

DEVELOPMENT AND VALIDATION OF A LARGE SCALE ICU SIMULATION MODEL WITH BLOCKING

Theologos Bountourelis
Louis Luangkesorn
Andrew Schaefer
Lisa Maillart

Spencer G. Nabors
Gilles Clermont

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

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

ABSTRACT

Intensive Care Units (ICUs) are specialized healthcare delivery units for patients that require the highest level of monitored care. ICUs are typically integrated into larger healthcare facilities and their operation is dependent on the operational status of other inpatient units and departments of the host facility. As patients transition between units, a lack of available beds in a requested unit may cause patients to stay in a level of care other than that which is clinically indicated, leading to unnecessary or unwarranted costs without improving medical outcomes. The simulation modeling work presented in this paper is part of a multidisciplinary research project aimed towards patient delays. We describe the design and validation of a large scale ICU simulation model that includes various inpatient units and departments of the hospital. We describe the (i) input data analysis, (ii) modeling of patient flow, and (iii) validation of the simulation model.

1 INTRODUCTION

An ICU is a limited-capacity, resource-intensive unit designed to deliver monitored care for a diverse set of clinical conditions. The daily cost of ICU operations is higher than that of standard care, non-monitored units due to the requirement for a higher nurse-patient ratio, invasive bedside procedures and expensive monitoring equipment (Groeger et al. 1993). The daily ICU operations are driven by a number of factors, such as managerial decisions, clinical practices, unit occupancy, staff scheduling and the operating status of other inpatient units and departments of the host healthcare facility. Patients are admitted to the ICU from the Emergency Department (ED), as transfers from various surgical and other inpatient units, and as transfers from other hospitals. A typical ICU patient is transferred to a less intensive *monitored* or a "step-down" bed and, from there, to a regular hospital bed from which he/she is discharged from the hospital. However, some patients may experience a delay transferring into, or out of, the ICU. For example, an arriving patient may experience an admission delay or a diversion if there is no available ICU bed. On the other hand, a patient that is medically able to leave the ICU might experience a prolonged stay due to unavailable beds downstream. We will refer to this phenomenon as patient *blocking*.

Patient blocking is one of the most important patient flow characteristics with an impact on operating costs and the timely delivery of healthcare services. Blocked patients occupy resources that are not clinically appropriate, which increases operating costs without improvement on healthcare outcomes. Furthermore, patient blocking might propagate and disrupt operations to other related inpatient units and departments.

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 severity of illness necessitating ICU levels of care may be forced to receive intensive care in non-monitored beds, or other presumably more stable patients may be discharged prematurely or diverted to another hospital usually at a high cost.

Identifying the factors that drive patient blocking is a challenging task given the inherent complexity of the healthcare facilities where ICU operate. While most researchers concentrate their modeling effort on some individual area of the hospital, models that consider ICUs as an integral part of a larger healthcare facility are relatively rare. This may be due to the unique challenges pertaining to the increased data requirements and the need for extensive clinical personnel involvement. Furthermore, there are only a few models which formally define and consider patient blocking as an output performance measure. In light of the above observations, the work presented herein includes (i) a formal definition of patient blocking, (ii) a large scale ICU model that includes a comprehensive set of inpatient units and departments, and (iii) a systematic model validation procedure. The rest of this paper is structured as follows: In Section 2 we present a literature review related to the ICU simulation models. In Section 3 we give an overview of the healthcare facility modeled in this work. Next, in Section 4 we describe the data analysis that preceded the model development and the relevant statistical measures derived from it. Subsequently, in Section 5 we present the simulation model developed with a focus on patient movement and blocking. In Section 6 we describe the model validation process and demonstrate that the model output closely resembles the behavior of the real system. Finally, in Section 7 we discuss the limitations of our study and outline a number of directions for future work.

2 LITERATURE REVIEW

Discrete-event simulation models of healthcare systems have been used for many purposes (e.g., estimating bed utilization, facility design/redesign, patient scheduling, staffing etc). The ICU is an ideal application domain for discrete-event simulation models since it can allow the modeling of a heterogeneous patient mix as well as complex patient admission, movement, and discharge policies. Next, we briefly sample and review a few representative articles relative to the ICU modeling of operations.

