OR Spectrum (2014) 36:585–605
DOI 10.1007/s00291-013-0323-x

REGULAR ARTICLE

Reducing the number of required beds by rearranging the OR-schedule

J. Theresia van Essen · Joël M. Bosch · Erwin W. Hans · Mark van Houdenhoven · Johann L. Hurink

Published online: 2 April 2013

Abstract After surgery most of the surgical patients have to be admitted in a ward in the hospital. Due to financial reasons and a decreasing number of available nurses in the Netherlands over the years, it is important to reduce the bed usage as much as possible. One possible way to achieve this is to create an operating room (OR) schedule that spreads the usage of beds nicely over time, and thereby minimizes the number of required beds. An OR-schedule is given by an assignment of OR-blocks to specific days in the planning horizon and has to fulfill several resource constraints. Due to the stochastic nature of the length of stay of patients, the analytic calculation of the number of required beds for a given OR-schedule is a complex task involving the convolution of discrete distributions. In this paper, two approaches to deal with this complexity are presented. First, a heuristic approach based on local search is given that takes into account the detailed formulation of the objective. A second approach reduces the complexity by simplifying the objective function. This allows modeling and solving the resulting problem as an ILP. Both approaches are tested on data provided by Hagaziekenhuis in the Netherlands. Furthermore, several what-if scenarios are evaluated. The computational results show that the approach that uses the simplified objective function provides better solutions to the original problem for instances based on the situation in HagaZiekenhuis. By using this approach, the number of required beds for the considered instance of HagaZiekenhuis can be reduced by almost 20 %.

J. T. van Essen (✉) · J. M. Bosch · E. W. Hans · J. L. Hurink
Center of Healthcare Operations Improvement and Research (CHOIR),
University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands
e-mail: j.t.vanessen@utwente.nl

J. T. van Essen · J. M. Bosch · M. van Houdenhoven
HagaZiekenhuis, P.O. Box 40551, 2504 LN Den Haag, The Netherlands

Keywords Operating room scheduling · Ward occupancy · Simulated annealing · Integer programming

1 Introduction

Due to an ageing population and increased health care costs, hospitals are forced to use their resources more efficiently, meaning that the same amount of patients has to be treated with less resources or more patients with the same amount of resources. One of the resources used in hospitals are the beds on the nursing wards. The cost for acquiring these beds is not substantial, however, the costs for maintaining and cleaning the beds, and the labour costs for treating the admitted patients are significantly high. Also, more and more costly technical appliances, such as interactive screens, are available at each bed. In addition, the number of available nurses in the Netherlands has been decreasing significantly over the years and will further decrease the coming years. Therefore, it is important to reduce the number of required beds as much as possible.

The starting point of this research was a request from HagaZiekenhuis, a hospital in the Netherlands, to get more insight in the factors that influence the bed occupancy. The operating room (OR) schedule is one of the most important factors that influence the bed occupancy, since most of the surgical patients have to be admitted at one of the wards after surgery. Therefore, it is important to consider the required bed capacity when creating the OR-schedule, which is the topic of this paper. First, we analyze this problem and investigate several approaches to solve it, and second, we show what improvements can be made in HagaZiekenhuis.

There is a vast amount of literature on OR planning and scheduling. Cardoen et al. (2010) and Hulshof et al. (2011) provide an overview of papers that address this topic. The problems discussed in these papers can be divided into two groups. The first group considers OR planning on the tactical level and the second group considers OR planning on the operational level. The tactical level is concerned with allocating available resources to groups of patients that share the same characteristics from a medical and logistics point of view in a time-horizon of a few weeks up to a few months. The operational level is concerned with planning and scheduling a given demand of elective patients within a time-horizon of a few days up to a few weeks while taking into account several uncertainties such as arriving emergency patients and stochastic treatment durations.

Several of the papers mentioned by Cardoen et al. (2010) and Hulshof et al. (2011) address the issue of considering the wards when creating an OR-schedule. The first paper that considers this topic is the work of Beliën and Demeulemeester (2007). They schedule blocks of elective surgeries of the same type by assigning them to a day in the planning horizon while minimizing the number of required beds. They assume that the length of stay (LOS) is given by a multinomial distribution that differs per surgery type. The number of required beds resulting from an OR-schedule is approximated in several ways, however, no exact formulation is used. Beliën et al. (2009) extend this approach by including multiple wards instead of one, allowing different block lengths and by scheduling individual surgeons instead of surgeon groups. In addition, they develop a decision support system that visualizes the OR-schedule and the resulting bed occupancy.

van Oostrum et al. (2008) schedule surgical procedures instead of OR-blocks by assigning the procedures to an OR and to a day in the planning horizon. The LOS of the patients is assumed to be deterministic and by using this deterministic LOS, the number of required beds is minimized. Adan et al. (2009, 2011) schedule surgical procedures by assigning them to a day in the planning horizon as is done by van Houdenhoven et al. (2008) and van Oostrum et al. (2008). But opposite to van Houdenhoven et al. (2008) and van Oostrum et al. (2008), they assume a fixed amount of beds is available at the hospital and minimize the over- and underutilization of these beds. Thus, the number of required beds is not minimized, but their use is optimized.

Chow et al. (2011) develop an integer linear programming model to generate improved OR-schedules in terms of the maximum expected bed occupancy. This expected bed occupancy is calculated by using the expected LOS of surgery types and after this, the average bed occupancy per day is determined by means of simulation.

As in Beliën and Demeulemeester (2007), Vanberkel et al. (2011a,b) schedule blocks of surgeries of the same type and assume a multinomial distribution of the LOS that differs per surgery type. However, in contrast to Beliën and Demeulemeester (2007), they analytically determine the complete probability distribution of the number of occupied beds for each day in the planning horizon. The goal of this approach is not to minimize the number of required beds, but to develop a model that can be used as an evaluation tool for the OR-schedule.

Bekker and Koeleman (2011) developed a method that determines the mean bed occupancy per day when a weekly admission pattern is given. In addition, they use a quadratic programming model to determine the optimal number of elective admissions per day such that an average desired bed occupancy per day is achieved.

For the operational planning level, (Cardoen et al. 2009a,b) propose a mixed integer linear programming approach and a column generation approach to determine the sequence in which patients must have surgery on a given day such that the peak use of recovery beds is minimized. They assume the LOS on the recovery to be deterministic and thus, no stochasticity is included. Fei et al. (2010) also focus on the sequence in which patients must have surgery, however, they consider the number of beds available at the recovery to be fixed, and therefore, do not focus on minimizing the peak use.

Many of the discussed papers consider the expected LOS of patients instead of the LOS probability distribution or focus on minimizing the maximum expected bed occupancy without considering the bed occupancy probability distribution. However, in practice, the LOS of patients is stochastic, and thus, it is important to also consider the variance in the bed occupancy. In this paper, we both incorporate the stochasticity of the LOS and of the bed occupancy to account for these variances. As in practice most hospitals use a cyclic OR-schedule, we develop an OR-schedule by assigning OR-blocks to a day in the planning horizon. We assume an OR-block consists of not only one but several surgical procedure types to make the problem more suitable for application. Because we schedule surgery types and not individual patients, this scheduling problem is considered to be on the tactical level. As in Beliën and Demeulemeester (2007) and Vanberkel et al. (2011b), we assume the LOS to be multinomially distributed. This distribution can easily be obtained from historical data. We use the analytical formulation of Vanberkel et al. (2011b) to determine the number of required beds, and minimize this number while taking into account several restrictions on the OR-schedule such as OR, surgeon, and instrument availability. As the problem originated in HagaZiekenhuis, we focus on resource constraints that are relevant in the setting of this hospital. However, it is possible to add additional constraints without destroying the structure of the developed model. Note that we only consider the scheduling of elective surgeries, but it is quite easy to also include emergency surgeries when determining the number of required beds.

