To avoid technical topological issues we can assume it to be piecewise linear (or we could take it to be smooth). Note that a line, α , cuts D into two connected components, or “halfspaces”. A “transverse orientation” on α consists of a choice of one of the halfspaces. We imagine the transverse orientation pointing towards this halfspace.

By an “arrangement” we mean a locally finite family of lines such that any two meet, if at all, transversely at a single point. We will also assume that no subset of three lines intersect. The set of lines with their transverse orientations then has the structure of a proset. The involution reverses the transverse orientation, and

we write $a < b$ to mean that the corresponding lines are disjoint, the orientation on a points away from b , and the orientation on b points towards a . In this way, $a \pitchfork b$ means that the underlying lines meet transversely at a single point.

A particular example is to take $D = \mathbb{H}^2$ to be the hyperbolic plane, and lines to be bi-infinite geodesics. We will return to this example in Subsection 9.4. One can also generalise to higher dimensions: see Example (Ex19.4) of Subsection 19.4.

9.2. Duality.

Let M be a median algebra. We define a map $\eta : M \rightarrow \mathcal{F}(\mathcal{H}(M))$ by setting $\eta(x) = \{H \in \mathcal{H}(M) \mid x \in H\}$. (Here, $\mathcal{H}(M)$ is the proset of halfspaces, as in Example (Ex9.3) of the previous subsection.)

Lemma 9.2.1. *The map η is a monomorphism of median algebras.*

Proof. The fact that η is a homomorphism is purely formal. Let $x, y, z \in M$. Then $H \in \eta(xyz) \Leftrightarrow xyz \in H \Leftrightarrow$ at least two of x, y, z lie in $H \Leftrightarrow H$ lies in at least two of $\eta(x), \eta(y), \eta(z) \Leftrightarrow H \in \eta(x)\eta(y)\eta(z)$.

To see that η is injective, let $x, y \in M$ with $x \neq y$. By Proposition 3.2.11, there is some $H \in \mathcal{H}(M)$ with $x \in H$ and $y \notin H$. In other words, $H \in \eta(x)$ and $H \notin \eta(y)$, so $\eta(x) \neq \eta(y)$. \square

In general η need not be surjective. (See the discussion of Roller boundaries in Subsection 11.12.) However:

Lemma 9.2.2. *If M is a finite median algebra, then $\eta : M \rightarrow \mathcal{F}(\mathcal{H}(M))$ is an isomorphism.*

Proof. It remains to show that η is surjective. Let $R \in \mathcal{F}(\mathcal{H}(M))$. Then R is a finite family of pairwise intersecting convex sets. By the Helly Property (Lemma 7.1.1), $\bigcap R \neq \emptyset$. Let $x \in \bigcap R$. Then $R \subseteq \eta(x)$. But R and $\eta(x)$ are both $*$ -transversals, so $R = \eta(x)$. \square

Now let X be any set, and let $F(X)$ be the free median algebra on X .

We have a map $[H \mapsto H \cap X] : \mathcal{H}(F(X)) \rightarrow \mathcal{P}_0(X)$. In fact, this is a bijection. Its inverse arises directly from the defining property of $F(X)$. Given any $A \in \mathcal{P}_0(X)$, its characteristic function is a surjective map $X \rightarrow \{0, 1\}$, which has a unique extension to an epimorphism, $\phi_A : F(X) \rightarrow \{0, 1\}$. Now $H = \phi_A^{-1}(1)$ is a halfspace with $H \cap X = A$. Its uniqueness comes from the uniqueness of the extension ϕ_A . Moreover, the map $[H \mapsto H \cap X]$ clearly respects inclusion and complements. We have shown:

Lemma 9.2.3. *The map $[H \mapsto H \cap X] : \mathcal{H}(F(X)) \rightarrow \mathcal{P}_0(X)$ is an isomorphism of prosets.*

This induces a canonical median isomorphism $\mathcal{F}(\mathcal{H}(F(X))) \rightarrow \mathcal{F}(\mathcal{P}_0(X)) = \Phi(X)$.

Now if X is finite, Lemma 9.2.2 tells us that $\eta : F(X) \rightarrow \mathcal{F}(\mathcal{H}(F(X)))$ is an isomorphism. The composition therefore gives us a canonical median isomorphism $F(X) \rightarrow \Phi(X)$.

In fact, the explicit construction of $F(X)$ gave us $F(X) \subseteq \Phi(X)$. We can now prove Theorem 5.2.3 which asserted that for X finite, we have, $F(X) = \Phi(X)$.

Proof of Theorem 5.2.3. It is sufficient to observe that $\#F(X) = \#\Phi(X)$, since these are isomorphic.

Alternatively, unravelling the definitions, one sees that the inclusion $F(X) \subseteq \Phi(X)$ agrees with the isomorphism constructed above: they both extend the natural inclusion of X into $\Phi(X)$. \square

So far, we have made no use of general posets, other than as a source of terminology.

We now want to show how to recover median algebras from posets, so as to obtain a complete duality in the finite case.

Let Ω be a poset. By a **partial flow** on Ω we mean a flow on some subposet of Ω . In other words, $R \subseteq \Omega$ is a partial flow if there does not exist $a \in \Omega$ with $a, a^* \in R$. nor do there exist $a, b \in R$ with $a < b^*$.

Lemma 9.2.4. *Every partial flow on a poset can be extended to a flow.*

In particular, every poset admits a flow.

The proof in general uses Zorn's Lemma, though for a finite poset it amounts to a simple induction.

Proof. Let $U \subseteq \Omega$ be a partial flow. Let \mathcal{R} be the set of all partial flows that contain U , ordered by inclusion. Clearly \mathcal{R} is closed under increasing union. Let R be a maximal element of \mathcal{R} as given by Zorn's lemma. We claim that R is a flow. In other words, we need to check that R is a $*$ -transversal.

Suppose, for contradiction, that $b \in \Omega$ with $b, b^* \notin R$. If $c < b$ for some $c \in R$, set $T = R \cup \{b\}$. If $d < b^*$ for some $d \in R$, set $T = R \cup \{b^*\}$. Note that these cannot both hold, as they would imply $c < b < d^*$, contrary to the assumption that R is a partial flow. If neither holds, we set $T = R \cup \{b\}$ (or $T = R \cup \{b^*\}$ arbitrarily).

We claim that T is a flow. In the first case, suppose $e \in R$ with $b < e^*$, then $c < b < e^*$, contrary to the assumption that R is a partial flow. A similar contradiction arises in the second case. In the third case, $b < e^*$ would imply $e < b^*$ contrary to our assumption.

This contradicts the maximality of R , and so the lemma follows. \square

As a variation on the last lemma, we have:

Lemma 9.2.5. *Suppose that Ω is a poset, and that $\Omega_0 \subseteq \Omega$ is a subposet such that either $a = b^*$ or $a \pitchfork b$ for all distinct $a, b \in \Omega_0$. Then there is a flow, R , on $\Omega \setminus \Omega_0$ such that for any flow U on Ω_0 , $R \cup U$ is a flow on Ω .*

(Note that in this case, any $*$ -transversal on Ω_0 will be a flow on Ω_0 .)

Proof. Let \mathcal{S} be the set of partial flows, S , on $\Omega \setminus \Omega_0$ such that $S \cup U$ is a flow on Ω for all flows U on Ω_0 . As with Lemma 9.2.4, \mathcal{S} has a maximal element, R . We claim that R is a flow on Ω .

For suppose $b \in \Omega$ with $b, b^* \notin R$. If $c < b$ for some $c \in R \cup \Omega_0$, set $T = R \cup \{b\}$. If $d < b^*$ for some $d \in R \cup \Omega_0$, set $T = R \cup \{b^*\}$. Again we claim that these cannot both hold. For if they did, we would get $c < d^*$. In this case, c, d cannot both lie in Ω_0 by hypothesis; neither can they both lie R , since R is a partial flow. If $c \in \Omega_0$ and $d^* \in R$, choose any flow U on Ω_0 containing c , and we contradict the assumption that $R \cup U$ is a flow. We get a similar contradiction if $c \in R$ and $d^* \in \Omega_0$. This proves the claim. If neither of these conditions holds, we (arbitrarily) set $T = R \cup \{b\}$.

Similarly as with Lemma 9.2.4, we see that T is a flow, contrary to maximality of R . \square

Given $a \in \Omega$, let $H(a) = \{R \in \mathcal{F}(\Omega) \mid a \in R\}$. Lemma 9.2.4 (applied to the partial flow $\{a\}$) tells us that $H(a) \neq \emptyset$. Note that $H(a)$ is convex in $\mathcal{F}(\Omega)$. We also have $H(a^*) = (H(a))^*$. In other words, $\Omega = H(a) \sqcup H(a^*)$. Moreover, if $a \leq b$, then $H(a) \leq H(b)$ (since if $R \in H(a)$, then $a \in R$, so $b \in R$ since R is a flow, and so $R \in H(b)$).

This construction gives us a map: $H : \Omega \longrightarrow \mathcal{H}(\mathcal{F}(\Omega))$.

Lemma 9.2.6. *H is a monomorphism of prosets.*

Proof. We have already observed that H is a proset morphism, so it remains to check that H is injective.

Let $a, b \in \Omega$ with $a \neq b$. If $b = a^*$, then $H(b) = (H(a))^*$, so certainly $H(a) \neq H(b)$. So suppose $b \neq a^*$. Let R be a flow on $\Omega \setminus \{a, a^*\}$ as given by Lemma 9.2.5 (with $\Omega = \{a, a^*\}$). Thus $S := R \cup \{a\}$ and $T := R \cup \{a^*\}$ are both flows on Ω . Now either b lies in both S and T , or in neither. In other words, either $S, T \in H(b)$ or else $S, T \in (H(b))^*$. However, $S \in H(a)$ and $T \in (H(a))^*$, and so $H(a) \neq H(b)$. \square

Proposition 9.2.7. *If Ω is a finite proset, then $H : \Omega \longrightarrow \mathcal{H}(\mathcal{F}(\Omega))$ is an isomorphism of prosets.*

Proof. It remains to check that H is surjective.

Let $G \in \mathcal{H}(\mathcal{F}(\Omega))$. That is, G is a halfspace of the finite median algebra, $\mathcal{F}(\Omega)$. By Lemma 7.3.6, there are mutual gates, $R \in G$, and $S \in G^*$. Since $R \neq S$, there is some $a \in \Omega$ with $a \in R$ and $a^* \in S$. We claim that $G = H(a)$.

To see this, let $T \in G$. Since R is a gate for S in G , we have $T.R.S$ in the median structure of $\mathcal{F}(\Omega)$. This is equivalent to the statement $R \subseteq S \cup T$. Now since $a \in R$ and $a^* \in S$, we must have $a \in T$. That is, $T \in H(a)$. This shows that $G \subseteq H(a)$. Similarly (swapping the roles of a and a^*) we have $G^* \subseteq H(a^*)$. It follows that $G = H(a)$ as required. \square

In summary, we have a complete duality between finite median algebras and finite prosets, given by the maps $[M \mapsto \mathcal{H}(M)]$ and $[\Omega \mapsto \mathcal{F}(\Omega)]$, where $\mathcal{H}(M)$ is the proset of halfspaces of M , and $\mathcal{F}(\Omega)$ is the median algebra of flows on Ω .

There are more general duality results for infinite median algebras and prosets, though one needs to equip the spaces with appropriate topologies. One such result, due to Roller [R], is given as Theorem 12.5.1 here. Given some straightforward facts about topological median algebras, the proof of this just involves adapting the arguments we have already given. We will postpone discussion of this until Subsection 12.5.

We proceed to make some further observations regarding the constructions we have described.

9.3. The rank of a proset.

First we can give some further description of intervals in $\mathcal{F}(\Omega)$ for a proset Ω . Given $R, S \in \mathcal{F}(\Omega)$, recall that $R \triangle S$ is a subproset of Ω . We have observed that $T \in [R, S]$ if and only if $R \triangle S = (R \triangle T) \sqcup (S \triangle T)$. Note that $T \cap (R \triangle S)$ is a flow on $R \triangle S$. Conversely, given any flow, U , on $R \triangle S$, $(R \cap S) \cup U$ is a flow on Ω . (To see this, note that $(R \cap S) \cup U$ is a $*$ -transversal. Suppose, for contradiction, that $a, b \in (R \cap S) \cup U$, with $a < b^*$. We cannot have $a, b \in R \cap S$ or $a, b \in U$, since these are all flows. If $a \in R \cap S$ and $b \in U$, then without loss of generality $b \in R$, contrary to the fact that R is a flow. We get a similar contradiction if $a \in U$ and $b \in R \cap S$.) This gives us a natural bijection between $[R, S]$ and $\mathcal{F}(R \triangle S)$.

Note that any proset with at least four elements admits at least three flows. It follows that $R, S \in \mathcal{F}(\Omega)$ are adjacent if and only if $S = (R \setminus \{a\}) \cup \{a^*\}$ for some $a \in R$. In other words, R, S differ by applying the involution to a single element. (We generally refer to such an operation as **flipping** that element.)

Definition. The **rank**, $\text{rank}(\Omega)$ of a proset, Ω is the maximal cardinality of a finite subset A such that $a \pitchfork b$ for all distinct $a, b \in A$. We deem it to be ∞ if there is no such bound.

Lemma 9.3.1. $\text{rank}(\Omega) = \text{rank}(\mathcal{F}(\Omega))$.

Proof. The fact that $\text{rank}(\Omega) \leq \text{rank}(\mathcal{F}(\Omega))$ follows easily from Lemma 9.2.5. (Note that any $*$ -transversal on $\Omega_0 := \{a, a^* \mid a \in A\}$ is a flow, and these flows form a cube of rank $\#A$, embedded in $\mathcal{F}(\Omega)$.)

To see that $\text{rank}(\mathcal{F}(\Omega)) \leq \text{rank}(\Omega)$, suppose that $Q \subseteq \mathcal{F}(\Omega)$ is an n -cube. Choose any $R \in Q$, and let $S_1, \dots, S_n \in Q$ be the adjacent vertices. For each i , choose some $a_i \in R \setminus S_i$. We claim that $a_i \pitchfork a_j$ for all $i \neq j$. Since R is a flow, we have $a_i \not\prec a_j^*$. Since $R \in [S_i, S_j]$ we have $R \subseteq S_i \cup S_j$, so $a_i \in R \setminus S_i \subseteq S_j$. Since $a_j^* \in S_j$, and S_j is a flow, we have $a_i \not\prec a_j^{**} = a_j$. Similarly, $a_j \not\prec a_i$. Finally, let $T \in Q$ be the fourth vertex of the square containing R, S_i, S_j . Now $S_i \in [R, T]$, so $S_i \subseteq R \cup T$, so $a_i^* \in S_i \setminus R \subseteq T$. Similarly, $a_j^* \in T$. Since T is a flow, we get $a_i^* \not\prec a_j^{**} = a_j$. This shows that $a_i \pitchfork a_j$, proving the claim. We set $A = \{a_1, \dots, a_n\}$. \square

One can define a notion of an “unoriented” proset. Such a structure arises as a quotient, $\Omega/*$, of a proset, Ω , by the involution, $[a \mapsto a^*]$. As such it inherits a binary “crossing” relation, and a ternary “betweenness” relation, as we describe below. One can axiomatise the class of such structures, so as to be able to recover a proset in a canonical way. We will not do this explicitly here, since we do not need it. We will just describe a few of the basic properties which will be useful later.

As particular examples, we can think of $\mathcal{P}_0(X)/*$ as the set of all non-trivial partitions of X into two non-empty subsets. If M is a median algebra, then we can identify $\mathcal{H}(M)/*$ with $\mathcal{W}(M)$. In the example of an arrangement of lines (Example (Ex9.5) above) we just forget the transverse orientations etc.

Let Ω be a proset. A typical element of $\Omega/*$ has the form $\alpha = \{a, a^*\}$ for $a \in \Omega$. Given $\alpha, \beta \in \Omega/*$, write $\alpha \pitchfork \beta$ to mean $a \pitchfork b$ for some (hence any) $a \in \alpha$ and $b \in \beta$. In this case, we say α **crosses** β . Given $\alpha, \beta, \gamma \in \Omega/*$, write $\alpha : \beta : \gamma$ to mean that $a < b < c$ for some $a \in \alpha$, $b \in \beta$ and $c \in \gamma$. Note that this implies $\gamma : \beta : \alpha$ (since $c^* < b^* < a^*$) and that α, β, γ are pairwise non-crossing. We think of β as lying strictly between α and γ . (This has an obvious interpretation in each of the examples of Subsection 9.1.) We say that α, β, γ are **nested** if one of the statements $\alpha : \beta : \gamma$, $\beta : \gamma : \alpha$ or $\gamma : \alpha : \beta$ holds. It is easily checked that these three cases are mutually exclusive.

We say that $\alpha, \beta, \gamma \in \Omega/*$ are **unnested** if $a^* < b$, $b^* < c$ and $c^* < a$, for some $a \in \alpha$, $b \in \beta$, $c \in \gamma$. Note that this is equivalent to $b^* < a$, $c^* < b$ and $a^* < c$, so it is independent of the order of α, β, γ . Again, α, β, γ are pairwise non-crossing, and not nested.

To justify the terminology, we note:

Lemma 9.3.2. *If $\alpha, \beta, \gamma \in \Omega/*$ are distinct and pairwise non-crossing, then α, β, γ are either nested or unnested.*

Proof. Since $\alpha \not\pitchfork \beta$, we have $a^* < b$ for some $a \in \alpha$ and $b \in \beta$. Also, we have either $c < a^*$ or $a^* < c$ for some $c \in \gamma$. Now one of the relations $b < c$, $c < b$, $b < c^*$ or $c^* < b$ must hold. These four cases give respectively $\alpha : \beta : \gamma$, $\alpha : \gamma : \beta$, the contradiction $a^* < c^*$, and that α, β, γ are unnested. \square

For future reference, we make the following observation.

Lemma 9.3.3. *Any infinite proset admits infinitely many flows.*

Proof. First suppose that $\text{rank}(\Omega) = \infty$. Then we can find subsets of Ω of arbitrarily large finite cardinality whose elements pairwise cross. We can arbitrarily swap the orientation on any subset of such a set to give a partial flow on Ω , which then extends to a flow. This gives rise to infinitely many distinct flows.

On the other hand, if $\text{rank}(\Omega) < \infty$, then we can find an infinite chain $\Phi \subseteq \Omega$ (for example, applying Dilworth’s Lemma). If $a \in \Phi$, we obtain a partial flow as $\{b^* \mid b \in \Phi, b \leq a\} \cup \{b \in \Phi \mid a < b\}$. Extending to Ω , these again give infinitely many distinct flows. \square

Since the interval $[R, S]$ is naturally identified with the set of flows on $R \triangle S$, as an immediate corollary we have:

Corollary 9.3.4. *If R, S are flows on a proset Ω , then $[R, S]$ is finite if and only if $R \triangle S$ is finite.*

9.4. The discrete case.

The following construction will be referred to again in Section 19. A **downward sequence** is a subset of Ω which is order isomorphic to $-\mathbb{N} \subseteq \mathbb{Z}$. In other words, we can index it as an infinite sequence $a_0 > a_1 > a_2 > \dots$. We write $\mathcal{F}_1(\Omega) \subseteq \mathcal{F}(\Omega)$ for the set of flows $\mathcal{F}(\Omega)$ which do not contain any downward sequence. We claim that $\mathcal{F}_1(\Omega)$ is convex in $\mathcal{F}(\Omega)$. To see this, suppose that $R, S \in \mathcal{F}_1(\Omega)$ and $T \in [R, S]_{\mathcal{F}(\Omega)}$. Then $T \subseteq R \cup S$, and so any infinite downward sequence in T would have to contain an infinite subset in either R or S , giving a contradiction. It follows that $T \in \mathcal{F}_1(\Omega)$ as required.

We remark that in general, $\mathcal{F}_1(\Omega)$ might be empty. For example, let $P = \{1/n \mid n \in \mathbb{Z} \setminus \{0\}\} \subseteq \mathbb{R}$ with the usual linear order, and let $\Omega = P \times \{-1, 1\}$. Let $(x, \delta)^* = (-x, -\delta)$ and write $(x, \delta) < (y, \epsilon)$ if $x < y$ and $\delta = \epsilon$. Then Ω is a proset with $\mathcal{F}_1(\Omega) = \emptyset$. We will give a condition under which $\mathcal{F}_1(\Omega) \neq \emptyset$ as Proposition 9.5.11.

We say that a median algebra is **discrete** if all intervals are finite. (This will be the topic of Section 11.)

We say that a proset is **subinfinite-rank** if any set of pairwise-crossing elements is finite. Clearly finite-rank implies subinfinite-rank.

Lemma 9.4.1. *Suppose that Ω is a subinfinite-rank proset. Then $\mathcal{F}_1(\Omega)$ is a discrete median algebra.*

Proof. Let $R, S \in \mathcal{F}_1(\Omega)$. Now $\Omega_0 := R \triangle S$ is a subproset of Ω . Moreover, $[R, S]_{\mathcal{F}(\Omega)}$ consists of those $T \in \mathcal{F}(\Omega)$ for which $\Omega_0 = (R \triangle T) \sqcup (S \triangle T)$. To show that $[R, S]_{\mathcal{F}(\Omega)}$ is finite, we therefore need to show that Ω_0 is finite.

First note that $\Omega_0/*$ has no unnested triples. For suppose $a, b, c \in \Omega_0$ with $a^* < b, b^* < c$ and $c^* < a$. Without loss of generality, we can suppose that $a, b \in R$. Then $a^*, b^* \in S$, and $a^* < b^{**}$, contradicting the fact that S is a flow.

Now suppose for contradiction that Ω_0 is infinite. We note that there is an infinite subproset, $\Omega_1 \subseteq \Omega_0$, such that $a \not\# b$ for all $a, b \in \Omega_1$. (This follows by applying Ramsey's Theorem for infinite graphs to the complete graph on $\Omega_0/*$, where we colour each edge, α, β , according to whether $\alpha \# \beta$ or $\alpha \not\# \beta$. By hypothesis there is no infinite monochromatic clique of the first colour, so there must be one of the second. This gives us $\Omega_1/*$.)

By Lemma 9.3.2, any three elements of Ω_1 are nested. It follows that we can partition Ω_1 into two totally ordered sets, swapped by the involution on Ω_1 . In particular, Ω_1 contains a downward sequence $a_0 > a_1 > a_2 > \dots$. Now $\Omega_1 \subseteq R \cup S$, and so this has an infinite subsequence in either R or S , contradicting $R, S \in \mathcal{F}_1(\Omega)$. \square

Note that the rank of Ω is also the maximal cardinality of any set of pairwise crossing elements of $\Omega/*$. By Lemma 9.3.1, we have $\text{rank } \mathcal{F}_1(\Omega) \leq \text{rank } \mathcal{F}(\Omega) = \text{rank } \Omega$.

Example. As an example, consider a locally finite family of bi-infinite geodesics in the hyperbolic plane \mathbb{H}^2 . Suppose that there is some $\theta > 0$ such that no two of these geodesics meet at an angle less than θ . A simple geometric argument shows that there is some $\nu = \nu(\theta) \in \mathbb{N}$, such that any subset of pairwise crossing geodesics has cardinality at most ν . It follows that $\mathcal{F}_1(\Omega)$ is a discrete median algebra of rank at most ν . (This is essentially the same as saying that it is a CAT(0) cube complex of dimension at most ν .)

Such a situation might arise, for example, if we have a finite set of closed geodesics on a closed hyperbolic (i.e. constant curvature -1) surface, Σ . The universal cover of Σ is isometric to \mathbb{H}^2 , and the preimages of the closed geodesics give rise to such a family of bi-infinite geodesics. We therefore get a finite-dimensional CAT(0) cube complex with a natural action of $\pi_1(\Sigma)$.

In higher dimensions, one can generalise to a family of codimension-1 hyperplanes: see Example (Ex19.4) of Subsection 19.4.

Let X be any set, and let $\mathcal{H} \subseteq \mathcal{P}_0(X)$ be a subproset (i.e. closed under the involution $*$). Let $\mathcal{F} = \mathcal{F}(\mathcal{H})$. We have a map $\eta : X \rightarrow \mathcal{F}$ defined by $\eta(x) = \{H \in \mathcal{H} \mid x \in H\}$ as in Subsection 9.2. We say that \mathcal{H} is *discrete* if $\eta(x) \Delta \eta(y)$ is finite for all $x, y \in X$. (This is a “space with walls” which we discuss further in Section 19.) For the remainder of this subsection, we will assume that \mathcal{H} is discrete.

We say that a flow, $R \in \mathcal{F}$, is *almost principal* if $R \Delta \eta(x)$ is finite for some (hence any) $x \in X$. We write $\mathcal{F}_0 \subseteq \mathcal{F}$ for the set of almost principal flows. Note that this is a subalgebra of \mathcal{F} . Also, if $R, S \in \mathcal{F}_0$, then $R \Delta S$ is finite. It follows that $[R, S]_{\mathcal{F}_0}$ is finite, so \mathcal{F}_0 is a discrete median algebra. We also note that $\mathcal{F}_0 \subseteq \mathcal{F}_1(\mathcal{H})$ as defined above.

We can get from S to R by a sequence of “flips”, replacing some element $A \in S$ by A^* at each stage. One way to see this is to note that two elements of \mathcal{F}_0 are adjacent if and only if they differ by a flip. This follows as in the proof of Lemma 5.3.2. (It was assumed there that X is finite, but the argument only really required that $R \Delta S$ be finite.) Alternatively, we can use induction on $\#(S \setminus R)$: if $A \in S \setminus R$ is minimal with respect to inclusion, then $(S \setminus \{A\}) \cup \{A^*\} \in \mathcal{F}_0$.

Let $\langle \eta(X) \rangle$ be the subalgebra of \mathcal{F} generated by $\eta(X)$. Clearly, $\langle \eta(X) \rangle \subseteq \mathcal{F}_0$. In fact:

Proposition 9.4.2. $\mathcal{F}_0 = \langle \eta(X) \rangle$.

Proof. If not, by the above observation, we can find $P \in \langle \eta(X) \rangle$ and $R \in \mathcal{F}_0 \setminus \langle \eta(X) \rangle$, with $P = (R \setminus \{A\}) \cup \{A^*\}$ for some $A \in R$. Let $\mathcal{S} = \{S \in \langle \eta(X) \rangle \mid A \in S\} \subseteq \mathcal{F}_0$. Note that $\mathcal{S} \neq \emptyset$ (since $\eta(x) \in \mathcal{S}$ for any $x \in A$). Choose $S \in \mathcal{S}$ with $\#(S \setminus P)$ minimal. We claim that $S = R$. Since R, S are both $*$ -transversals, it is enough to show that $S \subseteq R$.

Suppose to the contrary that there is some $B \in S \setminus R$. Now $A, B \in S$, so $B \neq A^*$. Therefore, we also have $B \in S \setminus P$. Since $A, B^* \in R$, and R is a flow, we have $A \cap B^* \neq \emptyset$. Choose $x \in A \cap B^*$ and let $T = \eta(x)$ and $U = PST$. Since $P, S, T \in \langle \eta(X) \rangle$, $U \in \langle \eta(X) \rangle$. Now $A \in S \cap T \subseteq U$, and so $U \in \mathcal{S}$. Since $U \subseteq P \cup S$, we have $U \setminus P \subseteq S \setminus P$. Also $B^* \in P \cap T \subseteq U$, so $B \notin U$. Thus, $U \setminus P$ is a proper subset of $S \setminus P$, contradicting the minimality of $\#(S \setminus P)$. We have shown that $S \subseteq R$, and so $R = S$ as claimed. Since $R \notin \mathcal{S}$, this is a contradiction. \square

In the case where X is finite and $\mathcal{H} = \mathcal{P}_0(X)$, then $\mathcal{F}_0 = \mathcal{F}$ is the superextension, $\Phi(X)$, of X , as defined in Subsection 5.1. This gives another proof of Theorem 5.2.3.

9.5. The structure of the space of flows.

We make some further observations about the structure of flows on a proset. We describe a quotient, $\hat{\mathcal{F}}$, of the set of flows. Under certain assumptions, this will be discrete (Proposition 9.5.4). The equivalence class of a flow can be described in terms of the set of downward sequences which it contains. This will be of relevance to the Roller boundary as we discuss in Subsection 11.12 (see Proposition 11.12.10 and related discussion). We will make frequent use of the fact (Lemma 9.2.4) that any partial flow can be extended to a flow.

Let Ω be a proset, and $\mathcal{F} = \mathcal{F}(\Omega)$ be the set of flows. We define an equivalence relation, \sim , on \mathcal{F} by writing $R \sim S$ if $\#(R \triangle S) < \infty$. Note that by Corollary 9.3.4, $R \sim S$ if and only if $\#[R, S] < \infty$. In other words, \sim is the same as the relation that we defined on an arbitrary median algebra in Example (Ex3.7) of Subsection 3.4.

Let $\hat{\mathcal{F}} = \hat{\mathcal{F}}(\Omega) = \mathcal{F}/\sim$ be the quotient median algebra. Denoting equivalence classes by $[\cdot]$, by definition we have $[R][S][T] = [RST]$ for $R, S, T \in \mathcal{F}$. Note that if $[R].[T].[S]$ holds, then $[T] = [RST]$, so (replacing T by RST) there is no loss in assuming that $R.T.S$ holds. As we have noted, the latter is equivalent to saying that $R \cap S \subseteq T$, or that $T \subseteq R \cup S$ in Ω .

Recall the notion of a “downward sequence” in Ω , as defined in the previous subsection. We say that $b \in \Omega$ is a **lower bound** for the sequence α if $b < a$ for all $a \in \alpha$. A subset Ω' of Ω is **boundless** if no downward sequence in Ω' has a lower bound in Ω . Of course we can apply this terminology to a downward sequence or to Ω itself. (Note that if Ω is subinfinite-rank, as defined in the previous subsection, then Ω is boundless if and only if $\{c \in \Omega \mid a \leq c \leq b\}$ is finite for all $a, b \in \Omega$.)

Let $\mathcal{L} = \mathcal{L}(\Omega)$ be the set of downward sequences in Ω . Given $\alpha, \beta \in \mathcal{L}$, write $\alpha \approx \beta$ to mean that they are cofinal in the sense that $(\forall a \in \alpha)(\exists b \in \beta)(b < a)$ and $(\forall b \in \beta)(\exists a \in \alpha)(a < b)$. Clearly, this is an equivalence relation, and we write $\hat{\mathcal{L}} = \hat{\mathcal{L}}(\Omega) = \mathcal{L}/\approx$. We again denote equivalence classes by $[\cdot]$. Note that if $\alpha \in \mathcal{L}$, then $\bigcup[\alpha] \subseteq \Omega$ is a partial flow. (In fact, $\bigcup[\alpha]$ is the set of $b \in \Omega$ with $b > a$ for some $a \in \alpha$.)

Suppose $R \in \mathcal{F}$. Note that if $\alpha \in \mathcal{L}$ with $\#(\alpha \cap R) = \infty$, then in fact $\alpha \subseteq R$. (For if $a \in \alpha$, there is some $b \in \alpha \cap R$ with $b < a$, so $a \in R$.) Write $\mathcal{L}(R) = \{\alpha \in \mathcal{L} \mid \alpha \subseteq R\}$. It is easily checked that if $\alpha \in \mathcal{L}(R)$ and $\beta \approx \alpha$, then $\beta \in \mathcal{L}(R)$. Moreover, if $R \sim S$, then $\mathcal{L}(R) = \mathcal{L}(S)$. This gives us a map,

$$\lambda : \hat{\mathcal{F}} \longrightarrow \mathcal{P}(\hat{\mathcal{L}}),$$

where $\lambda([R]) = \{[\alpha] \mid \alpha \in \mathcal{L}(R)\}$ for all $R \in \mathcal{F}$. In fact, if $[\alpha] \in \lambda([R])$, then $\alpha \in \mathcal{L}(R)$ for any representative α .

The image, $\lambda([R])$, will give us information about the structure of $[R]$ and its position in the quotient space $\hat{\mathcal{F}}$. In particular, we are aiming at Propositions 9.5.4, 9.5.11 and 9.5.14. We begin with:

Lemma 9.5.1. *λ is a median homomorphism.*

Proof. It is enough to show that if $[R].[T].[S]$ holds for $R, S, T \in \mathcal{F}$, then $\lambda([R]).\lambda([T]).\lambda([S])$ holds in $\mathcal{P}(\hat{\mathcal{L}})$. This is purely formal. We observed earlier that we can assume that $R.T.S$ holds in \mathcal{F} : in other words $R \cap S \subseteq T$. For the conclusion, we want to show that $\lambda([R]) \cap \lambda([S]) \subseteq \lambda([T])$. Now if $[\alpha] \in \lambda([R]) \cap \lambda([S])$, then $\alpha \in \mathcal{L}(R) \cap \mathcal{L}(S)$, so $\alpha \subseteq R \cap S \subseteq T$, so $\alpha \in \mathcal{L}(T)$, so $[\alpha] \in \lambda([T])$ as required. \square

Lemma 9.5.2. *If Ω is subinfinite-rank, then λ is injective.*

Proof. Suppose $R, S \in \mathcal{F}$ with $R \not\sim S$. The proset $R \Delta S$ is infinite, and as in the proof of Lemma 9.4.1, we can find a downward sequence, $\alpha \subseteq R \Delta S$. We can suppose that $\#(\alpha \cap R) = \infty$. But then $\alpha \subseteq R$. Thus, $\alpha \in \mathcal{L}(R) \setminus \mathcal{L}(S)$, so $[\alpha] \in \lambda([R]) \setminus \lambda([S])$. In particular $\lambda([R]) \neq \lambda([S])$. \square

The following construction will be used at various points. Suppose $A \subseteq \hat{\mathcal{L}}$. We can choose a representative, α , for each $[\alpha] \in A$. (In subsequent notation, we will generally assume that we have done so.) We set $U = U(A) = \bigcup_{[\alpha] \in A} \alpha \subseteq \Omega$. If Ω is boundless (no downward sequence has a lower bound) and A is finite, we can always choose our representatives such that there do not exist $a, b \in U$ with $a^* \leq b$. (For suppose $[\alpha], [\beta] \in A$. Let a be the initial (maximal) element of α . Since β has no lower bound, there must be some $b \in \beta$ such that $a^* \leq b$ fails. It follows that $c^* \leq d$ also fails for all $c \leq a$ and $d \leq b$. We can now replace β by $\{d \in \beta \mid d \leq b\}$. We can do this in turn for each pair of elements of A .)

Lemma 9.5.3. *Suppose Ω is boundless, and let $n \in \mathbb{N}$. Suppose that $R \in \mathcal{F}$ is such that every set of pairwise crossing elements of R has at most n elements. Then $\#\lambda([R]) \leq n$.*

Proof. Write $A = \lambda([R])$. Suppose, for contradiction, that $\#A > n$. Let $U = U(A)$ be as constructed above. Now $U \subseteq R$. Since R is a flow, we cannot have $a < b^*$ or $a = b^*$ for any $a, b \in U$. By construction of U , we cannot have $a^* \leq b$. Thus, if $a, b \in U$ are distinct, we must have $a \cap b$, $a < b$ or $b < a$. It follows by hypothesis that any antichain in U has cardinality at most n . By Dilworth's Lemma (cf. the

proof of Lemma 8.3.2) we can write U as a union of n chains: $U = U_1 \cup \dots \cup U_n$. If $[\alpha] \in A$, then $\#(\alpha \cap U_{i(\alpha)}) = \infty$ for at least one $i(\alpha) \in \{1, \dots, n\}$. Since $\#A > n$, there are distinct $[\alpha], [\beta] \in A$ with $i(\alpha) = i(\beta)$. In other words, $\alpha \cap V$ and $\beta \cap V$ are both infinite, where $V = U_{i(\alpha)}$. Since $\alpha \not\preceq \beta$, after swapping α and β , we can suppose that there is some $a \in \alpha$ such that $b \not\preceq a$ for all $b \in \beta$. Since V is a chain, we must have $a \leq b$. This implies that a is a lower bound for β , contrary to our assumption on R . \square

As a consequence, we get the following:

Proposition 9.5.4. *If Ω is a subinfinite-rank boundless proset, then $\hat{\mathcal{F}}(\Omega)$ is discrete.*

Proof. Suppose $[R], [S] \in \hat{\mathcal{F}}$. By Lemmas 9.5.1 and 9.5.2, λ maps the interval $[[R], [S]] \subseteq \hat{\mathcal{F}}$ injectively into the interval $[\lambda([R]), \lambda([S])] \subseteq \mathcal{P}(\hat{\mathcal{L}})$. Now any element of $[\lambda([R]), \lambda([S])]$ is a subset of $\lambda([R]) \cup \lambda([S]) \subseteq \hat{\mathcal{L}}$. Moreover, $\lambda([R]) \cup \lambda([S])$ is finite by Lemma 9.5.3, and so there are only finitely many possibilities for this element. In other words, $[\lambda([R]), \lambda([S])]$ is finite, so $[[R], [S]]$ is finite as required. \square

Given a subset, $A \subseteq \hat{\mathcal{L}}$, let $\mathcal{F}_A = \{R \in \mathcal{F} \mid A \subseteq \lambda([R])\}$. Clearly this is a union of \sim -classes. In fact:

Lemma 9.5.5. *\mathcal{F}_A is convex in \mathcal{F} .*

Proof. Let $R, S \in \mathcal{F}_A$ and $T \in [R, S] \subseteq \mathcal{F}$. If $[\alpha] \in A \subseteq \lambda([R]) \cap \lambda([S])$, then $\alpha \subseteq R \cap S \subseteq T$, so $[\alpha] \in \lambda([T])$. This shows that $A \subseteq \lambda([T])$, so $T \in \mathcal{F}_A$. \square

Let $U = U(A) = \bigcup_{[\alpha] \in A} \alpha \subseteq \Omega$, as defined earlier.

Lemma 9.5.6. *$\mathcal{F}_A \neq \emptyset$ if and only if U is a partial flow on Ω .*

Proof. If $R \in \mathcal{F}_A$, then $U \subseteq R$, so U is a partial flow. Conversely, if U is a partial flow, we can extend it to a flow, $R \subseteq \Omega$. If $[\alpha] \in A$, then $\alpha \subseteq U \subseteq R$, so $[\alpha] \in \lambda([R])$. This shows that $A \subseteq \lambda([R])$, so $R \in \mathcal{F}_A$, and so $\mathcal{F}_A \neq \emptyset$. \square

Suppose $\alpha \in \mathcal{L}$. It is easily seen that if an element, $b \in \Omega$, crosses infinitely many elements of α , then it crosses all but finitely many elements of α . We write $\Omega_\alpha \subseteq \Omega$ for the set of such b . We also easily check that this only depends on the \sim -class of α . (These statements are both based on the observation that if $a < a' < a''$ with $a \pitchfork b$ and $a'' \pitchfork b$, then $a' \pitchfork b$.) Given a subset, $A \subseteq \hat{\mathcal{L}}$ the set $\Omega_A := \bigcap_{[\alpha] \in A} \Omega_\alpha$ is therefore well defined. Note that Ω_A is a subproset of Ω . Define a map $\theta : \mathcal{F}_A \rightarrow \mathcal{F}(\Omega_A)$ by restricting flows to Ω_A . This is clearly a median homomorphism. Under the combined hypotheses of the following two lemmas, this is an isomorphism.

Lemma 9.5.7. *If Ω is boundless then θ is injective.*

Proof. Suppose for contradiction that $R, S \in \mathcal{F}_A$ are distinct with $\theta(R) = \theta(S)$. In other words, R, S agree on Ω_A . Since $R \neq S$, we can find $b \in R$ with $b^* \in S$, and so $b \notin \Omega_A$. Thus, there is some $[\alpha] \in A$ with $b \notin \Omega_\alpha$. We can suppose that $b \not\prec a$ for all $a \in \alpha$. Now $a \subseteq \lambda([R]) \cap \lambda([S])$, so $\alpha \subseteq R \cap S$, and it follows that neither $a < b$ nor $a < b^*$ can hold. Therefore we have either $b < a$ or $b^* < a$. It follows that either b or b^* is a lower bound for α , contrary to our hypothesis on Ω . \square

Lemma 9.5.8. *If $\mathcal{F}_A \neq \emptyset$, then θ is surjective.*

Proof. Let U be as in Lemma 9.5.6. Since $\mathcal{F}_A \neq \emptyset$, U is a partial flow. If $R \in \mathcal{F}(\Omega_A)$, then $R \cup U$ is also a partial flow. (Otherwise we would have $b \in R$ and $a < b^*$ for some $a \in \alpha$ where $[\alpha] \in A$. Then $c < b^*$ for all $c < a$, so $c \not\prec b$. But since $b \in R \subseteq \Omega_A \subseteq \Omega_\alpha$, this gives a contradiction.) We now extend $R \cup U$ to a flow, S , on Ω . Then $\theta(S) = R$ and $S \in \mathcal{F}_A$. \square

(Of course, it is possible that $\Omega_A = \emptyset$, in which case, $\mathcal{F}(\Omega_A) = \{\emptyset\}$.)

We can now apply this as follows:

Lemma 9.5.9. *Suppose Ω is boundless and $A \subseteq \hat{\mathcal{L}}$ with $\mathcal{F}_A \neq \emptyset$. Then $\text{rank}(\mathcal{F}_A) + \#A \leq \text{rank}(\mathcal{F}(\Omega))$.*

Here we can assume that all these numbers are finite: if $\#A$ were infinite, then applying the result for arbitrarily large finite A we get that $\text{rank}(\Omega)$ is infinite. That is all we are claiming in that regard.

Proof. By Lemmas 9.5.7 and 9.5.8, we have $\mathcal{F}_A \cong \mathcal{F}(\Omega_A)$, and so $\text{rank}(\mathcal{F}_A) = \text{rank}(\mathcal{F}(\Omega_A))$. (We only really need Lemma 9.5.7 for the current proof.) Therefore, by Lemma 9.3.1, the inequality is equivalent to asserting that $\text{rank}(\Omega_A) + \#A \leq \text{rank}(\Omega)$. Let $A = \{[\alpha_1], \dots, [\alpha_n]\}$ with $\alpha_1, \dots, \alpha_n \in \mathcal{L}$, with $\alpha_i \not\approx \alpha_j$ for $i \neq j$. (So $n = \#A$.) Let $U = \alpha_1 \cup \dots \cup \alpha_n$. As discussed before Lemma 9.5.3, we can assume that there do not exist $a, b \in U$ with $a^* \leq b$. By Lemma 9.5.6, U is a partial flow on Ω , so we cannot have $a < b^*$ either. Thus one of $a < b$, $b < a$ or $a \pitchfork b$ must hold. Let $b_1, \dots, b_m \in \Omega_A$ be a set of pairwise crossing elements of Ω_A . We want to show that $m + n \leq \text{rank}(\Omega)$. We can assume that $b_i \pitchfork a$ for all i and for all $a \in U$ (after removing finitely many elements from each α_j). If there were a set of n pairwise crossing elements of U , then taking this set together with $\{b_1, \dots, b_m\}$, we would get $\text{rank}(\Omega) \geq m + n$, as required. So suppose, for contradiction, that no such set exists. It follows that any antichain in U has cardinality at most $n - 1$. We can now proceed similarly as in the proof of Lemma 9.5.3. By Dilworth's Lemma, we can partition U into $n - 1$ chains. At least one of these chains meets two distinct α_i and α_j each in an infinite set. Since $\alpha_i \not\approx \alpha_j$, we contradict the hypothesis that Ω is boundless, similarly as before. \square

Corollary 9.5.10. *Suppose Ω is boundless. If $R \in \mathcal{F}$, then $\text{rank}([R]) + \#\lambda([R]) \leq \text{rank} \mathcal{F}(\Omega)$.*

Proof. Let $A = \lambda([R])$. Then $[R] \subseteq \mathcal{F}_A$, so $\text{rank}([R]) \leq \text{rank}(\mathcal{F}_A)$. The statement now follows from Lemma 9.5.9. \square

We next want to interpret the quantity $\#\lambda([R])$ in terms of the intrinsic structure of $\hat{\mathcal{F}}$.

For the remainder of this subsection, we will assume that Ω is non-empty, subinfinite-rank and boundless.

Since λ is injective, there is at most one $[R] \in \hat{\mathcal{F}}$ with $\lambda([R]) = \emptyset$. This class (if it exists) is precisely the convex subset of flows $\mathcal{F}_1(\Omega) \subseteq \mathcal{F}$, defined in the previous subsection. We observed there that in general this might be empty. However, under the present assumption, we have:

Proposition 9.5.11. *If Ω is subinfinite-rank and boundless, then $\mathcal{F}_1(\Omega) \neq \emptyset$.*

We will give this as a corollary of a more general result, namely Lemma 9.5.12. (For the application to the Roller boundary in Subsection 11.12, the existence of such flows is clear.)

First we give some more general discussion.

Given $\alpha, \beta \in \mathcal{L}$, write $\alpha \preceq \beta$ to mean $(\forall b \in \beta)(\exists a \in \alpha)(a < b)$. Note that \preceq is transitive, and by definition $\alpha \approx \beta \Leftrightarrow (\alpha \preceq \beta \ \& \ \beta \preceq \alpha)$. We can also write $[\alpha] \preceq [\beta]$ to mean $\alpha \preceq \beta$. In this way, \preceq is a partial order on $\hat{\mathcal{L}}$.

Suppose $\alpha \not\preceq \beta$. In other words, there is some $b \in \beta$ such that $a \not< b$ for all $a \in \alpha$. Note that $a \not< c$ for all $c \leq b$, and so replacing β by $\{c \in \beta \mid c \leq b\}$, we can suppose this holds for all $a \in \alpha$ and all $b \in \beta$. Since Ω is boundless, from an earlier discussion, we can also suppose $a^* \not\preceq b$ for all such a, b . Finally, if $\alpha \cup \beta$ is a partial flow then also $a \not< b^*$. Therefore, in the last case, we can assume (after removing finitely many elements from the sequences) that for all $a \in \alpha$ and all $b \in \beta$ we have either $b < a$ or $b \pitchfork a$.

Suppose $A, B \subseteq \mathcal{L}$ are finite and disjoint and suppose that $\alpha \not\preceq \beta$ for all $[\alpha] \in A$ and $[\beta] \in B$. Let $U(A \sqcup B) = U(A) \cup U(B)$ be as constructed earlier, so that $a^* \not\preceq b$ for all $a, b \in U(A \sqcup B)$. Suppose that $U(A \sqcup B)$ is partial flow. Applying the construction of the previous paragraph for all pairs α, β , we can suppose that if $a \in U(A)$ and $b \in U(B)$, then $b < a$ or $b \pitchfork a$.

Now suppose R is a flow with $\lambda([R]) = A \sqcup B$. Let $R' \subseteq R$ be the set of $d \in R$ such that $d \leq b$ for some $b \in U(B)$. Let $S = (R \setminus R') \cup \{b^* \mid b \in R'\}$. In other words, we “flip” all the elements of R' . Note that $U(A) \subseteq R \cap S$ and $U(B) \subseteq R' = R \setminus S$.

Lemma 9.5.12. *S is a flow with $\lambda([S]) = A$.*

Proof. To check that S is flow, suppose for contradiction that $c, d \in S$ with $c < d^*$. Since R is a flow, we can suppose that $c \notin R$. Thus, $c^* \in R \setminus S = R'$, so there is some $a \in U(B)$ with $c^* \leq a$, or equivalently, $a^* \leq c$. If also $d \notin R$, then similarly, we have $d^* \leq b$ for some $b \in U(B)$. But then $a^* \leq c < d^* \leq b$, so $a^* \leq b$, which we disallowed in our construction of $U(A \sqcup B)$. Thus, $d \in R$. But then, since $d < c^* \leq a \in U(B)$, we have $d \in R'$. Thus $d^* \in S$, contradicting our assumption that $d \in S$.

This shows that S is a flow. It remains to show that $\lambda([S]) = A$.

We have noted that $U(A) \subseteq S$ and $U(B) \cap S = \emptyset$. Therefore $A \subseteq \lambda([S])$ and $B \cap \lambda([S]) = \emptyset$. Since $\lambda([R]) = A \sqcup B$, it remains to show that $\lambda([S]) \subseteq \lambda([R])$.

Suppose, for contradiction, that $\gamma \in \mathcal{L}(S) \setminus \mathcal{L}(R)$. Then $\gamma \subseteq S$ and we can suppose that $\gamma \cap R = \emptyset$. Since Ω is boundless, we can suppose that we have $c^* \not\leq b$ for any $b \in \beta \in B$ and any $c \in \gamma$. But if $c \in \gamma$, then $c \in R \setminus S$, so $c^* \in S \setminus R = R'$, so there is some $b \in U(B)$ with $c^* \leq b$ giving a contradiction. \square

Now suppose $\#\lambda([R]) = n < \infty$. Now $\lambda([R])$ is partially ordered by \preceq . We can write $A = \{[\alpha_1], \dots, [\alpha_n]\}$ such that $\alpha_i \not\leq \alpha_j$ whenever $i < j$. (Inductively, choose $[\alpha_{i+1}]$ to be \preceq -maximal in $\lambda([R]) \setminus \{[\alpha_1], \dots, [\alpha_i]\}$.) Given $i \in \{0, 1, \dots, n\}$, let $A_i = \{[\alpha_j] \mid j \leq i\}$ and $B_i = \{[\alpha_j] \mid j > i\}$. Thus, $\lambda([R]) = A_i \sqcup B_i$. (Note that $A_0 = \emptyset$ and $A_n = A$.) Let S_i be the flow constructed as above, so that $\lambda([S_i]) = A_i$. In particular, we have $\lambda([S_0]) = \emptyset$. We immediately get:

Proof of Proposition 9.5.11. We start from any flow $R \in \mathcal{F}$. By Lemma 9.5.3, $\#\lambda([R]) < \infty$. We now apply the above to give us $S_0 \in \mathcal{F}_1$. (To be precise, we should also make the observation that if $\Omega = \emptyset$, then $\mathcal{F}_1 = \{\emptyset\} \neq \emptyset$.) \square

Now the sequence $(\lambda([S_i]))_i$ is strictly increasing. It is therefore monotone in the median structure on $\mathcal{P}(\hat{\mathcal{L}})$. (If $i \leq j \leq k$, then $\lambda([S_i]) \cap \lambda([S_k]) \subseteq \lambda([S_j])$, so $\lambda([S_i]) \cdot \lambda([S_j]) \cdot \lambda([S_k])$.) Since $\lambda : \hat{\mathcal{F}} \rightarrow \mathcal{P}(\hat{\mathcal{L}})$ is a monomorphism, it follows that $([S_i])_i$ is also monotone in $\hat{\mathcal{F}}$.

This shows (under the standing assumption that Ω is boundless and of subinfinite rank):

Lemma 9.5.13. *If $[R] \in \hat{\mathcal{F}}$, then there is a strictly monotone sequence of length $\#\lambda([R]) + 1$ from \mathcal{F}_1 to $[R]$ in \mathcal{F} .*

Given that \mathcal{F}_1 is non-empty, the following is straightforward:

Proposition 9.5.14. *Suppose Ω is boundless and of subinfinite rank. If $[R] \in \hat{\mathcal{F}}$, then any strictly monotone sequence from \mathcal{F}_1 to $[R]$ in $\hat{\mathcal{F}}$ has length at most $\#\lambda([R]) + 1$.*

Proof. Let $\mathcal{F}_1 = [R_0], [R_1], [R_2], \dots, [R_m] = [R]$ be a strictly monotone sequence in $\hat{\mathcal{F}}$ from \mathcal{F}_1 to $[R]$. We can choose representatives R_i so that R_0, R_1, \dots, R_m is monotone in $\mathcal{F}(\mathcal{H}(\Pi))$. (This can be done by backward induction on i , following the remark at the beginning of this subsection.) Since $\lambda([R_0]) = \emptyset$, we have $\lambda([R_i]) \subseteq \lambda([R_j])$ when $i < j$ (cf. the proof of Lemma 9.5.1). Since $R_i \not\sim R_j$, these inclusions are strict (see Lemma 9.5.2). Therefore $m \leq \#\lambda([R])$, and the statement follows. \square

In other words, this shows that $\lambda([R]) + 1$ is precisely the maximal length of a strictly monotone sequence from \mathcal{F}_1 to $[R]$ in $\hat{\mathcal{F}}$. It follows that $\lambda([R])$ is the distance between \mathcal{F}_1 and $[R]$ in the combinatorial metric on $\hat{\mathcal{F}}$, to be defined in Section 11.

9.6. Formulation in terms of ideals.

To finish this section, we briefly return to the case of a boolean algebra, B , viewed with its structure as a boolean ring, as discussed in Example (Ex3.5) of Subsection 3.4. Recall that convex subsets of B are precisely the translates of ideals. As noted there, we can put a new boolean ring structure on B by setting $x \oplus y = 1 + x + y$ and $x \cdot y = x + y + xy$. This swaps 0 and 1 (and interchanges meet and join).

Let $I \leq B$ be an ideal. Let $1 + I = \{1 + x \mid x \in I\}$ be its translate by 1. Clearly $1 + (1 + I) = I$. Note that $1 + I$ is an ideal of the ring (B, \oplus, \cdot) . This is because the map $[x \mapsto 1 + x]$ is an isomorphism between the two ring structures. (Or, more explicitly, if $x, y \in 1 + I$ and $r \in B$, then $x \oplus y = 1 + (1 + x) + (1 + y) \in 1 + I$, and $r \cdot x = r + x + rx = 1 + (1 + r)(1 + x) \in 1 + I$.) Note also that if $I \cap (1 + I) \neq \emptyset$, then $1 \in I$, so $I = B$.

Lemma 9.6.1. *Let $I \leq B$ be an ideal. The following are equivalent:*

- (1) $B = I \sqcup (1 + I)$,
- (2) I is a maximal ideal,
- (3) I is a prime ideal.

Proof. Suppose $B = I \sqcup (1 + I)$. Then certainly $I \neq B$. Moreover, if $x \in B \setminus I$, then $1 + x \in I$, so $1 = (1 + x) + x \in \langle I, x \rangle$, so $\langle I, x \rangle = B$. Therefore I is maximal.

Conversely, suppose that I is maximal. Since $I \neq B$, we have observed that $I \cap (1 + I) = \emptyset$. Suppose $x \in B \setminus I$. Then $B = \langle I, x \rangle = \{i + rx \mid i \in I, r \in B\}$. Thus $1 = i + rx$ for some $i \in I$ and $r \in B$. Now $x + ix = (1 + i)x = rx^2 = rx = 1 + i$, so $1 + x = ix + i = i(1 + x) \in I$. This shows that $B = I \sqcup (1 + I)$.

In any commutative ring, maximal implies prime. (Or more directly in this case: if $x, y \in B$, with $xy \in I$ and with $y \notin I$, then $1 + y \in I$, so $x = x(1 + y) + xy \in I$.) Conversely, if I is prime, and $x \in B$, then $x(1 + x) = 0 \in I$, so either $x \in I$ or $1 + x \in I$, and we see that I is maximal by (1). \square

Recall that I and $1 + I$ are both convex in the median structure on B . Therefore, I is a halfspace of B if and only if it is a maximal ideal. We see that there is a natural bijection between walls of B and prime ideals. We note that the complement of a prime ideal is also the same thing as an ultrafilter in the lattice structure on B . It also follows that any pair of distinct walls of B must cross: no maximal ideal can be strictly contained in another.

As an example, let Ξ be a non-empty compact totally disconnected hausdorff space (often called a **Stone space** in this context). Let $\mathcal{B} = \mathcal{B}(\Xi) \subseteq \mathcal{P}(\Xi)$ be the set of clopen subsets of Ξ . This is a subalgebra of $\mathcal{P}(\Xi)$, hence intrinsically a boolean algebra. Given $a \in \Xi$, let $P(a) = \{x \in \mathcal{B} \mid a \in x\}$ be the principal ultrafilter at a . Let $I(a) = 1 + P(a)$ and $\eta(a) = \{I(a), P(a)\}$. As noted above, $\eta(a)$ is a wall of \mathcal{B} , and we get an injective map $\eta : \Xi \rightarrow \mathcal{W}(\mathcal{B}(\Xi))$. In fact, this is surjective. For suppose I is a prime ideal of \mathcal{B} . Then $P := 1 + I$ is a family of non-empty closed subsets of Ξ , closed under finite intersection. By compactness,

there is some $a \in \bigcap P \subseteq \Xi$. Then $P \subseteq P(a)$, so $I(a) \subseteq I$, and so in fact, $I = I(a)$. In summary, this shows that there is a natural bijection from Ξ to the set of walls of $\mathcal{B}(\Xi)$. (Cf. the discussion of cubes in Subsection 11.11.)

We remark that the Stone Duality Theorem tells us that every boolean algebra canonically arises in this way from a Stone space. Given a boolean algebra \mathcal{B} , we can recover Ξ as the set of prime ideals equipped with the Zariski topology. By definition, this has base of closed sets given by the family $\{\mathcal{I}(x) \mid x \in \mathcal{B}\}$, where $\mathcal{I}(x)$ is the set of prime ideals containing x .

A related duality theorem will be discussed in Subsection 12.5.

10. HYPERCUBES

We describe some basic notions relating to hypercubes (both finite and infinite: in the finite case, we will generally use the term ‘‘cube’’). We introduce a general construction which can be used to describe realisations and subdivisions of hypercubes. These will be applied to discrete median algebras in Section 11. We also give some results about constructing cubes in a given median algebra.

10.1. Hypercubes and their faces.

We originally defined a ‘‘hypercube’’ to be a median algebra isomorphic to $\{0, 1\}^X$ for some set X , which we can identify with its power set, $\mathcal{P}(X)$. However, this is in some ways unnatural, in that it entails choosing a preferred element (where all the coordinates are 0) to be identified with $\emptyset \in \mathcal{P}(X)$.

So let us redefine a ‘‘hypercube’’ to be a median algebra, Ψ , isomorphic to a direct product $\Psi \cong \prod_{i \in \mathcal{I}} \delta_i$, where \mathcal{I} is an indexing set, and each δ_i is a 2-point median algebra (without a preferred element). This is clearly equivalent (at least given the Axiom of Choice) to the original formulation.

Let $\pi_i : \Psi \rightarrow \delta_i$ be the projection to the i th coordinate. This is an epimorphism and so determines a wall, $W_i \in \mathcal{W}(\Psi)$, by taking the preimages of the two points of δ_i . Note that $W_i \pitchfork W_j$ whenever $i \neq j$. More generally, if $\mathcal{J} \subseteq \mathcal{I}$, we have a projection, $\pi_{\mathcal{J}} : \Psi \rightarrow \Psi(\mathcal{J})$, where $\Psi(\mathcal{J})$ is the hypercube $\prod_{i \in \mathcal{J}} \delta_i$. Note that we can write $\Psi = \Psi(\mathcal{J}) \times \Psi(\mathcal{I} \setminus \mathcal{J})$.

Given any subset $\mathcal{J} \subseteq \mathcal{I}$, we have an involution, $\theta_{\mathcal{J}} : \Psi \rightarrow \Psi$, where $\theta_{\mathcal{J}}(a)$ swaps the i th coordinate of a whenever $i \in \mathcal{J}$, and leaves all the other coordinates alone. If $a, b \in \Psi$, then $b = \theta_{\mathcal{J}}a$, where \mathcal{J} is the set of coordinates for which a and b differ. In particular, we see that Ψ is homogeneous.

We refer to $\theta_{\mathcal{I}}$ as the ‘‘antipodal map’’. We say that $a, b \in \Psi$ are *antipodal* points if $b = \theta_{\mathcal{I}}a$. It is easily seen that this is equivalent to saying that $\Psi = [a, b]$.

Definition. A *face* of Ψ is a non-empty finite convex subset.

Let $R \subseteq \Psi$ be a face, and write $\mathcal{I}(R) = \{i \in \mathcal{I} \mid R \pitchfork W_i\}$. Referring back to the discussion in Section 8, we see that the walls, $\{W_i \mid i \in \mathcal{I}(R)\}$ restricted to R , pairwise cross (by the Helly Property, Lemma 7.1.1). Moreover, $\Psi \cap O(\epsilon) \neq \emptyset$ for each orthant, $O(\epsilon)$, for $\epsilon \in \{+, -\}^{\mathcal{I}(R)}$. Now any two points of R are separated by

at least one of the walls, W_i (since this is true in Ψ). It follows (as in Subsection 8.1) that $\#(R \cap O(\epsilon)) = 1$ for all ϵ , and that R is isomorphic to $\Psi(\mathcal{I}(R))$. In fact, it has the form $\Psi(\mathcal{I}(R)) \times \{a\}$ for some $a \in \Psi(\mathcal{I} \setminus \mathcal{I}(R))$, under the identification $\Psi = \Psi(\mathcal{I}(R)) \times \Psi(\mathcal{I} \setminus \mathcal{I}(R))$. We refer to R as an n -**face** where $n = \#\mathcal{I}(R)$. We write $\mathcal{C}_n(\Psi)$ for the set of n -faces, and $\mathcal{C}(\Psi) = \bigcup_{n=0}^{\infty} \mathcal{C}_n(\Psi)$ for the set of all faces.

We can view $\mathcal{C}(\Psi)$ as a subset of the set, $\mathcal{K}(\Psi)$, of all non-empty convex subsets of Ψ , as defined in Subsection 7.1. Recall that $\mathcal{K}(\Psi)$ is a median algebra, with the median defined by $RST = \{abc \mid a \in R, b \in S, c \in T\}$, for $R, S, T \in \mathcal{K}(\Psi)$. If $R, S, T \in \mathcal{C}(\Psi)$, then $RST \in \mathcal{C}(\Psi)$. (If R, S, T are all finite, then so is RST .) In other words, $\mathcal{C}(\Psi)$ is a subalgebra of $\mathcal{K}(\Psi)$, and so in particular, intrinsically a median algebra. We can in turn view Ψ as a subalgebra of $\mathcal{C}(\Psi)$ via the embedding $[x \mapsto \{x\}]$.

10.2. Construction of subdivisions and realisations.

There is a very general construction for embedding Ψ into a product of median intervals.

Suppose that for each $i \in \mathcal{I}$, we have associated a median algebra, Υ_i , with δ_i identified as a subset of Υ_i , and such that $\Upsilon_i = \text{hull}(\delta_i)$. In other words, Υ_i is a median interval with endpoints δ_i . (This is essentially the same thing as a bounded distributive lattice, with δ_i consisting of its maximum and minimum, though we don't in general have a preferred order on δ_i .) We write $\hat{\Upsilon}(\Psi) = \prod_{i \in \mathcal{I}} \Upsilon_i$ for the product median algebra. Thus Ψ is a subalgebra of $\hat{\Upsilon}(\Psi)$. We note that $\hat{\Upsilon}(\Psi)$ is the convex hull of Ψ in $\hat{\Upsilon}(\Psi)$. In fact, it is easily seen to be the median interval between any two antipodal points of Ψ .

(In most cases of interest to us, Υ_i will be a linearly ordered set. For the following discussion, it might be helpful to imagine it as a compact real interval, which is a particular case we will consider later.)

Given $R \in \mathcal{C}(\Psi)$, let $\Upsilon(R) \subseteq \hat{\Upsilon}(\Psi)$ be its convex hull. This can be described explicitly as follows. For each $i \in \mathcal{I}$, write $\Upsilon(\pi_i R)$ for the convex hull of $\pi_i R$ in Υ_i . Thus, if $i \in \mathcal{I}(R)$, $\Upsilon(\pi_i R) = \Upsilon_i$, whereas if $i \notin \mathcal{I}(R)$, then $\Upsilon(\pi_i R)$ is one of the endpoints of Υ_i . By Lemma 7.5.1, we see that $\Upsilon(R)$ is naturally isomorphic to $\prod_{i \in \mathcal{I}} \Upsilon(\pi_i R)$ (see also Lemma 10.3.5 below). By Lemma 7.4.2, we have $\Upsilon(RST) = \Upsilon(R)\Upsilon(S)\Upsilon(T)$ for all $R, S, T \in \mathcal{C}(\Psi)$.

Given any subset, $\mathcal{D} \subseteq \mathcal{C}(\Psi)$, let $\Upsilon(\mathcal{D}) = \bigcup_{Q \in \mathcal{D}} \Upsilon(Q)$. (Note that there is no loss in assuming \mathcal{D} to be closed under inclusion.)

We note:

Lemma 10.2.1. *If \mathcal{D} is a subalgebra of $\mathcal{C}(\Psi)$, then $\Upsilon(\mathcal{D})$ is a subalgebra of $\hat{\Upsilon}(\Psi)$. Moreover, $\mathcal{D}' \subseteq \mathcal{D}$ is convex in \mathcal{D} , then $\Upsilon(\mathcal{D}')$ is convex in $\Upsilon(\mathcal{D})$.*

Proof. Let $x, y, z \in \Upsilon(\mathcal{D})$. The $x \in \Upsilon(R)$, $y \in \Upsilon(S)$, $z \in \Upsilon(T)$, for some $R, S, T \in \mathcal{C}(\mathcal{D})$. Then $xyz \in \Upsilon(R)\Upsilon(S)\Upsilon(T) = \Upsilon(RST) \subseteq \Upsilon(\mathcal{D})$.

For the second statement, Let $x, y \in \Upsilon(\mathcal{D}')$ and $z \in \Upsilon(\mathcal{D})$. Then $x \in \Upsilon(R)$, $y \in \Upsilon(S)$, $z \in \Upsilon(T)$, where $R, S \in \mathcal{D}'$ and $T \in \mathcal{D}$. We have $RST \in \mathcal{D}'$. Therefore, $xyz \in \Upsilon(R)\Upsilon(S)\Upsilon(T) = \Upsilon(RST) \subseteq \Upsilon(\mathcal{D}')$. \square

Suppose $\Pi \subseteq \Psi$ is any subalgebra of Ψ . Let $\mathcal{C}(\Pi, \Psi) = \{R \in \mathcal{C}(\Psi) \mid R \subseteq \Pi\}$. This is subalgebra of $\mathcal{C}(\Psi)$ in its median structure. (If $R, S, T \subseteq \Pi$, then $RST \subseteq \Pi$.) We write $\Upsilon(\Pi, \Psi) = \Upsilon(\mathcal{C}(\Pi, \Psi))$. If $\Pi' \subseteq \Pi$ is convex in Π , then $\mathcal{C}(\Pi', \Psi)$ is convex in $\mathcal{C}(\Pi, \Psi)$. (In fact, it's not hard to see that any convex subset of $\mathcal{C}(\Pi, \Psi)$ which is closed under inclusion has the form $\mathcal{C}(\Pi', \Psi)$, where $\Pi' \subseteq \Pi$ is convex: namely the set of 0-cells.)

As an immediate consequence of Lemma 10.2.1, we get:

Lemma 10.2.2. *If $\Pi \leq \Psi$ is a subalgebra of Ψ , then $\Upsilon(\Pi, \Psi)$ is a subalgebra of $\hat{\Upsilon}(\Psi)$. Moreover, if Π' is convex in Π , then $\Upsilon(\Pi', \Psi)$ is convex in $\Upsilon(\Pi, \Psi)$.*

We will generally abbreviate: $\Upsilon(\Psi) := \Upsilon(\Psi, \Psi)$. If Ψ is finite, then $\Upsilon(\Psi) = \hat{\Upsilon}(\Psi)$.

A particular example of the above construction is “binary subdivision”. Let $\Psi = \prod_{i \in \mathcal{I}} \delta_i$ be a hypercube as above. For each $i \in \mathcal{I}$, we adjoin a “midpoint”, m_i , to δ to give a 3-point median algebra, $\Sigma_i = \delta_i \cup \{m_i\}$. Here Σ_i is isomorphic to the totally ordered set, $\{-1, 0, 1\}$, with m_i corresponding to 0. We write $\Sigma(\Psi) \leq \hat{\Sigma}(\Psi)$ etc. for the above constructions (with “ Σ ” replacing “ Υ ”). We refer to $\Sigma(\Psi)$ as the **binary subdivision** of Ψ .

Note that there is an isomorphism, $\phi : \Sigma(\Psi) \rightarrow \mathcal{C}(\Psi)$, where $\phi(a)$ is the smallest face of Ψ containing a . The inverse map sends a face of Ψ to its centre. If $\Pi \subseteq \Psi$ is a subalgebra, this restricts to an isomorphism of $\Sigma(\Pi, \Psi)$ to $\mathcal{C}(\Pi, \Psi)$.

There are some other natural choices for Υ_i . For example, we could more generally take it to be a linear $(p+1)$ -point median algebra with endpoints δ_i , in place of each Υ_i we would obtain the the “ **p -ary subdivision**” of Ψ . Indeed we could use any family of finite linear sets to give us some “subdivision” of Ψ . We will discuss subdivisions further in Subsection 11.10.

Of particular interest later, is the case where take each Υ_i to be a compact real interval, Δ_i , with endpoints, δ_i . In this case, we will write $\hat{\Delta}(\Psi) = \prod_i \Delta_i$ (for $\hat{\Upsilon}(\Psi)$). This gives us a **real hypercube**, $\hat{\Delta}(\Psi)$, with **corners** $\Psi \subseteq \hat{\Delta}$. Again Ψ is a subalgebra, with $\hat{\Delta}(\Psi)$ the convex hull of Ψ . We refer to $\hat{\Delta}(\Psi)$ as the **full realisation** of Ψ (viewed simply as a median algebra).

Given any subalgebra, $\Pi \subseteq \Psi$, we set $\Delta(\Pi, \Psi) = \bigcup_{R \in \mathcal{C}(\Pi, \Psi)} \Delta(R) \subseteq \hat{\Delta}(\Psi)$. Thus Lemma 10.2.2 tells us that $\Delta(\Pi, \Psi)$ is a subalgebra of $\hat{\Delta}(\Psi)$. In particular, we have $\Delta(\Psi) := \Delta(\Psi, \Psi) \leq \hat{\Delta}(\Psi)$. If Ψ is finite, then $\Delta(\Psi) = \hat{\Delta}(\Psi)$.

We will return to this construction in Subsection 11.2. For the remainder of this section, we restrict attention to finite hypercubes, or as we will call them, finite “cubes”.

10.3. Finite cubes.

Let Q be an n -cube. We write $I = \{1, \dots, n\}$. We have pairwise crossing walls $\{W_1, \dots, W_n\}$. As in the above discussion of faces, we see that in fact, $\mathcal{W}(Q) = \{W_1, \dots, W_n\}$ in this case. (This also follows from Lemma 8.1.3, as discussed there.)

Conversely, if Q is a finite median algebra whose walls pairwise cross, then the natural embedding of Q into the cube $\prod \mathcal{W}(Q)$ is an isomorphism. We conclude:

Lemma 10.3.1. *A finite median algebra is a cube if and only if all its walls cross.*

(For a more general statement which includes the infinite case, see Lemma 11.11.5.)

We have observed that hypercubes are homogeneous. Indeed the converse is true in the finite case:

Proposition 10.3.2. *Any finite homogeneous median algebra is a cube.*

This follows from a more general statement, namely Proposition 11.1.2, and we postpone the proof for the moment.

We also note:

Lemma 10.3.3. *Suppose that $\pi : Q \rightarrow P$ is an epimorphism to a median algebra P . Then P is a cube. In fact, there is some $J \subseteq I$ such that $\pi = \theta \circ \pi_J$, where $\pi_J : Q \rightarrow Q(J)$ is the projection to the quotient cube, $Q(J)$, and where θ is an isomorphism from $Q(J)$ to P .*

Proof. There is a natural injective map, $\mathcal{W}(P) \rightarrow \mathcal{W}(Q)$. Let $J \subseteq I$ be the set of $i \in I$ such that W_i lies in its image. Then π factors through π_J . If $\pi(x) = \pi(y)$, then any wall that separates x and y must lie in $I \setminus J$, and so $\pi_J(x) = \pi_J(y)$. Therefore the map from $Q(J)$ to P is injective, hence an isomorphism. \square

The following notation will be useful in dealing with finite cubes.

We choose some pair, a, b , of antipodal points of Q , so that $Q = [a, b]$. In this way, Q has the structure of a bounded distributive lattice, with minimum a , and maximum b . Let $e_1, \dots, e_n \in Q$ be the elements adjacent to a . Given $J \subseteq I$, write $e_J = \bigvee_{i \in J} e_i$, with the convention that $e_\emptyset = a$. Thus $e_{\{i\}} = e_i$ and $e_I = b$. We will abbreviate $e_{ij} := e_{\{i,j\}}$. Note that $e_{J \cap K} = e_J \wedge e_K$ and $e_{J \cup K} = e_J \vee e_K$. In other words, $[J \mapsto e_J] : \mathcal{P}(I) \rightarrow Q$ is an isomorphism.

Let $S = \{e_J \in Q \mid \#J \leq 2\}$.

Lemma 10.3.4. $Q = \langle S \rangle$.

Proof. For $n \leq 2$, there is nothing to prove. We prove the statement inductively for $n \geq 3$. By the inductive hypothesis, we know that every $(n - 1)$ -face of Q containing a lies in $\langle S \rangle$. The union of these faces is $Q \setminus \{e_I\}$. But $e_I = e_{I_1} e_{I_2} e_{I_3}$, where $I_i = I \setminus \{i\}$, so $e_I \in \langle S \rangle$ also. \square

Let M be a median algebra. Suppose $Q \subseteq M$ is an n -cube. Now the 1-cells of Q crossing a given wall, W_i , of Q , are all parallel in Q , hence also in M . (See the

definitions and discussion in Subsection 7.2.) Given any such edge, we can take the interval in M between them. These intervals will also all be parallel in M and the translations between them will all be median isomorphisms. In particular, they will all be isomorphic to $[a, e_i]$ in the above notation. We write D_i for this isomorphism class. Note that the gate map from M to D_i is well defined independently of which representative of the parallel class we choose. Let $D = \prod_{i=1}^n D_i$.

Lemma 10.3.5. *There is a natural isomorphism from D to $\text{hull}_M(Q)$.*

Proof. In the above notation, $\text{hull}_M(Q) = [a, b]_M$. Identifying D_i with $[a, e_i]_M$, we define a map, $\phi : D \rightarrow [a, b]_M$ by setting $\phi(x_1, \dots, x_n) = x_1 \vee \dots \vee x_n$. It has inverse given by $\phi^{-1}(y) = (y \wedge e_1, \dots, y \wedge e_n)$. It is easily checked that these are isomorphisms of distributive lattices, hence also median isomorphisms.

This construction involved choosing some $a \in Q$. However, we get the same maps regardless of this choice: note that the coordinates of $\phi^{-1}(y)$ are given by gate maps to the factors D_i . \square

In particular, we note that if $\text{rank}(M) = \nu$, then the convex hull of any ν -cube is intrinsically a direct product of totally ordered intervals (since each of the factors has rank 1).

We now move on to constructing homomorphisms from cubes to a general median algebra, M . We begin with a uniqueness statement.

Lemma 10.3.6. *A homomorphism $\phi : Q \rightarrow M$ is completely determined by the values of $\phi(a), \phi(e_1), \dots, \phi(e_n)$.*

Proof. In view of Lemma 10.3.4, it is enough to check this for $n = 2$.

Write $c = \phi(e_1)$, $d = \phi(e_2)$, and suppose that e, f are two possible values of $\phi(e_{12})$. Then $ae f = ae(cdf) = (aec)(aed)f = cdf = f$. Similarly, $a f e = e$, and so $e = f$. \square

We now consider existence. We assume $n \geq 3$. Given distinct $i, j \in I$, write Q_{ij} for the 2-face, $\{a, e_i, e_j, e_{ij}\}$. Recall that we defined $S = \{e_J \in Q \mid \#J \leq 2\} = \bigcup_{i,j} Q_{ij}$.

Lemma 10.3.7. *Let M be a median algebra, and let $\phi : S \rightarrow M$ be a map such that $\phi|_{Q_{ij}}$ is a homomorphism for all distinct $i, j \in I$. Then there is a unique extension of ϕ to a homomorphism, $\hat{\phi} : Q \rightarrow M$, such that $\hat{\phi}|_S = \phi$. Moreover, if $\phi(e_i) \neq \phi(a)$ for all i , then $\hat{\phi}$ is injective.*

Proof. We first consider the case where $M = \{0, 1\}$. We can assume that $\phi(a) = 0$. If $\phi(e_i) = 0$ for all i , then $\phi(e_{ij}) = 0$ for all i, j , and we set $\hat{\phi} \equiv 0$. So we assume there is some i with $\phi(e_i) = 1$. Then $\phi(e_j) = 0$ for all $j \neq i$ (since $e_i.a.e_j$). We see that $\phi(e_{jk}) = 1$ if and only if either $j = i$ or $k = i$. We now set $\hat{\phi}(e_J) = 1$ if $i \in J$ and $\hat{\phi}(e_J) = 0$ if $i \notin J$. This is clearly a homomorphism.

Now let M be any median algebra. By Proposition 3.2.13, M embeds into a hypercube, say Ψ . This gives us a map $\phi : S \rightarrow \Psi$, with $\phi|_{Q_{ij}}$ a homomorphism

for all i, j . On projecting to any coordinate of Ψ , we get a map $S \rightarrow \{0, 1\}$ as in the previous paragraph. We may therefore extend on each coordinate independently, and then combine them to give us a homomorphism $\hat{\phi} : Q \rightarrow \Psi$, extending ϕ . By Lemma 10.3.4, $Q = \langle S \rangle$, and so $\hat{\phi}(Q) \subseteq \langle \hat{\phi}S \rangle \subseteq M$. Thus $\phi : Q \rightarrow M$ is the required extension.

The uniqueness of the extension follows by Lemma 10.3.6.

The final statement regarding injectivity follows immediately from Lemma 10.3.3. □

As we have described it, the proof uses the Axiom of Choice via Proposition 3.2.13. However, this can be avoided by replacing M by the subalgebra generated by $\phi(S)$. This is finite by Proposition 3.3.3, a constructive proof of which was given in Subsection 6.3. The finite case of Proposition 3.2.13 can be proven constructively (see Lemma 3.2.12).

Here is another observation:

Lemma 10.3.8. *A map $\phi : Q \rightarrow M$ from a finite cube, Q , to a median algebra, M , is a homomorphism if and only if its restriction to each 2-face of Q is a homomorphism.*

Proof. We will abbreviate $x' := \phi(x)$ for $x \in Q$. Given $x, y \in Q$ for the number of walls separating x, y , i.e. the rank of the face $[x, y]$. Note that ρ is a metric on Q . (See Subsection 11.1 for a more general discussion.) It is enough to show that $a'.b'.c'$ holds in M whenever $a.b.c$ holds in Q . By hypotheses this hold when $\rho(a, c) \leq 2$, and we prove it by induction on $\rho(a, c)$ when $\rho(a, c) \geq 3$. We can suppose that $a \neq b$ and that $\rho(b, c) \geq 2$. Thus we can find $d, e \in [b, c]$ so that b, d, c, e is a 2-cube. We have $a.b.d.c, a.b.e.c$ and $d.c.e$. Thus, $\rho(a, d), \rho(a, e) < \rho(a, c)$ and $\rho(d, e) = \rho(b, c) < \rho(a, c)$. Thus the inductive assumption tells us that $a'.b'.d', a'.b'.e'$ and $d'.c'.e'$ hold. Therefore $a'b'c' = a'b'(d'e'c') = (a'b'd')(a'b'e')c' = b'b'c' = b'$ and so $a'.b'.c'$ as required. □

We note that all we have really used regarding Q is that it is a finite (or discrete) median algebra such that any two distinct non-adjacent points, $x, y \in Q$, are the antipodal vertices of 2-cube of Q . In fact, a slight elaboration on the argument shows that we can restrict this hypothesis to the case where x, y are both adjacent to a third point.

We remark that there is an analogous statement for quasimedial graphs, namely Lemma 23.3.4. This intersects with the current discussion in the context of median graphs. Lemma 23.4.4 is an analogue of Lemma 10.3.7, and suggests another more constructive proof thereof.

11. DISCRETE MEDIAN ALGEBRAS

Discrete median algebras are essentially the same structures as (combinatorially) CAT(0) cube complex, which we will return to in a more geometrical setting in

Section 18. There can also be reinterpreted as the vertex sets of median graphs, which we discuss in Section 13.

Here we develop the basic theory. We begin by describing the combinatorial metric on a discrete median algebra, Π . We associate to Π a “cell complex”, $\mathcal{C}(\Pi)$, which itself has the structure of a discrete median algebra (Lemma 11.2.2). This can be realised as a CW complex, $\Delta(\Pi)$. (This will be discussed further in Section 17.) We give a criterion for convexity in Π (Lemma 11.4.4) and observe that all non-empty convex subsets are gated (Lemma 11.3.3). We also complete the proof of Proposition 8.2.3 regarding convex hulls (see Proposition 11.5.4). We describe the notion of a “cube path” in Π , and relate this to the notion of a helly graph, which will be mentioned again in Subsection 25.2. Some properties of subalgebras are described in Subsections 11.7 and 11.8. We give some further discussion of the structure of a free median algebra (Subsection 11.9). We give a discussion of “subdivisions” of a discrete median algebra (see Lemma 11.10.1), and show that any finite median algebra is canonically a subdivision of a certain minimal subalgebra (Proposition 11.10.3). We give a construction, and various characterisations, of cubes (Lemmas 11.11.4 and 11.11.5). We go on to define the “Roller boundary” of Π and give a number of descriptions thereof. We also relate this to infinite cube paths in Π (Proposition 11.12.8). This can be viewed as a canonical combing of Π from a given basepoint. We finish the section with a discussion of “event structures”, which give another description of discrete median algebras relative to a given basepoint.

11.1. Some basic definitions.

We have already mentioned discrete median algebras in Subsection 9.4. Here is the definition again:

Definition. A median algebra, Π , is *discrete* if $[a, b]$ is finite for all $a, b \in \Pi$.

Note that by Corollary 3.2.14, this is equivalent to saying that $\mathcal{W}(a, b)$ is finite.

Given $a, b \in \Pi$, we write $\rho(a, b) = \rho_{\Pi}(a, b) = \#\mathcal{W}(a, b)$. Recall that for any $c \in \Pi$, we have $\mathcal{W}(a, b) \subseteq \mathcal{W}(a, c) \cup \mathcal{W}(c, b)$. It follows that ρ is a metric on Π . Moreover, $a.c.b$ holds if and only if $\mathcal{W}(a, c) \cap \mathcal{W}(c, b) = \emptyset$. Therefore, $[a, b] = \{c \in \Pi \mid \rho(a, b) = \rho(a, c) + \rho(c, b)\}$. (This means that ρ is a “median metric”, which we will define in Section 13.) We also note that $\rho(a, b)$ is the length, n , of any maximal chain, $a = a_0 < a_1 < \dots < a_n = b$ in $[a, b]$. We refer to such a maximal chain as a *geodesic path* in Π . (Various equivalent definitions will be given later.)

More generally we define a *1-path*, \underline{a} , of *length* n , in Π to be a sequence, a_0, a_1, \dots, a_n , in Π with a_{i+1} adjacent to a_i for all i . (It is convenient disallow $a_{i+1} = a_i$ in what we do, though that is not an essential point.) It is easily checked that a geodesic between a and b is precisely a 1-path of length $\rho(a, b)$. If we take $W_i \in \mathcal{W}(\Pi)$ to be the wall with $\mathcal{W}(a_{i-1}, a_i) = \{W_i\}$, then $\mathcal{W}(a, b) = \{W_1, \dots, W_n\}$.

Put another way, we see that ρ is the combinatorial metric on the adjacency graph $\Gamma(\Pi)$ defined in Subsection 5.1. Recall that this has vertex set Π , and edges

determined by adjacency in Π . Such a graph is a “median graph” as we will discuss in Section 16.

Note that if $a, b, c, d \in \Pi$ with c, d adjacent, then abc and abd are adjacent or equal. (For if $e \in \Pi$, then without loss of generality, we have $cde = c$, so $(abc)(abd)e = ab(cde) = abc$.)

Definition. A subset $A \subseteq \Pi$ is **1-path-connected** if any two points of A are connected by a 1-path in A .

Lemma 11.1.1. *If $A \subseteq \Pi$ is a 1-path-connected subalgebra, then $\Gamma(A)$ is an isometrically embedded subgraph of $\Gamma(\Pi)$.*

(In fact, $\Gamma(A)$ is a retract of $\Gamma(\Pi)$ — see Proposition 11.8.2.)

Proof. Given $a, b \in A$, we can connect a, b by a geodesic path, $a = a_0, a_1, \dots, a_n = b$ in Π . The projection $(aba_i)_i$ lies in A , and so $\rho_\Pi(a, b) = n = \rho_A(a, b)$. \square

Recall by Proposition 3.2.13 that any finite median algebra, M , embeds in a cube $Q := \{0, 1\}^{\mathcal{W}(M)}$. Moreover, it is immediate from the construction that any two adjacent points in M are also adjacent in Q . Thus, M is 1-path-connected in Q , and so $\Gamma(M)$ is isometrically embedded in the cubical graph, $\Gamma(Q)$.

From this we can get the following result of [Mul2]:

Proposition 11.1.2. *Let M be a finite median algebra such that $\Pi(M)$ is regular. Then M is a cube.*

(Recall that “regular” means that all vertices have the same valence.)

In particular, this implies Lemma 10.3.2, namely that any homogeneous finite median algebra is a cube.

Proof. Let M be n -regular (i.e. every vertex has valence n). We embed M in a cube, Q , so that $\Gamma(M)$ is isometrically embedded in $\Gamma(Q)$. We claim that M is face of Q .

To see this, let $a, b \in M$ with $\rho(a, b)$ maximal. Let C be the set of vertices of M adjacent to a . Let $F \subseteq Q$ be the face of Q containing a , and with adjacent vertices C . This is an n -cube. If $c \in M$ is adjacent to a , then $a.c.b$ (otherwise, we have $a.b.c$, giving the contradiction that $\rho(a, c) > \rho(a, b)$). Moreover, if c, d are distinct and adjacent to a , then $bcd \in M$, and a, c, bcd, d forms a 2-cube. By Lemma 10.3.4, F is generated by the union of these 2-cubes, so $F \subseteq M$. (One could also deduce this by a more direct induction argument.) Now $\Gamma(F)$ is embedded in $\Gamma(M)$, and both graphs are n -regular, and so $F = M$, as claimed. \square

We also note:

Lemma 11.1.3. *The convex hull of a finite subset of a discrete median algebra is finite.*

Proof. Let $A \subseteq \Pi$ be finite, and let $H = \text{hull}(A)$. We can identify $\mathcal{W}(H)$ with the set of walls of Π which cross H . This is equal to $\bigcup_{a, b \in A} \mathcal{W}(a, b)$, which is finite. Therefore, by Corollary 3.2.14, H is finite. \square

(Lemma 11.1.3 also follows from the fact that one can obtain the convex hull in this case by iterating the join a finite number of times: see Lemma 3.2.9 and the discussion in Subsection 7.4.)

11.2. Relation to cube complexes.

We describe how a discrete median algebra can be viewed as the vertex set of a cube complex.

Definition. An n -*cell* of Π is a convex n -cube in Π .

We write $\mathcal{C}_n(\Pi)$ for the set of n -cells of Π . (This is consistent with the notation of Section 10 in the case of a finite cube.) We write $\mathcal{C}(\Pi) = \bigcup_{n=0}^{\infty} \mathcal{C}_n(\Pi)$ for the set of all cells of Π . In this way, we can view Π as an abstract “cube complex”. (More precise definitions will be given in Section 16.)

Let $\Psi(\Pi) = \prod \mathcal{W}(\Pi) \equiv \{0, 1\}^{\mathcal{W}(\Pi)}$. We can identify $\mathcal{W}(\Pi)$ as a subset of $\mathcal{W}(\Psi(\Pi))$. We have a natural embedding of Π as a subalgebra of $\Psi(\Pi)$. The embedding has the property that each cell of Π is a face of $\Psi(\Pi)$. (Note that if $Q \in \mathcal{C}_n(\Pi)$, then Q is the intersection of the halfspaces of Π containing Q . These halfspaces account for all but n walls of Π , and so their intersection in $\Psi(\Pi)$ will be an n -cube containing Q , hence precisely, Q .) In the notation of Section 10, this says that, under this identification, we have $\mathcal{C}(\Pi) = \mathcal{C}(\Pi, \Psi(\Pi))$.

Recall from Subsection 7.1 that the set, $\mathcal{K}(\Pi)$, of all non-empty convex subsets of Π , has a natural structure as a median algebra. By definition, $\mathcal{C}(\Pi) \subseteq \mathcal{K}(\Pi)$. In fact:

Lemma 11.2.1. $\mathcal{C}(\Pi)$ is a subalgebra of $\mathcal{K}(\Pi)$.

Proof. Let $R, S, T \in \mathcal{C}(\Pi)$ and let $Q = RST$. We need to check that Q is a cube. By Lemmas 8.1.4 and 10.3.1, it is enough to check that if $W_1, W_2 \in \mathcal{W}(\Pi)$ are distinct with $Q \pitchfork W_1$ and $Q \pitchfork W_2$ then $W_1 \pitchfork W_2$.

Suppose not. Then without loss of generality, we have $W_1^- \subseteq W_2^-$ (and so $W_2^+ \subseteq W_1^+$). Since R is a cube, we cannot have both $R \pitchfork W_1$ and $R \pitchfork W_2$ (otherwise, this would give $W_1 \pitchfork W_2$). It follows that either $R \subseteq W_2^-$ or $R \subseteq W_1^+$ (or both). The same holds for S and T . Therefore, without loss of generality, we have $R, S \subseteq W_1^+$, and so $Q \subseteq W_1^+$, contradicting $Q \pitchfork W_1$. \square

If $R, S \in \mathcal{C}(\Pi)$, then $\bigcup [R, S]_{\mathcal{C}(\Pi)} \subseteq J(R, S)$. (Recall that $J(R, S) = \bigcup_{x \in R, y \in S} [x, y]$ is the join of R and S .) Since R, S are finite, so is $J(R, S)$. We deduce:

Lemma 11.2.2. $\mathcal{C}(\Pi)$ is a discrete median algebra.

We also note:

Lemma 11.2.3. If $R, S \in \mathcal{C}(\Pi)$ are adjacent in $\mathcal{C}(\Pi)$, then either R is a corank-1 face of S , or S is a corank-1 face of R .

Proof. Let $R, S \in \mathcal{C}(\Pi)$ be adjacent: that is, $[R, S] = \{R, S\}$.

First note that if $R \subseteq S$, then R is a corank-1 face of S . Otherwise, let Q be any face of S such that $R \subset Q \subset S$ are strict inclusions. Then $Q \in [R, S]$ giving a contradiction. We can therefore suppose that neither $R \subseteq S$ nor $S \subseteq R$ holds.

Let $R_S = \omega_R S$ and $S_R = \omega_S R$, where ω_R and ω_S are the respective gate maps as given by Lemma 7.3.7 (or see Lemma 11.6.2 below). Now R_S, S_R are convex, by Lemma 7.3.2 and so $R_S, S_R \in \mathcal{C}(\Pi)$. In fact, by Lemma 7.3.9, we have $R_S, S_R \in [R, S]$. Thus, $R_S, S_R \in \{R, S\}$. Now $R_S \subseteq R$, and so we must have $R_S = R$. Similarly, $S_R = S$. By Lemma 7.3.4, R and S are parallel. Let $\mathcal{R} \subseteq \mathcal{K}(\Pi)$ be the parallel class. Now $\mathcal{R} \subseteq \mathcal{C}(\Pi)$. By Lemma 7.2.6, $\bigcup \mathcal{R}$ is convex in Π . By Lemma 7.2.7, $\bigcup \mathcal{R}$ is isomorphic to $R \times \mathcal{R}$, where \mathcal{R} is given by the median structure induced from \mathcal{C} . Since R, S are adjacent in \mathcal{R} , $R \cup S$ is isomorphic to $R \times \{R, S\}$, hence a cube. Moreover it is convex in $\bigcup \mathcal{R}$, hence also in Π . Therefore, $R \cup S \in \mathcal{C}(\Pi)$. But now $R \cup S \in [R, S]$ giving a contradiction.

In summary, we conclude that either $R \subseteq S$ or $S \subseteq R$, and that these are corank-1 faces. \square

We can “realise” $\mathcal{C}(\Pi)$ as a genuine cube complex $\Delta(\Pi)$, built out of real cubes. We give a brief description of this here, which will serve in this section mainly to motivate a number of constructions. It is not logically essential to the development. We will return to it in more detail later (see Section 17).

Let $\Psi(\Pi) = \{0, 1\}^{\mathcal{W}(\Pi)}$. We can identify Π as a subalgebra of $\Psi(\Pi)$ (via the duality described in Subsection 9.2). By definition, we can identify $\hat{\Delta}(\Psi(\Pi))$ with $[0, 1]^{\mathcal{W}(\Pi)}$. In this way, $\Psi(\Pi)$ is the set of corners $\hat{\Delta}(\Psi(\Pi))$, and Π is a subalgebra. Let $\Delta(\Pi) = \Delta(\Pi, \Psi(\Pi)) = \bigcup_{R \in \mathcal{C}(\Pi)} \Delta(R)$, as defined in Subsection 10.2. By Lemma 10.2.2, $\Delta(\Pi)$ is a subalgebra of $\hat{\Delta}(\Pi, \Psi(\Pi))$, hence intrinsically a median algebra. We refer to it as the *realisation* of Π . Note that Π is also a subalgebra of $\Delta(\Pi)$.

We remark that one can equip $\Delta(\Pi)$ with various metrics and topologies, as we discuss in Section 17. One such topology gives $\Delta(\Pi)$ the structure of a CW complex, with 1-skeleton $\Gamma(\Pi)$ and with n -cells the realisations of elements of $\mathcal{C}_n(\Pi)$.

There are a number of variations on this construction, for example, to construct subdivisions of discrete median algebras, and cube complexes. This will be discussed further in Subsection 11.3. For the moment, we just mention the *binary subdivision*, $\Sigma(\Pi)$, of Π . This can intuitively be thought of as taking each cell Q of Π and subdividing $\Delta(Q)$ into $2^{\text{rank}(Q)}$ “equal” subcubes. The 0-skeleton of the resulting complex is a subalgebra of $\Delta(\Pi)$, namely $\Sigma(\Pi)$. Note that we can identify $\Delta(\Pi)$ with $\Delta(\Sigma(\Pi))$. More formally, we can construct $\Sigma(Q)$ similarly as for $\Delta(Q)$, by replacing $\hat{\Delta}(\Psi(\Pi))$ with $\hat{\Sigma}(\Psi(\Pi)) := \{0, \frac{1}{2}, 1\}^{\mathcal{W}(\Pi)}$.

Note that each cell, $Q \in \mathcal{C}(\Pi)$, has a natural “centre”, $c(Q)$, where all the coordinates are equal to $\frac{1}{2}$. (Identifying Π with $\mathcal{C}_0(\Pi)$, we take this to be the identity on Π .) It is readily checked that the map $c : \mathcal{C}(\Pi) \rightarrow \Sigma(\Pi)$ is a median isomorphism.

11.3. Some basic properties.

We now consider some basic properties of Π .

Lemma 11.3.1. *Let $W_1, \dots, W_n \in \mathcal{W}(\Pi)$ be pairwise crossing walls. Then there is some $Q \in \mathcal{C}_n(\Pi)$ with $Q \pitchfork W_i$ for all i .*

Proof. This is a slight variation on the proof of Lemma 8.1.5. To begin, we claim that there is some finite convex set, $C \subseteq \Pi$, which crosses all the walls W_i . To see this, choose any a and b in antipodal orthants of the family W_1, \dots, W_n (for example, $a \in \bigcap_{i=1}^n W_i^-$ and $b \in \bigcap_{i=1}^n W_i^+$). We can then set $C = [a, b]$. Note that we can identify $\mathcal{W}(C)$ with $\{W \in \mathcal{W}(\Pi) \mid C \pitchfork W\}$.

We now choose such a finite convex set, $Q \subseteq \Pi$, with $\#\mathcal{W}(Q)$ minimal. By the same argument as in the proof of Lemma 8.1.5, we see that $\#(Q \cap O(\epsilon)) = 1$ for all orthants, $O(\epsilon)$, and so Q is an n -cube. \square

In fact, we see that the orthants of the family of walls $\{W_1, \dots, W_n\}$ are precisely the preimages of the elements of Q under the gate map of Π to Q .

Putting Lemma 11.3.1 together with Lemma 8.2.1, we immediately get:

Lemma 11.3.2. $\text{rank}(\Pi) = \sup\{n \in \mathbb{N} \mid \mathcal{C}_n(\Pi) \neq \emptyset\}$.

In other words, $\text{rank}(\Pi)$ is the same as the dimension (possibly infinite) of the realisation $\Delta(\Pi)$, in the usual sense of a CW complex.

We also note the following generalisation of Lemma 7.3.7.

Lemma 11.3.3. *Any non-empty convex subset of a discrete median algebra is gated.*

Proof. Let $C \subseteq \Pi$ be convex, and let $a \in \Pi$. The proof follows exactly as in that of Lemma 7.3.7, since we only used the fact that $C \cap [a, b]$ is finite for any $b \in C$.

Alternatively, we could just take $c \in C$ so as to minimise $\rho(a, c)$. It is readily checked that c is a gate for a . \square

The following observation will be used in the proof of Lemma 11.4.2 below, and is also relevant to the discussion of median graphs in Section 16.

Lemma 11.3.4. *Suppose $p, a, b, c \in \Pi$ with a, b distinct, both adjacent to c , and with $\rho(p, a) = \rho(p, b) = \rho(a, c) - 1$. Then there is some $d \in \Pi$, adjacent to both a and b and with $\rho(p, d) = \rho(p, c) - 2$.*

Proof. If $a = b$, let d be the penultimate vertex of any geodesic path from p to a . If $a \neq b$, Let $d = abc$. Since $\rho(a, b) = 2$, it follows that $\rho(a, d) = \rho(b, d) = 1$, and that $\rho(p, d) = \rho(p, c) - 2$. \square

Note that if $a \neq b$, then d is unique and a, b, d, c is a 2-cell of Π (see Lemma 10.3.6).

11.4. Links and local convexity.

We next consider a couple of combinatorial properties of Π (the first local, and the second global) which will be central to the discussion of CAT(0) cube complexes in Section 16.

Definition. Given $a \in \Pi$, the *link*, $L(a) \subseteq \Pi$, is the set of elements adjacent to a .

In other words, $L(a) = \{x \in \Pi \mid \{a, x\} \in \mathcal{C}_1(\Pi)\}$.

We can equip $L(a)$ with the structure of an abstract simplicial complex by taking the simplices to be sets of the form $Q \cap L(a)$ for $Q \in \mathcal{C}(\Pi)$. (Thus the “link” of a in $\Delta(\Pi)$, in the usual sense of a polyhedral complex, is the realisation of the abstract simplicial complex we have just defined.)

Recall that a simplicial complex is **flag** if every finite clique (i.e. complete subgraph) in the 1-skeleton is the 1-skeleton of a simplex of the complex.

