European Journal of Operational Research 236 (2014) 298–312

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Optimal control policies for ambulance diversion

Adrian Ramirez-Nafarrate [a], A. Baykal Hafizoglu [b], Esma S. Gel [b,*], John W. Fowler [b]

[a] Department of Industrial and Operations Engineering, Instituto Tecnologico Autonomo de Mexico, Mexico, DF 01080, Mexico
[b] School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85287, United States

ARTICLE INFO

Article history:
Received 19 November 2012
Accepted 16 November 2013
Available online 26 November 2013

Keywords:
Ambulance diversion
Emergency departments
Markov decision processes

ABSTRACT

Ambulance diversion (AD) is used by emergency departments (EDs) to relieve congestion by requesting ambulances to bypass the ED and transport patients to another facility. We study optimal AD control policies using a Markov Decision Process (MDP) formulation that minimizes the average time that patients wait beyond their recommended safety time threshold. The model assumes that patients can be treated in one of two treatment areas and that the distribution of the time to start treatment at the neighboring facility is known. Assuming Poisson arrivals and exponential times for the length of stay in the ED, we show that the optimal AD policy follows a threshold structure, and explore the behavior of optimal policies under different scenarios. We analyze the value of information on the time to start treatment in the neighboring hospital, and show that optimal policies depend strongly on the congestion experienced by the other facility. Simulation is used to compare the performance of the proposed MDP model to that of simple heuristics under more realistic assumptions. Results indicate that the MDP model performs significantly better than the tested heuristics under most cases. Finally, we discuss practical issues related to the implementation of the policies prescribed by the MDP.

1. Introduction

Media and academic papers have highlighted the prevalence of the overcrowding problem in emergency departments (EDs) in the United States (US) during recent years (Associated Press, 2006). One of the major negative impacts of congestion in EDs is the long time that patients have to wait before starting to receive treatment, resulting in seriously adverse events, including death (CNN U.S., 2008; News, 2010). The risk of such adverse events increases when the condition of the patient is severe and when waiting times extend beyond a Recommended Safety Time Threshold (RSTT), which is used by the Center for Disease Control and Prevention (CDC) based on patient severity (which is assessed by various indicators of the health condition of the patient such as vital signs and stability), and the amount of resources required. Reports from the United States General Accounting Office (2009) have highlighted the poor performance of EDs around the US where patients have to wait beyond their RSTT.

Researchers and managers of EDs have explored different strategies to reduce congestion and avoid potential implications of long wait times. Hoot and Aronsky (2008) cite *ambulance diversion* as one of the main approaches hospitals use for demand management in EDs. In particular, hospitals, at times of excessive congestion, may divert ambulances to other neighboring hospitals by requesting emergency medical services to bypass their facilities. This strategy is implemented relatively frequently in US hospitals. According to the United States General Accounting Office, in 2003, 25% or more of the hospitals in several US metropolitan areas were on diversion more than 10% of the time. During 2006, 27.3% of hospitals reported going on diversion, and the average number of hours on diversion during that year were 473 hours (United States General Accounting Office, 2009).

Although EDs often divert ambulances to tackle overcrowding, this approach can have negative consequences when AD policies are not properly designed. For instance, Yankovic, Glied, Green, and Grams (2010) and Shen and Hsia (2011) indicate that AD might increase mortality among patients with acute myocardial infarction that are transported by ambulance. Consequently, AD decisions should consider various factors such as the current congestion at the ED, severity of the patients, and the status of neighboring hospitals. For example, if a neighboring hospital is relatively near and currently less crowded, then it is more likely that an arriving patient in an ambulance can start receiving appropriate treatment earlier if he/she is diverted from an overcrowded facility. On the other hand, while ambulances can be diverted, EDs do not have control over walk-in arrivals, who, by law, have to be accepted and treated. Therefore, while on diversion, EDs still accept walk-in patients; these patients also contribute to congestion.

Empirical studies on AD focus on the effectiveness of various AD approaches, and provide some managerial recommendations to re-

* Corresponding author. Tel.: +1 480 965 2906; fax: +1 480 965 2751.
E-mail address: esma.gel@asu.edu (E.S. Gel).

http://dx.doi.org/10.1016/j.ejor.2013.11.018
0377-2217/

duce or minimize AD, such as restricting the number of hours spent on diversion by hospitals serving a specific geographic region. The implementation of these guidelines have resulted in significant decreases in the number of hours on AD in the regions of study; this includes San Diego and Sacramento, California (Asamoah, Weiss, Ernst, Richards, & Sklar, 2008; Patel, Derlet, Vinson, Williams, & Wills, 2006; Vilke et al., 2004). Unfortunately, these studies do not discuss the effect of avoiding AD on other performance measures, such as the average patient waiting time.

Analytical studies on AD, on the other hand, suggest that appropriate policies could significantly improve the logistical performance of an emergency care system. For example, Deo and Gurvich (2011) modeled the decisions of two EDs using game theoretic approaches with the objective of minimizing the average patient waiting time for each hospital in a system with two EDs. The authors found that a centralized design of diversion policies is Pareto-improving compared to a decentralized strategy, which leads to a defensive equilibrium. The authors also proposed a threshold-type AD policy, but they did not explore the optimality of this type of control policy. Using similar approaches, Hagtvedt, Griffin, Keskinocak, Ferguson, and Jones (2009) analyzed AD, and pointed out the need for a central agent that coordinates AD. Ramirez, Fowler, and Wu (2011) presented a simulation model of an emergency care delivery system to analyze the effectiveness of diversion and destination policies. They evaluated the use of an effective combination of diversion–destination policies as an ambulance flow control mechanism in order to reduce the average time spent in activities with inappropriate level of care, which includes transportation to ED, waiting, and boarding in the ED.

