Reducing Surgical Ward Congestion through Improved Surgical Scheduling and Uncapacitated Simulation

Vincent S. Chow[1], Martin L. Puterman[2], Neda Salehirad[3], Wenhai Huang[1], Derek Atkins[3]

1. British Columbia Cancer Agency.600 West 10th Avenue, Vancouver BC, Canada V5Z 4E6

2. Centre for Health Care Management, University of British Columbia. 209 - 2053 Main Mall, Vancouver BC, Canada V6T 1Z2

3. Centre for Operations Excellence, University of British Columbia. 2053 Main Mall, Vancouver BC, Canada V6T 1Z2

Contact Information

Vincent S. Chow	*vchow@bccancer.bc.ca*	*604-877-6000*
Martin L. Puterman	*martin.puterman@sauder.ubc.ca*	*604-822-8388*
Neda Salehirad	*neda_sr@yahoo.com*	*604-822-1800*
Wenhai Huang	*khuang@bccancer.bc.ca*	*604-877-6000*
Derek Atkins	*derek.atkins@sauder.ubc.ca*	*604-822-9665*

0 Abstract

High surgical bed occupancy levels often result in heightened staff stress, frequent surgical cancellations and long surgical wait times. This congestion is in part attributable to surgical scheduling practices which often focus on the efficient use of operating rooms but ignore resulting downstream bed utilization. This paper describes a transparent and portable approach to improve scheduling practices which combines a Monte Carlo simulation model and a mixed integer optimization model. For any surgical schedule, the simulation samples from historical case records and predicts bed requirements assuming no resource constraints. The mixed integer optimization model compliments the simulation model by scheduling both surgeon blocks and patient types to reduce peak bed occupancies. Scheduling guidelines were developed from the optimized schedules to provide surgical planners with a simple and implementable alternative to the optimization model. This approach has been tested and delivered to planners at a health authority in British Columbia, Canada. The models have been used to propose new surgical schedules and to evaluate the impact of proposed system changes on ward congestion. Insights from the models are guiding future surgical schedule revisions.

Key Words: Surgical Scheduling, Mixed Integer Optimization, Monte Carlo Simulation, Scheduling Guidelines, Hospital Bed Management

1 Introduction

In recent years, national governments have placed a high priority on reducing surgical wait times; Postl (2006) provides a Canadian perspective on this issue. This pressure, in combination with the increasing demands of a growing and aging population has led to a significant rise in elective surgical volumes. This increase in surgical volumes has led to high utilization rates of hospital and specifically surgical beds. High surgical bed utilization poses a challenge to management because:

- Peaks in demand can result in periods of insufficient capacity.
- Lack of capacity in speciality specific surgical wards may result in sub-optimal treatment and extended lengths of stay for surgical patients.
- Available capacity on surgical wards may be occupied by *medical* (non-surgical) patients limiting future access to surgical patients.

The result of this is often cancellations and long wait times for surgeries.

The day on which a surgeon operates is determined by the *surgical block schedule (SBS)*. The SBS assigns full or half-day blocks of operating room (OR) time to surgeons over a multi-week cycle. Usually, surgical planners consult with surgeons when constructing the SBS, while surgeons schedule patients within the blocks. Since each surgery type has a specific post-operative bed need, the SBS provides a significant lever for managing bed utilization. Creating an SBS is challenging, especially in the absence of planning software. Surgical planners must account for staff, room and equipment availability, and surgeon preferences when developing an SBS and often require many revisions just to take these considerations into account. However, being able to assess the impact of this schedule on downstream bed utilization is beyond the scope of most scheduling activities.

The main contribution of this paper is the development of a transparent, portable, and rigorous approach to improve surgical scheduling. It consists of two components, a trace-driven Monte Carlo simulation model and a mixed integer optimization model. The simulation enables surgical planners to predict the impact of a SBS on surgical ward occupancies. The optimization model schedules both surgeon blocks and patient mix within each block to help planners create surgical schedules with minimal bed requirements. Since managers and planners will not be able to run the optimization model, a set of scheduling guidelines was developed to support ongoing scheduling revisions.

Our paper is organized as follows. Section 2 provides a review of recent literature. Section 3 provides a description of the models. Section 4 provides a case study that illustrates the application of our approach. Section 5 presents study results and illustrations of how the models can be used and Section 6 concludes with general insights and future research.

2 Related Literature

Operating room planning and scheduling has been widely studied; a review paper by Cardoen et al. (2008) describes recent advances in this area. We describe other relevant research below.

Researchers frequently use Discrete Event Simulation (DES) to model complex patient flow and to estimate resource utilization in hospitals that cannot be accurately characterized using queuing models. Both Jun et al. (1999) and Jacobson et al. (2006) provide a review of DES applications in health care. In Surgical Services, Everett (2002) developed a decision support tool to evaluate various policies on wait lists and bed occupancies. VanBerkel and Blake (2007) developed a DES model to evaluate surgical wait times and support capacity planning decisions. Blake et al. (1995) proposed a surgical process management tool that utilizes summarized historical patient records directly in the model. For general patient flow within hospitals, Isken and Rajagopalan (2002) modelled obstetrical and gynaecological patient flow. Both Harper (2002) and Pitt (1997) developed generic frameworks for portable DES models of patient flows.

Accurate DES patient flow models often require a significant amount of time and resources to develop as they are often highly specialized for each system. Harper (2002) and Pitt (1997) have tried to address these issues but other challenges, such as the creation of clinically relevant groupings to define parameters within the model can be difficult (Isken and Rajagopalan, 2002). A retrospective paper by Carter and Blake (2005) also noted the challenges of creating a flexible DES model as in Blake et al. (1995).

Other methods of simulation have also been used. Henderson and Mason (2005) utilized a trace-driven simulation model to support ambulance-planning. The paper argued that the ability to maintain correlations between various parameters within the simulation outweighs the limitations of using trace data.

A large body of research concerns development of methods to improve OR scheduling. The scheduling process can be divided into three main scheduling levels: the discipline level, the surgeon level, and the patient level (Cardoen et al. 2008). Discipline level scheduling determines the amount of operating room time to allocate to each surgical specialty. Blake and Donald (2002) developed an integer programming model to produce equitable master surgical schedules. Beliën and Demeulemeester (2007) devised a mixed integer programming (MIP) model to minimize the expected bed shortages by relocating specialty blocks. Santibanez et al. (2007) optimally allocated surgical specialties across a system of hospitals to explore wait lists and resource utilization under different objective functions.

Surgeon level scheduling determines the allocation of surgeon blocks to the available operating room time within each specialty. Beliën et al. (2009) expanded their previous model (2007) to level bed occupancy and variance by reassigning surgeons to different ORs on different days and integrated a deterministic support tool developed by Beliën et al. (2006) to visualize planned surgical occupancies under various surgical schedules. The support tool reports the average occupancy, standard deviation, and expected total bed shortage on each day.

Patient level scheduling involves the allocation of patients within the surgical blocks. Van Oostrum et al. (2008) used a two-phase decomposition approach with column generation and a MIP to determine the patient mix and the OR scheduled each day. Vissers et al. (2005) determined the OR time and the number of each patients in each patient category scheduled on each day of the week in a cardiothoracic surgery department. Other work, based on heuristics, includes maximizing operating room use within the day (Dexter and Traub, 2002) and determining various surgical sequencing procedures to reduce idling and overtime (Denton et al. 2007).

Various papers have integrated simulation and optimization techniques. Testi et al. (2007) used a three-phase approach to scheduling ORs. The first two phases involved the allocation of OR time and OR blocks to specialties. The third phase involved in the execution of the schedule using a DES model to evaluate the sequencing of surgical activities. Persson and Persson (2009) developed a DES model to test policy changes so as to ensure elective surgeries wait no more than 90 days. An optimization component was integrated into the model to decide the scheduling of patients each week.

Our approach uses two models: the Bed Utilization Simulator (BUS), a trace-driven unconstrained Monte Carlo simulation model that predicts surgical bed occupancy and the

Surgical Schedule Optimizer (SSO), a MIP model to level surgical bed occupancy. Advantages of BUS over other patient flow models include a short development time and increased portability. The base formulation of the SSO model is similar to Beliën et al. (2009) but differs from all other surgical scheduling models by scheduling both surgeon blocks and the patient mix within each block. Both of these models are used in an integrated approach to improve surgical scheduling. Finally, our paper seeks to use the models to develop general scheduling guidelines for smoother ward occupancies.

3 Models

Figure 1 gives an overview of the proposed framework. The SSO model helps develop improved surgical schedules. Surgical planners can either obtain a SBS directly from the SSO, or use surgical scheduling guidelines derived from the SSO to help construct the block schedules. Surgical planners would then revise these schedules as necessary and test the schedules in BUS to capture both unplanned patients and variability in the surgical system. Results from BUS are analyzed and the schedules can be readjusted and re-simulated until a final SBS is obtained. Both BUS and SSO models are described below. Some surgical scheduling guidelines are described in Section 5.