Lowery (1992) presents a simulation model for the surgical suite and critical care area of a large hospital. The model simulates the flow of patients through a hospital ICUs, including the operating room, post anesthesia recovery unit, surgical ICU, intermediate surgical care unit, coronary care unit, intermediate coronary care unit, monitored unit, medical intensive care unit, and ventilator unit. In this model, when a patient attempts to move to a full unit, either the patient moves to an alternative unit, or an empty bed is created by bumping a patient who is considered almost ready to be discharged. The primary objective of this work is to quantify the effect of various bed configurations to performance measures such as utilization rates and number of patients that are turned away or accommodated on alternative units. Harper and Shahani (2002) demonstrate the importance of modeling the various types of patient flows when simulating bed occupancies and patient rejection rates. They provide evidence that the explicit modeling of the patient mix results in higher fidelity models able to capture the bed occupancy fluctuations over time. Cochran and Bharti (2006) present a methodology behind the compilation and validation of a large scale simulation of an obstetrics hospital that includes ICU and monitored units. The model is used to identify bed allocation schemes that satisfy given constraints and result in balanced bed utilizations while minimizing patient blocking across different units. Shahani, Ridley, and Nielsen (2008) present a simulation model for a single critical care unit. Emphasis is given to the modeling of the patient mix using the classification and regression tree (CART) method to obtain LOS distributions for statistically different patient categories. The model is used to implement changes in bed numbers, patient LOS, discharges, and the transfer of long-stay patients outside the critical care unit in order to explore their effects on bed occupancy and refused admissions.

In the existing literature, models that consider a network of clinical units and patient blocking are relatively rare (Cochran and Bharti (2006), Koizumi, Kuno, and Smith (2005)). Similarly, few papers document the validation procedure for their simulation model (Lowery (1993), Cochran and Bharti (2006)).

In this work we develop a discrete event simulation model of a large scale integrated hospital system that focuses on patient blocking as an output measure to elucidate patient flow dynamics. However, we take this analysis further with a systematic validation procedure which builds upon patient blocking.

3 ICU FACILITY

This simulation study is part of a collaborative project to examine issues related to the modeling, analysis, and management of ICUs. The project is focused on the ICUs of the Veterans Affairs (VA) Pittsburgh Healthcare System (VAPHS), a major healthcare provider in the Pittsburgh area. The hospital serves as the acute care facility and has 146 operating beds distributed among medicine, surgery, neurology, cardiology and critical care. This facility provides a full range of inpatient and outpatient care and it is the regional medical center of a system that covers western Pennsylvania, and parts of West Virginia and Ohio.

3.1 Facility

The VAPHS facility is comprised by several ICUs, inpatient units and departments. The units represented in our model are the following:

1. the Medical ICU (MICU) containing 9 beds;
2. the Surgical ICU (SICU) containing 12 beds;
3. the Coronary Care Unit (CCU) containing 18 beds;
4. a Step Down Unit (SDU) containing 9 beds;
5. a Monitored Medical unit containing 15 beds;
6. and a Monitored Surgical unit containing 12 beds.

After examining patient data files we compiled a list of physical locations from which patients apply for admission to the above inpatient units. Subsequently, after consulting with the clinical advisory board, we aggregated those locations into categories referred to as *patient sources* defined as:

7. Non-Monitored Medical and Surgical inpatient units (71 beds)
8. Emergency Department (ED)
9. Operating Room (OR)
10. Post-Anesthesia Recovery (PAR) room
11. Other Floors
12. Direct Admissions

Note that we consider the non-monitored medical and surgical units as patient sources. This is consistent with our intention not to model the non-monitored units explicitly but rather model them as sources and sinks of patient flows.

4 AVAILABLE RESOURCES AND DATA ANALYSIS

The main resources for this simulation model were (i) the VAPHS clinical staff that provided valuable feedback on the day-to-day clinical practices and policies governing patient movement, and (ii) the VAPHS data warehouse that contained explicit patient movement information. The VAPHS data warehouse contained a number of patient information databases from which we chose the (i) the Bed Management System (BMS), and (ii) the Patient Movement (PM). These databases contain patient movement and transition data. Both have limitations when used alone but together they provide a more complete picture of patient movement.

