

MANAGING PATIENT BACKLOG IN A SURGICAL SUITE THAT USES A BLOCK-BOOKING SCHEDULING SYSTEM

Oleg V. Shylo
Louis Luangkesorn
Oleg A. Prokopyev
Jayant Rajgopal
Andrew Schaefer

University of Pittsburgh
3700 O'Hara St.
Pittsburgh, PA 15261

ABSTRACT

Effective scheduling of elective cases in an operating room suite is a challenging task due to inherent uncertainty and competing performance metrics. In this paper, we present a simulation model for the surgical suite within the VA Pittsburgh Health Care System (VAPHS) that is used to evaluate and optimize different scheduling policies. A flexible set of probabilistic scheduling rules is evaluated and a dynamic scheduling policy is proposed as an alternative to static strategies. The dynamic scheduling policy allows us to reduce the variance in patient waiting times and backlogs. The developed simulation model is based on the data collected at the VAPHS.

1 INTRODUCTION

In a hospital setting, the operating room (OR) suite can account for about 42% of total revenues (Healthcare Financial Management Association 2003). At the same time, the average OR runs at about 68 % of capacity; thus revenues could be made higher through improved utilization. However, this is complicated by the inherent difficulty in scheduling and managing operating room resources. Conflicting priorities, competing performance criteria for the operating room suite, uncertainty in surgery duration, uncertainty in patient recovery times and the limited capacity of operating room suites all contribute to this problem. Inefficient or inaccurate scheduling of OR time often results in delays of surgery or cancellations of procedures, which are costly to the patient and the hospital (Gordon et al. 1988).

Hospitals seek ways to efficiently manage operating rooms and schedule procedures. To analyze potential scheduling procedures, one can model the process of scheduling and managing operating rooms. In particular, simulation models can be used to evaluate proposed changes in the delivery of health care and in the management of patient flow in surgical care (Sobolev et al. 2011).

In this paper we develop a simulation model of the case backlog at the surgical suite within the U.S. Department of Veterans Affairs (VA) Pittsburgh Healthcare System (VAPHS). The model accounts for arrivals, scheduling, and the actual performance of the generated schedules. Currently, at the VAPHS many of the surgical specialties have backlogs that are greater than what is desirable. In this work we define a measure of the case backlog, and then demonstrate that the case backlog and the corresponding overtime are dependent on the settings of two management parameters.

Using this model, we propose a method to manage the case backlog and overtime using an upper and a lower control limit to trigger policy adjustments. Then we will use the model to demonstrate the effects of different parameter settings in controlling backlog and overtime. While control charts have been used in

health care management to monitor clinical outcomes and in-hospital waiting (Winkel and Zhang 2007), this is the first known application of control charts to monitor the time required to make an appointment for surgery.

2 LITERATURE REVIEW

The complexity and cost of operating room planning and scheduling make this a natural topic for analysis using operations research methods. Cardoen et al. (2010) reviewed recent research that explicitly incorporates planning and scheduling considerations. They categorize the 247 manuscripts according to patient characteristics, performance measures, decision delineation, research methodology, uncertainty, and applicability of research. Our work addresses elective patients, with case backlog and overtime as performance measures. We address assignment of individual patients to a particular day, develop heuristics for decision making and use a combination of discrete event and Monte Carlo simulation to evaluate the heuristic. The procedure accounts for duration uncertainty, and uses data from actual hospital operations with the objective of implementing these methods in practice.

2.1 Simulation Modeling of Surgical Suites

A systematic review of simulation models of surgical suites was done by Sobolev et al. (2011). Through a set of systematic searches they identified 131 publications on the application of simulation methods to study surgery. After a full text review, they found 34 publications that described a computer simulation of patient flow in surgical care. Of these, 31 discussed the utility of simulation for analyzing changes in delivery of surgical care, 17 discussed its utility for analysis of different policies, seven discussed the utility of wait-list management and one publication discussed the utility of changes in management.

One use of simulation is to demonstrate the tradeoffs between performance measures of interest in the surgical suite. Denton et al. (2006) use simulation as a tool to evaluate the tradeoff between patient waiting time, OR team waiting time, OR idling and overtime in a multi-room surgical suite.

