An exact approach for relating recovering surgical patient workload to the master surgical schedule

PT Vanberkel[1,2]*, RJ Boucherie[1], EW Hans[1], JL Hurink[1], WAM van Lent[1,2] and WH van Harten[1,2]

[1]University of Twente, Enschede, the Netherlands; and [2]Netherlands Cancer Institute-Antoni van Leeuwenhoek Hospital, the Netherlands

No other department influences the workload of a hospital more than the Department of Surgery and in particular, the activities in the operating room. These activities are governed by the master surgical schedule (MSS), which states which patient types receive surgery on which day. In this paper, we describe an analytical approach to project the workload for downstream departments based on this MSS. Specifically, the ward occupancy distributions, patient admission/discharge distributions and the distributions for ongoing interventions/treatments are computed. Recovering after surgery requires the support of multiple departments, such as nursing, physiotherapy, rehabilitation and long-term care. With our model, managers from these departments can determine their workload by aggregating tasks associated with recovering surgical patients. The model, which supported the development of a new MSS at the Netherlands Cancer Institute–Antoni van Leeuwenhoek Hospital, provides the foundation for a decision support tool to relate downstream hospital departments to the operating room.

Introduction

Driven by an ageing population, public opinion, increased health expenditures and long waiting lists, a flood of changes in the health-care system have been set into motion. Many of these changes aim to improve efficiency and are planned by considering patient interactions within a single department. By limiting the scope of projects to a single department, the complexity and uncertainty which is inherent in health care, becomes more manageable for the problem solver. However, suboptimal conclusions may be drawn when only considering a single department, particularly if the influences of other services are ignored or if the impact of the change on the overall care chain is overlooked (Vanberkel et al, 2010).

No other single hospital department influences the workload of other departments more than the Department of Surgery, and in particular, the activities in the operating room. This influence depends directly on what types of patients receive surgery and when. Generally speaking, more invasive surgeries require more care during a patient's recovery. Giving consideration to this downstream effect of the operating room is essential for balancing the workload

of the hospital. The planning and scheduling of operating room time is discussed by many authors (Blake and Carter, 1997; Blake and Donald, 2002; Beliën and Demeulemeester, 2007; Van Houdenhoven et al, 2007; Wachtel and Dexter, 2008; Cardoen et al, 2010) and is often described as a multiple stage process.

The multiple stage process is used by many hospitals and starts with the long-term allocation of operating room time to the surgical specialties, for example the number of surgery hours per year. This allocation, referred to as Stage 1, is a strategic decision that reflects patient demand patterns and the priorities defined by hospital management. From this strategic decision a master surgical schedule (MSS) is developed for a shorter time horizon, which divides operating room time (aggregated into blocks) amongst the specialties; known as Stage 2. The specific assignment of patients to operating room blocks within the MSS is commonly referred to as Stage 3. A fourth stage 'addresses the monitoring and control of the operating room activities' (Van Oostrum et al, 2010) on the day of surgery. In this paper we focus on the development of a MSS in Stage 2.

The MSS is often specialty specific (Beliën and Demeulemeester, 2007), meaning operating room time is dedicated to a surgical specialty. In these MSSs, the decision of which patients (and consequently which

*Correspondence: PT Vanberkel, University of Twente, Captiool 15, Room E106, PO Box 217, Enschede, 7500 AE, the Netherlands.

surgeries) to schedule within each operating room block, is determined by the surgical specialty through consultation with the operating room manager. Other MSSs are more specific with operating room blocks being allotted to specific surgical procedures (Van Oostrum *et al*, 2008; Van Houdenhoven *et al*, 2008). Instead of using the term MSS, other authors refer to the distribution of operating room time among surgical specialties as a surgical block schedule (Santibáñez *et al*, 2007) and a timetable of operations (Harris, 1986).

The development of a MSS is often a complex balancing act. As the operating room is one of the hospital's most expensive resources, hospitals wish to maximize its performance through high resource utilization, minimal overtime, minimal cancellations and the elimination of conflicting equipment needs between rooms. Many authors describe methods for developing the MSS taking into account various resources within the operating room such as staff, equipment and instrument trays, for a review see (Cardoen *et al*, 2010). Furthermore, the operating room is often described as the engine that drives the hospital (Litvak and Long, 2000), implying many other departments are impacted by the MSS. The effect of the MSS on ward occupancy (Harris, 1986; Adan and Vissers, 2002; Beliën *et al*, 2006; Beliën and Demeulemeester, 2007; Vanberkel and Blake, 2007; Van Oostrum *et al*, 2008; Van Houdenhoven *et al*, 2008; Beliën *et al*, 2009), critical care resources (McManus *et al*, 2003; Griffiths *et al*, 2005; Beliën *et al*, 2006; Van Houdenhoven *et al*, 2008; Van Oostrum *et al*, 2008) and waiting lists (Santibáñez *et al*, 2007; Vanberkel and Blake, 2007) has notably been studied. Three of the mentioned papers emulate the relationship with deterministic models while the remaining consider at least one variable as stochastic. The stochastic models are either simulations used to evaluate 'what if' scenarios or mathematical programming (MP) models used to optimize certain objectives. In the following paragraph, we relate the model presented in this paper to these two techniques, thereby highlighting the contribution of our approach to MSS development.

