Five axioms for location functions on median graphs

F.R. McMorris
Department of Applied Mathematics, Illinois Institute of Technology
Chicago, IL 60616 USA
and
Department of Mathematics, University of Louisville
Louisville, KY 40292 USA
e-mail: mcmorris@iit.edu

Henry Martyn Mulder
Econometrisch Instituut, Erasmus Universiteit
P.O. Box 1738, 3000 DR Rotterdam, The Netherlands
e-mail: hmmulder@few.eur.nl

Beth Novick
Department of Mathematical Sciences, Clemson University
Clemson, SC 29634 USA
e-mail: nbeth@clemson.edu

R.C. Powers
Department of Mathematics, University of Louisville
Louisville, KY 40292 USA
e-mail: robert.powers@louisville.edu

May 15, 2014

Econometric Institute Report EI 2014-10

Abstract

In previous work, two axiomatic characterizations were given for the median function on median graphs: one involving the three simple and natural axioms anonymity, betweenness and consistency; the other involving faithfulness, consistency and $\frac{1}{2}$-Condorcet. To date, the independence of these axioms has not been a serious point of study. The aim of this paper is to provide the missing answers. The independent subsets of these five axioms are determined precisely and examples provided in each case on arbitrary median graphs. There are three cases that stand out. Here non-trivial examples and proofs are needed to give a full answer. Extensive use of the structure of median graphs is used throughout.
Keywords: axioms; location function; median function; median graph

Mathematics Subject Classification 2010: 05C90, 90B80, 05C75

1 Introduction

The notions of location and consensus can often be considered the same formally. To illustrate this suppose G is a finite connected graph and we have a set $\{1, \ldots, k\}$ of clients (voters). Each client i selects a most suitable, or preferred, location x_i in V, and it is the task of a location (consensus) function to return those vertices that best satisfy various constraints and properties deemed appropriate for the particular problem at hand. In location problems the constraints are usually in the form of optimizing certain criteria. In consensus problems one usually requires certain simple and acceptable rules or axioms that make the voting a reasonable and rational procedure.

One of the early papers on location problems is the classical paper of Witzgall in 1965 [34]. Since then hundreds of papers have been written about location problems on graphs using the geodesic metric, see for example the reference lists in [8, 19, 28, 29, 30, 31, 32]. The earliest paper on the axiomatic study of consensus is the classical paper of Arrow in 1951 [1]. This was the beginning of a fruitful and rich area of research, see for example [2, 3, 5, 6, 7]. Holzman [9] was the first to study a location function as a consensus problem, that is, finding axiomatic characterizations of location functions, thus combining the areas of location and consensus. For some recent work in this area, see [12, 13, 15, 17, 18].

A popular location property is to require the function to return vertices "closest" to the input list of votes or desired locations. Both processes of finding an optimal location or achieving consensus via simple and rational rules can be modeled as a function, which we call a *location function* on the graph G. It is a mapping $L : V^* \longrightarrow 2^V - \{\emptyset\}$, where $2^V - \{\emptyset\}$ denotes the set of all nonempty subsets of V, and V^* is the set of all finite sequences of elements in V. Such a sequence is called a *profile* and denoted by $\pi = (x_1, x_2, \ldots, x_k)$, with k the *length* of the profile. In this paper we focus on the *median function.* Let d be the usual distance function on G, where $d(x, y)$ is the length of a shortest path between x and y. The median function, denoted by

Med, is defined as $Med(\pi) = \{x \in V \mid \sum_{i=1}^{k} d(x, x_i) \text{ is minimum}\}$, for any profile π in V^*. Other location functions that are defined by using different optimization criteria can be found in [4, 12, 13, 18].

A *median graph* is a graph in which $|Med(u, v, w)| = 1$, for every three vertices u, v, w. Prime examples are the trees and the hypercubes. There exists a rich structure theory for these graphs, see [11, 20, 21, 24, 27]. The median function behaves very nicely on median graphs, so these graphs have been a point of focus in the study of the median function.

In [26] an axiomatic characterization of the median function on median graphs is given that involves three simple and basic axioms: anonymity (A) (the location function does not distinguish between clients), betweenness (B) (every location between two clients is equally preferred) and consistency (C) (if two profiles both agree on some output x, then the combined profile has x in its output as well). In the literature simple examples are available to show that (A) & (B) do not imply (C), and that (A) & (C) do not imply (B). But, surprisingly, it was not known whether (A) was independent from (B) & (C). This gap in our knowledge was the primary motivation for this paper. A second axiomatic characterization on median graphs was given in [14]. Besides consistency, it involves the axioms faithfullness (F) (if there is only one client, then the preferred location of this client is returned by the location function) and $\frac{1}{2}$-Condorcet $(Cond)$. For an explanation of this last axiom see Section 4. Because the two sets both characterize the median function on median graphs, these five axioms are not independent. Besides the independence of (A), (B) and (C), we determine exactly which other subsets of the five axioms are actually independent. We do this by providing an example in all relevant cases. We could have restricted ourselves to giving examples on K_2 in each case. But, because of the rich structure theory on median graphs, we aimed at giving an example in every case on an arbitrary median graph.

Three cases stand out. First, it turns out that showing that (B) & (C) do not imply (A) is far from trivial. Here we really have to make full use of the structure of median graphs. Second, the case, whether (C) & $(Cond)$ imply (A) or not, is exceptional. In [22] the 'Meta-Conjecture' was proposed that "every property shared by the trees and the hypercubes is shared by all median graphs". This Meta-Conjecture has proven itself quite fruitful: many interesting results motivated by it were found. The case of (C) & $(Cond)$ related to (A) is the first 'negative' result underscoring this Meta-Conjecture: on arbitrary median graphs (C) & $(Cond)$ do *not* imply (A); on hypercubes they do, but on trees they do not. Again these results were not trivial and lean heavily on the structure of median graphs.

Sections 2 and 3 provide the necessary background. Section 4 discusses the five axioms. Section 5 lists the subsets of independent axioms, and provides an example on every median graph in all cases but three. These three cases are then discussed separately in Sections 6, 7, and 8.

2 Basic Definitions

In this short section we review some of the essential definitions that are needed, some already given in the Introduction. Throughout this paper $G = (V, E)$ is a connected, simple graph. For any two vertices u, v in V, we denote the distance (i.e., the length of a shortest path) between u and v by $d(u, v)$. The *interval* between u and v in G is the set

$$I(u, v) = \{w \mid d(u, w) + d(w, v) = d(u, v)\},$$

in other words, the set of all vertices 'between' u and v.

Let W be a subset of V. Then W is *convex* in G if it contains the vertices of all shortest paths between pairs of vertices in W, that is, $I(u, v)$ is contained in W, for any two vertices u, v in W. Trivially, the intersection of two convex subsets is again convex. The *convex closure* $Con(W)$ of W is the smallest convex set containing W. A subgraph of G is a *convex subgraph* if it is induced by a convex set in G. Let v be a vertex of G. If there is a unique vertex x in W such that x lies in $I(v, w)$, for all w in W, then x is a *gate* for v in W. Note that, if v has a gate in W, then it is the unique vertex in W closest to v. The converse need not be true. Clearly, if v lies in W, then v is its own gate. The subset W is called *gated* if each vertex v of G has a gate in W. A subgraph is a *gated subgraph* if it is induced by a gated set. A gated set is necessarily convex. For arbitrary graphs the converse is not true.

The *hypercube* of dimension n, or n-cube, has the 0,1-vectors of length n as vertices, and two vertices are adjacent if, as vectors, they differ in exactly one position.

A *profile* π on G of length k is a nonempty sequence $\pi = (x_1, x_2, \dots, x_k)$ of vertices of V with repetitions allowed. We call $x_1, x_2, \dots, x_k$ the *elements of* π. Note that π thus has k distinct elements. We call the vertices that occur as elements in π the *vertices of* π. So, if at least one vertex occurs more than once in π, then π has less than k vertices. We denote its length by $k = |\pi|$. When $|\pi|$ is odd, we call π an *odd profile*, otherwise an *even profile*. The *carrier set* $\{\pi\}$ of π is the set of all vertices in π. Let V^* be the set of all profiles. The *concatenation* of the profiles $\pi = (x_1, x_2, \dots, x_k)$ and $\rho = (y_1, y_2, \dots, y_s)$ is the profile $(x_1, x_2, \dots, x_k, y_1, y_2, \dots, y_s)$, and will be denoted by $\pi\rho$. We refer to a profile π, of which the elements are contained in a subgraph H, as a profile *contained in* H, and, abusing notation slightly, we write $\pi \subseteq H$.

A *location function* on G is a function $L : V^* \to 2^V - \{\emptyset\}$, where $2^V - \{\emptyset\}$ denotes the family of all nonempty subsets of V. For convenience, we will usually write $L(x_1, x_2, \dots, x_k)$ instead of $L((x_1, x_2, \dots, x_k))$, for any function L defined on profiles, but will keep the brackets where needed.

A *median* of a profile $\pi = (x_1, x_2, \dots, x_k)$ is a vertex x in V minimizing the distance sum $\sum_{i=1}^{k} d(x, x_i)$. The *median set* $Med(\pi)$ of π is the set of all medians of π. Note that, since G is connected, this defines a location function, namely the *median function* $Med : V^* \to 2^V - \{\emptyset\}$. Trivially, we have $Med(x) = \{x\}$, and $Med(x, y) = I(x, y)$. Moreover, if $I(u, v) \cap I(v, w) \cap I(w, u) \neq \emptyset$, then $Med(u, v, w) = I(u, v) \cap I(v, w) \cap I(w, u)$. The median function behaves very nicely on trees and hypercubes.

3 Preliminaries on Median Graphs

A median graph is a graph G for which $|I(u,v) \cap I(v,w) \cap I(w,u)| = 1$, for any three vertices u, v, w in G. Trees and hypercubes are prime examples of median graphs. Clearly, median graphs are connected and it is a simple exercise to prove that they are bipartite. In median graphs, convex sets are gated, which follows easily from the definition of median graph, using results from [21]. Median graphs possess a beautiful structure and elegant characterizations abound, see e.g. the surveys in [11, 24]. One such characterization is that they are precisely the graphs in which every profile of length 3 has a unique median. The most useful and insightful characterization of median graphs might be the Expansion Theorem in [20]: A graph G is a median graph if and only if G can be obtained from the one-vertex graph K_1 by successive 'convex expansions', explained below. See also [21, 22, 24].

At first sight one might think that median graphs are quite esoteric. But in [10] a one-to-one correspondence was established between the class of connected triangle-free graphs and a special subclass of the class of median graphs. Hence, as median graphs are triangle-free and connected, this implies that in the universe of all graphs, there are as many median graphs as there are connected triangle-free graphs.

Because we make extensive use of the ideas of the Expansion Theorem and the notation therein, we give a survey of the necessary details. These are summarized in Figure 1.

On the left we see the graph G', which is covered by two convex subgraphs G'_1 and G'_2, with intersection the subgraph G'_0, such that there are no edges between $G'_1 - G'_2$ and $G'_2 - G'_1$. Note that G'_0 is convex as well, being the intersection of two convex subgraphs. We call such a covering of G' a *convex cover*. The right hand graph G is constructed from G' as follows. The subgraph G_i is an isomorphic copy of G'_i, for $i = 1, 2$, where λ_i is the isomorphism. We call λ_1 and λ_2 the *lift maps* that lift G'_1 and G'_2 up to G_1 and G_2 respectively. Thus G'_0 is lifted up to G_{01} and G_{02}, respectively. Now we insert edges between the corresponding vertices of G_{01} and G_{02}, thus producing the matching F_{12}, which induces an isomorphism between G_{01} and G_{02}. Note that the subgraphs G_1, G_2, G_{01} and G_{02} are all convex subgraphs of G. We call G the *convex expansion* of G' with respect to the convex cover G'_1, G'_2. It is straightforward to prove that, if G' is a median graph, then G is a median graph as well. Hence all graphs obtained from the one-vertex graph K_1 by a succession of convex expansions are median graphs.