The BMS is a collection of entries corresponding to patient's *transfer requests*. As depicted in Figure 1, each patient in BMS is represented by a sequence of entries corresponding to the transfer requests logged in when the patient is evaluated as ready to be transferred to another unit. Each entry includes (i) the

Patient ID	Requested	Removed	Requested Unit	Current Unit
123xxxx	02/05/10 10:35	02/05/10 18:35	MICU	ED
123xxxx	02/08/10 11:35	02/08/10 17:35	Monitored Medical	MICU
123xxxx	02/10/10 12:30	02/10/10 13:20	Non-Monitored Medical	Monitored Medical

Figure 1: Bed Management System entries.

Patient ID	Date and Time	Type of Entry	Unit
123xxxx	02/05/10 18:32	Admission	MICU
123xxxx	02/08/10 17:40	Transfer	Monitored Medical
123xxxx	02/10/10 13:25	Transfer	Non – Monitored Medical
123xxxx	02/11/10 14:40	Discharge	Non – Monitored Medical

Figure 2: Patient Movement entries.

patient identifier, (ii) the time and date this transfer was “requested”, (iii) the time the transfer request “removed” (due to the transfer occurrence or request withdrawal), (iv) the requested unit for transfer, and (v) the current unit of the patient. Given the BMS data, we have been able to identify the patient sources and calculate the corresponding, time and date dependent, arrival rates. Furthermore, using these data, we estimated the patient blocking times experienced at each unit by calculating the time difference between a patient transfer and the corresponding transfer request.

The PM database is a collection of entries corresponding to patients *realized transitions*. Each entry corresponds to a patient admission, transfer or discharge and contains the relevant unit and the time this transition took place. As depicted in Figure 2, each patient is represented by a sequence of entries describing the patient’s unit transitions from admission to discharge. From this information, we estimated a number of patient flow statistics such as LOS, transition probabilities and patient throughput for every unit. Using the entire set of PM entries we calculated the bed occupancy rates for any given unit and date.

The two data files contain partial information when viewed in isolation. By combining the information residing in them, we were able to compile a complete picture of the arrival, the LOS, the discharge and the delays the patient experienced due to blocking. Furthermore, we coupled our data analysis with the experience of our clinical board that highlighted those data entry practices that may bias the statistics obtained from them. For example, not all blocking times were attributed to a shortage of available beds. The clinical advisory board pointed out that blocking instances of less than 3 hours can be attributed to the patient’s need to be reevaluated by a doctor and not to bed availability issues. Furthermore, for patient sources like the OR and PAR, patient blocking times are rather biased indicators of actual patient delays. This is due to the clinical practice that patient transfer requests in these locations are logged long before the patient is medically cleared for transfer. For example, a transfer request may be logged from OR even before a surgery operation has started. Therefore, a number of adjustments had to be made in order to alleviate the biasing effects and ensure that the derived statistics could be used in the validation process of the simulation model.

In the next section, we refine the definition of LOS in order to account for the patient blocking phenomenon.

4.1 Patient Length of Stay and Blocking

We define the *blocking time* as the elapsed time between a transfer request and the time this transfer was realized. The analysis of LOS data without considering patient blocking will result in estimates of the