2.2 Scheduling Surgeries

Denton et al. (2010) consider assigning a surgery to operating rooms on the day of the surgery. They use both a two-stage stochastic linear program with recourse as well as a robust heuristic. Min and Yih (2010) examine this problem while accounting for downstream capacity constraints and develop a stochastic programming model; they use a sample average approximation algorithm to solve this within a reasonable amount of time.

2.3 Use of Analysis for Management of Health Care

Some models look at the overall management of health care and the surgical suite as opposed to modeling the clinical practice of health care. Dexter et al. (1999) describe the simulation model of patient scheduling using bin-packing algorithms, where the patients were modeled as having a given average duration. They find that the time until a patient surgery is scheduled depends on the scheduling method used and the length of the average surgery. Hans et al. (2008) look at the robust surgery loading problem, which assigns surgeries to operating room days while incorporating sufficient slack. Their objectives are to maximize utilization while minimizing risk of overtime. They propose heuristics and local search methods that allow for swapping surgeries between days or rooms.

Using administrative data to manage surgical suites is explored by Beaulieu et al. (2010). They work with a major medical center to utilize administrative data for tracking various clinical performance measures in real time and providing this in a useful form to frontline care teams. Different clinical interventions and practices were tested using this feedback.

More generally, Winkel and Zhang (2007) describe the implementations of control charts in health care settings. They include examples of tracking waiting room times, procedure complications, inter-unit transport time, error/defect rates, infectious rates and patient satisfaction.

3 SIMULATION MODEL OF ONLINE SCHEDULING

Currently, the surgical suite at the VAPHS consists of nine operating rooms (OR1-OR9). Under the block-booking scheduling system that is in use, each medical specialty (e.g., ophthalmology, cardiac, orthopedics) is assigned a recurrent (repeating monthly) schedule of reserved times in specific operating rooms. For example, the block schedule for the ophthalmology department at the VAPHS in 2009 consisted of either two or three blocks per week in a specific operating room: i) a time block from 8 am until 5 pm in OR9 every Tuesday; ii) a time block from 8 am until 3 pm in OR9 every Thursday; and iii) a time block from 8 am until 12 noon in OR9 every other Friday. Other specialties, depending on the specifics of the surgical procedures and current demand trends, might be assigned multiple blocks on a given day across a number of operating rooms. Two different specialties can have consecutive time blocks on the same day in the same operating room (e.g., cardiac from 8am to 12:30pm, followed by orthopedics from 12:30pm to 3pm every Monday in OR4). The scheduling personnel work with the medical departments to provide a feasible allocation of surgeries to the corresponding scheduling blocks. The current scheduling procedure can be described as a four-step process:

- The medical specialty submits a surgery request, which contains the details of the surgical procedure and a preferred date. Most such requests are based on an agreement between the surgeon and the patient to perform a specific surgical procedure in the future (elective surgeries).
- The scheduling personnel at the VAPHS accommodate the request by assigning the specific date and operating room based on the current schedule and future blocks available to the specialty. Local improvements can be made to this allocation subject to cancellations or other considerations (some mix of surgeries is more predictable in terms of processing times than others, when assigned to the same scheduling block).
- The exact sequence of the surgeries within each block is determined 1-3 days prior to the date of the actual procedures.
- Short-term changes to the surgery sequence and operating room assignment can take place even on the surgery day. For example, the surgery that is scheduled for the end of the day can be moved to an earlier time slot if there is an opening in some other operating room due to cancellations or early completions. According to the scheduling personnel at the VAPHS, such adjustments are fairly common.

The scheduling problem can be decomposed into a set of non-overlapping sub-problems, each corresponding to a patient stream within a single medical department. Based on our conversations with the management team of the OR suite at the VAPHS, such a decomposition is valid since the availability of the downstream resources such as cleaning crews or recovery beds is not a pressing issue within the present system.