The converse in the Expansion Theorem is the difficult part of the proof. The main steps in this proof are as follows. Let G be a median graph, see Figure 1. Take an arbitrary edge u_1u_2 in G. Let G_1 be the subgraph consisting of all vertices closer to u_1 than to u_2, and let G_2 be the subgraph consisting of all vertices closer to u_2 than u_1. Since G is bipartite, these two subgraphs partition the vertex set of G. Then we prove that these two subgraphs are convex. Let G_{01} be the subgraph of G_1 consisting of the vertices having a neighbor in G_2, and let G_{02} be the subgraph of G_2 consisting of the vertices having a neighbor in G_1. Let F_{12} be the set of edges between G_{01} and

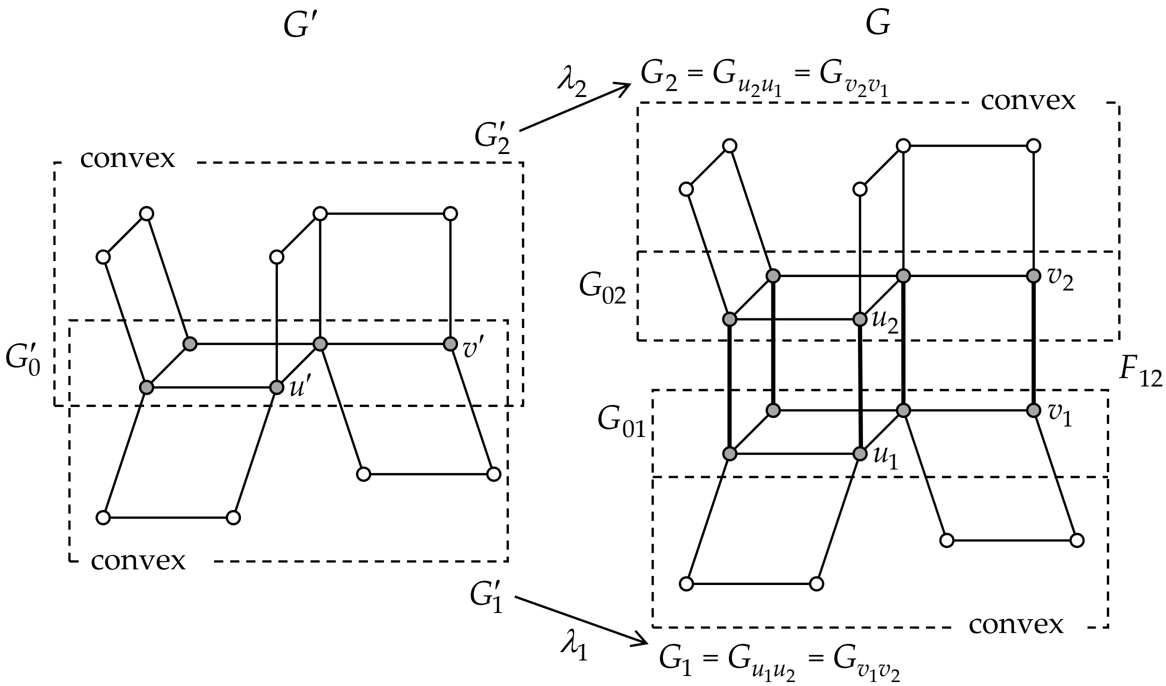

Figure 1: Expansion

G_{02}. Next we prove that F_{12} is a matching that induces an isomorphism between G_{01} and G_{02} as depicted in Figure 1. Moreover, G_{01} and G_{02} are convex subgraphs as well. We call G_1, G_2 a *split* of G with *split sides* G_1 and G_2. We will call G_1 and G_2 *opposides* of each other. For the origin of this neologism see [24].

Next we prove that any other edge of F_{12}, say v_1v_2, defines the same split G_1, G_2, that is, G_1 consists of all vertices closer to v_1 than to v_2 and G_2 consists of all vertices closer to v_2 than to v_1. Note that therefore we do not have to refer to a specific edge when we consider a split G_1, G_2. Now we contract the edges in F_{12} and identify the corresponding vertices in G_{01} and G_{02}, thus producing the graph G' on the left. We can state this more formally by defining the graph homomorphism $\kappa : G \to G'$ by setting $\kappa|_{G_1} = \lambda_1^{-1}$ and $\kappa|_{G_2} = \lambda_2^{-1}$. The mapping κ is the *contraction map*. We call G' the *contraction* of G with respect to the split G_1, G_2.

The last step of the proof is then to show that G' is a median graph with convex cover G'_1, G'_2. It is clear that then G is the convex expansion of G' with respect to this cover. Hence *any* median graph can be obtained this way. In constructing G from K_1 by convex expansions, we may take the expansions in any order. The Expansion Theorem gives us a very strong tool: on median graphs we can use induction on the number of expansions, or, equivalently, the number of splits.

We need some more facts that follow from the result above. Let G_1, G_2 be a split in a median graph G, and let uv be an edge between the two sides with u in G_1 and v in G_2. Then v is the gate for u in G_2 and u is the gate for v in G_1. We call u and v *mates*. If x is in, say, G_2, and its gate in G_{02} is v, then its gate in G_1 is the mate u of v in G_{01}. For an edge uv, we will denote the side of the split closer to u than to v by G_{uv} and the side closer to v than to u by G_{vu}.

Let H_1, H_2 be another split. If $G_i \cap H_j \neq \emptyset$, for every i and j, then we say that the two splits *cross*, or that they are *crossing splits*. In Figure 2 in Section 6, we see two crossing splits in the left hand graph. Note that, in the case of crossing splits, there must be a 4-cycle on which the two splits “cross”, that is, two opposite edges of the 4-cycle are between G_1 and G_2 and the other two opposite edges are between H_1 and H_2. We can see this easily in Figure 2. We also say that the split G_1, G_2 *crosses* H_1 as well as H_2.

Let $\pi = (x_1, x_2, \ldots, x_k)$ be a profile on a median graph G, and let G_1, G_2 be a split in G. Denote by π_i the subprofile of elements in π that lie in G_i, for $i = 1, 2$. We call G_1, G_2 a *balanced split* of π if $|\pi_1| = |\pi_2|$, and say that π is *balanced* on that split. We call the split *unbalanced* if $|\pi_1| \neq |\pi_2|$, or equivalently, that π is *unbalanced* on the split. In this case, G_1 is the *majority side* if $|\pi_1| > |\pi_2|$. Now contract G with respect to the split G_1, G_2 to obtain the median graph G' with convex cover G'_1, G'_2, using the contraction map κ. The *contracted profile* π' is the profile $\pi' = (x'_1, x'_2, \ldots, x'_k)$ in G', where $x'_i = \kappa(x_i)$, for $i = 1, 2, \ldots, k$. Then π'_i is the subprofile of π' consisting of the elements in π' that lie in G'_i, for $i = 1, 2$. If the split is defined with respect to an edge uv, so it is G_{uv}, G_{vu}, then we denote the subprofile of π in G_{uv} by π_{uv}. Similarly for π_{vu}.

In the sequel we will use the notation and ideas introduced above without further mention.

An important consequence of the Expansion Theorem was proved in [23] and [16]. The median set of a profile is always contained in the majority side of an unbalanced split, and it intersects both sides of a balanced split. This is made more precise in the following theorem, which is basic for almost all the proofs in this paper. It says that the median set of a profile is just the intersection of the majority sides with respect to the profile. For convenience, we abuse notation and let G_i indicate the set of vertices of the side G_i as well. We use the following convention: taking an empty intersection results in the whole set, that is, $\bigcap \emptyset = V$.

Theorem 1 *Let G be a median graph and let π be a profile on G. Then*

$$Med(\pi) \;=\; \bigcap \{\, G_1 \mid G_1, G_2 \textit{ is a split with } |\pi_1| > |\pi_2| \}$$

Note that, split sides being convex, this means that median sets are necessarily convex. It is a well-known fact that odd profiles have a unique median in median graphs. This also follows easily from Theorem 1. Let π be an odd profile. Then there are no balanced splits for π. So, if u is a median vertex and v is a neighbor of u, then G_{uv} must be the majority side. Hence G_{vu} is a minority side in the split G_{uv}, G_{vu}, and thus v is not in $Med(\pi)$. Since $Med(\pi)$ is convex, it consist only of u.

A simple corollary of Theorem 1 concerns intervals. This had already been proved in [27] but without any reference to profiles and the median function. Recall that $I(x, y) = Med(x, y)$ in any connected graph.

Corollary 2 *Let G be a median graph and let x and y be vertices of G. Then*

$$I(x, y) \;=\; \bigcap\{\, G_1 \mid G_1, G_2 \textit{ is a split with } x, y \in G_1\}$$

In the sequel we need a result that is essentially contained in the theory developed in [16], but has never been stated explicitly. So we present it here as a lemma with a proof.

Lemma 3 *Let $G = (V, E)$ be a median graph, and let u and v be two vertices of G. Then $V = I(u, v)$ if and only if the profile $\pi = (u, v)$ is balanced on all splits.*

Proof. First let $V = I(u, v)$. Take any split G_1, G_2. If this split were unbalanced, then, by Theorem 1, this would contradict $Med(u, v) = V = I(u, v)$.

Next assume that $\pi = (u, v)$ is balanced on all splits. Then, in Theorem 1, we have the intersection $\cap\emptyset$. So $V = Med(\pi) = I(u, v)$. □

4 The Axioms and Fundamental Theorems

The following are some axioms that we might expect a well-behaved location function to enjoy. The first four below are easily seen to be satisfied by Med on any finite connected graph G, while the last one requires G to be a median graph in order to make the definition well-defined. Below each definition we give a brief description. Let L be a location function on $G = (V, E)$.

(F) **Faithfulness:** $L(x) = \{x\}$, for all $x \in V$.

If there is only one client, then the preferred location of this client is returned by the location function.

(A) **Anonymity:** For any profile $\pi = (x_1, x_2, \ldots, x_k)$ on V and any permutation σ of $\{1, 2, \ldots, k\}$, we have $L(\pi) = L(\pi^\sigma)$, where $\pi^\sigma = (x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(k)})$.

The location function cannot distinguish among the group of clients.

(C) **Consistency:** If $L(\pi) \cap L(\rho) \neq \emptyset$ for profiles π and ρ, then

$L(\pi\rho) \;=\; L(\pi) \;\cap\; L(\rho).$

Loosely speaking, if two profiles agree on some output x, then x is in the output of the concatenation as well.

(B) **Betweenness:** $L(u,v) = I(u,v)$, for $u, v \in V$.

All locations between exactly two preferred locations are equally good.

$(Cond)$ $\frac{1}{2}$**-Condorcet:** $u \in L(\pi)$ if and only if $v \in L(\pi)$, for each profile π on G and any edge uv of G with $|\pi_{uv}| = |\pi_{vu}|$.

The location function does not make a distinction between mates that are equally preferred.

It is a simple exercise to prove that axioms (B) and (C) imply (F) on any graph: just observe that (B) and (C) imply that $L(x) = L(x) \cap L(x) = L(x,x) = I(x,x) = \{x\}$.