Simulation models, which are well equipped to capture the broad scope of complex systems, are inexact and typically require a great deal of time to develop. As such, when analytical approaches can be used as an alternative, they have distinct advantages in terms of precision and development time. The analytical model presented in this paper has these advantages and for specific objectives can be used as an alternative to simulation. MP models on the other hand are capable of optimizing for certain objectives but require a more limited scope in order to be solvable. As such, many of the factors that may be included in a simulation model are ignored by the MP model. To broaden the scope of MP models, our model can be used to quickly evaluate proposed MSS solutions for additional factors. Such evaluations can be automated and a feedback loop can be created allowing the original MP model to account for additional factors.

Using the presented model, downstream workload distributions can be computed as a function of the MSS for all departments that provide care for recovering surgical patients. Specifically, the model computes the ward occupancy distributions, the patient admission/discharge distributions and the distributions for the ongoing interventions/treatments required by recovering patients. Furthermore the cumulative influence of multiple MSS cycles are considered. Since the MSS is identical from cycle to cycle, the overlapping of patients from one cycle to the next can be anticipated. A single MSS design is expected to remain in place for a long period of time leading to 'steady-state' workload distributions for each day of the MSS cycle. From a tactical point of view, this allows managers to see how balanced the expected workload is for a given MSS. Alternatively, the same model can support decisions at an operational level. Instead of computing the expected patients in recovery, the realized patients in recovery can be used as input. By aggregating this with the expected new arrivals from the operating room, real-time workload projections can be used to identify upcoming staffing shortages.

Such a model is valuable for department managers who want a means to relate their department's workload to the activities of the Department of Surgery. The motivation for this model comes from a real case at the Netherlands Cancer Institute–Antoni van Leeuwenhoek Hospital (NKI–AVL). The NKI–AVL planned to open an additional operating room and wanted to know the impact of this decision on the hospital as a whole, not only on the Department of Surgery. The model presented in this paper was applied at the NKI–AVL to support the development of their new MSS.

The paper is organized as follows. Initially, we describe the model and how to derive workload metrics. The next section describes the application and implication of the model results at the NKI–AVL. The final section, discusses the potential for the model as a decision support tool.

Model

This section describes a model to determine the workload placed on hospital departments by recovering surgical patients. In the same way as a MSS describes resource demands within the operating room, we show how the resources of other departments can be seen as a function of the MSS. The method relies only on data, which are easily extractable from typical patient management systems.

The model is most easily described from a queueing theory perspective. The kernel of the model is a single operating room block and its expected impact on the arrival rate to the hospital wards. The number of cases
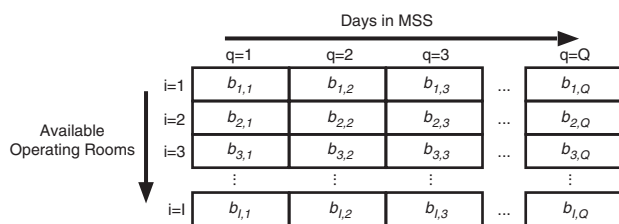
scheduled in such a block varies per specialty and is modelled as a specialty-specific random variable. This variable also represents the number of patients arriving to the ward (batch size). At the ward, each patient directly occupies a bed for a certain period of time. In the queueing model, the ward is seen as an infinite server system in which the patients occupy a server (ward bed) without delay. The time spent occupying a bed is the service time, which is modelled as a random variable of the Length of Stay (LOS). Again this random variable is specific to the surgical specialty. Since patients occupying a server do not interfere with each other during their recovery, the aggregate number of patients for all operating room blocks can be computed by adding the individual effects of all operating room blocks. Finally, since the MSS is cyclical, the cumulative number of patients from recurring MSS cycles can be computed.

The main output of the model is the distribution for the number of patients anticipated in the system on each day of the MSS. The model used for these calculations is explained in the following subsection. The three subsequent subsections explain how to modify this model to obtain the distributions for (1) ward occupancies, (2) admissions/discharges and (3) the number of patients in a specific day of their recovery. The time scale in the model is days; therefore all metrics are considered on a daily, not hourly, basis.

Model inputs

A MSS represents a repetitive pattern over a certain number of days (say Q). For each day $q \in \{1, 2, \ldots, Q\}$ in the MSS each of the I available operating rooms has to be assigned to one of the available surgical specialties. More precisely, the MSS is described by the assignment of a surgical specialty j to each operating room block $b_{i,q}$, where $i \in \{1, 2, \ldots, I\}$ and $q \in \{1, 2, \ldots, Q\}$. Using this notation, an empty MSS (ie before specialties have been assigned operating room blocks) is shown in Figure 1, where each cell represents an operating room block. It is common that multiple blocks are assigned to a single specialty on the same day.

