

A self-stabilizing algorithm for the median problem in partial rectangular grids and their relatives

Victor Chepoi, Tristan F evat, Emmanuel Godard, and Yann Vax es

LIF-Laboratoire d'Informatique Fondamentale de Marseille, UMR 6166, Marseille, France
{chepoi,fevat,godard,vaxes}@lif.univ-mrs.fr

Abstract. Given a graph $G = (V, E)$, a vertex v of G is a *median vertex* if it minimizes the sum of the distances to all other vertices of G . The median problem consists in finding the set of all median vertices of G . In this note we present a self-stabilizing algorithm for the median problem in partial rectangular grids. Our algorithm is based on the fact that partial rectangular grids can be isometrically embedded into the Cartesian product of two trees, to which we apply the algorithm proposed by Antonoiu, Srimani (1999) and Bruell, Ghosh, Karaata, Pemmaraju (1999) for computing the medians in trees. Then we extend our approach from partial rectangular grids to a more general class of plane quadrangulations.

1 Introduction

Given a (finite, connected) graph G one is sometimes interested in finding the vertices minimizing the total distance $\sum_u d(u, x)$ to the vertices u of G , where the distance $d(u, x)$ between u and x is the length of a shortest path connecting u and x . A vertex x minimizing this expression is called a *median* (vertex) of G . The median problem consists in finding the set $\text{Med}(G)$ of all medians vertices. The median problem arises with one of the basic models in discrete facility location [34, 38] and with majority consensus in classification and data analysis [7, 9, 10]. This is also a classical topic in graph theory [11, 31, 38]. Linear time algorithms for computing medians are known for several classes of graphs: among them are trees [38], planar quadrangulations and triangulations with degree-constraints, partial rectangular grids [18], and a few other classes of graphs. A distributed algorithm for computing medians in graphs is given in [35]. Finally notice that some papers [7, 37] investigate the structural properties of median sets in special classes of graphs.

In distributed systems, the median is a suitable location for information exchange and communication. Indeed, to place a common resource at a median site minimizes the cost of sharing the resource with other sites. Note also that [26] shows that, given a network of a tree topology, choosing a median and then routing all the information through it minimizes the number of messages sent during any execution of any distributed sorting algorithm. Moreover, partial rectangular grids and trees are among the most used topologies in the design of microprocessors and distributed architectures. It is therefore of important practical interest to solve the median problem in a distributed setting on such a topology.

A distributed system can be defined as a set of processors exchanging information between neighbors. The system state, called *global state* or *configuration*, is the union of all the local states. A processor only has a partial knowledge of the system, this knowledge varying according to the system connectivity for example. It is in general desirable to maintain the system in a certain set of states, *the legitimate states*. An algorithm running in a system is said to be *self-stabilizing* if any execution has a suffix in the set of the legitimate states [22]. Self-stabilization is very desirable and useful in distributed systems because it provides immunity to transient failures and can even make possible in some cases a dynamical and transparent modification of the system topology. Moreover self-stabilizing algorithms are often elegant and simple. The self-stabilizing paradigm was introduced by Dijkstra in 1973 [21]. He gave three self-stabilizing mutual exclusion algorithms on rings, opening a field of research still extremely dynamic today in distributed calculus. Self-stabilizing algorithms have been conceived to answer problems of routing [19, 20], synchronization [6, 32], leader election [4, 23], spanning tree construction [2, 5], maximum

flow and maximum matching [27, 30], graph coloring [28], and mutual exclusion in several kinds of networks [36].

Antonoiu, Srimani [3] and Bruell, Ghosh Karaata, Pemmaraju [13] proposed a strikingly simple and nice self-stabilizing algorithm for computing the median set of a tree T . The state of each node is an integer s . At each step, the algorithm updates the s -value of the currently active vertex v : it sets $s(v) = 1$ if v is a leaf, otherwise the algorithm computes the sum of the s -values of all neighbors of v minus the largest s -value of a neighbor and then add 1 (denoted by $1 + \sum(N_s^-(v))$). If the current s -value of v is different from $1 + \sum(N_s^-(v))$, then this value becomes the new s -value of v . The algorithm terminates when there is no s -value to modify. Interestingly, this happens to be a “valid” global state as the medians of T are the vertices with maximum s -values. The authors of [3, 13] establish that the algorithm stabilizes in a polynomial number of steps. See also [3, 13, 12, 33] for self-stabilizing algorithms solving other facility location problems on tree-networks.

In this note, we propose self-stabilizing median computation algorithms for two classes of plane graphs: partial rectangular grid and even squaregraphs. Our algorithms are based on the fact that such graphs isometrically embed into the Cartesian product of two trees and that the median in the initial graph G can be derived from the medians in two spanning trees of G closely related to the two tree-factors. Using the sense of direction in the grid, the algorithm computes the tree-factors and apply the algorithm of [3, 13] to compute the medians of both spanning trees. This computation is performed anonymously. For an even squaregraph G , the algorithm first computes a spanning tree of G using the self-stabilizing spanning tree algorithm of Afek, Kutten, and Yung [1]. This algorithm needs unique identities to be available for every node of the network. Then the algorithm “repairs” this spanning tree in order to produce the two tree-factors and their spanning trees relatives, to which the median computation algorithm of [3, 13] is applied. The algorithms have a round complexity of $O(e)$ and $O(n^2)$ respectively, where e is the maximum distance from a vertex to a median vertex. By self-stabilization, our algorithms are resilient to transient failures. Concerning non-transient failures like the permanent crash of a link or a processor, if the resulting topology is still a partial grid or an even squaregraph, then the algorithm will dynamically adjust to the changes. If not, and the crash creates a hole in the partial grid, then it is of course likely that our algorithm will not find the correct medians or even will not converge, because the computed factors could then contain cycles.

The article is organized as follows. In the first part, we investigate the properties of partial grids, squaregraphs and their medians which are used in the algorithm. In the second part, we describe the algorithms for computing their medians and give a proof of the correctness as well as an upper bound of the time and round complexity. This is, to our knowledge, the first self-stabilizing algorithm for location problems on non-tree networks.

