



Strategy-proof Mechanisms for Facility Location Games with Many Facilities

Bruno Escoffier, Laurent Gourvès, Kim Thang Nguyen, Fanny Pascual,
Olivier Spanjaard

► To cite this version:

Bruno Escoffier, Laurent Gourvès, Kim Thang Nguyen, Fanny Pascual, Olivier Spanjaard. Strategy-proof Mechanisms for Facility Location Games with Many Facilities. 2nd International Conference on Algorithmic Decision Theory (ADT'11), Oct 2011, Piscataway, NJ, United States. pp.67-81, 10.1007/978-3-642-24873-3_6 . hal-01285708

HAL Id: hal-01285708

<https://hal.science/hal-01285708>

Submitted on 10 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Strategy-proof Mechanisms for Facility Location Games with Many Facilities^{*}

Bruno Escoffier¹, Laurent Gourvès¹, Thang Nguyen Kim¹, Fanny Pascual²,
and Olivier Spanjaard²

1. Université Paris-Dauphine, LAMSADE-CNRS, UMR 7243, F-75775 Paris, France

2. UPMC, LIP6-CNRS, UMR 7606, F-75005 Paris, France

{bruno.escoffier, laurent.gourves, kim-thang.nguyen}@lamsade.dauphine.fr,
{fanny.pascual, olivier.spanjaard}@lip6.fr

Abstract. This paper is devoted to the location of public facilities in a metric space. Selfish agents are located in this metric space, and their aim is to minimize their *own* cost, which is the distance from their location to the nearest facility. A central authority has to locate the facilities in the space, but she is ignorant of the true locations of the agents. The agents will therefore report their locations, but they may lie if they have an incentive to do it. We consider two social costs in this paper: the sum of the distances of the agents to their nearest facility, or the maximal distance of an agent to her nearest facility. We are interested in designing *strategy-proof* mechanisms that have a small approximation ratio for the considered social cost. A mechanism is strategy-proof if no agent has an incentive to report false information. In this paper, we design strategy-proof mechanisms to locate $n - 1$ facilities for n agents. We study this problem in the general metric and in the tree metric spaces. We provide lower and upper bounds on the approximation ratio of deterministic and randomized strategy-proof mechanisms.

Key words: Facility location games; Strategy-proof mechanisms; Approximation guarantee

1 Introduction

We study *Facility Location Games* that model the following problem in economics. Consider installation of public service facilities such as hospitals or libraries within the region of a city, represented by a metric space. The authority announces that some locations will be chosen within the region and runs a survey over the population; each inhabitant may declare the spot in the region that she prefers some facility to be opened at. Every inhabitant wishes to minimize her individual distance to the closest facility, possibly by misreporting her preference to the authorities. The goals of the authority are twofold: avoiding such

^{*} This work is supported by French National Agency (ANR), project COCA ANR-09-JCJC-0066-01.

misreports and minimizing some social objectives. The authority needs to design a *mechanism*, that maps the reported preferences of inhabitants to a set of locations where the facilities will be opened at, to fulfill the purposes. The mechanism must be *strategy-proof*, i.e., it ensures that no inhabitant can benefit by misreporting her preference. At the same time, the mechanism should guarantee a reasonable approximation to the optimal social cost. The model has many applications in telecommunication networks where locations may be easily manipulated by reporting false IP addresses, false routers, etc.

1.1 Facility Location Games

We consider a metric space (Ω, d) , where $d : \Omega \times \Omega \rightarrow \mathbb{R}$ is the metric function. Some usual metrics are the line, circle and tree metrics where the underlying spaces are an infinite line, a circle and an infinite tree, respectively. The distance between two positions in such metrics is the length of the shortest path connecting those positions. Let n be the number of agents, each agent i has a location $x_i \in \Omega$. A *location profile* (or *strategy profile*) is a vector $\mathbf{x} = (x_1, \dots, x_n) \in \Omega^n$. Let k be the number of facilities that will be opened. A *deterministic* mechanism is a mapping f from the set of location profiles Ω^n to k locations in Ω . Given a reported location profile \mathbf{x} the mechanism's output is $f(\mathbf{x}) \in \Omega^k$ and the individual cost of agent i under mechanism f and profile \mathbf{x} is the distance from its location to the closest facility, denoted by $c_i(f, \mathbf{x})$:

$$c_i(f, \mathbf{x}) := d(f(\mathbf{x}), x_i) := \min\{d(F, x_i) : F \in f(\mathbf{x})\}$$

A *randomized* mechanism is a function f from the set of location profiles to $\Delta(\Omega^k)$ where $\Delta(\Omega^k)$ is the set of probability distributions over Ω^k . The cost of agent i is now the expected distance from its location to the closest facility over such distribution:

$$c_i(f, \mathbf{x}) := \mathbb{E}[d(f(\mathbf{x}), x_i)] := \mathbb{E}[\min\{d(F, x_i) : F \in f(\mathbf{x})\}]$$

We are interested in two standard *social objectives*: (i) the *utilitarian objective* defined as the total individual costs (total individual expected cost for a randomized mechanism), i.e., $C(f, \mathbf{x}) = \sum_{i=1}^n c_i(f, \mathbf{x})$; and (ii) the *egalitarian objective* defined as the maximal individual cost (expected maximal individual cost for a randomized mechanism), i.e., $C(f, \mathbf{x}) = \mathbb{E}[\max_{1 \leq i \leq n} d(f(\mathbf{x}), x_i)]$. This is thus simply $\max_{1 \leq i \leq n} c_i(f, \mathbf{x})$ for deterministic mechanisms.

We say that a mechanism f is *r-approximate* with respect to profile \mathbf{x} if

$$C(f, \mathbf{x}) \leq r \cdot OPT(\mathbf{x})$$

where $OPT(\mathbf{x})$ is the social cost of an optimal facility placement (for the egalitarian or utilitarian social cost). Note that since for a randomized mechanism the social cost is the expectation of the social cost on each chosen set of locations, there always exists an optimal deterministic placement.

We will be concerned with *strategy-proof* (SP) mechanisms, which render truthful revelation of locations a dominant strategy for the agents.

Definition 1. (Strategyproofness) Let $\mathbf{x} = (x_1, \dots, x_n)$ denote the location profile of n agents over the metric space (Ω, d) . A mechanism f is strategy-proof (SP) if for every agent $1 \leq i \leq n$ and for every location $x'_i \in \Omega$, $c_i(f, (x'_i, x_{-i})) \geq c_i(f, \mathbf{x})$ where x_{-i} denotes the locations of the agents other than i in \mathbf{x} .

