



Hospitalization admission control of emergency patients using markovian decision processes and discrete event simulation

Martin Prodel, Vincent Augusto, Xiaolan Xie

► To cite this version:

Martin Prodel, Vincent Augusto, Xiaolan Xie. Hospitalization admission control of emergency patients using markovian decision processes and discrete event simulation. 2014 Winter Simulation Conference, Dec 2014, Savannah, United States. pp.Pages 1433-1444. emse-01108932

HAL Id: emse-01108932

<https://hal-emse.ccsd.cnrs.fr/emse-01108932>

Submitted on 4 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HOSPITALIZATION ADMISSION CONTROL OF EMERGENCY PATIENTS USING MARKOVIAN DECISION PROCESSES AND DISCRETE EVENT SIMULATION

Martin Prodel
Vincent Augusto
Xiaolan Xie

Center for Biomedical and Healthcare Engineering
École Nationale Supérieure des Mines de Saint-Étienne
158 cours Fauriel
42023 Saint-Étienne CEDEX 2, FRANCE

ABSTRACT

This paper addresses the hospitalization admission control policies of patients from an emergency department that should be admitted shortly or transferred. When an emergency patient arrives, depending on his/her health condition, a physician may decide to hospitalize him/her in a specific department. Patient admission depends on the availability of beds, the length of stay (LOS) and the reward of hospitalization which are both patient-class specific. The problem consists in determining patient admission policies in order to maximize the overall gain. We first propose a Markov Decision Process (MDP) Model for determination of the optimal patient admission policy under some restrictive and necessary assumptions such as exponentially distributed LOS. A simulation model is then built to assess MDP admission policies under realistic conditions. We show that MDP policies significantly improve the overall gain for different types of facilities.

1 INTRODUCTION

Health-care facilities such as hospitals and clinics are facing multiple challenges: increasing demand due to an aging population, cost reduction pressure due to increasing health expenditures and decreasing government funding, and increasing expectations of patients. As a consequence, an increasing number of French hospitals, both public and private, are losing money. Budget control in health-care has become a key component in the organization of care, and will ultimately impact the quality of care. Health authorities have been cutting beds and hospital personnel. The delivery of care, the organization of the resources and the financial results of a facility are closely related. The current situation requires new policies to be implemented and pursued, which considers all those aspects.

Emergency Departments (ED) are one of the most complex parts of hospitals to manage, and yet a major entry point for patients. It deals with patients without an appointment and with a wide range of illnesses. Even if most patients arriving to an ED leave the hospital after having seen a physician at the ED, a significant part of them need to be hospitalized. In many hospitals, finding available beds for unscheduled patients is extremely complicated. Even if all patients arriving at the ED do not require the same level of care, many hospitals proceed with the following policy: accept any patient until no bed is available. However, more sophisticated policies, including bed booking strategies and dynamic decisions, can lead to significant improvement of overall hospital performance.

The aim of this paper is to develop methods to identify efficient hospitalization admission control policies for the ED patients. By efficient, we mean a hospital which is able to admit all the patients entering the ED without increasing its deficit. Motivated by observations of our partner hospital (CH-Firminy) with over 70% of inpatients admitted through the ED, we consider the hospitalization admission control of a

hospital with all patients admitted through the ED, i.e. patients are admitted immediately or transferred to other hospitals. Admission is mainly constrained by bed availability in the medical departments. Patients are grouped into many different groups, the diagnosis-related groups (DRG). These groups are then combined in classes.

The main contributions include (i) a Markov Decision Process (MDP) model for computation of optimal admission policies under some restrictive assumptions, combined with (ii) a DES model for simulation and evaluation of ED hospitalization admission policies under relaxed assumptions.

The remaining part of the paper is organized as follows. The next section compares this paper with relevant literature. Section 3 provides a detailed description of the system under consideration. In section 4, a Markov Decision Process (MDP) model is built and the optimal policy is computed using a value iteration algorithm. In Section 5, simulation is used to bypass assumptions of the MDP model. Simulation enables the validation of MDP policies on realistic assumptions and for different scenarios. Conclusions and perspectives are given in Section 6.

2 LITERATURE REVIEW

Markov Decision Processes (MDPs) provide a strong mathematical framework for modeling dynamics of systems that include decision making. Under strict assumptions, evaluation and prediction of system performances are possible. MDP theory and concepts are entirely described by Puterman (2005). There is rather extensive literature on the admission of emergency and scheduled patients in health-care facilities, and also on MDPs applied to Admission Control Policy (ACP) in various domains. Our study focuses on using MDPs to find the optimal ACP in an health-care facility.

