HagaZiekenhuis van Den Haag

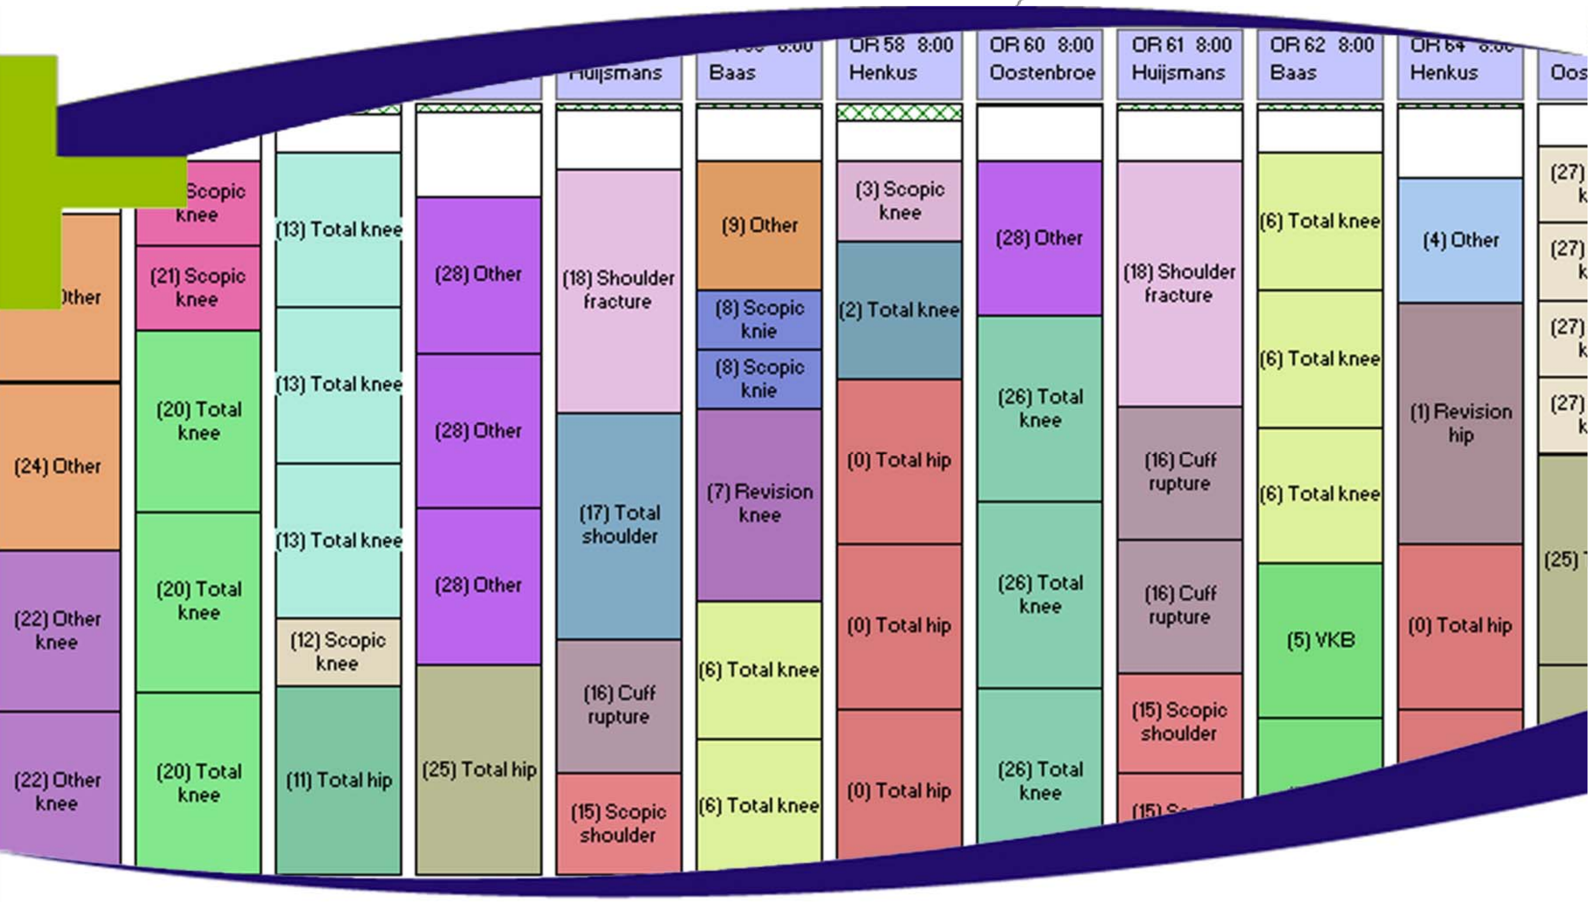

Better utilisation of the OR with less beds

A tactical surgery scheduling approach to improve OR utilisation and the required number of beds in the wards

Author :
J.M. (Joël) Bosch, BSc

August 2011
Public version

Supervisors:
Dr. Ir. E.W. Hans , University of Twente
Ir. J.T. van Essen, University of Twente
A.A.J. van der Zalm, HagaZiekenhuis

UNIVERSITY OF TWENTE.