The way specialty j fills in an operating room block is described by two parameters, c^j and d_n^j. c^j is a discrete distribution for the number of surgeries carried out in one block, that is $c^j(k)$ is the probability of k surgeries, $k \in \{0, 1, \ldots, C^j\}$, where C^j is the maximum number of surgeries of specialty j that can be completed in one block. Specialties independently decide which patients to operate on during each block, meaning that the number of surgeries completed in one block, does not influence the number of surgeries completed in another. The second parameter d_n^j is the probability that a patient, who is still in the ward on day n, is to be discharged that day ($n \in \{0, 1, \ldots, L^j\}$, where L^j is the maximum LOS for specialty j). A finite LOS is required for numerical purposes but also to ensure convergence to a steady-state result. Note that d_0^j is the probability that the patient is discharged on the same day as surgery (ie an outpatient surgery or day-case surgery) and $d_{L^j}^j = 1$. The parameter d_n^j is computed by dividing the probability that a patient's total stay is exactly n days by the probability that the patient was not yet discharged before day n. Let $P^j(n)$ be the probability that the LOS of a patient from specialty j is exactly n days long, then formally d_n^j is computed as follows:

$$d_n^j = \frac{P^j(n)}{\sum\limits_{k=n}^{L^j} P^j(k)}. \tag{1}$$

Recovering patients in the hospital

Using c^j and d_n^j as model inputs, for a given MSS the probability distribution for the number of recovering patients on each day q can be computed. The three steps used for the calculation are described in the following three subsections. Step 1 computes the distribution of recovering patients from a single operating room block of a specialty j, that is we essentially precalculate the distribution of recovering patients expected from an operating room block of a specialty. In Step 2, we consider a given MSS and use the result from Step 1 to compute the distribution of recovering patients given a single cycle of the MSS. Finally in Step 3, we incorporate recurring MSSs and compute the probability distribution of recovering patients on each day.

Step 1: *Distribution of recovering patients from specialty j following from a single operating room block*

In Step 1, we ignore the MSS and consider a single specialty j operating in a single operating room block. The patient flow process is as follows. During the operating room block patients receive surgery. The number of patients who undergo surgery in one operating room block is given by the random variable c^j. After surgery each patient still on the ward on day n has the probability d_n^j of being discharged that day. In the following we compute

Figure 1 Empty MSS illustrating paper notation.

the probability $h_n^j(x)$ that n days after carrying out a block of specialty j, x patients of the block are still in recovery. Note that $n \in \{0, 1, \ldots, L^j\}$ and $x \in \{0, 1, \ldots, C^j\}$ and that, for example, $h_3^j(5) = 0.25$ means that 3 days after surgery there is a 25% chance that five patients are still recovering in the hospital. For the probabilities h_n^j we have $\sum_{x=0}^{C^j} h_n^j(x) = 1$, $\forall n \in \{0, 1, \ldots, L_j\}$.

Day $n = 0$ is defined as the day of surgery and it is assumed that patients occupy a bed all day on the day of surgery even though they may physically be in the operating room. This is consistent with practice where patients have a recovery bed reserved for them before surgery. As such the number of patients in recovery from specialty j on day $n = 0$ is by definition, the number of surgeries performed that day by specialty j. It follows that the distribution for the number of recovering patients on day $n = 0$ is $h_0^j = c^j$.

Note that on day n, each patient still in the hospital has a probability d_n^j of being discharged that day and $(1 - d_n^j)$ of staying. If there are k patients in recovery on day n, then the probability of s patients in recovery (where $s \leqslant k$) on day $n+1$ is computed as $\binom{k}{s}(d_n^j)^{k-s}(1 - d_n^j)^s$. Since we know the probability distribution for the number of patients at the end of day $n = 0$, we can iteratively use this formula to compute the probability of k patients at the end of all days $n > 0$. Summarizing, the distribution for the number of recovering patients on day n is recursively computed by

$$h_n^j(x) = \begin{cases} c^j(x) & \text{when } n = 0 \\ \sum_{k=x}^{C^j} \binom{k}{x}(d_{n-1}^j)^{k-x}(1 - d_{n-1}^j)^x h_{n-1}^j(k) & \text{otherwise} \end{cases}$$

(2)

Step 2: Aggregate distribution of recovering patients following from a single MSS cycle

In this subsection, we consider the previously computed probability distribution h_n^j and a given MSS as input. Although the MSS is cyclical and repeats after Q days, in this subsection we consider only a single MSS in isolation. The MSS states when each specialty is assigned an operating room block which consequently defines the days when patients of specialty j arrive to the ward. Based on these, we compute the total number of patients in recovery by means of discrete convolutions.

To calculate the overall distribution of recovering patients, we first have to identify for each block $b_{i,q}$ the impact this block has on the number of recovering patients in the hospital on days $(q, q+1, \ldots)$. If j denotes the specialty assigned to block $b_{i,q}$, then the distribution $\bar{h}_m^{i,q}$ for the number of recovering patients of block $b_{i,q}$ on day m ($m \in \{1, 2, \ldots, Q, Q+1, Q+2, \ldots\}$) is given by

$$\bar{h}_m^{i,q} = \begin{cases} h_{m-q}^j & \text{if } q \leqslant m < L^j + q, \\ \mathbf{0} & \text{otherwise} \end{cases}$$

(3)

where $\mathbf{0}$ means $\bar{h}_m^{i,q}(0) = 1$ and all other probabilities $\bar{h}_m^{i,q}(l), l > 0$ are 0.