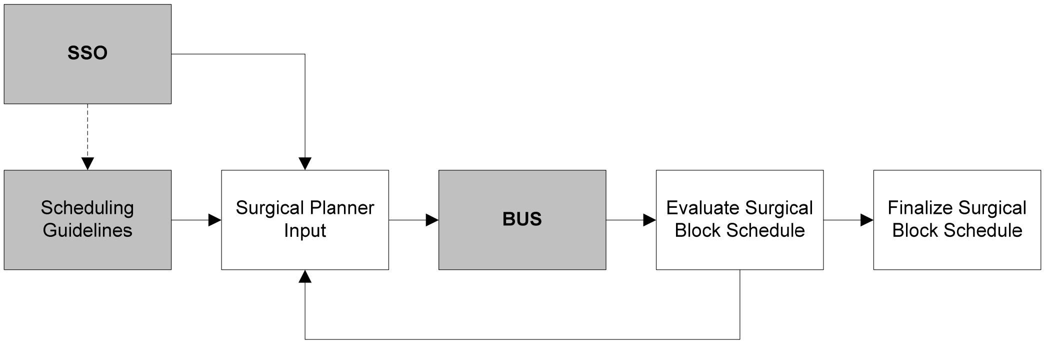

Figure 1. The proposed surgical scheduling framework. Models are shaded in grey.

3.1 The Bed Utilization Simulator (BUS)

BUS was designed to provide surgical planners with an easy-to-use tool to investigate the impact of an SBS on downstream ward volumes. Unlike other models which include “hard” ward capacities and decision rules to handle situations when wards are full, BUS predicts the *daily demand* for ward beds in an *uncapacitated system.* Hence, the primary output of BUS is the predicted daily bed occupancy in each downstream surgical ward, ignoring all other competitive demands for beds. This enables planners to determine the true needs of their system and set capacities and schedules to reflect these needs. This approach also allows for a simple model logic that can be broadly applied.

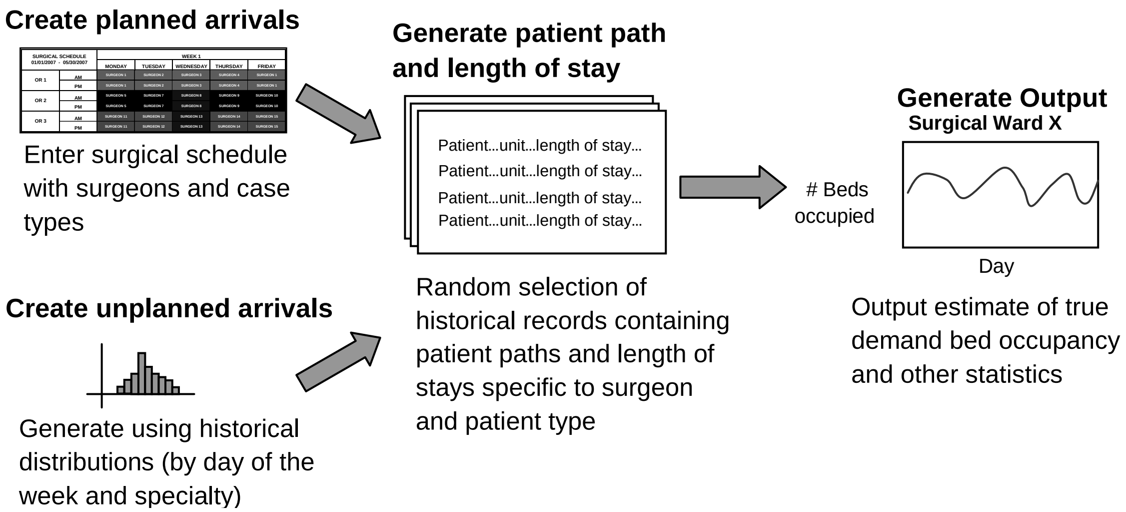

Figure 2. An overview of the BUS model logic.

Figure 2 illustrates the model logic. There are two sources of arrivals, planned patient arrivals are generated according to a user provided SBS and unplanned patient arrivals are generated to follow historical arrival rates. Planned patients are further classified into patient types based on their downstream bed requirements. The number of patient types can be adjusted to reflect variation in operation between various hospitals. For each surgical block, the number of patients of each patient type may be either specified or sampled from historical surgeon *slates*. A slate is a combination of patient types performed during a block. One slate may consist of one “Patient Type 1” case and three “Patient Type 2” cases while another slate may consist of one “Patient Type 1”, one “Patient Type 2” case and one “Patient Type 3” case. These slates would be obtained from historical slates, specific to each surgeon block. Our model assumes that all planned surgeries take place. The arrival of unplanned patients can be generated from a Poisson or empirical distribution with day of week specific arrival rates.

The model then generates ward occupancy patterns for each patient by randomly sampling an appropriate patient record from historical data. Planned patients are selected from records with the same surgeon and patient type, while unplanned patients are selected from records with the same specialty and patient type. For surgeons with limited data, records are selected from a similar surgeon or from a pool of records of the same specialty and patient type. Each record contains information on the sequence of surgical wards and length of stay in each ward the patient visited relative to the day of surgery. Finally, lengths of stays are aggregated for each ward to generate the daily demand for ward beds.

Figure 3 illustrates the structure of a historical patient record, and how the information is used to update ward occupancies. The attributes are read sequentially and written into tables corresponding to surgical wards. In this example, a patient of Dr. Smith spends one night in the ICU and the next three nights in Ward 1. The "*Rel.Start*" fields indicate the first night a patient occupies the unit relative to the operation date. Day *t* represents the day in which the operation is performed. It is possible to occupy beds prior to the operation. In such cases, some of the "*Rel.Start*" field would be negative. This process is repeated for each day of the simulated period.

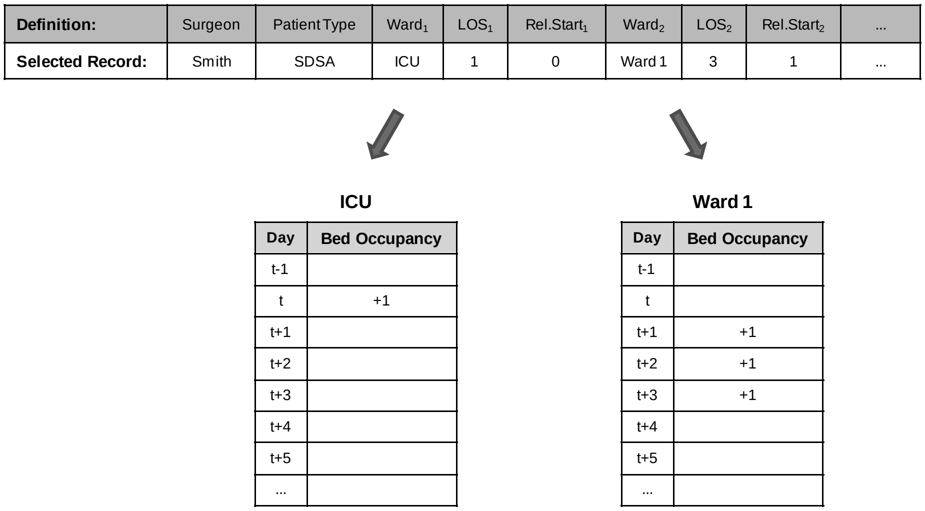

Definition:	Surgeon	Patient Type	Ward$_1$	LOS$_1$	Rel.Start$_1$	Ward$_2$	LOS$_2$	Rel.Start$_2$	...
Selected Record:	Smith	SDSA	ICU	1	0	Ward 1	3	1	...

ICU

Day	**Bed Occupancy**
t-1	
t	+1
t+1	
t+2	
t+3	
t+4	
t+5	
...	

Ward 1

Day	**Bed Occupancy**
t-1	
t	
t+1	+1
t+2	+1
t+3	+1
t+4	
t+5	
...	

Figure 3. Structure of a historical record how a patient is used to update ward occupancies

The metric *bed-days over capacity* summarizes a surgical schedule’s performance. This metric captures and quantifies the frequency at which patient off-servicing and/or surgery cancellation would be necessary if the system were capacitated. The metric is computed by summing the total number of beds in excess of a user-specified bed capacity for each day of the scheduling period within each surgical ward. The user-defined capacities are an input to the model and can vary by day of the week if required. We emphasize that the only purpose for including these capacities is to derive the metric; they don’t impact patient flow in any way since the model is uncapacitated. The model also produces distributions and descriptive statistics of ward bed occupancies.

Since this model relies on patient specific input data, the approach to selecting historical cases can highly influence results. Therefore, it is crucial that only cases with paths and length of stays that reflect clinical necessity be included in the historical record database. Historical data should be screened to remove planned cases with patient paths that are inconsistent with known patient requirements.

BUS was developed in MS Excel using Visual Basic for Applications (VBA) to facilitate portability, reduce software cost and reduce the need for simulation software expertise. The model interface was designed in collaboration with surgical planners and executed through a series of VBA forms and spreadsheets. BUS includes an enhanced spreadsheet to define surgical schedules (Figure 4A) and an output spreadsheet to display the simulation results (Figure 4B). The numbers of user-definable simulation parameters are kept to a minimum to facilitate use by non-experts. The two main user-definable parameters are the number of warm up days and the number of replications. Recommended values are set as defaults.

