European Journal of Operational Research 217 (2012) 270–277

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Discrete Optimization

Cost-sharing mechanisms for scheduling under general demand settings

Sindhura Balireddi [a], Nelson A. Uhan [b,*]

[a] Amazon.com, 2201 Westlake Ave, Seattle, WA 98121, USA
[b] School of Industrial Engineering, Purdue University, Grissom Hall 308, West Lafayette, IN 47904, USA

ARTICLE INFO

Article history:
Received 26 November 2010
Accepted 13 September 2011
Available online 25 September 2011

Keywords:
Game theory
Cost sharing
Scheduling
Combinatorial optimization

ABSTRACT

We investigate cost-sharing mechanisms for scheduling cost-sharing games. We assume that the demand is general—that is, each player can be allocated one of several levels of service. We show how to design mechanisms for these games that are weakly group strategyproof, approximately budget-balanced, and approximately efficient, using approximation algorithms for the underlying scheduling problems. We consider scheduling cost-sharing games in single machine, parallel machine, and concurrent open shop environments.

1. Introduction

In a *general demand cost-sharing game*, there is a set of customers who is interested in receiving service from a service provider. The customers' demand for service is *general*—that is, each customer can receive various levels of service. Each customer has its own private valuation for the various levels of service. In order to determine which customers to serve, the service provider solicits bids from each customer. Based on these bids and the cost of providing service, the service provider determines the level of service it provides each customer, and how the cost of providing this service is shared by the customers—in other words, how much each customer has to pay. The algorithm that the service provider uses to determine these quantities is called a *cost-sharing mechanism*.

To illustrate, consider the following example, based on a scheduling problem. In this setting, there is a set of customers, each of which has a number of identical jobs it would like to be processed on a single machine. This machine is maintained by a service provider. Here, the level of service that a customer receives corresponds to the number of its jobs processed on the machine. The service provider solicits bids from the customers, and based on these bids and the processing costs, determines how many of each customer's jobs to process, and the price to charge each customer for processing its jobs.

There has been considerable work on cost-sharing mechanisms, focusing on designing mechanisms with various desirable properties, such as: (i) *truthfulness*, the idea that it is optimal for each customer to bid its private valuation, (ii) *budget-balance*, the notion that the service provider recovers the cost of providing the service, and (iii) *efficiency*, the idea that the total social welfare of the customers is maximized. Most of the work so far has been on binary demand cost-sharing games—that is, when customers either receive one level of service or none at all. For this case, Moulin (1999) and Moulin and Shenker (2001) proposed a class of cost-sharing mechanisms, known as *Moulin mechanisms*, and showed that these mechanisms achieve a notion of truthfulness known as *group strategyproofness*. Several researchers have studied the design of approximately budget-balanced Moulin mechanisms for a variety of cost-sharing games, such as those arising from network design (e.g., Jain and Vazirani, 2001; Archer et al., 2004; Gupta et al., 2007a,b; Roughgarden and Sundararajan, 2007), facility location (e.g., Pál and Tardos, 2003; Leonardi and Schäfer, 2004; Könemann et al., 2005; Immorlica et al., 2008), and logistics (e.g., Xu and Yang, 2009). Motivated by impossibility results on the existence of simultaneously truthful, budget-balanced, and efficient mechanisms (Green et al., 1976; Roberts, 1979), Roughgarden and Sundararajan (2009) developed an alternate framework to quantify efficiency in cost-sharing mechanisms. Mehta et al. (2009) proposed a generalization of Moulin mechanisms, called *acyclic mechanisms*, and showed that they achieve a weaker notion of truthfulness known as *weak group strategyproofness* for cost-sharing games with binary demand, as well as for those with general demand. Brenner and Schäfer (2008b) developed a framework for obtaining approximately budget-balanced and approximately efficient acyclic mechanisms for binary demand cost-sharing games. Brenner and Schäfer (2010) studied cost-sharing mechanisms in an online setting, in which players arrive over time and reveal their characteristics only at the time of arrival.