1.2 Previous Work

The facility locations game where only *one* facility will be opened is widely-studied in economics. On this topic, Moulin [6] characterized all strategy-proof mechanisms in the line metric space. Subsequently, Schummer and Vohra [10] gave a characterization of strategy-proof mechanisms for the circle metric space. More recently, Procaccia and Tennenholtz [9] initiated the study of approximating an optimum social cost under the constraint of strategy-proofness. They studied deterministic and randomized mechanisms on the line metric space with respect to the utilitarian and egalitarian objectives. Several (tight) approximation bounds for strategy-proof mechanisms were derived in their paper. For general metric space, Alon et al. [1] and Nguyen Kim [7] proved randomized tight bounds for egalitarian and utilitarian objectives, respectively.

Concerning the case where *two* facilities are opened, Procaccia and Tennenholtz [9] derived some strategy-proof mechanisms with guaranteed bounds in the line metric space for both objectives. Subsequently, Lu et al. [5] proved tight lower bounds of strategy-proof mechanisms in the line metric space with respect to the utilitarian objective. Moreover, they also gave a randomized strategy-proof mechanism, called Proportional Mechanism, that is 4-approximate for general metric spaces. It is still unknown whether there exists a deterministic strategy-proof mechanism with bounded approximation ratio in a general metric space.

Due to the absence of any positive result on the approximability of multiple facility location games for *more than two* facilities, Fotakis and Tzamos [3] considered a variant of the game where an authority can impose on some agents the facilities where they will be served. With this restriction, they proved that the Proportional Mechanism is strategy-proof and has an approximation ratio linear on the number of facilities.

1.3 Contribution

Prior to our work, only extreme cases of the game where the authority opens one or two facilities have been considered. No result, positive or negative, has been known for the game with three or more facilities. Toward the general number of facilities, we need to understand and solve the extreme cases of the problem. We consider here the extreme case where many facilities will be opened.

This type of situation occurs when every agent would like to have its own personal facility. The problem becomes interesting when it lacks at least one facility to satisfy everyone, i.e. $k = n - 1$. For instance, consider a blood collection agency that wishes to install 19 removable collection centers in the city of Paris, which consists of 20 districts. The agency asks every district council for the most

Objective	Tree metric space	General metric space
Utilitarian	UB: $n/2$ (rand)	UB: $n/2$ (rand)
	LB: $3/2$ (det), 1.055 (rand)	LB: 3 (det), 1.055 (rand)
Egalitarian	UB: $3/2$ (rand)	UB: n (rand)
	LB: $3/2$ (rand) [9]	LB: 2 (det)

Table 1. Summary of our results. In a cell, UB and LB mean the upper and lower bounds on the approximation ratio of strategy-proof mechanisms. Abbreviation *det* (resp. *rand*) refers to deterministic (resp. randomized) strategy-proof mechanisms.

frequented spot in the district, and will place the facilities so as to serve them at best (minimize the sum of the distances from these spots to the nearest centers).

Another example, more related to computer science, is the service of k servers for online requests in the metric of n points. This issue, which is the k -servers problem [4], has been extensively studied and plays an important role in Online Algorithms. The special case of k servers for the metric of $(k+1)$ points is widely studied [2]. Similar problematics have also been addressed in Algorithmic Game Theory for the replication of data in a network, from the viewpoint of Price of Anarchy and Stability [8]. These issues are also interesting from the viewpoint of strategy-proofness. Assume that each server replicates some data to optimize the requests of the clients, but the positions of the clients in the network are private. The efficiency of the request answer depends on the distance from the client to the nearest server. The clients are thus asked for their positions, and one wishes to minimize the sum of the distances from the clients to the nearest servers.

In this paper, we study strategy-proof mechanisms for the game with n agents and $n-1$ facilities in a general metric space and in a tree metric space. Our main results are the following ones. For general metric spaces, we give a randomized strategy-proof mechanism, called *Inversely Proportional Mechanism*, that is an $n/2$ -approximation for the utilitarian objective and an n -approximation for the egalitarian one. For tree metric spaces, we present another randomized strategy-proof mechanism that particularly exploit the property of the metric. This mechanism is also an $n/2$ -approximation under the utilitarian objective but it induces a $3/2$ -approximation (tight bound) under the egalitarian objective.

Besides, several lower bounds on the approximation ratio of deterministic/randomized strategy-proof mechanisms are derived (see Table 1 for a summary). We proved that any randomized strategy-proof mechanism has ratio at least 1.055 even in the tree metric space. The interpretation of this result is that no mechanism, even randomized one, is both socially optimal and strategy-proof. Moreover, deterministic lower bounds of strategy-proof mechanisms are shown to be: at least $3/2$ in a tree metric space, utilitarian objective; at least 3 in a general metric space, utilitarian objective; and at least 2 in a general metric space, egalitarian objective. Note that the lower bounds given for a tree metric space hold even for a line metric space.

Organization We study the performance of randomized SP mechanisms in general metric spaces and in tree metric spaces in Section 2, and Section 3, respectively. Due to lack of space, some claims are only stated or partially proved.

2 SP Mechanisms for General Metric Spaces

2.1 Inversely Proportional Mechanism

Consider the setting of n agents whose true locations are $\mathbf{x} = (x_1, \dots, x_n)$. For each location profile $\mathbf{y} = (y_1, \dots, y_n)$, define $P_i(\mathbf{y})$ as the placement of $(n-1)$ facilities at the reported locations of all but agent i , i.e., $P_i(\mathbf{y}) = \{y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n\}$. Moreover, $d(y_i, P_i(\mathbf{y}))$ is the distance between y_i and her closest location in $P_i(\mathbf{y})$. The idea of the mechanism is to choose with a given probability a location y_i where no facility is open (and to put $n-1$ facilities precisely on the $n-1$ locations of the other agents), i.e., to choose with a given probability the placement $P_i(\mathbf{y})$. The main issue is to find suitable probabilities such that the mechanism is strategy-proof, and such that the expected cost is as small as possible.

