Medians in Median Graphs and Their Cube Complexes in Linear Time

Laurine Bénéteau
Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France
laurine.beneteau@lis-lab.fr

Jérémie Chalopin
Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France
jeremie.chalopin@lis-lab.fr

Victor Chepoi
Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France
victor.chepoi@lis-lab.fr

Yann Vaxès
Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France
yann.vaxes@lis-lab.fr

Abstract

The median of a set of vertices P of a graph G is the set of all vertices x of G minimizing the sum of distances from x to all vertices of P. In this paper, we present a linear time algorithm to compute medians in median graphs, improving over the existing quadratic time algorithm. We also present a linear time algorithm to compute medians in the ℓ_1-cube complexes associated with median graphs. Median graphs constitute the principal class of graphs investigated in metric graph theory and have a rich geometric and combinatorial structure. Our algorithm is based on the majority rule characterization of medians in median graphs and on a fast computation of parallelism classes of edges (Θ-classes or hyperplanes) via Lexicographic Breadth First Search (LexBFS). To prove the correctness of our algorithm, we show that any LexBFS ordering of the vertices of G satisfies the following *fellow traveler property* of independent interest: the parents of any two adjacent vertices of G are also adjacent.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Theory of computation → Computational geometry

Keywords and phrases Median Graph, CAT(0) Cube Complex, Median Problem, Linear Time Algorithm, LexBFS

Digital Object Identifier 10.4230/LIPIcs.ICALP.2020.10

Category Track A: Algorithms, Complexity and Games

Related Version A full version of this paper is available at `https://arxiv.org/abs/1907.10398`.

Funding The research on this paper was supported by the ANR project DISTANCIA (ANR-17-CE40-0015).

Acknowledgements The authors are grateful to the anonymous referees for useful remarks that helped improving the presentation of this work.

1 Introduction

The median problem (also called the Fermat-Torricelli problem or the Weber problem) is one of the oldest optimization problems in Euclidean geometry [43]. The *median problem* can be defined for any metric space (X, d): given a finite set $P \subset X$ of points with positive weights, compute the points x of X minimizing the sum of the distances from x to the points of P multiplied by their weights. The median problem in graphs is one of the principal models in network location theory [35, 62] and is equivalent to finding nodes with largest closeness centrality in network analysis [14, 15, 57]. It also occurs in social group choice as the Kemeny median. In the consensus problem in social group choice, given individual rankings one has to compute a consensual group decision. By Arrow's impossibility theorem, there is no consensus function satisfying natural "fairness" axioms. It is also well-known that the majority rule leads to Condorcet's paradox, i.e., to the existence of cycles in the majority relation. In this respect, the Kemeny median [39, 40] is an important consensus function and corresponds to the median problem in graphs of permutahedra.

The median problem in Euclidean spaces cannot be solved in a symbolic form [6], but it can be solved numerically by Weiszfeld's algorithm [66] and its convergent modifications (see e.g. [51]) and can be approximated in nearly linear time with arbitrary precision [26]. For the ℓ_1-metric the median problem becomes much easier and can be solved by the majority rule on coordinates, i.e., by taking as median a point whose ith coordinate is the median of the list consisting of ith coordinates of the points of P. This kind of majority rule was used in [38] to define centroids of trees (which coincide with their medians [32, 62]). For graphs with n vertices, m edges, and standard graph distance, the median problem can be trivially solved in $O(nm)$ time by solving the All Pairs Shortest Paths (APSP) problem. One may ask if APSP is necessary to compute the median. However, by [1] APSP and median problem are equivalent under subcubic reductions. It was also shown in [2] that computing the medians of sparse graphs in subquadratic time refutes the HS (Hitting Set) conjecture.

In this paper, we show that the medians in median graphs can be computed in optimal $O(m)$ time (i.e., without solving APSP). Median graphs are the graphs in which each triplet u, v, w of vertices has a unique median, i.e., a vertex metrically lying between u and v, v and w, and w and u. They originally arise in universal algebra [4, 18] and their properties have been first investigated in [45, 49]. It was shown in [24, 54] that the cube complexes of median graphs are exactly the CAT(0) cube complexes, i.e., cube complexes of global non-positive curvature. CAT(0) cube complexes, introduced and nicely characterized in [33] in a local-to-global way, are now one of the principal objects of investigation in geometric group theory [59]. Median graphs also occur in Computer Science: by [3, 13] they are exactly the domains of event structures (one of the basic abstract models of concurrency) [50] and median-closed subsets of hypercubes are exactly the solution sets of 2-SAT formulas [48, 60]. The bijections between median graphs, CAT(0) cube complexes, and event structures have been used in [20, 21, 25] to disprove three conjectures in concurrency [56, 63, 64]. Finally, median graphs, viewed as median closures of sets of vertices of a hypercube, contain all most parsimonious (Steiner) trees [12] and as such have been extensively applied in human genetics. For a survey on median graphs and their connections with other discrete and geometric structures, see the books [36, 42], the surveys [10, 41], and the paper [22].

As we noticed, median graphs have strong geometric and metric properties. For the median problem, the concept of Θ-classes is essential. Two edges of a median graph G are called opposite if they are opposite in a common square of G. The equivalence relation Θ is the reflexive and transitive closure of this oppositeness relation. Each equivalence class of Θ is called a Θ-class (Θ-classes correspond to hyperplanes in CAT(0) cube complexes [58] and to events in event structures [50]). Removing the edges of a Θ-class, the graph G is split into two connected components which are convex and gated. Thus they are called halfspaces of G. The convexity of halfspaces implies via [29] that any median graph G isometrically embeds into a hypercube of dimension equals to the number q of Θ-classes of G.

Our results and motivation. In this paper, we present a linear time algorithm to compute medians in median graphs and in ℓ_1-cube complexes associated to median graphs. Our main motivation and technique stem from the majority rule characterization of medians in median graphs and the unimodality of the median function [8, 61]. Even if the majority rule is simple to state and is a commonly approved consensus method, its algorithmic implementation is less trivial if one has to avoid the computation of the distance matrix. On the other hand, the unimodality of the median function implies that one may find the median set by using local search. More generally, consider a partial orientation of the input median graph G, where an edge uv is transformed into the arc $\overrightarrow{uv}$ iff the median function at u is larger than the median function at v (in case of equality we do not orient the edge uv). Then the median set is exactly the set of all the sinks in this partial orientation of G. Therefore, it remains to compare for every edge uv the median function at u and at v in constant time. For this we use the partition of the edge-set of a median graph G into Θ-classes and; for every Θ-class, the partition of the vertex-set of G into complementary halfspaces. It is easy to notice that all edges of the same Θ-class are oriented in the same way because for any such edge uv the difference between the median functions at u and at v, respectively, can be expressed as the sum of weights of all vertices in the same halfspace as v minus the sum of weights of all vertices in the same halfspace as u.

Our main technical contribution is a new method for computing the Θ-classes of a median graph G with n vertices and m edges in linear $O(m)$ time. For this, we prove that Lexicographic Breadth First Search (LexBFS) [55] produces an ordering of the vertices of G satisfying the following *fellow traveler property*: for any edge uv, the parents of u and v are adjacent. With the Θ-classes of G at hand and the majority rule for halfspaces, we can compute the weights of halfspaces of G in optimal time $O(m)$, leading to an algorithm of the same complexity for computing the median set. We adapt our method to compute in linear time the median of a finite set of points in the ℓ_1-cube complex associated with G. The method can be applied to compute the total distance (the Wiener index) in optimal $O(m)$ time and the distance matrix of G in optimal $O(n^2)$ time (see the full version [16]).

Related work. The study of medians in median graphs originated in [8, 61] and continued in [7, 44, 46, 47, 53]. Using different techniques and extending the majority rule for trees [32], the following *majority rule* has been established in [8, 61]: a halfspace H of a median graph G contains at least one median iff H contains at least one half of the total weight of G; moreover, the median of G coincides with the intersection of halfspaces of G containing strictly more than half of the total weight. Hence the median set is always convex. It was shown in [61] that the median function of a median graph is weakly convex (an analog of a discrete convex function). This convexity property characterizes all graphs in which all local medians are global [9]. A nice axiomatic characterization of medians of median graphs via three basic axioms has been obtained in [47]. More recently, [53] characterized median graphs as *closed Condorcet domains*, i.e., as sets of linear orders with the property that, whenever the preferences of all voters belong to the set, their majority relation has no cycles and also belongs to the set. In the full version [16] we show that the median graphs are the bipartite graphs in which the medians are characterized by the majority rule.

Prior to our work, the best algorithm to compute the Θ-classes of a median graph G has complexity $O(m \log n)$ [34]. It was used in [34] to recognize median graphs in subquadratic time. The previous best algorithm for the median problem in a median graph G with n vertices and q Θ-classes has complexity $O(qn)$ [7] which is quadratic in the worst case. Indeed q may be linear in n (as in the case of trees) and is always at least $d(\sqrt[d]{n}-1)$ as shown below

(d is the largest dimension of a hypercube which is an induced subgraph of G). Additionally, [7] assumes that an isometric embedding of G in a q-hypercube is given. The description of such an embedding has already size $O(qn)$. The Θ-classes of a median graph G correspond to the coordinates of the smallest hypercube in which G isometrically embeds (this is called the *isometric dimension* of G [36]). Thus one can define Θ-classes for all partial cubes, i.e., graphs isometrically embeddable into hypercubes. An efficient computation (in $O(n^2)$ time) of all Θ-classes was the main step of the $O(n^2)$ algorithm of [30] for recognizing partial cubes. The fellow-traveler property (which is essential in our computation of Θ-classes) is a notion coming from geometric group theory [31] and is a main tool to prove the (bi)automaticity of a group. In a slightly stronger form it allows to establish the dismantlability of graphs (see [19, 23, 24] for examples of classes of graphs in which a fellow traveler order was obtained by BFS or LexBFS). LexBFS has been used to solve optimally several algorithmic problems in different classes of graphs, in particular for their recognition (for a survey, see [28]).