Lemma 11.4.1. *For all $a \in \Pi$, the link, $L(a)$, is a flag complex.*

Proof. Suppose $A \subseteq L(a)$ be the vertex set of some finite clique in the 1-skeleton of $L(a)$. By definition, this means that if $b, c \in L(a)$, then there is some $d \in \Pi$ such that $\{a, b, c, d\}$ is a 2-cell of Π .

We claim that there is some cube Q in Π with $Q \cap L(a) = A$. To see this, let $Q_0 = \prod_{i=1}^n \{a, e_i\}$. In the notation of Subsection 10.3, for all $i \neq j$, we have a monomorphism of the 2-face, $\{a, e_i, e_j, e_{ij}\}$, into Π . By Lemma 10.3.7, this extends to a monomorphism of Q_0 into Π , and we let Q be its image. This proves the claim.

We also claim that Q is convex in Π . One way to see this is to apply Lemma 10.3.5. In the notation there, $D_i \equiv [a, e_i]_{\Pi} = \{a, e_i\}$, and so $D = Q$. Lemma 10.3.5 therefore tells us that $\text{hull}_{\Pi}(Q) = D = Q$.

This shows that Q is a cell of Π . Therefore $A = Q \cap L(a)$ is a simplex of the link $L(a)$. □

Our global property relates to 1-paths in Π .

Let $\underline{a} = a_0, a_1, \dots, a_n$ be a 1-path. We consider the following two “moves”.

- (1): If $a_{i-1} = a_{i+1}$, we replace a_{i-1}, a_i, a_{i+1} with a_{i-1} .
- (2): If there is some $b \in \Pi$ such that a_{i-1}, a_i, a_{i+1}, b is a 2-cell of Π (with b antipodal to a_i), we replace a_i with b .

Lemma 11.4.2. *Let $a, b \in \Pi$, and let \underline{a} be 1-path from a to b in Π . Then we can reduce \underline{a} to a geodesic path from a to b , by applying a finite sequence of moves of type (1) and (2) described above.*

Proof. Let $A(\underline{a}) = \sum_{i=0}^n \rho(a_0, a_n)$. We claim that if \underline{a} is not geodesic, then there is a move which strictly reduces $A(\underline{a})$. We use the fact that $\mathcal{W}(a_{i-1}, a_i) = \{W_i\}$ for some $W_i \in \mathcal{W}(\Pi)$.

To verify the claim, let $m > 0$ be minimal such that $\rho(a, a_{m+1}) \leq m$. Then $\rho(a, a_m) = m$ and $\rho(a, a_{m-1}) = \rho(a, a_{m+1}) = m - 1$. (Note that $\rho(a, a_i)$ changes by -1 or $+1$ at each step, depending on whether or not $W_i \in \mathcal{W}(a, a_{i-1})$.) We see that $a_{m-1}, a_{m+1} \in [a, a_m]$. If $a_{m-1} = a_{m+1}$ we do Move (1). So suppose $a_{m-1} \neq a_{m+1}$. Let $d = aa_{m-1}a_{m+1}$ (as in Lemma 11.3.4). Then a_{m-1}, a_m, a_{m+1}, d is a 2-cell of Π , with d antipodal to a_m . We now do Move (2). Since $\rho(a, d) = m - 2$, we have reduced $A(\underline{a})$ (by at least 2).

After a finite number of such steps, we arrive at a geodesic as required. \square

We say that a path, \underline{a} , is **closed** if $a_n = a_0$.

Clearly, the only geodesic from a_0 to itself is the constant path, a_0 . Therefore as an immediate consequence of Lemma 11.4.2, we have:

Lemma 11.4.3. *Let \underline{a} be a closed path in Π . Then we can reduce \underline{a} to the constant path a_0 , by applying a finite sequence of moves of type (1) and (2) described above.*

Note that in the above procedure, $A(\underline{a})$, reduces by at least 2 at each step. If \underline{a} has length n , then (in the notation of the proof of Lemma 11.4.2) we have $A(\underline{a}) \leq n^2/2$. It therefore takes us at most $n^2/4$ steps to reduce \underline{a} to a constant path.

In terms of the realisation, $\Delta(\Pi)$, we see that a path of length n in the 1-skeleton, $\Gamma(\Pi)$, bounds a singular disc in $\Delta(\Pi)$ comprising at most $n^2/4$ 2-cells. This is a form of the quadratic isoperimetric inequality. Note in particular that it implies that $\Delta(\Pi)$ is simply connected. We will say more about this in Section 16.

We can also describe convexity in these terms.

Definition. A subset $A \subseteq \Pi$ is **locally convex** if $\#(Q \cap A) \neq 3$ for all $Q \in \mathcal{C}_2(\Pi)$.

This can be phrased more intuitively in terms of links. Let $L(a) \subseteq \Pi$ be the link in $a \in A$. Recall that a simplex of $L(a)$ has the form $Q \cap L(a)$ for some $Q \in \mathcal{C}(\Pi)$. We define a subcomplex of $L(a)$ taking as simplices those simplices $Q \cap L(a)$ for which $Q \subseteq A$. This has vertex set A . In these terms, local convexity asserts that if $b, c \in A \cap L(a)$ are adjacent in $L(a)$ (that is, if $\{b, c\}$ is a 1-simplex) b, c are also adjacent in the subcomplex. (In fact, using Lemma 10.3.7, we see that the subcomplex is a full subcomplex of $L(a)$.) This therefore corresponds to a geometric notion of local convexity, which we will return to in Subsection 18.2.

Note that if $a, b, c, d \in \Pi$ with c, d adjacent, Recall the definition of “1-path-connected” from Subsection 11.1.

Lemma 11.4.4. *A subset $A \subseteq \Pi$ is convex if and only if it is 1-path-connected and locally convex.*

Proof. Suppose A is convex. Then any geodesic path between $a, b \in A$ lies in $[a, b] \subseteq A$, so A is 1-path-connected. If $Q \in \mathcal{C}_2(\Pi)$, then $Q \cap A$ is convex, so $\#(Q \cap A) \neq 3$.

For the converse, suppose A is 1-path-connected and locally convex. Let $a, b \in A$ be distinct.

We first claim that there is some geodesic path in Π from a to b which lies entirely in A . To see this, connect a to b by some 1-path, \underline{a} , and apply the procedure in the proof Lemma 11.4.2. Note that at each stage of the process, \underline{a} remains a subset of A . This is because in each application of Move (2), $a_{i-1}, a_i, a_{i+1} \in A$, so by local convexity, $b \in A$. When the process terminates, we end up with the required geodesic.

The geodesic we have constructed lies in $[a, b]$. It now follows in particular, that $L(a) \cap [a, b] \cap A \neq \emptyset$.

We finally want to show that $[a, b] \subseteq A$. We proceed by induction on $\rho(a, b)$. We first show that $L(a) \cap [a, b] \subseteq A$. By the previous paragraph, there is certainly some $d \in L(a) \cap [a, b] \cap A$. Suppose that $e \in L(a) \cap [a, b]$ with $e \neq d$. Now $\rho(d, b) = \rho(a, b) - 1$, so by the inductive hypothesis, $[d, b] \subseteq A$. Let $f = bde$. Then $f \in [d, b] \subseteq A$. Also a, d, f, e is a 2-cell of Π , with e antipodal to d . Since $a, d, f \in A$, by local convexity, we have $e \in A$. We have shown that $L(a) \cap [a, b] \subseteq A$ as claimed.

Finally, we observe that $[a, b] = \{a\} \cup \bigcup_{e \in L(a) \cap [a, b]} [e, b]$, and so $[a, b] \subseteq A$, again by the inductive hypothesis. \square

11.5. Parallel edges and convex hulls.

Next we consider convex hulls in Π .

Let $W \in \mathcal{W}(\Pi)$. Let $\mathcal{E} = \mathcal{E}(W) \subseteq \mathcal{C}_1(\Pi)$ be the set of 1-cells of Π that cross W .

Lemma 11.5.1. *\mathcal{E} is a parallel class in Π .*

Proof. Let $\epsilon, \epsilon' \in \mathcal{E}$. Then the gate map $\omega_\epsilon | \epsilon'$ must be injective, hence an isomorphism, and so $\epsilon \parallel \epsilon'$. Conversely, if $\epsilon \parallel \epsilon''$, then $\mathcal{W}(\epsilon) = \mathcal{W}(\epsilon'')$, so $\epsilon'' \in \mathcal{E}$. \square

It follows from the discussion of parallelism in Subsection 7.2 that \mathcal{E} has a natural structure as a median algebra. We write $\Pi_W = \bigcup \mathcal{E}$. Then Π_W is convex in Π and is naturally isomorphic to $\mathcal{E} \times \{0, 1\}$. (See Lemmas 7.2.6 and 7.2.7.) We write $\Pi_W^\pm = \Pi_W \cap W^\pm$, with the convention that Π_W^- gets identified with $\mathcal{E} \times \{0\}$ and Π_W^+ with $\mathcal{E} \times \{1\}$. We note that Π_W^- and Π_W^+ are also parallel in Π .

Remark. It might be helpful to think of walls of Π in terms of “hyperplanes” in the realisation $\Delta(\Pi)$ in the following sense. If $W \in \mathcal{W}(\Pi)$, then we can embed $\Delta(\mathcal{E})$ in $\Delta(\Pi)$ as $\Delta(\mathcal{E}) \times \{\frac{1}{2}\} \subseteq \Delta(\mathcal{E}) \times [0, 1] \equiv \Delta(\Pi_W) \subseteq \Delta(\Pi)$. This is a convex subset of $\Delta(\Pi)$ which cuts $\Delta(\Pi)$ into two pieces, one containing $\Delta(W^-)$ and the other containing $\Delta(W^+)$. We will say more about this in Subsection 18.2.

Lemma 11.5.2. $\text{rank}(\mathcal{E}(W)) \leq \text{rank}(\Pi) - 1$.

Proof. If $Q \subseteq \mathcal{E}$ is a cube in \mathcal{E} , then $Q \times \{0, 1\}$ is a cube in $\mathcal{E} \times \{0, 1\} \equiv \Pi_W \subseteq \Pi$. \square

Let $\omega = \omega_{\Pi_W} : \Pi \rightarrow \Pi_W$ be the gate map to Π_W . It is easily seen that $\omega(W^\pm) = \Pi_W^\pm$, and that $\omega|W^\pm$ is the gate map to Π_W^\pm in W^\pm . We can postcompose ω with the projection of Π_W to \mathcal{E} , to give us an epimorphism, $\pi_W : \Pi \rightarrow \mathcal{E}$. This induces an

inclusion of $\mathcal{W}(\mathcal{E})$ into $\mathcal{W}(\Pi)$. Its image is $\mathcal{W}(W, \Pi) := \{W' \in \mathcal{W}(\Pi) \mid W' \pitchfork W\}$. Note that two such walls cross in Π if and only if they cross in \mathcal{E} .

Suppose that $a \in W^-$ and $b \in W^+$. Then $[a, b] \cap \Pi_W \neq \emptyset$. (If $a = x_0, \dots, x_n = b$ is a geodesic sequence from a to b , then $\{x_i, x_{i+1}\} \in \mathcal{E}$ for some i , so $x_i, x_{i+1} \in [a, b] \cap \Pi_W$.) Since gate maps are epimorphisms (Lemma 7.3.1) we have $\omega([a, b]) = [\omega a, \omega b]$. Identifying $\Pi_W \equiv \mathcal{E} \times \{0, 1\}$, we have $\omega a \in \mathcal{E} \times \{0\}$ and $\omega b \in \mathcal{E} \times \{1\}$. It follows that $\omega([a, b]) \equiv [\pi_W a, \pi_W b]_{\mathcal{E}} \times \{0, 1\}$ (see for example, Lemma 7.5.1).

Let $A \subseteq \Pi$. Recall that the join of A is defined by $J(A) := \bigcup_{a, b \in A} [a, b]$. We use $J_{\mathcal{E}}$ to denote join in \mathcal{E} . From the above observations we see:

Lemma 11.5.3. *Suppose $A \subseteq \Pi$ and $A \pitchfork \Pi$. Then $\omega(A) \subseteq J(A)$. Moreover, $J(A) \cap \Pi_W \equiv J_{\mathcal{E}}(\pi_W A) \times \{0, 1\}$.*

Recall that in general $\text{hull}(A) = \bigcup_{i=0}^{\infty} J^i(A)$. We can now reduce this to a finite union in the finite-rank case as follows:

Proposition 11.5.4. *Let Π be a discrete median algebra of rank at most ν . If $A \subseteq \Pi$, then $\text{hull}(A) = J^{\nu}(A)$.*

Proof. We need to show that $\text{hull}(A) \subseteq J^{\nu}(A)$. We proceed by induction on ν . Note that if $\nu = 0$, then Π is a singleton. This case is trivial, so we assume $\nu > 0$. We can also assume that $\#A \geq 2$.

To begin, let $W \in \mathcal{W}$ be any wall with $A \pitchfork W$. Let $\omega : \Pi \rightarrow \Pi_W$ be the gate map. Since ω is an epimorphism, $\omega^{-1}(\text{hull}(\omega A))$ is convex in Π and contains A . Therefore $\text{hull}(A) \subseteq \omega^{-1}(\text{hull}(\omega A))$, so $\omega(\text{hull}(A)) \subseteq \text{hull}(\omega A)$. (In fact, one can check that these sets are equal.) Moreover, since $A \pitchfork W$, we have $\text{hull}(A) \cap \Pi_W \neq \emptyset$, and so, $\omega(\text{hull}(A)) = \text{hull}(A) \cap \Pi_W$ (see Lemma 7.3.5). Therefore, $\text{hull}(A) \cap \Pi_W \subseteq \text{hull}(\omega A)$. By Lemma 11.5.3, $\omega A \subseteq J(A) \cap \Pi_W$, so $\text{hull}(A) \cap \Pi_W \subseteq \text{hull}(J(A) \cap \Pi_W)$. Again by Lemma 11.5.3, $J(A) \cap \Pi_W$ is just a product $B \times \{0, 1\}$ in $\mathcal{E} \times \{0, 1\}$, where $\mathcal{E} = \mathcal{E}(W)$ and $B = J_{\mathcal{E}}(\pi_W A)$. By Lemma 11.5.2, $\text{rank}(\mathcal{E}) \leq \nu - 1$. Therefore, by the inductive hypothesis in \mathcal{E} , we have $\text{hull}_{\mathcal{E}}(B) = J_{\mathcal{E}}^{\nu-1}(B)$, and so in Π_W , we have $\text{hull}(J(A) \cap \Pi_W) = J^{\nu-1}(J(A) \cap \Pi_W)$. Therefore, $\text{hull}(A) \cap \Pi_W \subseteq \text{hull}(J(A) \cap \Pi_W) \subseteq J^{\nu-1}(J(A) \cap \Pi_W) = J^{\nu-1}(J(A)) = J^{\nu}(A)$.

We now apply this to prove the proposition. Recall that we are assuming that $\#A \geq 2$. In this case, we claim that $\text{hull}(A) \subseteq \bigcup \{\Pi_W \mid W \pitchfork A\}$. To see this, let $a \in \text{hull}(A)$. Choose any $b \in A \setminus \{a\}$. Choose $c \in [a, b]$ adjacent to a . Then $\mathcal{W}(a, c) = \{W\}$ for some $W \in \mathcal{W}(\Pi)$. Now $\{a, c\} \in \mathcal{E}(W)$, and so $a \in \Pi_W$, proving the claim.

It now follows that $\text{hull}(A) \subseteq \bigcup_{W \pitchfork A} (\text{hull}(A) \cap \Pi_W) \subseteq J^{\nu}(A)$ as required. \square

As noted in Subsection 8.2, Proposition 11.5.4 holds without the assumption that Π is discrete, since it can be reduced to the case of finite median algebras which we have now proven. This therefore completes the proof of Proposition 8.2.3.

We also note:

Lemma 11.5.5. *Let \mathcal{C} be any non-empty family of pairwise intersecting bounded convex subsets of Π . Then $\bigcap \mathcal{C} \neq \emptyset$.*

Proof. By the Helly Property (Lemma 7.1.1), we can assume that \mathcal{C} is closed under finite intersection. Choose any $a \in \Pi$. Note that $\rho(a, C)$ is bounded above for all $C \in \mathcal{C}$ (by $\text{diam}(\{a\} \cup B)$ for any choice of $B \in \mathcal{C}$). Choose $C \in \mathcal{C}$ so as to maximise $\rho(a, C)$. Let $c \in C$ be the gate for a in C (which exists by Lemma 11.3.3: it is the unique $c \in C$ with $\rho(a, c) = \rho(a, C)$). We claim that $c \in D$ for all $D \in \mathcal{C}$. To see this, let $d \in C \cap D$ be the gate for a in $C \cap D \in \mathcal{C}$. Since $d \in C$, we have $a.c.d$, and so $\rho(a, d) = \rho(a, c) + \rho(c, d)$. But $\rho(a, d) = \rho(a, C \cap D)$, so by maximality, we have $\rho(c, d) = 0$, so $c = d \in D$ as claimed. \square

(Of course, we only really need one of the elements of \mathcal{C} to be bounded.)

We will give a generalisation of this statement in Subsection 13.3 (see Lemma 13.3.9). Related results for complete metric spaces are given by Lemmas 22.2.1 and 22.2.2.

11.6. Canonical paths.

We now consider certain canonical paths in a discrete median algebra.

Let $a, b \in \Pi$. Recall that $\mathcal{W}(a, b) = \{W \in \mathcal{W}(P) \mid a|_W b\}$, and that $\rho(a, b) = \#\mathcal{W}(a, b)$. Given $W \in \mathcal{W}(a, b)$, we direct W so that $a \in W^-$ and $b \in W^+$. Given $W_1, W_2 \in \mathcal{W}(a, b)$, we write $W_1 \leq W_2$ to mean that $W_1^- \subseteq W_2^-$. This is equivalent to $W_2^+ \subseteq W_1^+$. We also note that this is intrinsic to $[a, b]$: it is equivalent to $W_1^- \cap [a, b] \subseteq W_2^- \cap [a, b]$ (since $x \in W_i^\pm$ if and only if $abx \in W_i^\pm$). Clearly \leq is transitive. Moreover, if $W_1, W_2 \in \mathcal{W}(a, b)$, then exactly one of the relations, $W_1 = W_2$, $W_1 < W_2$, $W_2 < W_1$ or $W_1 \pitchfork W_2$ holds. We say that a sequence W_1, W_2, \dots, W_n in $\mathcal{W}(a, b)$ is a **chain** (of length n) if $W_1 < W_2 < \dots < W_n$. We write $\sigma(a, b)$ for the maximal length of a chain in $\mathcal{W}(a, b)$. Clearly, $\sigma(a, b) \leq \rho(a, b)$.

Lemma 11.6.1. *σ is a metric on Π .*

Proof. Clearly $\sigma(a, b) = \sigma(b, a)$, and $\sigma(a, b) = 0$ if and only if $a = b$ (by Lemma 3.2.12). To check the triangle inequality, let $c \in \Pi$, and let $W_1 < W_2 < \dots < W_n$ be a chain in $[a, b]$ with $n = \sigma(a, b)$. Let $m \in \{0, 1, \dots, n\}$ be maximal such that $c \notin W_m^-$ (with the convention that $W_0^- = \emptyset$). Then $W_1 < \dots < W_m$ and $W_{m+1} < \dots < W_n$ are (possibly empty) chains in $\mathcal{W}(a, c)$ and $\mathcal{W}(b, c)$ respectively, so $\sigma(a, c) \geq m$ and $\sigma(c, b) \geq n - m$. Thus $\sigma(a, b) \leq \sigma(a, c) + \sigma(c, b)$. \square

Note that if $a \neq b$, then $\sigma(a, b) = 1$ if and only if $[a, b]$ is a cube (since the walls of $\mathcal{W}(a, b)$ pairwise cross).

Definition. A sequence $a = a_0, \dots, a_n = b$ in Π is a **cube-path** from a to b if a_i, a_{i+1} lie in a cell of Π for all i .

This is the same as saying that $\sigma(a_i, a_{i+1}) \leq 1$. Therefore, $\sigma(a, b) \leq n$. In fact:

Lemma 11.6.2. *$\sigma(a, b)$ is the minimal length of a cube-path from a to b . Moreover, such a minimal cube-path can be constructed in $[a, b]$.*

In fact, there is a canonical cube-path from a to b constructed as follows.

Let $\mathcal{W}_0(a)$ be the set of all walls of Π adjacent to a . In other words, $\mathcal{W}_0(a) = \mathcal{W}(L(a))$ is the set of walls which cross edges of $\Gamma(\Pi)$ incident on a . Now the walls of $\mathcal{W}_0(a) \cap \mathcal{W}(a, b)$ pairwise cross. Therefore, there is a unique $p = p(a, b) \in [a, b]$, such that $[a, p]$ is a cell of Π with $\mathcal{W}([a, p]) = \mathcal{W}(a, p) = \mathcal{W}_0(a) \cap \mathcal{W}(a, b)$. (Writing $[a, b] \cap L(a) = \{x_1, \dots, x_m\}$, we can set $p = x_1 \vee \dots \vee x_m$ in the lattice structure on $[a, b]$.) Note that $\mathcal{W}(p, b) = \mathcal{W}(a, b) \setminus \mathcal{W}_0(a)$, and so $\sigma(a, b) \leq \sigma(p, b) + 1$. We claim also that $\sigma(a, b) \geq \sigma(p, b) + 1$. For let $W_1 < W_2 < \dots < W_m$ be a chain in $\mathcal{W}(p, b)$ with $m = \sigma(p, b)$. Now $W_1 \notin \mathcal{W}_0(a)$, and so there is some $W_0 \in \mathcal{W}_0(a) \cap \mathcal{W}(a, b)$ with $W_0 < W_1$. Thus $W_0 < W_1 < \dots < W_m$ is a chain in $\mathcal{W}(a, b)$, so $\sigma(a, b) \geq m + 1$ as claimed. Thus, $\sigma(a, b) = \sigma(p, b) + 1$. (We remark that $\rho(a, b)$ can be defined as the projection of b to the ‘‘star’’, $S(a)$, as defined in Subsection 11.7, under the gate map to $S(a)$.)

We now define a cube-path inductively by setting $a_0 = a$ and $a_{i+1} = p(a_i, b)$. This gives a cube path of length $\sigma(a, b)$ from a to b in $[a, b]$.

This proves Lemma 11.6.2.

We also note that from the construction, a_0, \dots, a_m is also the canonical cube-path from a to a_m for all $m \leq n$.

Given $a \in \Pi$ and $r \in \mathbb{N}$, let $N_\sigma(a, r) = \{x \in \Pi \mid \sigma(a, x) \leq r\}$.

Lemma 11.6.3. $N_\sigma(a, r)$ is convex.

Proof. Suppose $b \notin N_\sigma(a, r)$. Let $W_1 < W_2 < \dots < W_{r+1}$ be a chain in $\mathcal{W}(a, b)$. Thus $b \notin W_{r+1}^-$. We claim that $N_\sigma(a, r) \subseteq W_{r+1}^-$. For if $c \in W_{r+1}^+$, then $W_1 < W_2 < \dots < W_{r+1}$ is also a chain in $\mathcal{W}(a, c)$, so $\sigma(a, c) \geq r + 1$. We see that W_{r+1}^- is a halfspace containing $N_\sigma(a, r)$ but not b . This shows that $N_\sigma(a, r)$ is an intersection of halfspaces, hence convex. \square

Lemma 11.6.4. Let $(a_i)_{i \in \mathcal{I}}$ and $(r_i)_{i \in \mathcal{I}}$ be families in Π and \mathbb{N} respectively, both indexed by some set \mathcal{I} . Suppose that $\sigma(a_i, a_j) \leq r_i + r_j$ for all $i, j \in \mathcal{I}$. Then $\bigcap_{i \in \mathcal{I}} N_\sigma(a_i, r_i) \neq \emptyset$.

Proof. By Lemma 11.6.3, each $N_\sigma(a_i, r_i)$ is convex. For each $i, j \in \mathcal{I}$, $N_\sigma(a_i, r_i) \cap N_\sigma(a_j, r_j) \neq \emptyset$. (Take a cube-path of length $\sigma(a_i, a_j)$ from a_i to a_j . Its r_i th element lies in the intersection.)

By Lemma 11.5.5, any pairwise intersecting family of bounded convex sets in discrete median algebra has non-empty intersection.

The statement now follows. \square

The above can be reinterpreted as follows.

Let $\Gamma = \Gamma(\Pi)$ be the graph with vertex set, $V(\Gamma) = \Pi$, and edges given by adjacency in Π , as defined above. Let $\Gamma^\Delta \supseteq \Gamma$ be the graph with vertex set $V(\Gamma^\Delta) = V(\Gamma) = \Pi$, and where a, b are deemed adjacent in Γ^Δ if they both lie in some cell of Π (in other words, $\sigma(a, b) \leq 1$). By Lemma 11.6.2, σ is precisely the combinatorial metric induced from Γ^Δ . In these terms, the conclusion of Lemma

11.6.4 is the statement that Γ^Δ is a “helly graph”: something we will mention again in Subsection 25.2.

Given any $a, b \in \Pi$, the canonical cube-path from a to b gives us a geodesic, $\alpha(a, b)$, in Γ^Δ . Let $T_a = \bigcup_{b \in \Pi} \alpha(a, b)$. Then T_a is a maximal subtree of Γ^Δ . This is the combing of Γ^Δ , as discussed in [NibR2]. It is also related to the Roller boundary of Π , see Proposition 11.12.8.

11.7. Generating subalgebras.

We next turn to the proof of Proposition 8.2.4. Recall that this asserts that in the finite-rank case, the median algebra generated by a subset can be obtained by iterating the median operation a bounded number of times.

We first define the notion of a “star” related to the notion of a link of a point defined earlier.

Let Π be a discrete median algebra and let $a \in \Pi$. Let $L(a)$ be the link of a . Let $S(a)$ be the union of all cells of Π which contain a . We refer to $S(a)$ as the **star** of a . Thus $b \in S(a)$ if and only if $[a, b]$ is a cube, and if and only if the walls of $\mathcal{W}(a, b)$ pairwise cross. We note:

Lemma 11.7.1. *$S(a)$ is convex.*

Proof. Let $b, c \in \Pi$ and let $x \in [b, c]$. If $x \notin S(a)$, then there are distinct walls, $W_1, W_2 \in \mathcal{W}(a, x)$ which do not cross. We can assume that $x \in W_2^+ \subseteq W_1^+$. Up to swapping b with c , we can also assume that $b \in W_2^+$, so $W_1, W_2 \in \mathcal{W}(a, b)$, so $b \notin S(a)$. \square

We therefore have a gate map, $\omega = \omega_S : \Pi \rightarrow S$. One sees easily that $\omega^{-1}(a) = \{a\}$. (In fact, as we noted in Subsection 11.6, if $b \neq a$, then ωb is the first vertex of the canonical cube path from a to b .)

Let $G = G(a)$ be the graph with vertex set $V(G) = L(a)$, where $b, c \in L(a)$ are deemed adjacent in $G(a)$ if a, b, c all lie in a 2-cell of Π . Let $S = S(G) \subseteq \mathcal{P}(L(a))$ be the discrete median algebra described by Example (Ex3.2) of Subsection 3.4. Let $\alpha = \emptyset \in S$ be the central element. Given $\beta \in S$, write $Q(\beta) = [\alpha, \beta]$ for the cube with antipodal vertices α, β . Then the map $[\beta \mapsto Q(\beta)]$ is a bijection from S to the set, $\mathcal{C}(S)$, of cells of S . In fact this is a median isomorphism between their respective median structures.

We now move on to the proof of Proposition 8.2.4. Recall that if $A \subseteq M$ is any subset of a median algebra, M , then we defined $T(A)$ by $T(A) := \{abc \mid a, b, c \in A\}$. We will need the following simple observation. If $\phi : M \rightarrow N$ is a homomorphism of median algebras, and $A \subseteq M$, then $T(\phi(A)) = \phi(T(A))$. By iterating (or more directly) we get $\phi(\langle A \rangle) = \langle \phi(A) \rangle$.

Lemma 11.7.2. *Let Q be a n -cube and $A \subseteq Q$ with $Q = \langle A \rangle$. Then $Q = T^n(A)$.*

Proof. We can assume that $n \geq 3$ (otherwise any subset is a subalgebra). Let $a \in Q$ and let e_1, \dots, e_n be the adjacent vertices. By collapsing the parallel class

of the edge $\{a, e_i\}$, we can assume inductively that there is some $c_i \in \{a, e_i\}$ with $c_i \in T^{n-1}(A)$. Now $a = c_1 c_2 c_3$, and so the statement follows by induction. \square

Next we have following lemma regarding simplex graphs. Let G be a graph, and let $S = S(G)$, and let $\alpha \in S$ be the central vertex.

Lemma 11.7.3. *Suppose that $\text{rank}(S) \leq \nu < \infty$, and that $S = \langle A \rangle$ for some $A \subseteq S$. Then $S = T^{2\nu}(A)$.*

Proof. Let $S_0 \subseteq S$ be the set of maximal elements of S with respect to inclusion (in other words, the vertex sets of maximal cliques of G). Since $\text{rank}(S) < \infty$, every element of S is contained in an element of S_0 . Let $S_1 \subseteq S$ be the set of those $\gamma \in S$ such that there is precisely one $\beta \in S_0$ with $\gamma \subseteq \beta$. We write $\beta(\gamma) = \beta$. Note that $S_0 \subseteq S_1$. In fact, if $\gamma \in S_0$, then $\beta(\gamma) = \gamma$. Note also that if $\gamma \in S_1$, and $\delta \in S$ with $\delta \supseteq \gamma$, then $\delta \in S_1$ and $\beta(\delta) = \beta(\gamma)$.

Let $\gamma \in S_1$. Let $\beta = \beta(\gamma)$ and let $Q = Q(\beta) \subseteq S$. Given $\delta \in S$, write $\omega(\delta) = \beta \cap \delta$. Thus, $\omega : S \rightarrow Q$ is an epimorphism. (In fact, it is the gate map to Q .) We note that $\omega^{-1}(\gamma) = \{\gamma\}$. (For if $\omega(\delta) = \beta \cap \delta = \gamma$, then $\gamma \subseteq \delta$. Since $\gamma \in S_1$, we have $\delta \subseteq \beta$, so $\delta = \beta \cap \delta = \gamma$.) Now $Q = \omega S = \omega \langle A \rangle = \langle \omega A \rangle$, so by Lemma 11.7.2, $\gamma \in T^\nu(\omega A) = \omega(T^\nu(A))$. Thus $\gamma = \omega \epsilon$ for some $\epsilon \in T^\nu(A)$, and so in fact $\gamma = \epsilon \in T^\nu(A)$.

We now claim that if $n \leq \nu$ and $\gamma \in S$ with $\#\gamma \leq n$, then $\gamma \in T^{2\nu-n}(A)$. We prove this by downward induction on n .

By the above, the claim holds when $\gamma \in S_1$. In particular, it holds for all $\gamma \in S$ with $\#\gamma = \nu$ (since then $\gamma \in S_0 \subseteq S_1$), and so the induction starts.

We can therefore assume that $\gamma \notin S_1$. This implies that there are distinct elements $d_1, d_2 \in V(G) \setminus \gamma$ such that $\delta_1 := \gamma \cup \{d_1\}$ and $\delta_2 := \gamma \cup \{d_2\}$ both lie in S .

We claim that there is some $\epsilon \in A$ with $d_1, d_2 \notin \epsilon$. For suppose not. Let $F = \{\zeta \in S \mid \zeta \cap \{d_1, d_2\} \neq \emptyset\}$. Now $F \leq S$ is a subalgebra and $A \subseteq F$. Thus $F = \langle A \rangle = S$. In particular $\alpha = \emptyset \in F$ giving a contradiction.

Now $\delta_1 \delta_2 \epsilon \subseteq \delta_1 \cup \delta_2$ and $d_1, d_2 \notin \delta_1 \delta_2 \epsilon$. Therefore, $\delta_1 \delta_2 \epsilon = \delta_1 \cap \delta_2 = \gamma$. By the inductive hypothesis, $\delta_i \in T^{2\nu-(n+1)}(A)$, and so $\gamma \in T^{2\nu-n}(A)$, proving the claim.

In particular we see that $\gamma \in T^{2\nu-0}(A)$ as required. \square

We can now finally prove Proposition 8.2.4:

Proof of Proposition 8.2.4. We have observed that we can assume that A is finite and that $M = \langle A \rangle$. Let $a \in M$, let $S = S(a)$ be the star at a , and let $\omega : M \rightarrow S$ be the gate map. Now $S = \langle \omega A \rangle$. Also S is isomorphic to $S(G(a))$ by an isomorphism taking a to the central vertex, α , of $S(G(a))$. Therefore, by Lemma 11.7.3, we have $a \in T^{2\nu}(\omega A) = \omega(T^{2\nu}(A))$. Now $a = \omega b$ for some $b \in T^{2\nu}(A)$, so in fact, $a = b \in T^{2\nu}(A)$ as required. \square

11.8. Retracts.

In this subsection, we prove a result of [Ban] which (in our terminology) relates 1-path-connected subalgebras of a discrete median algebra to metric retracts.

Suppose $A \subseteq \Pi$ is (a-priori) any subset of a discrete median algebra Π . A **(metric) retraction** to A is a 1-lipschitz map, $f : \Pi \rightarrow \Pi$, such that $f(\Pi) = A$ and f is the identity on A . (We are not assuming here that f is median homomorphism.) If such a map exists we refer to A as a **(metric) retract** of Π . In such a case, we can extend f to a retraction of the adjacency graph, $\Gamma(\Pi)$, to the full subgraph, $\Gamma(A)$, with vertex set A . (This is in the graph-theoretical sense of a map which sends each edge to a vertex or edge.) Note in particular, that A is 1-path-connected.

In fact, A is a subalgebra of Π . To see this, suppose $a, b \in A$ and $x \in [a, b]_\Pi$. Then $\rho(a, fx) + \rho(b, fx) \leq \rho(a, x) + \rho(b, x) = \rho(a, b)$, and so $fx \in [a, b]_\Pi$. It follows that if $a, b, c \in A$, then $f(abc) = abc$, so $abc \in A$ as required.

Note that by Lemma 11.1.1, $\Gamma(A)$ is an isometrically embedded copy of the adjacency graph of A .

In particular, we have shown:

Lemma 11.8.1. *Any metric retract of a discrete median algebra is a subalgebra.*

The converse also holds. In fact, we can make a stronger statement. We say that a map $f : \Pi \rightarrow \Pi$ is a **folding** to A if it is the identity on A , and if fx is adjacent to fy whenever x is adjacent to y . We say that A is a **folding** of Π . Clearly any folding is a retraction. In this case, the extension of f to a map $f : \Gamma(\Pi) \rightarrow \Gamma(\Pi)$ sends edges to edges.

Note that, if A is convex, then the gate map to A is a retraction, but not a folding (unless $A = \Pi$).

We claim:

Proposition 11.8.2. *Let $A \subseteq \Pi$ be a non-empty subset of a discrete median algebra, Π . The following are equivalent:*

- (1) *A is a 1-path-connected subalgebra of Π ,*
- (2) *A is a retract of Π ,*
- (3) *A is a singleton or a folding of Π .*

In view of Lemma 11.8.1, we just need to prove (1) \Rightarrow (3).

Let $A \leq \Pi$ be a non-empty 1-path-connected subalgebra, with $A \neq \Pi$. Write $A^* = \Pi \setminus A$. Given $a \in \Pi$ write $A(a) \subseteq A$ for the set of $b \in A$ which are adjacent to a in Π . This must be non-empty for at least one element $a \in A^*$. Moreover, $\#A(a) \leq 2$ for all $a \in A^*$ (for if $x, y, z \in A(a)$ were distinct, we would have $a = xyz \in A$). Also, by Lemma 11.4.4, if $\#A(a) \leq 1$ for all $a \in A^*$, then A is convex.

Let us suppose that $a \in A^*$ with $A(a) = \{a_1, a_2\}$ where $a_1 \neq a_2$. Since A is 1-path-connected, there is some $a' \in A$ adjacent to both a_1 and a_2 . In fact, a' is unique, and $\{a, a_1, a_2, a'\}$ is a 2-cube. We write $\mathcal{W}(a, a_i) = \{W_i\}$, so that $\mathcal{W}(a, a') = \{W_1, W_2\}$, and $W_1 \pitchfork W_2$.

Let B be the set of $b \in \Pi$ such that there exist $b_1, b_2 \in A$ with $bb_i \parallel aa_i$. (Recall the discussion of parallel sets in Subsection 7.2.) Note that $b \in A^*$ (otherwise $a = a_1a_2b \in A$). In particular, $A \cap B = \emptyset$. We also see that $A(b) = \{b_1, b_2\}$, and that $\mathcal{W}(b, b_i) = \{W_i\}$. Setting $b' = a'b_1b_2$, we have $b' \in A$ adjacent to both b_1 and b_2 . In fact, if $b \neq a$, then $\{a, a_1, a_2, a', b, b_1, b_2, b'\}$ is a 3-cube. Note also that if $a, b \in B$ are adjacent in Π , then a', b' are adjacent in A .

Lemma 11.8.3. *$A \cup B$ is a 1-path-connected subalgebra of Π .*

Proof. Certainly $A \cup B$ is 1-path-connected. We need to check that $A \cup B \leq \Pi$. Let $a, b, c \in A \cup B$, and set $d = abc$. We can suppose a, b, c do not all lie in A . There are thus three cases to consider.

Case (1): $a, b, c \in B$. Let $d_i = a_ib_ic_i \in A$. By Lemma 7.2.3, $dd_i \parallel aa_i$, so $A(d) = \{d_1, d_2\}$ and $d \in B$.

Case (2): $a, b \in B, c \in A$. Let $d_i = a_ib_ic \in A$. By Lemma 7.2.2, $dd_i \parallel aa_i$, so again, $d \in B$.

Case (3): $a \in B, b, c \in A$. Let $d_i = a_ibc$. Then $\mathcal{W}(d, d_i) \subseteq \mathcal{W}(a, a_i) \cap \mathcal{W}(b, c) = \{W_i\} \cap \mathcal{W}(b, c)$. If $W_i \notin \mathcal{W}(b, c)$, then $d = d_i \in A$. On the other hand, if $d \notin \{d_1, d_2\}$, then $W_1, W_2 \in \mathcal{W}(b, c)$, so $\mathcal{W}(d, d_i) = \{W_i\} = \mathcal{W}(a, a_i)$, so $dd_i \parallel aa_i$, so $d \in B$. \square

Proof of Proposition 11.8.2. As noted earlier, we just need to prove (1) \Rightarrow (3). So let $A \leq \Pi$ be a 1-path-connected subalgebra with $\#A \geq 2$. We want to construct a folding $f : \Pi \rightarrow A$.

For simplicity, we first consider the case where Π is finite. We will construct a strictly increasing sequence of 1-path-connected subalgebras, $A = A_0 \leq A_1 \leq \dots \leq A_n = \Pi$, and for $i < n$, a map $g_{i+1} : A_{i+1} \rightarrow A_i$, which is an intrinsic folding of A_{i+1} to A_i . We can then set $f = g_1 \circ g_2 \circ \dots \circ g_n$.

We begin by setting $A_0 = A$. Suppose we have found $A_i \neq \Pi$.

If A_i is convex, then we can find $a \in A^*$ and $a_0 \in A_i$ with a, a_0 adjacent in Π . We set $A_{i+1} = A_i \cup \{a\}$. This is a subalgebra of Π (being the union of two convex sets). Since $\#A_i \geq 2$, there is some $a' \in A_i$ adjacent to a_0 . We define g_{i+1} by setting $g_{i+1}(a) = a'$.

If A_i is not convex, there is some $a \in A^*$ with $\#A(a) = 2$. We now set $A_{i+1} = A_i \cup B$, where B is as defined earlier (substituting $A = A_i$). By Lemma 11.8.3, $A_{i+1} \leq \Pi$. We define g_{i+1} by setting $g_{i+1}(b) = b'$ (in our earlier notation) for all $b \in B$.

Since $\#A_i$ is strictly increasing, this process must terminate at some $A_n = \Pi$. We define f as above. This proves the result in the finite case.

(We remark that if we just wanted a retraction, we could stop as soon as we arrive at a convex set, and precompose with the gate map.)

If Π is infinite, we use transfinite induction. We construct subalgebras, A_α , indexed by ordinals α up to some ordinal γ with $A_\gamma = \Pi$, and such that $A_\beta \leq A_\alpha$ whenever $\beta \leq \alpha$. To do this, we again set $A_0 = A$. For a limit ordinal, α , we set $A_\alpha = \bigcup_{\beta < \alpha} A_\beta$. For a successor ordinal $\alpha + 1$, we construct $A_{\alpha+1}$ from A_α exactly

as in the finite case. Since the sequence is strictly increasing, it must terminate at some $A_\gamma = \Pi$.

For $\alpha > 0$, we define foldings $f_\alpha : A_\alpha \rightarrow A$, with $f_\alpha|_{A_\beta} = f_\beta$ whenever $\beta \leq \alpha$ as follows. If α is a limit ordinal, f_α is just the union of f_β for $\beta < \alpha$. For a successor ordinal $\alpha + 1$, we set $f_{\alpha+1} = f_\alpha \circ g_{\alpha+1}$, where $g_{\alpha+1} : A_{\alpha+1} \rightarrow A_\alpha$ is defined as in the finite case.

We finally set $f = f_\gamma$. □

We remark that in the case where Π is countable, one could use ordinary induction over the natural numbers, rather than transfinite induction. (Note that in any countable connected graph there is a finite or semi-infinite path which visits every vertex.)

In the finite case we get:

Proposition 11.8.4. *A non-empty finite graph is a median graph if and only if it is the retract of a cubical graph.*

Recall that a “median graph” is one which has the form $\Gamma(\Pi)$ for a discrete median algebra, Π . It is “cubical” if Π is a finite cube. (In the literature, a cubical graph is often referred to as a “hypercube”, though we have used this term here with a different meaning.) Median graphs will be the topic of Section 16.

Proof. This follows by Proposition 11.8.3, since we can embed any finite median algebra as a 1-path-connected subalgebra of a finite cube Ψ : see Proposition 3.2.13, and the discussion of Subsection 11.1. □

There is a similar statement in the infinite case, where we replace “a cubical graph” with “the adjacency graph of a cube” as we define in Subsection 11.11. For this, we use can Lemma 11.11.4 in place of Proposition 3.2.13.

There is also a version of Proposition 11.8.4 for quasimedial graphs: see the Notes to Section 23.

11.9. More about free median algebras.

We return to the discussion of free median algebras in the finite case.

Let X be a finite set. Recall from Subsection 5.1 that the free median algebra on X can be identified with the superextension, $\Phi(X)$, of X , with the natural inclusion, $\iota : X \hookrightarrow \Phi(X)$. A wall of $\Phi(X)$ determines a bipartition of X (taking the preimages of the halfspaces under ι). This gives a natural bijection between $\mathcal{W}(\Phi(X))$ and the set of such bipartitions. If $\mathcal{A} \in \Phi(X)$ and $W \in \mathcal{W}(\Phi(X))$, then precisely one element of the bipartition corresponding to W is contained in the family \mathcal{A} . This is the halfspace of W which contains \mathcal{A} in $\Phi(X)$. We have noted (Lemma 5.3.2) that each element of $\Phi(X)$ adjacent to \mathcal{A} is obtained by flipping a minimal element of \mathcal{A} . Thus the walls of $\Phi(X)$ adjacent to \mathcal{A} correspond exactly to the minimal elements of \mathcal{A} (together with their complements in X).

If $\mathcal{A}, \mathcal{B} \in \Phi(X)$, then $\mathcal{A} \triangle \mathcal{B}$ is a subproset of $\mathcal{P}_0(X)$ (that is, invariant under $*$). The set of walls, $\mathcal{W}(\mathcal{A}, \mathcal{B})$, separating \mathcal{A} and \mathcal{B} are precisely those which are

contained in $\mathcal{A} \triangle \mathcal{B}$. Thus $\rho(\mathcal{A}, \mathcal{B}) = \frac{1}{2} \#(\mathcal{A} \triangle \mathcal{B})$. This is maximised precisely when \mathcal{A}, \mathcal{B} are distinct and lie in $\iota(X)$. In this case, $\rho(\mathcal{A}, \mathcal{B}) = 2^{\#X-2}$. In other words, the ρ -diameter of $\Phi(X)$ is $2^{\#X-2}$.

On the other hand, clearly any chain of proper subsets of X has length at most $\#X - 1$. Moreover, this is attained (by a chain of sets all of which contain one element of X and exclude another). Therefore the σ -diameter of $\Phi(X)$ is $\#X - 1$. (Recall that σ is metric defined in Subsection 11.6.)

We can say more about σ as follows.

Given $\mathcal{A} \in \Phi(X)$, write $h(\mathcal{A}) = \min\{\#A \mid A \in \mathcal{A}\}$. Thus, $h(\mathcal{A}) = 1$ if and only if $\mathcal{A} \in \iota X$. If $\#X = 2N - 1$ is odd, then $h(\mathcal{A})$ attains the maximum, N , precisely at the central vertex. If $\#X = 2N$ is even, then $h(\mathcal{A})$ attains the maximum, N , precisely when \mathcal{A} lies in the central cube.

Let $\mathcal{A} \in \Phi(X)$. We say that a subset $\mathcal{L} \subseteq \mathcal{A}$ is **flippable** if each element of \mathcal{L} is minimal in \mathcal{A} , and $A \cup B \neq X$ for all $A, B \in \mathcal{L}$. In this case, let $\mathcal{A}[\mathcal{L}] = (\mathcal{A} \setminus \mathcal{L}) \cup \{A^* \mid A \in \mathcal{L}\}$. It is readily checked that $\mathcal{A}[\mathcal{L}] \in \Phi(X)$. It is the result of flipping each element of \mathcal{L} simultaneously, or in any sequence. Clearly, any subset of \mathcal{L} is also flippable, and we get a map $\phi : \mathcal{P}(\mathcal{L}) \rightarrow \Phi(X)$ by setting $\phi(\mathcal{N}) = \mathcal{A}[\mathcal{N}]$. This is a median monomorphism. Its image, $\phi(\mathcal{P}(\mathcal{L}))$, is thus a $(\#\mathcal{L})$ -cube in $\Phi(X)$. In fact, it is a $(\#\mathcal{L})$ -cell: it is convex, since it is precisely the interval $[\mathcal{A}, \mathcal{A}[\mathcal{L}]]$ in $\Phi(X)$.

Let $N = \lfloor \#X/2 \rfloor$. Given $\mathcal{A} \in \Phi(X)$ with $h(\mathcal{A}) < N$, let $\mathcal{L}(\mathcal{A})$ be the set of minimal elements of \mathcal{A} which have size at most N . Then $\mathcal{L}(\mathcal{A})$ is flippable. Set $\mathcal{A}' = \mathcal{A}[\mathcal{L}(\mathcal{A})]$. Then $h(\mathcal{A}') = h(\mathcal{A}) + 1$. By iterating this, we get a sequence, $\mathcal{A} = \mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_m$, with $m = N - h(\mathcal{A})$ and with $h(\mathcal{A}_m) = N$. If $\#X$ is odd, then \mathcal{A}_m is the central vertex of $\Phi(X)$. If $\#X$ is even, then \mathcal{A}_m lies in the central cube of $\Phi(X)$. Thus $N - h(\mathcal{A})$ is the σ -distance of \mathcal{A} from the central vertex or cube. Moreover, the path $\mathcal{A}_0, \dots, \mathcal{A}_m$ is the canonical cube path from \mathcal{A} to \mathcal{A}_m (as defined in Subsection 11.6).

One can also calculate the rank of $\Phi(X)$. We have noted that a wall of $\Phi(X)$ corresponds to a bipartition of X . We say that two bipartitions “cross” if all four intersections arising from these bipartitions are non-empty. This is equivalent to saying that the corresponding pair of walls cross. (Note that the union of two halfspaces of $\Phi(X)$ is a subalgebra thereof, so if this contains every element of X , it must be all of $\Phi(X)$.) Therefore, by Lemma 8.2.1, the rank of $\Phi(X)$ is the maximal cardinality of a family of pairwise crossing bipartitions.

The situation is perhaps more transparent if we consider the even and odd cases separately. Suppose $\#X = 2N$ is even. Note that the collections of all sets arising from a pairwise crossing set of bipartitions is a Sperner family on X . By Sperner’s Lemma (Lemma 2.2.3), it must have at most $\binom{2N}{N}$ elements. Therefore, there are at most $\frac{1}{2} \binom{2N}{N} = \binom{2N-1}{N}$ bipartitions. In fact, this is realised by taking the set of all equal bipartitions. In that case, the resulting walls are precisely those which cross the central cube (which we we have previously noted has this rank). We therefore

see that $\text{rank}(\Phi(X)) = \binom{2N-1}{N}$. (In fact, Sperner’s Lemma is sharp, which means that the central cube is the unique cube of this rank.)

Now suppose that $\#X = 2N - 1$ is odd. In this case, we need a variation of Sperner’s lemma as given in [ErKR]. Given any set of pairwise crossing bipartitions, select the smaller set from each bipartition. Each set of this family has size at most $N - 1$, and any two of them intersect. Now Theorem 1 of [ErKR] implies that such a family has at most $\binom{2N-2}{N-2} = \binom{2N-2}{N}$ elements. (See the Notes to this section.) Again, this is realised. To see this, choose any $a \in X$, and consider the set of all subsets of $X \setminus \{a\}$ of size N . This has $\binom{2N-2}{N}$ elements, and the resulting bipartitions of X all cross. In fact, the corresponding walls all cross a cube of this rank: this cube contains the central element of $\Phi(X)$ and its antipodal vertex is obtained by flipping each of these sets. (There are $\#X$ cubes which arise in this way.) We therefore see that $\text{rank}(\Phi(X)) = \binom{2N-2}{N}$.

We can summarise the two cases together as:

Proposition 11.9.1. $\text{rank}(\Phi(X)) = \binom{\#X-1}{\lceil \#X/2 \rceil}$.

(Here $\lceil \cdot \rceil$ denotes the ceiling function.)

Note that any permutation of X extends to an automorphism of $\Phi(X)$.

Lemma 11.9.2. *Suppose $Q \subseteq \Phi(X)$ is a cell of $\Phi(X)$ which is invariant under all permutations of X . If $\#X$ is odd, then Q is the central element of $\Phi(X)$. If $\#X$ is even, then Q is the central cube of $\Phi(X)$.*

Proof. Suppose $\#X$ is odd. It is easy to see that the central element is the only vertex of $\Phi(X)$ which is invariant under all permutations of X . Consider the gate of this vertex in Q together with the antipodal vertex of Q . These are also invariant under all permutations, hence all three points are equal. In other words, Q is just the central element.

Suppose $\#X$ is even. Let Q_0 be the central cube, and let $F \subseteq Q_0$ be the image of Q under the gate map to Q_0 . Thus, F is a face of Q_0 . Suppose $Q \neq Q_0$. Since $\text{rank}(Q_0) = \text{rank}(\Phi(X))$, F is proper face. Suppose that $\{A, A^*\}$ is an equal bipartition of X . This corresponds to a wall of $\Phi(X)$ crossing Q_0 . Let $\pi : X \rightarrow X$ be any permutation of X with $\pi A = A^*$. Choose any $\mathcal{A} \in F$. By construction of Q_0 , we can suppose that $A \in \mathcal{A}$. Now $A^* \in \pi \mathcal{A} \in F$. This shows that any wall of $\Phi(X)$ which crosses Q_0 also crosses F . This contradicts the fact that F is proper face. Therefore $Q = Q_0$ as required. \square

The above justifies our use of the term “central”. Let $\Delta = \Delta(\Phi(X))$ be the realisation of $\Phi(X)$ as a cube complex. Since the construction is canonical, any permutation of X extends to an automorphism of Δ . Lemma 11.9.2 tells us that any canonically determined point of Δ must be the central vertex, or the centre of the central cube (depending on the parity of $\#X$).

Suppose for example, we equip Δ with the standard euclidean metric, so that each cube is isometric to a euclidean cube with unit side lengths. As explained in

Section 18, this metric is CAT(0). In any complete CAT(0) space, any finite set of points has a canonical centre: the unique point which minimises the maximal distance to any of these points. Applying this to $X \subseteq \Delta$, we see that this is the central point of Δ .

11.10. Subdivisions.

We now consider subdivisions of a discrete median algebra. We have already mentioned subdivisions in Subsection 10.2. If we think of a discrete median algebra, Λ , in terms of its realisation, $\Delta(\Lambda)$, then the process of subdivision cuts each cell of $\Delta(\Lambda)$ into smaller cubes. In this way, Λ can be identified as a subalgebra of a discrete median algebra, Π , namely the set of 0-cells of the subdivision. Going in the opposite direction, we want to characterise the subalgebras, Λ , of Π that can arise in this way. We begin with a formal definition, and then relate it to a more geometrical interpretation.

Definition. A subalgebra $\Lambda \subseteq \Pi$ is *subdividing* if every wall of Π crosses Λ , and every pair of crossing walls in Π also cross in Λ . We also say that Π is a *subdivision* of Λ .

Note that the first statement is equivalent to saying that $\Pi = \text{hull}_{\Pi}(\Lambda)$.

Given any $Q \in \mathcal{C}(\Lambda)$ let $\Sigma(Q) = \text{hull}_{\Pi}(Q)$.

Suppose that $a, b \in \Lambda$ are adjacent in Λ , that is $\{a, b\} \in \mathcal{C}_1(\Lambda)$. Then $\Sigma(\{a, b\}) = [a, b]_{\Pi}$. This has rank 1. (For suppose W, W' are walls of Π crossing $\{a, b\}$. If W, W' cross in Π then they cross in Λ . In particular, there is some $c \in \Lambda$ with $c|_W a$ and $c|_{W'} b$. It follows that $abc \in [a, b]_{\Lambda} \setminus \{a, b\}$, contradicting the fact that a, b are adjacent in Λ .) We see that, $[a, b]_{\Pi}$ is a finite totally ordered set, and therefore median isomorphic to $I_p = \{1, \dots, p\}$ for some $p > 2$.

Now suppose $Q \in \mathcal{C}_n(\Lambda)$. The walls, W_1, \dots, W_n , of Q , each determine a parallel class of 1-cells of Λ which cross that wall. If $\epsilon_i \in \mathcal{C}_1(Q)$ crosses W_i , then ϵ_i determines an interval, $\text{hull}_{\Pi}(\epsilon_i) \cong I_{p_i}$. By Lemma 10.3.5, $\Sigma(Q)$ is naturally isomorphic to $\prod_{i=1}^n \Sigma(\epsilon_i) \cong \prod_{i=1}^n I_{p_i}$. This accords with the more intuitive notion of what we mean by a ‘‘subdivision’’ of Q .

Suppose that $R, S \in \mathcal{C}(\Lambda)$. Now R, S are convex in Λ , and $\Sigma(R), \Sigma(S)$ are respectively, their convex hulls in Π . Therefore, by Proposition 7.4.11, we have $\Sigma(R) \cap \Sigma(S) = \Sigma(R \cap S)$. In this way we can think of $\{\Sigma(Q)\}_{Q \in \mathcal{C}(\Lambda)}$ as a complex of subdivided cubes.

We can interpret this in terms of the construction of Subsection 10.2. We embed Λ into $\Psi(\Lambda) = \{0, 1\}^{\mathcal{W}(\Lambda)}$. For each $W \in \mathcal{W}(\Lambda)$ let $\Upsilon(W)$ be a copy of $\Sigma(\{a, b\}) = [a, b]_{\Pi}$, where $\{a, b\}$ is any 1-cell of Λ crossing W . We can in turn embed $\Psi(\Lambda)$ into $\hat{\Upsilon}(\Psi(\Lambda)) := \prod_{W \in \mathcal{W}(\Lambda)} \Upsilon(W)$. Given $Q \in \mathcal{C}(\Lambda)$, let $\Upsilon(Q) = \text{hull}_{\hat{\Upsilon}(\Psi(\Lambda))}(Q)$. Let $\Upsilon(\Lambda) = \bigcup_{Q \in \mathcal{C}(\Lambda)} \Upsilon(Q)$. By Lemma 10.2.2, $\Upsilon(\Lambda)$ is a subalgebra of $\hat{\Upsilon}(\Lambda)$. For each $Q \in \mathcal{C}(\Lambda)$ we can naturally identify $\Sigma(Q)$ with $\Upsilon(Q)$. We can therefore identify $U := \bigcup_{Q \in \mathcal{C}(\Lambda)} \Sigma(Q) \subseteq \Pi$ with $\Upsilon(\Lambda)$. As noted in Subsection 10.2, we

have $\Sigma(RST) = \Sigma(R)\Sigma(S)\Sigma(T)$ for all $R, S, T \in \mathcal{C}(\Lambda)$, using the natural median structure on $\mathcal{C}(\Lambda)$ described earlier.

Finally we claim that $U = \Pi$. For this we show that U is 1-path-connected and locally convex in Π . By Lemma 11.4.4 it then follows that U is convex, hence all of Π .

The fact that U is 1-path-connected is easy to see, since each of the sets $\Sigma(Q)$ is 1-path-connected, and Λ is intrinsically 1-path-connected.

To see that U is locally convex, let $Q \in \mathcal{C}_2(\Pi)$, and suppose that $a, b \in Q$ are antipodal corners of Q , with $a, b \in U$. We claim that $Q \subseteq U$. (This is stronger than what we need.) To see this, let $c \in Q \setminus \{a, b\}$. Now the two walls of Q cross also in Λ . Choose any $d \in \Lambda$ in the intersection of halfspaces containing c . Then $c = abd$. Let $a \in \Sigma(R)$, $b \in \Sigma(S)$ and let $T = \{d\} = \Sigma(\{d\})$, where $R, S, T \in \mathcal{C}(\Lambda)$. Now $c \in \Sigma(R)\Sigma(S)\Sigma(T) = \Sigma(RST) \subseteq U$ as required.

In summary we have shown:

Lemma 11.10.1. *Let Λ be a subdividing subalgebra of a discrete median algebra, Π . Given $Q \in \mathcal{C}(\Lambda)$, let $\Sigma(Q) = \text{hull}_{\Pi}(Q) \subseteq \Pi$. Then, $\Sigma(Q)$ is isomorphic to a direct product of finite totally ordered sets with its corners at Q . Moreover, $\Pi = \bigcup_{Q \in \mathcal{C}(\Lambda)} \Sigma(Q)$, and $\Sigma(R) \cap \Sigma(S) = \Sigma(R \cap S)$ for all $R, S \in \mathcal{C}(\Lambda)$.*

This accords with the more intuitive notion that Π is a subdivision of Λ .

In the special case of a finite median algebra, Π , there is a unique minimal subdividing subalgebra, which can be described in terms of extreme points.

If $a, b \in \Pi$ are adjacent, we write $W(a, b)$ be the unique wall of Π separating a, b . Given $a \in \Pi$, we refer to the walls $W(a, b)$ for $b \in L(a)$ as the **adjacent** walls to a .

Definition. We say $a \in \Pi$ is **extreme** if its adjacent walls pairwise cross.

In other words, $W(a, b) \pitchfork W(a, c)$ for all distinct $b, c \in L(a)$.

We write $\text{ext}(\Pi)$ for the set of extreme points. We will assume for the following discussion that $\Pi \neq \emptyset$. Then it is easy to see that extreme points exist — any point b which maximises $\rho(a, b)$ for any fixed $a \in \Pi$ will be extreme. In fact:

Lemma 11.10.2. *Suppose Π is finite and non-empty. Let $W_1, \dots, W_n \in \mathcal{W}(\Pi)$ be a set of pairwise crossing walls, and let $\epsilon \in \{+, -\}^n$. Then $\text{ext}(\Pi) \cap O(\epsilon) \neq \emptyset$.*

Proof. By Lemma 11.3.1, there is some n -cell, $Q \subseteq \Pi$, with $Q \pitchfork W_i$ for all i . Thus $Q \cap O(\epsilon)$ consists of a single point, a , and $O(\epsilon) = \omega^{-1}(a)$, where $\omega : \Pi \rightarrow Q$ is the gate map. Choose $b \in O(\epsilon)$ so as to maximise $\rho(a, b)$. We claim that $b \in \text{ext}(\Pi)$. To see this, suppose for contradiction that $c, d \in L(b)$ with $c \neq d$ and $W(b, c) \pitchfork W(b, d)$ (otherwise, a, b, c, d would all lie in distinct orthants of $\{W(b, c), W(b, d)\}$). Then $\rho(a, c)$ and $\rho(a, d)$ cannot both be less than $\rho(a, b)$, so without loss of generality, $\rho(a, c) = \rho(a, b) + 1$. By maximality of $\rho(a, b)$ we see that $c \notin O(\epsilon)$, so $W(b, c) = W_i$ for some i . It follows that $W(b, d) \notin \{W_1, \dots, W_n\}$, and so $W(b, d)$ separates b from

Q . Since $W_i \not\cap W(b, d)$, W_i cannot cross Q , giving a contradiction. Thus $b \in \text{ext}(\Pi)$ as claimed. \square

In particular, this shows that $\langle \text{ext}(\Pi) \rangle$ is a subdividing subalgebra of Π .

Conversely we claim that any subdividing subalgebra, Λ , contains $\text{ext}(\Pi)$. To see this, let $a \in \text{ext}(\Pi)$. For each $b \in L(a)$, let $H(b)$ be the halfspace of $W(a, b)$ containing a . By hypothesis $\Lambda \cap H(b) \neq \emptyset$. Also, if $b, c \in L(a)$ with $b \neq c$, then $\Lambda \cap H(b) \cap H(c) \neq \emptyset$. Thus, $\{\Lambda \cap H(b)\}_{b \in L(a)}$ is a pairwise intersecting finite family of convex sets. Thus, by the Helly Property (Lemma 7.1.1) $\Lambda \cap \bigcap_{b \in L(a)} H(b) \neq \emptyset$. But $\bigcap_{b \in L(a)} H(b) = \{a\}$, and so $a \in \Lambda$ as required.

In summary, this shows:

Proposition 11.10.3. *Any finite median algebra contains a unique minimal subdividing subalgebra, namely that generated by the set of extreme points.*

11.11. Infinite cubes.

We continue this section with a discussion of (possibly infinite) cubes. These can be thought of as universal discrete median algebras (see Lemma 11.11.4). We first return to a brief discussion about hypercubes.

Let X be any set, and let $\mathcal{P}(X)$ be its power set. We view $\mathcal{P}(X)$ as a median algebra, naturally isomorphic to the cube $\{0, 1\}^X$. By Lemma 2.1.3, if $A, B, C \in \mathcal{P}(X)$, then the statement $A.C.B$ is equivalent to $A \cap B \subseteq C \subseteq A \cup B$. Given $A \in \mathcal{P}(X)$, we write $A^* = X \setminus A$. Given $A, B \in \mathcal{P}(X)$, we write $A \triangle B = (A \cup B) \setminus (A \cap B) = (A \setminus B) \sqcup (B \setminus A)$ for the symmetric difference. (Thus, $(\mathcal{P}(X), \Delta, \cap)$ is a boolean ring, as discussed in Subsection 3.4.) Given any $P \in \mathcal{P}(X)$, the involution $[A \mapsto A \triangle P] : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ is a median automorphism. Under the identification with $\{0, 1\}^X$, it corresponds to swapping 0 and 1 in each of the p -coordinates where $p \in P$. Note also that any permutation, $\theta : X \rightarrow X$, of X also induces an automorphism, $[A \mapsto \theta A]$, of $\mathcal{P}(X)$.

Let $\mathcal{T}(X) \subseteq \mathcal{P}(X)$ the set of all finite subsets of X . Then $\mathcal{T}(X)$ is convex in $\mathcal{P}(X)$, hence intrinsically a median algebra. Note that if P is finite, then the map $[A \mapsto A \triangle P]$ preserves $\mathcal{T}(X)$. We refer to a map of this sort as a **coordinate flip**. Given that $A \triangle (A \triangle B) = B$, we see that $\mathcal{T}(X)$ is homogeneous. We also note that any permutation of X induces an automorphism of $\mathcal{T}(X)$. If $Z \subseteq X$, then we can naturally identify $\mathcal{T}(Z)$ as a convex subset of $\mathcal{T}(X)$. Note that $\mathcal{T}(X)$ is naturally isomorphic to the direct product, $\mathcal{T}(Z) \times \mathcal{T}(Z^*)$.

Clearly $\mathcal{T}(X)$ only depends up to isomorphism on the cardinality, $\#X$, of X . Writing $\aleph = \#X$, we can write $\mathcal{T}_\aleph \equiv \mathcal{T}(X)$. Thus $\mathcal{T}_{\aleph+\beth} \cong \mathcal{T}_\aleph \times \mathcal{T}_\beth$ for any cardinals, \aleph and \beth . If \aleph is finite, then \mathcal{T}_\aleph is an \aleph -cube. If \aleph is infinite, then assuming the Axiom of Choice, $\#\mathcal{T}_\aleph = \aleph$. Also $\mathcal{T}_\aleph \cong \mathcal{T}_\aleph \times \mathcal{T}_\aleph \cong \mathcal{T}_\aleph \times \{0, 1\}$.

Definition. We will refer to a median algebra isomorphic to \mathcal{T}_\aleph for a cardinal \aleph as an \aleph -**cube**. A **cube** is an \aleph -cube for some \aleph .

Note that this agrees with previous terminology when \aleph is finite.

Given $a \in X$, let $W_a^+ = \{A \in \mathcal{T}(X) \mid a \in A\}$ and $W_a^- = \{A \in \mathcal{T}(X) \mid a \notin A\}$, and set $W_a = \{W_a^-, W_a^+\}$. Then W_a is a wall of $\mathcal{T}(X)$. (It corresponds to the projection of $\{0, 1\}^X$ to the a -coordinate.) We will see (Lemma 11.11.3) that every wall has this form. Note that if $a, b \in X$ are distinct, then $W_a \pitchfork W_b$ (since $\emptyset, \{a\}, \{b\}, \{a, b\}$ lie in different orthants of $\{W_a, W_b\}$).

Suppose that $Y \subseteq Z \subseteq X$. Let $\mathcal{A}(Y, Z) = \{C \in \mathcal{T}(X) \mid Y \subseteq C \subseteq Z\}$. (In this notation, $W_a^+ = \mathcal{A}(\{a\}, X)$ and $W_a^- = \mathcal{A}(\emptyset, \{a\}^*)$.) Note that $\mathcal{A}(Y, Z) \neq \emptyset$ if and only if Y is finite. If $A, B \in \mathcal{T}(X)$, then (by Lemma 2.1.3) $[A, B] = \mathcal{A}(A \cap B, A \cup B)$. In particular, $[A, B]$ is finite, and so $\mathcal{T}(X)$ is a discrete median algebra.

Lemma 11.11.1. *If $Y \subseteq Z \subseteq X$, then $\mathcal{A}(Y, Z)$ is convex in $\mathcal{T}(X)$.*

Proof. If $A, B \in \mathcal{A}(Y, Z)$ and $C \in [A, B]$, then $Y \subseteq A \cap B \subseteq C \subseteq A \cup B \subseteq Z$, so $C \in \mathcal{A}(Y, Z)$. \square

Note that there is a natural isomorphism, $\mathcal{T}(Z \setminus Y) \longrightarrow \mathcal{A}(Y, Z)$ given by $[A \mapsto A \sqcup Y]$. In particular, if Z is finite, then $\mathcal{A}(Y, Z)$ an n -cube, where $n = \#(Z \setminus Y)$.

It turns out that all convex subsets of $\mathcal{T}(X)$ have this form:

Lemma 11.11.2. *Let $\mathcal{B} \subseteq \mathcal{T}(X)$ be convex. Let $Y = \bigcap \mathcal{B}$ and $Z = \bigcup \mathcal{B}$. Then $\mathcal{B} = \mathcal{A}(Y, Z)$.*

Proof. Clearly $\mathcal{B} \subseteq \mathcal{A}(Y, Z)$. We claim that $\mathcal{A}(Y, Z) \subseteq \mathcal{B}$.

Let $A \in \mathcal{A}(Y, Z)$. If $B \in \mathcal{B} \subseteq \mathcal{A}(Y, Z)$, then $A \triangle B$ is finite. We choose $B \in \mathcal{B}$ with $\#(A \triangle B)$ minimal. We claim that $A \triangle B = \emptyset$. For suppose $a \in A \triangle B$. Since $A, B \in \mathcal{A}(Y, Z)$, we have $Y \subseteq A \cap B$ and $A \cup B \subseteq Z$. Thus, $a \in Z \setminus Y$. If $a \in A \setminus B$, let $B' = B \cup \{a\}$. Since $a \in Z = \bigcup \mathcal{B}$, there is some $C \in \mathcal{B}$, with $a \in C$. Now $B \cap C \subseteq B' \subseteq B \cup C$, so $B \cdot B' \cdot C$, so $B' \in \mathcal{B}$. If $a \in B \setminus A$, let $B' = B \setminus \{a\}$. Since $a \notin Y = \bigcap \mathcal{B}$, there is some $C \in \mathcal{B}$ with $a \notin C$. Again, $B \cap C \subseteq B' \subseteq B \cup C$, so $B \cdot B' \cdot C$, so again $B' \in \mathcal{B}$. Either way, $A \triangle B' = (A \triangle B) \setminus \{a\}$, so $\#(A \triangle B') < \#(A \triangle B)$, contradicting minimality. Thus $A \triangle B = \emptyset$ as claimed. Thus $A = B \in \mathcal{B}$. \square

Since $\mathcal{A}(Y, Z)$ is median isomorphic to $\mathcal{T}(Z \setminus Y)$, it follows that any convex subset of $\mathcal{T}(X)$ is intrinsically a cube.

In particular, we see that every finite convex subset of $\mathcal{T}(X)$ is a cube, hence a cell of $\mathcal{T}(X)$.

Lemma 11.11.3. *Let W be a wall of $\mathcal{T}(X)$. Then $W = W_a$ for some $a \in X$.*

Proof. Let $W = \{\mathcal{B}, \mathcal{B}'\}$. By Lemma 11.11.2, we have $\mathcal{B} = \mathcal{A}(Y, Z)$ and $\mathcal{B}' = \mathcal{A}(Y', Z')$, where $Y \subseteq Z \subseteq X$ and $Y' \subseteq Z' \subseteq X$, and Y, Y' are finite. We can suppose that $\emptyset \in \mathcal{B}'$, so $Y' = \emptyset$ and $Y \neq \emptyset$. Now $Z' \neq X$. Let $a \in (Z')^*$, so $\{a\} \notin \mathcal{B}'$, so $\{a\} \in \mathcal{B}$, so $Y \subseteq \{a\}$, so $Y = \{a\}$. If $b \in \{a\}^*$, then $\{b\} \notin \mathcal{B}$, so $\{b\} \in \mathcal{B}'$, so $b \in Z'$, so $Z = \{a\}^*$. It now follows that $\mathcal{B} = \mathcal{A}(\{a\}, X) = W_a^+$ and $\mathcal{B}' = \mathcal{A}(\emptyset, \{a\}^*) = W_a^-$. Thus $W = W_a$. \square

It now follows that the walls of $\mathcal{T}(X)$ pairwise cross.

We also note that $\mathcal{W}(A, B) = \{W_a \mid a \in A \triangle B\}$, where $\mathcal{W}(A, B) \subseteq \mathcal{W}(\mathcal{T}(X))$ is the set of walls separating A and B . In particular, the combinatorial metric, $\rho_{\mathcal{T}(X)}$, on $\mathcal{T}(X)$, is given by $\rho_{\mathcal{T}(X)}(A, B) = \#(A \triangle B)$.

We next observe that any discrete median algebra can be embedded in a median algebra of this form.

Let Π be a discrete median algebra, and let $X = \mathcal{W}(\Pi)$. Choose any basepoint, $p \in \Pi$. Given any $W \in \mathcal{W}(\Pi)$, direct W so that $p \in W^-$. Given any $a \in \Pi$, let $A(a) = \{W \in \mathcal{W}(\Pi) \mid a \in W^+\} \in \mathcal{P}(X)$. Note that $A(a)$ is precisely the set of walls separating a and p , so $A(a) \in \mathcal{T}(X)$. Now $[a \mapsto A(a)]$ is a median monomorphism, so its image, $\mathcal{A} = \{A(a) \mid a \in \Pi\}$ is a subalgebra of $\mathcal{T}(X)$. (Note that if $a.b.c$ holds in Π , then $A(a) \cap A(b) \subseteq A(c) \subseteq A(a) \cup A(b)$: cf. the discussion of hypercubes in Subsection 10.2.) In fact, \mathcal{A} is 1-path-connected, and $\bigcup \mathcal{A} = X$. This shows:

Lemma 11.11.4. *Any discrete median algebra embeds as a 1-path-connected subalgebra of \mathcal{T}_{\aleph} , where $\aleph = \#\mathcal{W}(\Pi)$.*

We note (given the Axiom of Choice) that if Π is infinite, then $\#\mathcal{W}(\Pi) = \#\Pi$.

Also, by Proposition 11.8.2, one can say in addition that the subalgebra is a retract of \mathcal{T}_{\aleph} .

We can now give a number of characterisations of \mathcal{T}_{\aleph} .

Lemma 11.11.5. *Let Θ be a discrete median algebra. Then the following are equivalent.*

(U1): Θ is a cube.

(U2): The walls of Θ pairwise cross.

(U3): If $a, b, c \in \Theta$ are distinct with b, c adjacent to a , then $\{a, b, c\}$ lies in a 2-cell of Θ .

(U4): Every interval in Θ is a (finite) cube.

(U5): Every finite convex subset of Θ is a (finite) cube.

Proof. We have seen that (U1) implies (U2), (U4) and (U5). Also (U3) follows from (U4) applied to the interval $[b, c]$. Since intervals are convex, (U5) implies (U4).

To see that (U2) \Rightarrow (U3), let a, b, c be as in (U3). We can suppose that $b \neq c$. Let $W_1 \in \mathcal{W}(\Theta)$ be the wall separating a and b , and let $W_2 \in \mathcal{W}(\Theta)$ be the wall separating a and c . Now $W_1 \pitchfork W_2$. We can suppose that $a \in W_1^- \cap W_2^-$. Let $e \in W_1^+ \cap W_2^+$, and let $d = bce$. Then $\{a, b, c, d\}$ is a 2-cell of Θ .

We finally show that (U3) \Rightarrow (U1). By Lemma 11.11.4, we can embed Θ as a 1-path-connected subalgebra of $\mathcal{T} = \mathcal{T}_{\aleph}$. By Lemma 11.11.2 and the subsequent remark, it is enough to show that Θ is convex in \mathcal{T} . By Lemma 11.4.4, it is in turn enough to show that it is locally convex. To this end, suppose that $Q \subseteq \mathcal{T}$ is a 2-cell with $\#(Q \cap \Theta) = 3$. We have $Q \cap \Theta = \{a, b, c\}$ where $a \in [b, c]$. By (U3), there is some $d \in \Theta$ such that $\{a, b, c, d\}$ is a 2-cell of Θ , hence also of \mathcal{T} . There

can only be one such 2-cell containing $\{a, b, c\}$ (see Lemma 10.3.6). It now follows that $d \in Q$, giving a contradiction. Thus, Θ is convex in \mathcal{T} as claimed. \square

We can say a bit more about the structure of such median algebras.

Write $\text{Aut}(\mathcal{T}(X))$ for the automorphism group of $\mathcal{T}(X)$. We have noted that $\text{Aut}(\mathcal{T}(X))$ contains all coordinate flips and permutations of X . In fact, the set of all coordinate flips (on finite subsets of X) form a normal subgroup, $N \triangleleft \text{Aut}(\mathcal{T}(X))$. We claim that the quotient, $\text{Aut}(\mathcal{T}(X))/N$, can be naturally identified with the permutation group of X .

To see this, let $\phi \in \text{Aut}(\mathcal{T}(X))$. After composing with a coordinate flip (on $\phi(\emptyset)$) we can assume that $\phi(\emptyset) = \emptyset$. Now ϕ permutes the walls of $\mathcal{T}(X)$. Therefore, by Lemma 11.11.3, there is some permutation, $\theta : X \rightarrow X$, such that ϕ sends the wall W_a to $W_{\theta a}$, for all $a \in X$. Therefore, after composing again by the automorphism of $\mathcal{T}(X)$ induced by θ , we can suppose that ϕ fixes $\mathcal{W}(\mathcal{T}(X))$. Now, given $A \in \mathcal{T}(X)$, we have $\mathcal{W}(\emptyset, A) = \{W_a \mid a \in A\}$. Therefore, since ϕ fixes \emptyset and $\mathcal{W}(\emptyset, A)$, it must fix A . In other words, ϕ is the identity. This shows that every element of $\text{Aut}(\mathcal{T}(X))$ is a product of a coordinate flip and a permutation of X .

Another observation is that $\text{Aut}(\mathcal{T}(X))$ acts transitively on the set of n -cells of $\mathcal{T}(X)$, for all $n \in \mathbb{N}$. To see this, recall that by Lemma 11.11.2, any n -cell, Q , of $\mathcal{T}(X)$ has the form $\mathcal{A}(Y, Z)$, where $Y \subseteq Z \subseteq X$, and Y, Z are finite. After applying the coordinate flip corresponding to Y , we can suppose that $Y = \emptyset$. Now there is a permutation, $\theta : X \rightarrow X$, sending Z to a fixed subset, say Z_n , of X of cardinality n . The induced automorphism of $\mathcal{T}(X)$ now sends Q to the n -cell $\mathcal{A}(\emptyset, Z_n)$.

For future reference, we make the following definitions. Let Π be a discrete median algebra.

Definition. An \aleph -*cell* of Π is a convex subset median isomorphic to \mathcal{T}_\aleph .

If \aleph is a finite cardinal, then this is consistent with earlier terminology. By default, “cells” are assumed to be finite.

Definition. We say that Π is *small* if it contains no \aleph -cell for any infinite cardinal \aleph .

Note that (assuming the Axiom of Choice) this is equivalent to saying that Π contains no \aleph_0 -cell. It is also equivalent to saying that Π contains no increasing union of (finite) cells.

Clearly if $\text{rank}(\Pi) < \infty$, then it is small (by Lemma 11.3.2). More generally, if Π has no infinite family of pairwise crossing walls, then it is small. The converse is not true however. (Consider the subalgebra of $\mathbb{N} \times \mathcal{T}(\mathbb{N})$ consisting of those pairs, (n, A) , for which $A \subseteq \{0, \dots, n\}$.)

11.12. The Roller boundary.

We next describe the “Roller boundary” of a discrete median algebra, Π .

Recall from Subsection 9.2 that we have a monomorphism, $\eta : \Pi \longrightarrow \mathcal{F}(\mathcal{H}(\Pi))$ defined by setting $\eta(a) = \{H \in \mathcal{H}(\Pi) \mid a \in H\}$. Here $\mathcal{H}(\Pi)$ is the proset of halfspaces of Π , and $\mathcal{F}(\mathcal{H}(\Pi))$ is the median algebra of flows on $\mathcal{H}(\Pi)$.

Lemma 11.12.1. *If Π is discrete, then $\eta(\Pi)$ is convex in $\mathcal{F}(\mathcal{H}(\Pi))$.*

Proof. Let $a, b \in \Pi$, and let R lie in the median interval, $[\eta(a), \eta(b)]_{\mathcal{F}}$ in $\mathcal{F}(\mathcal{H}(\Pi))$. We want to find some $c \in \Pi$ such that $R = \eta(c)$.

From the definition of the median structure on $\mathcal{F}(\mathcal{H}(\Pi))$, we have $R \subseteq \eta(a) \cup \eta(b)$. In other words, if $H \in R$, then $\{a, b\} \cap H \neq \emptyset$, and so $[a, b] \cap H \neq \emptyset$. If $H, H' \in R$, then (from the definition of a flow) $H \cap H' \neq \emptyset$. Therefore, by the Helly Property (Lemma 7.1.1) we have $[a, b] \cap H \cap H' \neq \emptyset$. Since $[a, b]$ is finite, $\{[a, b] \cap H \mid H \in R\}$ is a finite family of pairwise intersecting convex sets. So by the Helly Property again, we see that $[a, b] \cap \bigcap R \neq \emptyset$. Let $c \in \bigcap R$. Then $R \subseteq \eta(c)$. Since R and $\eta(c)$ are both flows, it follows that $R = \eta(c)$, as we wanted.

In summary, we have shown that $[\eta(a), \eta(b)]_{\mathcal{F}} \subseteq \eta(\Pi)$, so $\eta(\Pi)$ is convex. \square

We remark that this can also be understood in terms of the construction of Example (Ex3.7) of Subsection 3.4: in fact $\eta(\Pi)$ is a \sim -class, where two points are \sim -related if the interval between them is finite. We say a bit more about this at the end of this subsection.

We write $\partial_R \Pi = \mathcal{F}(\mathcal{H}(\Pi)) \setminus \eta(\Pi)$. We refer to $\partial_R \Pi$ as the ***Roller boundary*** of Π . We mention a few simple examples.

(Ex11.1): If Π is finite, then by Lemma 9.2.2, $\partial_R \Pi = \emptyset$.

(Ex11.2): Suppose that Π has rank 1. Then the adjacency graph, $\Gamma(\Pi)$, is a simplicial tree. We discussed this case in Example (Ex9.4) of Subsection 9.1. There is a natural boundary, $\partial \Gamma(\Pi)$, which can be defined in many equivalent ways. (See Subsection 15.2 for some discussion in the context of \mathbb{R} -trees.) Here a flow can be thought of as assigning a direction to each edge of $\Gamma(\Pi)$ such that there are no sources. There is at most one sink. If $a \in \Pi$ is a sink, then the flow must be $\eta(a)$. If there is no sink then the flow converges on some unique point in $\partial \Gamma(\Pi)$. Conversely, every point $\partial \Gamma(\Pi)$ has a unique flow converging on it. In this way, we can naturally identify $\partial_R \Pi$ with $\partial \Gamma(\Pi)$.

(Ex11.3): As a special case of (Ex11.2), the integers, \mathbb{Z} , naturally form a rank-1 discrete median algebra with $\Gamma(\mathbb{Z})$ identified with the real line, \mathbb{R} . In this case, we can write $\partial \mathbb{R} = \{-\infty, +\infty\}$. So $\partial_R \mathbb{Z} = \{-\infty, +\infty\}$.

(Ex11.4): If Π, Π' are discrete median algebras, one can check that there is a natural isomorphism $\mathcal{F}(\mathcal{H}(\Pi \times \Pi')) \cong \mathcal{F}(\mathcal{H}(\Pi)) \times \mathcal{F}(\mathcal{H}(\Pi'))$. (See Example (Ex9.2) of Subsection 9.1, given that $\mathcal{H}(\Pi \times \Pi') \cong \mathcal{H}(\Pi) \sqcup \mathcal{H}(\Pi')$.) Thus $\partial_R(\Pi \times \Pi') = (\partial_R \Pi \times \Pi') \cup (\Pi \times \partial_R \Pi')$. For example, $\partial_R(\mathbb{Z}^2)$ consists of a “square” with four corners, $\{(-\infty, \infty), (-\infty, -\infty), (\infty, -\infty), (\infty, \infty)\}$, together with a copy of \mathbb{Z} inserted

into each of its four sides.

(Ex11.5): Let X be any set, and let $\Theta = \mathcal{T}(X)$ be the cube as defined in the previous subsection. Then $\mathcal{H}(\Theta) = \{W_a^+ \mid a \in X\} \sqcup \{W_a^- \mid a \in X\}$. Given $Y \in \mathcal{P}(X)$, let $R(Y) = \{W_a^+ \mid a \in Y\} \sqcup \{W_a^- \mid a \notin Y\} \subseteq \mathcal{H}(\Theta)$. Since the walls of Θ pairwise cross, this is a flow, i.e. $R(Y) \in \mathcal{F}(\mathcal{H}(\Theta))$. Conversely, if $R \in \mathcal{F}(\mathcal{H}(\Theta))$, let $Y(R) = \{a \in X \mid W_a^+ \in R\} \in \mathcal{P}(X)$. We see that the maps $[Y \mapsto R(Y)]$ and $[R \mapsto Y(R)]$ give us inverse bijections between $\mathcal{P}(X)$ and $\mathcal{F}(\mathcal{H}(\Theta))$. If $A \in \Theta = \mathcal{T}(X)$, then $A \in W_a^+ \Leftrightarrow a \in A$ and $A \in W_a^- \Leftrightarrow a \notin A$. Thus, $\eta(A) = \{A \in \mathcal{H}(\Theta) \mid A \in H\} = R(Y)$. It follows that $\eta(\Theta) = \{R(A) \mid A \in \mathcal{T}(X)\}$, which is identified with $\mathcal{T}(X)$ under the above bijections. In this way we can naturally identify $\partial_R \Theta$ with $\mathcal{P}(X) \setminus \mathcal{T}(X)$; that is, the set of infinite subsets of X .

There are several equivalent ways of describing the Roller boundary. Here is another way of interpreting the above construction. A variation on this will be described in Subsection 12.6: see Lemma 12.6.2.

Let Π be a discrete median algebra. Recall that we can identify Π as a subalgebra of the hypercube $\Psi = \Psi(\Pi) := \prod \mathcal{W}(\Pi)$. A wall $W \in \mathcal{W}(\Pi)$ determines a wall of Ψ . Given $p \in \Psi$, we write $H(W, p) \in \mathcal{H}(\Psi)$ for the corresponding halfspace of Ψ which contains p . By construction, $\Pi \cap H(W, p) \in \mathcal{H}(\Pi)$ is the halfspace of Π determined by the W -coordinate of p in Ψ . Note that $\bigcap_{W \in \mathcal{W}(\Pi)} H(W, p) = \{p\}$. We equip Ψ with the product topology, which is compact by Tychonoff's Theorem. A neighbourhood base of p in Ψ is given by the family $(\bigcap_{W \in \mathcal{V}} H(W, p))_{\mathcal{V}}$ as \mathcal{V} ranges over all finite subsets of $\mathcal{W}(\Pi)$.

We can view Ψ as the set of all $*$ -transversals to $\mathcal{H}(\Pi)$. In this way, $\mathcal{F}(\mathcal{H}(\Pi)) \subseteq \Psi$. By definition, if $p \in \Psi$, then $p \in \mathcal{F}(\mathcal{H}(\Pi))$ if and only if $\Pi \cap H(W, p) \cap H(W', p) \neq \emptyset$ for all $W, W' \in \mathcal{W}(\Pi)$. This is a closed property, so $\mathcal{F}(\mathcal{H}(\Pi))$ is a closed subset of Ψ . Under the above identification, we have $\Pi \subseteq \mathcal{F}(\mathcal{H}(\Pi))$.

Now if $p \in \mathcal{F}(\mathcal{H}(\Pi)) \subseteq \Psi$, and $\mathcal{V} \subseteq \mathcal{W}(\Pi)$ is finite, then by the Helly Property (Lemma 7.1.1) in Π , we have $\Pi \cap \bigcap_{W \in \mathcal{V}} H(W, p) = \bigcap_{W \in \mathcal{V}} (\Pi \cap H(W, p)) \neq \emptyset$. We therefore see that p lies in the closure of Π in Ψ .

In summary, we have shown:

Lemma 11.12.2. *Let Π be a discrete median algebra. Let $\Psi(\Pi) = \prod \mathcal{W}(\Pi)$, and identify $\Pi \subseteq \mathcal{F}(\mathcal{H}(\Pi)) \subseteq \Psi(\Pi)$ as above. Then $\mathcal{F}(\mathcal{H}(\Pi))$ is the closure, $\bar{\Pi}$, of Π in $\Psi(\Pi)$.*

We can thus identify $\partial_R \Pi$ with $\bar{\Pi} \setminus \Pi$.

For example, consider the case of a cube, Θ , as in Example (Ex11.5) above. There is a natural bijection, $\mathcal{P}(X) \equiv \Psi(\Theta) = \prod \mathcal{W}(\Theta) = \prod_{a \in X} \{W_a^-, W_a^+\}$. (The a -coordinate of $A \in \mathcal{P}(X)$ is W_a^+ if and only if $a \in A$.) In this case, Θ is dense in $\Psi(\Theta)$, and so we see once more that $\partial_R \Theta \equiv \Psi \setminus \Theta$.

Another way to express this is in terms of “Busemann cocycles”. We briefly describe this as follows. Let \mathcal{R} be the set of 1-lipschitz cocycles on Π . That is, functions, $f : \Pi^2 \rightarrow \mathbb{Z}$ satisfying $|f(x, y)| \leq \rho(x, y)$ and $f(x, y) + f(y, z) + f(z, x) = 0$ for all $x, y, z \in \Pi$. Then, \mathcal{R} is compact in the subspace topology of the product topology on \mathbb{Z}^{Π^2} . Define $\zeta : \Pi \rightarrow \mathcal{R}$, by setting $\zeta(a)(x, y) = \rho(a, y) - \rho(a, x)$. This is injective, and we can identify $\partial_R \Pi$ with $\overline{\zeta(\Pi)} \setminus \zeta(\Pi)$. (Such cocycles are generally termed “Busemann cocycles” with respect to the metric, ρ .) To relate this to the above description, note that we have a map $\xi : \mathcal{F}(\mathcal{H}(\Pi)) \rightarrow \mathcal{R}$, given by setting $\xi(R)(x, y) = \#(R \cap \mathcal{H}(x, y)) - \#(R \cap \mathcal{H}(y, x))$, where $\mathcal{H}(x, y) = \{\mathcal{H}(\Pi) \mid x \in H^*, y \in H\}$. In this way, $\zeta = \xi \circ \eta$. We leave the details as an exercise.

In the case of small discrete median algebras there is yet another, perhaps more intuitive, description of a Roller boundary given by Proposition 11.12.8 below. For the moment, Π can be any discrete median algebra.

Recall that $\mathcal{H}(\Pi)$ is the set of halfspaces of Π . Given $a, b \in \Pi$, let $\mathcal{H}(a, b) = \{H \in \mathcal{H}(\Pi) \mid a \in H^*, b \in H\}$. Thus $\mathcal{H}(a, b) \subseteq \mathcal{H}(a, c) \cup \mathcal{H}(c, b)$ for any $c \in \Pi$. Also $c \in [a, b]$ if and only if $\mathcal{H}(a, c) \sqcup \mathcal{H}(c, b) = \emptyset$.

Recall that a sequence $\underline{a} = (a_i)_{i \in \mathbb{N}}$ is **monotone** if $a_i.a_j.a_k$ holds whenever $i \leq j \leq k$. Note that \underline{a} is eventually constant if and only if it is bounded, i.e. there is some $b \in \Pi$ such that $a_0.a_i.b$ for all i . On the other hand, if the a_i are all distinct, then \underline{a} is unbounded. We write \mathcal{S} for the set of all monotone sequences in Π .

Suppose $\underline{a} \in \mathcal{S}$. Then \underline{a} can cross any given wall, $W \in \mathcal{W}(\Pi)$, at most once. In particular, it must eventually lie in either W^- or W^+ . Let $R = R(\underline{a}) \subseteq \mathcal{H}(\Pi)$ be the set of halfspaces, $H \in \mathcal{H}(\Pi)$, such that \underline{a} eventually lies in H . Clearly R is a flow on $\mathcal{H}(\Pi)$, that is, $R \in \mathcal{F}(\mathcal{H}(\Pi))$. We say that \underline{a} **converges** on R .

Recall that we have an injective map, $\eta : \Pi \rightarrow \mathcal{F}(\mathcal{H}(\Pi))$. Note that $R \in \eta(\Pi)$ if and only if $\bigcap R \neq \emptyset$, in which case, $\bigcap R = \{a\}$ and $R = \eta(a)$. Let $\partial_R \Pi = \mathcal{F}(\mathcal{H}(\Pi)) \setminus \eta(\Pi)$ be the Roller boundary.

Lemma 11.12.3. *Let $\underline{a} \in \mathcal{S}$ and let $R = R(\underline{a})$. Then $R \in \eta(\Pi)$ if and only if \underline{a} is bounded.*

Proof. If \underline{a} is bounded, then it is eventually constant and equal to some $a \in \Pi$. In this case, $R = \eta(a)$.

Conversely, suppose that $R = \eta(a)$ for some $a \in \Pi$. We claim that $a_i \in [a_0, a]$ for all i . For if not, there is some $H \in \mathcal{H}(\Pi)$, with $a_0, a \in H$ and $a_i \in H^*$. Since $R = \eta(a)$, we have $H \in R$. Therefore, there is some $j \geq i$ with $a_j \in H$. But $a_i \in [a_0, a_j] \subseteq H$, giving a contradiction. Thus, \underline{a} is contained in $[a_0, a]$, hence bounded (and is eventually a). \square

Note that if \underline{b} is any subsequence of \underline{a} , then $R(\underline{a}) = R(\underline{b})$. This leads to the following definition.

We let \sim be the smallest equivalence relation on \mathcal{S} , such that any $\underline{a} \in \mathcal{S}$ is \sim -related to any subsequence of \underline{a} . Note that if $\underline{a} \sim \underline{b}$, then \underline{a} is bounded if and

only if \underline{b} is, in which case, they eventually stabilise on the same element of Π . Also by the above observation, if $\underline{a} \sim \underline{b}$, then $R(\underline{a}) = R(\underline{b})$.

We claim that the converse is also true. To see this, we introduce the following notation.

Given $R \in \mathcal{F}(\mathcal{H}(\Pi))$ and $a, b \in \Pi$, we write $a \leq_R b$ to mean that $\mathcal{H}(a, b) \subseteq R$. If $a \leq_R c$ and $c \leq_R b$, then $a \leq_R b$ and $a.c.b$. (The former statement follows since $\mathcal{H}(a, b) \subseteq \mathcal{H}(a, c) \cup \mathcal{H}(c, b)$. For the latter, if $a.c.b$ fails, let H be a halfspace with $a, b \in H^*$ and $c \in H$. Then $H \in \mathcal{H}(a, c) \subseteq R$, and $H^* \in \mathcal{H}(c, b) \subseteq R$, giving a contradiction.)

Lemma 11.12.4. *Let $\underline{a} \in \mathcal{S}$ and $R = R(\underline{a})$. Then $a_i \leq_R a_j$ whenever $i \leq j$.*

Proof. Suppose not. There is some $H \in \mathcal{H}(a_i, a_j)$ with $H^* \in R$. Thus $a_k \in H^*$ for some $k \geq j$. But $a_j \in [a_i, a_k] \subseteq H^*$, contradicting $H \in \mathcal{H}(a_i, a_j)$. \square

Lemma 11.12.5. *Let $\underline{a} \in \mathcal{S}$ and $R = R(\underline{a})$. Given any $b \in \Pi$, we have $b \leq a_i$ for all sufficiently large i .*

Proof. Note that by Lemma 11.12.4 and transitivity of \leq_R , it is enough to find some i for which $b \leq_R a_i$. Now $\mathcal{H}(a_0, b)$ is finite, and so there is some i such that $a_i \in H$ for all $H \in \mathcal{H}(a_0, b) \cap R$. Now $\mathcal{H}(b, a_i) \subseteq \mathcal{H}(b, a_0) \cup \mathcal{H}(a_0, a_i)$. By Lemma 11.12.4, $\mathcal{H}(a_0, a_i) \subseteq R$.

Suppose, for contradiction that there is some $H \in \mathcal{H}(b, a_i) \setminus R$. Then $H \notin \mathcal{H}(a_0, a_i)$, so $H \in \mathcal{H}(b, a_0)$. Then $H^* \in \mathcal{H}(a_0, b) \cap R$, so $a_i \in H^*$. But this contradicts $H \in \mathcal{H}(b, a_i)$. We have shown that $\mathcal{H}(b, a_i) \subseteq R$, in other words $b \leq_R a_i$ as required. \square

Now suppose that $\underline{a}, \underline{b} \in \mathcal{S}$ with $R(\underline{a}) = R(\underline{b}) = R$, say. Applying Lemma 11.12.5 to \underline{a}, b_0 , we can find some $i_1 > 0$ with $b_0 \leq_R a_{i_1}$. Applying it again to \underline{b}, a_{i_1} we find some $j_1 > 0$ with $a_{i_1} \leq_R b_{j_1}$. We similarly get $i_2 > i_1$ with $b_{j_1} \leq_R a_{i_2}$. Continuing in this manner we find an infinite sequence, $b_0 \leq_R a_{i_1} \leq_R b_{j_1} \leq_R a_{i_2} \leq_R b_{j_2} \leq_R \dots$ with $(i_n)_n$ and $(j_n)_n$ both strictly increasing. This sequence is monotone, and shares some common subsequence with both \underline{a} and \underline{b} . Therefore $\underline{a} \sim \underline{b}$. We have shown:

Lemma 11.12.6. *If $\underline{a}, \underline{b} \in \mathcal{S}$, then $R(\underline{a}) = R(\underline{b})$ if and only if $\underline{a} \sim \underline{b}$.*

In fact, there is some $\underline{c} \in \mathcal{S}$ which shares a common subsequence with both \underline{a} and \underline{b} .

This shows we have natural injective map $\mathcal{S}/\sim \rightarrow \mathcal{F}(\mathcal{H}(\Pi))$. In general, this is not surjective (as the example of the \aleph_0 -cube \mathcal{T}_{\aleph_0} shows). However, recall that Π is “small” if it contains no infinite increasing sequence of cells.

Lemma 11.12.7. *If Π is small and $R \in \mathcal{F}(\mathcal{H}(\Pi))$, then there is some $\underline{a} \in \mathcal{S}$, with $R = R(\underline{a})$.*

If $R = \eta(a)$, then we could just take \underline{a} to be the constant sequence a . So we might as well take $R \in \partial_R \Pi$.

In this case, we can say more. Given any $a \in \Pi$, there is a canonical monotone cube-path, \underline{a} , with $a_0 = a$ and with $R = R(\underline{a})$. This is essentially the same construction of a canonical cube-path between two points, which we discussed earlier (see Lemma 11.6.2).

To describe this, let $\mathcal{H}_0(a)$ be the set of $H \in \mathcal{H}(\Pi)$ adjacent to a , with $a \in H^*$. In other words, there is some $b \in \Pi$ (adjacent to a) with $\mathcal{H}(a, b) = \{H\}$. Now the elements of $\mathcal{H}_0(a) \cap R$ pairwise cross (since, by definition of a flow, $H_1 \cap H_2 \neq \emptyset$ for all $H_1, H_2 \in R$). Since Π is small, it follows that $\mathcal{H}_0(a) \cap R$ is finite. Therefore, there is some $p = p(a, R) \in \Pi$, with $[a, p]$ a cube, such that $\mathcal{H}(a, p) = \mathcal{H}_0(a) \cap R$. Note that $a <_R p$. We now proceed inductively, setting $a_0 = a$, and $a_{i+1} = p(a_i, R)$. This gives us a monotone sequence, \underline{a} , with $a_0 <_R a_1 <_R a_2 <_R \dots$.