2 Partial grids, squaregraphs and their medians

2.1 Preliminaries

In a graph $G = (V, E)$, the *length* of a path from a vertex x to a vertex y is the number of edges in the path. The *distance* $d_G(x, y)$ (or $d(x, y)$ if G is obvious from the context) between x to y is the length of a shortest path connecting x and y . The interval $I(u, v)$ between two vertices u, v of G is the set $I(u, v) = \{x \in V : d(u, v) = d(u, x) + d(x, v)\}$. A subset $S \subseteq V$ is called *gated* [29] if for each $v \notin S$ there exists a unique vertex $v' \in S$ (the gate of v in S) such that $v' \in I(v, u)$ for every $u \in S$. For every edge uv of G , define $W(u, v) = \{x \in V : d(u, x) < d(v, x)\}$. Given two connected graphs $G = (V(G), E(G))$ and $H = (V(H), E(H))$, we say that G admits an *isometric embedding* into H if there exists a *mapping* $\alpha : V(G) \rightarrow V(H)$ such that $d_H(\alpha(x), \alpha(y)) = d_G(x, y)$ for all vertices $x, y \in V(G)$.

The *Cartesian product* $H = H_1 \times \dots \times H_m$ of connected graphs H_1, \dots, H_m is defined upon the Cartesian product of the vertex sets of the corresponding graphs (called *factors*),

i.e., $V(H) = \{u = (u_1, \dots, u_m) : u_i \in V(H_i), i = 1, \dots, m\}$. Two vertices $u = (u_1, \dots, u_m)$ and $v = (v_1, \dots, v_m)$ are adjacent in H if and only if the vectors u and v coincide except at one position i , in which we have two vertices u_i and v_i adjacent in H_i . The *distance* $d_H(x, y)$ between two vertices $x = (x_1, \dots, x_m)$ and $y = (y_1, \dots, y_m)$ of H is $\sum_{i=1}^m d_{H_i}(x_i, y_i)$.

2.2 Medians

In the following, we consider graphs with weighted vertices. A *weight function* is any mapping π from the vertex set to the positive real numbers. The total weighted distance of a vertex x in G is given by $M_\pi(x) = \sum_u \pi(u)d(u, x)$. A vertex x minimizing this expression is a *median* (vertex) of G with respect to π , and the set of all medians is the *median set* $\text{Med}_\pi(G)$. For a subset of vertices $S \subseteq V$, denote by $\pi(S) = \sum_{s \in S} \pi(s)$ the weight of S . We continue with the following property of median functions:

Lemma 1. *For each edge uv in a graph G , $M_\pi(u) - M_\pi(v) = \pi(W(v, u)) - \pi(W(u, v))$.*

Proof. Indeed, since u and v are adjacent, we have $d(u, x) - d(v, x) = d(v, y) - d(u, y) = 1$ for any vertices $x \in W(v, u)$ and $y \in W(u, v)$. Since $d(x, z) = d(y, z)$ for any $z \notin W(u, v) \cup W(v, u)$, we conclude that $M_\pi(u) - M_\pi(v) = \sum_{x \in W(u, v)} \pi(x)(d(u, x) - d(v, x)) + \sum_{x \in W(v, u)} \pi(x)(d(u, x) - d(v, x)) = \pi(W(v, u)) - \pi(W(u, v))$. \square

Goldman and Witzgall [29] established that if the weight of a gated set S of G is larger than one half of the total weight, then $\text{Med}_\pi(G) \subseteq S$. In trees, in partial rectangular grids, in squaregraphs, and, more generally, in all median graphs, for each edge uv , the sets $W(u, v)$ and $W(v, u)$ are gated. Recall that G is a *median graph* if for each triplet u, v, w the intersection $I(u, v) \cap I(v, w) \cap I(w, u)$ consists of a single vertex. We can then infer that these graphs satisfy the following *majority rule* (which is a folklore for trees):

Lemma 2. *[7, 37] If G is a median graph, then $u \in \text{Med}_\pi(G)$ iff $\pi(W(u, v)) \geq \pi(W(v, u))$ for each neighbor v of u . If T is a tree, then $\pi(W(u, v)) = \pi(W(v, u))$ if and only if $\text{Med}_\pi(T) = \{u, v\}$.*

It is well-known since C. Jordan (1869) that the median set of a tree consists of one or two adjacent vertices. For median graphs, the following generalization holds:

Lemma 3. *[7, 37] If G is a median graph, then $\text{Med}_\pi(G)$ induces a hypercube. In particular, if G is a squaregraph or a partial grid, then $\text{Med}_\pi(G)$ is a vertex, an edge, or a square.*

2.3 Partial grids and squaregraphs

A *rectangular system* or a *partial rectangular grid* is the subgraph of the regular rectangular grid which is induced by the vertices of the grid lying either on a simple circuit of the grid (with possibly some vertices visited more than once) or inside the region bounded by this circuit. Every partial rectangular grid is a connected plane graph with inner faces of length four and inner vertices of degree four (the converse in general is not true). More generally, a *squaregraph* [17] is a plane graph with inner faces of length four and inner vertices of degree at least four. An *even squaregraph* is a squaregraph in which all inner vertices have even degrees (see Fig. 1). Squaregraphs constitute a particular subclass of median graphs. Median graphs arise in several areas of discrete mathematics, geometry, and theoretical computer science. One of the basic properties of median graphs is that for any edge uv the sets $W(u, v)$ and $W(v, u)$ constitute a partition of G into two gated subgraphs.




Fig. 1. A rectangular grid (a) and an even squaregraph (b)

2.4 Isometric embedding into products of two trees

Now, we will describe the isometric embedding of partial rectangular grids and even squaregraphs $G = (V, E)$ into the Cartesian product of two trees. For this we will use the notations in [8, 17]. For a squaregraph or a partial grid G denote by ∂G the bounding cycle of G .

First, let $G = (V, E)$ be a partial rectangular grid bounded by the cycle ∂G . Denote by E_1 the set of vertical edges of G and consider the graph $G_1 = (V, E_1)$. It is clear that the connected components of G_1 are paths of G with end-vertices on ∂G . Define the graph $T_1 = (V(T_1), E(T_1))$ whose vertices are the connected components of G_1 and two components P' and P'' are adjacent if and only if there exists an edge of G with one end in P' and another one in P'' . In the same way we can define the set E_2 of horizontal edges, the graph G_2 , and the tree $T_2 = (V(T_2), E(T_2))$.

We obtain the following canonical embedding α of G into the Cartesian product $T_1 \times T_2$. For any vertex v of G , we set $\alpha_1(v)$ (resp. $\alpha_2(v)$) to be the connected component of v in G_1 (resp. G_2). The embedding is defined by $\alpha(v) = (\alpha_1(v), \alpha_2(v))$. It can be verified that α provides an isometric embedding of G into $T_1 \times T_2$. For all vertices x, y of G , we have $d_G(x, y) = d_{T_1}(\alpha_1(x), \alpha_1(y)) + d_{T_2}(\alpha_2(x), \alpha_2(y))$. From now on, we will identify a vertex of G to the couple of vertical and horizontal paths to which it belongs. We call such path P a *fiber*, as P is equal to the subgraph induced by $\alpha_1^{-1}(P)$.