Management summary

HagaZiekenhuis has an increased incentive to perform their processes more efficiently. Since the operating room (OR) capacity is the most expensive resource, management and surgeons believe that focussing on OR planning and processes will yield the highest improvements on efficiency problems. Also, HagaZiekenhuis faces the problem that health insurers force them to reduce the number of beds. In this research we focus on efficient use of OR capacity and reducing the total number of required beds. In addition, we want see if it is feasible to close wards in the weekend.

Our research objective is *to develop a tactical surgery scheduling approach that maximises the OR utilisation while minimising the required number of beds in the wards.*

To develop a tactical surgery scheduling approach, we perform a case study on the orthopaedic department. A performance analysis shows that the OR utilisation in 2010 was 78%. The wards required at most 52 beds on 95% of all the days in 2010. Only 45 beds are physically available in the wards because some patients can occupy the same bed on a day. For example, a bed from a patient who is discharged in the morning can be reused for another patient who is hospitalised in the afternoon.

Approach

Based on a literature review, we propose a non-cyclic master surgery scheduling (MSS) approach that maximises the OR utilisation while minimising the required number of beds. The MSS approach is based on the paper of Van Oostrum et al. [15]. We used the model of Vanberkel et al. [16] to determine the length of stay. Our surgery scheduling approach contributes to literature because we model both the surgery durations and the length of stay as random variables. Next to that, our MSS is non-cyclic. Also, we do not only schedule surgery types, but also surgeons. Furthermore, the number of surgeries per surgery type is variable while, in most models from literature, a fixed number of surgeries is scheduled. We also introduce a new method to linearise the constraint to limit the probability of overtime. The method finds a minimal number of piecewise linear functions to approximate the probability of overtime.

The surgery scheduling approach consists of two phases. In the first phase, we apply a column generation technique to generate a set of operating room day schedules (ORDSs) that optimise OR utilisation. In the second phase, we apply a simulated annealing approach that swaps the assignment of ORDSs to OR-days around in order to optimise the number of required beds at a 95% quantile. The required number of beds at a 95%-quantile means that the probability of having sufficient beds on a day is at least 95%.

Results

We show that OR utilisation highly depends on the amount of slack to buffer against overtime. The amount of slack is determined based on the maximum probability of overtime, the variance of surgery durations, and the number of surgeries per ORDS. Furthermore, OR utilisation improves when we increase opening hours or when we define more surgery types. The required number of beds increases with the number of surgeries per ORDS. Adjusting the current set of instruments, surgeon availability, and available OR-days does not influence on OR utilisation and the required number of beds.

The results show that OR utilisation improves from 78% to 90% by using the our surgery scheduling approach. We show that is not feasible to close wards during the weekends. Neither can the number of physical beds be reduced for all days of the week. Our model assumes that patients occupy a bed for the whole day of hospitalisation and discharge. In practice, however, some patients can occupy the same bed on a day. Therefore, the required number of beds that results from our model is a worst-case scenario for bed occupancy. For that reason, it is likely that beds can be closed in practice.

Recommendations

Our main recommendation is to apply our tactical surgery scheduling approach. A tactical surgery schedule should be generated every 4-week period, 3 months in advance. We give general recommendations and tips for implementation.

General recommendations

- *Schedule surgeries based on surgery durations.* To improve the estimation of intervention and changeover time, surgeries should be scheduled based on surgery durations from historical data.

- *OR capacity is known before constructing the MSS.* We recommend that the OR department finds ways to increases the reliability of assigning OR capacity to specialisms, since knowing the available OR capacity is a prerequisite of our surgery scheduling approach.

- *(Re-)define the start of OR-days.* Differentiating the start of OR-days may reduce the surgery duration of the first surgery on an OR-day. Also, the OR department should define what activity marks the start of an OR-day.

- *Improve quality of data.* It is of major importance to have reliable data to find bottlenecks in processes, to evaluate performance, and to determine the effect of interventions.

Implementation

- *Phase 1: Adjust the information system.* Instead of manual processing the data, surgery durations and length of stay per surgery type should be provided by the information system. Furthermore, diagnose codes should be re-defined in correspondence with the surgery types. A diagnose code should be assigned to each surgery request that is made by a surgeon. Also, a surgery request should be created in the information system directly instead of filling in a paper ‘admission form’.

- *Phase 2: Determine the surgery type demand per 4-week period.* The surgery scheduling approach requires to define a minimum and maximum demand per surgery type. Prior to applying the surgery scheduling approach, management should decide how the production agreements and waiting lists should relate to the minimum and maximum demand per 4-week period. The minimum and maximum demand per surgery type should be refined after each planning horizon based in an evaluation of planned and actual production.