We claim that $R(\underline{a}) = R$. To see this, let $H \in R$. Suppose that $a_0 \in H^*$. Let $H_1 > H_2 > \dots > H_n = H$ be a chain in $\mathcal{H}(\Pi)$, with $a_0 \in H_1^*$, and with n maximal. (Recall that $>$ here means proper superset.) Since $H_n \in R$, $H_m \in R$ for all m . In particular, $H_1 \in R$. But also $H_1 \in \mathcal{H}_0(a)$. (For otherwise, there would be some $H_0 \in \mathcal{H}(\Pi)$, with $a \in H_0^*$ with $H_0 > H_1$, contradicting maximality.) By construction, $a_1 = p(a_0, R) \in H_1$. Moreover, the maximal length of a chain of halfspaces between a_1 and H is at most $n - 1$. Continuing inductively, we see that $a_n \in H$. It then follows that $a_i \in H$ for all $i \geq n$. Therefore $H \in R(\underline{a})$. In other words, $R \subseteq R(\underline{a})$. Since these are both flows, $R = R(\underline{a})$, as claimed.

This proves Lemma 11.12.7.

In summary we have shown:

Proposition 11.12.8. *If Π is a small discrete median algebra, then its Roller boundary, $\partial_R \Pi$, can be naturally identified with the set of \sim -classes of unbounded monotone sequences in Π . Moreover, given any basepoint in Π , there is a canonical monotone cube-path converging on any given element of $\partial_R \Pi$.*

Under a slightly stronger assumption, one can say a bit more about the structure of the Roller boundary. We say that Π is **subinfinite-rank** if every set of pairwise crossing walls of Π is finite. This is the same as saying that the proset, $\mathcal{H}(\Pi)$, is subinfinite-rank, in the terminology introduced in Subsection 9.5. Note that finite-rank implies subinfinite-rank, which, in turn, implies small. (However, small does not imply subinfinite rank — as noted at the end of Subsection 11.11.)

As in Example (Ex3.7) of Subsection 3.4, we define an equivalence relation, \sim , on $\mathcal{F}(\mathcal{H}(\Pi))$ by setting $R \sim S$ if the interval $[R, S]$ is finite (or equivalently, $R \triangle S$ is finite). We write $\hat{\mathcal{F}} = \hat{\mathcal{F}}(\mathcal{H}(\Pi)) = \mathcal{F}(\mathcal{H}(\Pi))/\sim$ for the quotient median algebra. This was discussed in a more general context in Subsection 9.5. Note that, in the terminology there, the proset $\mathcal{H}(\Pi)$ is boundless. (In fact, if $H_0 > H_1 > H_2 > \dots$ is any downward sequence in $\mathcal{H}(\Pi)$ then $\bigcap_{i=0}^{\infty} H_i = \emptyset$.) Therefore, an immediate consequence of Proposition 9.5.4 is:

Proposition 11.12.9. *If Π is subinfinite-rank, then the quotient median algebra $\hat{\mathcal{F}}$ is discrete.*

We write $\hat{\rho}$ for the combinatorial metric on $\hat{\mathcal{F}}$.

Recall that in Subsection 9.4, we defined subalgebras, $\mathcal{F}_0 \leq \mathcal{F}_1 \leq \mathcal{F}(\mathcal{H}(\Pi))$. In the present set-up we have $\eta\Pi = \mathcal{F}_0$ (essentially by definition). Now \mathcal{F}_1 is a \sim -class, and any two elements of $\eta\Pi$ are \sim -related, and so we get $\eta\Pi = \mathcal{F}_0 = \mathcal{F}_1 \in \hat{\mathcal{F}}$.

Suppose now that $\text{rank}(\Pi)$ is finite.

Proposition 11.12.10. *If $[R] \in \hat{\mathcal{F}}$, then $\text{rank}([R]) + \hat{\rho}(\eta\Pi, [R]) \leq \text{rank}(\Pi)$.*

Proof. As observed at the end of Subsection 9.5, we have $\hat{\rho}(\eta\Pi, [R]) = \lambda([R])$, where $\lambda : \hat{\mathcal{F}} \rightarrow \mathcal{P}(\hat{\mathcal{L}})$ is the homomorphism defined there. Now $\text{rank}(\Pi)$ is equal to the rank of $\mathcal{H}(\Pi)$ as a proset. The statement therefore follows from Corollary 9.5.10. (In fact, for this we only require that $\hat{\rho}(\eta\Pi, [R]) \leq \lambda([R])$, namely Proposition 9.5.14, which bypasses much of the argument.) \square

In particular, $\hat{\mathcal{F}}$ has diameter at most $2 \text{rank}(\Pi)$.

We can think of the above more geometrically. Let Γ be the adjacency graph of the median algebra $\mathcal{F}(\mathcal{H}(\Pi))$. Any element of $\hat{\mathcal{F}}$ is then the vertex set of some connected component of Γ (essentially by definition). Each such component is a median graph. One such component is the adjacency graph of Π . The remaining components constitute the Roller boundary. The set of all components is naturally the vertex set of a median graph $\hat{\Gamma}$, namely the adjacency graph of $\hat{\mathcal{F}}$. This has a “central” vertex corresponding to Π . The rank of any component of Γ is at most the rank of Π minus the distance of that component from the central vertex in $\hat{\Gamma}$.

11.13. Event structures.

The notion of an “event structure” originally arose in the theory of computation (see the Notes to this section). We note that discrete median algebras with a preferred basepoint can be described in terms of event structures. Here is the definition:

Definition. An *event structure* on a set E consists of a pair of binary relations, \leq and $\#$, on E such that \leq is a partial order, and $\#$ is symmetric and antireflexive (that is, $\neg(\alpha \# \alpha)$, for all α), and such that for all $\alpha, \beta, \gamma \in E$, we have:

(E1): $(\alpha \leq \beta \ \& \ \alpha \# \gamma) \Rightarrow \beta \# \gamma$, and

(E2): $\downarrow\alpha := \{\delta \in E \mid \delta \leq \alpha\}$ is finite.

Note that this implies that $\alpha \leq \beta$ and $\alpha \# \beta$ cannot hold simultaneously.

Definition. A *configuration* on E is a finite subset, $A \subseteq E$, such that:

(C1): if $\alpha, \beta \in A$, then $\neg(\alpha \# \beta)$, and

(C2): if $\alpha \leq \beta$ and $\beta \in A$, then $\alpha \in A$.

We write $\mathcal{D} = \mathcal{D}(E)$ for the set of configurations. It is easily seen that this is a subalgebra of the cube, $\mathcal{T}(E)$ (that is, the median algebra of all finite subsets

of E as discussed in Subsection 11.11). Note that $\emptyset \in \mathcal{D}$ and that $\downarrow\alpha \in \mathcal{D}$ for all $\alpha \in E$.

Event structures are essentially equivalent to pointed discrete median algebras. This equivalence can be described as follows.

Let Π be a discrete median algebra with preferred basepoint, $p \in \Pi$. Let $\mathcal{W} = \mathcal{W}(\Pi)$ be the set of walls of Π . Given $W \in \mathcal{W}$ we write $W = \{W^-, W^+\}$ with the convention that $p \in W^-$. Given $W, W' \in \mathcal{W}$, write $W \leq W'$ to mean that $W^- \subseteq (W')^-$ (or equivalently $(W')^+ \subseteq W^+$), and $W \# W'$ to mean that $W^- \cup (W')^- = \Pi$ (or equivalently $W^+ \cap (W')^+ = \emptyset$). It is easily seen that $(\leq, \#)$ is an event structure on \mathcal{W} .

Given $a \in \Pi$, write $D(a) = \mathcal{W}(p, a) = \{W \in \mathcal{W} \mid a \in W^+\}$. It is easily seen that $D(a) \in \mathcal{D}$. Note also that if $D(a) = D(b)$, then $\mathcal{W}(a, b) = \emptyset$, and so $a = b$. In fact:

Lemma 11.13.1. *If $D \in \mathcal{D}(\mathcal{W})$ then $D = D(a)$ for some (unique) $a \in \Pi$.*

Proof. Note that (by (C1)) the set of halfspaces, W^+ , for $W \in D$ pairwise intersect. Therefore, by the Helly Property (Lemma 7.1.1), we have $\bigcap_{W \in D} W^+ \neq \emptyset$. Choose $a \in \bigcap_{W \in D} W^+$ with $\rho(p, a) = \#D(a)$ minimal. Note that $D \subseteq D(a)$. We claim that $D = D(a)$.

To see this, suppose first that $W_0 \in D(a)$ is adjacent to a , that is $\mathcal{W}(a, b) = \{W_0\}$ for some $b \in \Pi$ adjacent to a . Thus, $b \in W^+$ for all $W \in D \setminus \{W_0\}$. If $W_0 \notin D$, then $b \in \bigcap_{W \in D} W^+$. But $\rho(p, b) = \rho(p, a) - 1$, contradicting the minimality of $\rho(p, a)$. This shows that $W_0 \in D$.

Now if $W \in D(a)$ is any element of $D(a)$, then $W \leq W'$ for some $W' \in D(a)$ adjacent to a . By the previous paragraph, $W' \in D$, and so it follows that by (C2) that $W \in D$. This shows that $D(a) \subseteq D$ as required. \square

In summary, this shows that the map $[a \mapsto D(a)] : \Pi \longrightarrow \mathcal{D}(\mathcal{W}(\Pi))$ is a bijection for any pointed discrete median algebra, Π .

In fact, any event structure arises in this way as follows.

Let $(\leq, \#)$ be an event structure on a set, E . Let $\mathcal{D} = \mathcal{D}(E)$ be the set of configurations, viewed as a subalgebra of $\mathcal{T}(E)$, and let $p = \emptyset \in \mathcal{D}$ be the basepoint. Let $\mathcal{W} = \mathcal{W}(\mathcal{D})$ be the set of walls of \mathcal{D} .

Given $\alpha \in E$, let $W_\alpha^+ = \{A \in \mathcal{T}(E) \mid \alpha \in A\}$, and $W_\alpha^- = \{A \in \mathcal{T}(E) \mid \alpha \notin A\}$. Then $W_\alpha := \{W_\alpha^-, W_\alpha^+\}$ is a wall of $\mathcal{T}(E)$. In fact, by Lemma 11.11.3, $\mathcal{W}(\mathcal{T}(E)) = \{W_\alpha \mid \alpha \in E\}$. Note that $\emptyset \in W_\alpha^- \cap \mathcal{D}$, and $\downarrow\alpha \in W_\alpha^+ \cap \mathcal{D}$. In particular, these sets are non-empty, and so $W_\alpha^{\mathcal{D}} := \{W_\alpha^- \cap \mathcal{D}, W_\alpha^+ \cap \mathcal{D}\}$ is a wall of \mathcal{D} . Indeed, by Lemma 11.11.3, every wall of \mathcal{D} has this form. (Recall that the natural map $\mathcal{W}(\mathcal{T}(E)) \longrightarrow \mathcal{W}(\mathcal{D})$ is surjective.)

In fact, this α is uniquely determined by the wall $W_\alpha^{\mathcal{D}} \in \mathcal{W}$. To see this, suppose $W_\alpha^{\mathcal{D}} = W_\beta^{\mathcal{D}}$. Now $\downarrow\alpha \in W_\alpha^+ \cap \mathcal{D}$, so $\alpha \in W_\beta^{\mathcal{D}}$, so $\beta \in \downarrow\alpha$. In other words, $\beta \leq \alpha$. Similarly, $\alpha \leq \beta$. Since \leq is a partial order, this shows that $\alpha = \beta$ as required.

We have shown that $[\alpha \mapsto W_\alpha^{\mathcal{D}}] : E \longrightarrow \mathcal{W}$ is a bijection. In fact, we claim it is an isomorphism of event structures.

For suppose $\alpha \leq \beta$. If $A \in W_\beta^+ \cap \mathcal{D}$, then $\beta \in A$. Since A is configuration on E , we get $\alpha \in A$, so $A \in W_\alpha^+$. This shows that $W_\beta^+ \cap \mathcal{D} \subseteq W_\alpha^+ \cap \mathcal{D}$, so by definition of \leq on \mathcal{W} , we have $W_\alpha^{\mathcal{D}} \leq W_\beta^{\mathcal{D}}$. Conversely, suppose $W_\alpha^{\mathcal{D}} \leq W_\beta^{\mathcal{D}}$. Then $\downarrow\beta \in W_\beta^- \cap \mathcal{D} \subseteq W_\alpha^+$, so $\alpha \in \downarrow\beta$, so $\alpha \leq \beta$.

Secondly, suppose $\alpha \# \beta$. If $A \in W_\alpha^+ \cap W_\beta^+ \cap \mathcal{D}$, we get the contradiction that $\alpha, \beta \in A$. Thus $W_\alpha^+ \cap W_\beta^+ \cap \mathcal{D} = \emptyset$, so $W_\alpha^{\mathcal{D}} \# W_\beta^{\mathcal{D}}$. Conversely, if $W_\alpha^+ \cap W_\beta^+ \cap \mathcal{D} = \emptyset$, then we cannot have $\alpha, \beta \in A$ for any $A \in \mathcal{D}$, so $\alpha \# \beta$.

This proves the claim. In summary, we have shown:

Proposition 11.13.2. *Let E be a set equipped with an event structure, and let $\mathcal{D} = \mathcal{D}(E)$ be the set of configurations with its structure as a discrete median algebra with basepoint \emptyset . Then the map $[\alpha \mapsto W_\alpha^{\mathcal{D}}] : E \rightarrow \mathcal{W}(\mathcal{D})$ is an isomorphism of event structures.*

This gives the natural bijection between event structures and pointed discrete median algebras which we alluded to earlier.

We can also describe event structures in terms of subalgebras of cubes as follows.

Let E be any set. We refer to a subalgebra \mathcal{A} of $\mathcal{T}(E)$ as a **configuration space** if it satisfies:

(A1): $\emptyset \in \mathcal{A}$, and

(A2): if $\alpha, \beta \in E$, then there is some $A \in \mathcal{A}$ with $\#(A \cap \{\alpha, \beta\}) = 1$.

Note that (setting $\alpha = \beta$), (A2) implies:

(A3): $\bigcup \mathcal{A} = E$.

Note that $\mathcal{D}(E)$ satisfies these conditions. (For example, if $\alpha \neq \beta$, then $W_\alpha^{\mathcal{D}} \neq W_\beta^{\mathcal{D}}$, so choose some $A \in W_\alpha^{\mathcal{D}} \Delta W_\beta^{\mathcal{D}}$.)

In fact, the converse also holds:

Proposition 11.13.3. *A subalgebra $\mathcal{A} \leq \mathcal{T}(E)$ is a configuration space if and only if $\mathcal{A} = \mathcal{D}(E)$ for some (unique) event structure on E .*

Proof. We have already observed the “if” direction. For the converse, suppose $\mathcal{A} \leq \mathcal{T}(E)$, satisfies (A1) and (A2). The definition of the event structure is forced: namely, we define $(\leq, \#)$ on E by setting $\alpha \leq \beta$ to mean that for all $A \in \mathcal{A}$, $\beta \in A \Rightarrow \alpha \in A$ and setting $\alpha \# \beta$ to mean that $\neg(\exists A \in \mathcal{A})(\alpha, \beta \in A)$.

Thus, in the earlier notation, we have

$$\alpha \leq \beta \Leftrightarrow W_\alpha^- \cap W_\beta^+ \cap \mathcal{A} = \emptyset$$

$$\alpha \# \beta \Leftrightarrow W_\alpha^+ \cap W_\beta^+ \cap \mathcal{A} = \emptyset.$$

By (A2), $\alpha \leq \beta$ & $\beta \leq \alpha \Rightarrow \alpha = \beta$, and so \leq is a partial order on E . It follows easily that $(\leq, \#)$ is an event structure on E . Now let $\mathcal{D} = \mathcal{D}(E)$ be the set of configurations on E .

Note that, by (A3), $W_\alpha^+ \cap \mathcal{A} \neq \emptyset$ for all $\alpha \in E$. Also $\emptyset \in W_\alpha^-$, and so by (A1), $W_\alpha^- \cap W_\beta^- \cap \mathcal{A} \neq \emptyset$ for all $\alpha, \beta \in E$.

Now let $\mathcal{D} = \mathcal{D}(E)$. Then $\mathcal{A} \leq \mathcal{D} \leq \mathcal{T}(E)$. We claim that $\mathcal{A} = \mathcal{D}$.

Suppose, for contradiction, that $D \in \mathcal{D} \setminus \mathcal{A}$. By Proposition 8.2.5, there are halfspaces, H, H' of \mathcal{D} , with $D \in H \cap H'$ and $H \cap H' \cap \mathcal{A} = \emptyset$.

Recall that $\mathcal{W}(\mathcal{D}) = \{W_\alpha^{\mathcal{D}} \mid \alpha \in E\}$. Now $W_\alpha^- \cap W_\beta^- \cap \mathcal{A} \neq \emptyset$, and so without loss of generality, we have either:

- (1): $H = W_\alpha^+ \cap \mathcal{D}$ and $H' = W_\beta^+ \cap \mathcal{D}$, or
- (2): $H = W_\alpha^- \cap \mathcal{D}$ and $H' = W_\beta^+ \cap \mathcal{D}$.

In case (1), we have $W_\alpha^+ \cap W_\beta^+ \cap \mathcal{A} = \emptyset$, so $\alpha \# \beta$. But $D \in W_\alpha^+ \cap W_\beta^+ \cap \mathcal{D}$, so $\alpha, \beta \in D$, contradicting the fact that D is a configuration.

In case (2), we have $W_\alpha^- \cap W_\beta^+ \cap \mathcal{A} = \emptyset$, so $\alpha \leq \beta$. Now $D \in W_\beta^+$, so $\beta \in D$, so $\alpha \in D$, so $D \in W_\alpha^+$. But $D \in W_\alpha^-$, again giving a contradiction.

This shows that $\mathcal{A} = \mathcal{D}$ as claimed. \square

We briefly mention a couple of particular cases of event structures.

(1): Given any set, E , with a partial order, \leq , define $\#$ by setting $\alpha \# \beta$ to mean that neither $\alpha \leq \beta$ nor $\beta \leq \alpha$ holds. Provided $\downarrow \alpha$ is finite for all $\alpha \in E$, then $(\leq, \#)$ is an event structure on E . In this case, $\mathcal{D}(E)$ has rank 1, so the adjacency graph, $\Gamma := \Gamma(\mathcal{D}(E))$, is a tree, with edge-set identified with E . Under this identification, $\alpha \leq \beta$ means that either $\alpha = \beta$ or α separates β from the basepoint, $\emptyset \in V(\Gamma)$.

(2): If \leq is empty, we can construct a graph, G , with vertex set, $V(G) = E$, by deeming distinct $\alpha, \beta \in E$ to be adjacent if $\neg(\alpha \# \beta)$. In this case, $\mathcal{D}(E) \subseteq \mathcal{T}(E)$ is the subalgebra described by Example (Ex3.2) of Subsection 3.4. This has basepoint \emptyset . It is the vertex of a cube complex, which is a union of cubes all containing \emptyset . A particular case is where $\#$ is also empty. In this case, $\mathcal{D}(E) = \mathcal{T}(E)$.

12. TOPOLOGICAL MEDIAN ALGEBRAS

Many median algebras come equipped with a natural topology: for example, median metric spaces (Section 13) and cube complexes (Section 17). We consider a number of convexity properties of topological median algebras, as well as dimension and connectedness of such spaces. If one assumes that all intervals are compact, then one can say more. This applies to a large class of examples. In such a case, one can construct a compactification, which reduces to the Roller boundary in the discrete case (Lemma 12.6.2). We finish the section a brief discussion of rank-1 topological median algebras.

12.1. Definition and examples.

We begin with the definition:

Definition. A *topological median algebra* is a hausdorff topological space, M , equipped with the structure of a median algebra, such that the median map, $[(x, y, z) \mapsto xyz] : M^3 \rightarrow M$ is continuous.

Here of course, we give M^3 the product topology.

Clearly this is closed under passing to a subalgebra with the subspace topology. Also the direct product of any family of topological median algebras is a topological median algebra.

Here are a few examples.

(Ex12.1): Any median algebra with the discrete topology.

(Ex12.2): For any set X , we can take the hypercube $\Psi = \mathcal{P}(X) \equiv \{0, 1\}^X$ with the product topology. We also have $\hat{\Delta}(\Psi) \equiv [0, 1]^X$ with the product topology. These are compact.

(Ex12.3): We can instead take the l^1 non-finite metric on $\hat{\Delta}(\Psi)$ (with the standard metric on $[0, 1] \subseteq \mathbb{R}$), and take the induced topology. (Recall that “non-finite” means it takes values in $[0, \infty]$.) Identifying $(\hat{\Delta}(\Psi))^3$ with $\Delta(\Psi^3)$, the l^1 metrics agree, and it is easily seen that the median map is 1-lipschitz. Therefore $\hat{\Delta}(\Psi)$ is also a topological median algebra in the induced topology. This is at least as fine as the product topology, and strictly finer if X is infinite. These spaces are discussed further in Section 17.

(Ex12.4): Let Π be any discrete median algebra. We can embed Π into $\Psi(\Pi) = \{0, 1\}^{\mathcal{W}(\Pi)}$. We can further embed $\Psi(\Pi)$ into $\hat{\Delta}(\Psi(\Pi))$ to construct the realisation, $\Delta(\Pi) = \Delta(\Pi, \Psi(\Pi))$, as described in Subsection 11.2. We can now put either the induced metric topology or the product topology on $\Delta(\Pi)$. In either case, $\Delta(\Pi)$ will be a topological median algebra. We remark that in the induced l^1 metric, $\Delta(\Pi)$ is an example of a “median metric space”, which will be the topic of Section 13.

(Ex12.5): Again, if Π is a discrete median algebra, there is a third topology we can put on $\Delta(\Pi)$, namely the CW topology. Thus, a subset $U \subseteq \Delta(\Pi)$ is open if and only if $U \cap \Delta(Q)$ is open in $\Delta(Q)$ for all $Q \in \mathcal{C}(\Pi)$. (Here $\Delta(Q)$ is a finite-dimensional real cube.) This is at least as fine as the metric topology, and strictly finer if Π is infinite. This case is more subtle. We can identify $(\Delta(\Pi))^3$ with $\Delta(\Pi^3)$ as median algebras. Suppose we put the CW topology on $\Delta(\Pi^3)$. Then again the median map is continuous. (It is enough to observe that if $R, S, T \in \mathcal{C}(\Pi)$, then $\Delta(R \times S \times T) = \Delta(R) \times \Delta(S) \times \Delta(T)$ is a cell of $\Delta(\Pi^3)$, and the median map is continuous there.) However, we need that the map be continuous with respect to the product of the CW topologies on $\Delta(\Pi)$. If Π is countable, then these two topologies agree [Mil]. However in general they do not. In summary, if Π is a countable discrete median algebra, then the realisation, $\Delta(\Pi)$ is a topological median algebra in the CW topology. The general situation for uncountable Π is unclear.

(See the Notes to this Section.)

(Ex12.6): An important class of topological median algebras are median metric spaces: the topic of Section 13. More generally, we have “lipschitz median algebras”, where the median map is assumed to be lipschitz with respect to some metric. Such spaces often arise as asymptotic cones. We will say more about these in Subsections 12.2, 13.4 and 24.3.

(Ex12.7): Let Ω be a proset, as defined in Subsection 9.1. Let T be the direct product $\prod\{a, a^*\}$ as $\{a, a^*\}$ ranges over $\Omega/*$. We can think of T as the set of $*$ -transversals of Ω , which we equip with the product topology. This is a compact topological median algebra. We can identify the set of flows, $\mathcal{F}(\Omega)$, as a closed subalgebra. As such, it is itself a topological median algebra, with the topology of a Stone space (that is a compact totally disconnected space). Such spaces arise in the duality result stated as Theorem 12.5.1 below.

(Ex12.8): Let L be a totally ordered set. As described in Example (Ex2.2) of Subsection 3.4, L is a median algebra with obvious notion of “betweenness”. (In other words, the median of three points, $x \leq y \leq z$, is y .) Given $a \in L$, we write $(-\infty, a) = \{x \in L \mid x < a\}$, for the open initial segment. We similarly define (a, ∞) to be the open final segment. Given $b > a$, let $(a, b) = \{x \in L \mid a < x < b\} = (a, \infty) \cap (-\infty, b)$, for the open interval. Recall that **order topology** on L is defined by taking as base the set of all open intervals, together with all open initial and final segments. (The last are only required when L has a minimum or a maximum respectively.)

We claim that the median map is continuous in the order topology. To see this, suppose $a_1, a_2, a_3 \in L$, with $a_1.a_2.a_3$. Suppose that V is an open set containing a_2 . We choose base elements U_1, U_2, U_3 , with $a_i \in U_i$ as follows. If a_1, a_2, a_3 are all distinct, we take the U_i to be pairwise disjoint. If $a_1 = a_2 \neq a_3$, we take $U_1 = U_2$ to be disjoint from U_3 . We do likewise if $a_1 \neq a_2 = a_3$. If $a_1 = a_2 = a_3$, we take $U_1 = U_2 = U_3$. Moreover, we can always assume that $U_2 \subseteq V$. (It is easily checked that one can always do this.) Note that if $x_i \in U_i$, then $x_1x_2x_3 \in V$. This shows that the median map is continuous as claimed. In fact, the argument shows that the median map is also continuous in any finer topology.

Conversely, any topology on L for which the median is continuous must be at least as fine as the order topology. To see this, suppose $x \in (a, \infty)$. Then the statement that $y \in (a, \infty)$ is equivalent to saying that $axy \neq a$, which is an open property in y . Therefore, (a, ∞) is an open subset of L . The same applies to $(-\infty, a)$. It follows also that if $a < b$, then (a, b) is also open in this topology. Since these form the base for the order topology, the statement follows. In general it may be strictly finer (for example, the discrete topology, as well as the Examples (Ex12.9) and (Ex12.10) below).

Such order topologies give rise to many exotic examples. For example, we could take L be the long line.

Another example would be to take $L = \{0, 1\}^\alpha$ for some infinite ordinal α , and equip it with the lexicographic order. (Of course, this is quite distinct from the hypercube $\{0, 1\}^\alpha$ with its product structure.) In this case, all non-trivial intervals have cardinality $2^{\#\alpha}$.

(Ex12.9): As another illustration of Example (Ex12.8), we could take the “K-topology” on the real line, \mathbb{R} , as defined by Munkres (see [Mun]). Here the base of the topology are all open intervals together with sets of the form $(a, b) \setminus \{1/n \mid n \in \mathbb{N}, n > 0\}$, for $a < b$. This is connected hausdorff, but not regular. With the standard median of betweenness, it is a topological median algebra. However, the interval $[-1, 1]$ is not compact.

(Ex12.10): In a similar vein, let $L \subseteq \mathbb{R}$ be any subset with the induced median, and let $f : L \rightarrow \mathbb{R}$ be any function. We can identify L with its graph: $\{(x, f(x)) \mid x \in L\} \subseteq \mathbb{R}^2$, and take the induced topology. Since the projection to the first coordinate is continuous, this is at least as fine as the order topology. Thus, as described in Example (Ex12.8) above, L is a topological median algebra. If the graph is a closed subset, then it is locally compact. One such example is given by $L = \{0\} \cup \{1/n \mid n \in \mathbb{N}\}$ and $f(0) = 0$, $f(1/2n) = 0$ and $f(1/(2n + 1)) = n$ for $n \in \mathbb{N}$. We mention example this again shortly as an example of a median algebra that is not weakly locally convex.

12.2. Gates and local convexity.

We now start on the general theory. Let M be a topological median algebra.

Given a subset, $A \subseteq M$, we write \bar{A} for the closure of A . It is easily seen that if A is a subalgebra, so is \bar{A} . If A is convex, so is \bar{A} .

Let $a, b \in A$. The property that $abx = x$ is closed in x , and so $[a, b] \subseteq M$ is closed. The gate map $[x \mapsto abx] : M \rightarrow [a, b]$ is continuous. So, for example, if M is connected, then so is $[a, b]$, and it follows that any convex subset of M is connected.

Lemma 12.2.1. *A gated convex subset of a topological median algebra M is closed.*

Proof. Let $C \subseteq M$ be convex and gated. Let $a \in M \setminus C$, and let $b \in C$ be a gate for a in C . Let $\omega : M \rightarrow [a, b]$ be the gate map to $[a, b]$. (That is, $\omega(x) = abx$.) Then $\omega(C) = \{b\}$. Since ω is a continuous, $\omega^{-1}([a, b] \setminus \{b\})$ is an open set containing a and disjoint from C . □

Lemma 12.2.2. *Suppose $C \subseteq M$ is convex, $a \in M$, $c \in C$, and $[a, c] \cap C$ is compact. Then there is a gate for a in C .*

Proof. Let $\mathcal{C} = \{[a, d] \cap C \mid d \in C\}$. Now \mathcal{C} is a family of convex subsets of M . Note that if $d, e \in C$, then $ade \in [a, d] \cap [a, e] \cap C$. Therefore, by the Helly Property

(Lemma 7.1.1), any finite subset of \mathcal{C} has non-empty intersection. At least one element, namely $A := [a, c] \cap C$, of \mathcal{C} is compact. Since intervals are closed, $A \cap B$ is compact for all $B \in \mathcal{C}$. Therefore $D := \bigcap \mathcal{C} \neq \emptyset$. Let $b \in D$. Thus $b \in [a, x]$ for all $x \in C$, and so b is a gate for a . \square

Lemma 12.2.3. *Any non-empty compact convex subset of a topological median algebra is gated. Moreover, the gate map is continuous.*

Proof. The existence of the gate map is an immediate consequence of Lemma 12.2.2. We need to check continuity.

Suppose, for contradiction, that the gate map, ω , is not continuous at some point, $a \in M$. Let $b = \omega a$. This means that there is some open set $U \subseteq M$ with $b \in U$, such that for every open set $V \subseteq M$ with $a \in V$ such that $\omega(V) \setminus U \neq \emptyset$. Now $C \setminus U$ is compact. Let $d \in C \setminus U$ be an accumulation point of the sets $\omega(V) \setminus U$, as V varies over the directed set of open sets containing a . By continuity, we have $a.d.b$. But since $d \in C$, and b is gate, we also have $a.b.d$. Therefore $d = b \in U$, giving a contradiction. \square

We introduce the following properties, which will be satisfied by most of the topological median algebras in which we are interested. (The terminology is not a standard one.)

Definition. M is **weakly locally convex** if for all $a \in M$ and any open neighbourhood, $U \ni a$, there is an open set $V \subseteq U$ with $a \in V$ such that $[x, y] \subseteq U$ for all $x, y \in V$.

Note that this is the same as saying that $J(V) \subseteq U$, where J denotes the join, as defined in Subsection 7.1.

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

Clearly this implies weakly locally convex. We have the following converse:

Lemma 12.2.4. *If M is weakly locally convex and has finite rank, then M is locally convex.*

Proof. Let $\nu = \text{rank}(M) < \infty$. Let $a \in M$, and let $O \neq a$ be an open neighbourhood. Define open sets, $U_i \neq a$, inductively by setting $U_0 = O$, and letting U_{i+1} be such that $J(U_{i+1}) \subseteq U_i$. Thus, $J^i(U_i) \subseteq O$ for all i . By Proposition 8.2.3, $\text{hull}(U_\nu) = J^\nu(U_\nu) \subseteq O$. In other words, $\text{hull}(U_\nu)$ is a convex neighbourhood of a contained in O . \square

We also observe that if M is locally convex, then so is any subalgebra.

Not all topological median algebras are weakly locally convex. For example, the linear median algebra described at the end of Example (Ex12.10) above is rank-1, locally compact and metrisable. However, it is not weakly locally convex: note that $1/2n \rightarrow 0$ in this topology, as $n \rightarrow 0$, but the median interval $[0, 1/2n]$ is unbounded, and so leaves every compact set.

However, there are various conditions which do imply (weak) local convexity.

Lemma 12.2.5. *If M is compact then M is weakly locally convex. If M is connected and locally compact then M is weakly locally convex.*

Proof. Let M be a topological median algebra. Let $U \subseteq M$ be open and $a \in U$. Given any $b \in M$, we have $aab = a$, and so there are open sets, $V_b \ni a$ and $O_b \ni b$ such that if $x, y \in V_b$ and $z \in O_b$, then $xyz \in U$.

Suppose $K \subseteq M$ is compact. We can find a finite set $b_1, \dots, b_n \in K$ such that $K \subseteq \bigcup_{i=1}^n O_{b_i}$. Let $V = \bigcap_{i=1}^n V_{b_i}$. Then if $x, y \in V$ and $z \in K$, then $xyz \in U$. We can suppose that $V \subseteq U$.

If M is compact, take $K = M$. Then $[x, y] \subseteq U$ for all $x, y \in V$ as required.

If M is connected and locally compact, we can suppose that ∂U is compact. Let $K = \partial U$. Then for any $x, y \in V \subseteq U$, $[x, y]$ is connected, and $[x, y] \cap \partial U \subseteq U \cap \partial U = \emptyset$. Therefore $[x, y] \subseteq U$ as required. \square

As another example, suppose that M admits a metric, ρ , inducing the given topology. We put the induced l^1 metric on M^3 . We say that M is ***k-lipschitz*** if the median map, $[(x, y, z) \mapsto xyz] : M^3 \rightarrow M$ is k -lipschitz. In particular, if $a, b \in M$, then the diameter of $[a, b]$ in this metric is at most $k\rho(a, b)$. From this it follows that M is weakly locally convex. We note that a median metric space (to be defined in Section 13) is 1-lipschitz (Lemma 13.2.2). More generally k -lipschitz median algebras arise as certain asymptotic cones (see Subsection 24.3). We will mention a general result about lipschitz median algebras in Subsection 13.4 (Theorem 13.4.1).

The following notion will appear frequently.

Definition. We say that M is ***interval-compact*** if $[a, b]$ is compact for all $a, b \in M$.

We will see some natural conditions under which this holds in Section 13 (see, for example, Lemma 13.2.10). It also applies to realisations of discrete median algebras (that is CAT(0) cube complexes) in either the metric topology or the CW topology.

The following is an immediate consequence of Lemma 12.2.2 above.

Lemma 12.2.6. *If M is interval-compact, then any non-empty closed convex subset of M is gated.*

We therefore have a gate map $\omega : M \rightarrow C$. It is not clear under what conditions this must be continuous. (We give one such condition in Lemma 13.3.1.)

12.3. Rank and dimension.

We now return to the general case.

Definition. A wall $W \in \mathcal{W}(M)$ ***strongly separates*** $a, b \in M$, if there are open neighbourhoods, $U \ni a$ and $V \ni b$, such that $U|_W V$.

In other words (up to swapping $-$ and $+$) a and b lie respectively in the interiors of W^- and W^+ .

Lemma 12.3.1. *Suppose that M is locally convex, and $a, b \in M$ are distinct, then a, b are strongly separated by a wall.*

Proof. Let A and B be disjoint convex neighbourhoods of a and b . By Theorem 8.1.2, A and B are separated by a wall. \square

Given $W \in \mathcal{W}(M)$, let $P(W) = \bar{W}^- \cap \bar{W}^+$. Note that, since \bar{W}^- and \bar{W}^+ are closed and convex, so is $P(W)$.

Lemma 12.3.2. *Suppose M is locally convex, and that $\text{rank}(M) = \nu < \infty$. Let $W \in \mathcal{W}(M)$. Then $\text{rank}(P(W)) \leq \nu - 1$.*

(The statement is vacuously true if $\nu = 0$, since then M is a singleton and $\mathcal{W}(M) = \emptyset$.)

Proof. Let $Q \subseteq P(W)$ be an n -cube. Choose any $a \in Q$, and let e_1, \dots, e_n be the adjacent vertices of Q . By Lemma 12.3.1, there is a wall $W_i \in \mathcal{W}(M)$ strongly separating a and e_i . We can suppose that $a \in W_i^-$. We see that $W_i \pitchfork W_j$ in M for all $i \neq j$ (since this holds in Q).

We also claim that $W \pitchfork W_i$ for all i . To see this, note that we can find $a^\pm \in W_i^- \cap W^\pm$ and $e_i^\pm \in W_i^+ \cap W^\pm$ (close to a and e_i). From this, it follows that the four orthants, $W_i^\mp \cap W^\pm$, are all non-empty, and so $W \pitchfork W_i$ as claimed.

Thus, $\{W, W_1, \dots, W_n\}$ is a set of $n+1$ pairwise crossing walls of M . By Lemma 8.2.1, we have $n+1 \leq \nu$, as required. \square

Definition. The *locally compact dimension* of a topological space is the maximal dimension of a locally compact subset.

For a locally compact topological space, all the standard definitions of dimension are equivalent. (See for example, [En]). For definiteness we could take it to mean the covering dimension. Note that a locally compact space has dimension at most ν if any two distinct points are separated by a closed subset of dimension $\nu - 1$. This leads to an inductive formulation of dimension where the empty set is deemed to have dimension -1 . (For non-locally compact spaces, one would need to separate disjoint closed sets, instead of just points, in order to arrive at the standard definition.)

Lemma 12.3.3. *Let M be a locally convex median algebra of rank at most $\nu < \infty$. Then M has locally compact dimension at most ν .*

Proof. We prove this by induction on ν . The case when $\nu = 0$ is trivial, since M is just a singleton.

Let $F \subseteq M$ be any locally compact subset, and let $a, b \in F$ be distinct. By Lemma 12.3.1, there is a wall $W \in \mathcal{W}(M)$, strongly separating a and b . By Lemma 12.3.2, $\text{rank}(P(W)) \leq \nu - 1$. Let $C = F \cap \bar{W}^-$ and $D = F \cap \bar{W}^+$. Then

$F = C \cup D$ and $C \cap D = F \cap P(W)$. Now $F \cap P(W)$ is locally compact, and so by the inductive hypothesis, $\dim(C \cap D) \leq \nu - 1$. This shows that any two distinct points of F are separated by a closed subset of dimension at most $\nu - 1$. Therefore $\dim(F) \leq \nu$. \square

A certain converse to Lemma 12.3.3 will be a consequence of Lemma 12.4.9. Lemma 12.3.3 finds application to coarse geometry via asymptotic cones, as we discuss in Subsection 24.3.

Lemma 12.3.4. *Let $A \leq M$ be a topologically dense subalgebra of a topological median algebra, M . Then $\text{rank}(A) = \text{rank}(M)$.*

Proof. Certainly, $\text{rank}(A) \leq \text{rank}(M)$.

For the reverse inequality, let $Q \leq M$ be an n -cube. Let $p = 2^{2^{2^n}}$. Given $a \in Q$ and $i \in \mathbb{N}$, define a neighbourhood, U_a^i , of a in M as follows. First choose U_a^0 so that $U_a^0 \cap U_b^0 = \emptyset$ for all distinct $a, b \in Q$. Given $(U_a^i)_{a \in Q}$ choose $(U_a^{i+1})_{a \in Q}$ such that $U_a^{i+1} U_b^{i+1} U_c^{i+1} \subseteq U_{abc}^i$ for all $a, b, c \in Q$. Note that this implies that $U_a^{i+1} \subseteq U_a^i$. Given $a \in Q$ choose any $x_a \in A \cap U_a^{p+1}$, and let $\Pi = \langle \{x_a \mid a \in Q\} \rangle \leq A$. We claim that if $x \in \Pi$, then $x \in U_{\phi(x)}^1$ for a unique $\phi(x) \in Q$. Uniqueness is clear. For existence, note that we can write x as a median expression in the elements x_a which involves applying the median operation at most p times (since, by Proposition 3.3.3, $|\Pi| \leq 2^{2^{\#Q}} = p$). Now $\phi(x)$ is obtained by substituting each x_a with a in this expression, and evaluating the result in Q . Note that $\phi(x_a) = a$. This gives us a surjective map $\phi : \Pi \rightarrow Q$. In fact, this is a homomorphism. To see this, note that if $x, y, z \in \Pi$, then $xyz \in U_{\phi(x)}^1 U_{\phi(y)}^1 U_{\phi(z)}^1 \subseteq U_{\phi(x)\phi(y)\phi(z)}^0$. But also $xyz \in U_{\phi(xyz)}^1 \subseteq U_{\phi(xyz)}^0$, and so $\phi(x)\phi(y)\phi(z) = \phi(xyz)$. Thus $\phi : \Pi \rightarrow Q$ is an epimorphism, and so $n \leq \text{rank}(Q) \leq \text{rank}(\Pi) \leq \text{rank}(A)$, as required. \square

12.4. Totally ordered sets and connectedness.

We now make a slight digression concerning total orders.

Let L be a set with a total order, \leq , and with the standard median of betweenness. Then L is a topological median algebra of rank 1. As observed above (Example (Ex12.8)), the topology is at least as fine as the order topology. If L is compact (in the given topology) then these two topologies must agree. In this case, we see that L has both a minimum and a maximum. (For example, if L had no minimum, then the sets (a, ∞) for $a \in L$ would give an open cover with no finite subcover.)

We note:

Lemma 12.4.1. *Suppose that L is compact. Then L is connected if and only if it has no adjacent pair of points.*

Proof. Suppose that $a, b \in L$ were adjacent with $a < b$. Then $[\min L, a] \sqcup [b, \max L]$ would give a partition of L into two disjoint non-empty closed sets.

Conversely, suppose that $L = A \sqcup B$, where A, B are closed and non-empty. We can suppose that $\max L \in B$. Let $a = \max(A)$. Then $B \cap [a, \max L]$ is closed, hence compact. Let $b = \min(B \cap [a, \max L])$. Then $a < b$, and a, b adjacent in L . \square

Lemma 12.4.2. *Suppose that L is compact connected and metrisable. Then L is isomorphic, as a topological median algebra, to the real interval $[0, 1]$.*

Proof. This is one of the standard topological characterisations of the real interval with endpoints, a, b say: it is a metrisable continuum such that any point other than a, b separates a from b (see for example, Theorem 2-27 of [HoY]). Note also that the betweenness relation is determined by the topology: z lies strictly between x and y if and only if z separates x from y in L . Therefore, any homeomorphism must be a median isomorphism. \square

We remark that we could equivalently substitute a number of other conditions in place of metrisability in the statement of Lemma 11.5.2. For example, we could instead assume L to be separable or second countable.

Now let M be any topological median algebra. Let $a, b \in M$. Recall from Subsection 3.2 that $[a, b]$ is partially ordered by the relation \leq , where $x \leq y$ is equivalent to $a.x.y$ (or to $x.y.b$). A **chain** in $[a, b]$ is a subset totally ordered by \leq . Note that a chain is a subalgebra of $[a, b]$, hence of M . As such it is intrinsically a topological median algebra of the type we have been describing.

Lemma 12.4.3. *The closure of a chain is a chain.*

Proof. This is an immediate consequence of the fact that the statement $axy = x$ is a closed property in x and y . \square

Let \mathcal{L} be the set of chains in $[a, b]$. We order \mathcal{L} by inclusion. Note that an increasing union of chains is a chain. Certainly, $\mathcal{L} \neq \emptyset$, so by Zorn's Lemma, \mathcal{L} contains a maximal element, L . By Lemma 12.4.3, we see that L is closed. In particular, if $[a, b]$ is compact, then so is L .

We next consider connectedness.

Lemma 12.4.4. *A connected topological median algebra has no adjacent pairs.*

Proof. Suppose $a, b \in M$ are adjacent. Then $M = \{x \in M \mid a.b.x\} \sqcup \{x \in M \mid b.a.x\}$ is a partition of M into two disjoint non-empty closed subsets. \square

Conversely we have:

Lemma 12.4.5. *If M is interval-compact and has no adjacent pairs, then M is connected.*

Proof. Let $a, b \in M$, and let L be a maximal chain in $[a, b]_M$. Clearly, $a, b \in L$. Also, L has no adjacent pairs. (For if $c, d \in L$ were adjacent, then we could find

some $x \in [c, d]_M \setminus \{c, d\}$, and then $L \cup \{x\} \subseteq [a, b]_M$ would be a strictly larger chain, giving a contradiction.) Since L is compact, by Lemma 12.4.1 it is connected. It now follows that M is connected. \square

By a **path** in M we mean a continuous map $\gamma : I \rightarrow M$, where $I \subseteq \mathbb{R}$ is connected. An **arc** in M is an injective path.

Definition. A path, $\gamma : I \rightarrow M$, is **monotone** if for all $t \leq u \leq v$ in I , we have $\gamma(t) \cdot \gamma(u) \cdot \gamma(v)$.

In other words, γ is a median homomorphism.

Lemma 12.4.6. *Let M be connected, metrisable and interval-compact. Then any two distinct points of M are connected by a monotone arc.*

Proof. Let $a, b \in M$, and let $L \subseteq [a, b]$ be a maximal chain, as in the proof of Lemma 12.4.5. Again, L has no adjacent pairs, and so by Lemma 12.4.2, there is an isomorphism of topological median algebras, $\gamma : [0, 1] \rightarrow L$. In other words, γ is a monotone arc in M . \square

In summary, if M is metrisable and interval-compact, then M is connected if and only if it is path-connected, and if and only if it has no adjacent pairs.

In this context, we also mention the following fact (pointed out to me by Elia Fioravanti).

Proposition 12.4.7. *Let M be a topological median algebra. Then every continuous map of a sphere S^n into M is homotopic to a constant map.*

In other words, every singular sphere bounds a singular ball. This can be expressed by saying that all the homotopy groups of each path-connected component of M are trivial.

Proof. The proof is based on the following general construction. Let X be a topological space, and choose any $p \in X$. Let $S(X)$ be the smash product of two copies of X ; that is the quotient of X^2 after collapsing the wedge, $W(X) := (X \times \{p\}) \cup (\{p\} \times X) \subseteq X^2$. We can embed X into $S(X)$ via the diagonal $[x \mapsto (x, x)]$.

Let $f : X \rightarrow M$ be any continuous map. The map $[(x, y) \mapsto pxy] : X^2 \rightarrow M$ sends $W(X)$ to p , so we get a continuous map $F : S(X) \rightarrow M$. Identifying $X \subseteq S(X)$ via the diagonal, this restricts to f . In other words, any continuous map of X into M extends to a map of $S(X)$.

Now if X is an n -sphere, S^n , then $S(X)$ is homeomorphic to S^{2n} . The result now follows from the fact that $\pi_n(S^{2n})$ is trivial for $n > 0$. \square

(Note the result is very general: we have only used axiom (M1) of a median algebra here.)

Moving on, we have the following.

Lemma 12.4.8. *Let M be connected, metrisable and interval-compact. If $a, b \in M$ and $\text{rank}([a, b]) = 1$, then $[a, b]$ is homeomorphic to $[0, 1]$ via a median isomorphism.*

Proof. This follows directly from Lemma 12.4.2.

Alternatively, note that a monotone arc as given by Lemma 12.4.6 will be surjective to $[a, b]$ in this case. \square

Lemma 12.4.9. *Let M be connected, metrisable and interval-compact. Suppose $\text{rank}(M) = \nu < \infty$, and let $Q \subseteq M$ be a ν -cube. Then there is a homeomorphism of $[0, 1]^\nu$ onto $\text{hull}(Q)$ which is also a median isomorphism.*

Proof. Let $a \in Q$, and let a_1, \dots, a_ν be the adjacent points of Q . By Lemma 10.3.5, there is a median isomorphism $\phi : \prod_{i=1}^\nu [a, e_i] \rightarrow \text{hull}(Q)$. Note that this map and its inverse were defined using medians, and so both are continuous. Lemma 12.4.8 gives us a monotone isomorphism $\theta_i : [0, 1] \rightarrow [a, e_i]$. These combine to give an isomorphism, $\theta : [0, 1]^\nu \rightarrow \prod_{i=1}^\nu [a, e_i]$. Then $\phi \circ \theta$ is the required map. \square

Note that it follows that M has locally compact dimension at least ν . So if M is also (weakly) locally convex, then by Lemma 12.3.3, it has locally compact dimension exactly ν .

Recall that a (**connected**) **component** of a topological space is a maximal connected subset.

We can also define a relation, \sim , on M by writing $a \sim b$ if we can partition M into two closed subsets, $M = A \sqcup B$, with $a \in A$ and $b \in B$. This is easily seen to be an equivalence relation. A \sim -class is called a **quasicomponent**. Each quasicomponent is a union of components.

Lemma 12.4.10. *Let M be a topological median algebra. Every component of M is convex, and every quasicomponent of M is convex.*

Proof. Let C be a component of M . Let $a, b \in C$, and let $c \in [a, b]$. Write $\omega : M \rightarrow [a, c]$ for the gate map to $[a, c]$. Now $\omega a = a$, $\omega b = c$ and $\omega(C)$ is connected. Therefore c lies in the same component as a , namely C .

Let D be a quasicomponent of M . Let $a, b \in D$, and let $c \in [a, b]$. Write $\omega : M \rightarrow [a, c]$ for the gate map to $[a, c]$. Suppose $c \notin D$. Then we can write $M = A \sqcup B$, with A, B closed and with $a \in A$ and $c \in B$. This gives another partition, $M = \omega^{-1}(A \cap [a, c]) \sqcup \omega^{-1}(B \cap [a, c])$, into disjoint closed sets, giving the contradiction that $a \not\sim b$. \square

(In general, a component need not be a quasicomponent, but I do not have a counterexample in the case of a topological median algebra.)

We also note:

Lemma 12.4.11. *Let M be an interval-compact topological median algebra. Suppose $a, b \in M$ lie in different quasicomponents. Then there is a clopen halfspace of M containing a but not b .*

Proof. As in the proof of Lemma 12.4.5, there is an adjacent pair, c, d , in M with $a.c.d.b$. (Take an adjacent pair in any maximal chain in $[a, b]$.) Let $\omega : M \longrightarrow \{c, d\}$ be the gate map. Then $\omega^{-1}(c)$ is a clopen halfspace as required. \square

Most of our discussion of connectedness has focused on interval-compact spaces. It is natural to ask what one can say without this hypothesis. For example, if M is connected, are any two distinct points of M connected by a monotone arc? (Lemma 12.4.6 tells that this is true in the interval-compact case.) One certainly needs to add some further assumptions, such as metrisability and (weak) local convexity — as illustrated by some of the linear examples (Ex12.8) and (Ex12.9) of Subsection 12.1. It is not clear if this is sufficient. However, it is if we assume in addition that M has finite rank. We finish this subsection with a (somewhat involved) proof of this fact.

Proposition 12.4.12. *Let M be a connected, metrisable, locally convex, finite-rank topological median algebra. Then any two distinct points of M are connected by a monotone arc.*

Note that by Lemma 12.2.4, it is enough to assume that M is weakly locally convex (which is what we actually use). Moreover, one could substitute a number of hypotheses for “metrisable”, for example, by assuming instead that M be second countable.

Proof. The overall argument will proceed by induction on the rank of M : the case of rank 0 being trivial. Also note that the hypotheses are inherited by any convex subset of M .

Given any $x, y \in M$, write $x \leftrightarrow y$ to mean that either $x = y$ or there is a monotone arc from x to y . Note that \leftrightarrow is reflexive and symmetric. Moreover, $x.y.z \ \& \ x \leftrightarrow y \ \& \ y \leftrightarrow z \Rightarrow x \leftrightarrow z$.

We first observe that if x, y are the antipodal elements of some 2-cube, say Q , and $z, w \in [x, y]$, then $z \leftrightarrow w$. To see this, recall by Lemma 10.3.4, that $[x, y] = \text{hull}(Q)$ is intrinsically a direct product, $D_1 \times D_2$, of two convex subsets (intervals) $D_1, D_2 \subseteq M$. Note that $\text{rank}(D_i) < \text{rank}(M)$. Write $z = (z_1, z_2)$ and $w = (w_1, w_2)$, with $z_i, w_i \in D_i$. By the inductive hypothesis, $z_i \leftrightarrow w_i$ in D_i . We see that $z \leftrightarrow (z_1, w_2) \leftrightarrow w$, so $z \leftrightarrow w$ as claimed.

Now fix distinct $p, q \in M$. Our eventual aim is to show that $p \leftrightarrow q$. To simplify notation we can assume that $M = [p, q]$.

Recall that M is a distributive lattice with $x \wedge y = pxy$ and $x \vee y = qxy$. Given $A, B \subseteq M$, write $A < B$ to mean that $x < y$ for all $x \in A$ and all $y \in B$. We abbreviate $\{a\} < B$ to $a < B$ etc.

Given $x, y \in M$, write $x \sim y$ to mean that neither $x < y$ nor $y < x$ holds. If $x \sim y$ and $x \neq y$, then x, y are the antipodal elements of 2-cube (with remaining elements $x \wedge y$ and $x \vee y$), and so by the earlier observation, if $z, w \in [x, y]$ then $z \leftrightarrow w$. Of course, the same holds if $x = y$.

Let \approx be the transitive closure of \sim . Write $[x] \subseteq M$ for the \approx -class of $x \in M$, and let $\mathcal{A} = \{[x] \mid x \in M\}$. Note that if $A, B \in \mathcal{A}$, and there exist $x \in A$ and $y \in B$ with $x < y$, then $A < B$. (For if $z \sim y$, then $z \not\sim x$. We cannot have $z < x$, otherwise $z < x < y$ would contradict $z \sim y$. Therefore $x < z$. Since \approx is the transitive closure of \sim , we therefore have $x < z$ for all $z \in B$. Similarly, we now have $w < z$ for all $w \in A$.) This shows that $<$ is a strict total order in \mathcal{A} . It has minimum $\{p\}$, and maximum $\{q\}$. Let $S = \{a \in M \mid [a] = \{a\}\}$. Thus $\mathcal{A}_S := \{\{a\} \mid a \in S\} \subseteq \mathcal{A}$ is the set of singleton elements of \mathcal{A} .

Suppose that $a \in S$. Then a separates p from q : since we have $M = [p, a] \cup [a, q]$ and $[p, a] \cap [a, q] = \{a\}$. Note also that $[p, a] = \bigcup\{A \in \mathcal{A} \mid \{a\} \geq A\}$, and that $[a, q] = \bigcup\{A \in \mathcal{A} \mid \{a\} \leq A\}$.

Suppose that $A \in \mathcal{A} \setminus \mathcal{A}_S$. Now A is open. (For if $x \in A$, there is some $y \in A \setminus \{x\}$ with $y \sim x$, and then $z \sim y$ for all $z \in M$ sufficiently close to x .) In fact, we claim that the closure, \bar{A} , of A is equal to $\bar{A} = A \cup \{a, b\}$, with $a, b \in S$ and $a < A < b$. To see this, suppose $a \in \bar{A} \setminus A$. Now $[a]$ cannot be open, so it must be equal to $\{a\}$. In other words $a \in S$. Either $a < A$ or $A < a$. Suppose $a < A$. There can be at most one such a . For if $a' \in \bar{A}$ with $a' < A$, then $a \not\sim a'$ so (up to swapping a, a'), we can assume $a' < a$, so $a' \notin [a, q] \supseteq A$. Since $[a, q]$ is closed, this contradicts $a' \in \bar{A}$. In fact, such an a must exist. For suppose not. Choosing any $x \in A$, we see that $[p, x] \cap A = [p, x] \cap \bar{A}$ is both open and closed in $[p, x]$. Since $[p, x]$ is connected, we get $[p, x] \subseteq A$. Since $\{p\} \in \mathcal{A}$, this give the the contradiction that $A = \{p\} = \{x\} \in \mathcal{A}_S$. Similarly, there is exactly one element $b \in \bar{A} \setminus A$ with $b > A$. Therefore $\bar{A} = A \cup \{a, b\}$, as claimed. Note that if $x, y \in A$ then $x \wedge y = axy$ (since $axy \in [x, y]$, $axy \in [a, x] \subseteq [p, x]$ and $axy \in [a, y] \subseteq [p, y]$, and so $axy = pxy$). Similarly $x \vee y = bxy$.

We next aim to show that $a \leftrightarrow b$. It suffices to show that $a \leftrightarrow x$ for some $x \in A$. We begin with the following observation.

Suppose we have a sequence, x_0, x_1, \dots, x_n , in A with $x_i \sim x_{i+1}$ for all i . We claim that $x_0 \wedge \dots \wedge x_{n-1} \leftrightarrow x_0 \wedge \dots \wedge x_n$. We prove this by induction on n simultaneously for all such sequences $(x_i)_i$. In fact, we use the stronger inductive hypothesis that $z \leftrightarrow x_0 \wedge \dots \wedge x_n$ for all $z \in [x_0 \wedge \dots \wedge x_{n-1}, x_0 \wedge \dots \wedge x_n]$. This holds for $n = 1$ by an earlier observation, since $x_0 \wedge x_1 \in [x_0, x_1]$ and $x_0 \sim x_1$. For the inductive step, suppose $n \geq 2$, and set $t = x_1 \wedge \dots \wedge x_{n-1}$, $u = x_0 \wedge \dots \wedge x_{n-1}$, $v = x_1 \wedge \dots \wedge x_n$ and $w = x_0 \wedge \dots \wedge x_n$. Let $z \in [u, w]$. We want to show that $z \leftrightarrow w$. Note that $w = u \wedge v$, u, z, w, v , $w \leq z \leq u \leq t$ and $w \leq v \leq t$. There are three cases. If $u \sim v$, then since $z, w \in [u, v]$, we have $z \leftrightarrow w$ by an earlier observation. If $u \leq v$, then $w = u$, so $z = w$, so $z \leftrightarrow w$. Finally, if $v \leq u$, then $w = v$, so $z \in [u, w] = [u, v] \subseteq [t, v]$, and so by the inductive hypothesis, applied to the sequence x_1, \dots, x_n , we have $z \leftrightarrow v = w$, as required. This completes the inductive step, proving the claim.

Now, there is an infinite sequence, $x = x_0, x_1, x_2, \dots$, in A , with $x_n \sim x_{n+1}$ for all n , and such that some subsequence $(x_{n_i})_i$ converges to a . (Since $a \in \bar{A}$, there is some sequence, $z_i \in A$ with $z_i \rightarrow a$. Since $z_i \approx z_{i+1}$, we can interpolate by elements

x_n so that $z_i = x_{n_i}$.) We can assume that $n_0 = 0$ and $x_0 = x$. Let $y_n = x_0 \wedge \dots \wedge x_n$, so that $x = y_0 \geq y_1 \geq y_2 \dots$. By the previous paragraph, we have $y_n \leftrightarrow y_{n+1}$ for all n , and so $y_{n_i} \leftrightarrow y_{n_{i+1}}$. Let α_i be (the image of) a monotone arc from y_{n_i} to $y_{n_{i+1}}$, and let $\alpha = \bigcup_{i=0}^{\infty} \alpha_i$. This is again a monotone arc. We claim that it converges on a . To see this, note that $y_{n_j} \in [a, x_{n_i}]$ for all $j \geq i$. In particular, $\alpha_i \subseteq [y_{n_i}, y_{n_{i+1}}] \subseteq [a, x_{n_i}]$. Since $x_{n_i} \rightarrow a$, by (weak) local convexity, α_i eventually lies in any neighbourhood of a in A . This justifies the claim. We similarly get a monotone arc from x to converging on b . Let $\beta(A)$ be their concatenation. Then $\beta(A) \cup \{a, b\}$ is a monotone arc from a to b . Thus $a \leftrightarrow b$, as claimed.

Let $L = \bigcup_{A \in \mathcal{A}} \beta(A) \subseteq M$. By construction, this is a totally ordered subset (hence a subalgebra) with $\min L = p$ and $\max L = q$.

We claim that L is compact. Since it is metrisable, it is enough to show that any sequence $(z_i)_i$ in L has a convergent subsequence. After passing to a subsequence (and maybe reversing the order), we can suppose that $(z_i)_i$ is strictly increasing. Let $A_i = [z_i] \in \mathcal{A}$. Since the closure of each $\beta(A_i)$ is an arc (hence compact) we can suppose that $(A_i)_i$ is also strictly increasing in \mathcal{A} . Let $b_i = \max A_i \in S$, and let \mathcal{B} be the set of $B \in \mathcal{A}$ such that $B < A_i$ for some i . Since $b_i \in S$, each $[p, b_i] \setminus \{b_i\}$ is open, so $\bigcup \mathcal{B} = \bigcup_{i=0}^{\infty} ([p, b_i] \setminus \{b_i\})$ is open. By connectedness of M , it cannot be closed, so there is some sequence $w_n \in \bigcup \mathcal{B}$ converging to some $w \in M \setminus \bigcup \mathcal{B}$. Suppose $U \ni w$ is open. By (weak) local convexity, $[w_n, w] \subseteq U$ for some (indeed all sufficiently large) n . By definition of \mathcal{B} , we have $A_i > w_n$ for all sufficiently large i . Also $A_i < w$. Thus, $z_i \in A_i \subseteq [w_n, w] \subseteq U$. This shows that $z_i \rightarrow w$ as required.

Now L has no adjacent pairs. (For suppose $a < b$ were adjacent in L . Then certainly, $a, b \in S$. But then $M = [p, a] \cup [b, q]$ contradicting connectedness of M .) Therefore, by Lemma 12.4.1, L is connected, and by Lemma 12.4.2, it is homeomorphic to a compact real interval with endpoints, p, q . In other words, it is a monotone arc from p to q as required. \square

12.5. Totally disconnected median algebras and duality.

We give a duality result due to Roller which generalises from the finite case.

Definition. A *Stone median algebra* is a topological median algebra which is compact and totally disconnected.

Let M be a Stone median algebra. By a *clopen* wall we mean a wall such that both of its halfspaces are open (hence also closed). This corresponds to a continuous epimorphism to a two-point median algebra. By Lemma 12.4.11, any two distinct points of M are separated by a clopen wall. We write $\mathcal{H}_0(M) \subseteq \mathcal{H}(M)$ for the set of clopen halfspaces. This is a subproset of the proset, $\mathcal{H}(M)$, of all halfspaces of M . The inclusion $\mathcal{H}_0(M) \hookrightarrow \mathcal{H}(M)$ induces a median homomorphism, $\mathcal{F}(\mathcal{H}(M)) \longrightarrow \mathcal{F}(\mathcal{H}_0(M))$ (which by Lemma 9.2.4 is surjective).

Given any proset, Ω , the set of flows, $\mathcal{F}(\Omega)$, is naturally a Stone median algebra (see Example (Ex12.7) of Subsection 12.1). If $a \in \Omega$, then $H(a) := \{R \in \mathcal{F}(\Omega) \mid$

$a \in R\}$ is clopen. This tells us that the image of the map, $\eta : \Omega \longrightarrow \mathcal{H}(\mathcal{F}(\Omega))$, described in Subsection 9.2, lies in $\mathcal{H}_0(\mathcal{F}(\Omega))$.

Putting the above observations together, we see that there are natural maps, $M \longrightarrow \mathcal{F}(\mathcal{H}_0(M))$ and $\Omega \longrightarrow \mathcal{H}_0(\mathcal{F}(\Omega))$. Note that, in the case where M or Ω are finite, these are the same as the maps defined in Subsection 9.2.

The following result of [R] generalises Propositions 9.2.2 and 9.2.7.

Theorem 12.5.1.

(1) *Let M be a Stone median algebra. The natural map $M \longrightarrow \mathcal{F}(\mathcal{H}_0(M))$ is an isomorphism of topological median algebras.*

(2) *Let Ω be a proset. The natural map $\Omega \longrightarrow \mathcal{H}_0(\mathcal{F}(\Omega))$ is an isomorphism of prosets*

Proof.

(1) In Subsection 9.2, we described the monomorphism, $M \longrightarrow \mathcal{F}(\mathcal{H}(M))$. Here, we have postcomposed this with the homomorphism $\mathcal{F}(\mathcal{H}(M)) \longrightarrow \mathcal{F}(\mathcal{H}_0(M))$.

The fact that this remains injective follows exactly as in Lemma 9.2.1. Here we use the observation (Lemma 12.4.11) that any two distinct points of M are separated by a clopen wall.

The fact that the map is surjective follows as in Lemma 9.2.2. Here we are taking a family of pairwise intersecting closed convex sets (halfspaces). By compactness, and the Helly Property (Lemma 7.1.1), this again has non-empty intersection.

We claim that our map is continuous. Let $H \in \mathcal{H}_0(M)$. Recall that $\{H, H^*\}$ is a factor of the product space used to define the topology on $\mathcal{F}(\mathcal{H}_0(M))$. We therefore want to check that the postcomposition with projection to $\{H, H^*\}$ is continuous. But this is just the map $M \longrightarrow \{H, H^*\}$ which tells us in which halfspace a given element of M belongs. Since H is clopen, this map is continuous, as required. This proves the claim.

It follows that the map $M \longrightarrow \mathcal{F}(\mathcal{H}_0(M))$ is a continuous bijection between compact hausdorff spaces, hence a homeomorphism.

(2) This is the same as the map, $H : \Omega \longrightarrow \mathcal{H}(\mathcal{F}(\Omega))$, defined in Subsection 9.2, except that we have observed that its image lies in $\mathcal{H}_0(\mathcal{F}(\Omega))$, which we now view as the target. By Lemma 9.2.6, this a proset monomorphism, so we need to check that its image is precisely $\mathcal{H}_0(\mathcal{F}(\Omega))$. For this we follow the proof of Proposition 9.2.7. Here we take $G \in \mathcal{H}_0(\mathcal{F}(\Omega))$. Since G is clopen, it follows by Lemma 12.2.4 that both G and G^* are gated. By Lemma 7.3.6, there are mutual gates, $R \in G$ and $S \in G^*$. The argument can now be completed as before. \square

12.6. A compactification procedure.

We next describe a procedure for compactifying an interval-compact median algebra (though we don't need to assume interval-compactness for the moment). This is based on a construction in [Fi3].

Let M be a topological median algebra. Let $P = \prod_{a,b \in M} [a, b]$ be the direct product of all intervals in M . Then P is a topological median algebra equipped with the product median and topology. We define a map $\pi : M \rightarrow P$, by setting the $[a, b]$ coordinate of $\phi(x)$ to be $abx \in [a, b]_M$. Clearly ϕ is injective. Also ϕ is a continuous monomorphism (since each of the maps $[x \mapsto abx] : M \rightarrow [a, b]_M$ is a continuous homomorphism). We write \check{M} for the closure of $\phi(M)$ in P .

Given $p \in P$, write $[abp] \in [a, b]_M$ for the $[a, b]$ -coordinate of p . By construction, if $x \in M$, then $[ab(\phi x)] = abx$. If $c, d \in M$, and $p \in \phi(M)$, then $ab[cdp] = (abc)(abd)[abp]$. (Setting $p = \phi x$, this is just the long distributive law: $ab(cdx) = (abc)(abd)(abx)$ in M .) By continuity of the median in P , this also holds for all $p \in \check{M}$.

Lemma 12.6.1. *$\phi(M)$ is convex in \check{M} .*

Proof. Let $c, d \in M$ and let $p \in [\phi c, \phi d]_{\check{M}} = [\phi c, \phi d]_P \cap \check{M}$. We want to show that $p \in \phi(M)$. In fact, we claim that $p = \phi e$, where $e = [cdp] \in [c, d]_M$.

To see this, first note that by the definition of the median on P , we have $[abp] \in [abc, abd]_M$ for all $a, b \in M$. In other words, $[abp] = (abc)(abd)[abp]$. Moreover, since $p \in \check{M}$, we have $(abc)(abd)[abp] = ab[cdp] = abe$. Thus $[abp] = abe = [ab(\phi e)]$. In other words the $[a, b]$ -coordinates of p and of ϕe agree for all $a, b \in M$. Therefore $p = \phi e$ as claimed. \square

We can now identify M with $\phi(M)$ in \check{M} , except that the induced subspace topology on M might be coarser than the original. Note that by Lemma 12.3.4, we have $\text{rank}(\check{M}) = \text{rank}(M)$. We also note that \check{M} inherits certain properties from M . (For example, if M is locally convex, then so is \check{M} .)

The main interest in this construction is when M is interval-compact. In this case, P is compact by Tychonoff's Theorem, and so also is \check{M} .

In summary, any interval-compact median algebra has a continuous embedding as a dense convex subset of a compact median algebra. The induced subspace topology may be coarser than the original.

In the case of a discrete median algebra, Π , we recover the Roller boundary as follows.

Let $\Psi = \Psi(\Pi) := \prod \mathcal{W}(\Pi)$, and let $\Pi \subseteq \mathcal{F}(\mathcal{H}(\Pi)) \subseteq \Psi$, be as described in Lemma 11.12.2. Then, $\mathcal{F}(\mathcal{H}(\Pi))$ is the closure, $\bar{\Pi}$, of Π in Ψ , and $\partial_R \Pi \equiv \bar{\Pi} \setminus \Pi$.

Let $\mathcal{C}(\Psi)$ be the set of faces of Ψ , and let $\mathcal{B} \subseteq \mathcal{C}(\Psi)$ be some subset. Let $\Omega = \prod \mathcal{B}$. We equip Ω with the product topology, which is compact by Tychonoff's Theorem. We define a map $f : \Psi \rightarrow \Omega$ as follows. Given $p \in \Psi$ and $Q \in \mathcal{C}(\Psi)$, we let the Q -coordinate of $f(p)$ be $\pi_Q(p)$, where $\pi_Q : \Psi \rightarrow Q$ is projection to the face, Q . Thus, f is a continuous homomorphism. Moreover, if \mathcal{B} contains at least one 1-face crossing W for each $W \in \mathcal{W}(\Pi)$, then f is injective, hence a homeomorphism to its range, $f(\Psi) \subseteq \Omega$.

Now let $\mathcal{B} = \{[a, b]_\Psi \subseteq \Psi \mid a, b \in \Pi\}$. Note that $[a, b]_\Pi = \Pi \cap [a, b]_\Psi$. We can thus view $P = \prod_{a,b \in \Pi} [a, b]_\Pi$ as a closed subset of $\Omega = \prod_{a,b \in \Pi} [a, b]_\Psi$. We have

continuous monomorphisms, $\Pi \hookrightarrow P \hookrightarrow \Omega$ and $\Pi \hookrightarrow \Psi \hookrightarrow \Omega$, the last map being f constructed above. The two compositions agree. By definition, $\check{\Pi}$ is the closure of Π in P , and by Lemma 11.12.2, $\mathcal{F}(\mathcal{H}(\Pi))$ is the closure of Π in Ψ . After applying the maps $P \hookrightarrow \Omega$ and $\Psi \hookrightarrow \Omega$, we get a natural identification of $\check{\Pi}$ with $\mathcal{F}(\mathcal{H}(\Pi))$. In summary, this shows:

Lemma 12.6.2. *Let Π be a discrete median algebra. Then there is a natural identification of $\partial_R \Pi$ with $\check{\Pi} \setminus \Pi$.*

12.7. Rank-1 topological median algebras.

We finish with a brief discussion of the rank-1 case.

Lemma 12.7.1. *Let M be a locally convex median algebra of rank 1. Then M contains no topologically embedded circle.*

Proof. Suppose that $\sigma \subseteq M$ is homeomorphic to the circle. We choose any four distinct points a_1, a_2, a_3, a_4 in σ so as to cut it into four arcs: $\sigma = \sigma_1 \cup \sigma_2 \cup \sigma_3 \cup \sigma_4$ with $\sigma_{i-1} \cap \sigma_i = \{a_i\}$ (where we are taking indices modulo 4). Then $\sigma_1 \cap \sigma_3 = \sigma_2 \cap \sigma_4 = \emptyset$. Let $U_i \subseteq M$ be an open set containing σ_i such that $U_1 \cap U_3 = U_2 \cap U_4 = \emptyset$.

Now σ_i has an open cover by the interiors of convex sets which lie in U_i . By taking a finite subcover, we can find a sequence of points, $a_i = x_{i0}, x_{i1}, \dots, x_{ip_i} = a_{i+1}$ in σ_i such that for all j , there is some open set, V_{ij} with $x_{ij}, x_{i,j+1} \in V_{ij}$ and with \bar{V}_{ij} convex and contained in U_i . Let $X = \{x_{ij} \mid 1 \leq i \leq 4, 0 \leq j \leq p_i\}$, and let $\Pi = \langle X \rangle \subseteq M$ be the subalgebra generated by X .

Now Π has rank 1, and so is naturally the vertex set of a simplicial tree, $\Delta(\Pi)$. Let $\beta_i \subseteq \Delta(\Pi)$ be the arc connecting a_i to a_{i+1} in $\Delta(\Pi)$, and let α_{ij} be the arc connecting x_{ij} to $x_{i,j+1}$ in $\Delta(\Pi)$. Then $\beta_i \subseteq \bigcup_{j=0}^{p_i-1} \alpha_{ij}$. If $y \in \beta_i \cap \Pi$, then $y \in \alpha_{ij} \cap \Pi$ for some j . Now $x_{ij}, x_{i,j+1} \in V_{ij}$, so since $x_{ij} \cdot y \cdot x_{i,j+1}$ and \bar{V}_{ij} is convex, we have $y \in \bar{V}_{ij} \subseteq U_i$. In other words $\beta_i \cap \Pi \subseteq U_i$.

Now $\beta_1 \cup \beta_2 \cup \beta_3 \cup \beta_4$ is a closed path in $\Delta(\Pi)$. From this one sees easily that at least one of $\beta_1 \cap \beta_3 \cap \Pi$ or $\beta_2 \cap \beta_4 \cap \Pi$ is non-empty. But by the above observation, these are contained in $U_1 \cap U_3$ and $U_2 \cap U_4$ respectively, so we get a contradiction.

This shows that no such circle, σ , can exist. \square

Lemma 12.7.2. *Let M be locally convex, metrisable, interval-compact, and rank 1. Then any two points are connected by a unique arc (that is a subset homeomorphic to the closed real interval $[0, 1]$). Moreover, such any such arc is the median interval between its endpoints.*

Proof. Let $a, b \in M$. Then by Lemma 12.4.8, $[a, b]$ is homeomorphic to $[0, 1]$, via a median isomorphism. Thus a and b get sent to 0 and 1. By Lemma 12.7.1, M contains no topologically embedded circle. It is now a simple exercise to show that $[a, b]$ is the unique arc from a to b . \square

We can therefore think of such a space as a kind of “real tree”. Such objects have been studied by a number of authors (see [AdeN, Bo1], and the references

therein). For example, it is shown in [MayO] that such a space is homeomorphic to an \mathbb{R} -tree, which will be the main topic of Section 15. We also note Lemma 11.3.1 gives us a natural way to compactify such a space (modulo making the topology coarser). The resulting space is then a “dendron”.

13. MEDIAN METRIC SPACES

Median algebras often arise as median metric spaces (as illustrated by the examples in Subsection 13.1). The property is closed under completion (Lemma 13.2.6), so there is often no loss in assuming completeness. In this case, all non-empty closed convex subsets are gated (Lemma 13.3.7). If we also assume finite rank, then all intervals are compact (Lemma 13.2.10). A connected complete median metric space is geodesic (Lemma 13.3.2). There are also a number of results which allow us to construct or modify median metrics. A few of these are mentioned briefly at the end of the section (Theorems 13.4.1, 13.4.2 and 13.4.3). We begin with the definition.

13.1. Definition and examples.

Let (M, ρ) be a metric space. Given $a, b, x \in M$, write $a.x.b$ to mean $\rho(a, x) + \rho(x, b) = \rho(a, b)$. We write

$$[a, b]_\rho = \{x \in M \mid a.x.b\}.$$

Definition. ρ is a *median metric* if for all $a, b, c \in M$, we have

$$\#[a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho = 1.$$

We refer to M as a *median metric space*.

Writing $I(a, b) = [a, b]_\rho$, we see that M satisfies the Axioms (I1)–(I4) of Subsection 4.1: Axioms (I1)–(I3) are immediate from the metric space axioms, and (I4) is precisely the above definition. Therefore, by Sholander’s theorem (Theorem 4.1.1 here) the median map defined by (I4) gives M the structure of a median algebra such that $[a, b] = [a, b]_\rho$ for all $a, b \in M$. We see that the betweenness relation $a.x.b$ as defined above is consistent with its earlier definition.

The following gives us an equivalent way of defining a median metric.

Lemma 13.1.1. *Let M be a median algebra equipped with a metric ρ such that $\rho(a, c) + \rho(c, b) = \rho(a, b)$ for all $c \in [a, b]$. Then ρ is a median metric.*

Proof. We are assuming that $[a, b] \subseteq [a, b]_\rho$. We claim that $[a, b] = [a, b]_\rho$. To see this, let $c \in M$, and let $m = abc$. We have:

$$\begin{aligned} \rho(a, m) + \rho(m, c) &= \rho(a, c) \\ \rho(b, m) + \rho(m, c) &= \rho(b, c) \\ \rho(a, m) + \rho(m, b) &= \rho(a, b), \end{aligned}$$

and so

$$\rho(a, c) + \rho(c, b) = \rho(a, b) + 2\rho(m, c).$$

Therefore, if $c \in [a, b]_\rho$, then $\rho(m, c) = 0$ and so $c \in [a, b]$ as required.

Since $[a, b] \cap [b, c] \cap [c, a] = \{abc\}$, ρ satisfies (I4). □

Note that, in fact, we don't need to assume that ρ satisfies the triangle inequalities a-priori for this. These follow directly from the formulae given in the proof.

For some purposes it is convenient to allow for pseudometrics (i.e. where distinct points might be at distance 0). We will say that a pseudometric, ρ , on a median algebra is a **median pseudometric** if $[a, b] \subseteq [a, b]_\rho$ for all $a, b \in M$. Note that we cannot in general recover the median structure from the pseudometric. However, the hausdorffification of a median pseudometric is a median metric. (The hausdorffification process is discussed further in Subsection 15.3.) It can be seen that much of what we say about median metrics can be reinterpreted for pseudometrics. By default in this section, ρ will be a metric.

We note that properties which only refer to the betweenness relation in a median metric space (such as (Ex6.1) and (Ex6.2) of Subsection 6.2 etc.) pass directly to any subspace even if it is not itself median. In view of this, we say that a metric space is **submedian** if it isometrically embeds into a median metric space. Examples of such arise from spaces of measured walls (Section 19) and quasimedial graphs (Section 23).

We will see (Lemma 13.2.2) that a median metric space, M , is a topological median algebra in the induced metric.

We also note that any subalgebra of a median metric space is a median metric space in the induced metric.

Here are a few examples.

(Ex13.1): Let Π be a discrete median algebra, and let $w : \mathcal{W}(\Pi) \rightarrow (0, \infty)$ be any map. Given $a, b \in \Pi$, let $\rho(a, b) = \sum_{W \in \mathcal{W}(\Pi)} w(W)$. Then ρ is a median metric. This follows from Lemma 13.1.1, given that for any $a, b, c \in M$, we have $\mathcal{W}(a, b) \subseteq \mathcal{W}(a, c) \cup \mathcal{W}(c, b)$, and that $c \in [a, b]$ if and only if $\mathcal{W}(a, c) \cap \mathcal{W}(c, b) = \emptyset$. In fact, we will see (Lemma 13.2.5) that every discrete median metric space has this form. If we set $w \equiv 1$, then we recover the combinatorial metric defined in Subsection 11.2.

(Ex13.2): A particular case of (Ex13.1), which will be useful later, is when $\Pi = F(X)$ is the free median algebra on a finite set X . Here $\mathcal{W}(F(X))$ can be identified with the set of partitions of X into two disjoint non-empty subsets. A median metric on $F(X)$ arises from assigning a positive real number to each such partition.

(Ex13.3): Example (Ex13.1) can be viewed as a special case of a “space with measured walls”. These give rise to median (pseudo)metrics, which will be the topic of Section 19.

(Ex13.4): Let X be any set and let $\Psi = \{0, 1\}^X$. We can embed Ψ into the real hypercube, $\hat{\Delta}(\Psi) = [0, 1]^X$. Suppose $w : X \rightarrow (0, \infty)$. Then as a median algebra, $\hat{\Delta}(\Psi)$ is isomorphic to $\prod_{a \in X} [0, w(a)]$, where $[0, w(a)] \subseteq \mathbb{R}$ is the real interval. We can now equip $\hat{\Delta}(\Psi)$ with the l^1 non-finite metric. (Recall that this means that it may take infinite values.) It satisfies the hypotheses of Lemma 13.1.1 modulo this qualification. Therefore any subalgebra of $\hat{\Delta}(\Psi)$ on which the metric is finite will be a genuine median metric space. For example, if Π is a discrete median algebra, then we can embed its realisation, $\Delta(\Pi)$, as the subalgebra of $[0, 1]^{\mathcal{W}(\Pi)}$, as described in Subsection 11.2. If $w : \mathcal{W}(\Pi) \rightarrow (0, \infty)$, then this is median isomorphic to $\prod_{W \in \mathcal{W}(\Pi)} [0, w(W)]$. Let ρ be the induced l^1 metric on $\Delta(\Pi)$. Then $(\Delta(\Pi), \rho)$ is a median metric space. We discuss this example further in Section 17.

(Ex13.5): Such complexes as discussed in the previous example arise in group theory from right-angled Artin groups (often abbreviated to “RAAG”). By definition, such a group, G , has a presentation of the form $G = \langle X | R \rangle$, where $X = \{x_1, \dots, x_n\}$ is a finite set of generators, and R is the set of relators of the form $\{x_i x_j x_i^{-1} x_j^{-1} \mid (i, j) \in P\}$, for some subset $P \subseteq \{1, \dots, n\}^2$. (For example, if $P = \emptyset$, then G is free on X . If $P = \{1, \dots, n\}^2$, then $G \cong \mathbb{Z}^n$ is free abelian.) One can show that the word metric on G with respect to X is median. We can construct the cube complex $\Delta(G)$ as in example (Ex13.4). Its 1-skeleton is precisely the Cayley graph of G with generating set X .

(Ex13.7): Let Ω be a set equipped with a measure μ . Given a measurable function, $f : \Omega \rightarrow \mathbb{R}$, define its “ L^1 -norm” as $\|f\| := \int |f| d\mu$. Let \mathcal{L} be the set of those f for which $\|f\| < \infty$. Given $f, g \in \mathcal{L}$, let $\rho(f, g) = \|f - g\|$. Then (\mathcal{L}, ρ) is a median pseudometric space. In fact, the median, m , of $f, g, h \in \mathcal{L}$ can be defined by setting $m(x)$ to be the standard median of $f(x), g(x), h(x)$, for all $x \in \Omega$. Taking the hausdorffification of \mathcal{L} , we obtain a contractible geodesic median metric space. A standard result from measure theory tells us that this is complete (see for example [Ber]). In general, this may have infinite rank. Note that the set of characteristic functions of measurable subsets of Ω forms a topologically closed subalgebra of \mathcal{L} . We elaborate on a particular case of this in the next example.

(Ex13.8): Let Ω be a probability space: that is a set, Ω , equipped with a measure μ with $\mu(\Omega) = 1$. Let $\mathcal{M} \subseteq \mathcal{P}(\Omega)$ be the σ -algebra of measurable subsets. This a median subalgebra of $\mathcal{P}(\Omega)$, and intrinsically a boolean algebra. Given $A, B \in \mathcal{M}$, write $\rho(A, B) = \mu(A \Delta B)$. Since $A \Delta B \subseteq (A \Delta C) \cup (B \Delta C)$ for all $C \in \mathcal{M}$, we see that ρ is a pseudometric on \mathcal{M} . (This is a particular case of Example (Ex13.7) after identifying a set with its characteristic function.) Write $A \sim B$ to mean that $\rho(A, B) = 0$: in other words, A, B differ by a null set (a set of measure 0). Let $M = \mathcal{M}/\sim$. This is the hausdorffification of \mathcal{M} . We write ρ also for the induced metric on M .

Recall that, in the median structure induced from $\mathcal{P}(\Omega)$, we have $C \in [A, B]$ if and only if $A \cap B \subseteq C \subseteq A \cup B$. In this case, $A \triangle B = (A \triangle C) \sqcup (B \triangle C)$, so $\rho(A, B) = \rho(A, C) + \rho(B, C)$. It follows that ρ is a median pseudometric on \mathcal{M} , and that (M, ρ) is a median metric space.

In what follows, we will regard sets in \mathcal{M} to be defined up to \sim . Note that the operations of (finite or) countable union and intersection and (relative) complement are well defined modulo this equivalence. To work in M , we can choose arbitrary representatives in \mathcal{M} .

We note that (M, ρ) is complete. This follows from the general fact that L^1 spaces are complete. For a more direct argument, let $(A_n)_{n \in \mathbb{N}}$ be a Cauchy sequence. It is enough to show that $(A_n)_n$ subconverges, so after passing to a subsequence, we can suppose that $\rho(A_n, A_{n+1}) \leq \frac{1}{2^n}$ for all n . Let $B_n = \bigcup_{i=n}^{\infty} A_i$ and $B = \bigcap_{n=0}^{\infty} B_n$. We claim that $\rho(A_n, B) \leq \frac{2}{2^n}$. To this end, let $R_n = \bigcup_{i=0}^n (A_i \triangle A_{i+1})$. Thus, $\mu(R_n) \leq \sum_{i=0}^n \mu(A_i \triangle A_{i+1}) \leq \sum_{i=n}^{\infty} \frac{1}{2^i} = \frac{2}{2^n}$. Now for all $i \geq n$, we have $A_n \triangle A_i \subseteq R_n$, so $A_n \triangle B_i \subseteq \bigcup_{i=n}^{\infty} A_i \subseteq R_n$, and so, since $(B_n)_n$ is non-increasing, $A_n \triangle B \subseteq \bigcup_{i=n}^{\infty} (A_n \triangle B_i) \subseteq R_n$. Thus, $\mu(A_n, B) \leq \frac{2}{2^n}$ as claimed. This shows that M is complete.

If $P \in M$, then $(A \triangle P) \triangle (B \triangle P) = A \triangle B$, and so $\mu(A \triangle P, B \triangle P) = \mu(A, B)$. Thus $[A \mapsto A \triangle P]$ defines an isometric involution $\tau : M \rightarrow M$ (cf. the discussion of ternary boolean algebras in Example (Ex3.6) of Subsection 3.4). If we set $P = A \triangle B$, we get $\tau(A) = B$, showing that M is homogeneous. If we set $P = \Omega$, we get $\rho(A, \tau(A)) = 1$ for all $A \in M$. In the latter case, we can think of τ as an ‘‘antipodal map’’.

Suppose $\mathcal{I} \subseteq M$ is a family of pairwise disjoint non-null subsets of Ω . The map $[Q \mapsto \bigcup Q] : \mathcal{P}(\mathcal{I}) \rightarrow M$ is a median monomorphism, so its image in M is a $(\#\mathcal{I})$ -hypercube. In fact, any finite subset of M lies in a finite such cube: if $A_1, \dots, A_m \in M$, consider the set, \mathcal{I} , of non-null sets of the form $B_1 \cap \dots \cap B_m$, where each B_i is either A_i or $\Omega \setminus A_i$. Then each A_i is a disjoint union of elements of \mathcal{I} , and so lies in a $(\#\mathcal{I})$ -cube as described above.

(Ex13.9): As a particular case of Example (Ex13.8), let μ be Lebesgue measure on the unit real interval, $\Omega = [0, 1] \subseteq \mathbb{R}$. Given $A, B \in M$, and $t \in [0, 1]$, let $f(A, B, t) = (B \cap [0, t]) \cup (A \cap [t, 1])$. (This is well defined up to measure 0.) Note that $f(A, B, 0) = A$, $f(A, B, 1) = B$, and that $f(A, B, t) \in [A, B]$. If $C, D \in M$, and $u \in [0, 1]$, then $f(A, B, t) \triangle f(C, D, u) \subseteq (A \triangle C) \cup (B \triangle D) \cup [t, u]$, and so $\rho(f(A, B, t), f(C, D, u)) \leq \rho(A, C) + \rho(B, D) + |t - u|$. In other words, the map $f : M^2 \times [0, 1] \rightarrow M$ is 1-lipschitz with respect to the l^1 metric induced on $M^2 \times [0, 1]$. This shows that M is contractible. (For any basepoint $A \in M$, the map $[(X, t) \mapsto f(A, X, t)] : M \times [0, 1] \rightarrow M$ is a homotopy from the identity to the constant map at A .) Moreover, for any $A, B \in M$, the map $[t \mapsto f(A, B, t)] : [0, 1] \rightarrow M$ is a monotone path from A to B . This shows that M is a geodesic space (see Subsection 13.3).

Here is a more combinatorial way of describing M . We construct abstractly an increasing sequence, $Q_0 \hookrightarrow Q_1 \hookrightarrow Q_2 \hookrightarrow \dots$, of cubes Q_n of rank 2^n , and with each inclusion a median monomorphism. We equip each Q_n with a metric ρ_n , namely $\frac{1}{2^n}$ times the combinatorial metric (as defined in Subsection 11.2). In other words, if $\{a, b\}$ is a 1-face of Q_n , then $\rho_n(a, b) = \frac{1}{2^n}$. To get from Q_n to Q_{n+1} , we enlarge each such 1-face to a 2-cube, $\{a, c, b, d\}$, with a, b antipodal, so that Q_{n+1} is a direct product of 2^n such 2-cubes. Note that $\rho_{n+1}(a, b) = \rho_{n+1}(a, c) + \rho_{n+1}(b, c) = \frac{2}{2^{n+1}} = \frac{1}{2^n} = \rho_n(a, b)$. It follows that each inclusion is an isometric embedding. Let $Q_\infty = \bigcup_{n=0}^\infty Q_n$, with the induced median and metric. This is an infinite-rank median metric space. Let \hat{Q}_∞ be its metric completion. This is again a median metric space (see Lemma 13.2.6 below).

To relate this to the original description, let \mathcal{I}_n be a subdivision of $[0, 1]$ into 2^n disjoint intervals of length $\frac{1}{2^n}$. We can identify Q_n with the image of the monomorphism, $\mathcal{P}(\mathcal{I}_n) \rightarrow M$, as described in Example (Ex13.8). We have natural maps $Q_n \hookrightarrow Q_{n+1}$, and one can check that Q_∞ is dense in M . Thus, M is isometric to \hat{Q}_∞ , as in the combinatorial construction.

In summary, M is contractible, complete, homogeneous, geodesic and infinite-rank. Each finite subset of M lies in a finite cube, and M admits an isometric involution, $\tau : M \rightarrow M$, with $\rho(A, \tau(A)) = 1$ for all $A \in M$.

(Ex13.10): As noted earlier, median metric spaces arise from spaces with measured walls which are the topic of Section 19. They also arise as Guirardel cores (Subsection 15.4), and from asymptotic cones of certain spaces (Subsection 24.3).

We note that in Examples (Ex13.1) and (Ex13.2) above, we could allow w to take the value 0, in which case we would get a median pseudometric space.

13.2. Some basic properties.

We move on to consider some basic properties of median metric spaces.

The first observation holds in any metric space.

Lemma 13.2.1. *Let $a, c, d, b \in M$ with $a.c.d$ and $c.d.b$. Then $\rho(c, d) \leq \rho(a, b)$.*

Proof. We just add the inequalities:

$$\begin{aligned} \rho(a, c) + \rho(c, d) &= \rho(a, d) \leq \rho(a, b) + \rho(b, d) \\ \rho(b, d) + \rho(c, d) &= \rho(b, c) \leq \rho(a, b) + \rho(a, c), \end{aligned}$$

and subtract $\rho(a, c) + \rho(b, d)$. □

The following is an immediate consequence, but it is sufficiently fundamental to the subject to be worth recording explicitly.

Lemma 13.2.2. *Let (M, ρ) be a median metric space, and let $a, b, c, d \in M$. Then $\rho(abc, abd) \leq \rho(c, d)$.*

Note that Lemma 13.2.2 implies that M is 1-lipschitz in the sense defined in Subsection 12.2. That is to say, the median map $[(x, y, z) \mapsto xyz] : M^3 \rightarrow M$ is 1-lipschitz with respect to the l^1 metric on M^3 . In particular, M is a topological median algebra. We also note (since $\text{diam}([a, b]) \leq \rho(a, b)$ for all $a, b \in M$) that M is weakly locally convex, as defined in Subsection 12.2.

The following is also worth observing:

Lemma 13.2.3. *If $c, d \in [a, b]$, then $\rho(c, d) \leq \rho(a, b)$.*

Proof. $2\rho(c, d) \leq (\rho(c, a) + \rho(a, d)) + (\rho(c, b) + \rho(b, d)) = (\rho(a, c) + \rho(c, b)) + (\rho(a, d) + \rho(d, b)) = 2\rho(a, b)$. \square

Recalling the notion of parallelism defined in Subsection 7.2, the following is another immediate consequence of Lemma 13.2.1.

Lemma 13.2.4. *Suppose $a, a', b, b' \in M$ with $aa' \parallel bb'$. Then $\rho(a, a') = \rho(b, b')$.*

This in turn has an immediate consequence for a discrete median metric space, Π (by which we mean that Π is discrete as a median algebra). This is a kind of converse to the construction of Example (Ex13.1) above.

Recall that a wall, $W \in \mathcal{W}(\Pi)$, determines a parallel class, $\mathcal{E}(W)$, of 1-cells of Π which cross W (see Lemma 11.5.1). We write $w(W) = \rho(c, d)$, where $\{c, d\} \in \mathcal{E}(W)$. By Lemma 13.2.3, this is well defined, independently of the choice of 1-cell. We refer to $w(W)$ as the *width* of the wall W .

Suppose $a, b \in \Pi$. Recall from Subsection 11.2 that a “geodesic path” in M is a maximal chain, $a = a_0 < a_1 < \dots < a_n = b$ in $[a, b]$. We let $W_i \in \mathcal{W}(\Pi)$ be the wall separating a_{i-1} and a_i , i.e. for which $\{a_{i-1}, a_i\} \in \mathcal{E}(W_i)$. We then have $\mathcal{W}(a, b) = \{W_1, \dots, W_n\}$. Note also that $a_i a_j a_k$ whenever $i \leq j \leq k$.

Lemma 13.2.5. *Let (Π, ρ) be a discrete median metric space. If $a, b \in \Pi$, then $\rho(a, b) = \sum_{W \in \mathcal{W}(a, b)} w(W)$, where $w(W)$ is the width of W .*

Proof. By the above, we have

$$\rho(a, b) = \sum_{i=1}^n \rho(a_{i-1}, a_i) = \sum_{i=1}^n w(W_i) = \sum_{W \in \mathcal{W}(a, b)} w(W).$$

\square

We note that this is consistent with the discussion of Subsection 11.2, where the metric, ρ , there was defined by setting the width of each wall to be equal to 1.

We note that Lemma 13.2.5 applies equally well to a discrete median pseudo-metric space, where we allow walls to have width 0.

Since any finite subset of a metric median space, M , lies in a finite subalgebra, this is a useful way of understanding the distances between them.

Suppose, for example, $a, b, c, d \in M$. Let $F = F(\{a, b, c, d\})$ be the free median algebra on $\{a, b, c, d\}$. We gave a description of F in Subsection 5.4. It is the vertex set of a cube complex with one central cube, and four “free edges”

connecting alternate corners of the cube respectively to a, b, c, d . The walls of F correspond to the three partitions of $\{a, b, c, d\}$ into two disjoint non-empty subsets. The homomorphism from F into M induces a width to each wall (possibly 0), in the manner described above. We denote by $A, B, C, D \geq 0$, the widths of the walls crossing the free edges with endpoints a, b, c, d respectively. Thus, $A = \rho(a, (bcd|a))$ etc. We denote by R, S, T , the widths of the three remaining walls (crossing the central cube) corresponding respectively to the partitions $\{\{a, b\}, \{c, d\}\}$, $\{\{a, c\}, \{b, d\}\}$ and $\{\{a, d\}, \{b, c\}\}$. We can now read off the distances between points in $\langle\{a, b, c, d\}\rangle \subseteq M$ from this picture. For example, $\rho(d, abc) = D + R + S + T$: to get from d to abc , we cross the free edge at d and then three edges of the central cube.

To illustrate this, we note that Lemma 13.2.2 can be verified in these terms. In fact, we get $\rho(a, b) = \rho(abc, abd) + \rho(c, (abd|c)) + \rho(d, (abc|d))$, both sides equating to $C + S + T + D$. We give another application of this principle below.

First, we introduce the following notation. Given any $a, b, c \in M$, we write

$$\langle a, b : c \rangle := \frac{1}{2}(\rho(a, c) + \rho(c, b) - \rho(a, b)).$$

This is often referred to as the **Gromov product** of a, b based at c . Note that $\langle a, b : c \rangle = \rho(c, abc)$. In particular $\langle a, b : c \rangle = 0$ if and only if $a.c.b$.

Given $a, b, c, d \in M$, write

$$\begin{aligned} G(a, b, c|d) &:= \rho(a, d) + \rho(b, d) + \rho(c, d) - \frac{1}{2}(\rho(a, b) + \rho(b, c) + \rho(c, a)) \\ &= \langle a, b : d \rangle + \langle b, c : d \rangle + \langle c, a : d \rangle \geq 0. \end{aligned}$$

Note that $G(a, b, c|d) = 0$ if and only if $d = abc$. In fact,

$$G(a, b, c|d) = \rho(d, abc) + 2\rho(d, (abc|d)).$$

This can be seen by inspection of the picture of $F(\{a, b, c, d\})$ discussed above. In the above notation, we have $\rho(a, b) = S + T + A + B$, etc., and so

$$\begin{aligned} G(a, b, c|d) &= (2R + 2S + 2T + A + B + C + 3D) - (R + S + T + A + B + C) \\ &= R + S + T + 3D, \\ \rho(d, abc) &= R + S + T + D, \\ \rho(d, (abc|d)) &= D. \end{aligned}$$

The statement now follows. In particular, $\rho(d, abc) \leq G(a, b, c|d)$.

As an application, we have:

Lemma 13.2.6. *The metric completion of a median metric space is a median metric space of the same rank.*

Proof. Let \hat{M} be the completion of a median metric space, M . Let $a, b, c \in \hat{M}$. Let $(a_i)_i, (b_i)_i, (c_i)_i$ be sequences in M converging on a, b, c , respectively. Let $m_i = a_i b_i c_i$. Now $\langle a_i, b_i : m_i \rangle = 0$ for all i and so $\langle a, b : m \rangle = 0$, so $m \in [a, b]_\rho$. By symmetry, we have $m \in [a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho$. Suppose that $d \in [a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho$. Let $(d_i)_i$ be a sequence in M converging on d . Then, $G(a_i, b_i, c_i | d_i) \rightarrow G(a, b, c | d) = 0$. By the above observation, $\rho(d_i, m_i) \leq G(a_i, b_i, c_i | d_i)$ and so $\rho(d, m) = 0$, so $d = m$. We have shown that $[a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho = \{m\}$, and so ρ is a median metric.

The fact that $\text{rank}(\hat{M}) = \text{rank}(M)$ follows by Lemma 12.3.4. \square

We also note:

Lemma 13.2.7. *A median metric space of finite rank is locally convex.*

Proof. We have already noted that any median metric space is weakly locally convex, and so the statement follows from Lemma 12.2.4. \square

Note that any isometric embedding of one median metric space into another is necessarily a median monomorphism.