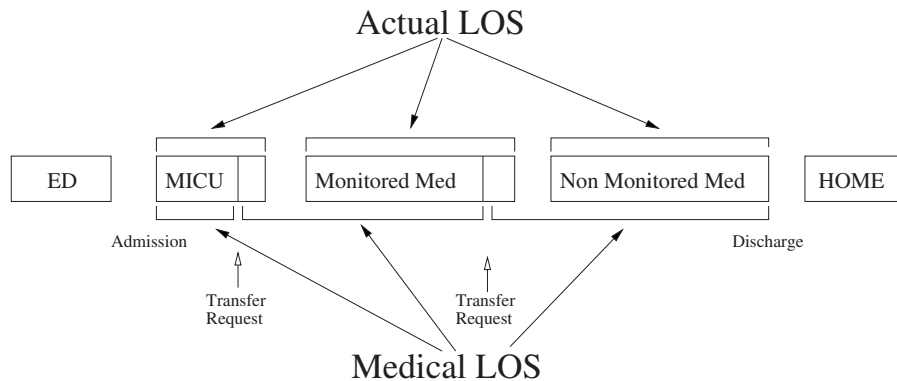


Figure 3: A graphical representation of the actual and the medical LOS for the patient depicted in Figures 1 and 2. The medical LOS for MICU is the time elapsed from admission to the transfer request. The medical LOS for the Monitored Medical unit is the time elapsed between the two transfer requests. The medical LOS for the last unit is the time elapsed from the last transfer request to discharge.

actual LOS, i.e., the total time duration a patient spent in a unit. However, in the presence of blocking, actual LOS may result in a rather biased estimator of the patient’s medical need for a bed. Therefore, we define the *medical LOS* as the clinically indicated time the patient had to spend in a unit before the patient is medically cleared for a transfer or a discharge. In order to obtain medical LOS estimates, we make use of the patient’s transfer requests found in the BMS database.

In particular, the medical LOS of a patient is determined using the (i) admission, (ii) transfer requests, and (iii) the discharge times. As seen in the patient example depicted in Figure 3, for a newly admitted patient coming from outside the hospital, the medical LOS for the first unit is the elapsed time from admission to the next transfer request. In subsequent units, observe that the medical LOS spans the time from the last transfer request, while the patient was at the previous unit, to the time of the current transfer request. This definition takes into account the fact that, while the patient is waiting to be transferred, his/her health condition is presumably improved and, therefore, this waiting period should be credited towards the medical LOS of the next unit. Finally, the medical LOS of the last unit is the elapsed time between the last transfer request and the discharge time.

5 SIMULATION MODEL

5.1 Simulation Software

The simulation model was 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 (e.g., the addition or removal of entire hospital units). This not only increases the fidelity of the model but will also allow for broad applicability with respect to the healthcare management field.

5.2 Simulation Model Description

In the majority of the ICU models in the literature, historical LOS data are used to fit a distribution that is used to sample patient LOS realizations as well as transition probabilities during simulation runs. In order to increase model validity, some authors propose the stratification of patients into different categories depending on the *treating specialty* (e.g., surgery, cardiology, neurology) and patient attributes (e.g., age or some selected clinical indicators). However, in a large scale simulation model, such an approach might not




Figure 4: Patient arrival rates from the ED.




Figure 5: Patient arrival rates from the PAR room.

provide the fidelity needed to capture the patient flow dynamics as patient entities occupy beds and transition from one unit to another. Patient transitions and medical LOS may be stochastically dependent on a number of factors including patient source, the sequence of prior visited units, and the corresponding medical LOS. Our approach towards the modeling of patients is to populate a *patient database* containing *actual* patient instances as observed from the data files. Each patient instance is described by (i) an arrival source (Items 7-12 of Section 3.1), (ii) a sequence of visited units and departments (Items 1-12 of Section 3.1), and (iii) a corresponding sequence of medical LOS values. Furthermore, for every patient source, we estimate the patient *arrival rates* (e.g., see Figures 4 and 5) which depend on the hour of the day and day of the week.

The program logic underlying patient generation and movement is described as follows: At every hour of the day and day of the week, the program retrieves the corresponding arrival rate for each patient source and generates the corresponding number of patients using the Poisson distribution. Subsequently, the program queries the patient database to randomly select the number of patients for each source. The selected patients are inserted into the simulation model and transition through the different units and sources based on their stored sequence of units and medical LOS values. A patient is deleted from the simulation upon completion of his unit visitation sequence. For example, a patient may originate from the ED, request and receive a bed at the MICU, and, subsequently, get admitted to the non-Monitored Medical unit before being discharged and deleted.

