Decision Support Systems 49 (2010) 463–473

Contents lists available at ScienceDirect

Decision Support Systems

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

Improving patient flow in a hospital through dynamic allocation of cardiac diagnostic testing time slots

Robert W. Day [a,1], Matthew D. Dean [b,*], Robert Garfinkel [a], Steven Thompson [c]

[a] School of Business, University of Connecticut, 2100 Hillside Road, Storrs, CT 06269, United States
[b] College of Business Administration, University of New Orleans, 2000 Lakeshore Drive, New Orleans, LA 70148, United States
[c] Robins School of Business, University of Richmond, 28 Westhampton Way, Richmond, VA 23173, United States

ARTICLE INFO

Article history:
Received 2 July 2009
Received in revised form 4 May 2010
Accepted 13 May 2010
Available online 21 May 2010

Keywords:
Health care
Patient flow
Network flow
Mathematical modeling

ABSTRACT

A cardiac diagnostic testing center (CDTC) makes real-time scheduling decisions that impact the use of its resources and the availability of telemetry-equipped beds within a hospital. Both inpatients and outpatients are frequent users of CDTC resources, and physicians prescribe one of several single-phase or multiple-phase test protocols. This complex online decision-making environment is modeled as a finite-horizon, discrete-time Markov decision process (MDP), but the growth of the state space motivates the introduction of a fast heuristic for real-time decision support. We therefore introduce a dynamic network scheduling tool which is both more flexible and more robust, making it applicable to the various configurations that may be found in practically any CDTC. We evaluate this new method computationally using simulation, comparing it to both an MDP model for small instances, and to the existing operational practice at our partner hospital for more realistic sized problems.

1. Introduction

Many hospitals in the United States provide a wide range of inpatient and outpatient health care services to the communities they serve. Outpatient services tend to be more lucrative because they require less infrastructural investment and have lower ancillary support costs. The profits obtained from outpatient services are crucial, because they enable the hospital to offset possible losses on certain inpatient services and invest in new medical technologies. Thus, the units that provide outpatient services are often treated as profit centers, each of which strives to maximize revenue. Inpatient services, with the exception of some surgical services, are typically not highly lucrative. For these services, the hospital typically receives a lump sum payment determined by the patient's classification to a particular Diagnostic Related Group (DRG) at the time of admission. However inpatient care requires a great deal of infrastructure and support personnel to be constantly in place. This results in very high fixed costs, so that hospital units that provide inpatient services must operate at high levels of utilization in order to break even. The problem of optimally balancing the levels of inpatient versus outpatient services is complicated by the fact that the demand for inpatient services is variable. The need to maintain high levels of utilization and meet uncertain demand makes the management of inpatient services challenging.

While low demand that results in poor utilization is a concern, high demand and the associated congestion can also present a significant problem. For example, if a given inpatient unit in a hospital does not have enough beds to accommodate all patients in need, then some of them will be "held" in the emergency department (ED) until a bed becomes available. When the ED becomes too crowded, it can no longer function effectively. In the extreme case the hospital may have to deny admission to potential inpatients. Often this takes the form of diversion of ambulances to other hospitals, since this is a common manner of inpatient arrival. The hospital then loses the revenue that would have resulted from treating those patients that were diverted, and its reputation as a reliable provider is also damaged.

Some hospital facilities are commonly termed "outpatient units", although that is not entirely accurate, as these units often provide services to inpatients as well. One typical example is the cardiac diagnostic testing center (CDTC) at Windham Community Memorial Hospital (WCMH) in Willimantic, CT. Like the CDTC at most hospitals, its primary role is to provide a range of cardiac screening tests to outpatients. These screenings are used for preventative care as opposed to the diagnosis of urgent, potentially life threatening conditions. Patients are typically referred to the CDTC by their physicians, and schedule appointments for prescribed tests, well in advance.

* Corresponding author.
E-mail addresses: bob.day@business.uconn.edu (R.W. Day), mddean@uno.edu (M.D. Dean), rgarfinkel@business.uconn.edu (R. Garfinkel), sthomps3@richmond.edu (S. Thompson).
[1] Authors are listed in alphabetical order.

0167-9236/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.dss.2010.05.007

The other role of the CDTC is to perform cardiac screenings on inpatients. For them, the stress test is generally the last test prior to discharge because it is the most aggressive in terms of how much strain is placed on the patient's heart. When a patient is admitted with chest pain a number of tests are done that do not place any strain on the heart (blood tests, EKGs, etc.). The purpose of those tests is to find evidence that the heart muscle is not receiving enough blood flow, causing the patient to experience chest pain. When these tests are negative, the final test is a stress test to push the heart and see if any symptoms arise. If the heart behaves normally under exercise conditions, then it is safe to send the patient home. Only if symptoms appear during this final screening, which in practice is uncommon, must the patient remain hospitalized. At WCMH in particular, even if a stress test is not normal, patients still leave the hospital. Rather than going home they are typically transferred to a larger facility that is able to perform complex cardiovascular procedures such as angioplasty, coronary artery stenting, coronary artery by-pass grafting, etc.