Most works using MDPs framework to study ACP are in the communication network field, where internal customers are often preferred to the admission of new customers. Altman (2000) provides an overview of the use of MDPs in this field. Here, unlike many studies using infinite queues with holding cost, we want to focus on loss queue system where clients (or patients) are lost if they are not immediately served (or admitted). Ku and Jordan (2002) and Tzeng (2009) studied control of parallel multiserver loss queues. Tzeng (2009) represented their one queue system as a two-dimensional Markov Chain with a variable number of servers. In our study, we intend to solve a problem with three input classes because it allows us to derive the case with any higher number of classes. The main contribution of Ku and Jordan (2002) is to consider more than two multiserver-queues. They demonstrated that the optimal admission policy for the considered parallel multiserver loss system is given by a set of thresholds. Each station has an arrival rate and service time. In contrast, we assume that arrival and departure rates depends on the arriving patient's class, and that all servers can be used for any class of patient.

In health-care, MDPs have been used to capture the stochastic aspects of patients and their impact on the utilization of the resources. Nunes, de Carvalho, and Rodrigues (2009) studied admission control policy of elective patients to promote an efficient utilization of hospital resources. They modeled different medical specialties and treatments in their MDP. Decisions consist in allocating a number of patients to each time period. The problem was solved using a value iteration algorithm, showing that the optimal policy can ensure a minimum utilization level while minimizing deviation costs. Our work differs on one basic assumption: we only consider emergency patients whereas they dealt with scheduled patients. Helm, AhmadBeygi, and Van Oyen (2011) took into account both unscheduled and elective patients. They especially focus on non-emergency patients as they assume that their bed admission can be delayed of a few days without deteriorating their health. They use a MDP to find an optimal control policy which balances canceling elective patients and scheduling later non-emergency patients. They found a threshold type optimum policy. Then they used simulation on a real-case hospital as a case study to validate their results. Our approach differs because we assume that any patient from ED needs to be hospitalized immediately. They deal with a scheduling problem with 3 sorts of patients, whereas we focused on a large number of DRG control policy.

In Pehlivan, Augusto, and Xie (2013), the authors studied admission control policies in a perinatal network. The main assumptions were a pure loss system of unscheduled patients, Poisson arrivals, same length of stay for all pregnant women, and only considering bed resources. In a network, each facility welcomes its own patients and overflowed patients from other facilities. Each class of patient induces a reward. In Pehlivan (2014), the dynamics of the network were deeply examined, considering hierarchical facilities, cyclical and non-cyclical overflows. She modeled her network with MDPs and solved it with an iteration value algorithm. She demonstrated that the optimum policy leads to a significant improvement in reducing overflows and in increasing total gain in the overall network. For the one-hospital case, she demonstrated that for m different classes of patients arriving at the ED and having the same length of stay law, the best control policy is of m -thresholds shape. The critical difference in our approach is that we consider that each class of patient has its own length of stay. Such an assumption requires one to rethink the modeling part and makes new challenges to solve.

Admission control of emergency patients has also been studied by Dormuth et al. (2008). Based on real data analysis, they implemented and compared two specific cost-sharing control policies. They used statistical methods to evaluate the impact of these policies on the number of emergency hospitalizations and physician visits.

Finally, discrete-event simulation (DES) has been widely used in the literature to model and assess the performance of health-care systems (Jun, Jacobson, and Swisher 1999; Augusto and Xie 2006; Augusto and Xie 2014). However, except in Pehlivan, Augusto, and Xie (2013), DES has not been used jointly with MDP analysis to assess the validity of theoretical models under many assumptions. In this paper, the main novelty relies on the application of DES to solve a realistic problem where actual patients' DRG are used to classify admissions.

The problem of finding an optimal admission policy considering the patients' DRG and the hospital economical balance is quite interesting, and to the best of our knowledge it has not been studied in the literature. Indeed, many of the problems dealing with patient admission in ED only consider objectives such as patient's LOS reduction, resources' sizing or bed occupancy, but the economical balance of the hospital and a DRG-based classification related to the stay of admitted patients was never taken into account. Moreover, MDP models require quite a number of restrictive assumptions, therefore it may not be realistic to use MDP results directly on a health-care problem due to its high complexity. Supporting and improving the decision making process with a simulation model which can better capture the complexity of the health-care system is the objective of this study.

3 HOSPITALIZATION ADMISSION OF EMERGENCY PATIENTS

A hospital often is composed of different several departments or medical units (cardiology, pneumology, geriatrics, pediatrics, ...). We assume that different departments of the hospital operate independently, that each department has a limited number of beds and it only takes care of specific classes of patients. All patients arrive through the ED (scheduled admissions are not taken into account in this study) and then they are admitted in different departments depending on their illness and whether hospitalization is needed. The main impact of patients on the hospital budget is the use of hospitalization beds, so occupancy of consultation rooms in the ED is not taken into account since the LOS is typically lower than a few hours. We focus on the hospitalization admission control in one department only, and we consider that it can be applied to any departments since the transfer of patients among departments is not possible.