Currently, the scheduler at the VAPHS uses a point estimate of the procedure duration of each case based on historical data, his own experience, and the estimate provided by the surgeon. Based on the current schedule, these point estimates are used to determine a “feasible” block assignment, i.e., one where the sum of the point estimates of durations of surgeries assigned to the same scheduling block and intervening cleaning times do not exceed the block length. In order to take into account the intrinsic uncertainty associated with surgery durations, we develop a prediction system based on generalized linear models (Firth 1991). The durations are modeled as Gamma variates, which is consistent with the literature on prediction of surgery durations. To predict the mean and variance of the new surgical cases, the model uses the specific procedure type as identified by a Common Procedure Terminology (CPT) code, the surgeon’s experience given by the total number of surgeries performed and an estimated duration provided by a surgeon. The prediction

model for the ophthalmology department was fitted using data from 2006-2008 and tested on historical data from 2009. It demonstrated a reasonable predictive potential: out of 619 cases in the validation set, 53% fall under the predicted 50th percentile, 79% fall under the predicted 80th percentile and 88% fall under the predicted 90th percentile (the percentiles differ across different cases).

Cancellations, changes in surgery dates, and the addition of emergency, add-on, urgent and emergent cases lead to alterations of the actual schedule on a daily basis. The decisions about such changes are almost impossible to predict, since they are made in real time based on expert opinions (both the surgeon and the scheduling team are involved), the current state of the schedule, the availability of add-on procedures and the willingness of the surgeon to accept such changes. In our model these situations are omitted from consideration, since they can be managed to a large extent by adjusting the scheduling parameters based on the actual cancellation rates and the arrival rates for add-on, emergent and urgent cases (e.g., overbooking for cancellations or underbooking for add-on cases). These rates differ substantially across the set of medical specialties, some of which (e.g. oral and ophthalmology) have a negligible cancellation rates.

Next, we present a simulation algorithm of the scheduling process described above:

1. *Initialize the current simulation date and the set B of blocks available for scheduling on the current date.*
The current date is incremented during the simulation and allows us to extract the day of the week and week numbers, which are used to identify clinic days and the blocks available for scheduling. Each block has a certain realization date and duration, which are determined by the block schedule.
2. *If the current date is a clinic date, generate a set of new surgeries S , otherwise proceed to Step 4.*
The number of arriving surgeries n is generated according to a Poisson distribution with a fixed arrival rate. A random sample of n surgery records is selected from the historical data. The Common Procedure Terminology (CPT) code of each surgery and the scheduled duration are used as inputs for the predictive model of surgery durations. This model provides the estimates of distribution parameters for each surgery from S .
3. *Sequentially assign each surgery from S to a block from B using a scheduling rule Π . If there is no feasible assignment for some surgery, add a new block to B according to the block schedule and assign the surgery to the new block.*
The most common choices for the scheduling rule Π include first-fit, and best-fit (see Dexter et al. 1999), possibly subject to certain probabilistic and block capacity constraints. A new block is added to the set B , if a new surgery cannot be booked into the current set of blocks without violation of scheduling constraints used by Π . This block is the next available block that is not in B according to the current date and the block schedule.
4. *Increment the current date (next day).*
5. *Process all the blocks from B that are scheduled for the current date and remove them from B .*
Different performance metrics (overtime, utilization, waiting time) are calculated for each removed block.
6. *Stop the simulation if the current date exceeds the end date of the simulation, otherwise proceed to Step 2.*

Two versions of scheduling rules were implemented for the purposes of the current study (see the description of Π in Step 3 of the simulation algorithm):

1. Π_{FFD} (*first-fit deterministic*): *First-fit scheduling rule using the mean values of surgery durations*
The surgery is assigned to the first available block for which the sum of mean durations is less than its length after the assignment.
2. Π_{FFP} (*first-fit probabilistic*): *First-fit scheduling rule under the probabilistic constraints*

This scheduling rule has two parameters: the threshold value L and the probability of overtime α . The surgery is assigned to the first available block for which the probabilistic capacity constraint (see the discussion below) is satisfied after the assignment.

Let $S^b = \{s_1, \dots, s_n\}$ denote a set of surgeries assigned to some scheduling block b with a given duration $l(b)$. The total duration associated with the surgery s_j is a random variable ξ_j , which includes the pre-processing and post-processing times associated with the procedure including the operating room preparation and cleaning times. The probabilistic constraint of the scheduling rule Π_{FFP} can be described as

$$Pr\left\{ \sum_{j: s_j \in S^b} \xi_j - l(b) > L \right\} \leq \alpha. \tag{1}$$

In the literature, the surgery durations are commonly modeled by a lognormal or gamma distribution; this has been validated in a variety of real hospital settings (Strum et al. 2000). Since there is no simple analytical expression for the convolution of gamma or lognormal variates, the only practical way of estimating the probability given by (1) is by using Monte Carlo simulation. Another alternative is to use the normal approximation for the convolution of gamma variates. In our simulation study, we investigate the accuracy of such approximation and demonstrate that the quality of the normal approximation is extremely accurate for our setting.

