

Subquadratic-Time Algorithm for the Diameter and All Eccentricities on Median Graphs

Pierre Bergé¹

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP, France
IRIF, CNRS, Université de Paris, France

Guillaume Ducoffe

National Institute of Research and Development in Informatics, Bucharest, Romania
University of Bucharest, Romania

Michel Habib

IRIF, CNRS; Université de Paris Cité, France

Abstract

On sparse graphs, Roditty and Williams [2013] proved that no $O(n^{2-\varepsilon})$ -time algorithm achieves an approximation factor smaller than $\frac{3}{2}$ for the diameter problem unless SE \overline{T} H fails. We answer here an open question formulated in the literature: can we use the structural properties of median graphs to break this global quadratic barrier?

We propose the first combinatorial algorithm computing exactly all eccentricities of a median graph in truly subquadratic time. Median graphs constitute the family of graphs which is the most studied in metric graph theory because their structure represents many other discrete and geometric concepts, such as CAT(0) cube complexes. Our result generalizes a recent one, stating that there is a linear-time algorithm for computing all eccentricities in median graphs with bounded dimension d , *i.e.* the dimension of the largest induced hypercube (note that 1-dimensional median graphs are exactly the forests). This prerequisite on d is not necessarily anymore to determine all eccentricities in subquadratic time. The execution time of our algorithm is $O(n^{1.6456} \log^{O(1)} n)$.

2012 ACM Subject Classification Theory of computation \rightarrow Parameterized complexity and exact algorithms; Theory of computation \rightarrow Data structures design and analysis

Keywords and phrases Diameter, Eccentricities, Metric graph theory, Median graphs, Hypercubes

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.9

Related Version *Full Version:* <https://arxiv.org/abs/2110.02709>

1 Introduction

Median graphs can be certainly identified as the most important family of graphs in metric graph theory. They are related to numerous areas: universal algebra [3, 16], CAT(0) cube complexes [5, 20], abstract models of concurrency [11, 37], and genetics [8, 10]. Let $d(a, b)$ be the length (*i.e.* the number of edges) of the shortest (a, b) -path for $a, b \in V$ and $I(a, b)$ be the set made up of all vertices u metrically between a and b , *i.e.* $d(a, b) = d(a, u) + d(u, b)$. Median graphs are the graphs such that for any triplet of distinct vertices $x, y, z \in V$, set $I(x, y) \cap I(y, z) \cap I(z, x)$ is a singleton, containing the *median* $m(x, y, z)$ of this triplet.

The purpose of this article is to break the quadratic barrier for the computation time of certain metric parameters on median graphs. In particular, we focus on one of the most fundamental problems in algorithmic graph theory related to distances: the *diameter*. Given an undirected graph $G = (V, E)$, the diameter is the maximum distance $d(u, v)$ for all $u, v \in V$. Two vertices at maximum distance form a *diametral pair*. An even more general

¹ Corresponding author



problem consists in determining all eccentricities of the graph. The eccentricity $\text{ecc}(v)$ of a vertex v is the maximum length of a shortest path starting from v : $\text{ecc}(v) = \max_{w \in V} d(v, w)$. The diameter is thus the maximum eccentricity.

1.1 State of the art

Executing a *Breadth First Search* (BFS) from each vertex of an input graph G suffices to obtain its eccentricities in $O(n|E|)$, with $n = |V|$. As median graphs are relatively sparse, $|E| \leq n \log n$, these multiple BFSs compute all eccentricities in time $O(n^2 \log n)$ for this class of graphs. Very efficient algorithms determining the diameter already exist on other classes of graphs, for example [2, 18, 24]. Many works have also been devoted to approximation algorithms for this parameter. Chechik *et al.* [19] showed that the diameter can be approximated within a factor $\frac{3}{2}$ in time $\tilde{O}(m^{\frac{3}{2}})$ on general graphs. On sparse graphs, it was shown in [36] that no $O(n^{2-\epsilon})$ -time algorithm achieves an approximation factor smaller than $\frac{3}{2}$ for the diameter unless the Strong Exponential Time Hypothesis (SETH) fails.

Median graphs are bipartite and can be isometrically embedded into hypercubes. They are the 1-skeletons of CAT(0) cube complexes [20] and the domains of event structures [11]. They admit structural properties, such as the Mulder's convex expansion [34, 35]. They are strongly related to hypercubes retracts [4], Cartesian products and gated amalgams [5], but also Helly hypergraphs [33]. They do not contain induced $K_{2,3}$, otherwise a triplet of vertices would admit at least two medians. The *dimension* d of a median graph G is the dimension of its largest induced hypercube. The value of this parameter is at most $\lceil \log n \rceil$ and meets this upper bound when G is a hypercube. Moreover, parameter d takes part in the sparsity of median graphs: $|E| \leq dn$.

An important concept related to median graphs is the equivalence relation Θ . This is the reflexive and transitive closure of relation Θ_0 , where two edges are in Θ_0 if they are opposite in a common 4-cycle. A Θ -class is an equivalence class of Θ . Each Θ -class of a median graph forms a matching cutset, splitting the graph into two convex connected components, called *halfspaces*. The number $q \leq n$ of Θ -classes corresponds to the dimension of the hypercube in which the median graph G isometrically embeds. Value q satisfies the Euler-type formula $2n - m - q \leq 2$ [29]. A recent LexBFS-based algorithm [13] identifies the Θ -classes in linear time $O(|E|) = O(dn)$.

There exist efficient algorithms for some metric parameters on median graphs. For example, the median set and the Wiener index can be determined in $O(|E|)$ [13]. Subfamilies of median graphs have also been studied. There is an algorithm computing the diameter and the radius in linear time for squaregraphs [21]. A more recent contribution introduces a quasilinear time algorithm - running in $O(n \log^{O(1)} n)$ - for the diameter on cube-free median graphs [23], using distance and routing labeling schemes proposed in [22]. Eventually, a linear-time algorithm [15] for the diameter on constant-dimension median graphs was proposed, *i.e.* for median graphs satisfying $d = O(1)$.

The existence of a truly subquadratic-time algorithm for the diameter on all median graphs is open and was recently formulated in [13, 15, 23]. An even more ambitious question can be asked. Can this subquadratic barrier be overpassed for the problem of finding all eccentricities of a median graph? As the total size of the output is linear and this problem generalizes the diameter one, this question is legitimate. More generally, the question holds for all metric parameters (except the median set and the Wiener index for which a linear-time algorithm was recently designed). We propose here the first subquadratic-time algorithm computing all eccentricities on median graphs.

1.2 Contributions

Our first contribution in this paper is the design of a quasilinear, *i.e.* $O((\log n)^{O(1)}n)$, time algorithm computing the diameter of simplex graphs. A simplex graph $K(G) = (V_K, E_K)$ of a graph G is obtained by considering the induced complete graphs (cliques) of G as vertices V_K . Then, two of these cliques are connected by an edge if they differ by only one element: one is C , the other is $C \cup \{v\}$. These edges form the set E_K . All simplex graphs are median [5, 12]. Moreover, we observe that simplex graphs fulfil an interesting property: they admit a central vertex - representing the empty clique - and every Θ -class has an edge incident to that vertex.

First, this algorithm extends the set of median graphs for which a quasilinear time procedure computing the diameter exists. Indeed, simplex graphs form a sub-class of median graphs containing instances with unbounded dimension d .

- There is a combinatorial algorithm determining the diameter and all eccentricities of simplex graphs in $O((d^3 + \log n)n)$: Theorem 3.8, Section 3.

Second, we remark that this method can be integrated to the algorithm already proposed in [15] to compute all eccentricities of median graphs in time $O(2^{O(d \log d)}n)$. This allows us to decrease this running time. Thanks to this modification, the new algorithm proposed computes all eccentricities of a median graph in $\tilde{O}(2^{2d}n)$, where notation \tilde{O} neglects poly-logarithmic factors. Even if the algorithm stays linear for constant-dimension median graphs, observe that the dependence on d decreases, from a slightly super-exponential function to a simple exponential one.

- There is a combinatorial algorithm determining all eccentricities of median graphs in $\tilde{O}(2^{2d}n)$: Theorem 4.16, Section 4.1.

The second and main contribution in this paper is the design of a subquadratic-time dynamic programming procedure which computes all eccentricities of any median graph. Here, the linear-time simple-exponential-FPT algorithm for all eccentricities presented above plays a crucial role: it is the base case. This framework consists in partitioning recursively the input graph G into the halfspaces of its largest Θ -class. With our construction, the leaves of this recursive tree are median graphs with dimension at most $\lfloor \frac{1}{3} \log n \rfloor$ and we can apply the former linear-time FPT algorithm.

- There is a combinatorial algorithm determining all eccentricities of median graphs in $\tilde{O}(n^{\frac{2}{3}})$: Theorem 4.22, Section 4.2.

We terminate by mentioning a possible improvement of this algorithm. Based on a faster enumeration of sets of pairwise orthogonal Θ -classes, its running time can be decreased to $\tilde{O}(n^\beta)$, where $\beta = 1.6456$. Due to page limit, this extra part is omitted. A long version of this paper presents this result [14].

All these outcomes put in evidence a relationship between the design of linear-time FPT algorithms and the design of subquadratic-time algorithms determining metric parameters on median graphs. We hope that the ideas proposed to establish all these results represent interesting tools to break the quadratic barrier on other open questions.