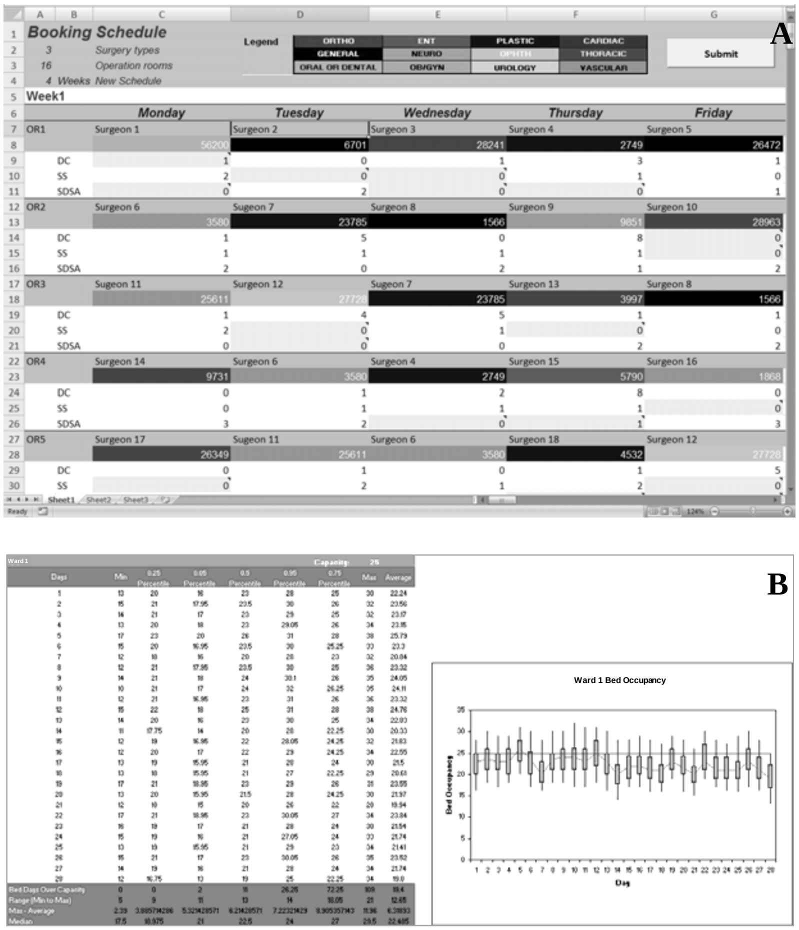

Booking Schedule
3 Surgery types
16 Operation rooms
4 Weeks New Schedule

ORTHO	ENT	PLASTIC	CARDIAC
GENERAL	NEURO	OPHTH	THORACIC
ORAL OR DENTAL	OB/GYN	UROLOGY	VASCULAR

Week1

		Monday	Tuesday	Wednesday	Thursday	Friday
OR1		Surgeon 1	Surgeon 2	Surgeon 3	Surgeon 4	Surgeon 5
		56200	6701	28241	2749	26472
	DC	1	0	1	3	1
	SS	2	0	0	1	0
	SDSA	0	2	0	0	1
OR2		Surgeon 6	Sugeon 7	Surgeon 8	Surgeon 9	Surgeon 10
		3580	23785	1566	9851	28963
	DC	1	5	0	8	0
	SS	1	1	1	1	0
	SDSA	2	0	2	1	2
OR3		Sugeon 11	Surgeon 12	Sugeon 7	Surgeon 13	Surgeon 8
		25611	27728	23785	3997	1566
	DC	1	4	5	1	1
	SS	2	0	1	0	0
	SDSA	0	0	0	2	2
OR4		Surgeon 14	Surgeon 6	Surgeon 4	Surgeon 15	Surgeon 16
		9731	3580	2749	5790	1868
	DC	0	1	2	8	0
	SS	0	1	1	1	0
	SDSA	3	2	0	1	3
OR5		Surgeon 17	Sugeon 11	Surgeon 6	Surgeon 18	Surgeon 12
		26349	25611	3580	4532	27728
	DC	0	1	0	1	5
	SS	0	2	1	2	0

Figure 4. A) A screenshot of the block schedule input interface. Here, the user can define surgeons, patient volumes and ward capacities. B) A screenshot of the output report. Calculated metrics are presented on the left hand table. The distribution of daily bed occupancies appears in the box plot on the right.

When one uses historical records directly in a simulation, the simulation is often said to be "trace-driven". Using historical records preserves some correlation patterns within the data. This can be either a disadvantage or an advantage. Law (2007) notes that when using trace-driven simulation, the model can only reproduce what has happened historically and may in fact be the result of system constraints or policy decisions that the simulation is being used to explore. On the other hand, one might wish to see how new system settings might have performed in the past so the trace driven approach will provide a reasonable base case. The BUS model has two types of correlation to consider: that within patients and that between patients. Within patient correlations are between patient type, length of stay and patient path. Ideally, these correlations should be preserved; especially if the data is cleaned so that non-desirable paths are removed. On the other hand, the correlation between patient arrivals should be avoided. Our approach preserves the within patient correlations and through random sampling of cases, eliminates between patient correlations. This allows us to capture the best of both worlds.

Other potential shortcomings of this modelling approach are that representative datasets are required for each surgeon and patient-type configuration, and that the model cannot be used to determine cancellation rates or numbers of patients off-serviced, i.e., those who are not assigned to the most appropriate ward. We addressed (see Section 4) the first shortcoming by careful design of data inputs and development of approaches for addressing cases with little data. We regard our choice of the metric “bed-days over capacity” as a surrogate for off-service and cancellation rates even though it cannot distinguish between the two. While discrete event simulations could determine these rates, strong assumptions would be required about cancellation and redirection policies which may not be explicit or available. Nonetheless varying the SBS to reduce bed-days over capacity should reduce cancellations and off-servicing which are two key objectives of these models.

3.2 The Surgical Slate Optimizer (SSO)

While BUS enables planners to investigate bed requirements of new surgical schedules, it remains extremely challenging for surgical planners to create a SBS that reduces day-to-day variability in bed occupancy across all surgical wards. Therefore, two integer optimization models were developed to generate improved surgical schedules. The *base model* assigns surgical blocks to days to achieve smoother bed occupancies. The *slate model* generalizes the base model by also determining the patient type mix (see above) within each block. Mathematical formulations of these models appear in Appendix 1 and 2. Each model assumes a multi-week block schedule.

In the base model, binary decision variables indicate whether a specific surgical block is scheduled on a specific weekday of a multi-week schedule. Each surgical block represents a unique a surgeon and duration combination that can be scheduled multiple times throughout the schedule. The duration of each block is expressed in terms of *OR-Days.* A surgeon using one OR for one day would be assigned 1 OR-Day while a surgeon using one OR for half a day would be assigned 0.5 OR-Days. It is important to note that for the same surgeon, a half day block is not equivalent to half of a full day block. Both the patient mix and length of stay can vary significantly between these two surgical blocks and they are modelled separately to capture this difference. Note, if each surgeon has only one block length, then the number of surgeons and surgeon blocks would be the same. We simplified the models by not including the decision of which operating room (OR) to assign to a surgical block; that is left for the surgical planners. If there are additional constraints where only specific types of specialties can be performed in specific ORs, then the decision variables must be expanded to assign surgical blocks to both days and ORs.

The objective of the model is to minimize the total maximum bed occupancy across all wards. By doing this we expect that day to day variation in ward occupancies will be minimized without resorting to nonlinear objectives. Base model constraints include:

- Total OR capacity cannot be exceeded each day.

- A surgeon cannot have more OR-Days than he/she can possibly do on one day. Usually this will be 0.5 or 1 but if a surgeon can operate in two ORs in sequence then this could be also 2.
- An upper bound on the number of each surgical block scheduled per week.
- An equality constraint which specifies the total number of blocks to be assigned to a surgeon over the entire scheduling cycle. Usually this will be chosen to be consistent with the “as is” frequencies at the time of the study.
- A bookkeeping constraint which defines the maximum number of beds needed in each ward on each day. Due to the cyclic nature of the schedule, any surgical blocks that are scheduled near the end of the schedule would result in additional bed demand in the beginning of the surgical schedule.

The slate model builds upon the base model by taking into account the patient mix within each surgical block. To reduce model size, we selected a few pre-defined slates from historical data to restrict the feasible patient mixes for each surgical block. Hence the binary decision variables determine whether a specific surgical block with a specific slate is assigned on a specific day. It was believed that with this added flexibility, downstream bed requirements could be reduced.

The slate model requires two additional constraints. The first constraint ensures that only one slate is chosen for each block. The second constraint ensures that the number of surgical cases performed by each surgeon is as least as great as the average number of surgical cases performed in the past. This ensures that no surgeon is penalized with fewer surgeries under the optimized schedules. The objective function of the slate model remains the same as the base model.

