

SENSITIVITY ANALYSIS OF AN ICU SIMULATION MODEL

Theologos Bountourelis
David Eckman
Louis Luangkesorn
Andrew Schaefer

Department of Industrial Engineering
University of Pittsburgh
1048 Benedum Hall
3700 O'Hara St.
Pittsburgh, PA 15261

Spencer G. Nabors
Gilles Clermont

Department of Veterans Affairs Medical Center
University Drive C
Pittsburgh, PA 15240

ABSTRACT

The modeling and simulation of inpatient healthcare systems comprising of multiple interconnected units of monitored care is a challenging task given the nature of clinical practices and procedures that regulate patient flow. Therefore, any related study on the properties of patient flow should (i) explicitly consider the modeling of patient movement rules in face of congestion, and (ii) examine the sensitivity of simulation output, expressed by patient delays and diversions, over different patient movement modeling approaches. In this work, we use a high fidelity simulation model of a tertiary facility that can incorporate complex patient movement rules to investigate the challenges inherent in its employment for resource allocation tasks.

1 INTRODUCTION

An Intensive Care Unit (ICU) is a limited-capacity, resource-intensive unit designed to deliver intensive monitored care for a diverse set of high acuity clinical conditions. The daily cost of ICU operations is higher than that of standard care nursing units due to the requirement for higher nurse-patient ratios, invasive bedside procedures, and expensive monitoring equipment (Groeger et al. 1993). The daily ICU operations are driven by a number of factors, such as clinical practices, unit occupancy, staff scheduling, the operating status of other units of the host healthcare facility, and resource management decisions.

At VA Pittsburgh Healthcare System (VAPHS), a large tertiary care Veterans Healthcare Administration (VHA) facility, the most intensive units in terms of human and technological resources include the medical intensive care unit (MICU), the coronary care unit (CCU), and the surgical intensive care unit (SICU). Step-down units are somewhat less intensive, while the monitored and telemetry units require still a lesser level of resources. Regular hospital beds are the least resource intensive. Patients can move from more intensive to less intensive levels of care (and the reverse) as their condition changes over the course of their hospital stay. However, due to capacity restrictions, patients may not be able to move to the requested unit and level of care. We will refer to this phenomenon as patient blocking. Furthermore, in the event of bed unavailability at VAPHS, patients may have to be diverted to another regional facility when medically able to do so. We will refer to this phenomenon as patient diversion.

Patient blocking is an important patient flow characteristic with an impact on operating costs and the timely delivery of healthcare services. Blocked patients can occupy resources meant for higher levels of care than clinically needed, thereby increasing operating costs without improvement on healthcare outcomes. We assume that a patient blocked in a unit with a higher level of care does not experience worse outcomes,

although this may not be true (Teng et al. 2009). Furthermore, patient blocking might propagate and disrupt operations in other related units. For example, patient blocking at an ICU may delay bed availability which, in turn, may affect the operations at the surgical theater. In some cases, patients with illnesses necessitating ICU level of care may be forced to receive intensive care in non-monitored beds, or other more stable patients may be discharged prematurely incurring a significant cost.

Patient diversion is another important patient flow statistic with a significant impact on VAPHS patient expenditures. In reality, the majority of ICU patients get diverted before their arrival at the VAPHS facility. An operator is responsible for orchestrating the transfer of the patient from the off-site facility to the VAPHS. Given the patient's medical condition and a waiting time estimate from the Emergency Department (ED), the operator decides whether the patient should be diverted to another facility. There is no generally accepted estimate on the time above which a patient is diverted and often diversion decisions are case specific. Furthermore, given the large risks and expense of diverting patients from the ICU, there is a strong bias against diversions whenever possible.

Identifying the factors that drive patient blocking and diversions is a challenging task given the inherent complexity of the healthcare facilities where ICUs operate. Given that most ICU patients experience multiple unit transitions, it is important to adopt a modeling approach that will consider the ICUs as an integral part of a larger healthcare facility. While most researchers concentrate their modeling effort on some individual area of the hospital, large scale simulation models are relatively rare. This may be due to the unique challenges pertaining to increased data requirements and the need for extensive clinical personnel involvement. Moreover, given the size of the facility and the complexity of rules (often case specific) dictating patient movement, it is important to develop simulation code able to accurately represent patient flow under a large number of model inputs pertaining to (i) resource allocation schemes, and (ii) possible sets of patient movement rules. This is particularly important for the sensitivity analysis of model output since the testing of various interventions will have to be performed over a variety of patient movement rules and policies that attempt to approximate, on average, the actual patient flow dynamics.