Assuming that ξ_j 's have a Gaussian distribution with mean μ_j and standard deviation σ_j for each $j: s_j \in S$, the block overtime $\sum_{j: s_j \in S^b} \xi_j - l(b)$ is a normally distributed random variable with mean $\mu_b = \sum_{j: s_j \in S^b} \mu_j - l(b)$ and variance $\sigma_b^2 = \sum_{j: s_j \in S^b} \sigma_j^2$. In this case, the constraint (1) can be simplified:

$$L \geq \mu_b + \phi^{-1}(1 - \alpha)\sigma_b, \tag{2}$$

where $\phi(p)$ denote the cumulative distribution function of the standard normal variate. In our simulation studies, the scheduling rule Π_{FFP} that uses the normal approximation simply checks for the violation of (2) every time a new surgery is assigned to the scheduling block.

4 COMPUTATIONAL RESULTS

The simulation algorithm described in the previous section was implemented using the OMNeT++ package (<http://www.omnetpp.org>), which is a C++ library for building simulation models. Our model implements the scheduling steps described in the previous section. It was used to simulate scheduling decisions for the ophthalmology department at the VAPHS. The length of each simulation run was set to one year. Patients seeking surgery are assumed to arrive according to a Poisson distribution with a mean of 13 per week. This figure is based on the average historical rate for the ophthalmology department during 2009. Each set of simulation parameters was used to generate 100 replications (each replicate simulates schedules for one year). The warm-up period was set to 20 days, providing an initial backlog of 20 days. The block schedule used in the simulation was identical to the one that was used by the ophthalmology department at the VAPHS in 2009 (see the description in Section 3).

Since current practice is to use mean surgery durations to determine the feasibility of scheduling assignments, the first-fit rule based on mean durations Π_{FFD} was used to simulate the current scheduling policy. The distribution of overtime obtained via simulation using Π_{FFD} is similar to the historical distribution (Figure 1). For each value of overtime, the plot shows the percentage of blocks that exceeded that value. Some of the discrepancies between the simulation and the historical distribution of overtime are mainly due to non-uniform arrival rates; this leads to low arrival rates during certain time intervals. Under the given block schedule, the average throughput for the Π_{FFD} is 660 patients per year, and the 95% CI for this is [659.31,661.13]. The average final number of backlog days for the Π_{FFD} after one year is 38.98




Figure 1: Historical distribution of the overtime within the blocks allocated for the ophthalmology department in 2009 compared to the simulation based on Π_{FFD} scheduling rule.

days, and the 95% CI is [36.43,41.53] (the initial backlog was set to 20 days). Finally, the total average overtime over during the year was 1284 minutes (95% confidence interval = [1271.59,1298.16]).

Table 1: Simulation results for probabilistic scheduling rule.

Threshold parameters	Overtime (minutes)	Throughput(surgeries)	Backlog(days)	Utilization
L=0, $\alpha=5\%$	78	563.92	106.87	64.26%
L=0, $\alpha=10\%$	162	584.49	90.71	66.56%
L=0, $\alpha=15\%$	253	599.05	77.6	68.11%
L=30, $\alpha=5\%$	283	601.89	75.86	68.61%
L=30, $\alpha=10\%$	514	622.77	60.66	70.60%
L=30, $\alpha=15\%$	747	638.04	52.44	71.86%
L=60, $\alpha=5\%$	833	640.99	48.75	72.20%
L=60, $\alpha=10\%$	1346	662.9	38.52	73.64%
L=60, $\alpha=15\%$	1837	676.94	28.68	74.45%

The probabilistic scheduling rule Π_{FFP} was evaluated using different values for the threshold parameter L and overtime probability α . Table 1 provides the average overtime, throughput, backlog and utilization performance metrics for different policies. The set of scheduling policies provides an easy way to control the output performance metrics based on management preferences. This is difficult to achieve using just the point estimates for scheduling. These results highlight the effect of competing priorities that are common to OR scheduling problems; the patient backlog and room utilization can be improved only at the cost of increased overtime.