When patients reach the end of their medical LOS they request a bed at the next assigned unit. If a bed is not available, the simulation program searches for a bed at an alternative unit based on some predetermined rules that correspond to actual clinical practices. If not found the patient is declared blocked and waits for the signal indicating an empty bed at the requested level of care. Upon transition, the accumulated blocking time is subtracted from the next unit's medical LOS value. Our model collects and reports patient blocking times and bed occupancy rates for all the inpatient units corresponding to Items 1-6 of Section 3.1.

6 VALIDATION

The validation process of a simulation model is based on the purpose and the questions the model is designed to answer (Sargent 1996). Our model validation process consisted of two steps: a *conceptual validation* which ensures the face validity of the model and its underlying logic and an *operational validation* focusing on the quantitative characteristics of the model's output.

6.1 Experimental Settings

The BMS data spanned a time period of 9 months and the PM data a period of 2 years. We extracted a database that covers nearly 9 months, after removing and cleaning incomplete records from BMS. We ran 25 independent replications of the model, using the first 5 months of each replication as a warm up period followed by 9 months of recorded results.

6.2 Conceptual Validation

For the conceptual validation we presented a high-fidelity simulation animation to clinical staff associated with this project. For this, we developed a simulation prototype using a commercial simulation package (ARENA) with animation capabilities (Rockwell Automation 2008). The simulation animation included (i) patient movements as they get admitted, transferred and blocked in the various monitored units, and (ii) the visual display of various performance measures (e.g., number of occupied beds and blocked patients per unit) to ensure that the observed patterns are consistent with the experience of the clinical staff. Furthermore, the simulation animation was used to facilitate our communication with the VAPHS clinical staff in order to get feedback and document those clinical policies that dictate patient movement.

6.3 Operational Validation and Calibration

As with all modeling, we must first make sure the model output has the appropriate qualitative characteristics. Then, once this process of *validation* is complete we must *calibrate* the model with respect to subject specific parameters that cannot be quantified a-priori. To establish the operational validation of our model we used two output statistical measures relevant to patient movement: patient blocking and bed occupancy. The initial simulation runs indicated that the simulation output corresponding to patient blocking times and the number of occupied beds is close with the corresponding statistics retrieved from the data files. However, there is a set of input parameters that play a pivotal role in the simulation but cannot be quantified from the available data. These parameters need to be calibrated, Kleijnen (1995), in order to further improve the statistical proximity of the simulation output to the output observed from data. The input parameters used for the calibration of the model correspond to the (i) the number of beds of each unit, (ii) a common probability of *alternative unit assignment* (i.e. the probability that a patient is diverted to another unit of the same level of care due to bed unavailability vs the patient getting blocked), and (iii) a set of *scaling factors* for modifying arrival rates. For example, the number of available beds in the real system is a variable parameter that depends on a number of factors pertaining to staff availability and schedules. Similarly, the probability of an alternative unit assignment depends on a number of patient specific attributes (e.g., age, severity, and treating specialty) and it is a difficult, if not impossible, task to obtain an accurate estimate of it. On the other hand, the use of factors for scaling the patient arrival rates is necessary in order to offset the errors and possible data discrepancies that might compromise the accuracy of the calculated arrival rate estimates.

Model calibration can be a time consuming process due to the many dimensions of the parameter space induced by the total number and range of values of the input parameters. There is no standardized calibration procedure in the literature and it appears there are no absolute “correct” answers when selecting a set of input parameters. After experimenting with various sets of parameters, we were able to identify a set of parameter values under which the model output closely resembles the output observed from data. For comparing the model and data output we used a combination of *subjective* and *objective* validation techniques.

6.3.1 Patient Blocking

The most important model output is patient blocking. As we described in Section 4, patient blocking is defined as the time elapsed between a patient transfer request and the patient transfer. Figure 6 depicts




Figure 6: The Q-Q plots for patient blocking - Simulation compared to Actual.

the observed patient blocking in MICU and SICU compared to the corresponding simulation output. We focused on those two units since these are the clinical areas where critical care patient blocking mostly occurs. In both cases, the Q-Q plot demonstrates that there is close linear relationship between the corresponding quantiles. Given the inherent complexity of the ICU operations and the necessary simplifying assumptions, we believe that our model performs satisfactorily in replicating critical care patient blocking.