We consider a given department with a fixed number N of beds dedicated to inpatients. One immediate consequence is that the department cannot admit more than N patients at the same time. Overflow is immediately lost and must be transferred to another hospital. Different classes of patients arrive at the emergency department. Classes are determined by a pathology-based classification. Let m be the number of different classes of patients to be admitted in the department. In this paper, we assume that all N beds of the department can be used by any of the m classes of patients. We do not consider ED patients who do not

need hospitalization because they do not impact the bed availability. Figure 1 illustrates this hospitalization admission process for ED patients.

The assumption of Poisson arrivals has been shown to be relevant for unscheduled hospital admissions (Brandeau, Sainfort, and Pierskalla 2004). We assume that each class i of patients arrive according to an independent Poisson process at rate λ_i ($1 \leq i \leq m$). The length of stay of each patient depends on his classes. We first assume that each class- i patient requires one bed for an i.i.d. exponentially distributed service time of rate μ_i . An exponential distribution is a restricted but necessary hypothesis in order to use the MDP model. The stochastic aspects of the problem are characterized by Poisson arrivals and Exponentially distributed service times. When the department admits a class- i patient, it generate a reward R_i corresponding to the financial and medical gain for having taken care of the patient. The reward depends on the DRG of the patient. Thus, the hospital admits patients with different arrival rates, different lengths of stay and different rewards, so that each patient impacts the system differently. The objective is to maximize the overall reward over a long-term horizon. To model a long term impact, we introduce a discount factor $\alpha \in [0, 1]$, and we select a large α very close to 1 to give much more importance to the future. When $\alpha = 0$ ($\alpha = 1$), any immediate reward is infinitely better than a later reward (one unit of immediate reward is as valuable as one unit of a later reward).

To sum up, there are five main assumptions in our problem setting:

1. m different classes of unscheduled patients arrives to the ED.
2. Any patient can be refused as an inpatient and then transferred to another hospital.
3. No transfer between departments is allowed.
4. Each class of patient arrives according to an independent Poisson process.
5. The length of stay is class-dependent and exponentially distributed.

Our system is entirely defined by the following notation:

- N Total number of beds,
- m Number of patient classes (indexed by i),
- λ_i Arrival rate for class- i patients,
- μ_i Service rate of class- i patients,
- R_i Reward of admitting a class- i patient,
- α Discounting factor.

The next section is dedicated to the use of the Markov Decision Process methodology to determine the optimal admission policy. A Value Iteration Algorithm is used to compute the optimal control policy.




Figure 1: Hospitalization admission process of ED patients.

4 MARKOV DECISION PROCESS MODEL

In this section we formulate a general MDP model for admission control of patients in any department of a hospital. The basics of MDP are available in the textbook Puterman (2005).

4.1 MDP Formulation

Although the MDP model of this paper holds for any number of patient classes, due to state space explosion, we limit ourselves to the case of three classes of patients with $m = 3$. For each state, there are $2m$ possible transitions to other states. Table 1 shows the impact of the number of patient classes and the number of beds on the number of states for the system.

A hospital department can be modeled has an M/M/N/N queue with no waiting room but with Poisson arrivals of multi-class of customers (patients), exponentially distributed service times, N identical servers (beds). The N servers are considered as staffed-beds, which is the combination of physical equipment (bed) with an appropriate coverage of nurses and physicians (Bruin et al. 2009). Each class i has specific arrival rate λ_i and service rate μ_i . A reward R_i is given for admitting a class- i patient. Without loss of generality, we assume $R_1 > R_2 > R_3$ so that admission of a class-1 patient is rewarded more than that of a class-2 or class-3 patient. Refusing the hospitalization of any patient generates no reward. Arrivals rates, service rate and rewards are constant over time.

The state of the system is modeled as a three-dimensional integer vector $X = (x_1, x_2, x_3) \in \Omega$, where x_i denotes the number of class- i patients in the system. Then $x_1 + x_2 + x_3$ is the total number of busy beds in the department. The state space of the department is finite and can be formulated as follows:

$$\Omega = \{(x_1, x_2, x_3) \mid 0 \leq x_i \leq N, i \in \{1, 2, 3\} \text{ and } (x_1 + x_2 + x_3) \leq N\} \quad (1)$$

When the system is in state X , a control action is defined by a vector $A(X) = [a_i(X)]$, where $a_i(X) = 1(0)$ indicates whether to admit (or refuse) class- i arrival in state X , for $i \in \{1, 2, 3\}$. For instance, the following is true for a policy that always accepts patients when the system is empty:

- $a_i(X) = 0 \forall i \in \{1, 2, 3\}$ at state $X = (x_1, x_2, x_3)$ where $(x_1 + x_2 + x_3) = N$: when department is full, no patient can be admitted.
- $a_i(X) = 1 \forall i \in \{1, 2, 3\}$ at state $X = (x_1, x_2, x_3)$ where $(x_1 + x_2 + x_3) = 0$: when department is empty, all patients are admitted.

