Games and Economic Behavior 110 (2018) 174–193

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Setting lower bounds on truthfulness

Ahuva Mu'alem [a,*], Michael Schapira [b]

[a] Computer Science Department, Holon Institute of Technology (HIT), Holon, Israel
[b] School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel

ARTICLE INFO

Article history:
Received 31 July 2015
Available online 6 March 2018

JEL classification:
C60
D82

Keywords:
Algorithmic Mechanism Design
Scheduling
Bayesian Incentive-Compatible mechanisms
Truthfulness-in-expectation mechanisms
Randomized truthful mechanisms

ABSTRACT

This paper presents inapproximability results for paradigmatic multi-dimensional truthful mechanism design problems.
We first show a lower bound of $2 - \frac{1}{n}$ for the scheduling problem with n unrelated machines (formulated as a mechanism design problem in the seminal paper of Nisan and Ronen on Algorithmic Mechanism Design). Our lower bound applies to universally-truthful randomized mechanisms, regardless of any computational assumptions on the running time of these mechanisms. Moreover, it holds even for the wider class of truthfulness-in-expectation mechanisms.
We then turn to Bayesian settings and show a lower bound of 1.2 for Bayesian Incentive-Compatible (BIC) mechanisms. No lower bounds for truthful mechanisms in multi-dimensional settings which incorporate randomness were previously known.
Next, we define the workload-minimization problem in networks. We prove lower bounds for the inter-domain routing setting presented by Feigenbaum, Papadimitriou, Sami, and Shenker.
Finally, we prove lower bounds for Max–Min fairness, Min–Max fairness, and envy minimization.

1. Introduction

1.1. Inapproximability issues in Algorithmic Mechanism Design

Mechanism Design is a field of economic theory and game-theory that deals with protocols for optimizing global goals that require interaction with selfish players (Mas-Collel et al., 1995; Osborne and Rubinstein, 1994). *Algorithmic Mechanism Design* (Nisan and Ronen, 2001) combines economic perspectives (e.g., strategic behaviour of the players) with theoretical computer-science perspectives (such as computational-efficiency and approximability).

More formally, Algorithmic Mechanism Design attempts to solve problems of the following nature: given a finite set of *alternatives* $A = \{a, b, c, \ldots\}$, and a set of *strategic players* $N = \{1, \ldots, n\}$, each player i has a *valuation function* $v_i : A \to \mathbb{R}$ that is its own private information. The players are self-interested and only wish to maximize their own utility. The global goal is expressed by a *social choice function* f that assigns to every possible n-tuple of players' valuations $(v_1, \ldots, v_n)$ an alternative $a \in A$. Mechanisms are said to *truthfully implement* a social choice function f, if their outcome for every n-tuple

* Corresponding author.
E-mail addresses: ahumu@yahoo.com (A. Mu'alem), schapiram@huji.ac.il (M. Schapira).

https://doi.org/10.1016/j.geb.2018.02.001

of players' valuations matches that of f, and if they enforce payments of the different players in a way that motivates truthful reporting of their valuations (no matter what the other players do).[1]

A canonical social choice function is the *utilitarian* function, which aims to maximize the *social welfare*, i.e., to find the alternative $a \in A$ for which the expression $\Sigma_i \ v_i(a)$ is maximized. Another canonical social choice function is the *Max–Min* function (based on the philosophical work of Rawls, 1971): For every n-tuple $(v_1, \ldots, v_n)$ of valuations, the Max–Min function assigns the alternative $a \in A$ that maximizes the expression $\min_i \ v_i(a)$. Intuitively, the Max–Min function chooses the alternative $a \in A$ in which the least satisfied player has the highest value.