Inversely proportional mechanism Let \mathbf{y} be a reported location profile. If there are at most $(n-1)$ distinct locations in profile \mathbf{y} then open facilities at the locations in \mathbf{y} . Otherwise, choose placement $P_i(\mathbf{y})$ with probability

$$p_i(\mathbf{y}) = \frac{\frac{1}{d(y_i, P_i(\mathbf{y}))}}{\sum_{j=1}^n \frac{1}{d(y_j, P_j(\mathbf{y}))}}$$

Lemma 1. *The Inversely Proportional Mechanism is strategy-proof in a general metric space.*

Sketch of the proof. Let $\mathbf{x} = (x_1, \dots, x_n)$ be the true location profile of the agents, and let $d_j := d(x_j, P_j(\mathbf{x}))$ for $1 \leq j \leq n$.

If there are at most $(n-1)$ distinct locations in profile \mathbf{x} then the mechanism locates one facility on each position: no agent has incentive to misreports its location. In the sequel, we assume that all the agent locations in \mathbf{x} are distinct. If all the agents report truthfully their locations, the cost of agent i is

$$c_i := c_i(f, \mathbf{x}) = \sum_{j=1}^n p_j(\mathbf{x}) \cdot d(x_i, P_j(\mathbf{x})) = p_i(\mathbf{x}) \cdot d_i = \frac{1}{\sum_{j=1}^n 1/d_j}$$

Thus $c_i < d_i$. Let us now suppose that i misreports its location and bids x'_i . Let $\mathbf{x}' = (x'_i, x_{-i})$ be the location profile when i reports x'_i and the other agents report truthfully their locations. Let $d'_j = d(P_j(x_j, \mathbf{x}'))$ for $j \neq i$ and $d'_i = d(P_i(x'_i, \mathbf{x}'))$. We will prove that $c'_i := c_i(f, \mathbf{x}') \geq c_i$. The new cost of agent i is:

$$c'_i = \sum_{j=1}^n p_j(\mathbf{x}') \cdot d(x_i, P_j(\mathbf{x}')) \geq p_i(\mathbf{x}') \cdot d_i + (1 - p_i(\mathbf{x}')) \min\{d_i, d(x_i, x'_i)\}$$

where the inequality is due to the fact that in $P_j(\mathbf{x}')$ (for $j \neq i$), agent i can choose either some facility in $\{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n\}$ or the facility opened at x'_i . Define $T := \{j : d'_j \neq d_j, j \neq i\}$. Note that

$$p_i(\mathbf{x}') = \frac{1/d'_i}{\sum_{j \notin T} 1/d_j + \sum_{j \in T} 1/d'_j + 1/d'_i}$$

Let $e := d(x_i, x'_i)$. Remark that i has no incentive to report its location x'_i in such a way that $e \geq d_i$ since otherwise $c'_i \geq p_i(\mathbf{x}') \cdot d_i + (1 - p_i(\mathbf{x}'))d_i = d_i > c_i$. In the sequel, consider $e < d_i$. In this case,

$$c'_i \geq p_i(\mathbf{x}') \cdot d_i + (1 - p_i(\mathbf{x}')) \cdot e$$

We also show that $e \geq |d'_i - d_i|$ by using the triangle inequality. Then, by considering two cases (whether $\sum_{j \in T} \frac{1}{d'_j}$ is larger than $\sum_{j \in T} \frac{1}{d_j}$ or not), we show that in both case $c'_i \geq c_i$ (technical details are omitted): any agent i has no incentive to misreport its location, i.e., the mechanism is strategy-proof. \square

Theorem 1. *The Inversely Proportional Mechanism is strategy-proof, an $n/2$ -approximation with respect to the utilitarian social cost and an n -approximation with respect to the egalitarian one. Moreover, there exists an instance in which the mechanism gives the approximation ratio at least $\frac{n}{2} - \epsilon$ for the utilitarian social cost, and $n - \epsilon$ for the egalitarian one, where $\epsilon > 0$ is arbitrarily small.*

Proof. By the previous lemma, the mechanism is strategy-proof. We consider the approximation ratio of this mechanism. Recall that $\mathbf{x} = (x_1, \dots, x_n)$ is the true location profile of the agents. Let $P_i := P_i(\mathbf{x})$, $d_i := d(x_i, P_i)$ and $p_i = p_i(\mathbf{x})$. Let $\ell := \arg \min\{d_i : 1 \leq i \leq n\}$. For the egalitarian social cost, due to the triangle inequality at least one agent has to pay $d_\ell/2$, while the optimal solution for the utilitarian objective has cost d_ℓ (placement P_ℓ for instance).

The mechanism chooses placement P_i with probability p_i . In P_i , agent i has cost d_i and the other agents have cost 0. Hence, the social cost induced by the mechanism (in both objectives) is $\sum_j p_j(\mathbf{x})d_j = \frac{n}{\sum_j 1/d_j}$. For the utilitarian objective, the approximation ratio is $\frac{n}{d_\ell \sum_j 1/d_j} < \frac{n}{2}$ since in the sum of the denominator, there are two terms $1/d_\ell$. Similarly, it is at most $\frac{2n}{d_\ell \sum_j 1/d_j} < n$ for the egalitarian objective.

We describe an instance on a line metric space in which the bounds $n/2$ and n are tight. Let M be a large constant. Consider the instance on a real line in which $x_1 = 1, x_2 = 2, x_{i+1} = x_i + M$ for $2 \leq i \leq n$. We get $d_1 = d_2 = 1$ and $d_i = M$ for $3 \leq i \leq n$. An optimal solution chooses to put a facility in each x_i for $i \geq 2$ and to put the last one in the middle of $[x_1, x_2]$. Its social cost is 1 for the utilitarian objective and $1/2$ for the egalitarian one. The cost (in both objectives) of the mechanism is

$$\frac{n}{\sum_{j=1}^n 1/d_j} = \frac{n}{2 + (n-2)/M} = \frac{nM}{2M + n - 2}$$

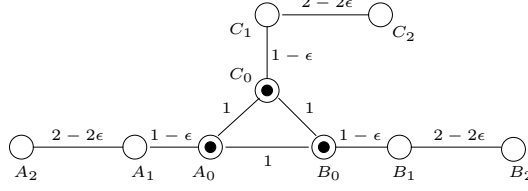


Fig. 1. Graph metric that gives a lower bound on the ratio of strategy-proof mechanisms in a general metric space (dots are the agents' locations in profile \mathbf{x}).

