

A STOCHASTIC CONTROL APPROACH TO AVOIDING EMERGENCY DEPARTMENT OVERCROWDING

Arun Chockalingam
Krishna Jayakumar

Mark A. Lawley

School of Industrial Engineering
Purdue University
315 N. Grant St
West Lafayette, IN 47907, USA

Weldon School of Biomedical Engineering
Purdue University
206 S. Martin Jischke Dr
West Lafayette, IN 47907, USA

ABSTRACT

Emergency Department (ED) overcrowding is a common problem in hospitals in the United States. Presenting a barrier to safe delivery of healthcare, hospitals address ED overcrowding by diverting ambulances to the nearest available facility, leading to delays in healthcare delivery and losses in revenue. Control policies on hospital resources could greatly improve healthcare delivery by preventing overcrowding and ambulance diversion. In this paper, we use Petri-nets (PNs) to model patient and resource flow in a hospital system. Simulating these PNs, we can observe changes in the availability of resources over time and obtain a stochastic differential equation (SDE) which models the hospitals proximity to entering a divert state (in a Euclidean sense). Likening the resource allocation problem to a stochastic control problem, we derive the related free-boundary problem. The solution of this problem is the optimal control policy that dictates when and how many resources should be added or removed.

1 INTRODUCTION

Ambulance diversion is a state which occurs when a hospital Emergency Department (ED) is overcrowded and cannot care for additional emergency patients. In this situation, it redirects incoming ambulances to other nearby hospitals or medical facilities. A number of reasons lead to hospital overcrowding. Lack of access to primary care physicians and facilities cause many patients to present themselves at the ED for treatment (Grumbach, Keane, and Bindman 1993). Furthermore, EDs are required by the Emergency Medical Treatment and Active Labor Act (EMTALA), passed in 1986, to screen incoming patients and to provide treatment if needed. Irrespective of the state of the hospital, these laws mandate providing patient care which further exacerbates the problem. A detailed discussion on primary causes of ED overcrowding can be found in Committee on Pediatric Emergency Medicine (2004).

A recent Institute of Medicine study shows that ED overcrowding presents a barrier to safe and timely delivery of healthcare. From 1993 to 2003, the number of ED visits have increased by 90.3 million while the number of hospitals has decreased by 703, the number of beds by 198,000, and the number of EDs by 425 (Institute of Medicine 2006). During this time, 45% of the hospitals reported some period of time in ambulance diversion (Burt, McCaig, and Valverde 2006). Studies have shown that ED overcrowding is a serious problem that can lead to adverse patient outcomes and impaired access to care.

1.1 Motivation

Ambulance diversion can have significant effects on patients, hospitals and the community, as stated by Pham and Patel (2006), Glushak, Delbridge, and Garrison (1997). Critically ill patients may face long delays to get hospital care. Overloading of resources may result in physician and staff dissatisfaction. Ambulance diversions to distant hospitals may lead to unavailability of ambulances to the community. There is also revenue loss associated with ambulance diversion. Each hour spent on diversion results in

a loss of \$1086 in revenue for hospitals (McConnell, Christopher, and Mohamud 2006). Furthermore, each patient not seen equates to \$8000 to \$10000 in lost revenues to the hospital (Jenson 2003). The problem of hospital overcrowding is a serious one that needs to be addressed.

1.2 Contribution and outline

In this paper, we propose a novel approach to derive control policies which hospitals may use to manage resources better to improve patient care and profit to the hospital. Specifically, these policies dictate when and how many additional resources need to be brought in to ease congestion. We develop a mathematical model of a hospital system using Petri-net theory, and convert this model to a simulation model. The novelty of the approach lies in the next step, where using the PN model, we develop a measure termed *distance to divert*, which measures the proximity of a hospital to a divert state, in an Euclidean sense. Output obtained from the simulation runs of the hospital system are used to obtain a stochastic representation for this measure, based on which we derive a stochastic impulse control problem. Stochastic impulse control theory has been widely applied to financial engineering problems, but to the best of our knowledge, it has not been applied in the context of healthcare problems. The solution to this stochastic control problem is the optimal control policy, which can be used by hospitals to better manage resources.

The paper is laid out as follows. In Section 2, we discuss our use of Petri-net theory to model a hospital system. In Section 3, we describe the simulation of the Petri-net model. In Section 4, we link the current problem to a class of problems known as stochastic impulse control problems. We provide a brief numerical example illustrating this approach in Section 5 before concluding in Section 6.

2 PETRI-NET MODELS OF PATIENT FLOW

The patient flow in a typical hospital system is shown in Figure 1. Solid lines in the figure represent patient flow and dotted lines represent information flow. Patient arrival to the ED is either by emergency medical services (EMS) or walk-ins. Physicians order tests as part of patient assessment and these tests are done in the labs. From the ED, patients may move to one of three units; the Operating Room (OR), the Intensive Care Unit (ICU) or the Medical Surgical Unit (Med/Surg). The path followed depends on the severity of their condition and the type of treatment required. Patients may move to the OR for required procedures. After emergency treatment, patients with higher acuity may move into the ICU while stable patients move to Med/Surg. A situation known as *access blocking* may arise, where patients from the ED cannot move to units such as the ICU and Med/Surg due to a lack of free inpatient beds. This situation prevents ED patients requiring inpatient-care from gaining access to appropriate hospital beds within a reasonable time frame and halts their progress. A propagation of these effects lead to congestion, and ultimately a state of ambulance diversion.

