European Journal of Operational Research 248 (2016) 192–203

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Design of automated negotiation mechanisms for decentralized heterogeneous machine scheduling

Fabian Lang[a,*], Andreas Fink[a], Tobias Brandt[b]

[a] Institute of Computer Science, Helmut Schmidt University Hamburg, Holstenhofweg 85, 22043 Hamburg, Germany
[b] Information Systems Research, Albert Ludwigs University Freiburg, Platz der Alten Synagoge, 79098 Freiburg i. Br., Germany

A R T I C L E I N F O

Article history:
Received 3 April 2014
Accepted 24 June 2015
Available online 2 July 2015

Keywords:
Group decisions and negotiations
Scheduling
Negotiation protocol
Mechanism design
Complex contracts

A B S T R A C T

We consider a hard decentralized scheduling problem with heterogeneous machines and competing job sets that belong to different self-interested stakeholders (agents). The determination of a beneficial solution, i.e., a respective contract in terms of a common schedule, is particularly difficult due to information asymmetry and self-interested behavior of the involved agents. The agents intend to minimize their individual costs that consist of tardiness cost and their share of the machine operating cost. The aim of this study is to find socially beneficial outcomes by means of negotiation mechanisms that comply with decentralized information and conflicting interests. For this purpose, we present an automated negotiation protocol, which is inspired by metaheuristics, along with a set of optional building blocks. In the protocol, new solutions are iteratively generated, as mutations of a single provisional contract, and proposed to the agents, while feasible rules with quotas restrict the acceptance decisions of the agents. The computational experiments show that the protocol—without central information and subject to strategic behavior—can achieve high quality solutions which are very close to results from centralized multi-criteria procedures. Particular building block configurations yield improved outcomes. Concluding, the considered scheduling problem enhances standard scheduling models by incorporating multiple stakeholders, nonlinear cost functions, and machine operating cost, whereas the presented negotiation approach contributes to the methodology and practice of collaborative decision making.

1. Introduction

Planning and scheduling face new challenges these days. As companies are often strongly integrated within supply networks, related decisions are not exclusively subject to a single company's own preferences but increasingly depend on the interaction with, e.g., suppliers, subcontractors, or customers (Dawande, Geismar, Hall, & Sriskandarajah, 2006). Consequently, decision making becomes a collaborative task. Since the stakeholders have their very own preferences and objectives, planning and scheduling must consider strategic behavior (Kersten & Mallory, 1990). For example, decision makers that compete for shared machine resources may not be willing to disclose private information or may provide biased information such as exaggerated needs to work toward higher individual profits to the detriment of the overall allocation efficiency (Klein, Faratin, Sayama, & Bar-Yam, 2003). Thus, revealed information can be incomplete or misleading and, hence, is a fruit of a poisonous tree for a central authority. As a consequence, a traditional centralized scheduling approach may not be effective in obtaining high quality solutions. Since Operational Research is generally concerned with analytical methods for decision making, the issue of collaborative decision making within a group of autonomous decision makers is highly relevant. In multi-party situations, there is usually no single fully-informed authority that is entitled to hierarchically allocate resources and determine the courses of action for the whole system. Thus, a decentralized group decision making procedure is needed which takes into account strategic behavior of self-interested parties and restricted availability of information.

In this paper, we focus on decentralized decision problems on the level of operations management (in particular scheduling problems), which are characterized by recurring decision tasks as well as the potential and need for a formalized and automated decision support procedure. Thus, it is appropriate to devise and apply an automated negotiation approach. At this, decision making entities are represented by autonomous non-cooperative software agents. In contrast to human negotiators, software agents can easily negotiate for mil-

* Corresponding author. Tel.: +494065412321.
E-mail address: fabian.lang@hsu-hamburg.de (F. Lang).

http://dx.doi.org/10.1016/j.ejor.2015.06.058
0377-2217/

lions of rounds to find a mutually agreed upon solution (also called contract). Respective negotiations may be regarded as decentralized search processes which aim to iteratively enhance a contract from a contract space (search space). A major challenge of negotiations are complex contract spaces, which are characterized by multiple interdependent issues that lead to many local optima as well as a vast nonlinear search space (Fink, 2004; Klein et al., 2003). A negotiation procedure is defined by a protocol that controls the observable actions of the agents and related interactions. Such procedures work without requiring the parties to reveal full information on preferences. The rules of the protocol should work toward finding a contract that achieves pursued criteria such as Pareto-efficiency and social welfare.