- *Phase 3: Develop a decision support (DSS) tool.* A DSS-tool can be developed in house or by a software development company. We estimate that in house development costs €100,000. This investment pays back within a year when the DSS-tool improves the average OR utilisation by at least 2.25%. This is feasible, since we show that our surgery scheduling approach improves OR utilisation by 12%.

- *Phase 4: Use the DSS-tool for evaluation of the current scheduling practice.* To gain commitment of the OR planners, the DSS-tool could first be used as an evaluation tool. This evaluation tool does not optimise the surgery schedule, but determines the expected OR utilisation and bed occupancy of a given surgery schedule. After a couple of weeks, the expected performance of the model could be compared with the actual performance to show that the forecasts of the model are accurate.

- *Phase 5: Perform a pilot in the orthopaedic department.* The tactical surgery schedule can be implemented by requiring the OR planners to assign patients to their corresponding surgery type slot. In the process, the model may need some fine-tuning of the surgery type definitions and instrument constraints. When the actual bed occupancy goes down, management could decide to close some beds. To account for urgent surgeries, not all surgery type slots should be filled with elective patients. The OR planners can still try to find an elective patient for surgery type slots that remains empty a couple of days before surgery, but this decreases patient service. Management should therefore decide upon a balance between OR utilisation and patient service by determining up to what point elective patients can still be scheduled.

- *Phase 6: Use the DSS-tool for other specialisms.* When the DSS-tool is in place at the orthopaedic department, the next step is to apply the DSS-tool for other specialisms.

Management samenvatting

Het HagaZiekenhuis heeft een toenemende stimulans om hun processen efficiënter in te richten. Aangezien de operatiekamer (OK) de duurste capaciteit is, denken management en specialisten dat het focussen op OK planning en processen de grootste efficiëntie verbeteringen zal opleveren. Verder heeft het HagaZiekenhuis te maken met een toenemende druk van zorgverzekeraars om het aantal bedden te verminderen. Dit onderzoek richten wij op efficiënt gebruik van OK capaciteit en het verminderen van het benodigd aantal bedden. Daarnaast willen we zien of het haalbaar is om een verpleegafdeling te sluiten in het weekend.

Ons onderzoeksdoel is *het ontwikkelen van een tactische operatie planning aanpak die OK benutting maximaliseert en het aantal benodigde bedden op de verpleegafdeling minimaliseert.*

Voor het ontwikkelen van een tactische operatie planning aanpak voeren we een casestudy uit bij de orthopedie afdeling. Analyses wijzen uit dat de OK benutting 78% was in 2010. De verpleegafdeling had maximaal 52 bedden nodig op 95% van alle dagen in 2010. Feitelijk heeft de verpleegafdeling maar 45 fysieke bedden omdat sommige patiënten hetzelfde bed gebruiken op een dag. Bijvoorbeeld, een bed van een patiënt die 's ochtends ontslagen wordt kan weer gebruikt worden door een andere patiënt die 's middags wordt opgenomen.

Aanpak

Op basis van een literatuur studie, stellen we voor om een non-cyclische Master Surgery Scheduling (MSS) aanpak toe te passen die de OK benutting maximaliseert en het aantal benodigde bedden minimaliseert. De MSS aanpak is gebaseerd op het artikel van Van Oostrum et al. [15]. We gebruiken het model van Vanberkel et al. [16] om de ligduur te modelleren. Onze operatie planning aanpak draagt bij aan de literatuur omdat we zowel de operatieduur als de ligduur modelleren als stochastische variabelen. Ook plannen wij niet alleen operatietypen, maar ook specialisten. Verder is het aantal geplande operaties per operatie type in onze aanpak variabel terwijl in de meeste modellen in de literatuur een vast aantal operaties wordt gepland. Ook introduceren we een nieuwe methode voor het lineariseren van de restrictie om de kans op overtijd te limiteren. The methode bepaald een minimum aantal stuksgewijze lineaire functies die de kans op overtijd benaderen.

De operatie planning aanpak bestaat uit twee fases. In de eerste fase passen we een kolom-generatie methodiek toe om een set van operatiekamer dagschema's (ORDSs) te genereren die de OK benutting maximaliseren. In fase twee voeren we een simulated annealing procedure uit. Tijdens de procedure wordt de toewijzing van ORDSs aan OK-dagen continu verwisselt om tot een oplossing te komen die het aantal benodigde bedden minimaliseert op een 95% kwantiel. Met het aantal bedden op een 95% kwantiel bedoelen we dat de kans minimaal 95% is dat er op een bepaalde dag voldoende bedden zijn.