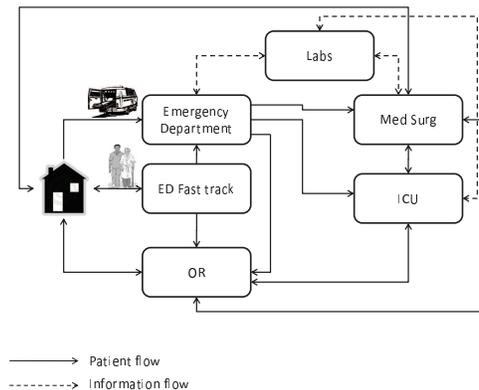


Figure 1: Hospital patient-flow diagram

In this paper, we use Petri-nets (PNs) to model patient flow in a hospital. Petri-nets are directed, bipartite graphs that can be used to model discrete distributed systems. A PN consists of a set of *places*, and a set of transitions and arcs that connect place(s) to a transition and vice-versa. Non-negative

integers assigned to every place in the net are known as *tokens*. A *marking* assigns to each place in the PN a number of tokens. The marking can be viewed as a ‘snapshot’ of how tokens are distributed in the PN. The PN is defined by its structure and its initial marking M_0 . A transition is said to be enabled when every input place of the transition has at least k tokens, where k is the weight of the arc from an input place to the transition. When a transition fires, tokens are removed from input places and deposited into output places. Input places are analogous to pre-conditions and output places are analogous to post-conditions. Transitions can be viewed as events. Thus, for an event to occur (transition to fire), pre-conditions (tokens in input places) must be satisfied. Occurrence of the event creates the post conditions. Thus, tokens are removed from input places and deposited in output places. A detailed tutorial on Petri-nets can be found in [Murata \(1989\)](#).

Our choice of using PNs to model patient flow in hospitals stems from the fact that PNs capture structural properties of the underlying system which we can study and use. Two properties we are interested in are *reachability* and *boundedness*.

A marking M_f is said to be *reachable* from a marking M_0 if there exists a sequence of firings that transform M_0 to M_f . So, if M_f is reachable from M_0 , and the net consists of m places and n transitions, we can write the state equation

$$M_f = M_0 + \phi X, \tag{1}$$

where $\phi \in \mathbb{Z}^m \times \mathbb{Z}^n$ is the *incidence matrix* and $X \in \mathbb{Z}_+^n$ is the *Parikh vector* (PV). Entry ϕ_{ij} in the incidence matrix describes the change in tokens in place $i \in [1, m]$ when transition $j \in [1, n]$ fires. The PV is the set of transition firings required for the system to evolve from M_0 to M_f . Each entry in the PV (x_j) describes the number of times transition j needs to fire in order for M_f to be reached from M_0 . In our context, using the reachability property, given a current state (or marking), we can determine the sequence of firings necessary for the hospital to enter a divert state.

A PN is said to be *bounded* if the number of tokens in each place does not exceed a finite number for any marking reachable M_0 . The boundedness property is used to verify and validate the PN model used to describe patient flow in the hospital.

An example of a PN with 5 places and 3 transitions is shown in Figure 2. The current marking

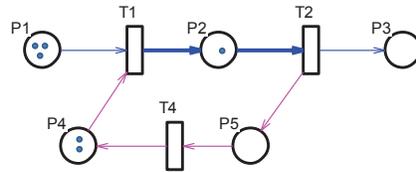


Figure 2: Example of a Petri-net

of this PN is $[3, 1, 0, 2, 0]^T$. Firing transition T1 removes 1 token from P1 and 1 token from P4 and deposits 1 token in P2. This results in a new marking $[2, 2, 0, 1, 0]^T$. Firing T2 now changes the marking to $[2, 1, 1, 1, 1]^T$. The PV for this sequence of transitions is $[1, 1, 0]^T$. The state equation for this example can now be written as

$$\begin{bmatrix} 2 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \\ 0 \\ 2 \\ 0 \end{bmatrix} + \begin{bmatrix} -1 & 0 & 0 \\ 1 & -1 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & -1 \end{bmatrix} \times \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \tag{2}$$

We develop patient-flow PN models, consisting of the units depicted in Figure 1, based on care delivery processes in the hospital. Figure 3 shows the complete PN for the hospital system. The models contain resources like beds, nurses, physicians etc., in the form of tokens and patients use these tokens according to model logic to receive care. These resource tokens are held by the patients until they move to the next stage. For instance, patients transferring from ED to the ICU will not release ED beds until they acquire an ICU bed. These create dependencies between the different hospital units and allow us to model access blocking and divert. We define divert state to be the state

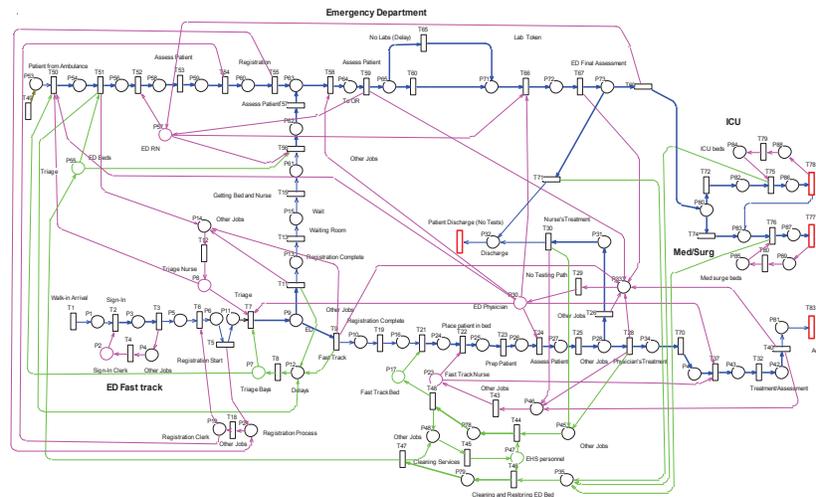


Figure 3: Petri-net model of the hospital

where there are no ED and ICU beds available. In the PN model, it can be viewed as “no available tokens” in places which hold resource tokens for ED and ICU beds. Using the number of beds to decide divert status is logical, since the hospital can instantly determine how many beds are available, whereas the problem of partial observability arises when trying to determine the state of other areas in the hospital, such as the waiting room.