This canonical embedding of partial grids can be generalized to all graphs isometrically embeddable into Cartesian products of two trees. It was established in [8] that a graph G can be embedded into the Cartesian product of two trees if and only if G is a cube-free median graph without odd bipartite wheels. In particular, from this characterization follows that even squaregraphs admit such embeddings. To derive the embedding, the edges of an even squaregraph $G = (V, E)$ are divided into two sets E_1 and E_2 subject to the constraint that two incident edges e_1 and e_2 of a common inner face of G belong to different edge-sets. Equivalently, if we define the *side-graph* of G as the graph having the edges of G as the vertex-set and two edges e_1, e_2 of G are adjacent in the side-graph if and only if e_1 and e_2 are incident sides of some inner face of G , then the side-graph is bipartite. Note that the bipartition $\{E_1, E_2\}$ of E satisfies the following two conditions: (i) all “parallel” edges of G , i.e., edges which belong to the same equivalence class of the transitive closure of the binary relation “to be opposite edges of a common inner face of G ” all belong to the same color-class, and (ii) if we consider all edges incident to an inner vertex v of G , and number them counterclockwise, starting with an arbitrary edge, then the edges having numbers of the same parity all belong to the same color-class E_1 or E_2 . Note also that the set of all edges parallel to a given edge e (i.e. all edges that belong to the equivalence class of e) constitutes a cut-set of the graph G .

Analogously to the case of partial grids, the connected components of the graphs $G_1 = (V, E_1)$ (resp. $G_2 = (V, E_2)$) are called the fibers of G_1 (resp. G_2). They are (gated) trees. Define the graphs $T_i = (V(T_i), E(T_i)), i = 1, 2$, whose vertices are the fibers of G_i and two fibers F' and F'' are adjacent if and only if there exists an edge of G with one end in F' and another one in F'' . Then T_1 and T_2 are trees. We obtain an isometric embedding α of G into the Cartesian product of the two trees $T_1 \times T_2$, so that for any vertex v of G ,

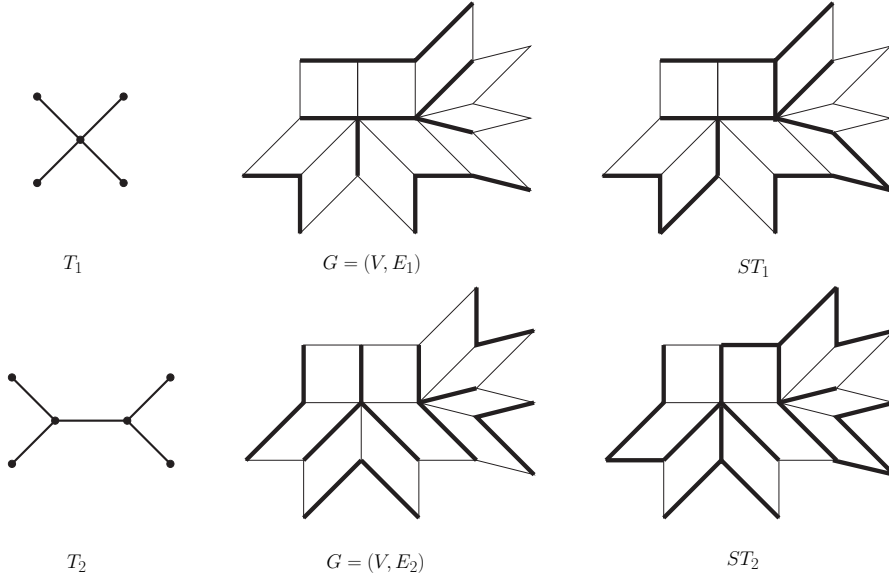


Fig. 2. An even squaregraph and its trees T_i and ST_i

$\alpha(v) = (P, Q)$, where P and Q are the fibers of the graphs G_1 and G_2 , sharing the vertex v [8].

This isometric embedding into the Cartesian product of two trees is used to establish one of the two properties on which our algorithms are based. Each vertex F_i of T_i is given the weight $\pi_i(F_i)$ equal to the total weight of vertices of G located in F_i .

Proposition 1. *Let G be an even squaregraph or a partial grid. A vertex $u = (\alpha_1(u), \alpha_2(u))$ is a median vertex of G if and only if $F_1 = \alpha_1(u)$ and $F_2 = \alpha_2(u)$ are median vertices of the trees T_1 and T_2 endowed with the weight functions π_1 and π_2 respectively.*

Proof. Denote by F_i a fiber of the graph G_i which is a median vertex of the tree T_i , $i = 1, 2$. We assert that $F_1 \cap F_2 \neq \emptyset$. Suppose by way of contradiction that F_1 and F_2 are disjoint. From the definition of F_1 and F_2 we immediately conclude that F_2 is completely contained in the same connected component of the graph $G \setminus F_1$ obtained from G by removing all vertices of F_1 . Thus all vertices x of F_2 have their image $\alpha_1(x)$ in the same connected component of $T_1 \setminus F_1$. Denote by F'_1 the neighbor of F_1 in the subtree of $T_1 \setminus F_1$ which contains F_2 . From Lemma 2 and the fact that F_1 is median in T_1 , we deduce that $\pi_1(W_{T_1}(F_1, F'_1)) \geq \frac{1}{2}\pi_1(V(T_1)) = \frac{1}{2}\pi(V)$. In the same way, defining F'_2 to be the neighbor of F_2 in the connected component of $T_2 \setminus F_2$ which contains F_1 , we obtain $\pi_2(W_{T_2}(F_2, F'_2)) \geq \frac{1}{2}\pi(V)$. The sets of vertices of G whose images are respectively in $W_{T_1}(F_1, F'_1)$ and $W_{T_2}(F_2, F'_2)$ are disjoint and do not entirely cover the graph G . Since $\pi(x) > 0$ for any vertex x of G , we obtain a contradiction with Lemma 2. Thus $F_1 \cap F_2 \neq \emptyset$, whence there exists a vertex m of G such that $\alpha_1(m) = F_1$ and $\alpha_2(m) = F_2$ are medians in T_1 and T_2 respectively. Thanks to the isometric embedding, we obtain

$$\begin{aligned} M_\pi(m) &= \sum_{x \in V} \pi(x) d_G(m, x) = \sum_{x \in V} \pi(x) d_{T_1}(F_1, \alpha_1(x)) + \sum_{x \in V} \pi(x) d_{T_2}(F_2, \alpha_2(x)) \\ &= \sum_{R \in V(T_1)} \pi_1(R) d_{T_1}(F_1, R) + \sum_{Q \in V(T_2)} \pi_2(Q) d_{T_2}(F_2, Q). \end{aligned}$$