Resultaten

We tonen aan dat OK benutting voornamelijk afhankelijk is van de hoeveelheid slack die nodig is om te bufferen tegen overtijd. De hoeveelheid slack wordt bepaald door de maximum kans op overtijd, de variantie van operatietypen en het aantal operaties per ORDS. Verder verbetert de OK benutting als we de openingstijden van de OK verlengen of als we meer operatietypen definiëren. Het aantal benodigde bedden stijgt met het aantal operaties per ORDS. Het veranderen van de huidige set van instrumenten, de beschikbaarheid van de specialisten en de beschikbare OK-dagen hebben geen invloed op de OK-benutting and het aantal benodigde bedden.

Met het toepassen van onze operatie planning aanpak kan de OK benutting verhoogd van 78% naar 90%. We tonen aan dat het niet haalbaar is om een verpleegafdeling te sluiten in het weekend. Ook het aantal bedden op de afdeling kan niet gereduceerd worden. In ons model nemen we aan dat patiënten een bed bezetten voor de volledige dag van opname en ontslag. In de praktijk kunnen sommige patiënten hetzelfde bed gebruiken op een dag. Daarom is het aantal benodigde bedden in ons model een worst-case scenario voor de bedbezetting. Om die reden zal het naar alle waarschijnlijkheid in de praktijk wel mogelijk zijn om bedden te sluiten.

Aanbevelingen

Onze voornaamste aanbeveling is om onze operatie planning aanpak toe te passen. Een tactische operatie planning voor een periode van 4 weken zou 3 maanden van tevoren opgesteld moeten worden. We geven algemene aanbevelingen en tips voor implementatie.

Algemene aanbeveling

- *Plan operaties op basis van operatieduur.* De schatting van snijtijden en wisseltijden kan worden verbeterd door operaties te plannen op basis van de operatie duur vanuit historische data.

- *OK capaciteit is bekend voor het bepalen van een MSS.* We bevelen aan dat de OK afdeling de betrouwbaarheid van toebedeelde OK-tijd aan specialismen verbeterd. Het is namelijk van belang om met zekerheid de OK capaciteit te weten voordat onze operatie planning aanpak wordt toegepast.

- *(Her-)definieer de start van een OK-dag.* De operatieduur van de eerste operatie van een OK-dag kan verlaagd worden door OK's op verschillende tijdstippen te laten starten. Verder zou de OK afdeling duidelijk moeten definiëren welke activiteit de start van een OK-dag markeert.

- *Verbeter de kwaliteit van data.* Het is van groot belang om betrouwbare data te hebben om bottlenecks in processen te vinden, om prestaties te bepalen en het effect van interventies te meten.

Implementatie

- *Fase 1: Aanpassing van het informatiesysteem.* In plaats van het handmatig verwerken van data, zou de duur van operaties en de ligduur van patiënten automatisch verwerkt moeten worden door het informatiesysteem. Verder zouden de diagnose codes opnieuw gedefinieerd moeten worden in overeenstemming met de operatietypen van de operatie planning aanpak. Een diagnose code zou toegewezen moeten worden aan ieder operatie aanvraag. Ook zou een operatie aanvraag aangemaakt moeten worden in het informatiesysteem in plaats van het invullen van een papieren 'opname formulier'.

- *Fase 2: Definieer de vraag per operatie type voor elke 4-weekse periode.* De operatie planning aanpak vereist een minimum en maximum vraag per operatie type. Voordat de operatie planning aanpak wordt toegepast, moet het management bepalen hoe de productieafspraken en wachtlijsten gerelateerd moeten worden aan de minimum en maximum vraag per operatie type voor elke periode van 4 weken. De minimum en maximum vraag per operatie type moet na elke planningshorizon verfijnd worden op basis van een evaluatie van de geplande en gerealiseerde productie.

- *Fase 3: Ontwikkel een beslissing ondersteuning tool.* Een 'decision support tool' (DSS) kan intern of door een software bedrijf worden ontwikkeld. We schatten dat interne ontwikkeling €100,000 kost. Deze investering betaalt zich binnen een jaar terug als de DSS-tool de gemiddelde OK benutting met tenminste 2.25% verhoogd. Dit is haalbaar omdat we aantonen dat onze operatie planning aanpak de OK benutting verbeterd met 12%.

- *Fase 4: Gebruik de DSS-tool voor de evaluatie van de huidige planning methodiek.* Om betrokkenheid te verkrijgen van de OK planners zou de DSS-tool gebruikt kunnen worden als evaluatie-instrument. Dit evaluatie-instrument optimaliseert de operatie planning niet, maar bepaald de OK-benutting en bedbezetting op basis van een gegeven operatie planning. Na een aantal weken kan de verwachte prestatie van het model vergeleken worden met de gerealiseerde prestatie om aan te tonen dat het model accuraat is.