4 Case Study

The approach was applied at the Royal Jubilee Hospital (RJH) in Victoria, B.C. RJH is the largest tertiary care hospital for Vancouver Island's 750,000 residents and as well serves as a regional hospital for approximately 350,000 residents in Victoria. In 2007, RJH ran 16 operating rooms (ORs) and 428 inpatient beds, of which about 100 are "reserved" for surgical patients. At the time of our study, high variability in demand for surgical beds had resulted in an increased risk of surgical cancellations during peak occupancy periods. Also, shortages of medical beds during times of increased emergency department congestion placed increasing pressure to use surgical beds to accommodate medical (non-surgical) patients who require a bed for monitoring and/or treatment. The consequence of this was that beds planned for use by surgical patients become blocked. The intent of the study was to reduce surgical cancellations through smoothing bed requirements for planned surgical cases. Other benefits of doing this include lower costs or increase access.

4.1 Background

Figure 5 provides an overview of surgical patient flow in the system studied. It shows two primary inputs: planned elective surgeries scheduled up to a week in advance and unplanned surgeries that result primarily from emergencies. After surgery, most patients stay temporarily in the Post Anaesthesia Recovery Room for monitoring prior to being transported to one of six surgical wards. Patients in unstable condition or who have received major cardiac surgery stay instead in the Intensive Care Unit or the Cardiac Care Unit prior to being redirected to surgical wards. The six surgical wards at RJH are categorized by the acuity of the surgery. Day Care or Short Stay Wards are for patients with minor surgeries or surgeries requiring overnight stay.

Both wards operated only on weekdays, and patients who did not meet discharge criteria are transferred to Wards 1-4 on weekends. Wards 1-4 are for patients with major surgeries and are assigned to specific surgical specialties.

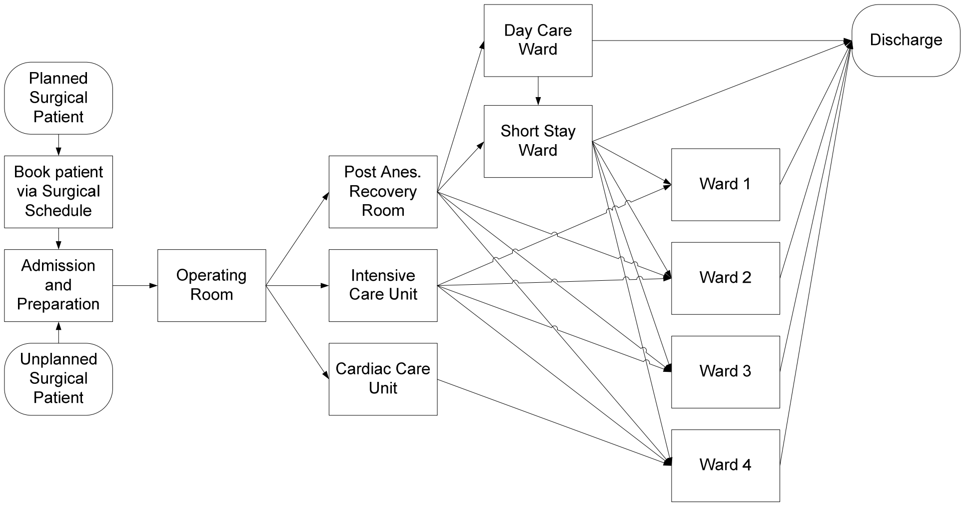

Figure 5. A high level process map of surgical patient flow at RJH.

RJH uses a 4-week cyclical SBS that assigns full or half-day blocks to surgeons. Planned patients are categorized into three patient types: Day Care (DC), Short Stay (SS) and Same Day Surgical Admit (SDSA). Post-surgery, DC and SS patients stay in corresponding wards. SDSA are inpatients that typically sojourn to Wards 1-4 depending on the type of surgery.

4.2 The Problem

Figure 6 provides a nine month time series of surgical bed occupancy, staffed surgical bed capacity, and surgical cancellations in Wards 1-4 and clearly illustrates the challenges health system managers faced in operating this system. Observe that:

- surgical volumes were highly variable and that the system was always operating close to or above capacity,
- a significant portion of the surgical capacity was occupied by non-surgical (medical) patients who were assigned to these units when space was available, and
- towards the end of 2006, there were frequent surgical cancellations when the system was over capacity and beds were not available for surgical patients in these wards. (Cancellations may also occur for patients who cannot be off-serviced to an alternate surgical ward when their target ward is at capacity)

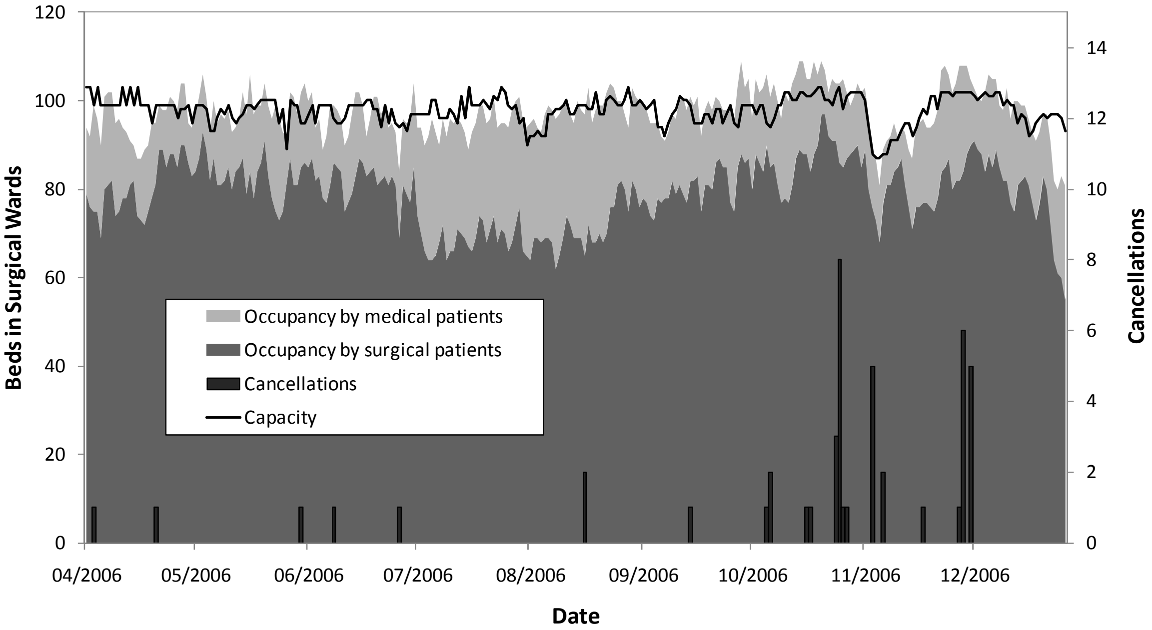

Figure 6. Time series of bed occupancy, staffed bed capacity, and surgical cancellations in Wards 1-4. The bed capacity is adjusted for bed closures due to staffing issues, maintenance, and outbreaks.

Further analysis of occupancy patterns within surgical wards reveals that unplanned patients occupy a larger proportion of surgical beds (54%) than planned patients (46%). However, when analyzing this occupancy across the days of the week, volumes of unplanned patients varied around a common mean, while planned patients exhibited systematic day to day variation. Figure 7 provides box plots of daily bed occupancy of planned patients in a typical ward which shows that median occupancy varied between 6 on Mondays to 10 on Fridays and also that there was considerable within day variability. Given the stochastic nature of unplanned arrivals and the imbalance in planned patient occupancy, we focused our efforts on improving the scheduling of planned patients; the only lever directly available to managers in this system.

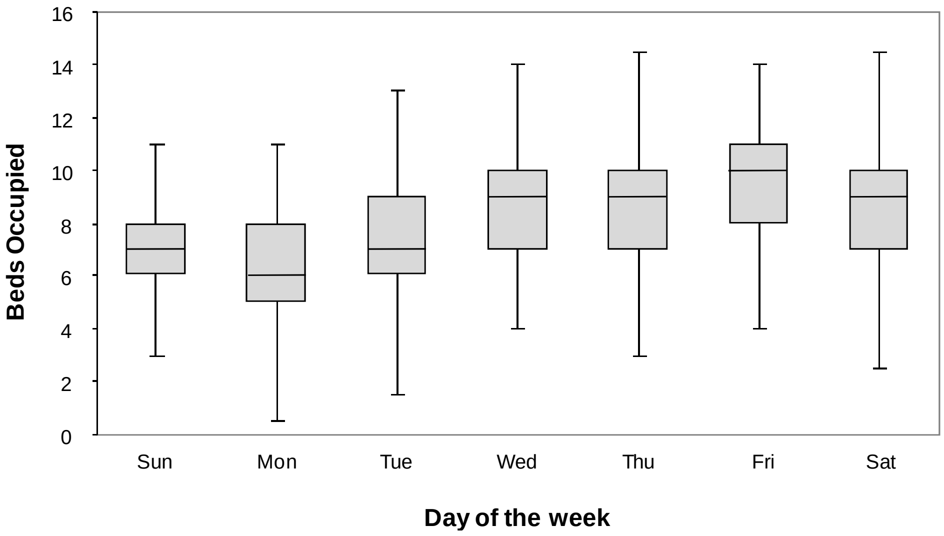

Figure 7. A box-plot of daily bed occupancy by planned patients in the surgical Ward 1.