In our model, patients can walk-in to the ED or arrive by ambulance. ED Fast Track is the part of ED that serves less critical patients and eases congestion in the ED by freeing up ED resources for more critical patients. For modeling purposes, we assume that all walk-in patients are directed to FT, and all ambulance arrivals are directed to the ED. A patient’s condition is assessed in the triage. Critical walk-in patients may be transferred to ED. Patients move into Med/Surg and ICU as the case may require. Stabilized patients are discharged home. The physicians may use the services of lab to perform various tests on the patient.

Along the patient-flow path, a patient will require and use resources like nurses, beds and physicians during their stay in the hospital. These resources are modeled as places with limited tokens in the PN. The delay in availability of a resource is represented as non-availability of tokens to advance the patient tokens through the net. This delay increases the number of patient tokens in the system waiting for a resource. Disruption of patient flow in the ED can be created not only by lack of resources in the ED, but also by a lack of ICU or Med/Surg resources. The extent of non-availability of resources and the number of patients in the system can be used to determine whether the hospital is close to entering a divert state.

By assigning probability distributions to the occurrence of events, modeled by transitions, and to the length of time taken for an event to complete, we can use discrete-event simulation software to study the hospital’s operations in detail. Using *AutoMod* to simulate the hospital, we can study measures such as time spent in divert state, utilization level of resources and total cost to the hospital. We have used data obtained from St. Vincent’s hospital in Birmingham, Alabama to keep our model as close to the real world as possible, and simulated hospital operations for 30 days, as in [Kokler \(2008\)](#).

Before we can simulate the PN model we have just developed, we need to verify that the net represents the real world process. The boundedness property of the Petri-net allows us to study the tendency of the system to overflow and thereby verify the accuracy of the model. Given an initial marking M_0 , we determine the maximum number of tokens that each place in the net can hold for all reachable markings. The bound for each place is obtained by solving an integer-programming (IP) problem using Equation (1), the state equation. For each place $i \in [1, m]$, the problem is formally

stated as

$$\begin{aligned} & \max y_i \\ \text{such that } & \mathbf{y} = \mathbf{M}_0 + \phi \mathbf{X} \\ & y_i \in \mathbb{Z}_+ \\ & x_j \in \mathbb{Z}_+ \quad \forall j \in [i, n], \end{aligned}$$

where $\mathbf{y} = (y_1, y_2, \dots, y_m)^T$. Solving this problem for each of the m places will provide us with the vector \mathbf{y} , the maximum number of tokens that can be held in each place.

Now, resources such as registration clerks, are also modeled as tokens in places. Clearly, all places dependent on this resource, such as “patient with the registration clerk”, are constrained by this resource. The maximum number of tokens that this constrained place can hold is therefore limited to the maximum number of available tokens in the constraining place. For example, if there are only 4 registration clerks, at any time, there can be at most 4 tokens in the place “patients with registration clerk”. On the other hand, for any place that is not constrained by other resources, the maximum number of tokens that can be held in this place is the number of tokens that the system is initialized with. The place “patients in the waiting room” (assumed to have infinite capacity) is not constrained by any resource, and depends only on the number of patients (tokens) in the system. The maximum number of tokens that can be held in this place would clearly be the number of patients in the system.

Our PN model consists of 90 places and 83 transitions. For the boundedness study, we load the system with 100 tokens, and solve the IP problem described above for each of the 90 places. We list the boundedness results for 6 of the 90 places in Table 1 as an illustration of the boundedness study. For each of the 6 places, the table lists the associated ID (or number), a description, the maximum number of tokens that can be held in that place, and the constraining resource (if any), together with the maximum number of tokens available at the constraining resource in parentheses.

Table 1: Boundedness property verification

Place ID	Description	Bound	Constraining resource (Max. no. of tokens for resource place)
P6	Patient with the Registration clerk	4	Registration clerk (4)
P9	Patient in the Triage Bay with Triage Nurse and Physician	2	Triage Bay/Triage Nurse (2)
P16	Patient in Waiting room to get into Fast Track	100	None
P53	Patient arriving by ambulance	100	None
P25	Patient in the Fast-track bed with Fast Track Nurse	2	Fast Track Nurse (2)
P86	Patient in the ICU bed	6	ICU Bed (6)

Looking at the table, we find that for places which are constrained by some resource, the bound is the maximum number of tokens available for that resource place. For other places, the bound is 100, the number of tokens the system was initially loaded with. A similar situation is observed with the remaining places, which we have omitted to show here. We therefore find that the net is indeed bounded, and can proceed with simulating the PN model.