Lemma 13.2.8. *Let Π be a discrete median metric space, and let $a, b \in \Pi$. Suppose $\text{rank}([a, b]) = \nu < \infty$. Then $[a, b]$ isometrically embeds in the l^1 product, $\prod_{i=1}^\nu [0, l_i]$, for some $l_i \geq 0$ with $\sum_{i=1}^\nu l_i = \rho(a, b)$.*

Proof. Recall that by Lemma 8.3.3, there is a median monomorphism of $[a, b]$ into $P := \prod_{i=1}^\nu I_{p_i}$, where $I_p = \{1, \dots, p\}$. By the construction, we have an identification of $\mathcal{W}(a, b)$ with $\mathcal{W}(P) = \mathcal{W}(I_{p_1}) \sqcup \dots \sqcup \mathcal{W}(I_{p_\nu})$. Given $W \in \mathcal{W}(a, b)$, let $w(W)$ be the width of W in Π . This gives rise to a median metric on P , by the construction of Example (Ex13.1) above. Moreover, the embedding of $[a, b]$ into P is isometric. (Since if $c, d \in [a, b]$, then $\mathcal{W}(c, d) \subseteq \mathcal{W}(a, b)$, and the distance between c and d , namely $\sum_{W \in \mathcal{W}(c, d)} w(W)$, is the same whether interpreted in $[a, b]$ or in P .)

We can further embed P isometrically into $\prod_{i=1}^\nu [0, l_i]$, where $l_i = \sum_{W \in \mathcal{W}(I_{p_i})} w(W)$. The composition of these maps gives an isometric embedding of $[a, b]$ into $\prod_{i=1}^\nu [0, l_i]$ as required. Note that $\sum_{i=1}^\nu l_i = \sum_{W \in \mathcal{W}(a, b)} w(W) = \rho(a, b)$. \square

We will give a generalisation of Lemma 13.2.8 as Lemma 13.3.5, below.

Recall that a metric space is **totally bounded** (or “precompact”) if for all $\delta > 0$ it has a finite covering by δ -balls. This is equivalent to saying that for any $\epsilon > 0$, there is a bound on the cardinality of any finite ϵ -separated subset (i.e. a subset such that any two distinct points thereof are a distance at least ϵ apart).

Lemma 13.2.9. *Let M be a median metric space, and let $a, b \in M$, with $\text{rank}([a, b]) < \infty$. Then $[a, b]$ is totally bounded.*

Proof. Let $\nu = \text{rank}([a, b]) < \infty$. Let $A \subseteq [a, b]$ be a finite ϵ -separated subset. Let $\Pi = \langle A \cup \{a, b\} \rangle \subseteq [a, b]$. Then Π is finite, and $\Pi = [a, b]_\Pi$. By Lemma 13.2.8, there is an isometric embedding of $[a, b]$ into $\prod_{i=1}^\nu [0, l_i]$, where $\sum_{i=1}^\nu l_i = \rho(a, b)$. It now follows easily that there is a bound on $\#A$ in terms of ν and $\rho(a, b)$. \square

We will say that a median metric space is **interval-complete** if $[a, b]$ is complete for all $a, b \in M$. Clearly this holds if M is either complete or interval-compact.

Lemma 13.2.10. *A finite-rank interval-complete median metric space is interval-compact.*

Proof. This follows from Lemma 13.2.9, given that any complete totally bounded metric space is compact. \square

Lemma 13.2.11. *Let M be an interval-complete median metric space of finite rank. If $A \subseteq M$ is compact, then $\text{hull}(A)$ is compact.*

Proof. Let $\nu = \text{rank}(M)$. By Proposition 8.2.3, $\text{hull}(A) = J^\nu(A)$. It is therefore sufficient to show that the join, $J(A)$, is compact. To this end, let $(x_i)_i$ be any sequence in $J(A)$. Then $x_i \in [a_i, b_i]$ for some $a_i, b_i \in A$. After passing to a subsequence, we have $a_i \rightarrow a$ and $b_i \rightarrow b$ with $a, b \in A$. Since $a_i b_i x_i = x_i$, we have $\rho(x_i, a b x_i) \leq \rho(a_i, a) + \rho(b_i, b) \rightarrow 0$. Now $a b x_i \in [a, b]$, and $[a, b]$ is compact by Lemma 13.2.10, so again after passing to subsequence, we have $a b x_i \rightarrow y \in [a, b] \subseteq J(A)$. Therefore, $x_i \rightarrow y$. This shows that $J(A)$ is sequentially compact, hence compact. \square

13.3. Connectedness, geodesics and embeddings.

In order to introduce connectedness properties, we make another digression into linear median algebras.

Let L be a totally ordered set, equipped with a metric which makes it into a median metric space with the standard betweenness relation. Given $x, y \in L$, let $\tau(x, y) = \rho(x, y)$ if $x \leq y$ and $\tau(x, y) = -\rho(x, y)$ if $y \leq x$. Then τ is a cocycle: $\tau(x, y) + \tau(y, z) + \tau(z, x) = 0$ for all $x, y, z \in L$. (Up to permuting x, y, z , we have $x \leq y \leq z$, so $\tau(x, y) + \tau(y, z) = \rho(x, y) + \rho(y, z) = \rho(x, z) = -\tau(z, x)$.) It follows that if we fix any $a \in L$, then the map $[x \mapsto \tau(a, x)] : L \rightarrow \mathbb{R}$ is an isometric embedding. This shows:

Lemma 13.3.1. *L isometrically embeds into the real line.*

Note that L is complete if and only if the image of L is closed in \mathbb{R} . In this case, L is connected if and only if it has no adjacent pairs. Note also that if L is complete and bounded, then it is compact.

We say that a subset, $L \subseteq M$, of a median metric space is **monotone** if it is a subalgebra and intrinsically rank-1. In particular, if $x, y, z \in L$, then at least one of $x.y.z, y.z.x$ or $z.x.y$ holds. (As observed in Example (Ex3.3) of Subsection 3.3, the converse also holds if we assume in addition that L is not a 2-cube.) By the above, L is isometric to a subset of the real line. Note that by the continuity of

the median, the closure of a monotone set is monotone. If M is complete, then a closed bounded monotone set is compact. In particular, it is intrinsically the median interval between its two endpoints.

Recall that a **geodesic** in a metric space, (M, ρ) , is a path $\gamma : I \rightarrow \mathbb{R}$, where $I \subseteq \mathbb{R}$ is connected, such that $\rho(\gamma(t), \gamma(u)) = |t - u|$ for all $t, u \in \mathbb{R}$. In other words, γ is an isometric embedding of I into M . If M is a median metric space, it follows that it is a median monomorphism. In other words, γ is a monotone arc in the sense defined in Subsection 12.4.

Conversely, any monotone arc, γ , can be reparameterised as a geodesic. To see this, let $\theta : I \rightarrow \mathbb{R}$ be the embedding given by Lemma 13.3.1 for the induced metric on I . Then $\theta(I) \subseteq \mathbb{R}$ is connected, and $\gamma \circ \theta^{-1} : \theta(I) \rightarrow M$ is a geodesic.

In other words in a median metric space, geodesics and monotone arcs are the same class of paths up to reparameterisation.

We also note that (for any metric space) a geodesic between two points is the same up to reparameterisation as an arc whose rectifiable length is equal to the distance between its endpoints.

Definition. A **geodesic space** is a metric space in which any two points are connected by a geodesic.

(This is also sometimes called a “length space”.) We will give more discussion of geodesics in Sections 19 and 24.

Lemma 13.3.2. *Let M be an interval-complete median metric space. Then the following are equivalent:*

- (1): M has no adjacent pairs.
- (2): M is connected.
- (3): M is a geodesic space.

Proof. (3) \Rightarrow (2) is immediate, and (2) \Rightarrow (1) follows from Lemma 12.4.4.

For (1) \Rightarrow (3), suppose that M has no adjacent pairs. Let $a, b \in M$, so that $[a, b]$ is complete. Let $L \subseteq [a, b]$ be a maximal chain in $[a, b]$. By Lemma 12.4.3, L is closed. It is therefore intrinsically a complete median metric space of the type described above. Since L is maximal, it has no adjacent pairs (since such would also be an adjacent pair in M). By Lemma 13.3.1, it isometrically embeds in \mathbb{R} , and is therefore isometric to a compact real interval. The isomorphism from the interval is therefore a geodesic. \square

Remark. We note that (1) \Rightarrow (3) is also a direct consequence of Menger’s theorem. This says that any complete metric space, (M, ρ) , with the property that if $a, b \in M$ are distinct there is some $c \in M \setminus \{a, b\}$, with $\rho(a, b) = \rho(a, c) + \rho(c, b)$, then M is a geodesic space.

The following can easily be seen by putting together observations given above:

Lemma 13.3.3. *Let M be a median metric space. Suppose $a, b \in M$ with $\text{rank}([a, b]) = 1$. Then $[a, b]$ isometrically embeds into a compact real interval. If M is connected, we can take the embedding to be an isometry to its range.*

Lemma 13.3.4. *Let M be a median metric space of rank $\nu < \infty$. Let $Q \subseteq M$ be a ν -cube. Then $\text{hull}(Q)$ isometrically embeds into the l^1 product of ν compact real intervals. If M is connected, we can take the embedding to be an isometry.*

Proof. This follows from Lemmas 10.3.5 and 13.3.3. □

We also note the following generalisation of Lemma 13.2.8.

Lemma 13.3.5. *Let M be a median metric space, and let $a, b \in M$. Suppose $\text{rank}([a, b]) = \nu < \infty$. Then $[a, b]$ isometrically embeds in the l^1 product, $\prod_{i=1}^{\nu} [0, l_i]$, for some $l_i \geq 0$ with $\sum_{i=1}^{\nu} l_i = \rho(a, b)$.*

This can be deduced by a compactness argument using the fact that it is true for all finite subalgebras of $[a, b]$, similarly as in the proof of Proposition 15.3.1 of Section 15. We will not give details here since it can also be seen as a corollary of that result, as we point out in Subsection 15.3.

For the record, we also note the following immediate consequence of Lemma 12.4.12, which does not require interval completeness.

Lemma 13.3.6. *A connected finite-rank median metric space is geodesic.*

(I don't know if this holds in the infinite-rank case.)

Lemma 13.3.7. *Let M be an interval-complete median metric space. Then any non-empty closed convex subset of M is gated.*

Proof. Let $C \subseteq M$ be closed and convex, and let $a \in M$. Choose any $c \in C$. Let L be any maximal chain in $[a, c] \cap C$. By Lemma 12.4.3, L is closed. As observed after Lemma 13.3.1, L is compact, and so has a minimum, b . We claim that b is a gate for a in C . To see this, let $d \in C$, and let $e = abd$. Then $e \leq b$ and $e \in C$. Therefore $L \cup \{e\}$ is a chain in $[a, c] \cap C$, so $e = b$, so $a.b.d$ as required. □

Note also that b is the unique point of C which minimises $\rho(a, c)$.

We also note:

Lemma 13.3.8. *Let C be a gated convex subset of a median metric space, M . Then the gate map, $M \rightarrow C$ is 1-lipschitz.*

Proof. Let $x, y \in M$. Then $x.\omega x.\omega y$ and $y.\omega y.\omega x$. It follows from Lemma 13.2.1 that $\rho(\omega x, \omega y) \leq \rho(x, y)$. □

The following result generalises Lemma 11.5.5.

Lemma 13.3.9. *Let M be a complete median metric space. Let \mathcal{C} be any family of pairwise intersecting closed bounded convex subsets. Then $\bigcap \mathcal{C} \neq \emptyset$.*

Of course, it is enough to assume that just one of the convex subsets is bounded.

Proof. By Lemma 13.3.7, each element of \mathcal{C} is gated. The result therefore follows by a much more general statement about gated sets in a complete metric space, namely Lemma 22.2.2. □

13.4. Connections with CAT(0) and injective metric spaces.

We mention a few additional results about median metric spaces, though we refer to the original papers for the proofs.

Recall that a k -lipschitz median algebra is one which admits a metric such that the median map is k -lipschitz. By a *k -lipschitz path* we mean a k -lipschitz map of a real interval into the space. (For most applications, we could take this to be a geodesic.)

Theorem 13.4.1. *Let $k, l \geq 0$ and $\nu \in \mathbb{N}$. Let (M, ρ) be a k -lipschitz median algebra of rank ν , such that any two points are connected by an l -lipschitz path. Then M admits a canonical median metric, ρ' , inducing the same median structure, such that $\rho \leq \rho' \leq K\rho$, where $K \geq 1$ depends only on k, l and ν .*

If one drops the word “canonical” from the statement, this is proven in [Bo3] (cf. Proposition 15.3.1 here). A variation on this construction in [Ze] shows that one can obtain a canonical median metric.

Theorem 13.4.2. *Let (M, ρ) be a complete connected median metric space of rank $\nu < \infty$. Then M admits a canonical CAT(0) geodesic metric, σ , with $\rho \leq \sigma \leq \rho\sqrt{\nu}$. Moreover geodesics in (M, σ) are monotone arcs with respect to the median structure.*

Theorem 13.4.3. *Let (M, ρ) be a complete connected median metric space of rank $\nu < \infty$. Then M admits a canonical injective metric, σ , with $\rho \leq \sigma \leq \nu\rho$.*

These are proven in [Bo4] and [Bo7] respectively. A somewhat different construction for the injective metric has been given independently by Miesche [Mie2].

We will discuss CAT(0) metrics further in Section 18, and injective metrics in Section 25. Both CAT(0) and injective metric spaces are necessarily contractible. It therefore follows from either Theorem 13.4.2 or Theorem 13.4.3 that M is contractible.

One motivation for the above results is their application to asymptotic cones of coarse median spaces, which we will discuss in Subsection 24.3.

14. SUBMEDIAN RELATIONS

In this section, we consider subsets of median algebras with their induced ternary relations. These are generally termed “submedian”. Examples arise as submedian metric spaces as we discuss in Section 19, and from quasimedial graphs, discussed in Section 23. Some aspects of more general ternary betweenness relations will be discussed in Section 22.

It does not seem that the notion of “rank” can be cleanly dealt with in this context, except in the rank-1 case. In rank-1, such structures have been called “pretrees” and will be discussed in Subsection 14.2.

14.1. **Characterisation of submedian relations.**

Recall that the structure of any median algebra, M , is determined by the induced ternary betweenness relation, which we have denoted by $x.y.z$ for $x, y, z \in M$. Moreover, any subset of M inherits such a ternary relation. One can ask which ternary relations arise in this way: we will refer to such as “submedian”. There are certainly obstructions: such as the linear interpolation rule, though this is not in general sufficient. In section 6, we gave a procedure for verifying “tautological” identities. We can apply this to any ternary relation (without introducing medians into the proceedings) and make it into a recursive axiom system. This gives rise to a characterisation of submedian relations (see Proposition 14.1.5). As we note at the end, however, there is no equivalent finite set of axioms.

First, we give some definitions. By a **ternary set**, we mean a set T equipped with a ternary relation, denoted $x.y.z$ for $x, y, z \in T$, and such that $x.y.z \Leftrightarrow z.y.x$. To abbreviate notation, we will sometimes write t° to mean $x.y.z$ where $t = (x, y, z) \in T^3$. By default, if $X \subseteq T$ is any subset, we will take the induced ternary relation on X : if $x, y, z \in X$ then $x.y.z$ holds in X if and only if it holds in T . A subset $C \subseteq T$ is **convex** if $y \in C$ whenever $x, z \in C$ and $x.y.z$ holds in T . A **wall** is a bipartition of T into two non-empty convex subsets. We write $\mathcal{W}(T)$ for the set of all walls. A map $\phi : S \rightarrow T$ between two ternary sets is a **homomorphism** if $t^\circ \Rightarrow (\phi t)^\circ$ for $t \in S^3$. It is **full** if, conversely, $(\phi t)^\circ \Rightarrow t^\circ$. We write I for some fixed two-point median algebra, say $\{0, 1\}$. If $\phi : T \rightarrow I$ is an epimorphism, then $\{\phi^{-1}0, \phi^{-1}1\}$ is a wall.

Definition. We say that ternary set, T , is **submedian** if it admits an embedding, $\iota : T \hookrightarrow M$ into a median algebra M , such that if $x, y, z \in T$, then $x.y.z$ holds in T if and only if $\iota x.\iota y.\iota z$ holds in M .

In other words, ι is a full homomorphism.

Recall the procedure described in Subsection 6.2 for verifying median identities in a median algebra. We can apply a similar principle to a ternary set. Suppose we have variables, $x_1, y_1, z_1, \dots, x_n, y_n, z_n, x, y, z$ (not in general distinct). We can think of a relation $x.y.z$ as a “consequence” of relations $x_1.y_1.z_1, \dots, x_n.y_n.z_n$, if whenever we evaluate x_1, \dots, z_n in I such that $x_i.y_i.z_i$ hold for all i , then $x.y.z$ holds there too. The ternary relation on a submedian set will always be closed under such consequences. We can systematically make a list of such rules, and (by universal quantification) view it as an axiom system. We can express the assertion that a given ternary set is a model for this axiom system as follows.

Definition. A ternary set T is **complete** if given any $t \in T^3$ such that $(\phi t)^\circ$ holds for all homomorphisms $\phi : T \rightarrow I$, then t° holds in T .

We say that T is **finitely complete** if every finite subset is complete.

Thus “finitely complete” formalises what we earlier referred to as being “closed under consequences”.

Note that (by restricting homomorphisms) any complete set is finitely complete. In fact, we will see that these are equivalent (Proposition 14.1.4).

Note that finitely complete implies that $x.x.y$ holds for all $x, y \in T$, since this relation always holds in I .

We begin by considering finite sets.

Let X be a finite ternary set. Let $F(X) \supseteq X$ be the free median algebra on (the underlying set) X . Recall that, via the inclusion $X \subseteq F(X)$, we can identify $\mathcal{W}(F(X))$ with the set of proper bipartitions of X . Those bipartitions which are walls of X , in its ternary structure, give us a subset $\mathcal{W}_0 \subseteq \mathcal{W}(F(X))$ (so we can identify $\mathcal{W}_0 \equiv \mathcal{W}(X)$). Let $\mathcal{W}_1 = \mathcal{W}(F(X)) \setminus \mathcal{W}_0$. In other words, $W \in \mathcal{W}(F(X))$ lies in \mathcal{W}_1 if and only if we can find $x, y, z \in X$ such that W separates y from $\{x, z\}$ in $F(X)$, whereas $x.y.z$ holds in X . Let $G(X)$ be the quotient median algebra obtained by collapsing each of the walls of \mathcal{W}_1 (i.e. $G(X) = F(X)/\sim$, where $a \sim b$ if and only if every wall of $F(X)$ which separates a, b lies in \mathcal{W}_1). The quotient map, $\theta : F(X) \rightarrow G(X)$ is a homomorphism. Taking inverses gives us a bijection, $\theta^{-1} : \mathcal{W}(G(X)) \rightarrow \mathcal{W}_0$. We write $\theta = \theta_X : X \rightarrow G(X)$ for the restriction to X . This is a homomorphism. (For suppose $x, y, z \in X$ with $x.y.z$ but $\neg\theta.x.\theta.y.\theta.z$. Then there is some $W \in \mathcal{W}(G(X))$ separating θy from $\{\theta x, \theta z\}$. Now $\theta^{-1}W \in \mathcal{W}(F(X))$ separates y from $\{x, z\}$, giving the contradiction that $\theta^{-1}W \in \mathcal{W}_1$.)

Lemma 14.1.1. *Let $\phi : X \rightarrow M$ be a homomorphism of X into a median algebra, M . Then there is a unique homomorphism, $\pi : G(X) \rightarrow M$ with $\phi = \pi \circ \theta$.*

Proof. First we extend ϕ to a homomorphism, $\phi : F(X) \rightarrow M$. We claim that if $a, b \in F(X)$ with $\theta a = \theta b$, then $\phi a = \phi b$. For if not, there is some wall $W \in \mathcal{W}(M)$ separating $\phi a, \phi b$. Let W_0 be the preimage of this wall in $F(X)$. Since $\theta a = \theta b$, we have $W_0 \in \mathcal{W}_1$. In other words, we can find $x, y, z \in X$ with W_0 separating y from $\{x, z\}$ in $F(X)$, and such that $x.y.z$ holds in X . Thus, W separates ϕy from $\{\phi x, \phi z\}$ in M . But since ϕ is a homomorphism, we get the contradiction that $\phi x.\phi y.\phi z$ holds in M . This proves the claim.

We now define π by setting $\pi(\theta a) = \phi a$ for $a \in F(X)$. Since θ and ϕ are homomorphisms, so is π . The uniqueness of the map π is clear. \square

Lemma 14.1.2. *Let $t \in X^3$. Then $(\theta t)^\circ$ holds in $G(X)$ if and only if $(\phi t)^\circ$ holds in I for every homomorphism, $\phi : X \rightarrow I$.*

Proof. Suppose $(\theta t)^\circ$ holds. Given a homomorphism, $\phi : X \rightarrow I$, let $\pi : G(X) \rightarrow I$ be as given by Lemma 14.1.1, so that $\phi = \pi \circ \theta$. Since π is a homomorphism, we have $(\phi t)^\circ$ as required.

Conversely, suppose $(\theta t)^\circ$ fails. Writing $t = (x, y, z)$, there is a wall separating θy from $\{\theta x, \theta z\}$ in $G(X)$ (i.e. separating θy from $\theta x \theta y \theta z$). This gives a homomorphism, $\pi : G(X) \rightarrow I$ with $\pi \theta x = \pi \theta z \neq \pi \theta y$. In other words, $(\phi t)^\circ$ fails in I , where $\phi := \pi \circ \theta : X \rightarrow I$. \square

Corollary 14.1.3. *The map $\theta : X \rightarrow G(X)$ is full if and only if X is complete.*

Proof. These statements can be paraphrased respectively as “ $(\theta t)^\circ \Rightarrow t^\circ$ ” and “ $((\phi t)^\circ \text{ holds for all homomorphisms } \phi : X \rightarrow I) \Rightarrow t^\circ$ ”. These are equivalent, by Lemma 14.1.2. \square

Note that any homomorphism, $\phi : X \rightarrow Y$ between finite ternary sets induces a homomorphism, $\hat{\phi} : G(X) \rightarrow G(Y)$ such that $\hat{\phi} \circ \theta_X = \theta_Y \circ \phi$. This process respects compositions.

Now, let T be any ternary set. Let \mathcal{X} be the directed set of finite subsets of T , partially ordered by inclusion. If $X \subseteq Y \in \mathcal{X}$, we have an induced homomorphism, $\phi_{XY} : G(X) \rightarrow G(Y)$, extending the inclusion $X \hookrightarrow Y$. We let $G(T)$ be the direct limit of this system. We have a natural map $\theta = \theta_T : T \rightarrow G(T)$.

We can define a ternary operation $[(a, b, c) \mapsto abc]$ on $G(T)$ as follows. By definition of direct limit, there is some $X \in \mathcal{X}$, with respective representatives, a', b', c' , in $G(X)$. We set abc to be the class of $a'b'c'$ in the median structure on $G(X)$. Again, by the definition of direct limit, and the fact that all the maps ϕ_{XY} are median homomorphisms, it follows that this is well defined, independently of the choice of X and representatives, a', b', c' . Moreover, we see that $G(X)$ is a median algebra in this structure. Since the maps θ_X are all homomorphisms, it follows that θ_T is also. (This notation is consistent with that already defined if T happens to be finite.)

If T is finitely complete, then (by Corollary 14.1.3), each of the maps, θ_X is full. It follows that θ_T is full. (For if $x, y, z \in T$ with $\theta_T x.\theta_T y.\theta_T z$ in $G(T)$, then we have $\theta_X x.\theta_X y.\theta_X z$ for some $X \in \mathcal{X}$ with $x, y, z \in X$, and so $x.y.z$ holds in X , hence also in T .)

We now have:

Proposition 14.1.4. *Let T be a ternary set. The following are equivalent.*

- (1): T is complete,
- (2): T is finitely complete,
- (3): T admits a full homomorphism to a median algebra.

Proof. We have already observed that (1) \Rightarrow (2) is trivial. For (2) \Rightarrow (3), we have just noted that if T is finitely complete, then $\theta_T : T \rightarrow G(T)$ is full.

For (3) \Rightarrow (1), suppose $\psi : T \rightarrow M$ is a full homomorphism to a median algebra M . Suppose that $x, y, z \in T$ and that $x.y.z$ fails in T . Then $\psi x.\psi y.\psi z$ fails in M . Thus, there is a median homomorphism, $\pi : M \rightarrow I$ with $\pi\psi x = \pi\psi y \neq \pi\psi z$, and so $\phi x.\phi y.\phi z$ fails in I , where $\phi = \pi \circ \psi : T \rightarrow I$. This shows that T is complete. \square

Now consider the following property of a ternary set:

(I): If $x, y \in T$ with $x.y.x$, then $x = y$.

Note that if $\phi : T \rightarrow M$ is a full map to a median algebra, M , and T satisfies (I), then T is injective. (Since, $\phi x = \phi y \Rightarrow \phi x.\phi y.\phi x \Rightarrow x.y.x \Rightarrow x = y$.)

In view of Proposition 14.1.4, we have now shown:

Proposition 14.1.5. *Let T be a ternary set. The following are equivalent:*

- (1): T is submedian,
- (2): every finite subset of T is submedian,
- (3): T is complete and satisfies (I).

It would be nice to have a simple axiomatisation of the submedian property. This can be achieved for small finite X . We assume that X is a ternary set which satisfies (I), and that $a.a.b$ holds for all $a, b \in X$. By Proposition 14.1.5, it is sufficient to find additional constraints which imply completeness.

For $\#X \leq 3$, there is no other constraint.

For $\#X = 4$, the linear interpolation rule is sufficient. Recall that this says that for $a, b, c, d \in X$,

$$a.b.c \ \& \ a.c.d \Rightarrow a.b.d \ \& \ b.c.d.$$

We claim this implies completeness in this case. To see this, write $X = \{x, y, z, p\}$, and suppose $\neg x.p.y$. We define $\phi : X \rightarrow I$ by first setting $\phi x = \phi y = 0$ and $\phi p = 1$. We want to define ϕz . If $\neg x.p.z \ \& \ \neg y.p.z$ we set $\phi z = 0$, and easily see that this is a homomorphism. Otherwise, we set $\phi z = 1$. To see this is a homomorphism, we can suppose (up to swapping x, y) that $x.p.z$ holds. Now $x.z.y$ would imply $x.p.z.y$, and $p.y.z$ would imply $x.p.y.z$, both giving $x.p.y$ contrary to our assumption. Thus, $\neg x.z.y \ \& \ \neg p.y.z$, and it now follows that ϕ is a homomorphism as claimed.

For $\#X = 5$, we have three rules, namely:

$$a.b.c \ \& \ b.c.d \ \& \ a.d.e \Rightarrow b.c.e$$

$$a.c.b \ \& \ a.d.b \ \& \ c.e.d \Rightarrow a.e.b$$

$$a.b.c \ \& \ a.b.d \ \& \ c.e.d \Rightarrow a.b.e.$$

Note that the first of these ((Ex6.1) of Subsection 6.2) implies the linear interpolation rule (respectively setting $a = b$ and $c = d$). A somewhat more laborious argument, along the lines of the previous case, shows that these three rules together imply completeness in this case.

For $\#X \geq 6$ more rules are needed. (For example, the rule given by Lemma 3.2.8 is not implied by any combination of the above — no two of the input relations have any two points in common.)

In fact, it is not possible to give any finite set of rules in general, as the following example illustrates.

Example. We construct a ternary set as follows. Let $n \in \mathbb{N} \setminus \{0\}$. Let $\Sigma_n = \{0, 1\}^n$ be the set of sequences in $\{0, 1\}$ of length n , and let $\Sigma_{<n} = \bigcup_{m=0}^{n-1} \Sigma_m$. (Here these just serve as indexing sets.) Let $X = X_n$ be a set of $2^{n+2} + 2$ elements written as $a, b, c, \{d_\sigma, e_\sigma \mid \sigma \in \Sigma_n\}$ and $\{p_\sigma \mid \sigma \in \Sigma_{<n} \cup \Sigma_n\}$. If σ has length 0, we write $p = p_\sigma$. Let \mathcal{R} be the set of relations of one of the following forms: $p.a.b, a.b.d_\sigma, b.c.e_\sigma, e_\sigma.d_\sigma.p_\sigma$ for $\sigma \in \Sigma_n, p_{\sigma 0}.p_\sigma.p_{\sigma 1}$ for $\sigma \in \Sigma_{<n}$, and finally $x.x.y$ for all $x, y \in X$.

(We are regarding $x.y.z$ as equivalent to $z.y.x$.) Let \mathcal{R}' be \mathcal{R} with the relation $a.b.c$ added.

One can check that (X, \mathcal{R}') is submedian (see below). However, (X, \mathcal{R}) is not. In fact, we claim that the relation $a.b.c$ is a consequence of \mathcal{R} .

To see this, suppose $\phi : (X, \mathcal{R}) \rightarrow I$ is a homomorphism. Write $x\downarrow, x\uparrow$ respectively for $\phi x = 0$ and $\phi x = 1$. Suppose, for contradiction, that $a\downarrow, b\uparrow, c\downarrow$. Now $a.b.d_\sigma$ gives $d_\sigma\uparrow$, $b.c.e_\sigma$ gives $e_\sigma\downarrow$, and $e_\sigma.d_\sigma.p_\sigma$ gives $p_\sigma\uparrow$ for all $\sigma \in \Sigma_n$. Now, by backwards induction on the length of σ , $p_{\sigma_0}.p_\sigma.p_{\sigma_1}$ gives $p_\sigma\uparrow$ for all $\sigma \in \Sigma_{<n}$. In particular, $p\uparrow$. Now $p.a.b$ gives the contradiction $a\uparrow$. This proves the claim.

Of course, we can swap \uparrow and \downarrow in the above. However, any other assignment of \uparrow and \downarrow to a, b, c does not, in itself, place any restriction on any p_σ . After considering various cases, one can see that any assignment of \uparrow and \downarrow to any three points of X not related by \mathcal{R}' extends to an assignment on all of X consistent with \mathcal{R}' . By Proposition 14.1.5, this then shows that (X, \mathcal{R}') is submedian.

However, if $\mathcal{R}_0 \subseteq \mathcal{R}$ is any proper subset, then the chain of reasoning is broken: one can check that the assignment $a\downarrow, b\uparrow, c\downarrow$ can be extended to an assignment on all of X consistent with \mathcal{R}_0 .

In particular, it now follows that any proper subset of X (with the ternary relation induced by \mathcal{R}) is submedian.

The above example shows that the class of ternary sets which are not submedian is not axiomatisable (in the first-order vocabulary with one ternary relation). This follows by a standard argument of model theory. Suppose that \mathcal{S} were a set of first-order sentences which characterise such ternary sets. Given any $i \in \mathbb{N}$, let θ_i be a sentence which asserts that any subset of cardinality at most i is submedian. Let \mathcal{S}' be \mathcal{S} together with all the sentences θ_i . Now any finite subset of \mathcal{S}' has a model, namely our example X_n for large enough n . Thus, any finite subset of \mathcal{S}' is consistent, and so \mathcal{S}' is consistent. By the Model Existence Theorem, \mathcal{S}' has a model. This is a ternary set which is not submedian, but every finite subset thereof is submedian. This contradicts Proposition 14.1.5.

In particular, it follows that the property of being submedian is not finitely axiomatisable.

14.2. Pretrees.

We now restrict to the rank-1 case. We first give another description of a rank-1 median algebra. The rank-1 hypothesis is equivalent to saying that every interval is totally ordered by the relation \leq . Such spaces have been studied for some time under a variety of different names. Here we will use the term *median pretree*. They are typically formulated by a different set of axioms. Here is one version.

Let T be a set equipped with a ternary relation, denoted $a.b.c$ for $a, b, c \in T$. Consider the following properties:

(R1): $(\forall a, b, c \in T) (a.b.c \Leftrightarrow c.b.a)$

(R2): $(\forall a, b \in T) a.a.b$

(R3): $(\forall a, b, c \in T) ((a.b.c \ \& \ a.c.b) \Rightarrow b = c)$

(R4): $(\forall a, b, c, d \in T) (a.b.c \Rightarrow a.b.d \text{ or } d.b.c)$

(R5): $(\forall a, b, c \in T)(\exists d \in T) (a.d.b \ \& \ b.d.c \ \& \ c.d.a)$

Suppose that T is a rank-1 median algebra. Interpreting $a.b.c$ as the induced betweenness relation, one can easily check that properties (R1)–(R5) are satisfied. (For (R4), set $e = acd$. Since $[a, b]$ is totally ordered, we have either $a.b.e.c$ or $a.e.b.c$. In the former case (given $a.e.d$), we have $a.b.e.d$ so $a.b.d$. In the latter case (given $c.e.d$), we have $c.b.e.d$ so $d.b.c$.)

Conversely, suppose that T satisfies (R1)–(R5). We claim that it admits a unique structure as a rank-1 median algebra inducing the given betweenness relation.

We first make the following observations which are simple consequences of Properties (R1)–(R4). Suppose $a, b, c, d \in T$. We have:

(Q1): $a.b.c \ \& \ b.c.d \ \& \ b \neq c \Rightarrow a.b.c.d,$

(Q2): $a.b.c \ \& \ a.c.d \Rightarrow a.b.c.d,$

(Q3): $a.b.d \ \& \ a.c.d \Rightarrow (a.b.c.d \text{ or } a.c.b.d).$

We set $I(a, b) = \{x \in T \mid a.x.b\}$. We need to check that the family $\{I(a, b)\}_{a, b \in T}$ satisfies the axioms (I1)–(I4) of Subsection 4.1. Properties (I1) and (I2) are immediate from (R1)–(R3), and Property (I3) follows from (Q3) above.

To verify (I4), we need to show that the point d given by (R5) is unique. So suppose e were another such point. Now $d, e \in I(a, b)$, so by (Q3), we can assume that $a.d.e.b$ holds. Since $d, e \in I(a, c)$, we must have $a.d.e.c$ (again using (Q3)). Also, since $d, e \in I(b, c)$, we must have $b.e.d.c$. Now $d.e.c \ \& \ e.d.c$ gives $d = e$ as required.

Theorem 4.1.1 now tells us that T has the structure of a median algebra, with $[a, b] = I(a, b)$. Moreover T must have rank 1, since intervals are totally ordered by the above. (More directly, any 2-cube, $\{a, b, c, d\}$ would violate axiom (R4).)

In summary, this shows that axioms (R1)–(R5) give us an equivalent formulation of the notion of a rank-1 median algebra.

More generally, we have the following notion.

Definition. A *pretree* is a set equipped with a ternary relation satisfying (R1)–(R4) above.

Clearly any subset of a pretree is a pretree with the induced ternary relation. In fact, any pretree arises in a canonical way as a subset of a median pretree. One statement of this is as follows.

Suppose that T is a median pretree, and that $P \subseteq T$. Consider the following properties.

(T1): If $x \in T$, then there exist $a, b, c \in P$ with $x = abc$.

(T2): If $x, y \in T$ are distinct, then $[x, y]_T \cap P \neq \emptyset$.

We claim:

Proposition 14.2.1. *Let P be a pretree. Then we can embed P as subset of a median pretree T , satisfying (T1) and (T2), such that if $a, b, c \in P$, then $a.b.c$ holds in P if and only if $a.b.c$ holds in T .*

Moreover, if $T' \supseteq P$ is another such median pretree, then there is a unique isomorphism from T to T' fixing P .

Note that if P is finite or countable, then (by (T1)) so is T .

We first consider the case of a finite pretree, Π . (This will also serve to illustrate a more general construction in Subsection 23.4.)

Given $a, b \in \Pi$, write $[a, b] = [a, b]_{\Pi} = \{x \in \Pi \mid a.x.b\}$. We say that a, b are **adjacent** if $[a, b] = \{a, b\}$. Applying (Q1) and (Q2), we see that, in general, we have $[a, b] = \{a_0, a_1, \dots, a_n\}$, where $a_0 = a$, $a_n = b$, $a_0.a_1 \dots .a_n$ and a_i is adjacent to a_{i+1} for all i . Conversely, suppose we have $a_0, a_1, \dots, a_n \in \Pi$ such that for all i , a_i is adjacent to a_{i+1} but not to a_{i+2} . Then (applying also (Q3)) we see that $a_0.a_1 \dots .a_n$ holds and that $[a, b] = \{a_0, a_1, \dots, a_n\}$.

Let $\Gamma = \Gamma(\Pi)$ be the adjacency graph; that is, it has vertex set Π and edges determined by adjacency. From the above, we see that Γ is connected. Let \mathcal{B} be the set of all maximal cliques of Γ . We construct a bipartite graph, $\Theta = \Theta(\Pi)$, with vertex set, $V(\Theta) = \Pi \sqcup \mathcal{B}$, and with $a \in \Pi$ deemed adjacent to $B \in \mathcal{B}$ if $a \in B$. From the above, one sees that Θ has no embedded cycles. In other words, Θ is a finite simplicial tree. As such, $V(\Theta)$ has the structure of a median pretree and induces the original pretree structure on Π . (This median pretree almost proves the existence part of Proposition 14.2.1 in the finite case, though we would need to remove the degree-2 vertices of $V(\Theta) \setminus \Pi$ in order to satisfy (T1). However, in the argument that follows, we will leave Θ as it is.)

Remark. We see, in fact, that the elements of \mathcal{B} are precisely the “blocks” of Γ ; that is, the maximal 2-vertex-connected subgraphs. For any connected graph, one can similarly define a bipartite tree from the set of blocks. If the blocks are all complete graphs, as in our case, the graph is called a **block graph**. Such a graph is an example of a “quasimedial graph”, to be defined in Subsection 23.1. To any quasimedial graph one can canonically associate a median graph (see Theorem 23.4.11), which in our case is precisely the bipartite tree.

Suppose $\Pi' \subseteq \Pi$ is a non-empty subset. Then we have a retraction $\Theta(\Pi) \rightarrow \Theta(\Pi')$. This is obtained, first by collapsing every edge of $\Theta(\Pi)$ disjoint from Π' , and then collapsing those edges which are incident on degree-1 vertices which do not lie in Π' . Note that the induced map on vertex sets, $V(\Theta(\Pi)) \rightarrow V(\Theta(\Pi'))$, is a median homomorphism which restricts to the identity on Π' .

We now move on to the existence part of Proposition 14.2.1 in the general case.

Let P be any pretree. Let \mathcal{A} be the set of all finite subsets of P . To any $A \in \mathcal{A}$ we can associate a bipartite tree, $\Theta(A)$, taking the induced pretree structure on A . To avoid ambiguity, given any $a \in A$, we will write $a_A \in V(\Theta(A))$ for the

corresponding vertex of $\Theta(A)$. Note that if $A \subseteq B \in \mathcal{A}$, then we have a retraction $\Theta(B) \rightarrow \Theta(A)$.

We define a relation, \sim , on the set of ordered triples, $\mathcal{T}(P)$ by setting $(a, b, c) \sim (a', b', c')$ to mean that $a_A b_A c_A = a'_A b'_A c'_A$ for all $A \in \mathcal{A}$ with $a, b, c, a', b', c' \in A$. This is an equivalence relation on P^3 (since $a_A b_A c_A = a'_A b'_A c'_A$ implies $a_B b_B c_B = a'_B b'_B c'_B$ for any $B \supseteq A$). We write $[(a, b, c)]$ for the equivalence class of (a, b, c) . (Note that $(a, b, c) \sim (b, a, c) \sim (b, c, a)$, so the order of the entries is unimportant.) We write $T = T(P) = P^3 / \sim$ for the quotient space.

Given $x \in T$, let $\mathcal{A}(x)$ be the set of $A \in \mathcal{A}$ such that $x = [(a, b, c)]$ for some $a, b, c \in A$. In this case, we write $x_A = a_A b_A c_A$. This is well defined by the definition of \sim . If $X \subseteq T$, we write $\mathcal{A}(X) = \bigcap_{x \in X} \mathcal{A}(x)$.

Given $a \in P$, let $a_T = [(a, a, a)] \in T$. This gives us an injective map, $[a \mapsto a_T] : P \rightarrow T$. Note that $(a_T)_A = a_A a_A a_A = a_A$. Also, $\mathcal{A}(a_T) = \{A \in \mathcal{A} \mid a \in A\}$.

Note that by construction, if $a, b, c \in A \in \mathcal{A}$, then $a.b.c$ holds in P if and only if $a_A.b_A.c_A$ holds in $\Theta(A)$. For later reference, we note:

Lemma 14.2.2. *Suppose $x \in T$, $A \in \mathcal{A}(x)$ and $x_A = d_A$ for some $d \in P$. Then $x = d_T$.*

Proof. Let $x = [(a, b, c)]$ with $a, b, c \in A$. Now $a_A.d_A.b_A$, $b_A.d_A.c_A$ and $c_A.d_A.a_A$ all hold in $\Theta(A)$, so $a.d.b$, $b.d.c$ and $c.d.a$ hold in P , so $a_B.d_B.b_B$, $b_B.d_B.c_B$ and $c_B.d_B.a_B$ hold for any $B \in \mathcal{A}$ with $a, b, c, d \in B$. It follows that $(d, d, d) \sim (a, b, c)$, and so $d_T = [(d, d, d)] = [(a, b, c)] = x$. \square

We now define a ternary relation on T by writing $x.y.z$ to mean that $x_A.y_A.z_A$ holds in $\Theta(A)$ for all $A \in \mathcal{A}(\{x, y, z\})$. Note that if $a, b, c \in P$, then $a.b.c$ holds in P if and only if $a_T.b_T.c_T$ holds in T .

Taking the contrapositive, if $x, y, z \in T$ fails, then there is some $A \in \mathcal{A}(\{x, y, z\})$ for which $x_A.y_A.z_A$ fails. Moreover, if $B \in \mathcal{A}$ with $A \subseteq B$, then $x_B.y_B.z_B$ also fails (since the retraction $\Theta(B) \rightarrow \Theta(A)$ is a median homomorphism). We see that we can take $A \in \mathcal{A}(\{x, y, z\})$ above to contain any given finite subset of P .

Lemma 14.2.3. *With the above ternary relation, T is a median pretree.*

Proof. We need to verify properties (R1)–(R5). Properties (R1)–(R3) are immediate from the corresponding statements for $\Theta(A)$ for all $A \in \mathcal{A}$.

To verify (R4), suppose for contradiction that $x, y, z, w \in T$ with $x.y.z$, but that $x.y.w$ and $w.y.z$ both fail. Therefore $x_A.y_A.w_A$ and $w_B.y_B.z_B$ both fail for some $A, B \in \mathcal{A}$. Let $C = A \cup B \in \mathcal{A}$. Then $x_C.y_C.w_C$ and $w_C.y_C.z_C$ both fail, but $x_C.y_C.z_C$ holds, contradicting (R4) for $\Theta(C)$.

To verify (R5), let $x, y, z \in T$. We want to find some $w \in T$, such that $x.w.y$, $y.w.z$ and $z.w.x$ all hold. We can suppose that none of the relations $x.y.z$, $y.z.x$ or $z.x.y$ hold. Thus, there is some $A \in \mathcal{A}(\{x, y, z\})$ for which $x_A.y_A.z_A$, $y_A.z_A.x_A$ and $z_A.x_A.y_A$ all fail. We have $x_A = a_A^1 a_A^2 a_A^3$, $y_A = b_A^1 b_A^2 b_A^3$, and $z_A = c_A^1 c_A^2 c_A^3$ where $a^i, b^j, c^k \in A$ for all i, j, k . Let $m = x_A y_A z_A$ in $\Theta(A)$. Since $\Theta(A)$ is a tree, one can certainly find i, j, k so that $m.x_A.a_A^i$, $m.y_A.b_A^j$ and $m.z_A.c_A^k$ all hold. Write

$a = a^i$, $b = b^j$ and $c = c^k$, and set $w = [(a, b, c)] \in T$. By definition, we have $w_A = a_A b_A c_A = m$ in $\Theta(A)$, and so $w_A \cdot x_A \cdot a_A$, $w_A \cdot y_A \cdot b_A$ and $w_A \cdot z_A \cdot c_A$ all hold in $\Theta(A)$. Therefore, $w_A = x_A y_A z_A$.

We claim that $x \cdot w \cdot y$ holds in T . For if not, there is some $B \in \mathcal{A}$ for which $x_B \cdot w_B \cdot y_B$ fails. We can assume that $A \subseteq B$. By definition of w , we have $w_B = a_B b_B c_B$ in $\Theta(B)$. Since $\Theta(B)$ is a tree, we see easily that we must have either $a_B \cdot w_B \cdot x_B$ or $b_B \cdot w_B \cdot y_B$ (or both). Without loss of generality, we can assume $a_B \cdot w_B \cdot x_B$. Since $A \subseteq B$, we then get $a_A \cdot w_A \cdot x_A$. Together with $w_A \cdot x_A \cdot a_A$, this implies $w_A = x_A$. But now $y_A \cdot w_A \cdot z_A$, so we get $y_A \cdot x_A \cdot z_A$, contradicting an earlier assumption. This shows that $x \cdot w \cdot y$ holds as claimed.

Similarly, we have $y \cdot w \cdot z$ and $z \cdot w \cdot x$. This proves (R5). □

We can now use the usual notation for medians in T .

Lemma 14.2.4. *Suppose $x, y \in T$ are distinct. Then there is some $c \in P$ with $x \cdot c_T \cdot y$.*

Proof. Since $x \neq y$, there is some $A \in \mathcal{A}(\{x, y, z\})$ with $x_A \neq y_A$. Since $\Theta(A)$ is a tree, we can find $a, b \in A$ with $a_A \cdot x_A \cdot y_A \cdot b_A$. Moreover, since $\Theta(A)$ is bipartite, there is some $c \in A$ with $x_A \cdot c_A \cdot y_A$. Thus, $a_A \cdot c_A \cdot b_A$, and so $a \cdot c \cdot b$ holds in P .

We claim that $x \cdot c_T \cdot y$ holds in T . For if not, there is some $B \in \mathcal{A}$ for which $x_B \cdot c_B \cdot y_B$ fails. We can assume that $A \subseteq B$. Now $a_B \cdot c_B \cdot b_B$ holds, so (since $\Theta(B)$ is a tree) we see easily that either $a_B \cdot c_B \cdot x_B$ or $y_B \cdot c_B \cdot b_B$ (or both) hold in $\Theta(B)$. These respectively imply $a_A \cdot c_A \cdot x_A$ or $y_A \cdot c_A \cdot b_A$. Given $a_A \cdot x_A \cdot c_A \cdot y_A \cdot b_A$, these in turn imply $x_A = c_A$ or $y_A = b_A$. Lemma 14.2.2 now gives $x = c_T$ or $y = c_T$. Either way, we have $x \cdot c_T \cdot y$ (formally a contradiction) proving the claim. □

Now if $a, b, c \in P$, then $a \cdot b \cdot c$ hold in P if and only if $a_A \cdot b_A \cdot c_A$ holds in $\Theta(A)$ for all $A \in \mathcal{A}$ with $a, b, c \in A$, which in turn holds if and only if $a_T \cdot b_T \cdot c_T$ holds in T .

Also note that if $x = [(a, b, c)] \in T$, then $x_A = a_A b_A c_A$ for all $A \in \mathcal{A}$, and so $x = a_T b_T c_T$ in T .

We can now identify a with a_T for all $a \in \Pi$, so that $P \subseteq T$, and T induces the original ternary relation on P . In view of the above, we see:

Lemma 14.2.5. *$P \subseteq T$ satisfies properties (T1) and (T2).*

This proves the existence part of Proposition 14.2.1. For uniqueness, suppose conversely that $P \subseteq T$ satisfies (T1) and (T2). The following observations are simple consequences of the fact that any finite subset of T embeds in the vertex set of a finite simplicial tree.

Suppose $a_1, a_2, a_3, b_1, b_2, b_3, c \in P$. Let $x = a_1 a_2 a_3$ and $y = b_1 b_2 b_3$. Then $x \cdot c \cdot y$ holds in T if and only if there exist $i \neq j$ and $k \neq l$ such that $a_i \cdot c \cdot b_k$, $a_i \cdot c \cdot b_l$, $a_j \cdot c \cdot b_k$ and $a_j \cdot c \cdot b_l$ all hold in P . Also, by (T2), if $x, y \in T \setminus P$, then $x \neq y$ if and only if there is some $d \in P$ with $x \cdot d \cdot y$. Note that all the above are all recognisable in P .

Suppose that $x, y \in T$ and $z \in T \setminus P$. Then $x \cdot z \cdot y$ fails in T if and only if there is some $d \in P$ with $x \cdot d \cdot z$ and $y \cdot d \cdot z$. (Take d so that $z \cdot d \cdot x y z$ by (T2).)

Now suppose that $T' \supseteq P$ also satisfies (T1) and (T2). If $a_1, a_2, a_3, b_1, b_2, b_3 \in P$, then $a_1 a_2 a_3 = b_1 b_2 b_3$ holds in T if and only if it holds in T' . In view of (T1), this gives a bijection between T and T' , fixing P . From the previous paragraph, this respects the ternary relation, and is therefore a median isomorphism. It is clearly the unique such isomorphism.

This completes the proof of Proposition 14.2.1.

15. \mathbb{R} -TREES

The notion of an \mathbb{R} -tree was introduced in [MoS]. It has since become a central notion in geometric group theory.

Here we set some of the basic ideas in the context of median metric spaces. We go on to show that any finitely colourable median metric space isometrically embeds into a direct product of \mathbb{R} -trees, and complete the proof of Lemma 13.3.5. We also give a brief account of Guirardel cores at the end of the section.

15.1. Characterisations of \mathbb{R} -trees.

There are many equivalent ways of defining a \mathbb{R} -tree. For example:

Definition. An \mathbb{R} -tree is a connected median metric space of rank 1.

We will relate this to some more common descriptions below. (See Lemmas 15.1.2 and 15.1.4.) First we make a few basic observations.

Let (T, ρ) be an \mathbb{R} -tree. Let $a, b \in T$. By Lemma 13.3.3, $[a, b] \subseteq T$ is isometric to the real interval $[0, \rho(a, b)]$. By Lemma 12.7.2, $[a, b]$ is the unique topological arc in T from a to b . Note that it follows that any connected subset of T is convex, and intrinsically an \mathbb{R} -tree. We also note that T is a “median pretree”, that is, it satisfies the axioms (R1)–(R5) given in Subsection 14.2. Also, by Lemma 13.2.6, the metric completion of an \mathbb{R} -tree is an \mathbb{R} -tree (though one can give more direct proofs of this, for example, using Lemma 15.1.4 below).

Here is a basic construction that can be used to simplify arguments about \mathbb{R} -trees.

Let T be an \mathbb{R} -tree, and let $A \subseteq T$ be a finite subset. Let $\Pi = \langle A \rangle \subseteq T$, and let $\Delta(\Pi)$ be its realisation. Then $\Delta(\Pi)$ is a simplicial tree, with edge set naturally identified with $\mathcal{W}(\Pi)$. In fact, if $W \in \mathcal{W}(\Pi)$, then $\mathcal{E}(W)$ (as defined in Subsection 11.5) consists of a single 1-cell, $\{c, d\}$, of Π , and we set $w(W) = \rho(c, d)$ to be the width of W . In this way, we can put a median metric on $\Delta(\Pi)$ as described in Subsection 13.1, so that the edge of $\Delta(\Pi)$ with endpoints $\{c, d\}$ is isometric to $[0, \rho(c, d)] \subseteq \mathbb{R}$. This is also isometric to $[c, d]_T$, so combining these identifications, we get a natural map of $\Delta(\Pi)$ into T which is an isometry on each edge. In fact, this map of $\Delta(\Pi)$ is an isometric embedding. To see this, note that if $a, b \in \Pi$, then $[a, b]_{\Delta(\Pi)}$ is by construction isometric to $[0, \rho(a, b)]$ and so also to $[a, b]_T$. We also see that $\Delta(\Pi) \subseteq T$ is precisely the join, $J(A)$, of A in T . By Proposition 8.2.3, $\text{hull}(A) = J(A)$ (though of course, this can be seen much more directly, in this case). In summary, we have shown:

Lemma 15.1.1. *Let T be an \mathbb{R} -tree, and let $A \subseteq T$ be finite. Then $\text{hull}(A)$ can be naturally identified with the simplicial tree, $\Delta(\langle A \rangle)$, with vertex set $\langle A \rangle \subseteq T$.*

(See Proposition 17.4.1 for a related result in a more general setting.)

Here is an equivalent formulation of the notion of an \mathbb{R} -tree.

Lemma 15.1.2. *Let T be a metric space. Then T is an \mathbb{R} -tree if and only if for all $a, b \in T$ there is a unique topological arc in T from a to b , and this arc is isometric to a real interval.*

Proof. We have already seen than any \mathbb{R} -tree satisfies these conditions.

Conversely, suppose that (T, ρ) is any metric space with this property.

Given $a, b \in T$, we write $I(a, b)$ for the unique arc between them. Since this is isometric to a real interval, we must have $I(a, b) \subseteq [a, b]_\rho := \{x \in T \mid \rho(a, b) = \rho(a, x) + \rho(x, b)\}$. Now T contains no topologically embedded circle. From this, one can see easily that if $a, b, c \in T$, then $I(a, b) \cap I(b, c) \cap I(c, a) = \{d\}$ for some $d \in T$. It now follows that if $c \in [a, b]_\rho$, then $I(a, b) = I(a, c) \cup I(b, c)$, and so in that case $c = d$. In other words $[a, b]_\rho = I(a, b)$. Therefore, T is a median metric space, and $[a, b] = [a, b]_\rho$. From this it follows that T has rank 1. By definition, T is an \mathbb{R} -tree. □

Another characterisation of \mathbb{R} -trees arises from the following.

Definition. A metric space (M, ρ) is **weakly 0-hyperbolic** if it satisfies:

(FP): For all $a, b, c, d \in M$, we have:

$$\rho(a, b) + \rho(c, d) \leq \max\{\rho(a, c) + \rho(b, d), \rho(a, d) + \rho(b, c)\}.$$

We refer to (FP) as the “four-point condition”. It can be equivalently expressed by the saying that the largest two of the three “distance sums”, $\rho(a, b) + \rho(c, d)$, $\rho(a, c) + \rho(b, d)$ and $\rho(a, d) + \rho(b, c)$ are equal.

It can also expressed in terms of Gromov products (as defined in Subsection 13.2). Namely, a metric space is weakly 0-hyperbolic if and only if for all $a, b, c, p \in M$, we have

$$\langle a, b : p \rangle \geq \min\{\langle a, c : p \rangle, \langle b, c : p \rangle\}.$$

Remark. We use the term “weakly” because “hyperbolic” is usually taken to imply that the space is also a geodesic space, which we are not generally assuming here. See Subsection 24.1 for more discussion of this.

It is easily seen that any \mathbb{R} -tree is (weakly) 0-hyperbolic. (For example, by Lemma 15.1.1, it is sufficient to check this for a simplicial tree with at most four extreme points.)

Lemma 15.1.3. *Any weakly 0-hyperbolic metric space isometrically embeds into an \mathbb{R} -tree.*

Proof. This is a well known standard result, and we only outline the argument.

Let M be weakly 0-hyperbolic, and fix any basepoint, $p \in M$. Given any $a \in M$, let $I(a)$ be a copy of the real interval, $[0, \rho(a, p)]$, let $D = \bigsqcup_{a \in M} I(a)$. We have a natural map, $\sigma : D \rightarrow [0, \infty)$, such that $\sigma|_{I(a)}$ is the identification $I(a)$ with $[0, \rho(a, p)]$. Given $a \in M$, write $\bar{a} \in I(a)$ for the final point of $I(a)$. Thus $\sigma(\bar{a}) = \rho(a, p)$. Note that $I(p) = \{\bar{p}\}$.

We now define a relation, \sim , on D , by writing $x \sim y$ to mean that $\sigma(x) = \sigma(y) \leq \langle a, b : p \rangle$, where $x \in I(a)$ and $y \in I(b)$. This is clearly symmetric. Since $\langle a, a : p \rangle = \rho(a, p)$, we have $x \sim x$. Moreover, from (FP) (formulated in terms of Gromov products) we see that \sim is transitive. Thus, \sim is an equivalence relation, and we write $T = D/\sim$. In other words, for each a, b , we have identified the initial segments of $I(a)$ and $I(b)$ corresponding to $[0, \langle a, b : p \rangle]$. Note that the initial points of these intervals all get identified so some basepoint of T .

Suppose that $A \subseteq T$ is finite. We can perform the above construction restricting to the set $A \cup \{p\} \subseteq T$. In other words, we can set $D(A) = \bigsqcup_{a \in A} I(a)$, and let $T(A) = D(A)/\sim \subseteq T$. In this case, $T(A)$ is combinatorially a simplicial tree. Therefore, if we are only considering a finite number of points at any given moment, we can pretend we are in a simplicial tree. This makes the observations below more transparent.

Let $x, y \in T$. Choose representatives $x' \in I(a)$ and $y' \in I(b)$, and set $\hat{\rho}(x, y) = \sigma(x') + \sigma(y') - 2\langle a, b : p \rangle \geq 0$. One checks that this is well defined, and that $\hat{\rho}$ is a metric on T . (We only need to consider finitely many intervals, $I(a)$, at a time, so we can pretend we are in a simplicial tree.)

Given $x, y \in T$, we again choose representatives, $x', y' \in D$, as above. Then $x, y \in T(\{a, b\}) \subseteq T$, and we let $J(x, y) \subseteq T(\{a, b\})$ be the arc between them in $T(\{a, b\})$. This gives us an arc, $J(x, y) \subseteq T$. This is again independent of the choice of x', y' . In fact, one checks that $J(x, y) = [x, y]_{\hat{\rho}}$. Moreover, if $x, y, z \in T$, then $\#(J(x, y) \cap J(y, z) \cap J(z, x)) = 1$. (Again, we only really need to check this holds in a simplicial tree.) It follows that $(T, \hat{\rho})$ is an \mathbb{R} -tree.

Finally, given $a \in M$, let \hat{a} be the image of $\bar{a} \in D$, in T . We check that the map $[a \mapsto \hat{a}] : M \rightarrow T$ is an isometric embedding. \square

Remark. In the above, we only used the four-point property for some fixed $p \in T$, so we could weaken the quantification to say that there exists $p \in M$, such that the inequality holds for all $a, b, c \in M$.

As we have noted, any connected subset of an \mathbb{R} -tree is an \mathbb{R} -tree. We therefore arrive at the following well known characterisation:

Lemma 15.1.4. *A metric space is an \mathbb{R} -tree if and only if it is connected and weakly 0-hyperbolic.*

(This gives a much easier way of seeing that the metric completion of an \mathbb{R} -tree is an \mathbb{R} -tree.)

In particular, an \mathbb{R} -tree is a geodesic weakly 0-hyperbolic metric space. This last condition is (equivalent to) the standard notion of a “0-hyperbolic space” (that is,

a “ k -hyperbolic space” with $k = 0$). This will be discussed further in Subsection 24.1.

15.2. Boundaries of \mathbb{R} -trees.

For future reference, we note that an \mathbb{R} -tree has associated with it a boundary ∂T , which can be defined in a number of equivalent ways. For example, it is the set of equivalence classes of rays in T (that is, isometrically embedded copies of $[0, \infty)$), where two rays are deemed equivalent if they intersect in another ray. One can put a natural hausdorff topology on $T \cup \partial T$. This is an example of a much more general construction of the boundary of a Gromov hyperbolic space as well as that of a CAT(0) space — these two constructions agree in the case of an \mathbb{R} -tree.

In the case of a simplicial tree, ∂T can be canonically identified with the Roller boundary, $\partial_R T$. However the canonical topology on $T \cup \partial T$ is not in general compact, and so may differ from the compact topology on $T \cup \partial_R T$ as described by Lemma 12.6.2. We have only defined the Roller boundary here for discrete median algebras, though there is a more general notion (see for example, [Fi3]). For a general \mathbb{R} -tree, ∂T and $\partial_R T$ may differ.

15.3. Embedding in products of \mathbb{R} -trees.

In what follows, it will be convenient to also allow for pseudometrics. Recall that, given any pseudometric space, (M, ρ) , we can define an equivalence relation, \sim , on M by writing $x \sim y$ to mean that $\rho(x, y) = 0$. The quotient, M/\sim , is then naturally a metric space called the *hausdorffification* of M . (Here we imagine the pseudometric as inducing a non-hausdorff topology on M .) We say that (M, ρ) is a “0-hyperbolic pseudometric space” if it satisfies the four-point condition (FP) for all quadruples of points of M . In this case, the hausdorffification will be a weakly 0-hyperbolic metric space.

Recall the notion of ν -colourability from Subsection 8.3. Any \mathbb{R} -tree is 1-colourable (since no pair of walls cross) and so a direct product of ν \mathbb{R} -trees is ν -colourable. So also is any subalgebra of such a product. Conversely, we have:

Proposition 15.3.1. *A ν -colourable median metric space isometrically embeds into an l^1 product of ν \mathbb{R} -trees.*

As noted earlier, such an embedding is necessarily a median monomorphism.

Proof. Let (M, ρ) be ν -colourable.

Our aim is to write $\rho = \sum_{i=1}^{\nu} \rho_i$, where each ρ_i is a 0-hyperbolic pseudometric on M . Given this, we can then let M_i be the hausdorffification of (M, ρ_i) . This is a 0-hyperbolic metric space, so by Lemma 15.1.3, it isometrically embeds into an \mathbb{R} -tree, T_i . We now embed M into M^ν via the diagonal map. We then postcompose this with the direct product of the respective hausdorffifications, to give us map into $\prod_{i=1}^{\nu} M_i$. We further postcompose by the product of the embeddings into $\prod_{i=1}^{\nu} T_i$. This then proves the result.

To decompose the metric in this way, we use a compactness argument. Let $D = \prod_{a,b \in M} [0, \rho(a, b)]$ be a direct product of compact real intervals. By Tychonoff's Theorem, D is compact in the product topology. Given $\sigma \in D$, we get a map, $\sigma : M^2 \rightarrow [0, \infty)$, where $\sigma(a, b)$ is the (a, b) -coordinate of σ . We think of σ as defining a distance on M , with $\sigma(a, b) \leq \rho(a, b)$ for all a, b .

Let \mathcal{A} be the set of finite subalgebras of M . Given $\Pi \in \mathcal{A}$, let $P(\Pi) \subseteq D$ be the set of $\sigma \in D$ such that $\sigma|_{\Pi^2}$ is a 0-hyperbolic pseudometric on Π . Thus $P(\Pi)$ is closed in D . Let $R(\Pi) \subseteq (P(\Pi))^\nu \subseteq D^\nu$ be the set of $(\rho_1, \dots, \rho_\nu) \in (P(\Pi))^\nu$ such that $\sum_{i=1}^\nu \rho_i(x, y) = \rho(x, y)$ for all $x, y \in \Pi$. Then $R(\Pi)$ is a closed subset of D^ν .

We claim that $(R(\Pi))_{\Pi \in \mathcal{A}}$ has the finite intersection property. Note that if $\Pi_1, \dots, \Pi_n \in \mathcal{A}$, then $\Pi := \langle \Pi_1 \cup \dots \cup \Pi_n \rangle \in \mathcal{A}$, and that $R(\Pi) \subseteq \bigcap_{i=1}^n R(\Pi_i)$. To verify the claim, it is therefore sufficient to show that $R(\Pi) \neq \emptyset$ for all $\Pi \in \mathcal{A}$.

Now Π is itself ν -colourable. In other words, we can write $\mathcal{W}(\Pi) = \mathcal{W}_1 \sqcup \dots \sqcup \mathcal{W}_\nu$, such that if $W, W' \in \mathcal{W}_i$ then $W \not\cap W'$. Given $W \in \mathcal{W}(\Pi)$, let $w(W)$ be the width of W in Π . Given $x, y \in \Pi$, write $\rho_{i,\Pi}(x, y) = \sum \{w(W) \mid W \in \mathcal{W}_i \cap \mathcal{W}(x, y)\}$. Thus, $\rho(x, y) = \sum_{i=1}^\nu \rho_{i,\Pi}(x, y)$. Now since no walls of \mathcal{W}_i cross, it is easy to see that $\rho_{i,\Pi}$ is a 0-hyperbolic pseudometric on Π . (The hausdorffification of $(\Pi, \rho_{i,\Pi})$ is a finite rank-1 median metric space, hence the vertex set of a simplicial tree.) Defining $\rho_{i,\Pi}$ arbitrarily elsewhere, see that $\rho_{i,\Pi} \in P(\Pi)$. Therefore $(\rho_{1,\Pi}, \dots, \rho_{\nu,\Pi}) \in R(\Pi)$. In particular, $R(\Pi) \neq \emptyset$, as claimed.

Now by Tychonoff's Theorem, we have $\bigcap_{\Pi \in \mathcal{A}} R(\Pi) \neq \emptyset$. Let $(\rho_1, \dots, \rho_\nu) \in \bigcap_{\Pi \in \mathcal{A}} R(\Pi)$. Then, $\rho_i \in \bigcap_{\Pi \in \mathcal{A}} P(\Pi)$, and $\sum_{i=1}^\nu \rho_i = \rho$.

We claim that ρ_i is a 0-hyperbolic pseudometric on M . To see this, it is enough to note that any finite subset, A , we have $\rho_i \in P(\langle A \rangle)$. By definition, this means that $\rho_i|_{P(\langle A \rangle)}$ is a 0-hyperbolic pseudometric. In this way we readily verify the pseudometric space axioms and the four-point property for ρ_i .

As explained in the first paragraph, the result now follows. □

Remark. In the above proof, we only used the fact that any finite subalgebra of M is ν -colourable. It can therefore be used to give another, though less direct, proof of Lemma 8.3.1.

Remark. A related but different embedding theorem can be found in [Bo4]. The hypotheses there are weaker in assuming we just have a lipschitz median algebra, but stronger in assuming that any two points are connected by a uniformly lipschitz path.

We can now use this to prove Lemma 13.3.5.

Proof of Lemma 13.3.5. Let M be a median metric space, and let $a, b \in M$ with $\text{rank}([a, b]) = \nu < \infty$. By Lemma 8.3.2, $[a, b]$ is ν -colourable, and so by Proposition 15.3.1, it isometrically embeds into an l^1 product of ν \mathbb{R} -trees. Now its image lies in the interval between the images of a and b . This interval is a product of ν real intervals. □

(Of course, one could apply the compactness argument in the proof of Proposition 15.3.1 more directly in this case.)

15.4. **Guirardel cores.**

An important example of a subalgebra of a product of (two) \mathbb{R} -trees is the Guirardel core. Before discussing this, we note that we can easily describe the walls of an \mathbb{R} -tree.

To this end, let T rank-1 median algebra and let $p \in T$. We define a relation, \sim , on $T \setminus \{p\}$ by writing $x \not\sim y$ if $x.p.y$. It is easily checked that \sim is an equivalence relation. A **branch** of T based at p is an equivalence class. If B a branch, then it is also easily checked that T and $T \setminus B$ are both convex, so $\{B, T \setminus B\}$ is a wall.

If T is an \mathbb{R} -tree then every wall arises in this way. To see this, suppose $W = \{W^-, W^+\} \in \mathcal{W}(T)$. Choose any $a \in W^-$ and $b \in W^+$. Now $\{[a, b] \cap W^-, [a, b] \cap W^+\}$ is a wall of the real interval $[a, b]$. Therefore, up to swapping a, b , we have $[a, b] \cap W^- = [a, c]$ and $[a, b] \cap W^+ = [c, b]$ for some $c \in [a, b]$. Thus W^- is the preimage of $[a, c]$ under the gate map to $[a, b]$, which is easily seen to be a branch of T based at c . Note also, in this case, a branch based at p could equivalently be defined as a connected component of $T \setminus \{p\}$.

We also note that a branch, B , of T determines a (possibly empty) subset, $\partial B \subseteq \partial T$, given by all rays which (eventually) lie in B .

The Guirardel core is an instance of a much more general construction, which we describe first.

Let M_1, M_2 be median algebras, and let $M = M_1 \times M_2$. Let $\mathcal{B}_i \subseteq \mathcal{P}(M_i)$ be a set of subsets of M_i whose complements in M_i are convex. Let $\mathcal{C} \subseteq \mathcal{B}_1 \times \mathcal{B}_2$ be any subset, and let $M(\mathcal{C}) = M \setminus \bigcup_{(B_1, B_2) \in \mathcal{C}} (B_1 \times B_2)$. We claim that $M(\mathcal{C})$ is a subalgebra of M . To see this, suppose $a, b, c \in M$ with $abc \notin M(\mathcal{C})$. Then $abc \in B_1 \times B_2$ for some $(B_1, B_2) \in \mathcal{C}$. Writing $a = (a_1, a_2)$ etc, we have $a_1 b_1 c_1 \in B_1$ and $a_2 b_2 c_2 \in B_2$. Now $M_i \setminus B_i$ is convex, so at least two of a_i, b_i, c_i lie in B_i . Without loss of generality, we have $a_1 \in B_1$ and $a_2 \in B_2$. Thus, $a \in B_1 \times B_2$, so $a \notin M(\mathcal{C})$. This proves the claim.

Now suppose we have isometric actions of the same group, Γ , on two \mathbb{R} -trees, T_1 and T_2 . We now let \mathcal{B}_i be the set of all branches of T_i at all points of T_i . Let $(B_1, B_2) \in \mathcal{B}_1 \times \mathcal{B}_2$. Let $(p_1, p_2) \in T_1 \times T_2$. We deem (B_1, B_2) to lie in \mathcal{C} if for any sequence of elements $(g_i)_{i \in \mathbb{N}}$ such that $g_i p_1 \in B_1$ and $g_i p_2 \in B_2$ for all i , then either $g_i p_1$ is bounded or $g_i p_2$ is bounded. (This is independent of the choice of (p_1, p_2) .) Then $M(\mathcal{C})$ is the **Guirardel core** of this pair of actions (at least in the case where it is non-empty). This construction has many applications as we mention in the Notes to this section.

16. MEDIAN GRAPHS

This section can be thought of as a continuation of the discussion of discrete median algebras in Section 11, though we approach it from a graph-theoretical viewpoint. Some of the definitions below repeat, or reinterpret, earlier ones. We

will see that they are consistent where they overlap. The more general notion of a “quasimedial graph” will be the topic of Section 23.

We will give a number of characterisations of a median graph. In particular, it is the adjacency graph of a discrete median algebra (Lemma 16.1.2). It can also be thought of as the 1-skeleton of a CCAT(0) cube complex (Proposition 16.3.1): a subject we will return to in Sections 17 and 18. One of the main results here (due to Chepoi) is that it can also be described by a local property, together with a “simple connectedness” assumption (Theorem 16.2.3).

16.1. Characterisations of median and modular graphs.

Let Γ be a connected graph. (By default, all graphs we consider in this section will be connected.) We assume that Γ has no loops or multiple edges. We write $V(\Gamma)$ and $E(\Gamma)$ for the vertex and edge sets. A **path** of length n from a to b in $V(\Gamma)$ is a sequence of vertices, $a = a_0, a_1, \dots, a_n = b$, such that a_i, a_{i+1} are adjacent for all i . We will typically denote it by $a_0 a_1 \dots a_n$, omitting the commas. We refer to a and b as the **initial** and **terminal** vertices. We say that α **emanates** from a and **terminates** at b . We say that α is **geodesic** if n is minimal. In this case, we write $\rho(a, b) = \rho_\Gamma(a, b) = n$. Thus ρ is a metric on $V(\Gamma)$: the **combinatorial metric**.

Given $a, b, c \in V(\Gamma)$, we write $a.c.b$ to mean that $\rho(a, b) = \rho(a, c) + \rho(c, b)$. In other words, c lies on some geodesic from a to b . We write $[a, b]_\rho = \{x \in V(\Gamma) \mid a.x.b\}$. Given $a, b, c \in V(\Gamma)$, let

$$\text{Med}(a, b, c) = [a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho.$$

An element of $\text{Med}(a, b, c)$ is called a **median** of a, b, c .

As usual, we write $x_1.x_2.\dots.x_n$ to mean that $x_i.x_j.x_k$ holds whenever $i \leq j \leq k$. As with any metric space, we have a linear interpolation rule (similarly as for a median algebra as discussed in Subsection 3.2).

An **n -cycle** in Γ is a closed path (i.e. with $a_0 = a_n$). (We do not, in general, assume that a cycle is embedded.) We say that Γ is **bipartite** if all cycles have even length. This is equivalent to saying that $V(\Gamma)$ can be partitioned into two subsets such that no two vertices in the same subset are adjacent (i.e. it is 2-colourable). This is in turn equivalent to saying that $\rho(a, b) + \rho(b, c) + \rho(c, a)$ is even for all $a, b, c \in V(\Gamma)$.

Definition. We say that Γ is **modular** if $\text{Med}(a, b, c) \neq \emptyset$ for all $a, b, c \in V(\Gamma)$.

Note that if $\text{Med}(a, b, c) \neq \emptyset$, then $\rho(a, b) + \rho(b, c) + \rho(c, a)$ is even. We deduce that any modular graph is bipartite.

Definition. We say that Γ is **median** if $\#\text{Med}(a, b, c) = 1$ for all $a, b, c \in V(\Gamma)$.

In other words, $(V(\Gamma), \rho)$ is a median metric space.

There are many ways of characterising median graphs (see for example, [KlaM, Chep2]). We begin with the following observation.

We write $K_{2,3}$ for the complete bipartite graph. In other words, it has five vertices, x_1, x_2, x_3, m_1, m_2 , and six edges connecting x_i to m_j for all i, j . Note that a median graph cannot contain a (subgraph isomorphic to) $K_{2,3}$. (Since, in that case, $m_1, m_2 \in \text{Med}(x_1, x_2, x_3)$.) In fact we have:

Lemma 16.1.1. *A modular graph which contains no $K_{2,3}$ is median.*

Proof. Let Γ be modular and contain no $K_{2,3}$. Suppose for contradiction, that there exist $x_1, x_2, x_3 \in V(\Gamma)$ with distinct $m, m' \in \text{Med}(x_1, x_2, x_3)$. Among all possibilities for x_1, x_2, x_3, m, m' , we choose one such that $\rho(m, m') > 0$ is minimal.

Choose $y_i \in \text{Med}(x_i, m, m')$. Note that for all $i \neq j$, we have $x_i \cdot y_i \cdot m \cdot y_j \cdot x_j$ (by the linear interpolation rule for graphs mentioned above), so $m \in \text{Med}(y_1, y_2, y_3)$. Similarly $m' \in \text{Med}(y_1, y_2, y_3)$.

Now $y_1 \cdot m \cdot y_2, y_1 \cdot m' \cdot y_2, m \cdot y_1 \cdot m'$ and $m \cdot y_2 \cdot m'$. From this, it follows that $\rho(y_1, m) = \rho(y_2, m')$ and $\rho(y_1, m') = \rho(y_2, m)$ (cf. Lemma 13.2.3). Similarly, $\rho(y_2, m) = \rho(y_3, m')$, $\rho(y_2, m') = \rho(y_3, m)$ and $\rho(y_3, m) = \rho(y_1, m')$, $\rho(y_3, m') = \rho(y_1, m)$. In other words, there is some $p > 0$ such that $\rho(y_i, m) = \rho(y_i, m') = p$ for all i . Also, $\rho(m, m') = 2p$ and $\rho(y_i, y_j) = 2p$ for all $i \neq j$. We must have $p \geq 2$ (otherwise y_1, y_2, y_3, m, m' would be the vertex set of a $K_{2,3}$).

Now let $z, z' \in V(\Gamma)$ be adjacent to y_3 , with $\rho(z, m) = \rho(z', m') = p - 1$. Thus, $m \cdot z \cdot y_3 \cdot z' \cdot m'$. Choose some $w \in \text{Med}(y_1, z, z')$ and $w' \in \text{Med}(y_2, z, z')$. Then w and w' are both adjacent to both z and z' . In fact, $w = w'$ (otherwise w, w', y_3, z, z' would be the vertex set of a $K_{2,3}$). Note that $m \cdot z \cdot w \cdot z' \cdot m'$, so $\rho(w, m) = \rho(w, m') = p$. Also $\rho(y_1, w) = \rho(y_2, w) = 2p - 2$.

Now choose $m'' \in \text{Med}(w, y_1, y_2)$, so $z \cdot w \cdot m'' \cdot y_1$ and $z \cdot w \cdot m'' \cdot y_2$. Now $\rho(y_1, m'') = \rho(y_2, m'') = p$, and $\rho(w, m'') = p - 2$. In particular, $m'' \neq m$. Also, $\rho(z, m'') = p - 1$. Recall that $\rho(y_1, m) = \rho(y_2, m) = p$ and that $\rho(y_1, y_2) = 2p$. Also, $\rho(z, y_1) = \rho(z, y_2) = 2p - 1$. It follows that $m, m'' \in \text{Med}(y_1, y_2, z)$. But $\rho(m, m'') \leq \rho(m, z) + \rho(z, m'') \leq 2p - 2 < 2p = \rho(m, m')$ contradicting minimality. \square

Next, we relate median graphs to discrete median algebras. This is relatively straightforward.

Let Γ be a median graph, and write $\Pi = V(\Gamma)$. By definition, Π is a median metric space, hence a median algebra. We can recover $\Gamma = \Gamma(\Pi)$ as the graph with vertex set Π and with adjacency determined by adjacency in Π . (This has already been defined in Subsection 5.1, and mentioned again in Subsection 11.4.)

We claim that Π is discrete. To see that, let $a, b \in \Pi$, and let $a = a_0, \dots, a_n = b$ be any path between them. For any i , a_{i-1} and a_i determine a wall, $W_i \in \mathcal{W}(\Pi)$ such that $\mathcal{W}(a_{i-1}, a_i) = \{W_i\}$. Then $\mathcal{W}(a, b) \subseteq \{W_1, \dots, W_n\}$. In particular, $\mathcal{W}(a, b)$ is finite. Therefore, $[a, b]$ is finite (see Corollary 3.2.14).

In summary, we have shown:

Lemma 16.1.2. *A graph Γ is median if and only if it has the form $\Gamma = \Gamma(\Pi)$ for a discrete median algebra, Π .*

We finish this subsection with another characterisation of modular graphs, namely Proposition 16.1.3.

Let Γ be a connected graph. Let $p \in V(\Gamma)$ be some basepoint. Suppose $a, b \in V(\Gamma)$ with $\rho(p, a) = \rho(p, b)$ and $\rho(a, b) = 2$. Then $d \in \text{Med}(a, b, c)$ if and only if d is adjacent to both a and b , and $\rho(p, d) = \rho(p, a) - 1$. Thus, if Γ is modular, then d always exists. If Γ is median, then it is also unique. (This accords with the notation introduced in Subsection 3.2 for a median algebra.)

Consider the following condition on Γ with respect to p :

$(\diamond(p))$: Suppose $a, b, c \in V(\Gamma)$ with a, b distinct, both adjacent to c , and with $\rho(p, a) = \rho(p, b) = \rho(a, c) - 1$. Then there is some $d \in V(\Gamma)$, adjacent to both a and b and with $\rho(p, d) = \rho(p, c) - 2$.

(If Γ has no embedded $K_{2,3}$, then d will be unique, though we are not assuming that for the moment.)

We will say that Γ “satisfies (\diamond) ” if $(\diamond(p))$ holds for all $p \in V(\Gamma)$.

(Note that this is precisely the conclusion of Lemma 11.3.4, for the adjacency graph associated with a discrete median algebra.)

Proposition 16.1.3. *A connected graph is modular if and only if it is bipartite and satisfies (\diamond) .*

Proof. We have already observed that the “only if” implication holds. So we assume that Γ is a connected bipartite graph satisfying (\diamond) . Let $x, y, z \in V(\Gamma)$. We want to show that $\text{Med}(x, y, z) \neq \emptyset$.

To this end, choose $a \in V(\Gamma)$ so as to minimise $\rho(x, a) + \rho(y, a) + \rho(z, a)$. Write $k = \rho(x, a)$, $l = \rho(y, a)$ and $m = \rho(z, a)$. We claim that $\rho(x, y) = k + l$, $\rho(y, z) = l + m$ and $\rho(z, x) = m + k$.

Suppose for contradiction that $\rho(x, y) < k + l$. Let $y_0 y_1 \dots y_l$ be a geodesic from $y_0 := a$ to $y_l := y$. Let i be minimal such that $\rho(x, y_i) < k + i$. Thus, $i > 0$. Now $\rho(x, y_{i-1}) = k + i - 1$. Since Γ is bipartite, we have $\rho(x, y_i) = k + i - 2$. Suppose that $i \geq 2$. Then $\rho(x, y_{i-2}) = \rho(x, y_i) = k + i - 2$. Applying $(\diamond(x))$, there is some $d \in V(\Gamma)$ adjacent to both y_i and y_{i-2} with $\rho(x, d) = k + i - 3$. We now replace y_{i+1} with d to give another geodesic from a to y . This time, $\rho(a, y_{i-1}) < k + i - 1$. Continuing inductively in this manner, we arrive at a geodesic for which $\rho(x, y_1) = k + 1 - 2 = k - 1$. (If $i = 1$ to begin with, we would also have $\rho(x, y_1) = k - 1$.) Now $\rho(y, y_1) = l - 1$ and $\rho(z, y_1) \leq m + 1$. Therefore $\rho(x, y_1) + \rho(y, y_1) + \rho(z, y_1) \leq k + l + m - 1$, contradicting the choice of a .

We conclude that $\rho(x, y) = k + l$. Similarly, $\rho(y, z) = l + m$ and $\rho(z, x) = m + k$, and so $a \in \text{Med}(x, y, z)$. \square

16.2. Equivalence with CAT(0) cube complexes.

Next, we give another description of median graphs which will relate them to CAT(0) cube complexes as discussed at the end of this section.

First, we make a few more definitions regarding a general connected graph, Γ .

Let $G \subseteq \Gamma$ be a connected subgraph. Note that the inclusion map is 1-lipschitz with respect to the metrics ρ_G and ρ_Γ . We say that G is **full** if every edge of Γ with endpoints in G lies in $E(G)$. We say that G is **isometrically embedded** if the inclusion $(G, \rho_G) \hookrightarrow (\Gamma, \rho_\Gamma)$ is an isometric embedding. Note that this implies that G is full.

A **square** in Γ is a full embedded 4-cycle viewed as a subgraph of Γ . Clearly, any square is isometrically embedded. (Note that if Γ has no 3-cycles, then every 4-cycle is full.) We write $\mathcal{S}(\Gamma)$ for the set of all squares in Γ .

Definition. A **square structure** on Γ is a set $\mathcal{S} \subseteq \mathcal{S}(\Gamma)$, of squares of Γ . A **square complex** consists of a graph, Γ , together with a square structure.

Remark. In fact, for most purposes, we could instead deal just with the square structure, $\mathcal{S}(\Gamma)$, consisting of all squares of Γ , so that it becomes an intrinsic feature of Γ . However, we may as well deal with the more general case, since it does not add any significant complications to the argument.

We will also refer to a path of length 4, say $x_0x_1x_2x_3x_4$, as a “square” if $x_0 = x_4$ and x_0, x_1, x_2, x_3 are all distinct. We view it as the same square as $x_1x_2x_3x_4x_1$ etc.

We can define moves on paths in Γ similarly as for discrete median algebras in Subsection 11.4. Namely we have moves:

- (1): If $a_{i-1} = a_{i+1}$, we replace a_{i-1}, a_i, a_{i+1} with a_{i-1} .
- (1'): If b is adjacent to a_i , then we replace a_i by a_i, b, a_i .
- (2): If there is some $b \in V(\Gamma)$ such that a_{i-1}, a_i, a_{i+1}, b is a square in \mathcal{S} , (with b antipodal to a_i), we replace a_i with b .

(Here, we are allowing an inverse of move (1), and we are reinterpreting a “2-cell” as an element of \mathcal{S} .)

Consider the following condition:

(S1): Any cycle can be reduced to a constant path by a finite sequence of moves of type (1), (1') and type (2).

We note:

Lemma 16.2.1. *Any graph which admits a square structure satisfying (S1) is bipartite.*

Proof. Note that the above moves do not change the length of a cycle modulo 2, and so every cycle has even length. □

As a converse, we have:

Lemma 16.2.2. *Suppose Γ is a bipartite connected graph satisfying $(\diamond(p))$ for some $p \in V(\Gamma)$. Then $\mathcal{S}(\Gamma)$ satisfies (S1).*

Proof. Let $\underline{a} = a_0 a_1 \dots a_n$ be a closed path with $n > 0$ (so $a_n = a_0$). Let $A(\underline{a}) = \sum_{i=1}^n \rho(p, a_i)$. Choose i so as to maximise $\rho(p, a_i)$. Since Γ is bipartite, we have $\rho(p, a_{i-1}) = \rho(p, a_{i+1}) = \rho(p, a_i) - 1$. If $a_{i-1} = a_{i+1}$, apply Move (1). If not, let $d = a_{i-1} \wedge a_{i+1}$ as given by $(\diamond(p))$, and apply Move (2) by replacing a_i by d . This reduces $A(\underline{a})$. Continuing in this manner, we eventually arrive at a constant path. \square

(Note that we did not need Move (1') for this.)

Property (S1) has an obvious topological interpretation as follows.

Let Γ be a graph with a square structure, \mathcal{S} . Let $\Delta = \Delta(\Gamma, \mathcal{S})$ be the 2-complex with 1-skeleton Γ obtained by gluing a disc (a “2-cell”) along its boundary to each element of \mathcal{S} . We can equip Δ with the CW topology (see below for a definition). It is not hard to see that Property (S1) above is equivalent to asserting that Δ is simply connected. We will not give a formal proof of this. It is not essential for what we do here, since everything can be interpreted combinatorially.

(In fact, we could equivalently put other topologies on Δ . For example there is the metric topology induced by giving each cell the structure of a unit euclidean square. Again, \mathcal{S} satisfies (S1) if and only if Δ is simply connected in this metric topology. This can be viewed in terms of Dowker’s Theorem, [Do], as we will discuss in more detail in Subsection 17.2.)

We now consider additional properties of a square structure, \mathcal{S} :

(S2): No two distinct elements of \mathcal{S} meet in three distinct vertices.

Note that if Γ contains no full subgraph isomorphic to $K_{2,3}$, then any square structure \mathcal{S} satisfies (S2). Conversely, if $\mathcal{S}(\Gamma)$ satisfies (S2) then Γ contains no full $K_{2,3}$.

Here is another condition.

A **hexagon** in Γ is an embedded graph whose vertex set is a 6-cycle. (We do not need to assume that it is isometrically embedded, though it will be in the main cases of interest.)

Definition. A **wheel** in a square complex, (Γ, \mathcal{S}) , is an embedded subgraph, $\Omega \subseteq \Gamma$, which consists of a hexagon, $R = R(\Omega) \subseteq \Omega$ (the **rim** of Ω), with $V(R) = \{a_1, a_2, a_3, a_4, a_5, a_6\}$, together with a seventh vertex $h = h(\Omega) \in V(\Omega)$ (the **hub** of Ω) which is connected to each of $a_1, a_3, a_5 \in V(R)$ by three additional edges (the **spokes** of Ω). We assume that the three squares, $\{a_1, a_2, a_3, h\}$, $\{a_3, a_4, a_5, h\}$ and $\{a_5, a_6, a_1, h\}$ all lie in \mathcal{S} . (Here, of course, the vertices a_i are cyclically ordered around the rim, with indices taken modulo 6.)

In other words, Ω is the union of three squares in \mathcal{S} , meeting at a common vertex, and with each pair sharing a common edge. Note that $R(\Omega)$ and $h(\Omega)$ are completely determined by Ω .

Definition. We say that two wheels, Ω and Ω' are **dual** if $R(\Omega) = R(\Omega') = \Omega \cap \Omega'$ and the spokes of Ω are disjoint from the spokes of Ω' .

We see that $\Omega \cup \Omega'$ is an embedded subgraph isomorphic to a 3-cubical graph: that is the adjacency graph, $\Gamma(Q)$, of a 3-cube Q . Note also that if (Γ, \mathcal{S}) satisfies (S2), then each wheel has at most one dual.

Consider the following property:

(S3): Every wheel has a dual.

In other words, we can complete every wheel to a 3-cubical graph.

Here is another characterisation of median graphs (see [Chep2]).

Theorem 16.2.3. *A connected graph Γ is median if and only if it admits a square structure, \mathcal{S} , satisfying (S1), (S2) and (S3). In this case, $\mathcal{S} = \mathcal{S}(\Gamma)$.*

The “only if” part of Theorem 16.2.3 has already been established, and only calls for brief comment. Suppose Γ is a median graph. By Lemma 16.1.2, $\Gamma = \Gamma(\Pi)$ for a discrete median algebra, Π . In the terminology of Section 11, the squares of Γ are precisely the 2-cells of Π . Let $\mathcal{S} = \mathcal{S}(\Gamma)$ be the set of such squares. Now Lemma 11.4.2 implies that $\mathcal{S}(\Gamma)$ satisfies (S1). Moreover, Lemma 11.4.1 implies that $\mathcal{S}(\Gamma)$ satisfies (S3). (For if $h, a_1, a_2, a_3, a_4, a_5, a_6$ is a wheel in the above notation, then $\{a_1, a_3, a_5\}$ is a 2-simplex in the link $L(h)$. Therefore the wheel lies in a 3-cell of Π . Let h' be vertex of the 3-cell antipodal to h . Then $h', a_2, a_3, a_4, a_5, a_6, a_1$ is a dual wheel.)

The “if” direction of Theorem 16.2.3 is more involved. There are a number of different approaches to this. (See the Notes on this section for more discussion.) For example, it can be viewed in terms of CAT(0) geometry, as we discuss in Subsection 18.3. In this section, we will give a combinatorial proof, which can broadly be seen as a particular case of the argument given in [ChalCHO]. For this, we need some more definitions.

Given $a \in V(\Gamma)$, let $L_\Gamma(a) \subseteq V(\Gamma)$ be the set of adjacent vertices. If \mathcal{S} is a square structure on Γ satisfying (S2), we can view $L_\Gamma(a)$ as the vertex set of a graph, where $x, y \in L_\Gamma(a)$ are deemed adjacent if a, x, y all lie in some element of \mathcal{S} . By (S2), this element is unique. We can think of this graph as the link of a in the 2-complex $\Delta(\Gamma, \mathcal{S})$ defined above.

Let Γ_0 be another graph. A **morphism** from Γ_0 to Γ is a map $f : \Gamma_0 \rightarrow \Gamma$ sending vertices to vertices and edges to edges. Thus $f(L_{\Gamma_0}(a)) \subseteq L_\Gamma(fa)$ for all $a \in V(\Gamma_0)$. Given a path $\alpha = x_0x_1 \dots x_m$ in Γ_0 , we get a path $fa := fx_0 \dots fx_m$ in Γ . We say that f is an **immersion** if $f|_{L_{\Gamma_0}(a)}$ is injective for all $a \in V(\Gamma_0)$. In other words, every path of length 2 in Γ_0 is mapped injectively to Γ .

Suppose $\mathcal{S}_0 \subseteq \mathcal{S}(\Gamma_0)$ and $\mathcal{S} \subseteq \mathcal{S}(\Gamma)$ are square structures. We say that an immersion, $f : \Gamma_0 \rightarrow \Gamma$, **respects** the square structures if $f(\sigma) \in \mathcal{S}$ for all $\sigma \in \mathcal{S}_0$. In such a case, if (Γ, \mathcal{S}) satisfies (S2), then so does $(\Gamma_0, \mathcal{S}_0)$. Moreover, we get an induced continuous locally injective map $f : \Delta(\Gamma_0, \mathcal{S}_0) \rightarrow \Delta(\Gamma, \mathcal{S})$. (In our argument below, we will just take $\mathcal{S}_0 = \mathcal{S}(\Gamma_0)$.)

We say that an immersion, $f : \Gamma_0 \rightarrow \Gamma$ is a **local isomorphism** if $f|_{L_{\Gamma_0}(a)}$ is a bijection from $L_{\Gamma_0}(a)$ to $L_\Gamma(fa)$ for all $a \in V(\Gamma_0)$. If $x \in V(\Gamma_0)$ and β is a path

in Γ emanating from fx , then there is a unique path, α , in Γ_0 emanating from x with $f\alpha = \beta$. We refer to α as the *lift* of α .

More generally, suppose $f : \Gamma_0 \rightarrow \Gamma$ is an immersion, and $H \leq \Gamma_0$ is a full subgraph. We say that f is a **local isomorphism over H** if $f|_{L_{\Gamma_0}(a)} : L_{\Gamma_0}(a) \rightarrow L_{\Gamma}(fa)$ is bijective for all $a \in V(H)$. In this case, we can lift a path β to a path α in Γ_0 with either $f\alpha = \beta$ or $f\alpha$ an initial segment of β terminating on a vertex of $V(\Gamma_0) \setminus V(H)$. Any maximal such lift is unique.

Let $f : \Gamma_0 \rightarrow \Gamma$ be locally injective. Consider the following property:

(*): If α is a path of length 4 in Γ_0 with $f(\alpha) \in \mathcal{S}$, then the initial and terminal vertices of α are equal.

In other words, $\alpha \in \mathcal{S}(\Gamma_0)$.

Lemma 16.2.4. *Let (Γ, \mathcal{S}) satisfy (S1). Let $f : \Gamma_0 \rightarrow \Gamma$ be a local isomorphism respecting the square structures, $\mathcal{S}(\Gamma_0)$ and \mathcal{S} , and which satisfies (*). Then the induced map $f : \Delta(\Gamma_0, \mathcal{S}(\Gamma_0)) \rightarrow \Delta(\Gamma, \mathcal{S})$ is a covering map.*

Proof. It is enough to check that f is a local homeomorphism at every vertex, $a \in V(\Gamma_0)$. It is in turn enough to check that f induces a homeomorphism of the link at a to the link at $a' := fa \in V(\Gamma)$. Since f is a local isomorphism, it restricts to a bijection from $L_{\Gamma_0}(a)$ to $L_{\Gamma}(a')$. Since f respects the square structure, it sends adjacent vertices of $L_{\Gamma_0}(a)$ to adjacent vertices of $L_{\Gamma}(a')$. We need to check conversely that if $x, y \in L_{\Gamma_0}(a)$ with $x' := fx$ and $y' := fy$ adjacent in $L_{\Gamma}(a')$, then x, y are adjacent in $L_{\Gamma_0}(a)$. By definition, $x'a'y'$ lies in some square, $x'a'y'z'x'$ in \mathcal{S} . Let $xayzv$ be the lift of this path (based at x) to Γ_0 . By (*) we have $v = x$, so $xayzx \in \mathcal{S}(\Gamma_0)$. Therefore x, y are adjacent in Γ_0 as claimed. \square

To properly justify the above proof, we would need to explain why the topologies agree locally. This is true because the CW topology on the open star of any cell only depends on the cells meeting that star. We omit the details of this, since it just a convenient way of expressing what is really a combinatorial argument. It can all be rephrased in term of lifting paths from Γ to Γ_0 .

The key to proving the “if” part of Theorem 16.2.3 is the following:

Lemma 16.2.5. *Let Γ be a graph with square structure \mathcal{S} satisfying (S2) and (S3). Let $q \in V(\Gamma)$. Then there is a bipartite graph, G , and some $p \in V(G)$ such that G satisfies $(\diamond(p))$; together with a local isomorphism $f : G \rightarrow \Gamma$, with $fp = q$, respecting the square structures, $\mathcal{S}(G)$ and \mathcal{S} , and satisfying (*).*

Given Lemma 16.2.5, we can complete the proof of Theorem 16.2.3 as follows.

Proof of Theorem 16.2.3. It remains to prove the “if” direction.

Let Γ be a graph with square structure, \mathcal{S} , satisfying (S1), (S2) and (S3). Since \mathcal{S} satisfies (S1), Γ is bipartite, by Lemma 16.2.1. We claim that Γ satisfies (\diamond) .

To this end, let $q \in V(\Gamma)$, and let $f : G \rightarrow \Gamma$ be as given by Lemma 16.2.5. This induces a map $f : \Delta(G, \mathcal{S}(G)) \rightarrow \Delta(\Gamma, \mathcal{S})$, which by Lemma 16.2.4 is a

covering map. Since Γ satisfies (S1), $\Delta(\Gamma, \mathcal{S})$ is simply connected, so this map is a homeomorphism. In particular, $f : G \rightarrow \Gamma$ is an isomorphism. Since G satisfies $(\diamond(p))$ and $fp = q$, we see that Γ satisfies $(\diamond(q))$. Since q was arbitrary, Γ satisfies (\diamond) as claimed. Thus, by Proposition 16.1.3, Γ is modular.

Moreover, f maps $\mathcal{S}(G)$ bijectively to \mathcal{S} , and so $\mathcal{S} = \mathcal{S}(\Gamma)$. Since \mathcal{S} satisfies (S2), we see that Γ contains no full $K_{2,3}$, hence no $K_{2,3}$ at all. Therefore, by Lemma 16.1.1, Γ is median. \square

It therefore remains to prove Lemma 16.2.5.

We need some more notation. Let G be a connected graph, and let $p \in V(G)$ be some basepoint. We define $h : V(G) \rightarrow \mathbb{N}$ by setting $h(x) = \rho_G(p, x)$. Note that G is bipartite if and only if $|h(x) - h(y)| = 1$ for all adjacent pairs, $x, y \in V(G)$. Given $x, y \in V(G)$, we write $x \prec y$ to mean that x, y are adjacent and $h(x) = h(y) - 1$. In this notation, the property $(\diamond(p))$ can be expressed by saying that if $a, b, c \in V(G)$ with $a \neq b$, $a \prec c$ and $b \prec c$, then there is some $d \in V(G)$ with $d \prec a$, $d \prec b$. We say that a square $\sigma \in \mathcal{S}(G)$ is **folded** if $|h(x) - h(y)| \leq 1$ for all $x, y \in \sigma$. Otherwise, it is **unfolded**.

Given $n \in \mathbb{N}$, let $V_n = h^{-1}(n)$ and let $V_{\leq n} = h^{-1}([0, n])$. Let $G_n = G_n(G)$ be the full subgraph on $V_{\leq n}$. Note that $(\diamond(p))$ holds for G if and only if it holds for G_n for all n .

We now return to our square complex (Γ, \mathcal{S}) satisfying (S2) and (S3), as in the hypotheses of Theorem 16.2.3. We aim to construct G inductively as an increasing union, $G = \bigcup_{n=0}^{\infty} G_n$, of graphs G_n . If $m \leq n$, we identify $G_m = G_m(G_n)$. We also inductively define morphisms, $f_n : G_n \rightarrow \Gamma$, such that $f_n|_{G_m} = f_m$ for $m \leq n$. These will assemble to give us our morphism $f : G \rightarrow \Gamma$.

We will assume inductively that:

- (G1): $h(x) = n$ for all $x \in V(G_n) \setminus V(G_{n-1})$.
- (G2): G_n is bipartite.
- (G3): G_n satisfies $(\diamond(p))$.
- (G4): f_n is an immersion of G_n .
- (G5): f_n is a local isomorphism over G_{n-1} .
- (G6): If $\sigma \in \mathcal{S}(G_n)$, then σ is unfolded in G_n and $f_n(\sigma) \in \mathcal{S}$.
- (G7): $f_n : G_n \rightarrow \Gamma$ satisfies $(*)$ above (with $\Gamma_0 = G_n$).