4.3 Applying the Models

A large portion of the study concerned obtaining and analyzing data. Our primary data sources were; the Admissions, Discharge, Transfer System (ADT) and the Operating Room Scheduling Office System (ORSOS). ADT is a patient tracking software that stores admission, transfer or discharge time stamps. ORSOS is a scheduling and management system for surgical services within the hospital. Both databases were linked via patient identification and surgery times to create the dataset for modelling. This data was then crosschecked with a patient discharge database and a manual nightly bed census count to insure data integrity. The final dataset contained the following information for each case:

- responsible surgeon and surgical specialty
- patient type (DC, SS, SDSA)
- method of entry to the hospital (elective or emergency surgery)
- surgery date
- sequence of post-surgical wards occupied, and the length of stay within each

We observed that planned cases with patients off-servicing to non-designated surgical wards were largely caused by operational as opposed to clinical reasons. Upon reviewing data with ward managers, 7.7% of the data had paths that were inconsistent. Removing these cases resulted in minor change to the average total length of stay of inpatients (1.5% decrease) with the difference being statistically insignificant at the 0.05 level. A comparison between each surgeon's mix of cases in the final database and the historical mix ensured that a similar proportion of procedures were represented.

The data was then entered into the BUS model and validated against historical records. A direct validation of the model-reported bed occupancies was not possible since BUS was an unconstrained model and intended to determine the impact of a SBS in an ideal world. Since the logic was simple and historical data was used, we had confidence in our results. On the other hand, we validated results graphically for a specific ward which had few patient relocations and cancellations (Figure 8). We observed that when using historical surgical schedules to generate planned arrivals and randomly generating unplanned arrivals using historical distributions, the simulated bed occupancy patterns are similar to historical patterns. Results for other wards were reviewed by hospital managers and planners to determine whether the model reported over/under capacity patterns reflected the level off-servicing and cancellations observed in reality.

This dataset was also used to compute the bed demand parameters and surgical block compositions for the SSO. However, there are several unique characteristics at RJH that required minor modifications to the model. ORs are not a limiting constraint at RJH and Urology surgeons are assigned two ORs when they operate. These surgeons alternate between both ORs to minimize their idle time caused by the setup and cleanup. Thus a urologist is assigned two blocks on days when he/she operates. In addition, two ORs were equipped to perform all specialties; while the remaining 14 ORs can be used by any specialty except urology. Therefore, a constraint is placed to limit the number of ORs allocated to Urology each day. See Appendix 3 for details.

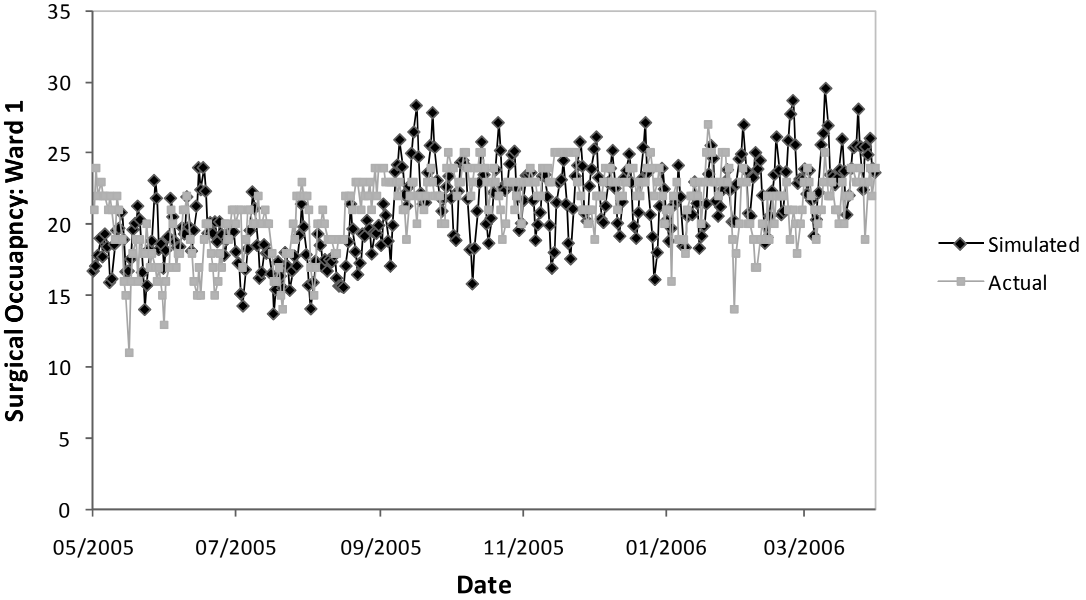

Figure 8. Actual vs. Simulated occupancies for Ward 1. Simulated data was generated from 20 runs using historical arrival times of planned patients and random arrival times of unplanned patients. Each simulated case was randomly selected by the model.

5 Case Study Results and Analysis

Both SSO and BUS were used to explore opportunities to reduce peak surgical bed utilization and smooth out day-to-day variability in bed occupancy. The impact of using the two versions of SSO will be discussed first. Next, the development of surgical scheduling guidelines and the resulting rules will be presented. Finally, the use of BUS as a standalone tool to evaluate various ‘what-if’ scenarios is illustrated with examples.

5.1 Optimized Surgical Schedules

In our application, SSO included 47 surgeons, 74 unique surgical blocks, and 8 surgical wards over a 4 week scheduling cycle. The base model consists of 1488 decision variables and 1574 constraints while the slate model with 2 slate choices in each block consists of 2968 decision variables and 3195 constraints. The models were solved with CPLEX 11 using GAMS on an Intel Q6600 with 2 GB of RAM. Both the base model and the slate model provided good feasible solutions (optimality gap ~5%) within 10 minutes. Optimality could not be reached after 6 hours of computation and only minor decreases were observed in the optimality gap (<1%). Methods to reduce computational time through heuristics and alternate model formulations are areas for future work.

The SSO base model output, prior to evaluation with BUS, suggests that significant reduction in average peak occupancy can be achieved in each ward. Figure 9 shows that prior to optimization there was considerable variability in average bed occupancies between days in all wards. For example, in Ward 3, the maximum average occupancy was 11.5 beds on Day 23 and the minimum average occupancy was 4 beds on Day 16. Post optimization, the maximum bed occupancy dropped to 9 and the minimum bed occupancy increased to 7. This effectively reduced peak occupancies by 2.5 beds and reduced the range of average bed occupancies from 7.5 to 2. Reductions in peak occupancies were observed across all wards and the total of these decreases was predicted to be 8.5 beds. Consequently there is strong evidence that variation caused by planned surgical patients can be reduced using the optimized schedule.