Writing up a similar expression for any other vertex v of G and using the fact that F_1 and F_2 are medians of T_1 and T_2 , respectively, we conclude that $M_\pi(m) \leq M_\pi(v)$, thus m is a median vertex of G . Conversely, the previous equality also shows that any median vertex of $\text{Med}_\pi(G)$ can be expressed as the intersection of two median paths, one of T_1 and another of T_2 . \square

Before proving the second property of medians of partial grids and even squaregraphs, we define two particular spanning trees ST_1 and ST_2 of an even squaregraph G . ST_1 contains all edges of E_1 plus exactly one edge running between each pair of incident fibers of the graph $G_1 = (V, E_1)$. The choice of this edge is arbitrary (we can also select an edge belonging to the bounding cycle of G). We call such extra-edges the *switch edges* of ST_1 and denote them by E'_1 . Clearly, since all fibers of G_1 are trees, the graph $ST_1 = (V, E_1 \cup E'_1)$ is indeed a spanning tree of G . Analogously, we define the spanning tree $ST_2 = (V, E_2 \cup E'_2)$.

Proposition 2. *A vertex F_i of T_i ($i = 1, 2$) is a median vertex with respect to the weight function π_i if and only if F_i contains a median vertex m of the tree ST_i with respect to the weight function π .*

Proof. First suppose that m is a median vertex of ST_i (i.e., $m \in \text{Med}_\pi(ST_i)$) and let F_i be the vertex of T_i such that $\alpha_i(m) = F_i$. In other words, F_i is the fiber of G_i containing m . Suppose by way of contradiction that F_i is not a median of T_i (i.e. $F_i \notin \text{Med}_{\pi_i}(T_i)$). Lemma 2 yields that $M_{\pi_i}(F') < M_{\pi_i}(F_i)$ for some vertex F' of T_i adjacent to F_i . By Lemma 1 and the construction of T_i we conclude that $M_{\pi_i}(F_i) - M_{\pi_i}(F') = \pi(W(x', x)) - \pi(W(x, x')) > 0$, where $x'x$ is any edge running between F' and F_i . If m has a neighbor m' in F' and mm' is a switch edge, then it can be easily seen from the definition of the tree ST_i that all vertices of $W(x', x)$ (this set is defined in G) are closer in ST_i to x' than to x . This implies $M_\pi(m) - M_\pi(m') \geq \pi(W(x', x)) - \pi(W(x, x')) > 0$, contrary to the assumption that m is a median of ST_i . On the other hand, if the switch between F_i and F' is the edge $p'p$ with $p' \in F'$ and $p \in F_i$, and we denote by m' the neighbor of m on the unique path connecting m with p in the tree-fiber F_i , then again, in the tree ST_i all vertices of $W(p', p)$ are closer to m' than to m . Since $\pi(W(p', p)) > \frac{1}{2}\pi(V)$, Lemma 2 yields that m is not a median vertex of ST_i , a contradiction.

Conversely, suppose that F_i is a median vertex of the tree T_i . We assert that $F_i \cap \text{Med}_\pi(ST_i) \neq \emptyset$. Remove from T_i all edges incident to F_i and denote by S_1, \dots, S_k the resulting subtrees of T_i not containing F_i . Then Lemmas 1 and 2 imply that $\pi_i(S_j) \leq \frac{1}{2}\pi(V)$ for any subtree S_j . Now, if we pick the switch edge x_jm_j running between S_j and F_i with $x_j \in S_j$ and $m_j \in F_i$, then from the definition of the spanning tree ST_i we infer that S_j coincides with the set of all vertices which are closer in ST_i to x_j than to m_j . The majority rule for trees implies that $M_\pi(m_j) \leq M_\pi(x_j)$ holds in ST_i . Now, the median function M_π on trees is convex [38]. Since $M_\pi(m_j) \leq M_\pi(x_j)$, this implies that $M_\pi(x_j) \leq M_\pi(y_j)$ for any vertex $y_j \in S_j \setminus \{x_j\}$. Since any vertex z outside F_i is located in some subtree S_j , we conclude that $M_\pi(m_j) \leq M_\pi(x_j) \leq M_\pi(z)$ holds in ST_i . This shows that indeed F_i must contain at least one median vertex of the tree ST_i . \square

We obtain the following corollary as a direct consequence of the two previous properties:

Corollary 1. *$m \in \text{Med}_\pi(G)$ if and only if $\alpha_i(m) \in \text{Med}_\pi(ST_i)$, for $i = 1, 2$.*

3 Algorithms for the median problem

In the introduction, we outlined the self-stabilizing algorithm for the median problem in a tree proposed in [3, 13]. We continue, with a more detailed account of the model used by this and our algorithms. Then we present the algorithm of [3, 13] and our algorithms for partial grids and even squaregraphs.

3.1 Computational model

The nodes of the graph $G = (V, E)$ are seen as processors executing the same algorithm. Each processor $v \in V$ has a memory whose value (its *state*) can be read by its neighbors, but can be changed only by v itself. A distributed algorithm is a set of rules (a pair of

precondition and *command*) that describe how a processor has to change its current state (the command) according to the state of all its neighbors (the precondition or guard). We say that a rule R is *activable* at a processor v if the neighborhood of v satisfies the precondition of R . In this case, the node v is also said to be *activable*. If a rule R is activable in v , an *atomic move* for v consists in reading the states of all its neighbors, computing a new value of its state according to the command of R , and writing this value to the local memory. An *execution* is a sequence of atomic moves. This is indeed an interleaved (central daemon) asynchronous model of computation. (Implicit) termination or stabilization is reached when there is no more activable rules. The described model is a standard model for distributed computing. It is exactly the same model as in [13], where the algorithms are expressed in the language of “guarded commands”. It is extensively studied in [14]. In his thorough study of computational models, it is coined by Chalopin as the “interleaved cellular model” [15, chapter 5]. In [1], it appears as the “local detection paradigm”.