3 SIMULATION OUTPUT

The calculations that can be performed on the PN are restricted to the deterministic model. To study the stochasticity of the patient flow process, the PN model is converted into a simulation model. Code templates were developed for places and transitions and these templates were modified to fit the logic of the Petri-net. The initial resource levels used in the simulation model are shown in Table 2.

A discrete-event simulation model of the PN shown in Figure 3 was created and simulated using *AutoMod* for a runtime of 30 days. During the simulation run, the hospital states (markings) are mapped every 10 minutes. Recall that the PN was validated using the *boundedness* property of Petri-nets. The simulation model of the PN is also validated in a similar fashion, by observing the number of available resources at all times. Instantaneous resource availability over the length of the simulation run is shown in Figure 4. The availability of ‘Fast track beds’, ‘Physicians’, ‘ICU beds’ and ‘ED beds’ is shown in the figure as an illustration. Statistics on all other resources, as well as patient-related metrics, can also be readily obtained from the simulation output.

Table 2: Initial resource-token distribution

Places	Resource	Max tokens
P7	Triage Bays	2
P8	Triage Nurse	2
P19	Registration clerk	4
P17	Fast-track Bed	5
P23	Fast track Nurse	2
P55	ED Beds	25
P30	Physician	3
P47	EHS	2
P84	ICU Beds	6
P85	Med/Surg Beds	26

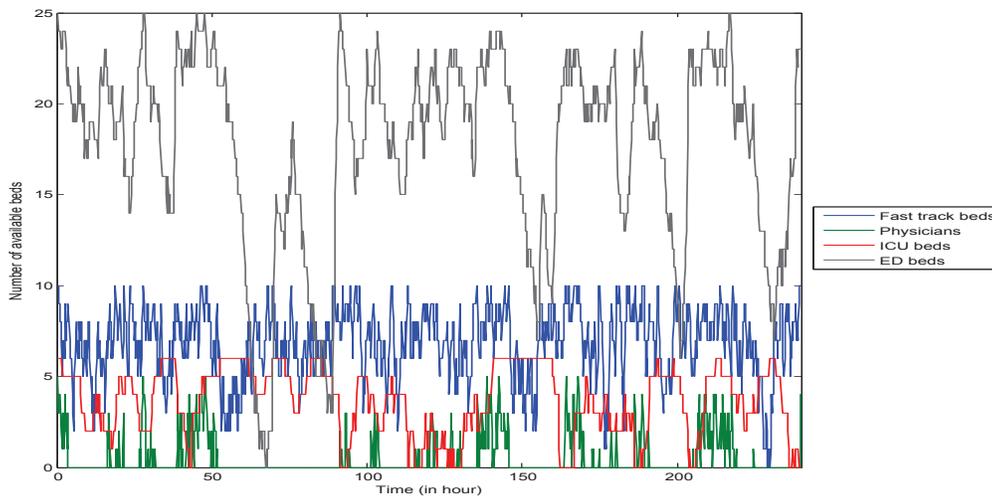


Figure 4: Simulation output: Instantaneous resource availability

It can be seen from Figure 4 that as the system is loaded with patients, the level of resources available varies due to randomness in the patient arrival process and the resource service times. As the system gets loaded and unloaded, it changes states. These states can be classified as good or bad states based on the PN marking.

4 OPTIMAL CONTROL POLICIES

In this section, we discuss the formulation of the underlying stochastic control problem using the PN model of the hospital. The solution of this problem will yield an optimal control policy as its solution. We then describe how this policy can be used by the hospital to manage resources, and prevent ED overcrowding.

Now, when a hospital is overcrowded, access blocking (defined in Section 2) occurs and restricts patient movement in the hospital. The hospital eventually enters a divert state when they are unable to accept and care for patients arriving by ambulance. They discharge patients as quickly as possible to free up resources and exit the divert state. Such situations could be avoided by monitoring the state of the hospital and appropriately bringing in extra resources, such as nurses and beds. These extra resources brought in to cope with overcrowding are referred to as *surge units*. Obviously, bringing in extra resources translates to additional costs, either fixed or proportional, or even both. These additional costs prevent the continuous addition of resources in the hospital. Rather, resources have to be brought in judiciously, to manage the patient flow, while at the same time controlling costs.

So, at certain critical times, surge units need to be added to the hospital to handle overcrowding. When the overcrowding subsides, these surge units need to be removed so that the hospital may remain cost efficient. The hospital therefore needs a control policy that dictates *when, which, and how many* resources should be brought in or removed. This problem of determining the optimal control policy can be cast as a *stochastic impulse control problem* though. Typically in such problems, the controller faces some underlying process, a reward (or cost) structure, and a infinite time horizon. There exist fixed and proportional costs associated with the controls. Exerting the controls moves the underlying process from one point to another and incurs both a fixed cost and a cost proportional to the change in the underlying process. The controller's objective is to maximize (minimize) the total discounted expected reward (cost). The control policies in these problems are characterized by four scalars, d, D, U and u , with $d < D < U < u$. Thus, when the underlying process strikes d , the control increases the process to D , and when the process strikes u , the control decreases the process to U . The region $[D, U]$ is the optimal no-control region. When the process is in this region, no control is exerted. When the process tries to leave this region, a control is exerted to move it back to the no-control region. The reader is referred to [Bensoussan and Lions \(1984\)](#) for a detailed treatment of typical stochastic impulse control problems.

The problem at hand is clearly equivalent to a stochastic impulse control problem. Once we cast our problem in the required form, an optimal control policy can be obtained that will help the hospital avoid overcrowding and entering a divert state. In order to derive the stochastic control problem though, we need a model that accurately describes the underlying process, and on which controls can be exerted.