The median function Med satisfies trivially both (A) and (B). It is probably now folklore that Med also satisfies (C), see e.g. [16, 17] for a simple proof. This raised the natural question for which connected graphs the median function Med is actually characterized by the three simple, basic axioms (A), (B) and (C). In general this seems to be a difficult question. In 1998 [16], as a corollary to their main result, McMorris, Mulder and Roberts established that the answer is affirmative for all trees (see Vohra [33] for a continuous version of this result). Mulder and Novick [25] proved that it is also the case on hypercubes using techniques that are specific for hypercubes. As remarked in the Introduction, back in 1990 Mulder [22] stated the following 'Meta-Conjecture': "Any (sensible) property that is shared by the trees and the hypercubes is shared by all median graphs". This provided the inspiration to prove the next result in [26].

Theorem 4 *Let L be a location function on the median graph G. Then $L = Med$ if and only if L satisfies (A), (B), and (C).*

Before Theorem 4 was proved the following companion theorem, found in [14], had been established.

Theorem 5 *Let L be a location function on the median graph G. Then $L = Med$ if and only if L satisfies (F), (C), and $(Cond)$.*

These two theorems imply that Med necessarily satisfies all five axioms on median graphs. It is interesting that two groups of three axioms provide sufficiency for a location function to be the median function on median graphs. Looking at Theorem 5, we see that (A) and (B) are not used explicitly, but follow obviously as a consequence of Theorem 4, because Med satisfies these axioms. Reversing the point-of-view, (F) and $(Cond)$ do not appear in Theorem 4, but follow as a consequence of Theorem 5. This prompts us to ask for direct proofs of these implications that do not use the median function as an intermediate step.

Theorem 6 *Let L be a location function on the median graph G.*

(i) *If L satisfies (A), (B) and (C), then L satisfies (F) and $(Cond)$.*

(ii) *If L satisfies (F), (C) and $(Cond)$, then L satisfies (A) and (B).*

Proof. As observed above, (B) and (C) imply (F). The direct proof that (A), (B) and (C) imply $(Cond)$ is the crucial Lemma 10 in [26]. Thus part (i) requires no new proof.

For part (ii), assume L satisfies (F), (C) and $(Cond)$. Let $\pi = (x_1, \ldots, x_k)$ be a profile on G and let $\rho = (x_{\sigma(1)}, \ldots, x_{\sigma(k)})$, where σ is a permutation of $\{1, 2, \ldots, k\}$. To show that L satisfies (A) we must show that $L(\pi) = L(\rho)$. Assume to the contrary that, after possibly interchanging the roles of π and ρ, there exist $u, v \in V$ such that $u \in L(\pi)$ but $v \in L(\rho) - L(\pi)$.

Since π and ρ share the same elements, any split of G that is balanced with respect to π, is balanced with respect to ρ, and vice versa. And, for any other (non-balanced) split the majority sides of π and ρ coincide.

Fix a uv-geodesic P and let y be the first vertex on P that is not in $L(\pi)$. Let x be the vertex preceding y on P and consider the split G_{xy}, G_{yx}. Note that $u \in G_{xy}$ and $v \in G_{yx}$. Suppose this split is not balanced. In the proof of Theorem 5 in [14] we showed, employing all three of the axioms for L, that if G_1, G_2 is an unbalanced split for π with G_1 the majority side, then $L(\pi) \subseteq G_1$. Since $u \in L(\pi)$ it now follows that G_{xy} is the majority side for π. Yet by the same argument, since $v \in L(\rho)$, it follows that G_{yx} is the majority side for ρ. This contradiction shows that this split is in fact balanced. But then $x \in L(\pi)$ and $y \notin L(\pi)$ contradicts the $(Cond)$ property of L. This contradiction settles the proof that L satisfies (A).

We now show that L satisfies axiom (B). Let u and v be vertices in V and $\pi = (u, v)$. If all splits are balanced, then, by Lemma 3, we have $V = I(u, v)$. So, by $(Cond)$, we have $L(\pi) = V = I(u, v)$. Let G_1, G_2 be an unbalanced split such that $u, v \in G_1$. Then, $L(\pi) \subseteq G_1$ again from the proof of Theorem 5 in [14]. By Corollary 2, we have $L(u, v) \subseteq I(u, v)$.

Suppose there exists $x \in I(u, v) - L(u, v)$, and consider a path $ux_1 \cdots x \cdots x_m v$. Each split determined by the edges $ux_1, x_1x_2, \ldots, x_m v$ is balanced with respect to $\pi = (u, v)$. Since $x \notin L(u, v)$, property $(Cond)$ implies that $u, v \notin L(u, v)$. Since $L(u, v) \cap I(u, v) \neq \emptyset$, take $y \in L(u, v) \cap I(u, v)$ and consider the path $uy_1 \cdots y \cdots y_n v$. Again, the splits determined by the edges in this path are balanced with respect to π so $(Cond)$ and $y \in L(u, v)$ give $u, v \in L(u, v)$, a contradiction. Therefore $L(u, v) = I(u, v)$ and the proof is complete. $\square$

Since (B) and (C) imply (F), we can state the following.

Corollary 7 *Let L be a location function on the median graph G. Then $L = Med$ if and only if L satisfies (B), (C), and $(Cond)$.*

5 Independent Sets Among the 5 Axioms

Theorems 4 and 5 and Corollary 7 provide three axiomatic characterizations of the median function on median graphs. For these results to be tight one would also like to know whether the three sets of axioms in these characterizations are independent sets of axioms. Clearly the set of all five axioms (A), (B), (C), (F) and $(Cond)$ does not form an independent set. We have the trivial implication that (B) and (C) imply (F). Moreover Theorems 4 and 5 provide the implications given in Theorem 6. What about other implications? Otherwise formulated: what are the independent sets of axioms among the five given ones?

In the literature it was already observed that (A) and (B) do not imply (C) and that (A) and (C) do not imply (B). But so far it was not known whether (B) and (C) would imply (A) or not. This will be answered in Section 6. In this paper we want to establish exactly which subsets of the five form an independent set of axioms. Because of the rich structure of median graphs, we wanted to have an example for *every* case and on *any* median graph.

Note that there is only one location function on the one vertex graph K_1, the *trivial median graph*. So in this case anything holds. Hence, throughout we let $G = (V, E)$ be a median graph with at least one split. The smallest such median graph, of course, is K_2. In many examples we consider a fixed split G_1, G_2 which we call the *determining split*. Note that in all the cases where we use a determining split, we have as many different examples as there are different splits in the median graph. L will always denote a location function on G, and $\pi = (x_1, x_2, \ldots, x_k)$ a profile of length k on G with $k \geq 1$. Recall that, abusing notation we are letting G_i denote the vertex set of the side G_i as well.

1. L **satisfies exactly one of the axioms.**

   (i) (A):

$$L(\pi) = \begin{cases} G_1 \text{ if } k = 2, \\ V \text{ if } k \neq 2. \end{cases}$$

   Clearly, L satisfies (A) but not (F). Take two mates x and y with respect to the determining split, with x in G_1 and y in G_2. Then $L(x, y) = G_1$, so L does not satisfy (B) or $(Cond)$. For any vertex x, we have $L(x) = V$; but $L(x, x) = G_1$. Thus L does not satisfy (C) either.

   (ii) (B):

$$L(\pi) = \begin{cases} V \text{ if } \pi = (x_1) \\ I(x_1, x_2) \text{ if } \pi = (x_1, x_2) \\ \{x_1\} \text{ if } k \geq 3. \end{cases}$$

Clearly, L satisfies (B) but not (F). The profiles of length at least 3 prohibit L to be anonymous. Take two neighbors a and b. The profile (a, a, b, b) makes L to fail $(Cond)$. Since $L(a, b) = I(a, b) = \{a, b\}$, the profile (a, b, a, b) makes L to fail (C).

(iii) (C):

$$L(\pi) = \begin{cases} G_1 \text{ if } x_1 \in G_2 \\ G_2 \text{ if } x_1 \in G_1. \end{cases}$$

Clearly, L does not satisfy (F) or (A). The profile (x_1, x_2), with both x_1 and x_2 on the same side of the determining split, shows that L does not satisfy (B). Let x_i be in G_i, for $i = 1, 2$, such that they are mates (i.e. adjacent). Then the profile $\pi = (x_1, x_2)$ is balanced on the determining split, x_2 is in $L(\pi)$, but x_1 is not. So L does not satisfy $(Cond)$. To show consistency, take two profiles ρ_1 and ρ_2 with $L(\rho_1) \cap L(\rho_2) \neq \emptyset$. Then their first element must be on the same side. Without loss of generality assume this is G_1. But then both concatenations $\rho_1\rho_2$ and $\rho_2\rho_1$ have their first element in G_1. So the output of L, for all four profiles ρ_1, ρ_2, $\rho_1\rho_2$ and $\rho_2\rho_1$ is G_2.

(iv) $(Cond)$:

$$L(\pi) = \begin{cases} V \text{ if } \pi \text{ is balanced on some split} \\ G_1 \text{ if } \pi \text{ is unbalanced on all splits and } x_1 \in G_2 \\ G_2 \text{ if } \pi \text{ is unbalanced on all splits and } x_1 \in G_1. \end{cases}$$

Clearly, L satisfies $(Cond)$, but not (F) or (A). Any profile of length two with both elements on one side of the determining split avoids L to satisfy (B). Take the profiles (a, b, b) and (a) with a in G_1 and b in G_2. Then $L(a, b, b) = L(a) = G_2$, but $L(a, b, b, a) = V$. So L is not consistent.

(v) (F):

$$L(\pi) = \begin{cases} \{x_1\} \text{ if } k \neq 3, \\ \{\pi\} \text{ if } k = 3. \end{cases}$$

Clearly, L satisfies (F), but not (A) or (B). Let a and b be two distinct vertices. Then we have $L(a) = L(a, b) = \{a\}$, but $L(a, a, b) = \{a, b\} \neq \{a\}$. This destroys consistency.

2. L **satisfies exactly two of the axioms.**
Because (B) and (C) imply (F), there are 9 possibilities for this case. Recall that G_1, G_2 is the determining split.

(i) (A) & (B):

$$L(\pi) = \begin{cases} I(x_1, x_2) \text{ if } \pi = (x_1, x_2) \\ G_1 \text{ if } k \neq 2. \end{cases}$$

Clearly, L satisfies (A) and (B), but not (F). Take x_1 and x_2 both in G_2. Then $L(x_1) = L(x_2) = G_1$, but $L(x_1, x_2) = I(x_1, x_2) \subseteq G_2$. So L is not consistent. Take adjacent vertices a and b with a in G_1 and b in G_2. Then $L(a, a, b, b) = G_1$, so it contains a but not b, whence L does not satisfy $(Cond)$.

(ii) (A) & (C):

$$L(\pi) = G_1 \text{ for all } \ \pi.$$

Clearly L satisfies (A) and (C), but does not satisfy (B), (F) or $(Cond)$.

(iii) (A) & $(Cond)$:

$$L(\pi) = \begin{cases} G_1 \text{ if the split } G_1, G_2 \text{ is unbalanced} \\ V \text{ if the split } G_1, G_2 \text{ is balanced.} \end{cases}$$

Clearly, L satisfies (A) and $(Cond)$, but not (F). For a in G_2, we have $L(a, a) = G_1$, so L does not satisfy (B). With x_i in G_i, for $i = 1, 2$, we have $L(x_1) = L(x_2) = G_1$, but $L(x_1, x_2) = V$. Hence L is not consistent.

(iv) (A) & (F):

$$L(\pi) = \begin{cases} G_1 \text{ if } \pi = (x_1, x_2) \\ \{\pi\} \text{ for } k \neq 2. \end{cases}$$

Clearly, L satisfies (A) and (F) but not (B). Take two mates a and b with a in G_1 and b in G_2. Then $L(b) = \{b\}$, which lies in G_2, but $L(b, b) = G_1$, whence L does not satisfy (C). Moreover, a lies in $G_1 = L(a, b)$, but b does not lie in $L(a, b)$, so that L does not satisfy $(Cond)$ either.