4.1 Analysis of the Approximation Within the Probabilistic Scheduling Rules

In a second set of experiments, we compared the scheduling decisions for the probabilistic rule Π_{FFP} (which uses Monte Carlo estimates of the probabilistic feasibility modeled by (1)) and the decisions based on a normal approximation that uses a simplified chance constraint given by (2). Once again, 100 one-year replications were simulated for the ophthalmology department. Decisions about whether or not to include a new arrival to the surgical unit in the block based on a sample of 10,000 realizations of random Gamma

Table 2: Comparison of normal approximation with Monte Carlo simulation.

Approximation	False positives	False negatives
Normal	0%	0.4%
MC 50 scenarios	1.1%	2.5%
MC 100 scenarios	0.7%	1.6%
MC 300 scenarios	0.3%	1.2%
MC 500 scenarios	0.2%	1.0%

distributed durations were considered to be “true” decisions in our tests. These “true” decisions were cross-checked with recommendations provided by the normal approximation scheme and the Monte Carlo method with varying sample sizes. The percentages of false positives and false negatives as compared to the “true” recommendations are presented in Table 2. It may be seen that the normal approximation provides extremely accurate decisions for a mix of surgeries and duration distributions. This approximation is often ignored in the research on operating room scheduling and preference is given to Monte Carlo estimates. The implications of the observed accuracy of the normal approximation are two-fold. First, expensive Monte Carlo calculations within a simulation model can be substituted with an much simpler evaluation of the formula (2) without any significant loss of accuracy. Second, the use of the normal approximation allows us to provide exact optimization models without the need to resort to scenario based approaches. This opportunity is often ignored by the current models of optimal operating room scheduling (Min and Yih 2010, Denton et al. 2010). To obtain performance comparable to the one using the normal approximation one needs a large number of samples (see Table 2). Therefore, the optimization models based on the normal approximation scheme have the potential to provide better solutions with much lower computational effort when compared to scenario based approaches.

However, it is important to stress that our results are based solely on the ophthalmology department at the VAPHS, and even though we make the reasonable assumption that similar results might hold across other specialties and hospitals, a proper verification is required before drawing any conclusion.

4.2 Analysis of Adaptive Scheduling Using a Control Chart Scheme

The simulation results of the previous section suggest that for any fixed scheduling policy the waiting time for a patient (time from an initial request to actual surgery day) is subject to large variability. This is illustrated in Figures 2 and 3. One can immediately see that the scheduling policy with overtime threshold $L = 30$ and probability parameter $\alpha = 5\%$ leads to a steady increase in average waiting times and large variability across different replications. On the other hand, an aggressive scheduling policy with overtime threshold $L = 60$ and probability parameter $\alpha = 10\%$ provides low levels of waiting times (at the cost of high overtime). Both of these policies are undesirable from a scheduling perspective: effective management of the surgical suite requires that waiting times do not rise to undesirable levels, while at the same time shorter waiting times do not result in unnecessarily high levels of overtime.

Previous discussion suggests that the choice of L and α can be used to manage performance metrics. In order to reduce the variance and provide stability with respect to the average waiting times, we investigate a scheduling policy that dynamically adjusts its scheduling settings. In our experiments, the dynamic approach chooses one of the scheduling parameters that were reported in Table 1. The adjustments of the scheduling parameters are made on a weekly basis. The average waiting time for all patients in the current waiting list is calculated and plotted on a control chart. The control chart is used to monitor the current state of patient waiting times; when predefined control limits are reached, interventions are triggered for bringing the process back under control. The upper and lower control limits of the control chart were set to 15 days and 5 days, respectively. If the new average value exceeds the upper control limit, the current scheduling parameters are changed to a more aggressive setting. On the other hand, the lower control limit is used to provide the transitions into less aggressive scheduling policies. The set of nine policies presented in Table 1 was sorted




Figure 2: Box plot of average patient waiting time for the scheduling policy with $L = 30$ and $\alpha = 5\%$.




Figure 3: Box plot of average patient waiting time for the scheduling policy with $L = 60$ and $\alpha = 10\%$.