We make the following observations. Property (G6) together with (S2) for Γ imply that G_n contains no full $K_{2,3}$. It follows that if $x, y \in V_m(G_n)$ are distinct, then there is at most one $t \in V_{m-1}(G_n)$ adjacent to both x and y (i.e. $t \prec x, y$). (For suppose u were another such. Let $v \prec t, u$, again as given by (G3). Now t, u, x, y, v form the vertices of a full $K_{2,3}$ in G_n giving a contradiction.) If such a t exists (for example by applying (G3)) we can unambiguously write $t = x \wedge y$.

We now begin by setting $G_0 = \{p\}$ and $f_0(p) = q$.

Suppose we have constructed $f_n : G_n \rightarrow \Gamma$, satisfying (G1)–(G7). We set about constructing $f_{n+1} : G_{n+1} \rightarrow \Gamma$. To simplify notation, we will write $x' = f_n x$ for $x \in V(G_n)$.

Let $x \in V_n(G_n)$. By (G4), $f_n|_{L_{G_n}(x)} : L_{G_n}(x) \rightarrow L_\Gamma(x')$ is injective. Let $D(x) = L_\Gamma(x') \setminus f_n(L_{G_n}(x))$, and let $D = \{(x, a) \mid x \in V_n(G_n), a \in D(x)\}$. We define a relation, \sim , on D , by writing $(x, a) \sim (y, b)$ if $a = b$ and (either $x = y$ or there is some $t \in V_{n-1}(G_n)$ adjacent to both x, y and with $ax't'y'a \in \mathcal{S}$). (We have already observed that such a t is unique.)

Lemma 16.2.6. \sim is an equivalence relation on D .

Proof. We need to check that \sim is transitive, so suppose $(x, a) \sim (y, a) \sim (z, a)$, with $x, y, z \in V_n(G_n)$ and with $a \in D(x) \cap D(y) \cap D(z)$.

Let $t = x \wedge y$ and $u = y \wedge z$. By the definition of \sim (and the uniqueness of t, u observed above) we have $ay't'x'a, ay'u'z'a \in \mathcal{S}$. Now $t, u \prec y$, and so by (G3), we can set $v = t \wedge u$. Now $y'tvuy \in \mathcal{S}(G_n)$, so by (G6), $y't'v'u'y' \in \mathcal{S}$. We see that $ax't'v'u'z'a$ is the rim of a wheel in Γ with hub at y' . By (S3), this has a dual wheel, and we set $b \in V(\Gamma)$ to be its hub. Now $v'b$ is an edge of Γ (a spoke of the dual wheel). Thus, by (G5), it lifts to an edge $vw \in G_n$, where $w \in V(G_n)$. In other words, $w' = b$. By (G5), we can further lift $v'w'x'$ to a path vws in G_n , with $s' = x'$. Now $xtvws$ is a path of length 4 in G_n with $x't'v'w'x' \in \mathcal{S}$. (This is a square of the dual wheel.) Therefore, by (G7), we must have $x = s$. This shows that w is adjacent to x in G_n . Since $h(v) = n - 2$, we have $h(w) = n - 1$. In other words, $w \prec x$. Similarly $w \prec z$. Now $ax'w'z'a \in \mathcal{S}$. (It is another square of the dual wheel.) Therefore, by definition of \sim , we have $(x, a) \sim (z, a)$ as required. \square

Now let $Y = D/\sim$. We write $[(x, a)] \in Y$ for the \sim -class of (x, a) .

We now construct a graph G_{n+1} with vertex set $V(G_{n+1}) = V(G_n) \sqcup Y$. We deem two vertices to be adjacent if they both lie in $V(G_n)$, or else, one, x , lies in $V(G_n)$ and the other has the form $[(x, a)] \in Y$ for some $a \in D(x)$. In this way, we can identify G_n as a full subgraph of G_{n+1} . We define $f_{n+1} : V(G_{n+1}) \rightarrow V(\Gamma)$ by setting $f_{n+1}|_{V(G_n)} = f_n$ and $f_{n+1}([(x, a)]) = a$ for $[(x, a)] \in Y$. This gives us a morphism $f_{n+1} : G_{n+1} \rightarrow \Gamma$.

We now set about verifying properties (G1)–(G7) for f_{n+1} . Similarly as before, we will write $x' = f_{n+1}x$ to simplify notation.

Now (G1) is immediate from the construction, and so $V_{n+1} = Y$. Also, no two vertices of V_m are adjacent for any m , and so G_{n+1} is bipartite, giving (G2).

To verify (G3), let $\theta \in V_{n+1}$. Suppose that $x, y \in V_n$ are adjacent to θ (that is $x, y \prec \theta$) and distinct. By construction, this means that $\theta = [(x, \theta')] = [(y, \theta')]$, where $(x, \theta) \sim (y, \theta)$. By definition of \sim , there is some $t \in V(G_n)$ with $t \prec x, y$ and with $\theta'x't'y'\theta' \in \mathcal{S}$. This is property $(\diamond(p))$, as required.

To verify (G4), there are three new cases to consider. Suppose first that $x \in V_n$ and $\eta, \theta \in V_{n+1}$ are adjacent to x , with $\eta' = \theta'$. Then $\eta = [(x, \eta')] = [(x, \theta')] = \theta$ as required. Secondly, suppose $\theta \in V_{n+1}$ and $t \in V_{n-1}$ are both adjacent to $x \in V_n$. Now $\theta = [(x, \theta')]$ with $\theta' \in D(x)$. By definition of $D(x)$, $\theta' \notin f_{n+1}(L_{G_n}(x))$, and so

$\theta' \neq t'$. Finally, suppose $x, y \in V_n$ are both adjacent to $\theta \in V_{n+1}$. Let $t = x \wedge y$ as given by (G3) for f_{n+1} . Now x, y are both adjacent to t in G_n . Therefore, by (G4) in G_n , if $x' = y'$ then $x = y$.

For (G5), let $x \in V_n$ and $a \in L_\Gamma(x')$. If $a \notin f_{n+1}(L_{G_n}(x))$, then $(x, a) \in D(x)$. Thus, $\theta := [(x, a)] \in V_{n+1}$ is adjacent to x and $\theta' = a$. This shows that $f_{n+1}(L_{G_{n+1}}(x)) = L_{G_{n+1}}(x')$ as required.

To continue, we will use the following terminology. We say that a path $x_0x_1x_2x_3x_4$ in V_{n+1} is of “type $r_0r_1r_2r_3r_4$ ” with $r_i \in \mathbb{N}$ to mean that $h(x_i) = n + 1 - r_i$ for all i .

For (G6), suppose that α is a cycle of length 4 in G_{n+1} meeting V_{n+1} . We can suppose that the basepoint lies in V_{n+1} , and so α has type 01010 or 01210, according to whether it is folded or unfolded. Suppose that $\alpha = \eta x t y \eta$ has type 01210. Thus $t = x \wedge y$. Now as observed earlier, t is the unique vertex of G_n with $t \prec x, y$. Therefore, since $(x, \theta') \sim (y, \theta')$, we must have $\eta' x' t' y' \eta' \in \mathcal{S}$, by definition of \sim . This proves (G6) in this case. We claim that the second case (type 01010) cannot occur. For suppose $\alpha = \eta x \theta y \eta$, with $\eta, \theta \in V_{n+1}$. Let $t = x \wedge y$. Similarly as before, we have $\eta' x' t' y' \eta', \theta' x' t' y' \theta' \in \mathcal{S}$. But $\eta' \neq \theta'$, contradicting (S2). In particular, this shows that no square in G_{n+1} is folded.

For (G7), let α be a path of length 4 in G_{n+1} with $f_{n+1}\alpha \in \mathcal{S}$. We can assume that α meets V_{n+1} . Therefore, up to reversing the order of the indices, we can assume that α is of one of the following types: 01010, 01012, 01210, 10101, 21012, 10121, 10123, 01212, 01232 or 01234.

Suppose $\alpha = \zeta x \eta y \theta$ is of type 01010. Then $\zeta' = \theta'$ and $\eta' x' \zeta' y' \eta' \in \mathcal{S}$. Let $t = x \wedge y$. By (G6), we have $\eta' x' t' y' \eta' \in \mathcal{S}$. Now $t' \neq \zeta'$ contradicting (S2), so this case cannot occur.

Suppose $\alpha = \eta x t y \theta$ is of type 01210. Since $\eta' = \theta'$, we have $(x, \eta') \sim (y, \theta')$ by the definition of \sim , and so $\eta = [(x, \eta')] = [(y, \theta')] = \theta$ as required.

Suppose that $\alpha = x \eta y \theta z$ is of type 10101. Then $x' = z'$ and $x' \eta' y' \theta' x' \in \mathcal{S}$. Let $t = x \wedge y$. Then (by (G6)) $x' \eta' y' t' x' \in \mathcal{S}$. But $t' \neq \theta'$ contrary to (S2). Therefore, this case cannot occur.

Suppose $\alpha = t x \theta y u$ is of type 21012. Let $v = x \wedge y$ (as given by (G3)). Then $t' x' \theta' y' t', v' x' \theta' y' v' \in \mathcal{S}$. By (S2), we get $t' = v'$. By (G4) applied to txv , we get $t = v$. Then by (G4) applied to uyv , we get $u = v$. Thus $t = u$ as required.

Suppose $\alpha = x \theta y t z$ is of type 10121 or 10123. Let $u = x \wedge y$. Then $x' \theta' y' t' x', x' \theta' y' u' x' \in \mathcal{S}$, so (S2) gives $t' = u'$. By (G4) applied to txu , we get $t = u$, and so x is adjacent to t . By (G4) applied to xtz we get $x = z$ as required. (In fact, this gives a contradiction for type 10123.)

Finally, suppose $\alpha = \theta x t y v$ is of type 01212, 01232 or 01234. Now $\theta' = v'$, so v', x' are adjacent in Γ . By (G5), we can lift $v'x'$ to an edge vy of G_n , where $y \in V(G_n)$ with $y' = x'$. Now consider the path $\beta := x t u v y$ in G_n . We have $f_{n+1}\beta = f_{n+1}\alpha \in \mathcal{S}$. By (G7) in G_n , we have $x = y$. Therefore v is adjacent to x . Now (G4) applied to θxv gives the contradiction $\theta' \neq v'$. Therefore, none of these three cases can occur.

We have therefore verified the inductive step, and hence constructed $f_n : G_n \rightarrow \Gamma$ for all n .

We now set $G = \bigcup_{n=0}^{\infty} G_n$, and set $f(x) = f_n(x)$ for all sufficiently large n . Properties (G2) and (G3) tell us that G is bipartite and satisfies $(\diamond(p))$. Property (G5) tells us that f is a local isomorphism. Property (G6) tells us that it respects the square structure, and (G7) that it satisfies $(*)$.

This proves Lemma 16.2.5, hence also concludes the proof of Theorem 16.2.3.

16.3. Cubical structures.

We finish this section by relating square structures to “cubical” structures, which we will discuss in more detail in the next section. First, we give some more definitions.

Let Q be a finite cube (as a median algebra). Let $\Gamma(Q)$ be the adjacency graph with vertex set Q . Note that we can recover the median structure on Q from the combinatorial metric, ρ_Q . In particular, any automorphism of $\Gamma(Q)$ is a median automorphism of Q . By a **face** of $\Gamma(Q)$ we mean the full subgraph on a face of Q . We say that a graph is **cubical** if it is isomorphic to $\Gamma(Q)$ for some cube Q . We say it is **n -cubical** if $Q = \{0, 1\}^n$, and we say that it has **dimension** n .

Let Γ be a connected graph. (As usual, we assume it has no loops or multiple edges.)

Definition. By a **cubical structure** on Γ , we mean a family, \mathcal{G} , of (embedded) cubical subgraphs of Γ , referred as **cells**, satisfying the following:

(C1): every edge of Γ is a cell,

(C2): every face of a cell is a cell, and

(C3): the intersection of two cells is a (possibly empty) disjoint union of cells.

We write \mathcal{G}_n for the set of n -cells of \mathcal{G} . Thus, \mathcal{G}_2 is a square structure on Γ satisfying (S2).

Given $a \in V(\Gamma)$, we have defined the **link**, $L(a) = L_{\Gamma}(a)$, of a as the set of adjacent vertices. Given a cubical structure, \mathcal{G} , on Γ , we give $L(a)$ the structure of a simplicial complex, where the simplices are sets of the form $P \cap L(a)$ for some $P \in \mathcal{G}$.

Recall that a simplicial complex is **flag** if every complete subgraph of its 1-skeleton lies in a simplex.

Definition. We say that (Γ, \mathcal{G}) is **locally CCAT(0)** if the link of every vertex is a flag complex.

Definition. We say that (Γ, \mathcal{G}) is **CCAT(0)** if (Γ, \mathcal{G}) is locally CCAT(0) and \mathcal{G}_2 satisfies (S1).

Here “CCAT(0)” stands for “combinatorially CAT(0)” (our terminology), where (locally) CAT(0) is the standard metric non-positive curvature condition (to be discussed in Sections 17 and 18). One often just says “CAT(0)” in the above definition. But we will be dealing with other metrics which are not CAT(0) in the

geometric sense. For clarity, we will sometimes use the term “globally CCAT(0)” to mean the same as CCAT(0).

Here is the main result relating to this:

Proposition 16.3.1. *A graph, Γ , is median if and only if it admits a cubical structure, \mathcal{G} , with respect to which it is CCAT(0). In this case, \mathcal{G} consists of all cubical subgraphs. Moreover, each cell is isometrically embedded in Γ .*

Proof. The “only if” direction follows directly from the results of Section 11. Suppose that Π is a discrete median algebra. Let $\Gamma = \Gamma(\Pi)$, and let $\mathcal{C}_n(\Pi)$ be the set of all n -cells of Π , and let \mathcal{G}_n be the set of all full subgraphs of Γ on elements of $\mathcal{C}_n(\Pi)$. The link condition is Lemma 11.4.1, and the fact that (Γ, \mathcal{G}_2) satisfies (S1) is Lemma 11.4.3. Note that, by Lemma 16.1.2, every median graph has this form.

For the converse, we just note that the link condition implies that the square complex (Γ, \mathcal{G}_2) satisfies (S3). It now follows that Γ is a median graph by Theorem 16.2.3. Moreover, by that result, $\mathcal{G}_2 = \mathcal{S}(\Gamma)$ consists of all squares of Γ .

Now if G is any cubical subgraph of $\Gamma(\Pi)$, then by Lemma 16.3.2 below, G is isometrically embedded in Γ , and $V(G)$ is a cell of Π . Moreover, we see that all squares of G lie in \mathcal{G}_2 , so by (S3), it follows that $G \in \mathcal{G}$. □

Lemma 16.3.2. *Let Π be a discrete median algebra, and let $G \leq \Gamma(\Pi)$ be a cubical subgraph. Then G is isometrically embedded in $\Gamma(\Pi)$, and $V(G) \subseteq \Pi$ is a cell of Π .*

Proof. First note that any two adjacent edges of G lie in a (unique) square of G , and that this is also a square of $\Gamma(\Pi)$.

Let $a, b \in V(G)$. We connect a, b by a path \underline{a} in G . We now apply the procedure in the proof of Lemma 13.2.2 to reduce \underline{a} to a geodesic path in $\Gamma(\Pi)$, by applying moves of type (1) and (2). By the above observation, each intermediate path lies in \underline{a} , and so we end up connecting a to b by a geodesic in $\Gamma(\Pi)$ which lies entirely in G . This shows that G is isometrically embedded in $\Gamma(\Pi)$. In particular, $V(G)$ is a cube in Π .

Since G is full, the initial observation shows that we cannot have $\#(Q \cap V(G)) = 3$ for any 2-cell, Q , of Π . Therefore, by Lemma 11.4.4, $V(G)$ is convex, hence a cell of Π . □

Note that Proposition 16.3.1 implies that Γ has a unique structure as a CCAT(0) complex.

We remark that there is a version of Lemma 16.3.2 in the more general context of quasimedial graphs — see Lemma 23.3.4.

It is natural to interpret the above in terms of a topological realisation, $\Delta(\Gamma, \mathcal{G})$, of the cubical structure. This is a real cube complex, where each cubical graph $\Gamma(Q)$ in \mathcal{G}_n is the 1-skeleton of a real cube, $\Delta(Q) \cong [0, 1]^n$. We can give this complex the CW topology (or a number of other topologies). In these terms, \mathcal{G}_2 -connectedness is equivalent to saying that $\Delta(\Gamma, \mathcal{G})$ is simply connected. The link at each vertex corresponds to the usual notion in a polyhedral complex. We will

return to this in Section 18. In the meantime, in the next section, we study this complex when $\Gamma = \Gamma(\Pi)$ is a median graph. In this case, $\Delta(\Gamma, \mathcal{G})$ will be the same space as $\Delta(\Pi)$ discussed in Sections 10 and 11.

17. CUBE COMPLEXES

The “cube complexes” we discuss in this section are (mostly) assumed to be CCAT(0) (though some of the discussion applies more generally). Such complexes are built out of euclidean cubes (that is, finite direct products of compact real intervals). They have natural structures as median algebras. They can also be given a number of different metrics and topologies. The main results, Propositions 17.1.1–17.1.8 are stated at the beginning. We then proceed to the proofs. We include a brief discussion of subdivisions of cube complexes. We finish the section with a result which shows that complete connected median algebras contain isometric copies of cube complexes (Proposition 17.4.1).

17.1. Topologies, metrics, and statement of main results.

Let Π be a discrete median algebra. In Section 10 and Subsection 11.2, we constructed a complex, $\Delta(\Pi)$, as the “realisation” of Π . Informally it can be described by taking a “real cube”, that is a copy of $[0, 1]^n$, for each cell of Π , and gluing them together in such a way as to respect the inclusion of cells. We refer to these real cubes as the “cells” of $\Delta(\Pi)$.

We saw (Lemma 10.2.2 and subsequent discussion) that $\Delta(\Pi)$ is naturally a median algebra with Π as a subalgebra. It can be given additional structure. For example, we can put a topology on $\Delta(\Pi)$ by deeming a subset to be open if its intersection with each cell of $\Delta(\Pi)$ is open in the standard topology on that cell — that is the direct product of the real intervals $[0, 1]$. To be more specific, we say that such a set is **CW-open**. With this topology and cell structure, $\Delta(\Pi)$ is a CW complex in the usual sense. We observed in Subsection 5.1, that the 1-skeleton of $\Delta(\Pi)$ is connected graph (see Lemma 5.1.1). Since the cells of $\Delta(\Pi)$ are connected, it follows easily that $\Delta(\Pi)$ is CW-connected.

In Subsection 12.1, we noted that, at least if Π is countable, then the median operation is continuous, and so $\Delta(\Pi)$ is a topological median algebra. It is unclear if this holds without the countability assumption. In any case, regardless of cardinality, one can make a number of statements regarding this topology. For example, we will show that following three statements hold.

By a **(finite) subcomplex** of $\Delta(\Pi)$ we mean a (finite) union of cells. We can assume this collection of cells to be closed under inclusion.

Proposition 17.1.1. *If $A \subseteq \Delta(\Pi)$ is finite, then $\text{hull}(A)$ is CW-compact. In fact, A lies in a finite convex subcomplex of $\Delta(\Pi)$.*

It follows immediately from Proposition 17.1.1 that $\Delta(\Pi)$ is interval-compact in the CW topology. Since the CW topology is finest among the topologies that we will be considering, this statement will hold for all the others.

We also note:

Proposition 17.1.2. *Suppose that $C \subseteq \Delta(\Pi)$ is CW-closed and convex. Then C is gated. Let Π' be the image of Π under the gate map to C . Then Π' is discrete, and C is homeomorphic in the subspace topology to $\Delta(\Pi')$ via a median isomorphism.*

We will also see:

Proposition 17.1.3. *$\Delta(\Pi)$ is contractible in the CW topology.*

This last statement will be a consequence of the fact that there is a homotopy equivalence between the CW and CAT(0) topologies on $\Delta(\Pi)$. The CAT(0) topology will be discussed below.

All the topologies we will consider will have the following property:

Definition. A topology on $\Delta(\Pi)$ is **cell-compatible** if the induced subspace topology on each cell is standard.

In other words, a topology is cell-compatible if it is at least as coarse as the CW-topology.

One can put various other cell-compatible topologies on $\Delta(\Pi)$. For example, one can take the subspace topology arising from the embedding of $\Delta(\Pi)$ in the cube $\hat{\Delta}(\Psi(\Pi))$, with the product topology. This was mentioned in Subsection 13.1, where we observed that $\Delta(\Pi)$ is a topological median algebra in this topology. (It turns out that this is the same as the compactification of $\Delta(\Pi)$ as discussed at the end of Section 12.) However, we won't have much use for this topology in the present discussion.

Of more interest here are topologies arising from various natural metrics on $\Delta(\Pi)$. To describe these, we should begin by recalling some basic metric-space theory.

Let (D, ρ) be a metric space. A **(finite) path** in D is a continuous map, $\alpha : [a, b] \rightarrow D$, to D , where $[a, b] \subseteq \mathbb{R}$ is a compact real interval with $a < b$. Its **rectifiable length** can be defined as the supremum of the sums $\sum_{i=1}^n \rho(\alpha(t_{i-1}), \alpha(t_i))$, as $(t_i)_i$ varies over all sequences, $a = t_0 \leq t_1 \leq \dots \leq t_n = b$ for all $n \in \mathbb{N}$. We write $\text{length}(\alpha) \in [0, \infty]$ for its rectifiable length. We say that α is **rectifiable** if $\text{length}(\alpha) < \infty$. Given $x, y \in D$, we write $\sigma_\rho(x, y)$ for the infimum of $\text{length}(\alpha)$ as α varies over all rectifiable paths from x to y . We write $\sigma_\rho(x, y) = \infty$ if there is no such path. It is easily checked that σ_ρ is a non-finite metric on D . (Recall that "non-finite" means that it takes values in $[0, \infty]$.) We refer to it as the **induced path-metric** on D . Since clearly, $\sigma_\rho \geq \rho$, the induced topology is at least as fine as that induced by ρ . It is also easily checked that $\sigma_{\sigma_\rho} = \sigma_\rho$.

In Subsection 13.3, we defined a "geodesic" to be a path $\alpha : [a, b] \rightarrow D$, such that $\rho(\alpha(t), \alpha(u)) = |t - u|$ for all $t, u \in [a, b]$. One checks that any geodesic is rectifiable, and that $\text{length}(\alpha) = \rho(x, y)$ where $x = \alpha(a)$ and $y = \alpha(b)$. Conversely, if $\alpha : [a, b] \rightarrow D$ is a rectifiable arc from x to y with $\text{length}(\alpha) = \rho(x, y)$, then we can reparameterise α so that it is geodesic. Recall that, by definition, (D, ρ) is a **geodesic space** if every two points of D are connected by a geodesic.

Suppose that D is compact in the topology induced by ρ , and that $\rho(x, y) < \infty$ for all $x, y \in D$. Then the infimum is attained: that is, there is a path from x to y of length equal to $\sigma_\rho(x, y)$. We can assume it to be an arc, and so, after reparameterisation, a geodesic. Given that $\sigma_{\sigma_\rho} = \sigma_\rho$, it follows that (D, σ_ρ) is a geodesic space.

We now return to our discussion of the cell complex, $\Delta(\Pi)$.

Suppose we are given a function, $w : \mathcal{W}(\Pi) \rightarrow (0, \infty)$. (We saw in Subsection 13.2, that this is equivalent to putting a median metric on Π .) Suppose we are also given some $p \in [1, \infty]$. We can put a metric, $\sigma = \sigma_{w,p}$, on $\Delta(\Pi)$ as follows.

Let $Q \subseteq \Pi$ be an n -cell, and let $\Delta(Q) \subseteq \Delta(\Pi)$ be the corresponding cell of $\Delta(\Pi)$. We write $\mathcal{W}(Q) = \{W_1, \dots, W_n\}$. As a median algebra, we can identify $\Delta(Q)$ with $\prod_{i=1}^n [0, w(W_i)] \subseteq \mathbb{R}^n$. We restrict the l^p -metric on \mathbb{R}^n to give us a metric, ρ_Q , on $\Delta(Q)$. In other words, if $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ lie in $\Delta(Q)$, then $\rho_Q(x, y)^p = \sum_{i=1}^n |x_i - y_i|^p$ if $p \in [1, \infty)$, and $\rho_Q(x, y) = \max_{1 \leq i \leq n} |x_i - y_i|$ if $p = \infty$. Note that the metrics thus defined agree on intersections of cells.

Suppose that $\alpha : [a, b] \rightarrow \Delta(\Pi)$ is a path. We say that α is *piecewise cellular* if we can write α as a concatenation of paths, $\alpha = \alpha_1 \cup \dots \cup \alpha_m$, where α_j maps into some cell, $\Delta(Q_j)$, of $\Delta(\Pi)$. We write $\text{length}(\alpha_j)$ for the length of α_j in the metric ρ_{Q_j} . We set $\text{length}(\alpha) = \sum_{j=1}^m \text{length}(\alpha_j)$. Given that the metrics agree on intersections of cells, it is easily checked that this is well defined, no matter how we write α as such a concatenation. Suppose that $x, y \in \Delta(\Pi)$. It is easily seen (and we will later verify more formally) that x and y are connected by such a path α , with $\text{length}(\alpha) < \infty$. We write $\sigma_{w,p}(x, y)$ for the infimum of $\text{length}(\alpha)$ as α varies over all such paths. It is easily checked that $\sigma := \sigma_{w,p}$ is metric on $\Delta(\Pi)$. We omit technical details for the moment, since we will later describe σ in a slightly different, but equivalent way.

We will show:

Proposition 17.1.4. *Let Π be a discrete median algebra. Let $p \in [1, \infty]$, let $w : \mathcal{W}(\Pi) \rightarrow (0, \infty)$ be any map, and let $\sigma = \sigma_{w,p}$ be the metric defined on $\Delta(\Pi)$ as above. Then any two points of $\Delta(\Pi)$ are connected by a geodesic which lies in the median interval between them. Moreover, the induced subspace topology on each cell is the standard one. If $p = 1$, then $(\Delta(\Pi), \sigma)$ is a median metric space.*

In particular, $(\Delta(\Pi), \sigma)$ is a geodesic space, and induces a cell-compatible topology. This is at least as coarse as the CW topology. Therefore, in this metric topology, closed sets are CW-closed, and CW-compact sets are compact.

The construction in the l^1 case is relatively straightforward: see Example (Ex13.4) of Subsection 13.1.

In what follows, “convex” will always mean convex in the median structure.

Proposition 17.1.5. *With respect to the metric $\sigma = \sigma_{w,p}$, any closed convex subset of $\Delta(\Pi)$ is gated, and the gate map is 1-lipschitz. The convex hull of any finite set is compact, and lies in a finite union of cells of $\Delta(\Pi)$.*

Apart from the 1-lipschitz statement, this follows immediately from Propositions 17.1.1 and 17.1.2. In fact, we will see that the gate map is nearest-point projection to the convex set.

In particular, if $a, b \in \Delta(\Pi)$, then the gate map $[x \mapsto abx] : \Delta(\Pi) \rightarrow [a, b]$, is 1-lipschitz. Putting the l^1 -product of the metric σ on $\Delta(\Pi)^3$, then the median operation, $\Delta(\Pi)^3 \rightarrow \Delta(\Pi)$, is 1-lipschitz. In particular, we deduce:

Proposition 17.1.6. *In the metric topology induced by σ , $\Delta(\Pi)$ is an interval-compact topological median algebra.*

Propositions 17.1.4 to 17.1.6 hold without any restriction on w . If one places a positive lower bound on w , that is there is some $\eta > 0$ such that $w(W) > \eta$ for all $W \in \mathcal{W}$, then one can say more:

Proposition 17.1.7. *Given a positive lower bound on w , $\Delta(\Pi)$ is contractible.*

(I don't know if this holds without the lower bound.)

Recall, from Subsection 11.11 that Π is said to be “small” if it contains no \aleph -cell for any infinite cardinal \aleph . This is equivalent to saying that $\Delta(\Pi)$ contains no infinite increasing union of cells. (By default, a “cell” is taken to be finite-dimensional.)

Proposition 17.1.8. *Suppose there is a positive lower bound on w and that Π is small. Then the metric σ is complete.*

This clearly fails in general. For example, suppose $\Pi = \mathbb{N}$ with the standard median, and suppose that the wall separating i and $i + 1$ has width 2^{-i} . Then $\Delta(\Pi)$ is isometric to $[0, 2) \subseteq \mathbb{R}$. Similarly, if Π is an \aleph -cube for some infinite cardinal \aleph (as defined in Subsection 11.11), and each wall has unit width, then again $\Delta(\Pi)$ is not complete.

Certain values of p are of particular interest. If $p = 1$, then σ is a median metric. If $p = 2$, it is CAT(0). If $p = \infty$ and there is a positive lower bound on w , then σ is an injective metric. We will say more about these cases later.

17.2. Proofs.

We now set about the proofs of the above statements. We briefly recall the construction of Subsection 10.2.

Let $\Psi = \prod_{i \in \mathcal{I}} \delta_i$, be a hypercube. Here \mathcal{I} is some indexing set, and each δ_i is a 2-point median algebra. Let $\mathcal{C}(\Psi)$ be the set of faces of Ψ .

For each $i \in \mathcal{I}$, let Δ_i be a non-trivial compact real interval with endpoints identified with δ_i . Let $\hat{\Delta} = \hat{\Delta}(\Psi) = \prod_{i \in \mathcal{I}} \Delta_i$. Thus, $\hat{\Delta}$ is a median algebra with $\Psi \subseteq \hat{\Delta}$. Let $\pi_i : \hat{\Delta} \rightarrow \Delta_i$ be the projection map. Given $Q \in \mathcal{C}(\Psi)$, then $\Delta(Q) = \text{hull}_{\hat{\Delta}}(Q) \subseteq \hat{\Delta}$. If $\Pi \subseteq \Psi$ is a subalgebra, we write $\Delta(\Pi, \Psi) = \bigcup_{Q \in \mathcal{C}(\Pi, \Psi)} \Delta(Q)$, where $\mathcal{C}(\Pi, \Psi) = \{Q \in \mathcal{C}(\Psi) \mid Q \subseteq \Pi\}$. By Lemma 10.2.2, $\Delta(\Pi, \Psi)$ is a subalgebra of $\hat{\Delta}$.

Given $x, y \in \hat{\Delta}$, let $\mathcal{I}(x, y) = \{i \in \mathcal{I} \mid \pi_i x \neq \pi_i y\}$. We can define an equivalence relation on $\hat{\Delta}$ by deeming x, y to be equivalent if $\mathcal{I}(x, y)$ is finite. We restrict this to an equivalence relation on Ψ , and let $\Theta \subseteq \Psi$ be an equivalence class. It is easily checked that Θ is convex and 1-path-connected. We write $\Delta(\Theta) = \Delta(\Theta, \Psi)$. Clearly, up to isomorphism, the choice of equivalence class does not matter. In fact, Θ is isomorphic to the cube, \mathcal{T}_{\aleph} , where $\aleph = \#\mathcal{I}$. Also $\Delta(\Theta)$ is its realisation as a cell complex, in the sense already defined. One explicit description of this is as follows.

We choose any basepoint, $\underline{0} \in \Theta$. Given any $i \in \mathcal{I}$, we identify δ_i with $\{0, 1\}$ and Δ_i with $[0, 1]$ in such a way that $\pi_i \underline{0} = 0$. Thus, $\Psi = \{0, 1\}^{\mathcal{I}}$ and $\hat{\Delta} = [0, 1]^{\mathcal{I}}$. Also $\Delta(\Theta) = \{x \in \hat{\Delta} \mid \#\mathcal{I}(x, \underline{0}) < \infty\}$, and $\Theta = \Psi \cap \Delta(\Theta)$. The map $[x \mapsto \mathcal{I}(x, \underline{0})] : \Theta \rightarrow \mathcal{T}(\mathcal{I})$ gives us an isomorphism from Θ to the cube, $\mathcal{T}(\mathcal{I})$. Given $x, y, z \in \hat{\Delta}$ with $x.z.y$, we have $\mathcal{I}(z, \underline{0}) \subseteq \mathcal{I}(x, \underline{0}) \cup \mathcal{I}(y, \underline{0})$. It follows that $\Delta(\Theta)$ is convex in $\hat{\Delta}$.

Note that if we have a partition, $\mathcal{I} = \mathcal{I}_1 \sqcup \mathcal{I}_2$, then $\Psi = \Psi_1 \times \Psi_2$, where $\Psi_j = \prod_{i \in \mathcal{I}_j} \delta_i$. We also get a splitting, $\Theta = \Theta_1 \times \Theta_2$, with $\Theta_j \subseteq \Psi_j$, and $\Delta(\Theta) = \Delta(\Theta_1) \times \Delta(\Theta_2)$. In particular, if $\mathcal{I}_2 = \{i\}$, we get $\Delta(\Theta) = \Delta(\Theta_1) \times \Delta_i \cong \Delta(\Theta_1) \times [0, 1]$.

If $A \subseteq \Delta(\Theta)$ is finite, let $\mathcal{I}(A) = \bigcup_{x \in A} \mathcal{I}(x, \underline{0})$, and let $F(A) = \{x \in \Delta(\Theta) \mid (\forall i \in \mathcal{I} \setminus \mathcal{I}(A)) (\pi_i x = 0)\}$. Thus, $F(A)$ is a cell of $\Delta(\Theta)$ containing A . In particular, $\text{hull}_{\Delta(\Theta)}(A) \subseteq F(A)$ is compact.

We can give a general description of convex subsets of $\Delta(\Theta)$ as follows.

Suppose that P_i is a non-empty convex subset of Δ_i for all $i \in \mathcal{I}$. Let $\hat{P} = \prod_{i \in \mathcal{I}} P_i \subseteq \hat{\Delta}$. Then \hat{P} is convex in $\hat{\Delta}$ and so $P := \hat{P} \cap \Delta(\Theta)$ is convex in $\Delta(\Theta)$. In fact, all convex subsets of $\Delta(\Theta)$ arise in this way.

Lemma 17.2.1. *Let $P \subseteq \Delta(\Theta)$ be convex. For each $i \in \mathcal{I}$, let $P_i = \pi_i P \subseteq \Delta_i$, and let $\hat{P} = \prod_{i \in \mathcal{I}} P_i$. Then each P_i is convex in Δ_i , and $P = \hat{P} \cap \Delta(\Theta)$.*

Proof. The fact that P_i is convex follows immediately from the fact that $\pi_i : \Delta(\Theta) \rightarrow \Delta_i$ is an epimorphism.

The remainder follows by essentially the same argument as Lemma 11.11.2. Clearly $P \subseteq \hat{P}$, so we want to show that $\hat{P} \cap \Delta(\Theta) \subseteq P$. Let $x \in \hat{P} \cap \Delta(\Theta)$, and choose $y \in P$ so as to minimise $\#\mathcal{I}(x, y)$. We claim that $\mathcal{I}(x, y) = \emptyset$. For suppose that $i \in \mathcal{I}(x, y)$. Since $x \in \hat{P}$, there is some $z \in P$ with $\pi_i z = \pi_i x$. Let $y' \in \Delta(\Theta)$ be such that $\pi_i y' = \pi_i z = \pi_i x$ and $\pi_j y' = \pi_j y$ for all $j \in \mathcal{I} \setminus \{i\}$. Then $y.y'.z$. Since P is convex, $y' \in P$. But $\mathcal{I}(x, y') = \mathcal{I}(x, y) \setminus \{i\}$, contradicting the minimality of $\mathcal{I}(x, y)$. Therefore $\mathcal{I}(x, y) = \emptyset$. It follows that $x = y \in P$ as required. \square

In particular, given any $i \in \mathcal{I}$, we can write $\Delta(\Theta) = \Delta(\Theta_1) \times \Delta_i$, and $P = R_i \times P_i$, where $R_i \subseteq \Delta(\Theta_1)$ is convex. Given any $x \in R_i$, $\{x\} \times P_i = (\{x\} \times \Delta_i) \cap P$. If P is non-empty and closed in $\Delta(\Theta)$ it follows that P_i is closed in Δ_i . Let $\omega_i : \Delta_i \rightarrow P_i$ be the gate map. Taking a direct product of these we get a map, $\hat{\omega} : \hat{\Delta} \rightarrow \hat{P}$,

which is easily seen to be a gate map to \hat{P} . Since $\Delta(\Theta)$ is convex, this restricts to a gate map, $\omega : \Delta(\Theta) \rightarrow P$. In particular, this shows:

Lemma 17.2.2. *If $P \subseteq \Delta(\Theta)$ is non-empty, closed and convex, then $P_i := \pi_i P \subseteq \Delta_i$ is closed for all i , and P is gated.*

We next consider metrics on $\Delta(\Theta)$.

Let $\lambda : \mathcal{I} \rightarrow (0, \infty)$ be any map. We equip Δ_i with a metric, ρ_i , so that it is isometric to a compact real interval of length $\lambda(i)$. Given $p \in [1, \infty)$, let $\rho = \rho_{\lambda,p}$ be the non-finite l^p -metric on $\hat{\Delta}$. In other words, if $x, y \in \hat{\Delta}$, we have $\rho(x, y)^p = \sum_{i \in \mathcal{I}} \rho_i(\pi_i x, \pi_i y)^p$ for $p \in [0, \infty)$, and $\rho(x, y) = \sup_{i \in \mathcal{I}} \rho_i(\pi_i x, \pi_i y)$ for $p = \infty$. This restricts to a genuine metric on $\Delta(\Theta)$ (i.e. it takes finite values). It induces a cell-compatible topology on $\Delta(\Theta)$.

Note that if $\lambda' \leq \lambda$ and $p' \geq p$, then $\rho_{\lambda',p'} \leq \rho_{\lambda,p}$. In particular, the topology induced by $\rho_{\lambda',p'}$ is at least as coarse as that induced by $\rho_{\lambda,p}$.

Suppose that $P \subseteq \Delta(\Theta)$ is closed and convex. The maps $\omega_i : \Delta_i \rightarrow P_i$ as defined above are all 1-lipschitz and so the gate map, $\omega : \Delta(\Theta) \rightarrow P$ is 1-lipschitz. In particular, if we put the induced l^1 -metric on $\Delta(\Theta)$ ³, then the median map is 1-lipschitz. It follows that $\Delta(\Theta)$ is a topological median algebra in the induced metric topology.

The above statements all hold without restriction on λ . Let us now assume that λ is bounded below, that is, there is some $\eta > 0$ such that $\lambda(i) \geq \eta$ for all $i \in \mathcal{I}$.

Define a map, $\psi : [-2, 2] \rightarrow [-2, 2]$ by $\psi(x) = 2x$ if $|x| \leq 1$, $\phi(x) = 2$ if $x \geq 1$, and $\phi(x) = -2$ if $x \leq -1$. Given $i \in \mathcal{I}$, let $\phi_i : \Delta_i \rightarrow [-2, 2]$ be a linear homeomorphism, and set $\theta_i = \phi_i^{-1} \circ \psi \circ \phi_i : \Delta_i \rightarrow \Delta_i$. In other words, θ_i collapses each of the two end quarter-intervals of Δ_i to a point, and expands the middle half-interval to all of Δ_i . This is 2-lipschitz in the standard metric on Δ_i . Taking a direct product of these maps, we get a map $\hat{\theta} : \hat{\Delta} \rightarrow \hat{\Delta}$, which restricts to a map $\theta : \Delta(\Theta) \rightarrow \Delta(\Theta)$.

Lemma 17.2.3. *Suppose that $\lambda : \mathcal{I} \rightarrow (\eta, \infty)$ for some $\eta > 0$, and $p \in [1, \infty)$. Then $\theta : \Delta(\Theta) \rightarrow \Delta(\Theta)$ is continuous from the topology induced by $\rho = \rho_{\lambda,p}$ to the CW topology on $\Delta(\Theta)$.*

Proof. We can assume that $p = \infty$, since this gives the coarsest topology for any given λ . We show that θ is continuous at any given $x \in \Delta(\Theta)$.

Let $F = F(\{x\})$ be the cell containing x as defined earlier. (That is, the set of $y \in \Delta(\Theta)$ such that $\pi_i y = 0$ whenever $\pi_i x = 0$.) Note that $\theta(F) = F$. Let $U \subseteq \Delta(\Theta)$ be a CW-open neighbourhood of $\theta(x)$. Then $F \cap U$ is open, and so contains an ϵ -neighbourhood of $\theta(x)$ in the metric ρ for some $\epsilon > 0$. Let $y \in \Delta(\Theta)$ with $\rho(x, y) \leq \min(\eta/4, \epsilon/2)$. We claim that $\theta(y) \in U$. To see this, note that if $i \in \mathcal{I} \setminus \mathcal{I}(x, \underline{0})$ (that is $\pi_i x = 0$) then $\rho_i(\pi_i y, 0) \leq \eta/4$, so $\pi_i \theta y = \theta_i \pi_i y = 0$, so (by definition of F) $\theta(y) \in F$. On the other hand, if $i \in \mathcal{I}(x, \underline{0})$, then $\rho_i(\theta_i \pi_i x, \theta_i \pi_i y) \leq 2\rho_i(\pi_i x, \pi_i y) \leq 2(\epsilon/2) = \epsilon$. Thus, $\theta(y) \in F \cap U \subseteq U$, as claimed. This shows that θ is continuous at x as claimed. \square

Note that the identity map, $\Delta(\Theta) \rightarrow \Delta(\Theta)$ is continuous from the CW topology to the metric topology. We claim that this is a homotopy inverse to θ . This amounts to constructing a homotopy from θ to the identity which is continuous in both the CW topology and the metric topology, on taking a direct product with the standard topology on the real interval $[0, 1]$.

We can do this by linear isotopy. Define $\bar{\psi} : [-2, 2] \times [0, 1] \rightarrow [-2, 2]$ by $\bar{\psi}(x, t) = x(t + 1)$ for $|x| \leq 1$, $\bar{\psi}(x, t) = x + 2t - xt$ for $x \geq 1$, and $\bar{\psi}(x, t) = x - 2t - xt$ for $x \leq -1$. Define $\bar{\theta}_i : \Delta_i \times [0, 1] \rightarrow \Delta_i$ by $\bar{\theta}_i(x, t) = \phi_i^{-1}(\bar{\psi}(\phi_i(x), t))$. Taking direct products, we get a map $\hat{\Delta} \times [0, 1] \rightarrow \hat{\Delta}$, which restricts to a map $\bar{\theta} : \Delta(\Theta) \rightarrow \Delta(\Theta)$ which is a homotopy from the identity to θ . Clearly this is continuous on any cell of $\Delta(\Theta)$ times $[0, 1]$, and so $\bar{\theta}$ is continuous in the CW topology. Also, it is 2-lipschitz with respect to the ρ metric on $\Delta(\Theta)$ and its l^1 product with the standard metric on $[0, 1]$. Therefore it is continuous also with the metric topology. This shows that $\theta : \Delta(\Theta) \rightarrow \Delta(\Theta)$ is a homotopy inverse to the identity map as claimed.

We have shown the following special case of Dowker's Theorem [Do]:

Lemma 17.2.4. *The identity map, $\Delta(\Theta) \rightarrow \Delta(\Theta)$, is a homotopy equivalence from the CW topology to the metric topology.*

In particular, simple connectedness is well defined independently of which of these topologies we take.

We now apply the above to any discrete median algebra, Π .

As noted in Subsection 11.11, Π embeds in a cube, Θ . We can describe this using the above notation as follows. We write $\mathcal{W}(\Pi) = \{W_i \mid i \in \mathcal{I}\}$, where \mathcal{I} is some indexing set. Given $i \in \mathcal{I}$, let $\delta_i = \{W_i^-, W_i^+\}$, and let $\Psi = \prod_{i \in \mathcal{I}} \delta_i$. We can now identify Π as a 1-path-connected subalgebra of $\Theta \subseteq \Psi$ as defined above. We write $\Delta(\Pi) = \Delta(\Pi, \Psi) \subseteq \Delta(\Theta, \Psi) = \Delta(\Theta)$. This is a subalgebra and subcomplex of $\Delta(\Theta)$. Note that (by Lemma 11.11.3) there is a natural identification of $\mathcal{W}(\Pi)$ with $\mathcal{W}(\Theta)$.

Suppose $\Pi' \subseteq \Pi$ is convex in Π . Then we can identify $\mathcal{W}(\Pi')$ with a subset of $\mathcal{W}(\Pi)$, and so we can identify $\Psi(\Pi')$ with a face, Ψ' , of $\Psi(\Pi)$. In this way, we can identify $\Delta(\Pi') = \Delta(\Pi', \Psi')$ as a subalgebra and subcomplex of $\Delta(\Pi) = \Delta(\Pi, \Psi)$.

Let $i \in \mathcal{I}$ and \mathcal{E}_i be the set of 1-cells of Π which cross W_i . Let $\Pi_i = \bigcup \mathcal{E}_i$. We saw in Subsection 11.5, that \mathcal{E}_i is a parallel class in Π , with the intrinsic structure of a discrete median algebra. Moreover, Π_i is convex in Π and intrinsically isomorphic to $\mathcal{E}_i \times \delta_i$. (See Lemmas 7.2.6 and 7.2.7.) Thus, $\Delta(\Pi_i)$ is a convex subset and subcomplex of $\Delta(\Pi)$. It is intrinsically isomorphic to $\Delta(\mathcal{E}_i) \times \Delta_i$. Note that $\pi_i : \Delta(\Pi) \rightarrow \Delta_i$ restricted to $\Delta(\Pi_i)$ is just projection to the second factor.

Proof of Proposition 17.1.1. Let $A \subseteq \Delta(\Pi)$ be finite. By Lemma 7.4.6, $\text{hull}_{\Delta(\Pi)} A = \Delta(\Pi) \cap \text{hull}_{\Delta(\Theta)} A$. We saw above that $\text{hull}_{\Delta(\Theta)} A$ lies in a cell, $F(A)$, of $\Delta(\Theta)$. Now $\Delta(\Pi) \cap F(A)$ is a finite convex subcomplex of $\Delta(\Pi)$, so the statement follows. \square

Proof of Proposition 17.1.2. Let $C \subseteq \Delta(\Pi)$ be convex. Let $P = \text{hull}_{\Delta(\Theta)} C \subseteq \Delta(\Theta)$. By Lemma 7.4.6, $C = P \cap \Delta(\Pi)$. Let $P_i = \pi_i P \subseteq \Delta_i$, and let $\hat{P} = \prod_{i \in \mathcal{I}} P_i$. By Lemma 17.2.1, we have $P = \hat{P} \cap \Delta(\Theta)$. Thus, $C = \hat{P} \cap \Delta(\Pi)$.

Suppose that $i \in \mathcal{I}$. Then we can write $\Delta(\Theta)$ as a product $\Delta(\Theta_i) \times \Delta_i$ as described above. Let Π_i be as described above, so that $\Delta(\Pi_i) \subseteq \Delta(\Pi)$ is convex, and intrinsically isomorphic to $\Delta(\mathcal{E}_i) \times \Delta_i$ where $\Delta(\mathcal{E}_i)$ is a subcomplex of R_i . Note that $C \cap \Delta(\Pi_i) = P \cap \Delta(\Pi_i)$ is convex. If this is empty, then P_i is a single endpoint of Δ_i . If not, then it is direct product of a convex subset of $\Delta(\mathcal{E}_i)$ with P_i . Similarly as with Lemma 17.2.2, we see that P_i is closed. We can now define a gate map $\omega_i : \Delta_i \rightarrow P_i$. Taking the direct product over $i \in \mathcal{I}$, we get a gate map, $\omega : \hat{\Delta} \rightarrow \hat{P}$, which restricts to a gate map $\omega : \Delta(\Theta) \rightarrow P$. Its further restriction to $\Delta(\Pi)$ is a gate map to C (by Lemma 7.4.8).

Next we show that $\Pi' := \omega_C \Pi$ is discrete. Given $i \in \mathcal{I}$, write $\epsilon_i = \partial P_i \subseteq \Delta_i$. Suppose $Q \in \mathcal{C}(\Pi)$, so that $\Delta(Q)$ is a cell of $\Delta(\Pi)$. Write $\mathcal{I}_Q = \{i \in \mathcal{I} \mid Q \cap W_i \neq \emptyset\}$. We can write $\Delta(Q) = \prod_{i \in \mathcal{I}_Q} D_i$, where $\pi_i D_i = \Delta_i$. Suppose $C' := C \cap \Delta(Q) \neq \emptyset$. By Lemma 7.5.1, $C = \prod_{i \in \mathcal{I}_Q} L_i$, where $L_i \subseteq D_i$ is a closed interval. Note that $\pi_i L_i = P_i$ and $\pi_i(\partial L_i) = \epsilon_i$. Now $C' = \text{hull}(Q')$, where $Q' = \prod_{i \in \mathcal{I}_Q} \partial L_i$, and $\omega_{C'}(Q) = Q'$. By Lemma 7.3.5, we have $\omega_C \omega_{\Delta(Q)} = \omega_{\Delta(Q)} \omega_C = \omega_{C'}$. Thus $\omega_{C'} \Pi = \omega_C \omega_{\Delta(Q)} \Pi = \omega_C Q = \omega_{C'} Q = \omega_{\Delta(Q)} \omega_C \Pi$. In particular, $\Pi' \cap \Delta(Q) = \omega_{C'} Q = Q'$. We see that Π' meets every cell of $\Delta(\Pi)$ in a (possibly empty) finite set. It now follows from Proposition 17.1.1 that Π' is discrete as claimed.

Now let $\mathcal{I}(C) \subseteq \mathcal{I}$ be the set of indices, i , such that P_i is not a singleton. We can identify $\mathcal{W}(\Pi')$ with $\{W_i \mid i \in \mathcal{I}(C)\}$. Let $\Psi_C = \prod_{i \in \mathcal{I}(C)} \delta_i$, and $\tilde{\Delta}(\Psi_C) = \prod_{i \in \mathcal{I}(C)} \Delta_i$. The projection map $\tilde{\Delta}(\Psi) \rightarrow \tilde{\Delta}(\Psi_C)$ is injective on C . Write $\mathcal{C}(\Pi, C) = \{Q \in \mathcal{C}(\Pi) \mid \mathcal{I}_Q \subseteq \mathcal{I}(C)\}$. Then $\mathcal{C}(\Pi') = \{Q' \mid Q \in \mathcal{C}(\Pi, C)\}$. Writing $\Delta(Q') = \text{hull}(Q')$, we have $C = \bigcup_{Q \in \mathcal{C}(\Pi, C)} \Delta(Q') = \bigcup_{Q' \in \mathcal{C}(\Pi')} \Delta(Q')$ which we can identify with $\Delta(\Pi')$.

We need to check that the CW topology on $\Delta(\Pi') = C$ agrees with the subspace topology of the CW topology on $\Delta(\Pi)$. Suppose U is open in the subspace topology. Then $U = V \cap \Delta(\Pi')$ where $V \subseteq \Delta(\Pi)$ is open. For each cell $\Delta(Q')$ of $\Delta(\Pi')$, we have $U \cap \Delta(Q') = V \cap \Delta(Q')$ which is open in $\Delta(Q')$ since $V \cap \Delta(Q)$ is open in $\Delta(Q)$. Thus, U is open in the CW topology. Conversely, suppose U is open in the CW-topology. For each cell $\Delta(Q')$, $U \cap \Delta(Q')$ is open in $\Delta(Q')$, so $U \cap \Delta(Q') = V_Q \cap \Delta(Q')$ where $V_Q \subseteq \Delta(Q)$ is open in $\Delta(Q)$. Let V be the union of all such V_Q . Then V is open in $\Delta(\Pi)$ and $U = V \cap \Delta(\Pi')$. Thus, U is open in the subspace topology. \square

We will postpone the proof of Proposition 17.1.3 for the moment.

Now let $w : \mathcal{W}(\Pi) \rightarrow (0, \infty)$ be any map. Define $\lambda : \mathcal{I} \rightarrow (0, \infty)$. Let $\lambda(i) = w(W_i)$, and set $\rho = \rho_{w,p} = \rho_{\lambda,p}$ be the l^p metric defined above, restricted to $\Delta(\Pi)$. Note that if $C \subseteq \Delta(\Pi)$ is CW-closed and convex, then the gate map

$\omega : \Delta(\Pi) \rightarrow C$ is 1-lipschitz in the metric ρ , and it follows that $\Delta(\Pi)$ is a topological median algebra in the induced topology.

It is more usual to equip $\Delta(\Pi)$ with the path-metric, σ_ρ , induced by ρ . We will write it as $\sigma_{w,p}$, and abbreviate it to σ . We will see that this is a geodesic metric. In fact, we will see that a geodesic between two points can be written as a finite concatenation of paths, each of which maps into a cell of $\Delta(\Pi)$. From this, it follows that σ agrees with the definition (intrinsic to $\Delta(\Pi)$) which we gave at the beginning. One advantage of using σ is that it is invariant under subdivisions of the cube complex, as we will observe in Subsection 17.3.

Proof of Proposition 17.1.4. Let $x, y \in \Delta(\Pi)$. Then x, y lie in some cell, F , of $\Delta(\Theta)$. Now $G := F \cap \Delta(\Pi)$ is a finite union of faces, $G = \bigcup_j F_j$, of F . Now G is convex in $\Delta(\Pi)$, and hence connected. Since G is compact, we see that x, y are connected in G by an intrinsic geodesic, α , with respect to the metric ρ . Since the gate map to G is 1-lipschitz in the metric ρ , we see that this is also geodesic in $\Delta(\Pi)$. Moreover, since the gate map to each face, F_j , is 1-lipschitz, we can assume that α meets such each face (if at all) in a connected subpath. In this way, α is a finite concatenation of such paths. This shows that the induced metric σ agrees with the metric as previously defined. It also follows that α is geodesic also in the metric σ , and so $\Delta(\Pi)$ is a geodesic metric space. We also note that (since the gate map to the interval $[x, y]$ is 1-lipschitz) we can take this geodesic to lie in $[x, y]$.

Finally, if $p = 1$, then ρ is a median metric (see Example (Ex13.4) of Subsection 13.1). By Proposition 17.1.1, $\Delta(\Pi)$ is interval compact, so by Lemma 13.3.2, ρ is a geodesic metric. Therefore $\sigma = \rho$. \square

Proof of Proposition 17.1.5. The gate map to C is 1-lipschitz with respect to ρ , and therefore also with respect to σ . (The image of a rectifiable path under a 1-lipschitz map is rectifiable, and its length cannot increase.) As observed above, the remainder of Proposition 17.1.5 follows directly from Propositions 17.1.1 and 17.1.2. \square

Proof of Proposition 17.1.6. As observed earlier, this is an immediate consequence of Proposition 17.1.5. \square

Suppose now that we have a lower bound, say $\eta > 0$, on w . Disjoint cells of $\Delta(\Pi)$ are a distance at least η apart in the ρ -metric. Suppose $a, b \in \Delta(\Pi)$ with $\rho(a, b) < \eta$. Let F, F' be cells of $\Delta(\Pi)$ containing a, b respectively. From this we see that $F \cap F' \neq \emptyset$. In this case, it follows easily that $\sigma(a, b) \leq 2\rho(a, b)$. (Let c be the projection of a to F' . Then $c \in F \cap F'$ and $\sigma(a, c) = \rho(a, c) \leq \rho(a, b)$ and $\sigma(b, c) = \rho(b, c) \leq \rho(a, b)$.) Since clearly $\rho \leq \sigma$, it follows that ρ and σ induce the same topology on $\Delta(\Pi)$. Moreover, a sequence is convergent (respectively Cauchy) with respect to ρ if and only if it is convergent (respectively Cauchy) with respect to σ . In particular, completeness with respect to the two metrics is equivalent.

We have a map $\theta : \Delta(\Pi) \rightarrow \Delta(\Pi)$, continuous from the ρ metric topology to the CW topology, obtained by restricting the map $\theta : \Delta(\Theta) \rightarrow \Delta(\Theta)$. Since $\rho \leq \sigma$,

this is also continuous from the σ metric topology. We also have a homotopy, $\bar{\theta} : \Delta(\Pi) \times [0, 1] \rightarrow \Delta(\Pi)$, again restricting the map $\bar{\theta}$ defined earlier. This is a homotopy from θ to the identity on $\Delta(\Pi)$. It is continuous in the CW topology. It is also 2-lipschitz with respect to the metric ρ , hence also with respect to σ . It is thus a homotopy inverse to the identity map on $\Delta(\Pi)$, hence a homotopy equivalence.

It is well known that, in certain cases at least, $\Delta(\Pi)$, is contractible. For example, if $p = 2$ and w is identically 1, then we get the standard CAT(0) metric on $\Delta(\Pi)$. This contractible, as we note in Section 18.

We now finally get to:

Proof of Proposition 17.1.3. The CW topology and the standard CAT(0) metric topology are homotopy equivalent. Since the latter is contractible, so is the former. □

Proof of Proposition 17.1.7. The CW topology and the $\sigma_{w,p}$ metric topology are homotopy equivalent. Since the former is contractible by Proposition 17.1.3, so is the latter. □

Proof of Proposition 17.1.8. We suppose that $w(W_i) \geq \eta > 0$ for all $i \in \mathcal{I}$. In what follows it will be convenient to view $\Delta(\Pi)$ as a subset of $\hat{\Delta} \equiv \prod_{i \in \mathcal{I}} [0, w(W_i)]$, with $\underline{0} \in \Delta(\Pi)$.

Let $(x_n)_{n \in \mathbb{N}}$ be a cauchy sequence in $\Delta(\Pi)$ with respect to the metric σ , hence also with respect to ρ . We can assume that $\rho(x_m, x_n) \leq \eta/2$ for all m, n . Since $(\pi_i x_n)_n$ is cauchy, we have $\pi_i x_n \rightarrow y_i$ for some $y_i \in [0, w(W_i)]$. Let $y = (y_i)_{i \in \mathcal{I}} \in \hat{\Delta}$. Write $\mathcal{I} = \mathcal{I}_0 \sqcup \mathcal{I}_1 \sqcup \mathcal{I}_2$, where $\mathcal{I}_0 = \{i \in \mathcal{I} \mid y_i = 0\}$, $\mathcal{I}_1 = \{i \in \mathcal{I} \mid y_i = w(W_i)\}$ and $\mathcal{I}_2 = \mathcal{I} \setminus (\mathcal{I}_0 \cup \mathcal{I}_1)$. Note that \mathcal{I}_1 is finite (since $\rho_i(\pi_i x_0, y_i) < \eta$ for all i , so $\mathcal{I}_1 \subseteq \mathcal{I}(\underline{0}, x_0)$).

Let $\Theta \subseteq \Psi$ be the set of $z \in \Psi$ such that $\pi_i z = 0$ for all $i \in \mathcal{I}_0$ and $\pi_i z = w(W_i)$ for all $i \in \mathcal{I}_1$. Then Θ convex in Ψ , and isomorphic to $\mathcal{T}(\mathcal{I}_2)$.

We claim that $\Theta \subseteq \Pi$. Since Π is assumed to be small, it then follows that \mathcal{I}_2 is finite. It in turn follows that $\rho(x_n, y) \rightarrow 0$, and so $y \in \Delta(\Pi)$ and so $(x_n)_n$ is convergent.

To prove the claim, let $\mathcal{J} \subseteq \mathcal{I}_2$ be any finite subset. For all sufficiently large n , $\pi_i x_n$ lies in the interior of $[0, w(W_i)]$ for all $i \in \mathcal{J}$. Therefore the corresponding cell of $\Delta(\Psi)$ lies in $\Delta(\Pi)$. □

17.3. Subdivisions.

We make a few observations about subdivisions of cube complexes.

First consider a finite totally ordered set, $\Sigma = \{x_0, x_1, \dots, x_m\}$, with $m \geq 2$ and $x_0 < x_1 < \dots < x_m$. Write $\partial\Sigma = \{x_0, x_m\}$. Let W_j be the wall crossing $\epsilon_j := \{x_{j-1}, x_j\}$, so that $\mathcal{W}(\Sigma) = \{W_1, \dots, W_m\}$. Write $\Psi = \prod_{j=1}^m \epsilon_j$. Then $\Delta(\Sigma)$ is a path in the 1-skeleton of $\Delta(\Psi)$, which we can identify with the real interval

$\Delta(\partial\Sigma)$. Given any map, $w : \mathcal{W}(\Sigma) \rightarrow (0, \infty)$, and $p \in [1, \infty]$, let $\rho_{w,p}$ be the l^p -metric on $\Delta(\Sigma)$ defined above. Then the 1-cells of $\Delta(\Sigma)$ are respectively isometric to the real intervals $[0, w(W_j)]$. We see that in the induced path-metric, $\sigma_{w,p}$, $\Delta(\Sigma)$ is isometric to the real interval $[0, \sigma(x_0, x_m)]$, with $\sigma(x_0, x_m) = \sum_{j=1}^m w(W_j)$.

More generally, suppose $\Sigma = \prod_{i=1}^n \Sigma_i$, where Σ_i is a non-trivial finite totally ordered set. Then $Q := \prod_{i=1}^n \partial\Sigma_i \leq \Sigma$ is an n -cube. We have $\mathcal{W}(\Sigma) = \bigsqcup_{i=1}^n \mathcal{W}(\Sigma_i)$. We can identify $\Delta(\Sigma) = \Delta(Q)$, by using the above construction on each factor Σ_i . Given $w : \mathcal{W}(\Sigma) \rightarrow (0, \infty)$ and $p \in [1, \infty]$, we can define the metrics $\rho_{w,p}$ and $\sigma_{w,p}$. In the metric $\sigma_{w,p}$, $\Delta(\Sigma)$ is an l^p -product of the real intervals $\Delta(\Sigma_i) \cong \Delta(\partial\Sigma_i)$.

Finally suppose that Π is a discrete median algebra and $\Lambda \leq \Pi$ is a subdividing subalgebra, as defined in Subsection 11.10. By Lemma 11.10.1, we can write $\Pi = \bigcup_{Q \in \mathcal{C}(\Lambda)} \Sigma(Q)$, where each $\Sigma(Q)$ is a direct product of finite totally ordered sets. By the above, we can identify each $\Delta(\Sigma(Q))$ with $\Delta(Q)$, and so identify $\Delta(\Pi)$ with $\Delta(\Lambda)$. If Π is a discrete median metric space, then we get an induced median metric on Λ . If $p \in [1, \infty]$ we see that in the respective l^p path-metrics σ , the identification of $\Delta(\Pi)$ with $\Delta(\Lambda)$ is an isometry. In short, subdividing a cube complex does not change the induced path metric. (We have only described this for $\text{CCAT}(0)$ cube complexes, though a similar observation holds for cube complexes in general.)

17.4. Subalgebras.

We finish this section with following result which shows how cube complexes feature prominently in the theory of median metric spaces more generally.

Proposition 17.4.1. *Let M be a geodesic median metric space, and let $\Pi \leq M$ be a discrete subalgebra. Let $\Delta = \Delta(\Pi) \supseteq \Pi$ be the realisation of Π with the l^1 metric. Then the inclusion of Π into M extends to an isometric embedding of Δ into M .*

Note that such an isometric embedding is necessarily a median monomorphism.

Proof. We first extend the inclusion to the 1-skeleton of Δ . Let $a, b \in \Pi$ be adjacent in Π , so that $[a, b]_\Delta$ is isometric to a real interval of length $\rho(a, b)$. Since M is geodesic, there is an isometric embedding of $[a, b]_\Delta$ into $[a, b]_M$ fixing a and b . Moreover, we can choose these embeddings so that if $a', b' \in \Pi$ is parallel to a, b , then the respective embeddings commute with the translation of $[a, b]_\Delta$ to $[a', b']_\Delta$ and the translation of $[a, b]_M$ to $[a', b']_M$. (Just do this for one edge, a, b , in each parallel class to begin with, and then compose with translations.)

Now let Q be an n -cell of Π . By Lemma 10.3.5, we can identify $\Delta(Q) := \text{hull}_\Delta(Q) \equiv \prod_{i=1}^n D_i$ and $\text{hull}_M(Q) \equiv \prod_{i=1}^n D'_i$, where $D_i \equiv [a_i, b_i]_\Delta$ and $D'_i \equiv [a_i, b_i]_M$, and where a_i, b_i are 1-faces of Q . We have defined embeddings $D_i \hookrightarrow D'_i$, so these combine to give us an embedding $\Delta(Q) \hookrightarrow M$. Since the maps on the 1-skeleton commute with translation, these maps are consistent. Assembling these embeddings, we get a map, $f : \Delta \rightarrow M$, with $f|_\Pi$ the inclusion map. This is isometric on each cell, and hence 1-lipschitz. It remains to check that it is an isometric embedding.

Let $x, y \in \Delta$. We can find $a, b \in \Pi$ so that $a.x.y.b$ holds in Δ . Now $\rho_M(a, fx) \leq \rho_\Delta(a, x)$, $\rho_M(b, fy) \leq \rho_\Delta(b, y)$ and $\rho_M(fx, fy) \leq \rho_\Delta(x, y)$. Also:

$$\begin{aligned} \rho_\Delta(a, b) &= \rho_M(a, b) \leq \rho_M(a, fx) + \rho_M(fx, fy) + \rho_M(fy, b) \\ &\leq \rho_\Delta(a, x) + \rho_\Delta(x, y) + \rho_\Delta(y, b) = \rho_\Delta(a, b), \end{aligned}$$

so we have equality throughout, and $\rho_M(fx, fy) = \rho_\Delta(x, y)$ as required. \square

18. THE CAT(0) PROPERTY

We begin with a brief review the general notion of a CAT(0) (or locally CAT(0)) space. The idea is to capture the notion of non-positive curvature in purely metric terms.

We give some results relating median convexity to metric convexity. We finish with a geometric proof that the 1-skeleton of a CAT(0) cube complex is a median graph. Put together with the general Cartan-Hadamard Theorem for CAT(0) spaces, this gives another proof of Theorem 16.2.3.

18.1. Definition and basic facts.

In what follows we write σ_E for the euclidean (l^2) metric on \mathbb{R}^2 .

Let (M, σ) be a metric space. Given $a, b \in M$, write $I_\sigma(a, b) = \{x \in M \mid \sigma(a, b) = \sigma(a, x) + \sigma(x, b)\}$. A **kite** in M is a 4-tuple, (a, b, c, d) , of elements of M such that $d \in I_\sigma(a, b)$. (This is not standard terminology.) A **comparison kite** for (a, b, c, d) is a kite, (a', b', c', d') , in the euclidean plane, (\mathbb{R}^2, σ_E) , such that $\sigma_E(a', b') = \sigma(a, b)$, $\sigma_E(b', c') = \sigma(b, c)$, $\sigma_E(c', a') = \sigma(c, a)$, $\sigma_E(a', d') = \sigma(a, d)$ and (then necessarily) $\sigma_E(b', d') = \sigma(b, d)$. It follows from the triangle inequality in M that such a comparison kite always exists. Moreover, it is unique up to isometry of \mathbb{R}^2 . In particular, the quantity, $\sigma_E(c', d')$ is well defined. (Of course, one could write down an explicit formula in terms of the other distances involved. But there is little point.)

Definition. We say that the kite (a, b, c, d) satisfies the **CAT(0)** condition if $\sigma_E(c', d') \leq \sigma(c, d)$.

Definition. A **CAT(0)** space is a geodesic metric space in which every kite satisfies the CAT(0) condition.

Examples of CAT(0) spaces are Hadamard manifolds (simply connected complete riemannian manifolds of non-positive curvature). One can also construct many examples as euclidean (or hyperbolic) polyhedral complexes, as we will mention.

We list a few basic properties of CAT(0) spaces.

Any CAT(0) space, (M, σ) , is uniquely geodesic — that is, there is a unique geodesic connecting any two points, $a, b \in M$. Its image is $I_\sigma(a, b)$. Any local geodesic (that is a path that is geodesic on some neighbourhood of any point in the domain) is a (global) geodesic. Any CAT(0) space is contractible. The completion of a CAT(0) space is CAT(0). One can associate a canonical boundary, ∂M , to M ,

and equip $M \sqcup \partial M$ with a canonical hausdorff topology, such that M is open and dense in $M \sqcup \partial M$, and the subspace topology on M agrees with the original metric topology. Since this space is not a disjoint union topologically, it is more usual to write it as $M \cup \partial M$. If M is proper (that is, complete and locally compact), then $M \cup \partial M$ (hence also ∂M) is compact.

As with median algebras, a notion of convexity plays a central role.

Definition. A subset, $C \subseteq M$, is *metrically convex* if $I_\sigma(a, b) \subseteq C$ for all $a, b \subseteq C$.

This is usually just termed “convex”, but will later want to distinguish it from median convexity. We will for the moment drop the “metrically”.

Clearly intersections of convex sets are convex, so it makes sense to talk of “convex hulls”. The closure of a convex set is convex. A CAT(0) space is locally convex: that is every point has a base of convex neighbourhoods (which we can take to be closed). More generally, we say that a subset $C \subseteq M$ is *locally convex* if every point of C has a convex neighbourhood U such that $C \cap U$ is convex. It turns out that a subset $C \subseteq M$ is convex if and only if it is connected and locally convex. If C is convex, and $x \in M$, then there is at most one “nearest” point $y \in C$ (that is which minimises $\rho(x, y)$). If such a point exists for all $x \in M$, then we have a nearest-point projection, $\omega : M \rightarrow C$. Such a projection is necessarily 1-lipschitz.

Definition. We say that a metric space, M , is *locally CAT(0)* if every point $x \in M$ has a neighbourhood $U \ni x$, such that any kite in U satisfies the CAT(0) condition.

One can show that any simply connected geodesic space which is locally CAT(0) is (globally) CAT(0). This is a version of the “Cartan-Hadamard Theorem”.

As an example, any non-positively curved riemannian manifold is locally CAT(0). Now any complete riemannian manifold is geodesic, and it follows that any Hadamard manifold is CAT(0), as we observed above.

Of more direct interest here is the fact that one can construct locally CAT(0) spaces by gluing together euclidean polyhedra so as to form a euclidean polyhedral complex. The local CAT(0) condition turns out to be equivalent to a condition on the link of every cell (viewed as a spherical polyhedral complex). If the complex is simply connected, then it will be (globally) CAT(0). We note that by Dowker’s Theorem [Do], simple connectedness is the same whether interpreted in the metric topology or the CW topology (see Lemma 17.2.4 here).

This is a very general construction, and there are numerous accounts of it — see the Notes to this section. Here we shall focus on the case of cube complexes.

Let Γ be a connected graph. Let \mathcal{G} be a cubical structure on Γ , as defined in Subsection 16.3. We can construct a cell complex, $\Delta(\Gamma, \mathcal{G})$, as follows. For each $G \in \mathcal{G}$, we take a copy of $\Delta(G) := \Delta(Q)$, where Q is the vertex set of \mathcal{G} , and glue these together according to the combinatorial structure. More precisely, if G' is a

face of G , then we identify $\Delta(G')$ as a subset of $\Delta(G)$. The transitive closure of this relation is an equivalence relation on the disjoint union $\bigsqcup_{G \in \mathcal{G}} \Delta(G)$, and we take the quotient. The quotient topology is the same as the CW topology, but we are more interested here in $\Delta(\Gamma, \mathcal{G})$ as a metric space.

Let us suppose that we have a metric, σ , on Δ such that $(\Delta(\Gamma, \mathcal{G}), \sigma)$ is a geodesic space, and such that the induced path-metric, σ_G , on any cell, $G \in \mathcal{G}$, $\Delta(G) \subseteq \Delta$, is euclidean. More precisely, if $G \in \mathcal{G}_n$, then $(\Delta(G), \sigma_G)$ is isometric to the euclidean metric, σ_E , on $\prod_{i=1}^n [0, r_i] \subseteq \mathbb{R}^n$, for some $r_1, \dots, r_n > 0$.

The link condition alluded to above can now be made more precise:

Theorem 18.1.1. *$(\Delta(\Gamma, \mathcal{G}), \sigma)$ is locally CAT(0) if and only if (Γ, \mathcal{G}) is locally CCAT(0).*

Recall that “locally CCAT(0)” was defined in Subsection 16.3, to mean that the link of every vertex is a flag simplicial complex. (This is, in fact, equivalent to saying that the link of every cell is a flag simplicial complex.)

Theorem 18.1.1 was proven by Gromov. There are now numerous accounts of it (see the Notes to this section). We will not reproduce this here.

As we observed in Subsection 16.2, simple connectedness of Δ is equivalent to a combinatorial condition. We defined “CCAT(0)” to mean that Δ is simply connected and locally CCAT(0). We therefore have:

Theorem 18.1.2. *$(\Delta(\Gamma, \mathcal{G}), \sigma)$ is CAT(0) if and only if (Γ, \mathcal{G}) is CCAT(0).*

Now by Proposition 16.3.1, we see that Γ is a median graph, and so by Lemma 16.1.2, it has the form $\Gamma = \Gamma(\Pi)$ for a discrete median algebra Π . Retrospectively, we see that $\Delta(\Gamma, \mathcal{G})$ is the realisation, $\Delta(\Pi)$, of a discrete median algebra. Moreover, Π has median metric, ρ , for which that σ has the form $\sigma_{w,2}$ as discussed in Section 17, where w is the width function.

Note that the above applies to any such metric, σ , on $\Delta(\Gamma, \mathcal{G})$ (and so, the property of being CAT(0) is in fact independent of this choice). One particular such metric is obtained by giving each cell the structure of a unit euclidean cube: the “standard” CAT(0) structure. Here the width function is identically 1.

18.2. Hyperplanes and convexity in CAT(0) cube complexes.

So let us suppose that we are in the situation of Theorem 18.1.2. Thus, $\Delta = \Delta(\Pi)$ is a CAT(0) cube complex, with CAT(0) metric σ , and median metric ρ . By default, the metric we use will be σ , and the topology will be that induced by σ .

A key notion in the study of such cube complexes is that of a “hyperplane”, which cuts the space into two “halfspaces”. This can be thought of as a geometric interpretation of a wall in this context.

Let $W \in \mathcal{W}(\Pi)$, and let $\mathcal{E}(W)$ be the parallel class consisting of all those 2-cells of Π which cross W . Given $e \in \mathcal{E}(W)$, we can write $\Delta(e) = [e^-, e^+]$, where $e \cap W^\pm = \{e^\pm\}$. There is an isometry, $\iota : \Delta(e) \rightarrow [0, w(W)]$, with $\iota(e^-) = 0$ and $\iota(e^+) = w(W)$, where $w(W) = \rho(e^-, e^+) = \sigma(e^-, e^+)$ is the width of the wall W .

Let $\pi_W = \iota \circ \omega_{\Delta(e)} : \Delta \rightarrow [0, w(W)]$, where $\omega_{\Delta(e)}$ is the gate map to $\Delta(e)$. This is a median epimorphism. Since all cells of $\mathcal{E}(W)$ are parallel, this is well defined, independently of the choice of e (see Lemma 11.5.1). Recall that $\Delta(\Pi)$ embeds in to $\hat{\Delta} = \prod_{W \in \mathcal{W}(\Pi)} [0, w(W)]$, so that π_W is the projection to the $[0, w(W)]$ factor. The median metric ρ is the induced l^1 metric.

Let $t \in (0, w(W))$. We write $H^- = \pi_W^{-1}[0, t]$, $\bar{H}^- = \pi_W^{-1}[0, t]$, $H^+ = \pi_W^{-1}(t, w(W)]$ and $\bar{H}^+ = \pi_W^{-1}(t, w(W)]$. These are all median convex. We generally refer to them as **generic halfspaces**. In particular, $\{H^-, \bar{H}^+\}$ and $\{\bar{H}^-, H^+\}$ are both walls of Δ . It is easily seen that \bar{H}^\pm is the closure of H^\pm . (It is also the closure with respect to the metric ρ , and also with respect to the CW topology.) We write $P = \bar{H}^- \cap \bar{H}^+$. This is referred to as a **hyperplane**. It is also convex, and the common boundary of H^- and H^+ .

Lemma 18.2.1. *The halfspaces, H^\pm and \bar{H}^\pm are metrically convex.*

Proof. We just outline the proof, since this is a standard fact about CAT(0) cube complexes.

It's enough to deal with H^+ . Since H^+ is median convex, it is connected. By the general property of convexity in CAT(0) spaces mentioned earlier, it's enough to show that H^+ is locally metrically convex. Again, it's enough to check this at points in its boundary, namely P .

So let $x \in P$. Now x has a neighbourhood, U , isometric to an l^2 product, $D \times (-\epsilon, \epsilon)$, where D is a geodesic space and $(-\epsilon, \epsilon)$ is a real interval, and such that $H^+ \cap U$ corresponds to $D \times (0, \epsilon)$ under this isometry. (Consider a neighbourhood of x in $\hat{\Delta} \supseteq \Delta$, where $(-\epsilon, \epsilon)$ arises from an interval in the $[0, w(W)]$ factor.) In the l^2 metric, $D \times (0, \epsilon)$ is metrically convex in $D \times (-\epsilon, \epsilon)$. \square

Lemma 18.2.2. *If $a, b \in \Delta$ are distinct, there is a generic halfspace, $H \subseteq \Delta$, with $a \in H$ and $b \in H^* := \Delta \setminus H$.*

Proof. Consider the embedding $\Delta \subseteq \hat{\Delta}$. There is some $W \in \mathcal{W}(\Pi)$ such that $\pi_W a \neq \pi_W b$. Choose some $t \in (\pi_W a, \pi_W b) \subseteq [0, w(W)]$, and let $H = \pi_W^{-1}[0, t)$. \square

Lemma 18.2.3. *If $a, b \in \Delta$, then $I_\sigma(a, b) \subseteq [a, b]$.*

Proof. Let $c \in \Delta \setminus [a, b]$. By Lemma 18.2.2, there is a generic halfspace, H , with $abc \in H$ and $c \in H^*$. Since H^* is median convex, $a, b \in H$. Since H is metrically convex, $I_\sigma(a, b) \subseteq H$. Therefore, $c \notin I_\sigma(a, b)$. \square

This can be rephrased by saying:

Corollary 18.2.4. *Any median convex subset of Δ is metrically convex.*

It also follows that any σ -geodesic in Δ is monotone (with respect to the median) hence also a ρ -geodesic up to reparameterisation.

There is a partial converse to Corollary 18.2.4. By a ‘‘subcomplex’’ of Δ , we mean a subset of the form $\Delta(\mathcal{D}) := \bigcup_{Q \in \mathcal{D}} \Delta(Q)$, where $\mathcal{D} \subseteq \mathcal{C}(\Pi)$. We can assume

that \mathcal{D} is closed under inclusion. It's easily seen that any subcomplex is closed in the metric topology. We note:

Lemma 18.2.5. *A geodesically convex subcomplex of Δ is median convex.*

Proof. We write the subcomplex as $\Delta(\mathcal{D})$ as above. Recall from Subsection 10.1 that we can embed Π as a subalgebra of a hypercube, Ψ .

We first claim that $\mathcal{D} = \mathcal{C}(\Pi', \Psi)$, where $\Pi' \subseteq \Pi$ is the set of 0-cells (singleton elements) of \mathcal{D} . (As in Subsection 10.2, $\mathcal{C}(\Pi', \Psi)$, denotes the set of cells of Ψ contained in Π' .) Certainly $\mathcal{C}(\Pi', \Psi) \subseteq \mathcal{D}$. For the reverse inclusion, let $Q \in \mathcal{D}$. Consider the euclidean geodesic connecting two antipodal corners of Q in $\Delta(Q)$. This is a local σ -geodesic hence a global σ -geodesic (since σ is CAT(0)). This must lie in $\Delta(\mathcal{D})$ by metric convexity. Since its midpoint lies in the interior of $\Delta(Q)$, it follows that $Q \in \mathcal{D}$. This proves the claim.

We now claim that Π' is convex in Π . By Lemma 11.4.4, it is enough to show that Π' is 1-path-connected and locally convex in Π .

The fact that \mathcal{D} is 1-path-connected is a simple consequence of the fact that $\Delta(\mathcal{D})$ is path-connected: take any path between any two given 0-cells in Δ , and consider the sequence of cells of Δ which it passes through.

For local convexity, suppose that $Q = \{a, b, c, d\}$ is a 2-cell of Π with $a, b, c \in \Pi'$ and with d antipodal to a . As above, the euclidean geodesic from b to c in $\Delta(Q)$ lies in $\Delta(\mathcal{D})$ and so $Q \in \mathcal{D}$, so $Q \subseteq \Pi'$. This proves local convexity.

We have shown that $\Delta(\mathcal{D}) = \Delta(\Pi', \Psi)$, and so it is convex in $\Delta = \Delta(\Pi, \Psi)$ by Lemma 10.2.2. □

An important special case of the above is the “standard” CAT(0) metric on Δ , already alluded to in the previous subsection. This is defined by deeming the width of each wall to be equal to 1, so that $\hat{\Delta} = [0, 1]^{\mathcal{W}(\Pi)}$ and each cell of Δ is isometric to a unit euclidean cube in the induced path metric (and indeed in the induced metric). Here is usual to take hyperplanes to be the mid-hyperplanes: that is taking $t = \frac{1}{2}$. In this way, the mid-hyperplanes are in bijective correspondence with the set of walls of Δ , and the family of such hyperplanes is preserved under any automorphism of Δ (that is, induced by an automorphism of Π).

18.3. Another description of the equivalence with median graphs.

Let us now go back to the earlier discussion, where we started from a connected graph, Γ , with a cubical structure, \mathcal{G} , and constructed a complex $\Delta(\Pi)$. In showing that the 1-skeleton is a median graph, we invoked Proposition 16.3.1, whose proof (in this direction) relied on Theorem 16.2.3. The combinatorial proof of Theorem 16.2.3 was a bit of a rigmarole, which occupied all of Subsection 16.2. However, with the help of CAT(0) geometry, we can now give another, quite different, proof. This shows directly that the combinatorial metric on $V(\Gamma)$ is a median metric, without invoking Lemma 16.1.1 or Proposition 16.1.3.

To this end, we give another criterion for a graph to be median.

Let Γ be a connected graph. Let $\Pi = V(\Gamma)$ be the vertex set, and $E(\Gamma)$ be edge set.

Suppose we have a partition of $E(\Gamma)$ as $E(\Gamma) = \bigsqcup_{W \in \mathcal{W}} \mathcal{E}(W)$, indexed by some set \mathcal{W} . We refer to the elements of \mathcal{W} as **walls**. We assume:

(G1): For each $W \in \mathcal{W}$, no two elements of $\mathcal{E}(W)$ are adjacent. Moreover, \mathcal{E} separates Γ into two connected subgraphs.

We denote the subgraphs (arbitrarily) as $\Gamma^-(W)$ and $\Gamma^+(W)$. We write $W^\pm = \Pi \cap \Gamma^\pm(W)$. Thus, $W = W^- \sqcup W^+$. We refer to W^- and W^+ as **halfspaces**. We assume:

(G2): For each $W \in \mathcal{W}$, there are retractions, $\omega_W^- : \Gamma \rightarrow \Gamma^-(W)$ and $\omega_W^+ : \Gamma \rightarrow \Gamma^+(W)$, both of which collapse each edge of $\mathcal{E}(W)$ to a point.

By a ‘‘retraction’’ we mean that ω_W^\pm is the identity of $\Gamma^\pm(W)$ and sends each edge of Γ to either a vertex or an edge. In particular, it restricts to a map $\omega_W^\pm : \Pi \rightarrow W^\pm$. This is 1-lipschitz with respect to the combinatorial metric, $\rho := \rho_\Gamma$. Note that any geodesic in Γ meets each $\mathcal{E}(W)$ in at most one edge, in which case we say that it ‘‘crosses’’ W . In particular, we see that W^\pm is ‘‘convex’’ in the sense that any geodesic in Γ with both endpoints in W^\pm lies entirely in $\Gamma^\pm(W)$. It follows that $\rho_{\Gamma^\pm(W)}$ is the metric ρ restricted to $\Gamma^\pm(W)$.

Given $a, b \in \Pi$, we write $a|_W b$ to mean that $\{a, b\}$ meets both W^- and W^+ . We write $\mathcal{W}(a, b) = \{W \in \mathcal{W} \mid a|_W b\}$. Note that any geodesic from a to b crosses precisely the walls of $\mathcal{W}(a, b)$. In particular, we have $\rho(a, b) = \#\mathcal{W}(a, b)$, and so if $a \neq b$, $\mathcal{W}(a, b) \neq \emptyset$. Note also that $\mathcal{W}(a, b) \subseteq \mathcal{W}(a, c) \cup \mathcal{W}(c, b)$ for all $a, b, c \in \Pi$.

We write $[a, b]_\rho = \{x \in \Pi \mid \rho(a, b) = \rho(a, x) + \rho(x, b)\}$. From the above, we see that $c \in [a, b]_\rho$ if and only if $\mathcal{W}(a, c) \cap \mathcal{W}(c, b) = \emptyset$, and if and only if $\mathcal{W}(a, b) = \mathcal{W}(a, c) \cap \mathcal{W}(c, b)$. Also, $[a, b]_\rho$ is the intersection of all halfspaces containing a and b .

We need one additional assumption:

(G3): Suppose $W_1, W_2 \in \mathcal{W}$ and $W_1^- \cap W_2^- \neq \emptyset$. Then $w_{W_1^-}(W_2^-) = W_1^- \cap W_2^-$.

Note that if $\Gamma = \Gamma(\Pi)$ is a median graph, then the above conditions are satisfied with $\mathcal{W}(\Pi)$ the set of walls, $\mathcal{E}(W)$ the set of edges crossing W , and ω_W^\pm the gate map to W^\pm . There were all described in Subsection 11.5.

We want to prove the converse. In other words, we assume (G1)–(G4) and want to show that $\#[a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho = 1$ for all $a, b, c \in \Pi$.

To show uniqueness, suppose $m, m' \in [a, b]_\rho \cap [b, c]_\rho \cap [c, a]_\rho$. There is some $W \in \mathcal{W}$ with $m \in W^-$ and $m' \in W^+$. We can suppose that $a, b \in W^-$. But then $m' \subseteq [a, b]_\rho \subseteq W^-$, giving a contradiction.

For existence, we first note the following. Suppose $a, b \in \Pi$ and $W_1 \in \mathcal{W}(a, b)$. Then $\omega_{W_1}^-([a, b]_\rho) \subseteq [a, b]_\rho$. To see this, let $m \in [a, b]_\rho$. We can suppose $a \in W_1^-$. If $\omega_{W_1}^- m \notin [a, b]_\rho$, there is a wall $W_2 \in \mathcal{W}$ with $a, b \in W_2^-$ and $\omega_{W_1}^- m \in W_2^+$. But now $W_1^- \cap W_2^- \neq \emptyset$, and so by (G3), $\omega_{W_1}^- m \in W_1^- \cap W_2^-$, giving a contradiction.

Now suppose $a, b, c \in \Pi$. Choose $m \in [a, b]_\rho$ so as to minimise $\rho(c, m)$. We claim that $m \in [a, c]_\rho$. For suppose not. Then there is some wall $W \in \mathcal{W}$ with $a, c \in W^-$ and $m \in W^+$. Thus, $b \in W^+$. By the previous paragraph, $\omega_W^- c \in [a, b]_\rho$. Since ω_W^- is a retraction, and collapses each edge of $\mathcal{E}(W)$, we have $\rho(c, \omega_W^- m) < \rho(c, m)$ contradicting minimality. This proves the claim. Similarly, $m \in [b, c]_\rho$. This shows existence.

We have verified that $V(\Gamma)$ is a median metric space, and so Γ is median graph.

In summary, we have shown:

Lemma 18.3.1. *Let Γ be a connected graph. Let $E(W) = \bigsqcup_{W \in \mathcal{W}} \mathcal{E}(W)$ be a partition of the edge set, together with a family of maps, ω_W^\pm for $W \in \mathcal{W}$, satisfying (G1)–(G3) above. Then Γ is a median graph.*

In fact, we see retrospectively that $\mathcal{W} = \mathcal{W}(\Pi)$ is the set of walls of Π as originally defined, \mathcal{E} is the set of edges crossing W , and ω_W^\pm is the gate map to W^\pm .

We now apply this to graphs with a CCAT(0) structure.

Let (Γ, \mathcal{G}) be CCAT(0). We put a path metric σ , on $\Delta := \Delta(\Gamma, \mathcal{G})$ using piecewise cellular paths, as described in Section 17, so that each cell is isometric to a unit euclidean cube. This is a geodesic metric. By Theorem 18.1.2, (Δ, σ) is CAT(0). We now invoke some standard constructions from the geometry of CAT(0) cube complexes.

We first note that if $K \subseteq \Delta$ is any metrically convex subcomplex, we have a nearest-point retraction, $\omega_K : \Delta \rightarrow K$, which is 1-lipschitz and sends cells to cells. In particular, it restricts to a 1-lipschitz retraction, $\Gamma \rightarrow \Gamma \cap K$. If K' is another such metrically convex subcomplex with $K \cap K' \neq \emptyset$, then $\omega_K(K') = K \cap K'$.

For example, we can construct a family of mid-hyperplanes. Each 1-cell of Δ meets exactly one such hyperplane (at its midpoint). Such a hyperplane, P , is metrically convex, and cuts Δ into two convex “halfspaces”. The union of all cells of Δ contained in such a halfspace is a convex subcomplex of Δ . (We have removed a “collar” of the form $P \times (0, \frac{1}{2})$ from the halfspace.)

We now see that Γ satisfies properties (G1)–(G3) above, and conclude, by Lemma 18.3.1, that Γ is a median graph.

19. SPACES WITH MEASURED WALLS

There is a kind of duality between median metric spaces and “spaces with measured walls”. These generalise the notion of “spaces with walls” (the discrete case) discussed Subsection 9.4. To some extent this is an elaboration on the duality we described there. The main results here in this regard are Theorem 19.2.8 and Propositions 19.2.9 and 19.2.10.

The notion of a space with measured walls was introduced in [CherMV], and the connection with median metric spaces is explored in [ChatteDH]. One of the main motivations has been its connection with the Haagerup property of groups. (See the Notes to this section.) Some of the constructions have their origins in the work of Sageev [Sa], which introduces a generalisation of the Bass-Serre theory of group splittings, where trees are replaced by CAT(0) cube complexes. We give more background to these ideas in the Notes to this section.

There are many natural examples of spaces with measured walls, in addition to those arising directly from median metric spaces. We give a few examples at the end of this section. First, we recall some basic definitions from measure theory.

19.1. Definitions.

Let \mathcal{W} be, for the moment, any set. (We will eventually think of it as a set of “walls” in some sense.) Let $\mathcal{P}(\mathcal{W})$ be its power set, thought of as a boolean algebra. Recall that a subset, $\mathcal{R} \subseteq \mathcal{P}(\mathcal{W})$ is a **ring** (in the sense of measure theory) if $A \triangle B \in \mathcal{R}$ and $A \cap B \in \mathcal{R}$ for all $A, B \in \mathcal{R}$. (That is to say, it is a subring of $\mathcal{P}(\mathcal{W})$ with its structure as a boolean ring.) Note that $A \cup B = A \triangle B \triangle (A \cap B)$, and so any ring is a sublattice of $\mathcal{P}(\mathcal{W})$, hence also a median subalgebra of $\mathcal{P}(\mathcal{W})$. We will generally assume that $\mathcal{R} \neq \emptyset$. (Since $A \triangle A = \emptyset$, this is equivalent to saying that $\emptyset \in \mathcal{R}$.)

Let $\mu : \mathcal{R} \rightarrow [0, \infty]$. We say that μ is **finitely additive** if $\mu(\emptyset) = 0$ and $\mu(A \cup B) = \mu(A) + \mu(B)$ for all disjoint $A, B \in \mathcal{R}$. (Here we adopt the convention that $t + \infty = \infty$ for all $t \in [0, \infty]$.) We say that it is **σ -additive** if in addition, $\mu(A_n) \rightarrow \mu(A)$ for any increasing sequence, $(A_n)_{n \in \mathbb{N}}$ in \mathcal{R} , for which $A := \bigcup_{n=0}^{\infty} A_n \in \mathcal{R}$. We say that μ is **finite** if $\mu(A) < \infty$ for all $A \in \mathcal{R}$. (The requirement that $\mu(\emptyset) = 0$ is almost redundant here. It is just to rule the case that $\mu(A) = \infty$ for all A .)

A **σ -ring**, \mathcal{M} , is a ring which is also closed under countable increasing unions. We say that a function $\mu : \mathcal{M} \rightarrow [0, \infty]$ is a **measure** defined on \mathcal{M} if it is σ -additive.

Given a ring, \mathcal{R} , let $\mathcal{M}(\mathcal{R})$ be the σ -ring generated by \mathcal{R} . One can show that any finite σ -additive function on \mathcal{R} extends uniquely to a measure defined on $\mathcal{M}(\mathcal{R})$.

We note that one could go one step further and extend μ to a **complete** measure, $\mu : \mathcal{M}' \rightarrow [0, \infty]$, where \mathcal{M}' is a σ -ring containing \mathcal{M} . This means that if $A \in \mathcal{M}'$ with $\mu(A) = 0$ and $B \subseteq A$, then $B \in \mathcal{M}'$ (so also, $\mu(B) = 0$).

Of course, the above statements can be made more general, but this is all we will need here.

Now let X be any set, and let $\mathcal{L}(X)$ be the set of proper partitions of X into two non-empty subsets. Let $\mathcal{W} \subseteq \mathcal{L}(X)$ be any subset. (We think of elements of \mathcal{W} as “walls”.) We say that $W \in \mathcal{W}$ **separates** $x, y \in X$ if x and y lie in different elements of the partition. We write $\mathcal{W}(x, y) \subseteq \mathcal{W}$ for the set of walls separating x and y .

Definition. A *space with measured walls* $(X, \mathcal{W}, \mathcal{M}, \mu)$ consists of a set X , a subset $\mathcal{W} \subseteq \mathcal{L}(X)$, a σ -ring $\mathcal{M} \subseteq \mathcal{P}(\mathcal{W})$ containing $\mathcal{W}(x, y)$ for all $x, y \in X$, and a measure $\mu : \mathcal{M} \rightarrow [0, \infty]$ defined on \mathcal{M} such that $\mu(\mathcal{W}(x, y)) < \infty$ for all $x, y \in X$.

We will refer to an element of a wall, $W \in \mathcal{W}$, as a *halfspace*. We will generally denote the set of all halfspaces by \mathcal{H} . This is a subproset of the proper power set, $\mathcal{P}_0(X) := \mathcal{P}(X) \setminus \{\emptyset, X\}$.

19.2. Duality with median metric spaces.

Our first aim is to show that a median metric space canonically admits such a structure, where \mathcal{W} is the set of all walls as we originally defined them. To this end, we will first construct a finitely additive function on the ring generated by all sets of the form $\mathcal{W}(x, y)$, where $\mathcal{W}(x, y) \subseteq \mathcal{W}$ is the set of walls separating x, y in the original sense defined in Subsection 8.1. (In fact, this is sufficient for most of what we do here.) We can then apply general results to extend it to a genuine (complete) measure on a σ -ring.

In the following discussion, it is convenient to adjoin the trivial partition, $\{\emptyset, M\}$, to the set of walls, $\mathcal{W}(M)$ of a median algebra, M . We write $\mathcal{W}^\infty(M) = \mathcal{W}(M) \sqcup \{\{\emptyset, M\}\}$. In this way, if $N \leq M$ is any subalgebra, we have a natural map $\mathcal{W}^\infty(M) \rightarrow \mathcal{W}^\infty(N)$, just by intersecting the elements of a partition with N . (This is just for expository convenience. The trivial partition does not play any significant role in what follows.)

First, we consider the discrete case.

Recall (from Subsection 13.2) that if (Π, ρ) is a discrete median metric space, then to each wall $W \in \mathcal{W}(\Pi)$ we have associated a “width”, $w_\Pi(W)$. Let $\mathcal{R}(\Pi)$ be the ring of finite subsets of $\mathcal{W}(\Pi)$. Given a finite subset, $A \subseteq \mathcal{W}(\Pi)$, we set $\mu_\Pi(A) = \sum_{W \in A} w_\Pi(W)$. Clearly this is finitely additive (hence also σ -additive). Note that if $x, y \in \Pi$, then $\rho(x, y) = \mu_\Pi(\mathcal{W}(x, y))$. In this case, $\mathcal{M}(\Pi)$ consists of the set of all countable subsets of $\mathcal{W}(\Pi)$. (A special case is where ρ is the combinatorial metric on Π . In this case, $\mu(A) = \#A$ for all $A \in \mathcal{R}(\Pi)$. This extends to the counting measure $\mu : \mathcal{W}(\Pi) \rightarrow \mathbb{N} \cup \{\infty\}$, by setting $\mu(A) = \infty$ for all infinite A .)

Now let Θ be a finite median metric space, and $\Pi \leq \Theta$ be a subalgebra. We have a natural map, $r : \mathcal{W}^\infty(\Theta) \rightarrow \mathcal{W}^\infty(\Pi)$. If $W \in \mathcal{W}(\Pi)$, then by definition, $w_\Pi(W) = \rho(x, y)$ for any 1-cell, $\{x, y\}$ of Π , which crosses W . By construction, $r^{-1}(W) = \mathcal{W}_\Theta(x, y)$, and so $w_\Pi(W) = \sum \{w_\Theta(W') \mid pW' = W\} = \mu_\Theta(r^{-1}(W))$. From this, we get:

Lemma 19.2.1. $\mu_\Pi(A) = \mu_\Theta(r^{-1}A)$ for all $A \subseteq \mathcal{W}(\Pi)$.

Now let M be any non-empty median algebra. A *wall-interval* is a subset of $\mathcal{W}(M)$ of the form $\mathcal{W}(x, y)$ for some $x, y \in M$. (This includes \emptyset , setting $x = y$.) Given a subset $Y \subseteq M$, let $\mathcal{R}(M, Y)$ be the ring of $\mathcal{P}(\mathcal{W}(M))$ generated

by $\{\mathcal{W}_M(x, y) \mid x, y \in Y\}$. Let $\mathcal{R}(M) := \mathcal{R}(M, M)$ be the ring generated by all wall-intervals.

(Note that this agrees with the earlier definition when M is discrete — in this case, every singleton of $\mathcal{W}(M)$ is a wall-interval.)

Lemma 19.2.2. *If $A \in \mathcal{R}(M)$, there is a finite subalgebra, $\Pi \subseteq M$ such that $A \in \mathcal{R}(M, \Pi)$.*

Proof. We can write A as a finite expression involving the binary operations, \cap and Δ , with arguments of the form $\mathcal{W}(x, y)$ with $x, y \in M$. Let $Y \subseteq M$ be the set of such x, y featuring in this expression. Let $\Pi = \langle Y \rangle$ be the subalgebra of M generated by Y . Then $A \in \mathcal{R}(M, Y) \subseteq \mathcal{R}(M, \Pi)$. \square

Suppose $A \in \mathcal{R}(M, \Pi)$, and let $p : \mathcal{W}^\infty(M) \rightarrow \mathcal{W}^\infty(\Pi)$ be the natural map.