In this paper, we are interested in studying effective ambulance diversion control policies that optimize the time to patients' access to the needed healthcare (at either the considered hospital or at the neighboring hospitals), and address the following research questions: (i) Can optimal AD policies significantly increase the safety of patients by minimizing the time that patients wait beyond their RSTT? (ii) What is the structure of optimal AD policies? (iii) What are the impacts of patient traffic and severity mix on optimal AD decisions? (iv) What is the value of information of knowing severity level of ambulance patients, and about the time to start treatment in the neighboring hospital(s) on the performance of the optimal policy? and (v) How do policies applied in practice perform compared to optimal AD policies?

In particular, we explore optimal AD policies for EDs with two treatment areas: an emergent care area and a fast-track treatment area. This is a layout commonly found in EDs around the world (Combs, Chapman, & Bushby, 2006, 2007). We study AD policies characterized by a two-dimensional diversion trigger that accounts for crowding in both of these treatment areas. This structure allows us to explore the impact of various parameters (such as those pertaining to the frequency and severity of patients arriving by ambulance) on optimal ambulance diversion decisions, and study various inherent dependencies in such systems. To the best of our knowledge, this is the first paper that takes into account patient severity levels, regardless of arrival mode (i.e., walk-in or by ambulance), for constructing AD policies.

Even though admission control methods are commonly used in various manufacturing and service systems, the AD literature has not considered the use of such methods in the control of ambulance arrivals to date. Early studies on admission control typically focus on the control of a single customer class using $M/M/1$ queueing models (see Stidham (1985) for a survey). More recently, studies consider control of several demand classes requiring different levels of service. Ha (1997) discusses an inventory control problem of N demand classes that incur different lost sales costs when customers are not admitted into the system. Similar to our setting, Carr and Duenyas (2000) discuss two demand classes, where one of the classes is always accepted into the system (similar to the walk-ins in our model), and the company has an option to reject the arrivals from the other class (similar to the ambulance arrivals in our model). Gupta and Wang (2007) consider one contracted demand class whose orders are always accepted, and one transactional demand class whose orders can be rejected. Similarly, Feng and Pang (2010) consider a long-term contract market whose orders are always accepted, and the spot market whose orders may be subject to rejection. In the recent work of Chen, Chen, Parlar, and Xiao (2011), the authors discuss the admission control problem of the orders coming from an online retailer. All of above discussed studies control demand using accept/reject decisions, similar to our accept/divert decisions.

In particular, this paper contributes to the existing literature on AD by presenting a Markov Decision Process (MDP) formulation to obtain optimal AD control policies for a hospital. The objective is to minimize the long-run average expected tardiness per patient, where tardiness is defined as the length of time that a patient waits beyond his/her RSTT, before starting to receive treatment. Hence, we assume that the hospital sets the ambulance diversion policy to optimize the patient welfare (rather than a profit or cost measure) by considering the patient's options of being treated at their facility versus at a neighboring hospital. We then explore the effectiveness of the obtained policies with respect to other performance measures of interest to the hospital management, such as the average fraction of time on diversion.

Assuming Poisson arrivals, exponential length of stay in the ED and two patient severity levels, we analyze the structure of optimal policies using both theoretical and computational analysis. Using theoretical analysis, we show that the optimal diversion policy can be characterized by a threshold curve. Using computational analysis, we further study the structure of optimal AD policies by observing the impact of (i) patient arrival rates, (ii) the severity mix of the patient population, and (iii) the "amount" of available information on the time to start treatment at the neighboring hospital(s). We next present a simulation study, where various modeling assumptions are relaxed to represent more realistic scenarios, and compare the performance of optimal policies with that of other simpler policies used in practice, such as not diverting at all and diverting only when there are no available beds. Computational analysis verifies the superior performance of the optimal policies obtained using our MDP model. In addition, a simple policy that diverts ambulances when there are no available beds for critical patients is shown to yield somewhat satisfactory results in certain cases. We finally discuss possible drawbacks of our approach in practice, and provide some recommendations to overcome these.

This study has two main contributions to the health-care literature. First, to the best of our knowledge, it is the first paper discussing optimal control of AD using an MDP formulation. Second, it considers a novel objective that minimizes the time that patients wait beyond a RSTT before starting to receive treatment. Although the AD literature includes various studies that discuss the minimization of time spent in ED, this objective does not take into account the severity of more critical patients, whose treatment delays may result in death. Since the RSTT depends on the severity level of patients, our objective considers the safety of the patient as a performance measure for AD policies, which is a significant measure to evaluate the effectiveness of AD policies according to Asplin (2003). In addition, because our objective is in time units, it does not require any cost parameterizations that have plagued the previous studies.

The remaining sections of the paper are organized as follows. Section 2 introduces the model. Section 3 presents some properties of an optimal AD policy. Section 4 analyzes the impact of the level of information on the time to start treatment in a neighboring hos-