Cube complexes of median graphs with ℓ_1-metric have been investigated in [65]. The same complexes but endowed with the ℓ_2-metric are exactly the CAT(0) cube complexes. As we noticed above, they are of great importance in geometric group theory [59]. The paper [17] established that the space of trees with a fixed set of leaves is a CAT(0) cube complex. A polynomial-time algorithm to compute the ℓ_2-distance between two points in this space was proposed in [52]. This result was recently extended in [37] to all CAT(0) cube complexes. A convergent numerical algorithm for the median problem in CAT(0) spaces was given in [5].

2 Preliminaries

All graphs $G = (V, E)$ in this paper are finite, undirected, simple, and connected; V is the vertex-set and E is the edge-set of G. We write $u \sim v$ if $u, v \in V$ are adjacent. The *distance* $d(u, v) = d_G(u, v)$ between two vertices u and v is the length of a shortest (u, v)-path, and the *interval* $I(u, v) = \{x \in V : d(u, x) + d(x, v) = d(u, v)\}$ consists of all the vertices on shortest (u, v)–paths. A set H (or the subgraph induced by H) is *convex* if $I(u, v) \subseteq H$ for any two vertices u, v of H; H is a *halfspace* if H and $V \setminus H$ are convex. Finally, H is *gated* if for every vertex $v \in V$, there exists a (unique) vertex $v' \in V(H)$ (the *gate* of v in H) such that for all $u \in V(H)$, $v' \in I(u, v)$. The *k-dimensional hypercube* Q_k has all subsets of $\{1, \ldots, k\}$ as the vertex-set and $A \sim B$ iff $|A \triangle B| = 1$. A graph G is called *median* if $I(x, y) \cap I(y, z) \cap I(z, x)$ is a singleton for each triplet x, y, z of vertices; this unique intersection vertex $m(x, y, z)$ is called the *median* of x, y, z. Median graphs are bipartite and do not contain induced $K_{2,3}$. The *dimension* $d = \dim(G)$ of a median graph G is the largest dimension of a hypercube of G. In G, we refer to the 4-cycles as *squares*, and the hypercube subgraphs as *cubes*.

A map $w : V \to \mathbb{R}^+ \cup \{0\}$ is called a *weight function*. For a vertex $v \in V$, $w(v)$ denotes the weight of v (for a set $S \subseteq V$, $w(S) = \sum_{x \in S} w(x)$ denotes the weight of S). Then $F_w(x) = \sum_{v \in V} w(v) d(x, v)$ is called the *median function* of the graph G and a vertex x minimizing F_w is called a *median vertex* of G. Finally, $\mathrm{Med}_w(G) = \{x \in V : x \text{ is a median of } G\}$ is called the *median set* (or simply, the *median*) of G with respect to the weight function w.

3 Facts about median graphs

We recall the principal properties of median graphs used in the algorithms. Some of those results are a part of folklore for the people working in metric graph theory and some other results can be found in the papers [45, 46] by Mulder.

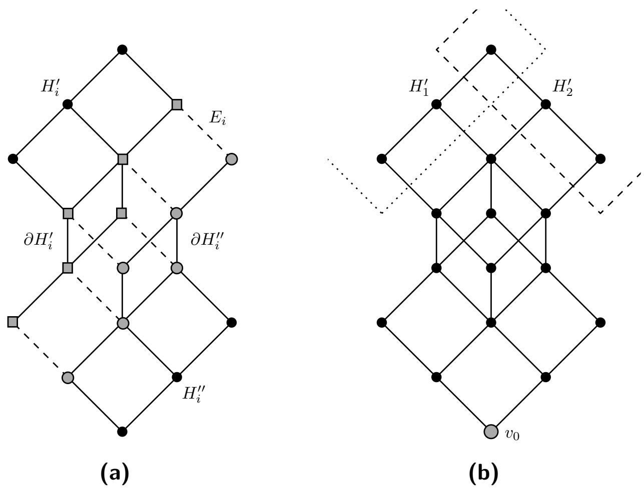

Figure 1 (a) In dashed, the Θ-class E_i of D, its two complementary halfspaces H'_i and H''_i and their boundaries $\partial H'_i$ and $\partial H''_i$, (b) two peripheral halfspaces of D, and (c) a LexBFS ordering of D.

From now on, $G = (V, E)$ is a median graph with n vertices and m edges. The first three properties follow from the definition.

▶ **Lemma 1** (Quadrangle Condition). *For any vertices u, v, w, z of G such that $v, w \sim z$ and $d(u, v) = d(u, w) = d(u, z) - 1 = k$, there is a unique vertex $x \sim v, w$ such that $d(u, x) = k - 1$.*

▶ **Lemma 2** (Cube Condition). *Any three squares of G, pairwise intersecting in three edges and all three intersecting in a single vertex, belong to a 3-dimensional cube of G.*

▶ **Lemma 3** (Convex=Gated). *A subgraph of G is convex if and only if it is gated. Each convex subgraph S of G is the intersection of all halfspaces containing S.*

Two edges uv and $u'v'$ of G are in relation Θ_0 if $uvv'u'$ is a square of G and uv and $u'v'$ are opposite edges of this square. Let Θ denote the reflexive and transitive closure of Θ_0. Denote by $E_1, \ldots, E_q$ the equivalence classes of Θ and call them Θ-*classes* (see Fig. 1(a)).

▶ **Lemma 4.** [45] (Halfspaces and Θ-classes). *For any Θ-class E_i of G, the graph $G_i = (V, E \setminus E_i)$ consists of exactly two connected components H'_i and H''_i that are halfspaces of G; all halfspaces of G have this form. If $uv \in E_i$, then H'_i and H''_i are the subgraphs of G induced by $W(u, v) = \{x \in V : d(u, x) < d(v, x)\}$ and $W(v, u) = \{x \in V : d(v, x) < d(u, x)\}$.*

Two Θ-classes E_i and E_j are *crossing* if each halfspace of the pair $\{H'_i, H''_i\}$ intersects each halfspace of the pair $\{H'_j, H''_j\}$; otherwise, E_i and E_j are called *laminar*.

▶ **Lemma 5** (Crossing Θ-classes). *Any vertex $v \in V(G)$ and incident edges $vv_1 \in E_1, \ldots, vv_k \in E_k$ belong to a single cube of G if and only if $E_1, \ldots, E_k$ are pairwise crossing.*

The *boundary* $\partial H'_i$ of a halfspace H'_i is the subgraph of H'_i induced by all vertices v' of H'_i having a neighbor v'' in H''_i. A halfspace H'_i of G is *peripheral* if $\partial H'_i = H'_i$ (See Fig. 1(b)).

▶ **Lemma 6** (Boundaries). *For any Θ-class E_i of G, $\partial H'_i$ and $\partial H''_i$ are isomorphic and gated.*

From now on, we suppose that G is rooted at an arbitrary vertex v_0 called the *basepoint*. For any Θ-class E_i, we assume that v_0 belongs to the halfspace H''_i. Let $d(v_0, H'_i) = \min\{d(v_0, x) : x \in H'_i\}$. Since H'_i is gated, the gate of v_0 in H'_i is the unique vertex of H'_i

at distance $d(v_0, H_i')$ from v_0. Since median graphs are bipartite, the choice of v_0 defines a canonical orientation of the edges of G: $uv \in E$ is oriented from u to v (notation $\overrightarrow{uv}$) if $d(v_0, u) < d(v_0, v)$. Let $\overrightarrow{G}_{v_0}$ denote the resulting oriented pointed graph.

▶ **Lemma 7.** [46] (Peripheral Halfspaces). *Any halfspace H_i' maximizing $d(v_0, H_i')$ is peripheral.*

For a vertex v, all vertices u such that $\overrightarrow{uv}$ is an edge of $\overrightarrow{G}_{v_0}$ are called *predecessors* of v and are denoted by $\Lambda(v)$. Equivalently, $\Lambda(v)$ consists of all neighbors of v in the interval $I(v_0, v)$. A median graph G satisfies the *downward cube property* if any vertex v and all its predecessors $\Lambda(v)$ belong to a single cube of G.

▶ **Lemma 8.** [45] (Downward Cube Property). *G satisfies the downward cube property.*

Lemma 8 immediately implies the following upper bound on the number of edges of G:

▶ **Corollary 9.** *If G has dimension d, then $m \le dn \le n \log n$.*

We give a sharp lower bound on the number q of Θ-classes, which is new to our knowledge.

▶ **Proposition 10.** *If G has q Θ-classes and dimension d, then $q \ge d(\sqrt[d]{n} - 1)$ and this bound is tight.*

Proof. Let $\Gamma(G)$ be the crossing graph $\Gamma(G)$ of G: $V(\Gamma(G))$ is the set of Θ-classes of G and two Θ-classes are adjacent in $\Gamma(G)$ if they are crossing. Note that $|V(\Gamma(G))| = q$. Let $X(\Gamma(G))$ be the clique complex of $\Gamma(G)$. By the characterization of median graphs among ample classes [11, Proposition 4], the number of vertices of G is equal to the number $|X(\Gamma(G))|$ of simplices of $X(\Gamma(G))$. Since G is of dimension d, by [11, Proposition 4], $\Gamma(G)$ does not contain cliques of size $d+1$. By Zykov's theorem [68] (see also [67]), the number of k-simplices in $X(\Gamma(G))$ is at most $\binom{d}{k}\left(\frac{q}{d}\right)^k$. Hence $n = |V(G)| = |X(\Gamma(G))| \le \sum_{k=0}^{d} \binom{d}{k}\left(\frac{q}{d}\right)^k = \left(1 + \frac{q}{d}\right)^d$ and thus $q \ge d(\sqrt[d]{n} - 1)$. Let now G be the Cartesian product of d paths of length $(\sqrt[d]{n} - 1)$. Then G has $(\sqrt[d]{n} - 1 + 1)^d = n$ vertices and $d(\sqrt[d]{n} - 1)$ Θ-classes (each Θ-class of G corresponds to an edge of one of factors). ◀