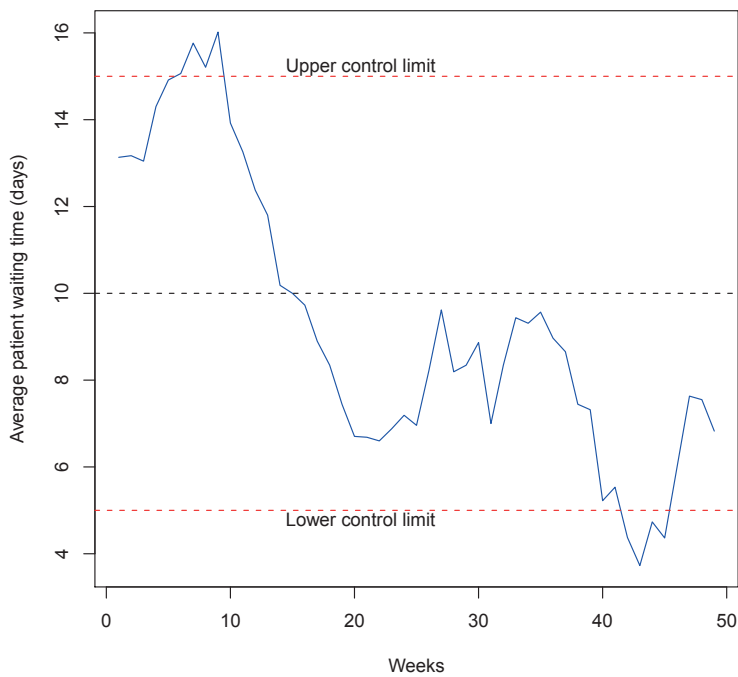


Figure 4: An example of the dynamic scheduling policy based on control chart.

by their average total overtime values and the aggressiveness of each policy was determined by this ordering. Transitions are made between policies that have consecutive positions on the list. For example, if the current setting is $(L = 30, \alpha = 15\%)$, then we either switch to a more aggressive setting $(L = 60, \alpha = 5\%)$, or to a less aggressive $(L = 30, \alpha = 10\%)$ setting. An example of this process is presented in Figure 4.

Figures 5 and 6 demonstrate the effect of the dynamic policy on the average waiting times and their variability, compared to the stable static policy for a given arrival rate. Each rectangle in the box plots




Figure 5: Box plot of average patient waiting time for the scheduling policy with $L = 30$ and $\alpha = 15\%$.




Figure 6: Box plot of average patient waiting time for the dynamic scheduling policy.

shows the interquartile range (IQR) for the average waiting time for the corresponding week and is based on 100 simulation runs. The IQR range goes from the first quartile (the 25th percentile) to the third quartile (the 75th percentile), whereas the whiskers correspond to the 1.5 IQR Tukey whiskers. One can see that the average waiting times for both policies (static and dynamic) are quite stable. However, the range of waiting times experienced across the replications is much smaller in the case of the dynamic strategy, demonstrating a process that has been brought under control. The dynamic policy allows waiting times to be managed, but without more overtime penalties than is necessary to accomplish that goal.

5 CONCLUSIONS

This article describes a discrete event simulation model of the operating room suite that allows us to study various scheduling policies and their effect on the performance metrics. Flexible probabilistic scheduling policies are proposed and evaluated using the simulation model. A dynamic scheduling policy based on Shewhart charts demonstrates that the variability and scale of the patients' waiting times can be effectively reduced and managed within the predefined limits. We compare the performance of Monte Carlo simulation with a Gaussian approximation to a convolution of gamma variates for modeling surgery durations. Our analysis suggests that the approximation approach provides excellent accuracy within the scheduling framework. These results point to the potential for exact optimization models - something that is often ignored in the literature on operating room scheduling - since most existing approaches lean towards Monte Carlo simulation. Our results also point to the need for future work on optimization models based on Gaussian approximations and the use of simulation techniques for long term planning and management of surgical suites.

ACKNOWLEDGMENTS

This material is based upon work supported by the Department of Veterans Affairs, Veterans Health Administration, Office of Research and Development. Our research was supported in part by the University of Pittsburgh Center for Simulation and Modeling through the supercomputing resources provided. We specifically acknowledge the assistance of Dr. Levent Yilmaz. The authors would like to thank the staff of the VA Pittsburgh Healthcare System (VAPHS) for their assistance and feedback.