A distributed algorithm is said to be *self-stabilizing* if an execution starting from any arbitrary global state has a suffix belonging to the set of legitimate states. For our purposes, we additionally suppose that the state variable of each node has a specific bit named the *median flag*. Then a global state is *legitimate*, if the median flag of a node v is set up if and only if v is a median vertex of G . The *time complexity* of a self-stabilizing algorithm is the maximum number of moves that are performed until stabilization. A *round* [22] is a sequence of moves such that each activable node at the beginning of the round is activated at least once. The *round complexity* of a self-stabilizing algorithm is the maximum number of rounds required by an execution to reach a legitimate state. Whenever we compose two self-stabilizing algorithms, then we assume that this is done as a “fair composition”; this guarantees the self-stabilization of the resulting composite algorithm; see Subsection 2.7 of [22]. In all three algorithms described below, only $O(\log n)$ bits are used to store the state of each node.

We specify now the structural information that is used by each of three median computation algorithms. In the algorithm for trees, neither nodes nor edges have identifiers. In this case, the system is said to be *anonymous*. On the other hand, in the algorithm for partial rectangular grids, nodes are anonymous but edges are not: for each node, there exists a labeling of the outgoing edges that has the property of (weak) “sense of direction”, allowing to compute the second neighborhood of each node in a self-stabilizing manner. Informally, a system represented by a directed graph is said to have *sense of direction* if it is possible to know, from the labels associated to the edges, whether different walks from any given node v end up in the same node or not. The use of sense of direction in a distributed system often leads to significant improvements on computability and complexity [24]. Finally, in the algorithm for even squaregraphs, each processor has a unique identifier (in fact, it is unclear for us whether this is a necessary structural information).

3.2 Trees

Let $T = (V, E)$ be a tree with n vertices and let π be a weight function on V . We need the following notations:

- v .*s*-value is the local value of the vertex v . It is also called the *s*-value of v ;
- $\gamma_1(v) = \{u \in V : uv \in E\}$ is the set of neighbors of the vertex v in T ;
- $N_s(v) = \{u$.*s*-value : $u \in \gamma_1(v)\}$ is a multiset;
- $N_s^-(v) = N_s(v) \setminus \{\max(N_s(v))\}$.

The main result of [3, 13] is the proof of self-stabilization in polynomial time for the following algorithm (which we slightly modify to capture weighted medians as well). Algorithms are described by a list of rules, and to simplify the formulation of preconditions, we assume in all the following that a rule is not activable if its execution would not change the state of the node.

MEDIAN TREE	
$(v \text{ is a leaf})$	$\longrightarrow v.s\text{-value} = \pi(v)$
$(v \text{ is not a leaf})$	$\longrightarrow v.s\text{-value} = \pi(v) + \sum(N_s^-(v)).$

In a stabilized state, the median vertices are those vertices whose s -value is greater than the s -values of all their neighbors. It is shown in [13] that MEDIAN TREE stabilizes in $O(e)$ rounds, where e is the maximum distance of a median to a leaf. Moreover, the algorithm makes $O(n^3 \cdot c_s)$ moves in the worst case, where c_s is the maximum initial s -value of any processor [13].

3.3 Partial rectangular grids

Let $G = (V, E)$ be a partial grid with n vertices bounded by the cycle ∂C . A vertex of G having Cartesian coordinates (x, y) is denoted by $v_{x,y}$. The first neighborhood $\gamma_1(v)$ of a vertex $v = v_{x,y}$ of G is the set $V \cap \{v_{x,y+1}, v_{x,y-1}, v_{x+1,y}, v_{x-1,y}\}$, where $v_{x,y+1}$, $v_{x,y-1}$, $v_{x+1,y}$ and $v_{x-1,y}$ are the vertices of the square grid located at the North, the South, the East, and the West of v , respectively. The second neighborhood $\gamma_2(v)$ of v in G is the set $V \cap \{v_{x+1,y+1}, v_{x-1,y+1}, v_{x+1,y-1}, v_{x-1,y-1}\}$, where $v_{x+1,y+1}$, $v_{x-1,y+1}$, $v_{x+1,y-1}$ and $v_{x-1,y-1}$ are the vertices of the square grid which are located at the North East, the North West, the South East and the South West of v respectively (see Figure 3).




Fig. 3. The first and second neighborhood of a vertex.

The model used by the algorithm MEDIANPARTIALGRID is similar to the one of MEDIAN TREE except that each processor v has four variables whose values can be read by its four possible neighbors, but can be changed only by v . A processor v should also know the existence of the edges between its first neighborhood $\gamma_1(v)$ and its second neighborhood $\gamma_2(v)$. This structural information will be provided by communicating with the first neighborhood under the assumption that the system has “sense of direction”. Each processor v has also two boolean variables $v.b_1$ and $v.b_2$. The variable $v.b_i$ will be set if the path $\alpha_i(v)$ is a median vertex of T_i . The algorithm MEDIANPARTIALGRID consists of three phases.

Phase 1. In this phase, with the knowledge of $\gamma_2(v)$ (it is obvious that the sense of direction enables, for each node v , to compute $\gamma_1(v) \cup \gamma_2(v)$ in a self-stabilizing manner), each processor $v = v_{x,y}$ computes the sets $\gamma_{ST_1}(v)$ and $\gamma_{ST_2}(v)$ of its neighbors in the spanning trees ST_1 and ST_2 , respectively. For example, for the tree ST_1 , this can be done by communicating with its first neighborhood $\gamma_1(v)$ in G according to the following rules:

$$\begin{aligned}
v_{x,y-1} \in \gamma_{ST_1}(v) &\text{ iff } v_{x,y-1} \in \gamma_1(v), \\
v_{x,y+1} \in \gamma_{ST_1}(v) &\text{ iff } v_{x,y+1} \in \gamma_1(v), \\
v_{x-1,y} \in \gamma_{ST_1}(v) &\text{ iff } v_{x-1,y} \in \gamma_1(v) \wedge [v_{x-1,y} \notin \gamma_1(v_{x-1,y-1}) \vee v_{x-1,y-1} \notin \gamma_1(v_{x,y-1}) \vee v_{x,y-1} \notin \gamma_1(v)], \\
v_{x+1,y} \in \gamma_{ST_1}(v) &\text{ iff } v_{x+1,y} \in \gamma_1(v) \wedge [v_{x+1,y} \notin \gamma_1(v_{x+1,y-1}) \vee v_{x+1,y-1} \notin \gamma_1(v_{x,y-1}) \vee v_{x,y-1} \notin \gamma_1(v)].
\end{aligned}$$

Phase 2. In this phase, each processor v runs the algorithm MEDIAN TREE in parallel on each of the spanning trees ST_1 and ST_2 . The variable $v.s_i\text{-value}$ is the s -value of v for the tree ST_i . Then a median of ST_i can be identified by the fact that the s_i -value of respective processor is maximum in its neighborhood on ST_i , $i = 1, 2$.

