OR Spectrum (2008) 30:355–374
DOI 10.1007/s00291-006-0068-x

REGULAR ARTICLE

A master surgical scheduling approach for cyclic scheduling in operating room departments

Jeroen M. van Oostrum · M. Van Houdenhoven · J. L. Hurink · E. W. Hans · G. Wullink · G. Kazemier

Published online: 21 September 2006

Abstract This paper addresses the problem of operating room (OR) scheduling at the tactical level of hospital planning and control. Hospitals repetitively construct operating room schedules, which is a time-consuming, tedious, and complex task. The stochasticity of the durations of surgical procedures complicates the construction of operating room schedules. In addition, unbalanced scheduling of the operating room department often causes demand fluctuation in other departments such as surgical wards and intensive care units. We propose cyclic operating room schedules, so-called master surgical schedules (MSSs) to deal with this problem. In an MSS, frequently performed elective surgical procedure types are planned in a cyclic manner. To deal with the uncertain duration of procedures we use planned slack. The problem of constructing MSSs is modeled as a mathematical program containing probabilistic constraints. Since the resulting mathematical program is computationally intractable we propose a column generation approach that maximizes the operation room utilization and levels the requirements for subsequent hospital beds such as wards and

J. M. van Oostrum (✉) · M. Van Houdenhoven · G. Wullink · G. Kazemier
Department of operating rooms, Anesthesiology, and Intensive Care,
Erasmus MC, P.O. box 2040, 3000 CA Rotterdam, The Netherlands
e-mail: j.vanoostrum@erasmusmc.nl

G. Kazemier
Department of Surgery, Erasmus MC, P.O. box 2040, 3000 CA Rotterdam, The Netherlands

J. L. Hurink
Department of Electrical Engineering, Mathematics and Computer Science,
University of Twente, P.O. box 217, 7500 AE Enschede, The Netherlands

E. W. Hans
School of Business, Public Administration and Technology, University of Twente,
P.O. box 217, 7500 AE Enschede, The Netherlands

intensive care units in two subsequent phases. We tested the solution approach with data from the Erasmus Medical Center. Computational experiments show that the proposed solution approach works well for both the OR utilization and the leveling of requirements of subsequent hospital beds.

Keywords Scheduling · Master surgical schedules · Healthcare planning · Mathematical modeling

Mathematics Subject Classification (2000) 90B35

1 Introduction

Increasing costs of health care imply pressure on hospitals to make their organization more efficient. Recent studies show that operations research provides powerful techniques in this context (Carter 2002). One of the most expensive resources in a hospital is the operating room (OR)[1] department. Since up to 70% of all hospital admissions involve a stay in an OR department (OECD 2005), optimal utilization of OR capacity is of paramount importance.

Operating room utilization is typically jeopardized by numerous factors and various players are active in OR planning, such as individual surgeons, OR managers, and anesthesiologists (Weissman 2005). All players have autonomy, and can have conflicting objectives with respect to productivity, quality of care, and quality of labor (Glouberman and Mintzberg 2001). As a result, OR planning is constantly under scrutiny and pressure of potentially competing objectives.

A further complicating factor of the OR planning is the stochastic nature of the process. There are many uncertainties, such as stochastic durations of surgical procedures, no-shows of patients, personnel availability, and emergency surgical procedures. In addition, because surgeons tend to plan their procedures independently from others, this results in peak demands at subsequent hospital resources such as intensive care units (ICU). As a result, unavailability of for example ICU bed capacity can result in cancelation of surgical procedures (McManus et al. 2003).

In this paper we consider the problem of scheduling elective procedures, which is an operational planning problem that concerns the assignment of elective procedures to ORs over the days of the week. Due to the aforementioned difficulties, the planning process is complex, time consuming, and often under a lot of pressure. However, a lot of elective procedures tend to be identical during consecutive weeks in the year. In a regional hospital it is not uncommon that this is for more than 80% of the total volume the case (Bakker and Zuurbier 2002). In manufacturing as well as in health care, repetitive production is common practice. In such environments a cyclic planning approach is often used (e.g., Tayur 2000; Schmidt et al. 2001; Millar and Kiragu 1998). This reduces

[1] Note that "OR" in this context does not mean "operations research", but "operating room".

planning efforts considerably, and leads to reduced demand fluctuations within the supply chain, and higher utilization rates.

We propose in this paper a model for a cyclic scheduling approach of elective surgical procedures. We refer to such a cyclic surgical schedule as a *master surgical schedule* (MSS). An MSS specifies for each "OR-day" (i.e. operating room on a day) of the planning cycle a list of recurring surgical procedure types that must be performed. We demonstrate that our approach is generic: it not only allows to level and control the workload of the involved surgical specialties, but also from succeeding departments such as ICUs and surgical wards. It optimizes OR utilization without increasing overtime and cancelations. Furthermore, our approach accounts for the stochastic nature of the surgical process, such as stochastic durations of surgical procedures.

The approach for generation of MSSs was tested with data from the Erasmus Medical Center in Rotterdam, The Netherlands, which is a large university hospital. Approximately 15,000 patients annually undergo surgery in the OR departments of Erasmus MC. Since 1994, Erasmus MC has collected their surgical data in a database of 180,000 surgical procedures. The hospital actively supported the research project and affirms the applicability of this study.

The remainder of the paper is structured as follows. Section 2 presents an overview of studies related to the problem of construction MSSs. Section 3 presents a base model that represents the problem of constructing MSSs. Section 4 proposes a solution approach to solve the problem. In Sect. 5 we evaluate the solution approach. Section 6 draws conclusions from this research.