- *Fase 5: Voer een pilot uit in de orthopedische afdeling.* Een tactische operatie planning kan geïmplementeerd worden door de OK planners patiënten in te laten plannen in het slot van een overeenkomstig operatie type. Bij het eerste gebruik kan blijken dat de definities van operatietypen en/of instrumenten restricties bijgesteld moet worden. Wanneer na verloop van tijd de bedbezetting daadwerkelijk omlaag gaat, kan het management besluiten om een aantal bedden te sluiten. Om rekening the houden met urgente operaties moeten niet alle slots in de planning worden gevuld met electieve patiënten. De OK planners kunnen proberen om alsnog electieve patiënten in te plannen voor slots die een paar dagen voor een operatiedag nog leeg zijn, maar dit verlaagd de patiënt vriendelijkheid. Het management moet zich daarom uitspreken over een balans tussen OK benutting en patiënt vriendelijkheid door te bepalen tot welk punt electieve patiënten nog ingepland mogen worden.

- *Fase 6: Gebruik de DSS-tool voor andere specialismen.* Zodra de DSS-tool volledig functioneert op de orthopedie afdeling kan deze ook worden toegepast bij andere specialismen.

Preface

Sitting behind my desk, I see people come and go from the parking lot to the 'HagaZiekenhuis' hospital. An occasional ambulance disrupts the peaceful scene on this sunny day. Behind the parked cars, I see a flat that you would expect to encounter in an Eastern European country. It is the flat where I lived during the last stage of my study. Walking from my home to my office only took me a couple of minutes each day since they lie 300 meters apart. This 300 meter radius is the place where I have spent most of my time working on my master thesis. A master thesis that I idealised as the grand finale of my Industrial Engineering & Management study, but appeared to be more cumbersome than I had expected. Moments of enthusiasm and frustration alternate frequently. After all, I am proud of the result of my hard work in the form of this report.

In the first few months, I was the only graduate from University of Twente at HagaZiekenhuis. I appreciated the company of Erik and Frank who joined the graduate group later on. During a workday we always had a few discussions about the research of one of us to help each other out. Twice a week, we changed into our running clothes to run through the park or to the beach. Almost every evening, we had dinner together with some other flatmates. Being together with this group of like-minded people was a great motivation during my master thesis.

I thank Mark van Houdenhoven for the weekly sessions we had with him. Even though he has a busy schedule, he paid interest in our research and gave us useful feedback. I thank Theresia van Essen for her support at HagaZiekenhuis. The meetings with her helped me to keep my project on track. I greatly appreciated it to have her around as a PhD candidate for quick academic feedback. I thank Arnoud van der Zalm for his practical feedback as he has many years of experience in HagaZiekenhuis and healthcare in general. Also, I thank Erwin Hans for triggering my interest of doing my master thesis in healthcare. For me, he is one of the most inspiring teachers of my master study. During my master thesis project, his enthusiasm and pep-talks helped me to keep motivated and his methodical thinking and his knowledge helped me to structure my research.

During my time at HagaZiekenhuis, I was involved with several side-projects concerning logistical issues. It made me realise that there is much to gain from improving health care logistics. I hope my research contributes to a more efficient health care system that remains affordable in the future. My ambition is to continue contributing to the improvement of health care logistics in my working life.

Joël Bosch

Den Haag, July 2011

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this report, we perform research on how to improve the OR utilisation and the required number of beds at HagaZiekenhuis. Based on a case-study at the orthopaedic department, we develop an tactical surgery scheduling approach. Our surgery scheduling approach is based on the research of Van Oostrum et al. [15]. The most important difference between his and our model is that we propose a non-cyclic master surgery schedule approach in which we do not only schedule surgery types, but also surgeons. Furthermore, we assume both surgery types and length of stay to be random variables and we let the number of scheduled surgeries per surgery type be variable.

Section 1.1 describes the context of our research. Section 1.2 discusses the research motivation. After that, we describe the problem description in Section 1.3. In Section 1.4, we discuss the research objective that follows from the problem description. In Section 1.5, we conclude with the research questions that structure the outline for the rest of this report.

1.1 Context

The research in this thesis is performed in HagaZiekenhuis in The Hague, The Netherlands. HagaZiekenhuis is a top-clinical teaching hospital with 245 specialists, 729 beds and 35.571 admissions in 2010 (Table 1.1). The hospital has three locations, because it originated in 2004 as a merger between three hospitals. The 'Leyweg' location is the largest one, where the focus is on urgent and complex care. At the 'Sportlaan' location, the emphasis is on elective care. The third location is 'Juliana Kinderziekenhuis' which is a children's hospital.