Phase 3. In this phase, a classical broadcasting self-stabilizing algorithm [22, chap.4, p97] is replicated in both directions on every vertical (respectively, horizontal) path to set the boolean variables $v.b_1$ and $v.b_2$. Once the “vertical” and “horizontal” broadcasting algorithms stabilize, the median set of G is formed by all vertices v of the partial grid G for which $v.b_1 \wedge v.b_2$ is true.

Theorem 1. *The algorithm MEDIANPARTIALGRID computes the median set of a partial grid G with n vertices in $O(n)$ rounds and $O(c_s n^3)$ moves.*

Proof. First we show that the algorithm stabilizes. Indeed, by simulating the execution of MEDIANTREE on the two spanning trees ST_1 and ST_2 , we obviously maintain the self-stabilization because the read and written variables are distinct (we use s_1 for the tree ST_1 and s_2 for ST_2). By composing these algorithms with the broadcasting algorithm, the self-stabilization is still maintained according to Theorem 2.2 of [22]. Indeed, each algorithm MEDIANTREE runs independently of the broadcasting algorithm and thus can reach a legitimate global state. The broadcasting algorithm can then start to stabilize. This shows that our algorithm stabilizes.

As to the time complexity, notice that the algorithm makes $O(n^3 \cdot c_s)$ moves in the worst case, where c_s is the maximum initial s -value of any processor. Since the time complexity of MEDIANTREE is superior to the time complexity of classical self-stabilizing broadcasting algorithm, the time complexity of MEDIANPARTIALGRID is of the same order than the time complexity of MEDIANTREE which is given in [13], concluding the proof. In the same way, the round complexity of MEDIANTREE is $O(e)$ where $e \leq n$ is the maximum distance from a median to a leaf in the spanning tree yielding a $O(n)$ round complexity for MEDIANPARTIALGRID.

Finally, we show the partial correctness of our algorithm, i.e. that in a global stabilized state the processors v that have their two boolean variables $v.b_1$ and $v.b_2$ set to true are medians of the partial grid G . Indeed, by Corollary 1 a vertex (processor) v is median in G if and only if it is on the vertical path of a median of ST_1 (thus variable $v.b_1$ set to true) and on the horizontal path of a median of ST_2 (variable $v.b_2$ set to true). Since the algorithm MEDIANTREE correctly computes the median set of each tree ST_1, ST_2 , we are done. \square

3.4 Even squaregraphs

Let $G = (V, E)$ be an even squaregraph. The “irregular” structure of G does not allow an easy use of the sense of direction as in the case of partial squaregrids. To obtain a bipartition of edges used in the construction of the trees T_1, T_2 and ST_1, ST_2 , we will use the self-stabilizing algorithm for constructing a spanning BFS-tree of a graph designed by Afek, Kutten, and Yung [1] (for a survey on other related algorithms for this problem, see [25]). This algorithm requires that each vertex v has a unique identifier $v.Id$. This extra-information allows to break the symmetry in order to select as the root of the spanning tree the vertex having the highest identifier.

Our median self-stabilizing algorithm MEDIANEVEN SQUAREGRAPH consists of four phases. In each phase, we present the specific conditions which allow to test if the state of a vertex is legal or not for the current phase. If the respective condition is not satisfied, then we describe the modifications which must be undertaken in order to return the system in a legal state. We establish that after a finite number of activations, the corresponding conditions of the current phase are satisfied by all vertices, thus the next phase can start.

Phase 1. In this phase, we construct a spanning tree using the algorithm of Afek and al [1]. When this phase terminates, each vertex v has computed the identifier $v.Root$ of the root node of the resulting spanning BFS-tree, the identifier $v.Parent$ of its father in this tree and an integer $v.Distance$ which is the tree-distance between v and the root. This phase ends if the following condition holds in each node v :

Condition $st(v)$:

$$\{[(v.\text{Root} = v.\text{Id}) \wedge (v.\text{Parent} = v.\text{Id}) \wedge (v.\text{Distance} = 0)] \vee [(v.\text{Root} > v.\text{Id}) \wedge (v.\text{Parent} \in v.\text{Edge-list}) \wedge (v.\text{Root} = v.\text{Parent}.\text{Root}) \wedge (v.\text{Distance} = v.\text{Parent}.\text{Distance} + 1)]\} \wedge (v.\text{Root} \geq \max_{x \in v.\text{Edge-list}} x.\text{Root})$$

The algorithm, which is executed by each processor so that the system stabilizes in a state in which all these conditions are satisfied for each node, is described in details and analyzed in [1].

Phase 2. In this second phase of the algorithm, we aim to partition the edges of G into two subsets E_1 and E_2 so that two incident edges belonging to a common square-face are included in different sets E_1 and E_2 . For a vertex v , we assume that the edges containing v as an end-vertex are numbered $0, \dots, \text{deg}(v) - 1$ in the order in which they appear in the counterclockwise traversal so that two incident to v edges which belong to a common square have numbers of different parity. For each $v \in V$, the variable $v.\text{Color}$ equals i if all edges incident to v which appear at even positions in the adjacency list of v belong to E_i and the remaining edges belong to E_{3-i} . Each vertex v runs the following algorithm:

$st(v) \wedge (v.\text{Id} = v.\text{Root})$	$\longrightarrow v.\text{Color} := 1$
$st(v) \wedge (v.\text{Id} \neq v.\text{Root}) \wedge$ (the edge connecting the vertex v with $v.\text{Parent}$ occurs with the same parity in the adjacency list of v as in the adjacency list of $v.\text{Parent}$)	$\longrightarrow v.\text{Color} := v.\text{Parent}.\text{Color}$
$st(v) \wedge (v.\text{Id} \neq v.\text{Root}) \wedge$ (the edge connecting the vertex v to $v.\text{Parent}$ occurs with different parities in the adjacency lists of v and $v.\text{Parent}$)	$\longrightarrow v.\text{Color} := 3 - v.\text{Parent}.\text{Color}$

We say that the condition $col(v)$ is satisfied by a vertex $v \in V$, if none of the preconditions of the three previous actions is satisfied. Then Phase 2 terminates if the condition $col(v)$ is satisfied by all vertices $v \in V$ of G . Then each vertex v knows the list $v.\text{Edge-list}_i$ of its neighbors in G_i .