4.1 Distance to Divert

The underlying process on which the stochastic control problem is based on is derived using the Parikh vector described earlier. Recall that the PV is a set of transition firings that transform the system from marking M_0 to marking M_f . These markings can be thought of as points in m -dimensional place, and represent the current state of the hospital. Using the PV X , we can compute how far M_f is from M_0 in Euclidean space, based on the transitions that need to be fired. The distance γ is given by

$$\gamma = \sum_{j=1}^n \left[\sum_{i=1}^m \phi_{ij}^2 \right]^{1/2} x_j \tag{3}$$

Equation (3) can be written as $\gamma = AX$ for notational convenience, where $A \in \mathbb{R}^n$, with the j^{th} entry of A given by $a_j = \left[\sum_{i=1}^m \phi_{ij}^2 \right]^{1/2}$.

We can now define the concept central to the underlying process we will use for our stochastic control problem, the *distance to divert* Y . Given a definition for divert, we can compute the distance between the current state and the nearest divert state. The definition of divert is context-specific. Recall that in our context, divert occurs when there are no available ED and ICU beds. Therefore, with a given current state, we can compute Y , the distance to the nearest divert state. Clearly, the hospital is in a divert state when $Y = 0$. Given a current state M_0 , the measure Y is obtained by solving the following IP problem:

$$\begin{aligned} & \min_{X \in \mathbb{Z}_+} AX (= Y) \\ & \text{such that } M'_f = M_0 + \phi X, \end{aligned} \tag{4}$$

where the marking M'_f has a zero entry for the places associated with ED beds and ICU beds.

Alternatively, the distance to divert measure can also be computed using the Minimum Parikh Vector (MPV) concept developed in [Criswell et al. \(2007\)](#). The MPV is an $n \times 1$ vector which describes the minimum number of times each transition must fire for the hospital to reach the *nearest* divert state. The MPV is obtained by replacing AX in Equation (4) with $\sum_{j=1}^n x_j$, and solving the IP problem. With the MPV solved for, Y is obtained using Equation (3).

Simulating the PN model representing the hospital allows us to compute and plot Y as it varies with time. The simulation run allows us to observe the state of the hospital at various times. Recall that we have used a 30-day period for our simulation runs. At 10-minute intervals, we capture the

state of the hospital and compute Y . Figure 5 plots Y from 1 simulation run over a 30-day period in.

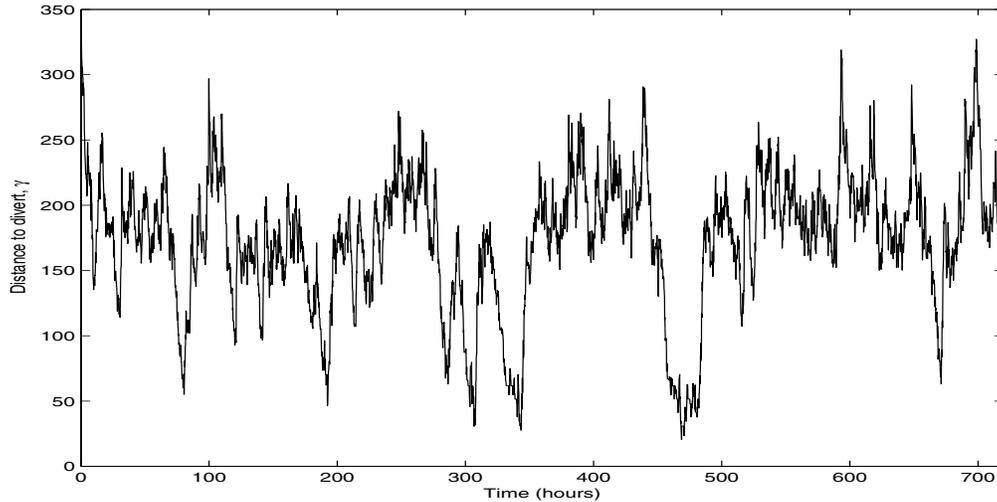


Figure 5: Distance to divert, Y

Looking at Figure 5, we find that Y resembles a stochastic process. We postulate that a geometric Brownian motion of the form

$$dY_t = \mu Y_t dt + \sigma Y_t dW_t, \tag{5}$$

can be used to model the process Y , where $\mu, \sigma \in \mathbb{R}$. Other stochastic processes, such as a standard Brownian motion or a mean-reverting process, may be used to model Y . The measure Y though, represents a distance and as such, has to be non-negative always. For this reason, we have chosen to model Y as a geometric Brownian motion process.

Running the simulation multiple times, we obtain multiple time series for the Y process. Conducting an augmented Dickey Fuller test on these time series, and by analyzing these time series, we find that they are stationary and independent. The geometric Brownian motion process can therefore be used to model the distance to divert measure Y .

4.2 The Stochastic Impulse Control Problem

We now formalize the stochastic impulse control problem which will yield the optimal control policy as its solution. The process Y_t models the hospital's proximity to a divert state at any point in time. Our aim is to derive a control policy based on the evolution of Y_t that will help the hospital avoid overcrowding and divert states. Now, at any point in time, the state of the system can be increased (bring in surge units) or decreased (remove excess resources). These changes carry with them fixed and proportional costs. Let the fixed and proportional costs of increasing (decreasing) the system be K (L) and k (l) respectively. Then, the cost of increasing the state from y to $y + \delta$ is $K + k\delta$ and the cost of decreasing the state from y to $y - \delta$ is $L + l\delta$. In our context, K and L might be administrative costs associated with bringing in or removing resources, for example. The costs k and l might be the cost associated with adding or removing an ED bed. These costs are problem-specific, depending heavily on the hospital system being modeled. There is also a profit function $h(\cdot)$ associated with the system, such that at any point in time, the instantaneous profit is $h(Y_t)$. We take h to be a convex function that achieves its maxima at a finite point.

