European Journal of Operational Research 218 (2012) 239–250

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Approximate dynamic programming for capacity allocation in the service industry

Hans-Jörg Schütz, Rainer Kolisch *

Technische Universität München, TUM School of Management Arcisstr. 21, D-80333 München, Germany

ARTICLE INFO

Article history:
Received 14 June 2010
Accepted 5 September 2011
Available online 4 November 2011

Keywords:
Capacity allocation
Services
Health care operations
Approximate dynamic programming
Reinforcement learning
Semi-Markov decision process

ABSTRACT

We consider a problem where different classes of customers can book different types of service in advance and the service company has to respond immediately to the booking request confirming or rejecting it. The objective of the service company is to maximize profit made of class-type specific revenues, refunds for cancellations or no-shows as well as cost of overtime. For the calculation of the latter, information on the underlying appointment schedule is required. In contrast to most models in the literature we assume that the service time of clients is stochastic and that clients might be unpunctual. Throughout the paper we will relate the problem to capacity allocation in radiology services. The problem is modeled as a continuous-time Markov decision process and solved using simulation-based approximate dynamic programming (ADP) combined with a discrete event simulation of the service period. We employ an adapted heuristic ADP algorithm from the literature and investigate on the benefits of applying ADP to this type of problem. First, we study a simplified problem with deterministic service times and punctual arrival of clients and compare the solution from the ADP algorithm to the optimal solution. We find that the heuristic ADP algorithm performs very well in terms of objective function value, solution time, and memory requirements. Second, we study the problem with stochastic service times and unpunctuality. It is then shown that the resulting policy constitutes a large improvement over an "optimal" policy that is deduced using restrictive, simplifying assumptions.

1. Introduction

In many service industries, such as legal, financial, and health services, customers are serviced in sequence using a limited resource, e.g. a lawyer, a bank counselor or a diagnostic device. We consider the problem of allocating the scarce capacity of this resource to stochastic demand of different customer classes and service types when there are cancellations, no-shows, and overbooking. In addition, the service times can be stochastic and customers might arrive tardy. We assume that the service provider has to decide instantly if a request from a customer is accepted (i.e., capacity is allocated) or rejected. As demand from customers with high priority might arrive after low priority demand, the allocation decision has to be based both on the capacity that is still available and on the potential future arrival of customers.

We model this problem as a continuous-time Markov decision process (CTMDP) and solve it using simulation-based approximate dynamic programming (ADP) combined with a discrete event simulation of the service period. ADP is also known as reinforcement learning (see [39]). The main advantage of this heuristic algorithm is that it can be applied to complex, realistic models with a very large state space. In order to determine optimal solutions the model has to be simplified using restrictive assumptions so that standard dynamic programming can be applied.

More specifically, we consider the following problem. Capacity is given by the total time the resource operates a day (e.g. 9 hours) and the service times vary according to the class-type combination (e.g. between 10 and 80 minutes). For the capacity of a specific day the booking requests arrive during a so-called booking horizon. The latter is defined by the time span between its opening time and its closure time where the latter coincides with the day of service. The length of the booking horizon depends on the industry. While it is up to one year in the airline industry, one month is common for radiology services in health care. Furthermore, there is a last minute (same-day) demand immediately before the start of service. Each booking request is defined by a class-type combination and the amount of capacity required. In case of acceptance, a class-type specific revenue is obtained while in case of rejection a class-type specific penalty cost might accrue. Customers, whose booking requests have been accepted, might cancel their appointments before the day of service or do not show up at the day of service. In this case the customers might get refunds. To counter the negative effects of cancellations and no-shows, the service provider can overbook available capacity. At the day of service the customers may arrive tardy and service times are stochastic. The objective of the service provider is to maximize the expected profit obtained from the limited capacity. We will relate to a profit maximizing decision maker noting that we can use a utility

* Corresponding author. Tel.: +49 89 28925161; fax: +49 89 2 8925166.
E-mail addresses: rainer.kolisch@wi.tum.de, rainer.kolisch@web.de (R. Kolisch).

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.09.007

function in the case of nonprofit institutions (see [28]). A specific feature of the problem is that capacity is allocated to demand in an aggregated fashion (i.e., a day) but that the service will be supplied in a sequential fashion where each accepted customer is served during a specific time slot within the day. Because of no-shows, the possibility of overbooking, tardy arrivals, and stochastic service times this feature requires decisions on the underlying appointment schedule in order to calculate the expected amount of overtime and customer waiting time which results from a given number of accepted customers. The decision on the capacity allocation has to be linked with information on the appointment schedule in order to derive good decisions.