Phase 3. In this phase, the algorithm constructs the spanning trees ST_i , $i = 1, 2$. For sake of simplicity, these trees will be rooted at the same vertex as the root of the spanning BFS-tree computed in Phase 1. To encode the tree ST_i , we introduce for each $v \in V$ a new variable $v.\text{Parent}_i$ which will denote the father of the vertex v in the tree ST_i . The algorithm consists in “correcting” the spanning tree of Phase 1 by replacing $v.\text{Parent}$ by a vertex belonging to $v.\text{Edge-list}_i$, if this list contains a vertex located at the same distance to the root as the vertex $v.\text{Parent}$. Since all fibers of the graphs G_i are gated sets (in G), the tree obtained in this way has all edges of E_i and exactly one switch edge of E_{3-i} running between each pair of neighboring fibers of G_i . Actually, if $v'v$ is a switch edge with $d_G(r, v') < d_G(r, v)$ and v belongs to the fiber F , then necessarily v is the gate of r in the fiber F and v' is the father of v in the BFS-tree. Indeed, the root r can be connected with any vertex u of F of G_i by a shortest path passing via the gate v of r in F , thus the edge vv' will be included in the BFS-tree before the edge running from u to its father. Each processor $v \in V$ runs at this phase the following algorithm:

$col(v) \wedge (v.\text{Id} = v.\text{Root})$	$\longrightarrow v.\text{Parent}_i := v.\text{Id}$
$col(v) \wedge (v.\text{Id} \neq v.\text{Root}) \wedge$ (the edge connecting v to $v.\text{Parent}$ belongs to E_i)	$\longrightarrow v.\text{Parent}_i := v.\text{Parent}$
$col(v) \wedge (v.\text{Id} \neq v.\text{Root}) \wedge$ (the edge connecting the vertex v to $v.\text{Parent}$ belongs to E_{3-i}) $\wedge (v.\text{Parent}.\text{Distance} < \min_{x \in v.\text{Edge-list}_i} x.\text{Distance})$	$\longrightarrow v.\text{Parent}_i := v.\text{Parent}$
$col(v) \wedge (v.\text{Id} \neq v.\text{Root}) \wedge$ (the edge connecting the vertex v with $v.\text{Parent}$ belongs to E_{3-i}) $\wedge (v.\text{Parent}.\text{Distance} \geq \min_{x \in v.\text{Edge-list}_i} x.\text{Distance})$	$\longrightarrow v.\text{Parent}_i := \text{argmin}_{x \in v.\text{Edge-list}_i} x.\text{Distance}$

We say that the condition $st_i(v)$ is satisfied by the vertex v if none of the preconditions of the previous actions is satisfied. Since the spanning trees ST_1 and ST_2 are constructed at the same time, Phase 3 terminates if the condition $st_i(v)$ is satisfied at each vertex $v \in V$ and for each index $i \in \{1, 2\}$. At the end of this phase, each node $v \in V$ knows its father $v.\text{Parent}_i$ in the tree ST_i .

Phase 4. With trees ST_1 and ST_2 at hand, the algorithm is similar to that for partial grids. First, the algorithm MEDIAN TREE is used to compute the medians of the trees ST_1 and ST_2 . Then, using a self-stabilizing broadcasting algorithm, we set to “true” the variable $v.b_i$ of each vertex v belonging to the same connected component of the graph G_i as a median vertex of the tree ST_i . Once this broadcasting algorithm stabilizes, the median set $\text{Med}_\pi(G)$ of G is formed by all vertices v of G for which $v.b_1 \wedge v.b_2$ is true.

Theorem 2. *The algorithm MEDIANEVEN SQUAREGRAPH computes the median set of an even squaregraph G with n vertices in $O(n^2)$ rounds.*

Proof. The algorithm given by Afek et al. [1] for constructing a spanning tree stabilizes in $O(n^2)$ rounds. This shows that Phase 1 stabilizes in $O(n^2)$ rounds. The Phase 2 needs $O(n)$ rounds in the worst case. By induction we can show that when a vertex v located at distance i from the root is activated during the round $i + 1$, the variable $w.\text{Color}$ of its father w , which is at distance $i - 1$ from the root, has been already correctly computed (by induction hypothesis). Thus the rules of Phase 2 correctly compute the value of $v.\text{Color}$ using that of $w.\text{Color}$. Since two edges incident to the same vertex v and belonging to the same square have numbers of different parity in the degree list of v , they will be included in different sets E_1 and E_2 by the algorithm. It remains to show that any edge uv is inserted in the same edge-list by both vertices u and v . For this we proceed by induction on $k := d_G(v, r) < d_G(u, r)$. Our previous argument shows that the assertion holds when v is the father of u in the BFS-tree. So, suppose that v' is the father of v and u' is the father of u in the BFS-tree and that $u' \neq v$. It was shown in [16] for graphs more general than squaregraphs, that, if u and v are adjacent in G , then their fathers u' and v' are also adjacent. Since $d_G(v', r) = k - 1$, by induction hypothesis we conclude that u' and v' inserted the edge $u'v'$ in the same set, say E_1 . Now, since each vertex inserts two incident edges of a square in different edge-lists and the edge connecting a vertex with its father is set in the same list by the two ends, we conclude that the edges $u'u$ and $v'v$ are in the lists E_2 of their extremities. This implies that the edge uv will be put in the list E_1 by both vertices u and v , thus establishing our assertion.

The rules of Phase 3 depend only of the information computed at Phases 1 and 2, thus in order to correctly compute in Phase 3 the trees ST_1 and ST_2 , it suffices that each vertex is activated at this phase once. As to the Phase 4, the algorithm MEDIAN TREE stabilizes in $O(e)$ rounds, where e is the largest eccentricity of a median vertex of G [13], thus its complexity is the same as that of broadcasting. Summarizing, we conclude that the construction in $O(n^2)$ rounds of a spanning tree of G dominates the overall complexity of our algorithm. Most importantly, Proposition 2 establishes that MEDIANEVEN SQUAREGRAPH correctly computes $\text{Med}(G)$. \square

Remark 1. Note that, given a spanning tree of an even squaregraph and the unique identifiers, it would be possible to reconstruct locally a map representing the topology of the graph and then compute *internally* the median set. But it would require much more than $O(\log n)$ bits per node and therefore is not a satisfactory solution.