Hence, for any $\epsilon > 0$, one can choose M large enough such that the approximation ratio is larger than $\frac{n}{2} - \epsilon$ for the utilitarian objective and to $n - \epsilon$ for the egalitarian one. \square

2.2 Lower Bounds on the Approximation Ratio for SP Mechanisms

Proposition 1. *Any deterministic strategy-proof mechanism has approximation ratio at least $3 - 2\epsilon$ for the utilitarian objective and $2 - 2\epsilon$ for the egalitarian objective where $\epsilon > 0$ is arbitrarily small.*

Proof. We consider the metric space induced by the graph in Figure 1. Note that this is a discrete space where agents and possible locations for facilities are restricted to be on vertices of the graph, i.e., $\Omega = V$. There are three agents and two facilities to be opened. Let f be a deterministic strategy-proof mechanism. Let \mathbf{x} be a profile where $x_1 = A_0, x_2 = B_0, x_3 = C_0$. For any (deterministic) placement of two facilities, there is one agent with cost at least 1. By symmetry of the graph as well as profile \mathbf{x} , suppose that agent 1 has cost at least 1.

Consider another profile \mathbf{y} where $y_1 = A_1, y_2 = B_0, y_3 = C_0$ (\mathbf{y} and \mathbf{x} only differ on the location of agent 1). In this profile, no facility is opened neither at A_0 nor at A_1 since otherwise agent 1 in profile \mathbf{x} could report its location as being A_1 and reduce its cost from 1 to $1 - \epsilon$ or 0. We study two cases: (i) in profile $f(\mathbf{y})$, there is a facility opened at A_2 ; and (ii) in profile $f(\mathbf{y})$, no facility is opened at A_2 .

In the former, a facility is opened at A_2 , no facility is opened at A_0, A_1 . For the egalitarian objective, the social cost is at least $2 - 2\epsilon$. For the utilitarian objective, the total cost of agents 2 and 3 is at least 1 and the cost of agent 1 is $2 - 2\epsilon$, that induces a social cost at least $3 - 2\epsilon$. An optimal solution has cost 1 (for both objective) by opening a facility at A_1 and a facility at B_0 .

In the latter, the cost of agent 1 is at least $2 - \epsilon$ (since no facility is opened at A_0, A_1, A_2). Consider a profile \mathbf{z} similar to \mathbf{y} but the location of agent 1 is now at A_2 . By strategy-proofness, no facility is opened at A_0, A_1, A_2 in $f(\mathbf{z})$ (since otherwise, agent 1 in profile \mathbf{y} can decrease its cost by reporting its location as A_2). So, the social cost induced by mechanism f in \mathbf{z} is at least $4 - 3\epsilon$ (for both objectives), while optimal is 1 (for both objectives) by placing a facility at A_2 and other at B_0 .

Therefore, in any case, the approximation ratio of mechanism f is at least $3 - 2\epsilon$ for the utilitarian objective and $2 - 2\epsilon$ for the egalitarian objective. \square

3 Randomized SP Mechanisms on Trees

We study in this section the infinite tree metric. This is a generalization of the (infinite) line metric, where the topology is now a tree. Infinite means that, like in the line metric, branches of the tree are infinite. As for the line metric, the locations (reported by agents or for placing facilities) might be anywhere on the tree. We first devise a randomized mechanism. To achieve this, we need to build a partition of the tree into subtrees that we call components, and to associate a status even or odd to each component. This will be very useful in particular to show that the mechanism is strategy-proof. In the last part of this section, we propose a lower bound on the approximation ratio of any strategy-proof mechanism.

3.1 Preliminary Tool: Partition into Odd and Even Components

Partition procedure Given a tree \mathcal{T} and a set of vertices V on this tree, we partition \mathcal{T} into subtrees with respect to V . For the ease of description, consider also some virtual vertices, named ∞ , which represent the extremities of the branches in \mathcal{T} . We say that two vertices i and j are *neighbor* if the unique path in \mathcal{T} connecting i and j contains no other vertex ℓ . A *component* T_t is a region of the tree delimited by a *maximal* set of pairwise neighbor vertices (see below for an illustration). The maximality is in the sense of inclusion: T_t is maximal means that there is no vertex $i \notin T_t$ such that vertex i is a neighbor of all vertices in T_t . The set $\{T_1, \dots, T_m\}$ of all components is a cover of the tree \mathcal{T} . Note that a vertex i can appear in many sets T_t . As \mathcal{T} is a tree, the set of all T_t 's is well and uniquely defined.

For instance, in Figure 2, the components are the subtrees delimited by the following sets of vertices: $\{1, 2, 3\}$, $\{1, 4\}$, $\{2, 5\}$, $\{2, 6\}$, $\{6, 10\}$, $\{4, 7\}$, $\{4, 8, 9\}$, $\{3, \infty\}$, $\{5, \infty\}$, $\{8, \infty\}$, $\{9, \infty\}$, $\{10, \infty\}$.

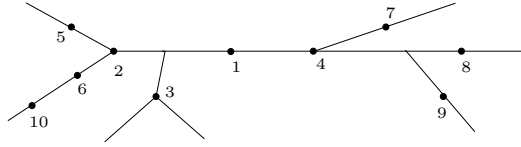


Fig. 2. An illustration of the partition procedure

Odd and even components Root the tree at some vertex i_0 , and define the depth of a vertex j as the number of vertices in the unique path from i_0 to j (i_0 has depth 1). Then each component T corresponds to the region of the tree between a vertex j (at depth p) and some of its sons (at depth $p+1$) in the tree. We say that T is odd (resp. even) if the depth p of j is odd (resp. even). This obviously depends on the chosen root.

For instance, in Figure 2 vertices of the same depth are in the same horizontal position (the tree is rooted at vertex 1). Then the components corresponding

to $\{1, 2, 3\}, \{1, 4\}, \{5, \infty\}, \{6, 10\}, \dots$ are odd while the ones corresponding to $\{2, 5\}, \{2, 6\}, \{3, \infty\}, \{4, 8, 9\}, \dots$ are even.

Note that each vertex except the root — and the ∞ -vertices — is both in (at least) one even component and in (at least) one odd component. The root is in (at least) one odd component.

3.2 A Randomized Mechanism

Given a reported profile \mathbf{y} and a tree \mathcal{T} as a metric space, let $2\alpha = 2\alpha(\mathbf{y})$ be the minimum distance between any two neighbor agents. Let $i^* = i^*(\mathbf{y})$ and $j^* = j^*(\mathbf{y})$ be neighbor agents such that $d(y_{i^*}, y_{j^*}) = 2\alpha$ (if there are more than one choice, break ties arbitrarily). We partition \mathcal{T} into its components as described previously, considering as vertices the set of locations \mathbf{y} . Let T^* be the component containing y_{i^*} and y_{j^*} , and let U be the set of agents in T^* . For instance, in Figure 3, the components are $\{7, 10, 11, 12\}, \{4, 6, 7, 8\}, \{6, 13\}, \{13, \infty\}, \dots$. Suppose that $i^* = 4$ and $j^* = 7$. Then T^* is the component whose set of agents is $U = \{4, 6, 7, 8\}$.