In this study, we use event-based dynamic programming (DP). Two classes of events can occur: arrival or departure of a patient. Let $V_i^a(X)$ and $V_i^d(X)$ be the optimal discounted value reward function starting at state S with respectively a class- i arrival or departure, which are evaluated in equations (3) and (4). Our objective is to find an optimal control policy π^* that maximizes the expected total α -discounted reward over an infinite horizon ($0 < \alpha < 1$).

The model described above is a continuous-time Markov decision process. We convert it to a discrete-time MDP by uniformization (Lippman 1975). This method leads to an equivalent discrete-time MDP by allowing fictitious events. We use uniformization rate $\gamma = \sum_{i=1}^m (\lambda_i + \mu_i N)$. In the equivalent discrete-time MDP, events occur with probability $p_i^a = \frac{\lambda_i}{\gamma}$ for class- i arrivals and $q_i^d = \frac{\mu_i N}{\gamma} = \frac{\mu_i x_i}{\gamma} + \frac{\mu_i (N - x_i)}{\gamma}$ for class- i departures that are either real or artificial. For a given state $X = (x_1, x_2, x_3)$, then $\forall i \in \{1, 2, 3\}$, actual class- i departures occur with probability $\frac{\mu_i x_i}{\gamma}$, and artificial ones with probability $\frac{\mu_i (N - x_i)}{\gamma}$. Artificial departures are introduced so that $p_i^a + q_i^d = 1$, and they correspond to a self-loop on state X since no present patient actually leave the department.

Using uniformization (Ross 1989), the optimal value function $V(x)$, i.e. total discounted reward by starting at state X , is the unique solution of the following optimality equation:

$$V(X) = \sum_{i=1}^m p_i^a V_i^a(X) + \sum_{i=1}^m q_i^d V_i^d(X) \quad (2)$$

where, $\forall i \in \{1, 2, 3\}$, $\forall X = (x_1, x_2, x_3) \in \Omega$:

$$V_i^a(X) = \begin{cases} \max(R_i + \alpha V(X + e_i), \alpha V(X)) & \text{if } (x_1 + x_2 + x_3) < N \\ \alpha V(X) & \text{if } (x_1 + x_2 + x_3) = N \end{cases} \quad (3)$$

$$V_i^d(X) = \begin{cases} \alpha(\frac{x_i}{N}V(X - e_i) + \frac{N-x_i}{N}V(X)) & \text{if } x_i \geq 1 \\ \alpha V(X) & \text{if } x_i = 0 \end{cases} \quad (4)$$

- $\gamma = \sum_{i=1}^m (\lambda_i + \mu_i N)$ is the uniformization rate,
- $p_i^a = \lambda_i / \gamma$ is the probability of a class i arrival,
- $q_i^d = \mu_i N / \gamma$ is the probability of a real or artificial class i departure,
- $m = 3$ is the number of patient classes,
- V_i^a is the optimal value function with a class- i arrival,
- V_i^d is the optimal value function with a tpe- s departure,
- e_i is the m -dimensional vector with i th entry equal to 1 and other entries equal to 0,
- R_i is the reward for accepting a class i patient
- α is the discount rate, ($0 < \alpha < 1$),

The optimal admission policy can be described as follows: in each state X , a class- i patient is accepted if its immediate reward R_i is higher than the loss of opportunity $\alpha(V(X) - V(X - e_i))$.

4.2 Computation of an Optimal Policy with MDP

According to the above description, the Markov chain model has a finite state space since the total number of beds is fixed (it is also irreducible and homogeneous). Thus, there is a unique equilibrium solution. The number of possible states for our system increases exponentially with the number of classes and beds (Table 1). Hence the analytic model focuses only on the case with three classes of patients, with a moderate number of beds (less than one hundred). In the next section, we show that this limit is not a problem when applying efficient control strategies in a real hospital, considering more than three diagnosis-related groups.

Table 1: Number of possible states in the Markov Chain, depending on m and N .

Number of classes (m)	Number of beds (N)	Number of states
2	10	66
2	30	496
2	100	5151
3	10	287
3	30	5457
3	100	176852

In our case study with three classes of patients, any state has six neighboring states corresponding to one of the six possible events: one arrival and one departure of any of the three classes. We use an efficient computational algorithm, Value Iteration (VI) Algorithm, to reach the optimal policy for our case study. It relies on an iterative procedure which computes the expected gain of any possible state, using gains of neighboring states. It stops when the gain difference between two successive steps is very low. The aim is to find the policy which leads to to maximization of every state's gain. The proposed VI algorithm was implemented in C++ as presented in Figure 2.

As a result, the algorithm returns $\pi^* = \{\pi^*(X), X \in \Omega\}$, the optimal admission control policy. For any state $X \in \Omega$, for any arrival of a patient, $\pi^*(X)$ contains the information about the best decision to make. In the general case with m classes of patients, there are 2^m possible policies for each state. Indeed, in a given state, each class of patient can be either Admitted or Refused. It makes two choices for each patient class, combined by the number of classes. The four first columns of Table 2 enumerates the 2^3 possible policies for the three-classes problem, the decision for each class is represented as a binary (Admit or Refuse).