2 Related literature

There exist a strong interest in OR scheduling problems, resulting in a wide range of papers on this subject. These studies can be separated into short-term operating room scheduling (e.g., Gerhak et al. 1996; Sier et al. 1997; Ozkaraham 2000; Lamiri et al. 2005; Jebali et al. 2006) and mid-term planning and control (e.g., Guinet and Chaabane 2003; Ogulata and Erol 2003; Kim and Horowitz 2002). Studies about MSS are, however, scarce. Moreover, various definitions of a MSS are used. Blake and Donald (2002) construct MSSs that specify the number and type of operating rooms, the hours that ORs are available, and the specialty that has priority at an operating room. They use an integer programming formulation for the assignment of specialties to operating rooms. The objective function minimizes penalties related to the total under-supply of operating rooms to specialties. The authors implement a straightforward enumerative algorithm, which results in considerable improvements. Beliën and Demeulemeester (2005) use a nonlinear integer programming model to construct MSSs. The model assigns blocks of OR time to specialties in such a way, that the total expected bed shortage on the wards is minimized. After linearization of the model the authors examine and compare several heuristics to solve the resulting mixed integer program. They conclude that a simulated annealing approach yields the best results, but since this heuristic requires much computation time they propose a hybrid algorithm that combines simulated annealing with a quadratic programming model. This approach yields the best results concerning solution quality and computation times. Vissers et al. (2005) propose an MSS approach for a cardiothoracic department. At an aggregate level they form surgical procedure types and level resource requirements such as bed requirements. The objective of their approach is to minimize the deviation of target utilization rates for the OR, the ICU, and the wards. The approach focuses on capacity planning and does not account for the stochastic nature of health care processes.

The aforementioned authors propose various approaches for cyclic OR planning, some of them taking into account succeeding or preceding hospital departments. These approaches are designed for a higher level of aggregation than what we focus on. None actually constructs OR schedules in which actual surgical procedures or procedure types and their stochasticity are incorporated.

3 Problem description

The aim of this paper is to develop methods to generate MSSs, i.e., OR schedules that are cyclically executed in a given planning period. The cyclic nature of an MSS requires that not surgical procedures of concrete patients but surgical procedures of a certain type are scheduled. The concrete assignment of patients to the planned procedure types has to be done in a latter stage. To make such an approach applicable, the types of surgical procedures must represent surgical procedures, which are medically homogeneous in the sense that they share the same diagnosis and are performed by the same surgical department. In most hospitals there are three categories of types of procedures:

- Category A: elective procedures that occur quite frequent,
- Category B: elective procedures that occur rather seldom,
- Category C: emergency procedures.

Following the above discussion, an MSS can concern only Category A procedures. More precisely, we define Category A procedures as elective procedure types, which have a frequency such that they occur at least once during the cycle time of the MSS. The chosen cycle length thus determines the number of surgical procedure types incorporated in an MSS. Category B procedures consist of all other elective procedures and cannot be planned in an MSS, whereas Category C procedures cannot be planned due to their nature. However, in the construction of an MSS, capacity for the procedures of types B and C will be reserved.

An MSS is part of a cyclic OR planning strategy, which has three stages. First, clinicians and managers determine the MSS cycle length. Correspondingly, they determine how the OR capacity is divided over the three categories. Second, before each cycle, clinicians assign actual Category A patients to the procedure type "slots" in the MSS, and Category B procedures to their reserved capacity. Third, during execution of the elective schedule, Category C (emergency) procedures are scheduled. Widely used approaches are to assign these to reserved capacity (Goldratt 1997), or to capacity obtained by canceling elective procedures (Jebali et al. 2006).

In this paper we propose a model for the construction of MSSs for Category A procedures. Scheduling Category B and C procedures is beyond the scope of this paper. An MSS can be used repetitively by a hospital until the size and the content of the three categories change. Then, the MSS must be reoptimized.

The goal of our MSS is to generate a cyclic schedule, in which all Category A procedures are scheduled according to their expected frequency, in such a way that the workload of subsequent departments like wards and IC is leveled as much as possible. This leveling results in reduction of peak demands on hospital bed departments caused by elective surgical procedures and, as such positively influences resource shortages and minimizes the number of cancelation of surgical procedures McManus et al. 2003. The number of available ORs restricts constructing the MSS as well as the available operating time and the capacity of succeeding departments (i.e., number of available beds). Personnel restrictions are not taken into account. We assume that sufficient flexibility remains for personnel scheduling at the operational level when the scheduling of Category B procedures is done. To avoid the probability of overtime, planned slack is included in the construction of MSSs. The amount of slack depends on the accepted probability that overtime occurs, which is determined by the management, and the variance of procedure durations. We use the portfolio effect to minimize the total amount of required slack (Hans et al. 2006). The portfolio effect is the tendency for the risk of a well-diversified range of stochastic variables to fall below the risk of most and sometimes, all of its individual components. This principle can be applied with respect to the stochastic surgery durations. Exploiting the portfolio effect can thus reduce the required amount of slack.

3.1 Formal problem description

The surgical procedures to be incorporated into an MSS (Category A procedures) are categorized into I different types of medical and logistical similar procedures. From type i, $i = 1, \ldots, I$ we have s_i procedures to be added in the MSS. The duration of a surgical procedure of type i is a stochastic variable ξ_i, and based on Strum et al. (2000). we assume that ξ_i has a lognormal distribution. Let B be the number of different hospital bed types. The various hospital bed types differ in importance and to indicate the relative importance of hospital bed type b we introduce priority factor c_b. The duration of hospital bed requirements of type b for a procedure of type i is denoted by $l_{ib} \in \mathbb{N}, i = 1, \ldots, I; b = 1, \ldots, B$. We assume that only one patient per day can use a bed.

The MSS has a fixed duration, the cycle length T. This cycle length is measured in days and typically is a multiple of 7 days. The given surgical procedures have to be carried out in J identical ORs, where OR j on day t has a capacity of o_{jt}, $j = 1, \ldots, J; t = 1, \ldots, T$. For creating an MSS, procedures have to be assigned to the ORs. The total sum of the duration of procedures assigned on a single OR on a specific day may not exceed the available capacity with probability α, i.e., with probability α that no overtime occurs. We refer to OR j on day t as OR-day (j, t).

The combined objective of the problem is to construct MSSs such that both the required OR capacity is minimized and the hospital bed requirements are leveled over the cycle.

3.2 Base model

In this subsection we give a base model of the MSS problem. The aim of the model is to create a precise description of the objectives and the constraints.

To distinguish between minimization of OR capacity and hospital bed requirement leveling we define a weighted objective function, in which θ_1 is the weight of minimization of the required OR capacity and θ_2 is the weight of the hospital bed leveling. The weights may for example be related to the costs of the reduction of required OR capacity relative to the costs of peak demand on hospital beds.