We design a mechanism made of four deterministic placements P_1, P_2, P_3 and P_4 ; each P_i occurs with probability $1/4$. Intuitively, the mechanism satisfies the following properties: (i) all agents have the same expected cost α , and (ii) for any component in \mathcal{T} , with probability $1/2$, no facility is opened inside the component (but possibly at its extremities). To get this, each agent i different from i^* and j^* will have its own facility F_i open at distance α , while i^* and j^* will “share” a facility open either at y_{i^*} , or at y_{j^*} , or in the middle of the path between y_{i^*} and y_{j^*} . However, to ensure strategy-proofness, we need to carefully combine these positions.

If we remove the component T^* (while keeping its vertices) from \mathcal{T} , we now have a collection of subtrees \mathcal{T}_i for $i \in U$, where \mathcal{T}_i is rooted at y_i (the location of agent i). For each rooted-subtree \mathcal{T}_i , assign the status odd or even to its components according to the procedure previously defined. In Figure 3 (B) if we remove T^* we have four subtrees rooted at 4, 6, 7 and 8. Bold components are odd.

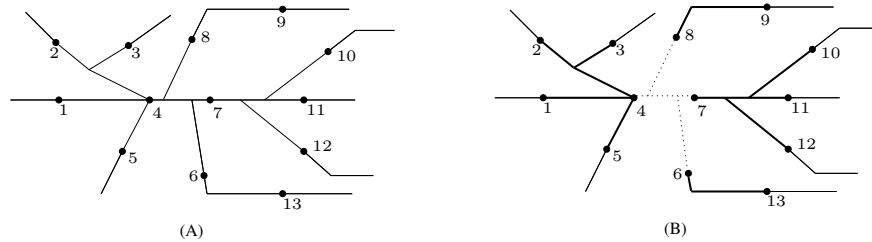


Fig. 3. (A) A tree \mathcal{T} and a profile \mathbf{y} where agents' locations are dots. (B) The four subtrees obtained after removing T^* . Bold components are the odd ones.

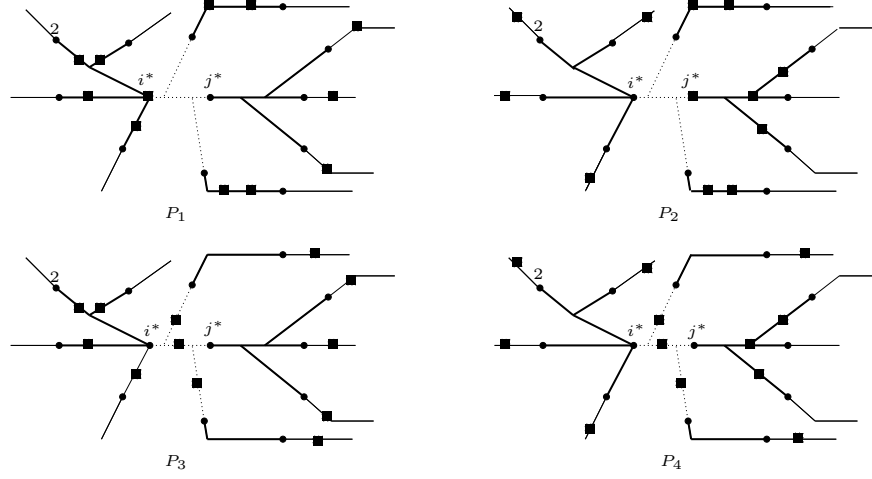


Fig. 4. Placements P_1, P_2, P_3, P_4 for the instance in Figure 3. Agents i^*, j^* are 4, 7. Facilities are represented by squares.

We are now able to define the four placements P_1, P_2, P_3, P_4 . Nevertheless, recall that a node is in at least one odd component and at least one even component. Each agent $i \neq i^*, j^*$ is associated with a facility F_i , while i^* and j^* share a common facility. We describe in the following the placements of these facilities. We distinguish the agents with respect to the subtree \mathcal{T}_i where they are.

Placement	i^*	$i \in \mathcal{T}_{i^*}$	j^*	$i \in \mathcal{T}_{j^*}$	$i \in U \setminus \{i^*, j^*\}$	$i \in \mathcal{T}_\ell \setminus U$ for $\ell \in U$
P_1	at y_{i^*}	<i>O</i>	no facility	<i>E</i>	<i>O</i>	<i>O</i>
P_2	no facility	<i>E</i>	at y_{j^*}	<i>O</i>	<i>O</i>	<i>O</i>
P_3	mid. y_{i^*}, y_{j^*}	<i>O</i>	no facility	<i>E</i>	T^*	<i>E</i>
P_4	no facility	<i>E</i>	mid. y_{i^*}, y_{j^*}	<i>O</i>	T^*	<i>E</i>

Table 2. Placements of facilities associated with agents

In Table 2, *E* (resp. *O*) means that we open a facility F_i in an even component (resp. odd component) at distance α of y_i for agent i ; T^* means that the facility F_i is opened in the component T^* , with distance α from y_i . For the location of any facility, if there are several choices, pick one arbitrarily. In placements P_3 and P_4 “mid. i^*, j^* ” means that the position is the middle of the path connecting y_{i^*} and y_{j^*} . We denote by $F^*(\mathbf{y})$ the facility opened at this position. In this case, i^* and j^* share the same facility $F^*(\mathbf{y})$.

An illustration is shown in Figure 4. For instance, since y_2 is in the subtree $\mathcal{T}_4 = \mathcal{T}_{i^*}$, the facility F_2 associated with agent 2 is opened in an odd (bold) component in placements P_1 and P_3 and in an even one in placements P_2 and P_4 .

Analysis By definition, all the placements P_1, P_2, P_3, P_4 are well defined, i.e., there are at most $n - 1$ opening facilities in each placement (one associated to

each agent $i \neq i^*, j^*$, plus only one shared by i^* and j^*). The following lemma shows some properties of the mechanism.

