Theoretical Computer Science 636 (2016) 27–46

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Heterogeneous facility location without money ☆

Paolo Serafino, Carmine Ventre *

School of Computing, Teesside University, UK

ARTICLE INFO

Article history:
Received 14 May 2015
Received in revised form 4 March 2016
Accepted 27 April 2016
Available online 30 April 2016
Communicated by P. Spirakis

Keywords:
Algorithmic mechanism design
Facility location
Approximate mechanism design without money

ABSTRACT

The study of the facility location problem in the presence of self-interested agents has recently emerged as the benchmark problem in the research on mechanism design without money. In the setting studied in the literature so far, agents are single-parameter in that their type is a single number encoding their position on a real line. We here initiate a more realistic model for several real-life scenarios. Specifically, we propose and analyze *heterogeneous facility location without money*, a novel model wherein: (*i*) we have multiple heterogeneous (i.e., serving different purposes) facilities, (*ii*) agents' locations are disclosed to the mechanism and (*iii*) agents bid for the set of facilities they are interested in (as opposed to bidding for their position on the network).

We study the heterogeneous facility location problem under two different objective functions, namely: *social cost* (i.e., sum of all agents' costs) and *maximum* cost. For either objective function, we study the approximation ratio of both deterministic and randomized truthful algorithms under the simplifying assumption that the underlying network topology is a line. For the social cost objective function, we devise an $(n-1)$-approximate deterministic truthful mechanism and prove a constant approximation lower bound. Furthermore, we devise an *optimal* and *truthful* (in expectation) randomized algorithm. As regards the maximum cost objective function, we propose a 3-approximate deterministic strategyproof algorithm, and prove a 3/2 approximation lower bound for deterministic strategyproof mechanisms. Furthermore, we propose a 3/2-approximate randomized strategyproof algorithm and prove a 4/3 approximation lower bound for randomized strategyproof algorithms.

1. Introduction

Mechanism design without money is a relatively recent and challenging research agenda introduced by Procaccia and Tennenholtz in [16]. It is mainly concerned with the design of *truthful* (or *strategyproof*, *SP* for short) *mechanisms* in scenarios where monetary compensation cannot be used as a means to realign the agents' interest to the mechanism designer's objective (as, e.g., done by VCG mechanisms). It has been noticed that such a circumstance occurs very frequently in real-life scenarios, as payments between agents and the mechanism are either illegal (e.g., organ transplant) or unethical (e.g., in the case of political decision making). To circumvent the impossibility of utilizing payments to enforce truthfulness, Procaccia and Tennenholtz propose instead to leverage the *approximation ratio* of the mechanism in those cases where the optimal outcome is not truthful. The facility location problem is arguably the archetypal problem in mechanism design without

☆ Preliminary versions of this paper have appeared as [19,18,20]. This work is partially supported by EPSRC grant EP/M018113/1.
* Corresponding author.
E-mail addresses: p.serafino@tees.ac.uk (P. Serafino), c.ventre@tees.ac.uk (C. Ventre).

http://dx.doi.org/10.1016/j.tcs.2016.04.033
0304-3975/

Table 1
Summary of our results.

	Social cost		Maximum cost	
	LB	UB	LB	UB
Deterministic	9/8	$n-1$	3/2	3
Randomized	1		4/3	3/2

money [16]. It demands locating a set of facilities on a network, on input the bids of the agents for their locations, in such a way as to optimize a given *objective function* that depends on agents' costs. If we regard the problem of locating facilities as a political decision (e.g., a city council locating facilities of public interest on the basis of the population residing in a certain area), the impossibility to utilize payments and the need to locate the facilities to minimize the social cost (e.g., the traffic in the city) in this context becomes immediately apparent.

Another application scenario that can be envisaged is *big data distribution in cloud networks*. Consider a multinational company having to decide how to distribute the data contained in its databases over its data network. Not all the various offices working for the company need access to the whole dataset, e.g., a payroll office arguably needs access to employees' data but not to customers', whilst sales offices need customers' data but not employees'. Thus, a demand-based allocation seems a sensible approach. However, due to well-known issues such as space consumption, consistency and query latency, it might be impractical to allow replication of the requested data at all the demanding offices' sites. Fast data access becomes then competitive and, guided by their willingness to have prompt access to the data they need, offices might strategize and amend their demands accordingly. The company, however, wants to minimize the maximum access time in order to guarantee a decent quality of service so that each office can work efficiently.

1.1. Our contribution

Inspired by the work on facility location without money, and aiming at analyzing a richer and more realistic setting, we introduce and study the problem of *heterogeneous facility location without money*. With respect to the main stream of works on facility location, our model features *heterogeneous* (i.e., serving different purposes) as opposed to the *homogeneous* (i.e. serving the same purpose) facilities. Allowing heterogeneous facilities influences the agent cost model as in our setting the cost of an agent is the cost to access the set of facilities she is interested in, rather than accessing (as in the traditional setting) the nearest facility. Furthermore, we assume in our model that agents' locations are disclosed to the mechanism. This assumption fits many real-life applications (e.g., for the aforementioned examples, the city council can ask for proof of residence whilst the multinational company knows where its payroll offices are located).

In more detail, we focus on the heterogeneous facility location problem in the case in which the agents are on a discrete *line* and we have *two* facilities to locate. Despite its apparent simplicity, this class of instances already encodes many intricacies and showcases the tension between truthfulness without money and approximation. Moreover, these instances model the aforementioned content distribution scenario (the linear network being the backbone of the company's data network; facilities being employee and customer records). We study both utilitarian (i.e., social cost) and non-utilitarian (i.e., max-cost) objective functions. Under either objective function, we analyze both deterministic and randomized algorithms (see Table 1), prove that in both cases the optimal allocation does not preserve truthfulness, and provide lower and upper bounds for the approximation of truthful mechanisms.

As regards the social cost objective function, we prove a 9/8 lower bound for the approximation of deterministic strategyproof algorithms. We then propose a *truthful* $(n-1)$*-approximate deterministic algorithm* named TwoExtremes, an adaptation to our model of a mechanism already proposed in [16], that assigns each facility to an extreme of the subnetwork of nodes requesting it. In order to provide better approximation guarantees, we then turn our attention to randomized algorithms and devise an *optimal randomized algorithm* that is *truthful in expectation*. At intuitive level, the reason for which deterministic optimal algorithms are not truthful resides in the richness of optimal solutions in very symmetric instances. For each way a deterministic optimum can break these ties, one side of the network will be disadvantaged and will then be able to manipulate the algorithm. The idea behind our randomized algorithm is to take care of these symmetries with randomization so that in expectation agents on either sides of the network are "happy". The technical challenge is that, in some cases, there are not enough optimal solutions to randomize upon and therefore a careful combination of deterministic and randomized solutions is designed and shown to preserve truthfulness.

As regards the maximum cost objective function, we prove a lower bound of 3/2 on the approximation guarantee of deterministic *SP* mechanisms. The proof connects three different instances and uses truthfulness constraints on two agents to establish the lower bound. This is somehow more complex than typical lower bounds in literature wherein two instances and one lying agent are normally considered. We then analyze TwoExtremes for maximum cost and prove it is 3-approximate. We observe that TwoExtremes retains its strategyproofness as the latter is independent from the objective function of the mechanism but depends solely on the agents' cost function. Regarding randomized mechanisms we first prove a lower bound of 4/3 and then design a 3/2-approximate randomized *SP* mechanism. This algorithm is mainly based on the idea of allocating (in expectation) each facility on the average position of the subgraph comprised of agents requesting it. This way truthfulness is guaranteed since there is no advantage in hiding one's own requested facilities as