* Corresponding author. Tel.: +1 765 496 7241; fax: +1 765 494 1299.
E-mail addresses: ramaba@amazon.com (S. Balireddi), nuhan@purdue.edu (N.A. Uhan).

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.09.030

One type of cost-sharing game that has received a fair amount of attention is the kind that arises from scheduling problems. Like in the illustrative example described above, in a *general demand scheduling cost-sharing game*, each customer has a number of jobs that it would like to be processed by a set of machines, maintained by a service provider. Each customer requires that its jobs must be processed in a certain order, and the service level that a customer receives corresponds to the number of its jobs that are processed. The service provider's cost of processing these jobs is given by the optimal value of an associated scheduling problem (e.g., the minimum sum of weighted completion times). Brenner and Schäfer (2008a) and Bleischwitz and Monien (2009) studied the design of approximately budget-balanced and approximately efficient Moulin mechanisms for binary-demand cost-sharing games arising from various scheduling problems. In addition, Brenner and Schäfer (2008b) applied their framework for approximately budget-balanced and approximately efficient acyclic mechanisms to binary-demand scheduling cost sharing games. Finally, Brenner and Schäfer (2010) investigated cost-sharing mechanisms for scheduling in online settings.

In this work, we study how to design cost-sharing mechanisms for general demand cost-sharing games. We extend the framework of Brenner and Schäfer (2008b) and show how to use an approximation algorithm for the service provider's underlying optimization problem to obtain an acyclic mechanism that is approximately budget-balanced and approximately efficient (Theorems 3.1 and 3.2). Then, we give acyclic mechanisms for general demand cost-sharing games that arise from a variety of scheduling problems with *concave regular sum objectives*. This class of objective functions includes the classic total weighted completion time objective. We consider scheduling cost-sharing games in single machine (Theorem 4.3), identical parallel machine (Theorem 4.7), and concurrent open shop environments (Theorem 4.11). We accomplish this by using the framework with a *list scheduling algorithm*—an algorithm that schedules according to a permutation of the jobs—for each of the underlying scheduling problems. The budget-balance and efficiency guarantees dictated by the framework hold for these mechanisms, as long as the list scheduling algorithm used is *compatible* with the customers' service levels: that is, as long as the customers require their jobs be processed in the same order as the list scheduling algorithm.

2. Preliminaries on general demand cost-sharing games

In this section, we give a brief introduction to general demand cost-sharing games and cost-sharing mechanisms as studied and presented by Mehta et al. (2009).

Consider a setting with a service provider and a universe $U = \{1,2,\ldots,n\}$ of players. Every player $i \in U$ is interested in a set of services $\{1,2,\ldots,R_i\}$ where R_i is the publicly known maximum service level of player i. These services are ordered so that player i has to obtain all the services $\{1,2,\ldots,j-1\}$ before obtaining service j. Hence, the set of services can also be thought of as levels of service. For the sake of compactness, if player i is served service levels $\{1,2,\ldots,s_i\}$ for some $0 \leqslant s_i \leqslant R_i$, we say s_i is the *service level* of player i.

An *allocation* $S = (s_1, s_2, \ldots, s_n) \in \mathbb{Z}^U_{\geqslant 0}$ describes the level of service offered, or allocated, to each player: in allocation S, s_i is the service level allocated to player $i \in U$. The minimal allocation is $\emptyset = (0,0,\ldots,0)$ and the maximal allocation is $R_{\max} = (R_1,R_2,\ldots,R_n)$. Alternatively, an allocation $S = (s_1,s_2,\ldots,s_n)$ can be viewed as a closed set of (player, service level) pairs $S = \bigcup_{i=1}^n \{(i,j) : j = 1,2,\ldots,s_i\}$. [1] In this work, we use both notions of an allocation interchangeably. We denote the set of all allocations by $\mathcal{A}$.

[1] A set P of pairs of positive integers is *closed* if $(i,j) \in P$ implies $(i,j-1) \in P$.