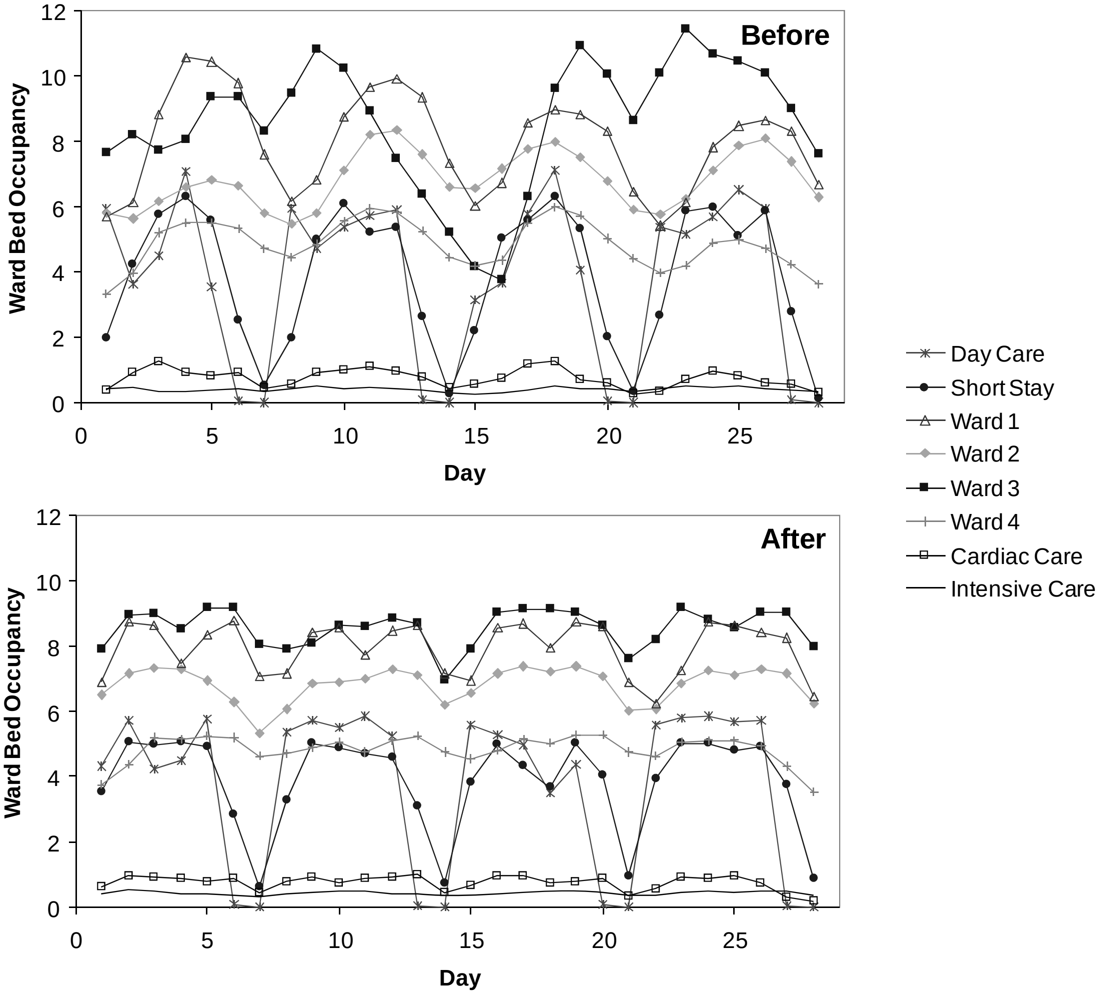

Figure 9. Modelled average occupancy by planned patients in each surgical ward before and after optimization using the base case optimization method.

The optimized schedule was then entered into BUS to assess the impact of variability and unplanned patients on ward occupancies. Wards were initialized to zero at the beginning of every run/replication. A warm up was used to allow the wards to reach a pseudo “steady state” prior to collecting performance statistics. To determine an adequate duration for the warm up period, the ward occupancies of several runs were analyzed graphically to determine the number of simulated days necessary to achieve stable bed occupancies. It was found that 6 scheduling cycles (24 weeks) was sufficient under all scenarios. Statistics were then collected from the seventh surgical schedule cycle. This entire process is repeated for subsequent replications.

In our study, the total number of bed-days over capacity prior to optimization was 80 over a total of 277 inpatient cases in a 4 week period. The optimization resulted in a reduction of 16% or 13 bed-days over capacity. The practical significance of this is that up to 13 instances of patient relocation to inappropriate wards and/or surgery cancellations could be avoided over a 4 week period.

The same analysis was repeated using the SSO slate model. The total number of cases performed with the optimized schedule was higher than the initial schedule due to the lower bound constraint on the number of cases performed by each surgeon. The optimized model resulted in an increase of 15 surgical cases (4 SDSA and 11 SS) per 4 week period. These additional cases increased surgical volume by 5%. The total number of bed-days over capacity using this schedule was reduced by 9% to 7. This result is significant, demonstrating that the hospital could both serve additional surgical cases and reduce the number of bed-days over capacity simultaneously.

Our decision to limit choice to two slates per surgeon was driven primarily by computational issues. However because of the limited number of slates, the resulting model might be infeasible because the minimum demand constraint could not be satisfied. Slates were pre-screened to ensure a feasible solution could be found. An alternative method to address this issue would have been to make the demand constraint a soft constraint.

5.2 Surgical Scheduling Guidelines

One of the limitations of SSO is its inaccessibility to hospital planners. Past experience suggests that optimization models are difficult to operate without expert knowledge; especially when dealing with issues regarding infeasible solutions. Also, modification of constraints and continual upkeep of the model to changing operational parameters would be difficult for planners to perform. Therefore, it was concluded that the development of scheduling guidelines would be beneficial to planners.

Optimized schedules from both model formulations exhibited characteristics that were significantly different than surgical schedules in use at the time of the study. This was especially evident in the slate optimized schedules where the selection of surgical blocks and slates exhibited recurring patterns. These patterns have been analyzed and formulated into the following scheduling guidelines:

1. Group surgical blocks with similar ward requirements together. Usually a group will consist of one surgical specialty, but a group can include several specialties if these specialties share the same ward.
2. Within each group, schedule surgeon blocks with high patient volumes and long length of stay requirements (i.e. SDSA patients) at the beginning and the end of the week. This would increase the occupancy of under-utilized surgical wards at the beginning of the week and over the weekend. Surgeons with lower inpatient bed requirements can be scheduled in the middle of the week to maintain utilization rates across the week.

3. For wards that close on weekends, schedule surgeons with high demand for short length of stay cases (2 days) on Mondays and Wednesdays. Scheduling primarily on Mondays and Wednesdays maximizes ward utilization and minimizes patient off-servicing to inpatient wards on the weekend.

The main principle behind these guidelines is that surgical blocks should be scheduled based on both surgical ward requirements and patient mix. Blocks of the same speciality tend to have similar ward requirements but their patient mix can vary significantly. Some surgeons may operate exclusively on outpatient cases while others on inpatient cases, both of which have highly different downstream requirements. Only by understanding the differences in these requirements can surgical planners make more informed scheduling choices to smooth occupancies.

The proposed scheduling guidelines are general enough to apply to other hospitals with similar ward configurations. While true optima are not achieved using these guidelines, surgical planners can still create improved and implementable surgical schedules. In this way the optimization results *guide* practice, but does not *dictate* it.

5.3 Scenario Analysis Using the Bed Utilization Simulator

The Bed Utilization Simulator is also a powerful standalone tool that can be used on a “what-if?” basis to investigate schedule changes. We now describe how it was used to evaluate new surgical schedules, assess the impact of increasing knee replacement cases, assess the impact of changing ward capacities, and determine the number of surgical beds to protect in a shared surgical and medical ward. For most scenarios, two hundred replications were required to insure that the 95% confidence intervals for daily mean bed occupancies and bed-days over capacity were within acceptable levels (<0.2 beds, <0.5 bed-days over capacity). Additional replications were performed if these levels were not achieved.

5.3.1 Introducing a New Surgical Schedule

Hospital management was interested in determining the impact of a new surgical schedule that was planned to be implemented shortly after our study was completed. This schedule was evaluated using BUS; some results appear in Figure 10.

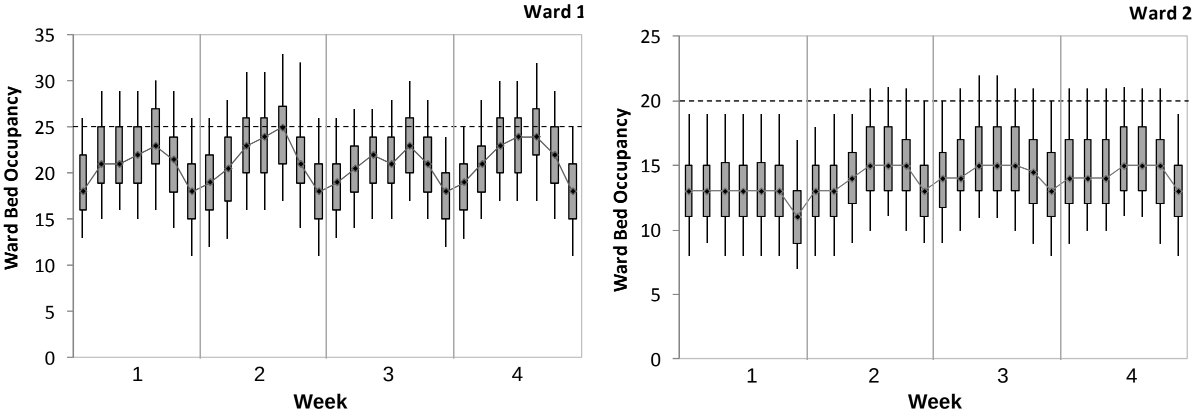

Figure 10. Simulated bed occupancies for Ward 1 and Ward 2. The dotted line represents the approved surgical capacity for the ward.

BUS identified several unique characteristics of the new surgical schedule. Results from Ward 1 show a consistent weekly pattern of low bed utilization at the beginning of the week and high bed utilization at end of the week. This result was expected as the new surgical schedule did not differ significantly from previous schedules under which similar patterns had been observed. The results also suggest that bed occupancies were slightly higher and more variable on the second and fourth week of the schedule. Closer inspection of the new surgical schedule showed that two orthopaedic surgeons whose cases required long lengths of stay were scheduled to operate on the same day during these two weeks. Simulation results for Ward 2 reveals relatively smooth surgical bed occupancies and lower utilization in comparison to Ward 1. Results also indicate that the first week has lower bed occupancies and upon reviewing the surgical schedule, this was attributed to a lack of surgeons with high inpatient demands scheduled on week 1.

5.3.2 Increasing the weekly volume of Knee Replacement Cases

Hospital management wished to investigate the impact of scheduling two additional knee replacement procedures every Monday to reduce current knee replacement wait times. To represent this scenario in BUS, a new surgeon who only performs knee replacements was created and added to the schedule. Results showed that scheduling two additional knee replacement cases would increase bed utilization across the week and as well, increase the likelihood that Ward 1, an orthopaedics ward, would be over capacity. A comparison of average median bed occupancies also indicates a significant change, increasing from 21.1 beds to 22.4 beds, while bed-days over capacity increased by 8. To reduce these effects, management could either add one additional bed to this ward to keep bed-days over capacity near previous levels or reschedule other surgeons to decrease the anticipated peak bed occupancies in this ward (increase bed-days over capacity by 3).