The developed model is discussed in Sect. 2 and it consists of linear constraints and a complex non-linear objective function that involves the convolution of discrete distributions. To deal with this complexity, we introduce in Sect. 3 two different approaches to approximate the optimal solution. The first approach is a local search approach that takes the complex formulation of the objective function into account. We have chosen to use Simulated Annealing (SA) since this approach is easy to implement and has proven to be successful for other combinatorial optimization problems. The second approach reduces the complexity of the problem by linearizing the objective function. Although we prove the resulting problem to be NP-hard, we model and solve the resulting problem as an Integer Linear Program (ILP), because our considered instances are small enough to be solved within a reasonable amount of time. By comparing these two different approaches, we can determine whether it is better to not fully search the solution space with a complete evaluation of the objective function or to approximate the objective function and search the complete solution space. In fact, it is investigated if it is necessary to model the problem in full detail to be able to achieve a good solution.

The comparison is performed on data provided by HagaZiekenhuis. However, we use this data not just for comparing the two contrasting approaches, but we also aim to support HagaZiekenhuis by determining which resources are a bottleneck for minimizing the number of required beds. We do this by considering several what-if scenarios that relax some or all of the resource constraints. The computation results of the comparison and the what-if scenarios are given in Sect. 4. Section 5 presents conclusions and gives recommendations for further research.

2 Problem formulation

Hospitals aim to use as few beds as possible. When less beds are used, as a consequence less personnel is needed and less money is spent on cleaning and maintaining these beds. Another effect of using less beds is that also the bed occupancy during the week is better levelled and this reduces stress on the wards.

In hospitals, the number of beds occupied during the week is mostly determined by the OR-schedule. In general, a patient is admitted on the day of surgery, and after surgery, the patient must stay in the hospital for a few extra days. Thus, in order to influence the number of beds used, we should create an OR-schedule that minimizes the number of required beds, and thereby levels the amount of occupied beds as much as possible. HagaZiekenhuis, like many other hospitals, uses a cyclic OR-schedule that repeats every $\mathscr{T}$ days. This means that we have to develop such a cyclic OR-schedule for $\mathscr{T}$ days and not an OR-schedule for a whole year.

An OR-schedule consists of OR-blocks that are assigned to days of the planning cycle. Each OR-block is dedicated to a specific specialism or specialist and is filled with several surgery types chosen by this specialism or specialist. Thus, each specialism or specialist provides a list containing as many OR-blocks as this specialism or specialist gets during a period of $\mathscr{T}$ days. It only remains to assign these OR-blocks to a specific day in the planning horizon to create an OR-schedule.

2.1 Restrictions

In this section, we discuss several restrictions on the OR-schedule that are relevant for HagaZiekenhuis. Although we only provide these specific constraints, it is possible to add additional constraints without destroying the structure of the chosen approach.

Let K be the given set of OR-blocks. To each OR-block $k \in K$, we have to assign a specific day $t \in T = \{1, \ldots, \mathscr{T}\}$. For this, we define binary decision variables X_{kt} that are one when OR-block $k \in K$ is assigned to day $t \in T$, and zero otherwise. Then, the following constraints ensure that all OR-blocks $k \in K$ are assigned to a day in the OR-schedule:

$$\sum_{t \in T} X_{kt} = 1, \quad \forall k \in K. \tag{1}$$

The assignment of OR-blocks to days is limited by several constraints. First, some OR-blocks can only be performed in a subset of the available ORs, because, for example, special equipment is needed that is not available in all ORs. To model this, we define a set J of different OR types and for OR type $j \in J$, we denote by the subset $K_j \subseteq K$ the OR-blocks that can be performed in OR type $j \in J$. In addition, the number of available ORs of type $j \in J$ on day $t \in T$ is limited and denoted by a_{jt}. The following constraints ensure that the assignment of OR-blocks to days fulfils these limitations:

$$\sum_{k \in K_j} X_{kt} \leq a_{jt}, \quad \forall j \in J, t \in T. \tag{2}$$

Each OR-block is allocated to a specific surgeon type, because most surgeons in HagaZiekenhuis are specialized in a certain set of surgery types. The surgeon types are given by set S, and the OR-blocks that have to be performed by surgeon type $s \in S$ are given by subset $K_s \subseteq K$. The number of available surgeons of type $s \in S$ on day $t \in T$ is limited and denoted by b_{st}. The following constraints ensure that the assignment of OR-blocks to days fulfils these restrictions:

$$\sum_{k \in K_s} X_{kt} \leq b_{st}, \quad \forall s \in S, t \in T. \tag{3}$$

Each OR-block consists of several surgeries that must be performed consecutively. The total set of possible surgery types is defined by I, and the number of surgeries of a specific type $i \in I$ performed in OR-block $k \in K$ is given by o_{ik}. For each surgery type $i \in I$, a specific set of instruments is needed to perform the surgery. The set of all available instrument sets is given by set R, and w_{kr} denotes how many instrument sets $r \in R$ are needed for OR-block $k \in K$. Because a limited number of instrument sets is available and the instrument sets have to be sterilized after surgery, the number of surgeries that need instrument set $r \in R$ scheduled per day is limited by q_r. This is ensured by the following constraints:

$$\sum_{k \in K} X_{kt} w_{kr} \leq q_r, \quad \forall r \in R, \forall t \in T. \tag{4}$$

Note that the o_{ik} values are not used explicitly, but are covered in the w_{kr} values. However, we have introduced the values since they are needed in the next section.

2.2 Objective function

Constraints (1)–(4) are the restrictions on the decision variables X_{kt}, and therefore, describe the set Φ of feasible solutions. In this subsection, we specify the quality of a feasible solution $\phi \in \Phi$ given by the maximum number of beds needed during the entire planning horizon. To determine this number for a proposed OR-schedule, we have to determine the bed occupancy for each day. If we would specify the bed occupancy by a deterministic measure (e.g., maximum or expected number of used beds), we do not take the stochastic nature of the LOS into account. Using the expected bed occupancy per day results in canceling patients for surgery because quite often not enough beds are available to admit them after surgery. Using the maximum number of beds needed leads to a solution for which almost always too much beds are available. Therefore, we choose to calculate the complete bed occupancy probability distribution per day and afterwards take the p-percentile of these probability distributions to ensure that sufficient beds are available with p percent chance. Since these percentiles represent the number of beds needed on day $t \in T$ of the planning horizon, we obtain the number of beds needed in the wards by taking the maximum over all days.