Let H_m be a discrete distribution for the total number of recovering patients on day m resulting from a single MSS cycle. Since recovering patients do not interfere with each other we can simply iteratively add the distributions of all the blocks corresponding to the day m to get H_m. Adding two independent discrete distributions is done by discrete convolutions which we indicated by $*$. Let A and B be two independent discrete distributions. Then $C = A * B$, is computed by
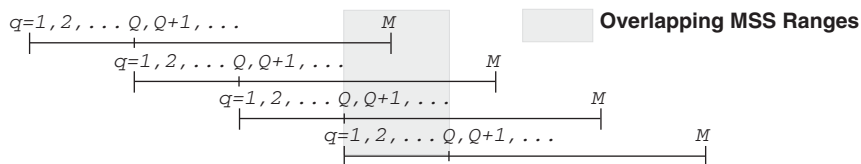
$$C(x) = \sum_{k=0}^{\tau} A(k)B(x-k)$$

where τ is equal to the largest x value with a positive probability that can result from $A * B$. Using this notation, H_m is computed by

$$H_m(x) = \bar{h}_m^{1,1} * \bar{h}_m^{1,2} * \ldots * \bar{h}_m^{1,Q} * \bar{h}_m^{2,1} * \ldots * \bar{h}_m^{I,Q}.$$

(4)

Step 3: Steady-state distribution of recovering patients

In Step 3, we consider a series of MSSs to compute the steady-state probability distribution of recovering patients. The cyclic structure of the MSS implies that patients receiving surgery during one cycle may overlap with patients from the next cycle. In the case of a small Q, patients from many different cycles may overlap.

In Step 2, we have computed H_m for a single MSS in isolation. Let M be the last day where there is still a positive probability that a recovering patient is present in H_m. This $M = \max_j\{L^j + x^j\}$ (where x^j is the latest day q of a block assigned to specialty j) indicates the range of the MSS. To calculate the overall distribution of recovering patients when the MSS is repeatedly executed we must take into account $\lceil M/Q \rceil$ consecutive MSS (see Figure 2). Let

Figure 2 Consecutive MSSs illustrating overlapping recovering patients.

Table 1 Model summary

Step 1   *Input*:
$c^j(x)$: probability distribution of specialty j completing x surgeries in one operating room block
d_n^j: is the probability that a patient, who is still in the ward on day n, is to be discharged that day
Output:
$h_n^j(x)$: probability distribution of x patients of a single operating room block of specialty j still in recovery n days after surgery

$$h_n^j(x) = \begin{cases} c^j(x) & \text{when } n = 0 \\ \sum_{k=x}^{C^j} \binom{k}{x} (d_{n-1}^j)^{k-x} (1 - d_{n-1}^j)^x h_{n-1}^j(k) & \text{otherwise.} \end{cases} \quad (2)$$

Step 2   *Input*:
$h_n^j(x)$: From Step 1
A Single MSS: Defines which day q specialty j operates in operating room I
Output:
$\bar{h}_m^{i,q}(x)$: probability distribution of x patients of the MSS still in recovery on day m.

$$\bar{h}_m^{i,q} = \begin{cases} h_{m-q}^j & \text{if } q \leqslant m < L^j + q, \\ \mathbf{0} & \text{otherwise} \quad \text{where } \mathbf{0}, \text{ means } \bar{h}_m^{i,q}(0) = 1 \end{cases} \quad (3)$$

$H_m(x)$: probability distribution of x patients still in recovery on day m (result for a single MSS in isolation). Let * indicate a convolution then,

$$H_m(x) = \bar{h}_m^{1,1} * \bar{h}_m^{1,2} * \ldots * \bar{h}_m^{1,Q} * \bar{h}_m^{2,1} * \ldots * \bar{h}_m^{I,Q} \quad (4)$$

Step 3   *Input*:
$H_m(x)$: From Step 2
Repeating MSSs
Output:
$H_q^{SS}(x)$: the steady-state probability distribution of recovering patients on day q.

$$Hq^{SS}(x) = H_q * H_{q+Q} * H_{q+2Q} \ldots * H_{q+yQ} \quad (5)$$

H_q^{SS} denote the probability distribution of recovering patients on day q of the MSS cycle, resulting from $\lceil M/Q \rceil$ consecutive MSSs. Since the MSS does not change from cycle to cycle, H_q^{SS} is the same for all MSS cycles. Such a result, where the probabilities of various states remain constant overtime, is referred to as a steady-state result. Using discrete convolutions, H_q^{SS} is computed by

$$H_q^{SS}(x) = H_q * H_{q+Q} * H_{q+2Q} * \ldots * H_{q+\lceil M/Q \rceil Q} \quad (5)$$

The relationship between the distribution H_q^{SS} and the workload associated with recovering patients is discussed in detail in the following three subsections. A summary of the three steps is provided in Table 1.

Ward occupancy

Perhaps the most common measure of inpatient workload is ward occupancy. Ward occupancy, the distribution of the number of inpatients on a ward, follows easily from the basic model where we compute the distribution of all

recovering patients. In practice, patients tend to be segregated into different wards depending on the type of surgery they received. To incorporate this segregation into the model and to consequently have recovering patient distributions for each ward, a minor modification needs to be made to the model. Let W_k be the set of specialties j whose patients are admitted to Ward k. Then in Step 2 we only have to consider those operating room blocks assigned to a specialty in W_k.