Hence, this paper introduces a novel decentralized scheduling problem that incorporates competing job sets that are connected to multiple self-interested agents, machine operating costs, particularly energy costs, as well as tardiness costs. Extending conventional scheduling assumptions, we argue that both cost functions may be nonlinear and decisions are to be made by finding a multilateral contract. The assumption of nonlinearity in tardiness costs rests upon several considerations: firstly, in some scenarios, tardiness may induce a base cost—e.g., due to missing a strict deadline. Secondly, it might be that a small tardiness can mostly be offset by countermeasures, but the more tardiness occurs the higher the cost are (increasing marginal cost). Thirdly, the opposite can occur; missing a due date induces an immediate high penalty but the cost converges to an upper bound (decreasing marginal cost)—e.g., a product cannot be delivered to a customer anymore. Furthermore, energy costs have drawn the interest of managers and researchers in recent years. For instance, while they have always constituted a large part of operational costs for energy-intensive industries, rising energy prices and increasing power consumption of hardware have turned energy into a major cost driver for providers of IT infrastructure as well (Chen et al., 2005). The fact that energy consumption does not necessarily increase linearly with workload adds an additional computational challenge to scheduling approaches that include energy costs (Bodenstein, Schryen, & Neumann, 2012), as do dynamic tariff schemes that compensate industrial customers for reducing their consumption during peak times (Braithwait & Hansen, 2011; Mohsenian-Rad, Wong, Jatskevich, Schober, & Leon-Garcia, 2010). We include these determinants of energy costs into our negotiation scheme to provide a realistic representation of total expenditures. Concerning the nonlinearity of machine operating costs, empirical work has shown that, e.g., computing components' energy consumption is nonlinearly increasing or decreasing subject to their utilization (Bodenstein et al., 2012). Similarly, vessels' fuel consumption follows an approximate cubic function subject to their speed (Tierney, Áskelsdóttir, Jensen, & Pisinger, 2014). Furthermore, machine pools tend to be heterogeneous and of mixed efficiency, as broken machines are replaced by more advanced ones, while other older machines are still in place (Heath, Diniz, Carrera, Meira, & Bianchini, 2005).

The aim of this paper is twofold. Firstly, we want to advance the knowledge on generic decentralized negotiation procedures for complex contract spaces. For this purpose we describe, enhance, and analyze negotiation mechanisms with different building blocks for achieving beneficial outcomes. Secondly, we introduce and solve the considered multi-agent machine scheduling problem by means of a negotiation protocol with problem-specific operators. For this purpose, we incorporate characteristics of restricted information availability and strategic behavior of autonomous agents in a decentralized scheduling situation. The findings from this study contribute to the methodology and practice of decision support for hard decentralized scheduling problems with information asymmetry and multiple self-interested agents. For such kinds of problems, for the first time, we devise a rich framework of mainly generic negotiation mechanisms which is applied for challenging problem instances with up to 19 agents.

The remainder of this paper is structured as follows: After this introduction, we formally define the considered multi-agent scheduling problem and give illustrative examples for applications. Afterward, we give an overview of related work and introduce the devised negotiation-based solution mechanism. The protocol is evaluated in a computational study which is subsequently presented together with a discussion of results. Finally, we conclude the paper and present future work.