Requirements	$\Pi = \{\pi(X), X \in \Omega\}$ with $\pi(X)$ optimal policy to apply in state X
Input	$N, m, \alpha, \varepsilon, R_i, \lambda_i, \mu_i$
1. Initialization	$t \leftarrow 0, \varepsilon \leftarrow 10^{-3}, it \leftarrow 0, MAX_IT \leftarrow 10^4$ $\forall X \in \Omega \forall t \in \mathbb{N}, 0 \leq V_0(X) \leq V_t(X)$ $\Omega_N = \{X_C = (x_1, \dots, x_m) \in \Omega \mid \sum_{i=1}^m x_i = N\}$ $\forall X \in \Omega_N, \Pi(X) = 0$
2. $t \leftarrow t + 1, it \leftarrow it + 1$	$\forall X \in \Omega - \Omega_N$ $V_t(X) = \max_{a \in A(X)} \left\{ \sum_{i=1}^m p_i^a V_{i,t-1}^a(X) + q_i^d V_{i,t-1}^d(X) \right\}$ $\Pi_t(X) = \operatorname{argmax}_{a \in A(X)} \left\{ \sum_{i=1}^m p_i^a V_{i,t-1}^a(X) + q_i^d V_{i,t-1}^d(X) \right\}$ $\forall X \in \Omega_N$ $V_t(X) = \sum_{i=1}^m (p_i^a \alpha V_{t-1}(X) + q_i^d V_{i,t-1}^d(X))$
3. Compute	$DIF_MIN = \min_{Y \in \Omega} \{V_t(Y) - V_{t-1}(Y)\}$ $DIF_MAX = \max_{Y \in \Omega} \{V_t(Y) - V_{t-1}(Y)\}$ if $DIF_MAX - DIF_MIN \leq \varepsilon \cdot DIF_MIN$ or $it > MAX_IT$ $\Pi = \{\pi_t(X), X \in \Omega\}$ is an optimal policy. STOP else go to step 2

Figure 2: Pseudo-code of the Value Iteration algorithm for the admission control policy optimization.

4.3 Analytic Results

We tested the VI algorithm on several instances of the 3-classes problem ($m = 3$). Algorithm's performances depend on both the total number of beds and the maximum number of iterations we allow for the dynamic programming part. The performance is measured by the execution time and the reach of the best solution.

Table 2 shows the results provided by the algorithm on a certain case study. The execution time is less than one second. For each of the eight possible policies, it shows the states in which such policy must be applied to maximize the gain. Recall that a state is given by the triplet (x_1, x_2, x_3) , so that there are 56 states for $N = 5$. The first immediate inference we can make from those results, and which is intuitive, is that the best policy never refuses the patient which the highest reward by unit time (Class-1 patients in our case), except when all beds are already occupied. This numerical result is coherent with theoretical demonstrations of (Pehlivan 2014). Then, we see that when only few beds are occupied ($(x_1 + x_2 + x_3) \leq 5$), all classes patients are admitted. Nevertheless, as $(x_1 + x_2 + x_3)$ increases, best policy starts refusing class-2 and class-3 patients: for instance, policy 4 in Table 2 shows that class-1 patients are the only patients admitted in 13 different states.

Figure 3 shows the convergence speed of the VI algorithm on a simple case study. Input parameters are the following: $N = 5, \lambda_1 = 4, \lambda_2 = 5, \lambda_3 = 3, \mu_1 = 1.1, \mu_2 = 1, \mu_3 = 1.2, R_1 = 4, R_2 = 2, R_3 = 1$ and $\alpha = 0.99$. Finding the admission control policy which brings the maximum reward is our critical purpose. Figure 3 illustrates the fact that choosing a small number of iterations will not allow to find a policy which maximizes the total gain. Indeed, as we are giving much more importance to the future (discount rate $\alpha = 0.99$ very close to 1), dynamic programming must compute a lot of steps to know the impact of a current decision on the future. Numerical experiments lead us to assume that choosing at least 2000 iterations is sufficient to find optimal policies in our case studies. The number of iterations has a heavy impact on the execution time, so finding such a lower limit is of interest to ensure computational success. For instance, when choosing $N = 20$, execution time increases from 0.05 second for 100 iterations to 4.15 seconds for 3000 iterations.

Table 2: For $m = 3$, enumeration of the height possible policies for each state (1=Admit, 0=Refuse).