We introduce an integer decision variable V_{ijt} to indicate the number of surgical procedures of type i that is assigned to OR-day (j, t), and an auxiliary binary variable W_{jt} to indicate whether an OR j is used on day t. An OR is considered to be used on day t if at least one surgical procedure is assigned to this OR-day. The total amount of OR capacity that is made available on day t is the sum of the available capacity of all used ORs. This is given by

$$\sum_{t=1}^{T} \sum_{j=1}^{J} o_{jt} \cdot W_{jt}.$$

To calculate the number of beds that is required from hospital bed type b, we introduce parameters $\psi_{t\tau ib}$ that denotes the requirements for hospital bed type b on day τ for a surgical procedure of type i, if this procedure is scheduled on day t. More specific, parameter $\psi_{t\tau ib}$ is $\left\lceil \frac{l_{ib}}{T} \right\rceil$ if $\min\{(t-1) \mod T, (t+l_{ib}-2) \mod T\} \leq (\tau-1) \leq \max\{(t-1) \mod T, (t+l_{ib}-2) \mod T\}$ and $\left\lfloor \frac{l_{ib}}{T} \right\rfloor$ otherwise. To illustrate this expression, suppose an MSS has cycle length $T = 7$ days. On day $t = 5$, a procedure of type i is scheduled that subsequently requires an IC bed for 8 days ($l_{ib} = 8$). This results in the requirement of two ICU beds on day $\tau = 5$ of the cycle and one IC bed on all other days. On day 5 the requirement is two beds, because the patient of the previous cycle is still occupying an ICU bed.

To level the hospital bed requirements, we minimize the maximum demand for hospital beds during an MSS cycle. This min–max type of resource leveling objective is generally used for problems where resource usage is very expensive (for this and other types, see: Brucker et al. 1999; Neumann and Zimmermann

2000). The presented approach is not specific for beds but can be used similarly for other types of hospital resources.

The maximum demand for hospital bed type b in a cycle is: $\max_{\tau \in T} \sum_{i=1}^{I} \sum_{j=1}^{J} \sum_{t=1}^{T} \psi_{t\tau ib} \cdot V_{ijt}$. To ensure that the objective function is not influenced by the total requirement of different hospital bed types, but only by their relative importance, we normalize the maximum demand for any hospital bed. The normalization factor is the total demand for an hospital bed type b during one cycle: $\left(\sum_{i=1}^{I} l_{ib} \cdot s_i\right) / T$. This yields the normative sum of the maximum demand of all hospital bed types:

$$\sum_{b=1}^{B} \left[\frac{c_b}{\left[\sum_{i=1}^{I} l_{ib} \cdot s_i\right] / T} \right] \cdot \max_{\tau \in T} \sum_{i=1}^{I} \sum_{j=1}^{J} \sum_{t=1}^{T} \psi_{t\tau ib} \cdot V_{ijt}$$

The overall objective function consisting of the weighted sum of needed OR capacity and the peak demands of hospital beds is given by formula (1) in the base model presented below.

To ensure that an operating room is considered to be used if at least one procedure is assigned to that operating room, constraints (2) are introduced. Constraints (3) ensure that all surgical procedures of all types are assigned. To model the bound on the probability that overtime occurs, we introduce a function $f_{jt}(V)$. It denotes the probability distribution of the total duration of all procedures that are scheduled on OR-day (j,t) by V, where V is the vector of all variables V_{ijt} (a possible way to deal with this function, is given in the following section). Using the function $f_{jt}(V)$, the restriction that the total duration of procedures on an OR-day may not exceed the available capacity with probability α, can be expressed by the probabilistic constraints (5). We refer to Charnes et al. (1964) for detailed information on probabilistic constraints. Summarizing, the base model becomes:

$$\min \theta_1 \cdot \sum_{t=1}^{T} \sum_{j=1}^{J} o_{jt} \cdot W_{jt} + \theta_2 \cdot \sum_{b=1}^{B} \left[\frac{c_b}{\left[\sum_{i=1}^{I} l_{ib} \cdot s_i\right] / T} \cdot \max_{\tau \in T} \sum_{i=1}^{I} \sum_{j=1}^{J} \sum_{t=1}^{T} \psi_{t\tau ib} \cdot V_{ijt} \right] \tag{1}$$

subject to

$$V_{ijt} \leq s_i \cdot W_{jt}, \quad i = 1, \ldots, I, \ j = 1, \ldots, J, \ t = 1, \ldots, T \tag{2}$$

$$\sum_{t=1}^{T} \sum_{j=1}^{J} V_{ijt} = s_i, \quad i = 1, \ldots, I \tag{3}$$

$$\begin{aligned} \Pr[f_{jt}(V) \le o_{jt}] \ge \alpha, &\quad j = 1,\dots,J,\ \ t = 1,\dots,T \qquad (4)\\ V_{ijt} \in \mathbb{N}, &\quad i = 1,\dots,I,\ \ j = 1,\dots,J,\ \ t = 1,\dots,T\\ W_{jt} \in \{0,1\}, &\quad i = 1,\dots,I,\ \ t = 1,\dots,T. \end{aligned}$$

The min–max objective can be reformulated (see Williams 1999, p. 23) such that the base model is an integer linear program (ILP) with additional probabilistic constraints. The size of instances from practice gets extremely large (the Erasmus MC instances approximately have 1.9×10^5decision variables), such that even without the probabilistic constraints this is far too large to solve the model to optimality within reasonable computation time. The MSS problem itself is NP-hard even if the probabilistic effects are neglected. The first part of the objective function together with the packing constraints contains e.g. the bin-packing problem and the second part of the objective function contains e.g. the three-partitioning problem. Based on this, we concentrate on a heuristic approach to solve the MSS problem.

4 Solution approach

The main decision in the MSS problem is to fill OR-days (j, t) according to the imposed restrictions. Since in practice the given capacities o_{jt} are often the same for different ORs and for different days, we introduce the concept of so-called *operating room day schedule* (ORDS). An ORDS for capacity o is a set of surgical procedures of various types, which is feasible with respect to the OR-capacity constraint (5) with $o_{jt} = o$. As a consequence, an ORDS for capacity o can be assigned to all OR-days (j, t) with $o_{jt} = o$. MSS comprises of assigning one ORDS to each OR-day $(j, t,)$ in the cycle, such that the objective function (1) is minimized.

We propose a two-phase decomposition approach. In Phase 1 hospital bed requirement leveling is ignored, and a set of ORDSs that covers all procedures is selected. These ORDSs have capacities fitting to the capacities of the OR-days, and minimize the required OR capacity. We discretize the probabilistic OR capacity constraints, and formulate an ILP that we solve with an implicit column generation approach. In Phase 2 we assign ORDSs to concrete OR-days in such a way, that the hospital bed capacity demand is leveled. For this purpose, the problem is formulated as mixed integer linear program (MILP).

