Games and Economic Behavior 88 (2014) 29–46

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Fair by design: Multidimensional envy-free mechanisms

Ahuva Mu'alem [1]

Software Engineering Dept., Ort Braude College of Engineering, Karmiel, Israel

ARTICLE INFO

Article history:
Received 30 November 2010
Available online 15 August 2014

JEL classification:
D02
D40
C70
C61

Keywords:
Envy-free allocations
Optimization
Profit maximization
Mechanism design

ABSTRACT

We address the common scenario where a group of agents wants to divide a set of items fairly, and at the same time seeks to optimize a global goal. Suppose that each item is a task and we want to find an allocation that minimizes the completion time of the last task in an envy-free manner, where no agent prefers anyone else's allocated task bundle over its own. This optimization goal is called *makespan minimization*, and the agents are often treated as machines. We give *tight* deterministic bounds for: (1) two unrelated machines; and (2) $m \geq 2$ related machines.

A natural question to ask is whether envy-free pricing techniques can improve the current known bounds for *truthful* mechanisms for the task-scheduling problem studied in the seminal paper of Nisan and Ronen (2001). We find that for two unrelated machines, envy-free in-expectation is a far weaker constraint (i.e. less restrictive) than truthful in-expectation.

Published by Elsevier Inc.

1. Introduction

Traditionally fairness analysis focuses on the individual performance each participant receives, while allocation algorithms consider optimizing overall criteria. This paper formulates the *fair by design approach* to combine these two point of views.

Consider a project with different tasks to be assigned to a heterogeneous group of employees. As a motivating scenario assume that Alice is a technician who specializes in repairing antennas and Bob specializes in repairing battery charges. Suppose that Carol the customer has two antennas and two battery chargers for repair. Consider two possible allocations: (1) the allocation in which Alice receives two antennas and Bob receives two battery charges; and (2) the allocation in which each technician receives an antenna and a battery charger. If the goal of the manager is to complete the repair as soon as possible to please Carol, then the first allocation is preferable. However, if the manager's goal is to expand the expertise of Alice and Bob then the second allocation is preferable. Observe that both allocations are *fair* from the point of view of Alice and Bob. This is not the case in general.

A natural challenge is determining a *fair* allocation such that the last task completes as soon as possible. In general, no such allocation may exist if the tasks are indivisible. For instance, in a project with a single task, the fastest employee should be assigned the task. However, this allocation would not be considered fair from the perspective of the fastest employee. This suggests that some (additional) reward should be allowed to guarantee a fair division of the tasks. It is convenient to assume that rewards are granted in the form of monetary payments. In the scheduling literature, the above optimization goal is called *makespan* minimization of unrelated machines. In the economics literature, an allocation algorithm coupled with a price function is called a *mechanism*.

E-mail address: ahumu@yahoo.com.
URL: http://sites.google.com/site/ahuvamualem.

[1] This research was done while the author was a Post-Doctoral Fellow at the Social and Information Sciences Laboratory, California Institute of Technology.

http://dx.doi.org/10.1016/j.geb.2014.08.001
0899-8256/Published by Elsevier Inc.

We treat each machine as a distinct economic agent and say that an allocation is *envy-free* if no agent would prefer to exchange its assigned set of tasks with those of any other agent. Informally, a heterogeneous group of machines is called *unrelated* in the scheduling literature and *multidimensional* in the economics literature, while a homogeneous group is called *related* and *single-dimensional*, respectively.

The design of mechanisms for our general problem can be regarded as solving an optimization problem with binding envy-freedom constraints. However, even if we allow monetary payments, it could be that no feasible solution exists: we are faced with an inherent clash between global optimization goals and envy-free pricing constraints (regardless of any computational considerations). Thus the need to consider approximations motivates the following definition: we say that an allocation a is a *ρ-approximation* if the makespan of allocation a is no more than a factor of ρ times the makespan of an optimal allocation, where $\rho \geq 1$.[2] Our main question is how well can fundamental global goals be approximated in an envy-free manner?