Number of employees	3,763
Number of specialists	245
Number of beds	729
Number of admissions	35,571
First outpatient visits	209,500
Number of day care patients	28,808
Average length of stay	5.2 days

Table 1.1: Key figures HagaZiekenhuis (source: annual report 2010)

1.2 Research motivation

From a historical perspective, the health care industry has paid relatively little attention to logistic issues compared to other industries like manufacturing [5]. The board of directors at HagaZiekenhuis believes that there is much to gain from improving their healthcare logistics. Since the operating room (OR) capacity is the most expensive resource and most problems are perceived there, management and surgeons believe that focussing on OR planning and processes will yield the highest efficiency improvements. Also, HagaZiekenhuis faces the problem that health insurers force them to reduce the number of beds. Academic research and mathematical techniques seem a promising solution to tackle these problems. In this research we focus on efficient use of OR capacity and reducing the total number of required beds.

1.3 Problem description

There are multiple stakeholders who are involved in the functioning of the OR department. It is important to consider which problems are perceived by these stakeholders in an early stage because stakeholder dynamics are generally more complex in health care than in manufacturing environments [6]. We distinguish the following stakeholders: patients, hospital management, speciality management, surgeons, OR personnel, OR planners and nurses. We subsequently describe the problems they perceive.

Patients Perhaps the most important issue for patients who undergo surgery, aside from the quality of care, is access time. Access time for orthopaedic surgery is long in comparison with other hospitals, but according to the doctors this is primarily caused by the positive reputation of the department in the region. Still, patients would prefer to have shorter access times.

Hospital management In general, hospital management is interested in increasing their margins, for example by improving the efficiency of the use of OR capacity. Apart from that, bed occupancy is currently a main issue because insurance companies force hospitals to decrease the number of beds. Therefore, hospital management would like to know how surgery scheduling can contribute to minimising the required number of beds by the wards.

Speciality management Each speciality has a fixed number of operating rooms they can use each day from 8.00 to 16.00. Since this OR capacity is an expensive resource and a major bottleneck in the admission of patients, it is in the best interest of the speciality management to use the OR capacity as efficiently as possible. This generates more revenues while using the same amount of resources. Management perceives that a lot of OR capacity is underutilised due to late starts of the first surgery of the day, long changeovers between surgeries and early endings of OR programs.

Surgeons Like management, surgeons want to work more efficiently. However, their focus is on their own time instead of the OR capacity as a whole. Surgeons are compensated per surgery and do not directly experience expenses within the OR department.

OR personnel There are multiple groups of OR personnel like surgery assistants, anaesthetists, anaesthesia assistants, holding and recovery personnel and cleaners. We consider them as one stake holding group because they have similar interests in how the OR functions. Unlike surgeons, OR personnel is employed by the hospital and therefore they are more inclined than doctors to claim their breaks and not to work overtime. This is also caused by a high percentage of OR personnel being hired from outside contractors. The main problem OR personnel encounter, is that they feel rushed by the surgeons when changeovers between surgeries take place. This is understandable from the surgeon's point of view because he does not like to wait (he wants to use his own time as efficiently as possible). Consequently, OR personnel claims that surgeons are not present during activities that have to be carried out when a changeover takes place.

OR planners Some speciality departments have secretaries who schedule surgeries and perform related administrative tasks. They attempt to make a good surgery schedule but they have to deal with several parties who have conflicting interests. The main conflict is to differentiate between the access time of patients for financial reasons. Management may want to perform more knee than hip surgeries, but this may increase the access time of patients who need a hip replacement. The OR planners have trouble to explain to patients why some have to wait four months for hip surgery while others only need to wait a month to have a knee replaced. Another major difficulty is that the number of available operating rooms may be reduced only two weeks in advance due to lack of OR personnel. Also, urgent surgeries could be added to the schedule for medical reasons. Due to this uncertainty about the available OR capacity and surgery demand, the OR planners schedule surgeries for a short planning horizon (2 to 3 weeks). This results in ad-hoc way of planning which leads to a suboptimal OR utilisation and stress for the OR planners.

Nurses The wards are no bottlenecks because the nurses always manage to have a bed available for a patient. This already implies that, most of the time, more beds are available than needed. However, due to the variability in bed occupancy sometimes all the beds are occupied. The variability is not a problem for the nurses, because they can adjust the number of personnel to the number of patients and the care that they need. The problem is that the personnel rosters are made 2 months in advance because of employees agreements. In that stage, nurses do not know how many and what kind of patients to expect because the surgery schedule is not fully known. Therefore, they need to make some ad-hoc changes in the personnel roster when there is a gap between the expected and the actual number of patients, but there is limited flexibility in making these last-minute adjustments.

1.4 Research objective

From the problem descriptions we conclude that the stakeholders have three major objectives:

To maximise OR utilisation Hospital and speciality management, and to some lesser extent the surgeons, would like to improve the efficiency of the OR capacity use. As described earlier, this means increasing revenues with the same amount of resources. Formulated in a more measurable metric, improving the efficiency can be expressed as maximising the OR utilisation. We give a more formal definition of the OR utilisation in Chapter 2. Furthermore, the access time for patients decreases when a higher OR utilisation is achieved.

To minimise the total number of required beds This objective is primarily based on the problem description of the hospital management. Since they are at the highest level in the organisational hierarchy, this single objective is an important one.