(v) (B) & $(Cond)$:

$$L(\pi) = \begin{cases} I(x_1, x_2) \text{ if } k = 2 \\ V \text{ if } k \neq 2 \text{ and the split } G_1, G_2 \text{ is balanced} \\ G_2 \text{ if } k \neq 2 \text{ and the split } G_1, G_2 \text{ is unbalanced, with } x_1 \in G_1 \\ G_1 \text{ if } k \neq 2 \text{ and the split } G_1, G_2 \text{ is unbalanced, with } x_1 \in G_2. \end{cases}$$

Clearly L satisfies (B), but neither (A) nor (F). Take a in G_1. Then $\{a\} = I(a, a) = L(a, a) = L(a, a) \cap L(a, a) \neq G_2 = L(a, a, a, a)$. So L is not consistent. For $\pi = (x_1, x_2)$, the balanced splits are precisely those that cross the interval $I(x_1, x_2)$. So L satisfies $(Cond)$ for profiles of length 2. For any profile of different length that is balanced on the split G_1, G_2, we have $L(\pi) = V$. So L satisfies $(Cond)$ for these profiles trivially. For any profile of different length, for which the split G_1, G_2 is unbalanced, the output $L(\pi)$ is one of the split sides, so again L satisfies $(Cond)$.

(vi) (B) & (F):

$$L(\pi) = \begin{cases} Med(\pi) \text{ if } k = 1, 2 \\ G_1 \text{ if } k \geq 3 \text{ and } x_1 \in G_2 \\ G_2 \text{ if } k \geq 3 \text{ and } x_1 \in G_1. \end{cases}$$

Clearly, L satisfies (B) and (F), but not (A). Take a in G_1. Then $\{a\} = L(a) = L(a,a) = L(a) \cap L(a,a) \neq G_2 = L(a,a,a)$. So L is not consistent. Let a and b be mates with a in G_1 and b in G_2. Then $L(a,a,b,b) = G_2$, whence it contains b, but it does not contain a. So L fails $(Cond)$.

(vii) (C) & (F):

$$L(\pi) = \{x_1\}.$$

That is, L is the projection on the first position. Clearly, L is faithful, but not anonymous. For distinct vertices a and b, we have $L(a,b) = \{a\} \neq I(a,b)$. So L does not satisfy (B), but also not $(Cond)$. If ρ_1 and ρ_2 two profiles with $L(\rho_1) \cap L(\rho_2) \neq \emptyset$, then their first element must be the same, say a. But then both concatenations $\rho_1\rho_2$ and $\rho_2\rho_1$ have a as their first element. So the output of L, for all four profiles ρ_1, ρ_2, $\rho_1\rho_2$ and $\rho_2\rho_1$ is $\{a\}$, by which consistency is proved.

(viii) $(Cond)$ & (F):

$$L(\pi) = \begin{cases} \{x_1\} \text{ if } k = 1 \\ V \text{ if } k \geq 2 \text{ and the split } G_1, G_2 \text{ is balanced} \\ G_1 \text{ if } k \geq 2 \text{ and the split } G_1, G_2 \text{ is unbalanced, with } x_1 \in G_2 \\ G_2 \text{ if } k \geq 2 \text{ and the split } G_1, G_2 \text{ is unbalanced, with } x_1 \in G_1. \end{cases}$$

Clearly, L is faithful, but not anonymous. Let a be a vertex in G_1. Then $L(a) = \{a\}$, but $L(a,a) = G_2$, so L is not consistent, and does not satisfy betweenness. Profiles of length 1 have no balanced splits, so $(Cond)$ is satisfied for such profiles. Let π be a profile of length at least 2. If G_1, G_2 is balanced, then $L(\pi) = V$. Hence for such profiles $(Cond)$ is satisfied. If G_1, G_2 is unbalanced, then this split cannot carry any output to the other side using $(Cond)$. So again $(Cond)$ is satisfied.

(ix) (C) & $(Cond)$: This is an exceptional case which is discussed in Sections 7 and 8.

3. L **satisfies exactly three of the axioms.**
   Due to Theorem 6 and Corollary 7, there are only seven combinations of three axioms that do not imply a fourth one.

(i) (A) & (B) & $(Cond)$:

$$L(\pi) = \begin{cases} I(x_1, x_2) \text{ if } k = 2 \\ V \text{ otherwise.} \end{cases}$$

Clearly, L satisfies (A) and (B) but not (F). For $\pi = (x_1, x_2)$, the balanced splits are precisely those that cross the interval $I(x_1, x_2)$. So $(Cond)$ satisfies for profiles of length two. For the other profiles of different length $(Cond)$ is satisfied trivially. For any vertex a, then $L(a) = V$ but $L(a, a) = \{a\}$, and so (C) is not satisfied.

(ii) (A) & (B) & (F):

$$L(\pi) = \begin{cases} M(\pi) \text{ if } k = 1, 2 \\ G_1 \text{ otherwise.} \end{cases}$$

Clearly, L satisfies (A), (B) and (F). Take two mates a and b with a in G_1 and b in G_2. Then $\{b\} = L(b) = L(b, b)$, which lies in G_2 but $L(b, b, b) = G_1$, so L does not satisfy (C). Moreover, the profile (a, b, a, b) is balanced on G_1, G_2, but $L(a, b, a, b) = G_1$, so L does not satisfy $(Cond)$.

(iii) (A) & (C) & $(Cond)$:

$$L(\pi) = V.$$

Clearly, L satisfies (A), (C) and $(Cond)$, but not (F). Take any vertex a. Then $L(a, a) = V \neq \{a\}$, so that L does not satisfy (B).

(iv) (A) & (C) & (F): Fix a vertex y.

$$L(\pi) = \begin{cases} \{x_1\} \text{ if } |\{\pi\}| = 1 \\ \{y\} \text{ if } |\{\pi\}| \geq 2. \end{cases}$$

Clearly L satisfies (A) and (F), but not (B) or $(Cond)$. We call a profile ρ with $|\{\rho\}| = 1$ a *constant profile*. To prove consistency, let ρ_1 and ρ_2 be profiles with $L(\rho_1) \cap L(\rho_2) \neq \emptyset$. Then there are three possibilities. First, both ρ_1 and ρ_2 are constant containing the same vertex x. Then the concatenation is of the same type, and we are done. Second, both ρ_1 and ρ_2 are not constant, whence also the concatenation is also not constant, and again we are done. Finally, one is constant and the other is not. To have a non-empty intersection, the constant profile must consist of y's, and again we are done.

It is interesting to note here that Roberts in [30] showed that the axioms of anonymity, consistency, and faithfulness, along with another simple axiom that is not discussed in our paper, completely characterize the location function based on plurality-rule. This was done on *any* finite connected graph. See [30] for the definition of the plurality function and for the statement and proof of this characterization result.

(v) (A) & $(Cond)$ & (F):

$$L(\pi) = \begin{cases} \{x_1\} \text{ if } k = 1 \\ V \text{ otherwise.} \end{cases}$$

Clearly, L satisfies (A) and (F). Since $(Cond)$ always involves two distinct vertices, $(Cond)$ is also trivially satisfied. For any vertex a, we have $L(a) = \{a\}$ and $L(a,a) = V \neq \{a\}$. So (B) and (C) are not satisfied.

(vi) (B) & $(Cond)$ & (F):

$$L(\pi) = \begin{cases} M(\pi) \text{ if } k = 1,2 \\ G_1 \text{ if } k \geq 3, \text{ and } G_1, G_2 \text{ is unbalanced, and } x_1 \in G_1 \\ G_2 \text{ if } k \geq 3, \text{ and } G_1, G_2 \text{ is unbalanced, and } x_1 \in G_2 \\ V \text{ if } k \geq 3 \text{ and } G_1, G_2 \text{ is balanced.} \end{cases}$$

Note that this example is a slight variation of Example 2(viii) above. The only difference is that now we need L to satisfy (B). Clearly, L satisfies (B) and (F), but not (A). That L satisfies $(Cond)$ follows as in previous cases. Take two mates a and b with a in G_1 and b in G_2. Then $L(a,b) = \{a,b\}$, and $L(b) = \{b\}$, but $L(a,b,b) = G_1$, so that L does not satisfy (C).

(vii) (B) & (C) & (F): This case requires a much more elaborate example and accompanying proof, which is given in Section 6.

4. L **satisfies exactly four of the axioms.**
There is only one possibility here: (C) is excluded. For convenience, we write $Con(\pi)$ instead of $Con(\{\pi\})$, for the convex closure of $\{\pi\}$.

(i) (A) & (B) & (F) & $(Cond)$:

$$L(\pi) = Con(\pi).$$

Clearly L satisfies (A). We have $L(x) = Con(x) = \{x\}$ and $L(x,y) = Con(x,y) = I(x,y)$. So L satisfies (F) and (B). To prove that L satisfies $(Cond)$, let $k \geq 3$, and let H_1, H_2 be any split crossing $Con(\pi)$. If π is unbalanced on this split nothing has to be proved, so let π be balanced on this split. Let a and b be mates with a in H_1 and b in H_2, say, with a in $Con(\pi)$. Since the split crosses $Con(\pi)$, there is an element y of π in H_2. Since the gate for a in H_2 is b, it follows that b lies in $I(a,y) \subseteq Con(\pi)$, and we are done. That (C) fails follows from $L(a) = \{a\}$ and $L(a,b) = \{a,b\} = L(a,a,b)$.

6 (B) and (C) do not imply (A) on median graphs with at least one split

In this section we give an example of a location function on an arbitrary, non-trivial median graph that satisfies (B) and (C) but not (A). As observed above in Section 4, such a location function satisfies (F) but not $(Cond)$.

Let $G = (V, E)$ be a median graph with at least one split, so that it is not K_1. Let $\pi = (x_1, x_2, \ldots, x_k)$ be a profile on G. A *pair* of π is a pair of consecutive elements of π of the form (x_{2j+1}, x_{2j+2}), so an element with an odd index followed by the next element in π. Now take a split G_1, G_2. We call a pair *balanced* on this split if one of the elements is in G_1 and the other is in G_2, otherwise *unbalanced*, that is, both elements are on one side of the split. If π is even and contains only balanced pairs, then we call the split G_1, G_2 a *strongly balanced* split with respect to π, and π strongly balanced on the split.

If π is not strongly balanced on G_1, G_2, then π must be odd and/or must contain at least one unbalanced pair. Assume that π contains an unbalanced pair. We go from left to right through the profile. Let (x_{2j+1}, x_{2j+2}) be the first unbalanced pair that we encounter, so that both elements are on one side of the split, say G_1. Note that now G_1 is a majority side for the subprofile $(x_1, x_2, \ldots, x_{2j+2})$. We call this pair a *winning pair* and we call the side G_1, in which they are, the *winning side* of the split (with respect to π). If π does not contain an unbalanced pair, then π must be odd, say $k = 2\ell + 1$, and π contains ℓ balanced pairs and one extra element right at the end: $x_{2\ell+1} = x_k$, say with x_k in G_1. Now we call x_k *winning*, and we call G_1 the *winning side* of the split.

We define the function $L : V^* \to 2^V$ on G as follows:

$$L(\pi) = \bigcap \{G_1 \mid G_1 \text{ is the winning side of a non-strongly balanced split of } \pi\}.$$

Note that, by our convention that $\bigcap \emptyset = V$, this definition implies that, if π is strongly balanced on all splits, then $L(\pi) = V$.