Providing screenings to inpatients presents a challenge because, unlike those for outpatients, the requests for screening often come with little or no notice. If at the time of the request there is no immediate opening, a scheduling decision must be made. In essence, this involves choosing which among various classes of patients, some of whom are inpatients and others outpatients, to service at a given time period. A serious area of concern observed at WCMH prior to our study was that some outpatients do indeed balk, presumably choosing to be screened elsewhere, if they are rescheduled from their prearranged time slot. Further, in many cases physicians are affiliated with multiple hospitals in a given area, and as a result, hospitals face competition from other hospitals when the physicians feel that one hospital is providing better service. Hospitals can also face competition from the physicians that are affiliated with the hospital. For example, one of the cardiologists affiliated with WCMH made the decision a few years prior to our study to stop working with the CDTC and open her own screening clinic. It follows that, considering only the profit to the CDTC, part of an intuitively dominant decision rule would be to give priority to outpatients. Thus an ad hoc strategy used by many hospitals, including WCMH, is to give priority to outpatients except in some time slots that are reserved for inpatients.

On the other hand, the CDTC can play a key role in maintaining patient flow and helping the hospital avoid the negative consequences of an overcrowded ED. By providing an inpatient screening the CDTC effectively "opens a bed" because that patient can now almost always be discharged, either to home or to another hospital, after its completion. This provides the hospital with the ability to accommodate an additional patient, should one arrive. Note that any patient whose screening is delayed must also be assigned a subsequent time slot. Thus, because there are two potential classes of costs, we choose to follow many authors (e.g. [8]) and optimize a composite "social welfare" function of the profits to the hospital from providing patient services, costs to the hospital of lost opportunity, and "inconvenience" costs to the patients of delays in treatment. It should be stressed that the resulting algorithm is a decision support system, so that its users can certainly input any costs that they deem appropriate, including the possibility of not considering patient inconvenience costs by setting those values to zero.

The scheduling problem is complicated by the fact that there are different types of cardiac screenings. Some have multiple phases, some span a two-day period, and the phases in a multiple-phase test must be done in lockstep fashion over time. That is, once a screening begins, or resumes on its second day, subsequent steps in that day must be done in the immediately following time periods. The problem can be modeled as a finite-horizon, discrete-time Markov decision process (MDP) as described in Appendix A. However, one quickly finds that the computational methods based on this model tend to be impractically slow and memory intensive when compared to the pace of the actual decision making environment, motivating the introduction of a computationally tractable scheduling tool.

Therefore, we introduce a decision support system based on a Multi-commodity Flow Snapshot Model (MFSM), grounded in the theory of dynamic network optimization. This approach is both flexible and robust, making it applicable to the various configurations of patient and test-protocol types that may be found in practically any CDTC. Based on historical data from WCMH, our experiments indicate that the MFSM approach results in drastically improved quality of service for inpatients, with little or no reduction in the quality of service to outpatients relative to the existing policy at WCMH. The bottom line of the hospital will also show a healthy improvement. Our methods provide both an online decision support tool, and a framework by which simulation results may be used as guidance for more profitable long-term capacity planning.

The rest of the paper is organized as follows. The relevant related literature is reviewed in Section 2 and in Section 3 the problem is formally defined. In Section 4 we provide an overview of the MDP approach with details in Appendix A. The dynamic network model is described in Section 5. In Section 6 we use simulation and real data from WCMH to evaluate the results of the network model relative to the MDP and to the prior decision making policy utilized by WCMH. Conclusions, with managerial implications and directions for future research, are addressed in Section 7.

2. Literature review

Efficiency of health care operations has been widely studied in the management science literature, and detailed histories of the application of operations research models to problems in the health care industry can be found in [6] and [3]. A number of techniques have been employed to help improve patient flow including dynamic allocation of patients to the inpatient units [17], bedside registration [14], and applications of simulation and queueing theory to identify and remove patient flow bottlenecks [12].

The problem of allocating CDTC time slots shares some traits with a number of stochastic control applications. Allocating service capacity over time among several competing customer classes has been studied in various settings outside of health care, including hotel management [2,11], car rentals [4,7], and airline yield management [1]. These articles fall into the category of revenue management. One distinguishing feature between revenue management and our diagnostic allocation problem is that in many instances of revenue management, customers can choose which "fare class" they belong to, whereas in our setting, the screening class is determined by a physician.

The more closely related problem of appointment scheduling and real-time capacity allocation in an MRI setting was addressed by [8]. The use of approximate dynamic programming to solve the problem of dynamically allocating diagnostic imaging resources to multiple patient priority classes, in order to achieve targeted wait times was investigated in [15]. While [15] and [8] consider allocation of time slots among multiple patient classes, the problem addressed here is different for two reasons. First, we expand the state space definition to include the number of available inpatient beds for different services. This increases the size and complexity of the problem, but thereby more accurately captures the costs that CDTC decisions impose on the rest of the hospital. Second, we consider time slot allocation decisions for both multiple-phase and single-phase tests. The distinction is that a multiple-phase test involves performing a set of procedures at separate times. To the best of our knowledge a setting involving a mixture of single- and multiple-phase tests with a mixture of patient classes in the operations research health care literature, has not been investigated prior to this study. Our proposed approach to solve the problem uses a dynamic network flow model. A recent example