4 Computation of the Θ-classes via LexBFS

The *Breadth-First Search (BFS)* refines the basepoint order and defines the same orientation $\overrightarrow{G}_{v_0}$ of G. BFS uses a queue Q and the insertion in Q defines a total order $<$ on the vertices of G: $x < y$ iff x is inserted in Q before y. When a vertex u arrives at the head of Q, it is removed from Q and all not yet discovered neighbors v of u are inserted in Q; u becomes the *parent* $f(v)$ of v; for any vertex $v \ne v_0$, $f(v)$ is the smallest predecessor of v. The arcs $\overrightarrow{f(v)v}$ define the *BFS-tree* of G. The *Lexicographic Breadth-First Search (LexBFS)*, proposed in [55], is a refinement of BFS. In BFS, if u and v have the same parent, then the algorithm order them arbitrarily. Instead, the LexBFS chooses between u and v by considering the ordering of their second-earliest predecessors. If only one of them has a second-earliest predecessor, then that one is chosen. If both u and v have the same second-earliest predecessor, then the tie is broken by considering their third-earliest predecessor, and so on (See Fig. 1(c)). The LexBFS uses a set partitioning data structure and can be implemented in linear time [55]. In median graphs, the next lemma shows that it suffices to consider only the earliest and second-earliest predecessors, leading to a simpler implementation of LexBFS:

▶ **Lemma 11.** *If u and v are two vertices of a median graph G, then $|\Lambda(u) \cap \Lambda(v)| \le 1$.*

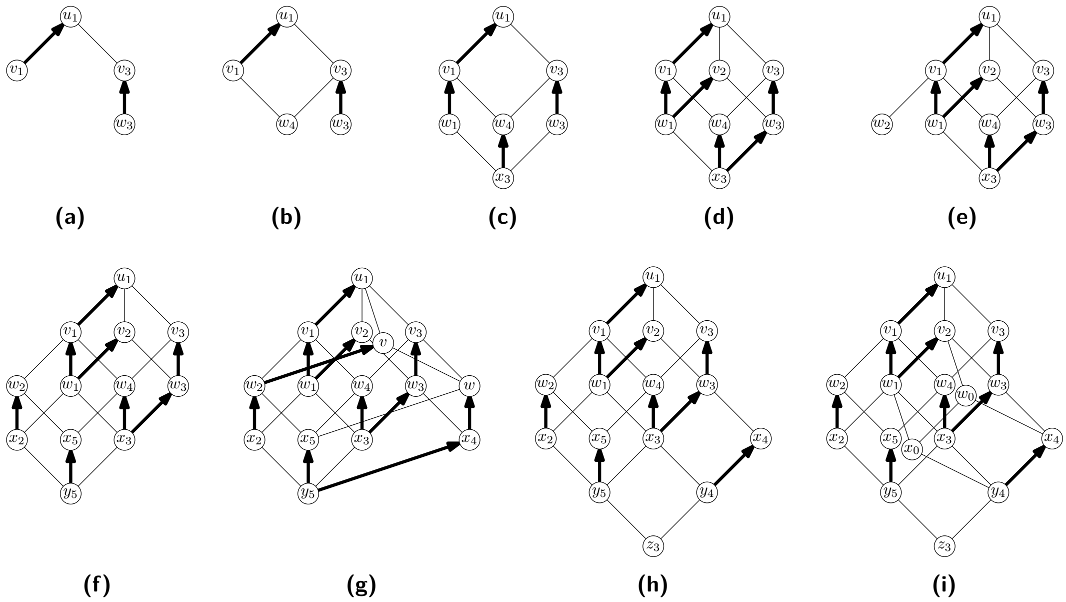

Figure 2 Animated proof of Theorem 13.

During the execution of BFS (or LexBFS), one can assume that for each vertex v the set $\Lambda(v)$ of its predecessors is computed as an ordered list $\Lambda_<(v)$ (ordered by $<$). By Lemma 8, $|\Lambda_<(v)| \leq d := \dim(G)$. Note that $<$ also gives rise to a total order on the edges of G: for two edges uv and $u'v'$ with $u < v$ and $u' < v'$ we have $uv < u'v'$ iff $u < u'$ or if $u = u'$ and $v < v'$. The next lemma characterizes the first edge in the order $<$ of each Θ-class E_i.

▶ **Lemma 12.** *An edge $uv \in E_i$ with $d(v_0, u) < d(v_0, v)$ is the first edge of E_i iff $\Lambda_<(v) = \{u\}$.*

A graph G satisfies the *fellow-traveler property* if for any LexBFS ordering of the vertices of G, for any edge uv with $v_0 \notin \{u, v\}$, the parents $f(u)$ and $f(v)$ are adjacent.

▶ **Theorem 13.** *Any median graph G satisfies the fellow-traveler property.*

Proof. Let $<$ be an arbitrary LexBFS order of the vertices of G and f be its parent map. Since any LexBFS order is a BFS order, $<$ and f satisfy the following properties of BFS:

(BFS1) if $u < v$, then $f(u) \leq f(v)$;
(BFS2) if $f(u) < f(v)$, then $u < v$;
(BFS3) if $v \neq v_0$, then $f(v) = \min_<\{u : u \sim v\}$;
(BFS4) if $u < v$ and $v \sim f(u)$, then $f(v) = f(u)$.

Notice also the following simple but useful property:

▶ **Lemma 14.** *If $abcd$ is a square of G with $d(v_0, c) = k$, $d(v_0, b) = d(v_0, d) = k+1$, $d(v_0, a) = k + 2$ and $f(a) = b$, and the edge ad satisfies the fellow-traveler property, then $f(d) = c$.*

We prove the fellow-traveler property by induction on the total order on the edges of G defined by $<$. The proof is illustrated by several figures (the arcs of the parent map are represented in bold). We use the following convention: all vertices having the same distance to the basepoint v_0 will be labeled by the same letter but will be indexed differently; for example, w_1 and w_2 are two vertices having the same distance to v_0.

Suppose by way of contradiction that $e = u_1v_3$ with $v_3 < u_1$ is the first edge in the order $<$ such that the parents $f(u_1)$ and $f(v_3)$ of u_1 and v_3 are not adjacent. Then necessarily $f(u_1) \neq v_3$. Set $v_1 = f(u_1)$ and $w_3 = f(v_3)$ (Fig. 2a). Since $d(v_0, v_1) = d(v_0, v_3)$ and $u_1 \sim v_1, v_3$, by the quadrangle condition v_1 and v_3 have a common neighbor at distance

$d(v_0, v_1) - 1$ from v_0. This vertex cannot be w_3, otherwise $f(u_1)$ and $f(v_3)$ would be adjacent. Therefore there is a vertex $w_4 \sim v_1, v_3$ at distance $d(v_0, v_1) - 1$ from v_0 (Fig. 2b). By induction hypothesis, the parent $x_3 = f(w_4)$ of w_4 is adjacent to $w_3 = f(v_3)$. Since $u_1 \sim v_1 = f(u_1), v_3$ and $v_3 \sim w_3 = f(v_3), w_4$, by (BFS3) we conclude that $v_1 < v_3$ and $w_3 < w_4$. By (BFS2), $f(v_1) \leq f(v_3)$, whence $f(v_1) \leq w_3$ and since $f(v_1) \neq f(v_3)$ (otherwise, $f(u_1) \sim f(v_3)$), we deduce that $f(v_1) < w_3 < w_4$. Hence $f(v_1) \neq w_4$. Set $w_1 = f(v_1)$. By the induction hypothesis, $f(v_1) = w_1$ is adjacent to $f(w_4) = x_3$ (Fig. 2c). By the cube condition applied to the squares $w_4v_1w_1x_3$, $w_4v_1u_1v_3$, and $w_4v_3w_3x_3$ there is a vertex v_2 adjacent to u_1, w_1, and w_3. Since $u_1 \sim v_2$ and $f(u_1) = v_1$, by (BFS3) we obtain $v_1 < v_2$. Since v_2 is adjacent to w_1 and $w_1 = f(v_1)$, by (BFS4) we obtain $f(v_2) = f(v_1) = w_1$, and by (BFS2), $v_2 < v_3$. Since $f(v_2) = w_1$, by Lemma 14 for $v_2w_1x_3w_3$, we obtain $f(w_3) = x_3$ (Fig. 2d). Since $v_1 < v_2$, $f(v_1) = f(v_2) = w_1$, and $v_2 \sim w_1, w_3$, by LexBFS v_1 is adjacent to a predecessor different from w_1 and smaller than w_3. Since $w_3 < w_4$, this predecessor cannot be w_4. Denote by w_2 the second smallest predecessor of v_1 (Fig. 2e) and note that $w_1 < w_2 < w_3 < w_4$.

By the quadrangle condition, w_2 and w_4 are adjacent to a vertex x_5, which is necessarily different from x_3 because G is $K_{2,3}$-free. By the induction hypothesis, $f(w_2)$ and $f(v_1) = w_1$ are adjacent. Then $f(w_2) \neq x_3, x_5$, otherwise we obtain a forbidden $K_{2,3}$. Set $f(w_2) = x_2$. Analogously, $f(x_5) = y_5$ and $f(w_2) = x_2$ are adjacent as well as $f(x_5) = y_5$ and $f(w_4) = x_3$ (Fig. 2f). By (BFS1), $x_2 = f(w_2) < f(w_3) = x_3$ and by (BFS3), $x_3 = f(w_4) < x_5$. Since $w_3 < w_4$ with $f(w_3) = f(w_4)$ and w_4 is adjacent to x_5, by LexBFS w_3 must have a predecessor different from x_3 and smaller than x_5. This vertex cannot be x_2 by (BFS3) since $f(w_3) = x_3$. Denote this predecessor of w_3 by x_4 and observe that $x_2 < x_3 < x_4 < x_5$. By the induction hypothesis, the parent of x_4 is adjacent to $f(w_3) = x_3$. Let $y_4 = f(x_4)$.