Lemma 2. *Given a reported profile \mathbf{y} , the expected distance between y_i and its closest facility equals $\alpha(\mathbf{y})$ for $1 \leq i \leq n$. Moreover, for any component, there are at least two placements in $\{P_1, P_2, P_3, P_4\}$ where the component does not contain any facility (but facilities can be at the extremities of the component).*

Proof. Consider an agent $i \neq i^*(\mathbf{y}), j^*(\mathbf{y})$ where we recall that $i^*(\mathbf{y}), j^*(\mathbf{y})$ denote the two players whose reported locations are at minimum distance. In any placement, the closest facility is opened at distance $\alpha(\mathbf{y})$ from y_i . For agent $i^* = i^*(\mathbf{y})$, the distance from y_{i^*} to the closest facility is: 0 in P_1 , $2\alpha(\mathbf{y})$ in P_2 , $\alpha(\mathbf{y})$ in P_3 and P_4 . Hence, the average is $\alpha(\mathbf{y})$, and similarly for agent $j^*(\mathbf{y})$.

Let T be the component containing the locations of agents $i^*(\mathbf{y})$ and $j^*(\mathbf{y})$. No facility is opened inside T under placements P_1 and P_2 . Besides, by the definition of the mechanism, there are at least two placements in $\{P_1, P_2, P_3, P_4\}$ where a component does not contain a facility¹. \square

Now we prove the strategy-proofness of the mechanism. Suppose that an agent i strategically misreports its location as x'_i (while other agents' locations remain unchanged). Let $\mathbf{x}' = (x'_i, x_{-i})$, where $\mathbf{x} = (x_1, \dots, x_n)$ is the true location profile. Define the parameters $2\alpha := 2\alpha(\mathbf{x}), i^* := i^*(\mathbf{x}), j^* := j^*(\mathbf{x})$. For every agent i , $\mathcal{N}(i, \mathbf{x})$ denotes the set of its neighbors in profile \mathbf{x} ($\mathcal{N}(i, \mathbf{x})$ does not contain i). The strategy-proofness is due to the two following main lemmas.

Lemma 3. *No agent i has incentive to misreport its location as x'_i such that $\mathcal{N}(i, \mathbf{x}) \neq \mathcal{N}(i, \mathbf{x}')$.*

Proof. Suppose that $\mathcal{N}(i, \mathbf{x}) \neq \mathcal{N}(i, \mathbf{x}')$. In this case, the locations of agents in $\mathcal{N}(i, \mathbf{x})$ form a component T' of tree \mathcal{T} with respect to profile \mathbf{x}' . By Lemma 2, with probability at least $1/2$, no facility is opened in T' , i.e., in those cases agent i is serviced by a facility outside T' . Note that the distance from x_i to the location of any agent in $\mathcal{N}(i, \mathbf{x})$ is at least 2α . Therefore, the new cost of agent i is at least α , meaning i has no incentive to report x'_i . \square

Lemma 4. *Agent i cannot strictly decrease its cost by reporting a location $x'_i \neq x_i$ such that $\mathcal{N}(i, \mathbf{x}) = \mathcal{N}(i, \mathbf{x}')$.*

Proof. As $\mathcal{N}(i, \mathbf{x}) = \mathcal{N}(i, \mathbf{x}')$, the path connecting x_i and x'_i contains no other agent's location. Hence, there is a component T'_i in the partition of \mathcal{T} with respect to \mathbf{x}' such that $x'_i \in T'_i$ and $x_i \in T'_i$. Let $2\alpha'$ be the minimum distance between two neighbors in \mathbf{x}' . Also let $e = d(x_i, x'_i)$.

¹ There are facilities in T under P_3 and P_4 but facilities are put on the extremities under placements P_1 and P_2 . Notice that a component may never receive a facility if there are two components named $\{i, \infty\}$ and i is located at the intersection of two branches of the tree, see location 3 in Figure 2.

Case 1: Consider the case where, with the new location x'_i , i is neither $i^*(\mathbf{x}')$ nor $j^*(\mathbf{x}')$. Hence, $\alpha' \geq \alpha$. By Lemma 2, with probability at least $1/2$, no facility is opened inside T'_i . In this case, the distance from x_i to the closest facility is at least $\min\{d(x_i, x'_i) + d(x'_i, F'_i), d(x_i, x_\ell) + d(x_\ell, F'_\ell)\}$ where: $\ell \in \mathcal{N}(i, \mathbf{x})$ and F'_ℓ is its associated facility; and F'_i is the facility opened at distance α' from x'_i , F'_i is in a component different from T'_i . In other words, this distance is at least $\min\{e + \alpha', 2\alpha\}$ since $d(x'_i, F'_i) = \alpha'$ and $d(x_i, x_\ell) \geq 2\alpha$. Besides, with probability at most $1/2$, the closest facility to x_i is either F'_i (the facility opened in component T'_i at distance α' from x'_i) or some other facility F'_ℓ in T'_i for some $\ell \in \mathcal{N}(i, \mathbf{x})$. The former gives a distance $d(x_i, F'_i) \geq \max\{d(x'_i, F'_i) - d(x'_i, x_i), 0\} = \max\{\alpha' - e, 0\}$ (by triangular inequality). The latter gives a distance $d(x_i, F'_\ell) \geq \max\{d(x_i, x_\ell) - d(x_\ell, F'_\ell), 0\} \geq \max\{2\alpha - \alpha', 0\}$. Hence, the cost of agent i is at least

$$\frac{1}{2} (\min\{e + \alpha', 2\alpha\} + \min\{\max\{\alpha' - e, 0\}, \max\{2\alpha - \alpha', 0\}\}) \geq \alpha$$

where the inequality is due to $\alpha' \geq \alpha$. Indeed, this is immediate if $e + \alpha' \geq 2\alpha$. Otherwise, the cost is either at least $e + \alpha' + \alpha' - e = 2\alpha'$, or $e + \alpha' + 2\alpha - \alpha' \geq 2\alpha$. Hence, $c_i(\mathbf{x}') \geq c_i(\mathbf{x})$.

Case 2: Consider the case where with the new location x'_i agent $i = i^*(\mathbf{x}')$ (the case where $i = j^*(\mathbf{x}')$ is completely similar)². Let $j = j^*(\mathbf{x}')$. Let d_1, d_2, d_3, d_4 be the distance from x_i to the closest facility in placements P_1, P_2, P_3, P_4 (in \mathbf{x}'), respectively. Let T' be the component in \mathcal{T} with respect to \mathbf{x}' that contains x'_i and x_j . By the triangle inequality, we know that

$$e + 2\alpha' = d(x_i, x'_i) + d(x'_i, x_j) \geq d(x_i, x_j) \geq 2\alpha \quad (1)$$

We study the two sub-cases and prove that $\sum_{t=1}^4 d_t \geq 4\alpha$ always holds, meaning that agent i 's deviation cannot be profitable since its cost is α when it reports its true location x_i .