4.1 Phase 1

The problem in Phase 1 consists of selecting a set of ORDSs that covers all surgical procedures and all OR-day capacities and minimizes the required OR capacity. In Sect. 4.1.1 we formalize the problem as an ILP problem where the variables correspond to ORDSs of given capacities. Afterwards, in Sect. 4.1.2 we propose a column generation approach to generate possible ORDSs. In this part we discretize the probabilistic constraints on the ORDSs.

4.1.1 Phase 1 model

The available capacity of ORs in the MSS cycle may differ from day to day. Let R be the number of different OR capacity sizes (sorted in non-decreasing order). The actual capacity of an OR of capacity size type r is given by $d_r,\ r = 1, \ldots, R$. Let U be the set of possible ORDSs, and let U_r be the subset of U that contains all the ORDSs that belong to the rth capacity size. In this context an ORDS u belongs to U_r if the rth capacity size is the smallest available capacity size where the ORDS fits in. Hence, $U = \cup_{r=1}^{R} U_r$. Let m_r be the number of OR-days within one cycle length that have the rth capacity size and let φ_r be the set of corresponding tuples (j, t). For a given ORDS $u \in U$ we denote the number of surgical procedures of type i that are scheduled in u by $a_{iu} \in \mathbb{N}$.

To formulate the Phase 1 model, we introduce integer decision variables X_u ($u \in U$) that represent the number of times that ORDS u is selected. The objective function (5) corresponds to the first part of the objective function (1) of the base model: minimization of the required OR capacity. Constraints (6) impose that all procedures are selected. The number of ORDSs generated for every OR capacity size that we can select is restricted by the number of available OR-days m_r of capacity type r. This restriction is imposed by constraints (7). Summarizing, in Phase 1 we must solve the following ILP:

$$\min \sum_{r=1}^{R} \sum_{u \in U_r} d_r \cdot X_u \tag{5}$$

subject to

$$\sum_{r=1}^{R} \sum_{u \in U_r} a_{iu} \cdot X_u \geq s_i \qquad i = 1, \ldots, I \tag{6}$$

$$\sum_{u \in U_r} X_u \leq m_r \qquad r = 1, \ldots, R \tag{7}$$

$$X_u \in \mathbb{N} \qquad u \in U.$$

This model has two main drawbacks. The set of possible ORDSs U grows exponentially with the number of procedure types, and due to the probabilistic constraints, the identification of all possible elements of U is difficult. To overcome this, a column generation approach for this problem is presented where furthermore the check on containment of an ORDS in a set U_r is discretized.

4.1.2 Column generation

Column generation is an often-used approach to solve complex optimization problems with a large number of variables (e.g. cutting stock, capacity planning, and crew scheduling, e.g., Barnhart et al. 1998; Pinedo 2005). The outline of our approach is as follows. We use column generation to solve the LP relaxation of the Phase 1 model, and round this solution to obtain a feasible solution. In the column generation procedure we iteratively generate subsets of U (i.e., subsets of ORDSs) and solve the Phase 1 model for these subsets. The Phase 1 model restricted to such a subset of U is called the restricted master problem. In each iteration, solving the restricted LP-relaxation (i.e. the LP-relaxation of the restricted master problem) yields shadow prices. These are used as input for the sub-problem (the pricing problem), which revolves around generating ORDSs that are not included in the restricted master problem, but that may improve its solution. The reduced costs of the corresponding variables X_u are negative. These ORDSs are added to the restricted master problem, and the LP-relaxation is re-optimized. This procedure stops if no ORDSs exist that may improve the restricted LP-relaxation solution. The restricted LP-relaxation solution is then optimal to the LP-relaxation. We then apply a rounding procedure to obtain a feasible Phase 1 solution.

Initialization We use an initialization heuristic to generate subsets of U_r for all OR capacity sizes $r = 1, \ldots, R$. More precisely, for each $r = 1, \ldots, R$ we generate subsets $\bar{U}_r \subset U$ of ORDSs that cover all surgical procedures. This initial set of ORDSs serves as a starting point for the column generation procedure.

Let the variable $Z_i^r \in \mathbb{N}$, $(i = 1, \ldots, I,)$ denote the number of procedures of type i that is scheduled in an ORDS for OR capacity size r. Any vector $Z^r = (Z_1^r, \ldots, Z_I^r)$ must satisfy the probabilistic bin-packing constraint (8) to be a feasible ORDS for capacity size r, where $f(Z^r)$ denotes the distribution function that represents the stochastic sum of the duration of all surgical procedures in the ORDS.

$$\Pr[f(Z^r) \leq d_r] \geq \alpha \tag{8}$$

The probabilistic constraints (8) impose difficulties on the generation of ORDSs. We discretize constraints (8) using prediction bounds. A prediction bound n_i^α denotes that the duration ξ_i of procedure type i is smaller than or equal to n_i^α with a probability α. These prediction bounds are used to replace the stochastic variables ξ_i, and can be calculated using the primitive of the distribution function of ξ_i. The total required OR capacity for an ORDS given by the vector Z^r is given by $\sum_{i=1}^{I} n_i^\alpha \cdot Z_i^r$. The difference between the value of a prediction bound and the average surgical procedure duration is used to compute the planned slack.

As discussed by Hans et al. (2006) the total amount of planned slack for a multiple of surgical procedures is reduced by the portfolio effect. This portfolio effect may be approximated by a function g, which only depends on the number of procedures that are scheduled in the operating room and on the average standard deviation of all types of surgical procedures. The reduction of required planned slack $g(\sum_{i=1}^{I} Z_i^r)$, as a result of the portfolio effect, is subtracted from the sum of the prediction bounds. This results in the following OR capacity constraints:

$$\left(\sum_{i=1}^{I} n_i^{\alpha} \cdot Z_i^r\right) - g\left(\sum_{i=1}^{I} Z_i^r\right) \leq d_r \qquad (9)$$

All vectors $\left(Z_1^r, \ldots, Z_I^r\right)$ that satisfy constraints (9) are possible elements of U_r. Since the generation of ORDS is basically a bin-packing problem, we may apply bin-packing heuristics such as First Fit Decreasing (FFD), Best Fit Decreasing (BFD) and Minimum Bin Slack (MBS) (Gupta and Ho 1999) or a heuristic such as Randomized List Scheduling Heuristic (van den Akker et al. 1999) to generated initial set of ORDSs. Since in a study of off-line bin-packing algorithms by Dell'Olmo and Speranza (1999) Longest Processing Time (LPT) performs well, we use this heuristic for the generation of an initial set of ORDSs for an OR capacity size r. LPT first sorts all procedures of all types in decreasing order of their prediction bound n_i^{α} and then it creates an ORDS in which it plans the longest procedure that fit, i.e., that satisfy constraints (9). If the heuristic reaches the end of the ordered list it closes the ORDS. This is repeated until no surgical procedures remain in the ordered list. The heuristic is executed for all OR capacity sizes.