For a given OR-schedule, the probability distribution of the bed occupancy can be obtained by using the LOS distribution of all surgery types scheduled in the OR-blocks. The LOS distribution of each surgery type $i \in I$ is given by a multinomial distribution that can be obtained from historical data. After the LOS distributions are obtained from the historical data, we can compute the probability distribution of the bed occupancy for each day as in Vanberkel et al. (2011b) by taking discrete convolutions of the LOS distributions. In the following paragraphs, we shortly explain this method. For a more detailed description of this method, we refer to Vanberkel et al. (2011b).

Before we explain the method into more detail, we want to note that the LOS of a patient in the historical data might be influenced by the bed occupancy on the wards. When the wards are more crowded, it is likely that a patient is discharged sooner than would have been the case when the ward was less crowded. However, because our approach is on the tactical level and not on the operational level, we do not include these dependencies.

The probability distribution of the LOS of surgery type $i \in I$ is given by values l_n^i, which denote the probability that the LOS of a surgery type $i \in I$ is exactly n days ($n \in \{1, \ldots, \mathscr{L}_i\}$), where $\mathscr{L}_i$ is the maximum LOS of surgery type $i \in I$. From this, we can determine the conditional probability that a patient who is still admitted on day n is discharged that day, which is denoted by d_n^i. Note that d_1^i denotes the probability that a patient is discharged on the day of surgery (i.e., an outpatient surgery) and $d_{\mathscr{L}_i}^i = 1$. The value of d_n^i is given by

$$d_n^i = \frac{l_n^i}{\sum_{m=n}^{\mathscr{L}_i} l_m^i}. \tag{5}$$

From these values, we can calculate the probability distribution $h_n^{ik}(x)$ that n days after carrying out OR-block $k \in K$, x patients of surgery type $i \in I$ are still in recovery. Recall that o_{ik} denotes the number of patients of type $i \in I$ assigned to OR-block $k \in K$. Therefore, these probabilities are computed recursively as follows:

For $n = 1$:

$$h_1^{ik}(x) = \begin{cases} 1 & \text{when } x = o_{ik}, \\ 0 & \text{otherwise}. \end{cases} \tag{6}$$

For $n > 1$:

$$h_n^{ik}(x) = \sum_{y=x}^{o_{ik}} \binom{y}{x} \left(d_{n-1}^i\right)^{y-x} \left(1 - d_{n-1}^i\right)^x h_{n-1}^{ik}(y). \tag{7}$$

Next, we take discrete convolutions of $h_n^{ik}(x)$ over all $i \in I$ to determine the bed occupancy caused by OR-block $k \in K$. This gives the probability $\tilde{h}_n^k(x)$ that n days after carrying out OR-block $k \in K$, x patients are still in recovery:

$$\tilde{h}_n^k(x) = h_n^{1k}(x) * h_n^{2k}(x) * \cdots * h_n^{\mathscr{I}k}(x). \tag{8}$$

Because we use a cyclic OR-schedule, which repeats every $\mathscr{T}$ days, patients who had surgery in one cycle may still be admitted in the next cycle. Therefore, we must take into account $\lceil \mathscr{N}_k / \mathscr{T} \rceil$ consecutive cycles, where $\mathscr{N}_k$ denotes the maximum LOS of the surgeries scheduled in OR-block $k \in K$, i.e, $\mathscr{N}_k = \max_{i \in I | o_{ik} \geq 1} \mathscr{L}_i$. In other words, $\mathscr{N}_k$ represents the range of one cycle of the OR-schedule. Now, by again using discrete convolutions, we can compute the probability distribution $H_t^k(x)$ of recovering patients on day $t \in T$ of the cycle induced by OR-block $k \in K$ as follows:

$$H_t^k(x) = \tilde{h}_t^k(x) * \tilde{h}_{t+\mathscr{T}}^k(x) * \tilde{h}_{t+2\mathscr{T}}^k(x) * \cdots * \tilde{h}_{t+\lceil \mathscr{N}_k/\mathscr{T} \rceil \mathscr{T}}^k(x). \tag{9}$$

The last step in calculating the probability distribution of the bed occupancy is to combine the probability distributions $H_t^k(x)$ for all OR-blocks. To do this, we first have to shift the distribution $H_t^k(x)$ such that the patients who have surgery in OR-block $k \in K$ are admitted on the day they have surgery, i.e., the day $t \in T$ for which $X_{kt} = 1$. The shifted probability distribution is denoted by $\bar{H}_t^k(x)$ and is defined as follows:

$$\bar{H}_t^k(x) = \begin{cases} H_{t-\hat{t}+1}^k(x) & \text{for } \hat{t} \text{ with } X_{k\hat{t}} = 1 \text{ and } \hat{t} \leq t, \\ H_{t-\hat{t}+\mathscr{T}+1} & \text{otherwise .} \end{cases} \tag{10}$$

By taking the discrete convolutions of $\bar{H}_t^k(x)$ over $k \in K$, we now determine the probability distribution of the bed occupancy for each day $t \in T$ denoted by H_t, which is computed by

$$H_t(x) = \bar{H}_t^1(x) * \bar{H}_t^2(x) * \cdots * \bar{H}_t^{\mathscr{K}}(x). \tag{11}$$

Thus, the number of required beds $\gamma(\phi)$ for a given solution $\phi \in \Phi$ is given by:

$$\gamma(\phi) = \max_{t \in T} \min \left\{ x \,\middle|\, \sum_{y=0}^{x} H_t(y) \geq \frac{p}{100} \right\}. \tag{12}$$

Note, that the calculations for (5)–(9) can be performed beforehand. Nevertheless, determining the objective function for a new OR-schedule still involves convoluting several probability distributions as shown by Eq. (11). Therefore, it is not straightforward to quantify or predict the effect in the objective function when changing an OR-schedule and it is hard to approximate the objective function. Moreover, calculating the objective function takes a lot of computation time. To reduce this computation time, we can either choose to not fully search the solution space, or to approximate the objective function. This leads to the following two solution approaches: (1) use a local search heuristic based on the given constraints and objective function and (2) approximate the objective function and incorporate this approximation in an ILP that includes the given constraints of the OR-schedule. For the second approach, the original objective value is determined afterwards to determine the number of beds needed in practice and to make a fair comparison between the two approaches. The comparison is used to determine whether it is better to not fully search the solution space with a complete evaluation of the objective function or to approximate the objective function and search the complete solution space. The two approaches are discussed in more detail in the following section.

3 Solution approaches

Because of the complex objective function, we cannot solve reasonable sized instances of the problem to optimality. Therefore, we have to make a choice between using a heuristic procedure to solve the original problem and using a global approach to solve a simplified version of the problem. Both approaches do not guarantee to find the optimal solution, therefore, we want to investigate which of these two methods leads to better solutions. The first approach is based on SA, which is a local search method.

The second approach is an ILP that uses an approximation of the objective function. In the following, these two approaches are discussed in more detail.