Policy	Class 1	Class 2	Class 3	Case study: list of states in which policy applies
1	1	1	1	(0,0,0) (0,0,1) (0,0,2) (0,1,0) (0,1,1) (1,0,0) (0,2,0) (0,3,0) (1,0,1) (1,1,0) (1,2,0) (2,0,0) (2,0,1) (2,1,0) (3,0,0) (0,0,3) (0,1,2) (0,2,1) (1,0,2) (1,1,1) (4,0,0)
2	1	1	0	(0,4,0)
3	1	0	1	(0,0,4) (0,1,3) (0,2,2) (0,3,1) (1,0,3) (1,1,2) (1,2,1) (1,3,0) (2,0,2) (2,1,1) (2,2,0) (3,0,1) (3,1,0)
4	1	0	0	Never
5	0	1	1	Never
6	0	1	0	Never
7	0	0	1	Never
8	0	0	0	When full : (0,0,5) (0,1,4) (0,2,3) (0,3,2) (0,4,1) (0,5,0) (1,2,2) (1,3,1) (1,4,0) (2,0,3) (2,1,2) (2,2,1) (2,3,0) (3,0,2) (3,1,1) (3,2,0) (4,0,1) (4,1,0) (5,0,0) (1,0,4) (1,1,3)




Figure 3: Optimum value returned by the VI algorithm depending on the number of iterations.

Table 3 shows the execution time. Tests were ran on a cluster with up to 24GB of RAM. The main limitation of our approach is related with memory usage to record the states of the Markov chain. We were able to solve problems up to 50 beds. Note that once the memory is allocated, execution time is fast.

Table 3: Execution time of the algorithm in seconds, depending on the number of beds.

Number of beds (N)	10	20	30	40	50	100
Execution time (sec)	0.48	3.19	4.55	7.68	13.35	NA

5 DISCRETE EVENT SIMULATION

5.1 Data Collection

All the data we used for the simulation model come from a local middle-size hospital, the hospital of Firminy (Loire, France). From all patients who arrived at the ED for the year 2010, we selected the 14,735 who needed to be hospitalized. For each patient, the database provides the arrival and departure dates, from which we derive the length of stay, the medical unit and the diagnosis-related group (DRG). There is a list published from the French Ministry of health with the corresponding price of each DRG. We observed 453 different DRG among all the patients, which we clustered in three theoretical same size groups of patients according to their ratio:

$$\frac{\text{Price of the DRG}}{\text{Average LOS of the DRG}} \quad (5)$$

From the three main classes generated this way, each simulated patient will be assigned a specific DRG among the 453 according to a probabilistic distribution. The following attributes are assigned to each patient in our simulation model: (i) unique id, (ii) real DRG (corresponding to the main diagnosis), (iii) theoretical DRG (three classes with three theoretical costs) corresponding to the cost classes used in

the last section. Control policies will be applied to each of the three theoretical groups, instead of the 453 initial groups because of the computing time limitation. The stochastic part of the length of stay of each patient is modeled with a Log-normal distribution, where the average is the length of stay of the DRG and the standard deviation is forty percent of the mean. This distribution is the best fit for our real-case data while looking for a single realistic law for all the classes of patients.

5.2 Model Description

Using the Arena software (Altiok and Melamed 2007), the simulation model has been built after the model described in Figure 1. *Create* blocks model patient arrivals in the hospital. Several attributes define patient pathology and severity, following the distribution described in the last section: a unique id, a *DRG* (Diagnosis-related group) to calculate the actual cost and a fictional class $c \in \{1, 2, 3\}$ to take the decision at the control point modeled using a *Decide* block. Finally, the stay in the hospital is modeled using a *Delay* block, with a probabilistic length depending on actual patient's *DRG*. Patient stay benefit b for the hospital is calculated using the actual DRG agreed sum g_{DRG} , the effective length of stay (*LOS*), the DRG upper bound (*UB*), and a variable positive gain b_{DRG} used to compensate money loss related with the additional length of stay:

$$b = g_{DRG} + \max(0, LOS - UB) \times b_{DRG} \quad (6)$$

Relevant performance indicators are gathered when entities exit the model: (i) total gain of the hospital (depending on actual patient classification) and (ii) number of transferred (refused) patients from the hospital. In order to ensure a 95 % confidence interval for such indicators, 20 initial replications were launched. The minimal number of replications n is calculated using the following formula (Law and Kelton 2004).

$$\bar{X}(n) \pm t_{n-1, 1-\alpha/2} \sqrt{\frac{S^2(n)}{n}} \quad (7)$$

$S^2(n)$ is the estimated standard deviation and $\bar{X}(n)$ is the estimated average for the X_1, X_2, \dots, X_n values of an indicator for n replications. The value $t_{n-1, 1-\alpha/2}$ is the critical point of the distribution t with $n - 1$ degrees of liberty and a covering equal to $\frac{1-\alpha}{2}$. In our case, $\alpha = 0.05$. The minimum number of replications n^* to ensure a relative error $\gamma = \frac{\bar{X} - \mu}{\mu}$ is calculated using an iterative procedure where the number of replications is incremented until we reach an admissible confidence interval.

A warm-up period is also defined to allow the system to enter its steady-state. The length of the warm-up is estimated following the procedure described in Law and Kelton (2004). In our case, the warm-up length is 30 days. Replication length is set to 1000 days, which is long enough to assess the proposed methods.