In light of the above, the work presented herein concerns our extensive simulation work on a resource allocation project on the VAPHS. Although the work is specific to the VAPHS facility we believe that the depth of the analysis and the empirical findings can provide useful feedback to future large scale healthcare projects. The rest of this paper is structured as follows: In Section 2 we give an overview of the challenges inherent in the modeling of ICUs. Next, in Section 3 we describe the facility that is the center of our analysis. Subsequently, in Section 4 we present the simulation model developed with a focus on patient blocking and diversions. In Section 5 we describe the simulation work and the sensitivity analysis of our results. Finally, in Section 6 we summarize our work and outline a number of directions for future research.

2 CHALLENGES IN ICU MODELING

Patient Blocking, Diversions, and Other Issues: In this section, we take a closer look at the factors that make modeling, analyzing, and managing ICUs challenging. Clinical units and particularly ICUs are specialized units that admit critically ill patients and may vary significantly across different hospitals. Some hospitals utilize ICUs that are dedicated to specific patient types. For example, a single hospital might possess a medical ICU as well as a surgical ICU that only handles surgery patients. Smaller hospitals may have only a single ICU that handles all patient types and illnesses in which case prioritizing patients is a critical issue. Yet still other hospitals may have ICUs dedicated to a subset of patients, for example a neurosurgical ICU. Irrespective of the configuration, all ICUs have only a finite (and relatively small) number of beds. Several types of illnesses can send a patient to the ICU, and depending on the type and acuity of the illness, the regimen prescribed to the patient, and physiological differences between patients, it may be difficult to model a "typical" ICU patient's flow through the system. Even patients suffering from the same illness can have vastly different experiences in the ICU. This inherent heterogeneity makes it difficult to assess an "average" patient's length of stay or the overall impact on other units of the hospital.

As the most resource-intensive unit in a healthcare facility, the ICU cannot be analyzed in isolation from other upstream or downstream inpatient units and departments. The majority of ICU patients follow a similar flow through the system. As soon as the patient is medically ready to leave the ICU, a request is made by the attending physician to transfer the patient to a telemetry, step-down, or regular hospital bed that is less intensive than the care received in the ICU. From the step-down bed, the patient is usually transferred to a regular hospital bed from which he or she is ultimately discharged. However, patients may have trouble getting into, and out of, the ICU due to the blocking phenomenon.

Blocking in the ICU can happen in the following ways: First, an arriving patient who requires intensive care may be denied admission to the ICU due to bed unavailability (i.e., all ICU beds are occupied, or the unoccupied beds may not be suitable for the new patient). In such cases, an existing ICU patient may be (prematurely) transferred to a lower level of care to accommodate the new patient; otherwise, the arriving patient must be diverted to another hospital or treated at a lower level of care, thereby delaying the critical care they need. In extreme cases, patients must receive ICU-level care in non-monitored beds. For instance, equipment and personnel may need to be moved to the emergency department to administer ICU-level care and stabilize the patient until an ICU bed is available. Blocking has devastatingly negative effects on the operation and flow of the ICU and other hospital units. For example, unavailability of an ICU bed may cause ED overcrowding, surgery cancellations for procedures that require post-surgery intensive care, and inadequate delivery of high acuity care.

The second type of blocking is experienced by existing ICU patients who cannot be transferred out of the ICU due to unavailability of beds at lower levels of care. That is, patients who are medically able to leave the ICU might be forced to stay there until a downstream bed becomes available. Obviously, this type of blocking serves to exacerbate the former type as fewer ICU patients can be admitted. It is important to note that patients who are prematurely transferred out of the ICU may be forced to re-enter it if their medical condition worsens at any downstream stage. Blocking also stems from a shortfall of external nursing care facilities, or due to institutional rules and procedures. For example, if transfers to external facilities are prohibited during the weekend, step-down and hospital beds can all become occupied over the weekend, thereby blocking the transfer of ICU patients.

Clinical Practices and Procedures: Modeling an ICU (and its interacting units) is complicated by clinical practices and procedures that are often employed to control the flow of patients into and out of the ICU. These practices can be viewed as “rules of thumb” for daily ICU operations, and they can vary dramatically across hospitals or across doctors within the same hospital. For example, a medical ICU (MICU) might accept post-surgery patients if no beds are available in the surgical ICU (SICU) or in the post-anesthesia room (PAR) of the hospital. As noted earlier, some ICUs may discharge patients early if the ICU occupancy reaches its peak. Patient demographics, proximity to other clinical facilities, and the level of outpatient care available may also affect the length of stay and the patient mix. Ultimately, a universal set of rules for transferring or discharging patients does not exist, thereby complicating the task of modeling these important dynamics.

The simulation modeler must decide which rules to include and which to exclude, for each rule added to the model can potentially add one or more control parameters that need to be estimated from data. For example, it is clear that not all surgery patients are diverted to the MICU following surgery, but the exact proportion that are diverted may be difficult to estimate, even for clinical staff members who are involved in day-to-day decision making. Out of necessity, these parameters can be assigned initial values based on an educated guess and subsequently updated when the (simulation or analytical) model is calibrated. But if the number of unknown parameters is significant, the process of calibrating the model can be nontrivial and very time-consuming. In the next section we provide a description of the clinical facility that is the focus of this work.