5.3.3 Changing Ward Capacities

BUS can also be used to evaluate whether the current allocation of beds across wards can be improved. Using bed-days over capacity as the main metric, sensitivity analysis can be performed on the current bed allocations to determine the impact of redistributing bed capacity.

Results from three wards appear in Table 1. Change in bed-days over capacity given an increase or decrease change in current bed capacity. They show that the largest reduction in bed-days over capacity can be achieved by adding one additional bed to Ward 1. If a bed were to be eliminated, removing it from Ward 2 would have the least impact on bed-days over capacity. Similar sensitivity results (less than 0.5 bed-days over capacity difference) are observed when changing capacities while using optimized schedules.

It is important to note that the actual change in capacity would also depend on the cost of operating one bed in each unit and whether there was staff and physical capacity to do so. This information can be used as one measure by managers to help redistribute current beds or plan for future bed capacities across wards.

Table 1. Change in bed-days over capacity given an increase or decrease change in current bed capacity

Surgical Ward	Bed-Days Over Capacity		
	Current	Increase One Bed	Decrease One Bed
Ward 1	32.4	-8.4	+10.5
Ward 2	11.7	-4.1	+5.7
Ward 3	27.8	-7.3	+9.2

5.3.4 Setting Surgical Bed Protection Levels

Many hospitals use utilization rates to determine target ward capacities. However, this provides no indication on the accessibility of this ward to patients. Instead, it has been suggested that operating targets (Green, 2002) or specific access levels (Proudlove et al., 2006) be established for wards. We use BUS to explore these approaches. We define the access level as the probability that the demand for beds on a given day is less than or equal to the staffed ward capacity. Thus, the higher the access level, the less the chance of the ward being over capacity.

In the context of our case study we explore improving access for surgical patients in Ward 2. At the time of the study, Ward 2 was a shared ward with 31 beds, 20 of which were protected for surgical patients. Assuming that other patients can be properly managed, the number of protected beds for surgical patients can be determined using BUS. If a 75% access level for surgical patients is desired, the model suggests that 18 beds should be protected. If a 95% access level is desired, then 25 beds should be protected.

6 Conclusions

This paper describes our development and use of the Bed Utilization Simulator and the Surgical Schedule Optimizer, to support surgical scheduling and bed management. These tools are transparent to users and are easily adaptable to other hospital systems. Compared to other simulation models described in the literature, the benefits of BUS include simple logic, short development time, an intuitive user interface and that it MS Excel based. Despite its simplicity, BUS is capable of analyzing many important "what-if?" questions that continually challenge surgical managers. In addition, bed management concepts such as the impact of variability on ward accessibility can be captured and conveyed to managers and planners. Thus, BUS doubles as a valuable teaching tool.

The SSO model demonstrates that significant reduction in off-servicing and surgical cancellations can be achieved through enhanced SBSs. Results from using the slate model shows that surgical throughput can be increased while decreasing peak bed occupancies. However, SSO is challenging for non-technical users for many reasons so scheduling guidelines were developed to allow for incorporating the optimization results into practice.

The results of our study have influenced surgical operations at Royal Jubilee Hospital. The planners are using results from the project to support future surgical schedule revisions.

Surgical planners recognize that BUS can assist them in testing new schedules on a "what-if?" basis. The scheduling guidelines also challenged previous assumptions. Prior to the project, planners believed scheduling specialties evenly across the days of the week would smooth bed occupancy; our models clearly showed this was not the case. These guidelines are now being used on a daily basis to support ongoing scheduling decisions. A newly established Operations Research department has since been placed in charge of the surgical scheduling project and is evaluating the use of these tools at other facilities within the health authority.

Ongoing challenges exist in ensuring that models such as these are used in the future. While BUS requires relatively basic data for its operation, hospitals also need to have an IM/IT infrastructure that collects the appropriate information to support it. A substantial amount of time was required to clean and validate data prior to generating the BUS input database. Therefore to allow for the dissemination of BUS, ongoing work to integrate operations based data into current information systems is required to lay the foundation for current and future operations research based studies.

There are several fruitful areas of research arising from this work. Investigations into methods for improving the slate version of the SSO to increase the number of possible slates and reduce computation time are warranted. In addition, there is an opportunity to improve management of unplanned surgical patients. While little variation in mean bed occupancies of unplanned patients was observed across the week, large variations in bed occupancies exist within each day. Further analysis into possible options of managing unplanned surgeries by the division of patients into urgent and emergent should be considered.

Acknowledgements

We wish to thank Eileen Goudy, David McCoy, Andrea Boardman, Sheri Yager and all Vancouver Island Health Authority staff involved in this study for their help and continued support. Also, we want to acknowledge Yixin Chen and Mahesh Nagarajan for their earlier work on surgical scheduling. The results reported here were developed in the Centre for Operations Excellence in the University of British Columbia's Sauder School of Business.

References

Blake J., M. Carter, L. O'Brien-Pallas, L. McGillis-Hall. 1995. A surgical process management tool. Medinfo, 8(2), 1527-1531.

Blake J., J. Donald. 2002. Mount Sinai Hospital Uses Integer Programming to Allocate Operating Room Time. Interfaces, 32(2), 63-73.

Beliën J., E. Demeulemeester. 2007. Building cyclic master surgery schedules with levelled resulting bed occupancy. European Journal of Operational Research, 176(2), 1185-1204.

Beliën J., E. Demeulemeester, B. Cardoen. 2009. A decision support system for cyclic master surgery scheduling with multiple objectives. Journal of Scheduling, 12(2), 147-161.

Beliën J., E. Demeulemeester, B. Cardoen. 2006. Visualizing the demand for various resources as a function of the master surgery schedule: A case study. Journal of Medical Systems, 30(5), 343-350.

Carter M., J. Blake. 2005. Using Simulation in an Acute-care Hospital: Easier Said Than Done. Operations Research and Health Care: A Handbook of Methods and Applications, Ed. M. Brandeau, F. Sainfort, W. Pierskalla. Kluwer, New York, 191-215.

Cardoen B., E. Demeulemeester, J. Beliën. 2008. Operating room planning and scheduling: A literature review. Technical Report KBI 0807, Katholieke Universiteit Leuven, Belgium.

Denton B., J. Viapiano, A. Vogl. 2007. Optimization of surgery sequencing and scheduling decisions under uncertainty. Health Care Management Science, 10(1), 13-24.

Dexter F., R.D. Traub. 2002. How to Schedule Elective Surgical Cases into Specific Operating Rooms to Maximize the Efficiency of Use of Operating Room Time. Anesthesia and Analgesia, 94(4), 933-942.

Everett J.E. 2002. A Decision Support Simulation Model for the Management of an Elective Surgery Waiting System. Health Care Management Science, 5(2), 89-95.

Green L.V. 2002. How many hospital beds? Inquiry, 39(4), 400-412.

Harper P.R. 2002. A Framework for Operational Modelling of Hospital Resources. Health Care Management Science, 5(3), 165-173.

Harper P.R., A.K. Shahani. 2002. Modelling for the planning and management of bed capacities in hospitals. The Journal of the Operational Research Society, 53(1), 11-18.

Henderson S.G., Mason A.J. 2005. Ambulance Service Planning: Simulation and Data Visualization. Operations Research and Health Care: A Handbook of Methods and Applications, Ed. M. Brandeau, F. Sainfort, W. Pierskalla. Kluwer, New York, 77-102.

Isken M.W., B. Rajagopalan. 2002. Data Mining to Support Simulation Modeling of Patient Flow in Hospitals. Journal of Medical Systems, 26(2), 179-197.

Jacobson S., S. Hall, R. Swisher. 2006. Discrete-Event Simulation of Health Care Systems. Patient Flow: Reducing Delay in Healthcare Delivery. Ed. R.W. Hall. Springer, Los Angeles, 211-252.

Jun J.B., S.H. Jacobson, J.R. Swisher. 1999. Application of discrete-event simulation in health care clinics: A survey. The Journal of the Operational Research Society, 50(2), 109-123.

Persson M., J.A. Persson. 2009. Health economic modeling to support surgery management at a Swedish hospital. Omega, 37(4), 853-863.

Pitt M. 1997. A Generalised Simulation System to Support Strategic Resource Planning In Healthcare. Proceedings of the 1997 Winter Simulation Conference, 1155-1162.

Postl B.D. 2006. Final Report of the Federal Advisor on Wait Times. http://www.hc-sc.gc.ca/hcs-sss/pubs/system-regime/2006-wait-attente/index-eng.php. Accessed Nov 08.

Proudlove N.C., S. Black, A. Fletcher. 2006. OR and the challenge to improve the NHS: modelling for insight and improvement in in-patient flows. The Journal of the Operational Research Society, 58, 145-158.

Santibanez P., M. Begen, D. Atkins. 2007. Surgical block scheduling in a system of hospitals: an application to resource and wait list management in a British Columbia health authority. Health Care Management Science, 10(3), 269-282.