Lemma 19.2.3. $p^{-1}pA = A$.

Proof. If $x, y \in \Pi$, then $\mathcal{W}_M(x, y) = p^{-1}\mathcal{W}_\Pi(x, y)$, and so $p^{-1}p(\mathcal{W}_M(x, y)) = \mathcal{W}_M(x, y)$. Note also that the property that $p^{-1}pA = A$ is closed under the operations \cap and Δ . \square

(Note that this implies that any element of $\mathcal{R}(M)$ is a finite disjoint union wall-intervals. This is because, in the above notation, $A = \bigsqcup_{W \in pA} p^{-1}W$, and $p^{-1}W = \mathcal{W}_M(x, y)$ where $\{x, y\}$ is any 1-cell of Π crossing W .)

Now suppose that (M, ρ) is a median metric space.

If $A \in \mathcal{R}(M, \Pi)$, for $\Pi \leq M$ finite, then we set $\mu(A, \Pi) := \mu_\Pi(pA)$, as defined above.

Lemma 19.2.4. *Suppose $\Pi, \Pi' \leq M$ are finite subalgebras. If $A \in \mathcal{R}(M, \Pi) \cap \mathcal{R}(M, \Pi')$, then $\mu(A, \Pi) = \mu(A, \Pi')$.*

Proof. Let $\Theta = \langle \Pi \cup \Pi' \rangle$. Then $A \in \mathcal{R}(M, \Theta)$. We have maps $q : \mathcal{W}^\infty(M) \rightarrow \mathcal{W}^\infty(\Theta)$ and $r : \mathcal{W}^\infty(\Theta) \rightarrow \mathcal{W}^\infty(\Pi)$. The composition, rq , is the natural map $\mathcal{W}^\infty(M) \rightarrow \mathcal{W}^\infty(\Pi)$. By Lemma 19.2.3, $A = p^{-1}pA = (rq)^{-1}pA = q^{-1}r^{-1}pA$, and so $qA = r^{-1}(pA)$. By Lemma 19.2.1, we have $\mu(A, \Pi) = \mu_\Pi(pA) = \mu_\Theta(r^{-1}(pA)) = \mu_\Theta(qA) = \mu(A, \Theta)$. Similarly, $\mu(A, \Pi') = \mu(A, \Theta)$. \square

Given $A \in \mathcal{R}(M)$, we can now define $\mu(A) := \mu(A, \Pi)$, where $\Pi \leq M$ is any finite subalgebra with $A \in \mathcal{R}(M, \Pi)$, as given by Lemma 19.2.2.

Lemma 19.2.5. *Let $A, B \in \mathcal{R}(M)$ be disjoint. Then $\mu(A \cup B) = \mu(A) + \mu(B)$.*

Proof. Let $A \in \mathcal{R}(M, \Pi)$, $B \in \mathcal{R}(M, \Pi')$ and set $\Theta = \langle \Pi \cup \Pi' \rangle$. Then $A, B, A \cup B \in \mathcal{R}(M, \Theta)$. By definition, $\mu(A) = \mu_\Theta(qA)$, $\mu(B) = \mu_\Theta(qB)$ and $\mu_\Theta(A \cup B) = \mu_\Theta(q(A \cup B))$. Now $q^{-1}qA = A$ and $q^{-1}qB = B$, and so $qA \cap qB = \emptyset$, so $\mu(A \cup B) = \mu_\Theta(qA \cup qB) = \mu_\Theta(qA) + \mu_\Theta(qB) = \mu(A) + \mu(B)$. \square

(Another way to view this is that, following the observation after Lemma 19.2.3, we can write $A = \mathcal{W}_M(x_1, y_1) \sqcup \cdots \sqcup \mathcal{W}_M(x_n, y_n)$, where $x_i, y_i \in M$. For any such

expression, we have $\mu(A) = \sum_{i=1}^n \rho(x_i, y_i)$. This also implies that μ is the unique extension to a finitely additive function on $\mathcal{R}(M)$.)

Finally note that if $x, y \in M$, then $\mathcal{W}_M(x, y) \in \mathcal{R}(M, \Pi)$, where $\Pi = \{x, y\}$. Now the natural map from $\mathcal{W}^\infty(M)$ to $\mathcal{W}^\infty(\Pi)$ sends all of $\mathcal{W}_M(x, y)$ to the single element of $\mathcal{W}(\Pi)$. This has width $\rho(x, y)$. Thus, $\mu(\mathcal{W}_M(x, y)) = \rho(x, y)$.

In summary, we have shown:

Proposition 19.2.6. *Let (M, ρ) be a median metric space. Let $\mathcal{R}(M) \subseteq \mathcal{P}(\mathcal{W}(M))$ be the ring generated by wall-intervals. Then there is a unique map $\mu : \mathcal{R}(M) \rightarrow [0, \infty)$, such that $\mu(A \cup B) = \mu(A) + \mu(B)$ for all disjoint $A, B \in \mathcal{R}(M)$, and such that $\mu(\mathcal{W}_M(x, y)) = \rho(x, y)$ for all $x, y \in M$.*

To obtain a genuine measure, we need to verify σ -additivity.

Suppose again that M is any median algebra. Recall from Subsection 7.5 that $\mathcal{W}^\infty(M)$ has a natural topology as a compact totally disconnected space. In this topology, $\mathcal{W}(x, y)$ is a clopen subset for all $x, y \in M$. It follows that all elements of $\mathcal{R}(M)$ are clopen. In particular, we deduce:

Lemma 19.2.7. *Let $(A_n)_{n \in \mathbb{N}}$ be a decreasing family of non-empty elements of $\mathcal{R}(M)$. Then $\bigcap_{n=0}^\infty A_n \neq \emptyset$.*

In particular, it follows that if $(A_n)_{n \in \mathbb{N}}$ an increasing union of sets in $\mathcal{R}(M)$ with $A = \bigcup_{n=0}^\infty A_n \in \mathcal{R}(M)$, then $A_n = A$ for all sufficiently large n . This is much stronger than what we need to see that σ -additivity holds. Therefore we get:

Theorem 19.2.8. *Let (M, ρ) be a median metric space. Let $\mathcal{M} \subseteq \mathcal{P}(\mathcal{W}(M))$ be the σ -ring generated by wall-intervals. Then there is a unique measure $\mu : \mathcal{M} \rightarrow [0, \infty]$, such that $\mu(\mathcal{W}_M(x, y)) = \rho(x, y)$ for all $x, y \in M$.*

As noted earlier, we could further extend μ to a complete measure.

We also note that this construction can be adapted to a submedian space: that is a metric space S that can be isometrically embedded into a median metric space M . Let $\mathcal{W}(S)$ be the set of bipartitions of S which are induced by some wall $W \in \mathcal{W}(M)$. (That is to say, it has the form $\{S \cap W^-, S \cap W^+\}$, with both these sets non-empty.) Given $A \subseteq \mathcal{W}(S)$, let $A_M \subseteq \mathcal{W}(M)$ be the set of walls inducing some element of A . We deem A_M to be measurable if A is, and set their measures to be equal. This gives a measure on $\mathcal{W}(S)$, though it might depend on the embedding into M .

We now consider how to go in the opposite direction, and construct a median pseudometric space from a space with measured walls.

Let X be a set, and let $\mathcal{P}_0(X) := \mathcal{P}(X) \setminus \{\emptyset, X\}$ be the proper power set with its structure as a proset (as defined in Subsection 9.1). Given $P \in \mathcal{P}_0(X)$, let $\pi(P) = \{P, P^*\}$. Let $\mathcal{L}(X) = \{\pi(P) \mid P \in \mathcal{P}_0(X)\}$ be the set of proper binary partitions of X mentioned earlier. Thus $\pi : \mathcal{P}_0(X) \rightarrow \mathcal{L}(X)$ is the quotient map under the involution $[P \mapsto P^*]$.

Let $\mathcal{W} \subseteq \mathcal{L}(X)$ be some subset. Let $\mathcal{H} = \mathcal{H}(\mathcal{W}) = \pi^{-1}\mathcal{W} \subseteq \mathcal{P}_0(X)$. This is a subproset of $\mathcal{P}_0(X)$. In this context, we think of an element of \mathcal{W} as a “wall” and an element of \mathcal{H} as a “halfspace”.

As before, we say that $W = \{P, P^*\}$ *separates* $x, y \in X$ if $\{x, y\} \cap P \neq \emptyset$ and $\{x, y\} \cap P^* \neq \emptyset$. We write $\mathcal{W}(x, y) \subseteq \mathcal{W}$ for the set of walls separating x and y . Let $\mathcal{H}(x, y) = \{H \in \mathcal{H} \mid x \in H, y \in H^*\}$. Thus $\mathcal{H}(x, y) \sqcup \mathcal{H}(y, x) = \pi^{-1}\mathcal{W}(x, y)$.

Recall that $\mathcal{P}(\mathcal{W})$ is a median algebra, with the usual power-set median. The following is all we will need to apply our construction.

Definition. A *space with weakly measured walls*, $(X, \mathcal{W}, \mathcal{M}, \mu)$, consists of a set X , a subset $\mathcal{W} \subseteq \mathcal{L}(X)$, a median subalgebra, $\mathcal{M} \subseteq \mathcal{P}(\mathcal{W})$ containing $\mathcal{W}(x, y)$ for all $x, y \in X$, and a function, $\mu : \mathcal{M} \rightarrow [0, \infty)$ such that if $A, B \in \mathcal{M}$ are disjoint and $A \cup B \in \mathcal{M}$, then $\mu(A \cup B) = \mu(A) + \mu(B)$.

Clearly this is weaker than the earlier definition.

The aim now is to construct a canonical median algebra, $\Lambda(X)$, and a median pseudometric, ρ , on $\Lambda(X)$, together with a natural map $\eta : X \rightarrow \Lambda(X)$ such that $\rho(\eta(x), \eta(y)) = \mu(\mathcal{W}(x, y))$ for all $x, y \in X$.

The construction of $\Lambda(X)$ follows from the discussion of Subsection 9.2. Let $\mathcal{F}(\mathcal{H})$ be the set of flows on the proset \mathcal{H} . This is a subalgebra of the power set, $\mathcal{P}(\mathcal{H})$. Given $x \in X$, let $\eta(x) = \{H \in \mathcal{H} \mid x \in H\}$. Thus $\eta(x) \in \mathcal{F}(\mathcal{H})$. Let $\Lambda(X) = \langle \eta(X) \rangle$ be the subalgebra of $\mathcal{F}(\mathcal{H})$ generated by the image $\eta(X)$. We thus get a map $\eta : X \rightarrow \Lambda(X)$. Note that if $x, y \in X$, then by definition, $\mathcal{W}(x, y) = \eta(x) \Delta \eta(y)$. We aim to put a median pseudometric on $\Lambda(X)$.

Given $R, S \in \mathcal{F}(\mathcal{H})$, then $R \Delta S$ has the form $\pi^{-1}\mathcal{L}(R, S)$ for some subset, $\mathcal{L}(R, S) \subseteq \mathcal{W}$. Note that if $T \in [R, S]_{\mathcal{F}(\mathcal{H})}$, then $R \cap S \subseteq T \subseteq R \cup S$, and so $R \Delta S = (R \Delta T) \sqcup (S \Delta T)$, and so $\mathcal{L}(R, S) = \mathcal{L}(R, T) \sqcup \mathcal{L}(S, T)$.

Given $R_1, R_2, R_3, S \in \mathcal{P}(\mathcal{H})$ we have

$$(R_1 R_2 R_3) \Delta S = (R_1 \Delta S)(R_2 \Delta S)(R_3 \Delta S).$$

(Recall that taking the symmetric difference with a fixed set gives a median automorphism of a power set.) Therefore

$$\mathcal{L}(R_1 R_2 R_3, S) = \mathcal{L}(R_1, S) \mathcal{L}(R_2, S) \mathcal{L}(R_3, S)$$

in $\mathcal{P}(\mathcal{W})$. Since \mathcal{M} is a subalgebra of $\mathcal{P}(\mathcal{W})$, it follows that if $R_1, R_2, R_3, S_1, S_2, S_3 \in \mathcal{F}(\mathcal{H})$ with $\mathcal{L}(R_i, S_j) \in \mathcal{M}$ for all i, j , then $\mathcal{L}(R_1 R_2 R_3, S_1 S_2 S_3) \in \mathcal{M}$. Also if $x, y \in X$, then $\eta(x) \Delta \eta(y) = \mathcal{H}(x, y) \cup \mathcal{H}(y, x)$, so $\mathcal{L}(\eta(x), \eta(y)) = \mathcal{W}(x, y) \in \mathcal{M}$ by hypothesis. By iterating the median operation, we see that $\mathcal{L}(R, S) \in \mathcal{M}$ for all $R, S \in \langle \eta(X) \rangle = \Lambda(X)$. We can therefore set $\rho(R, S) = \mu(\mathcal{L}(R, S))$. Note that $\rho(\eta(x), \eta(y)) = \mu(\mathcal{L}(\eta(x), \eta(y))) = \mu(\mathcal{W}(x, y))$ for all $x, y \in X$. Moreover, if $R, S \in \Lambda(X)$ and $T \in [R, S]_{\Lambda(X)}$, we have observed that $\mathcal{L}(R, S) = \mathcal{L}(R, T) \sqcup \mathcal{L}(S, T)$, so by the additive property of μ we have $\rho(R, S) = \rho(R, T) + \rho(S, T)$. Thus ρ is a median pseudometric on $\Lambda(X)$ (see Lemma 13.1.1 and the subsequent remark).

In summary, we have shown:

Proposition 19.2.9. *Let $(X, \mathcal{W}, \mathcal{M}, \mu)$ be a space with weakly measured walls. Then there is a canonical map $\eta : X \rightarrow \Lambda$ into a median pseudometric space, Λ , such that $\mu(\mathcal{W}(x, y)) = \rho(\eta(x), \eta(y))$ for all $x, y \in X$.*

A particular case of interest is when $\mathcal{W}(x, y)$ is finite for all $x, y \in X$, and where \mathcal{M} consists of all finite subsets of \mathcal{W} . (This is called a “space with walls”.) Given $A \in \mathcal{M}$, we can set $\mu(A) = \#A$. In this case, iterating the median similarly as above, we see that $R \triangle S$ is finite for all $R, S \in \Lambda(X)$. If $T \in [R, S]_{\Lambda(X)}$ we have $R \triangle S = (R \triangle T) \sqcup (S \triangle T)$. There are only finitely many possibilities for such T , and so $[R, S]_{\Lambda(X)}$ is finite. In other words, $\Lambda(X)$ is a discrete median algebra. Moreover, ρ is the discrete combinatorial metric on $\Lambda(X)$.

As a slight generalisation of the above, suppose that $w : \mathcal{W} \rightarrow [0, \infty)$ is some function. Then given $A \in \mathcal{M}$, we can set $\mu(A) = \sum_{W \in A} w(W)$. This gives rise to a median pseudometric on the same discrete median algebra $\Lambda(X)$. Indeed by Lemma 13.2.5 and the subsequent remark, every median (pseudo)metric on $\Lambda(X)$ arises in this way.

The two constructions we have described are inverses in the following sense.

Suppose that M is a median algebra. In Subsection 9.2, we defined a monomorphism $\eta : M \rightarrow \mathcal{F}(\mathcal{H}(M))$. This agrees with the map defined above on setting $X = M$ and $\mathcal{W} = \mathcal{W}(M)$. In this case, $\eta(M)$ is already a subalgebra of $\mathcal{F}(\mathcal{H}(M))$, so by construction, $\Lambda(M) = \eta(M)$. We can therefore identify M with $\Lambda(M)$ via η . If (M, ρ) is a median metric space, then for all $x, y \in M$ we have $\rho(\eta(x), \eta(y)) = \mu(\mathcal{W}(x, y)) = \rho(x, y)$, so this is also an isometry.

Conversely, suppose that X is any set, and $\mathcal{W} \subseteq \mathcal{L}(X)$. Let $\mathcal{H} \subseteq \mathcal{P}_0(X)$ be the corresponding set of halfspaces. In Subsection 9.2, we defined a proset monomorphism $\zeta : \mathcal{H} \rightarrow \mathcal{H}(\mathcal{F}(\mathcal{H}))$ by setting $\zeta(H) = \{R \in \mathcal{F}(\mathcal{H}) \mid H \in R\}$. Since this commutes with the involutions of taking complements in X and in $\mathcal{H}(\mathcal{F}(\mathcal{H}))$ respectively, it descends to an injective map $\mathcal{W} \rightarrow \mathcal{W}(\mathcal{F}(\mathcal{H}))$.

In fact, some of the argument of Subsection 9.2 can be bypassed here. To see that $\zeta(H) \neq \emptyset$ for any $H \in \mathcal{H}$, choose any $x \in H$ and note that $\eta(x) \in \zeta(H)$. To see that ζ is injective, suppose that $H, H' \in \mathcal{H}$ are distinct. Up to swapping H and H' we can suppose that there is some $x \in H \setminus H'$. Then $\eta(x) \in \zeta(H) \setminus \zeta(H')$, and so $\zeta(H) \neq \zeta(H')$. (In other words we can bypass Lemmas 9.2.4 to 9.2.6.)

Now $\Lambda := \Lambda(X) \leq \mathcal{F}(\mathcal{H})$ is a subalgebra, and we similarly have a map $\theta : \mathcal{H} \rightarrow \mathcal{H}(\Lambda)$ given by $\theta(H) = \zeta(H) \cap \Lambda = \{R \in \Lambda \mid H \in R\}$. The argument of the previous paragraph shows that indeed $\theta(H) \neq \emptyset$, and that θ is injective. This time we get an injective map $\phi : \mathcal{W} \rightarrow \mathcal{W}(\Lambda)$. (This is the map $\mathcal{W} \rightarrow \mathcal{W}(\mathcal{F}(\mathcal{H}))$ previously defined postcomposed with the natural map $\mathcal{W}^\infty(\mathcal{F}(\mathcal{H})) \rightarrow \mathcal{W}^\infty(\Lambda)$.)

Let $H \in \mathcal{H}$, and $R, S \in \mathcal{F}(\mathcal{H})$. Note that $\theta(H) \in \mathcal{H}_\Lambda(R, S) \Leftrightarrow R \in \theta(H) \ \& \ S \notin \theta(H) \Leftrightarrow H \in S \setminus R$. Thus, $\theta(H) \in \mathcal{H}_\Lambda(R, S) \cup \mathcal{H}_\Lambda(S, R) \Leftrightarrow H \in R \triangle S$. We see that if $W \in \mathcal{W}(\Lambda)$, then $\theta(W) \in \mathcal{W}_\Lambda(R, S) \Leftrightarrow W \in \mathcal{L}(R, S)$. In other words, $\mathcal{L}(R, S) = \phi^{-1}(\mathcal{W}_\Lambda(R, S))$.

Now suppose that we have a measure μ on \mathcal{W} , defined on a σ -algebra in $\mathcal{P}(\mathcal{W})$ containing $\mathcal{W}(x, y)$ for all $x, y \in X$. This gives us a median metric, ρ , on Λ , which in turn gives us a measure, μ_Λ , on $\mathcal{W}(\Lambda)$ defined on the σ -algebra, $\mathcal{M}(\Lambda)$, generated by all wall-intervals of Λ . If $R, S \in \Lambda$, then $\mu_\Lambda(\mathcal{W}_\Lambda(R, S)) = \rho(R, S) = \mu(\mathcal{L}(R, S)) = \mu(\phi^{-1}(\mathcal{W}_\Lambda(R, S)))$. It now follows that $\mu_\Lambda(A) = \mu(\phi^{-1}A)$ for all $A \in \mathcal{M}(\Lambda)$.

We note that one can give a variation on the above construction, which gives rise to a complete median pseudometric. (In particular, its hausdorffification is a complete median metric.)

Suppose that $(X, \mathcal{W}, \mathcal{M}, \mu)$ is a space with measured walls. Let $\hat{\Lambda}(X)$ be the set of flows, R , on \mathcal{H} such that $\mathcal{L}(R, \eta(x)) \in \mathcal{M}$ and $\mu(\mathcal{L}(R, \eta(x))) < \infty$ for some, hence any, $x \in X$. Thus, $\hat{\Lambda}(X)$ is a subalgebra of $\mathcal{F}(\mathcal{H})$ containing $\Lambda(X)$ (cf. the discussion of ‘‘almost principal flows’’ in Subsection 9.4). Note that if $R, S \in \hat{\Lambda}(X)$, then $\mathcal{L}(R, S) \in \mathcal{M}$ (this time, because \mathcal{M} is, by hypothesis, a subring of $\mathcal{P}(\mathcal{W})$). We again set $\rho(R, S) = \mu(\mathcal{L}(R, S))$, so that ρ is a pseudometric on $\hat{\Lambda}(X)$ for the same reason as before. It agrees with that previously defined on $\Lambda(X)$.

We claim that if μ is a complete measure on \mathcal{W} , then ρ is a complete pseudometric. To see this, let $(R_n)_{n \in \mathbb{N}}$ be a cauchy sequence in $\hat{\Lambda}(X)$. We can suppose that $\rho(R_n, R_{n+1}) \leq 1/2^n$ for all n . Let $D_n = \bigcup_{m=n}^{\infty} \mathcal{L}(R_m, R_{m+1})$ and $E = \bigcap_{n=0}^{\infty} D_n$. Thus, $D_n, E \in \mathcal{M}$. Now $\mu(D_n) \leq \sum_{m=n}^{\infty} \mu(\mathcal{L}(R_m, R_{m+1})) \leq 2/2^n$, and $\mu(E) = 0$. Let $A \in \mathcal{H} \setminus \pi^{-1}E$. Then for all sufficiently large n , $A \notin R_n \triangle R_{n+1}$. In other words, A eventually lies in R_n or in $\mathcal{H} \setminus R_n$. Let S be the set $A \in \mathcal{H} \setminus \pi^{-1}E$ which lie in R_n for all sufficiently large n . This is a flow on $\mathcal{H} \setminus R_n$ (since the restriction of R_n to this subproset is a flow). By Lemma 9.2.4, S extends to a flow, R , on \mathcal{H} . Now $\mathcal{L}(R_n, R) \subseteq D_n$, and so $\mathcal{L}(R_0, R) \setminus D_n = \mathcal{L}(R_0, R_n) \setminus D_n$. Thus $\mathcal{L}(R_0, R) \setminus D_n \in \mathcal{M}$, and it follows that $\mathcal{L}(R_0, R) \setminus E \in \mathcal{M}$. Also, since μ is complete, $E \cap \mathcal{L}(R, R_0) \in \mathcal{M}$ and $\mu(E \cap \mathcal{L}(R, R_0)) = 0$. In particular (since $E \in \mathcal{M}$) we have $\mathcal{L}(R_0, R) \in \mathcal{M}$, so $R \in \hat{\Lambda}(X)$. Moreover, $\rho(R_n, R) = \mu(\mathcal{L}(R_n, R)) \leq \mu(D_n) \rightarrow 0$, as required.

In summary, we get:

Proposition 19.2.10. *Let $(X, \mathcal{W}, \mathcal{M}, \mu)$ be a space with measured walls. Then $\hat{\Lambda}(M)$ is a median pseudometric space containing $\Lambda(M)$. If the measure, μ is complete, then $\hat{\Lambda}(M)$ is a complete pseudometric space.*

The above can be illustrated using \mathbb{R} with the standard metric. This is a median metric space. Here, halfspaces are of the form (∞, x) , $(-\infty, x]$, (x, ∞) or $[x, \infty)$ for $x \in \mathbb{R}$. Given $x \in \mathbb{R}$, write $W_-(x) = \{(-\infty, x), [x, \infty)\}$ and $W_+(x) = \{(-\infty, x], (x, \infty)\}$. Thus $\mathcal{W}(\mathbb{R}) = \{W_-(x), W_+(x) \mid x \in \mathbb{R}\}$. If $x < y$, then $\mathcal{W}(x, y) = \{W_+(x), W_-(y)\} \cup \{W_-(z), W_+(z) \mid x < z < y\}$. We can think of $\mathcal{W}(\mathbb{R})$ as two disjoint copies of \mathbb{R} , and the measure induced is half the usual Lebesgue measure (restricted to the σ -ring of borel sets).

We can describe flows on $\mathcal{W}(\mathbb{R})$ as follows. Let $S(x) = \{(-\infty, y), (-\infty, y] \mid y > x\} \cup \{(y, \infty), [y, \infty) \mid y < x\}$. Let $R_0(x) = S(x) \cup \{(-\infty, x], [x, \infty)\}$, $R_-(x) = S(x) \cup \{[x, \infty), (x, \infty)\}$, and $R_+(x) = S(x) \cup \{(-\infty, x], (-\infty, x)\}$. Then $\mathcal{F}(\mathcal{H}(\mathbb{R})) = \{R_0(x), R_-(x), R_+(x) \mid x \in X\}$. As a median algebra, this is isomorphic to the direct product $\mathbb{R} \times \{-, 0, +\}$ via the map $[(x, \epsilon) \mapsto R_\epsilon(x)]$ for $\epsilon \in \{-, 0, +\}$ (where (x, ϵ) denotes ordered pair). The map $\mathbb{R} \rightarrow \mathcal{F}(\mathcal{H}(\mathbb{R}))$ is then given by $[x \mapsto (x, 0)]$ and $\Lambda(\mathbb{R}) \equiv \mathbb{R} \times \{0\}$. One can also check that $\hat{\Lambda}(\mathbb{R}) = \mathcal{F}(\mathbb{R})$, with $\rho((x, \pm), (x, 0)) = \rho((x, -), (x, +)) = 0$ for all $x \in \mathbb{R}$. Its hausdorffification is naturally identified with \mathbb{R} .

We can do a similar construction with the rational numbers \mathbb{Q} . In this case, $\mathcal{W}(\mathbb{Q})$ consists of two copies of \mathbb{Q} and one copy of $\mathbb{R} \setminus \mathbb{Q}$. The measure is now Lebesgue measure. We can identify $\mathcal{F}(\mathcal{H}(\mathbb{Q}))$ with the median algebra $(\mathbb{R} \times \{0\}) \cup (\mathbb{Q} \times \{-, +\})$, and $\Lambda(\mathbb{Q})$ with $\mathbb{Q} \times \{0\}$.

19.3. A construction of a median metric.

One interesting application of these ideas is described in [Ge2]. (Our account below is slightly different.)

Let M be a median algebra. Given any $C \subseteq M$, write $\mathcal{W}(C) = \{W \in M \mid W \pitchfork C\} = \bigcup_{a,b \in C} \mathcal{W}(a, b)$. Note that $\mathcal{W}(\text{hull}(C)) = \mathcal{W}(C)$.

Let $(M, \mathcal{W}, \mathcal{M}, \mu)$ be a space with measured walls, \mathcal{W} . We say that $A \subseteq M$ is **wall-finite** if $\mathcal{W}(A) \in \mathcal{M}$ and $\mu(\mathcal{W}(A)) < \infty$. We write \mathcal{S} for the set of all non-empty wall-finite subsets. From the above observation, this is closed under taking convex hulls. It is also closed under finite union. For suppose $A, B \in \mathcal{S}$ are non-empty. Choose any $a \in A$ and $b \in B$. Then $\mathcal{W}(A \cup B) = \mathcal{W}(A) \cup \mathcal{W}(B) \cup \mathcal{W}(a, b)$, and by definition, $\mathcal{W}(a, b) \in \mathcal{M}$ and $\mu(\mathcal{W}(a, b)) < \infty$.

Given $A \in \mathcal{S}$, we abbreviate $\mu(A) = \mu(\mathcal{W}(A))$.

Suppose $A, B \in \mathcal{S}$. We write

$$\sigma(A, B) = 2\mu(A \cup B) - \mu(A) - \mu(B).$$

(It is not hard to check that $\sigma(A, B) = \mu(\mathcal{W}(A) \triangle \mathcal{W}(B)) + 2\mu(\mathcal{W}(A|B))$, where $\mathcal{W}(A|B)$ is the set of walls separating A and B , though we won't be using this fact directly.)

Suppose $A, B, C \subseteq M$ are non-empty. Then $\mathcal{W}(C) \subseteq \mathcal{W}(A \cup C) \cap \mathcal{W}(B \cup C)$, and $\mathcal{W}(A \cup B) \subseteq \mathcal{W}(A \cup C) \cup \mathcal{W}(B \cup C)$. Thus, if $A, B, C, A \cup B, A \cup C, B \cup C \in \mathcal{S}$, we have $\mu(A \cup B) + \mu(C) \leq \mu(A \cup C) + \mu(B \cup C)$, and so $\sigma(A, B) \leq \sigma(A, C) + \sigma(B, C)$.

Now let $\mathcal{K} = \mathcal{K}(M)$ be the median algebra of all non-empty convex subsets of M , where $ABC = \{abc \mid a \in A, b \in B, c \in C\}$. (This was example (Ex3.4) of Subsection 3.4.) Let \mathcal{G} be a median subalgebra of $\mathcal{S} \cap \mathcal{K}$, such that $\text{hull}(A \cup B) \in \mathcal{G}$ for all $A, B \in \mathcal{G}$.

Suppose that $A, B, C \in \mathcal{G}$ and that $A.C.B$ holds. Then $C \subseteq \text{hull}(A \cup B)$ and so $\mathcal{W}(A \cup C) \cup \mathcal{W}(B \cup C) = \mathcal{W}(A \cup B)$. Moreover, $\mathcal{W}(A \cup C) \cap \mathcal{W}(B \cup C) = \mathcal{W}(C)$. (For suppose $W \in \mathcal{W}(A \cup C) \cap \mathcal{W}(B \cup C) \setminus \mathcal{W}(C)$. We can suppose that $C \subseteq W^-$. Now choose any $a \in A \cap W^+$, $b \in B \cap W^+$ and $c \in C$. Then $abc \in C \cap W^+$ giving

a contradiction.) It now follows that $\mu(A \cup B) + \mu(C) = \mu(A \cup C) + \mu(B \cup C)$, and so $\sigma(A, B) = \sigma(A, C) + \sigma(B, C)$.

By Lemma 13.1.1, we have shown:

Proposition 19.3.1. *σ is a median pseudometric on the median algebra, \mathcal{G} .*

In particular, if $\mu(\mathcal{W}(a, b)) > 0$ for all distinct $a, b \in M$, then (\mathcal{G}, σ) is a median metric space, inducing the given subalgebra median on \mathcal{G} .

If (M, ρ) is a median metric space, and μ is the measure on $\mathcal{W}(M)$ as given by Theorem 19.2.8, then $\sigma(\{a\}, \{b\}) = 2\rho(a, b)$ for all $a, b \in M$. Note that the map $[a \mapsto \{a\}] : M \rightarrow \mathcal{K}(M)$ is a median monomorphism. Therefore, if \mathcal{G} includes all singletons, this gives an isometric embedding of (M, ρ) into $(\mathcal{G}, \sigma/2)$.

It is unclear in general which subsets of M are wall-finite. However, we could certainly take \mathcal{G} to be the set of convex hulls of finite sets, which we have noted is subalgebra. It is this space which is studied in [Ge2].

19.4. Examples.

Here are some examples of spaces with measured walls.

(Ex19.1): Let \mathbb{E}^n be euclidean n -space with the standard geodesic metric, σ . Let $\mathcal{W} = \mathcal{W}(\mathbb{E}^n)$ be the set of partitions of \mathbb{E}^n of the form $\{H, \mathbb{E}^n \setminus H\}$, where H is a closed euclidean halfspace. Let Γ_n be the isometry group of \mathbb{E}^n . Then Γ_n acts transitively on \mathcal{W} . Note that (on identifying \mathbb{E}^n as $\mathbb{E}^{n-1} \times \mathbb{R}$) we have $(\mathbb{E}^{n-1} \times [0, \infty), \mathbb{E}^{n-1} \times (-\infty, 0)) \in \mathcal{W}$. Its stabiliser can be identified with Γ_{n-1} , and so we can identify \mathcal{W} with the coset space Γ_n/Γ_{n-1} . In this way, \mathcal{W} , acquires a natural Γ_n -invariant measure, μ , coming from the Haar measure on the Lie group, Γ_n . Since Γ_n acts transitively on pairs, we see that there is a continuous function $\lambda : [0, \infty) \rightarrow [0, \infty)$, such that $\mu(\mathcal{W}(x, y)) = \lambda(\sigma(x, y))$ for all $x, y \in \mathbb{E}^n$. In fact, $\lambda(t) = kt$ for some constant $k > 0$. (To see this, let z lie on the euclidean geodesic from x to y . Then $\mathcal{W}(x, y) = \mathcal{W}(x, z) \sqcup \mathcal{W}(z, y)$ and $\sigma(x, y) = \sigma(x, z) + \sigma(z, y)$. We see that $\lambda(t + u) = \lambda(t) + \lambda(u)$ for all $t, u \geq 0$, and the statement follows easily.) We can therefore normalise so that $\mu(\mathcal{W}(x, y)) = \sigma(x, y)$ for all x, y . Applying the above construction, we can embed \mathbb{E}^n into a median metric space, (Λ, ρ) , such that $\rho(x, y) = \sigma(x, y)$ for all $x, y \in \mathbb{E}^n$.

In fact, Λ is connected. To see this, recall that, by definition, each element, $a \in \Lambda$, has the form $E(\eta(x_1), \dots, \eta(x_m))$, where $x_1, \dots, x_m \in \mathbb{E}^n$, and E is a median expression. Let $p \in \mathbb{E}^n$ be any basepoint. Let $\alpha_i : [0, 1] \rightarrow \mathbb{E}^n$ be a continuous path from p to x_i in \mathbb{E}^n . Then $[t \mapsto E(\eta(\alpha_1(t)), \dots, \eta(\alpha_m(t)))]$ is a continuous path from $\eta(p)$ to a in Λ . In fact, by taking each α_i to be geodesic, one can see that Λ is contractible. (Note the path is independent of the choice of median expression: it can be alternatively be described by applying dilations to \mathbb{E}^n about p .)

The construction of Λ is canonical, so the action of Γ_n gives rise to an isometric action on Λ . This shows that we have an equivariant isometric embedding of

euclidean space into a median metric space. In particular, the euclidean space \mathbb{E}^n is (canonically) submedian.

Note that in the case where $n = 1$, we can identify $\Lambda(\mathbb{E}^1)$ with $\Lambda(\mathbb{R}) \cong \mathbb{R}$, as described in Subsection 9.1.

(Ex19.2): Here is another embedding of \mathbb{E}^n into a median metric space, which gives rise to a more explicit description of $\Lambda(\mathbb{E}^n)$.

Let $\mathbb{P} = \mathbb{P}(\mathbb{E}^n)$ be real projective space: the ideal sphere of directions in \mathbb{E}^n factored out by the antipodal map. By a “hyperplane” in \mathbb{E}^n we mean a codimension-1 totally geodesic subspace. Let $G = G(\mathbb{E}^n)$ be the set of such hyperplanes. We give G a structure as a line bundle, $\omega : G \rightarrow \mathbb{P}$, over \mathbb{P} , where the fibre $G_\xi := \omega^{-1}\xi$ is the set of hyperplanes orthogonal to ξ . We can identify the set, Σ , of (set theoretical) sections of G with the direct product $\prod_{\xi \in \mathbb{P}} G_\xi$. Now we can view each G_ξ as a median algebra isomorphic to \mathbb{R} . This gives Σ the structure of a median algebra as a direct product. (In other words, we take the median pointwise on each fibre.) Note that the set, $\Sigma_0 = \Sigma_0(\mathbb{E}^n)$, of continuous sections of G , is a subalgebra of Σ . We have a compatible median metric, ρ , on Σ_0 defined by the integral $\rho(\sigma, \tau) = \int_{\mathbb{P}} |\sigma(\xi) - \tau(\xi)| d\xi$ with respect to Lebesgue measure on \mathbb{P} . (Here $|\cdot - \cdot|$ is just the usual distance on $G_\xi \cong \mathbb{E}^1$.) Given any $x \in \mathbb{E}^n$, we have a continuous section, $\sigma(x) \in \Sigma_0$, where $\sigma(x)$ is the hyperplane in G_ξ containing x . The map $\sigma : \mathbb{E}^n \rightarrow \Sigma_0$ is an isometric embedding, modulo some fixed multiplicative factor depending on n .

We can relate this to Example (Ex19.1) as follows.

Let $\mathcal{H} = \mathcal{H}(\mathbb{E}^n)$ be the set of halfspaces of \mathbb{E}^n with its structure as a space with measured walls. Thus, an element of \mathcal{H} is either a closed or an open euclidean halfspace in the usual sense. Its boundary is a hyperplane. This gives a natural map, $\mathcal{H} \rightarrow G$. Postcomposing with ω , we get a map, $\beta : \mathcal{H} \rightarrow \mathbb{P}$. Write $\mathcal{H}_\xi = \beta^{-1}(\xi) \subseteq \mathcal{H}$. This is a subproset of \mathcal{H} , isomorphic to $\mathcal{H}(\mathbb{E}^1)$. As such we can identify \mathcal{H}_ξ with $\mathcal{H}(G_\xi)$. As described in Subsection 9.1, we can also identify G_ξ with $\Lambda(G_\xi) \subseteq \mathcal{F}(\mathcal{H}_\xi)$.

As a proset, \mathcal{H} is a disjoint union $\bigsqcup_{\xi \in \mathbb{P}} \mathcal{H}_\xi$. As in Example (Ex9.2) of Subsection 9.1, we see that $\mathcal{F}(\mathcal{H})$ is median isomorphic to $\prod_{\xi \in \mathbb{P}} \mathcal{F}(\mathcal{H}_\xi)$: an uncountable direct product of copies of $\mathcal{F}(\mathcal{H}(\mathbb{E}^1))$. Under the identification of each G_ξ as a subalgebra of $\mathcal{F}(\mathcal{H}_\xi)$, we get an identification of G as a subalgebra of $\mathcal{F}(\mathcal{H})$. Recall that, given $x \in \mathbb{E}^n$, we have a section, $\sigma(x) : \mathbb{P} \rightarrow G \leq \mathcal{F}(\mathcal{H})$. Its image in G is the set of hyperplanes containing x . Identified as a subset of $\mathcal{F}(\mathcal{H})$ is precisely the flow $\eta(x) \in \Lambda(\mathbb{E}^n) \subseteq \mathcal{F}(\mathcal{H})$.

More generally, suppose $a \in \Lambda(\mathbb{E}^n) \subseteq \mathcal{F}(\mathcal{H})$. We claim that a is also the image of a continuous section of $G \leq \mathcal{F}(\mathcal{H})$, which is piecewise of the above form. More precisely, by definition, we can write $a = E(\eta(x_1), \dots, \eta(x_m))$ where $x_1, \dots, x_m \in \mathbb{E}^n$, and E is a median expression. Given $\xi \in \mathbb{P}$, we can evaluate $E(\sigma(x_1)\xi, \dots, \sigma(x_m)\xi)$ in G_ξ to give us some $\tau(\xi) \in G_\xi$. Since every subset of G_ξ is a subalgebra, we have

$\tau(\xi) = \sigma(x_{i(\xi)})$ for some $i(\xi) \in \{1, \dots, m\}$. The map $\tau : \mathbb{P} \rightarrow G$ is continuous section of G , and its image in $G \leq \mathcal{F}(\mathcal{H})$ is precisely the flow, a . In fact, let $U \subseteq \mathbb{P}$ be the open dense subset of \mathbb{P} consisting of those $\xi \in \mathbb{P}$ for which the map $[i \mapsto \sigma(x_i)\xi]$ is injective. Then $[\xi \mapsto i(\xi)]$ is well defined and constant on each component of U . Note that, while there is a choice of the median expression E , the map $\tau_a := \tau$ is well defined.

As an example, suppose p is odd, and E happens to be majority vote. (We can suppose the x_i are all distinct.) Then a corresponds to those hyperplanes which intersect $\{x_1, \dots, x_m\}$, and which have more than half of these points contained in each of the closed halfspaces of \mathbb{E}^n which the hyperplane bounds. This defines a continuous section of G . (See the discussion of majority vote in Subsection 5.3.)

(Ex19.3): The above description was a bit clumsy, in that each hyperplane of \mathbb{E}^n corresponds to two walls of $\mathcal{W}(\mathbb{E}^n)$. This was necessary to conform with the general definition of a space with measured walls. In practice, we don't really care very much which of the two halfspaces of such a wall is closed in \mathbb{E}^n , and which is open.

Instead, let us take $\mathcal{T} = \mathcal{T}(\mathbb{E}^n)$ be the the set of closed halfspaces of \mathbb{E}^n . This has the structure of a proset, where \leq is defined by inclusion as before, and where H^* is now the closure of $\mathbb{E}^n \setminus H$. We can identify $\mathcal{T}/*$ with G , and equip it with the usual measure. Informally, we can think of an element of \mathcal{T} as a transversely oriented hyperplane (oriented towards the halfspace) cf. Example (Ex9.4) of Subsection 9.1.

We can now proceed similarly as before. Given $x \in \mathbb{E}^n$, we define a flow, $\eta(x) \in \mathcal{F}(\mathcal{T})$, by orienting each hyperplane towards x . Note that the set of hyperplanes containing x has measure 0, so we can orient these arbitrarily. We can now identify $\Lambda(\mathbb{E}^n)$ with $\langle \eta(X) \rangle \leq \mathcal{F}(\mathcal{T})$. In these terms $\hat{\Lambda}(\mathbb{E}^n)$ corresponds to measurable sections of G .

(Ex19.4): The construction of Example (Ex19.1) applies equally well to real hyperbolic n space, \mathbb{H}^n , with the usual geodesic metric. As before, a ‘‘halfspace’’ of \mathbb{H}^n is a non-empty geodesically convex subset whose complement is non-empty and convex. A ‘‘wall’’ is then an unordered pair, $\{H, \mathbb{H}^n \setminus H\}$, where H is a closed halfspace. The space, $\mathcal{W}(\mathbb{H}^n)$, of such walls similarly admits a canonical complete measure. We can now construct spaces, $\Lambda(\mathbb{H}^n) \leq \hat{\Lambda}(\mathbb{H}^n)$, with \mathbb{H}^n canonically isometrically embedded. As described by Example (Ex19.3), it is perhaps easier to think in terms of the proset, \mathcal{T} , of transversely oriented hyperplanes (similarly defined).

There are few qualitative differences. In this case, the set of hyperplanes separating a point of \mathbb{H}^n from a given hyperplane, has positive measure, and this measure tends to infinity as the distance between them goes to infinity. Suppose that $R \in \hat{\Lambda}(\mathbb{H}^n)$, thought of as a flow on \mathcal{T} . For any basepoint, $x \in \mathbb{H}^n$, the set of hyperplanes of R which are oriented away from x all lie at most a bounded bounded distance, say r , from x in \mathbb{H}^n . Let R_t be obtained by reversing the orientation on such hyperplanes of R which are a distance greater than t from x . Then

$[t \mapsto R_t] : [0, r] \longrightarrow \hat{\Lambda}(\mathbb{H}^n)$ is a path from $\eta(x)$ to R in $\hat{\Lambda}(\mathbb{H}^n)$. This shows that the hausdorffification of $\hat{\Lambda}(\mathbb{H}^n)$ is connected. In fact, combining such paths one can see that it is contractible.

By related arguments it is shown in [ChatteD] that $\hat{\Lambda}(\mathbb{H}^n)$ is locally compact (assuming we take the complete measure on the set of walls).

(Ex19.5): We can put a different metric on \mathbb{H}^n , namely the square root of the usual geodesic metric. It is shown in [FaH] that this space embeds in a Hilbert space. It follows that it also embeds in an L^1 space. It is therefore submedian and admits a structure as a space of measured walls (see Remark 3.23 of [ChatteDH]). The same applies to complex hyperbolic n -space with the square root of the geodesic metric (again by [FaH]). These constructions are quite different from our earlier examples, and do not seem to have such a transparent geometric interpretation.

One cannot apply the same construction to quaternionic hyperbolic space (in quaternionic dimension at least 2) nor to the Cayley hyperbolic plane, since their isometry groups do not have the Haagerup property: see the Notes to this Section. At least, such a construction cannot be equivariant, for this reason.

Various other examples are given in [CherMV] and [ChatteDH].

The median algebras that arise in this context typically have infinite rank, and so the subject has a somewhat different flavour to much of what we have discussed in earlier sections. In contrast, the examples that arise as asymptotic cones, as described in Section 24, are typically of finite rank.

20. BOOLEAN FUNCTIONS AND MAJORITY VOTE

Median algebras arise in the theory of boolean functions, which explains some of their interest in computer science. They also give us another way of viewing Theorem 5.2.3: see Proposition 20.1.2. We describe a connection with logic circuits in Subsection 20.2. The particular case of “majority vote” is discussed in Subsection 20.3.

20.1. Definitions and connections with median algebras.

Let $[x \mapsto x^*]$ be the involution on $\{0, 1\}$, defined by $0^* = 1$ and $1^* = 0$, commonly referred to as “inversion”. Given $\underline{x} = (x_1, \dots, x_n) \in \{0, 1\}^n$, write $\underline{x}^* = (x_1^*, \dots, x_n^*)$. Given $\underline{y} \in \{0, 1\}^n$, write $\underline{x} \preceq \underline{y}$ to mean that $x_i \leq y_i$ for all i . Clearly this implies $\underline{y}^* \preceq \underline{x}^*$.

An ***n -ary boolean function*** is a map $f : \{0, 1\}^n \longrightarrow \{0, 1\}$. The ***dual***, \hat{f} , of f is defined by setting $\hat{f}(\underline{x}) = (f(\underline{x}^*))^*$. We say that f is ***self-dual*** if $\hat{f} = f$. We say that f is ***monotone*** if $f(\underline{x}) \leq f(\underline{y})$ whenever $\underline{x} \preceq \underline{y}$. Note that the dual function, \hat{f} , is also monotone.

There is a natural correspondence between boolean functions and subsets of $\mathcal{P}(I_n)$, where $I_n = \{1, \dots, n\}$. Namely, to a boolean function, f , we associate

the family \mathcal{F} , where $F \in \mathcal{F}$ if $f(\underline{x}(F)) = 1$, where $x_i(F) = 1$ if $x \in F$, and $x_i(F) = 0$ if $x \in F^* := I_n \setminus F$. Under this correspondence, f is self-dual if and only if \mathcal{F} is a $*$ -transversal: that is exactly one of F or F^* lies in \mathcal{F} for all $F \subseteq I_n$. Also, f is monotone, if and only if \mathcal{F} satisfies property (P2) of Subsection 5.2: namely if $A \in \mathcal{F}$ and $B \in \mathcal{P}(I_n)$ with $A \subseteq B$, then $B \in \mathcal{F}$. Therefore self-dual monotone boolean functions correspond to “flows” as defined in Subsection 5.2. In particular, the set of all self-dual monotone boolean functions corresponds to the superextension of I_n .

When $n = 2$, there are six monotone functions, namely, the constant functions: $f \equiv 0$, $f \equiv 1$, the projection maps: $f(a, b) = a$, $f(a, b) = b$, and the minimum and maximum: $f(a, b) = a \wedge b$, $f(a, b) = a \vee b$. Of these, only the projection maps are self-dual.

In Subsection 5.4 we gave a complete list of all self-dual monotone n -ary boolean functions for $n \leq 5$ (in their guise as flows). In particular, when $n = 3$, these are just the projection functions ($f(a, b, c) = a$, $f(a, b, c) = b$ and $f(a, b, c) = c$) together with the median function: $f(a, b, c) = abc$.

Here is a key observation:

Lemma 20.1.1. *Let $f : \{0, 1\}^3 \rightarrow \{0, 1\}$ be monotone. Then for all $a, b, c \in \{0, 1\}$, we have $f(a, b, c) = f(a, a, c)f(a, b, b)f(c, b, c)$.*

Here, the right-hand side represents the usual median operation in $\{0, 1\}$.

Proof. After permuting a, b, c and replacing f by its dual, we can assume that $a = b = 0$ and $c = 1$. Now $(0, 0, 0) \preceq (0, 0, 1) \preceq (1, 0, 1)$, so $f(0, 0, 0) \leq f(0, 0, 1) \leq f(1, 0, 1)$ and so $f(0, 0, 1) = f(0, 0, 1)f(0, 0, 0)f(1, 0, 1)$ as required. \square

Note that if we hold some of the arguments of a monotone boolean function constant, then the function is monotone in the remaining arguments. We can therefore apply Lemma 20.1.1 to any three arguments in any n -ary monotone function for $n \geq 3$.

We can also repeat arguments in a monotone function. More formally, suppose $\sigma : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined by setting the i th coordinate of $\sigma(\underline{x})$ to be equal to $x_{j(i)}$ for some fixed $j(i)$. Then σ is monotone and $\sigma(\underline{x}^*) = (\sigma(\underline{x}))^*$. Therefore, if $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is monotone, then so is $f \circ \sigma : \{0, 1\}^m \rightarrow \{0, 1\}$. If f is self-dual, so is $f \circ \sigma$.

Now suppose $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is monotone, and $n \geq 3$. Then applying Lemma 20.1.1 to suitable sets of three variables, and repeating, we can eventually express f as a median expression, where each entry of the expression has the form $f \circ \sigma$ for some σ with $m = 3$ as above. (Here we are using the term “median expression” in the sense of Subsections 3.3 and 6.1, namely an iterated composition of median functions.) In other words, in each entry we have substituted the original arguments, x_1, \dots, x_n , of f , so that it features at most three of the x_i , with repetitions. Note that each of these entries is monotone. This shows that f can be expressed a median expression in ternary monotone boolean functions.

As observed in the previous paragraph, if f is also self-dual, then these ternary functions are themselves self-dual. Therefore each one of these ternary operations is either a projection function or the median map. In summary, we have shown:

Proposition 20.1.2. *Any self-dual monotone boolean function can be written as a median expression in its arguments.*

We note that Proposition 20.1.2 is equivalent to Theorem 5.2.3, namely that the superextension of a finite set is generated as a median algebra by the principal ultrafilters. This boils down to the same thing, given the correspondence described above.

An example to illustrate this is the “majority vote” on five variables, a, b, c, d, e , which we denote by $\langle abcde \rangle$. This is a self-dual monotone quinary function, which we met in Subsection 5.4.

We have

$$\begin{aligned} \langle abcde \rangle &= \langle aacde \rangle \langle abbde \rangle \langle cbcde \rangle \\ &= (\langle aacce \rangle \langle aacdd \rangle \langle aaede \rangle) \\ &\quad (\langle abbae \rangle \langle abbdd \rangle \langle ebbde \rangle) \\ &\quad (\langle cbcbe \rangle \langle cbcdd \rangle \langle cecde \rangle) \\ &= ((ace)(acd)(ade))((abe)(abd)(bde))((bce)(bcd)(cde)). \end{aligned}$$

Of course, one can do better. A more efficient procedure for majority functions is given in [TSAHD], where the authors derive the formula $\langle abcde \rangle = a(bcd)((bce)de)$, which we mentioned in Subsections 5.4 and 6.2. For the majority on seven variables, a, b, c, d, e, f, g , they give the formula: $((abc)((abc)de)f)((def)((def)ab)c)g$. We will give a general formula in Subsection 20.3.

We remark that one can also express these constructions in terms of the structure of $\{0, 1\}$ as a boolean ring, mentioned in Subsection 3.4. In these terms, the median operation is given by $ab + bc + ca$. One can express any self-dual monotone function as an iteration of this. For example, one sees that $\langle abcde \rangle = abc + abd + abe + acd + ace + ade + bcd + bce + bde + cde + abcd + abce + abde + acde + bcde$. (Here, of course, “ abc ” etc. denotes the product in the ring, not the median operation.) From this one can verify the formulae given earlier by direct expansion, though clearly this is a rather laborious way of going about it.

20.2. Clones.

One can express the above statements more formally in terms of “clones” of boolean functions.

Let Ω_n be the set of all n -ary boolean functions, and let $\Omega = \bigcup_{n=0}^{\infty} \Omega_n$ be the set of all boolean functions. (This includes the two constant 0-ary functions, 0 and 1.) If $n \geq i \geq 1$, we write $\pi_{i,n} \in \Omega_n$ for projection to the i th coordinate. We say that a subset $\Psi \subseteq \Omega$ is **closed** if given any $f \in \Psi \cap \Omega_m$ and any $g_1, \dots, g_m \in \Psi \cap \Omega_n$, the composition $f(g_1, \dots, g_m) \in \Omega_n$ also lies in Ψ .

Definition. A subset, $\Psi \subseteq \Omega$ is a (*boolean*) *clone* if it is closed and $\pi_{i,n} \in \Psi$ for all $n \geq i \geq 1$.

Clearly any intersection of clones is a clone. Given any subset $B \subseteq \Omega$, the clone *generated* by B is the intersection, Ψ , of all clones containing B . Less formally, $f \in \Psi$ if and only if it can be written as an expression involving only functions from the family B . (Here, in general, we need to take account of the order of the arguments in such an expression.) The set B is referred to a *basis* for the clone.

Around 1920, Post gave a classification of all (boolean) clones. He showed that every clone is finitely generated. They form a countably infinite family, which can be arranged as a partially ordered set ordered by inclusion. This is commonly known as “Post’s lattice”. It is described in [P]. See [Zv] for a more recent account, and shorter proof.

In these terms, Proposition 20.1.2 can be expressed by saying that the ternary median operation is a basis for the clone consisting of all self-dual monotone functions.

We mention a few other examples. The clone, Ω , of all boolean function is generated by the binary minimum function, $[(a, b) \mapsto a \wedge b]$, together with inversion ($[a \mapsto a^*]$). The clone of all 0,1-preserving functions (i.e. sending $(0, \dots, 0)$ to 0, and $(1, \dots, 1)$ to 1) is generated by the binary minimum and maximum functions. The clone of all self-dual functions is generated by the ternary median operation together with inversion.

One can equivalently express all this in terms of logical circuits. Such a circuit consists of a directed graph with no directed cycles. It has a set of input nodes (source vertices) and a terminal output node (sink vertex). All other vertices are “logic gates”. To each logic gate is associated a fixed boolean function: the “connective function” of that gate. It has an incoming edge for each of its arguments, and a number of outgoing edges, each of which outputs the (same) result of computing this function. By composing these connective functions, the circuit outputs a value in $\{0, 1\}$, for any assignment of its arguments. Thus the set of all possible functions one can compute on some machine with a given family of logic gate types is precisely the clone generated by these types.

In this context, 0, is commonly interpreted as “false” and 1 as “true”. The inversion map is “not”, minimum is “and”, and maximum is “or”. A gate that computes the median is sometimes called a “(ternary) majority” gate. Thus one can compute all boolean functions just with “and” and “not” gates, and one can compute all self-dual monotone boolean functions just with majority gates.

Traditional semiconductor (“CMOS”) chips typically use unary and binary logic gates, though there is now much interest in investigating the potential for ternary gates, for example for reversible or quantum computing.

One example we have already alluded to is the case where we allow two gates: the ternary majority gate and the (unary) inverter gate. These generate all self-dual boolean functions. These satisfy the “laws” of a ternary boolean algebra as

defined by Grau [Gra], and discussed in Subsection 3.4. Logical circuits using these two kinds of gates have been called “majority inverter graphs”. In [ChattoASGD] one can find an account of them, expressed using a system of “laws” of the type described above.

20.3. Majority consensus algebras.

We have referred to “majority vote” at a number of places (see Subsections 5.3, 5.4, 6.2 and 20.1), and it will appear again in Subsection 21.1. Such a structure arises as a symmetric n -ary operation on a median algebra, M , for any odd $n \geq 3$. (For $n = 3$ it is simply the original median.) It is natural to ask which n -ary operations arise in this way. Such a structure has been called an “algebra of majority consensus”, which we will abbreviate here to “consensus algebra”. In this subsection, we will explore this subject more systematically.

Let us begin with a special case. Let $I = \{0, 1\}$ be a 2-point set. Any symmetric monotone boolean function, $\mu : I^n \rightarrow I$, is a “threshold function”; that is to say there is some $r \in \{0, 1, 2, \dots, n\}$ such that $\mu(0^p 1^q) = 0$ if and only if $p \geq r$. (In writing arguments of operations, we will frequently abbreviate m consecutive occurrences of x as x^m .) If μ is also self-dual, then n is odd, and $r = \lceil n/2 \rceil$ (so $n = 2r - 1$). We refer to this as the “(standard) majority vote”, and denote it by $\langle x_1 x_2 \dots x_n \rangle$ for $x_1, x_2, \dots, x_n \in I$.

In Subsection 20.1, we saw that this can be written in terms of a median expression in the standard median on I (that is, the ternary majority function). Any two such expressions, E, E' , will necessarily be tautologically equivalent, in the sense defined in Section 6. Let us arbitrarily choose some such expression, E_n . (Some explicit formulae for such expressions are described at the end of this subsection.)

Given any set, X , we can identify its power set, $\mathcal{P}(X)$, with the “hypercube” I^X , via characteristic functions. Taking the direct product of the majority votes on each factor, we get the “standard” majority function on $\mathcal{P}(X)$, again denoted by $\langle \rangle$. Thus, if $A_1, A_2, \dots, A_n \in \mathcal{P}(X)$, then $\langle A_1 A_2 \dots A_n \rangle$ is the set of elements of X which lie in a majority of the sets A_i . Note that $\langle \rangle$ can again be written by applying the expression E_n on the standard median on $\mathcal{P}(X)$ (since this holds independently on each factor).

More generally, let M be any median algebra. We define an n -ary operation, $\langle \rangle$, on M by writing it as the median expression, E_n , applied to the median operation on M . (Note that the choice of E_n does not matter, since any two such expressions are tautologically equivalent.) In fact, we already gave an equivalent, and perhaps more natural, description of this operation in the terms of the centre of a free median algebra — see the remark at the end of Subsection 5.3. We refer to the operation, $\langle \rangle$, as the **(standard) majority vote** on the median algebra, M . Note that we can recover the original median of $x, y, z \in M$ as $xyz = \langle x^{r-1} y^{r-1} z \rangle$. In particular, if $n = 3$, then $xyz = \langle xyz \rangle$.

Definition. A *consensus algebra* is a set, M , equipped with a (necessarily symmetric) n -ary operation which arises as the standard majority vote derived from some (necessarily unique) median algebra structure on M .

Here, as usual, $n \geq 3$, is assumed odd. We write $r = \lceil n/2 \rceil$. For $n = 3$, a consensus algebra is the same thing as a median algebra.

As we will see shortly, it is not hard to give an axiomatisation of the class of consensus algebras. Indeed they are defined by a finite set of identities.

First note that we can talk about “expressions” involving a symmetric n -ary operation in some formal alphabet, just as we did (for $n = 3$) in Section 6. We write \mathcal{E}_n for the set of such expressions. Let $\langle \rangle$ be a symmetric n -ary operation on a set, M . If $E \in \mathcal{E}_n$ has formal arguments x_1, x_2, \dots, x_n , and $a_1, a_2, \dots, a_n \in M$, then we can substitute a_i for x_i in E , and evaluate in M to give us an element $E(a_1, \dots, a_n) \in M$. If $E, E' \in \mathcal{E}_n$, we say that E, E' *coevaluate* in M if $E(a_1, \dots, a_n) = E'(a_1, \dots, a_n)$ for all possible substitutions a_1, \dots, a_n in M . More generally, if M also has an m -ary operation, and $E \in \mathcal{E}_n$ and $E' \in \mathcal{E}_m$, we again say that E, E' *coevaluate* in M if they both evaluate to equal elements under every substitution.

Given formal expressions, $E, E' \in \mathcal{E}_n$, we say that E, E' are *tautologically equivalent* (or that the identity $E = E'$ is *tautological*) if E, E' coevaluate in I equipped with the standard n -ary majority vote. (When $n = 3$ this agrees with the notion already defined in Section 6: if they coevaluate in I then they coevaluate in any median algebra.)

We can recursively enumerate all tautological identities, and we claim that this serves as an axiom system for majority consensus algebras. (In fact, at most three such identities will do.)

Note that we can formally define a map, $\alpha : \mathcal{E}_n \rightarrow \mathcal{E}_3$, by recursively replacing each occurrence of the n -ary operation with the expression E_n with the same arguments. By definition, E_n coevaluates with the median operation in I , and so, for any $E \in \mathcal{E}_n$, E and αE coevaluate in I . It follows that the identity $E = E'$ is tautological only if the identity $\alpha E = \alpha E'$ is tautological. If M is a median algebra, then all tautological median identities hold in M , so by the definition of the standard majority vote on M , all tautological n -ary identities also hold in M with this structure. In other words, M satisfies all our axioms.

Conversely, let M be a set with a symmetric n -ary operation, $\langle \rangle$. Define a ternary operation on M by $xyz := \langle x^{r-1}y^{r-1}z \rangle$. We define a map $\beta : \mathcal{E}_3 \rightarrow \mathcal{E}_n$ by recursively replacing the ternary operation with the above expression in \mathcal{E}_n which was used to define it. Thus, E and βE coevaluate in M , with these respective structures. Moreover, since the expression $\langle x^{r-1}y^{r-1}z \rangle$ also describes the standard median in I , we see that E and βE also coevaluate in I with its standard majority structures. It follows that the identity $E = E'$ is tautological if and only if $\beta E = \beta E'$ is tautological.

Now suppose M satisfies our axioms: in other words all tautological n -ary identities hold in M . Then all tautological ternary identities also hold in M . In particular, this applies to the identities (M1) and (M2) in the definition of a median algebra. Thus M is a median algebra with this ternary operation. In fact, for this we only need the axioms “(M1(n))” and “(M2(n))” obtained by applying β to both sides of (M1) and (M2) respectively.

Let $N(n)$ be the formal identity $\langle x_1 \dots x_n \rangle = \beta E_n$ in \mathcal{E}_n , where x_1, \dots, x_n are the arguments of E_n . Since E_n describes majority vote in I , the two sides coevaluate in I . In other words, the identity is tautological in \mathcal{E}_n , and so, by hypothesis it holds in M . Thus, $\langle x_1 \dots x_n \rangle$ and E_n coevaluate in M . In other words, $\langle \rangle$ is indeed the standard majority on M derived from the above structure as a median algebra. Thus, M is a consensus algebra.

In fact, only need the axioms (M1(n)), (M2(n)) and (N(n)) for the above to hold. This gives us a finite axiomatisation.

Of course, these axioms are a bit clumsy. A more elegant system is given in [BanMe], which we now go on to describe.

Let $\langle \rangle$ be a formal symmetric n -ary operation. (As usual, $n = 2r - 1 \geq 3$ is odd.) Consider the following identities, with arguments x_i, y, y_i, z_i :

(J1(n)): $\langle x_1 \dots x_{r-1} y^r \rangle = y,$

(J2(n)):

$$\begin{aligned} &\langle x_1 \dots x_{n-1} \langle y_1 \dots y_r z_1 \dots z_{r-1} \rangle \rangle \\ &= \langle \langle x_1 \dots x_{n-1} y_1 \rangle \langle x_1 \dots x_{n-1} y_2 \rangle \dots \langle x_1 \dots x_{n-1} y_r \rangle z_1 \dots z_{r-1} \rangle. \end{aligned}$$

(Note that, when $n = 3$, these correspond respectively to properties (M1) and (M3) of a median algebra.)

Now let $\langle \rangle$ be a symmetric n -ary operation on a set, M . The following is shown in [BanMe]:

Theorem 20.3.1. *($M, \langle \rangle$) is a consensus algebra if and only if (J1(n)) and (J2(n)) hold for all $x_i, y, y_i, z_i \in M$.*

(Of course, we could retrospectively treat this as a formal definition of a consensus algebra.)

One first checks that these axioms indeed hold in a consensus algebra. As described earlier, this is equivalent to showing that they are tautological: in other words hold in I with its standard majority vote. We leave this verification as a straightforward exercise.

Conversely, suppose $(M, \langle \rangle)$ satisfies (J1(n)) and (J2(n)).

We can immediately define the ternary operation on M by $xyz := \langle x^{r-1} y^{r-1} z \rangle$.

The proof in [BanMe] begins with:

Lemma 20.3.2. *With the above ternary operation, M is a median algebra.*

The proof of this in [BanMe] proceeds by induction over n . For simplicity, we will restrict here to the case where $n = 5$; in other words, the first stage of the induction.

In this case, we have $xyz = \langle x^2y^2z \rangle$ and

(J1(5)): $\langle xyz^3 \rangle = z$, and

(J2(5)): $\langle xyzw\langle abcde \rangle \rangle = \langle \langle xyzwa \rangle \langle xyzwb \rangle \langle xyzwc \rangle de \rangle$.

We first check that the ternary operation thus defined is symmetric. It is enough to show that $\langle ab^2c^2 \rangle = \langle ba^2c^2 \rangle$ for all $a, b, c \in M$. Now, $\langle ab^2c^2 \rangle = \langle b^2c^2\langle a^2ca^2 \rangle \rangle = \langle \langle b^2c^2a \rangle^2 \langle b^2c^2c \rangle a^2 \rangle = \langle \langle b^2c^2a \rangle^2 ca^2 \rangle = \langle \langle bc^2ab \rangle^2 \langle bc^2ac \rangle a^2 \rangle = \langle bc^2a \langle b^2ca^2 \rangle \rangle = \langle abc^2 \langle a^2b^2c \rangle \rangle$. The last expression is symmetric in a and b , so reversing the argument, we see that it is also equal to $\langle ba^2c^2 \rangle$ as required.

Now property (M1) of a median algebra is immediate from (J1(5)). It remains to verify (M2). Making use of the symmetry we have just established, for all $a, b, c, d \in M$, we have: $\langle a^2b^2\langle a^2cd^2 \rangle \rangle = \langle \langle a^2b^2a \rangle^2 \langle a^2b^2c \rangle d^2 \rangle = \langle a^2 \langle a^2b^2c \rangle d^2 \rangle = \langle a^2 \langle a^2c^2b \rangle d^2 \rangle = \langle \langle a^2c^2a \rangle^2 \langle a^2c^2b \rangle d^2 \rangle = \langle a^2c^2 \langle a^2bd^2 \rangle \rangle$. In other words, $ab(acd) = ac(abd)$, which is (M2) as required.

This proves Lemma 20.3.2 when $n = 5$.

It remains to show that the n -ary operation, $\langle \rangle$, is precisely the majority vote derived from this particular median algebra.

We return to the general case. We will use the terms “subalgebra”, “homomorphism” etc, to refer to general n -ary operations, with the obvious meanings. To avoid ambiguity, we will sometimes use “ n -subalgebra”, “ n -homomorphism” etc. We will take for granted (Lemma 20.3.2) that M is a median algebra.

We begin with:

Lemma 20.3.3. *For all $x_1, \dots, x_{n-1}, a, b, c \in M$, we have*

$$\langle x_1 \dots x_{n-1}(abc) \rangle = \langle x_1 \dots x_{n-1}a \rangle \langle x_1 \dots x_{n-1}b \rangle c.$$

Proof. Abbreviating $x_1 \dots x_{n-1}$ to \underline{x} , and applying (J2(n)), we have: $\langle \underline{x}(abc) \rangle = \langle \underline{x} \langle a^{r-1}bc^{r-1} \rangle \rangle = \langle \langle \underline{x}a \rangle^{r-1} \langle \underline{x}b \rangle c^{r-1} \rangle = \langle \underline{x}a \rangle \langle \underline{x}b \rangle c$. \square

Lemma 20.3.4. *$(M, \langle \rangle)$ is n -isomorphic to a subalgebra of $\mathcal{P}(X)$ for some set X .*

As usual, we are taking the standard majority function on $\mathcal{P}(X)$, identified with I^X . Here the subalgebra of $\mathcal{P}(X)$ will be a median subalgebra, hence also an n -subalgebra with the standard majority (since the latter is expressible in terms of the median).

Proof. By Lemma 20.3.2 and Proposition 3.2.13, we know that there is a 3-monomorphism $\phi : M \rightarrow \mathcal{P}(X)$, for some set X . In other words, M is isomorphic to the 3-subalgebra, ϕM , of $\mathcal{P}(X)$ with its standard median. We need to check that ϕ is also an n -homomorphism. It is enough to check that, for each $x \in X$, the map $f := \pi \circ \phi : M \rightarrow I$ is an n -homomorphism, where $\pi : \mathcal{P}(X) \rightarrow I$ is projection to the x -coordinate. We already know that f is a 3-homomorphism, and we can assume that it is surjective.

We define another n -ary operation, $\langle \langle \rangle \rangle$, on I , by setting $\langle \langle f(a_1) \dots f(a_n) \rangle \rangle := f(\langle a_1 \dots a_n \rangle)$, for $a_1, \dots, a_n \in M$. We first check that this is well defined. It is enough to show that, for all $x_1, \dots, x_{n-1}, a, b \in M$ with $f(a) = f(b)$, we have $f(\langle \underline{x}a \rangle) = f(\langle \underline{x}b \rangle)$, where \underline{x} is an abbreviation for $x_1 \dots x_{n-1}$. Using Lemma 20.3.3 and the fact that f is a 3-homomorphism, we have $f(\langle \underline{x}a \rangle) = f(\langle \underline{x}(aba) \rangle) = f(\langle \underline{x}a \rangle \langle \underline{x}b \rangle a) = f(\langle \underline{x}a \rangle) f(\langle \underline{x}b \rangle) f(a) = f(\langle \underline{x}a \rangle) f(\langle \underline{x}b \rangle) f(b) = f(\langle \underline{x}(abb) \rangle) = f(\langle \underline{x}b \rangle)$ as required.

In fact, we claim that $\langle \langle \rangle \rangle$ is precisely the standard majority on I . To see this, choose $a, b \in M$ with $f(a) = 0$ and $f(b) = 1$. If $p > q$, then (J1(n)) gives $\langle \langle 0^p 1^q \rangle \rangle = f(\langle a^p b^q \rangle) = f(a) = 0 = \langle 0^p 1^q \rangle$, so the claim follows by symmetry.

By construction, f is an n -isomorphism to its range with respect to the structure $\langle \langle \rangle \rangle$ on the range, which we now know is also the standard majority. \square

We observed earlier that the standard majority on $\mathcal{P}(X)$ can be expressed by applying the expression E_n to the standard median on $\mathcal{P}(X)$. Since ϕ is an isomorphism to its range, the same is also true in M . Therefore, by definition, the operation $\langle \rangle$ is the derived majority vote, as required.

This completes the proof of Theorem 20.3.1 (at least in the case $n = 5$, as we have presented it).

The proof, as we have presented it, is non-constructive (unlike the previous axiomatisation we gave). In particular it made use of the Axiom of Choice, via the application of Proposition 3.2.13 to Lemma 20.3.4. In fact, this can be circumvented: one can work in the subalgebra generated by any finite subset (see Subsection 6.3). In any case, it would be nice to have a simple derivation of the identities needed for this argument to work.

We end this subsection with an explicit expression for majority vote in terms of the median, which is described in [BanMe].

We write $\langle \rangle_n$ for the n -ary majority vote, so that $\langle xyz \rangle_3 = xyz$. For $n \geq 3$ consider the identity

$$\langle a_1 \dots a_{r-1} b_1 \dots b_r c_1 c_2 \rangle_{n+2} = \langle a_1 \dots a_{r-1} d_1 \dots d_r \rangle_n$$

where $d_i = (c_1 c_2 b_1 b_2 \dots b_{i-1} | b_i)$ in the notation introduced in Subsection 5.1 (so $d_1 = c_1 c_2 b_1$).

Proposition 20.3.5. *The above identity is tautological for every odd $n \geq 3$.*

Proof. We need to check that both sides formula coevaluate in I . Using the notation of Subsection 6.2, we assume for contradiction that $\text{LHS}\uparrow$ and $\text{RHS}\downarrow$.

Write $A = \#\{i \mid a_i\uparrow\}$, $B = \#\{i \mid b_i\uparrow\}$, $C = \#\{i \mid c_i\uparrow\}$ and $D = \#\{i \mid d_i\uparrow\}$. Now $\text{LHS}\uparrow$ and $\text{RHS}\downarrow$ respectively give $A + B + C \geq r + 1$ and $A + D \leq r - 1$. Thus $D + 2 \leq B + C$, so certainly $D \leq B$.

We say that an index i is “bad” if $b_i\uparrow$ and $d_i\downarrow$. Note that this implies $c_1\downarrow$, $c_2\downarrow$ and $b_j\downarrow$ for all $j < i$. In particular, there can be at most one bad index. Thus $B \leq D + 1$, and so $C \geq 1$. In other words, $c_1\uparrow$ or $c_2\uparrow$, and so in fact, there are no bad indices. Thus, $B \leq D$, so $C = 2$, i.e. $c_1\uparrow$ and $c_2\uparrow$. Since there are no bad indices, $b_i\uparrow \Rightarrow d_i\uparrow$, and since $B = D$, we also have $d_i\uparrow \Rightarrow b_i\uparrow$. Now $d_1 = c_1c_2b_1\uparrow$, so $b_1\uparrow$, and so $d_2 = (c_1c_2b_1|b_2)\uparrow$. Continuing, we get $d_2\uparrow \Rightarrow b_2\uparrow \Rightarrow d_3\uparrow \Rightarrow \cdots \Rightarrow d_r\uparrow$, giving the contradiction that $\text{RHS}\uparrow$. \square

We can view the formula as an inductive prescription for writing majority vote as a median expression. In fact, one can get lots of different formulae by permuting the arguments at each stage of the recursion.

For example, for $n = 5$, we get $\langle abcde \rangle = \langle adefc \rangle = a(bcd)(bcd|e) = a(bcd)((bce)de)$, which is the formula we saw earlier.

A different general formula is also described in [BanMe].

21. GROUP ACTIONS

The motivation for much of the study of median algebras and metric spaces comes from group theory, and applications thereof. For example, there is a vast literature concerning group actions on cube complexes, \mathbb{R} -trees, etc. Much of this has generalisations to median algebras and median metric spaces. In this section, we give a sample of results in this context. In some respects this account is unsatisfactory, in that it is often unclear whether a particular hypothesis is essential, or if the conclusion is optimal. A more systematic account, mostly in the finite-rank case, can be found, in [Fi4]. Further results can be found in other papers of Fioravanti, as we mention in the Notes to this section. An extensive study in the discrete case (i.e. median graphs) can be found in [Ge3].

21.1. Fixed points.

In this subsection we will mostly focus on fixed-point sets of automorphisms and group actions. We begin in a very general setting.

Let M be a median algebra. Let $P \subseteq M$ be a finite subset, and let $n = \#P$. Let Φ be the free median algebra on P , and let $\phi : \Phi \rightarrow M$ be the homomorphism extending the inclusion, $P \hookrightarrow M$. Recall from Subsections 5.4 and 11.9, that if n is odd, then Φ has a central element, $c \in \Phi$; and if n is even, it has a central cube, $Q \subseteq \Phi$, with $\text{rank}(Q) = \frac{1}{2}\binom{n}{n/2}$. Let $C(P) \subseteq M$ be $\{\phi c\}$ or ϕQ in the respective cases. This set is canonically determined by P — in the odd case a singleton, and in the even case, a cube of rank at most $\frac{1}{2}\binom{n}{n/2}$.

Suppose G is a group acting by automorphisms on M , with a finite orbit, A . Then $C(A)$ is a G -invariant subset. In summary, this shows:

Lemma 21.1.1. *Let G be a group acting by automorphism on M , with finite orbit, A . Let $n = \#A$. If n is odd, then G has a fixed point in M . If n is even, then G has a G -invariant cube of rank at most $\frac{1}{2}\binom{n}{n/2}$.*

Clearly, this applies to any finite group, G , and we could set $n = \#G$.

(Ex21.1): In general, an automorphism of even order need not have a fixed point. For example, consider a cube, Q , of finite rank, n . Let $\mathcal{H}(Q)$ be the set of halfspaces of Q , with the involution, $[H \mapsto H^*]$, defined by taking complements. Let $\pi : \mathcal{H}(Q) \rightarrow \mathcal{H}(Q)$ be a permutation such that $\pi(H^*) = (\pi(H))^*$ for all $H \in \mathcal{H}(Q)$. This is a proset automorphism, and induces a median automorphism, ϕ , of Q . (See Section 9.) As a particular example, denoting the halfspaces by H_i and H_i^* for $i \in \{1, \dots, n\}$, let π be the cyclic permutation, $(H_1, \dots, H_n, H_1^*, \dots, H_n^*)$. Note that $\pi^n(H) = H^*$. In this case, ϕ has order $2n$, and ϕ^n is the antipodal map. In particular, ϕ has no fixed point.

Given any automorphism, $f : M \rightarrow M$, write $\text{fix}(f) = \text{fix}(f, M) = \{x \in M \mid fx = x\}$. If G is a group action on M , let $\text{fix}(G) = \bigcap_{f \in G} \text{fix}(f)$. These are subalgebras of M .

Suppose that G is finite, and write $F = \text{fix}(G)$. Suppose $F \neq \emptyset$. Choose any $a \in F$. Given any $x \in M$, let Gx be its G -orbit, and let $r(x) = (Gx|a)$ (in the notation introduced in Subsection 5.1). Note that $r(x) \in F$ for all $x \in M$, and $r(x) = x$ for all $x \in F$. In other words, $r : M \rightarrow F$ is a retraction (not necessarily a homomorphism).

Suppose now that M is a topological median algebra, and that G acts by continuous automorphism. In this case, F is (topologically) closed in M , and the map, r , is continuous. In particular, F is a topological retract of M . Therefore, if M is connected, so is F . This shows:

Lemma 21.1.2. *If G is a finite group acting by continuous automorphisms on a connected topological median algebra, M , then $\text{fix}(G)$ is a closed connected subalgebra of M .*

(This admits the possibility that $\text{fix}(G)$ may be empty.)

Next we prove the following:

Proposition 21.1.3. *Let M be a connected interval-compact median metric space. Then any finite-order isometry of M has a fixed point.*

Let $f : M \rightarrow M$ be an isometry of order n . Note that it is enough to verify the statement when n is prime. For suppose $n = st$, where $s, t > 1$. Now f^s has order t , so by induction on n , we can suppose that $\text{fix}(f^s) \neq \emptyset$. By Lemma 21.1.2, $\text{fix}(f^s)$ is connected and interval-compact. Now $f|_{\text{fix}(f^s)}$ has order s , so again by induction, has a fixed point.

Now we know by Lemma 21.1.1 that any odd-order isometry has a fixed point. Therefore we can assume that $n = 2$. Moreover, we can assume that M is compact: choose any $a \in M$, and replace M by the interval $[a, fa]$. Proposition 21.1.3 therefore follows from the following:

Lemma 21.1.4. *Any isometric involution of a compact connected median metric space has a fixed point.*

For the proof, we note that any two points, a, b in a connected median metric space, a, b have a midpoint c : that is such that $\rho(a, b) = 2\rho(a, c) = 2\rho(b, c)$. This is because the map $[x \mapsto \rho(a, abx)] : M \rightarrow [0, \rho(a, b)]$ is continuous, hence the range is connected, and hence all of the real interval $[0, \rho(a, b)]$. (In the compact case, this also follows immediately from the fact that M is geodesic, by Lemma 13.3.2. However, we will need the more general observation later: see Lemma 21.3.3.)

Proof. Suppose, for contradiction, that $f : M \rightarrow M$ is an involution with no fixed point. After rescaling the metric, we may as well assume that $\min\{\rho(x, fx) \mid x \in M\} = 1$. We aim to construct a sequence of f -invariant cubes, $Q_0 \hookrightarrow Q_1 \hookrightarrow Q_2 \hookrightarrow \dots$, similarly as in Example (Ex13.9) of Subsection 13, and thereby contradicting the compactness of M . We can begin by choosing any $x \in M$ with $\rho(x, fx) = 1$, and setting $Q_0 = \{x, fx\}$. The inductive step can be described as follows.

Suppose $Q \subseteq M$ is a p -cube with $p \geq 1$, with $fQ = Q$ and with $f|_Q$ the antipodal map. Also suppose that the metric on Q is $1/p$ times the combinatorial metric, so that antipodal points are distance 1 apart. We construct a $(2p)$ -cube, $Q' \supseteq Q$, as follows.

Recall, by Lemma 10.3.5, that we can identify $\text{hull}(Q)$ with a direct product, $\prod_{i=1}^p D_i$, where $D_i \cong [a_i, b_i]$ and $\{a_i, b_i\}$ is a 1-face of Q .

Fix, for the moment, some 1-cell, $\{a, b\}$, of Q . Let $a' = fb$ and $b' = fa$. Then $\{b', a'\}$ is a parallel face of Q , and $f|[a, b]$ is an isometry of $[a, b]$ to $[a', b']$. We also have a translation, $\tau : [a, b] \rightarrow [a', b']$, with $\tau a = a'$, $\tau b = b'$, and $\rho(x, \tau x) = \frac{p-1}{p}$, for all $x \in [a, b]$ (see the discussion of parallel sets in Subsection 7.2). Since M is connected, we can choose a midpoint, $c \in [a, b]$, of a, b . In other words, $\rho(a, b) = 2\rho(a, c) = 2\rho(b, c) = 1/p$. We claim that c has an ‘‘opposite point’’, $d \in [a, b]$, namely such that $\{a, c, b, d\}$ is a 2-cube. Note that, if such a point exists, it is unique, by Lemma 10.3.6. In fact, we can set $d = abd' = \tau^{-1}d'$, where $d' = fc \in [a', b']$. Now d is also a midpoint of a, b , so it remains to check that $\rho(c, d) = 1/p$. For one direction, we have $\rho(c, d) \leq \rho(c, a) + \rho(a, d) = 2(1/2p) = 1/p$. For the other direction, we have $1 \leq \rho(c, d') \leq \rho(c, d) + \rho(d, d') = \rho(c, d) + \frac{p-1}{p}$, so $\rho(c, d) \geq 1/p$ as required. Note also that d' is the unique point opposite $c' := \tau c$ in $[a', b']$.

To construct Q' , we now choose any midpoint, c_i , of a_i, b_i in $[a_i, b_i]$, and let $d_i \in [a_i, b_i]$ be the unique point opposite c_i . Let $Q' = \prod_{i=1}^p \{a_i, c_i, b_i, d_i\}$, which is a $(2p)$ -cube in $\text{hull}(Q) \subseteq M$. By construction, if $\{a, b\}$ is any 1-cell of Q , then, in the

above notation, we have $fc = d'$ and $fd = c'$. In other words, we have $fQ' = Q'$, and $f|_{Q'}$ is the antipodal map.

Let $h(Q') = (c_1, c_2, \dots, c_p) \in Q'$. Note that $\rho(h(Q'), x) = p(1/2p) = 1/2$ for all $x \in Q$.

We now construct Q_n inductively, by setting $Q_{n+1} = Q'_n$, so that Q_n is a 2^n -cube. Let $h_n = h(Q_n) \in Q_n$. By construction, $\rho(h_m, h_n) = 1/2$ for all $m \neq n$. This contradicts the compactness of M . \square

This concludes the proof Proposition 21.1.3.

We note that Lemma 21.1.4 (hence Proposition 21.1.3) does not in general hold if we replace “compact” with “complete”: see Example (Ex13.8) of Subsection 13.1.

It is unclear how much more one can say regarding actions on (possibly) infinite-rank median metric spaces. However, there is a lot more one can say in the finite-rank case.

We note, for example, that Proposition 21.1.3 holds if we replace “interval-compact” with “finite-rank”: we similarly derive a contradiction, in the proof of Lemma 21.1.4.

In the complete case, one can say more. Note that, by Lemma 13.2.10, any complete finite-rank median metric space is interval-compact. Invoking some additional machinery, we have:

Proposition 21.1.5. *Let M be a complete connected finite-rank median metric space, and let G be a group acting by isometries on M with a bounded orbit. Then G has a fixed point.*

Proof. By the same principle as applied to Lemma 21.1.1, it’s enough to show that one can associate a canonical “centre” to any bounded subset, $P \subseteq M$. To this end one can invoke Theorem 13.4.2 to obtain a canonical CAT(0) metric, σ , on M . We now apply a well known fact regarding complete CAT(0) metrics, namely that there is a unique point, $c \in M$ which minimises $\sup\{\sigma(x, c) \mid x \in P\}$, for any finite (or indeed bounded) subset $P \subseteq M$ (see, for example, Proposition II.2.7 of [BriH], as well as Corollary II.2.8 thereof). \square

In fact, Proposition 21.1.5 still holds if we drop the assumption that M is complete. This is a consequence of Corollary A of [Fi4], which uses a quite different argument.

We finish this subsection with:

Proposition 21.1.6. *Let M be a complete connected finite-rank median metric space, and let G be a group acting by isometries on M . If $\text{fix}(G)$ is non-empty, then it is contractible, and M retracts onto $\text{fix}(G)$.*

Proof. Let σ be the canonical CAT(0) metric, as in the proof of Proposition 21.1.5. Any two points, $x, y \in M$, are connected by a unique σ -geodesic, $\alpha \subseteq M$. If $x, y \in \text{fix}(G)$, then $\alpha \subseteq \text{fix}(G)$. In other words, $\text{fix}(G)$ is convex in the metric σ . Thus, M is intrinsically CAT(0), hence contractible. Moreover, M retracts onto $\text{fix}(G)$ by nearest-point retraction. \square

21.2. Actions on cube complexes.

We move on to consider isometries (generally) without fixed points. In this subsection, we will mostly restrict to cube complexes. Some results on more general spaces will be discussed in Subsection 21.3. First, we make some simple observations about isometries of general metric spaces.

Let (M, ρ) be a metric space and let $f : M \rightarrow M$ be an isometry. We have a spectrum of translation distances, namely $S(f, M) := \{\rho(x, fx) \mid x \in M\}$. Let $L = L(f, M) = \inf S(f, M)$, and let $R(x, fx) = \{x \in M \mid \rho(x, fx) = L\}$. In other words, $R(f, M)$ is the (possibly empty) set of points moved a minimal distance.

We let $R_0(f, M)$ be the set of $x \in M$ such that for all $n \in \mathbb{N}$ we have $\rho(x, f^n x) = n\rho(x, fx)$. Note that, for all $n \in \mathbb{N}$, we have $L(f^n, M) \leq nL(f, M)$ and $R_0(f, M) \subseteq R_0(f^n, M)$.

Lemma 21.2.1. *We have $R_0(f, M) \subseteq R(f, M)$. Moreover, if $R_0(f, M) \neq \emptyset$, then for all $n \in \mathbb{N}$, $L(f^n, M) = nL(f, M)$ and $R_0(f^n, M) = R(f^n, M)$.*

Proof. We can assume that $R_0(f, M) \neq \emptyset$. Let $x \in R_0(f, M)$ and $y \in M$. For all $n \in \mathbb{N}$, we have $n\rho(x, fx) = \rho(x, f^n x) \leq \rho(y, f^n y) + 2\rho(x, y) \leq n\rho(y, fy) + 2\rho(x, y)$. Letting $y \rightarrow \infty$, we see that $\rho(x, fx) \leq \rho(y, fy)$. Therefore, $L(f, M) = \rho(x, fx)$, and so $x \in R(f, M)$. This shows that $R_0(f, M) \subseteq R(f, M)$.

We also have $x \in R_0(f^n, M)$, so the above also gives $R_0(f^n, M) \subseteq R(f^n, M)$, and $L(f^n, M) = \rho(x, f^n x) = n\rho(x, fx) = nL(f, M)$. We also get $\rho(y, f^n y) \geq nL(f, M)$, so if $y \in R(f, M)$, we also have $\rho(y, f^n y) \leq nL(f, M)$, so $y \in R_0(f, M)$. \square

Definition. We say that f is **semisimple** if $R(f, M) \neq \emptyset$, and **simple** if $R_0(f, M) \neq \emptyset$.

(The term “semisimple” has been used elsewhere to mean either or both of these things. The term “simple” is not a standard one, but we want to make a clear distinction.)

Note that if $\text{fix}(f) \neq \emptyset$, then $\text{fix}(f) = R(f, M) = R_0(f, M)$, so f is simple.

By Lemma 21.2.1, simple implies semisimple. We will see various situations where the converse holds, but it does not in general. Here is one example, which is discussed in [Fi4] (Example 2.6 thereof).

(Ex21.2): Let $f : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ be a parabolic isometry of the hyperbolic plane, \mathbb{H}^2 . That is, f has no fixed point in \mathbb{H}^2 but fixes precisely one point of the ideal boundary. Then $L(f, \mathbb{H}^2) = 0$, but $R(f, \mathbb{H}^2) = \emptyset$. Also orbits are unbounded, i.e. $\rho(x, f^n x) \rightarrow \infty$ for all $x \in \mathbb{H}^2$.

Recall by Example (Ex19.4) of Subsection 19.4, that \mathbb{H}^2 canonically isometrically embeds into a connected complete locally compact median metric space, M . Thus, f extends to an isometry $f : M \rightarrow M$, with $L(f, M) = 0$, but with no fixed point.

Note that by taking a direct product with a translation of \mathbb{R} , we can construct a non-semisimple isometry, f , of $M \times \mathbb{R}$, with $L(f, M \times \mathbb{R})$ taking any non-negative value we want.

Note that here $\text{rank}(M) = \infty$. It is shown in [Fi4] that this cannot happen in the finite-rank case: see Theorem 21.3.4 below.

We note that if M is a geodesic metric space, and $L(f, M) > 0$, then $x \in R_0(f, M)$ if and only if x lies in an f -invariant bi-infinite geodesic (a concatenation of geodesic segments connecting $f^i x$ to $f^{i+1} x$ for all $i \in \mathbb{Z}$). This geodesic is translated a distance $L(f, M)$ by f .

We now consider various specific cases.

Let Π be a discrete median algebra (as described in Section 11). Recall from Subsection 13.3 that there is a natural bijection between compatible median metrics on Π and maps $w : \mathcal{W}(\Pi) \rightarrow (0, \infty)$, where $\mathcal{W}(\Pi)$ is the set of walls. The median metric, ρ , is given by $\rho(x, y) = \sum_{W \in \mathcal{W}(x, y)} w(W)$, where $\mathcal{W}(x, y)$ is the set of walls separating x, y . We refer to a space arising in this way as a **discrete median metric space**. When $w \equiv 1$, we get the standard combinatorial metric on Π , which we will denote here by ρ_0 . Note that a sequence in Π is geodesic with respect to ρ (or to ρ_0) if and only if it is monotone. In particular a geodesic sequence in one compatible median metric will also be geodesic in any other.

Clearly, any automorphism, $f : \Pi \rightarrow \Pi$, is an isometry of (Π, ρ_0) . Note that, in this metric, $S(f, \Pi) \subseteq \mathbb{N}$. In particular, we see that f is semisimple with respect to this metric.

Definition. Let $f : \Pi \rightarrow \Pi$ be an automorphism. We say that f acts **without inversion** if for every halfspace, $H \subseteq \Pi$, and every $n \in \mathbb{Z}$, $f^n H \neq \Pi \setminus H$.

The following result is proven in [Hag] for the combinatorial metric, ρ_0 . In fact, essentially the same argument works for any compatible median metric. This is the basis of the proof we present here, though we have expressed it a bit differently.

Theorem 21.2.2. *Let (Π, ρ) be a discrete median metric space, and let $f : \Pi \rightarrow \Pi$ be an isometry acting without inversion. Then $R_0(f, \Pi) = R(f, \Pi)$.*

In view of Lemma 21.2.1, this is equivalent to asserting that any semisimple isometry acting without inversion on a discrete median metric space is simple.

For the proof, we use the following construction. The idea will be used again in the proof of Proposition 21.3.1.

Let $f : \Pi \rightarrow \Pi$ be an isometry of a median metric space, Π . Write $L = L(f, \Pi)$. Given $x \in \Pi \setminus R_0(f, \Pi)$, let $n = n(x)$ be maximal such that $\rho(x, f^n x) = Ln$. Clearly $n(x) > 0$.

Lemma 21.2.3. *Suppose $R(f, \Pi) \neq \emptyset$ and $R_0(f, \Pi) = \emptyset$. Let $x \in R(f, \Pi)$ minimise $n := n(x)$ among all such x , and let $p = x(f^{-1}x)(f^n x)$. Then $x \neq p$. Moreover, if $y \in [x, p]$ then we have $\rho(y, f^n y) = \rho(x, f^n x)$, $x.y.f^n x$, $y.x.f^n y$ and $x.f^n y.f^n x$.*

Note that this lemma does not require Π to be discrete. The fact that $x.f^n y.f^n x$ holds will only be needed for the proof of Proposition 21.3.1.

Proof. By the definition of n , and the triangle inequality, we have $\rho(f^{-1}x, f^n x) = \rho(x, f^{n+1}x) < L(n+1)$. In particular, $x \notin [f^{-1}x, f^n x]$, so $x \neq p$. Since $y \in [x, p] \subseteq [x, f^{-1}x]$, we have $\rho(y, fy) \leq \rho(x, y) + \rho(x, fy) = \rho(x, y) + \rho(f^{-1}x, y) = \rho(x, f^{-1}x) = \rho(x, fx) = L$, and so $\rho(y, fy) = L$. Thus, $y \in R(f, \Pi)$ and $x \in [y, fy]$. Also, since $y \in [x, p] \subseteq [x, f^n x]$ we have $x.y.f^n x$.

Now by minimality of n , we have $n(y) \geq n$, and so $\rho(y, f^n y) = Ln = \rho(y, fy) + \rho(fy, f^n y)$. Thus, $fy \in [y, f^n y]$ and so $x \in [y, fy] \subseteq [y, f^n y]$. It follows that $\rho(x, f^n y) + \rho(f^n y, f^n x) = \rho(x, f^n y) + \rho(x, y) = \rho(y, f^n y) = Ln = \rho(x, f^n x)$, and so $f^n y \in [x, f^n x]$. \square

Proof of Theorem 21.2.2. By Lemma 21.2.1, $R_0(f, \Pi) \subseteq R(f, \Pi)$, so we can assume that $R(f, \Pi) \neq \emptyset$. We can also assume that $L := L(f, \Pi) > 0$. (Otherwise $R_0(f, \Pi) = R(f, \Pi)$ is the fixed-point set.) By Lemma 21.2.1 again, it is enough to show that $R_0(f, \Pi) \neq \emptyset$.

Suppose, for contradiction that $R_0(f, \Pi) = \emptyset$. As in Lemma 21.2.3, choose $x \in R(f, \Pi)$ so as to minimise $n := n(x)$ among such x , and set $p = x(f^{-1}x)(f^n x) \neq x$. Let y be adjacent to x in $[x, p]$. Then $\mathcal{W}(x, y) = \{W\}$ for some $W \in \mathcal{W}(\Pi)$. We can suppose that $x \in W^-$ and $y \in W^+$. By $y.x.f^n y$, we have $f^n y \in W^-$, and by $x.y.f^n x$, we have $f^n x \in W^+$. Since $f^n x, f^n y$ are adjacent, $\mathcal{W}(f^n x, f^n y) = \{W\}$. Thus $f^n W^- = W^+$ and $f^n W^+ = W^-$, contrary to our hypothesis that f acts without inversion. \square

We have observed that if $\rho = \rho_0$, then f is necessarily semisimple. From this we see:

Corollary 21.2.4. *Let $f : \Pi \rightarrow \Pi$ be an automorphism of a discrete median algebra acting without inversion. Then f either has a fixed point, or it preserves a bi-infinite monotone path.*

Proof. Choose any $x \in R(f, \Pi)$. If $fx \neq x$, let π be a monotone path from x to fx . Since $x \in R_0(f, \Pi)$, the bi-infinite path, $\bigcup_{i \in \mathbb{Z}} f^i \pi$ is also monotone. \square

One can get information about more general automorphisms by passing to the binary subdivision — such an automorphism necessarily acts without inversion. Let us first recall some facts from Section 11.

Let Π be a discrete median algebra, and let $\Sigma = \Sigma(\Pi)$ be its binary subdivision. We have a natural isomorphism, $[x \mapsto Q(x)] : \Sigma \rightarrow \mathcal{C}(\Pi)$, where $\mathcal{C}(\Pi)$ is the set of cells of Π . We can think of $x \in \Sigma$ as the central point of the cell $Q(x)$ of Π . We can define a “2-colouring” of Σ by assigning to x the value of $\text{rank}(Q(x))$ modulo 2. Clearly Π is coloured 0, and adjacent vertices of Σ have opposite colours. Any automorphism f of Π extends to an automorphism of Σ , and respects this colouring. It follows that for any $x \in \Sigma$, the combinatorial distance from x to fx in Σ is even.

We also have a natural surjection, $\mathcal{W}(\Sigma) \rightarrow \mathcal{W}(\Pi)$. The preimage of any wall, $W \in \mathcal{W}(\Pi)$ consists of two parallel walls of Σ . In fact, if $\{x, y\} \in \mathcal{C}_1(\Pi)$ is a 1-cell of Π crossing W , and $z \in \Sigma$ is its central point (so that $[x, y]_\Sigma = \{x, z, y\}$) then these are the walls of Σ which respectively cross the 1-cells $\{x, z\}$ and $\{y, z\}$ of Σ .

Note that each wall of Σ has a preferred “large” halfspace: in the above notation these are the halfspaces containing z . (It is easily checked that this is independent of our choice of $\{x, y\}$.) Now any automorphism of Π must preserve the family of large halfspaces of Σ , and so f acts on Σ without inversion, as we claimed earlier.

If ρ is a compatible median metric on Π , there is an obvious way to extend this to a compatible median metric, ρ^Σ , on Σ ; namely, such that if $x, y \in \Pi$ are adjacent and $z \in \Sigma$ is their central element, then $\rho^\Sigma(x, y) = \rho(x, y) = 2\rho^\Sigma(x, z) = 2\rho^\Sigma(y, z)$. Clearly, if ρ_0 is the combinatorial metric on Π , then ρ_0^Σ is $\frac{1}{2}$ times the combinatorial metric on Σ . From an earlier observation, we see that if f is any automorphism of Π and $x \in \Sigma$, then $\rho_0^\Sigma(x, fx) \in \mathbb{N}$.

In summary, this shows:

Lemma 21.2.5. *Any automorphism, $f : \Pi \rightarrow \Pi$, of a discrete median algebra Π extends to an automorphism, $f : \Sigma \rightarrow \Sigma$ of the binary subdivision acting without inversion on Σ . Moreover, with respect to the metric ρ_0^Σ , we have $L(f, \Sigma) \in \mathbb{N}$.*

In particular, by Corollary 21.2.4, f either fixes a point of Σ , or preserves a bi-infinite monotone path in Σ .

We can obtain a similar statement concerning real cube complexes as follows.