3 FACILITY

The VAPHS hospital serves as an acute care facility and has 146 inpatient beds distributed among medicine, surgery, neurology, cardiology, and critical care. The VAPHS facility is comprised of several ICUs, inpatient units, and departments. Table 1 summarizes the units (modules) explicitly modeled. After examining patient

Table 1: VAPHS Unit Bed Capacity and Staffing

Inpatient Units Explicitly Modeled		
Unit Name	# of Beds	Nurse-Patient Ratio
Medical ICU (MICU)	8	1:2
Surgical ICU (SICU)	9	1:2
Coronary Care Unit (CCU)	12	1:2
Step Down Unit (SDU)	9	1:3
Monitored Medical unit	13	1:4
Monitored Surgical unit	12	1:4
Non-Monitored Medical and Surgical units	71	1:5

data files, we compile a list of locations from which patients apply for admission to the inpatient units listed in Table 1. Subsequently, with the help of the Subject Matter Experts (SME), we aggregate those locations into five *patient sources* listed below:

1. Emergency Department (ED)
2. Operating Room (OR)
3. Post-Anesthesia Recovery (PAR) Room
4. Other Floors
5. Direct Admissions

4 SIMULATION MODEL

4.1 Simulation Software

The simulation model is coded in the OMNeT++ simulation package, an event-driven C++ simulation library designed for building network simulators (OMNeT++ Community 2001) for telecommunication applications. OMNeT++ provides versatile tools for configuring simulation models and collecting simulation output. This platform adopts a modular approach for building simulation models that makes it easy to implement complex scenarios and configurations. OMNeT++ allows the separation of model and input data through the use of a configuration text file. Furthermore, we can implement complex patient movement policies by describing them as text in the aforementioned configuration file. As a result of the above, the use of OMNeT++ can result in a simulation model that can be easily configured for thousands of simulation runs which can then be performed in parallel on a computer cluster. This is a necessary requirement for evaluating model inputs and performing sensitivity analysis as we will see in the subsequent chapters of this article.

4.2 Modeling the Facility

The ICU patient flow simulation model used in this work is an updated version of the model presented in Bountourelis et al. (2011). The simulation model is comprised of a set of modules corresponding to the patient sources (Items 1-5) and units (items of Table 1). The modules have a number of attribute variables pertinent to the unit or department they represent (e.g., number of beds).

An integral feature of the model is the use of a database containing patient instances that correspond to real de-identified cases of patients admitted and treated at the VAPHS. The patient instances were compiled

from patient movement data retrieved from the VAPHS data warehouse. Each patient instance is described by (i) an arrival source (Items 1-5), (ii) a sequence of visited units and departments (items of Table 1 and Items 1-5), and (iii) a corresponding sequence of length of stay (LOS) values. Furthermore, using the initial database entries for every patient we estimated a series of patient arrival rates. That is, for every patient source, hour of the day, and day of the week, we were able to estimate a patient arrival rate. For more details regarding the model structure and the processing of VAPHS patient movement data we refer the reader to Bountourelis et al. (2011).

The program logic underlying patient generation and movement is as follows: Every hour of the day and day of the week, the program retrieves the corresponding arrival rate for each patient source and generates a number of patient arrivals using the Poisson distribution. Subsequently, the program queries the patient database and randomly selects, with replacement, the corresponding number of patients from each source. The patient entities are inserted into the generating module in order to initiate their transition through the different modules of the simulation program. When a patient's assigned LOS expires, a request is made for a bed in the unit listed next in the unit sequence. In case there is no available bed in the next assigned unit, the patient entities can be moved to a different simulation module (representing a unit or department) according to a selected set of movement rules. After the next LOS expires, the patient resumes the unit sequence. If no bed is found, the patient is declared blocked until a suitable bed becomes available. When the patient entity completes the assigned unit sequence, the patient is discharged and output data is extracted.

4.3 Patient Movement Rules

As discussed in Section 2, patient flow is regulated by a collection of rules (often undocumented) that can vary even between doctors of the same hospital. Generally speaking, the transferring of patients between units is based on the grouping of monitored units by level of care. Units with similar levels of care (or higher) are often considered as alternatives for blocked patients. For example, a medical telemetry unit might accept post-surgery patients if no beds are available in the surgical telemetry, or the CCU might accept patients that otherwise would be admitted at the congested MICU. While a collection of these rules are an important component for the simulation work, it is impossible to account for all of the possible scenarios that might be encountered. Therefore, it is important for a large scale simulation model to read and implement sets of patient movement rules in an attempt to capture the underlying trends of patient flow. The objective is to run the model over different sets of rules so that the modeler can conduct sensitivity analysis and examine the robustness of resource allocation schemes under a variety of different scenarios. In this simulation work, we consider two approaches for modeling patient movement rules:

Explicit Rules: In this approach, patient movement rules are described as a set of *directives*. A set of directives is illustrated in Table 2 where every row is a collection of alternatives for patients who cannot find a bed at their next assigned unit. The first column denotes the patient's current location and the second column his or her assigned unit. The rest of the entries denote the units (in order of preference) where the simulation program will look for an available bed.