REFERENCES

- Beaulieu, P. A., J. H. Higgins, L. J. Dacey, W. C. Nugent, G. R. DeFoe, and D. S. Likosky. 2010, October. "Transforming administrative data into real-time information in the Department of Surgery." *Quality & Safety In Health Care* 19 (5): 399–404.
- Cardoen, B., E. Demeulemeester, and J. Beliën. 2010, March. "Operating room planning and scheduling: A literature review". *European Journal of Operational Research* 201 (3): 921–932.
- Denton, B. T., A. J. Miller, H. J. Balasubramanian, and T. R. Huschka. 2010, July. "Optimal Allocation of Surgery Blocks to Operating Rooms Under Uncertainty". *Operations Research* 58 (4-Part-1): 802–816.
- Denton, B. T., A. S. Rahman, H. Nelson, and A. C. Bailey. 2006, December. "Simulation of a multiple operating room surgical suite". In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, WSC '06, 414–424. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Dexter, F., A. Macario, R. Traub, M. Hopwood, and D. Lubarsky. 1999. "An operating room scheduling strategy to maximize the use of operating room block time: computer simulation of patient scheduling and survey of patients' preferences for surgical waiting time." *Anesthesia and analgesia* 89 (1): 7–20.
- Firth, D. 1991. "Generalized linear models". In *Statistical Theory and Modelling. In Honour of Sir David Cox*, edited by D. Hinkley, N. Reid, and E. Snell, Chapter 3, 55–82. FRS. London: Chapman & Hall.
- Gordon, T., S. Paul, A. Lyles, and J. Fountain. 1988. "Surgical unit time utilization review: Resource utilization and management implications". *Journal of Medical Systems* 12 (3): 169–179.
- Hans, E., G. Wullink, M. van Houdenhoven, and G. Kazemier. 2008. "Robust surgery loading". *European Journal of Operational Research* 185 (3): 1038 – 1050.
- Healthcare Financial Management Association 2003. "Achieving operating room efficiency through process integration". *Healthcare Financial Management* 57 (3): suppl 1–7. PMID: 12793445.
- Min, D., and Y. Yih. 2010, November. "Scheduling elective surgery under uncertainty and downstream capacity constraints". *European Journal of Operational Research* 206 (3): 642–652.
- Sobolev, B. G., V. Sanchez, and C. Vasilakis. 2011, February. "Systematic review of the use of computer simulation modeling of patient flow in surgical care." *Journal Of Medical Systems* 35 (1): 1–16.
- Strum, D., J. May, and L. G. Vargas. 2000. "Modeling the uncertainty of surgical procedure times: comparison of log-normal and normal models". *Anesthesiology* 92 (4): 1160–1167.
- Winkel, P., and N. F. Zhang. 2007, June. *Statistical Development of Quality in Medicine*. 1 ed. Wiley.

AUTHOR BIOGRAPHIES

OLEG V. SHYLO is a Visiting Assistant Professor in the Department of Industrial Engineering at the University of Pittsburgh. He received a B.S. in Applied Mathematics from the National Technical University of Ukraine in Kiev in 2004, and a PhD in Industrial Engineering from the University of Florida in 2009. His research interests include discrete optimization, parallel computing and optimization in health care. His email address is olegio@gmail.com.

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 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.

OLEG A. PROKOPYEV is an Assistant Professor in the Department of Industrial Engineering at the University of Pittsburgh. He received his B.S. and M.S. in applied mathematics and physics from Moscow

Institute of Physics and Technology (Russia), and his Ph.D. in industrial and systems engineering from the University of Florida. His primary research interests are in the areas of combinatorial optimization, stochastic programming, integer programming, and applications in healthcare and bioinformatics. His email address is prokopyev@engr.pitt.edu.

JAYANT RAJGOPAL has been a faculty member in the Department of Industrial Engineering at the University of Pittsburgh since 1986. He received a B.Tech degree in Chemical Engineering from the University of Madras in India, and MS and PhD degrees in Industrial and Management Engineering from the University of Iowa. His interests are in theoretical, computation and applied aspects of optimization. In particular, he works on applications to problems in production, operations, and healthcare delivery. His email address is rajgopal@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.