Admission rate/discharge rate

Bed occupancy alone does not fully account for the workload associated with caring for recovering patients. During patient admissions and discharges the nursing workload can increase. As such in this subsection we explain how to derive the probability distribution for daily admissions and discharges.

The admission rate is the rate of arriving patients which we previously defined as the number of surgeries completed

on day $n = 0$. For this metric we are only interested in a patient on the day of admission and wish to ignore them afterwards. To modify the model to reflect this new purpose replace (2) with

$$h_n^j(x) = \begin{cases} c^j(x) & \text{when } n = 0 \\ h_n^j(0) = 1 & \text{otherwise.} \end{cases} \quad (6)$$

With this modification, the resulting $H_q^{SS}(x)$ represents the distribution for daily admission for each day q of the MSS. To have ward specific results we again can restrict this to blocks belonging to specialties of the specific ward.

The discharge rate is the rate at which patients leave the ward and can be computed by adding an additional calculation in Step 1. The number of patients in recovery on day n is distributed according to $h_n^j(x)$, see (2). On day n each patient has the probability d_n^j of being discharged and the probability $(1-d_n^j)$ of staying. Let D_n^j be a discrete distribution for the number of discharges from specialty j on day n. Given $h_n^j(x)$ and d_n^j, $D_n^j(x)$ can be computed with a binomial distribution as follows:

$$D_n^j(x) = \sum_{k=x}^{C^j} \binom{k}{x} (d_n^j)^x (1 - d_n^j)^{k-x} h_n^j(k). \quad (7)$$

Finally, after computing $D_n^j(x)$, one can set $h_n^j(x) = D_n^j(x)$ and continue with Step 2. By doing so, the resulting $H_q^{SS}(x)$ represents the distribution for daily discharges for each day q of the MSS. As with admissions, ward specific results can be achieved.

Patients in day n of their recovery

The final workload metric we consider is the distribution of patients in day n of their recovery. This is relevant for predicting workload for the many hospital departments who treat recovering patients. For example, a patient recovering from hip surgery may need to see a physiotherapist every other day during their recovery. This metric states the frequency of such visits. The analogy holds true for all types of services that take place on specific intervals during a patient's recovery (eg chemotherapy, diagnostics, social work, discharge planning). In

particular, this can help dimension capacity for clinical pathways' patients whose recovery is intended to follow a strict regime.

This metric requires the most modification to the original model, as we now must carry an index (n) for the 'day of recovery' throughout the three steps. Let $\bar{h}_{m,n}^{i,q}$ be a discrete distribution for the number of recovering patients from block $b_{i,q}$ on day m in day n of their recovery. To compute $\bar{h}_{m,n}^{i,q}$ replace (3) with the following:

$$\bar{h}_{m,n}^{i,q} = \begin{cases} \mathbf{0} & \text{if } m - q \neq n \\ h_{m-q}^j & \text{otherwise,} \end{cases} \quad (8)$$

and we replace (4) with the following:

$$H_{m,n}(x) = \bar{h}_{m,n}^{1,1} * \bar{h}_{m,n}^{1,2} * \ldots * \bar{h}_{m,n}^{1,Q} * \bar{h}_{m,n}^{2,1} * \ldots * \bar{h}_{m,n}^{I,Q} \quad (9)$$

where $H_{m,n}(x)$ now denotes the number of patients from a single MSS on day m in day n of their recovery.

This alteration to the model eliminates the need for convolutions in Step 3. Since patients are indexed by their recovery day, patients from one MSS cycle are not aggregated with patients from the next. As such we need to replace (5) with

$$H_{q,n}^{SS}(x) = H_{q+Q\lfloor n/(Q+1) \rfloor, n}. \quad (10)$$

To help to interpret this metric, we consider the following fictitious example for patients who require a chemotherapy treatment on day 3 of their recovery. The Chemotherapy Department would like to know how frequently they need to provide this service. Example results for $H_{q,n}^{SS}(x)$ are illustrated in Table 2.

Assumptions

Inherent to the model are a number of assumptions, which are discussed in this subsection. One assumption resulting from the use of the infinite server system, is that there has to always be a bed available for a patient after surgery. This implies that surgeries are never cancelled due to bed shortages. In practice, this means that there is not a physical bed restriction and that additional staff can be called in when demand is higher than expected. The frequency of this occurring follows from the model. For example, if a hospital staffs 50 beds, then the probability

Table 2 Example results for the frequency of inpatient chemotherapy treatments

Example results	Parameters	Interpretation
$H_{1,3}^{SS}(2) = 0.3$	$n = 3, q = 1$	30% probability that exactly 2 treatments are required on the first day of the MSS cycle
$H_{1,3}^{SS}(3) = 0.5$	$n = 3, q = 1$	50% probability that exactly 3 treatments are required on the first day of the MSS cycle
$H_{1,3}^{SS}(4) = 0.2$	$n = 3, q = 1$	20% probability that exactly 4 treatments are required on the first day of the MSS cycle
$H_{2,3}^{SS}(2) = 0.4$	$n = 3, q = 2$	40% probability that exactly 2 treatments are required on the second day of the MSS cycle
$H_{2,3}^{SS}(3) = 0.4$	$n = 3, q = 2$	40% probability that exactly 3 treatments are required on the second day of the MSS cycle
$H_{2,3}^{SS}(4) = 0.2$	$n = 3, q = 2$	20% probability that exactly 4 treatments are required on the second day of the MSS cycle

of an additional staffed bed being needed on day q is H_q^{SS} (51).