1.3 Organization

In Section 2, we remind the definition of median graphs. The well-known properties and concepts related to them are listed, among them Θ -classes, signature, and POFs. Section 3 summarizes our contribution on simplex graphs. In Section 4, we show how to obtain a linear-time simple-exponential-FPT algorithm for all eccentricities of a median graph, parameterized

by the dimension d . Thanks to it, we propose a dynamic programming procedure to reduce the computation of eccentricities of any median graph to the same problem on a collection of median subgraphs of sub-logarithmic dimension.

2 Median graphs

In this section, we recall some notions related to distances in graphs, and more particularly median graphs. Two important tools are presented: the Θ -classes, which are equivalence classes over the edge set, and the *Pairwise Orthogonal Families* (POFs) characterizing Θ -classes belonging to a common hypercube.

2.1 Θ -classes

All graphs $G = (V, E)$ considered in this paper are undirected, unweighted, simple, finite and connected. We denote by $N(u)$ the *open neighborhood* of $u \in V$, *i.e.* the set of vertices adjacent to u in G . We extend it naturally: for any set $A \subseteq V$, the neighborhood $N(A)$ of A is the set of vertices outside A adjacent to some $u \in A$.

Given two vertices $u, v \in V$, let $d(u, v)$ be the *distance* between u and v , *i.e.* the length of the shortest (u, v) -path. The *eccentricity* $\text{ecc}(u)$ of a vertex $u \in V$ is the length of the longest shortest path starting from u . Put formally, $\text{ecc}(u)$ is the maximum value $d(u, v)$ for all $v \in V$: $\text{ecc}(u) = \max_{v \in V} d(u, v)$. The diameter of graph G is the maximum distance between two of its vertices: $\text{diam}(G) = \max_{u \in V} \text{ecc}(u)$.

We denote by $I(u, v)$ the *interval* of pair u, v . It contains exactly the vertices which lie metrically between u and v : $I(u, v) = \{x \in V : d(u, x) + d(x, v) = d(u, v)\}$. The vertices of $I(u, v)$ are lying on at least one shortest (u, v) -path.

We say that a set $H \subseteq V$ (or the induced subgraph $G[H]$) is *convex* if $I(u, v) \subseteq H$ for any pair $u, v \in H$. Moreover, we say that H is *gated* if any vertex $v \notin H$ admits a *gate* $g_H(v) \in H$, *i.e.* a vertex that belongs to all intervals $I(v, x)$, $x \in H$. For any $x \in H$, we have $d(v, g_H(v)) + d(g_H(v), x) = d(v, x)$. Gated sets are convex by definition.

Given an integer $k \geq 1$, the hypercube of dimension k , Q_k , is a graph representing all the subsets of $\{1, \dots, k\}$ as the vertex set. An edge connects two subsets if one is included into the other and they differ by only one element. Hypercube Q_2 is a *square* and Q_3 is a 3-cube.

► **Definition 2.1** (Median graph). *A graph is median if, for any triplet x, y, z of distinct vertices, the set $I(x, y) \cap I(y, z) \cap I(z, x)$ contains exactly one vertex $m(x, y, z)$ called the median of x, y, z .*

Observe that certain well-known families of graphs are median: trees, grids, square-graphs [7], and hypercubes Q_k . Median graphs are bipartite and do not contain an induced $K_{2,3}$ [5, 26, 34]. They can be obtained by Mulder's convex expansion [34, 35] starting from a single vertex.

Now, we define a parameter which has a strong influence on the study of median graphs. The dimension $d = \dim(G)$ of a median graph G is the dimension of the largest hypercube contained in G as an induced subgraph. In other words, G admits Q_d as an induced subgraph, but not Q_{d+1} . Median graphs with $d = 1$ are exactly the trees. Median graphs with $d \leq 2$ are called *cube-free* median graphs.

Figure 1 presents three examples of median graphs. (a) is a tree: $d = 1$. (b) is a cube-free median graph: it has dimension $d = 2$. To be more precise, it is a squaregraph [7], which is a sub-family of cube-free median graphs. The last one (c) is a 4-cube: it has dimension $d = 4$.

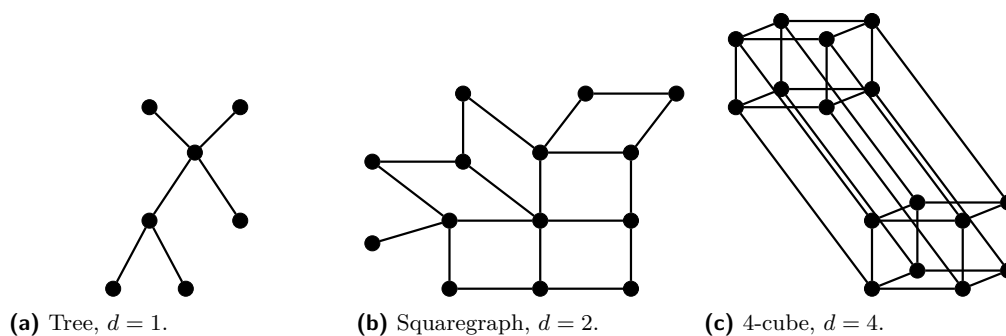


Figure 1 Examples of median graphs.

We provide a list of properties satisfied by median graphs. In particular, we define the notion of Θ -classes which is a key ingredient of several existing algorithms [13, 25, 27].

In general graphs, all gated subgraphs are convex. The reverse is true in median graphs.

► **Lemma 2.2** (Convex \Leftrightarrow Gated [5, 13]). *Any convex subgraph of a median graph is gated.*

To improve readability, edges $(u, v) \in E$ are sometimes denoted by uv . We remind the notion of Θ -class, which is well explained in [13], and enumerate some properties related to it. We say that the edges uv and xy are in relation Θ_0 if they form a square $uvyx$, where uv and xy are opposite edges. Then, Θ refers to the reflexive and transitive closure of relation Θ_0 . Let q be the number of equivalence classes obtained with this relation. The classes of the equivalence relation Θ are denoted by E_1, \dots, E_q . Concretely, two edges uv and $u'v'$ belong to the same Θ -class if there is a sequence $uv = u_0v_0, u_1v_1, \dots, u_rv_r = u'v'$ such that $u_i v_i$ and $u_{i+1} v_{i+1}$ are opposite edges of a square. We denote by \mathcal{E} the set of Θ -classes: $\mathcal{E} = \{E_1, \dots, E_q\}$. To avoid confusions, let us highlight that parameter q is different from the dimension d : for example, on trees, $d = 1$ whereas $q = n - 1$. Moreover, the dimension d is at most $\lfloor \log n \rfloor$ in general.


► **Lemma 2.3** (Θ -classes in linear time [13]). *There exists an algorithm which computes the Θ -classes E_1, \dots, E_q of a median graph in linear time $O(|E|) = O(dn)$.*

In median graphs, each class E_i , $1 \leq i \leq q$, is a perfect matching cutset and its two sides H'_i and H''_i verify nice properties, that are presented below.

► **Lemma 2.4** (Halfspaces of E_i [13, 25, 35]). *For any $1 \leq i \leq q$, the graph G deprived of edges of E_i , i.e. $G \setminus E_i = (V, E \setminus E_i)$, has two connected components H'_i and H''_i , called halfspaces. Edges of E_i form a matching: they have no endpoint in common. Halfspaces satisfy the following properties.*

- Both H'_i and H''_i are convex/gated.
- If uv is an edge of E_i with $u \in H'_i$ and $v \in H''_i$, then $H'_i = W(u, v) = \{x \in V : d(x, u) < d(x, v)\}$ and $H''_i = W(v, u) = \{x \in V : d(x, v) < d(x, u)\}$.

We denote by $\partial H'_i$ the subset of H'_i containing the vertices which are adjacent to a vertex in H''_i : $\partial H'_i = N(H''_i)$. Put differently, set $\partial H'_i$ is made up of vertices of H'_i which are endpoints of edges in E_i . Symmetrically, set $\partial H''_i$ contains the vertices of H''_i which are adjacent to H'_i . We say these sets are the *boundaries* of halfspaces H'_i and H''_i respectively. Figure 2 illustrates the notions of Θ -class, halfspace and boundary on a small example. In this particular case, an halfspace is equal to its boundary: $\partial H''_i = H''_i$. The vertices of $\partial H'_i$ are colored in blue.



■ **Figure 2** A class E_i with sets $H'_i, H''_i, \partial H'_i, \partial H''_i$.

► **Lemma 2.5** (Boundaries [13, 25, 35]). *Both $\partial H'_i$ and $\partial H''_i$ are convex/gated. Moreover, the edges of E_i define an isomorphism between $\partial H'_i$ and $\partial H''_i$.*

As a consequence, suppose uv and $u'v'$ belong to E_i : if uu' is an edge and belongs to class E_j , then vv' is an edge too and it belongs to E_j . We terminate this list of lemmas with a last property dealing with the orientation of edges from a canonical basepoint $v_0 \in V$. The v_0 -orientation of the edges of G according to v_0 is such that, for any edge uv , the orientation is \vec{uv} if $d(v_0, u) < d(v_0, v)$. Indeed, we cannot have $d(v_0, u) = d(v_0, v)$ as G is bipartite. The v_0 -orientation is acyclic.

► **Lemma 2.6** (Orientation [13]). *All edges can be oriented according to any canonical basepoint v_0 .*

From now on, we suppose that an arbitrary basepoint $v_0 \in V$ has been selected and we refer automatically to the v_0 -orientation when we mention incoming or outgoing edges.

2.2 Shortest paths and signature

We fix an arbitrary canonical basepoint v_0 and for each class E_i , we say that the halfspace containing v_0 is H'_i . Given two vertices $u, v \in V$, we define the set which contains the Θ -classes separating u from v .