Pricing problem An optimal solution of the LP relaxation of the restricted problem is optimal for the LP relaxation of the complete master problem if the corresponding dual solution is feasible for the dual problem of the LP relaxation of the master problem. The pricing problem is thus to determine whether there exist ORDSs that are not in the restricted LP relaxation that violate the dual constraints from the LP relaxation of the master problem. Such ORDSs are added to the restricted LP relaxation and a next iteration starts. If such ORDSs do not exist, column generation terminates, and the current restricted LP relaxation solution is optimal to the LP relaxation of the master problem.

The dual constraints of the LP relaxation of the Phase 1 model are:

$$\begin{aligned} \pi_r + \sum_{i=1}^{I} \lambda_i \cdot a_{iu} \leq d_r & \qquad r = 1, \ldots, R \\ \pi_r \leq 0 & \qquad r = 1, \ldots, R \\ \lambda_i \geq 0 & \qquad i = 1, \ldots, I, \end{aligned} \qquad (10)$$

where λ_i are the dual variables corresponding to constraints (6), and π_r the dual variables corresponding to constraints (7) of the Phase 1 LP.

As input for the pricing problem we obtain two vectors $(\bar{\pi}, \bar{\lambda})$ of shadow prices from the restricted LP relaxation. The pricing algorithm now examines whether for this solution $(\bar{\pi}, \bar{\lambda})$ an ORDS $u \in U_r$, represented by $a_{1u}, \ldots, a_{Iu}$, exists that violates the dual constraint (10), i.e. values $a_{1u}, \ldots, a_{Iu}$, with:

$$d_r - \bar{\pi}_r - \sum_{i=1}^{I} \bar{\lambda}_i \cdot a_{iu} < 0 \qquad (11)$$

The left-hand side of constraints (11) are the reduced costs for variable X_u ($u \in U_r$). We evaluate each OR capacity size r separately to determine whether an ORDS exists, formed by a vector $\left(Z_i^r, \ldots, Z_I^r\right)$, that violates the dual constraints (10). In the rth problem we thus need to maximize

$$\sum_{i=1}^{I} \bar{\lambda}_i \cdot Z_i^r$$

over all vectors $\left(Z_i^r, \ldots, Z_I^r\right)$ representing a new ORDS, i.e. satisfying constraint (9).

To solve the pricing problem as an ILP we write the term: $g(\sum_{i=1}^{I} Z_i^r)$ as a telescopic sum. For this purpose, we introduce additional notation. The binary variable A_e indicates whether there are at least e procedures in an ORDS ($e \leq E$, where E is the maximum number of procedures that can be performed during 1 day in one operating room). The function $g(e) := g_1 + \cdots + g_e$ provides the correction for the portfolio effect for e surgical procedures. Using this function and the binary variables A_e, the rth pricing problem ILP becomes:

$$\max \sum_{i=1}^{I} \bar{\lambda}_i \cdot Z_i^r$$

subject to

$$\begin{aligned} \left(\sum_{i=1}^{I} n_i^{\alpha} \cdot Z_i^r\right) - \sum_{e=1}^{E} g_e \cdot A_e \leq d_r \qquad & r = 1, \ldots, R \\ \sum_{i=1}^{I} Z_i^r = \sum_{e=1}^{E} A_e \qquad & \\ A_e \geq A_{e+1} \qquad & e = 1, \ldots, E \\ A_e \in \{0, 1\} \qquad & e = 1, \ldots, E \\ Z_i^r \in \mathbb{N} \qquad & i = 1, \ldots, I. \end{aligned}$$

After this problem is solved for all capacity sizes r, the resulting ORDSs with negative reduced costs are added to the restricted LP relaxation of the Phase 1 model. This model is reoptimized to obtain new shadow prices. Column generation stops if no such ORDSs are found any more. If in practice this process takes very long and generates a large number of extra columns, one might incorporate some of the stopping criteria like the amount of improvement in the LP resulting from the newly generated columns. This may have some effect on the quality of the LP-solution, but since afterwards still an integer solution has to be constructed, the effect on the solution after Phase 2 might be only marginal. In our test instances, we always were able to solve the LP-relaxation to optimality.

Rounding heuristic The solution to the restricted LP relaxation does not directly lead to a starting point for the second phase, since ORDSs may have been selected fractionally. To obtain an integer solution we use a rounding heuristic that rounds down the fractional solution. This results in an integer solution with a small number of surgical procedures that are not assigned to selected ORDSs. These procedures are assigned to newly created ORDSs using an LPT heuristic. There may also be some redundant surgical procedures due to the "$\geq$" sign in constraints (6). We remove these redundant procedures randomly. In general, this approach does not guarantee to result in a feasible solution. However, for the tested instances a quite large fraction of procedures was planned before rounding, only a fraction had to be planned by the LPT heuristic. We never got stuck with infeasible solutions at this stage. If infeasibility might get an issue, the simple rounding heuristic leave room for algorithmic improvements and may be replaced by more elaborate approaches. Summarizing, the output of Phase 1 consists of a set of ORDSs that cover the set of all surgical procedures to be assigned within the MSS.

4.2 Phase 2

In Phase 2 the actual MSS cycle is constructed. We propose an ILP in which the set of ORDSs is assigned to OR-days such that the hospital bed requirements are leveled over the days.