Clearly, L does not satisfy (A). It is straightforward to show that L satisfies (B). Let $\pi = (x, y)$. Then the winning sides are precisely the sides that contain both x and y, so a majority of the profile. Hence the winning sides are the majority sides, and we have $L(x, y) = Med(x, y) = I(x, y)$, by Corollary 2. Take any split G_1, G_2 and two mates a and b with a in G_1 and b in G_2. Then the profile $\pi = (a, b, b, b, a, a)$ is balanced on G_1, G_2 but not strongly balanced. So $L(\pi) = \{b\}$ and does not contain a, whence L does not satisfy $(Cond)$.

Two things remain to be proved. First, L is indeed a location function, that is, $L(\pi) \neq \emptyset$, for all profiles. Second, L satisfies (C). This we will do in Theorems 8 and 9.

To make full use of the power of induction on the number of expansions we now have a closer look at a specific type of split.

We call a split G_1, G_2 a *pendant split* if, say, $G_1 = G_{01}$. Note that this is equivalent to every vertex in G_1 having a mate in G_2. Then G_1 is the *pendant side*. In Figure 2

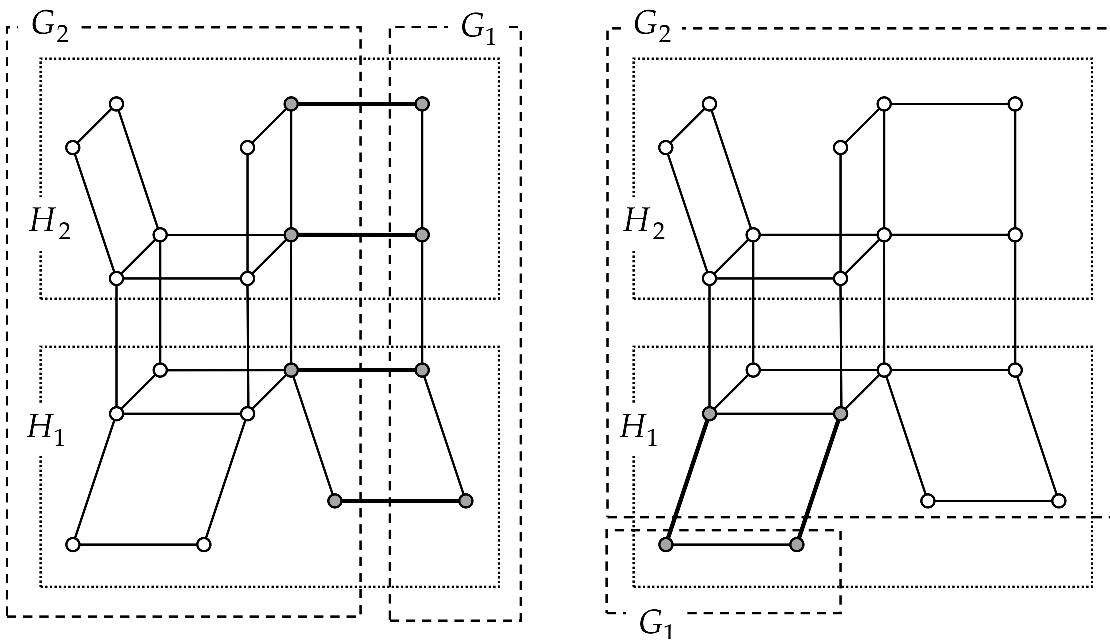

Figure 2: A pendant split, with a crossing split (left) and a non-crossing split (right)

we see two examples of pendant splits G_1, G_2. In a tree a pendant split corresponds to a pendant vertex as the pendant side, and the rest of the tree as the opposide. In an n-cube every $(n-1)$-cube is a pendant side. In [22] it was proved that in every non-trivial median graph there is always a pendant split and there are always at least two pendant sides. In our induction step below we will use a pendant split. Therefore we want to know the interplay between a pendant split and the other splits.

Let G_1, G_2 be a pendant split with pendant side $G_1 = G_{01}$. Clearly the contraction G', with respect to the pendant split G_1, G_2 with pendant side G_1, is now isomorphic to G_2, so, abusing notation, we may take G'_2 to be just G_2, and G'_1 to be just G_{02} and G' to be just G_2 as well.

Let H_1, H_2 be another split of G. First we consider the case that this split does not cross with G_1, G_2. Then the pendant side G_1 is contained in, say, H_1. Since the edges between G_1 and G_2 all determine the same split G_1, G_2, it follows that all these edges are in H_1 as well. So $G_1 \cup G_{02}$ is contained in H_1, whence H_2 is contained in G_2, see the splits in the right hand graph of Figure 2. In the contraction G' with respect to G_1, G_2, the split H_1, H_2 becomes H'_1, H'_2. The side H_2 that is contained in G_2 remains unaffected, while the other side H'_1 is obtained by "moving" the vertices in $G_1 = G_{01}$ to their mates in G_{02} (and deleting G_{01} altogether). So we have $H'_2 = H_2$ and $H'_1 = H_1 - G_1 = H_1 \cap G_2$. In getting back to G, by making the expansion, everything is restored.

If the split H_1, H_2 crosses G_1, G_2 (take the splits in the left hand graph in Figure 2), then we have we have $H_i' = H_i \cap G_2$, for $i = 1, 2$. In this case there are edges between the split sides H_1 and H_2 in $G_1 = G_{01}$ as well as G_{02}. So in the contraction $G' = G_2$ the corresponding edges are in G_{02}. We say that in G' this split crosses G_{02}. Then, in the expansion of G' with respect to the cover G_1', G_2', to get G back again, both sides H_1', H_2' are expanded to the sides H_1, H_2, respectively, such that both split sides contain parts of G_{01}, and thus are enlarged.

We may abuse notation in the following way. If W is a subset of vertices in G, and W' is the set obtained by taking the contraction with respect to G_1, G_2, then W' is a subset of vertices in G_2', which is isomorphic to G_2. We denote the subset of vertices in G_2 corresponding to W' by W' as well.

Now we are ready to show that the above defined function L involving strongly balanced splits is indeed a location function satisfying consistency. For convenience we recall the definition:

$$L(\pi) = \bigcap \{G_1 \mid G_1 \text{ is the winning side of a non-strongly balanced split of } \pi\}.$$

Theorem 8 *$L(\pi) \neq \emptyset$, for any profile π.*

Proof. First note that, if all splits are strongly balanced, then, by definition, $L(\pi) = V \neq \emptyset$. So in the sequel we will assume implicitly that there is always a split that is not strongly balanced.

The proof is by induction on the number of splits. The basis of the induction is L on K_2. Now there is only one split, so $L(\pi)$ is either the whole vertex set or consists of one of the split sides, whence $L(\pi)$ is trivially nonempty.

So let G be a median graph with n splits with $n \geq 2$, and assume that $L(\pi)$ is nonempty for any profile π on any median graph with fewer than n splits.

Fix a pendant split G_1, G_2, with G_1 being the pendant side (so $G_1 = G_{01}$). Let G' be the contraction with respect to this split, with proper covering G_1', G_2', so that we may take $G_1' = G_{02}$ and $G_2' = G_2$. Note that $G' = G_2$ is a median graph with one split less than G.

Now take any profile π on G with π' its contraction to $G' = G_2$. In the contraction π_2 is not affected and π_1 is moved over to $G_1' = G_{02}$ by taking mates (= gates). Let H_1, H_2 a be split distinct from G_1, G_2, and let H_1', H_2' be the contraction of that split. Due to the properties of splits and contractions, profile elements remain on the same side while taking the contraction. So H_1, H_2 is strongly balanced for π if and only if H_1', H_2' is strongly balanced for π'. Similarly, H_i is a winning side for π if and only if H_i' is a winning side for π'.

Let L' be the function on $G' = G_2$ defined by taking the intersection of the winning sides of the profile. By the induction hypothesis, $L'(\pi')$ is nonempty, for any profile π' on G'.

We fix a profile π on G. First note that, for any split H_1, H_2 distinct from G_1, G_2, we have $L'(\pi') \subseteq H_i'$ if and only if $L(\pi) \subseteq H_i$. Let W be the intersection of the winning sides of π of the splits distinct from G_1, G_2. Let a and b be mates with a in G_1 and b in G_2. Then a and b are on the same side of any other split H_1, H_2. So either both a and b are in W, or neither of a and b is in W. Let W' be the contraction of W, that is, mates with respect to G_1, G_2 in W are identified. So we may take $W' = W - G_1 = W \cap G_2$. In the contraction, we have $L'(\pi') = W'$.

With respect to the fixed split G_1, G_2, we consider three cases, depending on the behaviour of π on the split.

Case 1. In G, the side G_2 is a winning side of π.

Now $L(\pi) = W \cap G_2 = W' = L'(\pi')$, and we are done.

Case 2. The split G_1, G_2 is strongly balanced with respect to π.

Now no side of the split G_1, G_2 is involved in determining $L(\pi)$. So $L(\pi) = W$. Then in the contraction, $L'(\pi') = W' = W \cap G_2$. So $L'(\pi') \subset L(\pi)$, and we are done.

Case 3. In G, side G_1 is a winning side of π.

Now $L(\pi) = W \cap G_1$. The side G_1 contains a winning pair (y_{2j+1}, y_{2j+2}), or, in the case π is odd and contains only balanced pairs, a winning last element y_k.

Take any split H_1, H_2 that does not cross G_1, G_2. Then, say, G_1 is contained in H_1. So $G_1 \cup G_{02}$ is contained in H_1. Hence the pair (y_{2j+1}, y_{2j+2}), (or the last element y_k in case π is odd with only balanced pairs) is in H_1, and (y'_{2j+1}, y'_{2j+2}), (or the last element y'_k in case π is odd with only balanced pairs) is in H'_1 . This implies that the split H_1, H_2 cannot be strongly balanced. If $H'_2 = H_2$ were a winning side for π' in G', then it would have a winning pair that precedes (y'_{2j+1}, y'_{2j+2}) (or y'_k). This would be a pair of π in G_2 that precedes (y_{2j+1}, y_{2j+2}), or y_k, in the case π is odd and contains only balanced pairs. This is impossible. So H'_1 is a winning side in G' and H_1 is a winning side in G.

Next we show that $W' = L'(\pi')$ contains vertices in G_{02}. Assume the contrary. Take a vertex z in W', and, in G, let p be its gate in G_{02}. Take any geodesic P between z and p. Note that, p being a gate, it is the only vertex of P in G_{02}. Because of the convexity of G_2, the path P lies completely in G_2. Let q be the neighbor of p on P. Then q is in $G_2 - G_{02}$, so the edge pq is in G_2 but not in G_{02}. Hence the split H_{pq}, H_{qp} does not cross G_{02}, and $G_1 \cup G_{02}$ is contained in H_{pq}. Note that z is in H_{qp}. Above we showed that H_{pq} is a winning side for π. This yields a contradiction. So, indeed, W' contains a vertex b of G_{02}. Let a be the mate of b in G_1. Clearly, for any pair of mates in the split G_1, G_2, we have that either none or both are in the winning sides distinct from G_1. So a is in W, whence in $L(\pi) = W \cap G_1$. This settles Case 3 and concludes the proof. $\square$

The next thing we prove is that L is consistent.

Theorem 9 *L is consistent.*

Proof. Take two profiles ρ and σ with $L(\rho) \cap L(\sigma) \neq \emptyset$. Let $\pi = \rho\sigma$ be the concatenation of ρ and σ. We have to prove that $L(\pi) = L(\rho) \cap L(\sigma)$. Write $\rho = (x_1, x_2, \ldots, x_k)$ and $\sigma = (y_1, y_2, \ldots, y_m)$.

First we prove that $L(\rho) \cap L(\sigma) \subseteq L(\pi)$.

Let G_1, G_2 be a split of G. We next explore the role of this split in determining the sets $L(\rho)$, $L(\sigma)$, $L(\rho) \cap L(\sigma)$ and $L(\pi)$. Consider the four cases:

Case 1.1. ρ and σ are both strongly balanced.