5.3 Simulation Results

Five scenarios corresponding to five control policies have been tested: All, accept everyone; Best1 and Best2 use the control policy provided by the optimal control policy considering arrival rates 1 and 0.6 (respectively); C1C2 always refuses class 3 patients and always accept class 1 and class 2 patients; C1 always accepts class 1 patients and refuses everyone else. Actual life scenario observed in most hospitals are either corresponding to All or C1C2.

Results are presented in Figures 4, 5 and 6. Figure 4 shows the total theoretical reward (by using the same value as in the analytical approach) obtained using DES for different scenarios. We observe that Best1 and Best2 policies are better when considering an arrival rate lower than 0.4. For higher arrival rates, Best1 and C1C2 are the best scenarios. Similar results may be observed in Figure 5 when actual DRG rewards are considered: in that case, Best2 is the best strategy except for very high and unrealistic arrival rates greater than 1. When we consider a "regular" arrival rate of 0.4 or a high arrival rate of 0.8 (modeling winter epidemics), our optimal control policy remains better than any other empirical approach.

Finally, Figure 6 shows the total refusal rate for all policies. If the benefit of an optimal control policy compared to empirical approaches was not so important considering total gain, refusal rates remain much lower: the total area under the curves presented in Figure 6 for policies Best1 and Best2 are lower than any other, meaning our optimal control policies have good results for all possible arrival rates.




Figure 4: Total theoretical reward regarding arrival rate using DES.




Figure 5: Total actual DRG reward regarding arrival rate using DES.

6 CONCLUSION AND FUTURE WORKS

In this paper, we considered the problem of finding an optimal admission control policy for a hospital while considering diagnosis-related groups (DRG) and associated gains depending on length of stay. In order to model the system analytically, a Markovian assumption was necessary to make. Under this assumption, the system is modeled as an MDP and an MDP optimal policy is computed using a value iteration algorithm. Afterwards, a discrete-event simulation model is used to strengthen the decision making process by incorporating the complexity in the model which cannot be captured by the MDP and measure the impact of Markovian assumptions on total reward when considering actual DRG. To do so, the MDP optimal policy is considered as a base scenario and various control scenarios are tested considering the

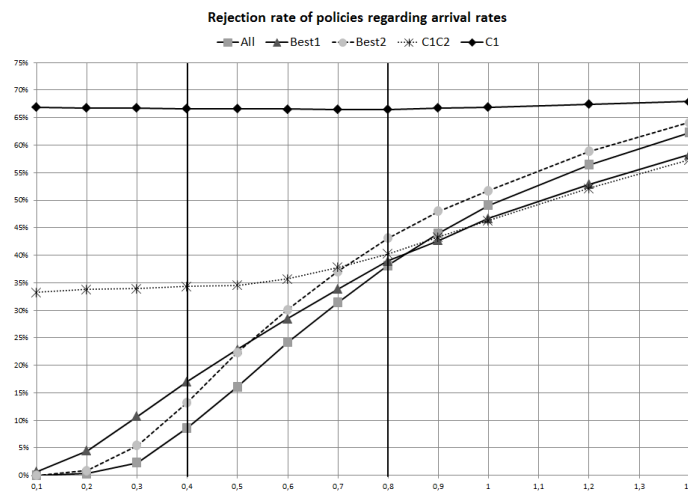


Figure 6: Refusal rates regarding arrival rate using DES.

total gain and the refusal rates as performance indicators. It has been observed that the optimal control policies outperform any other empirical scenarios that may be observed in actual hospitals.

In this study, overflowed patients are assumed to arrive in another hospital in the network. In a future work, we believe it is scientifically interesting to consider a pure-loss health-care network where overflows are less restricted as in reality and deal with the problem of finding optimal admission control policies for all hospitals in the network. Moreover bed capacities of each service in the hospital should be considered in order to improve the control strategy with more realistic decisions such as bed transfer.