4.2.1 Phase 2 model

Given is a set $\bar{U}$ of ORDSs to be assigned to the OR-days of the MSS. Let $\bar{U}_r \subset \bar{U}$ denote the ORDSs which are of capacity size r. To model the assignment of an ORDS u to an OR-day (j, t) we introduce binary decision variables Y_{ujt} for all $u \in \bar{U}_r$ and $(j, t) \in \varphi_r$. We ensure that the OR capacity sizes match and that at most one ORDS is assigned to an OR on a day. The objective function takes into account the requirements for all hospital beds for all days within one MSS cycle, thus also requirements of surgical procedures that have taken place in previous cycles. Corrected by a normalized priority factor (see Sect. 3.2), we minimize the maximum requirements for hospital beds. The objective function is the second term of the objective function (1) of the base model. This objective function is a minimax objective and can be rewritten to Eq. (12) and constraints (13) in which HB_b is the maximum requirement of hospital bed type b on a given day in the cycle.

All selected ORDSs from Phase 1 must be assigned to an operating room and a day. This is ensured by constraints (14). No more than one ORDS can be assigned to an operating room on a day, which is imposed by constraints (15). Summarizing, the model of Phase 2 is the following ILP:

$$\min \sum_{b=1}^{B} \left[\frac{c_b}{\left\lceil \sum_{i=1}^{I} l_{ib} \cdot s_i \right\rceil / T} \right] \cdot HB_b \tag{12}$$

$$\sum_{r=1}^{R} \sum_{u \in \bar{U}_r} \sum_{(j,t) \in \varphi_r} \sum_{i=1}^{I} \sum_{t=1}^{T} \psi_{t\tau ib} \cdot a_{iu} \cdot Y_{ujt} \leq HB_b \qquad \tau = 1, \ldots, T, \quad b = 1, \ldots, B \tag{13}$$

$$\sum_{(j,t) \in \varphi_r} Y_{ujt} = X_u \qquad r = 1, \ldots, R; \quad u \in \bar{U}_r \tag{14}$$

$$\sum_{u \in \bar{U}_r} Y_{ujt} \leq 1 \qquad r = 1, \ldots, R; \quad (j,t) \in \varphi \tag{15}$$

$$Y_{ujt} \in \{0, 1\} \qquad u \in \bar{U}_r; \quad (j,t) \in \varphi_r$$

$$z_b \geq 0 \qquad b = 1, \ldots, B.$$

4.2.2 Solving the Phase 2 model

We solve the Phase 2 model using the commercial solver ILOG CPLEX 9.0. We use lower bound on the values HB_b to determine the quality of an intermediate solution and to speed up the computation. These lower bounds are calculated by rounding up the sum of the total requirements of hospital beds during one cycle divided by the cycle length:

$$\left\lceil \frac{\sum_{i=1}^{I} l_{ib} \cdot s_i}{T} \right\rceil$$

This represents a theoretical minimum of the maximum requirements for hospital bed type b on 1 day in a cycle. The lower bounds are multiplied by the normative sum used in the objective (1) of the base model:

$$\sum_{b \in B} \left(\frac{c_b}{\left(\sum_{i=1}^{I} ql_{ib} \cdot s_i\right) / T} \right) \cdot \left\lceil \frac{\sum_{i=1}^{I} l_{ib} \cdot s_i}{T} \right\rceil \tag{16}$$

This overall lower bound (16) is given as an initial lower bound to CPLEX to speed up the branch-and-bound process.

5 Computational experiments

We implemented the two-phase approach in the AIMMS mathematical modeling-language 3.5 (Bisschop 1999), which interfaces with the ILOG CPLEX 9.0 LP/ILP solver. We test our approach with realistic data instances from the Erasmus MC based on the available database of surgical procedures that has been collected from 1994 until 2004. This data consists of the frequency of surgical procedures, procedure durations, and data about the usage of hospital beds after surgical procedures.

5.1 Instance generation

Since 1994 Erasmus MC has been collecting data on the frequency of surgical procedures, the duration of procedures, and standard deviation of the duration of procedures. In cooperation with surgeons we defined procedure types by grouping medically homogeneous procedures, which results in the Erasmus MC instance. The data consist for each surgical procedure type i of the frequency of a surgical procedure type during one cycle s_i, the prediction bound n_i^α, and the length of a request of a hospital bed l_{ib}. We vary the parameter values of the cycle length T, the number of operating rooms J, and the number of hospital bed types B (see Table 1), which results in 36 instances types. For each parameter combination 9 additional instances are generated, this yields a total of 360 instances. The additional instances are generated by randomly drawing data from the intervals in Table 2 and rounding them to the nearest integers (the values with a tilde in the table represent the values of the parameters resulting from the Erasmus MC instance).

The cycle length influences the number of procedure types and the number of surgical procedures that can be incorporated into the MSS (Category A procedures). Table 3 shows the dependency between the cycle length and the number of surgical procedure types in Category A together with their numbers and total duration.

We assume that all ORs are available during weekdays and are closed for elective procedures in weekends. For the computational experiments in this paper we use one OR capacity size ($R = 1$) of 450 min ($d_r := 450$). Furthermore, we assume that procedures are finished before their prediction bound in 69% of the cases, i.e., $\alpha := 69\%$. This value is taken from the current practice

Table 1 Parameter values for the instances

Cycle length in days	$T \in \{7, 14, 28\}$
Number of operating rooms	$J \in \{5, 10, 15, 20\}$
Number of hospital bed types	$B \in \{1, 2, 3\}$

Table 2 Intervals for creating instances

$s_i \in [0.9 \cdot \tilde{s}_i, 1.1 \cdot \tilde{s}_i]$
$n_i^\alpha \in [0.9 \cdot \tilde{n}_i^\alpha, 1.1 \cdot \tilde{n}_i^\alpha]$
$l_{ib} \in [0.5 \cdot \tilde{l}_{ib}, 1.5 \cdot \tilde{l}_{ib}]$

Table 3 The relation between the cycle length and procedures in Category A

Cycle length in days	Number of procedure types	Total number of procedures	Total duration of all procedures (in hours)
7	42	56	126
14	109	177	398
28	203	423	952

Table 4 Parameter values for function g, to model the portfolio effect

e	1	2	3	4	5
$g(e)$	$0.00 \cdot \bar{\sigma}$	$0.10 \cdot \bar{\sigma}$	$0.22 \cdot \bar{\sigma}$	$0.36 \cdot \bar{\sigma}$	$0.48 \cdot \bar{\sigma}$

of Erasmus MC. The priority factors of hospital beds are given by: $c(1) := 5$ $c(2) := 2$ $c(3) := 1$.