► **Definition 2.7** (Signature $\sigma_{u,v}$). *We say that the signature of the pair of vertices u, v , denoted by $\sigma_{u,v}$, is the set of classes E_i such that u and v are separated in $G \setminus E_i$. In other words, u and v are in different halfspaces of E_i .*

The signature of two vertices provide us with the composition of any shortest (u, v) -path. Indeed, all shortest (u, v) -paths contain exactly one edge for each class in $\sigma_{u,v}$.

► **Lemma 2.8** ([15]). *For any shortest (u, v) -path P , the edges in P belong to classes in $\sigma_{u,v}$ and, for any $E_i \in \sigma_{u,v}$, there is exactly one edge of E_i in path P . Conversely, a path containing at most one edge of each Θ -class is a shortest path between its departure and its arrival.*

This result is a direct consequence of the convexity of halfspaces. By contradiction, a shortest path that would pass through two edges of some Θ -class E_i would escape temporarily from an halfspace, say w.l.o.g H'_i , which is convex (Lemma 2.4).

Definition 2.7 can be generalized: given a set of edges $B \subseteq E$, its signature is the set of Θ -classes represented in that set: $\{E_i : uv \in E_i \cap B\}$. The signature of a path is the set of classes which have at least one edge in this path. In this way, the signature $\sigma_{u,v}$ is also the

signature of any shortest (u, v) -path. The signature of a hypercube is the set of Θ -classes represented in its edges: the cardinality of the signature is thus equal to the dimension of the hypercube.

2.3 Orthogonal Θ -classes and hypercubes

We present now another important notion on median graphs: *orthogonality*.

► **Definition 2.9** (Orthogonal Θ -classes). *We say that classes E_i and E_j are orthogonal (denoted by $E_i \perp E_j$) if there is a square $uvyx$ in G , where $uv, xy \in E_i$ and $ux, vy \in E_j$.*

We say that E_i and E_j are *parallel* if they are not orthogonal, that is $H_i \subseteq H_j$ for some $H_i \in \{H'_i, H''_i\}$, $H_j \in \{H'_j, H''_j\}$. We define the sets of pairwise orthogonal Θ -classes.

► **Definition 2.10** (Pairwise Orthogonal Family). *We say that a set of classes $X \subseteq \mathcal{E}$ is a Pairwise Orthogonal Family (POF for short) if for any pair $E_j, E_h \in X$, we have $E_j \perp E_h$.*

For any induced hypercube of G , its *basis* (resp. *anti-basis*) is the closest vertex (resp. farthest) to v_0 in it. All edges of the hypercube incident to the basis are outgoing from it in the v_0 -orientation. Hypercubes are in bijection with pairs (u, L) , where u is a vertex (the basis of the hypercube) and L is a POF outgoing from u (the *signature* of the hypercube).

The full version of this subsection is put in Appendix A.

3 Simplex graphs

Due to page limit, the proofs in this Section are omitted. Given any undirected graph G , the vertices of the simplex graph $K(G)$ associated to G represent the induced cliques (not necessarily maximal) of G . Two of these cliques are connected by an edge if they differ by exactly one element.

► **Definition 3.1** (Simplex graphs [9]). *The simplex graph $K(G) = (V_K, E_K)$ of $G = (V, E)$ is made up of the vertex set $V_K = \{C \subseteq V : C \text{ induced complete graph of } G\}$ and the edge set $E_K = \{(C, C') : C, C' \in V_K, C \subsetneq C', |C'| - |C| = 1\}$.*

Simplex graphs can be characterized as particular median graphs.

► **Theorem 3.2.** *Let G be a median graph. The following statements are equivalent:*

- (1) *G is a simplex graph.*
- (2) *There is a vertex $v_0 \in V(G)$ such that each Θ -class of G is adjacent to v_0 , i.e. $\forall 1 \leq i \leq q, \exists v_i \in V(G), v_0 v_i \in E_i$.*
- (3) *There is a vertex $v_0 \in V(G)$ contained in any maximal hypercube of G .*

In this section only, on simplex graphs, the canonical basepoint v_0 is not selected arbitrarily. We fix v_0 as a vertex adjacent to all Θ -classes, as put in evidence by Theorem 3.2. We call v_0 the *central vertex* of the simplex graph.

► **Definition 3.3** (Crossing graphs [9, 28]). *Let G be a median graph. Its crossing graph $G^\#$ is the graph obtained by considering Θ -classes as its vertices and such that two Θ -classes are adjacent if they are orthogonal.*

Restricted to simplex graphs, this transformation is the reverse of K : indeed, as stated in [28], $G = K(G)^\#$. The clique number of $G^\#$ is exactly the dimension of median graph G . For example, the crossing graph of a cube-free median graph contains no triangle. Each simplex graph admits a central vertex (v_0 in Theorem 3.2) which represents the empty clique of $G^\#$.

Now, we focus on the problem of determining a diametral pair of a simplex graph G and more generally all eccentricities. Observe that the distance between the central vertex v_0 and any vertex u of G can be deduced directly from the set $\mathcal{E}^-(u)$ of Θ -classes incoming into u . We state that $\sigma_{v_0,u} = \mathcal{E}^-(u)$. This is a consequence of Theorem 3.2: all Θ -classes of $\mathcal{E}^-(u)$ are adjacent to v_0 , so v_0 is the basis of the hypercube with signature $\mathcal{E}^-(u)$ and anti-basis u . A shortest (v_0, u) -path is thus made up of edges of this hypercube. The distance $d(v_0, u)$ is equal to its dimension: $d(v_0, u) = |\mathcal{E}^-(u)|$.

A key result is the fact that the central vertex v_0 of the simplex graph belongs to the interval $I(u, v)$ of any pair u, v satisfying $d(u, v) = \text{ecc}(u)$.

► **Lemma 3.4.** *Let $u, v \in V(G)$ such that $d(u, v) = \text{ecc}(u)$. Then, $v_0 \in I(u, v)$.*

Two vertices u, v forming a diametral pair cannot share a common incoming Θ -class E_i , in other words $\mathcal{E}^-(u) \cap \mathcal{E}^-(v) = \emptyset$, otherwise $m = m(u, v, v_0) \in I(u, v) \subseteq H'_i$ and $v_0 \in H'_i$. Moreover, the distance $d(u, v)$ is exactly $|\mathcal{E}^-(u)| + |\mathcal{E}^-(v)|$ because $|\mathcal{E}^-(u)| = d(v_0, u)$ and $|\mathcal{E}^-(v)| = d(v_0, v)$. So, determining the diameter of a simplex graph G is equivalent to maximizing the sum $|X| + |Y|$, where X and Y are two POFs of G that are disjoint. Computing the diameter is equivalent to find the largest pair of disjoint cliques in the crossing graph $G^\#$. Similarly, the eccentricity of a vertex u is exactly the size $|\mathcal{E}^-(u)| + |\mathcal{E}^-(v)|$ of the largest pair of disjoint POFs $(\mathcal{E}^-(u), \mathcal{E}^-(v))$. Now, we can define the notion of *opposite*.

► **Definition 3.5.** *Let G be a simplex graph and X a POF of G . We denote by $\text{op}(X)$ the opposite of X , i.e. the POF Y disjoint from X with the maximum cardinality.*

$$\text{op}(X) = \underset{Y \cap X = \emptyset}{\text{argmax}} |Y|.$$

With this definition, the eccentricity of a vertex u , if we fix $X_u = \mathcal{E}^-(u)$, is written $\text{ecc}(u) = |X_u| + |\text{op}(X_u)|$. Hence, the diameter of the simplex graph G can be written as the size of the largest pair POF-opposite: $\text{diam}(G) = \max_{X \in \mathcal{L}} (|X| + |\text{op}(X)|)$.

We propose now the definition of two problems on simplex graphs. The first one, called OPPOSITES (OPP) consists in finding all pairs POF-opposite. Its output has thus a linear size. Given the solution of OPP on graph G , one can deduce both the diameter and all eccentricities in $O(n)$ time with the formula: $\text{ecc}(u) = |X_u| + |\text{op}(X_u)|$.

► **Definition 3.6 (OPP).**

Input: Simplex graph G , central vertex v_0 .

Output: For each POF X , its opposite $\text{op}(X)$.

We define an even larger version of the problem where a positive integer weight is associated with each POF. We call it WEIGHTED OPPOSITES (WOPP).

► **Definition 3.7 (WOPP).**

Input: Simplex graph G , central vertex v_0 , weight function $\omega : \mathcal{L} \rightarrow \mathbb{N}^+$.

Output: For each POF X , its weighted opposite Y maximizing $\omega(Y)$ such that $X \cap Y = \emptyset$.

Obviously, OPP is a special case of WOPP when ω is the cardinality function. We show that WOPP can be solved in quasilinear time $O((d^3 + \log n)n)$, as $d \leq \lfloor \log n \rfloor$. As a consequence, all eccentricities of a simplex graph G can also be determined with such time complexity.

► **Theorem 3.8.** *There is a combinatorial algorithm solving WOPP in time $O((d^3 + \log n)n)$. Consequently, it determines all eccentricities of a simplex graph with this running time.*

4 Subquadratic-time algorithm for all eccentricities on median graphs

This section introduces the design of algorithms computing all eccentricities for the whole class of median graphs (not only simplex graphs). We begin in Section 4.1 with the proposal of a linear-time FPT algorithm, parameterized by the dimension d , running in $2^{O(d)}n$. It is based on some techniques of a paper of the literature [15] which provides a slightly super-exponential time algorithm - running in $2^{O(d \log d)}n$ - for the same problem. We prove that replacing one step of this procedure by the partitioning conceived in Section 3 allows us to decrease the exponential dependence on d .