If $y_4 = y_5$, applying the cube condition to the squares $x_3w_3x_4y_5$, $x_3w_4x_5y_5$, and $x_3w_4v_3w_3$ we find a vertex w adjacent to x_4, v_3, and x_5. Applying the cube condition to the squares $w_4v_3wx_5$, $w_4v_1w_2x_5$, and $w_4v_1u_1v_3$ we find a vertex v adjacent to u_1, w_2, and w. Since $v \sim w_2$, by (BFS3) $f(v) \leq w_2 < w_3 = f(v_3)$, hence by (BFS2) we obtain $v < v_3$. Therefore we can apply the induction hypothesis, and by Lemma 14 for $u_1v_1w_2v$, we deduce that $f(v) = w_2$. By Lemma 14 for $v_3w_3x_4w$, we deduce that $f(w) = x_4$ (Fig. 2g). Applying the induction hypothesis to the edge vw we have that $f(v) = w_2$ is adjacent to $f(w) = x_4$, yielding a forbidden $K_{2,3}$ induced by v, x_5, x_4, w, w_2 (Fig. 2g). All this shows that $y_4 \neq y_5$. By the quadrangle condition, y_5 and y_4 have a common neighbor z_3 (Fig. 2h).

Recall that $x_2 < x_3 < x_4 < x_5$, and note that by (BFS1), $y_4 = f(x_4) < f(x_5) = y_5$. We denote by H the subgraph of G induced by the vertices $V' = \{w_1, x_2, x_3, x_4, x_5, y_4, y_5, z_3\}$. The set of edges of H is $E' = \{z_3y_4, z_3y_5, y_4x_3, y_4x_4, y_5x_2, y_5x_3, y_5x_5, x_2w_1, x_3w_1\}$. To conclude the proof, we use the following technical lemma.

▶ **Lemma 15.** *Let $H = (V', E')$ (Fig. 3a) be an induced graph of G, where $d(v_0, w_1) = d(v_0, x_2) + 1 = \cdots = d(v_0, x_5) + 1 = d(v_0, y_4) + 2 = d(v_0, y_5) + 2 = d(v_0, z_3) + 3$ and $f(x_5) = y_5$ and $f(x_4) = y_4$, such that $x_2 < x_3 < x_4 < x_5$ and $y_4 < y_5$. If G satisfies the fellow-traveler property up to distance $d(v_0, w_1)$, then there exists a vertex x_0 such that $x_0 < x_2$ and $x_0 \sim w_1, y_4$ (Fig. 3b).*

Since G contains a subgraph H satisfying the conditions of Lemma 15, there exists a vertex x_0 such that $x_0 < x_2$ and $x_0 \sim w_1, y_4$ (Fig. 2i). By the cube condition applied to the squares $x_3w_1x_0y_4$, $x_3w_1v_2w_3$, and $x_3w_3x_4y_4$, there exists $w_0 \sim x_0, v_2, x_4$ (Fig. 2i). Since x_0 is adjacent to w_0, by (BFS3) $f(w_0) \leq x_0 < x_2 = f(w_2)$. By (BFS2), $w_0 < w_2$. Recall that $f(v_1) = w_1 = f(v_2)$ and that w_2 is the second-earliest predecessor of v_1. Since $w_0 < w_2$ and w_0 is a predecessor of v_2, by LexBFS we deduce that $v_2 < v_1$. Since v_1 and v_2 are

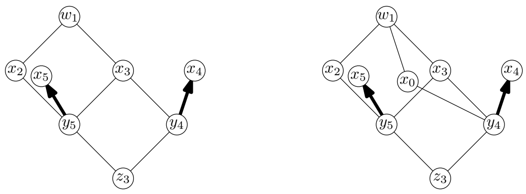

Figure 3 The induced subgraph H in Lemma 15.

both adjacent to u_1 we obtain a contradiction with $f(u_1) = v_1$. This contradiction shows that any median graph G satisfies the fellow-traveller property. This finishes the proof of Theorem 13. ◀

Now we use Theorem 13 to compute the Θ-classes of G. We run LexBFS and return a LexBFS-ordering of $V(G)$ and $E(G)$ and the ordered lists $\Lambda_<(v), v \in V$. Then consider the edges of G in the LexBFS-order. Pick the first unprocessed edge uv and suppose that $u \in \Lambda_<(v)$. If $\Lambda_<(v) = \{u\}$, by Lemma 12, uv is the first edge of its Θ-class, thus we create a new Θ-class E_i and insert uv as the first edge of E_i. We call uv the *root* of E_i and keep $d(v_0, v)$ as the distance from v_0 to H'_i. Now suppose $|\Lambda_<(v)| \geq 2$. We consider two cases: (i) $u \neq f(v)$ and (ii) $u = f(v)$. For (i), by Theorem 13, uv and $f(u)f(v)$ are opposite edges of a square. Therefore uv belongs to the Θ-class of $f(u)f(v)$ (which was already computed because $f(u)f(v) < uv$). In order to recover the Θ-class of the edge $f(u)f(v)$ in constant time, we use a (non-initialized) matrix A whose rows and columns correspond to the vertices of G such that $A[x, y]$ contains the Θ-class of the edge xy when x and y are adjacent and the Θ-class of xy has already been computed and $A[x, y]$ is undefined if x and y are not adjacent or if the Θ-class of xy has not been computed yet. For (ii), pick any $x \in \Lambda_<(v), x \neq u$. By Theorem 13, $uv = f(v)v$ and $f(x)x$ are opposite edges of a square. Since $f(x)x$ appears before uv in the LexBFS order, the Θ-class of $f(x)x$ has already been computed, and the algorithm inserts uv in the Θ-class of $f(x)x$. Each Θ-class E_i is totally ordered by the order in which the edges are inserted in E_i. Consequently, we obtain:

▶ **Theorem 16.** *The Θ-classes of a median graph G can be computed in $O(m)$ time.*

5 The median of G

We use Theorem 16 to compute the median set $\mathrm{Med}_w(G)$ of a median graph G in $O(m)$ time. We also use the existence of peripheral halfspaces and the majority rule.

5.1 Peripheral peeling

The order $E_1, E_2, \ldots, E_q$ in which the Θ-classes E_i of G are constructed corresponds to the order of the distances from v_0 to H'_i: if $i < j$ then $d(v_0, H'_i) \leq d(v_0, H'_j)$ (recall that $v_0 \in H''_i$). By Lemma 7, the halfspace H'_q of E_q is peripheral. If we contract all edges of E_q (i.e., we identify the vertices of $H'_q = \partial H'_q$ with their neighbors in $\partial H''_q$) we get a smaller median graph $\widetilde{G} = H''_q$; $\widetilde{G}$ has $q-1$ Θ-classes $\widetilde{E}_1, \ldots, \widetilde{E}_{q-1}$, where $\widetilde{E}_i$ consists of the edges of E_i in $\widetilde{G}$. The halfspaces of $\widetilde{G}$ have the form $\widetilde{H}'_i = H'_i \cap H''_q$ and $\widetilde{H}''_i = H''_i \cap H''_q$. Then $\widetilde{E}_1, \ldots, \widetilde{E}_{q-1}$ corresponds to the ordering of the halfspaces $\widetilde{H}'_1, \ldots, \widetilde{H}'_{q-1}$ of $\widetilde{G}$ by their distances to v_0. Hence the last halfspace $\widetilde{H}'_{q-1}$ is peripheral in $\widetilde{G}$. Thus the ordering $E_q, E_{q-1}, \ldots, E_1$ of the Θ-classes of G provides us with a set $G_q = G, G_{q-1} = \widetilde{G}, \ldots, G_0$ of median graphs such that G_0 is a single vertex and for each $i \geq 1$, the Θ-class E_i defines a peripheral halfspace in the graph G_i obtained after the successive contractions of the peripheral halfspaces of $G_q, G_{q-1}, \ldots, G_{i+1}$ defined by $E_q, E_{q-1}, \ldots, E_{i+1}$. We call $G_q, G_{q-1}, \ldots, G_0$ a *peripheral peeling* of G. Since each vertex of G and each Θ-class is contracted only once, we do not need to explicitly compute the restriction of each Θ-class of G to each G_i. For this it is enough to keep for each vertex v a variable indicating whether this vertex belongs to an already contracted peripheral halfspace or not. Hence, when the ith Θ-class must be contracted, we simply traverse the edges of E_i and select those edges whose both ends are not yet contracted.

5.2 Computing the weights of the halfspaces of G

We use a peripheral peeling $G_q, G_{q-1}, \ldots, G_0$ of G to compute the weights $w(H'_i)$ and $w(H''_i)$, $i = 1, \ldots, q$ of all halfspaces of G. As above, let $\widetilde{G}$ be obtained from G by contracting the Θ-class E_q. Consider the weight function $\widetilde{w}$ on $\widetilde{G} = H''_q$ defined as follows:

$$\widetilde{w}(v'') = \begin{cases} w(v'') + w(v') & \text{if } v'' \in \partial H''_q, v' \in H'_q, \text{ and } v'' \sim v', \\ w(v'') & \text{if } v'' \in H''_q \setminus \partial H''_q. \end{cases} \tag{5.1}$$

▶ **Lemma 17.** *For any Θ-class $\widetilde{E}_i$ of $\widetilde{G}$, $\widetilde{w}(\widetilde{H}'_i) = w(H'_i)$ and $\widetilde{w}(H''_i) = w(H''_i)$.*

By Lemma 17, to compute all $w(H'_i)$ and $w(H''_i)$, it suffices to compute the weight of the peripheral halfspace of E_i in the graph G_i, set it as $w(H'_i)$, and set $w(H''_i) := w(G) - w(H'_i)$.