3.1 Local search approach: simulated annealing

As a first approach to solve our problem we have chosen SA (Kirkpatrick et al. 1983). SA is a local search procedure that in each step moves from the current solution, denoted by ϕ_c, to a randomly selected neighbor solution, denoted by ϕ_n. A solution is represented by the assignment of OR-blocks to a day in the planning horizon and is considered to be feasible when it satisfies constraints (1)–(4). As neighbor solutions, we consider all feasible solutions that can be obtained by swapping two OR-blocks that are assigned to two different days. We do not consider swapping two OR-blocks assigned to the same day, because this does not affect the objective value. If the randomly selected neighbor solution has a lower objective function value than the current solution, i.e., $\gamma(\phi_n) \leq \gamma(\phi_c)$, the neighbor solution is accepted as the new current solution. Otherwise, the neighbor solution is accepted with a probability that depends on the objective value of the current and neigbhor solution and on a temperature parameter. This temperature parameter, denoted by Γ, gradually decreases during the search process, and therefore, also the acceptance probability of a worse solution decreases. The allowance of moving to worse solutions makes it possible to escape from a (poor) local minimum. For each temperature value, we perform ω iterations that together form a Markov chain, because the next state only depends on the current state. Also, during the entire process of SA, we keep track of the best solution found so far. A more detailed description of this method is given by Kirkpatrick et al. (1983).

Summarizing, our implementation of SA is as follows, where $\bar{\phi}$ denotes the current best solution:

[Step 1.] Start with the initial solution ϕ_c given by the OR-schedule currently used at the hospital. Set $\bar{\phi} := \phi_c$ and determine the objective function $\gamma(\phi_c)$. Set the initial temperature, i.e., $\Gamma := \Gamma_s$, and a reduction factor α.

[Step 2.] Repeat ω times:

(a) Randomly select a neighbor solution ϕ_n of the current solution and determine $\gamma(\phi_n)$.
(b) If $\gamma(\phi_n) \leq \gamma(\phi_c)$, set $\phi_c := \phi_n$, and if $\gamma(\phi_n) \leq \gamma\left(\bar{\phi}\right)$, set $\bar{\phi} := \phi_n$. Otherwise, set $\phi_c := \phi_n$ with probability $e^{\frac{\gamma(\phi_c)-\gamma(\phi_n)}{\Gamma}}$.

[Step 3.] Set $\Gamma = \alpha\Gamma$. If $\Gamma < \Gamma_f$, the final temperature, then stop; else, go to Step 2.

We choose the initial temperature Γ_s such that an increase of the objective value at the beginning of the procedure is accepted with a relatively high probability. This is needed to easily escape from a local minimum. We observe that the maximum increase of the objective value equals the maximum over the number of surgeries assigned to an OR-block minus the minimum over the number of surgeries assigned to an OR-block, i.e., $\max_k \sum_i o_{ik} - \min_k \sum_i o_{ik}$, because all patients are admitted on the day of surgery. We want to accept this maximum increase at the start of the procedure with probability 0.95, thus the initial temperature is given by

$$\Gamma_s = \frac{\max_k \sum_i o_{ik} - \min_k \sum_i o_{ik}}{\ln 0.95} \tag{13}$$

Using the same approach, we determine the final temperature Γ_f. This temperature is chosen such that the probability of accepting the minimum increase of the objective value is very low. This means that at the end of the procedure almost no worse solution is accepted, and thus, the procedure converts to a local minimum. Since our objective function returns an integer amount of beds, the minimum increase is one bed. Thus, we set the threshold temperature Γ_f such that an increase of one bed is accepted with probability 0.001, i.e.,

$$\Gamma_f = \frac{-1}{\ln 0.001} \tag{14}$$

We set the number of iterations for each temperature value equal to the number of neighbour solutions that can be achieved by one swap of the current solution. This number is equal to the total amount of OR-blocks, i.e, ω equals the cardinality of set K, because in theory each OR-block can be swapped with one of the other OR-blocks. However, due to the restrictions described in Sect. 2.1, some swaps are prohibited. The reduction factor α is set to 0.95.

During preliminary runs, we tested the effect of varying the values for the initial temperature Γ_s, the final temperature Γ_f, the length of the Markov chain ω and the reduction factor α. We evaluated different values for the initial temperature by determining the performance of the SA approach when the acceptance probability of the maximum increase is set to 0.75 and 0.5. This means that we limit the increase in the objective function at the start of the SA approach. The final temperature is varied by increasing the acceptance probability of the minimum increase to 0.1, 0.01 and 0.005. When the computational time of the SA approach is increased, the approach might provide better results. Increasing the computational time can be achieved by increasing the length of the Markov chain or by increasing the reduction factor. For the Markov chain, we have evaluated the values 100 and 500 and for the reduction factor the value 0.97. To further investigate the effect of the reduction factor, we have also evaluated the performance of the SA approach when the reduction factor is set to 0.9. None of these mentioned values lead to a significant change of the performance of the SA approach. Only when all parameters are set such that the computing time is maximal, i.e., Γ_s and Γ_f are set at their original values, and ω and α are set to 500 and 0.97, respectively, the performance slightly improves. However, when using these parameter settings the computational time increases by a factor 20. Therefore, we use the values of the parameters as described above.

3.2 Global approach: linearization of objective function

The local search approach described in the previous subsection incorporates the complete evaluation of the objective function but only searches a part of the solution space. The global approach described in this subsection searches the entire solution space, but for such an approach the relation between a solution and the objective function must

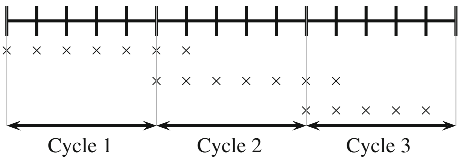

Fig. 1 Overlap of patients' multiple cycles

be evident. However, Sect. 2 shows that there is no straightforward and direct relation between a given OR-schedule and the resulting required number of beds. Therefore, we choose to linearize the objective function by replacing it with the maximum over the expected number of occupied beds per day. For calculating the expected number of occupied beds per day, we follow the approach of Beliën and Demeulemeester (2007). However, their formula does not work properly when the LOS of a patient exceeds the planning horizon. Figure 1 shows that when the maximum LOS $\mathscr{L}_i$ of a surgery type $i \in I$ exceeds the length $\mathscr{T}$ of the planning horizon, patients of two cycles of the OR-schedule may be admitted simultaneously, i.e., patients from different cycles may overlap. In Beliën and Demeulemeester (2007), it is assumed that this holds for all days in the planning horizon, however, this only holds for a few days as shown in Fig. 1, where a situation is sketched with $\mathscr{T} = 5$ and $\mathscr{L}_i = 7$. This deficiency can be accounted for by a small modification in the weight factor used in the formula defined in Beliën and Demeulemeester (2007). This modification ensures that patients are only counted multiple times on the days of the planning horizon that the LOS of several cycles overlap.

As a result, the expected number of occupied beds $\gamma_t(\phi)$ on day $t \in T$ of the planning horizon is given by:

$$\gamma_t(\phi) = \sum_{i\in I}\sum_{k\in K}\sum_{\tau\leq t}\left(\sum_{n=t-\tau+1}^{\mathscr{L}_i} l_n^i o_{ik}\left\lceil\frac{n-t+\tau}{\mathscr{T}}\right\rceil\right)X_{k\tau} + \sum_{i\in I}\sum_{k\in K}\sum_{\tau>t}\left(\sum_{n=\mathscr{T}+t-\tau+1}^{\mathscr{L}_i} l_n^i o_{ik}\left\lceil\frac{n-\mathscr{T}-t+\tau}{\mathscr{T}}\right\rceil\right)X_{k\tau} \quad (15)$$

Equation (15) determines for each day t in the planning horizon, the impact of all OR-blocks on the bed occupancy. Thus, for all OR-blocks it is determined whether patients operated on in this OR-block are still admitted in the hospital while taking into account overlapping cycles.