(a) The true location x_i belongs to T' .

For each agent $\ell \neq i, j$, let F'_ℓ be its associated facility. The facility opened in the middle of $[x'_i, x_j]$ is denoted by $F^*(\mathbf{x}')$. We have:

$$d_1 = \min\{d(x_i, x'_i), d(x_i, F'_\ell)\} = \min\{e, d(x_i, x_\ell) + d(x_\ell, F'_\ell)\} \geq \min\{e, 2\alpha + \alpha'\} \quad (2)$$

$$d_2 = \min\{d(x_i, x_j), d(x_i, F'_\ell)\} \geq \min\{d(x_i, x_j), 2\alpha + \alpha'\} \geq 2\alpha \quad (3)$$

$$d_3 = \min\{d(x_i, F^*(\mathbf{x}')), d(x_i, F'_\ell)\} \geq \min\{2\alpha - \alpha', e + \alpha', 2\alpha + \alpha'\} \quad (4)$$

$$d_4 = \min\{d(x_i, F^*(\mathbf{x}')), d(x_i, F'_\ell)\} \geq \min\{2\alpha - \alpha', e + \alpha', 2\alpha + \alpha'\} \quad (5)$$

where $\ell \neq i, j$ is some agent in $\mathcal{N}(i, \mathbf{x}')$ (note that agents ℓ in the expressions above are not necessarily the same). The first equality in (2) is due to the fact that in placement P_1 , agent i goes either to the facility opened at x'_i or

² Contrasting with Case 1, $\alpha' \leq \alpha$ does not necessarily hold.

to a facility (outside T') associated to some other agent. In placement P_2 , agent i can either choose a facility opened at x_j or another one outside T' , that is translated in the equality in (3). In placement P_3 and P_4 , agent i can go either to facility $F^*(\mathbf{x}')$ opened in the midpoint connecting x'_i and x_j , or to the facility associated with some agent ℓ (inside and outside T' respectively).

If $e + \alpha' < 2\alpha - \alpha'$ then $\sum_{t=2}^4 d_t \geq 2\alpha + 2e + 2\alpha' \geq 4\alpha$ (since $e + 2\alpha' \geq 2\alpha$). In the sequel, assume $e + \alpha' \geq 2\alpha - \alpha'$. If $e \geq 2\alpha + \alpha'$ then $d_1 + d_3 \geq 4\alpha$. Otherwise, $\sum_{t=1}^4 d_t \geq e + \min\{d(x_i, x_j), 2\alpha + \alpha'\} + 2\max\{2\alpha - \alpha', 0\}$. Note that by the triangle inequality $e + d(x_i, x_j) = d(x'_i, x_i) + d(x_i, x_j) \geq d(x'_i, x_j) = 2\alpha'$. Therefore, $\sum_{t=1}^4 d_t \geq \min\{2\alpha' + 4\alpha - 2\alpha', 2\alpha + \alpha' + 2\alpha - \alpha'\} = 4\alpha$. Hence, the new cost of i is at least α .

- (b) The true location x_i does *not* belong to T' .

Let T'_i be the component in \mathcal{T} with respect to profile \mathbf{x}' such that T'_i contains x_i and x'_i . Similar to the previous case, we have:

$$\begin{aligned} d_2 &= \min\{d(x_i, x_j), d(x_i, F'_\ell)\} = \min\{d(x_i, x'_i) + d(x'_i, x_j), d(x_i, x_\ell) + d(x_\ell, F'_\ell)\} \\ &\geq \min\{e + 2\alpha', 2\alpha + \alpha'\} \geq 2\alpha \end{aligned} \quad (6)$$

$$\begin{aligned} d_3 &= \min\{d(x_i, F^*(\mathbf{x}')), d(x_i, F'_\ell)\} \\ &\geq \min\{d(x_i, x'_i) + d(x'_i, F^*(\mathbf{x}')), d(x_i, x_\ell) - d(x_\ell, F'_\ell)\} = \min\{e + \alpha', 2\alpha - \alpha'\} \end{aligned} \quad (7)$$

$$d_4 = \min\{d(x_i, F^*(\mathbf{x}')), d(x_i, F'_\ell)\} \geq \min\{e + \alpha', 2\alpha + \alpha'\} \quad (8)$$

where $\ell \neq i, j$ is some agent in $\mathcal{N}(i, \mathbf{x}')$ (again agents ℓ in the expressions above are not necessarily the same). In placement P_2 , agent i can choose either a facility opened at x_j or another one outside T'_i . The last inequality of (6) is due to $e + 2\alpha' \geq 2\alpha$ (Inequality 1). In placement P_3 and P_4 , agent i can go either to facility $F^*(\mathbf{x}')$ opened in the midpoint connecting x'_i and x_j , or some facilities associated with some agent ℓ .

If $e + \alpha' < 2\alpha - \alpha'$ then $\sum_{t=2}^4 d_t \geq 2\alpha + 2e + 2\alpha' \geq 4\alpha$ (since $e + 2\alpha' \geq 2\alpha$). Otherwise, $\sum_{t=2}^4 d_t \geq \min\{e + 4\alpha, 4\alpha\} \geq 4\alpha$. Again, the new cost of agent i is at least α .

In conclusion, no agent has incentive to strategically misreport its location. \square

Theorem 2. *The mechanism is strategy-proof and it induces an $n/2$ -approximation according to the utilitarian objective and a tight $3/2$ -approximation according to the egalitarian objective.*

Proof. The mechanism is strategy-proof by previous lemmas. The cost of each agent is α , so in the utilitarian objective, the cost induced by the mechanism is $n\alpha$. An optimal placement is to open facilities at the locations of all agents but i^* , which induces a cost 2α . Hence, the mechanism is $n/2$ -approximation for the utilitarian objective.