REFERENCES

- Altioik, T., and B. Melamed. 2007. *Simulation modeling and analysis with Arena*. Academic Press Inc.
- Altman, E. 2000. "Applications of Markov decision processes in communication networks". *INRIA Research Report RR-3984*.
- Augusto, V., and X. Xie. 2006. "Modélisation et analyse de flux par la simulation en milieu hospitalier : état de l'art". In *Actes de la conférence GISEH*.
- Augusto, V., and X. Xie. 2014. "A Modeling and Simulation Framework for Health-care Systems". *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* 44 (1): 30–46.
- Brandeau, M. L., F. Sainfort, and W. P. Pierskalla. 2004. *Operations research and health care: A handbook of methods and applications*. Springer.
- Bruin, a. M., R. Bekker, L. Zanten, and G. M. Koole. 2009. "Dimensioning hospital wards using the Erlang loss model". *Annals of Operations Research* 178 (1): 23–43.
- Dormuth, C. R., M. Maclure, R. J. Glynn, P. Neumann, A. M. Brookhart, and S. Schneeweiss. 2008. "Emergency hospital admissions after income-based deductibles and prescription copayments in older users of inhaled medications.". *Clinical Therapeutics* 30 (1): 1038–1050.
- Helm, J. E., S. AhmadBeygi, and M. P. Van Oyen. 2011, May. "Design and Analysis of Hospital Admission Control for Operational Effectiveness". *Production and Operations Management* 20 (3): 359–374.
- Jun, J. B., S. H. Jacobson, and J. R. Swisher. 1999. "Application of discrete-event simulation in health care clinics: A survey". *Journal of the Operational Research Society* 50 (2): 109–123.
- Ku, C.-Y., and S. Jordan. 2002, December. "Access control of parallel multiserver loss queues". *Performance Evaluation* 50 (4): 219–231.
- Law, A. M., and W. D. Kelton. 2004. *Simulation Modeling and Analysis*. McGraw-Hill.
- Lippman, S. A. 1975. "Applying a new device in the optimization of exponential queuing systems". *Operations Research* 23 (4): 687–710.

- Nunes, L. G. N., S. V. de Carvalho, and R. D. C. M. Rodrigues. 2009. "Markov decision process applied to the control of hospital elective admissions." *Artificial Intelligence in Medicine* 47 (2): 159–171.
- Pehlivan, C. 2014. *Design and flow control of stochastic healthcare networks without waiting rooms: a perinatal application*. Ph. D. thesis, École nationale supérieure des Mines de Saint-Étienne.
- Pehlivan, C., V. Augusto, and X. Xie. 2013. "Admission Control in a Pure Loss Healthcare Network: MDP and DES approach". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 54–65. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Puterman, M. 2005. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, New York.
- Ross, K. W. 1989. "Randomized and pastdependent policies for Markov decision processes with multiple constraints". *Operations Research* 37 (3): 474–477.
- Tzeng, S.-S. 2009, December. "Call admission control policies in cellular wireless networks with spectrum renting". *Computer Communications* 32 (18): 1905–1913.

AUTHOR BIOGRAPHIES

MARTIN PRODEL is a Ph.D. Student at the company HEVA and the Center for Health Engineering at Ecole Nationale Supérieure des Mines de Saint Etienne (ENSM.SE), France. He received both his engineering degree and his Msc degree in Industrial Engineering from the ENSM.SE in 2013. His research interests include applications of modeling, simulation and Markov Theory in health care systems. His email address is martin.prodel@emse.fr.

VINCENT AUGUSTO received his Ph.D. in Industrial Engineering from the École Nationale Supérieure des Mines de Saint-Étienne, France, in 2008. He was a visiting scholar at CIRRELT (Centre Interuniversitaire de Recherche sur les Réseaux d'Entreprise, la Logistique et le Transport), University of Laval, Quebec, Canada in 2009. Currently, he is an assistant professor in the Department of Health care Engineering at the ENSM.SE. His research interests include modeling, simulation, optimization of health care systems and their supply chains. His e-mail and web addresses are augusto@emse.fr and <http://www.emse.fr/~augusto>.

XIAOLAN XIE received his Ph.D degree from the University of Nancy I, Nancy, France, in 1989, and the Habilitation à Diriger des Recherches degree from the University of Metz, France, in 1995. Currently, he is a distinguished professor of industrial engineering, the head of the department of Healthcare Engineering of the Center for Health Engineering and the head of IEOR team of CNRS UMR 6158 LIMOS, Ecole Nationale Supérieure des Mines (ENSMSE), Saint Etienne, France. He is also a chair professor and director of the Center for Healthcare Engineering at the Shanghai Jiao Tong University, China. Before Joining ENSMSE, he was a Research Director at the Institut National de Recherche en Informatique et en Automatique (INRIA) from 2002 to 2005, a Full Professor at Ecole Nationale d'Ingénieurs de Metz from 1999 to 2002, and a Senior Research Scientist at INRIA from 1990 to 1999. His research interests include design, planning and scheduling, supply chain optimization, and performance evaluation, of healthcare and manufacturing systems. He is author/coauthor of over 250 publications including over 90 journal articles and five books. He has rich industrial application experiences with European industries. He is PI for various national and international projects including ANR-TECSAN HOST on management of winter epidemics, NSF China key project on planning and optimization of health care resources, French Labex IMOBS3 project on home health cares, FP6-IST6 IWARD on swarm robots for health services. He has been an associate editor for *IEEE Transactions on Automation Science & Engineering*, *IEEE Transaction on Automatic Control*, *IEEE Transactions on Robotics & Automation* and *International Journal of Production Research*. He has been a Guest Editor of various special issues on healthcare engineering and manufacturing systems. He is general chair of ORAHS'2007 and IPC chair of the IEEE Workshop on Health Care Management WHCM'2010. His email address is xie@emse.fr.