Truthful Mechanisms for the Location of Different Facilities

(Extended Abstract)

Paolo Serafino
School of Computing
Teesside University
p.serafino@tees@ac.uk

Carmine Ventre
School of Computing
Teesside University
c.ventre@tees@ac.uk

ABSTRACT

In this paper we formalize and initiate the study of *heterogeneous k-facility location without money*, a problem akin to the classical k-facility location problem but encompassing a richer model and featuring multi-parameter agents. In particular, we consider *truthful mechanisms without money* for the problem in which *heterogeneous* (i.e. serving different purposes) facilities have to be located and agents are only interested in some of them. We study the approximation factor that can be achieved by truthful mechanisms in this setting and present some bounds which make a surprising parallel with our knowledge of truthfulness for the classical single-dimensional facility location problem.

Categories and Subject Descriptors

F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Theory, Economics

Keywords

Algorithmic Mechanism Design; Facility Location; Mechanisms without Money

1. INTRODUCTION

Mechanism Design is an established research field mainly concerned with optimization problems that have to operate under the assumption that their input is distributed across *selfish* agents. In this setting, *mechanisms* (i.e., typically *allocation algorithms*) have to elicit their input from the agents and have to ensure that agents report *truthfully*, i.e., in a *strategyproof* (*SP*) way, the part of input they possess. The challenge faced in this setting is that agents are not reliable, in the sense that they can misreport their private information (thus leading to a sub-optimum allocation) if they can gain by doing so. *Truthfulness* is typically achieved by means of suitable payment functions that realign, by means of financial compensation, the objective of the mechanism designer with the agents' personal objectives. Alas, in real-world scenarios it is often the case that monetary transfers between the mechanism designer and the agents cannot be performed, either because they are illegal (e.g., organ donations), unethical (e.g., political decision making) or difficult to implement. Motivated by this kind of considerations, Procaccia and Tennenholz [3] proposed the research agenda of *approximate mechanism design without money*, which aims at leveraging *approximation*, instead of payments, as a means of enforcing truthfulness.

In this line of work, the simple yet general and elegant problem of *facility location* has attracted much interest. This problem consists of locating a set of facilities on a given network on input the bids of the agents for their locations on the network (that is, agents are single-parameter). The agents' objective is to minimize their own connection costs (i.e. the distance to the nearest facility), whereas the designer's objective is to minimize the sum of the connection costs of all the agents.

Following this wake, we formulate and initiate the study of *heterogeneous facility location without money*, a problem akin to the traditional facility location problem but featuring multi-parameter agents. More specifically, we propose to study truthful mechanisms without money for the problem in which *heterogeneous* (i.e., serving different purposes) facilities have to be located and agents are only interested to some of them. We study the approximation ratio of truthful mechanisms in this setting, and derive some approximation bounds which make a surprising parallel with the classical single-dimensional facility location problem.

2. RELATED WORK

The facility location problem has attracted much attention from many diverse research areas.

The Social Choice community has been mostly concerned with the problem of locating a single facility on the line. In his classical paper [2] Moulin characterizes the class of *generalized median voter schemes* as the only deterministic *SP* mechanisms for *single-peaked* agents on the line. Schummer and Vohra [4] extend the result of Moulin to the more general setting where continuous graphs are considered, characterizing *SP* mechanisms on *continuous lines and trees.* They show that on circular graphs every *SP* mechanism must be dictatorial.

From an Algorithmic Game Theory perspective, Procaccia and Tennenholtz [3] initiated the field of *approximate mechanism design without money* by suggesting the idea of adopting *approximation* as a means of obtaining strategyproofness

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

in all those scenarios where (i) the optimal allocation is not *SP* and (ii) it is not possible to resort to monetary transfers. For the 2-facility location problem on the line, they propose the Two-Extremes algorithm, that places the two facilities in the leftmost and rightmost locations of the instance, and prove that it is group strategyproof and $(n-2)$-approximate, where n is the number of agents. Furthermore, they provide a lower bound of 3/2 for the approximation ratio of any *SP* algorithm and conjecture a lower bound of $\Omega(n)$. The latter conjecture was then proven by Fotakis et al. [1], along with a characterization of deterministic *SP* mechanisms with *bounded approximation ratio* for the 2-facility location problem on the line. In particular, they show that there exist only two such algorithms: (i) a mechanism that admits a unique dictator or (ii) the Two-Extremes mechanism proposed in [3].

3. MODEL AND RESULTS

We introduce a novel model, named *heterogeneous K-facility location*, where: (i) the location of the agents is common knowledge (i.e., not part of the agent's bid); (ii) agents request access to a certain set of *different*, *heterogeneous* facilities, each one serving a different purpose; and (iii) the cost of an agent is the sum of the connection costs she incurs when connecting to the facilities she bids for.