The function g, which we use to model the portfolio effect, depends on the number of procedures that is scheduled in an ORDS and the average standard deviation $\bar{\sigma}$ of all surgical procedures. We approximate the portfolio effect using the function $g(e)$ that takes the values indicated in Table 4. The value for the average surgical procedure standard deviation $\bar{\sigma}$ is 36, based on the database of the Erasmus MC.

5.2 Test results

In the tests we focus on three different aspects. Firstly, we study the dependencies of the computation times of both phases on the used parameter combinations. Secondly, we investigate the obtained results of the minimization of the required OR capacity. And finally, we address the hospital bed leveling. For this last issue, we have truncated computations that exceed 600 s and have used the best incumbent solutions as output. These incumbent solutions are, therefore, generally not optimal for the Phase 2 model.

5.2.1 Computation times

Table 5 presents the computation times in Phase 1 for all parameter combinations. The computation times in Phase 1 include the initialization and rounding heuristic.

The computation time increases with T, whereas B and J hardly influence the computation time. Similar results are obtained when computation times of the initialization heuristic are considered solely. Here the computation times vary from 0 to 6 s. We conclude that the initialization heuristic only needs a small fraction of time that is required by the complete Phase 1 computation. Table 6 presents the computation time in Phase 2 for all parameter combinations.

Table 5 Computation times of Phase 1 in relation with T, B and J

	$T \rightarrow$	7			14			28		
$J \downarrow$	$B \rightarrow$	1	2	3	1	2	3	1	2	3
5		15.10	17.08	13.76	43.91	47.00	45.90	80.96	78.78	74.56
10		15.29	16.59	13.36	47.12	44.28	45.62	80.24	83.90	87.01
15		16.29	16.12	13.17	47.24	44.70	44.03	80.20	75.96	95.17
20		15.01	16.73	14.35	48.01	45.94	42.39	81.07	75.00	89.70

Table 6 Computation times of Phase 2 in relation with T, B and J

	$T \rightarrow$	7			14			28		
$J\downarrow$	$B \rightarrow$	1	2	3	1	2	3	1	2	3
5		0.30	0.49	0.56	1.55	2.79	3.69	6.16	8.94	13.32
10		0.63	0.93	1.11	3.78	5.86	72.04	15.02	30.54	325.08
15		0.96	1.37	121.60	5.39	8.69	72.92	18.87	43.09	517.08
20		1.21	1.81	122.27	7.45	11.26	87.79	24.54	47.25	478.67

Table 7 Number of times that computation is truncated

	$T \rightarrow$	7			14			28		
$J\downarrow$	$B \rightarrow$	1	2	3	1	2	3	1	2	3
5		0	0	0	–	–	–	–	–	–
10		0	0	0	0	0	1	0	0	3
15		0	0	2	0	0	1	0	0	7
20		0	0	2	0	0	1	0	0	5

Table 6 shows that all three parameters have considerable impact on the computation time and in all cases the computations time increases with increasing parameter value. Table 7 shows the number of times that the calculation is truncated after 600 s for all parameter combinations. The '–' sign denotes that these test instances are infeasible due to the lack of operating rooms.

The extreme growth of the computation time for some of the test instances in Table 6 results mainly from hard instances, where the calculation is truncated (see Table 7). Computation times are not high and therefore allow use of the proposed approach in practice.

5.2.2 OR utilization

Table 8 shows the average number of required ORs per week in relation to the cycle length T. The number of required ORs increases if the cycle length increases, which may be expected since the total surgical procedure volume increases as well (see Table 3). The rounding gap between the integer solution

Table 8 Test results of Phase 1

	Initialization heuristic and column generation		Initialization heuristic only
T↓	Required number of operating rooms during 1 week	Rounding gap(%)	Required number of operating rooms during 1 week
7	16.50	1.25	16.50
14	27.80	0.9	27.80
28	34.18	0.6	34.33

of Phase 1 and the value after rounding up the optimal fractional solution of the LP relaxation denotes the quality of the rounding heuristic. We conclude that the rounding gap is small and decreases if more ORDSs are required. Thus, we may conclude that the achieved OR utilization after Phase 1 is close to the best possible utilization.

Table 8 gives the results of using only the ORDSs generated by the initialization heuristic. These values are found by solving the restricted LP using the initially generated ORDSs and applying the rounding heuristic. They are equal to the values of the complete column generation approach for the construction of MSSs with the cycle length of 7 and 14 days. For larger instances with the cycle length of 28 days, the complete column generation slightly improves the initialization heuristic. Thus, in most of the cases, the ORDSs generated by the initial heuristic already contain the ORDSs needed for the optimal fractional solution of the LP-relaxation of the Phase 1 model. But since an MSS is typically constructed once a year, the additional computational effort of the column generation approach should be used to try to improve the initial solution.

5.2.3 Hospital bed leveling

In this section we discuss the hospital bed leveling. The relative difference between the objective value of the Phase 2 model and the lower bound [see expression (16)] indicates the quality of the solutions found. Table 9 presents the relative differences.

The results in Table 9 show that the difference between the found solutions and the lower bound is small. Therefore, Phase 2 almost optimally levels the hospital bed requirements. This is the more surprising, since the ORDSs in Phase 1 have been generated with the only goal to optimize resource utilization not taking into account the subsequent problem of hospital bed leveling.

In 22 out of 360 experiments the computation of Phase 2 is truncated. Table 10 presents the relative differences between the found solution and the lower bound for the 22 truncated instances.

Even for these instances the average gap is small; the maximum gap is 10.1 %. Based on the presented results we conclude that the constructed MSSs level the hospital bed requirements of the incorporated surgical procedures. This means that the requirements on one day rarely exceed the lower bound.