Generic Rules: In this modeling approach blocked patients seek an available bed at a unit of the same level of care. If a bed cannot be found, they seek a bed at a unit of an immediately higher level of care with a pre-specified probability. If beds are available in more than one unit of immediately higher level of care, the patient is transferred to the one with the most available beds.

Patient Diversions: In reality, the majority of diversion decisions for ICU patients are made before their arrival at the VAPHS facility. On the other hand, telemetry patients usually have to wait for an amount of time (depending on their condition) before being diverted to another facility. In this model we adopt the following rules governing patient diversions: (i) for an ICU patient arriving from the ED

or Direct Admissions, if a bed cannot be found using the patient movement directives, he or she is diverted immediately, and (ii) for a telemetry patient arriving from ED or Direct Admissions, if a bed cannot be found using the patient movement directives, he or she is diverted after a pre-specified waiting time.

Patient Blocking: Patient blocking can occur for patients residing in patient sources (Items 1-5) and units (items of Table 1) if their next assigned unit (items of Table 1) has no available beds. If no bed is found at the alternative units, or if there are no alternatives at all, the patient is declared blocked. Patient movement directives attempt to capture a wide range of possible scenarios. In some instances a post-surgery patient might be moved to an MICU instead of the SICU, or another patient initially assigned to a monitored unit might be moved upstream to a higher level of care.

Table 2: Movement Directives

Explicit Patient Movement Rules					
Current	Destination	Alternatives	Current	Destination	Alternatives
ED	MICU	CCU, SICU	PAR	MICU	SICU, CCU
ED	SICU	MICU, CCU	PAR	SICU	MICU, CCU
ED	CCU	MICU, SICU	PAR	CCU	SICU, MICU
ED	SDU	SICU, MICU	PAR	SDU	SICU, MICU
ED	Mon. Med.	Mon. Surg., SDU	PAR	Mon. Med.	Mon. Surg.
ED	Mon. Surg.	Mon. Med., SDU	PAR	Mon. Surg.	Mon. Med.
Outside	MICU	CCU, SICU	MICU	SDU	None
Outside	SICU	MICU, CCU	MICU	Mon. Med.	Mon. Surg.
Outside	CCU	MICU, SICU	MICU	Mon. Surg.	None
Outside	SDU	SICU, MICU	SICU	SDU	None
Outside	Mon. Med.	Mon. Surg.	SICU	Mon. Med.	Mon. Surg.
Outside	Mon. Surg.	Mon. Med.	SICU	Mon. Surg.	SDU
OR	MICU	SICU, CCU	CCU	SDU	None
OR	SICU	MICU, CCU	CCU	Mon. Med.	Mon. Surg.
OR	CCU	SICU, MICU	CCU	Mon. Surg.	SDU
OR	SDU	SICU, MICU	SDU	MICU	SICU, CCU
OR	Mon. Med.	Mon. Surg.	SDU	SICU	MICU, CCU
OR	Mon. Surg.	Mon. Med.	SDU	CCU	MICU, SICU
SDU	Mon. Med.	Mon. Surg.	SDU	Mon. Surg.	None

4.4 Measures of Interest

In the following we introduce our approach to evaluate the potential cost impact of patient blocking and diversions through the use of relevant measures that capture the overall monetary effects of these phenomena. We will refer to them as the cost of blocking and the cost of diversions.

Blocking and Associated Cost: The driving factor behind the definition of blocking cost is that blocked patients occupy expensive resources no longer clinically required; consequently, operating costs increase without improvement in healthcare outcomes. We retrieve a daily operating cost for each of the units included in the simulation model and we calculate the cost of blocking as follows: For each unit, the simulation model records all the experienced blocking times and the corresponding destination units. Subsequently, we obtain the monetary impact of the particular blocking sample by multiplying the dif-

ference in daily operating expenses between the current and the destination unit with the corresponding blocking time. Using the sum of blocking costs we calculate an average monthly blocking costs for each unit.

Diversions and Associated Cost: From an analysis of VAPHS patient diversion data, we conclude that there are two types of diversions that represent the most substantial cost to the hospital: ICU and monitored diversions. ICU diversions are instances when a patient cannot be admitted to an intensive care unit (MICU, SICU, or CCU) and is diverted to another facility for immediate care. Monitored diversions differ from ICU diversions in that the patient is seeking admission to a lower-level monitored unit (monitored medical or monitored surgical unit). ICU diversions are the more expensive type with an estimated cost per episode of \$21,000 versus \$6,600 for monitored diversions. For our analysis, monthly average diversion costs are calculated by multiplying the average number of ICU and general diversion episodes by their respective cost estimates and summing.