Testi A., E. Tanfani, and G. Torre. 2007. A three-phase approach for operating theatre schedules. Health Care Management Science, 10(2), 163-172.

Van Oostrum J.M., M. Van Houdenhoven, J.L. Hurink, E.W. Hans, G. Wullink, G. Kazemier. 2008. A master surgery scheduling approach for cyclic scheduling in operating room departments. OR Spectrum, 30(2), 355-374.

VanBerkel P.T., J.T Blake. 2007. A comprehensive simulation for wait time reduction and capacity planning applied to general surgery. Health Care Management Science, 10(4), 373-385.

Vissers J.M.H., I. Adan, J.A. Bekkers. 2005. Patient mix optimization in tactical cardiothoracic surgery planning: A case study. IMA Journal of Management Mathematics, 16(3), 281-304.

Appendix 1: Surgical Scheduling Optimizer Base Model

The binary decision variables $X_b^{i,w}$ are used to indicate whether a surgical block *b* is scheduled on a day *i* of week *w* in the schedule. Each surgical block represents a unique combination of a surgeon *d* and duration ($NumOR_b$). This duration is expressed in terms of OR-Days. A surgeon who requires one OR for half a day would require 0.5 OR-Days while a surgeon who requires two ORs for an entire day would require 2 OR-Days. Blocks corresponding to surgeon *d* are part of the set *B(d).* The model assumes that the assignment of blocks to ORs can be managed by surgical planners.

The $Bed_b^{p,u,j,i,w}$ parameter captures the expected bed demand (in nights) for a patient (of block *b* and type *p*) in each ward *u* and on each day *j*, given the block is scheduled on day *i* of week *w* in the cyclic surgical schedule. For example suppose that a surgical block *b* and patient type *p* combination always resulted in 50% of the patients staying for 3 nights in Ward 1, and the other 50% of the patients staying for 1 night in Ward 1 and 1 night in Ward 2 before discharge. Then $Bed_b^{p,u,j,i,w}$ would equal 1 in Ward 1 on day 1, 0.5 in Ward 1 and Ward 2 on day 2, and 0.5 in Ward 1 on day 3. In outpatient wards (i.e. Day Care Ward), $Bed_b^{p,u,j,i,w}$ is expressed as bed-days. A 6 hours stay post operation would be represented as 0.25 bed-days on day 1. To capture the expected occupancies of a block, this parameter is multiplied by the expected number of patients of each patient type in the block ($NumCases_b^p$).

The model includes several constraints. The first constraint (1) is an OR capacity constraint where the total OR-day requirements on any given day cannot exceed the total OR-days available that day ($ORperDay^{i,w}$). The second constraint (2) limits total number of OR-days allowed for each surgeon on each day ($ORperDaySurgeon_d^{i,w}$). The third constraint (3) limits the number of instances scheduled for each surgical block in each week ($WeekBlock_b^w$) to distribute blocks across the scheduling horizon. The fourth constraint (4) ensures that the total number of surgical blocks scheduled equals a predefined volume ($TotalBlock_b$). This is usually chosen to be consistent with the "as is" frequencies at the time of the study. An additional inequality (5) defines the maximum bed utilization in each ward MD_u. The objective of the model (6) is to minimize the sum of maximum bed occupancy over surgical wards.

Sets

b: Surgical blocks

i: Weekdays (1…5)

w: Weeks of the surgical schedule

j: Days of the surgical schedule ($1 \ldots 7 \cdot w$)

u: Wards

d: Surgeons

p: Patient types

$B(d)$: Blocks associated with surgeon d

Parameters

$Bed_b^{p,u,j,i,w}$: Expected number of bed-nights/days used by one patient of type p in ward u on day j due to surgical block b scheduled on day i of week w.

$NumOR_b$: OR-days required for each surgical block b

$ORperDay^{i,w}$: OR-days available on day i of week w

$ORperDaySurgeon_d^{i,w}$: OR-days available on day i of week w for surgeon d

$WeekBlock_b^w$: Number of Blocks b available in week w

$TotalBlock_b$: Total number of blocks b in the surgical schedule

$NumCases_b^p$: Number of cases for each patient type p in surgical block b

Decision Variables

$X_b^{i,w}$: $\left\{ \begin{array}{l} 1 \text{ if block } b \text{ is scheduled on day } i \text{ of week } w \\ \end{array} \right.$

0 otherwise

MD_u : Maximum number of beds in use in ward u over the scheduling period

Constraints

Daily OR-day capacity: $\sum_{b} X_b^{i,w} \cdot NumOR_b \leq ORperDay^{i,w} \qquad \forall i, w$ (1)

Daily OR-day capacity for each surgeon:

$$\sum_{b \in B(d)} X_b^{i,w} \cdot NumOR_b \leq ORperDaySurgeon_d^{i,w} \qquad \forall i, w, d \tag{2}$$

Weekly surgical block capacity: $\sum_{i} X_b^{i,w} \leq WeekBlock_b^{w} \qquad \forall b, w$ (3)

Surgical blocks balance: $\sum_{w} \sum_{i} X_b^{i,w} = TotalBlock_b \qquad \forall b$ (4)

Definition of the maximum bed utilization across the scheduling period in each ward:

$$\sum_{w} \sum_{i} \sum_{b} \sum_{p} X_b^{i,w} \cdot NumCases_b^{p} \cdot Bed_b^{p,u,j,i,w} \leq MD_u \qquad \forall j, u \tag{5}$$

Objective

Minimize the summation of the maximum bed occupancy in each surgical ward:

$$Min \sum_{u} MD_u \tag{6}$$

Appendix 2: Surgical Scheduling Optimizer Slate Selection Model

In this model, the decision variables are modified to include slate choice s. The parameter $NumCases_b^{p,s}$ replaces $NumCases_b^{p}$ which now stores the number of each patient type for each slate. Two additional constraints are also added. The first constraint (7) ensures only one slate is selected for each instance of a surgical block and the second constraint (8) ensure the number of cases performed for each patient type is greater than historical demand for each surgeon ($Totalcases_d^p$).

Sets

s : Model addition - Slate choices

Parameters

$TotalCases_d^{\,p}$: Model addition - Number of historical cases that needs to be met for each surgeon d and patient type p

$NumCases_b^{\,p,s}$: Replaces $NumCases_b^p$ - Number of cases for each patient type p in surgical block b with slate s

Decision Variables

$$X_b^{i,w,s} : \text{Replaces } X_b^{i,w} \quad \begin{cases} 1 & \text{if block } b \text{ is scheduled on day } i \text{ of week } w \text{ with slate } s \\ 0 & \text{otherwise} \end{cases}$$

Constraints

Model addition - Choose at most one slate: $\sum_s X_b^{\,i,w,s} \le 1 \qquad \forall i, w, b$ (7)

Model addition - Number of cases must be at least equal to historical volumes:

$$\sum_w \sum_i \sum_{b \in B(d)} \sum_s X_b^{\,i,w,s} \cdot NumCases_b^{\,p,s} \ge TotalCases_d^{\,p} \qquad \forall p, d \tag{8}$$

Replaces (1): $\sum_b \sum_s X_b^{\,i,w,s} \cdot NumOR_b \le ORperDay^{i,w} \qquad \forall i, w$ (9)

Replaces (2): $\sum_{b \in B(d)} \sum_s X_b^{\,i,w,s} \cdot NumOR_b \le ORperDaySurgeon_d^{\,i,w} \qquad \forall i, w, d$ (10)

Replaces (3): $\sum_i \sum_s X_b^{\,i,w,s} \le WeekBlock_b^{\,w} \qquad \forall b, w$ (11)

Replaces (4): $\sum_w \sum_i \sum_s X_b^{\,i,w,s} = TotalBlock_b \qquad \forall b$ (12)

Replaces (5): $\sum_{w}\sum_{i}\sum_{b}\sum_{p}\sum_{s} X_b^{i,w,s} \cdot NumCases_b^{p,s} \cdot Bed_b^{p,u,j,i,w} \le MD_u \qquad \forall j,u$ (13)

Objective

The objective function remains the same as (6).

Appendix 3: Surgical Scheduling Optimizer Additional RJH Constraints

A constraint on the number of OR-days available on each day of the schedule for Urology ($ORperDayUrol^{i,w}$) is applied in the RJH setting. The following additions are made to both the base model and the slate model.

Sets

y: Model addition - Urology specialty

B(y): Model addition - Blocks associated with Urology

Parameter

$ORperDayUrol^{i,w}$: Model addition - OR-day available on day *i* of week *w* for Urology

Constraints

Base model addition - OR-day capacity for Urology specialty:

$$\sum_{b \in B(y)} X_b^{i,w} \cdot NumOR_b \le ORperDayUrol^{i,w} \qquad \forall i, w \tag{14}$$

Slate model addition - OR-day capacity for Urology specialty:

$$\sum_{b \in B(y)} \sum_{s} X_b^{i,w,s} \cdot NumOR_b \le ORperDayUrol^{i,w} \qquad \forall i, w \tag{15}$$