Let $\{\mathcal{F}_t\}$ be the filtration generated by the standard Brownian motion W_t . Let $0 \leq \tau_1 \leq \tau_2 \leq \dots$ be a sequence of stopping times adapted to $\{\mathcal{F}_t\}$ such that only a finite number of τ_i will occur in any bounded interval with probability one. Then, $\{\mathcal{F}_{\tau_i}\}$ is the σ -algebra of events up to τ_i . Let ξ_i be a random variable which is $\{\mathcal{F}_{\tau_i}\}$ -measurable. An impulse control ψ is defined as $\psi = (\tau_1, \xi_1; \tau_2, \xi_2; \dots)$. The τ 's represent stopping times at which the control is exerted. At time τ_i , the system is increased by ξ_i , i.e., $Y_{\tau_i} = Y_{\tau_i-} + \xi_i$, where Y_{τ_i-} is the value of the underlying process just before the control is exerted. In our context, this would imply bringing in surge units to increase the distance to divert by ξ_i . Note that ξ_i may be negative. Now, the total discounted expected profit associated with using

some control ψ is

$$\mathcal{J}(\psi) = \mathbb{E} \left\{ \int_0^\infty e^{-\theta t} h(Y_t) dt - \sum_i e^{-\theta \tau_i} [\mathbf{1}_{\xi_i > 0} \cdot (K + k|\xi_i|) + \mathbf{1}_{\xi_i < 0} \cdot (L + l|\xi_i|)] \right\}, \quad (6)$$

where $\mathbf{1}$ represents the indicator function and θ is the discount rate. As mentioned before, we are interested in determining the impulse control that will result in the greatest total expected discounted profit (value function), $V = \sup_\psi \mathcal{J}(\psi)$.

The dynamic programming principle is used in conjunction with Itô's lemma to obtain the differential equation problem, termed the Hamilton-Jacobi-Bellman (HJB) equation, from the stochastic control problem formulation. Say the system is uncontrolled for an infinitesimal amount of time, from 0 to Δt , before the optimal control policy is used. The resulting value function cannot be greater than the optimal value function, by definition of the optimal value function. Using this idea together with Itô's lemma, we obtain the following inequalities which V has to solve:

$$\frac{1}{2} \sigma^2 V''(y) + \mu V'(y) - \theta V(y) + h(y) \leq 0, \quad (7)$$

$$\sup_{\xi > y} \{V(\xi) + K + k(\xi - y)\} - V(y) \leq 0, \quad (8)$$

$$\sup_{\xi < y} \{V(\xi) + L + l(\xi - y)\} - V(y) \leq 0, \quad (9)$$

where superscripted dashes represent differentials and $Y(0) = y$.

Solving the inequalities in Equations (7) - (9) will provide us with both the optimal control $\hat{\psi}$ and the optimal value function V . This type of problem is known as a *free-boundary problem*, as the optimal controls, or boundaries (d , D , U , and u) are not known *a priori* and have to be solved for as part of the problem. This free-boundary problem can be solved using the the moving-boundary techniques developed in Chockalingam and Muthuraman (2009) and Feng and Muthuraman (2009). The scalars d , D , U , and u obtained by solving this problem will form the control policy the hospital should use.

4.3 Control Policy Translation

The control policy applies to the distance to divert measure. Thus, if $Y_t = d$, surge units should be brought in such that the distance to divert increases to D . Conversely, if $Y_t = u$, excess resources should be removed such that the distance to divert decreases to U . We need to determine the actions that need to be carried out to bring about the change in the distance to divert. For example, if we need to move Y_t from d to D , do we bring in an additional bed or an on-call physician? Is one bed sufficient, or do we need multiple beds? The control policy, as it is, provides limited information with respect to actions that need to be carried out. We need to be able to determine which resources have to be brought in or removed, and by how much, based on this control policy which is given in terms of the distance to divert.

Say the current state is M_0 . With this marking, we obtain a distance to divert Y_t , with the nearest divert state being M_f . Now say $Y_t = d$. This means we need to bring in surge units (increase tokens in some places) to obtain a new marking M'_0 that has a distance to divert of $Y_t = D$. Let $M'_0 = M_0 + y$. Then, $y \in \mathbb{Z}^m$ represents the change in the number of tokens. We define a decision vector $Z = (x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_m)^T$ which consists of the PV and the vector y . There might be an upper bound on surge units due to limited capacity. For example, there might only be 3 on-call physicians. Let the vector $C \in \mathbb{Z}^m$ represent this upper bound. The vector C can also be used to prevent tokens being added to places where changes cannot be made. An example of such a place would be 'patient with registration clerk'. There would be a zero entry in the vector C that is associated with this place to reflect that no changes may be made. With this setup in place, we need to solve the following IP problem to determine the changes in resources necessary to increase Y_t to D :