Thanks to this outcome, in Section 4.2, we are able to design a first subquadratic-time algorithm for all median graphs running in $\tilde{O}(n^{\frac{5}{3}})$. The proof of the Lemmas which are not given in this subsection are put in Appendix B.

4.1 Linear FPT algorithm for constant-dimension median graphs

We remind in this subsection the different steps needed to obtain a linear-time algorithm computing all eccentricities of a median graph with constant dimension, $d = O(1)$. We show how Theorem 3.8 can be integrated to it in order to improve the dependence on d . Let us begin with a reminder of the former result.

► **Lemma 4.1** ([15]). *There is a combinatorial algorithm computing all eccentricities in a median graph G with running time $\tilde{O}(2^{d(\log d+1)}n)$.*

Some parts of this subsection are redundant with [15], however we keep this subsection self-contained. The new outcomes presented are Theorems 4.10 and 4.16.

The algorithm evoked by Lemma 4.1 consists in the computation of three kinds of labels: *ladder* labels φ , *opposite* labels **op** and *anti-ladder* labels ψ . The order in which they are given correspond to their respective dependence: **op**-labelings are functions of labels φ and ψ -labelings are functions of both labels φ and **op**. The definition of **op**-labelings on general median graphs is closely related to the computation of eccentricities on simplex graphs evoked in Section 3.


4.1.1 Ladder labels

Some preliminary work has to be done before giving the definition of labels φ . We introduce the notion of *ladder set*. It is defined only for pairs of vertices u, v satisfying the condition $u \in I(v_0, v)$. In this situation, the edges of shortest (u, v) -paths are all oriented towards v with the v_0 -orientation.

► **Definition 4.2** (Ladder set $L_{u,v}$). *Let $u, v \in V$ such that $u \in I(v_0, v)$. The ladder set $L_{u,v}$ is the subset of $\sigma_{u,v}$ which contains the Θ -classes admitting an edge adjacent to u .*

Figure 3 shows a small median graph with four vertices v_0, u, v, x such that $u \in I(v_0, v)$ and $u \in I(v_0, x)$. It gives the composition of ladder sets $L_{u,v}$ and $L_{u,x}$.

A key characterization on ladder sets states that their Θ -classes are pairwise orthogonal. In brief, every set $L_{u,v}$ is a POF. Let us remind that the adjacency of all Θ -classes of a POF L with the same vertex u implies the existence of a (unique) hypercube not only signed with this POF L but also containing u (Lemma A.5). If additionally POF L is *outgoing from* u - said differently, the edges adjacent to u belonging to a Θ -class of L leave u -, then u is the basis of the hypercube. As the Θ -classes of $L_{u,v}$ are adjacent to u by definition, there is a natural bijection between (i) hypercubes (ii) pairs made up of a vertex u and a POF L outgoing from u and (iii) pairs vertex-ladder set $(u, L_{u,\cdot})$.



■ **Figure 3** Examples of ladder sets: $L_{u,v} = \{E_2, E_3\}$, $L_{u,x} = \{E_1, E_2, E_3\}$.

► **Lemma 4.3** ([15]). *Every ladder set $L_{u,v}$ is a POF. For any ordering τ of the Θ -classes in $L_{u,v}$, there is a shortest (u, v) -path such that, for any $1 \leq i \leq |L_{u,v}|$, the i^{th} first edge of the path belongs to the i^{th} Θ -class of $L_{u,v}$ in ordering τ .*

The necessary background to introduce labels φ is now known.

► **Definition 4.4** (Labels φ [15]). *Given a vertex u and a POF L outgoing from u , let $\varphi(u, L)$ be the maximum distance $d(u, v)$ such that $u \in I(v_0, v)$ and $L_{u,v} = L$.*

Intuitively, integer $\varphi(u, L)$ provides us with the maximum length of a shortest path starting from u into “direction” L . Observe that the total size of labels φ on a median graph G does not exceed $O(2^d n)$, according to Lemma A.7. We provide another notion related to orthogonality which will be used in the remainder.

► **Definition 4.5** (L -parallelism). *We say that a POF L' is L -parallel if, for any $E_j \in L'$, $L \cup \{E_j\}$ is not a POF.*

When L' is a L -parallel POF, we have $L \cap L' = \emptyset$, otherwise $L \cup \{E_j\} = L$ for some $E_j \in L'$. Presented differently, a L -parallel POF is such that any of its Θ -classes is parallel to at least one Θ -class of L .

A combinatorial algorithm running in $\tilde{O}(2^{2d} n)$ which computes all labels $\varphi(u, L)$ was identified in [15]: we provide an overview of it. There is a crucial relationship between a label $\varphi(u, L)$ and the labels of (i) the anti-basis u^+ of the hypercube with basis u and signature L and (ii) the L -parallel POFs outgoing from u^+ .

► **Lemma 4.6** (Inductive formula for labels φ [15]). *Let $u \in V$, L be a POF outgoing from u and Q be the hypercube with basis u and signature L . We denote by u^+ the opposite vertex of u in Q : u is the basis of Q and u^+ its anti-basis. A vertex $v \neq u^+$ verifies $u \in I(v_0, v)$ and $L_{u,v} = L$ if and only if (i) $u^+ \in I(v_0, v)$ and (ii) ladder set $L_{u^+,v}$ is L -parallel.*

A consequence of the previous lemma is that we can distinguish two cases for the computation of $\varphi(u, L)$. In the first case, $\varphi(u, L) = |L|$: it occurs when the farthest-to- u vertex with ladder set L is u^+ (base case). Indeed, u^+ is a candidate as $\sigma_{u,u^+} = L$: shortest (u, u^+) -paths pass through hypercube Q . This situation happens when either no Θ -class is outgoing from u^+ or when all Θ -classes outgoing from u^+ are orthogonal to L . In the second case, there are vertices farther to u than u^+ with ladder set L . As announced in Lemma 4.6, $\varphi(u, L)$ is a function of labels $\varphi(u^+, \cdot)$.

$$\varphi(u, L) = \max_{\substack{L^+ \text{ POF outgoing from } u^+ \\ \forall E_j \in L^+, L \cup \{E_j\} \text{ not POF}}} (|L| + \varphi(u^+, L^+)). \tag{1}$$

If there exists such a POF L^+ , then the label $\varphi(u, L)$ is given by Equation (1). Otherwise, it is given by the first case. Briefly, the algorithm consists in listing all pairs vertex-ladder set $((u, L), (u^+, L^+))$ such that u^+ is the anti-basis of the hypercube of basis u and signature L . For each of it, we verify whether L^+ is L -parallel. If it is, we update $\varphi(u, L)$ if $|L| + \varphi(u^+, L^+)$ is greater than the current value. The total number of pairs $((u, L), (u^+, L^+))$ is upper-bounded by $2^{2d}n$: there are at most $2^d n$ pairs (u^+, L^+) (bijection with hypercubes) and, for each of them, there are at most 2^d compatible pair (u, L) such that u^+ is the anti-basis of (u, L) . Indeed, the number of edges incoming into u^+ is at most d (Lemma A.6). For this reason, the computation of φ -labelings takes $\tilde{O}(2^{2d}n)$.

► **Theorem 4.7** (Computation of labels φ [15]). *There is a combinatorial algorithm which determines all labels $\varphi(u, L)$ in $\tilde{O}(2^{2d}n)$. It also stores, for each pair (u, L) , a vertex v satisfying $L_{u,v} = L$ and $d(u, v) = \varphi(u, L)$, denoted by $\mu(u, L)$.*

4.1.2 Opposite labels

The second type of labels needed to compute all eccentricities of a median graph G are opposite labels. Given a vertex u and a POF L outgoing from u , let $\text{op}_u(L)$ denote a POF outgoing from u with maximum label φ which is disjoint from L . As for φ , the total size of op -labelings is $O(2^d n)$.

► **Definition 4.8** (Labels op [15]). *Let $u \in V$ and L be a POF outgoing from u . Let $\text{op}_u(L)$ be one of the POF L' outgoing from u , disjoint from L , which maximizes value $\varphi(u, L')$.*

On simplex graphs, the opposite function provides in fact the op -labelings of vertex v_0 : $\text{op}(X) = \text{op}_{v_0}(X)$. As all vertices belong to hypercubes with basis v_0 , the ladder set $L_{v_0,v}$ for any vertex $v \in V$ is exactly the set $\mathcal{E}^-(v)$ of Θ -classes incoming into v . So, value $\varphi(v_0, X)$ is the distance $d(v_0, v)$ between v_0 and the only vertex v with ladder set $L_{v_0,v} = X$.

On general median graphs, the opposite label $\text{op}_u(L)$ allows us to obtain the maximum distance $d(s, t)$ such that $u = m(s, t, v_0)$ and the ladder set $L_{u,s}$ is L .

► **Lemma 4.9** (Relationship between medians and disjoint outgoing POFs [15]). *Let L, L' be two POFs outgoing from a vertex u . Let s (resp. t) be a vertex such that $u \in I(v_0, s)$ (resp. $u \in I(v_0, t)$) and $L_{u,s} = L$ (resp. $L_{u,t} = L'$). Then, $u \in I(s, t)$ if and only if $L \cap L' = \emptyset$. Therefore, given a single vertex s such that $u \in I(v_0, s)$ and $L_{u,s} = L$, the maximum distance $d(s, v)$ we can have with median $m(s, v, v_0) = u$ is exactly $d(u, s) + \varphi(u, \text{op}_u(L))$.*

Going further, given a vertex $u \in V$, the maximum distance $d(s, t)$ such that $u = m(s, t, v_0)$ is the maximum value $\varphi(u, L) + \varphi(u, \text{op}_u(L))$, where L is a POF outgoing from u .