2. Problem

2.1. Problem definition

A set of non-preemptable jobs $\mathcal{J} = \{1, \dots, j, \dots, J\}$ which originate from a set of competing agents $\mathcal{I} = \{1, \dots, i, \dots, I\}$ has to be scheduled on a set of machines $\mathcal{M} = \{1, \dots, m, \dots, M\}$. Each job j has an associated agent a_j ($\in \mathcal{I}$), a standardized processing time p^s_j, a resource consumption per time slot r_j (in terms of a single resource type), a release time s_j, and a due date d_j ($p^s_j, r_j, s_j, d_j \geq 0$). Without loss of generality, we assume a discrete planning horizon with $\mathcal{T} = \{1, \dots, t_{\max}\}$ as set of all time slots and p^s_j, s_j, and d_j as nonnegative integer values. We assume that p^s_j, r_j, and s_j are publicly known, whereas d_j is just known to agent a_j.

The M machines are heterogeneous and have three relevant characteristics: (1) a capacity cap_m, which is the maximal resource provision by a machine m per time slot, (2) an operating speed os_m, which determines the speed of job processing, and (3) an operating cost function $E_{t,m}$, which determines the cost at time t subject to the machine's utilization $u_{t,m}$ (as described below). All machines have to fulfill the capacity constraint which is $0 \leq u_{t,m} \leq 1, \forall t \in \mathcal{T}, \forall m \in \mathcal{M}$. The machines' parameters are public information, i.e., all parties are aware of them.

The key decision variable of the problem is the schedule π which determines the start time (σ_j) and assigned machine (μ_j) for each job:

$$\pi = \{(\sigma_1, \mu_1), \dots, (\sigma_J, \mu_J)\}. \tag{1}$$

The time $p_j(\pi)$ required for the processing of a job j on a machine $m = \mu_j$ is determined by a machine's operating speed os_m, the job's standardized processing time p^s_j, and a standardized operating speed $\overline{\text{os}}$: $p_j(\pi) = \lceil \frac{\overline{\text{os}}}{\text{os}_{\mu_j}} * p^s_j \rceil$ (if $\text{os}_{\mu_j} = \overline{\text{os}}$, then $p_j(\pi) = p^s_j$). The completion time f_j of a job j is determined by its start and processing time: $f_j = \sigma_j + p_j(\pi)$. As the problem includes release times, the start time must not be earlier than the release time of a job, that is $\sigma_j \geq s_j, \forall j \in \mathcal{J}$.

The objective of an agent i is the minimization of his or her total cost which consists of two components: machine operating costs and tardiness costs. The agents use the same measuring commodity (numéraire; e.g., monetary units).

The operating cost $E_{t,m}$ for a given machine m at a given time slot t is subject to the three non-negative parameters $\alpha^E_m, \beta^E_m, \gamma^E_m$, the machine utilization $u_{t,m}(\pi) = \frac{\sum_{k \in \mathcal{J}: \mu_k = m \wedge \sigma_k \leq t < f_k} r_k}{\text{cap}_m}$, and a tariff $\Gamma(t)$: $E_{t,m}(u_{t,m}) = [\alpha^E_m * (u_{t,m})^{\beta^E_m} + \gamma^E_m * \Theta(u_{t,m})] * \Gamma(t)$, with Θ as the Heaviside function: $\Theta(\bullet) = 1 : \bullet > 0; \Theta(\bullet) = 0 : \bullet \leq 0$.

The overall operating costs of the machines are given by

$$\text{EC}(\pi) = \sum_{t \in \mathcal{T}} \sum_{m \in \mathcal{M}} E_{t,m}(u_{t,m}). \tag{2}$$

The overall operating costs EC have to be split among the set of agents $\mathcal{I}$, i.e., each agent i has to cover a share EC_i with $\sum_{i \in \mathcal{I}} \text{EC}_i = \text{EC}$. Different cost allocation schemes are discussed later on.

The tardiness T_j of a job j depends on the completion time f_j and the due date d_j: $T_j = \max\{f_j - d_j; 0\}$. Only if a job is tardy, there arises a tardiness cost $\text{TC}_j(T_j)$. This function $\text{TC}_j(T_j)$ can be nonlinear and is represented by $\text{TC}_j(T_j) = \alpha^w_j * (T_j)^{\beta^w_j} + \gamma^w_j * \Theta(T_j)$, with