The objective of this paper is to propose a new model for the complex problem stated above. In contrast to similar models in the literature that also take the future arrival of customers into account, it includes a model for sequential delivery of appointment based services with multiple types of service, stochastic service times, and tardiness of clients. In our view this is an important extension as these features can be found in many applications. Employing the ADP algorithm of Gosavi et al. [19] that is adapted to our problem and combining it with a simulation of the service period we investigate on the benefits of ADP on capacity allocation in the service industry. First, we study a simplified problem with deterministic service time and punctual arrival of clients and compare the solution from the ADP algorithm to the optimal solution. We find that the heuristic ADP algorithm performs very well in terms of objective function, solution time, and memory requirements. Second, we study the problem with stochastic service times and tardiness. Because of the good performance results of the ADP algorithm that are obtained using the simplified problem, we believe that the policy we obtain is close to optimal. It is then shown that this policy constitutes a large improvement over an "optimal" policy that is deduced using restrictive, simplifying assumptions.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. In Section 3 the booking process is modeled as a continuous-time Markov decision process and the simulation model of the service period is presented. The approximate dynamic programming algorithm that is used to solve the problem is presented in Section 4. The results of the experimental investigation that is based on data from a university hospital are shown in Section 5. Section 6 concludes the paper.

2. Literature review

This section is divided into two parts. First, we give an overview on previous work on capacity allocation and appointment scheduling. Second, we review literature in the field of approximate dynamic programming.

Capacity allocation and appointment scheduling. The problem of scheduling clients to appointments can be divided into three problems: capacity allocation, appointment scheduling, and short-term decisions on the day of service.

Capacity allocation or advanced scheduling decides how many customers should be admitted into the service and how the available capacity should be divided among different types of customers. Most authors focus on applications in health care. Among those Gerchak et al. [16], Gupta and Wang [22], Patrick et al. [30], Schütz and Kolisch [34], Liu et al. [27], Erdelyi and Topaloglu [15], as well as Ayvaz and Huh [3] all model their problems as a Markov decision process (MDP).

Capacity allocation or capacity control also plays an important part in the field of revenue management. A detailed review is provided by Talluri and van Ryzin [41]. Lee and Hersh [25] and Subramanian et al. [38] use dynamic programming techniques to solve capacity control problems for flights that are closely related to our problem but do not include appointment scheduling.

Appointment scheduling deals with the problem of assigning appointment times to accepted customers. A review and classification of the appointment scheduling literature in health care is provided by Cayirli and Veral [7] and Gupta and Denton [21]. For literature that deals with clinical overbooking because of patient no-shows we refer to the literature review in Zeng et al. [46].

Short-term decisions on the day of the examination, i.e., after the appointments have been made, consider which customer should be examined next if more than one is present. This problem is considered in Green et al. [20] as well as Kolisch and Sickinger [24]. In a subsequent article Sickinger and Kolisch [35] also determine the optimal number of outpatients to be scheduled and assign the outpatients to an appointment schedule. In this paper we do not consider short-term decisions.

Most of the papers cited above decouple the capacity allocation and the appointment scheduling problem. However, performance measures such as cost, resource utilization, patient waiting time, and overtime depend on both allocation and scheduling decisions. A number of papers combine these two problems. Conforti et al. [9], Muthuraman and Lawley [29], Chakraborty et al. [8], and Zeng et al. [46] make joint decisions on admission and appointment times. However, they employ myopic algorithms that do not block capacity for patients arriving in the future. Vermeuelen et al. [43] present a simulation study of high complexity but have to compare a large number of scenarios and parameter settings in order to derive good results.

In this paper we study a capacity allocation model that takes into account the future arrival of patients. The appointment scheduling problem is still decoupled. However, compared to other papers we employ a more detailed appointment scheduling model that allows for a realistic feedback of the effect of allocation decisions within the service period and can be adapted to different settings. In our view this is an important step towards a model that combines allocation and scheduling decisions while also considering the impact on patients arriving in the future.

The model presented in this paper is closely related to the work of Patrick et al. [30], Gupta and Wang [22], and Schütz and Kolisch [34]. Patrick et al. [30] address the problem of dividing the capacity of a CT scanner between multiple outpatient classes. Their objective is to minimize the cost associated with the number of patients who fail to meet the waiting time targets of their class. A multi-period problem is considered where "batch"-decisions are taken once a day, i.e., customer requests are not processed individually and immediately. The authors show that the policy for each outpatient class can be determined by solving an approximate dynamic programming problem to optimality. The problem treated in our research differs from the one of Patrick et al. [30] because we consider different customer classes and treatment types where we decide immediately on each request. In addition, Patrick et al. do not take into account cancellations and no-shows. Gupta and Wang [22] consider the decision of accepting or rejecting individual booking requests for a single day in a clinic with multiple physicians. Each physician works using fixed time slots over the course of the day. Patients have preferences for particular combinations of physicians and time slots. The clinic can reject advance requests but has to serve same-day demand working overtime if necessary. The problem is modeled as a finite horizon MDP and solved by stochastic dynamic programming. The paper considers patient choice, which we do not, while we consider different class-type combinations of requests as well as cancellations and no-shows. As in Gupta and Wang we make a decision on each individual booking request. Schütz and Kolisch [34] model the problem presented in this paper as an MDP and assume deterministic service times and punctual arrivals of clients as well as other restrictive assumptions. They