An algorithm was initially proposed to compute all labels $\text{op}_u(L)$ consisting in a brute force bounded tree search [15]. Its execution time was $\tilde{O}(2^{O(d \log d)} n)$, leading to the global same asymptotic running time (Lemma 4.1) for finding all eccentricities.

Fortunately, the quasilinear time algorithm determining the eccentricities on simplex graphs (Theorem 3.8, Section 3) offers us the opportunity to decrease the exponential term to a simple exponential function 2^d . For any $u \in V$, let $G_u = G[V_u]$ be the *star* graph of u , using a definition from [22]. Its vertex set V_u is made up of the vertices belonging to a hypercube with basis u in G . Graph G_u is the induced subgraph of G on vertex set V_u (see Figure 8 for an example). Chepoi *et al.* [22] showed that graph G_u is a gated/convex subgraph of G . Applying the algorithm of Theorem 3.8 on the simplex graph G_u provides us with the opposite labels of u .

► **Theorem 4.10** (B.1, Computation of labels op). *There is a combinatorial algorithm which determines all labels $\text{op}_u(L)$ in $\tilde{O}(2^d n)$.*

4.1.3 Anti-ladder labels


We terminate with anti-ladder labels ψ which play the converse role of ladder labels φ . While $\varphi(u, L)$ is defined for POFs L outgoing from u , labels $\psi(u, R)$ are defined for POFs R incoming into u , *i.e.* every Θ -class of the POF has an edge entering u . As any such pair (u, R) can be associated with a hypercube of anti-basis u and signature R (Lemma A.6), the total size of ψ -labelings is at most $O(2^d n)$ too.

The notion of *milestone* intervenes in the definition of labels ψ . We consider two vertices u, v such that $u \in I(v_0, v)$. Milestones are defined recursively.

► **Definition 4.11** (Milestones $\Pi(u, v)$). *Let $L_{u,v}$ be the ladder set of u, v and u^+ be the anti-basis of the hypercube with basis u and signature $L_{u,v}$. If $u^+ = v$, then pair u, v admits two milestones: $\Pi(u, v) = \{u, v\}$. Otherwise, the set $\Pi(u, v)$ is the union of $\Pi(u^+, v)$ with vertex u : $\Pi(u, v) = \{u\} \cup \Pi(u^+, v)$.*

The milestones are the successive anti-bases of the hypercubes formed by the vertices and ladder sets traversed from u to v . Both vertices u and v are contained in $\Pi(u, v)$. The first milestone is u , the second is the anti-basis u^+ of the hypercube with basis u and signature $L_{u,v}$. The third one is the anti-basis u^{++} of the hypercube with basis u^+ and signature $L_{u^+,v}$, etc. All milestones are metrically between u and v : $\Pi(u, v) \subseteq I(u, v)$.

► **Definition 4.12** (Penultimate milestone $\bar{\pi}(u, v)$). *We say that the milestone in $\Pi(u, v)$ different from v but the closest to it is called the penultimate milestone. We denote it by $\bar{\pi}(u, v)$. Furthermore, we denote by $\bar{L}_{u,v}$ the anti-ladder set of u, v , *i.e.* the Θ -classes of the hypercube with basis $\bar{\pi}(u, v)$ and anti-basis v .*



■ **Figure 4** A pair u, v with $u \in I(v_0, v)$ and its milestones $\Pi(u, v)$ in red.

Figure 4 shows the milestones $\Pi(u, v) = \{u, u^+, u^{++}, v\}$. The hypercubes with the following pair basis-signature are highlighted with dashed edges: $(u, L_{u,v})$, $(u^+, L_{u^+,v})$, and $(u^{++}, L_{u^{++},v})$. We have $\bar{\pi}(u, v) = u^{++}$ and $\bar{L}_{u,v} = L_{u^{++},v}$ is drawn in purple.

Let R be a POF incoming to some vertex u and u^- be the basis of the hypercube with anti-basis u and signature R . Label $\psi(u, R)$ intuitively represents the maximum distance of a shortest path arriving to vertex u from “direction” R .

► **Definition 4.13** (Labels ψ [15]). *The label $\psi(u, R)$ is the maximum distance $d(u, v)$ we can obtain with a vertex v satisfying the following properties:*

- $m = m(u, v, v_0) \neq u$,
- the anti-ladder set of m, u is R : $\bar{L}_{m,u} = R$.

Equivalently, vertex u^- is the penultimate milestone of pair m, u : $u^- = \bar{\pi}(m, u)$.

As for the computation of labels φ , there is an induction process to determine all $\psi(u, R)$. As the base case, suppose that $u^- = v_0$. The largest distance $d(u, v)$ we can obtain with a

vertex v such that $v_0 \in I(u, v)$ consists in considering the opposite $\text{op}_{v_0}(R)$ of R which is outgoing from v_0 . Hence, $\psi(u, R) = |R| + \varphi(v_0, \text{op}_{v_0}(R))$.

For the induction step, we distinguish two cases. In the first one, assume that $m(u, v, v_0) = u^-$ - equivalently, $\Pi(m, u) = \Pi(u^-, u) = \{u^-, u\}$. A shortest (u, v) -path is the concatenation of the shortest (u, u^-) -path of length $|R|$ with a shortest (u^-, v) -path, and $u^- \in I(v_0, v)$. The largest distance $d(u, v)$ we can have, as for the base case, is $\psi(u, R) = |R| + \varphi(u^-, \text{op}_{u^-}(R))$.

In the second case, $m \neq u^-$, an inductive formula allows us to obtain $\psi(u, R)$. A consequence of Lemma 4.6 is that, for two consecutive milestones in $\Pi(u, v)$, say u and u^+ w.l.o.g, then $L_{u^+, v}$ is $L_{u, v}$ -parallel. This observation, applied to the penultimate milestone, provides us with the following theorem.

► **Lemma 4.14** (Inductive formula for labels ψ [15]). *Let $u, v \in V$ and $u \in I(v_0, v)$. Let L be a POF outgoing from v and w the anti-basis of hypercube (v, L) . The following propositions are equivalent:*

- (i) *vertex v is the penultimate milestone of (u, w) : $\bar{\pi}(u, w) = v$,*
- (ii) *the milestones of (u, w) are the milestones of (u, v) with w : $\Pi(u, w) = \Pi(u, v) \cup \{w\}$,*
- (iii) *the POF L is $\bar{L}_{u, v}$ -parallel.*

Set $\Pi(m, u)$ admits at least three milestones: m , u^- , and u . Let R^- be the POF incoming to u^- which is the ladder set (but also the signature) of (i) the milestone just before u^- and (ii) u^- . According to Lemma 4.14, vertex u^- is the penultimate milestone of (m, u) if and only if $R^- \cup \{E_j\}$ is not a POF, for each $E_j \in R$. For this reason, value $\psi(u, R)$ can be expressed as:

$$\psi(u, R) = \max_{\substack{R^- \text{ POF incoming to } u^- \\ \forall E_j \in R, R^- \cup \{E_j\} \text{ not POF}}} (|R| + \psi(u^-, R^-)) \quad (2)$$

Our algorithm consists in taking the maximum value between the two cases. The number of pairs $((u, R), (u^-, R^-))$ which satisfy the condition described in Equation (2) is at most $2^{2d}n$: it is identical to the one presented for φ -labelings.

► **Theorem 4.15** (Computation of labels ψ [15]). *There is a combinatorial algorithm which determines all labels $\psi(u, R)$ in $\tilde{O}(2^{2d}n)$.*

4.1.4 Better time complexity for all eccentricities

The computation of all labels $\varphi(u, L)$, $\text{op}_u(L)$ and $\psi(u, R)$ gives an algorithm which determines all eccentricities. Indeed, each eccentricity $\text{ecc}(u)$ is a function of certain labels φ and ψ . Let v be a vertex in G such that $\text{ecc}(u) = d(u, v)$. If $m = m(u, v, v_0) = u$, then $u \in I(v_0, v)$ and value $d(u, v)$ is given by a label $\varphi(u, L)$. Otherwise, if $m \neq u$, let u^- be the penultimate milestone in $\Pi(m, u)$ and R be the classes of the hypercube with basis u^- and anti-basis u . The eccentricity of u is given by a label $\psi(u, R)$. Conversely, each $\varphi(u, L)$ and $\psi(u, R)$ is the distance between u and another vertex by definition. Therefore, we have:

$$\text{ecc}(u) = \max \left\{ \max_{\substack{L \text{ POF} \\ \text{outgoing from } u}} \varphi(u, L), \max_{\substack{R \text{ POF} \\ \text{incoming to } u}} \psi(u, R) \right\} \quad (3)$$

In other words, the eccentricity of u is the maximum label φ or ψ centered at u . We can conclude with the main result of this subsection: the eccentricities of any median graph can be determined in linear time multiplied by a simple exponential function $2^{O(d)}$ of the dimension d .

