European Journal of Operational Research 240 (2015) 476–487

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Revenue management for operations with urgent orders

Philippe Chevalier [a,*], Alejandro Lamas [a], Liang Lu [b], Tanja Mlinar [a]

[a] Louvain School of Management – Center for Operations Research and Econometrics, Place des Doyens 1, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium
[b] Logistics Research Centre, School of Management and Languages, Heriot-Watt University, Edinburgh EH14 4AS, UK

ARTICLE INFO

Article history:
Received 6 September 2013
Accepted 9 July 2014
Available online 23 July 2014

Keywords:
Revenue management
Order acceptance
Markov Decision Process
Heuristics
Flexibility

ABSTRACT

This article is motivated by the case of a company manufacturing industrial equipment that faces two types of demand: on the one hand there are the so-called regular orders for installations or refurbishing of existing facilities, these orders have a relatively long lead time; on the other hand there are urgent orders mostly related to spare parts when a facility has a breakdown, the delay in such case is much shorter but higher margins can be obtained. We study the order acceptance problem for a firm that serves two classes of demand over an infinite horizon. The firm has to decide whether to accept a regular order (or equivalently how much capacity to set aside for urgent orders) in order to maximize its profit. We formulate this problem as a multi-dimensional Markovian Decision Process (MDP). We propose a family of approximate formulations to reduce the dimension of the state space via aggregation. We show how our approach can be used to compute bounds on the profit associated with the optimal order acceptance policy. Finally, we show that the value of revenue management is commensurate with the operational flexibility of the firm.

1. Introduction

This article is motivated by the case of a cast iron manufacturer. This company is specialized in the production of cast iron pieces for industrial equipments. Most of its orders are either for the preventive maintenance of installations or for the building of new facilities. Such projects are scheduled with long lead times, but it is very important that the pieces be delivered on time because the plant where the pieces have to be installed will have to be (at least partially) stopped for the maintenance or installation activities to take place and obviously the duration of such stoppage should be minimized. A different type of orders received by the company corresponds to corrective maintenance when a breakdown occurs in a plant, in those cases a new cast iron piece is needed to restart the facility. Given that production at the customer is stopped because of the breakdown, a much faster service is required but the company can charge higher prices for such "emergency" orders. Moreover, the bargaining power of the customer is much weaker in such circumstances. Given its finite production capacity, the company cannot always accept all orders and should sometimes forego a regular order in order to keep some possibility to accept an urgent order later on. This dilemma is faced by many suppliers confronted with urgent requests that are potentially very profitable but could be very disruptive if not taken into account in their planning. Typical examples from the service sector include heating ventilation air-conditioning companies. The installation of new systems is typically a large project with a relatively long lead time. In contrast, when a system fails it could block the operations of a customer that is then willing to pay a higher fee for speedy action. In a very different context, suppliers of the fashion industry are known to combine orders from large vendors and more profitable orders for high fashion clothes which often require fast delivery (see e.g. deB Harris & Pinder, 1995; Barut & Sridharan, 2005). The question facing the supplier is how much capacity should be set aside for the urgent high margin demand, given the inherent unpredictability of this type of orders.

To address this question, we build a model with a supplier that handles two demand classes, that we will refer to as regular and urgent respectively. The regular orders are typically characterized by longer processing times, longer lead times but lower margins, while the urgent orders have shorter processing times, shorter lead times and higher margins. If the supplier accepts orders without foresight it is likely that at some point, when an urgent order arrives, the supplier will be unable to accept this order as her short term capacity is already entirely committed for regular orders (that were booked earlier with a longer lead time). Given the difference in margin between the two classes, this situation causes some loss

---

[*] Corresponding author. Tel.: +32 10 478396; fax: +32 10 474301.
E-mail addresses: philippe.chevalier@uclouvain.be (P. Chevalier), alejandro.lamas@uclouvain.be (A. Lamas), nicowish@gmail.com (L. Lu), tanja.mlinar@uclouvain.be (T. Mlinar).

http://dx.doi.org/10.1016/j.ejor.2014.07.015

of revenue. On the other hand, rejecting a regular order in anticipation for potential urgent orders that do not materialize, also causes some revenue loss.

This trade-off has clear similarities with other revenue management problems. The distinct feature is that when an order is accepted the supplier keeps some flexibility. For example, if a regular order necessitates 10 days of work and the lead time is 20 days, in most cases the customer does not care about the days during which the order is effectively produced as long as it is finished on time. In revenue management terms, if we consider that the capacity available during each period is a distinct *product*, an order requiring more than one period of work is in fact reserving several *products*. But the supplier has some flexibility in assigning the *products* to the order and does not need to make a commitment at the time of reservation. Such consideration is reasonable and represent the reality at many firms in which jobs can be interrupted and resumed in later periods. One can draw a parallelism with the network revenue management problem but where the supplier can accept a reservation without committing to specific legs in the network, the only commitment is on the origin and destination points in the network.