In the current formulation the model ignores seasonality. Of course at certain times of the year surgical blocks are cancelled to accommodate vacations and slowdowns, representing a change in supply. In this case, a modified MSS is temporarily used breaking down the assumption that the same MSS repeats every Q days. However, given that the modifications to the MSS are simply the cancellation of certain operating room blocks, then the original result can act as an upper bound. Finally only elective surgeries are considered. To incorporate non-elective surgeries, it is possible to convolute an historic bed occupancy distribution for non-elective patients.

The model output reflects the variability in both c^j and d^j. Higher variability in these parameters leads to results with more possible states each with lower probabilities. As such the planning and scheduling becomes more difficult. In the NKI–AVL case study, for most operating room blocks the choice of which procedure to be carried out within the operating room block was left to the specialty. In hospitals where operating room blocks have fixed procedure assignments (Van Oostrum *et al*, 2008; Van Houdenhoven *et al*, 2008), one expects less variability in c^j and d^j and thus less variability in the model output.

Application

NKI–AVL is a comprehensive cancer centre providing hospital care and performance research, located in Amsterdam, the Netherlands. The hospital has 150 inpatient beds and sees about 24 000 new patients every year, making it approximately the size of a mid-sized general hospital. As with many Dutch hospitals, NKI–AVL is eager to improve access, and to this end, has expanded its surgical capacity by one operating room. The Department of Surgery was empowered to develop a new MSS encompassing the additional operating room, which fit within the capacity constraints of their department. The downstream departments, were asked to ramp up their staffing to accommodate the influx of extra patients. However, the extent to which they had to increase, and when they had to increase (in terms of which weekdays), was not known. In this section, we describe how the aforementioned model was implemented at NKI–AVL to address these concerns.

Implementation

The MSS at NKI–AVL repeats every two weeks and operating room blocks last for the complete day, which is typically eight hours. Each of the operating room blocks is assigned to either a surgical specialty, leaving the choice of procedure to the surgeon, or to a specific procedure. The total number of operating room blocks assigned to each specialty or procedure (Stage 1 of the planning process) has been defined by hospital management and does not exceed the total number of blocks available. As such the goal is to determine which of the operating room blocks during the 2-week cycle to assign to which specialty or procedure.

The data needed for the model were extracted from the hospital's patient management system for the calendar year 2008. The model calculations were facilitated with Microsoft Visual Basic.

The expansion of the operating room involved an increase in surgical capacity from five to six operating rooms. The former MSS for five operating rooms was completely replaced by a new MSS, which encompassed all six operating rooms. Initially, the Department of Surgery put forth a MSS proposal that accounted for physician schedules, available equipment, the Post Anaesthesia Recovery Unit and the Intensive Care Unit. However, in this proposal the impact of this MSS on the hospital's wards was not considered. Using the model described herein and working with managers from both the surgical department and the inpatient wards, this MSS was modified to balance the workload of recovering patients.

The modification of the initially proposed MSS to the final MSS, which was implemented in February of 2009, was completed over three months in an iterative manner. The team consisted of a team leader from the surgical department, a team leader from the inpatient wards, the manager of both groups and two authors who work in the hospital's operations improvement group. Using the model, the initially proposed MSS was deemed unacceptable due to high fluctuations in ward occupancy and concerns of bed shortages. Together a number of modifications to this MSS were proposed and scrutinized. The team members from the surgical department ensured these modifications did not cause conflicts within the operating room, such as with physician schedules and available equipment. After several weeks, a final version of the MSS was agreed upon which accounted for both the constraints of the surgery department and the needs of the inpatient wards.

The main benefit of using the model was to be able to quantify the concerns of ward staff, thereby providing a platform from which they could begin to negotiate a solution. Staff were quick to embrace the model output, particularly after seeing results from several MSS adaptations, at which point they were able to roughly predict the results. Once staff could foresee the impact of swapping one surgical block assignment with another, the final MSS came quickly. Although a local search heuristic may have found an acceptable (or optimal) solution more quickly, it would have required the many surgical department constraints to be modelled and would not have garnered

the same level of staff understanding and support as this more manual process. Example model outputs comparing the initially proposed MSS with the final MSS are provided in the following subsection.

Model output

The ward occupancy throughout the working week was the main metric used to evaluate each MSS proposal. At NKI–AVL there are two wards for treating surgical patients, Ward A and Ward B. Management strives to staff enough beds such that for 90% of the days there is sufficient coverage. In other words, they staff for the 90th percentile of demand. If demand is expected to exceed this staffing level, additional staff are brought in, making it rare that surgical procedures are cancelled due to a shortage in ward beds.
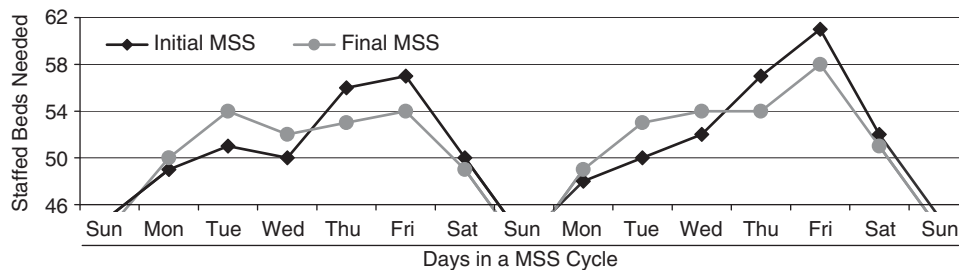
The 90th percentile of demand on Ward A is plotted in Figure 3 for both the initially proposed MSS and the final MSS. Formally, the 90th percentile is computed as follows. Let $H_q^{SS}(x)$ be the cumulative distribution function of $H_q^{SS}(x)$. Then the 90th percentile of demand for staffed ward beds is given by the minimum x such that $H_q^{SS}(x) \geqslant 0.9$.