Then, clearly, G_1, G_2 is also strongly balanced for π. So this split is ignored in determining the four sets $L(\rho)$, $L(\sigma)$, $L(\rho) \cap L(\sigma)$ and $L(\pi)$.

Case 1.2. ρ is strongly balanced but σ is not.

Clearly, in this case, the winning side of σ, say G_1, is also the winning side of π. Moreover, we have $L(\rho) \cap L(\sigma) \subseteq L(\sigma) \subseteq G_1$. So both $L(\pi)$ and $L(\rho) \cap L(\sigma)$ are on the same side G_1.

Case 1.3. ρ contains a winning pair.

Clearly this winning pair of ρ is also a winning pair of π. Say that G_1 is the winning side of ρ. Then G_1 is also the winning side of π. Moreover, we have $L(\rho) \cap L(\sigma) \subseteq L(\rho) \subseteq G_1$. Again both $L(\pi)$ and $L(\rho) \cap L(\sigma)$ are on the same side G_1.

Case 1.4. ρ is not strongly balanced but does not contain a winning pair.

This implies that ρ is odd and consists of balanced pairs and one extra element (to make it an odd profile). Say $\rho = (x_1, x_2, \ldots, x_{k-2}, x_{k-1}, x_k)$, where (x_{k-2}, x_{k-1}) is a balanced pair. Say x_k is in G_1. Then G_1 is a winning side of ρ, and we have $L(\rho) \cap L(\sigma) \subseteq L(\rho) \subseteq G_1$.

If G_1 is the winning side of π as well, then again we have both $L(\pi)$ and $L(\rho) \cap L(\sigma)$ on the same side.

Assume that G_2 is the winning side of π. We want to prove that in this case G_2 is a winning side of σ.

Recall that $\sigma = (y_1, y_2, \ldots, y_m)$. Note that (x_k, y_1) is a pair in π as are (y_{2j}, y_{2j+1}), for $j \geq 1$. Since x_k is in G_1 and G_2 is the winning side of $\pi = \rho\sigma$, we must have y_1 in G_2. If y_2 is in G_2, then (y_1, y_2) is a winning pair of σ in G_2, and we are done. If not, then y_2 is in G_1, but G_1 being a losing side of π, we deduce that y_3 must be in G_2. If y_4 is in G_2 we have a winning pair (y_3, y_4) in G_2 for σ, and we are done. If not, we can proceed until we find a winning pair of σ in G_2 or we end up with σ being odd and the elements of σ are alternatingly in G_2 and G_1, the first and the last being in G_2, so that indeed G_2 is the winning side of σ. So we have proved that in this case that G_1 is a winning side of ρ and G_2 is a winning side of σ, whence $L(\rho) \cap L(\sigma) = \emptyset$. This contradiction shows that G_2 cannot be the winning side of π. Hence we conclude that either G_1 is the winning side of π, or π is strongly balanced on the split, and the split is ignored in determining $L(\pi)$. Moreover, we have Ł$(\rho) \cap L(\sigma) \subseteq L(\rho) \subseteq G_1$. This concludes Case 1.4.

So in all four cases we get either that $L(\rho) \cap L(\sigma)$ and $L(\pi)$ are on the same side of the split or the split is ignored in determining $L(\pi)$. This implies that indeed $L(\rho) \cap L(\sigma) \subseteq L(\pi)$.

It still remains to be proved that $L(\rho) \cap L(\sigma) \subsetneq L(\pi)$ is not possible.

Assume that $L(\rho) \cap L(\sigma) \subsetneq L(\pi)$. Note that all sets involved are intersections of split sides, and therefore are gated and convex. So there is a vertex w in $L(\rho) \cap L(\sigma)$ and a neighbor z of w in $L(\pi) - [L(\rho) \cap L(\sigma)]$. Let G_1, G_2 be the split of edge wz with w in G_1 and z in G_2. Since w and z are both in $L(\pi)$ this split is strongly balanced with respect to π. We need to consider three cases:

Case 2.1. $z \notin L(\rho) \cup L(\sigma)$.

Now G_1 is a winning side for both ρ and σ. Since π is strongly balanced, this can only happen if both ρ and σ are odd and all pairs of π are balanced and x_k is in G_1. The remaining pairs of π are (x_k, y_1) and (y_{2j}, y_{2j+1}), for $1 \leq j \leq \lfloor \frac{m-1}{2} \rfloor$. For π to be strongly balanced, these pairs must all be balanced. Since x_k is in G_1, it follows that y_1 is in G_2. Now y_2 must be in G_1 again (otherwise G_2 would be winning for σ). So for (y_2, y_3) to be balanced, we have y_3 in G_1. Proceeding in this way we deduce that y_{m-1} is in G_1 and y_m is in G_2. But now all pairs of σ are balanced and y_m is in G_2. So G_2 should be winning for σ, which is impossible. This settles Case 2.1.

Case 2.2. $z \in L(\rho) - L(\sigma)$.

Now w and z are both in $L(\rho)$. So ρ is strongly balanced on the split G_1, G_2, whence ρ is even and all pairs of ρ are balanced on this split. On the other hand, G_1 is a winning side of σ. But this implies that G_1 is winning for $\pi = \rho\sigma$, which contradicts that π is strongly balanced on G_1, G_2. This settles Case 2.2.

Case 2.3. $z \in L(\sigma) - L(\rho)$.

Now G_1 is a winning side for ρ, whereas σ is strongly balanced on the split G_1, G_2. In particular, σ is even. Now, for π to be strongly balanced, ρ must be even. Hence ρ contains a winning pair in G_1, which is then also a winning pair for π. This contradicts the fact that π is strongly balanced, whence this case is settled.

Thus consistency is proved. □

7 (C) and $(Cond)$ does not imply (A) on median graphs with non-crossing splits

The combination of the axioms (C) and $(Cond)$ is an exceptional case. It turns out that on hypercubes these two axioms imply (A), as we show in Section 8. But on median graphs that are not hypercubes, we have an example of a location function that satisfies (C) and $(Cond)$ but *not* (A). Note that the smallest median graph that is not a hypercube is the path on three vertices. To help the reader to understand our example, one might restrict the location function L to the case where the median graph is the path $a \to x \to b$ on three vertices. In the general case vertex a is replaced by the subgraph G_a, vertex x by the subgraph G_x and vertex b by the subgraph G_b.

Let $G = (V, E)$ be a median graph that is not a hypercube. Then G contains a path $a \to x \to b$ such that the splits defined by the edges ax and xb are non-crossing (so the path is a convex subgraph). Let G_a be the subgraph of G induced by vertices closer to a than to x. Let G_b be the subgraph of G induced by the vertices closer to b than to x. Finally, let G_x be the subgraph induced by the remaining vertices, that is, the vertices that are closer to x than to a and b. Then the split defined by edge ax is $G_a, G_x \cup G_b$, and the split defined by the edge xb is $G_a \cup G_x, G_b$. Here it is understood that, when taking the union of two subgraphs, the edges between the two subgraphs are included as well. We denote each split by the edge that defines it. The output $L(\pi)$ of the location function L with respect to profile π depends on whether the two splits ax and xb are balanced or not. Furthermore, we have to take care that L does not satisfy anonymity, so in one case we let the order of π be involved in deciding what the output will be. We distinguish seven types of profiles. In each case we give the output for L and a label, and then which case it is. For each type we also give some helpful facts to be used in the proofs. As outputs we give subgraphs, but these should be read as the vertex sets of the subgraphs.

Type [axb]: $L(\pi) = V$.
Both splits ax and xb are balanced. Note that this means that exactly half of the elements of π is in G_a and exactly half of the elements of π is in G_b, so that π does not contain any element of G_x.

Type [axe]: $L(\pi) = G_a \cup G_x$.
Split ax is balanced, split xb is unbalanced, the majority side of this split is $G_a \cup G_x$, and the first element of G_x in π occurs in an *even* position. Note that, split ax being balanced, half of π is in $G_x \cup G_b$. But to avoid that split xb is balanced, less than half is in G_b, so G_x actually contains elements of π.

Type [axo]: $L(\pi) = G_b$.
Split ax is balanced, split xb is unbalanced, the majority side of this split is $G_a \cup G_x$, and the first element of G_x in π occurs in an *odd* position. Note as above that, split ax being balanced, half of π is in $G_x \cup G_b$. But to avoid that split xb is balanced, less than half is in G_b, so G_x actually contains elements of π.

Type [xb]: $L(\pi) = G_a$
Split ax is unbalanced, split xb is balanced, and the majority side of split ax is $G_b \cup G_x$.

Type [a]: $L(\pi) = G_b$.
Both splits are unbalanced, and the majority of π is in G_a.

Type [b]: $L(\pi) = G_a$.
Both splits are unbalanced, and the majority of π is in G_b.

Type [x]: $L(\pi) = G_a$.
Both splits are unbalanced and both G_a and G_b contain a minority of π, so that π contains elements of G_x.

Theorem 10 *L satisfies (C) and $(Cond)$ but not (A) or (B) or (F).*

Proof. The profiles of Type [axe] and [axo] guarantee that L is not anonymous. For any vertex x in G_a, the output $L(x)$ is G_b, so L is not faithful. For any two vertices x and y in G_a, the interval $I(x, y)$ is contained in G_a, but $L(x, y) = G_b$. So also betweenness is not satisfied.

To verify that L satisfies $(Cond)$, let π be any profile, and H_1, H_2 be a balanced split different from ax and xb. It is straightforward to check in the above cases that for any pair of mates u and v of this split, either both or none of the pair are in $L(\pi)$. Again, for the splits ax and xb, in the case they are balanced, it is straightfoward to check that both or none af any pair of mates are in $L(\pi)$.

Now we prove that L is consistent. Let π and ρ be two profiles with $L(\pi) \cap L(\rho) \neq \emptyset$. We distinguish a number of cases depending on what $L(\pi) \cap L(\rho)$ is.

Case 1. $L(\pi) \cap L(\rho) = V$
In this case we have $L(\pi) = L(\rho) = V$. So both π and ρ are of Type [axb]. Hence also $\pi\rho$ is of Type [axb], and so $L(\pi\rho) = V = L(\pi) \cap L(\rho)$.

Case 2. $L(\pi) \cap L(\rho) = G_a \cup G_x$.
Now one of the two profiles is of Type [axe] whereas the other is of Type [axb] or [axe]. So both profiles are balanced on the split ax and at least one is unbalanced on split xb with $G_a \cup G_x$ as majority side. So $\pi\rho$ is unbalanced on split xb with $G_a \cup G_x$ as majority side. If π is of Type [axb], then half of π lies in G_a and the other half lies in G_b. Moreover π does not contain an element in G_x. So the first element of $\pi\rho$ in G_x is the first element of ρ in G_x, and it is in an even position of ρ. Since π is even, the posititon of that element in $\pi\rho$ is also even. Hence $\pi\rho$ is of Type [axe]. If π is of Type [axe], then its first element in G_x is also the first element of $\pi\rho$ in G_x. So again $\pi\rho$ is of Type [axe], which settles Case 2.

Case 3. $L(\pi) \cap L(\rho) = G_a$.
Now one of the two profiles must be of Type [x] or [b] or [xb], while the other may be of Type [x] or [b] or [xb] or [axe] or [axb]. If both are of the same type, so either [x] or [b] or [bx], then $\pi\rho$ is also of the same type, and we are done. We need to check the other possible combinations of types.

First let one of the profiles be of Type [axb], so that half of the profile is in G_a and the other half in G_b. Then the other profile must be of Type [x] or [b] or [xb]. In all three cases the concatenation is of the same type as the latter profile, and we are done.