Table 9 Average gap between the lower bound and the Phase 2 solution

	$T \rightarrow$	7			14			28		
$J \downarrow$	$B \rightarrow$	1	2	3	1	2	3	1	2	3
5		0.0%	0.0%	0.5%	–	–	–	–	–	–
10		0.0%	0.0%	0.5%	0.0%	0.0%	0.2%	0.0%	0.0%	1.3%
15		0.0%	0.0%	0.5%	0.0%	0.0%	0.2%	0.0%	0.0%	2.4%
20		0.0%	0.0%	0.5%	0.0%	0.0%	0.2%	0.0%	0.0%	1.5%

Table 10 Average gap between the lower bound and the Phase 2 solution for truncated instances

	$T \rightarrow$	7			14			28		
$J \downarrow$	$B \rightarrow$	1	2	3	1	2	3	1	2	3
5										
10							1.9%			4.5%
15				2.7%			1.9%			3.4%
20				2.7%			1.9%			3.0%

6 Conclusions and further research

The computational experiments show that generation of MSSs is well possible within acceptable time bounds by the proposed two-phase decomposition approach. The proposed solution approach generates MSSs that minimize the required OR capacity for a given set of procedures and level the hospital bed requirements well. The chosen solution approach makes it possible to add restrictions imposed by personnel and to consider other types of hospital resources than beds. This flexibility is required to implement an OR planning strategy that includes an MSS. The approach has been successfully tested on real data from Erasmus MC. The hospital management is pleased with the outcomes, and encourages and initiates further research into implementing the MSS-approach in practice.

In further research we will investigate implementation aspects, and scheduling of Category B and C procedures as such is required to determine the overall benefits of cyclic scheduling of OR departments. This research should also provide insight into the benefits of a cyclic OR planning approach for hospitals with various patient mixes. Furthermore, we will investigate the leveling of hospital beds when the length of request for beds is assumed to be stochastic.

The repetitive nature of our cyclic surgical planning approach yields that it reduces the overall management effort. In addition, it not only optimizes OR utilization but also levels the output towards wards and ICU. This results in less surgery cancelations, and thus a reduction of the lead-time of the patient's care pathway. Therefore, MSS contributes to an improved integral planning of hospital processes. The intensive cooperation with clinicians and OR managers has lead to a framework for cyclic OR planning and a method for construction of MSSs that can handle constraints imposed by health care processes. This flexibility ensures the applicability of the developed method in OR departments and hospitals.

References

Bakker H, Zuurbier JJ (2002) Pareto's law (In Dutch). Zorgvisie 7

Barnhart C, Johnson EL, Neuhauser GL, Savelbergh MWP, Vance PH (1998) Branch-and-price: column generation for solving huge integer programs. Oper Res 46(3):316–329

Beliën J, Demeulemeester E (2005) Building cyclic master surgery schedules with leveled resulting bed occupancy. Eur J Oper Res (in press)

Bisschop J (1999) AIMMS – optimization modeling. Paragon Decision Technology B.V., Harlem
Blake JT, Donald J (2002) Mount sinai hospital uses integer programming to allocate operating room time. Interfaces 32(2):63–73
Brucker P, Drexl A, Mohring R, Neumann K, Pesch E (1999) Resource-constrained project scheduling: notation, classification, models, and methods. Eur J Oper Res 112:3–41
Carter M (2002) Diagnosis: mismanagement of resources. Oper Res Manage Sci Today April 26–32
Charnes A, Cooper W, Thompson G (1964) Critical path analyses via chance constrained and stochastic programming. Oper Res 12:460–470
Dell'Olmo P, Speranza MG (1999) Approximation algorithms for partitioning small items in unequal bins to minimize the total size. Discret Appl Math 94:181–191
Gerhak Y, Gupta D, Henig M (1996) Reservation planning for elective surgery under uncertain demand for emergency surgery. Manage Sci 42(3):321–334
Glouberman S, Mintzberg H (2001) Managing the care of health and the cure of disease Part i: Differentiation. Health Care Manage Rev 26(1):56–69
Goldratt EM (1997) Critical chain. The North River Press, Great Barrington
Guinet A, Chaabane S (2003) Operating theatre planning. Int J Prod Econ, 85:69–81
Gupta JND, Ho JC (1999) A new heuristic algorithm for the one-dimensional bin-packing problem. Prod Plan Control 10(6):598–603
Hans EW, Wullink G, Van Houdenhoven M, Kazemier G (2006) Robust surgery loading. Eur J Oper Res (accepted for publication)
Jebali A, Hadj Alouane AB, Ladet P (2006) Operating room scheduling. Int J Prod Econ 99:52–62
Kim S, Horowitz I (2002) Scheduling hospital services: the efficacy of elective-surgery quotas. Omega 30:335–346
Lamiri M, Xie X, Dolgui A, Grimaud F (2005) A stochastic model for operating room planning with elective and emergency surgery demands. In: Conference Proceedings ORAHS, Southampton, UK
McManus ML, Long MC, Copper A, Mandell J, Berwick DM, Pagano M, Litvak E (2003) Variability in surgical caseload and access to intensive care services. Anesthesiology 98:1491–1496
Millar HH, Kiragu M (1998) Cyclic and non-cyclic scheduling of 12 h shift nurses by network programming. Eur J Oper Res 104:582–592
Neumann K, Zimmermann J (2000) Procedures for resource leveling and net present value problems in project scheduling with general temporal and resource constraints. Eur J Oper Res 127:425–443
OECD (2005) Oecd health data 2005 – statistics and indicators for 30 countries
Ogulata SN, Erol R (2003) A hierarchical multiple criteria mathematical programming approach for scheduling general surgery operatons in large hospitals. J Med Syst 27(3):259–270
Ozkaraham I (2000) Allocation of surgeries to operating rooms by goal programming. J Med Syst 24(6):339–378
Pinedo M (2005) Planning and scheduling in manufacturing and services. Springer Series in Operations Research and Financial Engineering. Springer, Berlin Heidelberg New York
Schmidt E, Dada M, Adams D (2001) Using cyclic planning to manage capacity at alcoa. Interfaces 31(3):16–27
Sier D, Tobin P, McGurk C (1997) Scheduling surgical procedures. J Oper Res Soc 48:884–891
Strum DP, May JH, Vargas LG (2000) Modeling the uncertainty of surgical procedure times: comparison of log-normal and normal models. Anesthesiology 92:1160–1167
Tayur S (2000) Improving operations and quoting accurate lead times in a laminate plant. Interfaces 30(5):1–15
van den Akker JM, Hoogeveen JA, van de Velde SL (1999) Parallel machine scheduling by column generation. Oper Res 47(6):862–872
Vissers JMH, Adan IJBF, Bekkers JA (2005) Patient mix optimisation in cardiothoracic surgery planning: a case study. IMA J Manage Math 16(3):281–304
Weissman C (2005) The enhanced postoperative care system. J Clin Anesth 17:314–322
Williams HP (1999) Model building in mathematical programming, 4th edn. Wiley, New York