The first part of the equation only considers patients that are operated before or on the considered day $t \in T$ by integrating the binary variable $X_{k\tau}$, which is one when OR-block $k \in K$ is scheduled on day $\tau \in T$, and summing over all $\tau \leq t$. The expected number of patients of type i operated in block k and still admitted n days after surgery is given by $l_n^i o_{ik}$. As we only want to include the patients that are still admitted on day t, we only include l_n^i for $n \in \{t-\tau+1, \ldots, \mathscr{L}_i\}$. A patient should be counted only once if its LOS lies between $t-\tau+1$ and $\mathscr{T}-\tau+t$ and counted twice if its LOS lies between $\mathscr{T}+t-\tau+1$ and $2\mathscr{T}-\tau+t$, etc., which is represented

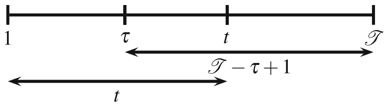

Fig. 2 Situation for first part of Eq. (15)

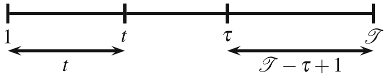

Fig. 3 Situation for second part of Eq. (15)

by $\left\lceil \frac{n-t+\tau}{\mathscr{T}} \right\rceil$. Figure 2 depicts the minimum LOS of a patient that should be counted twice when determining the expected number of occupied beds.

The second part of the equation only considers patients that are still admitted in one of the cycles after the cycle in which they had surgery. This means that we only include patients who have a LOS of $(\mathscr{T} - \tau + 1) + (t - 1 + 1) = \mathscr{T} - \tau + 1 + t$ or more days. This minimum LOS is depicted in Fig. 3. This means that when the patient's LOS is between $\mathscr{T} - \tau + 1 + t$ and $2\mathscr{T} - \tau + t$, the patient should only be counted once. When the LOS is between $2\mathscr{T} - \tau + t$ and $3\mathscr{T} - \tau - 1 + t$, the patient should be counted twice. This is represented by $\left\lceil \frac{n-\mathscr{T}-t+\tau}{\mathscr{T}} \right\rceil$.

Note that the expected value associated with the probability distribution of the bed occupancy as given in Sect. 2.2 corresponds to $\gamma_t(\phi)$. We can incorporate the linearized objective function given by Eq. (15) in an ILP that includes the constraints given in Sect. 2.1. Then, the resulting ILP is:

$$\begin{aligned} &\min_{\phi \in \Phi} \bar{\gamma}(\phi) \\ &\text{ s.t. } (1) - (4), (15) \\ &\bar{\gamma}(\phi) \geq \gamma_t(\phi), \forall t \in T \\ &X_{kt} \in \{0, 1\} \end{aligned} \tag{16}$$

This resulting problem is strongly NP-hard. For this, consider an instance with 3 ORs, a planning horizon of $\mathscr{T}$ days and thus $3\mathscr{T}$ OR-blocks. Each of the $3\mathscr{T}$ OR-blocks consists of a_k patients with $k = 1, \ldots, 3\mathscr{T}$ and $\sum_k a_k = \mathscr{T} b$, where b is a given positive integer. In addition, each of the patients has a LOS of exactly one day. For this instance, determining whether there exists an OR-schedule that requires b beds is equivalent to determining whether there are $\mathscr{T}$ pairwise disjoint subsets $R_l \subset \{1, \ldots, 3\mathscr{T}\}$ such that $\sum_{R_l} a_k = b$ for $l = 1, \ldots, \mathscr{T}$, which is known as the 3-partition problem (Garey and Johnson 1979).

The ILP given by (16) consists of $|K|\mathscr{T}$ binary variables and $(|J| + |S| + |R| + 2)\,\mathscr{T} + |K|$ constraints. After solving the ILP, each OR-block is assigned to a day in the planning horizon leading to a solution $\phi \in \Phi$. For this solution, the real objective value, i.e., the maximum over the p-percentiles of the resulting bed occupancy probability distribution, can be determined using the method described in Sect. 2.2.

The ILP only provides an optimal solution to the original problem when for each pair of solutions $\phi, \phi' \in \Phi$ the following holds: $\bar{\gamma}(\phi) \leq \bar{\gamma}(\phi') \Leftrightarrow \gamma(\phi) \leq \gamma(\phi')$, i.e., the ordering of the solutions in set Φ according to the expected number of beds should be the same as the ordering according to the number of beds needed for the

p-percentile. In general, the validity of this relation depends on the input data of the LOS distributions and cannot be guaranteed for all pairs of solutions. However, if it holds for most pairs, good solutions to the original problem may be obtained by solving the ILP.

4 Results

The purpose of this section is twofold. First, we compare the SA and ILP approach in Sect. 4.1 for 100 random generated instances based on data from HagaZiekenhuis. For the global approach, the original objective value for the resulting OR-schedule is determined afterwards such that a fair comparison can be made between the local and global approach. The results are used to determine whether it is better to not fully search the solution space with a complete evaluation of the objective function or to approximate the objective function and search the complete solution space. Second, we consider several what-if scenarios for HagaZiekenhuis with the solution approach that performed best in Sect. 4.1. We use these scenarios to determine whether the resource availability in HagaZiekenhuis limits the reduction of the number of required beds.

As mentioned in the introduction, the goal of the research is to give HagaZiekenhuis more insight in the factors that influence their bed occupancy. Therefore, the data used in the following subsections is based on data of HagaZiekenhuis. HagaZiekenhuis provided us with an OR-schedule of the orthopedics department with a planning horizon of 28 days, where up to three ORs and nine surgeons are available. The exact availability of the ORs and surgeons is given for each day in the planning horizon. The OR-schedule consists of 49 unique OR-blocks that have to be scheduled exactly once during the planning horizon. In total, 43 different surgery types are scheduled and the LOS per surgery type varies from 1 to 59 days with an average LOS of 3.7 days. For each surgery type, it is denoted which instrument sets are needed and for each of the ten available instrument sets it is given how many are available each day. As the number of required beds, we take the maximum of the 95-percentile of the probability distribution of the bed occupancy over the 28 days.

4.1 Comparing local and global approach

To determine which of the two considered approaches performs better, we have generated 100 random instances based on the data of HagaZiekenhuis. For the original data set, each OR-block had to be performed exactly once. To get different instances having similar characteristics as the original data, we vary the number of times a certain OR-block has to be performed during the planning horizon, i.e., some OR-blocks are not performed at all and some are performed multiple times in one cycle. To make sure there exists a solution that satisfies the fixed OR, surgeon and instrument sets availabilities, we generate the instances as follows: first, we randomly select for each available OR for each day in the planning horizon a surgeon available on that day. After this, we randomly select one of the OR-blocks that can be performed in the considered OR by the selected surgeon. During this selection process, we also consider the availability

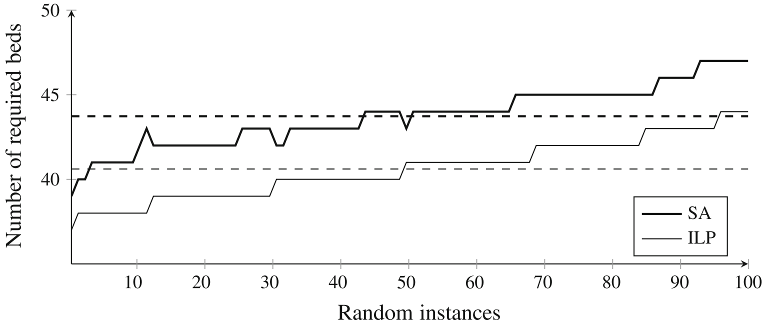

Fig. 4 Results of 100 random instances for 95-percentile

of instrument sets. Because the number of times an OR-block is performed varies for the generated instances, we create instances that vary among the number of surgeries and the average LOS of the patients. Therefore, we analyze a broad range of different instances. The number of surgeries varies between 201 and 236 and the average LOS over all surgeries scheduled in a random instance varies between 3.75 and 4.11 days.