More formally, we are given: a set of agents $N = \{1, \ldots, n\}$; an undirected weighted graph $G = (V, E, w)$, where $V \supseteq N$ and $w : E \to \mathbb{R}^+ \cup \{0\}$; a set of facilities $\mathfrak{F} = \{\mathsf{F}_1, \ldots, \mathsf{F}_K\}$; and the nodes' capacity C. Agents' *types* are subsets of $\mathfrak{F}$, called their *facility set*. We denote the true type of agent i as $T_i \subseteq \mathfrak{F}$. A mechanism M for the heterogeneous K-facility location problem takes as input a vector of types $\mathcal{T} = (T_1, \ldots, T_n)$ and returns as output a feasible allocation, consisting of a location (i.e., vertex) of G for each of the facilities without violating the nodes' capacity, i.e., $M(\mathcal{T}) = (F_1, \ldots, F_K)$, such that $|\{1 \leq j \leq K | F_j = v\}| \leq C$ for any $v \in V$.

Given a feasible allocation $\mathcal{F} = (F_1, \ldots, F_K)$, agent i has a cost defined as $cost_i(\mathcal{F}) = \sum_{j \in T_i} d_G(i, F_j)$, where $d_G(i, F_j)$ denotes the length of the shortest path from i to F_j in G. Naturally, agents seek to minimize their cost. Therefore, they could misreport their facility sets to the mechanism if this reduces their cost. We let $T'_i \subseteq \mathfrak{F}$ denote a declaration of agent i to the mechanism. We are interested in the following class of mechanisms.

Definition 3.1. *A mechanism M is* truthful *(or* strategyproof*,* SP*, for short) if for any i, declarations of the other agents, denoted as $\mathcal{T}_{-i}$, and T'_i, we have $cost_i(\mathcal{F}) \leq cost_i(\mathcal{F}')$, where $\mathcal{F} = M(\mathcal{T})$ and $\mathcal{F}' = M(T'_i, \mathcal{T}_{-i})$.*

We want SP mechanisms that return an allocation minimizing the *social cost*, i.e. $M(\mathcal{T}) \in \underset{\mathcal{F} \text{ feasible}}{\operatorname{argmin}} \sum_{i=1}^{n} cost_i(\mathcal{F})$. We call these mechanisms *optimal*. However, in all those cases where the optimal allocation is not strategyproof we have to content ourselves with sub-optimal mechanisms. In particular, we say that a mechanism is *α-approximate* if the allocation it returns is α-approximate for the underlying optimization problem.

We studied different variations of the heterogeneous facility location problem, firstly tackling the problem where the underlying network has an arbitrary topology, and then moving to more restricted network structures. For *general graphs* (i.e., the topology of G is unrestricted) with *capacious nodes* (i.e., $C \geq K$) we proved that the optimal allocation is SP and can be computed in polynomial time.

Theorem 3.2. *For any graph G and any $C \geq K$, an algorithm minimizing the social cost is truthful.*

Theorem 3.3. *For any graph G and any $1 \leq C \leq K$, a solution minimizing the social cost can be computed in polynomial time.*

We then considered the restricted setting where: (i) agents are located on an *unweighted linear graph*, (ii) nodes have *unary capacity* (i.e., only one facility can be located at each node) and (iii) just *two facilities* need to be located. In this restricted model, which hereinafter will be referred to as 2-facility location, we proved that the optimal allocation is not strategyproof. Furthermore, we proved that there is no *SP* α-approximate mechanism with $\alpha < 9/8$ for 2-facility location.

Theorem 3.4. *No deterministic α-approximate truthful mechanism can obtain an approximation ratio $\alpha < 9/8$ for the heterogeneous 2-location problem.*

Inspired by [3], we devised an adaptation of Two-Extremes algorithm for the 2-facility location. Algorithm TwoExtreme locates the two facilities respectively in the leftmost and rightmost location of the subgraphs induced by the agents requesting each facility. We proved that the TwoExtremes algorithm is *SP* and $(n-1)$-approximate.

Theorem 3.5. *Algorithm* TwoExtremes *is SP and $(n-1)$-approximate.*

In order to provide better approximation guarantees still preserving truthfulness, we turned our attention to randomized algorithms. We devised a randomized algorithm, named RandOpt, which randomizes over optimal outcomes in such a way that the expected position of a facility is the average location of the set of optimal allocation points. We proved that RandOpt is optimal and *SP* in expectation.

Theorem 3.6. *Algorithm* RandOpt *is optimal and SP in expectation.*

We notice that this stands in stark contrast with both the deterministic case for the heterogeneous facility location and with other models proposed in the literature. Finally, we observe that there is a quite big gap between the *linear* (in the number of agents) upper bound and *constant* lower bound of deterministic mechanisms. Narrowing this gap, along with the study of other cost functions, will be the object of future research.

4. REFERENCES

[1] D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. In *ICALP*, pages 449–460, 2013.

[2] H. Moulin. On strategy-proofness and single-peakedness. *Public Choice*, 35:437–455, 1980.

[3] A. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *EC*, pages 177–186, 2009.

[4] J. Schummer and R. V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104:405–428, 2002.