Let G be the current median graph, let H'_q be a peripheral halfspace of G, and $\widetilde{G} = H''_q$ be the graph obtained from G by contracting the edges of E_q. To compute $w(H'_q)$, we traverse the vertices of H'_q (by considering the edges of E_q). Set $w(H''_q) = w(G) - w(H'_q)$. Let $\widetilde{w}$ be the weight function on $\widetilde{G}$ defined by Equation 5.1. Clearly, $\widetilde{w}$ can be computed in $O(|V(H'_q)|) = O(|E_q|)$ time. Then by Lemma 17 it suffices to recursively apply the algorithm to the graph $\widetilde{G}$ and the weight function $\widetilde{w}$. Since each edge of G is considered only when its Θ-class is contracted, the algorithm has complexity $O(m)$.

5.3 The median $\mathrm{Med}_w(G)$

We start with a simple property of the median function F_w that follows from Lemma 4:

▶ **Lemma 18.** *If $xy \in E_i$ with $x \in H'_i$ and $y \in H''_i$, then $F_w(x) - F_w(y) = w(H''_i) - w(H'_i)$.*

A halfspace H of G is *majoritary* if $w(H) > \frac{1}{2}w(G)$, *minoritary* if $w(H) < \frac{1}{2}w(G)$, and *egalitarian* if $w(H) = \frac{1}{2}w(G)$. Let $\mathrm{Med}^{\mathrm{loc}}_w(G) = \{v \in V : F_w(v) \leq F_w(u), \forall u \sim v\}$ be the set of local medians of G. We continue with the majority rule:

▶ **Proposition 19.** [8, 61]. *$\mathrm{Med}_w(G)$ is the intersection of all majoritary halfspaces and $\mathrm{Med}_w(G)$ intersects all egalitarian halfspaces. If H'_i and H''_i are egalitarian halfspaces, then $\mathrm{Med}_w(G)$ intersects both H'_i and H''_i. Moreover, $\mathrm{Med}_w(G) = \mathrm{Med}^{\mathrm{loc}}_w(G)$.*

We use Proposition 19 and the weights of halfspaces computed above to derive $\mathrm{Med}_w(G)$. For this, we define a new orientation of the edges $v'v''$ of each Θ-class E_i of G as follows. If $v' \in H'_i$ and $v'' \in H''_i$, then we direct $v'v''$ from v' to v'' if $w(H''_i) > w(H'_i)$ and from v'' to v' if $w(H'_i) > w(H''_i)$. If $w(H'_i) = w(H''_i)$, then the edge $v'v''$ is not directed. We denote this partially directed graph by $\overrightarrow{G}$. A vertex u of G is a *sink* of $\overrightarrow{G}$ if there is no edge uv directed in $\overrightarrow{G}$ from u to v. From Lemma 18, u is a sink of $\overrightarrow{G}$ if and only if u is a local median of G.

By Proposition 19, $\mathrm{Med}^{\mathrm{loc}}_w(G) = \mathrm{Med}_w(G)$ and thus $\mathrm{Med}_w(G)$ coincides with the set $S(\overrightarrow{G})$ of sinks of $\overrightarrow{G}$. Note that in the graph induced by $\mathrm{Med}_w(G)$, all edges are non-oriented in $\overrightarrow{G}$. Once all $w(H'_i)$ and $w(H''_i)$ have been computed, the orientation $\overrightarrow{G}$ of G can be constructed in $O(m)$ by traversing all Θ-classes E_i of G. The graph induced by $S(\overrightarrow{G})$ can then be found in $O(m)$.

▶ **Theorem 20.** *The median* $\mathrm{Med}_w(G)$ *of a median graph* G *can be computed in* $O(m)$ *time.*

▶ **Corollary 21.** *If* $w(G) > 0$*, we can find* $u, v \in V(G)$ *in* $O(m)$ *such that* $\mathrm{Med}_w(G) = I(u, v)$*.*

6 The median problem in the cube complex of G

We describe a linear time algorithm to compute medians in cube complexes of median graphs.

6.1 The main result

The problem. Let $\mathcal{G}$ be the *cube complex* of a median graph G obtained by replacing each graphic cube of G by a unit solid cube and by isometrically identifying common subcubes. We refer to $\mathcal{G}$ as to the *geometric realization* of G (see Fig. 4(a)). We suppose that $\mathcal{G}$ is endowed with the intrinsic ℓ_1-metric d_1. Let P be a finite set of points of $(\mathcal{G}, d_1)$ (called *terminals*) and let w be a weight function on $\mathcal{G}$ such that $w(p) > 0$ if $p \in P$ and $w(p) = 0$ if $p \notin P$. The goal of the *median problem* is to compute the set $\mathrm{Med}_w(\mathcal{G})$ of median points of $\mathcal{G}$, i.e., the set of all points $x \in \mathcal{G}$ minimizing the function $F_w(x) = \sum_{p\in\mathcal{G}} w(p)d_1(x, p) = \sum_{p\in P} w(p)d_1(x, p)$.

The input. The cube complex $\mathcal{G}$ is given by its 1-skeleton G. Each terminal $p \in P$ is given by its coordinates in the smallest cube $Q(p)$ of $\mathcal{G}$ containing p. Namely, we give a vertex $v(p)$ of $Q(p)$ together with its neighbors in $Q(p)$ and the coordinates of p in the embedding of $Q(p)$ as a unit cube in which $v(p)$ is the origin of coordinates. Let δ be the sum of the sizes of the encodings of the points of P. Thus the input of the median problem has size $O(m + \delta)$.

The output. Unlike $\mathrm{Med}_w(G)$ (which is a gated subgraph of G), $\mathrm{Med}_w(\mathcal{G})$ is not a subcomplex of $\mathcal{G}$. Nevertheless we show that $\mathrm{Med}_w(\mathcal{G})$ is a subcomplex of the box complex $\widehat{\mathcal{G}}$ obtained by subdividing $\mathcal{G}$, using the hyperplanes passing via the terminals of P. The output is the 1-skeleton $\widehat{M}$ of $\mathrm{Med}_w(\widehat{G})$ and $\mathrm{Med}_w(\widehat{\mathcal{G}})$, and the local coordinates of the vertices of $\widehat{M}$ in $\mathcal{G}$. We show that the output has linear size $O(m)$.

▶ **Theorem 22.** *The 1-skeleton* $\widehat{M}$ *of* $\mathrm{Med}_w(\mathcal{G})$ *can be computed in linear time* $O(m + \delta)$*.*

6.2 Geometric halfspaces and hyperplanes

In the following, we fix a basepoint v_0 of G. For each point x of $\mathcal{G}$, let $Q(x)$ be the smallest cube of $\mathcal{G}$ containing x and let $v(x)$ be the gate of v_0 in $Q(x)$. For each Θ-class E_i defining a dimension of $Q(x)$, let $\epsilon_i(x)$ be the coordinate of x along E_i in the embedding of $Q(x)$ as a unit cube in which $v(x)$ is the origin. For a Θ-class E_i and a cube Q having E_i as a dimension, the *i-midcube* of Q is the subspace of Q obtained by restricting the E_i-coordinate of Q to $\frac{1}{2}$. A *midhyperplane* $\mathfrak{h}_i$ of $\mathcal{G}$ is the union of all i-midcubes. Each $\mathfrak{h}_i$ cuts $\mathcal{G}$ in two components [58] and the union of each of these components with $\mathfrak{h}_i$ is called a *geometric halfspace* (see Fig. 4(b)). The *carrier* $\mathcal{N}_i$ of E_i is the union of all cubes of $\mathcal{G}$ intersecting $\mathfrak{h}_i$; $\mathcal{N}_i$ is isomorphic to $\mathfrak{h}_i \times [0, 1]$. For a Θ-class E_i and $0 < \epsilon < 1$, the *hyperplane* $\mathfrak{h}_i(\epsilon)$ is the set of all points $x \in \mathcal{N}_i$ such that $\epsilon_i(x) = \epsilon$. Let $\mathfrak{h}_i(0)$ and $\mathfrak{h}_i(1)$ be the respective geometric

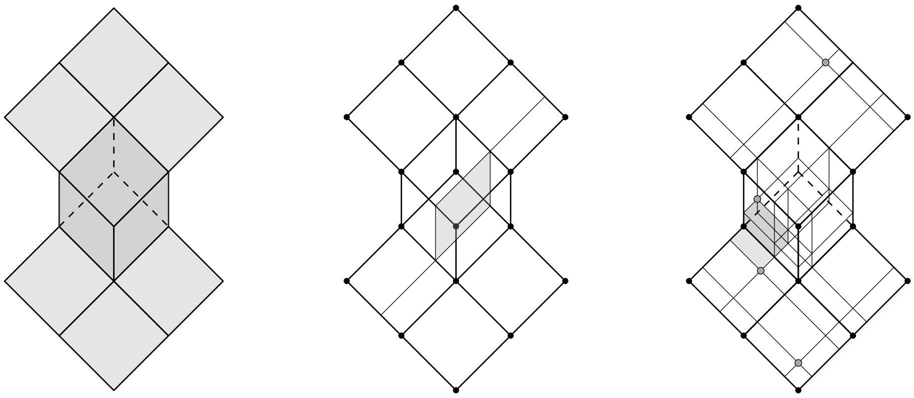

Figure 4 (a) The cube complex $\mathcal{D}$ of D, (b) a hyperplane of $\mathcal{D}$, and (c) the box complex $\widehat{\mathcal{D}}$ and $\mathrm{Med}_w(\mathcal{D})$ (in gray) defined by 4 terminals of weight 1.

realizations of $\partial H_i''$ and $\partial H_i'$. Note that $\mathfrak{h}_i(\epsilon)$ is obtained from $\mathfrak{h}_i$ by a translation. The *open carrier* $\mathcal{N}_i^\circ$ is $\mathcal{N}_i \setminus (\mathfrak{h}_i(0) \cup \mathfrak{h}_i(1))$. We denote by $\mathcal{H}_i'(\epsilon)$ and $\mathcal{H}_i''(\epsilon)$ the geometric halfspaces of $\mathcal{G}$ defined by $\mathfrak{h}_i(\epsilon)$. Let $\mathcal{H}_i'' := \mathcal{H}_i''(0)$ and $\mathcal{H}_i' := \mathcal{H}_i'(1)$; they are the geometric realizations of H_i' and H_i''. Note that $\mathcal{G}$ is the disjoint union of $\mathcal{H}_i'$, $\mathcal{H}_i''$, and $\mathcal{N}_i^\circ$.