Relation to truthful mechanisms

The design of *envy-free* mechanisms is intimately connected to the well-studied class of *truthful* mechanisms (Demange and Gale, 1985; Edelman et al., 2007; Varian, 2007; Hartline and Yan, 2011). A mechanism is truthful if no agent can ever improve its utility by misreporting his valuation or cost. In many interesting settings, truthful mechanisms are essentially equivalent to mechanisms that select envy-free allocations with the *smallest* supporting price vectors (Demange and Gale, 1985).

The seminal paper of Nisan and Ronen (2001) initiates the study of truthful mechanisms in computerized settings, examining the problem of how well the makespan goal on unrelated machines can be approximated. They showed that truthful mechanisms for two machines exhibit a *tight* bound of 2 (Nisan and Ronen, 2001), and they circumvent this impossibility result by employing randomization.

In a *truthful in-expectation mechanism*, each bidder prefers to truthfully report his value to the mechanism since this gives him higher *expected* utility. Nisan and Ronen presented a truthful in-expectation mechanism for two machines that exhibits an upper bound of $\frac{7}{4}$ (their upper bound was later improved to 1.5963 by Lu and Yu, 2008). The lower bound known for truthfulness in-expectation is $\frac{3}{2}$ (Mu'alem and Schapira, 2007).

As tightening the bounds for unrelated machines is a central problem of algorithmic mechanism design, this raises the question of whether envy-free techniques can tighten the current upper and lower bounds (1.5963 vs. $\frac{3}{2}$) known for truthful in-expectation mechanisms for two unrelated machines.

In this paper we study envy-free in-expectation mechanisms to minimize the makespan. We present a $\frac{4}{3}$-approximation envy-free in-expectation mechanism for two unrelated machines. By the simple fact that the envy-free in-expectation *upper bound* of $\frac{4}{3}$ is smaller than the truthful in-expectation *lower bound* of $\frac{3}{2}$ (Mu'alem and Schapira, 2007), we get that our $\frac{4}{3}$-approximation envy-free in-expectation mechanism is not truthful in-expectation. Intuitively, this suggests that envy-free bounding techniques cannot be applied straightforwardly to tighten the current truthful in-expectation bounds for minimizing the makespan on two unrelated machines. We conclude that in multidimensional settings, truthful in-expectation is a far more restrictive constraint than envy-free in-expectation.

Overview and results

In this paper we study two canonical objectives over multidimensional domains: profit-maximizing combinatorial auctions for general bidders and makespan-minimizing scheduling for unrelated machines.

We start by formally defining the notion of an envy-free allocation mechanism. In Section 2, we briefly state a known characterization of envy-free allocation mechanisms in terms of locally-efficient bundle assignments (Haake et al., 2002). Importantly, this characterization does not involve price functions.

In Section 3, we study envy-free combinatorial auctions for general bidders. In this scenario, a profit-maximizing auctioneer has a collection of items for sale, and bidders compete for subsets of items. Envy-free prices can be interpreted as anonymous non-discriminatory prices. We describe an envy-free mechanism that requires polynomial communication and achieves $(\min\{n, O(\sqrt{k}\log(\min\{k,n\}))\})$-approximation with respect to the maximum profit, where k is the number of items and n is the number of bidders. On the negative side, we show that any envy-free profit-maximizing mechanism with approximation ratio strictly better than n requires exponential communication.

In Section 4, we study envy-free scheduling mechanisms. There are k tasks that are to be scheduled on m unrelated machines. The total cost of a subset of tasks on machine i is the additive sum of the costs of the individual tasks on that machine. The global goal is minimizing the makespan of the chosen schedule; i.e., assigning the tasks to the machines in a way that minimizes the finishing time of the last task. This canonical optimization problem was extensively studied by Lenstra et al. (1990).

[2] Additionally, if we consider a maximization problem (such as profit maximization) we say that an allocation a is a *ρ-approximation* if the value of allocation a is at least a factor of $\frac{1}{\rho}$ times the value of an optimal allocation, where $\rho \geq 1$.

A fair-design problem exhibits an upper bound of ρ_U and a lower bound of ρ_L if there exists an envy-free ρ_U-approximation mechanism and if an envy-free $(\rho_L - \epsilon)$-approximation mechanism is impossible for every $\epsilon > 0$, respectively. If $\rho_U = \rho_L$ we say that the bounds are *tight*.