As illustrated in Figure 3, the initially proposed MSS caused high fluctuations in Ward A occupancy throughout the working week. The occupancy was relatively low from Monday to Wednesday and relatively high on Thursday and Friday. Four of the 10 working days required more than 54-staffed beds with the maximum requirement being 61. Two of the 10 working days required less than 50-staffed beds and the remaining four required between 50 and 54. This resulting demand for staffed beds concerned the ward manager as such an unbalanced demand profile makes staff scheduling difficult. The ward staff were also concerned that late in the week they would be overworked and patient care could be compromised. Upon seeing these results the operating room staff were in agreement that this initially proposed MSS, although appropriate for the operating room, was not feasible for the wards. A number of modifications were made to this MSS and eventually an improved and final version was agreed upon.

The final version of the MSS dampened the ward occupancy fluctuation by lowering it on Thursday and Friday and increasing it from Monday to Wednesday. With this MSS only one working day required more than 54 staffed beds and the maximum requirement was reduced from 61 to 58. A further reduction of this peak was not possible due to physician schedule restrictions caused by their commitments elsewhere in the hospital. Eight of the remaining nine working days required between 50 and 54 staffed beds. One day required less than 54 staffed beds. By changing the assignment of specialties to blocks, thereby reordering the arrival of different patient types, the fluctuation in required staffed beds was greatly reduced. Similar results were achieved for Ward B but are excluded from the text for brevity.

In the remainder of this subsection, we want to illustrate how the model can also be used to compute admission/discharge rates and patients in recovery day n. The interpretation of these metrics is discussed and example model output is provided from the NKI–AVL case study. The admission and discharge rates are also useful for managing the wards and to illustrate this metric, the additional admissions and discharges expected from the sixth operating room, are plotted in Figure 4.

To illustrate that knowing the number of patients in day n of their recovery can be helpful for managing a ward, we consider patients from the Lung Cancer specialty whose surgical recovery is managed by a clinical pathway. For each day of the patient's recovery, starting on the day of surgery ($n = 0$) until the eighth day of their recovery ($n = 8$), the clinical pathway defines all the activities involved with caring for this patient type. For any additional recovery days the activities defined for day eight are repeated. In the final version of the MSS, Lung surgeries consume two surgical blocks, one on each Wednesday. Using the model, the expected number of patients in day n of their recovery is provided for each day of the MSS. Such information, as shown in Table 3, is helpful in predicting the workload for those departments involved with the Lung Cancer patient clinical pathway. For example, the clinical pathway states that formal patient education only occurs on days $n = 4$ and $n = 5$. From Table 3, the department responsible for educating Lung Patients can see that their services are required each week on Sunday and Monday. To get an

Figure 3 Ninetieth percentile of demand on Ward A for each day of the MSS.
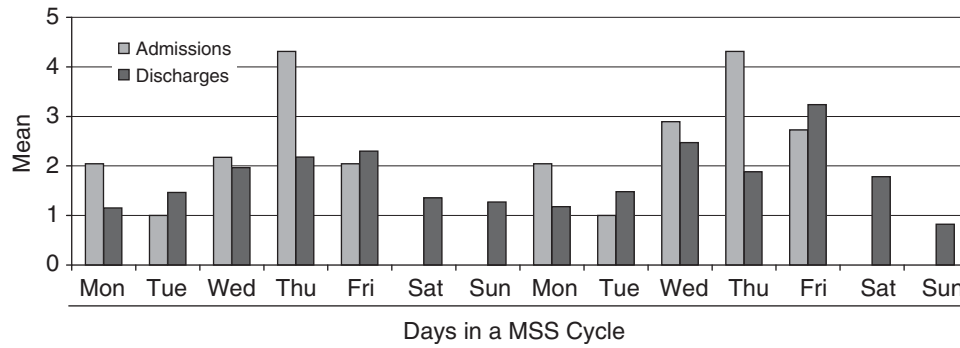
Figure 4 Additional admissions and discharges resulting from Operating Room 6 (both wards).