To structure the scheduling process Most of the problems are due to the fact that surgery scheduling is largely done based on intuition and experience. This causes the surgery scheduling activity to be perceived as a 'black box' by the other stakeholders. Also, this way of scheduling results in a limited number of rules of thumb that can be taken into account. A more automated and/or formalised way of scheduling could contribute to a more efficient surgery schedule because all relevant factors can be included. Furthermore, more formalised rules create awareness about the impact of scheduling on the OR processes and may increase the predictability of how many and which kind of patients are hospitalised at the wards. Also, it enables management to have better control on the bed occupancy.

We focus on both maximising the OR utilisation and minimising the total number of beds. We do not put particular emphasis on the objective to have a structured planning process since it is hard to quantify. Furthermore, we believe that a structured planning process is a by-product of the other two objectives.

We observe that there are two main ways of intervention to attain the objectives: improve surgery scheduling and improve OR processes. Surgery scheduling is about *what* to do while OR processes are about *how* to do it. It is likely that OR processes have more impact on OR utilisation than surgery scheduling, but improving the OR processes will not have any influence on the other objectives. On the other hand, improving surgery scheduling has impact on all the objectives. Figure 1.1 visualises this line of reasoning.

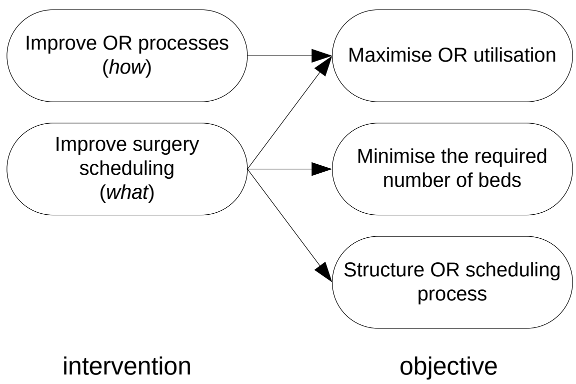

Figure 1.1: Objectives and interventions for this research

To demarcate the scope of this research, we focus on intervening on surgery scheduling. First, improving surgery scheduling has impact on all the objectives while improving OR processes only influences the OR utilisation. Second, there is sufficient data available to study the surgery scheduling process and performance while researching the OR processes requires new measurements. Third, we believe that solutions for improving the OR processes have a more organisational nature while a quantitative approach is more suitable to improve the scheduling process. Both improving the OR processes and surgery scheduling is needed, but since the researcher's interest and expertise is more in the field of quantitative analysis, we focus on improving the surgery scheduling process.

We position this research based on the healthcare framework for planning and control as proposed by Hans et al. [7]. This framework spans four hierarchical levels of control and four managerial areas as displayed in Figure 1.2. Since this research focuses on surgery scheduling, it is identified as resource capacity planning on a tactical level. This implies that we assume strategical decisions to be fixed (the number or ORs per day, amount of OR personnel) and that we do not include operational scheduling (case scheduling and emergency coordination) in our research. Likewise, we do not intervene on the medical planning (what kind of patients to treat), materials planning (what types of materials to have in stock) and financial planning.

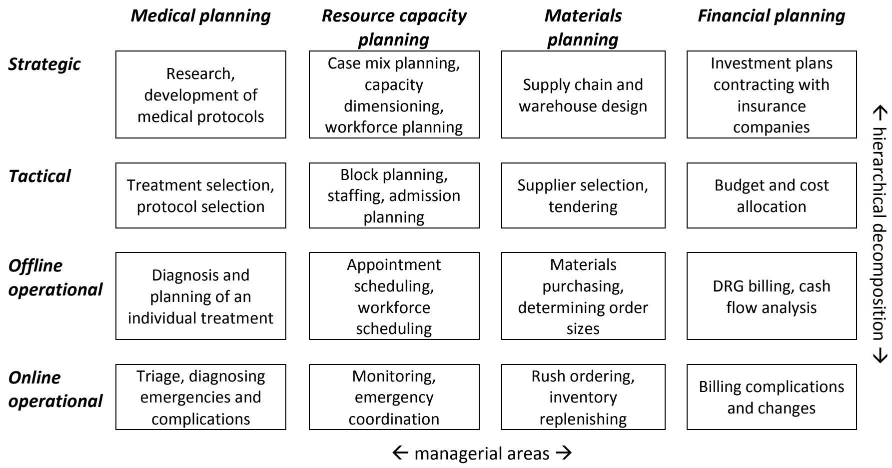

Figure 1.2: A framework for health care planning and control [7]

In conclusion, we state to following research objective:

Research objective

To develop a tactical surgery scheduling approach that maximises the OR utilisation while minimising the required number of beds in the wards.

1.5 Research questions

Based on our research objective, we pose the following research questions. Each chapter of our report corresponds to a research question. Figure 1.3 visualises the structure of our report.