6.3 The majority rule for $\mathcal{G}$

The box complex $\widehat{\mathcal{G}}$. By [65, Theorem 3.16], $(\mathcal{G}, d_1)$ is a median metric space (i.e., $|I(x,y) \cap I(y,z) \cap I(z,x)| = 1\ \forall x,y,z \in \mathcal{G}$). For each $p \in P$ and each coordinate $\epsilon_i(p)$, consider the hyperplane $\mathfrak{h}_i(\epsilon_i(p))$. All such hyperplanes subdivide $\mathcal{G}$ into a box complex $\widehat{\mathcal{G}}$ (see Fig. 4(c)). Clearly, $(\widehat{\mathcal{G}}, d_1)$ is a median space. By [65, Theorem 3.13], the 1-skeleton $\widehat{G}$ of $\widehat{\mathcal{G}}$ is a median graph and each point of P corresponds to a vertex of $\widehat{G}$. The Θ-classes of $\widehat{G}$ are subdivisions of the Θ-classes of G. In $\widehat{\mathcal{G}}$, all edges of a Θ-class of $\widehat{G}$ have the same length. Let $\widehat{G}_l$ be the graph $\widehat{G}$ in which the edges have these lengths. $\widehat{G}_l$ is a median space, thus $\mathrm{Med}_w(\widehat{G}_l) = \mathrm{Med}_w(\widehat{G})$ by [61]. By Proposition 19, $\mathrm{Med}_w(\widehat{G}_l)$ is the intersection of the majoritary halfspaces of $\widehat{G}$.

▶ **Proposition 23.** *$\mathrm{Med}_w(\mathcal{G})$ is the subcomplex of $\widehat{\mathcal{G}}$ defined by $\widehat{M} := \mathrm{Med}_w(\widehat{G}_l)$.*

The E_i-median problems. We adapt now Proposition 19 to the continuous setting. For a Θ-class of G, the E_i-*median* is the median of the multiset of points of the segment $[0,1]$ weighted as follows: the weight $w_i(0)$ of 0 is $w(\mathcal{H}_i'')$, the weight $w_i(1)$ of 1 is $w(\mathcal{H}_i')$, and for each $p \in P \cap \mathcal{N}_i^\circ$, there is a point $\epsilon_i(p)$ of $[0,1]$ of weight $w_i(\epsilon_i(p)) = w(p)$. It is well-known that this median is a segment $[\varrho_i'', \varrho_i']$ defined by two consecutive points $\varrho_i'' \le \varrho_i'$ of $[0,1]$ with positive weights, and for any $p \in P$, $\epsilon_i(p) \le \varrho_i''$ or $\epsilon_i(p) \ge \varrho_i'$. *Majoritary*, *minoritary*, and *egalitarian* geometric halfspaces of $\mathcal{G}$ are defined in the same way as the halfspaces of G.

▶ **Proposition 24.** *Let E_i be a Θ-class of G. Then the following holds:*

1. *$\mathrm{Med}_w(\mathcal{G}) \subseteq \mathcal{H}_i''$ (resp., $\mathrm{Med}_w(\mathcal{G}) \subseteq \mathcal{H}_i'$) if and only if $\mathcal{H}_i''$ is majoritary (resp., $\mathcal{H}_i'$ is majoritary), i.e., $\rho_i'' = \rho_i' = 0$ (resp. $\rho_i'' = \rho_i' = 1$);*
2. *$\mathrm{Med}_w(\mathcal{G}) \subseteq \mathcal{H}_i'' \cup \mathcal{N}_i^\circ$ (resp., $\mathrm{Med}_w(\mathcal{G}) \subseteq \mathcal{H}_i' \cup \mathcal{N}_i^\circ$) and $\mathrm{Med}_w(\mathcal{G})$ intersects each of the sets $\mathcal{H}_i''$ (resp., $\mathcal{H}_i'$) and $\mathcal{N}_i^\circ$ if and only if $\mathcal{H}_i''$ (resp. $\mathcal{H}_i'$) is egalitarian and $\mathcal{H}_i'$ (resp., $\mathcal{H}_i''$) is minoritary, i.e., $0 = \rho_i'' < \rho_i' < 1$ (resp. $0 < \rho_i'' < \rho_i' = 1$);*
3. *$\mathrm{Med}_w(\mathcal{G}) \subseteq \mathcal{N}_i^\circ$ if and only if $\mathcal{H}_i'$ and $\mathcal{H}_i''$ are minoritary, i.e., $0 < \rho_i'' \le \rho_i' < 1$;*
4. *$\mathrm{Med}_w(\mathcal{G})$ intersects the three sets $\mathcal{H}_i, \mathcal{H}_i''$, and $\mathcal{N}_i^\circ$ if and only if $\mathcal{H}_i'$ and $\mathcal{H}_i''$ are egalitarian, i.e., $0 = \rho_i'' \le \rho_i' = 1$ (and thus $w(\mathcal{N}_i^\circ) = 0$).*

6.4 The algorithm

Preprocessing the input. We first compute the Θ-classes $E_1, E_2, \ldots, E_q$ of G ordered by increasing distance from v_0 to H_i'. Using this, we can modify the input of the median problem in linear time $O(m+\delta)$ in such a way that for each terminal $p \in P$, $v(p)$ is the gate of v_0 in $Q(p)$. In this way, the local coordinates of the terminals of P coincide with the coordinates $\epsilon_i(p)$ defined in Section 6.2. For each Θ-class E_i, let $P_i = P \cap \mathcal{N}_i^\circ = \{p \in P : 0 < \epsilon_i(p) < 1\}$, and for each point $v \in V(G)$, let $P_v = \{p \in P : v(p) = v\}$. By traversing the points of P, we can compute all sets P_i, $1 \leq i \leq q$ and P_v, $v \in V$ and the weights of these sets in time $O(\delta)$.

Computing the E_i-medians. We first compute the weights $w_i(0) = w(\mathcal{H}_i'')$ and $w_i(1) = w(\mathcal{H}_i')$ of the geometric halfspaces $\mathcal{H}_i'', \mathcal{H}_i'$ of G. For each vertex v of G, let $w_*(v) = w(P_v)$. Note that $w_*(V) = w(P)$. Since $v_0 \in H_i''$, $w_*(H_i') = w(\mathcal{H}_i')$ and $w_*(H_i'') = w(\mathcal{H}_i'') + w(\mathcal{N}_i^o)$ for each Θ-class E_i. We apply the algorithm of Section 5.2 to G with the weight function w_* to compute the weights $w_*(H_i')$ and $w_*(H_i'')$ of all halfspaces of G. Since $w(\mathcal{N}_i^o) = w(P_i)$ is known, we can compute $w(\mathcal{H}_i') = w_*(H_i')$ and $w(\mathcal{H}_i'') = w_*(H_i'') - w(P_i)$. This allows us to complete the definition of each E_i-median problem which altogether can be solved linearly in the size of the input [27, Problem 9.2], i.e., in time $O(\Sigma_{i=1}^q(|P_i|+2)) = O(\delta + m)$.

Computing $\widehat{M}$. To compute the 1-skeleton $\widehat{M}$ of $\mathrm{Med}_w(\mathcal{G})$ in $\widehat{G}$, we orient the edges of E_i according to the weights of $\mathcal{H}_i'$ and $\mathcal{H}_i''$: $v'v'' \in E_i$ with $v' \in \mathcal{H}_i'$ and $v'' \in \mathcal{H}_i''$ is directed from v'' to v' if $\varrho_i' = \varrho_i'' = 1$ ($\mathcal{H}_i'$ is majoritary) and from v' to v'' if $\varrho_i' = \varrho_i'' = 0$ ($\mathcal{H}_i''$ is majoritary), otherwise the edges of E_i are not oriented. Denote this partially directed graph by $\overrightarrow{G}$ and let $S(\overrightarrow{G})$ be the set of sinks of $\overrightarrow{G}$. A non-directed edge $v'v'' \in E_i$ defines a *half-edge with origin* v'' if $\varrho_i'' > 0$ and a *half-edge with origin* v' if $\varrho_i' < 1$ (an edge $v'v''$ such that $0 < \varrho_i'' \leq \varrho_i' < 1$ defines two half-edges).

▶ **Proposition 25.** *For any vertex v of $\overrightarrow{G}$, all half-edges with origin v define a cube Q_v of $\mathcal{G}$.*

Proof. For any vertex v and two Θ-classes E_i, E_j defining half-edges with origin v, let v_i and v_j be the respective neighbors of v in $\widehat{G}$ along the directions E_i and E_j. By Proposition 24, vv_i and vv_j point to two majoritary halfspaces of $\widehat{G}$ (and $\mathcal{G}$). Since those two halfspaces cannot be disjoint, E_i and E_j are crossing. The proposition then follows from Lemma 5. ◀

For any cube Q of $\mathcal{G}$, let $B(Q) \subseteq Q$ be the subcomplex of $\widehat{\mathcal{G}}$ that is the Cartesian product of the E_i-medians $[\varrho_i'', \varrho_i']$ over all Θ-classes E_i wich define dimensions of Q. By the definition of the E_i-medians, $B(Q)$ is a single box of $\widehat{\mathcal{G}}$ and its vertices belong to $\widehat{G}$.