► **Theorem 4.16** (All eccentricities in $\tilde{O}(2^{2d}n)$ -time for median graphs). *There is a combinatorial algorithm computing the list of all eccentricities of a median graph G in time $\tilde{O}(2^{2d}n)$.*

4.2 Tackling the general case

Our new FPT algorithm for computing the list of eccentricities in a median graph has a runtime in $2^{\mathcal{O}(d)}n$, with d being the dimension (Theorem 4.16). This runtime stays subquadratic in n as long as $d < \alpha \log n$, for some constant $\alpha < 1$. In what follows, we present a simple partitioning scheme for median graphs into convex subgraphs of dimension at most $\alpha \log n$, for an arbitrary value of $\alpha \leq 1$. By doing so, we obtain (in combination with Theorem 4.16) the first known subquadratic-time algorithm for computing all eccentricities in a median graph.

We start with a simple relation between the eccentricity function of a median graph and the respective eccentricity functions of any two complementary halfspaces.

► **Lemma 4.17** (B.2). *Let G be a median graph. For every $1 \leq i \leq q$, let $v \in V(H'_i)$ be arbitrary, and let $v^* \in \partial H''_i$ be its gate. Then, $\text{ecc}(v) = \max\{\text{ecc}_{H'_i}(v), d(v, v^*) + \text{ecc}_{H''_i}(v^*)\}$.*

We will use this above Lemma 4.17 later in our proof in order to compute in linear time the list of eccentricities in a median graph being given the lists of eccentricities in any two complementary halfspaces.

Next, we give simple properties of Θ -classes, to be used in the analysis of our main algorithm in this section (see Lemma 4.21).

► **Lemma 4.18** (B.3). *Let H and G be median graphs. If H is an induced subgraph of G then, every Θ -class of H is contained in a Θ -class of G .*

This above Lemma 4.18 can be strengthened in the special case of *isometric* subgraphs.

► **Lemma 4.19** (B.4). *Let H and G be median graphs, and let E_1, E_2, \dots, E_q denote the Θ -classes of G . If H is an isometric subgraph of G then, the Θ -classes of H are exactly the nonempty subsets among $E_i \cap E(H)$, for $1 \leq i \leq q$.*

An important consequence of Lemma 4.18 is the following relation between the dimension d of a median graph and the cardinality of its Θ -classes.

► **Lemma 4.20**. *Let G be a median graph, and let $D := \max\{|E_i| \mid 1 \leq i \leq q\}$ be the maximum cardinality of a Θ -class of G . Then, $d = \dim(G) \leq \lfloor \log D \rfloor + 1$.*

Proof. Any induced d -dimensional hypercube of G contains exactly 2^{d-1} edges of its Θ -classes, so $2^{d-1} \leq D$. ◀


We are now ready to present our main technical contribution in this section.

► **Lemma 4.21**. *If there is an algorithm for computing all eccentricities in an n -vertex median graph of dimension at most d in $\tilde{O}(c^d \cdot n)$ time, then in $\tilde{O}(n^{2 - \frac{1}{1+\log c}})$ time we can compute all eccentricities in any n -vertex median graph.*

Proof. Let G be an n -vertex median graph. We compute its Θ -classes E_1, E_2, \dots, E_q , that takes linear time (Lemma 2.3). For some parameter D (to be fixed later in the proof), let us assume without loss of generality E_1, E_2, \dots, E_p to be the subset of all Θ -classes of cardinality $\geq D$, for some $p \leq q$. Note that we have $p \leq |E|/D = \tilde{O}(n/D)$.

We reduce the problem of computing all eccentricities in G to the same problem on every connected component of $G \setminus (E_1 \cup E_2 \cup \dots \cup E_p)$. More formally, we construct a rooted binary tree T , whose leaves are labelled with convex subgraphs of G . Initially, T is reduced to a single node with label equal to G . Then, for $i = 1 \dots p$, we further refine this tree so that, at the end of any step i , its leaves are labelled with the connected components of $G \setminus (E_1 \cup E_2 \cup \dots \cup E_i)$. An example of T is shown in Figure 5 with $D = 3$ and two Θ -classes reaching this cardinality bound.

For that, we proceed as follows. We consider all leaves of T whose label H satisfies $E(H) \cap E_i \neq \emptyset$. By Lemma 4.19, $E(H) \cap E_i$ is a Θ -class of H . Both halfspaces of E_i become the left and right children of H in T . Recall that the leaves of T at this step i are the connected components of $G \setminus (E_1 \cup E_2 \cup \dots \cup E_{i-1})$, and in particular that they form a partition of $V(G)$. Therefore, each step takes linear time by reduction to computing the connected components in vertex-disjoint subgraphs of G . Overall, the total time for constructing the tree T is in $O(pm) = \tilde{O}(n^2/D)$.



■ **Figure 5** An example of tree T associated with a graph G for $D = 3$: here, $p = 2$.

Then, we compute the list of eccentricities for all the subgraphs labelling a node, by dynamic programming on T . In particular, doing so we compute the list of eccentricities for G because it is the label of the root. There are two cases:

- If H labels a leaf (base case) then, we claim that we have $\dim(H) \leq \lceil \log D \rceil + 1$. Indeed, by Lemma 4.18, every Θ -class of H is contained in a Θ -class of G . Since we removed all Θ -classes of G with at least D edges, the claim now follows from Lemma 4.20. In particular, we can compute the list of all eccentricities for H in $\tilde{O}(c^{\lceil \log D \rceil + 1} |V(H)|) = \tilde{O}(D^{\log c} |V(H)|)$ time. Recall that the leaves of T partition $V(G)$, and therefore, the total runtime for computing the list of eccentricities for the leaves is in $\tilde{O}(D^{\log c} n)$.
- From now on, let us assume H labels an internal node of T (inductive case). Let H'_i, H''_i be its children nodes, obtained from the removal of $E(H) \cap E_i$ for some $1 \leq i \leq p$. – For convenience, we will say later in the proof that H is an i -node. – Recall that $E(H) \cap E_i$ is a Θ -class of H . In particular, H'_i, H''_i are gated subgraphs. By Lemma 4.17, we can compute in $O(|V(H'_i)|)$ time the eccentricities in H of all vertices in H'_i if we are given as input: the list of eccentricities in H'_i , the list of eccentricities in H''_i , and for every $v \in V(H'_i)$ its gate $v^* \in \partial H''_i$ and the distance $d(v, v^*)$. The respective lists of eccentricities for H'_i and H''_i were pre-computed by dynamic programming on T . Furthermore, we can compute the gate v^* and $d(v, v^*)$ for every vertex $v \in V(H'_i)$, in total $\tilde{O}(|V(H)|)$ time, by using a modified BFS rooted at H''_i (we refer to [22, Lemma 17] for a detailed description of this standard procedure). Overall (by proceeding the same way for H''_i as for H'_i) we can compute the list of eccentricities for H in $\tilde{O}(|V(H)|)$ time. This is in total $\tilde{O}(n)$ time for the i -nodes (*i.e.*, because they were leaves of T at step i , and therefore, they are vertex-disjoint), and so, in total $\tilde{O}(pn) = \tilde{O}(n^2/D)$ time for all the internal nodes.

The total runtime for our algorithm is $\tilde{O}(n^2/D + D^{\log c} n)$, and optimized for $D = n^{\frac{1}{\log c + 1}}$. ◀

► **Theorem 4.22.** *There is an $\tilde{O}(n^{5/3})$ -time algorithm for computing all eccentricities in median graphs.*

Proof. This result directly follows from Theorem 4.16 with Lemma 4.21 (for $c = 4$). ◀

As a natural extension of this work, the question of designing a linear-time or quasilinear-time algorithm to compute the diameter and all eccentricities of median graphs is now open. With the recursive splitting procedure of Lemma 4.21, unfortunately, the best execution time we could obtain is $\tilde{O}(n^{\frac{3}{2}})$. Reaching this bound could represent a first reasonable objective: it would “suffice” to propose a FPT combinatorial algorithm which computes all eccentricities in $\tilde{O}(2^d n)$ in order to obtain such time complexity.

Eventually, we note two lines of research on which this paper could have some influence: (i) the study of efficient algorithms for the computation of other metric parameters on median graphs (perhaps, the *betweenness centrality* [1]) and (ii) the design of subquadratic-time algorithms for the diameter and all eccentricities on larger families of graphs (*almost-median* or *semi-median* graphs [17, 30] for example).

References

- 1 A. Abboud, F. Grandoni, and V. V. Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *Proc. of SODA*, pages 1681–1697, 2015.
- 2 A. Abboud, V. V. Williams, and J. R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proc. of SODA*, pages 377–391, 2016.
- 3 S. P. Avann. Metric ternary distributive semi-lattices. *Proc. Amer. Math. Soc.*, 12:407–414, 1961.
- 4 H. Bandelt. Retracts of hypercubes. *Journal of Graph Theory*, 8(4):501–510, 1984.
- 5 H. Bandelt and V. Chepoi. Metric graph theory and geometry: a survey. *Contemp. Math.*, 453:49–86, 2008.
- 6 H. Bandelt, V. Chepoi, A. W. M. Dress, and J. H. Koolen. Combinatorics of lopsided sets. *Eur. J. Comb.*, 27(5):669–689, 2006.
- 7 H. Bandelt, V. Chepoi, and D. Eppstein. Combinatorics and geometry of finite and infinite squaregraphs. *SIAM J. Discret. Math.*, 24(4):1399–1440, 2010.
- 8 H. Bandelt, L. Quintana-Murci, A. Salas, and V. Macaulay. The fingerprint of phantom mutations in mitochondrial dna data. *Am. J. Hum. Genet.*, 71:1150–1160, 2002.
- 9 H. Bandelt and M. van de Vel. Embedding topological median algebras in products of dendrons. *Proc. London Math. Soc.*, 58:439–453, 1989.
- 10 H. J. Bandelt, P. Forster, B. C. Sykes, and M. B. Richards. Mitochondrial portraits of human populations using median networks. *Genetics*, 141(2):743–753, 1995.
- 11 J. Barthélemy and J. Constantin. Median graphs, parallelism and posets. *Discret. Math.*, 111(1-3):49–63, 1993.
- 12 J. Barthélemy, B. Leclerc, and B. Monjardet. On the use of ordered sets in problems of comparison and consensus of classifications. *Journal of Classification*, 3:187–224, 1986.
- 13 L. Bénéteau, J. Chalopin, V. Chepoi, and Y. Vaxès. Medians in median graphs and their cube complexes in linear time. In *Proc. of ICALP*, volume 168, pages 10:1–10:17, 2020.
- 14 P. Bergé, G. Ducoffe, and M. Habib. Subquadratic-time algorithm for the diameter and all eccentricities on median graphs. *CoRR*, abs/2110.02709, 2021.
- 15 P. Bergé and M. Habib. Diameter, radius and all eccentricities in linear time for constant-dimension median graphs. In *Proc. of LAGOS*, 2021.
- 16 G. Birkhoff and S. A. Kiss. A ternary operation in distributive lattices. *Bull. Amer. Math. Soc.*, 53:745–752, 1947.
- 17 B. Bresar. Characterizing almost-median graphs. *Eur. J. Comb.*, 28(3):916–920, 2007.