Our contributions are summarized as follows:

- We derive a Markov Decision Process (MDP) formulation of the order acceptance problem of two customer classes with lead time constraints. Likewise for the network revenue management problem, the size of this formulation quickly excludes the possibility of solving it exactly for larger size instances. We develop a family of approximate formulations parameterizable to range from a coarse approximation to the original full formulation. This makes it possible to choose between speed of solution and precision of result. What is of particular interest is that for each formulation, we can compute an upper and lower bound on the exact result. This last feature is rather uncommon for revenue management problems and is particularly interesting to make sure the adequate level of approximation is chosen in the proposed family of formulations.
- Through a numerical study we show that the proposed heuristics allow to obtain near-optimal solutions in a tractable time. We also show how the potential benefit of revenue management is commensurate with operational flexibility. In our setting operational flexibility consists in the slack between the promised lead time for an order class and the processing time needed for such order.

The remainder of the paper is structured as follows. In Section 2 we discuss related work. In Section 3 we provide a detailed description of our problem. In Section 4 we introduce an MDP formulation and discuss its resolution to obtain the optimal admission policy. In Section 5 we propose two heuristic formulations of the problem based on different levels of state aggregation and report on a numerical study of our proposed formulations in Section 6. Section 7 investigates the impact of operational flexibility on the benefit of revenue management. Finally, Section 8 summarizes the main conclusions and identifies future research directions.

2. Literature review

Our work belongs to the growing literature of Perishable Asset Revenue Management (PARM) which deals with the problem of allocation of scarce resources to different demand classes. Talluri and van Ryzin (2004) give a comprehensive overview of this topic. The first applications were for the airline industry by Littlewood (1972), and extended by Belobaba (1987), Wollmer (1992), and Brumelle and McGill (1993). In addition to airlines, typical service applications are in hotel management and car rental (Kimes, 1989; Bertsimas & Popescu, 2003; Talluri & van Ryzin, 2004; Bitran & Mondschein, 1995; Geraghty & Johnson, 1997). Gradually, new applications appeared for very different environments such as: MTO manufacturing (Balakrishnan, Sridharan, & Patterson, 1996; Barut & Sridharan, 2005; Spengler, Rehkoopf, & Thomas, 2007), project management (Herbots, Herroelen, & Leus, 2007; Herbots, Herroelen, & Leus, 2010) and health care (Gupta & Wang, 2008; Dobson, Hasija, & Pinker, 2011).

A stream of literature related to lead time decisions focuses on due-date quotation and scheduling problems in order to allocate the available capacity to incoming orders (see e.g. Kaminsky & Hochbaum, 2004; Keskinocak & Tayur, 2004, for an extensive literature review). Most of these works assign dynamically lead times to incoming orders depending on the state of the system and sequencing policies (see e.g. Duenyas, 1995; Duenyas & Hopp, 1995; Kapuscinski & Tayur, 2007). Kapuscinski and Tayur (2007) propose a dynamic programming approach to address the problem of lead-time quotation for two demand classes when customers are not equally sensitive to waiting. Lead time quotation is used to ensure that the capacity is allocated in such a way that all demands can be delivered on time. So, firms can change the quoted lead time based on the system state. Motivated by the prevalence of static lead time policies (see e.g. Cheng & Gupta, 1989; Hopp & Sturgis, 2000; Keskinocak & Tayur, 2004), we consider a different problem in which lead times are constant and exogenously given. Since the capacity may not be enough to cater all demands, the decision becomes accepting or rejecting orders depending on the system state.

Gupta and Wang (2007) consider an order acceptance problem in which the lead time requirement for regular orders is modeled as a soft operational constraint. It is assumed that tardiness cost is incurred if regular orders are not filled within their lead time window while urgent orders must be filled in the current period once accepted. The authors propose a multi-dimensional MDP whose optimal solution turns out to be a threshold based policy. This solution property is a consequence of the simplicity of their model setting, which leads to a well-structured value function. In contrast, our model is more general, assuming some flexibility in catering urgent orders – the lead time for urgent order does not necessarily need to be one. The state space in our problem is defined in a different way than in Gupta and Wang (2007), because their representation involves tracking the backlogging information for every demand class and becomes particularly inefficient when there are multiple demand classes, which greatly limits its application. Our representation “encodes” in itself the capacity allocation decisions and therefore is more efficient.

The following references focus on acceptance decision problems where the lead times must be strictly respected, as in our case. Germs and Van Foreest (2011) study an order acceptance problem with multiple customer classes with a common lead time, setup times and scheduling constraints. The problem is modeled as a Markov chain controlled by a threshold policy. The authors provide a numerical study for small instances which are computationally tractable. In contrast to their work, we provide efficient alternative methods to treat the state space explosion. Barut and Sridharan (2005) study an order acceptance problem involving multiple demand classes that differ in terms of price, lead time and demand pattern. The authors propose a nested rationing policy which fulfills incoming orders as much as possible while preserving a certain level of capacity for more profitable future orders. The proposed policy is computed using a myopic heuristic method that does not take the evolution of the capacity into account. Consequently, the efficiency of the heuristic is hurt by the simplified estimation of the future available capacity. Our formulation keeps track more