▶ **Proposition 26.** *For any cube Q of $\mathcal{G}$, if $Q \cap \mathrm{Med}_w(\mathcal{G}) \neq \emptyset$, then $B(Q) = \mathrm{Med}_w(\mathcal{G}) \cap Q$.*

Proof. If a vertex x of $B(Q)$ is not a median of $\widehat{G}$, by Proposition 19, x is not a local median of $\widehat{G}$. Thus $F_w(x) > F_w(y)$ for an edge xy of $\widehat{G}$. Suppose that xy is parallel to the edges of E_i of G. Then $\epsilon_i(x)$ coincides with ϱ_i'' or ϱ_i'. Since $F_w(x) > F_w(y)$, the halfspace $W(y,x)$ of $\widehat{G}$ is majoritary, contrary to the assumption that $\epsilon_i(x)$ is an E_i-median point. Thus all vertices of $B(Q)$ belong to $\widehat{M}$ and by Proposition 23, $B(Q) \subseteq \mathrm{Med}_w(\mathcal{G})$. It remains to show that any point of $Q \setminus B(Q)$ is not median. Otherwise, by Proposition 23 and since $\widehat{M}$ is convex, there exists a vertex $y \notin B(Q)$ of $(\widehat{M} \cap Q) \setminus B(Q)$ adjacent to a vertex x of $B(Q)$. Let xy be parallel to E_i. Then $\epsilon_i(x)$ coincides with ϱ_i'' or ϱ_i' and $\epsilon_i(y)$ does not belong to the E_i-median $[\varrho_i'', \varrho_i']$. Hence the halfspace $W(y,x)$ of $\widehat{G}$ is minoritary, contrary to $F_w(y) = F_w(x)$. ◀

For a sink v of $\overrightarrow{G}$, let $g(v)$ be the point of Q_v such that for each Θ-class E_i of Q_v, $\epsilon_i(g(v)) = \varrho'$ if $v \in \mathcal{H}'_i$ and $\epsilon_i(g(v)) = \varrho''$ if $v \in \mathcal{H}''_i$. Note that $g(v)$ is the gate of v in $B(Q_v)$ and $g(v)$ is a vertex of $\widehat{M}$. Conversely, let $x \in \widehat{M}$ and consider the cube $Q(x)$. Since $B(Q(x))$ is a cell of $\widehat{\mathcal{G}}$, for each Θ-class E_i of $Q(x)$, we have $\epsilon_i(x) \in \{\varrho'_i, \varrho''_i\}$. Let $f(x)$ be the vertex of $Q(x)$ such that $f(x) \in \mathcal{H}''_i$ if $\epsilon_i(x) = \varrho''_i$ and $f(x) \in \mathcal{H}'_i$ otherwise.

▶ **Proposition 27.** *For any $v \in S(\overrightarrow{G})$, $g(v)$ is the gate of v in $\widehat{M}$ and $\mathrm{Med}_w(\mathcal{G})$. For any $x \in \widehat{M}$, $x = g(f(x))$ is the gate of $f(x)$ in $\widehat{M}$ and $\mathrm{Med}_w(\mathcal{G})$.*

Furthermore, for any edge uv of G with $u, v \in S(\overrightarrow{G})$, either $g(u) = g(v)$ or $g(u)g(v)$ is an edge of $\widehat{M}$. Conversely, for any edge xy of $\widehat{M}$, $f(x)f(y)$ is an edge of G.

Proof. By Proposition 24 applied to $\mathcal{G}$, Proposition 19 applied to $\widehat{G}$, and the definition of sinks of $\overrightarrow{G}$, $g(v)$ is a sink of $\overrightarrow{G}$, thus $g(v)$ is a median of $\widehat{G}$ and $\mathcal{G}$. Since $B(Q_v) = \mathrm{Med}_w(\mathcal{G}) \cap Q_v$ is gated and non-empty, the gate of v in $\mathrm{Med}_w(\mathcal{G})$ belongs to $B(Q_v)$ and thus the gate of v in $\mathrm{Med}_w(\mathcal{G})$ is the gate of v on $B(Q_v)$. Conversely, $\epsilon_i(x) \notin \{0, 1\}$ for any E_i defining a dimension of $Q(x)$, thus there is an E_i-half-edge with origin $f(x)$. Pick now any E_j-edge incident to v such that E_j does not define a dimension of $Q(x)$. Without loss of generality, assume that $f(x) \in \mathcal{H}'_j$. Then $x \in \mathcal{H}'_j$, yielding $w(\mathcal{H}'_j) \geq \frac{1}{2}w(P)$. By Proposition 24, $\varrho'_j = 1$ and thus $f(x)$ is not the origin of an E_j-edge or E_j-half-edge. Consequently, $Q_{f(x)} = Q(x)$ by Proposition 25 and by the definition of $f(x)$ and $g(f(x))$, we have $x = g(f(x))$.

Let $v'v''$ be an E_i-edge between two sinks of $\overrightarrow{G}$ with $v' \in \mathcal{H}'_i$ and $v'' \in \mathcal{H}''_i$. Let $x' = g(v')$ and $x'' = g(v'')$ and assume that $x' \neq x''$. Let u', u'' be the points of $v'v''$ such that $\epsilon_i(u') = \varrho'_i$ and $\epsilon_i(u'') = \varrho''_i$. Note that u' and u'' are adjacent vertices of $\widehat{G}$ and that $u' \in I_{\widehat{G}}(v', x')$ and $u'' \in I_{\widehat{G}}(v'', x'')$. In $\widehat{G}$, x'' is the gate of u'' (and x' is the gate of u') in $\widehat{M}$. Since $d_{\widehat{G}}(u', x') + d_{\widehat{G}}(x', x'') = d_{\widehat{G}}(u', x'') \leq d_{\widehat{G}}(u'', x'') + 1$ and $d_{\widehat{G}}(u'', x'') + d_{\widehat{G}}(x', x'') = d_{\widehat{G}}(u'', x') \leq d_{\widehat{G}}(u', x') + 1$, we obtain that $d_{\widehat{G}}(x', x'') \leq 1$.

Any edge $x'x''$ of $\widehat{M}$ is parallel to a Θ-class E_i of G. For any Θ-class E_j of $Q(x')$ (resp. $Q(x'')$) with $j \neq i$, E_j is a Θ-class of $Q(x'')$ (resp. $Q(x')$) and $\epsilon_j(x') = \epsilon_j(x'')$. By their definition, $f(x')$ and $f(x'')$ can be separated only by E_i, i.e., $d_G(f(x'), f(x'')) \leq 1$. Since f is an injection from $V(\widehat{M})$ to $S(\overrightarrow{G})$, necessarily $f(x')$ and $f(x'')$ are adjacent. ◀

The algorithm computes the set $S(\overrightarrow{G})$ of all sinks of $\overrightarrow{G}$ and for each sink $v \in S(\overrightarrow{G})$, it computes the gate of $g(v)$ of v in $\widehat{M}$ and the local coordinates of $g(v)$ in $\mathcal{G}$. The algorithm returns $\left\{g(v) : v \in S(\overrightarrow{G})\right\}$ as $V(\widehat{M})$ and $\left\{g(u)g(v) : uv \in E \text{ and } u, v \in S(\overrightarrow{G})\right\}$ as $E(\widehat{M})$. Proposition 27 implies that $V(\widehat{M})$ and $E(\widehat{M})$ are correctly computed and that $\widehat{M}$ contains at most n vertices and m edges. Moreover each vertex x of $\widehat{M}$ is the gate $g(f(x))$ of the vertex $f(x)$ of $Q(x)$ that has dimension at most $\deg(f(x))$. Hence the size of the description of the vertices of $\widehat{M}$ is at most $O(m)$. This finishes the proof of Theorem 22.

References

1 A. Abboud, F. Grandoni, and V. Vassilevska Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *SODA 2015*, pages 1681–1697. SIAM, 2015. doi:10.1137/1.9781611973730.112.

2 A. Abboud, V. Vassilevska Williams, and J. R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *SODA 2016*, pages 377–391. SIAM, 2016. doi:10.1137/1.9781611974331.ch28.

3 F. Ardila, M. Owen, and S. Sullivant. Geodesics in CAT(0) cubical complexes. *Adv. in Appl. Math.*, 48(1):142–163, 2012. doi:10.1016/j.aam.2011.06.004.

4 S. P. Avann. Metric ternary distributive semi-lattices. *Proc. Amer. Math. Soc.*, 12(3):407–414, 1961.

5 M. Bacák. Computing medians and means in Hadamard spaces. *SIAM Journal on Optimization*, 24(3):1542–1566, 2014. doi:10.1137/140953393.

6 C. Bajaj. The algebraic degree of geometric optimization problems. *Discrete Comput. Geom.*, 3(2):177–191, 1988. doi:10.1007/BF02187906.

7 K. Balakrishnan, B. Brešar, M. Changat, S. Klavžar, M. Kovše, and A. R. Subhamathi. Computing median and antimedian sets in median graphs. *Algorithmica*, 57(2):207–216, 2010. doi:10.1007/s00453-008-9200-4.

8 H.-J. Bandelt and J.-P. Barthélémy. Medians in median graphs. *Discrete Appl. Math.*, 8(2):131–142, 1984. doi:10.1016/0166-218X(84)90096-9.

9 H.-J. Bandelt and V. Chepoi. Graphs with connected medians. *SIAM J. Discrete Math.*, 15(2):268–282, 2002. doi:10.1137/S089548019936360X.

10 H.-J. Bandelt and V. Chepoi. Metric graph theory and geometry: a survey. In J. E. Goodman, J. Pach, and R. Pollack, editors, *Surveys on Discrete and Computational Geometry: Twenty Years Later*, volume 453 of *Contemp. Math.*, pages 49–86. Amer. Math. Soc., Providence, RI, 2008. doi:10.1090/conm/453/08795.

11 H.-J. Bandelt, V. Chepoi, A. W. M. Dress, and J. H. Koolen. Combinatorics of lopsided sets. *European J. Combin.*, 27(5):669–689, 2006. doi:10.1016/j.ejc.2005.03.001.

12 H.-J. Bandelt, P. Forster, and A. Röhl. Median-joining networks for inferring intraspecific phylogenies. *Mol. Biol. Evol.*, 16(1):37–48, 1999. doi:10.1093/oxfordjournals.molbev.a026036.

13 J.-P. Barthélemy and J. Constantin. Median graphs, parallelism and posets. *Discrete Math.*, 111(1-3):49–63, 1993. Graph Theory and Combinatorics (Marseille-Luminy, 1990). doi:10.1016/0012-365X(93)90140-O.