- 18 S. Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. In *Proc. of SODA*, pages 2143–2152, 2017.
- 19 S. Chechik, D. H. Larkin, L. Roditty, G. Schoenebeck, R. E. Tarjan, and V. V. Williams. Better approximation algorithms for the graph diameter. In *Proc. of SODA*, pages 1041–1052, 2014.
- 20 V. Chepoi. Graphs of some CAT(0) complexes. *Adv. Appl. Math.*, 24(2):125–179, 2000.
- 21 V. Chepoi, F. F. Dragan, and Y. Vaxès. Center and diameter problems in plane triangulations and quadrangulations. In *Proc. of SODA*, pages 346–355, 2002.
- 22 V. Chepoi, A. Labourel, and S. Ratel. Distance labeling schemes for cube-free median graphs. In *Proc. of MFCS*, volume 138, pages 15:1–15:14, 2019.
- 23 G. Ducoffe. Isometric embeddings in trees and their use in distance problems. In *Proc. of MFCS*, volume 202 of *LIPICs*, pages 43:1–43:16, 2021.
- 24 G. Ducoffe, M. Habib, and L. Viennot. Diameter computation on H -minor free graphs and graphs of bounded (distance) VC-dimension. In *Proc. of SODA*, pages 1905–1922, 2020.
- 25 J. Hagauer, W. Imrich, and S. Klavzar. Recognizing median graphs in subquadratic time. *Theor. Comput. Sci.*, 215(1-2):123–136, 1999.
- 26 R. Hammack, W. Imrich, and S. Klavzar. *Handbook of Product Graphs, Second Edition*. CRC Press, Inc., 2011.
- 27 W. Imrich, S. Klavzar, and H. M. Mulder. Median graphs and triangle-free graphs. *SIAM J. Discret. Math.*, 12(1):111–118, 1999.
- 28 S. Klavzar and H. M. Mulder. Partial cubes and crossing graphs. *SIAM J. Discret. Math.*, 15(2):235–251, 2002.
- 29 S. Klavzar, H. M. Mulder, and R. Skrekovski. An Euler-type formula for median graphs. *Discret. Math.*, 187(1-3):255–258, 1998.
- 30 S. Klavzar and S. V. Shpectorov. Characterizing almost-median graphs II. *Discret. Math.*, 312(2):462–464, 2012.
- 31 M. Kovse. Complexity of phylogenetic networks: counting cubes in median graphs and related problems. *Analysis of complex networks: From Biology to Linguistics*, pages 323–350, 2009.
- 32 F. R. McMorris, H. M. Mulder, and F. S. Roberts. The median procedure on median graphs. *Discret. Appl. Math.*, 84(1-3):165–181, 1998.
- 33 H. M. Mulder and A. Schrijver. Median graphs and Helly hypergraphs. *Discret. Math.*, 25(1):41–50, 1979.
- 34 M. Mulder. The structure of median graphs. *Discret. Math.*, 24(2):197–204, 1978.
- 35 M. Mulder. The interval function of a graph. *Mathematical Centre Tracts, Mathematisch Centrum, Amsterdam*, 1980.
- 36 L. Roditty and V. V. Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proc. of STOC*, pages 515–524, 2013.
- 37 V. Sassone, M. Nielsen, and G. Winskel. A classification of models for concurrency. In *Proc. of CONCUR*, volume 715 of *Lecture Notes in Computer Science*, pages 82–96, 1993.
- 38 Peter M Winkler. Isometric embedding in products of complete graphs. *Discrete Applied Mathematics*, 7(2):221–225, 1984.

A Orthogonal Θ -classes and hypercubes

We present now another important notion on median graphs: *orthogonality*. In [31], Kovse studied a relationship between *splits* which refer to the halfspaces of Θ -classes. It says that two splits $\{H'_i, H''_i\}$ and $\{H'_j, H''_j\}$ are *incompatible* if the four sets $H'_i \cap H'_j$, $H'_i \cap H''_j$, $H''_i \cap H'_j$, and $H''_i \cap H''_j$ are nonempty. Another definition was proven equivalent to this one.

► **Definition A.1** (Orthogonal Θ -classes). *We say that classes E_i and E_j are orthogonal ($E_i \perp E_j$) if there is a square $wvxy$ in G , where $uv, xy \in E_i$ and $ux, vy \in E_j$.*

Indeed, classes E_i and E_j are orthogonal if and only if the splits produced by their halfspaces are incompatible.

► **Lemma A.2** (Orthogonal \Leftrightarrow Incompatible [15]). *Given two Θ -classes E_i and E_j of a median graph G , the following statements are equivalent:*

- *Classes E_i and E_j are orthogonal,*
- *Splits $\{H'_i, H''_i\}$ and $\{H'_j, H''_j\}$ are incompatible,*
- *The four sets $\partial H'_i \cap \partial H'_j$, $\partial H''_i \cap \partial H''_j$, $\partial H'_i \cap \partial H''_j$, and $\partial H''_i \cap \partial H'_j$ are nonempty.*

The concept of *orthogonality* is sometimes described with different words in the literature depending on the context: *incompatible*, *concurrent* or *crossing*. We say that E_i and E_j are *parallel* if they are not orthogonal, that is $H_i \subseteq H_j$ for some $H_i \in \{H'_i, H''_i\}$ and $H_j \in \{H'_j, H''_j\}$.

We pursue with a property on orthogonal Θ -classes: if two edges of two orthogonal classes E_i and E_j are incident, they belong to a common square.

► **Lemma A.3** (Squares [11, 15]). *Let $xu \in E_i$ and $uy \in E_j$. If E_i and E_j are orthogonal, then there is a vertex v such that uyv is a square.*

Pairwise orthogonal families. We focus on the set of Θ -classes which are pairwise orthogonal.

► **Definition A.4** (Pairwise Orthogonal Family). *We say that a set of classes $X \subseteq \mathcal{E}$ is a Pairwise Orthogonal Family (POF for short) if for any pair $E_j, E_h \in X$, we have $E_j \perp E_h$.*

This notion is not completely new, since it implicitly appears in certain properties established on median graphs (for instance, the *downward cube* property in [13]). The empty set is considered as a POF. We denote by \mathcal{L} the set of POFs of the median graph G . The notion of POF is strongly related to the induced hypercubes in median graphs. First, observe that all Θ -classes of a median graph form a POF if and only if the graph is a hypercube of dimension $\log n$ [31, 32]. Secondly, the next lemma precises the relationship between POFs and hypercubes.

► **Lemma A.5** (POFs adjacent to a vertex [15]). *Let X be a POF, $v \in V$, and assume that for each $E_i \in X$, there is an edge of E_i adjacent to v . There exists a hypercube Q containing vertex v and all edges of X adjacent to v . Moreover, the Θ -classes of the edges of Q are the classes of X .*

There is a natural bijection between the vertices of a median graph and the POFs. The next lemma exhibits this relationship.

► **Lemma A.6** (POFs and hypercubes [6, 8, 31]). *Consider an arbitrary canonical basepoint $v_0 \in V$ and the v_0 -orientation for the median graph G . Given a vertex $v \in V$, let $N^-(v)$ be the set of edges going into v according to the v_0 -orientation. Let $\mathcal{E}^-(v)$ be the classes of the edges in $N^-(v)$. The following propositions are true:*

- *For any vertex $v \in V$, $\mathcal{E}^-(v)$ is a POF. Moreover, vertex v and the edges of $N^-(v)$ belong to an induced hypercube formed by the classes $\mathcal{E}^-(v)$. Hence, $|\mathcal{E}^-(v)| = |N^-(v)| \leq d$.*
- *For any POF X , there is a unique vertex v_X such that $\mathcal{E}^-(v_X) = X$. Vertex v_X is the closest-to- v_0 vertex v such that $X \subseteq \mathcal{E}^-(v)$.*
- *The number of POFs in G is equal to the number n of vertices: $n = |\mathcal{L}|$.*

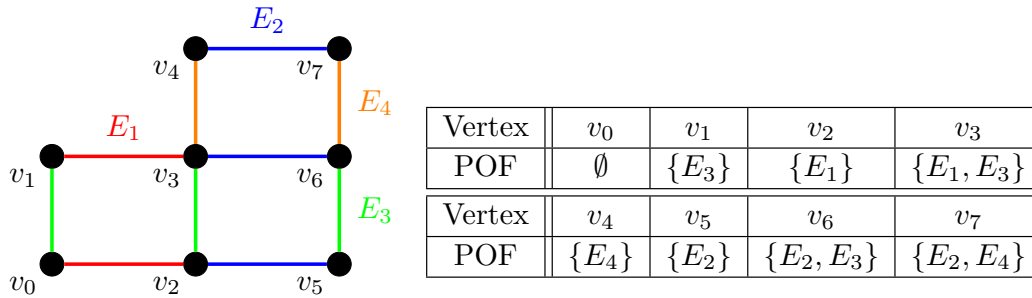


Figure 6 Illustration of the bijection between V and the set of POFs.