$$\min_{Z \in \mathbb{Z}^{n+m}} A_1 Z, \quad (10)$$

such that

$$M_f = M_0 + A_2Z + \phi A_3Z, \tag{11}$$

$$AA_3Z \geq D, \tag{12}$$

$$A_2Z \leq C, \tag{13}$$

$$A_1Z > \sum_{j=1}^n x_j, \tag{14}$$

$$Z_i \geq 0 \text{ for } i \in [1, m], \tag{15}$$

where A_1 is a $1 \times (n + m)$ vector with n 1's followed by m 0's, A_2 is a $m \times (n + m)$ matrix and A_3 is a $n \times (n + m)$ matrix, both consisting of 1's and 0's. Equation (10), the objective function of the IP problem, is concerned with minimizing the Parikh vector. The first constraint of the IP problem, Equation (11), is the state equation between the new marking M'_0 and the nearest divert state M_f . The second constraint, Equation (12), states that the new distance to divert has to be at least D . The vector A in this equation is the vector previously defined in Equation (3). Equation (13) handles the surge capacity constraint, and the condition that tokens may not be added to all places. Equation (14), states that the number of transitions required to enter the divert state from the new marking M'_0 has to be greater than the number of transitions required to enter the divert state from the previous marking M_0 . Equation (15) simply states that the number of transitions cannot be negative. If $Y_y = u$, an IP problem with similar structure to Equations (10) - (15) needs to be solved to determine the actions that need to be carried out.

Therefore, at each time there is a change in the state of the hospital, the distance metric Y needs to be computed and compared against the (d, D, U, u) control policy. If $Y \leq d$ or $Y \geq u$, the IP problem described in Equations (10) - (14) needs to be solved to determine the number and types of resources that have to be brought in or removed, to prevent hospital ED from overcrowding, and to avoid the hospital from entering a divert state.

5 A NUMERICAL EXPERIMENT

We now illustrate the approach proposed in this paper with a quick numerical experiment, using the PN model described in Section 2. To conduct the experiment, the parameters for the geometric Brownian motion process, μ and σ , together with the profit function h , are needed. The parameters μ and σ are computed directly from the time series obtained for the process Y from the simulation runs and are found to be 1 and 0.0016, respectively.

For the profit function, cost and revenue estimates for Indiana hospitals were used in the simulation model. At each 10-minute interval, the instantaneous cost and revenue were estimated based on the distribution of resources and patients in the system at that time. The instantaneous profit was then calculated using these estimates. The distance to divert at that time is also recorded. Plotting instantaneous profit against distance to divert provides us with instantaneous profit as a function of distance to divert. Now, a single distance to divert could correspond to multiple markings, which implies that a single distance to divert could result in different instantaneous profits, depending on the markings. For each distance to divert, we take instantaneous profit to be the minimum of all the instantaneous profits associated with that distance. Our aim in choosing the profit in this fashion is to determine an optimal control policy that is robust, i.e., we are interesting in maximizing profit in the worst-case scenario. With the profit function so chosen, we fit a quadratic curve to obtain the instantaneous profit function shown in Figure 6. We can now compute the (d, D, U, u) control policy described in Section 4.2 using the moving-boundary technique. Omitting details of the computation, we find the following: $d = 37.264$, $D = 44.949$, $U = 79.708$ and $u = 198.915$. We can now use the control policy translation described in Section 4.3 to determine what changes in resources need to be made in order to change the distance to divert metric. For example, say the current marking yielded $Y = 37.264$. The hospital has to bring in resources such that the new marking yields $Y = 44.949$. Solving the IP problem described above, we find that the hospital needs to add 1 triage nurse, 1 ED nurse, and 1 ED bed in order to effect this change.

Now we simulate and compare an uncontrolled and a controlled hospital system based on two metrics; profit at the end of the 30 days, and the leave without being seen (LWBS) rate. The LWBS rate captures the number of patients who check-in to the hospital for treatment, and leave without seeing a physician due to long wait times. Table 3 shows performance measures obtained with each system.

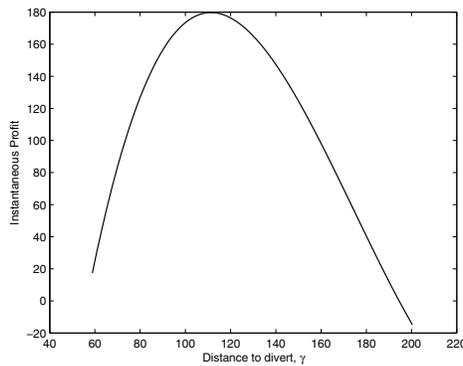


Figure 6: Instantaneous profit function

Table 3: Comparison of performance measures

Performance measure	Uncontrolled	Controlled
LWBS rate (%)	2.6	1.9
Profit (\$)	653,179	700,220

As can be seen, using the optimal control policy yields significant performance improvement. By optimally managing resources, the hospital accrues a larger profit, and also improves patient welfare, as signified by the lower LWBS rate.

6 CONCLUSION

The problem of hospital overcrowding is a serious one and has significant effects, one of which is hospitals entering divert a state. In this paper, we propose a novel approach to derive control policies with which hospitals could prevent overcrowding, and thereby prevent divert states from occurring. We have developed the *distance to divert* measure to describe the hospital’s proximity to entering a divert state. Using a stochastic process to describe this measure, we have derived a stochastic impulse control problem, which can be solved numerically using moving-boundary techniques. The optimal control policy obtained as a solution to this problem dictates when the hospital should add or remove resources (such as beds and nurses), and also how many of these resources should be added or removed. The control policy does not explicitly state which resources should be added or removed, rather only prescribing a change in the distance to divert. We therefore derive an integer-programming problem, based on the required change in the distance to divert measure, which explicitly states the change in resource levels necessary to change the distance to divert. The use of these optimal control policies leads to significant improvements in performance, as illustrated by a numerical experiment. This work focuses on developing control policies for a single hospital system. In reality, hospitals in a vicinity coordinate their divert strategies. An obvious direction for future research is therefore to develop control policies for hospitals that coordinate their strategies.