Second, let one of the profiles be of Type [axe], so that half of it is in G_a, the other half in $G_x \cup G_b$, and it contains elements of G_x, the first of which must be in an even position. The other profile must be of Type [x] or [b] or [xb]. If the latter profile is of Type [x], then the concatenation must have a minority in G_a and a minority in G_b, so that it is of Type [x], and we are done. If the latter profile is of Type [xb], the half of it is in G_b and the other half is in $G_a \cup G_x$, and it contains elements of G_x. So the concatenation contains elements of G_x and both G_a and G_b contain a minority, whence the concatenation is of Type [x], and we are done. Finally, if the latter profile is of Type [b], then the majority of this profile is in G_b. So the concatenation is unbalanced on split [ax], and the majority of the concatenation is in $G_x \cup G_b$. Hence the concatenation is of Type [x] or [b] or [xb], and again we are done.

Now we only have to check the cases where the two profiles are of distinct type, and both profiles are of Type [x], [b] or [xb].

Let one be of Type [x], so that of this profile there is a minority in G_a as well as in G_b. If the other is of Type [xb], so half is in G_b and the other half is in $G_a \cup G_x$, then the concatenation still has a minority in G_a as well as in G_b. So it is of Type [x], and we are done. If the other profile is of Type [b], then the concatenation has a majority in $G_x \cup G_b$, so that it is of Type [x] or [b] or [xb], and we are done.

Finally, let one profile be of Type [b] and the other of Type [xb]. Then, clearly, the concatenation has a majority in G_b, so that it is of Type [b], and we are done. This settles Case 3.

Case 4. $L(\pi) \cap L(\rho) = G_b$.
So the profiles are of Type [axb], [axo] or [a]. If one is of type [axb], then the other must be of type [axo] or [a], hence the concatenation is of that type to, and we are done. So we may take the profiles to be of type [axo] or [a]. If they are of the same type, then the concatenation is of that type too, and we are done. If they are of different type, then one profile has exactly half of its elements in G_a, whereas the other has a majority in G_a. Hence the concatenation has a majority in G_a as well. Again we are done. This completes the proof. $\square$

8 (C) and $(Cond)$ implies (A) on hypercubes

We need to first prove four lemmas that hold for all median graphs. Let $G = (V, E)$ be a median graph. Again we use the convention that $\cap\emptyset = V$. Let $\mathcal{S}$ be a set of splits of G. For v in V, we define the *spread* of v with respect to $\mathcal{S}$ to be

$$[v]_{\mathcal{S}} = \cap\{V(G_1) \mid G_1, G_2 \text{ is a split with } v \in G_1 \text{ and } G_1, G_2 \text{ not in } \mathcal{S}\}. \qquad (1)$$

Hence, if $\mathcal{S}$ consists of all splits in G, then, by the above convention that $\cap\emptyset = V$, we have $[v]_{\mathcal{S}} = V$, for every vertex v in V. Note that $[v]_{\mathcal{S}}$ is convex and v lies in $[v]_{\mathcal{S}}$. Moreover, for any edge yz with both ends in $[v]_{\mathcal{S}}$, it follows that the split defined by yz is in $\mathcal{S}$.

Lemma 11 *Let G be a median graph, $\mathcal{S}$ a set of splits of G, and v a vertex in G. Then*

$$w \in [v]_{\mathcal{S}} \Leftrightarrow [w]_{\mathcal{S}} = [v]_{\mathcal{S}}.$$

Proof. Note that w lies in $[v]_{\mathcal{S}}$ if and only if w is on the same side as v with respect to all splits not in $\mathcal{S}$. And this holds if and only if $[w]_{\mathcal{S}} = [v]_{\mathcal{S}}$. □

Note that Lemma 11 says that the spread operation induces, for every set $\mathcal{S}$ of splits, an equivalence relation on V.

Let W be a non-empty subset of V. We call W a *mating set* with respect to the set of splits $\mathcal{S}$ if the following holds: for any split G_1, G_2 in $\mathcal{S}$ and any pair of mates v_1, v_2 with v_1 in G_1 and v_2 in G_2 we have

$$v_1 \in W \Leftrightarrow v_2 \in W.$$

In the next lemma we establish the relationship between spread and mating set.

Lemma 12 *Let G be a median graph, $\mathcal{S}$ a set of splits of G, and W a mating set with respect to $\mathcal{S}$. Then*

$$v \in W \Rightarrow [v]_{\mathcal{S}} \subseteq W.$$

Proof. Let v be a vertex in W, and let z be any vertex in $[v]_{\mathcal{S}}$. Let P be a geodesic between v and z. By the convexity of $[v]_{\mathcal{S}}$, the path P is fully contained in $[v]_{\mathcal{S}}$. So, for any edge on P, the split defined by that edge is in $\mathcal{S}$. Hence, by the mating property of W all vertices of P are in W, in particular z. □

The next lemma is a first step in the study of location functions on median graphs that satisfy (*Cond*). Denote the set of balanced splits of a given profile π by $\mathcal{S}(\pi)$.

Lemma 13 *Let L be a location function on a median graph $G = (V, E)$. Then L satisfies* (*Cond*) *if and only if $v \in L(\pi) \Rightarrow [v]_{\mathcal{S}(\pi)} \subseteq L(\pi)$, for every profile π on G.*

Proof. First assume that L satisfies (*Cond*). Let π be a profile on G. Since L satisfies (*Cond*), the set $L(\pi)$ is a mating set. By Lemma 12, we have $v \in L(\pi) \Rightarrow [v]_{\mathcal{S}(\pi)} \subseteq L(\pi)$.

Next assume that L fails to satisfy (*Cond*). Then, for some profile π and some vertex v in $L(\pi)$, there is a neighbor w of v with w not in $L(\pi)$ such that π is balanced on the split G^v_{vw}, G^w_{vw}. Then w is contained in the spread of v with respect to $\mathcal{S}(\pi)$. Hence the implication fails to hold. □

The concatenation of k copies of a profile π is denoted by π^k. We refer to the following lemma as the Mixing Lemma.

Lemma 14 (Mixing Lemma) *Let $G = (V, E)$ be a median graph. Let π and ρ be profiles on G. Assume that G_1, G_2 is a split of G for which the majority side of π is G_1 and the majority side of ρ is G_2. Then there are positive integers a and b for which $\pi^a \rho^b$ is balanced on the split G_1, G_2.*

Proof. Recall that π_i is the subprofile of π in G_i and ρ_i that of ρ in G_i, for $i = 1, 2$. Set $a = |\rho_2| - |\rho_1|$ and set $b = |\pi_1| - |\pi_2|$. A straightforward counting argument shows that the profile $\pi^a \rho^b$ is balanced on the split G_1, G_2. □

Consistent with our above mentioned convention, we have, by Theorem 1, that $Med(\pi) = \cap\emptyset = V$, in the case that π is balanced on all splits.

Recall that the hypercube Q_n of dimension n is the graph whose vertex set is in one-to-one correspondence with n-tuples consisting of 0's and 1's (i.e., subsets of an n-element set) and two vertices of Q_n are adjacent if and only if they differ at exactly one coordinate. It is easy to see that Q_n can be constructed recursively from two copies of Q_{n-1} by adding edges from one Q_{n-1} to the corresponding vertices in the other. From the Expansion Theorem, a hypercube is a median graph G such that every vertex in G has a mate with respect to every split of G.

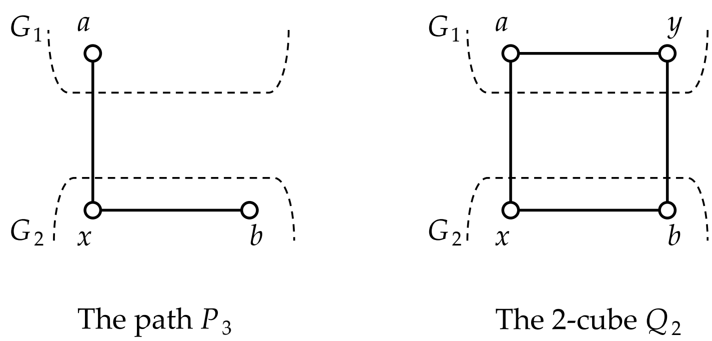

Figure 3: Difference between non-hypercube and hypercube

Our next result is specific for hypercubes. It states that the axioms ($Cond$) and (C) together imply (A) when G is a hypercube. As we have seen in Section 7, this fails to be true even for the path P_3 on 3 vertices. In order to clarify precisely where in the proof we use the fact that G is a hypercube, we first give an example, see Figure 3. Note that, in this proof, we consider two profiles π' and ρ. In P_3 as well as in Q_2 we have indicated a split G_1, G_2. Now take in each graph the profiles $\pi' = (a)$ and $\rho = (x, b)$. Note that in both graphs we have $Med(\rho) = G_2$, and ρ has exactly one unbalanced split, viz. the split G_1, G_2. Let L be a location function satisfying ($Cond$) and (C). We take $L(\pi') = \{b\}$ and $L(\rho) = \{x, b\}$. Then we have $L(\pi') \cap L(\rho) = L(\pi'\pi') \cap L(\rho) = \{b\}$, so that, by consistency, we have $L(a, a, x, b) = \{b\}$. Note that we have used the Mixing Lemma here to create the profile $(\pi')^2\rho^1$, which is balanced on split G_1, G_2. In the P_3-case everything is still fine up to this point. But in the Q_2-case we have a problem: the profile (a, a, x, b) is balanced on the split G_1, G_2. Since b is in $L(a, a, x, b)$, we deduce, by ($Cond$), that y must also be in $L(a, a, x, b)$. This contradiction is necessary in the proof below. The crucial point is that, in a hypercube, every vertex has a mate in the opposide in every split.

Theorem 15 *Let $G = (V, E)$ be a hypercube of dimension n. Let $L : V^* \to 2^V - \emptyset$ be a location function on G that satisfies (C) and $(Cond)$. Then L also satisfies (A).*

Proof. Recall that, G being a hypercube, G has n splits, and each split side is a hypercube of dimension $n-1$. Moreover, each vertex has a mate in all splits.

We need to show, for any profile $\pi = (x_1, x_2, \ldots, x_k)$ on G and any permutation σ of $\{1, 2, \ldots, k\}$, that $L(\pi) = L(\pi_\sigma)$, where $\pi_\sigma = (x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(k)})$. To achieve this, we prove something stronger, namely, that $L(\pi) = L(\pi')$ provided $Med(\pi) = Med(\pi')$. Since the median function respects anonymity, we have $Med(\pi) = Med(\pi_\sigma)$, from which our desired result then follows.

First we make a preliminary observation. Let ρ be a profile that is balanced on all splits *but* the split G_1, G_2. Then we have

$$L(\rho) = V(G_1), \quad \text{or} \quad L(\rho) = V(G_2), \quad \text{or} \quad L(\rho) = V(G_1) \cup V(G_2) = V. \tag{2}$$

Indeed, the set $\mathcal{S}(\rho)$ consists of all splits *but* the split G_1, G_2. So, for $i = 1, 2$, the spread $[v]_{\mathcal{S}(\rho)}$ of v is G_i, for any vertex v in G_i. When $L(\rho)$ has non-empty intersection with both $V(G_1)$ and $V(G_2)$, then Lemma 13 implies that $L(\rho) = V$. Otherwise, Lemma 13 implies that $L(\rho) = V(G_1)$ or $V(G_2)$.