6.3.2 Bed Occupancy

Bed occupancy is another important simulation output measure. It is an indicator of the actual patient LOS which is the sum of patient medical LOS and blocking time. Replicating the bed occupancy for the inpatient units is another indication that the simulation model faithfully represents key aspects of the patient flow in the actual system. Figure 7, depicts the average inpatient unit occupancy as retrieved from data and produced by the simulation model. The graph on the left depicts the average occupancy rates in the form of adjacent bar plots and their corresponding percentage difference. The graph on the right depicts the confidence interval produced by the independent replications. Note that the observed occupancies from data fall in the corresponding confidence intervals for all units except SICU and CCU. But even for those units, we believe that the simulation occupancy rates are satisfactorily close to the actual occupancy rates. Based on this analysis, we conclude that the simulation model can faithfully represent the level of blocking and bed occupancies in the current system. This allows us to analyze various policy options based on the current system using blocking as a performance measure.

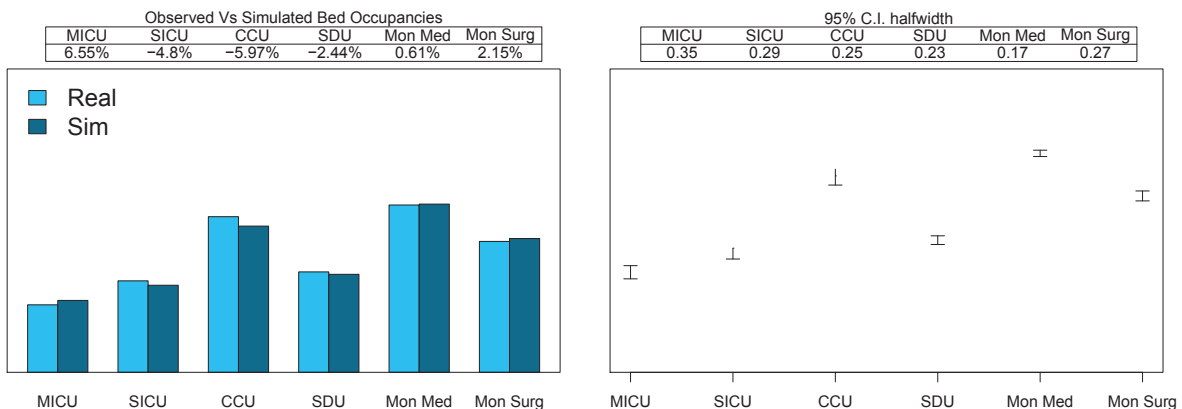


Figure 7: Simulated and real occupancy rates.

7 DISCUSSION AND FUTURE WORK

7.1 Discussion

For the validation of our simulation model we used methods such as Q-Q plots and confidence intervals to evaluate the simulation output with the output observed from real data. However, the model might be compromised by factors related to underlying assumptions and simplifications. For example, the non-monitored units are not modeled with the same fidelity as the ICU and monitored units. That is, they are considered to be “black boxes” where a patient might transfer when requesting a bed without considering availability of a bed. As a result, the model does not reflect blocking that can occur due to lack of available beds in the non-monitored units. Furthermore, we assumed that the number of unit beds remains constant during simulation runs and did not consider bed closures due to staffing shortages or scheduling changes (we attempt to alleviate this assumption by introducing a calibration parameter on the number of beds). These assumptions have a biasing effect on blocking and the actual length of stay at all levels. We successfully managed to alleviate those effects by the introduction and calibration of the input parameters of Section 6.3.

Finally, because of the limitations in accessibility of historical data sets, we are currently limited in the amount of 9 months of usable BMS data. As we collect and augment the patient database with more patient instances, this will increase the validity and ability of our model to represent finer grained patient populations.

7.2 Next Steps

Our next step is to obtain additional hospital operating data to proceed in a more complete validation procedure. In particular, we seek to obtain patient movement data that span the time periods before and after a major bed reconfiguration event. We seek to observe the effects of the event on patient blocking and investigate the ability of our model to replicate the patient blocking effects induced by the new configuration.