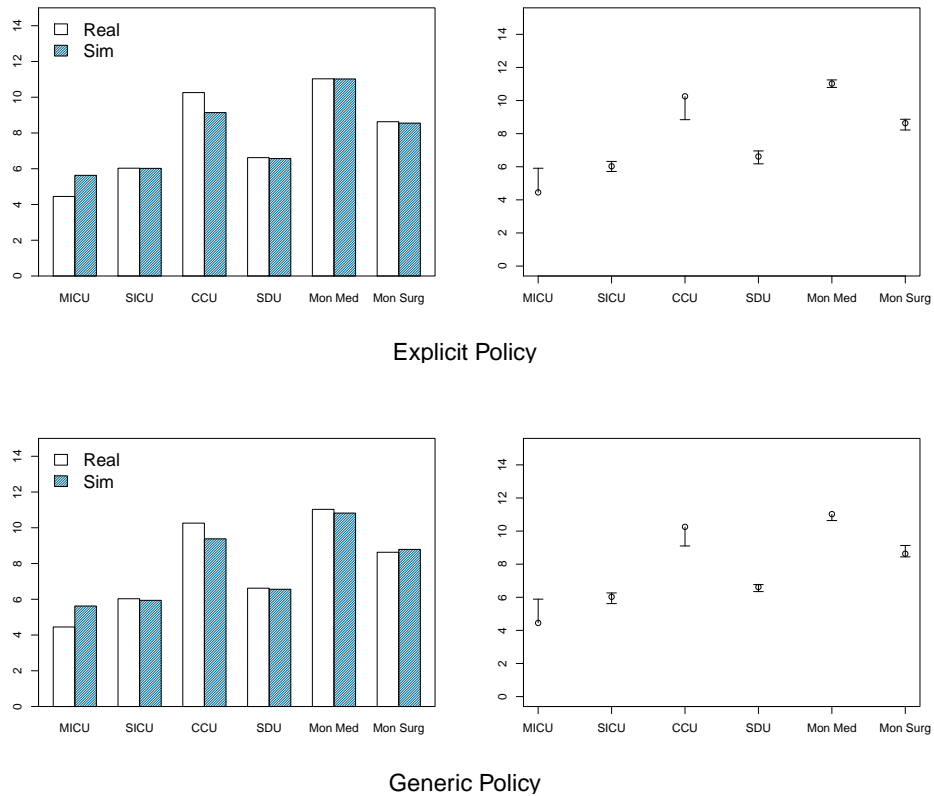


Figure 1: Real vs Simulated occupancies for “Explicit” and “General” patient movement rules.

5 CASE STUDY

The case study presented in this section concerns the construction of an additional floor at the VAPHS and its impact on the blocking and diversion costs under different numbers and mixtures of beds. We allow the additional space to include a mixture of beds belonging to the ICU and monitored units of Table 1. It is evident that the evaluation of all possible configurations for a given number of beds can be a computationally demanding proposition. For example, the addition of nine beds of (potentially) six different kinds can result in a total of 2002 configurations. Furthermore, the model will have to be tested over different patient

movement rules thus increasing significantly the total number of simulation runs. Therefore, to alleviate the computational effort we consider possible configurations with beds distributed between a maximum of two units; this results in *pure* and *pairwise* configurations for the total number of added beds (i.e., a pure configuration consists of a single bed type whereas a pairwise configuration consists of some combination of two bed types, for example, SDU and monitored beds).

To evaluate the sensitivity of simulation output with respect to patient movement rules, we perform our simulation runs over two sets of patient movement rules derived from the aforementioned modeling approaches, the explicit and generic movement rules. The first set is illustrated in Table 2 and will be referred to as “Explicit”. The second set will be referred to as “General”. The starting point for our analysis is that both of the considered patient movement rules result in a *valid* simulation model with respect to the average unit bed utilizations. As we illustrate in Figure 1, the average unit utilizations resulting from both models are close to the real bed utilizations observed from the analysis of patient flow data. In particular, for the model implementing the “Explicit” set of rules, the real utilization falls within the confidence interval produced by the model for all units except MICU and CCU. However we should note that the combined MICU and CCU model utilizations agree with the corresponding combined real utilizations. This is due to the fact that patient movement rules hardly discriminate between MICU and CCU thus leading to the observed imbalance. Similar observations pertain to the model with the “General” set of rules.