An example is given in Figure 6 with a small median graph of dimension $d = 2$. v_0 is the canonical basepoint and edges are colored according to their Θ -class. For example, $v_1v_3 \in E_1$. We associate with any POF X of G the vertex v_X satisfying $\mathcal{E}^-(v_X) = X$ with the v_0 -orientation. Obviously, the empty POF is associated with v_0 which has no incoming edges.

A straightforward consequence of this bijection is that parameter q , the number of Θ -classes, is less than the number of vertices n . But it can be used less trivially to enumerate the POFs of a median graph in linear time [8, 31]. Given a basepoint v_0 , we say that the *basis* (resp. *anti-basis*) of an induced hypercube Q is the single vertex v such that all edges of the hypercube adjacent to v are outgoing from (resp. incoming into) v . Said differently, the basis of Q is its closest-to- v_0 vertex and its anti-basis is its farthest-to- v_0 vertex. What Lemma A.6 states is also that we can associate with any POF X a hypercube Q_X which contains exactly the classes X and admits v_X as its anti-basis. This observation implies that the number of POFs is less than the number of hypercubes in G . Moreover, the hypercube Q_X is the closest-to- v_0 hypercube formed with the classes in X . Figure 7a shows a vertex v with its incoming and outgoing edges with the v_0 -orientation. The dashed edges represent the hypercube with anti-basis v and POF $\mathcal{E}^-(v)$.

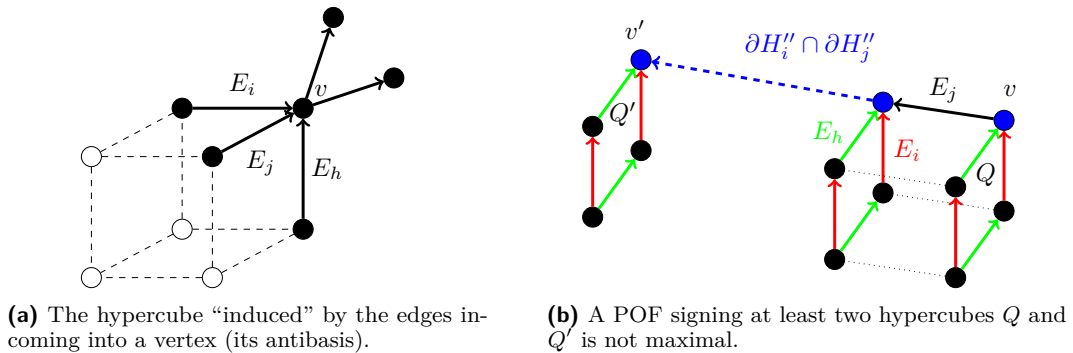


Figure 7 Properties of POFs.

Number of hypercubes. We remind a formula establishing a relationship between the number of POFs and the number of hypercubes in the literature. Let $\alpha(G)$ (resp. $\beta(G)$) be the number of hypercubes (resp. POFs) in G . Let $\beta_i(G)$ be the number of POFs of cardinality $i \leq d$ in G . According to [8, 31], we have:

$$\alpha(G) = \sum_{i=0}^d 2^i \beta_i(G). \tag{4}$$

Equation (4) produces a natural upper bound for the number of hypercubes.

► **Lemma A.7** (Number of hypercubes). $\alpha(G) \leq 2^d n$.

Value $\alpha(G)$ consider all hypercubes, in particular those of dimension 0, *i.e.* vertices. From now on, the word “hypercube” refers to the hypercubes of dimension at least one.

Each hypercube in the median graph G can be defined with only its anti-basis v and the edges \widehat{N} of the hypercube that are adjacent and going into v according to the v_0 -orientation. These edges are a subset of $N^-(v)$: $\widehat{N} \subseteq N^-(v)$. Conversely, given a vertex v , each subset of $N^-(v)$ produces a hypercube which admits v as an anti-basis (this hypercube is a sub-hypercube of the one obtained with v and $N^-(v)$, Lemma A.6). Another possible bijection is to consider a hypercube as a pair composed of its anti-basis v and the Θ -classes $\widehat{\mathcal{E}}$ of the edges in \widehat{N} (its signature).

As a consequence, a simple graph search as BFS enables us to enumerate the hypercubes in G in time $O(d2^d n)$.


► **Lemma A.8** (Enumeration of hypercubes [15]). *We can enumerate all triplets $(v, u, \widehat{\mathcal{E}})$, where v is the anti-basis of a hypercube Q , u its basis, and $\widehat{\mathcal{E}}$ the signature of Q in time $O(d2^d n)$. Moreover, the list obtained fulfils the following partial order: if $d(v_0, v) < d(v_0, v')$, then any triplet $(v, u, \widehat{\mathcal{E}})$ containing v appears before any triplet $(v', u', \widehat{\mathcal{E}}')$ containing v' .*

The enumeration of hypercubes is thus executed in linear time for median graphs with constant dimension. In summary, given any median graph, one can compute the set of Θ -classes and their orthogonality relationship (for each E_i , the set of Θ -classes orthogonal to E_i) in linear time, and the set of hypercubes with its basis, anti-basis and signature in $\tilde{O}(2^d n)$.

B Proofs of Section 4

► **Theorem B.1** (Computation of labels op). *There is a combinatorial algorithm which determines all labels $op_u(L)$ in $\tilde{O}(2^d n)$.*

Proof. Let $u \in V$: we denote by N_u the number of hypercubes of G with basis u . Convex subgraphs of median graphs are also median by considering the original definition of median graphs (Definition 2.1). Consequently, star graph G_u is median and all its maximal hypercubes contain a common vertex u . From Theorem 3.2, G_u is a simplex graph.



(a) A v_0 -oriented median graph G and a vertex $u \in V$. (b) Star graph G_u .

■ **Figure 8** Example of star graph G_u .

Any pair (u, L) of G , where L is a POF outgoing from u in G , can be associated to a unique hypercube with signature L and basis u . Thus, there is a natural bijection between (i) the POFs of G_u (ii) the vertices of G_u and (iii) the POFs L of G outgoing from u . Hence, $|V_u| = N_u$.

We associate with any POF L of G_u the weight $\omega_u(L) = \varphi(u, L)$. We apply the algorithm evoked in Theorem 3.8. The opposite computed with that configuration correspond exactly to the labels $\text{op}_u(L)$: a POF L' disjoint from L and maximizing $\varphi(u, L')$ among all POFs outgoing from u . The running time of the algorithm is $O((d^3 + \log |V_u|) |V_u|) = O((d^3 + \log n)N_u)$. Doing it for every vertex u of G , we obtain all opposite labels of G in $O(d^3 + \log n)2^d n$ as $\sum_{u \in V} N_u = 2^d n$ (Lemma A.7). ◀

► **Lemma B.2.** *Let G be a median graph. For every $1 \leq i \leq q$, let $v \in V(H'_i)$ be arbitrary, and let v^* be its gate in $\partial H''_i$. Then, $\text{ecc}(v) = \max\{\text{ecc}_{H'_i}(v), d(v, v^*) + \text{ecc}_{H''_i}(v^*)\}$.*

Proof. We have $\text{ecc}(v) = \text{ecc}_G(v) = \max\{d(u, v) \mid u \in V(H'_i)\} \cup \{d(w, v) \mid w \in V(H''_i)\}$. Since H'_i is convex, we have $\max\{d(u, v) \mid u \in V(H'_i)\} = \text{ecc}_{H'_i}(v)$. In the same way, since H''_i is gated (and so, convex), we have $\max\{d(w, v) \mid w \in V(H''_i)\} = d(v, v^*) + \max\{d(v^*, w) \mid w \in V(H''_i)\} = d(v, v^*) + \text{ecc}_{H''_i}(v^*)$. ◀

► **Lemma B.3.** *Let H and G be median graphs. If H is an induced subgraph of G then, every Θ -class of H is contained in a Θ -class of G .*

Proof. Every square of H is also a square of G . In particular, two edges of H are in relation Θ_0 if and only if, as edges of G , they are also in relation Θ_0 . Since the Θ -classes of H (resp., of G) are the transitive closure of its relation Θ_0 , it follows that every Θ -class of H is contained in a Θ -class of G . ◀

► **Lemma B.4.** *Let H and G be median graphs, and let E_1, E_2, \dots, E_q denote the Θ -classes of G . If H is an isometric subgraph of G then, the Θ -classes of H are exactly the nonempty subsets among $E_i \cap E(H)$, for $1 \leq i \leq q$.*

Proof. It is known [38] that two edges uv, xy of G are in the same Θ -class if and only if $d_G(u, x) + d_G(v, y) \neq d_G(u, y) + d_G(v, x)$. In particular, since H is isometric in G , two edges of H are in the same Θ -class of H if and only if they are in the same Θ -class of G . ◀