7 ACKNOWLEDGEMENTS

We would like to thank Manan Javeri for providing feedback and computational assistance on this paper.

REFERENCES

Bensoussan, A., and J. L. Lions. 1984. *Impulse Control and Quasi-variational Inequalities*. Gauthier-Villars.

- Burt, C. W., L. F. McCaig, and R. H. Valverde. 2006. Analysis of Ambulance Transports and Diversions Among US Emergency Departments. *Annals of Emergency Medicine* 47:317–326.
- Chockalingam, A., and K. Muthuraman. 2009. American Option Pricing under Stochastic Volatility. *Working paper*.
- Committee on Pediatric Emergency Medicine 2004. *Overcrowding Crisis in Our Nation's Emergency Departments: Is Our Safety Net Unravelling?* Committee on Pediatric Emergency Medicine.
- Criswell, M., I. Hasan, R. Kopach, S. Lambert, M. Lawley, D. McWilliams, G. Trupiano, and N. Varadarajan. 2007. Emergency Department Divert Avoidance Using Petri Nets. *Proceedings of IEEE Second International Conference on System of Systems Engineering (ICSoSE)*:1–6.
- Feng, H., and K. Muthuraman. 2009. A Computational Method for Stochastic Impulse Control Problems. *Working paper*.
- Glushak, C., T. R. Delbridge, and H. G. Garrison. 1997. Ambulance Diversion. Standards and Clinical Practices committee, National Association of EMS Physicians. *Prehospital emergency care* 1:100–103.
- Grumbach, K., D. Keane, and A. Bindman. 1993. Primary Care and Public Emergency Department Overcrowding. *American Journal of Public Health* 83:372–378.
- Institute of Medicine 2006. *The Future of Emergency Care in the United State Health System*. Institute of Medicine.
- Jenson, J. 2003. Going with the Flow: Tracking System Helps Midwest Hospital Streamline Patient Flow and Lower Emergency Room Divert Rate. *Health Management Technology* 24:43–45.
- Kokler, A. 2008. Process Modeling of Emergency Department Patient Flow: Effect of Length of Stay on ED Diversion. *Journal of Medical Systems* 32:389–401.
- McConnell, J. K., F. R. Christopher, and D. Mohamud. 2006. Ambulance Diversion and Lost Hospital Revenues. *Annals of Emergency Medicine* 48:702–710.
- Murata, T. 1989. Petri Nets: Properties, Analysis and Applications. *Proceedings of the IEEE* 77:541–580.
- Pham, J. C., and R. Patel. 2006. The Effects of Ambulance Diversion: A Comprehensive Review. *Academic Emergency Medicine* 13:1220–1227.

AUTHOR BIOGRAPHIES

ARUN CHOCKALINGAM is a Visiting Assistant Professor in the School of Industrial Engineering at Purdue University. His research interests lie in the areas of decision-making under uncertainty, stochastic control and stochastic optimization. In particular, he is interested in formulating decision-making problems as stochastic control problems, and constructing computational methods to solve these problems and obtain optimal policies. He is currently working on problems in two application areas, healthcare systems engineering and financial engineering.

KRISHNA JAYAKUMAR is a graduate student in the School of Industrial Engineering at Purdue University. He will receive his Masters degree in May 2010. His research interests lie in the area of simulation modeling and stochastic optimization, applied to healthcare systems.

MARK LAWLEY is a Professor in the Weldon School of Biomedical Engineering at Purdue University. Before joining Biomedical Engineering in 2007, he served nine years as Assistant and Associate Professor of Industrial Engineering, also at Purdue, two years as Assistant Professor of Industrial Engineering at the University of Alabama, and he has held engineering positions with Westinghouse Electric Corporation, Emerson Electric Company, and the Bevill Center for Advanced Manufacturing Technology. As a researcher in academics, he has authored approximately 100 technical papers including book chapters, conference papers, and refereed journal articles, and has won three best paper awards for his work in systems optimization and control. In January 2005, he was appointed Regenstrief Faculty Scholar in support of Purdues Regenstrief Center for Health Care Engineering. He is particularly interested in developing optimal decision policies for system configuration and resource allocation in large healthcare systems. As a Regenstrief Scholar, he has focused on research initiatives with Wishard Hospital, the Regenstrief Institute of Indianapolis, the Richard L. Roudebush Veterans Administration Medical Center, Ascension Health, St. Vincent Hospitals, and the American College of Physicians. His recent research has focused on problems in scheduling in primary care and chemotherapy clinics, financial sustainability for the patient centered medical home, diagnosis-based modeling of hospital patient flow, stochastic control approaches for emergency room divert and hospital surge capacity implementation, optimal fistula monitoring and maintenance for hemodialysis, hospital

stockpiling and optimal control theory approaches for non-pharmaceutical interventions for pandemic influenza. His research has been supported by the National Science Foundation, Union Pacific Railroads, Consilium Software, General Motors, Ascension Health, the Indiana State Department of Health, the Regenstrief Foundation, the St. Vincent Ministry, and many others. He received the PhD in Mechanical Engineering from the University of Illinois at Urbana Champaign in 1995 and is a registered Professional Engineer in the State of Alabama.