The SA approach is implemented in CodeGear Delphi, and the ILP is solved with CPLEX 12.3. Both methods are executed on a Intel Core2 Duo CPU P8600 2.40 GHz with 3.45 GB RAM. Since proving the optimality of a solution by CPLEX takes quite some time, we interrupt the solver after 10 min. The maximum achieved integrality gap for the 100 random instances is 1.41 %. The results for the 100 instances can be found in Fig. 4 where the dashed lines denote the average of the objective values for the two approaches. Note that the random instances are sorted according to the objective function values to clarify the differences between the two approaches.

Figure 4 shows that the global approach performs better than the local search approach for all of the 100 random instances. In addition, we see that the difference between the global and local objective value is almost everywhere the same. The difference in the objective values is two beds for 8 % of the instances, three beds for 73 % of the instances, four beds for 18 % of the instances, and five beds for 1 % of the instances. Note that both objective values represent the maximum of the 95-percentile of the probability distribution of the number of required beds over all days.

Figure 5 shows one of the random instances for which the difference between the two approaches can be explained nicely. The peaks for the global approach are flat, which results in a constant number of occupied beds. This flat bed occupancy can be achieved because the OR schedule leaves enough room for improvement. The constraints do not restrict the solution that much that a flat bed occupancy cannot be achieved. The peaks of the local search approach fluctuate, which results in a higher number of required beds. Note that because of these fluctuations, zero patients are admitted on the second Sunday of the cycle. When the peaks in bed occupancy are decreased, the number of admitted patients on this day would likely increase.

The solution time needed for the local search approach varies between 32 and 74 s with an average of 42 s. The solution time needed for the global approach is set to 600 s. Therefore, as expected, the global approach takes longer than the local search approach, but 10 min is still a reasonable amount of time.

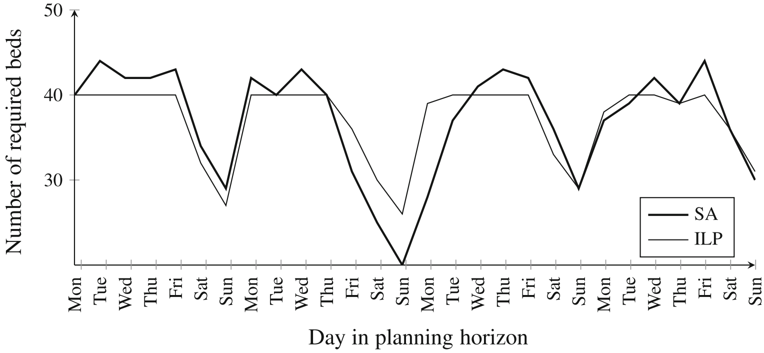

Fig. 5 Difference in levelled bed occupancy

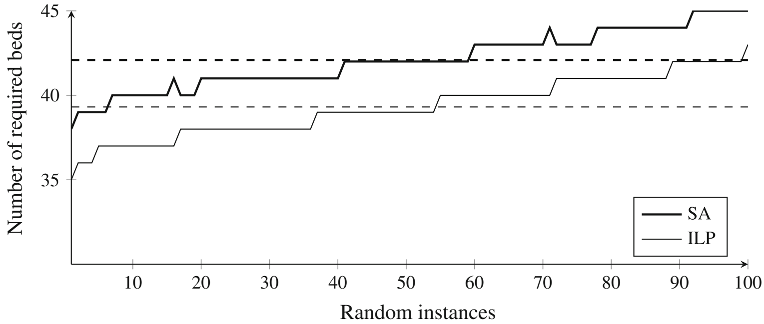

Fig. 6 Results of 100 random instances for 90-percentile

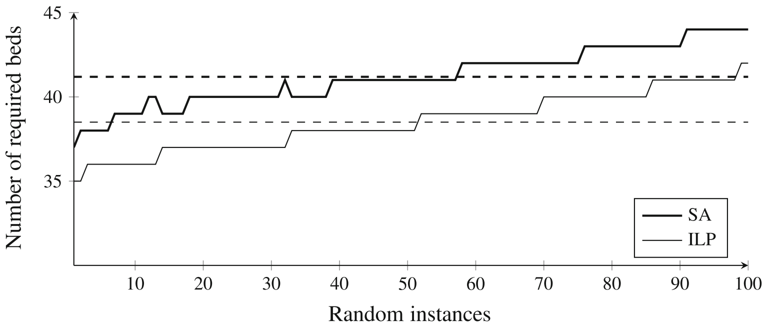

Fig. 7 Results of 100 random instances for 85-percentile

The ILP also outperforms the SA approach when we look at the 90-percentile and the 85-percentile. As the objective function of the ILP only depends on the expected number of beds needed, we do not have to solve the ILP again, but only have to determine the 90-percentile and 85-percentile for the solutions found by the ILP. The SA does consider the chosen percentile during the procedure, and therefore, we have to run SA again to determine new solutions for these percentiles. The results for the 100 random instances are depicted in Figs. 6 and 7.

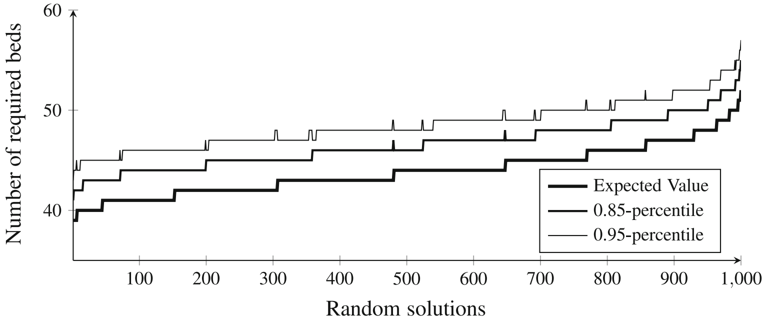

Fig. 8 Relation between expected value, 0.85-percentile and 0.95-percentile of 1,000 random solutions for instance based on data HagaZiekenhuis

The given results already indicate that for the considered instances based on the characteristics of the data from HagaZiekenhuis, it is not necessary to include the detailed objective function to determine a good OR-schedule as the ILP provides better results than the SA approach. As stated in Sect. 3.2, the ILP provides the optimal solution to the original problem when the ordering of the solutions in set Φ according to the expected number of beds is the same as the ordering according to the number of beds needed for the p-percentile. To investigate to which extend this holds for the considered instances, we have plotted the expected number of beds, the 85-percentile and the 95-percentile of beds needed in Fig. 8 for 1,000 random solutions to the original instance provided by HagaZiekenhuis. Note that the random instances are sorted in increasing order of the expected number of beds needed.