The scope of this resource allocation case study can be summarized as follows: (i) examine the marginal difference of different bed levels with respect to the cost of blocking and diversions, (ii) highlight the trade-offs of each combination of bed type and level, i.e., increases and/or decreases of blocking and diversion costs that are observed by varying the combination of different bed types, and (iii) perform a sensitivity analysis that will explore the robustness of the aforementioned observations over the two sets of the considered patient movement rules.

5.1 Simulation Output Analysis

We simulated all possible bed combinations resulting from the addition of 3 to 13 beds of pure and pairwise combinations of ICU, SDU and monitored units. Each configuration was simulated for 120 replications over a period of 12 months. The simulation output consisted of the following measures (i) an average monthly blocking cost, and (ii) the average monthly cost of ICU and general diversion episodes. An illustrative picture of the expected costs under the addition of different levels of beds is given by the plots of Figure 2. For each considered set of patient movement rules and additional bed levels, we demonstrate an “efficient frontier” of the configurations, i.e., the configurations with the best empirical blocking and diversion costs.

For each bed level and output measure we identify the set of the best configurations using a standard single-stage *subset selection* procedure. The procedure returns a subset (whose size can be random) that contains the best configurations with probability ≥ 0.95 . For more details on the procedure and its implementational details we refer the reader to Boesel et al. (2003). We illustrate the most favorable configurations for the case of nine additional beds (expressed as a vector of additional beds to the current base configuration) at Tables 3 and 4.

The first general observation comes from Figure 2 and pertains to the marginal difference of the different levels of additional beds. The difference is greater for the “General” policy, a set of movement rules that result in greater movement flexibility when compared to “Explicit” policy. This observed difference of scale (particularly in terms of divergence cost) is a result of the increased freedom of movement made possible by the “General” set of movement rules. We notice that even if the shape of the “efficient frontiers” remain largely unchanged under the two sets of patient movement rules, their positioning in the cost scale is different. In particular, the curves corresponding to “General” rules indicate considerable savings with respect to diversion costs. This observation supports the conjecture that there is more room for patients moving upstream when we relax patient movement rules. Furthermore, it highlights another important attribute of large scale simulations; patient movement rules can greatly affect a particular statistic (in this case patient diversion costs) with its values changing up to 50%.

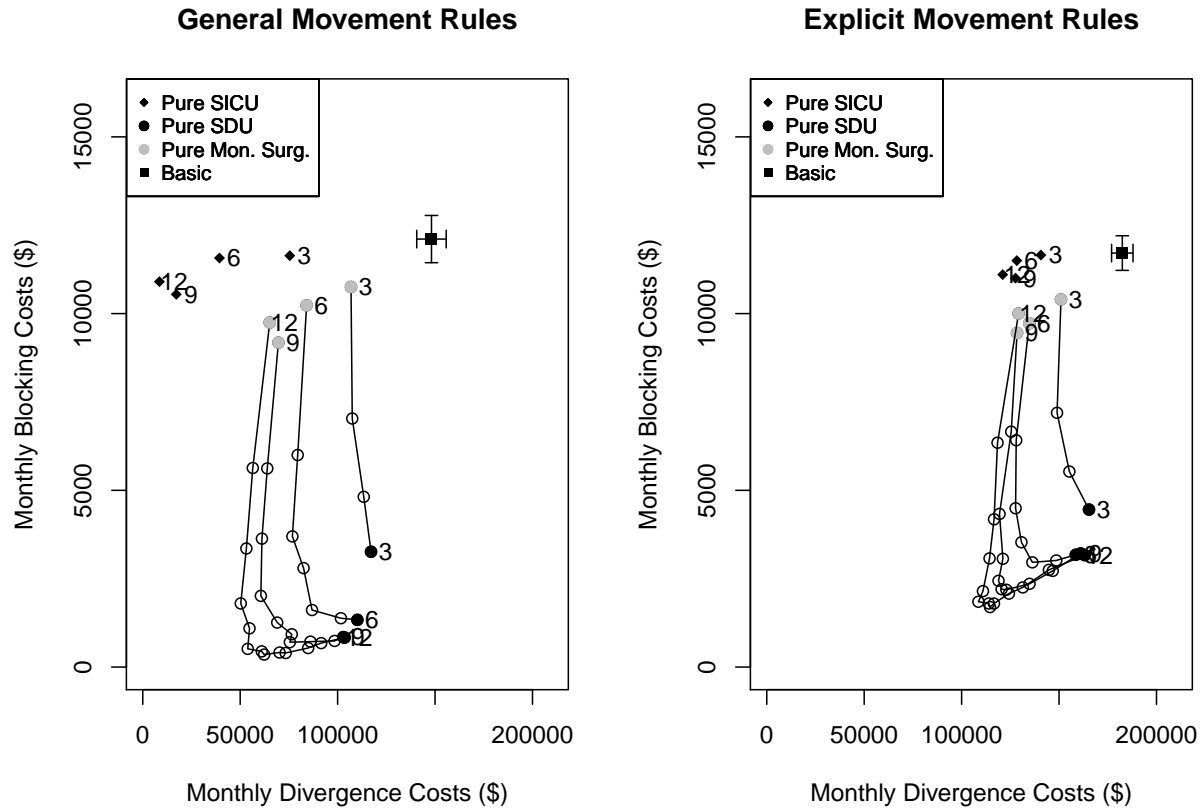


Figure 2: Frontiers for general and explicit movement rules. Solid lines connect solutions with the same number of added beds. The Basic configuration corresponds to the base bed configuration used in this experiment and the depicted bars to the confidence intervals for the considered output measures.