A *cost function* $c : \mathcal{A} \to \mathbb{R}$ describes the cost of providing service: $c(S)$ is the cost of providing the service levels in allocation S. By assumption, $c(\emptyset) = 0$, and $c(S)$ is nondecreasing in every component of S. In this work, we assume that the cost $c(S)$ for an allocation S is the value of a minimum-cost solution to an *underlying optimization problem* that models the service provider's problem of providing S.

Each player $i \in U$ has a private type $V_i = (v_i(1), v_i(2), \ldots, v_i(R_i))$ called a *valuation*. The value $v_i(j)$ is player i's marginal valuation of service level j; that is, the amount player i is willing to pay for service level j after receiving service levels $1,2,\ldots,j-1$. By assumption, $v_i(j)$ is nonincreasing in j. Player i's total valuation for service level j is $\sum_{k=1}^j v_i(k)$. In addition, each player $i \in U$ announces a *bid* $B_i = (b_i(1), b_i(2), \ldots, b_i(R_i))$. The value $b_i(j)$ is player i's announced marginal bid for service level j, that is, the amount player i announces it is willing to pay for service level j after receiving service levels $1,2,\ldots,j-1$. Player i's total bid for service level j is $\sum_{k=1}^j b_i(k)$. By assumption, $b_i(j)$ is also nonincreasing in j.

Before we move on, a note about the assumptions made here. Note we assume that $v_i(j)$ is nonincreasing in j for any player $i \in U$; in other words, each player derives nonincreasing marginal value for each additional level of service it obtains. This is the common "diminishing marginal utility" assumption in economics. This assumption is reasonable in many contexts. For example, consider a general demand cost-sharing game in which the service levels correspond to the number of connections made between a player and a server. In this case, the first connection is arguably the most important and valued most highly, while the subsequent connections are useful but redundant, and hence valued less. The same reasoning can be applied to the assumption that $b_i(j)$ is nonincreasing in j for each player $i \in U$; that is, each player's marginal bids are nonincreasing in the level of service.

A *cost-sharing mechanism* collects a bid B_i from each player $i \in U$, and determines an allocation $S = (s_1, s_2, \ldots, s_n) \in \mathcal{A}$ to serve and a price q_i for each player $i \in U$. We restrict our attention to mechanisms that satisfy the following standard assumptions: (1) *individual rationality*, meaning that for every player $i \in U, q_i = 0$ if $s_i = 0$ and $q_i \leqslant \sum_{j=1}^{s_i} b_i(j)$ if $s_i > 0$; and (2) *no positive transfers*, meaning that prices are always nonnegative. We also make the standard assumption that players maximize quasilinear utilities; in other words, each player $i \in U$ aims to maximize $u_i(S, q_i) = \sum_{j=1}^{s_i} v_i(j) - q_i$.

Although the primary role of a cost-sharing mechanism is to select an allocation to be served, we follow recent work and provide cost-sharing mechanisms that also produce a feasible way of serving the chosen allocation. In particular, all of the mechanisms reported in this work produce a feasible solution to the underlying optimization problem for providing allocation S. The cost of this feasible solution is denoted by $c_M(S)$, and is permitted to exceed the optimal cost $c(S)$. It is necessary to allow sub-optimal solutions in order to implement mechanisms efficiently. For example, if the underlying optimization problem is NP-hard, then computing the optimal cost $c(S)$ for a given allocation S cannot be accomplished in polynomial time, unless $P = NP$.

A mechanism can satisfy different notions of truthfulness. In this work, we focus on the notion of *weak group strategyproofness* (Devanur et al., 2005). A mechanism is said to be weakly group strategyproof if no coordinated false bid by a subset of players can ever strictly increase the utility of every one of its members. Thus, in a weakly group strategyproof mechanism, every defecting coalition has at least one indifferent member.

We evaluate a cost-sharing mechanism against two types of metrics: (1) revenue and (2) economic efficiency. A mechanism is said to be *β-budget-balanced* for some $\beta \geqslant 1$ if

$$c_M(S) \leqslant \sum_{i=1}^n q_i \leqslant \beta c(S)$$