References

1. Y. Afek, S. Kutten, and M. Yung. The Local Detection Paradigm and its Applications to Self-Stabilization. *Theoretical Computer Science*, 186(1-2):199–229, 1997.

2. S. Aggarwal and S. Kutten, Time optimal self-stabilizing spanning tree algorithm, In *FSTTCS93, Springer LNCS*, vol. 761, 1993, pp. 400–410.
3. G. Antonoiu and P. K. Srimani, A self-stabilizing distributed algorithm to find the median of a tree graph, *J. Comput. Syst. Sci.*, 58(1999), 215–221.
4. G. Antonoiu and P.K. Srimani, A self-stabilizing leader election algorithm for tree graphs, *J. Parallel Distrib. Comput.*, 34 (1996), 227–232.
5. G. Antonoiu and P.K. Srimani, Distributed self-stabilizing algorithm for minimum spanning tree construction, In *Euro-Par*, 1997, pp. 480–487.
6. B. Awerbuch, S. Kutten Y. Mansour B. Patt-Shamir, and G. Varghese, Time optimal self-stabilizing synchronization, In *STOC93 Proc. 25th Annu. ACM Symp. Theory of Comput.*, 1993, pp. 652–661.
7. H.-J. Bandelt and J.P. Barthélemy, Medians in median graphs, *Discrete Appl. Math.* 8 (1984), 131-142.
8. H.-J. Bandelt, V. Chepoi, and D. Eppstein, Ramified rectilinear polygons (in preparation).
9. J.P. Barthélemy and M.F. Janowitz, A formal theory of consensus, *SIAM J. Discrete Math.* 4 (1991), 305-322.
10. J.P. Barthélemy and B. Monjardet, The median procedure in cluster analysis and social choice theory, *Math. Social Sci.* 1 (1981), 235-268.
11. F. Buckley and F. Harary, *Distances in Graphs*, Redwood City, CA: Addison-Wesley, 1990.
12. J.R.S. Blair and F. Manne. Efficient self-stabilizing algorithms for tree networks. Technical Report 232, Departement of Informatics, University of Bergen, Norway, 2002.
13. S.C. Bruell, S. Ghosh M.H. Karaata, and S.V. Pemmaraju, Self-stabilizing algorithms for finding centers and medians of trees, *SIAM J. Computing*, 29(1999), 600–614.
14. P. Boldi and S. Vigna. An effective characterization of computability in anonymous networks. In Jennifer Welch, editor, *Proceedings of DISC 2001*, Lecture Notes in Computer Science 2180 (2001), 33-47.
15. J. Chalopin. Algorithmique Distribuée, Calculs Locaux et Homomorphismes de Graphes. PhD Thesis Université Bordeaux I (2006).
16. V. Chepoi, Graphs of some CAT(0) complexes, *Adv. Appl. Math.*, 24(2000), 125-179.
17. V. Chepoi, F. Dragan, and Y. Vaxès, Addressing, distances and routing in triangular systems with applications in cellular and sensor networks, *Wireless Networks*, 12 (2006) 671-679.
18. V. Chepoi, C. Fanciullini, and Y. Vaxès, Median problem in some plane triangulations and quadrangulations, *Comput. Geom.* 27(2004), 193-210.
19. A.K. Datta and J.L. Derby, J.E. Lawrence and S. Tixeuil. Stabilizing hierarchical routing, *J. Interconnexion Networks*, 1(2000), 283–302.
20. S. Delat and D. Nguyen and S. Tixeuil. Stabilité et auto-stabilisation du routage inter-domaine dans internet. In *RIVF*, pages 139–144, 2003.
21. E.W. Dijkstra, Self-stabilizing systems in spite of distributed control, *Comm. ACM*, 17 (1974),:643–644.
22. S. Dolev. *Self-stabilization*, MIT Press, 2000.
23. S. Dolev and A. Israeli and S. Moran. Self-stabilization of dynamic systems assuming only read/write atomicity. In *Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing*, pages 103–118, 1990.
24. P. Flocchini and B. Mans and N. Santoro. Sense of direction: formal definitions and properties. In *Proceedings of 1st international conference on structural information and communication complexity*, pages 9–34. Carleton university press, 1995.
25. F. Gärtner, A survey of self-stabilizing spanning-tree construction algorithms (2003).
26. O. Gerstel and S. Zaks, A new characterization of tree medians with applications to distributed algorithms, *Networks*, 24(1994), 23-29.
27. S. Ghosh and A. Gupta and S.V. Pemmaraju. A self-stabilizing algorithm for the maximum flow problem. *Distributed Computing*, 10(4):167–180, 1997.
28. S. Ghosh and M.H. Karaata. A self-stabilizing algorithm for coloring planar graphs. *Distributed Computing*, 7(1):55–59, 1993.
29. A.J. Goldman and C.J. Witzgall. A localization theorem for optimal facility placement. *Transportation Science*, 4:406–409, 1970.
30. S.T. Hedetniemi and D.P. Jacobs and P.K. Srimani. Maximal matching stabilizes in time $o(m)$. *Information Processing Letters*, 80(5):221–223, 2001.
31. G.R.T. Hendry, On graphs with prescribed medians, I, *J. Graph Theory* 9 (1985), 477–481.
32. T. Herman and S. Ghosh. Stabilizing phase-clocks. *Information Processing Letters*, 54:259–265, 1995.
33. T.C. Huang J. Lin and H. Chen. A self-stabilizing algorithm which finds a 2-center of a tree. *Computer and Mathematics with Applications*, 40:607–624, 2000.
34. O. Kariv, S.L. Hakimi, An algorithmic approach to network location problems, I,II, *SIAM J. Appl. Math.* 37 (1979) 513–538, 539–560.
35. E. Korach, D. Rotem, and N. Santoro, Distributed algorithms for finding centers and medians in networks, *ACM Transactions on Programming Languages and Systems*, 6(1984), 380-401.
36. M. Nesterenko and M. Mizuno. A quorum-based self-stabilizing distributed mutual exclusion algorithm. *Journal of Parallel and Distributed Computing*, 62(2):284–305, 2002.
37. P. Soltan and V. Chepoi, The solution of the Weber problem for discrete median metric spaces (in Russian), *Trudy Tbilisskogo Mat. Inst.* 85 (1987), 52–76.
38. B.C. Tansel R.L. Francis and T.J. Lowe. Location on networks: a survey. *Management Science*, 29:482–511, 1983.