Consider the egalitarian objective. By the mechanism, in P_3 and P_4 the maximum cost of an agent is α , while in P_1 and P_2 it is 2α . The average maximum cost of the mechanism is $3\alpha/2$. An optimal solution is to open facilities at locations of agents other than i^*, j^* and open one facility at the midpoint of the path connecting x_{i^*} and x_{j^*} ; that gives a cost α . So, the approximation ratio is $3/2$ and this ratio is tight, i.e., no randomized strategy-proof mechanism can do better [9, Theorem 2.4]. \square

3.3 Lower Bounds on the Approximation Ratio of SP Mechanisms

In this section, we consider only the utilitarian objective (as the tight bound for the egalitarian objective has been derived in the previous section). The proof of Proposition 2 is omitted.

Proposition 2. *No deterministic strategy-proof mechanism on a line metric space has an approximation ratio smaller than $3/2$.*

The following proposition indicates that even with randomization, we cannot get an optimal strategy-proof mechanism for the utilitarian objective.

Proposition 3. *No randomized strategy-proof mechanism on a line metric space has an approximation ratio smaller than $10 - 4\sqrt{5} \approx 1.055$.*

Proof. Let f be a randomized strategy-proof mechanism with an approximation ratio strictly better than $1 + \varepsilon > 1$. Consider a profile \mathbf{x} where the positions of the agents are $x_1 = A, x_2 = B, x_3 = C, x_4 = D$ (Figure 5). For any placement of three facilities, the total cost is at least 1. Hence, there exists an agent with (expected) cost at least $1/4$. Without loss of generality, suppose that agent 1 (with $x_1 = A$) has cost $c_1(f, \mathbf{x}) \geq 1/4$.

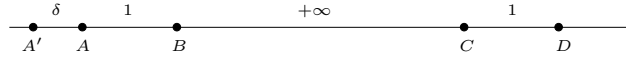


Fig. 5. Instance which gives the lower bound on the ratio of a randomized strategy-proof mechanism in a line metric space.

Let $0 < \delta < 1/4$ be a constant to be defined later. Let $A' \notin [A, B]$ be a location at distance δ from A . Let \mathbf{y} be the profile in which agent 1 is located at $y_1 = A'$ and the other agents' locations are the same as in \mathbf{x} . By strategy-proofness, $c_1(f, \mathbf{x}) \leq \delta + c_1(f, \mathbf{y})$. Hence, $c_1(f, \mathbf{y}) \geq 1/4 - \delta$. In \mathbf{y} , an optimal solution has cost 1 (e.g. place the facilities at the locations of the agents other than agent 4). As f is a $(1 + \varepsilon)$ -approximation, the total cost of the solution returned by the mechanism is $c_1(f, \mathbf{y}) + c_2(f, \mathbf{y}) + c_3(f, \mathbf{y}) + c_4(f, \mathbf{y}) \leq 1 + \varepsilon$. Thus, $c_3(f, \mathbf{y}) + c_4(f, \mathbf{y}) \leq 3/4 + \varepsilon + \delta$.

In outcome $f(\mathbf{y})$, let p be the probability that the closest facility of agent 3 is also the closest facility of agent 4 (in other words, agents 3 and 4 share one facility with probability p ; and with probability $(1 - p)$ there is at most one facility between A' and B). We have $c_3(f, \mathbf{y}) + c_4(f, \mathbf{y}) \geq p \cdot 1 = p$. Therefore, $p \leq 3/4 + \varepsilon + \delta$.

Besides, the social cost of $f(\mathbf{y})$ is at least $p + (1 - p)(1 + \delta) = 1 + \delta - p\delta$. This is lower bounded by $1 + \delta - (3/4 + \varepsilon + \delta)\delta$. Hence, $1 + \delta - (3/4 + \varepsilon + \delta)\delta \leq C(f, \mathbf{y}) \leq 1 + \varepsilon$. We deduce that $\varepsilon \geq \frac{\delta/4 - \delta^2}{1 + \delta}$.

The function $\frac{\delta/4 - \delta^2}{1 + \delta}$ for $\delta \in (0, \frac{1}{4})$ attains maximal value $9 - 4\sqrt{5}$ for $\delta = \sqrt{5}/2 - 1$. Thus the approximation ratio is at least $1 + \varepsilon \geq 10 - 4\sqrt{5} \approx 1.055$. \square

4 Discussion and Further Directions

The results presented in this paper are a first step toward handling the general case where one wishes to locate k facilities in a metric space with n agents (for $1 \leq k \leq n$). The general case is widely open since nothing on the performance of strategy-proof mechanisms is known. Any positive or negative results on the problem would be interesting. We suggest a mechanism based on the *Inversely Proportional Mechanism* in which the k facilities are put on reported locations. Starting with the n reported locations the mechanism would iteratively eliminate a candidate until k locations remain. We do not know whether this mechanism is strategy-proof. For restricted spaces such as line, cycle or tree metric spaces, there might be some specific strategy-proof mechanisms with guaranteed performance which exploits the structures of such spaces. Besides, some characterization of strategy-proof mechanisms (as done by Moulin [6] or Schummer and Vohra [10]), even not a complete characterization, would be helpful.

References

1. Alon, N., Feldman, M., Procaccia, A.D., Tennenholtz, M.: Strategyproof approximation of the minimax on networks. *Math. Oper. Res.* **35** (2010) 513–526
2. Coppersmith, D., Doyle, P., Raghavan, P., Snir, M.: Random Walks on Weighted Graphs and Applications to On-line Algorithms. *J. of ACM* **40**(3) (1993) 421–453
3. Fotakis, D., Tzamos, C.: Winner-imposing strategyproof mechanisms for multiple facility location games. In: *Proc. 6th WINE* (2010) LNCS 6484, 234–245
4. Koutsoupias, E.: The k -server problem. *Comp. Science Rev.* **3**(2) (2009) 105–118
5. Lu, P., Sun, X., Wang, Y., Zhu, Z.A.: Asymptotically optimal strategy-proof mechanisms for two-facility games. In: *ACM Conf. on Electronic Com.* (2010) 315–324
6. Moulin, H.: On strategy-proofness and single peakedness. *Public Choice* **35** (1980) 437–455
7. Nguyen Kim, T.: On (group) strategy-proof mechanisms without payment for facility location games. In: *Proc. 6th WINE* (2010) LNCS 6484, 531–538
8. Pollatos, G.G., Telelis, O.A., Zissimopoulos, V.: On the Social Cost of Distributed Selfish Content Replication. *Proc. 7th Int. IFIP Networking Conf.* (2008) 195–206
9. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: *ACM Conference on Electronic Commerce* (2009) 177–186
10. Schummer, J., Vohra, R.V.: Strategy-proof location on a network. *Journal of Economic Theory* **104** (2001)