Now fix a profile π. Let us denote the unbalanced splits of π by

$$G_{i_1}, G_{i_2} \quad i = 1, \ldots, k.$$

If there are no unbalanced splits we set $k = 0$. If all splits are unbalanced, then $k = n$. Without loss of generality, let G_{i_1} be the majority side of π, for any unbalanced split G_{i_1}, G_{i_2}. We denote the remaining, balanced splits by

$$G_{i_1}, G_{i_2} \quad i = k+1, \ldots, n.$$

For each unbalanced split G_{i_1}, G_{i_2} with $i = 1, \ldots, k$, we fix a profile ρ_i with $Med(\rho_i) = V(G_{i_2})$. Such a profile exists: G_{i_2} being an $(n-1)$-dimensional hypercube, we may take a profile of length 2 consisting of any two vertices at distance $n-1$ in G_{i_2}. Note that the majority side of ρ_i of this split is precisely the opposide G_{i_2} of the majority side G_{i_1} of π of this split, and that ρ_i is balanced on all other splits. So, by (2), we have $L(\rho_i)$ is either $V(G_{i_1})$ or $V(G_{i_2})$ or V.

Now let π' be any profile with $Med(\pi') = Med(\pi)$. Possibly $\pi' = \pi$. We use the values of $L(\rho_1)$, ..., $L(\rho_k)$ to determine $L(\pi')$. Let j be a number with $1 \leq j \leq k$. By the Mixing Lemma, there are numbers a and b such that the profile $\pi'^a \rho_j^b$ is balanced on the split G_{j_1}, G_{j_2}. Hence, by $(Cond)$, it follows that v lies in $L(\pi'^a \rho_j^b)$ if and only if the mate of v with respect to the split G_{j_1}, G_{j_2} also lies in $L(\pi'^a \rho_j^b)$. Since G is a hypercube, every vertex in one side of this split has a mate in the opposide. Therefore $L(\pi'^a \rho_j^b)$ is *not* contained entirely in either side of the split G_{j_1}, G_{j_2}. See the example above in Figure 2 for a clarification of this point.

First assume that $L(\rho_j) = V(G_{j_1})$. Then we must have $L(\pi') \cap L(\rho_j) = \emptyset$. For, otherwise, by consistency, we would have

$$L(\pi'^a \rho_j^b) = L(\pi'^a) \cap L(\rho_j^b) = L(\pi') \cap L(\rho_j) \subseteq L(\rho_j) = V(G_{j_1}),$$

which contradicts the above property of $L(\pi'^a\rho_j^b)$ that it intersects both sides of the split G_{j_1}, G_{j_2}. Hence, it follows that $L(\pi') \subseteq V(G_{j_2})$, that is, $L(\pi')$ is contained in the opposide of $L(\rho_j)$. Similarly, if $L(\rho_j) = G_{j_2}$, then we have $L(\pi') \subseteq V(G_{j_1})$. In both cases, for any vertex in $L(\pi')$ its mate with respect to the split G_{j_1}, G_{j_2} is *not* in $L(\pi')$.

Next we consider the case that $L(\rho_j) = V$. Then, by consistency, we have

$$L(\pi') = L(\pi') \cap L(\rho_j) = L(\pi'^a) \cap L(\rho_j^b) = L(\pi'^a\rho_j^b).$$

Hence, in this case, we have that v lies in $L(\pi')$ if and only if its mate with respect to the split G_{j_1}, G_{j_2} also lies in $L(\pi')$. We could say that π' acts with respect to L on G_{j_1}, G_{j_2} as if it were balanced on this split.

Recapitulating: if $L(\rho_j)$ equals V, then $L(\pi')$ is a mating set with respect to split G_1, G_2, and if $L(\rho_j)$ equals one of the split sides of G_1, G_2, then $L(\pi)$ is contained in the opposide.

Since $Med(\pi') = Med(\pi)$, the profiles π and π' are balanced on the same splits. So $\mathcal{S}(\pi') = \mathcal{S}(\pi)$. Let $\mathcal{S}'$ be the set of splits among the first k splits with $L(\rho_j) = V$. These are precisely the splits on which π' is not balanced but acts with respect to L as if it were balanced. Set $\mathcal{S} = \mathcal{S}' \cup \mathcal{S}(\pi)$. So $L(\pi')$ is a mating set with respect to $\mathcal{S}$. The remaining splits are then the splits on which π' is "compulsory unbalanced", and we call the sides in which $L(\pi')$ is contained the *compulsory sides*. Note that, since $L(\pi')$ is contained in the intersection of the compulsory sides, this intersection is nonempty. Clearly, the spread of any vertex in this intersection with respect to $\mathcal{S}$, is precisely this intersection.

In the previous paragraph we have established the following two facts:

(i) $L(\pi')$ is contained in the spread of any vertex in the intersection of the compulsory sides, i.e. the opposides of $L(\rho_j)$, for which $L(\rho_j) \neq V$,

(ii) $L(\pi')$ is a mating set with respect to $\mathcal{S}$.

Hence, by Lemmas 12 and 13, we conclude that $L(\pi')$ is precisely the intersection of the opposides of the $L(\rho_j)$ with $L(\rho_j) \neq V$ (the compulsory sides). Since π' was any profile such that $Med(\pi') = Med(\pi)$ we have now proved our stronger result. □

9 Concluding remarks

Theorems 4 and 5 give two elegant axiomatic characterizations of the median function Med on median graphs. This prompted the question of the independence of the five axioms involved in these Theorems. Apart from the obvious dependencies implied by these theorems, and the trivial fact that (F) follows from (B) and (C), we have examples on any non-trivial median graph that show that no other implications are possible. There is one exceptional case: (C) and $(Cond)$ do imply (A) on hypercubes, but do not on median graphs with non-crossing splits. For the examples and proofs we made extensive use of the rich structure theory that is available for median graphs.

A main open problem is to find other classes of graphs on which the median function is characterized by the three basic axioms (A), (B) and (C); axioms that are necessarily satisfied by Med on any finite connected graph. To date, we have found an infinite family of complete bipartite graphs, namely $K_{m,2}$ where $m \geq 3$, on which the median function is characterized by (A), (B) and (C). Also, for other non-median graphs, can we characterize the median function by adding one or two other simple axioms? This is certainly a point for further research. And then, of course, independence of sets of axioms would again be a pertinent question.

References

[1] K.J. Arrow, *Social Choice and Individual Values*, no. 12 in Cowles Commission for Research in Economics: Monographs, Wiley, New York, first ed. 1951.

[2] K.J. Arrow, A.K. Sen and K. Suzumura (eds), *Handbook of Social Choice and Welfare*, Volumes 1, North Holland, Amsterdam, 2002.

[3] K.J. Arrow, A.K. Sen and K. Suzumura (eds), *Handbook of Social Choice and Welfare*, Volumes 2, North Holland, Amsterdam, 2005.

[4] K. Balakrishnan, M. Changat, H.M. Mulder and A.R. Subhamathi, Axiomatic characterization of the antimedian function on paths and hypercubes, *Discrete Math. Algor. Appl.* **12** (2012) no. 4, 1250054, 20 pp.

[5] J.P. Barthélemy and M.F. Janowitz, A formal theory of consensus, *SIAM J. Discrete Math.* **4** (1991) 305 – 322.

[6] J.P. Barthélemy and B. Monjardet, The median procedure in cluster analysis and social choice theory, *Math. Soc. Sci.* **1** (1981) 235 – 268.

[7] W.H.E. Day and F.R. McMorris, *Axiomatic Consensus Theory in Group Choice and Biomathematics*, Frontiers in Appl. Math., SIAM, Philadelphia, 2003.

[8] D.P. Foster and R. Vohra, An axiomatic characterization of a class of location in tree networks, *Oper. Res.* **46** (1998) 347 – 354.

[9] R. Holzman, An axiomatic approach to location on networks, *Math. Oper. Res.* **15** (1990), 553 – 563.

[10] W. Imrich, S. Klavžar and H.M. Mulder, Median graphs and triangle-free graphs, *SIAM J. Discrete Math.* **12** (1999) 111 – 118.

[11] S. Klavžar and H.M. Mulder, Median graphs: characterizations, location theory and related structures, *J. Combin. Math. Combin. Comput.* **30** (1999) 103 – 127.

[12] F.R. McMorris, H.M. Mulder and O. Ortega, Axiomatic characterization of the mean function on trees, *Discrete Math. Algor. Appl.* **2** (2010) 313 – 329.

[13] F.R. McMorris, H.M. Mulder and O. Ortega, Axiomatic characterization of the ℓ_p-function on trees, *Networks* **60** (2012) 94 – 102.

[14] F.R. McMorris, H.M. Mulder and R.C. Powers, The median function on median graphs and semilattices, *Discrete Appl. Math.* **101** (2000) 221 – 230.

[15] F.R. McMorris, H.M. Mulder and R.C. Powers, The median function on distributive semilattices, *Discrete Appl. Math.* **127** (2003) 319 – 324.

[16] F.R. McMorris, H.M. Mulder and F.S. Roberts, The median procedure on median graphs, *Discrete Appl. Math.* **84** (1998) 165 – 181.

[17] F.R. McMorris, H.M. Mulder and R.V. Vohra, Axiomatic characterization of location functions, in H. Kaul and H.M. Mulder, eds, *Advances in Interdisciplinary Discrete Applied Mathematics*, Interdisciplinary Mathematical Sciences, Vol. 11, World Scientific, Singapore, 2010, pp. 71 – 91.

[18] F.R. McMorris, F.S. Roberts and C. Wang, The center function on trees, *Networks* **38** (2001) 84 – 87.

[19] P.B. Mirchandani and R.L. Francis, eds, *Discrete location theory*, Wiley, New York, 1990.

[20] H.M. Mulder, The structure of median graphs, *Discrete Math.*, **24** (1978) 197 – 204.

[21] H.M. Mulder, *The Interval Function of a Graph*, Mathematical Centre Tract 132, Mathematisch Centrum, Amsterdam, 1980, 459 – 477.

[22] H.M. Mulder, The expansion procedure for graphs, in *Contemporary Methods in Graph Theory*, BI Wissenschaftsverlag, Mannheim/Wien/Zürich, 1990, pp. 459 – 477.

[23] H.M. Mulder, The majority strategy on graphs, *Discrete Appl. Math.* **80** (1997) 97 – 105.

[24] H.M. Mulder, Median graphs. A structure theory, in H. Kaul and H.M. Mulder, eds, *Advances in interdisciplinary discrete applied mathematcis*, Interdisplinary Mathematical Sciences, Vol. 11, World Scientific, Singapore, 2010, pp. 93 – 125.

[25] H.M. Mulder and B.A. Novick, An axiomatization of the median procedure on the n-cube, *Discrete Appl. Math.* **159** (2011) 939 – 944.

[26] H.M. Mulder and B.A. Novick, A tight axiomatization of the median procedure on median graphs, *Discrete Appl. Math.* **161** (2013) 838 – 846.

[27] H.M. Mulder and A. Schrijver, Median graphs and Helly hypergraphs, *Discrete Math.* **25** (1979) 41 – 50.

[28] S. Nickel and J. Puerto, *Location Theory: A Unified Approach*, Springer, Berlin, 2005.

[29] K.B. Reid, Centrality measures in trees, in H. Kaul and H.M. Mulder, eds, *Advances in interdisciplinary discrete applied mathematcis*, Interdisplinary Mathematical Sciences, Vol. 11, World Scientific, Singapore, 2010, pp. 167 – 197.

[30] F. Roberts, Characterization of the plurality function, *Math. Social Sciences* **21** (1991) 101 – 127.

[31] C. Smart and P.J. Slater, Center, median, and centroid subgraphs, *Networks* **34** (1999) 303 – 311.

[32] P.J. Slater, A survey of sequences of central subgraphs, *Networks* **34** (1999) 244 – 249.

[33] R. Vohra, An axiomatic characterization of some location in trees, *Europ. J. Oper. Res.* **90** (1996) 78 – 84.

[34] C.J. Witzgall, Optimal location of a central facility, *Mathematical Models and Concepts*, National Bureau of Standards, Washington D.C., 1965.