Recall from Subsection 11.2, that a discrete median algebra, Π , embeds as a subalgebra of its realisation, $\Delta(\Pi) = \bigcup_{Q \in \mathcal{C}(\Pi)} \Delta(Q)$. Given a finite cube, Q , write $\partial\Delta(Q)$ for the union of all proper faces of $\Delta(Q)$, and write $\text{int}(\Delta(Q)) = \Delta(Q) \setminus \partial\Delta(Q)$ for its “interior”. (Thus, if Q is a 0-cell, then $\text{int}(\Delta(Q)) = \Delta(Q) = Q$.) Set-theoretically, $\Delta(\Pi)$ is a disjoint union of interiors, $\text{int}(\Delta(Q))$, as Q ranges over $\mathcal{C}(\Pi)$. Given $x \in \Delta(\Pi)$, write $Q(x)$ for the unique $Q \in \mathcal{C}(\Pi)$ with $x \in \text{int}(\Delta(Q))$. Note that we can also embed $\Sigma(\Pi)$ as a subalgebra of $\Delta(\Pi)$. If $x \in \Delta(Q)$, then we can think of x as the “central element” of $\Delta(Q)$. (This is consistent with the earlier notation.) We can also identify $\Delta(\Sigma(Q))$ with $\Delta(Q)$.

Recall that any compatible metric, ρ , on Π extends to a geodesic median metric in $\Delta(\Pi)$: see Example (Ex13.4) of Subsection 13.1, and Proposition 17.1.4. (This induces the metric ρ^Σ on $\Sigma(Q)$ as defined above.) Any isometry, $f : \Pi \rightarrow \Pi$ induces an isometry, $f : \Delta(\Pi) \rightarrow \Delta(\Pi)$. We refer to an isometry of $\Delta(\Pi)$ which arises in this way as a **cellular isometry**. If this has a fixed point, then $\text{fix}(f)$ is contractible, and $\Delta(\Pi)$ retracts onto $\text{fix}(f)$. (This follows exactly as with the proof of Proposition 21.1.6, on taking σ to be the l^2 metric on $\Delta(\Pi)$, which is CAT(0).) So we will assume that f has no fixed point in $\Delta(\Pi)$. This implies that $fQ \neq Q$ for all $Q \in \mathcal{C}(\Pi)$ (otherwise f would fix the central element of Q). It is sometimes helpful to think in terms of the quotient space, $\Delta(\Pi)/\langle f \rangle$, where $\langle f \rangle$ is the isometry group generated by f . Note that for all $Q \in \mathcal{C}(\Pi)$, the map of $\text{int}(\Delta(Q))$ to the quotient is injective, so $\Delta(\Pi)/\langle f \rangle$ is a disjoint union of such interiors.

Suppose α is an f -invariant bi-infinite rectifiable path in $\Delta(\Pi)$. Then α is proper, and translated some positive distance, say l , by f . (Thus α maps to a closed path

of length l in $\Delta(\Pi)/\langle f \rangle$.) We will abuse notation slightly and write $\alpha \subseteq \Delta(\Pi)$, even if α is not injective.

For our discussion, there will be no loss in assuming that α is “piecewise cellular”. This means that we can write $\alpha = \bigcup_{i \in \mathbb{Z}} \alpha_i$, as a concatenation of paths, where $\alpha_i \subseteq \Delta(Q_i)$ for some $Q_i \in \mathcal{C}(\Pi)$. Then $f\alpha_i = \alpha_{m+i}$ for all i , for some fixed period, $m \in \mathbb{N}$. We write $\hat{\alpha} = \alpha_1 \cup \dots \cup \alpha_m$, so that $\alpha = \bigcup_{j \in \mathbb{Z}} f^j \hat{\alpha}$, and f translates α a distance equal to $\text{length}(\hat{\alpha})$.

Note that if $x \in \Delta(\Pi)$, then x lies in such a bi-infinite path: connect x to fx by a geodesic π , and set $\alpha = \bigcup_{j \in \mathbb{Z}} f^j \pi$. This will be piecewise cellular. (Note that by Proposition 17.1.1, $[x, fx]$ lies in a finite convex subcomplex of $\Delta(\Pi)$.)

Definition. Let Q be a finite cube. A **crossing arc** of $\Delta(Q)$ is a monotone arc, $\delta \subseteq \Delta(Q)$, whose initial and terminal points lie respectively in $\text{int}(\Delta(P))$ and $\text{int}(\Delta(P'))$, where P, P' are antipodal corank-1 faces of Q . (In other words, $\{P, P'\}$ is a wall of Q .) We say that δ is a **centrally crossing arc** if the initial and final points are the central elements of $\Delta(P)$ and $\Delta(P')$ respectively.

Suppose ρ is a median metric on $\Delta(Q)$. Then a crossing arc will have length at least the width, w , of the wall $\{P, P'\}$. The length of the centrally crossing arc is equal to w . (Note that if Q is a 1-cell, then the image of a crossing arc is precisely Q .)

We will show:

Proposition 21.2.6. *Let $f : \Pi \rightarrow \Pi$ be an isometry of a discrete median metric space (Π, ρ) , and let $f : \Delta(\Pi) \rightarrow \Delta(\Pi)$ be the induced cellular isometry. Suppose f has no fixed point in $\Delta(\Pi)$. Let $l \in S(f, \Delta(\Pi))$. Then there is an f -invariant bi-infinite cellular path, $\alpha = \bigcup_{i \in \mathbb{Z}} \alpha_i$, translated a distance at most l , and with each α_i a centrally crossing arc of some $Q_i \in \mathcal{C}(\Pi)$. Moreover the ranks of the Q_i are equal, and the corank-1 faces $Q_i \cap Q_{i+1}$ are all parallel in $\Delta(Q)$.*

Note that if $\rho = \rho_0$ is the standard combinatorial metric, then each centrally crossing arc has length 1, so the translation distance of α lies in \mathbb{N} . In view of Theorem 21.2.2 and Lemma 21.2.5, we obtain:

Corollary 21.2.7. *Let $f : \Pi \rightarrow \Pi$ be an automorphism of a discrete median algebra Π . Let $f : \Delta(\Pi) \rightarrow \Delta(\Pi)$ be the induced isometry with the standard median metric. Then $L(f, \Delta(\Pi)) \in \mathbb{N}$, and $R(f, \Delta(\Pi)) = R_0(f, \Delta(\Pi)) \neq \emptyset$.*

For the proof of Proposition 21.2.6, we will use the following lemma:

Lemma 21.2.8. *Let Q be a finite cube. Then any two points of $\partial\Delta(Q)$ lie in a monotone arc of $\Delta(Q)$ which is either a crossing arc, or else lies in $\partial\Delta(Q)$.*

Proof. Let $a, a' \in \partial\Delta(Q)$. Choose corank-1 faces, $P, P' \subseteq Q$, with $a \in \Delta(P)$ and $a' \in \Delta(P')$. If $P \cap P' \neq \emptyset$, then $F := P \cap P'$ is a corank-2 face of Q . Let b, b' be respectively the projections of a and a' to F . Then $a.b.b'.a$, and we can interpolate by monotone paths lying in P, F and P' . Their concatenation lies in $\partial\Delta(Q)$ as

required. If $P \cap P' = \emptyset$, then P, P' are antipodal faces. If $a \in \partial\Delta(P)$, then there is another corank-1 face, P'' , with $a \in P'$ and $P'' \cap P' \neq \emptyset$. We can therefore replace P by P'' and apply the above. We can therefore assume that $a \in \text{int}(\Delta(P))$. Similarly, we can assume that $a' \in \text{int}(\Delta(P'))$. Now any monotone arc from a to a' is a crossing arc. \square

Proof of Proposition 21.2.6. As observed earlier, we can start with a piecewise cellular path, $\beta = \bigcup_{i \in \mathbb{Z}} \beta_i$ with $\beta_i \subseteq \Delta(Q_i)$, and with $\text{length}(\hat{\beta}) \leq l$, where $\hat{\beta} = \beta_1 \cup \dots \cup \beta_m$. Let $\nu_i = \text{rank}(Q_i)$. Given $\nu \in \mathbb{N}$, let $N_\nu = \#\{i \in \{1, \dots, m\} \mid \nu_i = \nu\}$. The sequence $\underline{N} := (N_\nu)_{\nu \in \mathbb{N}}$ is eventually 0. We order such sequences antilexicographically. (That is, $\underline{N} < \underline{N}'$ if there is some ν with $N_\nu < N'_\nu$ and $N_\omega = N'_\omega$ for all $\omega > \nu$.) This set is well ordered, and we choose our path β as above so as to minimise \underline{N} in this order.

We first note that for each i , $Q_i \cap Q_{i+1}$ is a proper face of both Q_i and Q_{i+1} . For if not, up to swapping $i, i + 1$, we can assume that $Q_i \subseteq Q_{i+1}$. But now we can join the paths β_i and β_{i+1} and eliminate Q_i , thereby reducing \underline{N} in the antilexicographic order. (It is assumed that we do the same for all indices $i \in \mathbb{Z}$, modulo the period, m , so that process is equivariant with respect to f .)

It now follows that the endpoints of each β_i lie in $\partial\Delta(Q_i)$. If we could connect them by a path, $\beta'_i \subseteq \partial\Delta(Q_i)$, then we could replace β_i by β'_i , again reducing \underline{N} . Therefore, by Lemma 21.2.8, we can replace each β_i by a crossing path γ_i with the same endpoints. It now follows that $Q_i \cap Q_{i+1}$ is a common corank-1 face, P_i . Moreover, all the P_i are parallel, and so the Q_i all have the same rank. We can now finally replace each γ_i with the centrally crossing arc, α_i , from $\Delta(P_{i-1})$ to $\Delta(P_i)$. We do this equivariantly: that is $\alpha_{m+i} = \alpha_i$ for all i . Note that $\text{length}(\alpha_i) \leq \text{length}(\gamma_i) = \text{length}(\beta_i)$, so $\text{length}(\hat{\alpha}) \leq \text{length}(\hat{\beta}) \leq l$, as required. \square

21.3. Some results regarding more general median metric spaces.

In this subsection, we mostly consider median metric spaces which are either finite-rank or interval-compact. In the finite-rank case, one can often give stronger statements, as in [Fi4], for example, though these require more sophisticated arguments than those presented here.

We begin with:

Proposition 21.3.1. *Let M be a connected median metric space, which is either finite-rank or interval-compact. Then any semisimple isometry of M is simple.*

For the proof, we will use the following observation.

Lemma 21.3.2. *Let M be a median metric space, and suppose $a, b \in M$ are distinct. Suppose $f : M \rightarrow M$ is an isometry with $fa = b$ and $fb = a$, and such that $\rho(x, fx) \geq \rho(a, b)$ for all $x \in [a, b]$. Then $f|[a, b]$ is an involution of $[a, b]$, and $\rho(x, fx) = \rho(a, b)$ for all $x \in [a, b]$.*

Proof. Since $f(\{a, b\}) = \{a, b\}$, we certainly have $f([a, b]) = [a, b]$. We also have $\rho(x, fx) \leq \rho(x, a) + \rho(a, fx) = \rho(x, a) + \rho(b, x) = \rho(a, b) \leq \rho(x, fx)$, so we have

equality throughout. Thus $a \in [x, fx]$. Similarly, $b \in [x, fx]$. Thus, $xa \parallel b(fx)$. Applying this to fx , and swapping a, b , we get $(fx)b \parallel a(f^2x)$, so $ax \parallel a(f^2x)$, so $f^2x = x$. \square

Lemma 21.3.3. *Suppose M is a connected median metric space which is either finite-rank or interval-compact, and let $f : M \rightarrow M$ be an isometry. Suppose $a, b \in M$ are distinct, and $ab \parallel (fb)(fa)$. Then there is some $x \in [a, b]$ with $\rho(x, fx) < \rho(a, fa)$.*

Proof. Suppose, for contradiction that $\rho(x, fx) \geq \rho(a, fa)$ for all $x \in [a, b]$. Let $\tau : [a, b] \rightarrow [fa, fb]$ be the translation, and let $g = \tau^{-1} \circ f : [a, b] \rightarrow [a, b]$. Then $g|_{[a, b]}$ is a self-isometry of $[a, b]$, which swaps a, b .

Let $x \in [a, b]$. Since τ^{-1} is the gate map to $[a, b]$, we have $gx = \tau^{-1}(fx) \in [x, fx]$. Since it is a translation, we have $\rho(gx, fx) = \rho(b, fa)$. Thus $\rho(a, b) + \rho(b, fa) = \rho(a, fa) \leq \rho(x, fx) = \rho(x, gx) + \rho(gx, fx) = \rho(x, gx) + \rho(b, fa)$, so $\rho(a, b) \leq \rho(x, fx)$. By Lemma 21.3.2, $g|_{[a, b]}$ is an involution of $[a, b]$ with $\rho(x, gx) = \rho(a, b)$ for all $x \in [a, b]$.

We can now apply the construction in the proof of Lemma 21.1.4 to give us cubes, $Q_n \subseteq [a, b]$, with $\text{rank}(Q_n) \rightarrow \infty$. (This only required that $[a, b]$ be connected.) This immediately shows that M has infinite rank. It also shows that M cannot be interval-compact, as in the proof of Lemma 21.1.4. \square

Proof of Proposition 21.3.1. We are assuming that $R(f, M) \neq \emptyset$. Suppose, for contradiction, that $R_0(f, M) = \emptyset$. We now apply Lemma 21.2.3 (noting that this made no use of discreteness).

Let $x \in R(f, M)$ and $n \in \mathbb{N}$ be as chosen there, and set $p = x(f^{-1}x)(f^n x)$. By Lemma 21.2.3, if $y \in [x, p]$ then we have $\rho(y, f^n y) = \rho(x, f^n x)$, $y.x.f^n y$, $x.y.f^n x$ and $x.f^n y.f^n x$.

Applying this to p itself, we get $\rho(p, f^n p) = \rho(x, f^n x)$, $p.x.f^n p$ and $x.f^n p.f^n x$. By definition of p , we also have $x.p.f^n x$. Now $\rho(p, f^n p) \leq \rho(p, f^n x) + \rho(f^n x, f^n p) = \rho(p, f^n x) + \rho(x, p) = \rho(x, f^n x) = \rho(p, f^n p)$, so we have equality throughout, and so $p.f^n x.f^n p$. Thus $xp \parallel (f^n p)(f^n x)$. But now by Lemma 21.3.3 applied to f^n , there is some $z \in [x, p]$ with $\rho(z, f^n z) < \rho(x, f^n x)$. But as noted in the previous paragraph, we must have $\rho(z, f^n z) = \rho(x, f^n x)$, giving a contradiction. \square

In fact, for a finite-rank space, the semisimple assumption is redundant. It is shown in [Fi4] (Corollary A thereof) that:

Theorem 21.3.4. *(Fioravanti) An isometry of a finite-rank connected median metric space is simple.*

It turns out that, at least in some cases, $R_0(f, M)$ is a subalgebra of M . For example, let us say that a median metric space M is **sensible** if it is finite-rank, or discrete, or is the realisation, $\Delta(\Pi)$, of a discrete space Π . We claim:

Proposition 21.3.5. *Let $f : M \rightarrow M$ be an isometry of a sensible median metric space, M . Then $R_0(f, M)$ is a subalgebra of M .*

First we make some general observations. Suppose $W = (W^-, W^+)$ is an oriented wall of M . Let $(x_i)_{i \in \mathbb{Z}}$ be a bi-infinite monotone sequence (i.e. $x_i.x_j.x_k$ whenever $i \leq j \leq k$). We say that $(x_i)_i$ “crosses W positively” if there is some $i \in \mathbb{Z}$ with $x_i \in W^-$ and $x_{i+1} \in W^+$. Note that this implies that $x_j \in W^-$ for all $j \leq i$ and $x_j \in W^+$ for all $j > i$. We similarly define crossing “negatively”.

Suppose $f : M \rightarrow M$ is an isometry. Let $C^+(W)$ be the set of $x \in R_0(f, M)$ such that $(f^i x)_{i \in \mathbb{Z}}$ crosses W positively. We similarly define $C^-(W)$.

Suppose $x \in C^+(W)$ and $y \in C^-(W)$. For all sufficiently large $n \geq 0$, we have $f^n x \in C^+(W)$, $f^{-n} x \in C^-(W)$, $f^n y \in C^-(W)$ and $f^{-n} y \in C^+(W)$. Thus, for all but finitely many $n \in \mathbb{Z}$, we have $f^n W \in \mathcal{W}(x, y)$.

We also claim that if $m \neq n$ then $f^m W \not\cap f^n W$. It is enough to show that $W \not\cap f^n W$ for all $n > 0$. Suppose $f^n W^+ \subseteq W^+$. Now there is some i with $f^i y \in W^+$. But then $f^{i+nj} y \in W^+$ for all $j \geq 0$, contradicting $y \in C^-(W)$. Suppose $f^n W^+ \subseteq W^-$. There is some i with $f^i x \in W^+$. But then $f^{i+n} x \in W^-$ contradicting $x \in C^+(W)$. We get similar contradictions to $f^n W^- \subseteq W^-$ and $f^n W^- \subseteq W^+$. This shows that $W \not\cap f^n W$ as claimed.

Note that this cannot arise if M is of finite rank (by Lemma 8.2.1). Also, if M is discrete, then $\mathcal{W}(x, y)$ is finite, so there must be some n with $W = f^n W$, contradicting $W \not\cap f^n W$. Similarly if M is the realisation of a discrete space, then it is easily seen that only finitely many distinct images of any given wall can lie in $\mathcal{W}(x, y)$. In all such cases, therefore, no such pair x, y can exist.

We see:

Lemma 21.3.6. *If M is sensible and $f : M \rightarrow M$ is an isometry, then for all $W \in \mathcal{W}(M)$, either $C^+(W) = \emptyset$ or $C^-(W) = \emptyset$ (or both).*

We can now give:

Proof of Proposition 21.3.5. Suppose to the contrary that we have $a, b, c \in R_0(f, M)$, but $p := abc \notin R_0(f, M)$. We can suppose that $f^{-i} p.p.f^j p$ fails for some $i, j > 0$. For $x \in M$, write $x^- = f^{-i} x$ and $x^+ = f^j x$. Thus, $p^- = a^- b^- c^-$ and $p^+ = a^+ b^+ c^+$. Note that we have $a^- . a . a^+$, $b^- . b . b^+$, $c^- . c . c^+$, but not $p^- . p . p^+$. Choose any wall, $W \in \mathcal{W}(p, p^- p p^+)$, and orient W so that $p \in W^+$. Given $x \in M$ write $x \uparrow$ to mean that $x \in W^+$, and $x \downarrow$ to mean that $x \in W^-$. Thus $abc \uparrow$, $a^- b^- c^- \downarrow$ and $a^+ b^+ c^+ \downarrow$. Up to permuting a, b, c , we can suppose that $c^- \downarrow$ and $c^+ \downarrow$. By $c^- . c . c^+$, we have $c \downarrow$. By $abc \uparrow$ we have $a \uparrow$ and $b \uparrow$. By $a^- b^- c^- \downarrow$, up to swapping a, b , we have $a^- \downarrow$. By $a^- . a . a^+$ and $a \uparrow$, we have $a^+ \uparrow$. By $a^+ b^+ c^+ \downarrow$, we have $b^+ \downarrow$. Now since $a^- \downarrow$ and $a \uparrow$, we have $a \in C^+(W)$, and since $b \uparrow$ and $b^+ \downarrow$, we have $b \in C^-(W)$. This contradicts Lemma 21.3.6. □

As a particular example, one can give a relatively straightforward account of isometries of an \mathbb{R} -tree T (see Section 15).

Note that, as for any rank-1 median algebra, if $a, b, c, d \in T$, then $a.b.c \ \& \ b.c.d \Rightarrow a.b.c.d$. In particular, if $f : T \rightarrow T$ is an isometry, and $x \in T$ satisfies $x.f x.f^2 x$, then $f^i x.f^j x.f^k x$ whenever $i \leq j \leq k$. If $f x \neq x$, then $\alpha := \bigcup_{i \in \mathbb{Z}} [f^i x, f^{i+1} x]$ is

a bi-infinite geodesic path. This is closed and convex, and we have a gate map, $\omega : T \rightarrow \alpha$. Note that $f \circ \omega = \omega \circ f$. If $y \in T$, then $\rho(\omega y, f\omega y) \leq \rho(y, fy)$, with equality if and only if $y \in \alpha$. We see that $\alpha = R(f, T) = R_0(f, T)$. We refer to α as the “axis” of f .

If $\text{fix}(f) = \emptyset$, then such an axis always exists. To see this, choose any $x \in T$, and let $p = x(fx)(f^2x)$. Now $f, fp \in [fx, f^2x]$, so either $fx.p.fp.f^2x$ or $fx.fp.p.f^2x$ holds.

In the former case, $p.fp.f^2x \Rightarrow fp.f^2p.f^3x$. Now $x.p.f^2x \Rightarrow fx.fp.f^3x$, which together with $fx.p.fp$ gives $p.fp.f^3x$ by interpolation. Again, by interpolation, we have $p.fp.f^2p$, so T has an axis, as above.

In the latter case, we have $p, f^2p \in [fp, f^2x]$. Since $\rho(p, fp) = \rho(f^2p, fp)$ we have $f^2p = p$, and f fixes the midpoint of $[p, fp]$.

22. GATES

In this section, we give a discussion of the notion of “gated” sets in a very general context. These were introduced in the context of median algebras in Subsection 7.3. (Here, of course, the word “gate” is used in a quite different sense from that of Section 20.) All we need for this is a notion of “betweenness” satisfying some basic conditions which one might expect to hold for any notion worthy of this name. In particular, it applies to median algebras, and to submedian relations, as discussed in Section 14. Our main use for these ideas here will be to the discussion of quasimedial graphs in Section 23.

22.1. General betweenness axioms and gates.

Let M be a set equipped with a ternary relation, \dots . As before, we write $x_1.x_2.\dots.x_n$ to mean that $x_i.x_j.x_k$ holds whenever $i \leq j \leq k$. Given $a, b \in M$, write $[a, b] = \{x \in M \mid a.x.b\}$. We assume the following hold for $a, b, c, d \in M$.

(B1): $[a, a] = \{a\}$,

(B2): $[a, b] = [b, a]$,

(B3): $a \in [a, b]$.

(B4): If $c \in [a, b]$ and $d \in [c, b]$, then $d \in [a, b]$ and $c \in [a, d]$.

We refer to Property (B4) as “interpolation”: it can be expressed more succinctly as $a.c.b \ \& \ c.d.b \Rightarrow a.c.d.b$. Note that (by interpolation, together with (B1)) we have $a.b.c \ \& \ a.c.b \Rightarrow b = c$.

A map $[(a, b) \mapsto [a, b]]$ satisfying these axioms has been called a “geometric interval operator”. (See the Notes to this section.)

Given $a, b, c, d \in M$, write $c, d \preceq a, b$ to mean $a.c.d \ \& \ b.d.c$, and write $a, b \parallel c, d$ to mean $a, b \preceq c, d \ \& \ c, d \preceq a, b$. (In general, these relations need not be transitive, but they will be in the main cases of interest to us: for example, if the ternary

relation is submedian: see Lemma 23.5.1.) Note that $a, b \parallel c, d \Leftrightarrow a, c \parallel b, d$. In this case, we say that a, b and c, d are **parallel**.

A map $\tau : A \rightarrow B$ between subsets, $A, B \subseteq M$, is a **translation** if it is bijective and $a, b \parallel \tau a, \tau b$ for all $a, b \in A$. The inverse of a translation is a translation. Also, any translation from a set A to itself is the identity (since $x, \tau x \parallel \tau^{-1}x, x$, so $x.\tau x.x$, so $\tau x = x$). Also, if two parallel sets intersect, then they are equal.

If $A \subseteq M$ and $x \in M$, we say that $a \in A$ is a **gate** for x in A if $x.a.b$ holds for all $b \in A$. If such a gate exists, then it is unique. Note that, by interpolation, if $y \in [x, a]$ then a is also a gate for y in A . The following is also a simple observation:

Lemma 22.1.1. *Suppose that $B \subseteq A \subseteq M$. Suppose $x \in M$, $a \in A$ and $b \in B$, that a is a gate for x in M , and b is a gate for a in B . Then b is a gate for x in B .*

Proof. Let $c \in B$. We have $x.a.c$ and $a.b.c$, so $x.a.b.c$. In particular, $x.b.c$. □

Definition. We say that $A \subseteq M$ is **gated** if each point of M has a gate in A .

In other words, there is a map $\omega_A : M \rightarrow A$ such that $\omega_A(x)$ is a gate for x . Note that $\omega_A|_A$ is the identity. Also, $\omega_A x, \omega_A y \preceq x, y$ for all $x, y \in M$.

We will henceforth assume that $M \neq \emptyset$, so that any gated set is necessarily non-empty.

Definition. We say that $A \subseteq M$ is **convex** if $[a, b] \subseteq A$ for all $a, b \in A$.

Note that convexity is closed under intersection, so we can define the **convex hull** of any subset of M as the smallest convex set containing it.

Lemma 22.1.2. *A gated set is convex.*

Proof. Let A be gated, $a, b \in A$ and $x \in [a, b]$. Set $y = \omega_A x$. We have $a.x.b$ & $x.y.a$ & $x.y.b \Rightarrow a.y.x.y.b \Rightarrow x = y$, so $x \in A$. □

Lemma 22.1.3. *Suppose $A, B \subseteq M$ are gated and $A \cap B \neq \emptyset$. Then $\omega_A B = A \cap B$, and $\omega_A \omega_B : M \rightarrow A \cap B$ is a gate map.*

Proof. Clearly $A \cap B \subseteq \omega_A B$. Let $b \in B$ and set $a = \omega_A b$. If $c \in A \cap B$, then $b.a.c$. Since B is convex, $a \in B$. This shows that $\omega_A B \subseteq A \cap B$.

Now given any $x \in M$, let $b = \omega_B x$, $a = \omega_A b$ and let $c \in A \cap B$. We have $x.b.c$ & $b.a.c \Rightarrow x.b.a.c \Rightarrow x.a.c$. Thus, a is a gate for x in $A \cap B$. □

Lemma 22.1.4. *Let \mathcal{A} be any non-empty finite family of pairwise intersecting gated subsets of M . Then $\bigcap \mathcal{A}$ is gated (in particular, non-empty).*

Proof. (cf. Lemma 7.1.1) By induction, it is enough to check this for a 3-element family, $\mathcal{A} = \{A, B, C\}$. Let $x \in B \cap C$. By Lemma 22.1.3, $\omega_A x \in A \cap B$ and $\omega_A x \in A \cap C$, so $A \cap B \cap C \neq \emptyset$. By Lemma 22.1.3, $A \cap B \cap C$ is gated. □

Under certain conditions, this can be extended to infinite families: see for example, Lemmas 22.2.1 and 22.2.2.

Here is a variation on Lemma 22.1.3. Given any two gated subsets, $A, B \subseteq M$, write $A_B = \omega_A B \subseteq A$ and $B_A = \omega_B A \subseteq B$.

Lemma 22.1.5. $\omega_A \omega_B \omega_A : M \longrightarrow A_B$ is a gate map to A_B .

Proof. Let $x \in M$, $y = \omega_A x$, $b = \omega_B y$ and $a = \omega_A b$. Suppose $c \in A_B$; in other words, $c = \omega_A d$ for some $d \in B$. We have $x.y.c$, $y.a.b$, $y.b.d$ and $a.c.d$. Now $y.a.b \ \& \ y.b.d \Rightarrow y.a.b.d \Rightarrow y.a.d$, $y.a.d \ \& \ a.c.d \Rightarrow y.a.c.d \Rightarrow y.a.c$ and $x.y.c \ \& \ y.a.c \Rightarrow x.y.a.c \Rightarrow x.a.c$. This shows that a is a gate for A_B as required. \square

Now $\omega_A \omega_B \omega_A|_{A_B} = \omega_A \omega_B|_{A_B}$ is the gate map from A_B to itself, hence the identity. Similarly, $\omega_B \omega_A|_{B_A}$ is the identity on B_A . Also, $\omega_B A_B \subseteq B_A$ and $\omega_A B_A \subseteq A_B$. We see that $\omega_A|_{B_A}$ and $\omega_B|_{A_B}$ are inverse bijections. Moreover, if $x, y \in A_B$ then $\omega_B x, \omega_B y \preceq x, y$ and $x, y = \omega_A \omega_B x, \omega_A \omega_B y \preceq \omega_B x, \omega_B y$, and so $\omega_B x, \omega_B y \parallel x, y$. It follows that $\omega_B|_{A_B}$ is a translation of A_B to B_A . In summary, we have shown:

Lemma 22.1.6. If $A, B \subseteq M$ are gated, then $\omega_A|_{B_A}$ and $\omega_B|_{A_B}$ are inverse translations between A_B and B_A .

We note that all the above applies to a median algebra, and we have seen most of it in some form before. (Note that conclusion of Lemma 22.1.5 is slightly weaker than that of Lemma 7.3.3: see the Notes to this section.)

22.2. Applications to metric spaces and graphs.

These results also apply to metric spaces.

Let (M, ρ) be a metric space. As in Section 13, we write $a.x.b$ to mean $\rho(a, b) = \rho(a, x) + \rho(x, b)$. One readily checks axioms (B1)–(B4). Moreover, $c, d \preceq a, b \Rightarrow \rho(c, d) \leq \rho(a, b)$ and so $c, d \parallel a, b \Rightarrow \rho(c, d) = \rho(a, b)$. (See Lemmas 13.2.1 and 13.2.4 which only really used the fact that M is a metric space.) In particular, gate maps are 1-lipschitz and translations are isometries. Moreover, if $\tau : A \longrightarrow B$ is a translation, then there is some **translation distance**, $r \geq 0$, such that $\rho(x, \tau x) = r$ for all $x \in A$. We also note if a is a gate for x in A , then it is the unique $a \in A$ for which $\rho(x, a) = \rho(x, A)$. In particular, we see that gated sets are closed.

Lemma 22.2.1. Suppose M is complete, and that \mathcal{A} is a non-empty family of gated subsets with $\bigcap \mathcal{A} \neq \emptyset$. Then $\bigcap \mathcal{A}$ is gated.

Proof. In view of Lemma 22.1.4, we can suppose that \mathcal{A} is closed under finite intersection. Let $x \in M$ and let $r = \sup\{\rho(x, A) \mid A \in \mathcal{A}\}$. Let $(A_n)_{n \in \mathbb{N}}$ be a sequence in \mathcal{A} with $\rho(x, A_n) \rightarrow r$. Let $B_n = \bigcap_{m=0}^n A_m$ and let $b_n = \omega_{B_n} x$. Since $\rho(x, B_n) \geq \rho(x, A_n)$, we have $\rho(x, b_n) \rightarrow r$. Moreover, if $m \leq n$, then $x.b_m.b_n$, and so $(b_n)_n$ is monotone (i.e. $b_m.b_n.b_p$ holds whenever $m \leq n \leq p$). It follows that $(b_n)_n$ is Cauchy, and so $b_n \rightarrow b \in M$.

Let $C \in \mathcal{A}$, let $C_n = C \cap B_n$, and let $c_n = \omega_{C_n}x$. Then $x.b_n.c_n$, so $\rho(x, b_n) \leq \rho(x, c_n) \leq r$, so $\rho(b_n, c_n) \rightarrow 0$, so $c_n \rightarrow b$. Since C is closed, $b \in C$. This shows that $b \in \bigcap \mathcal{A}$.

Suppose that $d \in \bigcap \mathcal{A}$. Then $d \in B_n$ for all n , so $x.b_n.d$. Since $b_n \rightarrow b$, we get $x.b.d$. Thus b is a gate for x in $\bigcap \mathcal{A}$. \square

In view of Lemma 22.2.1, we can define the **gated hull** of any non-empty subset, $A \subseteq M$, as the intersection of all gated sets containing A . This is the unique smallest gated set containing A . By Lemma 22.1.2 this contains the convex hull.

One can give a number of variations on Lemma 22.2.1. For example, the following immediately implies Lemma 13.3.9.

Lemma 22.2.2. *Suppose M is complete, and that \mathcal{A} is a non-empty family of pairwise intersecting gated subsets of M , at least one of which is bounded. Then $\bigcap \mathcal{A}$ is gated (in particular, non-empty).*

(Note that we can assume that all sets of \mathcal{A} are bounded: replace the elements of \mathcal{A} with their respective intersections with some bounded element.) The proof is the same as that of Lemma 22.2.1: this only used the fact that $\rho(x, A)$ is bounded as A varies over \mathcal{A} .

We can apply the above to a non-empty connected graph, Γ , via its vertex set, $V(\Gamma)$, with combinatorial metric, ρ . In this context, a subgraph $G \subseteq \Gamma$ is **convex** if any geodesic path in Γ (including its edges) with both endpoints in G lies entirely in G . Such a graph is isometrically embedded. A subgraph, $G \leq \Gamma$ is **gated** if it is connected and isometrically embedded and $V(G)$ is gated in $V(\Gamma)$ in the sense already defined. Such a graph is necessarily convex. Both properties are closed under finite intersection.

Suppose $G \leq \Gamma$ is gated. Let $\omega = \omega_G : V(\Gamma) \rightarrow V(G)$ be the gate map. Recall that gate maps are 1-lipschitz. In particular, ω extends to a retraction, $\omega : \Gamma \rightarrow G$, sending every vertex to a vertex, and every edge to either a vertex or an edge (cf. Subsection 11.8).

It is also easily checked that if x, y are adjacent and $\omega x \neq \omega y$, then $\omega x, \omega y \parallel x, y$.

23. QUASIMEDIAN GRAPHS

In this section, we describe the notion of a “quasimedial graph”. This is a generalisation of a median graph. In geometrical terms, such a graph can be thought of as the 1-skeleton of a CAT(0) “prism complex”: just like a cube complex, where euclidean cubes are replaced by direct products of simplices. (In general, one needs to allow for infinite-dimensional simplices.) Quasimedial graphs were originally introduced in [Mul1]. They are explored in some detail in [Ge1]. The notion has found some diverse applications, for example, to group theory and to phylogenetics (see the Notes to this section).

Recall that a “clique” is a complete subgraph of a graph. In a quasimedial graph, Γ , maximal cliques play a key role. Any two distinct maximal cliques intersect,

if at all, in a single point. What we call “gated prisms” are cartesian products of maximal cliques, where the cliques are the “sides” of the prism. These are generalisations of the cells of a median graph (in that case, the maximal cliques are single edges, and prisms are cubical subgraphs). The gated prisms form a complex, whose dual graph, $\Theta(\Gamma)$, we show to be median (Proposition 23.4.11). (This was proven in [Ge1] by a different argument.) The vertex set of Γ can be identified as a subset of that of $\Theta(\Gamma)$ (namely those vertices which correspond to 0-cells of the prism complex). In particular, Γ is submedian. This fact can be used to deduce various other properties of Γ . We will finish the section by relating this to partitions of sets.

We will use the notation $G \leq \Gamma$ to mean that G is a subgraph of Γ . Recall that G is **full** if any two adjacent vertices of G are also adjacent in Γ .

Unless otherwise stated, all graphs in this section are assumed to be connected.

23.1. Definitions.

We begin with the definition.

Consider the following properties of a connected graph Γ , with combinatorial metric ρ .

(∇): Suppose $p, a, b \in V(\Gamma)$ with a, b adjacent, and with $\rho(p, a) = \rho(p, b)$. Then there is some $d \in V(\Gamma)$, adjacent to both a and b and with $\rho(p, d) = \rho(p, a) - 1$.

(\diamond): Suppose $p, a, b, c \in V(\Gamma)$ with a, b distinct, both adjacent to c , and with $\rho(p, a) = \rho(p, b) = \rho(p, c) - 1$. Then there is some $d \in V(\Gamma)$, adjacent to both a and b and with $\rho(p, d) = \rho(p, c) - 2$.

(Recall that (\diamond) was previously defined in Subsection 16.2.)

Definition. A graph Γ is **weakly modular** if it is connected and satisfies (∇) and (\diamond).

(Recall, that if we replace (∇) with the stronger condition that Γ is bipartite, then we arrive at a characterisation of a modular graph: Proposition 16.1.3.)

Weakly modular graphs were originally introduced by Chepoi, and are studied at length in [ChalCHO]. In particular the authors prove a local-to-global characterisation analogous to Theorem 16.2.3 here.

Quasimedial graphs are a special case of weakly modular graphs.

Recall that $K_{2,3}$ denotes the complete bipartite graph. We write $K_{1,1,2}$ for the complete graph K_4 with one edge removed (i.e. two triangles meeting along an edge).

Definition. A graph is **quasimedial** if it is weakly modular and contains no full subgraph isomorphic to either $K_{2,3}$ or $K_{1,1,2}$.

Note that (by Lemma 16.1.1) a graph is median if and only if it is quasimedial and bipartite. In fact:

Lemma 23.1.1. *A triangle-free quasimedial graph is median.*

Proof. We just need to check that it is bipartite. Suppose for contradiction that we have a cycle of odd length. Let p be any vertex. There must be two consecutive vertices, a, b , of the cycle with $\rho(p, a) = \rho(p, b)$. Property (∇) gives us a vertex, d , such that $\{a, b, d\}$ forms a triangle, giving a contradiction. \square

Simple examples of quasimedial graphs are block graphs (i.e. graphs for which every block is a complete subgraph). These were mentioned in Subsection 14.2 in relation to pretrees. They serve to illustrate some of the constructions below.

It turns out that quasimedial graphs can be described in terms of “quasimedial triples”:

Definition. Let Γ be a graph, and (x, y, z) be a triple of vertices. An *r -central triple* of (x, y, z) is a triple, (a, b, c) , of vertices with $x.a.b.y, y.b.c.z, z.c.a.x$ and such that $\rho(a, b) = \rho(b, c) = \rho(c, a) = r$. A *quasimedial triple* is an r -central triple with r minimal such that an r -central triple exists for $\{x, y, z\}$.

Here we are using the standard betweenness relation in a metric space, as defined in Subsection 13.1.

The following is a key property of quasimedial graphs [BanMuW]:

Theorem 23.1.2. *Any triple of vertices in a quasimedial graph has a unique quasimedial triple.*

We will give a more geometric proof of this statement in Subsection 23.6.

We remark that one can characterise quasimedial graphs in these terms as described in [BanMuW]. In the above notation, we can set $xyz = a$, to give us a ternary operation on the vertex set. We refer to xyz as the *quasimedial* of x, y, z . Note that the definition refers to ordered triples. It is invariant under swapping x with y . (Clearly, if $a = b = c$, then a is a median, and so this agrees with the ternary operation already defined in a median graph.) It turns out that one can write down axioms which characterise ternary operations which arise as quasimedians in this way. Such structures are referred to as “quasimedial algebras” in [BanMuW]. (We won’t make explicit use of these structures here.)

A good way to understand quasimedial graphs geometrically stems from the fact that the vertex set can be isometrically embedded in a discrete median algebra. This is a consequence of the construction of [Ge1]. We will give a self-contained proof of this below (Proposition 23.4.11), based on the characterisation of median graphs in terms of wheels (Theorem 16.2.3).

23.2. Examples.

Before embarking on the general theory, let us consider a few examples of quasimedial graphs.

(Ex23.1): We have already noted that a median graph is quasimedial. The quasimedial map is the median map in this case.

(Ex23.2): Any complete graph is quasimedial. In this case, the quasimedial map is the **dual discriminator** function. That is to say, xyz is equal to x if $y \neq z$, and equal to y if $y = z$.

(Ex23.3): Recall from Subsection 14.2 that a “block graph” is a connected graph, all of whose blocks are complete. As noted there, such a graph is quasimedial.

(Ex23.4): Recall that the **cartesian product**, $G \square H$, of two graphs, G and H , is the 1-skeleton of the square complex $G \times H$. It is not hard to check that if G, H are quasimedial, then so is $G \square H$. The quasimedial map is determined separately on each coordinate of $V(G \square H) \equiv V(G) \times V(H)$. Of course, this extends to cartesian products of finitely many graphs.

(Ex23.5): In particular, a **prism** is a cartesian product, $\Pi = \Lambda_1 \square \Lambda_2 \square \cdots \square \Lambda_n$, of finitely many complete graphs, Λ_i . (These are allowed to be infinite.) We will generally assume that none of the Λ_i is a singleton. In this case, we refer to n as the **rank** of Π . Note that the combinatorial distance between two vertices of Π is just the number of coordinates, i , on which they differ. In this case, the quasimedial map is easy to understand. If (a, b, c) is the quasimedial triple for (x, y, z) , then for each i , the three coordinates a_i, b_i, c_i are either all equal or all distinct. The number of i for which they are distinct is thus equal to $\rho(a, b) = \rho(b, c) = \rho(c, a)$. The quasimedial operation, $[(x, y, z) \mapsto xyz]$, is the dual discriminator independently on each factor.

It is not hard to see that any automorphism of Π preserves the product structure up to permutation of the factors together with automorphisms of each of the factors. In fact, the factors can be recovered as the set of parallel classes of maximal cliques, where parallelism is defined as in Subsection 22.1. (Recall that a **clique** of a graph is a complete subgraph.) In view of this, one can unambiguously define a **face** of Π to be a subgraph of the form $\Lambda'_1 \square \Lambda'_2 \square \cdots \square \Lambda'_n$, where for each i either $\Lambda'_i = \Lambda_i$ or Λ'_i is a singleton. We refer to a rank-1 face as a **side** of Π . As noted, these are precisely the maximal cliques of Π .

It is not hard to check (cf. Lemma 7.5.1) that a subgraph of Π is convex in the sense of Subsection 22.1 if and only if it has the form $\Lambda'_1 \square \Lambda'_2 \square \cdots \square \Lambda'_n$ where $\Lambda'_i \subseteq \Lambda_i$. A subgraph of Π is gated if and only if it is a face.

Let $\Lambda \leq \Pi$ be a side of Π corresponding to the factor Λ_i . Then $V(\Lambda)$ consists of those elements of $V(\Pi)$ whose j th coordinate is fixed for all $j \neq i$. To obtain the gate in Λ for a given $x \in V(\Pi)$, we leave the coordinate x_i alone, and replace each coordinate x_j for $j \neq i$ by that determined by Λ .

(Ex23.6): Let I be a finite indexing set, and let $\Pi = \square_{i \in I} \Lambda_i$ be a prism. Given a subset $B \subseteq V(\Pi)$ we can define an equivalence relation, \sim_B , on I , by writing $i \sim_B j$ if, for all $x, y \in B$, it holds that $x_i = y_i \Leftrightarrow x_j = y_j$. We say subset $B \subseteq V(\Pi)$ is a **subalgebra** if it is closed under the quasimedial operation on $V(\Pi)$. Given any $B \subseteq V(\Pi)$, we let $A := \langle B \rangle$ be the smallest subalgebra containing B . This is sometimes called the **quasimedial hull**. It is easily seen that $\sim_A = \sim_B$.

Given a subalgebra $A \subseteq V(\Pi)$, we can delete all but one of the indices from any \sim_A -class (or equivalently, just replace the indices by the \sim_A -classes thereof). Note that A remains imbedded after projecting to the face of Π corresponding to the remaining indices, and that the quasimedial operation on A does not change. Thus, there is no loss in assuming that \sim_A is just equality. In other words, if $i \neq j$, then there exist $x, y \in A$ with $x_i = y_i$ and $x_j \neq y_j$.

In this case, we claim that if $a, b \in A$ with $\rho_\Pi(a, b) \geq 2$, then there is some $c \in A$ with $a.c.b$ and $c \neq a, b$. To see this, choose distinct i, j with $a_i \neq b_i$ and $a_j \neq b_j$. Let $x, y \in A$ be as above. Up to swapping a, b , we can assume that $a_i \neq x_i = y_i$. Let $d = axy \in A$. Then $d_i = x_i = y_i \neq a_i$ and $d_j = a_j$. Let $c = bad \in A$ be the quasimedial. Then $a.c.b$. Also $c_i = b_i$ and $c_j = a_j$. Thus $c \notin \{a, b\}$ as required. After repeated interpolation, we can now find a 1-path $a = a_0, a_1, \dots, a_m = b$ from a to b with $a_0.a_1 \cdots .a_m$. This shows that the full subgraph, Γ , of Π with vertex set $V(\Gamma) = A$ is connected and isometrically embedded in Π .

We can now easily check that Γ is a quasimedial graph which induces the given quasimedial operation. Since Π does not contain a full $K_{1,1,2}$ or $K_{2,3}$, neither does Γ . Properties (∇) and (\diamond) follow by setting $d = pab$ in both cases. Since betweenness in Γ agrees with betweenness on Π , it follows that the quasimedial operations agree.

One can always make the following simplifying assumption. There is no loss in assuming that the projection of B , or equivalently $A = \langle B \rangle$, to each factor $V(\Lambda_i)$ is surjective. To do this we just replace Π by the convex hull of B in Π . This is intrinsically a prism. (This corresponds to deleting those vertices of Λ_i which do not occur as the i th coordinate of any element of B , and then removing those factors, Λ_i , which are reduced to singletons.) In this case, the maximal cliques of Γ will all be sides of Π . (For suppose $a, b \in A = V(\Gamma)$ are adjacent. Let Λ be the side of Π containing a, b , and let $c \in V(\Lambda)$. Now the parallel class of Λ corresponds to a factor Λ_i of Π . By assumption, there is some $d \in A$ with $d_i = c_i$. Then $c = dab \in A$. This shows that $\Lambda \leq \Gamma$.) Note also that any two such maximal cliques are parallel in Γ if and only if they are parallel in Π , hence correspond to the same factor of Π .

(Ex23.7): Using (Ex23.6), we can associate a quasimedial graph to a family of partitions of a set in the following way. Let X be a set, and let $\mathcal{W} = (W_i)_{i \in I}$ be a family of non-trivial partitions of X , indexed by some finite set, I . Let Λ_i be the complete graph with vertex set W_i . Let Π be the prism $\Pi := \square_{i \in I} \Lambda_i$. There is a natural map $\eta : X \rightarrow V(\Pi)$. Here, the i th coordinate of $\eta(x)$ is defined

to be the element of the partition of W_i which contains x . Let us suppose that for all distinct x, y there are indices i, j such that x, y lie in the same element of W_i , but in different elements of W_j (as can be achieved by deleting redundant indices, as noted in (Ex23.6)). The subalgebra of $V(\Pi)$ generated by $\eta(X)$ is then naturally the vertex set of a quasimedial graph. This example has applications to phylogenetics: see the Notes to this Section.

23.3. Gates and prisms.

We now set about describing the properties of a general quasimedial graph, Γ .

First note that any non-trivial clique is contained in a unique maximal clique. A clique is gated if and only if it is maximal. Two maximal cliques intersect in at most one vertex. These observations all follow easily from the fact that any embedded $K_{1,1,2}$ is contained in a clique K_4 .

Together with the fact that there is no full $K_{2,3}$, one can also easily see that any subgraph of Γ isomorphic to $K_{2,3}$ is contained in a clique K_5 .

By a *square* in Γ we mean a full subgraph isomorphic to a 4-cycle. By the above, there is at most one way of “completing” a square in the following sense: if a, c, b, d and a, c, b, e are both squares (with a, b antipodal) then $d = e$. If such a d exists given a, b, c , we write $d = a[c]b$. In particular, in Axiom (\diamond) above, $d = a[c]b$ is unique.

In fact, the same is true of Axiom (∇) . For suppose d, d' both satisfy the conclusion for given p, a, b . Now d, d' must be adjacent, so applying (∇) to p, d, d' there is some $e \in V(\Gamma)$ adjacent to both d and d' , with $\rho(p, e) = \rho(p, a) - 2$. But now a and e are adjacent (indeed, a, b, d, d', e all lie in a K_5) and we get a contradiction.

In applying (∇) or (\diamond) , we write $d = pab$. This is the unique median of p, a, b .

Lemma 23.3.1. *Suppose $p, a, b, c \in V(\Gamma)$, with b, c adjacent to a , with $\rho(b, c) = 2$ and with $\rho(p, a) = \rho(p, b) = \rho(p, c) + 1$. Then there is some $d \in V(\Gamma)$ adjacent to both b and c and with $\rho(p, d) = \rho(p, c)$.*

Proof. Apply (∇) to give us $e := pab$. Now c cannot be adjacent to e (otherwise c, a, e, b would be a full $K_{1,1,2}$). Applying (\diamond) gives us $f := pce$. Applying (∇) to b, c, f we arrive at $d := bcf$. \square

We have following criterion for convexity (as defined in Subsection 22.1), cf. Lemma 11.4.4.

Lemma 23.3.2. *A subgraph, $G \leq \Gamma$, is convex if and only if it is connected and any triangle or square in Γ with at least two adjacent edges in G lies entirely in G .*

Proof. The “only if” direction is clear.

For the converse, we want to show that any geodesic in Γ connecting any two vertices $a, b \in V(G)$ lies entirely in G . We prove this by induction on $\rho(a, b)$.

First, let $a = a_0, a_1, \dots, a_n = b$ be any intrinsic geodesic in G from a to b (so that $\rho_G(a, b) = n$). We claim that this is also geodesic in Γ , i.e. $\rho(a, b) = n$. For

suppose not. Let m be minimal such that $\rho(a, a_{m+1}) \leq m$. We will construct a path $b_0, b_1, \dots, b_m = a_{m+1}$ in G , with b_i adjacent to a_i in G , and with $\rho(a, b_i) \leq i$ for all i . Note that it follows that $\rho_G(a, b_i) = i + 1$ (since $i + 1 = (m + 1) - (m - i) \leq \rho_G(a_0, b_m) - \rho_G(b_m, b_i) \leq \rho_G(a_0, b_m) \leq \rho_G(a_0, a_i) + \rho_G(a_i, b_i) \leq i + 1$). In particular, for $i > 0$, b_i is not adjacent to a_{i-1} in G , and so by the triangle hypothesis, these are not adjacent in Γ either.

To construct the path $(b_i)_i$, there are two cases to consider:

Case (1): $\rho(a, a_{m+1}) = m$.

We set $b_m = a_{m+1}$ and proceed by backward induction. By the triangle hypothesis on G , we must have $\rho(a_{m-1}, b_m) = 2$. Applying Lemma 23.3.1 to a, a_{m-1}, a_m, b_m , we get $b_{m-1} := a_{m-1}[a_m]b_m$, with $\rho(a, b_{m-1}) = m - 1$. Since the edges $a_{m-1}a_m$ and $a_m b_m$ both lie in G , the square hypothesis on G tells us that b_{m-1} and the edges $b_{m-1}a_{m-1}$ and $b_{m-1}b_m$ also lie in G . As noted above, a_{m-2} and b_{m-1} are not adjacent in Γ . We can now apply Lemma 23.3.1 to $a, a_{m-2}, a_{m-1}, b_{m-1}$ to give us $b_{m-2} := a_{m-2}[a_{m-1}]b_{m-1}$. We now continue inductively. We end up with b_0 adjacent to $a = a_0$, but with $\rho(a, b_0) = 0$, which is clearly a contradiction.

Case (2): $\rho(a, a_{m+1}) = m - 1$.

We again set $b_m = a_{m+1}$. This time, we apply (\diamond) to a, a_{m-1}, a_m, b_m to give us $b_{m-1} := a_{m-1}[a_m]b_m = aa_{m-1}b_m$. Then $\rho(a, b_{m-1}) = m - 2$. From the square condition again, we see that b_{m-1} and the edges b_{m-1}, a_{m-1} and $b_{m-1}b_m$ lie in G as before. We now continue by backward induction. This time we get $\rho(a, b_0) = -1$: an even more shocking contradiction.

This proves the claim that $\rho(a, b) = n$. In particular, a, b are connected by at least one geodesic, $a = a_0, a_1, \dots, a_n = b$ which lies entirely in G . By the inductive hypothesis on $\rho(a, b)$, any geodesic in Γ from a_1 to b also lies in G .

Now suppose that $a = a'_0, a'_1, \dots, a'_n = b$ is any geodesic in Γ from a to b . We want to show that this lies in G . By the inductive hypothesis again, it is enough to show that the edge aa'_1 lies in G . We can suppose that $a'_1 \neq a_1$. We now apply (\diamond) to b, a_1, a'_1 to give us $c := a_1[a]a'_1 = ba_1a'_1$. The edge a_1c lies in some geodesic from a_1 to b , hence lies in G . So does aa_1 . Therefore the square hypothesis on G tells us in particular, that the edge aa'_1 lies in G as required. \square

We also have a criterion for gated sets:

Lemma 23.3.3. *A subgraph, $G \leq \Gamma$, is gated if and only if it is convex and every triangle in Γ with at least one edge in G lies entirely in G .*

Proof. The “only if” part follows by Lemma 22.1.2, and the fact that gated graphs are full.

For the converse, let $p \in \Gamma$, and let $a \in V(G)$ be a vertex with $\rho(p, a) = \rho(p, G)$. We claim that a is a gate for p in G .

To this end let $b \in V(G)$. Let $r = \rho(p, G)$, and let $a = a_0, a_1, \dots, a_n$ be a geodesic in Γ from a to b . Since G is convex, this lies in G . Suppose for contradiction that $p.a.b$ does not hold. In other words, $\rho(p, b) < r + n$. Let m be minimal such that $\rho(p, a_{m+1}) \leq m$. We now construct a sequence, $b_0, b_1, \dots, b_m = a_{m+1}$,

with a_i adjacent to b_i and with $\rho(p, b_i) \leq \rho(p, a_i)$ similarly as in the proof of Lemma 23.3.2. This is again by reverse induction, using p in place of a . Note that $a = a_0, b_0, b_1, \dots, b_m = a_{m+1}$ is a path of length $m + 1$ in Γ from a to a_{m+1} , hence geodesic. By convexity, this lies in G . In particular, the edge ab_0 lies in G . Now $\rho(p, b_0) \leq \rho(p, a)$, so by minimality of $\rho(p, a)$, we have $\rho(p, a) = \rho(p, b_0) = r$. We now apply (∇) to give us $c := pab_0$. By the triangle hypothesis, $c \in V(G)$. But $\rho(p, c) = \rho(a, b) - 1$ contradicting minimality. We therefore must have $p.a.b$ as required. \square

In the following discussion, it will be convenient to work in terms of “submersions” of graphs, so that we don’t have to worry about them being (isometrically) embedded. In view of Lemma 23.3.4 below, that will retrospectively be automatic in the cases of interest to us.

Let G be a graph. A **morphism** to Γ is a map from G to Γ which sends vertices to vertices and edges to edges. We define an **submersion** to be a morphism which sends squares in G to squares in Γ .

Lemma 23.3.4. *Suppose that G is connected, and has the property that any two vertices distance 2 apart lie in a unique square in G . Suppose that $f : G \rightarrow \Gamma$ is a submersion. Then f is an isometric embedding, and $f(G)$ is convex in Γ .*

Proof. The proof follows that of Lemma 23.3.2. In fact, in view of Lemma 23.3.2, it is sufficient to show that f is an isometric embedding. To this end, let $a, b \in V(G)$, and let $a = a_0, a_1, \dots, a_n = b$ be a geodesic in G from a to b . Suppose for contradiction that $\rho(fa, fb) < n$. Let m be minimal such that $\rho(fa, fa_{m+1}) \leq m$. We now construct a path $a = b_0, b_1, \dots, b_m = a_{m+1}$, in G similarly as with Lemma 23.3.2. Again, b_i is adjacent to a_i and this time, $\rho(fa, fb_i) \leq i$. We begin by setting $b_m = a_{m+1}$. By applying Lemma 23.3.1 or (\diamond) , we can set $d = fa_{m-1}[fa_m]fb_m \in V(\Gamma)$. Since $\rho_G(a_{m-1}, b_m) = 2$, by hypothesis, there is a unique $b_{m-1} \in V(G)$ such that $a_{m-1}, a_m, b_m, b_{m-1}$ is a square in G . Since f is a morphism, $fa_{m-1}, fa_m, fb_m, fb_{m-1}$ is a square in Γ , and so $fb_{m-1} = d$. Therefore, $\rho(fa, fb_{m-1}) \leq m - 1$. We now proceed inductively. We end up with $b \in V(G)$ adjacent to $a = a_0$, with $\rho(fa, fb_0) \leq 0$, contradicting the assumption that f is a morphism. \square

Note that the above applies to any prism (as defined in Example (Ex23.5) above). By a **prism** in Γ we mean a subgraph isomorphic to a prism. We see that such a subgraph is necessarily isometrically embedded and convex. We also note:

Lemma 23.3.5. *A prism $\Pi \leq \Gamma$ is gated if and only if each side of Π is a maximal clique of Γ .*

Proof. Note that every edge of Γ lies in a unique maximal clique of Γ . We see that every side of Π is a maximal clique of Γ if and only if every triangle of Γ with at least one side in Π lies entirely in Π . We now apply Lemma 23.3.3. \square

We refer to such a subgraph as **gated prism**.

Note that any face of a gated prism is a gated prism. Also the intersection of any two gated prisms is a common face of each. This follows from the fact that the intersection of any two gated subgraphs is gated (Lemma 22.1.3), and the fact that a gated subgraph of a prism is a face.

We will eventually see (Corollary 23.4.7) that any prism in Γ lies in a unique gated prism of the same rank.

23.4. The prism complex.

We can associate to any quasimedial graph a canonical “prism complex” defined as follows.

Let $\Theta = \Theta(\Gamma)$ be the graph with vertex set identified with the set of all gated prisms in Γ . Given $a \in V(\Theta)$ we write $\Pi(a) \leq \Gamma$ for the corresponding gated prism. Two vertices, $a, b \in V(\Theta)$ are deemed adjacent if $\Pi(a)$ is a corank-1 face of $\Pi(b)$ or conversely. Note that we can identify a vertex $x \in V(\Gamma)$ with the rank-0 prism, $\{x\}$.

Let Σ be the binary subdivision of the graph Γ . In other words, a vertex of Σ is either a vertex of Γ , or else the midpoint of an edge of Γ . Giving Σ the combinatorial metric, ρ_Σ , we have $\rho_\Sigma(x, y) = 2\rho_\Gamma(x, y)$ for all $x, y \in V(\Gamma)$. Now each edge of Γ lies in a unique maximal clique of Γ : that is a rank-1 element of $V(\Theta)$. In this way, we can define a morphism $\theta : \Sigma \rightarrow \Theta$, which is the identity on Γ (under the identification of x with $\{x\}$ mentioned above). Note that $\rho_\Theta(\theta x, \theta y) \leq \rho_\Sigma(x, y)$ for all $x, y \in V(\Theta)$. In fact, we have:

Lemma 23.4.1. $\rho_\Theta(x, y) = \rho_\Sigma(x, y) = 2\rho_\Gamma(x, y)$ for all $x, y \in V(\Gamma)$.

Proof. Let $x, y \in V(\Gamma)$, so that $\theta x = x$ and $\theta y = y$. Let $x = a_0, a_1, \dots, a_n = y$ be a geodesic from x to y in Θ . Thus $\rho_\Theta(x, y) = n$. We want to show that $\rho_\Sigma(x, y) \leq n$. We write $\Pi_i = \Pi(a_i)$ and $r_i = \text{rank}(\Pi_i)$.

We claim that $2\rho_\Gamma(x, \Pi_i) + r_i \leq i$ for all i . This follows by induction. The induction starts since $r_0 = 0$. There are two cases for the inductive step. If $\Pi_i \leq \Pi_{i+1}$, then $\rho_\Gamma(x, \Pi_{i+1}) \leq \rho_\Gamma(x, \Pi_i)$ and $r_{i+1} = r_i + 1$, and so $2\rho_\Gamma(x, \Pi_{i+1}) + r_{i+1} \leq i + 1$ as required. If $\Pi_{i+1} \leq \Pi_i$, then Π_{i+1} is a corank-1 face of Π_i , and each point of Π_{i+1} is adjacent in Γ to a point of Π_i . Thus, $\rho_\Gamma(x, \Pi_{i+1}) \leq \rho_\Gamma(x, \Pi_i) + 1$. Also, $r_{i+1} = r_i - 1$, so again the inductive step follows. When $i = n$, we conclude that $2\rho_\Gamma(x, y) \leq \rho_\Sigma(x, y)$.

We have shown that $\rho_\Sigma(x, y) \leq \rho_\Theta(x, y)$. The reverse inequality is immediate from the fact that θ is a morphism, as observed above. □

(For future reference — see Lemma 23.6.2 — we note that the argument shows the following. Suppose $x \in \Gamma$ and $y \in \Theta$ then $2\rho_\Gamma(x, \Pi(y)) + \text{rank}(\Pi(y)) \leq \rho_\Theta(x, y)$. To see this, take a geodesic in Θ from x to y , and argue as before.)

Given $a \in V(\Theta)$, we abbreviate $\text{rank}(a) = \text{rank}(\Pi(a))$. The following is also easily verified.

Lemma 23.4.2. *If $a \in V(\Theta)$, then $\text{rank}(a) = \rho_\Theta(a, \theta(V(\Gamma)))$. Moreover, if $x \in V(\Gamma)$, then $x \in \Pi(a)$ if and only if $\rho_\Theta(\theta(x), a) = \text{rank}(a)$.*

The main aim now is to show that Θ is a median graph (Proposition 23.4.11 below).

For the following discussion we will refer to the prism $K_2 \square K_3$ as a **roof**. It is easily seen that if $R \leq \Gamma$ is a roof, then it is either convex or else contained in a clique $K_6 \leq \Gamma$. The full subgraph, R^- , on any five vertices of R consists of a triangle and a square meeting along an edge. It is easily checked that any submersion of R^- into Γ extends uniquely to a submersion of R . By Lemma 23.3.4, this is an embedding and its image is convex.

Let Q be the 3-cubical graph, $K_2 \square K_2 \square K_2$. Let $Q^- \subseteq Q$ be the full subgraph on any seven vertices of Q . In the terminology of Subsection 16.2, Q^- is a “wheel”. It consists of three squares meeting pairwise in single edges, and intersecting at a common vertex: the “hub” of the wheel. Again it is easily checked that any submersion of Q^- into Γ extends uniquely to a submersion of Q . By Lemma 23.3.4, this is an embedding and its image is convex.

These observations will be routinely used in the arguments below.

Let $\Pi = \Lambda_1 \square \Lambda_2 \square \cdots \square \Lambda_n$ be a prism. Choose some basepoint, $p \in V(\Pi)$. Given $x \in V(\Pi)$, we write x_i for its coordinate in $V(\Lambda_i)$. We write

$$h(x) := \rho_\Pi(p, x) = \#\{i \mid x_i \neq p_i\}.$$

Let P_m be the full subgraph of Π with vertex set $V(P_m) = \{x \in V(\Pi) \mid h(x) \leq m\}$.

In the following arguments, we use the following conventions. Given subsets, $A_i \subseteq V(\Lambda_i)$, we write (A_1, A_2, \dots, A_n) for $A_1 \times A_2 \times \cdots \times A_n \subseteq V(\Pi)$. In this context, we will abbreviate $\{x_i\}$ to x_i . We will abbreviate $(A_1, \dots, A_m, x_{m+1}, \dots, x_n)$ to $(A_1, \dots, A_m, -)$ if x_{m+1}, \dots, x_n are fixed and not directly relevant to the argument. For example, in this notation, if $x_1 \neq y_1$ and $x_2 \neq y_2$, then $(\{x_1, y_1\}, \{x_2, y_2\}, -)$ denotes (the vertex set of) a square in Π .

By Lemma 23.3.4, we know retrospectively that all the maps referred to in the following discussion are isometric embeddings. However, it is logically simpler to proceed without worrying about that. We therefore refer instead to submersions.

The next two results are analogues of Lemmas 10.3.6 and 10.3.7 respectively. (They overlap in the context of median graphs.)

Lemma 23.4.3. *Suppose that $f, f' : \Pi \rightarrow \Gamma$ are submersions, and that $f|_{P_1} = f'|_{P_1}$. Then $f = f'$.*

Proof. Suppose the conclusion fails. Choose $x \in V(\Pi)$ with $h(x)$ minimal such that $f(x) \neq f'(x)$. By hypothesis, $h(x) \geq 2$. Up to permuting the coordinates, we can write $x = (x_1, x_2, -)$, with $x_1 \neq p_1$ and $x_2 \neq p_2$. Consider the square σ with vertices $(\{p_1, x_1\}, \{p_2, x_2\}, -)$. By minimality of $h(x)$, f and f' agree on the three vertices $(p_1, p_2, -)$, $(p_1, x_2, -)$ and $(x_1, p_2, -)$. Since they are submersions, they agree on σ , so we get the contradiction that $f(x) = f'(x)$. \square

Lemma 23.4.4. *Suppose that $f : V(P_2) \rightarrow \Gamma$ is a map which sends every square or triangle in Π containing p to a square or triangle in Γ . Then f extends to a submersion $f : \Pi \rightarrow \Gamma$.*

Proof. For the proof, if we have f defined on a subset of $V(\Pi)$ we will say that this subset is “submersed” if f extends to a submersion on the full subgraph of Π with that vertex set.

By hypothesis, we have f already defined on $V(P_2)$, and we first check that $V(P_2)$ is submersed. By hypothesis, every triangle of the form $(\{p_1, x_1, y_1\}, p_2, p_3, \dots, p_n)$ and every square of the form $(\{p_1, x_1\}, \{p_2, x_2\}, p_3, \dots, p_n)$ is submersed, and similarly permuting the indices, i .

Now consider the square $(\{x_1, y_1\}, \{p_2, x_2\}, p_3, \dots, p_n)$, where $x_1, y_1 \neq p_1$. (It is implicitly assumed that $x_1 \neq y_1$ and $p_2 \neq x_2$.) This lies in the roof with vertices $(\{p_1, x_1, y_1\}, \{p_2, x_2\}, -)$, which contains the triangle $(\{p_1, x_1, y_1\}, p_2, -)$ and the squares $(\{p_1, x_1\}, \{p_2, x_2\}, -)$ and $(\{p_1, y_1\}, \{p_2, x_2\}, -)$, all of which meet at p . These triangles and squares are submersed by hypothesis, and so the roof is submersed. In particular, the square $(\{x_1, y_1\}, \{p_2, x_2\}, p_3, \dots, p_n)$, is submersed, and similarly permuting indices.

Now consider the square $(\{x_1, y_1\}, \{x_2, y_2\}, p_3, \dots, p_n)$, where $x_1, y_1 \neq p_1$ and $x_2, y_2 \neq p_2$. This lies in the roof $(\{x_1, y_1\}, \{p_2, x_2, y_2\}, -)$, which contains the triangle $(x_1, \{p_2, x_2, y_2\}, -)$ and the squares $(\{x_1, y_1\}, \{p_2, x_2\}, -)$ and $(\{x_1, y_1\}, \{p_2, y_2\}, -)$ (all meeting at $(x_1, p_2, -)$). The squares are submersed by the previous paragraph, so the roof is submersed, and in particular the square $(\{x_1, y_1\}, \{x_2, y_2\}, p_3, \dots, p_n)$ is submersed.

After permuting indices, this accounts for all the squares of P_2 . Moreover, every edge of P_2 lies in such a square. It follows that $V(P_2)$ is submersed as claimed.

We now extend f inductively over $V(P_m)$ and show that these sets are submersed. Let us assume that $V(P_{m-1})$ is submersed for $m \geq 3$. We want to define f on $V(P_m)$. Let $x \in V(P_m) \setminus V(P_{m-1})$, i.e. $h(x) = m$. After permuting indices, we can suppose that $x = (x_1, \dots, x_m, p_{m+1}, \dots, p_n)$ where $x_i \neq p_i$ for $i \leq m$.

Consider the 3-cubical graph with vertex set $(\{p_1, x_1\}, \{p_2, x_2\}, \{p_3, x_3\}, -)$. This has square faces $(p_1, \{p_2, x_2\}, \{p_3, x_3\}, -)$, $(\{p_1, x_1\}, p_2, \{p_3, x_3\}, -)$ and $(\{p_1, x_1\}, \{p_2, x_2\}, p_3, -)$. These all lie in P_{m-1} and meet at the vertex $(p_1, p_2, p_3, -)$. They form a wheel, which is submersed by hypothesis. The submersion of the wheel extends to a submersion of cube, which sends $x = (x_1, x_2, x_3, -)$ to some vertex, $y \in V(\Gamma)$.

If $m = 3$, then y is canonically determined, and we set $f(x) = y$.

Suppose $m \geq 4$. Consider the cube $(\{p_1, x_1\}, \{p_2, x_2\}, x_3, \{p_4, x_4\}, -)$. Similarly as before, this cube is submersed on sending x to some vertex $z \in V(\Gamma)$. These two cubes meet in the square $(\{p_1, x_1\}, \{p_2, x_2\}, x_3, x_4, -)$. This contains three vertices, $(p_1, p_2, x_3, x_4, -)$, $(p_1, x_2, x_3, x_4, -)$ and $(x_1, p_2, x_3, x_4, -)$ in $V(P_{m-1})$, where f is already defined. Therefore there is a unique extension to a submersion of the square. In other words, we must have $y = z$. Continuing in this manner, we see that any

cube of the form $(x_1, \dots, \{p_i, x_i\}, \dots, \{p_j, x_j\}, \dots, \{p_k, x_k\}, \dots, x_m, p_{m+1}, \dots, p_n)$ is submersed on setting $f(x) = y$.

After permuting indices, we can do this for all $x \in V(P_m) \setminus V(P_{m-1})$. This gives us a map $f : V(P_m) \rightarrow V(\Gamma)$. We need to check that $V(P_m)$ is submersed. Since every edge of P_m lies in a square in P_m , it is enough to check that every square in P_m is submersed.

To this end, given a square $\sigma \leq \Pi$, write $h(\sigma) = \sum_{x \in V(\sigma)} h(x)$. Suppose, for contradiction that P_m is not submersed. Let $\sigma \leq P_m$ be a square that is not submersed with $h(\sigma)$ minimal. By the inductive hypothesis, σ does not lie in P_{m-1} , so it contains at least one vertex, $x \in V(\sigma)$, with $h(x) = m$. After permuting indices, we can write $x = (x_1, \dots, x_m, p_{m+1}, \dots, p_n)$ and $V(\sigma) = (\{x_1, y_1\}, \{x_2, y_2\}, x_3, -)$ with $x_i \neq p_i$. (Possibly, $y_1 = p_1$ and/or $y_2 = p_2$.) Consider the cube $(\{x_1, y_1\}, \{x_2, y_2\}, \{p_3, x_3\}, -)$. This contains the squares, σ_1 , σ_2 and σ_3 , with vertex sets given respectively by $(\{x_1, y_1\}, \{x_2, y_2\}, p_3, -)$, $(\{x_1, y_1\}, x_2, \{p_3, x_3\}, -)$ and $(x_1, \{x_2, y_2\}, \{p_3, x_3\}, -)$, all meeting at $(x_1, x_2, p_3, -)$. Now $h(\sigma_i) < h(\sigma)$ for each i , so σ_i is submersed. Now $\sigma_1 \cup \sigma_2 \cup \sigma_3$ is a wheel of the cube, and so it follows that the cube is submersed. In particular, σ is submersed, giving a contradiction.

This shows that $V(P_m)$ is submersed as claimed. Proceeding inductively, we obtain a submersion of $\Pi = P_n$ as required. \square

Note that if F_1, F_2, \dots, F_m are pairwise intersecting corank-1 faces of Π , then $m \leq n$, and $\bigcap_{i=1}^m F_i$ is a corank- m face of Π .

Lemma 23.4.5. *Suppose that Π is a prism of rank $n \geq 3$. Suppose F_1, F_2, \dots, F_m are pairwise intersecting corank-1 faces of Π , with $m \geq 1$. Suppose $f_i : F_i \rightarrow \Gamma$ are submersions with, $f_i|(F_i \cap F_j) = f_j|(F_i \cap F_j)$ for all i, j . Then there is a unique submersion $f : \Pi \rightarrow \Gamma$ with $f|F_i = f_i$ for all i .*

Proof. We already have f defined on $\bigcup_{i=1}^m F_i \subseteq \Pi$. Choose any basepoint, $p \in \bigcap_{i=1}^m V(F_i)$, and define P_2 as above with respect to this basepoint. Since $n \geq 3$, we have $\text{rank}(F_i) \geq 2$ for each i , so $P_2 \subseteq \bigcap_{i=1}^m F_i$. In particular, $f|P_2$ is a submersion. By Lemma 23.4.4, this extends to a submersion of Π . By Lemma 23.4.3 applied to each F_i , the extension is unique on F_i , and so agrees with f_i . By Lemma 23.4.3, the extension of f to Π is unique. \square

Here is another corollary to Lemma 23.4.4.

Lemma 23.4.6. *Let $\Pi = \Lambda_1 \square \Lambda_2 \square \dots \square \Lambda_n$ be a prism of rank n . Let $\Lambda' \supseteq \Lambda_1$ be a complete graph containing Λ_1 , and let $\Pi' = \Lambda' \square \Lambda_2 \square \dots \square \Lambda_n \supseteq \Pi$. Then any submersion, $f : \Pi \cup \Lambda' \rightarrow \Gamma$ extends uniquely to a submersion $f : \Pi' \rightarrow \Gamma$.*

Proof. We can suppose $\Lambda' \neq \Lambda_1$. Choose any basepoint $p \in \Lambda' \subseteq \Pi$, and define $P_2 \leq \Pi'$ as above. By Lemma 23.4.4, it is enough to extend f to a map on $V(P_2)$ such that every square containing p is submersed. To this end, let $x \in V(P_2) \setminus V(\Pi)$ with $h(x) = 2$. After permuting indices, we can write $x = (x_1, x_2, p_3, \dots, p_n)$, with $x_1 \notin \Lambda_1$, and $x_2 \neq p_2$. Let $y = (y_1, p_2, -)$, where $y_1 \in \Lambda_1 \setminus \{x, p_1\}$. Consider

the prism $(\{p_1, x_1, y_1\}, \{p_2, x_2\}, -)$. (A “roof” in our earlier terminology.) This contains the triangle $(\{p_1, x_1, y_1\}, p_2, -)$ in Λ' , and the square $(\{p_1, y_1\}, \{p_2, x_2\}, -)$ in Π , meeting along the edge $(\{p_1, y_1\}, p_2, -)$. These are submersed, so it follows that f extends to a submersion of the whole roof. In particular, we can uniquely define $f(x)$ so that the square $(\{p_1, x_1\}, \{p_2, x_2\}, -)$ is submersed. \square

Corollary 23.4.7. *Any prism in Γ is contained in a unique gated prism of the same rank.*

Proof. Let $\Pi \leq \Gamma$ be a prism of rank n . Choose any basepoint $p \in V(\Pi)$, and let $\Lambda_1, \dots, \Lambda_n$ be the sides of Π containing p . Let Λ'_i be the maximal clique of Γ containing Λ_i , and let $\Pi' = \Lambda'_1 \square \dots \square \Lambda'_n \supseteq \Pi$. Applying Lemma 23.4.6 n times, we extend the inclusion of Π into Γ to a submersion of Π' . By Lemma 23.3.4, this is an isometric embedding. By Lemma 23.3.5 its image is a gated prism.

Uniqueness is clear from the fact that every non-singleton clique is contained in a unique maximal clique. \square

Note that any gated prism, $\Pi \subseteq \Gamma$, is intrinsically quasimedial, and we can identify $\Theta(\Pi)$ as a subgraph of $\Theta(\Gamma)$. We will give further discription of $\Theta(\Pi)$ later.

We next give a description of cubical subgraphs of $\Theta = \Theta(\Gamma)$.

Suppose $a, b \in V(\Theta)$, with $\Pi(b) \leq \Pi(a)$. Let $Q = \{c \in \Theta \mid \Pi(b) \leq \Pi(c) \leq \Pi(a)\}$. Then Q is the vertex set of a cubical subgraph of Θ . This can be described explicitly as follows. Let $e_1, \dots, e_m \in \Theta$ be the vertices such that $\Pi(e_1), \dots, \Pi(e_m)$ are the corank-1 faces of $\Pi(a)$ containing $\Pi(b)$. If $J \subseteq \{1, \dots, m\}$, then $\bigcap_{i \in J} \Pi(e_i)$ is a corank- $\#J$ face of Π , and corresponds to $\Pi(e_J)$ for some $e_J \in V(\Theta)$. Thus, $Q = \{e_J \mid J \subseteq \{1, \dots, m\}\}$ has the structure of a cube. In the notation of Subsection 10.3, we have $e_J = \bigvee_{i \in J} e_i$.

In fact, any cubical subgraph of Θ has this form. This is easily verified if $J = \{1, 2\}$. To see this, let a be a vertex of a 2-cubical graph (i.e. square) such that $\text{rank}(\Pi(a))$ is maximal. Let e_1, e_2 be the adjacent vertices, and e_{12} the antipodal vertex. Now $\Pi(e_1)$ and $\Pi(e_2)$ must be corank-1 faces of $\Pi(a)$. Since $\Pi(e_{12}) \neq \Pi(a)$, it must be a common corank-1 face of both $\Pi(e_1)$ and $\Pi(e_2)$, and so $\Pi(e_{12}) = \Pi(e_1) \cap \Pi(e_2)$ as required. Given this, the general case can be proven by induction on $\#J$, again starting at a vertex a with $\text{rank}(\Pi(a))$ maximal. We will not give details of the general case since we do not need it.

We can now set about verifying the hypotheses (S1)–(S3) of Theorem 16.2.3 for the graph Θ , where \mathcal{S} is the set of all squares of Θ .

First note that the above discussion shows that a square of Θ is determined by any three of its vertices. This is (S2).

Next we verify (S3):

Lemma 23.4.8. *Any wheel in Θ has a dual wheel.*

Proof. Let $x, x_1, x_2, x_3, x_{12}, x_{23}, x_{31}$ be a wheel in Θ , with hub, x , adjacent vertices x_1, x_2, x_3 , and such that x, x_i, x_{ij}, x_j is a square for all distinct i, j . We want to find a cube with eighth vertex x_{123} : the hub of the dual wheel.

We write $\Pi = \Pi(x)$, $\Pi_i = \Pi(x_i)$, $\Pi_{ij} = \Pi(x_{ij})$ and $\Pi_{123} = \Pi(x_{123})$. Let $r = \text{rank } \Pi$ and $r_i = \text{rank}(\Pi_i)$. Thus $|r - r_i| = 1$. We will specify the cube by describing prisms, Π_{\min} and Π_{\max} , of minimal and maximal rank, which correspond to the elements of Θ lying in the cube. There are three cases. The descriptions are all simple consequences of the fact that any square in Θ has the standard form described above.

Case (1): $r_1 = r_2 = r_3 = r - 1$.

We have $\Pi_{\max} = \Pi$ and Π_1, Π_2, Π_3 are corank-1 faces of Π . We set $\Pi_{123} = \Pi_{\min} = \Pi_1 \cap \Pi_2 \cap \Pi_3$.

Case (2): $r_1 = r_2 = r - 1, r_3 = r + 1$.

We have $\Pi_{\max} = \Pi_3$ and $\Pi_{\min} = \Pi_{12}$. Here Π, Π_{13} and Π_{12} are corank-1 faces of Π_3 , with $\Pi_1 = \Pi \cap \Pi_{13}$ and $\Pi_2 = \Pi \cap \Pi_{23}$. We set $\Pi_{123} = \Pi_{13} \cap \Pi_{23}$.

Case (3): $r_1 = r_2 = r + 1, r_3 = r - 1$.

We have $\Pi_{\max} = \Pi_{12}$ and $\Pi_{\min} = \Pi_3$. Here, Π_1 and Π_2 are corank-1 faces of Π_{12} and $\Pi = \Pi_1 \cap \Pi_2$. Also Π_{12} and Π_{23} are corank-1 faces of Π_1 and Π_2 respectively, and $\Pi_3 = \Pi_{13} \cap \Pi_{23}$. Then Π_{123} is the corank-1 face of Π_3 containing $\Pi_{13} \cup \Pi_{23}$.

Case (4): $r_1 = r_2 = r_3 = r + 1$.

Here $\Pi_{12}, \Pi_{23}, \Pi_{31}$ pairwise intersect in corank-1 faces, Π_1, Π_2, Π_3 , and $\Pi = \Pi_1 \cap \Pi_2 \cap \Pi_3$ is a common corank-2 face. We can therefore construct a prism, $\Pi_{123} \supseteq \Pi_{12} \cup \Pi_{23} \cup \Pi_{31}$, with $\Pi_{12}, \Pi_{23}, \Pi_{31}$ corank-1 faces of Π_{123} . By Lemma 23.4.5, the inclusion of $\Pi_{12} \cup \Pi_{23} \cup \Pi_{31}$ into Θ extends to a submersion of Π_{123} . By Lemma 23.3.4 this is an isometric embedding, and using Lemma 23.3.5 we see that its image is a gated prism. In this case, we have $\Pi_{\min} = \Pi$ and $\Pi_{\max} = \Pi_{123}$. \square

Next, we want to verify (S1): that is, the complex, $S(\Theta)$, obtained by gluing a 2-cell to each square of Θ is simply connected. Of course, this can equivalently be expressed in combinatorial terms, as we did in Subsection 16.2.

First, we consider the analogous complex, $S(\Gamma)$, obtained by gluing a 2-cell to each triangle and each square of Γ . We claim that this is also simply connected. In fact, we just need to consider the following moves on a path, a_0, a_1, \dots, a_n , in Γ :

(1): If $a_{i-1} = a_{i+1}$, remove a_i and a_{i+1} .

(2): If a_{i-1} is adjacent to a_{i+1} , remove a_i .

(3): If a_{i-1}, a_i, a_{i+1}, b form a square for some $b \in V(\Gamma)$, replace a_i by b .

(Of course, after Moves (1) and (2), we need to relabel the indices after $i - 1$.)

Lemma 23.4.9. *Any closed path $a_0, a_1, \dots, a_n = a_0$, can be reduced to a constant path at a_0 by a sequence of moves of type (1)–(3) above.*

Proof. Whenever possible, perform moves of type (1) and (2). This will reduce the total length, n , of the path. If neither such move is possible, let $i > 0$ be

minimal such that $\rho(a_0, a_{i+1}) \leq \rho(a_0, a_i)$. Now a_{i-1} cannot be adjacent to a_{i+1} . Therefore, by Lemma 23.3.1 or by Property (\diamond) , we can set $b = a_{i-1}[a_i]a_{i+1}$ so that $\rho(a_0, b) < \rho(a_0, a_i)$. We now perform move (3). This does not change the length of the path, but it reduces $\sum_{i=1}^n \rho(a_0, a_i)$. Therefore after a finite number of steps this process terminates with $n = 0$. \square

We have noted that there is a morphism, $\theta : \Sigma(\Gamma) \rightarrow \Theta$ from the binary subdivision of Γ to Θ . This extends to a map, $\theta : \Sigma(S(\Gamma)) \rightarrow S(\Theta)$, from the binary subdivision of the square complex $S(\Gamma)$ to $S(\Theta)$, which sends square cells to square cells. In view of Lemma 23.4.9, to show that $S(\Theta)$ is simply connected, it is enough to show that any closed path in Θ can be homotoped in $S(\Theta)$ to the θ -image of a closed path in $\Sigma(\Gamma)$.

To this end, note that if $\Lambda \leq \Gamma$ is a maximal clique, then $\Theta(\Lambda) \leq \Theta(\Gamma)$ consists of a central vertex corresponding to Λ which is connected by an edge to each vertex of Λ . In other words, it is a “star graph”, $K_{1,c}$, where c is the (possibly infinite) cardinality of Λ . More generally, if $\Pi = \Lambda_1 \square \cdots \square \Lambda_n$ is a gated prism in Γ , then $\Theta(\Pi) = \Theta(\Lambda_1) \square \cdots \square \Theta(\Lambda_n)$ is a cartesian product of star graphs. (In fact, one can readily verify that $\Theta(\Gamma_1 \square \Gamma_2) = \Theta(\Gamma_1) \square \Theta(\Gamma_2)$ for any quasimedial graphs, Γ_1, Γ_2 .) In particular, we see that the square complex, $S(\Theta(\Pi)) \subseteq S(\Theta)$ is simply connected.

Lemma 23.4.10. *The square complex $S(\Theta)$ is simply connected.*

Proof. Let α denote a closed path with vertices, $a_0, a_1, \dots, a_n = a_0$, in Θ . As observed above, it is enough to show that α homotopic in $S(\Theta)$ to $\theta(\gamma)$, for some closed path, γ , in $S(\Gamma)$.

Now $\Pi_i := \Pi(a_i)$ is a sequence of gated prisms in Γ , either containing or contained in the next. Also, as noted above each $S(\Theta(\Pi_i))$ is simply connected. Now choose any $b_i \in \Pi_i \subseteq \Theta(\Pi_i)$. If $\Pi_{i+1} \leq \Pi_i$, we connect b_i to b_{i+1} by a path in Π_i . If $\Pi_i \leq \Pi_{i+1}$, we connect them by a path in Π_{i+1} . Concatenating these paths, we get a closed path in Γ , hence a closed path, γ , in its binary subdivision, $\Sigma(\Gamma)$. Now α and $\theta(\gamma)$ pass through the same sequence of subsets, $S(\Theta(\Pi_i))$, of $S(\Theta)$. Since each of the $S(\Theta(\Pi_i))$ is simply connected, it is easily checked that α and $\theta(\gamma)$ are homotopic in $S(\Theta)$ as required. \square

We have now verified the three hypotheses, (S1)–(S3), of Theorem 16.2.3. We conclude:

Proposition 23.4.11. *If Γ is a quasimedial graph, then $\Theta(\Gamma)$ is a median graph.*

(Again, we have expressed it in topological terms, but it is easily reinterpreted combinatorially.)

Put together with Lemma 23.4.1, we see:

Proposition 23.4.12. *If Γ is a quasimedial graph, then the vertex set, $V(\Gamma)$, with the combinatorial metric isometrically embeds into a discrete median algebra with the standard metric.*

In particular, $V(\Gamma)$ is submedian.

23.5. Some consequences.

This immediately allows us to deduce certain facts. For example, it puts certain constraints on the betweenness relation. Among them, we have

$$a.c.d \ \& \ c.e.f \ \& \ d.f.e \Rightarrow a.e.f.$$

To verify this, we can use the procedure described in Subsection 6.2. If the statement fails, we can assume that $a \uparrow$, $f \uparrow$ and $e \downarrow$. Now $c.e.f$ gives $c \downarrow$, $a.c.d$ gives $d \downarrow$, and $d.f.e$ gives the contradiction that $f \downarrow$.

As an immediate consequence we have:

Lemma 23.5.1. *If $a, b, c, d, e, f \in V(\Gamma)$, with $e, f \preceq c, d \preceq a, b$ then $e, f \preceq a, b$.*

Here, \preceq is the relation on pairs defined in Subsection 22.1. Note that this also implies that parallelism is an equivalence relation.

Another application concerns convexity. This only makes reference to betweenness. Recall that in a discrete median algebra, the convex hull of any finite set is finite (Lemma 11.1.3). Therefore:

Lemma 23.5.2. *Let Γ be a quasimedial graph. The convex hull of any finite subset of $V(\Gamma)$ is a finite subgraph.*

In particular this applies to the subalgebra generated by the subset: that is, the smallest subset containing it which is closed under the quasimedial operation.

It does not in general apply to the gated hull of a finite set. For example, the gated hull of any edge is the maximal clique containing it. This might be infinite.

23.6. Existence and uniqueness of quasimedial triples.

We next set about the proof of Theorem 23.1.2. We will need the following.

Lemma 23.6.1. *Let $\Pi \subseteq \Gamma$ be a gated prism. Then $\Theta(\Pi)$ is the convex hull of Π in $\Theta(\Gamma)$.*

Proof. Certainly $\Theta(\Pi)$ is connected. Therefore, by Lemma 11.4.4, to show that $\Theta(\Pi)$ is convex, it is enough to verify that any square in $\Theta(\Pi)$ with three vertices in $\Theta(\Pi)$ lies entirely in $\Theta(\Pi)$. By the earlier discussion, any square has the form $\Pi_0, \Pi_1, \Pi_2, \Pi_{12}$, where Π_1, Π_2 are corank-1 faces of Π_0 , and $\Pi_{12} = \Pi_1 \cap \Pi_2$. The only case that calls for verification is when $\Pi_1, \Pi_2, \Pi_{12} \leq \Pi$. Now $\Pi_0 \cap \Pi$ is also a gated prism containing Π_1 and Π_2 . We note that coranks of $\Pi_1 \cap \Pi$ and $\Pi_2 \cap \Pi$ in Π are both at most 1, and therefore equal to 1. Thus, $\text{rank}(\Pi_0 \cap \Pi) = \text{rank}(\Pi_0)$, so $\Pi_0 \leq \Pi$, so Π_0 lies in $\Theta(\Pi)$ as required.

The fact that $\Theta(\Pi)$ is the convex hull now follows easily from the fact that it is the cartesian product of star graphs. \square

By Lemma 11.3.3, we see that $\Theta(\Pi)$ is also gated in $\Theta = \Theta(\Gamma)$.

To simplify notation in the following discussion, we identify $V(\Gamma)$ as a subset of $V(\Theta)$ via θ . As we have noted, the betweenness relation in $V(\Gamma)$ agrees with that in $V(\Theta)$.

Suppose $x \in V(\Gamma)$ and that $\Pi \leq \Gamma$ is a gated prism. Under our identification $V(\Gamma) \subseteq V(\Theta)$, we have $\Pi \subseteq \Theta(\Pi)$, and so by Lemma 23.6.1, $\Theta(\Pi) = \text{hull}_\Theta(\Pi)$. We have gates $\omega_\Pi x$ for x in Π , and $\omega_{\Theta(\Pi)} x$ for x in $\Theta(\Pi)$, constructed in Π and Θ respectively. In fact, these are the same:

Lemma 23.6.2. $\omega_\Pi x = \omega_{\Theta(\Pi)} x$.

Proof. We want to show that $\omega_{\Theta(\Pi)} x$ is a gate for x in $\Theta(\Pi)$. For this, it is enough to check that $\omega_{\Theta(\Pi)} x \in V(\Gamma)$. Write $y = \omega_{\Theta(\Pi)} x$ and $\Pi' = \Pi(y)$. Then $\Pi' \subseteq \Pi$. From the remark following Lemma 23.4.1, we have $2\rho_\Gamma(x, \Pi') + \text{rank}(\Pi') \leq \rho_\Theta(x, y)$. Thus, if $y \notin V(\Gamma)$, we get $2\rho_\Gamma(x, \Pi) \leq 2\rho_\Gamma(x, \Pi') < \rho_\Theta(x, y)$. But since $\omega_\Pi x \in \Pi \subseteq \Theta(\Pi)$, we have $\rho_\Theta(x, y) \leq \rho_\Theta(x, \omega_\Pi x) = 2\rho_\Gamma(x, \omega_\Pi x) = 2\rho_\Gamma(x, \Pi)$ (see Lemma 23.4.1). This gives a contradiction. \square

We are now ready for:

Proof of Theorem 23.1.2. Let $x_1, x_2, x_3 \in V(\Gamma)$. Let $m = x_1 x_2 x_3 \in V(\Theta)$ be the median in Θ . Let $\Pi = \Pi(m)$ be the gated prism of Γ corresponding to m , and let $r = \text{rank}(\Pi)$. Let $a_i = \omega_\Pi x_i = \omega_{\Theta(\Pi)} x_i$. We claim that (a_1, a_2, a_3) is the unique quasimedial triple of (x_1, x_2, x_3) in Γ .

If $i \neq j$, then $m \in [x_i, x_j] \cap \Theta(\Pi)$. We therefore have $x_i m x_j$, $x_i a_i m$ and $x_j a_j m$, and so $x_i a_i m a_j x_j$. Note that $x_i a_i a_j x_j$ holds also in Γ . It also follows that $m = a_1 a_2 a_3$ is the median of a_1, a_2, a_3 in Θ . Now by Lemma 23.4.2 we have $\rho_\Theta(a_i, m) = r$ for each i . Therefore, $\rho_\Theta(a_i, a_j) = 2r$, and so $\rho_\Gamma(a_i, a_j) = r$ for all $i \neq j$.

Now suppose $b_1, b_2, b_3 \in V(\Gamma)$ with $x_i b_i b_j x_j$ for all $i \neq j$. Let $m' = b_1 b_2 b_3$ in Θ . Then also $m' = x_1 x_2 x_3$, so $m' = m$. We therefore have $x_i b_i m b_j x_j$ for $i \neq j$. Now $\rho_\Theta(b_i, m) \geq r$, so $\rho_\Theta(b_i, b_j) \geq 2r$, so $\rho_\Sigma(b_i, b_j) \geq r$. If $\rho_\Sigma(b_i, b_j) = r$ for all $i \neq j$, then $\rho_\Theta(b_i, m) = r$ for all i . Thus, $\rho_\Theta(x_i, b_i) = \rho_\Theta(x_i, m) - r = \rho_\Theta(x_i, a_i)$, and so $\rho_\Sigma(x_i, b_i) = \rho_\Sigma(x_i, a_i)$. Moreover, since $\rho_\Theta(b_i, m) = r$, by Lemma 23.4.2, we have $b_i \in \Pi$. Since a_i is the gate for x_i in Π , it follows that $b_i = a_i$.

We have shown that (a_1, a_2, a_3) is the unique quasimedial triple for (x_1, x_2, x_3) , as required. \square

Note that, in the above proof, a_1, a_2, a_3 all lie in the gated prism Π , and that the distance between any two of them is equal to the rank of Π . It follows that Π is the gated hull of $\{a_1, a_2, a_3\}$. In summary, this shows:

Proposition 23.6.3. *Let $x_1, x_2, x_3 \in V(\Gamma)$ and let a_1, a_2, a_3 be their quasimedial triple. Then the gated hull of $\{a_1, a_2, a_3\}$ is a gated prism, and its convex hull is cartesian product of triangles.*

23.7. Some further observations.

We proceed with some more observations about a quasimedial graph.

It is easy to see that if $G \leq \Gamma$ is convex, then it is intrinsically quasimedial. If G is gated, then any gated prism in G is also a gated prism in Γ . In this way, we can identify $\Theta(G)$ as a full subgraph of $\Theta(\Gamma)$ (thereby generalising the case where G was a prism, described in Subsection 23.4). Under the identification of $V(\Gamma)$ as a subset of $V(\Theta)$, we have $G \subseteq \Theta(G)$.

The following generalises Lemmas 23.6.1 and 23.6.2:

Lemma 23.7.1. *If $G \leq \Gamma$ is gated, then $\Theta(G)$ is the convex hull of G in $\Theta(\Gamma)$. Moreover, $\omega_G|V(\Gamma) = \omega_{\Theta(G)}|V(\Gamma)$.*

Proof. First we show that $\Theta(G)$ is convex. Certainly $\Theta(G)$ is connected, so we want to check local convexity (Lemma 11.4.4). Using the notation of the proof of Lemma 23.6.1, let $\Pi_0, \Pi_1, \Pi_2, \Pi_{12}$ be the vertices of a square in $\Theta(\Gamma)$, three of which lie in $\Theta(G)$. We want to show they all do. Again, we need only consider the case where $\Pi_1, \Pi_2 \subseteq G$. Now Π_0 is the gated hull of $\Pi_1 \cup \Pi_2$ in Γ , so since G is gated, we have $\Pi_0 \subseteq G$, thus Π_0 lies in $\Theta(G)$, so $\Theta(G)$ is locally convex as required.

The fact that $\Theta(G)$ is the convex hull of G follows easily from the fact that this is true of gated prisms.

For the statement about gate maps, let $x \in V(\Gamma)$. As in the proof of Lemma 23.6.2, to verify that $\omega_G x = \omega_{\Theta(G)} x$, it is sufficient to show that $\omega_{\Theta(G)} x \in V(\Gamma)$. Let $\Pi = \Pi(\omega_{\Theta(G)} x)$. Now $\omega_{\Theta(G)} x = \omega_{\Theta(\Pi)} x$, and by Lemma 23.6.2, we have $\omega_{\Theta(\Pi)} x = \omega_{\Pi} x \in V(\Gamma)$ as required. \square

We have observed that the relation of parallelism is an equivalence relation in a quasimedial graph. We also note:

Lemma 23.7.2. *Suppose $G, G' \subseteq \Gamma$ are parallel subgraphs. Then G is gated if and only if G' is gated. Moreover, if $\tau : G \rightarrow G'$ is the translation between them, then $\omega_G = \tau \circ \omega_{G'}$.*

Proof. Suppose that G is gated. We just need to check that $\tau \circ \omega_G$ is a gate map to G' . We only need to consider vertex sets. Let $x \in V(\Gamma)$, $a = \omega_G x$ and $b = \tau a$. Let $c \in V(G')$, and set $d = \tau^{-1} c$. We have $x.a.d$ & $a.b.c$ & $b.c.d$. We claim that this implies $x.b.c$. For suppose not. The fact that $V(\Gamma)$ is submedial allows us to apply the reasoning of Subsection 6.2. We can assume $x \uparrow$, $c \uparrow$ and $b \downarrow$. By $a.b.c$, we have $a \downarrow$. By $b.c.d$, we have $d \uparrow$. Now $x.a.d$ gives the contradiction that $a \uparrow$. This shows that b is a gate for x in G' as required. \square

We also note that, in view of Lemma 23.7.2, if $G, G' \leq \Gamma$ are gated, then G, G' are parallel in Γ if and only if $\Theta(G)$ and $\Theta(G')$ are parallel in $\Theta(\Gamma)$.

Now suppose that $G \leq \Gamma$ is gated. Given $a \in V(G)$, let $H(G, a) = V(\Gamma) \cap \omega_G^{-1} a$. Note that $a \in H(G, a)$.

Lemma 23.7.3. *$H(G, a)$ is gated.*

Proof. Write $H = H(G, a)$. First note that H is 1-path-connected. To see this, note that if $x \in H$, then $[x, a] \subseteq H$. In particular, any geodesic path in $V(\Gamma)$ from x to a lies in H .

We now claim that if $x, y \in H$ are distinct and $z \in V(\Gamma)$ is adjacent to both x and y in Γ , then $z \in H$. For suppose not. Let $b = \omega_G z \neq a$. As observed at the end of Subsection 22.2, we have $x, z \parallel a, b \parallel y, z$, so by Lemma 23.5.1, we have $x, z \parallel y, z$ so $x = y$, giving a contradiction.

It now follows by Lemmas 23.3.2 and 23.3.3 that H is gated. □

(In fact, it is not hard to deduce the claim of the second paragraph directly from properties (∇) and (\diamond) without appeal to more sophisticated statements.)

In particular, Lemma 23.7.3 applies to gated prisms, including maximal cliques. In the latter case, we can say a bit more.

Given a gated subgraph, $G \leq \Gamma$, and a subset, $A \subseteq V(G)$, write $H(G, A) = V(\Gamma) \cap \omega_G^{-1}A$.

Lemma 23.7.4. *Let $\Lambda \leq \Gamma$ be a maximal clique, and let $A \subseteq V(\Lambda)$ be any non-empty set of vertices. Then $H(\Lambda, A)$ is convex in $V(\Theta)$.*

Proof. Again, $H := H(\Lambda, A)$ is 1-path-connected. Suppose $x, y \in H$ are distinct and non-adjacent in Γ , and that $z \in V(\Gamma)$ is adjacent in Γ to both x and y . Let $a = \omega_\Lambda x$, $b = \omega_\Lambda y$ and $c = \omega_\Lambda z$. Thus, $a, b \in A$. We claim $c \in A$. For if not, we have $a, x \parallel c, z \parallel b, y$, so $a, x \parallel b, y$. Since a, b are adjacent in Γ , the same is true of x, y , contrary to our assumption. This shows that $z \in H$. The statement now follows applying Lemma 23.3.2. □

(Again, we could alternatively apply (∇) and (\diamond) directly.)