A second observation pertains to the most favorable bed configurations (for the case of nine additional beds) under the “Explicit” policy illustrated in Tables 3 and 4. We notice that for both considered measures a combination of SDU and monitored beds is a well balanced robust solution that seems to alleviate both blocking and diversions. These two bed types serve as down-stream beds for patients from the SICU, the unit responsible for the greatest portion of blocking expenses according to historical VAPHS data. So a mixed configuration impacts the overall cost of blocking and diversions in the following ways: (i) by alleviating blocking cost at the SICU for patients going downstream, (ii) by indirectly reducing ICU diversions through facilitating SICU patient throughput to SDU, and (iii) significantly reducing monitored diversions with the direct addition of monitored beds. Although the above observations do not contradict with intuition, we are able to quantify the optimal mixture of beds by selecting the configurations that belong to both sets of favorable configurations.

However, the picture is different when we examine the most favorable configurations under the “General” set of rules. As illustrated in Table 3, the best configurations with respect to patient blocking is a combination of SDU and monitored beds. However, as illustrated in Table 4, the best configuration with respect to patient diversions allocates all nine beds to the SICU. This is due to the fact that the more relaxed rules of “General” policy tend to direct patient flow “upstream”, especially when there is enough capacity to do so. As a result, the increased SICU capacity minimizes ICU diversions whereas the upstream patient flow has an alleviating effect on the general diversion episodes originating from the monitored units. In reality,

such concentration of patients would be avoided since the daily cost of an ICU is greater than that of a less intensive care unit. This observation highlights one of the possible pitfalls when modeling complex healthcare systems; we might reach different conclusions under different but validated patient movement rules.

MICU	SICU	CCU	SDU	Mon Med	Mon Surg	Policy
0	0	0	6	0	3	Explicit
0	0	0	5	0	4	
0	0	0	7	0	2	
0	0	0	6	0	3	Generic
0	0	0	7	0	2	
0	0	0	8	0	1	

Table 3: Best solutions with respect to cost of patient blocking

MICU	SICU	CCU	SDU	Mon Med	Mon Surg	Policy
0	0	0	4	0	5	Explicit
0	0	0	5	0	4	
0	0	0	6	0	3	
0	0	0	3	0	6	
0	0	0	2	0	7	
0	9	0	0	0	0	Generic

Table 4: Best solutions with respect to cost of patient diversions

6 DISCUSSION

In this article we present our work regarding a resource allocation project using a large scale simulation model. The challenges pertain to (i) the complexity of patient flow dictated by clinical practices and procedures, (ii) the implementation of pertinent patient movement rules that capture patient flow characteristics, and (iii) the sensitivity of the resource allocation schemes with respect to different sets of patient movement rules. We compile a simulation model able to implement complex patient movement scenarios in the face of congestion and conduct a sensitivity analysis of the considered simulation output over two sets of patient movement rules.

We conclude that different patient movement rules can affect patient flow statistics as well as lead to different conclusions when used to support resource allocation decisions. Furthermore, the work presented herein highlights the increasing need for the employment of sophisticated simulation software when modeling large healthcare systems. In particular, the need for incorporating complex movement rules and extensive model testing translates to software that deviates from the standard market packages and can be executed in parallel in computer clusters.

Future work includes the utilization of the simulation model to explore interventions beyond the realm of resource allocation problems. An open question is whether simulation models can provide the necessary fidelity to explore, for example, the benefits of truncating the LOS of a particular class of patients or the impact of arrival stream variations.

ACKNOWLEDGMENTS

The authors would like to thank the staff of the VA Pittsburgh Healthcare System (VAPHS) Veterans Engineering Resource Center (VERC), Robert Monte, Matthew Jenkins, Scott Moore, Nicholas Katich

and Ian Joos for their assistance and feedback. We would also like to thank our SME, Joyce Ewing, Kathleen Fritz, Joan Garloff, Paul Rogers, MD, and John Hotchkiss, MD. This material is based upon work supported by the Department of Veterans Affairs (VA), Veterans Health Administration, Office of Research and Development. The contents of this document do not represent the views of the Department of Veterans Affairs or the United States Government.