14 A. Bavelas. Communication patterns in task-oriented groups. *J. Acoust. Soc. Am.*, 22(6):725–730, 1950. doi:10.1121/1.1906679SMASH.

15 M. A. Beauchamp. An improved index of centrality. *Behavorial Science*, 10:161–163, 1965.

16 L. Bénéteau, J. Chalopin, V. Chepoi, and Y. Vaxès. Medians in median graphs in linear time. *arXiv preprint*, 1907.10398, 2019. arXiv:1907.10398.

17 L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.*, 27(4):733–767, 2001. doi:10.1006/aama.2001.0759.

18 G. Birkhoff and S. A. Kiss. A ternary operation in distributive lattices. *Bull. Amer. Math. Soc.*, 53:749–752, 1947. doi:10.1090/S0002-9904-1947-08864-9.

19 B. Brešar, J. Chalopin, V. Chepoi, T. Gologranc, and D. Osajda. Bucolic complexes. *Adv. Math.*, 243:127–167, 2013. doi:10.1016/j.aim.2013.04.009.

20 J. Chalopin and V. Chepoi. A counterexample to Thiagarajan's conjecture on regular event structures. In *ICALP 2017*, volume 80 of *LIPIcs*, pages 101:1–101:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.ICALP.2017.101.

21 J. Chalopin and V. Chepoi. 1-safe Petri nets and special cube complexes: Equivalence and applications. *ACM Trans. Comput. Log.*, 20(3):17:1–17:49, 2019. doi:10.1145/3322095.

22 J. Chalopin, V. Chepoi, H. Hirai, and D. Osajda. Weakly modular graphs and nonpositive curvature. *Mem. Amer. Math. Soc.*, 2017. To appear.

23 V. Chepoi. On distance-preserving and domination elimination orderings. *SIAM J. Discrete Math.*, 11(3):414–436, 1998. doi:10.1137/S0895480195291230.

24 V. Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000. doi:10.1006/aama.1999.0677.

25 V. Chepoi. Nice labeling problem for event structures: a counterexample. *SIAM J. Comput.*, 41(4):715–727, 2012. doi:10.1137/110837760.

26 M. B. Cohen, Y. T. Lee, G. L. Miller, J. Pachocki, and A. Sidford. Geometric median in nearly linear time. In *STOC 2016*, pages 9–21. ACM, 2016. doi:10.1145/2897518.2897647.

27 T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* MIT Press, 3rd edition, 2009.

28 D. G. Corneil. Lexicographic breadth first search – A survey. In *WG 2004*, volume 3353 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2004. doi:10.1007/978-3-540-30559-0_1.

29 D. Ž. Djoković. Distance-preserving subgraphs of hypercubes. *J. Combin. Theory Ser. B*, 14(3):263–267, 1973. doi:10.1016/0095-8956(73)90010-5.

30 David Eppstein. Recognizing partial cubes in quadratic time. *J. Graph Algorithms Appl.*, 15(2):269–293, 2011. doi:10.7155/jgaa.00226.

31 D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston. *Word Processing in Groups.* Jones and Bartlett, Boston, MA, 1992.

32 A. J. Goldman and C. J. Witzgall. A localization theorem for optimal facility placement. *Transp. Sci.*, 4(4):406–409, 1970. doi:10.1287/trsc.4.4.406.

33 M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in Group Theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987. doi:10.1007/978-1-4613-9586-7_3.

34 J. Hagauer, W. Imrich, and S. Klavžar. Recognizing median graphs in subquadratic time. *Theoret. Comput. Sci.*, 215(1-2):123–136, 1999. doi:10.1016/S0304-3975(97)00136-9.

35 S. L. Hakimi. Optimum locations of switching centers and the absolute centers and medians of a graph. *Oper. Res.*, 12(3):450–459, 1964. doi:10.1287/opre.12.3.450.

36 R. Hammack, W. Imrich, and S. Klavšar. *Handbook of Product Graphs.* Discrete Math. Appl. CRC press, Boca Raton, 2nd edition, 2011. doi:10.1201/b10959.

37 K. Hayashi. A polynomial time algorithm to compute geodesics in CAT(0) cubical complexes. *Discrete Comput. Geom.*, (to appear), 2019. doi:10.1007/s00454-019-00154-2.

38 C. Jordan. Sur les assemblages de lignes. *J. Reine Angew. Math.*, 70:185–190, 1869. doi:10.1515/crll.1869.70.185.

39 J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959.

40 J. G. Kemeny and J. L. Snell. *Mathematical models in the social sciences.* MIT Press, Cambridge, Mass.-London, 1978. Reprint of the 1962 original.

41 S. Klavžar and H. M. Mulder. Median graphs: Characterizations, location theory and related structures. *J. Combin. Math. Combin. Comput.*, 30:103–127, 1999.

42 D. E. Knuth. *The Art of Computer Programming, Volume 4A, Fascicle 0: Introduction to Combinatorial Algorithms and Boolean Functions.* Addison-Wesley, 2008.

43 R. F. Love, J. G. Morris, and G. O. Wesolowsky. *Facilities location: Models and methods*, volume 7 of *Publ. Oper. Res. Ser.* North Holland, Amsterdam, 1988.

44 F. R. McMorris, H. M. Mulder, and F. S. Roberts. The median procedure on median graphs. *Discrete Appl. Math.*, 84(1-3):165–181, 1998. doi:10.1016/S0166-218X(98)00003-1.

45 H. M. Mulder. *The Interval Function of a Graph*, volume 132 of *Mathematical Centre tracts.* Mathematisch Centrum, Amsterdam, 1980.

46 H. M. Mulder. The expansion procedure for graphs. In *Contemporary methods in graph theory*, pages 459–477. Bibliographisches Inst., Mannheim, 1990.

47 H. M. Mulder and B. Novick. A tight axiomatization of the median procedure on median graphs. *Discrete Appl. Math.*, 161(6):838–846, 2013. doi:10.1016/j.dam.2012.10.027.

48 H. M. Mulder and A. Schrijver. Median graphs and Helly hypergraphs. *Discrete Math.*, 25(1):41–50, 1979. doi:10.1016/0012-365X(79)90151-1.

49 L. Nebeský. Median graphs. *Comment. Math. Univ. Carolinae*, 12:317–325, 1971.

50 M. Nielsen, G. D. Plotkin, and G. Winskel. Petri nets, event structures and domains, part I. *Theoret. Comput. Sci.*, 13(1):85–108, 1981. doi:10.1016/0304-3975(81)90112-2.

51 L. M. Ostresh. On the convergence of a class of iterative methods for solving the Weber location problem. *Oper. Res.*, 26(4):597–609, 1978. doi:10.1287/opre.26.4.597.

52 M. Owen and J. Scott Provan. A fast algorithm for computing geodesic distances in tree space. *IEEE/ACM Trans. Comput. Biology Bioinform.*, 8(1):2–13, 2011. doi:10.1109/TCBB.2010.3.

53 C. Puppe and A. Slinko. Condorcet domains, median graphs and the single-crossing property. *Econom. Theory*, 67(1):285–318, 2019. `doi:10.1007/s00199-017-1084-6`.

54 M. Roller. Poc sets, median algebras and group actions. Technical report, Univ. of Southampton, 1998.

55 D. J. Rose, R. E. Tarjan, and G. S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM J. Comput.*, 5(2):266–283, 1976. `doi:10.1137/0205021`.

56 B. Rozoy and P. S. Thiagarajan. Event structures and trace monoids. *Theoret. Comput. Sci.*, 91(2):285–313, 1991. `doi:10.1016/0304-3975(91)90087-I`.

57 G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966. `doi:10.1007/BF02289527`.

58 M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.*, s3-71(3):585–617, 1995. `doi:10.1112/plms/s3-71.3.585`.

59 M. Sageev. CAT(0) cube complexes and groups. In M. Bestvina, M. Sageev, and K. Vogtmann, editors, *Geometric Group Theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 6–53. Amer. Math. Soc., Inst. Adv. Study, 2012. `doi:10.1090/pcms/021/02`.

60 T. J. Schaefer. The complexity of satisfiability problems. In *STOC 1978*, pages 216–226. ACM, 1978. `doi:10.1145/800133.804350`.

61 P. S. Soltan and V. D. Chepoĭ. Solution of the Weber problem for discrete median metric spaces. *Trudy Tbiliss. Mat. Inst. Razmadze Akad. Nauk Gruzin. SSR*, 85:52–76, 1987.

62 B. C. Tansel, R. L. Francis, and T. J. Lowe. Location on networks: A survey. Part I: The p-center and p-median problems. *Manag. Sci.*, 29(4):482–497, 1983. `doi:10.1287/mnsc.29.4.482`.

63 P. S. Thiagarajan. Regular event structures and finite Petri nets: A conjecture. In *Formal and Natural Computing*, volume 2300 of *Lecture Notes in Comput. Sci.*, pages 244–256. Springer, 2002. `doi:10.1007/3-540-45711-9_14`.

64 P. S. Thiagarajan and S. Yang. Rabin's theorem in the concurrency setting: a conjecture. *Theoret. Comput. Sci.*, 546:225–236, 2014. `doi:10.1016/j.tcs.2014.03.010`.

65 M. van de Vel. *Theory of Convex Structures*, volume 50 of *North-Holland Math. Library*. North-Holland Publishing Co., 1993.

66 E. Weiszfeld. Sur le point pour lequel la somme des distances de n points donnés est minimum. *Tohoku Math. J.*, 43:355–386, 1937. (English transl.: *Ann. Oper. Res.*, 167:7–41, 2009).

67 D. R. Wood. On the maximum number of cliques in a graph. *Graphs Combin.*, 23(3):337–352, 2007. `doi:10.1007/s00373-007-0738-8`.

68 A. A. Zykov. On some properties of linear complexes. *Mat. Sb. (N.S.)*, 24(66)(2):163–188, 1949. (English transl.: *Amer. Math. Soc. Transl.* 79, 1952).