Here is another way of looking at this.

Let $\Lambda = \Lambda(p)$ be a maximal clique in Γ , corresponding to some $p \in V(\Theta)$. Let $\epsilon = \epsilon(\Lambda, a)$ be the edge of $\Theta(\Lambda)$ which connects a to the central vertex, p , of $\Theta(\Lambda)$. We have a gate map, $\omega_\epsilon : V(\Theta) \rightarrow \epsilon$. This determines a halfspace, $H_\Theta(\Lambda, a) := \omega_\epsilon^{-1}a$, and hence a wall, $W_\Theta(\Lambda, a) := \{H_\Theta(\Lambda, a), H_\Theta(\Lambda, a)^C\}$, of $V(\Theta)$. Since $\omega_\epsilon = \omega_\epsilon \circ \omega_{\Theta(\Lambda)}$, we see that $H(\Lambda, a) = V(\Gamma) \cap H_\Theta(\Lambda, a)$. Also, $H(\Lambda, V(\Lambda) \setminus \{a\}) = V(\Gamma) \cap H_\Theta(\Lambda, a)^C$.

Note that $W(\Lambda) := \{H(\Lambda, a) \mid a \in V(\Lambda)\}$ is a partition of $V(\Theta)$ into gated subsets. The set of partitions arising in this way satisfies a number of constraints.

Suppose Λ, Λ' are maximal cliques. By Lemma 22.1.2, $\omega_\Lambda \Lambda' \leq \Lambda$ and $\omega_{\Lambda'} \Lambda \leq \Lambda'$ are gated. In particular, each is either a maximal clique or a single vertex. There are two possibilities. Either $\omega_\Lambda \Lambda' = \Lambda$ and $\omega_{\Lambda'} \Lambda = \Lambda'$, and so Λ, Λ' are parallel. In this case, by Lemma 23.7.2, $W(\Lambda) = W(\Lambda')$. Or else, $\omega_\Lambda \Lambda' = \{a\}$ and $\omega_{\Lambda'} \Lambda = \{a'\}$ are both single vertices. We split this latter case in turn into two subcases as follows.

Maybe $H(\Lambda, a) \cup H(\Lambda', a') = V(\Gamma)$. In this case $H(\Lambda, b) \subseteq H(\Lambda, a')$ for all $b \in V(\Lambda) \setminus \{a\}$.

If not, then the walls $W_\Theta(\Lambda, a)$ and $W_\Theta(\Lambda', a')$ of $V(\Theta)$ cross. By Lemma 11.3.1, there is a square in $V(\Theta)$, crossed by each of these walls. This has the form $\Theta(\Pi)$, where $\Pi = \Lambda_0 \times \Lambda'_0$ is a rank-2 maximal prism, with sides Λ_0 and Λ'_0 , meeting at a vertex $a'' \in V(\Pi)$, such that $a.a''.a'$ holds in $V(\Theta)$ hence also in $V(\Gamma)$. Since a, Λ is parallel to a', Λ_0 in Θ , one can check that there is a translation of Λ to Λ_0 sending a to a'' . This is precisely the gate map from Λ to Π . Similarly the gate map to Π translates Λ' to Λ'_0 and sends a' to a'' . Therefore, up to translating Λ and Λ' , we can suppose that $\Lambda \cap \Lambda' = \{a\} = \{a'\}$, and that they are the sides of a rank-2 prism Π . In this case $H(\Lambda, b) \cap H(\Lambda', b') \neq \emptyset$ for all $b \in V(\Lambda)$ and all $b' \in V(\Lambda')$.

Let $\mathcal{W}(\Gamma)$ be the set of all partitions of the form $W(\Lambda)$ where $\Lambda \leq \Gamma$ is a maximal clique. Given distinct $W, W' \in \mathcal{W}(\Gamma)$, we say that W, W' are **nested** if there is some $H \in W$ and $H' \in W'$ with $H \cup H' = V(\Theta)$. We say that W, W' **cross** if $H \cap H' \neq \emptyset$ for all $H \in W$ and all $H' \in W'$.

The above discussion shows in particular:

Lemma 23.7.5. *Any two distinct elements of $\mathcal{W}(\Gamma)$ are either nested or cross.*

We can conversely use partitions to construct quasimedial graphs as discussed in Example (Ex23.7) above.

As with median algebras, this has an interpretation in terms of flows, as we briefly describe below.

Let Y be a set, and let $\mathcal{W} = (W_i)_{i \in I}$ be a family of non-trivial partitions of Y , indexed by some set, I . A **flow** on \mathcal{W} is a family, $R = (R_i)_{i \in I}$ with $R_i \in W_i$ for all i and such that $R_i \cap R_j \neq \emptyset$ for all i, j . Let $\mathcal{F} = \mathcal{F}(\mathcal{W})$ be set of all flows. There is a natural map, $\eta : Y \rightarrow \mathcal{F}$ defined by setting $\eta(x) = R$, where R_i is the element of the partition W_i containing x . Given three flows, $R, S, T \in \mathcal{F}$, define $RST := U$, where $U_i = R_i S_i T_i$ is the dual discriminator on W_i (i.e. $U_i = R_i$ if $S_i \neq T_i$ and $U_i = S_i = T_i$ if $S_i = T_i$). In general, RST need not be a flow. However, it will be if we assume:

(*) if $i \neq j$ then W_i, W_j are either nested or cross.

These terms were defined above for the partitions of the vertex set of a quasimedial graph arising as described above, and Lemma 23.7.5 tells us that (*) holds in this case.

More generally, let \mathcal{W} be a family of partitions satisfying (*), and let $R, S, T \in \mathcal{F}(\mathcal{W})$. To see that $U := RST$ is a flow in this case, suppose for contradiction that $U_i \cap U_j = \emptyset$ for some i, j . By (*), W_i, W_j are nested, i.e. there exist $H_i \in W_i$ and $H_j \in W_j$ with $H_i \cup H_j = Y$. Up to swapping i, j , we can assume that $U_i \neq H_i$. Since $U_i \subseteq H_i^C \subseteq H_j$ and $U_i \cap U_j = \emptyset$, we must have $U_j \neq H_j$ also. But now $R_i \cap R_j \neq \emptyset$ and so either $R_i = H_i$ or $R_j = H_j$ (or both). Similarly for S_i, S_j and T_i, T_j . Up to swapping i, j and S, T we therefore have either $R_i = S_i = H_i$ or $S_i = T_i = H_i$. These both give the contradiction that $U_i = H_i$.

This shows that \mathcal{F} is closed under the quasimedial operation.

In the construction of Example (Ex23.7) of Subsection 23.2, we started with an arbitrary finite family of partitions of a set X . (This need not satisfy (*).) We

embedded X into the vertex set, $Y := V(\Gamma)$, of a quasimedial graph Γ . This gives us a new family of partitions of this larger set Y , which does satisfy (*). Such constructions have applications to phylogenetics, as we mention in the Notes to this section.

24. COARSE GEOMETRY

The term “coarse geometry” refers to the large-scale structure of metric spaces. This typically means that we are only really interested in points defined up a bounded distance, and distances being specified up to an additive constant (or maybe within certain specified bounds). If we allow ourselves such freedom, then it turns out that many naturally occurring spaces admit a kind of “coarse median” structure. This has a number of applications.

24.1. Quasi-isometries and hyperbolicity.

We begin with some basic definitions from coarse geometry.

Recall that a metric space is “geodesic” if any two points are connected by a geodesic path — that is, a path whose length is equal to the distance between its endpoints.

Definition. Let (X, ρ) be (X', ρ') be geodesic metric spaces. We say that a map $\phi : X \rightarrow X'$ is **coarsely lipschitz** if there exist constants $k_1, k_2 \geq 0$ such that

$$\rho'(\phi(x), \phi(y)) \leq k_1 \rho(x, y) + k_2$$

for all $x, y \in X$.

We say that ϕ is a **quasi-isometric embedding** if it is coarsely lipschitz and also there exist constants $k_3, k_4 \geq 0$ such that

$$\rho(x, y) \leq k_3 \rho'(\phi(x), \phi(y)) + k_4$$

for all $x, y \in X$.

We say that ϕ is a **quasi-isometry** if it is a quasi-isometric embedding and also there is a constant $k_5 \geq 0$ such that $X' = N(\phi(X), k_5)$. In other words, the image of ϕ is cobounded in X' .

We say that X, X' are **quasi-isometric** if there is a quasi-isometry between them.

One can therefore think of a quasi-isometry as a bilipschitz map up to an additive constant. It should be noted that a quasi-isometry is not assumed to be continuous. Also a quasi-isometric embedding need not be injective (though if two points get mapped to the same point, then they must have been a bounded distance apart in the domain — in other words, it is “coarsely injective”). Of course, the above definitions make sense if the spaces are not assumed to be geodesic, though it is less useful in general. In what follows we will assume that we are dealing with geodesic spaces.

We write $X \sim Y$ to mean that X and Y are quasi-isometric. One can readily check that $X \sim X$, $X \sim Y \Leftrightarrow Y \sim X$ and that $X \sim Y \sim Z \Rightarrow X \sim Z$. One can therefore speak of “quasi-isometry classes” of spaces.

Here are some examples:

(Ex24.1a): Any bounded space is quasi-isometric to a point.

(Ex24.2a): If Δ is a connected finite-dimensional cube complex, built out of unit euclidean cubes and equipped with the geodesic metric, then the inclusion of its 1-skeleton, Γ , into Δ , is a quasi-isometry.

(Ex24.3a): Suppose that Γ is a connected locally finite graph, and G is a group which acts freely by graph automorphisms on Γ , with finite quotient graph, Γ/G . Then G is finitely generated. Moreover any finitely generated group acts on any graph in this way. (For example, take its Cayley graph with respect to some finite generating set.) Suppose the same group, G , also acts in this way on another graph, Γ' . Then $\Gamma \sim \Gamma'$. (Here, we are taking the combinatorial path-metrics, so that each edge has unit length.) It therefore makes sense to say that two finitely generated groups are “quasi-isometric”, or that a finitely generated group is “quasi-isometric” to a given geodesic space. For this we simply substitute such a graph, Γ , for the group G . As observed the choice of Γ does not matter.

(Ex24.4a): Examples (Ex24.2a) and (Ex24.3a) above, overlap in the case of right-angled Artin groups: see Example (Ex13.5) of Subsection 13.1.

(Ex24.5a): The fundamental group of a compact hyperbolic n -manifold is quasi-isometric to (constant curvature) hyperbolic n -space. In particular, this applies to fundamental groups of closed orientable surfaces of genus at least 2.

Remark. Examples (Ex24.3a) and (Ex24.5a) above are instances of the “Schwarz-Milnor Lemma”, which is a central result of coarse geometry. We say more about this in the Notes to this section.

Remark. Quasi-isometries do not, in general, respect any additional structure of the space. For example, a self-quasi-isometry of the euclidean plane, \mathbb{R}^2 , can be a mess. For example, it might send geodesic rays from the origin to logarithmic spirals. Nevertheless, there are also spaces for which any self-quasi-isometry is a bounded distance from an isometry. Such a space is sometimes called “quasi-isometrically rigid” (though this term has also been used to mean a number of different, but related, things).

Given geodesic spaces X, Y , we write $X \preceq Y$ to mean that there is a quasi-isometric embedding of X into Y . Clearly $X \sim Y \Rightarrow X \preceq Y$. Also, $X \preceq Y \preceq Z \Rightarrow X \preceq Z$. (However, $X \preceq Y$ & $Y \preceq X$ does not in general imply $X \sim Y$.)

For example, one can show that $\mathbb{R}^m \preceq \mathbb{R}^n$ implies $m \leq n$. It follows that $\mathbb{R}^m \sim \mathbb{R}^n$ if and only if $m = n$.

Another key notion in the subject is that of hyperbolicity as defined by Gromov.

Definition. A metric space X is *hyperbolic* if it is geodesic, and there exists a constant, k , such that for any geodesic triangle in X , there is some point, $m \in X$, within a distance k of each of the three sides of the triangle.

Here a *geodesic triangle* comprises three geodesics (its *sides*) cyclically connecting three points of X (its *vertices*).

It turns out that, up to bounded distance, m depends only on the three vertices, x, y, z , of a triangle. We choose such an m and write $\mu(x, y, z) = m$. This is sometimes referred to as a *centroid* of the triangle, or of the triple (x, y, z) .

There are many other equivalent definitions of hyperbolicity. We just mention here the “four-point condition”. We state it as a theorem:

Theorem 24.1.1. *A geodesic metric space is hyperbolic if and only if there is some constant $k \geq 0$ if for all $a, b, c, d \in M$, we have*

$$\rho(a, b) + \rho(c, d) \leq \max\{\rho(a, c) + \rho(b, d), \rho(a, d) + \rho(b, c)\} + k.$$

Note that this condition is the same as property (FP) described in Subsection 15.1, except with an additive constant. In particular, it follows that the notion of “0-hyperbolic” is equivalent to that defined there. Thus, a metric space is 0-hyperbolic if and only if it is an \mathbb{R} -tree.

Here is another key fact:

Theorem 24.1.2. *If X, Y are quasi-isometric geodesic spaces, then X is hyperbolic if and only if Y is.*

Remark. In the above (and subsequent) discussion there are various constants involved. Typically, the constants outputted only depend on those inputted. One can strengthen the statements so as to explicitly keep track of these constants, though we won’t bother with that here.

As an example, constant curvature (real) hyperbolic space (of any dimension) is hyperbolic. Hence, so is the fundamental group of any compact hyperbolic manifold. (See Examples (Ex24.3a) and (Ex24.5a) above.) In fact, the same is true of any compact manifold of strict negative curvature.

24.2. Coarse median spaces.

It is well known that a hyperbolic space has a certain “treelike structure”. From this, it follows that the centroid operation behaves like the median in a simplicial tree; that is, a rank-1 median algebra. One can formulate, and generalise, this idea by replacing a tree with a finite CAT(0) cube complex, or equivalently a finite median algebra. That is the idea behind the following definition.

Let (Λ, ρ) be a geodesic space.

Definition. A map $\mu : \Lambda^3 \rightarrow \Lambda$ is a *coarse median* if the following two conditions are satisfied:

(C1): there exist constants, $t, l \geq 0$ such that for all $a, b, c, a', b', c' \in \Lambda$, we have

$$\rho(\mu(a, b, c), \mu(a', b', c')) \leq t(\rho(a, a') + \rho(b, b') + \rho(c, c')) + l$$

and,

(C2): there is a function $h : \mathbb{N} \rightarrow [0, \infty)$ such that if $A \subseteq \Lambda$ with $\#A \leq p < \infty$, then there is a finite median algebra (Π, μ_Π) , and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$, such that for all $a \in A$, $\rho(a, \lambda\pi a) \leq h(p)$, and for all $a, b, c \in \Pi$,

$$\rho(\mu(\lambda a, \lambda b, \lambda c), \lambda\mu_\Pi(a, b, c)) \leq h(p).$$

We say that Λ has **coarse (median) rank** at most ν if (given some fixed map h) we can always take $\text{rank}(\Pi) \leq \nu$.

A **coarse median space** (of rank ν) is a geodesic metric space equipped with a coarse median (of rank ν).

In the above, we break with tradition, and denote the median on Π by μ_Π .

It is not hard to see that:

(C3): There is some $k \geq 0$ such that for all $a, b, c \in \Lambda$,

$$\rho(\mu(a, b, c), \mu(b, c, a)) \leq k \quad \text{and} \quad \rho(\mu(b, a, c), \mu(b, a, c)) \leq k.$$

In other words, μ is symmetric up to bounded distance. (Indeed, there is no essential loss in assuming it to be precisely symmetric.)

It also follows that for all $a, b \in \Lambda$, $\rho(a, \mu(a, a, b))$ is bounded. In other words, it satisfies axiom (M1) of a median algebra up to bounded distance.

In fact, an equivalent, and somewhat simpler, definition of a coarse median can be found in [NibWZ1]:

Theorem 24.2.1. *A geodesic metric space, (Λ, ρ) , is coarse median if and only if it satisfies (C3) and (C1) above, together with*

(C2a): *There is some constant $k \geq 0$ such that for all $a, b, c, d \in \Lambda$, we have:*

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

Note that axiom (C2a) is a weakening of axiom (M2) of a median algebra. In this way, we can think of Λ as having a median algebra structure up to bounded distance.

Theorem 24.2.1 does not in itself capture the notion of the rank of a coarse median space. However, this can be achieved via another result of [NibWZ1]. We just state this informally by noting that Λ has rank at most ν if any quasimorphism of a $(\nu+1)$ -cube into Λ is degenerate. A “quasimorphism” is a map which preserves the median operation up to bounded distance (as with the map λ of axiom (C2)). A quasimorphism of a cube is “degenerate” if two points (without loss of generality a 1-face of the cube) get mapped a bounded distance apart (depending on the constant of quasimorphism). The equivalence of these definitions is not hard to check (given appropriate quantification of the constants involved) though we will not give a proof here.

We give a few examples of coarse median spaces. (We will comment further on them in the Notes.) We say that a finitely generated group is “coarse median” if the associated graph, Γ (for example, its Cayley graph) admits a coarse median

(see Example (Ex24.3a) of the previous subsection). This does not depend on the choice of Γ . In each of the examples below, there is a canonical choice of coarse median operation (up to bounded distance).

(Ex24.1b): Any connected median metric space is coarse median with the same median operation. Here we can take the additive constants to be 0. This applies to CAT(0) cube complexes, and hence to right-angled Artin groups.

(Ex24.2b): The direct product of two coarse median spaces with the median defined independently on each coordinate and with l^2 product metric is coarse median.

(Ex24.b): A hyperbolic space is coarse median of rank 1, where the median is given by the centroid operation. Conversely, any coarse median space of rank 1 is hyperbolic, and the median equals the centroid operation up to bounded distance. In particular, this applies to hyperbolic groups.

(Ex24.4b): A group that is hyperbolic relative to a family of coarse median groups is coarse median. In particular, this applies to geometrically finite kleinian groups, as well as to “limit groups” as defined by Sela.

(Ex24.5b): The mapping class group of a compact orientable surface, Σ , is coarse median of rank $\xi(\Sigma)$. Here, the “complexity”, $\xi(\Sigma)$, of Σ is defined to be three times the genus plus the number boundary components minus 3.

(Ex24.6b): The Teichmüller space of a compact orientable surface, Σ , equipped with the Teichmüller metric, is also coarse median of rank $\xi(\Sigma)$.

(Ex24.7b) The Teichmüller space of a compact orientable surface, Σ , equipped with the Weil-Petersson metric, is coarse median of rank $\lfloor (\xi(\Sigma) + 1)/2 \rfloor$, where $\lfloor \cdot \rfloor$ denotes integer part.

To be precise, in Examples (Ex24.5b)–(Ex24.7b), we should assume that $\xi(\Sigma) \geq 2$. This rules out a finite number of exceptional surfaces.

24.3. An outline of some applications.

A useful fact about a coarse median space is that any tautological median identity holds in the coarse median space up to bounded distance, depending only on (the complexity of) the identity. This can be seen directly from the original definition we gave. Another way to think about it is to write a proof of the identity in terms of the axioms (M1) and (M2) (as is possible as we discussed in Subsection 6.1) and then rewrite the proof substituting the coarse median operation, and applying (C1) and (C2a) instead. A similar observation holds for conditional identities.

Given this, one can now mimic various median algebra constructions in this context.

For example, given $a, b \in \Lambda$, we can define the “coarse interval” between them as $[a, b]_\Lambda := \{\mu(a, b, c) \mid c \in \Lambda\}$. Up to bounded distance this is equal to $\{c \in \Lambda \mid \rho(c, \mu(a, b, c)) \leq k\}$ for some sufficiently large $k \geq 0$. (All constants can be chosen to

depend only on those featuring in the axioms for Λ as a coarse median space.) This can be viewed as “coarsifying” the corresponding statement in a median algebra, M : namely that a median interval $[a, b]_M$ can be equivalently defined as $\{abc \mid c \in M\}$ or as $\{c \in M \mid abc = c\}$. The fact that these are equivalent relies on the tautological median identity: $ab(abc) = abc$. The corresponding statement in Λ says that $\rho(\mu(a, b, \mu(a, b, c)), \mu(a, b, c))$ is bounded for all $a, b, c \in \Lambda$.

One can apply similar principles to other definitions and results. For example, one can define a subset $C \subseteq \Lambda$, to be “coarsely convex” if $[a, b]_\Lambda$ lies in a bounded neighbourhood of C for all $a, b \in C$. One can define a notion of “coarse gate map”, and of “coarse convex hull” etc. With this, one can prove coarse versions of results such as Lemma 7.4.4 etc. Of course, one needs to take some care to properly quantify these statements.

We mention one application of the coarse median property to illustrate how it can be used. (Some further applications are mentioned briefly in the Notes to this section.)

One can show:

Theorem 24.3.1. *Let Λ be a coarse median space of rank at most $\nu < \infty$. If $\mathbb{R}^n \preceq \Lambda$, then $n \leq \nu$.*

The proof of this (and lots of other results about such spaces) makes use of the “asymptotic cone” construction. We briefly give the idea as follows.

Let (X, ρ) be a metric space. Let $(r_i)_{i \in \mathbb{N}}$ be a sequence of positive numbers tending to ∞ . Let X_i be the rescaled metric space, $(X, \rho/r_i)$. The idea is to pass to a limiting space, (X_∞, ρ_∞) . This can be achieved by putting a non-principal ultrafilter on \mathbb{N} , and taking limits in the appropriate sense. This might depend on various choices, such as the choice of ultrafilter, or of the scaling factors, but the choice does not matter to us here. We refer to X_∞ as an “asymptotic cone” of X . It is a complete metric space. If X is a geodesic space, then so is X_∞ .

The general idea then is that additive constants disappear in the limit. For example, if X is a hyperbolic space, then X_∞ will be 0-hyperbolic: in other words an \mathbb{R} -tree (see Subsection 15.1).

Suppose that $\phi : X \rightarrow Y$ is a coarsely lipschitz map between two metric spaces, X, Y . Passing to the asymptotic cones, we get a lipschitz map $\phi_\infty : X_\infty \rightarrow Y_\infty$ — the additive constant has gone away. If ϕ is a quasi-isometric embedding, then ϕ_∞ is a bilipschitz map onto its range, $\phi_\infty(X_\infty)$, which is a closed subset of Y_∞ . If ϕ is a quasi-isometry, then ϕ_∞ is a bilipschitz map to Y_∞ .

Now suppose that Λ is a coarse median space. Property (C2) says that the map $\mu : \Lambda^3 \rightarrow \Lambda$ is coarsely lipschitz with respect to the l^1 product metric on Λ^3 . We therefore get a limiting lipschitz map, $\mu_\infty : \Lambda_\infty^3 \rightarrow \Lambda_\infty$. We have observed that (Λ, μ) satisfies the axioms of median algebra up to bounded distance. Therefore, $(\Lambda_\infty, \mu_\infty)$ is a median algebra. Indeed it is not hard to see that it is a lipschitz median algebra as defined in Subsection 12.2. Theorem 13.4.1 now tells us in the finite-rank case that Λ_∞ admits a median metric, ρ' , bilipschitz equivalent to ρ_∞ ,

and which induces the median structure μ_∞ . We can now invoke the theory of median metric spaces.

For example, to prove Theorem 24.3.1, let $\phi : \mathbb{R}^n \rightarrow \Lambda$ be a quasi-isometric embedding. Passing to asymptotic cones, we get a map $\phi_\infty : \mathbb{R}^n \rightarrow \Lambda_\infty$, bilipschitz onto its range, so in particular a topological embedding. (Taking an asymptotic cone of \mathbb{R}^n just gives us back \mathbb{R}^n , since rescaling the metric on \mathbb{R}^n does not change anything.) Now the rank of Λ_∞ as a median algebra is at most the coarse rank of Λ : by hypothesis at most ν . By Lemma 12.3.3, Λ_∞ has locally compact dimension at most ν . In particular, $\dim(\mathbb{R}^n) \leq \nu$, and so $n \leq \nu$, as required. (See [Bo2] for the details.)

In general asymptotic cones are quite complicated objects, and difficult to describe explicitly. Nevertheless, there are various regularity and structural theorems one can bring into play. For example, Theorems 13.4.2 and 13.4.3 tell us respectively that any asymptotic cone of a finite-rank coarse median space admits a bilipschitz equivalent CAT(0) metric and a bilipschitz equivalent injective metric. In particular, (by either of these statements) it is contractible.

24.4. Coarse median algebras.

We briefly mention the related notion of a “coarse median algebra” as defined in [NibWZ2]. This makes no direct reference to a metric. Let M be a set equipped with a symmetric ternary operation, μ .

Definition. We say that M is a *coarse median algebra* if it satisfies

(M1): for all $a, b \in M$, $\mu(a, a, b) = a$, and

(CM3): there is some $k \in \mathbb{N}$ such that for all $a, b, c, d, e \in M$,

$$\#[\mu(a, b, \mu(c, d, e)), \mu(\mu(a, b, c), \mu(a, b, d), e)] \leq k.$$

Here the interval, $[a, b]$, is defined to be $\{\mu(a, b, x) \mid x \in M\}$.

Note that (CM3) is a weakening of property (M3) of a median algebra. In particular, (M3) is obtained by setting $k = 1$.

A coarse median algebra is said to have *bounded valency* if for all $r \geq 0$, there is some $R \geq 0$ such that for all $a \in M$, $\#\{b \in M \mid \#[a, b] \leq r\} \leq R$.

An example of a bounded-valency coarse median algebra would be the vertex set of a bounded geometry graph with a coarse median structure with respect to its combinatorial metric. (By “bounded geometry” we mean that its vertex degrees are bounded by some finite number.) The definition might therefore be applied directly to a finitely generated group with the word metric.

It is shown in [NibWZ2] that a bounded-valency coarse median algebra admits a metric with respect to which it is a coarse median space, and that such a metric is unique up to quasi-isometry.

25. INJECTIVE METRIC SPACES AND HELLY GRAPHS

In this section, we briefly mention a couple of other structures which are not (in general) median algebras, but which have links to the general theory. The

first of these in an “injective metric space”. This has already been briefly alluded to in Subsections 13.4 and 24.3. One can loosely think of this as a third metric description of “non-positive curvature”, along with “median” and “CAT(0)”. We have observed that \mathbb{R}^n is respectively median, CAT(0), and injective in the l^1 , l^2 and l^∞ metrics. A similar statement holds for CCAT(0) cube complexes (see Example (Ex13.4) of Subsection 13.1, Theorem 18.1.2 and the discussion below). One can view the Helly Property of graphs as a combinatorial analogue of injectivity.

25.1. Injective metrics.

Here is the definition of “injective”:

Definition. A metric space, X , is *injective* if for every metric space, Y , and any 1-lipschitz map $f : A \rightarrow X$ defined on any subset, A , of Y , there is a 1-lipschitz extension $f : Y \rightarrow X$.

In other words, injective metric spaces are the injective objects in the category of metric spaces, where the morphisms are 1-lipschitz maps. One can show that any injective metric space is complete, geodesic and contractible.

As noted, examples include l^∞ spaces, in particular, the l^∞ metric on \mathbb{R}^n and on CCAT(0) cube complexes. Theorem 13.4.3 here gives a more general condition in which a median metric gives rise to a canonical injective metric. It also turns out that every metric space canonically isometrically embeds into an injective metric space, referred to as its “injective hull” [Is1]. See [Lan] for another proof of this.

A more geometric interpretation of injectivity is in terms of “hyperconvexity”.

Definition. A metric space, (X, σ) , is *hyperconvex* if given any family, $((x_i, r_i))_{i \in \mathcal{I}}$, in $X \times [0, \infty)$ indexed by some set \mathcal{I} , such that $\sigma(x_i, x_j) \leq r_i + r_j$ for all $i, j \in \mathcal{I}$, then $\bigcap_{i \in \mathcal{I}} N(x_i, r_i) \neq \emptyset$.

The following was proven in [ArP].

Theorem 25.1.1. *A metric space is injective if and only if it is hyperconvex.*

Note that in a geodesic metric space, the condition $\sigma(x_i, x_j) \leq r_i + r_j$ is equivalent to asserting that $N(x_i, r_i) \cap N(x_j, r_j) \neq \emptyset$.

The basic idea to relate this to median metrics is as follows.

First consider the case of \mathbb{R}^n . Here l^∞ balls are median convex in the l^1 metric. The Helly Property (Lemma 7.1.1) tells us that the intersection of any finite family of pairwise intersecting balls is non-empty. Since balls are compact, we see that this statement also applies to infinite such families.

This can be generalised in a number of ways. For example, we can take the l^∞ metric on a finite dimensional CCAT(0) cube complex. Again closed balls are median convex, and one can use this to deduce that such an l^∞ metric is hyperconvex, hence injective.

Such an idea is also the basis of Theorem 13.4.3. One can view the argument as approximating a median metric space by CCAT(0) complexes, by considering

subalgebras generated by finite subsets thereof. (A similar principle is used for Theorem 13.4.2, taking an l^2 metric instead of an l^∞ metric.)

25.2. Helly graphs.

We mentioned that injective metric spaces have a combinatorial analogue in “helly graphs”. Here is the definition:

Definition. A *helly graph* is a connected graph whose vertex set is hyperconvex in the combinatorial metric.

In other words, any family of pairwise intersecting balls has non-zero intersection.

One can give a number of equivalent characterisations, and they exhibit a number of useful combinatorial properties. These are exploited in [ChalCGHO].

A direct connection with median graphs was described at the end of Subsection 11.6. Given a graph Γ , recall that $\Gamma^\Delta \supseteq \Gamma$ is the graph with the same vertex set obtained by adding diagonals to cubes. (This can be thought of as analogous to replacing an l^1 metric with an l^∞ metric.) Lemmas 11.6.3 and 11.6.4 together tell us that if Γ is median, then Γ^Δ is helly — see the discussion at the end of Subsection 11.6.

We note that there is also a coarsification of this property described in [ChalCGHO], where the authors define the notion of a “coarse helly space”. A coarse version of Theorem 13.4.3 has been given in [HaeHP]. There, the authors construct a canonical new quasi-isometric metric on any coarse median space which has “quasicubical intervals”. Under some additional assumptions, they show that the new metric is coarse helly. All the conditions hold for most naturally occurring examples. They also derive a number of consequences: for example, such a space is “semihyperbolic”. This is yet another coarse non-positive curvature condition. It says loosely that the space admits a combing. This can be thought of as a coarse analogue of the cube-path construction described in Subsection 11.6.

26. NOTES

26.1. Notes on Section 1 (Introduction).

Some general references for median algebras are [Is2, BanHe, Verh, R]. They have been discovered or rediscovered in a variety of different contexts, and are known, at least in the earlier literature, by a variety of names. It is difficult to say precisely where the notion originates. Certainly, notions of “betweenness” are central to the axiomatisation of euclidean geometry (see for example, [Veb]). Various formal “postulates” and connections between them, appear in [HuK]. A notion of betweenness for metric spaces is explored in [Me]. More specifically, median operations are discussed in [BiK] and [Gra]. One of the first references to axioms in a form we give here (namely (M1) and (M3) of Subsection 3.2) appear in the abstract to a talk entitled “Ternary distributive semi-lattices” by Avann in [Gre] (see page 79 thereof). Shortly afterwards the notion was significantly

developed by work of Sholander [Sh] and others. They were later rediscovered by Nebeský [Ne]. More historical background can be found in [BanHe] and in [R].

26.2. Notes on Section 2 (Distributive lattices).

A general reference for lattice theory is [Bi].

There are several equivalent ways of presenting the axioms of a distributive lattice. For example, only one of the distributive laws is required: the other can be deduced given the other axioms.

It also turns out that, in a lattice, the distributive laws are equivalent to the statement that $(x \wedge y) \vee (y \wedge z) \vee (z \wedge x) = (x \vee y) \wedge (y \vee z) \wedge (z \vee x)$ for all x, y, z . This is Theorem 8 of Chapter II of [Bi]. We have proven one direction as Lemma 2.1.1. The argument in the reverse direction goes as follows. Suppose $p \leq x$. Then, using the absorption laws, for any r , we get $x \wedge (p \vee r) = x \wedge (x \vee r) \wedge (p \vee r) = (x \vee p) \wedge (x \vee r) \wedge (p \vee r) = (x \wedge p) \vee (x \wedge r) \vee (p \wedge r) = p \vee (p \wedge r) \vee (x \wedge r) = p \vee (x \wedge r)$. Applying this rule twice, we see that if $p, q \leq x$, then $x \wedge (p \vee q \vee r) = p \vee q \vee (x \wedge r)$. Now, let x, y, z be arbitrary. We apply the above, with $p = x \wedge y$, $q = x \wedge z$ and $r = y \wedge z$. We get $x \wedge (y \vee z) = x \wedge (x \vee y) \wedge (x \vee z) \wedge (y \vee z) = x \wedge ((x \wedge y) \vee (x \wedge z) \vee (y \wedge z)) = (x \wedge y) \vee (x \wedge z) \vee (x \wedge y \wedge z) = (x \wedge y) \vee (x \wedge z)$ again using the absorption laws. The other distributive law follows by symmetry.

In defining a free distributive lattice, some authors allow for meets and joins of the empty set. This entails adjoining maximum and minimum elements to the lattice, and allowing \emptyset and $\{\emptyset\}$ as Sperner families. However, this makes no essential difference to the discussion.

A short proof of Sperner's Lemma is given in [Lu].

The cardinality of the free distributive lattice on n elements is called the “ n th Dedekind number” (possibly adding 2 if we were to adjoin an additional maximum and minimum). Our argument shows that this is at most 2^{2^n} . One can certainly do better, though it is known to grow superexponentially in n . Some asymptotic estimates for Dedekind numbers are discussed in [KleM].

26.3. Notes on Section 3 (Basic facts about median algebras).

The binary join operation is often denoted by $A \circ B := J(A, B)$. We show in Section 7 (Lemma 7.1.4) that it is associative: $(A \circ B) \circ C = A \circ (B \circ C)$. Further discussion of the operation is given in [BanHe].

The proof of Lemma 3.2.10 is based on that of Nieminen [Niem] and simplified by Roller [R] (see the Notes to Section 8 below).

An immediate consequence of Proposition 3.3.3 is that iterating the median operation stabilises after a finite number of steps if we start with a finite set. This number is referred to as the “stabilisation degree” of the set. This is explored in [BanV2]. (See also Lemma 8.2.4 here.) In these terms, sets with stabilisation degree 0 are precisely linear median (sub)algebras or 2-cubes, as we discussed in Example (Ex3.3) of Subsection 3.4. This result can be found in [Ev].

The axioms of ternary boolean algebra, (M, μ) , as originally given by Grau [Gra], are in a slightly different form, namely:

$$\begin{aligned} \mu(a, b, \mu(c, d, e)) &= \mu(\mu(a, b, c), d, \mu(a, b, e)), \\ \mu(a, b, b) &= \mu(b, b, a) = b \quad \text{and} \\ \mu(a, b, b^*) &= \mu(b^*, b, a) = a \end{aligned}$$

for all $a, b, c, d, e \in M$. It is not assumed a-priori that the median is fully symmetric, nor that $a^{**} = a$. These are deduced. Note that, at least if we assume symmetry of the median, the first two statements are equivalent to the axioms, (M3) and (M1), of a median algebra.

In the rank-1 case, the construction of Example (Ex3.7) was used in [Bo1] as a stage in constructing an \mathbb{R} -tree, starting with a pretree satisfying certain hypotheses. In this case, the final relation \approx can be equivalently defined by saying that $a \not\approx b$ if the (totally ordered) interval $[a, b]$ contains a subset order isomorphic to the rationals, \mathbb{Q} . To be of any use, of course, one needs to find conditions under which M/\approx can be shown to be non-trivial.

26.4. Notes on Section 4 (Intervals and betweenness).

The fact that one can derive the identity (M3) from (M1) and (M2) appears in the paper [KoM]. The authors proceed by showing that this implies the interval conditions, essentially (I1)–(I4) here, and then referring to Sholander’s paper [Sh]. Sholander’s proof is quite ingenious, and somewhat involved. (See [Bo4] for some further comment on it.) It should also be noted that Sholander started from somewhat weaker hypotheses, which did not assume full symmetry of the median operation.

The proof given here is a slight rephrasing of the argument given in [VeroM]. It seems that Donald Knuth had taken an interest in this question, and set the challenge of finding an explicit derivation of this identity. (Of course, Sholander’s argument could, in principle, be reduced to such, though it would likely be extremely long.) In response to this, Veroff and McCune used their “Otter” software to generate such a derivation and posted the output. By identifying some key steps in this proof, Knuth then proposed a simplification based on his own arguments. This simplification was subsequently reported in [VeroM]. We note that Knuth also includes some discussion of median algebras from this perspective in his book [Kn].

We remark that a third way of proving (M3) would be to proceed via the “Pasch property”. This is a betweenness axiom which asserts that, for $a, b, c, x, y \in M$, $a.x.b$ & $a.y.c \Rightarrow [b, y] \cap [c, x] \neq \emptyset$. Note that this is implied by (M4) (since $abc \in [b, y] \cap [c, x]$). From this, one can in turn prove that any two distinct points of M are separated by a wall: see [Chep1]. (In fact, one can derive the stronger “Kakutani separation property”, given here as Theorem 8.1.2.) This means that it is sufficient to verify (M3) in a two-point median algebra (as in the general procedure described in Subsection 6.2). Note, however, that this argument is not

constructive: it makes use of Zorn's Lemma to deduce the Kakutani property. Some further discussion of the Pasch property can be found in [Vandev].

How these arguments compare may be debatable. In some ways, it is remarkable that in the seventy years since Sholander's original work, no-one has yet come up with a significantly shorter proof.

Our proof of Theorem 4.3.1 essentially condenses the relevant arguments found in [Is2]. The proof given there is distributed through a more general discussion, and somewhat complicated by the fact that the ternary operation is not generally assumed to be symmetric.

26.5. Notes on Section 5 (Free median algebras).

What we call a "flow" is commonly referred to as an "ultrafilter" elsewhere. We prefer the more intuitive term, which has a natural geometric interpretation (see Section 9), and has been used elsewhere in the context of various treelike structures.

The term "superextension" originates with work of de Groot in general topology. Here, we are just applying the term in the case of a finite set (i.e. a finite hausdorff topological space).

Some discussion of the structure of superextensions is given in [BanV1]. A picture of $\Phi(X)$ when $\#X = 5$ can also be found in [R]. The cardinality of $\Phi(X)$ for $\#X = n$ is sometimes called the n th "Hoşten-Morris number". For $n \leq 9$, these are calculated in [BrouMMV]. The sequence begins 0, 1, 2, 4, 12, 81, 2646, 1422564, ... and grows superexponentially (cf. the case of free distributive lattices mentioned in Section 2).

The operation we denote by $(A|p)$ is introduced in [SpW]. We make further use of it in Section 7.

26.6. Notes on Section 6 (Expressions and identities).

The method we describe for verifying identities in Subsection 6.2 is well established. For example, it is referred to as "halfspace reasoning" in [Verh]. We have introduced our own notation here.

Recall that, in the terminology of first-order logic, a "sentence" is a first-order formula with no free variables. The notation, $M \models \phi$, means that the sentence, ϕ , is valid in a particular set-theoretic interpretation, M . For us, median algebras are models of a first-order theory whose vocabulary consists of a single ternary operation (and the binary relation of equality). The axioms of the theory are the sentences $\theta_0, \theta_1, \theta_2$ which respectively assert that the ternary relation is symmetric and satisfies (M1) and (M2). In other words, M is a median algebra if and only if $M \models (\theta_0 \ \& \ \theta_1 \ \& \ \theta_2)$. In these terms, an "expression" can be thought of a first-order sentence of the form $(\forall x_1)(\forall x_2) \dots (\forall x_n)(E = E')$, where E, E' are median expressions in (at most) the variables, x_1, \dots, x_n . Thus, such an expression is "tautological" if $M \models \phi$ for every median algebra, M . (Note that the axioms of a median algebra are themselves tautological identities.)

We note that Gödel’s Completeness Theorem, applied to the theory of median algebras, tells us that any first-order sentence which is valid in every median algebra is provable in the theory. In particular, any (conditional) tautological identity admits such a proof. Note however, that a “proof” is more general than what we have called a “derivation”: the former might involve general first-order sentences along the way. The general Completeness Theorem relies on the Axiom of Choice. The fact identities admit primitive recursive derivations is therefore a quite different statement.

In fact, we remark that an analogous statement in group theory turns out to be false. There are examples of particular group presentations, which have word problems solvable in primitive recursive time, but which have non-primitive recursive Dehn functions. (See [CoMO, DisER].)

In general, the first-order theory of median algebras is undecidable. In other words, there is no algorithm to determine whether or not a given sentence is provable from the axioms, or equivalently, if it is valid in every median algebra. This was observed in [Is2] to be a consequence of the fact that the theory of bounded distributive lattices is undecidable. That result is generally attributed to [Grz]. Theorem 6 of that paper refers to “Brouwerian algebras”. It follows that the theory of Heyting algebras is undecidable. A Heyting algebra can be viewed as a bounded distributive lattice with an additional axiom, and it follows that the latter theory is also undecidable. Some more recent discussion of the undecidability of Heyting algebras is given in [IdI]. (In [Is2], it is suggested that the case of bounded distributive lattices follows “easily” from that of distributive lattices, though it is unclear to me why that is the case.)

To make the connection with median algebras, suppose that ϕ is a first-order sentence in the language of bounded distributive lattices: that is with binary function symbols, \wedge and \vee , and constants $0, 1$: the latter being interpreted respectively as the minimal and maximal elements. We can systematically convert it to a median-algebra sentence, ϕ^* , as follows. Let a, b be variables not featuring in ϕ . Replace all occurrences of $0, 1$ respectively by a, b , and replace all subformulae $p \wedge q$ and $p \vee q$ respectively by apq and bpq . Next, replace quantifiers, \forall, \exists , by quantifiers restricted by the predicate $(x = abx)$, where x is the quantified variable. This gives a formula, $\phi_{a,b}$. Now set $\phi^* = (\forall a)(\forall b)\phi_{a,b}$ (where \forall is unrestricted). We claim that ϕ is valid in every bounded distributive lattice if and only if ϕ^* is valid in every median algebra. For suppose L is a bounded distributive lattice, and write M for its usual structure as median algebra. Interpreting $a = 0$ and $b = 1$, we have $x = abx$ for all $x \in M$, and we see that $M \models \phi^*$ implies $L \models \phi$. Conversely, suppose M is a median algebra, and $a, b \in M$. Let $L = [a, b]_M$, with its structure as a bounded distributive lattice. Now $L \models \phi$ implies that $\phi_{a,b}$ holds in M . Therefore, if $L \models \phi$ holds for all L , then $M \models \phi^*$ as claimed. Thus, decidability for median algebras would imply decidability for bounded distributive lattices, contrary to known results.

26.7. Notes on Section 7 (Convex sets).

One can study median algebras from the point of view of convex structures. A convex structure is a set together with a family of subsets deemed “convex” which satisfy various general properties one might associate with convexity. In particular, it is closed under intersection and increasing union. A general reference to this subject, which includes a discussion of median algebras, is [Vandev]. Further discussion can be found in [Verh].

Convex sets in a median algebra are sometimes referred to as “ideals”. This is made explicit in the case of boolean rings as discussed in (Ex3.5) in Subsection 3.4 (see also Subsection 9.6). Gated convex sets are called “Čebyšev ideals” in [Is2]. Lemma 7.1.2 is due to Nieminen [Niem]. Its interpretation in terms of Lemmas 3.2.8 and 7.1.3 can be found in [BanHe]. A different proof of Lemma 7.4.6 can be found in [R].

26.8. Notes on Section 8 (Walls and rank).

Theorem 8.1.2 is due to Nieminen [Niem], using a result of Balbes. In [R], Roller states that this proof contains a “small gap”, and proceeds to give a self-contained argument. We have based our proof on the argument presented there.

In [Niem], everything is expressed in term of order structures. From this point of view, convex sets are “ideals” and halfspaces are “prime ideals”: see Subsection 9.6 and the Notes to Section 7.

A version of Proposition 8.2.4 is proven in [Fi5] by a rather different argument: see Proposition 4.3 thereof. (This makes use of the corresponding fact for convex hulls, namely Proposition 8.2.3 here.) More precisely, this states that $\langle A \rangle = T^{h(\nu)}(A)$ for some explicit function $h : \mathbb{N} \rightarrow \mathbb{N}$. It is not clear what the optimal function should be. For example one could certainly improve on the exponent n in Lemma 11.7.2. However, in [BanV2] it is shown that this exponent must in general be at least logarithmic in n . Further discussion of such matters can be found in that paper.

Proposition 8.2.5 is described in [Vandev]: see I.6.20 and II.4.25.7 thereof. Our argument essentially reproduces the proof there. (I thank Elia Fioravanti for bringing this fact to my attention.)

The use of Dilworth’s Lemma in the context of median algebras can be found in [BrodCGNW] where they prove a version of Lemma 8.3.3 for CAT(0) cube complexes.

26.9. Notes on Section 9 (Halfspaces and duality).

The notion of “poc set” was introduced in [R]. It is the same as what we have called a “proset”, except that it is assumed to have preferred “zero” element, 0 , minimal with respect to \leq , and in place of our Axiom (P2), it is assumed instead that if $a \leq a^*$, then $a = 0$. There is a simple canonical procedure to pass back and fore between a poc set and a proset: deleting 0 and 0^* or adjoining these

elements. For most of what we do, it is more convenient to exclude the zero element. However, in some contexts it is more natural to include it: for example when passing to subalgebras as in Subsection 19.2.

Our example (Ex9.5) of arrangements of lines (at least for finite poc sets) is described in [R], where they are called “poc pictures”. Not all finite poc sets can be represented in this way, and so they do not give a complete picture of what poc sets might look like.

The notion of a proset (or poc set) is a generalisation of the notion of a protree as defined by Dunwoody, and which corresponds to the “rank-1” case. These arise in various contexts: in particular, in relation to group splittings and accessibility of groups [Du].

In the discrete case, the duality between poc sets and median algebras was motivated by the construction of Sageev [Sa]. The approach there is somewhat different, and expressed in terms of cube complexes rather than median algebras.

As noted in Section 5, many sources talk in terms of “ultrafilters” rather than flows.

The fact that the superextension of a finite set is generated by the principal ultrafilters has an interpretation in terms of logical circuits as we discuss in Section 20.

Further discussion of duality is given in Subsection 12.5.

Proposition 9.4.2 is proven in [Nic]. The argument we have given is a slight variation on the proof there. This is really about spaces with walls, which we discuss in the more general context of spaces with measured walls in Section 19.

Some of the constructions of Subsection 9.5, such as the relations defined on downward sequences, can be found in [Gur]. (What we have called “subinfinite-rank” is termed “ ω -dimensional” there.)

26.10. Notes on Section 10 (Hypercubes).

Some further discussion of finite cubes can be found in [Bo4]. For example, one can show that for any given points, a, b in a median algebra M , there is at most one cube of maximal rank in M with antipodal points a, b .

26.11. Notes on Section 11 (Discrete median algebras).

Canonical cube paths of the type we describe are constructed in [NibR2].

Our proof of Proposition 11.8.4 broadly follows that of [Ban], though we have expressed it a bit differently. There the term “retract” is used for what we have called a “folding”. The latter term is not standard in general, but has been used with a similar meaning in the context of foldings of trees or metric 1-complexes.

Superextensions are discussed in [BanV2], where the authors attribute the calculation of the rank (given as Proposition 11.9.1 here) to Verbeek.

The Erdős-Ko-Rado Theorem (Theorem 1 of [ErKR]) asserts the following. Let X be a set with $\#X = n < \infty$. Let \mathcal{A} be a family of pairwise intersecting subsets, each of cardinality at most $r \leq n/2$, and such that no element of \mathcal{A} is properly

contained in any other. Then $\#\mathcal{A} \leq \binom{n-1}{r-1}$. A short proof of the theorem is given in [Ka]. (There it is assumed that the sets in \mathcal{A} all have size r , though the argument works to prove the more general statement.)

The notion of a helly graph is defined in Section 25. The fact that Γ^Δ is a helly graph (see Lemma 11.6.4 and subsequent discussion) can be found in [BanV1]. See [BanC2] for further discussion.

The original definition of the Roller boundary was given in [R]. There are now quite a few accounts of this, often giving equivalent constructions. They are discussed in some detail in [Gur]. That paper includes a discussion of canonical paths (there termed “canonical flows”), as well variations on some of the constructions we have described here. A general survey, with some more recent references, is included in [Fi3]. That paper also discusses variations for more general median algebras (see also Subsection 12.6 here). The description of the Roller boundary in terms of Busemann cocycles is discussed in [CaL], where it is attributed to unpublished work of Bader and Guralnik. Roller boundaries are also discussed in [Ge3] (phrased mostly in terms of median graphs). In that paper, it is shown that each component of the boundary has strictly smaller rank (cf. Proposition 11.12.10 here), and that the Roller boundary of each component can be naturally identified with a component of the original boundary.

Event structures were introduced in [NielpW], in the context of computational processes. The basic idea is to give a framework for describing the causal interdependence of processes, or “events”, occurring during a given computation. In this context, the relation $E \leq E'$ is meant to reflect a necessary “causal” link: event E must occur before event E' can occur; and the relation $E \# E'$ reflects “conflict”: events E and E' cannot both occur. A connection with median graphs, including an account of Proposition 11.13.2, is given in [BarC] (where the term “site” is used for an event structure). Further discussion of this connection is described in [ChalC]. (I thank Victor Chepoi bringing event structures to my attention.)

26.12. Notes on Section 12 (Topological median algebras).

The complicating issue regarding the CW topologies is this. If we take the product of two CW complexes (such as cube complexes) then in general the CW topology is finer than the product topology. There are various conditions under which they are the same. For example, this holds if both are countable [Mil]. A fairly complete account of the situation is given in [Broo]. However, this still does not answer the question of when $\Delta(\Pi)$ is a topological median algebra in the CW topology.

The proof of Lemma 12.2.2 is based on that of Lemma 2.6 of [Fi3].

One reference for the theory of dimension is [En]. A space has “covering dimension” at most ν if every open cover has a refinement for which every point lies in at most $\nu + 1$ elements of the refinement. One can also define an “inductive dimension” by saying that a space has dimension at most ν if any two disjoint closed subsets are separated by a closed subset of dimension at most $\nu - 1$, and deeming

the empty set to have dimension -1 . In the case of a locally compact hausdorff space, these are equivalent. In fact, in this case one can take the closed sets to be singletons. By any of these definitions, the dimension of \mathbb{R}^{ν} is ν .

Proposition 12.4.7 and its proof were shown to me by Elia Fioravanti. I thank him for his suggestion that I include it here. It is also natural to ask under what conditions a topological median algebra is contractible. (For example, is path-connected and second countable sufficient?) This certainly holds in some cases, such as cube complexes (Section 17).

What we call a “Stone median algebra” is called “Boolean median algebra” in [R] (though there is some potential for confusion with a boolean algebra viewed as a median algebra). Roller gives a second duality result, similar to Theorem 12.5.1. In that result, it is shown that a general median algebra is dual to a poc set equipped with a compact totally disconnected topology. In other words, we topologise the poc set rather than the median algebra. (This is a situation where one really wants to include the elements 0 and 0^* in the poc set so as to get a compact space.) The arguments are somewhat analogous, but more involved.

The compactificaton construction we describe (Lemma 12.6.2) can be found in [Fi3]. It is in turn motivated by a construction in Ward [Wa1], which serves to compactify certain real trees. There are many more results concerning Roller boundaries of topological median algebras in [Fi3], including a version of Lemma 12.6.2.

A (topological) “real tree” can be defined as a hausdorff topological space in which any pair of points are connected by a unique arc. (There is some variation in terminology in the literature.) A metrisable real tree admits a metric inducing the original topology, in which it is an \mathbb{R} -tree (as defined in Section 15). This is shown in [MayO]. A “dendron” is a compact real tree. A “dendrite” is a metrisable dendron. Such spaces arise in many contexts, and have various equivalent characterisations: see for example [Bo1] and the references therein.

26.13. Notes on Section 13 (Median metric spaces).

Some references for median metric spaces are [ChatteDH, Bo4, Fi1, Fi2, Fi3] and the citations therein. An exposition of right-angled Artin groups is given in [Char].

In [Fi1], it is shown that any finite-rank median metric space isometrically embeds into a connected complete median metric space of the same rank. Some further results on median metric spaces can be found in [Fi3]. For example, any halfspace in a complete finite-rank median metric space is either open or closed (or both).

Our discussion of connectedness leaves open a number of questions. For example, is every connected median metric space geodesic? We have seen that this holds if we make a few additional assumptions, but I know of no reference to this question in complete generality.

Another illustration of (the conclusion of) Theorem 13.4.2 relates to Guirardel cores (Subsection 15.4) as we mention in the Notes to Section 15.

Regarding Theorem 13.4.3, it was pointed out to me by Elia Fioravanti that one needs to add the hypothesis of “finite rank” to Corollary 3.3 of [Bo7]. This is to ensure that balls in the l^∞ metric are bounded in the l^1 metric. This is false in general, but is justified in the finite-rank case by Lemma 6.2 of that paper. This is sufficient to prove the main result, namely Theorem 7.8 thereof, which gives Theorem 13.4.3 here.

26.14. Notes on Section 14 (Submedian relations).

The Model Existence Theorem, generally attributed to Henkin, can be thought of as a variation of Gödel’s Completeness Theorem. It tells us, in particular, that any consistent countable theory has a (countable) model. In our argument at the end of Subsection 14.1, we could, instead, have constructed an (uncountable) model by taking an ultraproduct of the examples X_n .

One could restrict attention to submedian relations of bounded rank: that is to say, induced by embedding in a median algebra of rank bounded by some $\nu \in \mathbb{N}$. I do not know if this class is finitely axiomatisable for $\nu \geq 2$. (For $\nu = 1$, this is just the notion of a pretree.) One can check that the rank of our example, X_n , tends to ∞ as $n \rightarrow \infty$.

The axioms (R1)–(R4) of a pretree (or rather their equivalent for strict betweenness) appear in a paper of Ward [Wa2], where they are meant to capture the separation properties of points in a topological continuum. The term “pretree” to describe such a structure was proposed in [Bo1]. This term is meant to evoke the related notion of a “protree” due to Dunwoody [Du] (see the Notes to Section 9). Such structures were studied independently in [AdeN], where they are called “B-sets”.

If we assume also the existence of medians (R5), then the resulting structure has a longer history, and has been studied more extensively. In addition to “median pretree”, they have been known by variety of names, such as a “tree algebra” (for example in [BanHe]) or a “Herrlich tree” (for example in [Chi]). A variant of Proposition 14.2.1 is given in [Bo1]. This uses a somewhat more involved construction to canonically embed any pretree into a “complete” median pretree. Important examples of median pretrees are “ \mathbb{R} -trees” which we study in Section 15. It is shown in [Chi] that any countable median pretree (hence also any countable pretree) can be embedded in an \mathbb{R} -tree. We refer to [AdeN, Bo1] and references therein for further background to this topic.

In the case rank-1 median algebras, many properties can be deduced more simply. For example, one can give a relatively simple proof of the identity (M3) bypassing the rigmarole of Subsection 4.2. The key point is that any finite rank-1 median algebra is naturally the vertex set of a simplicial tree.

26.15. Notes on Section 15 (\mathbb{R} -trees).

The notion of an \mathbb{R} -tree was introduced in [MoS]. Since then, they have found numerous applications in low-dimensional topology and geometric group theory.

There is an extensive theory of isometric group actions on \mathbb{R} -trees, originating in the work of Rips. One way in which they arise is as asymptotic cones of Gromov hyperbolic spaces: a subject discussed briefly in Subsection 24.3. For a general survey of \mathbb{R} -trees, see [Bes].

In [Bo3], a variation on Proposition 15.3.1 was proven under the hypothesis that M is a ν -colourable k -lipschitz median algebra where any two points are connected by an l -lipschitz path, where $k, l \geq 0$ are fixed constants. This can be used to deduce Theorem 13.4.1, except that the median metric arising will not be canonical.

Guirardel cores were introduced in [Gui]. They arise naturally in a number of different contexts. They have found many applications, in particular to the study of group splittings. In [Gui] it is shown that if the core is connected, then it is CAT(0) in the induced path metric. In general, the \mathbb{R} -trees which arise in this context are not complete. But if we take their completions, then this can be viewed as an illustration of Theorem 13.4.2 here.

We remark that a different construction of “cores” associated to group actions on finite-rank median metric spaces are described in [Fi4]. These are different from Guirardel cores, being more closely related to convex cores rather than subalgebras.

Much of the basic theory of \mathbb{R} -trees can be generalised to “ Λ -trees”. The definition is similar, except that \mathbb{R} is replaced by an arbitrary ordered abelian group, Λ . As such, it is a Λ -metric space: it has a “metric” which takes values in the non-negative elements of Λ . Any Λ -tree is naturally a rank-1 median algebra (that is to say, a median pretree). Conversely, it is shown in [Chi] that any rank-1 median algebra can be embedded in a Λ -tree for some ordered abelian group, Λ . Some generalisations to median algebras of higher rank are described in [Bo8].

26.16. Notes on Section 16 (Median graphs).

A general reference for median graphs is [Ge3]. (See also the discussion of CAT(0) cube complexes in Section 18.)

Median graphs can be described in many equivalent ways, and there is now a vast literature on the subject. An early reference is [Av]. A more recent survey can be found [KlaM]. In particular, a version of Lemma 16.1.1 can be found in [Mul1].

The proof we have given for the “if” direction of Theorem 16.2.3 is along the lines of a more general result in [ChalCHO]. The idea is based on a similar principle to the usual proof of the Cartan-Hadamard theorem in riemannian geometry. This says that the exponential map to a complete simply connected riemannian manifold of non-positive curvature is a diffeomorphism. In our account, the map $f : G \rightarrow \Gamma$ is the analogue of the exponential map (see, for example, [BalGS]). We mention the result of [ChalCHO] again in Section 23, since it also applies to quasimedial graphs.

Another approach to this result is given in [Chep2]. The idea here is that any closed curve in Γ spans a (singular) disc in the 2-complex $\Delta(\Gamma, \mathcal{S})$. We can take this to have minimal area (having the minimal number of 2-cells). The condition (C2) then ensures that this has certain combinatorial properties. One can exploit this

in a number of ways. For example, given three points of Γ , connect them by three geodesics, and span the resulting geodesic triangle by a disc, such that the area is the least possible among all such geodesic triangles. If the disc were non-trivial, one can delete boundary squares, reducing the area. Therefore the minimal area is 0. In other words, the three geodesics all meet, and so the three points have median. The basic idea is fairly simple, but the details are somewhat involved. Another account of this approach is given in [Ge3].

We should observe that our statement is a bit more general than that given in [Chep2] in that we are not assuming a-priori that \mathcal{S} consists of all squares of Γ . However, it seems that the argument of [Chep2] could be used to prove the more general statement. One can also give a geometric interpretation in terms of the CAT(0) property, as discussed in Section 18.

The CCAT(0) condition can be equivalently phrased by saying that the link of every cell is a flag simplicial complex. We note that it is not sufficient just to require this for codimension-2 links. An interesting counterexample to this is given in [Adi]. That paper constructs a cubulation of the 3-sphere with at least four 3-cubes meeting around each 1-cell.

Median graphs (and more generally quasimedial graphs) have found application in phylogenetics: see the Notes to Section 23.

26.17. Notes on Section 17 (Cube complexes).

Some of the discussion of metrics on cube complexes applies to more general polyhedral complexes, at least in the l^2 case. Some general discussion of these matters can be found in [BriH]. A fairly detailed account of the standard l^2 and CW topologies on general cube complexes is given in the appendix to [Le].

26.18. Notes on Section 18 (The CAT(0) property).

The term “CAT(0)” (or more generally “CAT(κ)”) was introduced by Gromov. It is based on comparison theorems of Aleksandrov in riemannian geometry (see [BalGS]). Hence the “A”. The “C” and “T” refer to related work of Cartan and Toponogov. Some related notions of non-positive curvature can be found in the work of Busemann dating from the 1940s. A general reference to the topic is [BriH].

For any $\kappa \in \mathbb{R}$ one can define the notion of a “CAT(κ)” or a “locally CAT(κ)” space. One simply replaces the euclidean plane with the simply connected complete riemannian 2-manifold of constant curvature κ . A simply connected locally CAT(κ) space is (globally) CAT(κ). Up to rescaling the metric, there are three case, namely $\kappa = -1, 0, 1$: respectively “hyperbolic”, “euclidean” and “spherical”.

Suppose we construct a polyhedral complex out of convex hyperbolic, euclidean or spherical polyhedra. In the cases of interest to us, these polyhedra will be compact, so we get a genuine cell complex. One can put a path-metric on such a complex such the induced path metric on each polyhedral cell agrees with the original metric on that cell. There are criteria under which such a complex will be geodesic and/or complete. These are discussed in some detail in [BriH], and in

some greater generality in the appendix to [Le]. (For example, if there are only finitely many isometry classes of cells, the metric will be geodesic and complete.) The link of every cell has a natural structure as a spherical polyhedral complex. The local $\text{CAT}(\kappa)$ condition is then equivalent to saying that the link of each cell contains no closed locally geodesic path of length strictly less than 2π . This is in turn equivalent to saying that each component of the link of each 0-cell is (globally) $\text{CAT}(1)$. In the case of a cube complex, this is equivalent to saying that the link of each vertex is flag. This is a combinatorial condition which we have called “locally $\text{CCAT}(0)$ ”.

Examples of locally $\text{CAT}(\kappa)$ spaces are riemannian manifolds of sectional curvature at most κ . A standard reference, for $\kappa \leq 0$, is [BalGS]. Many of the constructions in that setting ultimately rely only on the metric geometry and generalise to $\text{CAT}(0)$ spaces.

A general reference for (locally) $\text{CAT}(0)$ cube complexes is [Ge3]. Further accounts can be found in [NibR1, NibR2]. They have played a major role in geometric group theory in recent years. One of particular note, is the contribution of Wise towards the proof of Thurston’s “virtual fibering conjecture” for hyperbolic 3-manifolds [Wis].

26.19. Notes on Section 19 (Spaces with measured walls).

The extension of a finitely additive function on a ring to a measure on a σ -ring is due to Carathéodory, see for example [Ber]. The completion of a measure can be described by taking the σ -algebra consisting of all sets of the form $A \triangle N$ where A is measurable, and N is a subset of a set of measure 0. The measure of $A \triangle N$ is defined to be equal to that of A . This is also described in [Ber].

The notion of a space with measured walls was defined in [CherMV]. They describe various other natural examples of such spaces. It is a generalisation of the discrete case, namely a “space with walls” as defined in [HagP]. In this case, the duality with discrete median metric spaces was described independently in [ChatteN] and [Nic]. It can be viewed in terms of the combinatorial construction in [Sa].

In the general case, the duality is described in [ChatteDH]. Our treatment is a bit different, but based on broadly similar principles.

One of the principal motivations for studying spaces with measured walls comes from the connection with the Haagerup property in group theory. This is a strong negation of the Kazhdan property (T), and it implies the Baum-Connes Conjecture. A detailed survey of the subject is [CherCJJV].

In particular, the following two results, (A) and (B), are proven in [ChatteDH].

Let G be a locally compact second countable topological group. The action of G on a (pseudo)metric space is “proper” if $\{g \in G \mid B \cap gB \neq \emptyset\}$ is finite for all bounded sets B .

(A) The following are equivalent:

- (1) G has Kazhdan property (T).
- (2) Any continuous isometric action of G on a median metric space has bounded orbits.
- (3) Any continuous action of G by automorphisms of a space with measured walls has bounded orbits in the induced pseudometric.

(B) The following are equivalent:

- (1) G has the the Haagerup property.
- (2) G admits a proper continuous isometric action on a median metric space.
- (3) G admits a proper continuous action by automorphisms on a space of measured walls.

The paper of Genevois [Ge2] considers the set of convex hulls of finite sets. The fact that this is a median metric space is used to show that the “diadem product” of two median metric spaces is a median metric space. This is in turn used, among other things, to give another proof of the result of Cornulier, Stalder and Valette that the Haagerup property is closed under wreath products for countable groups.

I do not know whether complex hyperbolic space in the geodesic metric is submedian (equivariantly or not).

26.20. Notes on Section 20 (Boolean functions and majority vote).

Some discussion of the use of boolean functions and median algebras in computer science is given in [Kn].

One can generalise this discussion to multivalued logic, where we allow a larger set, X , of “truth values”. In this case, a boolean function is generalised to a map $X^n \rightarrow X$ for $n \in \mathbb{N}$. One can define “clones” in a similar way, though the theory gets more complicated. If $\#X \geq 3$, it turns out that there are continuously many clones. One general reference to the theory of clones is [Lau].

Our account of Theorem 20.3.1 is taken from [BanMe], though we have rephrased things a little. That paper gives a more general result, which also allows for the possibility that n might be even. As mentioned in Subsection 20.2, some other calculations of majority vote are described in [TSAHD].

26.21. Notes on Section 21 (Group actions).

Various further results regarding isometric group actions on finite-rank median metric spaces are discussed in [Fi1, Fi2, Fi4]. We have already referred to [Fi4]. The paper [Fi1] gives a “Tits alternative”: either the group contains a nonabelian free subgroup, or else it is “small” in some sense. For example, if the action is free, then the group is virtually finite-by-abelian. The paper [Fi2] studies lattices in such groups. These papers make much use of the generalisation of the Roller boundary constructed in [Fi3], which was mentioned in Section 12.

Many of these results generalise analogous results for cube complexes, see for example [CaS, Hag].

26.22. Notes on Section 22 (Gates).

As mentioned in the Notes to Section 1, related axiom systems for betweenness have given by various authors. The particular formulation we give here (or more precisely, the map $[(a, b) \mapsto [a, b]]$) is often called a “geometric interval operator”. It was introduced in [BanC1], and is discussed in [Vandev] and [Verh].

Our proof of Lemma 22.1.5 is more or less the argument presented in [DreS] in the context of a metric space. There it is asserted that $\omega_A\omega_B$ is a gate map, though this is not necessarily true, as the following example shows.

Consider the 6-element metric space $M = \{a, b, c, d, x, y\}$ with distances 1,2,3 assigned as follows:

- 1 : ab, cd, ac, bd, xy, yb
- 2 : ad, bc, xb, xc, ya, yd
- 3 : xa, xd, yc .

Let $A = \{a, c\}$, $B = \{b, d, y\}$. One readily checks the following: A, B are gated with $A_B = A$ and $B_A = \{b, d\}$, $\omega_Ax = c$, $\omega_Bx = y$ and $\omega_Ay = a$. In particular, $\omega_A\omega_Bx = a \neq c = \omega_{A_B}x$.

We also remark that interpolation is not the only constraint on betweenness in a metric space. For example, suppose we have $2n$ points, a_i, b_i , for $i \in \{1, \dots, n\}$ cyclically ordered. Then the conjunction of $a_i.a_{i+1}.b_{i+1}$ & $b_i.b_{i+1}.a_{i+1}$ for all i implies $a_i.b_i.b_{i+1}$ & $b_i.a_i.a_{i+1}$ for all i . (Note that $\rho(a_{i+1}, b_{i+1}) \leq \rho(a_i, b_i)$, as in Lemma 13.2.1, and so these distances are all equal. Now for any i we have $\rho(b_i, b_{i+1}) + \rho(a_{i+1}, b_{i+1}) = \rho(b_i, a_{i+1}) \leq \rho(a_i, b_i) + \rho(a_i, a_{i+1})$, so $\rho(b_i, b_{i+1}) \leq \rho(a_i, a_{i+1})$. Swapping a, b we deduce that $\rho(a_i, a_{i+1}) = \rho(b_i, b_{i+1})$, and the statement follows.)

As another example, we have

$$x.z.y \ \& \ a.x.y \ \& \ b.y.x \ \& \ x.a.c \ \& \ z.c.a \ \& \ y.b.c \ \& \ z.c.b \Rightarrow a.c.b.$$

These statements also hold in any median algebra (see Examples (Ex6.1) and (Ex6.2) respectively of Subsection 6.2.) Indeed this must be the case: if such a statement failed in a median algebra, it would fail in a finite median algebra, and we could take the combinatorial metric thereon, giving a contradiction.

It is natural to ask what other relations hold in a metric space. I know of no reference to this question.

26.23. Notes on Section 23 (Quasimedial graphs).

Many of the results we give here can be found in some form in [Ge1]. There is some difference of terminology between our presentation and theirs. For example, a “clique” in [Ge1] is what we call here a “maximal clique”. Likewise, a “prism”

there is what we call a “gated prism”. Our strategy for proving Proposition 23.4.11 is somewhat different. The argument in [Ge1] proceeds by analysing the structure of halfspaces. I thank Anthony Genevois for suggesting Proposition 23.6.3 to me.

In the literature, prisms are often referred to as “Hamming graphs”.

In [Wil] it is shown that a graph is quasimedial if and only if it is a retract of a prism. This generalises the result of [Ban], given as Proposition 11.8.4 here.

What we have called a “quasimedial triple” is often simply called a “quasimedial”. To avoid confusion, we prefer to distinguish this from the ternary “quasimedial” operation. We should note that the version given in [BanMuW] is slightly different. There the condition “ $x.a.b.y$ ” is replaced by “ $x.a.y \ \& \ x.b.y$ ” etc. (in our notation in the definition of Subsection 23.1). It is not hard to get between these two formulations.

Quasimedial graphs from the point of view of partitions are explored in [BanHuM]. There the terms “strongly compatible” and “strongly incompatible” are used to mean what we call “nested” and “crossing” respectively. We have adopted terminology consistent with that of walls of a median algebra.

One can generalise the construction which gave rise to CAT(0) spaces starting with a median graph as follows. Let Γ be a quasimedial graph. We can construct a polyhedral complex, $\Delta(\Gamma)$, where the cells are all euclidean prisms in the following way. For each maximal clique, $\Lambda \leq \Gamma$, we construct a regular euclidean simplex of dimension $\#V(\Lambda) - 1$ with unit side-lengths. If this is infinite, we interpret it as a complex built out of all the finite dimensional subsimplices. For each gated prism, $\Pi \leq \Gamma$, we take the l^2 product of the simplices arising from its factors. This is a “euclidean prism”. We now glue these together as dictated by the structure of $\Theta(\Gamma)$. This gives us a polyhedral complex, $\Delta(\Gamma)$, with 1-skeleton, Γ , and 2-cells, $S(\Gamma)$. (If Γ is median, then this agrees with the construction of the cube complex in Section 17.) We can put a path metric on $\Delta(\Gamma)$, similarly as we did for cube complexes in Section 17. It follows from Lemma 23.4.9 that $\Delta(\Gamma)$ is simply connected (in either the metric, or the CW, topology). In fact, it is shown in [Ge1] that $\Delta(\Gamma)$ is CAT(0). (It remains to verify the link condition described in the Notes to Section 18.) In particular, $\Delta(\Gamma)$ is contractible. We remark that we can view $\Delta(\Theta(\Gamma))$ geometrically in these terms. For this, take the barycentric subdivision of each simplex of $\Delta(\Gamma)$. Now take the direct product subdivision for each euclidean prism. This gives us a subdivision of $\Delta(\Gamma)$. We can realise $\Delta(\Theta(\Gamma))$ as a subcomplex of this subdivision, onto which $\Delta(\Gamma)$ retracts.

Various applications of quasimedial graphs to group theory are discussed in [Ge2]. For example, the Cayley graphs of graph products of groups with suitable generating sets are quasimedial. They also arise from wreath products of groups. Some further applications of this theory can be found in [Val].

Quasimedial graphs have also found applications in phylogenetics (see, for example, [BanHuM]). One should imagine a population of individuals. The relevant types of these individuals are described by a finite number of traits (or “characters”), each of these falling into a number of possible classes. For example, the

individuals might be gene sequences, where the relevant traits to record are the bases at particular locations on the gene. In this case, at each location, there are four possible values of the trait (traditionally denoted “ A, G, C, T ”). In the notation of Example (Ex23.7) of Subsection 23.2, X can be thought of as the set of individuals in the population, and I is the set of traits that are being recorded. The partition W_i corresponds to the possible values of a particular trait: each element of W_i is the set of individuals in X which exhibit that given trait value. (So, for example, in the case of gene sequences, I is the set of locations on a gene. Given $i \in I$, W_i would be the partition $\{X_i^A, X_i^G, X_i^C, X_i^T\}$, where $X_i^B \subseteq X$ is the set of individuals with base B at the i th location.) We now take the cartesian product of $\#I$ copies of these traits, to give us a prism, Π , whose factors are indexed by I . (In the example, each factor is the complete graph on $\{A, G, C, T\}$.) The vertex set, $V(\Pi)$, of Π , represents all possible combinations of these traits. The image, $\eta(X) \subseteq V(\Pi)$, is the set of such combinations which actually occur in the population. One wants to study the structure of this set. In evolutionary biology it is reasonable to suppose that this might be well approximated by a tree. In particular, one might search for a “Steiner tree”: a tree which includes $\eta(X)$, which has vertices in $V(\Pi)$, and whose total length (the sum of the distances between each pair of adjacent vertices) is minimal. There are various algorithms for doing this. In [BanR] it is shown that one such algorithm gives a Steiner tree which is guaranteed to lie in the quasimedial hull. This, and other evidence, suggests that studying the structure of quasimedial graphs is a good way to understand the evolutionary relationships between individuals. In our discussion one can allow for infinitely many values of a particular trait. In practice however, there are typically only finitely many. If each trait takes only two possible values, then we are reduced to considering median graphs.

26.24. Notes on Section 24 (Coarse geometry).

A general reference for the background material in geometric group theory is [DruK]. The notion of a hyperbolic space was originally introduced in [Gro1].

A geodesic space is “proper” if it is complete and locally compact. The Schwarz-Milnor Lemma tells us that if a group acts properly discontinuously and cocompactly on a proper geodesic space (which essentially means that it has a compact hausdorff quotient) then the group is finitely generated and quasi-isometric to the original space. For example, the fundamental group of a compact riemannian manifold is quasi-isometric to its universal cover. If the manifold is strictly negatively curved, this will be (Gromov) hyperbolic. If it has constant curvature -1 , then this will be (the classical) hyperbolic space, \mathbb{H}^n . In particular, every closed orientable surface of genus at least 2 admits such a structure.

The notion of an asymptotic cone was introduced in [VandenW] as a means of interpreting Gromov’s theorem that groups of polynomial growth are virtually nilpotent. Gromov has elaborated on this theory in [Gro2]. A detailed account can be found in [DruK].

The notion of a coarse median space was introduced in [Bo2]. An exposition, with additional references, can be found in [Bo6], and some further general discussion is given in [Bo5]. One inspiration for the idea was the description of the centroid map in the mapping class groups as given in [BehM]. This, and other, applications to mapping class groups are based on the “hierarchy” machinery introduced in [MasM]. A number of variations and generalisations of this machinery have been described by various authors since. As we have noted, it can be applied to Teichmüller space in either the Teichmüller metric or the Weil-Petersson metric (see for example the references in [Bo2]). The median structure on the asymptotic cone of the mapping class group was first described in [BehDS].

Any asymptotic cone of a rank-1 coarse median space is 0-hyperbolic, hence an \mathbb{R} -tree. It is also well known that any geodesic metric space for which all asymptotic cones are \mathbb{R} -trees is hyperbolic. This shows that rank-1 coarse median is equivalent to hyperbolicity. A more concrete proof of this fact can be found in [NibWZ1].

Various applications of the theory make use of the notion of a “quasiflat”; that is, a quasi-isometrically embedded euclidean space. We noted that a quasiflat in a coarse median space has dimension at most the (coarse median) rank of the space. The geometry of such maximal dimensional quasiflats is useful for example in proving quasi-isometric rigidity. Passing to the asymptotic cone, one gets a bilipschitz embedding of euclidean space into a median metric space. If this has maximal dimension, one can prove a regularity theorem. This tells us that it has locally the structure of a cube complex (see [Bo5]). A related fact was used to prove the quasi-isometric rigidity of the mapping class group in [BehKMM]. This says that (with a few low-complexity exceptional cases) any self-quasi-isometry of a mapping class group is a bounded distance from the map induced by the left multiplication by some element. (In particular, it is a bounded distance from an isometry.) In fact, this is true if we only assume a-priori that the map is a quasi-isometric embedding [Bo5].

Some further results regarding group actions on coarse median spaces, with some applications, are described in [Fi2].

One can speak of a coarse median space as being “ ν -colourable”: in the definition we gave, we replace “ $\text{rank}(\Pi) \leq \nu$ ” by the requirement that Π be ν -colourable (in the original sense defined in Subsection 8.3). It turns out that the mapping class groups and Teichmüller spaces are finitely colourable in this sense (see [BesBF, Bo2]). It follows that their asymptotic cones are finitely colourable in the original sense of a median algebra. Given Theorem 13.4.1 and Proposition 15.3.1 we recover the fact, proven in [BehDS], that such spaces embed into a finite product of \mathbb{R} -trees. (It is unclear what the optimal value of ν would be for these spaces.)

We remark that one could similarly define a “coarse quasimedial space” using quasimedial graphs in place of median graphs. In view of Proposition 23.4.11, however, it seems that this would give us nothing essentially new, at least in the finite-rank case.

26.25. **Notes on Section 25 (Injective metric spaces and helly graphs).**

Some early references for injective metric spaces are [ArP, Is1]. A more recent account is given in [Lan]. In particular, [Lan] gives another proof of Theorem 25.1.1. The fact that l^∞ metrics on finite-dimensional CCAT(0) cube complexes are injective follows from the result of [MaiT]. See [Mie1] for more discussion.

It is shown in [Lan] that the injective hull of a hyperbolic group with the word metric is a finite dimensional complex, which comes equipped with a cocompact group action. In particular, it is contractible. This therefore gives an alternative construction to the “Rips complex” of a hyperbolic group.

Some equivalent descriptions of a helly graph are described in [BanC2]: see Theorem 3.1 thereof.

Properties of helly graphs are exploited in [ChalCGHO]. The authors define the notion of a helly group, and derive various properties of such groups.

The coarse helly property can be described in terms of injective hulls: see Proposition 3.12 of [ChalCGHO].

REFERENCES

- [AdeN] S.A.Adeleke, P.M.Neumann, *Relations related to betweenness: their structure and automorphisms* : Mem. Amer. Math. Soc. Vol. 623 (1998).
- [Adi] K.Adiprasito, *A note on the simplex-cosimplex problem* : European J. Combin. **66** (2017) 5–12.
- [ArP] N.Aronszajn, P.Panitchpakdi, *Extension of uniformly continuous transformations and hyperconvex metric spaces* : Pacific J. Math. **6** (1956) 405–439.
- [Av] S.P.Avann, *Metric ternary distributive semi-lattices* : Proc. Amer. Math. Soc. **12** (1961) 407–414.
- [BalGS] W.Ballmann, M.Gromov, V.Schroeder, *Manifolds of nonpositive curvature* : Progress in Mathematics Vol 61. Birkhäuser (1985).
- [Ban] H.-J.Bandelt, *Retracts of hypercubes* : J. Graph Theory **8** (1984) 501–510.
- [BanC1] H.-J.Bandelt, V.Chepoi, *A Helly theorem in weakly modular space* : Discrete Math. **160** (1996) 25–39.
- [BanC2] H.-J.Bandelt, V.Chepoi, *Metric graph theory and geometry: a survey* : in “Surveys on discrete and computational geometry” : Contemp. Math. **453** (2008) 49–86.
- [BanHe] H.-J.Bandelt, J.Hedlíková, *Median algebras* : Discrete Math. **45** (1983) 1–30.
- [BanHuM] H.-J.Bandelt, K.T.Huber, V.Moulton, *Quasi-median graphs from sets of partitions* : Discrete Appl. Math. **122** (2002) 23–35.
- [BanMe] H.-J.Bandelt, G.C.Meletioui, *The algebra of majority consensus* : Algebra Universalis **29** (1992) 546–555.
- [BanMuW] H.-J.Bandelt, H.M.Mulder, E.Wilkeit, *Quasi-median graphs and algebras* : J. Graph Theory **18** (1994) 681–703.
- [BanR] H.-J.Bandelt, A.Röhl, *Quasi-median hulls in Hamming space are Steiner hulls* : Discrete Appl. Math. **157** (2009) 227–233.
- [BanV1] H.-J.Bandelt, J.van de Vel, *Superextensions and the depth of median graphs* : J. Combin. Theory Ser. **57** (1991) 187–202.
- [BanV2] H.-J.Bandelt, J.van de Vel, *The median stabilization degree of a median algebra* : J. Algebraic Combin. **9** (1999) 115–127.

- [BarC] J.-P.Barthélemy, J.Constantin, *Median graphs, parallelism and posets* : Discrete Math. **111** (1993) 49–63.
- [BehDS] J.A.Behrstock, C.Druţu, M.Sapir, *Median structures on asymptotic cones and homomorphisms into mapping class groups* : Proc. Lond. Math. Soc. **102** (2011) 503–554.
- [BehKMM] J.A.Behrstock, B.Kleiner, Y.Minsky, L.Mosher, *Geometry and rigidity of mapping class groups* : Geom. Topol. **16** (2012) 781–888.
- [BehM] J.A.Behrstock, Y.N.Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.
- [Ber] S.K.Berberian, *Measure and integration* : Macmillan (1965).
- [Bes] M.Bestvina, *\mathbb{R} -trees in topology, geometry, and group theory* : in “Handbook of geometric topology” (ed. R.J.Daverman, R.B.Sher), North-Holland (2002) 55–91.
- [BesBF] M.Bestvina, K.Bromberg, K.Fujiwara, *Constructing group actions on quasi-trees and applications to mapping class groups* : Publ. Math. Inst. Hautes Études Sci. **122** (2015) 1–64.
- [Bi] G.Birkhoff, *Lattice theory, Third edition* : American Mathematical Society Colloquium Publications, Vol. XXV (1967).
- [BiK] G.Birkhoff, S.A.Kiss, *A ternary operation in distributive lattices* : Bull. Amer. Math. Soc. **53** (1947) 749–752.
- [Bo1] B.H.Bowditch, *Treelike structures arising from continua and convergence groups* : Memoirs Amer. Math. Soc. Volume 662 (1999).
- [Bo2] B.H.Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.
- [Bo3] B.H.Bowditch, *Embedding median algebras in products of trees* : Geom. Dedicata **170** (2014) 157–176.
- [Bo4] B.H.Bowditch, *Some properties of median metric spaces* : Groups, Geom. Dyn. **10** (2016) 279–317.
- [Bo5] B.H.Bowditch, *Large-scale rigidity properties of the mapping class groups* : Pacific J. Math. **293** (2018) 1–73.
- [Bo6] B.H.Bowditch, *Notes on coarse median spaces* : in “Beyond hyperbolicity”, London Math. Soc. Lecture Note Ser. Vol. 454, ed. M.Hagen, R.Webb, H.Wilton, Cambridge Univ. Press (2019) 3–24.
- [Bo7] B.H.Bowditch, *Median and injective metric spaces* : Math. Proc. Camb. Phil. Soc. **168** (2020) 43–55.
- [Bo8] B.H.Bowditch, *Embedding median algebras in median metric spaces* : preprint, Warwick 2022.
- [BriH] M.Bridson, A.Haefliger, *Metric spaces of non-positive curvature* : Grundlehren der Math. Wiss. No. 319, Springer (1999).
- [BrodCGNW] J.Brodzki, S.J.Campbell, E.Guentner, G.A.Niblo, N.J.Wright, *Property A and $CAT(0)$ cube complexes* : J. Funct. Anal. **256** (2009) 1408–1431.
- [Broo] A.D.Brooke-Taylor, *Products of CW complexes* : preprint (2018), posted at arXiv:1710.05296.
- [BrouMMV] A.E.Brouwer, C.F.Mills, W.H.Mills, A.Verbeek, *Counting families of mutually intersecting sets* : Electron. J. Combin. **20** (2013) Paper 8, 8 pp.
- [CaL] P.-E.Caprace, J.Lécureux, *Combinatorial and group-theoretic compactifications of buildings* : Ann. Inst. Fourier **61** (2011) 619–672.
- [CaS] P.-E.Caprace, M.Sageev, *Rank rigidity for $CAT(0)$ cube complexes* : Geom. Funct. Anal. **21** (2011) 851–891.
- [ChalC] J.Chalopin, V.Chepoi, *A counterexample to Thiagarajan’s conjecture on regular event structures* : J. Comput. System Sci. **113** (2020) 76–100.

- [ChalCHO] J.Chalopin, V.Chepoi, H.Hirai, D.Osajda, *Weakly Modular Graphs and Nonpositive Curvature* : Memoirs of the American Mathematical Society, Vol. 1309, American Mathematical Society (2020).
- [ChalCGHO] J.Chalopin, V.Chepoi, A.Gevenois, H.Hirai, D.Osajda, *Helly groups* : preprint (2020), posted at arXiv:2002.06895.
- [Char] R.Charney, *An introduction to right-angled Artin groups* : Geom. Dedicata **125** (2007) 141–158.
- [ChatteD] I.Chatterji, C.Druțu, *Median geometry for spaces with measured walls and for groups* : preprint (2017), posted at arXiv:1708.00254.
- [ChatteDH] I.Chatterji, C.Druțu, F.Haglund, *Kazhdan and Haagerup properties from the median viewpoint* : Adv. Math. **225** (2010) 882–921.
- [ChatteN] I.Chatterji, G.A.Niblo, *From wall spaces to CAT(0) cube complexes* : Internat. J. Algebra Comput. **15** (2005) 875–885.
- [ChattoASGD] A.Chattopadhyay, L.Amarù, M.Soeken, P.-E.Gaillardon, G.De Micheli, *Notes on majority Boolean algebra* : in “2016 IEEE 46th International Symposium on Multiple-Valued Logic”, IEEE Computer Soc. (2016) 50–55.
- [Chep1] V.Chepoi, *Separation of two convex sets in convexity structures* : J. Geom. **50** (1994) 30–51.
- [Chep2] V.Chepoi, *Graphs of some CAT(0) complexes* : Adv. in Appl. Math. **24** (2000) 125–179.
- [ChepH] V.Chepoi, M.F.Hagen, *On embeddings of CAT(0) cube complexes into products of trees via colouring their hyperplanes* : J. Combin. Theory **103** (2013) 428–467.
- [CherCJJV] P.-A.Cherix, M.Cowling, P.Jolissaint, P.Julg, A.Valette, *Groups with the Haagerup property* : Progress in Mathematics, Vol. 197. Birkhäuser (2001).
- [CherMV] P.-A.Cherix, F.Martin, A.Valette, *Spaces with measured walls, the Haagerup property and property (T)* : Ergodic Theory Dynam. Systems **24** (2004) 1895–1908.
- [Chi] I.M.Chiswell, *Generalised trees and Λ -trees* : in “Combinatorial and geometric group theory (Edinburgh, 1993)” London Math. Soc. Lecture Note Ser., Vol. 204, (ed. A.J.Duncan, N.D.Gilbert, J.Howie) Cambridge Univ. Press (1995) 43–55.
- [CoMO] D.E.Cohen, K.Madlener, F.Otto, *Separating the intrinsic complexity and the derivational complexity of the word problem for finitely presented groups* : Math. Logic Quart. **39** (1993) 143–157.
- [Dil] R.P.Dilworth, *A decomposition theorem for partially ordered sets* : Ann. of Math. **51** (1950) 161–166.
- [DisER] W.Dison, E.Einstein, T.Riley, *Taming the hydra: the word problem and extreme integer compression* : Internat. J. Algebra Comput. **28** (2018) 1299–1381.
- [Do] C.H.Dowker, *Topology of metric complexes* : Amer. J. Math. **74** (1952) 555–577.
- [DreS] A.Dress, R.Scharlau, *Gated sets in metric spaces* : Aequationes Math. **34** (1987) 112–120.
- [DruK] C.Druțu, M.Kapovich, *Geometric group theory* : AMS Colloquium Publications Vol. 63, American Mathematical Society (2018).
- [Du] M.J.Dunwoody, *Inaccessible groups and protrees* : J. Pure and Applied Algebra **88** (1993) 63–78.
- [En] R.Engelking, *Theory of dimensions finite and infinite* : Sigma Series in Pure Math. No. 10, Heldermann Verlag, Lemgo, (1995).
- [ErKR] P.Erdős, C.Ko, R.Rado, *Intersection theorems for systems of finite sets* : Quart. J. Math. **12** (1961) 313–320.
- [Ev] E.Evans, *Median lattices and convex subalgebras* : in “Universal algebra (Esztergom, 1977)” , Colloq. Math. Soc. János Bolyai, Vol. 29, North-Holland (1982) 225–240.
- [FaH] J.Faraut, K.Harzallah, *Distances hilbertiennes invariantes sur un espace homogène* : Ann. Inst. Fourier **24** (1974) 171–217.

- [Fi1] E.Fioravanti, *The Tits alternative for finite rank median spaces* : Enseign. Math. **64** (2018) 89–126.
- [Fi2] E.Fioravanti, *Superrigidity of actions on finite rank median spaces* : Adv. Math. **352** (2019) 1206–1252.
- [Fi3] E.Fioravanti, *Roller boundaries for median spaces and algebras* : Algebr. Geom. Topol. **20** (2020) 1325–1370.
- [Fi4] E.Fioravanti, *Convex cores for actions on finite-rank median algebras* : preprint (2021), posted at arXiv:2101.03060.
- [Fi5] E.Fioravanti, *Coarse-median preserving automorphisms* : preprint (2021), posted at arXiv:2101.04415.
- [Ge1] A.Genevois, *Cubical-like geometry of quasi-median graphs and applications to geometric group theory* : preprint (2017), posted at arXiv:1712.01618.
- [Ge2] A.Genevois, *Lamplighter groups, median spaces, and Hilbertian geometry* : preprint (2017), posted at arXiv:1705.00834.
- [Ge3] A.Genevois, *Algebraic properties of groups acting on median graphs* : book in preparation.
- [Gra] A.A.Grau, *Ternary Boolean algebra* : Bull. Amer. Math. Soc. **53** (1947) 567–572.
- [Gre] J.W.Green, *The November meeting in Pasedena* : Bull. Amer. Math. Soc. **54** (1948) 78–85.
- [Gro1] M.Gromov, *Hyperbolic groups* : in “Essays in group theory” Math. Sci. Res. Inst. Publ. No. 8, Springer (1987) 75–263.
- [Gro2] M.Gromov, *Asymptotic invariants of infinite groups* : “Geometric group theory, Vol. 2” London Math. Soc. Lecture Note Ser. No. 182, Cambridge Univ. Press (1993).
- [Grz] A.Grzegorzcyk, *Undecidability of some topological theories* : Fund. Math. **38** (1951) 137–152.
- [Gui] V.Guirardel, *Coeur et nombre d’intersection pour les actions de groupes sur les arbres* : Ann. Sci. École Norm. Sup. **38** (2005) 847–888.
- [Gur] D.P.Guralnik, *Coarse decompositions for boundaries of $CAT(0)$ groups* : Thesis, Technion, Haifa (2005), available at arXiv:math/0611006.
- [HaeHP] T.Haettel, N.Hoda, H.Petyt, *Coarse injectivity, hierarchical hyperbolicity, and semihyperbolicity* : Geom. Topol **27** (2023) 1587–1633.
- [Hag] F.Haglund, *Isometries of $CAT(0)$ cube complexes are semi-simple* : Ann. Math. Qué. **47** (2023) 249–261.
- [HagP] F.Haglund, F.Paulin, *Simplicité de groupes d’automorphismes d’espaces à courbure négative* : in “The Epstein birthday schrift” (ed. I.Rivin, C.Rourke, C.Series) Geom. Topol. Monogr. Vol. 1, Geom. Topol. Publ. (1998) 181–248.
- [HoY] G.J.Hocking, G.S.Young, *Topology* : Addison-Wesley, 1961.
- [HuK] E.V.Huntington, R.J.Kline, *Sets of independent postulates for betweenness* : Trans. Amer. Math. Soc. **18** (1917) 301–325.
- [IdI] K.Idziak, P.M.Idziak, *Decidability problem for finite Heyting algebras* : J. Symbolic Logic **53** (1988) 729–735.
- [Is1] J.R.Isbell, *Six theorems about injective metric spaces* : Comment. Math. Helv. **39** (1964) 65–76.
- [Is2] J.R.Isbell, *Median algebra* : Trans. Amer. Math. Soc. **260** (1980) 319–362.
- [Ka] G.O.H.Katona, *A simple proof of the Erdős-Chao Ko-Rado theorem* : J. Combinatorial Theory **13** (1972) 183–184.
- [KlaM] S.Klavžar, H.M.Mulder, *Median graphs: characterizations, location theory and related structures* : J. Combin. Math. Combin. Comput. **30** (1999) 103–127.
- [KleM] D.Kleitman, G.Markowsky, *On Dedekind’s Problem: The Number of Isotone Boolean Functions. II* : Trans. Amer. Math. Soc. **213** (1975) 373–390.

- [Kn] D.Knuth, *The art of computer programming, Volume 4A, Combinatorial algorithms* : Addison-Wesley (2011).
- [KoM] M.Kolibiar, T.Marciová, *On a question of J. Hashimoto* : Mat. Časopis Sloven. Akad. Vied **24** (1974) 179–185.
- [Lan] U.Lang, *Injective hulls of certain discrete metric spaces and groups* : J. Topol. Anal. **5** (2013) 297–331.
- [Lau] D.Lau, *Function algebras on finite sets. A basic course on many-valued logic and clone theory* : Springer Monographs in Mathematics, Springer-Verlag (2006).
- [Le] I.J.Leary, *A metric Kan-Thurston theorem* : J. Topol. **6** (2013) 251–284.
- [Lu] D.Lubell, *A short proof of Sperner’s lemma* : J. Combinatorial Theory **1** (1966) 299.
- [MaiT] J.H.Mai, Y.Tang, *An injective metrization for collapsible polyhedra* : Proc. Amer. Math. Soc. **88** (1983) 333–337.
- [MasM] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.
- [MayO] J.C.Mayer, L.G.Oversteegen, *A topological characterization of \mathbf{R} -trees* : Trans. Amer. Math. Soc. **320** (1990) 395–415.
- [Me] K.Menger, *Untersuchungen über allgemeine Metrik* : Math. Ann. **100** (1928) 75–163.
- [Mie1] B.Miesch, *Injective metrics in cube complexes* : preprint (2014), posted at arXiv:1411.7234.
- [Mie2] B.R.Miesch, *Gluing constructions and local-to-global results for hyperconvex metric spaces* : Thesis, ETH Zürich (2017).
- [Mil] J.Milnor, *Construction of universal bundles. I.* : Ann. of Math. **63** (1956) 272–284.
- [MoS] J.W.Morgan, P.B.Shalen, *Valuations, trees, and degenerations of hyperbolic structures. I.* : Ann. of Math. **120** (1984) 401–476.
- [Mul1] H.M.Mulder, *The interval function of a graph* : Mathematical Centre Tracts, 132, Mathematisch Centrum, Amsterdam (1980).
- [Mul2] M.Mulder, *n -cubes and median graphs* : J. Graph Theory **4** (1980) 107–110.
- [Mun] J.R.Munkres, *Topology: a first course* : Prentice-Hall (1975).
- [Ne] L.Nebeský, *Graphic algebras* : Comment. Math. Univ. Carolinae **11** (1970) 533–544.
- [NibR1] G.A.Niblo, L.D.Reeves, *Groups acting on $CAT(0)$ cube complexes* : Geom. Topol. **1** (1997) 1–7.
- [NibR2] G.A.Niblo, L.D.Reeves, *The geometry of cube complexes and the complexity of their fundamental groups* : Topology **37** (1998) 621–633.
- [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) Paper No. 20, 50pp.
- [Nic] B.Nica, *Cubulating spaces with walls* : Algebr. Geom. Topol. **4** (2004) 297–309.
- [NielPW] M.Nielsen, G.Plotkin, G.Winskel, *Petri nets, event structures and domains. I.* : Theoret. Comput. Sci. **13** (1981) 85–108.
- [Niem] J.Niemenen, *The ideal structure of simple ternary algebras* : Colloq. Math. **40** (1978/79) 23–29.
- [P] E.L.Post, *The two-valued iterative systems of mathematical logic* : Annals of Mathematics Studies, No. 5, Princeton University Press (1941).
- [R] M.A.Roller, *Poc-sets, median algebras and group actions, an extended study of Dunwoody’s construction and Sageev’s theorem* : Habilitationsschrift, Regensburg, 1998, posted at arXiv:1607.07747.
- [Sa] M.Sageev, *Ends of group pairs and non-positively curved cube complexes* : Proc. London Math. Soc. **71** (1995) 585–617.
- [Sh] M.Sholander, *Medians and betweenness* : Proc. Amer. Math. Soc. **5** (1954) 801–807.

- [SpW] J.Špakula, N.Wright, *Coarse medians and property A* : Algebr. Geom. Topol. **17** (2017) 2481–2498.
- [TSAHD] E.Testa, M.Soeken, L.Amarù, W.Haaswijk, G.De Micheli, *Mapping monotone Boolean functions into majority* : IEEE. Trans. Comput. **68** (2019) 791–797.
- [Val] M.Valiunas, *Acylindrical hyperbolicity of groups acting on quasi-median graphs and equations in graph products* : Groups Geom. Dyn. **15** (2021) 143–195
- [VandenW] 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.
- [Vandev] M.J.L.van de Vel, *Theory of convex structures* : North-Holland Mathematical Library, North-Holland Publishing (1993).
- [Veb] O.Veblen, *A system of axioms for geometry* : Trans. Amer. Math. Soc. **5** (1904) 343–384.
- [Verh] E.R.Verheul, *Multimedians in metric and normed spaces* : CWI Tract, 91, Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam (1993).
- [Verom] R.Veroff, W.McCune, *Otter Version of Don Knuth’s Hand Proof* : webpage, dated 2005, at: https://www.cs.unm.edu/~veroff/MEDIAN_ALGEBRA/knuth_proof.html
- [Wa1] L.E.Ward, *On dendritic sets* : Duke Math. J. **25** (1958) 505–514.
- [Wa2] L.E.Ward, *Axioms for cutpoints* : in “General topology and modern analysis”, Proceedings, University of California, Riverside (ed. L.F.McAuley, M.M.Rao) Academic Press (1980) 327–336.
- [Wil] E.Wilkeit, *The retracts of Hamming graphs* : Discrete Math. **102** (1992) 197–218.
- [Wis] D.T.Wise, *The structure of groups with a quasiconvex hierarchy* : Annals of Mathematics Studies, Vol. 209, Princeton University Press (2021).
- [Ze] R.Zeidler, *Coarse median structures and homomorphisms from Kazhdan groups* : Geom. Dedicata **180** (2016) 49–68.
- [Zv] I.E.Zverovich, *Characterizations of closed classes of Boolean functions in terms of forbidden subfunctions and Post classes* : Discrete Appl. Math. **149** (2005) 200–218.

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