In many computational and economic settings we may wish to implement a utilitarian social choice function in a truthful manner. In such cases we can rely upon a classic result of mechanism design which states that for every utilitarian social choice function there exists a mechanism that truthfully implements it – namely, a member of the celebrated family of *VCG mechanisms* (Vickrey, 1961; Clarke, 1971; Groves, 1973). However, no such general technique is known for non-utilitarian social choice functions such as revenue maximization in auctions (e.g., Fiat et al., 2002), minimizing the makespan in scheduling (e.g., Nisan and Ronen, 2001; Archer and Tardos, 2001; Andelman et al., 2007; Dhangwatnotai et al., 2011; Christodoulou and Kovács, 2013), fair allocation of resources (e.g., Bikhchandani et al., 2006; Bezáková and Dani, 2005; Lipton et al., 2004), etc. In fact, some non-utilitarian social-choice functions cannot be truthfully implemented (Bikhchandani et al., 2006; Nisan and Ronen, 2001). Hence, it is natural to ask how well non-utilitarian social choice functions can be *approximated* in a truthful manner.

In their seminal paper on Algorithmic Mechanism Design (Nisan and Ronen, 2001), Nisan and Ronen formulated the following natural scheduling problem as a mechanism design problem: The global goal is to minimize the makespan of the chosen schedule; i.e., to assign the tasks to the unrelated machines in a way that minimizes the latest finishing time. Obviously, the makespan-minimization social choice function is non-utilitarian and hence *may* not be truthfully implemented by any mechanism. Nisan and Ronen prove that not only is it impossible to minimize the makespan in a truthful manner, but that *any approximation strictly better than* 2 *cannot be achieved by a truthful deterministic mechanism*. Since a non-truthful $(1+\epsilon)$-approximation exists (Horowitz and Sahni, 1976) (assuming constant number of machines), this raises a natural question:

Can optimality be achieved by incorporating randomness in a truthful manner?

More formally, can near-optimal $(1+\epsilon)$-approximation *truthful* mechanisms *which incorporate randomness* exist for multi-dimensional non-utilitarian settings?

1.2. Our results

In this paper we present lower bounds on the approximability of truthful mechanisms. We obtain the first lower bounds for several canonical non-utilitarian multi-dimensional settings which incorporate randomness.

Section 3 proves several lower bounds for the scheduling problem studied by Nisan and Ronen. In particular, we prove that no universally-truthful randomized mechanism can achieve an approximation ratio better than $2 - \frac{1}{n}$. This nearly matches the known truthful upper bound of 1.58606 for the case in which there are only two machines (Chen et al., 2015). Surprisingly, this lower bound applies even for the substantially weaker notion of truthfulness for randomized mechanisms – truthfulness-in-expectation. Furthermore, we show a lower bound of 1.2 for Bayesian Incentive Compatible mechanisms (also known as Bayesian truthful mechanisms).

Hence, truthful $(1+\epsilon)$-approximation with randomness is ruled out for the canonical unrelated machines problem (regardless of computational efficiency).

These are the first lower bounds for multi-dimensional settings which incorporate randomness. In fact, to the best of our knowledge these are the first lower bounds for universally truthful mechanisms, truthful-in-expectation mechanisms and Bayesian Incentive Compatible mechanisms in *multi-dimensional* settings in general.

In addition, we show how to prove lower bounds for the important class of *strongly-monotone* deterministic mechanisms. The strongly-monotone property (Lavi et al., 2003) is essentially similar to Arrow's Independence of Irrelevant Alternatives (IIA). Lavi et al. (2003) show that in several canonical domains this property can be assumed without loss of generality. This natural property says that the social choice between two alternatives depends only on the individual valuation difference between these two alternatives.[2] This is another step towards proving the long-standing conjecture of Ronen and Nisan that *no truthful deterministic mechanism can obtain an approximation ratio better than n*.

[1] It is well known (e.g., Mas-Collel et al., 1995) that, without loss of generality, we can limit ourselves to only considering direct-revelation truthful mechanisms. In such mechanisms participants are always rationally motivated to correctly report their private information.

[2] Together with decisiveness, strong-monotonicity essentially implies affine maximization in general combinatorial auctions domains and multi-unit domains (Lavi et al., 2003). In several discrete domains (such as unrestricted integer domains), strong-monotonicity is sufficient for truthful implementability, while weak-monotonicity is not (Mu'alem and Schapira, 2008). For a recent characterization of strongly-monotone scheduling mechanisms see Kovács and Vidali (2015).