Table 3 Expected number of lung surgery patients in recovery day *n* for each day of the MSS

		Mon	*Tue*	*Wed*	*Thu*	*Fri*	*Sat*	*Sun*	*Mon*	*Tue*	*Wed*	*Thu*	*Fri*	*Sat*	*Sun*
Recovery day (*n*)	>7	0.6	0.5		1.2	1.0	0.9	0.7	0.6	0.5		1.2	1.0	0.9	0.7
	7			1.5							1.5				
	6		1.6							1.6					
	5	1.8							1.8						
	4							2.2							2.2
	3						2.4							2.4	
	2					2.5							2.5		
	1				2.8							2.8			
	0			3.0							3.0				

overall impression of the patient education workload, this metric could be computed for each operating room block. Alternatively, if this information was available during the development of the clinical pathway and if appropriate, efforts could be taken to limit the amount of educational activities that fall on Saturday and Sunday.

Discussion

In many hospitals, the surgical department is the foremost department in driving the hospital's workload. As such its activities (or lack of activities) cause a ripple effect elsewhere in the hospital. Upstream processes are less sensitive to changes as there is often a waiting list for surgical operations which acts as a buffer, dampening the effect. For downstream process this is different, as a buffer of post-surgery patients waiting to be admitted to a ward does not exist. Since post-surgical activities are quite sensitive to the activities in the operating room, it is important to derive one in terms of the other. As described in this paper, the workload for downstream departments can be described as a function of the MSS.

In the presented application, the model calculations were not completed by clinical staff but rather by the authors working in the operations improvement group. This was to accommodate the evolving nature of the model and the changing data requirements. However, with the formal model description it is easy to automate these computations into a decision support tool allowing staff to autonomously derive workload predictions of interest. Furthermore, the robust nature of the formulation allows the data to be aggregated in many ways, thereby making it possible to accommodate the requests of various departments.

The application of such a decision support tool is appropriate for both tactical and operational level decisions. At a tactical level, the probability of recurring downstream capacity conflicts can be computed. Using this information, a MSS which leads to the least likelihood of conflicts can be negotiated between the various stake holders. Of course, reducing the chance of problems occurring does not guarantee they will not occur. As such the same model can be applied at the hospital's operational level. By inputting the current state of inpatients alongside of the MSS, upcoming resources conflicts can be anticipated and efforts can be made to eliminate them beforehand.

References

Adan IJBF and Vissers JMH (2002). Patient mix optimisation in hospital admission planning: A case study. *Int J Opns Prod Mngt* **22**: 445–461.

Beliën J and Demeulemeester E (2007). Building cyclic master surgery schedules with leveled resulting bed occupancy. *Eur J Opl Res* **176**: 1185–1204.

Beliën J, Demeulemeester E and Cardoen B (2006). Visualizing the demand for various resources as a function of the master surgery schedule: A case study. *J Med Syst* **30**: 343–350.

Beliën J, Demeulemeester E and Cardoen B (2009). A decision support system for cyclic master surgery scheduling with multiple objectives. *J Sched* **12**: 147–161.

Blake JT and Carter MW (1997). Surgical process scheduling: A structured review. *J Soc Health Syst* **5**(3): 17–30.

Blake JT and Donald J (2002). Mount Sinai Hospital uses integer programming to allocate operating room time. *Interfaces* **32**(2): 63–73.

Cardoen B, Demeulemeester E and Beliën J (2010). Operating room planning and scheduling: A literature review. *Eur J Opl Res* **201**: 921–932.

Griffiths JD, Price-Lloyd N, Smithies M and Williams JE (2005). Modelling the requirement for supplementary nurses in an intensive care unit. *J Opl Res Soc* **56**: 126–133.

Harris RA (1986). Hospital bed requirements planning. *Eur J Opl Res* **25**: 121–126.

Litvak E and Long MC (2000). Cost and quality under managed care: Irreconcilable differences. *Am J Manag C* **6**: 305–312.

McManus ML, Long MC, Cooper A, Mandell J, Pagano DMBM and Litvak E (2003). Variability in surgical caseload and access to intensive care services. *Anesthesiology* **98**: 1491–1496.

Santibáñez P, Begen M and Atkins D (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 Mngt Sci* **10**: 269–282.

Van Houdenhoven M, van Oostrum JM, Hans EW, Gerhard W and Kazemier G (2007). Improving operating room efficiency by applying bin-packing and portfolio techniques to surgical case scheduling. *Anesth Analg* **105**: 707–714.

Van Houdenhoven M, van Oostrum JM, Wullink G, Hans E, Hurink JL, Bakker J and Kazemier G (2008). Fewer intensive care unit refusals and a higher capacity utilization by using a cyclic surgical case schedule. *J Crit Care* **23**: 222–226.

Van Oostrum JM, van Houdenhoven M, Hurink JL, Wullink EWHG and Kazemier G (2008). A master surgical scheduling approach for cyclic scheduling in operating room departments. *OR Spectrum* **30**: 355–374.

Van Oostrum JM, Bredenhoff E and Hans EW (2010). Suitability and managerial implications of a master surgical scheduling approach. *Ann Opns Res* **178**: 91–104.

Vanberkel PT and Blake JT (2007). A comprehensive simulation for wait time reduction and capacity planning applied in general surgery. *Health Care Mngt Sci* **10**: 373–385.

Vanberkel PT, Boucherie RJ, Hans EW and Hurink J (2010). A survey of health care models that encompass multiple departments. *Int J Health Mngt Inform* **1**: 37–69.

Wachtel RE and Dexter F (2008). Tactical increases in operating room block time for capacity planning should not be based on utilization. *Anesth Analg* **106**: 215–226.