REFERENCES

- Boesel, J., B. L. Nelson, and S.-H. Kim. 2003. "Using Ranking and Selection to "Clean Up" after Simulation Optimization". *Operations Research* 51 (5): 814–825.
- Bountourelis, T., L. Luangkesorn, A. Schaefer, L. Maillart, S. Nabors, and G. Clermont. 2011. "Development and validation of a large scale ICU simulation model with blocking". In *Proceedings of the 43rd Winter Simulation Conference, WSC '11*, 1143–1153.
- Groeger, J. S., K. K. Guntupalli, M. Strosberg, N. Halpern, R. C. Raphaely, F. Cerra, and W. Kaye. 1993. "Descriptive analysis of critical care units in the United States: patient characteristics and intensive care unit utilization". *Critical Care Medicine* 21 (2): 279–291.
- OMNeT++ Community 2001. "OMNeT++ Network Simulation Framework". <http://www.omnetpp.org>.
- Teng, S.-O., W.-S. Lee, T.-Y. Ou, Y.-C. Hsieh, W.-C. Lee, and Y.-C. Lin. 2009. "Bacterial contamination of patients' medical charts in a surgical ward and the intensive care unit: impact on nosocomial infections.". *Journal of Microbiology, Immunology and Infection* 42 (1): 86–91.

AUTHOR BIOGRAPHIES

THEOLOGOS BOUNTOURELIS is a Research Assistant Professor at the Department of Industrial Engineering at the University of Pittsburgh. He received a B.S. in Mathematics from the Aristotle University of Thessaloniki, Greece, and a PhD in Operations Research from the Georgia Institute of Technology. His research interests, are in the area of Markov Decision processes, Machine Learning theory, Simulation and its applications in various technological contexts including HealthCare applications. His email address is thb28@pitt.edu.

SPENCER G. NABORS is a Post Doctoral Scholar and Clinical Care Physician in the Department of Critical Care Medicine at the University of Pittsburgh Medical Center and the Department of Veterans Affairs (VA) Pittsburgh Healthcare System. He received his BA in Bio-Chemistry and Philosophy from New York University, College of Arts and Sciences, his MA in Bio-Clinical ethics from New York University, Graduate School of Arts and Sciences, his MPH in Health Policy & Management from Columbia University, Mailman School of Public Health and his MD from SUNY Downstate, College of Medicine. His research interests include Process Modeling, Simulation and its applications in various clinical practices, Management Sciences and Critical Care Ultrasound technologies. His work is supported by a training grant from the National Institutes of Health (HL07820). His email address is naborssg@upmc.edu.

DAVID ECKMAN is an Undergraduate Research Assistant in the Department of Industrial Engineering at the University of Pittsburgh. His research interests include logistics and scheduling in health care and transportation settings. His email address is dje14@pitt.edu.

LOUIS LUANGKESORN is a Research Assistant Professor in the Department of Industrial Engineering at the University of Pittsburgh. He received a B.S. in General Engineering and a B.A. in Political Science from the University of Illinois-Urbana, an M.A. in Science, Technology and Public Policy from The George Washington University, and a Ph.D. in Industrial Engineering and the Management Sciences from Northwestern University. His research interests include logistics and resource management in health care and emergency response settings, as well as the use of simulation and models for policy analysis and

evaluation. His email address is lol11@pitt.edu.

ANDREW SCHAEFER is an Associate Professor of Industrial Engineering and Wellington C. Carl Fellow at the University of Pittsburgh. He has courtesy appointments in Bioengineering, Medicine, and Clinical and Translational Science. He received his PhD in Industrial and Systems Engineering from Georgia Tech in 2000. His research interests include the application of stochastic optimization methods to health care problems, as well as stochastic optimization techniques, in particular stochastic integer programming. He is interested in patient-oriented decision making in contexts such as end-stage liver disease, HIV/AIDS, sepsis, and diabetes. He is also interested in health care systems, including operating rooms and intensive care units. He is an Associate Editor for INFORMS Journal on Computing and IIE Transactions. His email is schaefer@pitt.edu.

GILLES CLERMONT is an Associate Professor of Critical Care Medicine at the University of Pittsburgh. He earned his undergraduate and medical degree from McGill University in Montréal, Quebec, Canada. From there he went to the University of Montréal, where he earned his Master of Science in Physics. After six years of private practice, Dr. Clermont returned to the University of Montréal to serve as Chief Resident at the Hôpital Notre-Dame. He then became a Research Fellow in Critical Care Medicine at the University of Pittsburgh, School of Medicine. Dr. Clermont is on the editorial boards for several peer-reviewed journals, such as Drug Discovery Today, Journal of Statistical Physics, and Critical Care, amongst others. He is also the Vice-President of the Society for Complexity in Acute Illness (SCAI). His current research interests are Complexity in Critical Illness, Epidemiology of Critical Illness, and Cost Effectiveness Analysis. His email address is clermontg@upmc.edu.