The wiggliness of the graphs of the 85-percentile and 95-percentile shows that the stated property does not hold for all 1,000 solution pairs. However, the small number of wiggles indicate that this property does hold for most pairs. This gives an explanation why the ILP provides good solutions to the original problem. Even though we cannot guarantee that the ILP finds the optimal solution, the ILP performs better than the SA approach when we compare both methods in the ability to reach a good feasible solution for the considered instances. Figure 8 shows that the error made by using the expected number of beds instead of the 95-percentile is at most two beds for the considered 1,000 random solutions. If we, for example, choose one of the solutions with expected value 39, the 95-percentile varies from 43 to 45. Thus, although two solutions may be considered to be equally good based on the expected value, one of the solutions might outperform the other when considering the 95-percentile. However, if this error is limited by two beds, the ILP still provides good solution to the original problem if only the expected values are used.

Although the ILP performs good for the considered instances, we cannot guarantee that this method also works on instances from other hospitals. For each considered setting, first the relation between the expected value and the chosen p-percentile should be investigated. When the property stated above holds for most solution pairs, the ILP can be used to find good solutions to the original problem. Else, it might be better to use the SA approach to solve the problem. For instances arising from practice, we believe that the ILP will outperform the SA approach as we expect that the LOS distributions will not differ much from the LOS distributions used in the instance provided by HagaZiekenhuis. Nevertheless, it would be interesting to determine what conditions would result in completely or at least regularly fulfilling the stated property. These conditions can then be used to determine beforehand which of the two proposed methods would be most suitable. However, determining these conditions is outside the scope of this paper.

4.2 What-if scenarios

The starting point of this research was the request from HagaZiekenhuis to get more insight in the factors that influence the bed occupancy. Therefore, we use the global approach to show the reduction in the number of required beds when the OR-schedule is changed, and we investigate whether the resource availability at HagaZiekenhuis limits this reduction. The hospital provided us an OR-schedule with a planning horizon of 28 days used by the orthopedics department. For the OR-schedule as provided by HagaZiekenhuis, 48 beds were needed to admit all surgical patients. We determined a new OR-schedule by solving the ILP and interrupting the solver after 10 min. To determine whether one or more of the constraints limit the improvement of the OR-schedule, we also consider the following scenarios:

- **Relax the number of available ORs per day**: The number of ORs of type $j \in J$ that are available each day is given by a_{jt}. By relaxing constraint (2), we do not restrict the model to schedule a fixed amount of OR-blocks per day. Since we relax the problem, we expect to come up with a schedule that requires less beds on the wards. We allow a maximum of 5 OR-blocks scheduled per day since 5 ORs are physically available at the operating department. Note that the number of available ORs in the weekend is still set to 0 as usually no surgeries are performed during this time.
- **Relax the surgeon availability**: As with the previous scenario, the surgeon availability corresponds to a constraint in the model. To determine what restriction this constraint imposes on the resulting OR-schedule, and thus on the required number of beds, we solve the model while relaxing constraint (3). Note that the surgeons are not available during the weekend as usually no surgeries are performed during this time.
- **Relax the instrument availability**: For each instrument set $r \in R$, q_r denotes the number of instrument sets available per day. By omitting constraint (4), we can determine the impact of this constraint on the number of required beds.
- **Relax all constraints**: By solving the model without all the above-mentioned constraints, we can determine the number of required beds when all resource capacities are unrestricted.
- **Relax all constraints including weekends**: Typically, no elective surgeries are performed during the weekends. However, it might be interesting to see which restriction this imposes on the objective function. Therefore, we also relaxed the availability of the ORs and surgeons during the weekends, i.e., the OR availability is set to 5 and all surgeons are available during the weekend.

Table 1 Results of scenarios

	# Expected beds	# 95-perc. beds	# Times peak achieved	Int. gap (%)
Original	43.2	48	1	–
Global approach	34.7	40	14	1.06
Relax OR availability	34.4	40	5	1.93
Relax surgeon availability	34.1	40	5	1.08
Relax instrument availability	34.7	40	15	1.16
Relax all constraints	33.9	40	2	1.68
Relax all constraints including weekends	32.7	38	12	2.14

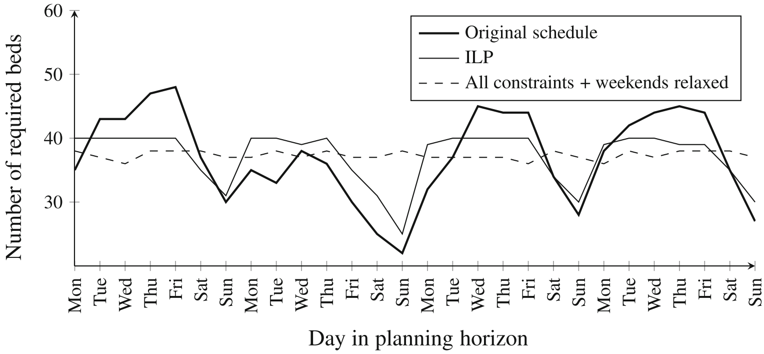

Fig. 9 Resulting bed occupancies of three scenarios

The results of the considered scenarios are given in Table 1. This table shows the expected number of required beds and the number of beds required when the 95-percentile is considered. In addition, we show the number of days in the planning horizon for which the maximum number of beds is achieved, which is denoted by '# Times Peak Achieved'. We also provide the integrality gap denoted by 'Int. Gap'.

Table 1 shows that the global approach reduces the number of required beds from 48 to 40 by reassigning the OR-blocks while taking into account all resource constraints. The results for the other scenarios show that the OR and surgeon availability during the week and the number of available instrument sets do not influence the required number of beds, because for all considered scenarios, except the last one, the number of beds needed equals 40. However, the expected number of required beds and the number of times the peak bed occupancy is achieved indicate a slight improvement for the scenarios where the OR availability, the surgeon availability or all resource constraints are relaxed. The scenario where the constraints are also relaxed during the weekends decreases the number of required beds from 40 to 38 beds. The resulting bed occupancies over the entire planning horizon for the original OR-schedule used in HagaZiekenhuis, the OR-schedule obtained by the global approach and the OR-schedule for the last mentioned scenario are given in Fig. 9.

The differences between the bed occupancy for the original OR-schedule and the one resulting from using the global approach might be explained by the differing number of surgeries scheduled per day. For the original OR-schedule, there is a peak in the number of surgeries scheduled per day at the start of the week and halfway through the week, while for the OR-schedule created by the global approach, there is only a peak at the start of the week. The global approach also schedules OR-blocks with a high average LOS at the end of the week. Note that the bed occupancy, shown in Fig. 9, for the global approach is rather flat during the week, however, during the weekends the bed occupancy is rather low. To flatten out these peaks, HagaZiekenhuis should consider to open the OR for elective surgeries during the weekends, because then, the number of beds needed can be reduced by two extra beds.