Another line of research concerns the modeling of patient entities. In the above described work we populated and utilized a database of actual patient entities compiled from data to produce a patient-derived model. An ongoing research effort will outline patient modeling procedures in order to derive transition and LOS distributions capturing the complexity of multiple unit transitions. Our goal is to use the aforementioned distributions for random patient generation so that the resulting model has the same level of fidelity as the model presented herein. Such a model would be a powerful tool allowing for the analysis of complex patient movement scenarios.

Finally, the goal of this ongoing research effort is to incorporate the simulation model into a decision making tool to answer various policy questions and alternatives. As a first step towards this direction, we have incorporated various cost variables into the model in order to quantify the monetary effects of patient blocking at the ICUs and patient diversions at the ED. Furthermore, we are developing *simulation optimization* methods to be used with this model in order to identify efficient sets of bed configurations, staffing variables and policy alternatives.

ACKNOWLEDGEMENTS

The authors would like to thank the staff of the VA Pittsburgh Healthcare System (VAPHS) for their assistance and feedback. We also thank our colleagues Dr. Winston Yang and Dr. Jeffrey Kharoufeh. This material is based upon work supported by the Department of Veterans Affairs (VA), Veterans Health Administration, Office of Research and Development.

REFERENCES

Cochran, J., and A. Bharti. 2006. “Stochastic bed balancing of an obstetrics hospital”. *Health Care Management Science* 9 (1): 31–45.

- Groeger, J., K. Guntupalli, M. Strosberg, N. Halpern, R. 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.
- Harper, P., and A. Shahani. 2002. "Modelling for the planning and management of bed capacities in hospitals". *Journal of the Operational Research Society* 53 (1): 11–18.
- Kleijnen, J. P. C. 1995, November. "Statistical validation of simulation models". *European Journal of Operational Research* 87 (1): 21–34.
- Koizumi, N., E. Kuno, and T. E. Smith. 2005. "Modeling Patients Flows Using a Queueing Network with Blocking". *Health Care Management Science* 8 (1): 49–60.
- Lowery, J. 1992, December. "Simulation of a hospital surgical suite and critical care area". In *Proceedings of the 1992 Winter Simulation Conference*, edited by J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson, WSC '92, 1071–78. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Lowery, J. 1993, December. "Multi-hospital validation of critical care simulation model". In *Proceedings of the 1993 Winter Simulation Conference*, edited by E. C. Russell, W. E. Biles, G. W. Evans, and M. Mollaghasemi, WSC '93, 1207–1215. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- OMNeT++ Community 2001. "OMNeT++ Network Simulation Framework". <http://www.omnetpp.org>.
- Rockwell Automation 2008. "Arena Version 12".
- Sargent, R. G. 1996, December. "Verifying and validating simulation models". In *Proceedings of the 1996 Winter Simulation Conference*, edited by J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain, WSC '96, 55–64. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Shahani, A., S. Ridley, and M. Nielsen. 2008. "Modelling patient flows as an aid to decision making for critical care capacities and organisation". *Anesthesia* 63 (10):1074 – 80.

AUTHOR BIOGRAPHIES

THEOLOGOS BOUNTOURELIS is a Post Doctoral Associate with 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 Fellow 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.

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 loll1@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.

LISA MAILLART Lisa Maillart is an Associate Professor in the Industrial Engineering Department at the University of Pittsburgh. Prior to joining the faculty at Pitt, she served on the faculty of the Department of Operations in the Weatherhead School of Management at Case Western Reserve University. She received her M.S. and B.S. in industrial and systems engineering from Virginia Tech, and her Ph.D. in industrial and operations engineering from the University of Michigan. Her primary research interest is in sequential decision making under uncertainty, with applications in medical decision making and maintenance optimization. She is a member of INFORMS, SMDM and IIE. Her email is maillart@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 Montral, Quebec, Canada. From there he went to the University of Montral, where he earned his Master of Science in Physics. After six years of private practice, Dr. Clermont returned to the University of Montral to serve as Chief Resident at the Hpital 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.