5 Conclusions

In this paper, we developed two approaches to improve the OR-schedule such that the number of required beds is reduced. The first approach incorporates the analytical formulation of the probability distribution of the bed occupancy and improves the OR-schedule by using a local search procedure. The second approach approximates the required number of beds by the expected bed occupancy, which enables us to solve the problem as an ILP. Both approaches are tested on 100 random instances to determine which of the two approaches provides the best solution to the original problem. The computational results show that the ILP with the simplified objective function performs the best for instances based on the situation in HagaZiekenhuis. Note that after solving the ILP, the number of required beds is still determined by using the analytic formulation. The computational results show that the number of required beds at the orthopedic department of HagaZiekenhuis can be reduced by almost 20 % when the ILP is used. None of the resources used at HagaZiekenhuis restrict the improvement that can be made to the OR-schedule; however, the number of required beds can be reduced slightly when the OR is also available for elective surgeries during the weekends.

Beliën and Demeulemeester (2007) considered a similar problem as discussed in this paper, however, they focused on minimizing the total expected bed shortage instead of minimizing the number of required beds. They compared an SA approach that considers the original objective function, and an ILP that considers an approximation of the objective function. The approximation used in the ILP is given by the minimization of the maximum expected bed occupancy which is quite different from the original objective function that indirectly focuses on minimizing the expected bed occupancy for all days in the planning horizon. Opposite to our results, Beliën and Demeulemeester (2007) conclude that the SA approach performs better than the ILP when the outcome of both approaches is compared based on the original objective function. This can be explained by the fact that the original and approximated objective function used by Beliën and Demeulemeester (2007) differ significantly, while in our case, both objective functions are quite similar. Therefore, we conclude that approximating the objective function only provides good solutions to the original problem when the approximated and original objective function leads to approximately the same ordering of feasible solutions. Therefore, when using the proposed solution approach in practice, it should be verified that most feasible solutions for the considered instance are ordered in the same way by the approximated and original objective function. Further research is needed to determine which conditions of an instance leads to entirely the same ordering of feasible solution.

The approach developed in this paper only considers elective surgeries, because only these surgeries can be scheduled in advance. However, patients who have to undergo surgery immediately, and as a consequence, their surgery cannot be scheduled beforehand, also have to be admitted at one of the wards after surgery. By using the model of Vanberkel et al. (2011b), we can incorporate these emergency surgeries by introducing dummy OR-blocks that are already fixed to a specific day in the planning horizon and contain the expected number of emergency surgeries. In this way, the arrival and admission of emergency patients is considered while determining a new OR-schedule for the elective surgeries, and thus, the total number of required beds is minimized and both elective and emergency patients can be admitted after surgery. However, this approach only considers the expected number of emergency patients and does not take into account the stochastic nature of the arrival process of emergency patients. Incorporating the stochastic arrival process of emergency surgeries would be an interesting topic for further research.

The developed approach can also be used to determine the admission schedule for non-surgical patients. To achieve this, we should schedule individual admissions instead of OR-blocks. This increases the complexity of the ILP as the number of variables increases. In addition, for the case of non-surgical patients, it is not defined how many admissions can be scheduled per day as this number may be unlimited. This also increases the complexity of the ILP due to the increasing solution space. Therefore, it might be needed to improve the solution approach to guarantee a reasonable computation time.

In the considered model, we assumed that the assignment of surgery types to OR-blocks is determined beforehand by the specialism of the surgeon. However, this assignment also influences the number of required beds on the wards. Therefore, it would be interesting to also incorporate this assignment when creating an OR-schedule such that the number of required beds can be reduced even further. Note that this also imposes some extra constraints on the model, because we also have to consider the stochastic duration of the surgeries such that the required surgical time does not exceed the available surgical time. Thus, it would be interesting to investigate this problem in future research.

Another interesting topic for future research is to take the available bed capacity at the wards into account when minimizing the number of required beds. For example, when the available bed capacity at the ward equals 40, it is not necessary to reduce the number of required beds further to 38. In addition, it might be beneficial to free as many wards as possible during the weekends to reduce the number of staff needed during the costly weekends.

Acknowledgments This research is supported by the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs. Furthermore, we would like to thank the anonymous referees for their useful comments which helped to improve the paper.

References

Adan I, Bekkers J, Dellaert N, Vissers J, Yu X (2009) Patient mix optimisation and stochastic resource requirements: a case study in cardiothoracic surgery planning. Health Care Manag Sci 12(2):129–141
Adan I, Bekkers J, Dellaert N, Jeunet J, Vissers J (2011) Improving operational effectiveness of tactical master plans for emergency and elective patients under stochastic demand and capacitated resources. Eur J Oper Res 213(1):290–308
Bekker R, Koeleman PM (2011) Scheduling admissions and reducing variability in bed demand. Health Care Manag Sci 14(3):237–249
Beliën J, Demeulemeester E (2007) Building cyclic master surgery schedules with leveled resulting bed occupancy. Eur J Oper Res 176(2):1185–1204
Beliën J, Demeulemeester E, Cardoen B (2009) A decision support system for cyclic master surgery scheduling with multiple objectives. J Sched 12(2):147–161
Cardoen B, Demeulemeester E, Beliën J (2009) Optimizing a multiple objective surgical case sequencing problem. Int J Prod Econ 119(2):354–366
Cardoen B, Demeulemeester E, Beliën J (2009) Sequencing surgical cases in a day-care environment: an exact branch-and-price approach. Comput Oper Res 36(9):2660–2669
Cardoen B, Demeulemeester E, Beliën J (2010) Operating room planning and scheduling: a literature review. Eur J Oper Res 201(3):921–932
Chow VS, Puterman ML, Salehirad N, Huang W, Atkins D (2011) Reducing surgical ward congestion through improved surgical scheduling and uncapacitated simulation. Prod Oper Manag 20(3):418–430
Fei H, Meskens N, Chu C (2010) A planning and scheduling problem for an operating theatre using an open scheduling strategy. Comput Ind Eng 58(2):221–230
Garey MR, Johnson DS (1979) Computers and intractability: a guide to the theory of NP-completeness. W. H. Freeman & Co Ltd, San Francisco
Hulshof PJH, Boucherie RJ, van Essen JT, Hans EW, Hurink JL, Kortbeek N, Litvak N, Vanberkel PT, van der Veen E, Veltman B, Vliegen IMH, Zonderland ME (2011) Orchestra: an online reference database of or/ms literature in health care. Health Care Manag Sci 14(4):383–384
Kirkpatrick S, Gelatt CD, Vecchi MP (1983) Optimization by simulated annealing. Science 220(4598): 671–680
van Houdenhoven M, van Oostrum JM, Wullink G, Hans E, Hurink JL, Bakker J, Kazemier G (2008) Fewer intensive care unit refusals and a higher capacity utilization by using a cyclic surgical case schedule. J Crit Care 23(2):222–226
van Oostrum JM, van Houdenhoven M, Hurink JL, Hans EW, Wullink G, Kazemier G (2008) A master surgical scheduling approach for cyclic scheduling in operating room departments. OR Spectr 30(2): 355–374
Vanberkel PT, Boucherie RJ, Hans EW, Hurink JL, van Lent WAM, van Harten WH (2011) Accounting for inpatient wards when developing master surgical schedules. Anesth Analg 112(6):1472–1479
Vanberkel PT, Boucherie RJ, Hans EW, Hurink JL, van Lent WAM, van Harten WH (2011) An exact approach for relating recovering surgical patient workload to the master surgical schedule. J Oper Res Soc 62(10):1851–1860