How is the system organised and how does it perform? (Chapter 2)
We answer this research question based on a case study of the orthopaedic department. We describe the system and how it is controlled. After that, we perform a data-analysis to assess how the current system performs and to pinpoint the causes of poor performance.

Which kind of approaches could be used to optimise the surgery schedule? (Chapter 3)
We carry out a literature review to come up with several approaches to optimise the surgery schedule.

How should the surgery scheduling approach be modelled? (Chapter 4)

We model a surgery scheduling approach by first formulating a conceptual model. Next, we gather data and formulate a technical model. Finally, we validate the model by comparing the output of our model with a Monte Carlo simulation.

How does the proposed surgery scheduling approach perform? (Chapter 5)
We perform numerical experiments to determine how the model performs under different circumstances. We introduce a number of scenarios and adjustment of experimental factors and discuss the results.

What are the main findings and how should they be implemented? (Chapter 6)
We summarise the findings of this report and discuss recommendations for HagaZiekenhuis. The recommendations include general recommendations and tips on how to implement the tactical surgery scheduling approach.

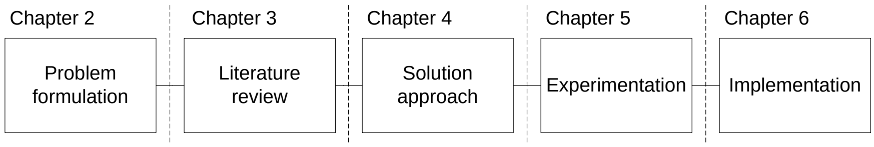

Figure 1.3: Structure of this report

Chapter 2

Context analysis

To develop a solution approach for tactical surgery scheduling, we conduct a case study on a specific speciality: the orthopaedics department. To perform the context analysis, we use data from the hospital database (SAP). We first describe the system of the orthopaedic and OR department in Section 2.1. Section 2.2 describes how the system is controlled by explaining the surgery scheduling process. In Section 2.3, we introduce performance indicators to analyse the performance of the system and to determine the causes of poor performance. Finally, we conclude this chapter in Section 2.4.

2.1 System description

Two departments are relevant for this case study. First, we describe the orthopaedic department in Section 2.1.1. Next, we describe the OR department in section 2.1.2.

2.1.1 Orthopaedic department

Table 2.1 displays key figures of the orthopaedic department. We define an OR day as a combination of an operating room and a day. For example, 2 operating rooms on each day of the week is equal to $2 \times 5 = 10$ OR days. The department is large in comparison with other orthopaedic departments in Dutch hospitals. Due to the size of the department, orthopaedic surgeons have a high degree of specialisation. The department is located at the 'Sportlaan' location and is only assigned to OR capacity at this location. The 'Sportlaan' location does not have an intensive care unit, so complex patients are operated at the 'Leyweg' location. Since the orthopaedic department is not assigned to OR capacity at the 'Leyweg' location, each complex patient is scheduled in an operating room of another specialism. Since the fraction of patient who undergo surgery at the 'Leyweg' is small (1%), we exclude surgery scheduling at the 'Leyweg' in this research.

Note that the number of surgeries and admissions is almost equal because nearly all admissions in the wards result from surgeries. The difference is caused by patients from other departments or patients who undergo multiple surgeries during their stay due to, for example, an infection.

Figure 2.1 shows the surgery categories in 2008 and 2009. These categories are based on diagnose codes that surgeons assign to surgery requests. Most surgery categories consist of multiple surgical procedure types. For example, the 'shoulder'-category includes a total shoulder prothesis but also a shoulder arthroscopy. On the contrary, the 'total hip', 'total knee' and 'scopic knee' are single surgical procedure types. Note that these surgical procedure types make up half of all the surgeries that are performed by the orthopaedic

Number of surgeons	9
OR days per week (normal: 40 weeks)	13
OR days per week (reduction: 12 weeks)	8
Number of beds	45
Number of surgeries	3,164
Number of admissions	3,142
Average length of stay	3.8 days

Table 2.1: Key figures orthopaedic department 2009 (source: SAP)

Available ORs	6
Available OR days per week	30
Office hours (5 days/week)	8.00 - 16.00
Number of elective surgeries	6,488 (90%)
Number of emergency surgeries	751 (10%)
Total number of surgeries	7,239

Table 2.2: Key figures OR department 'Sportlaan' 2009 (source: SAP)

department. A high percentage (27%) of the surgeries has no specific diagnose code. It could either be that no specific diagnose code exists (e.g. for feet surgeries), that the surgeon considers a surgery as a special case, or that he or she forgot to specify the diagnose code. The reason for the high percentage of surgeries with no specific diagnose code, is because most surgeries are not scheduled by their diagnose code. Instead, surgeries are scheduled by their surgery description which is a free text field in the database and therefore usually differs for each surgery.

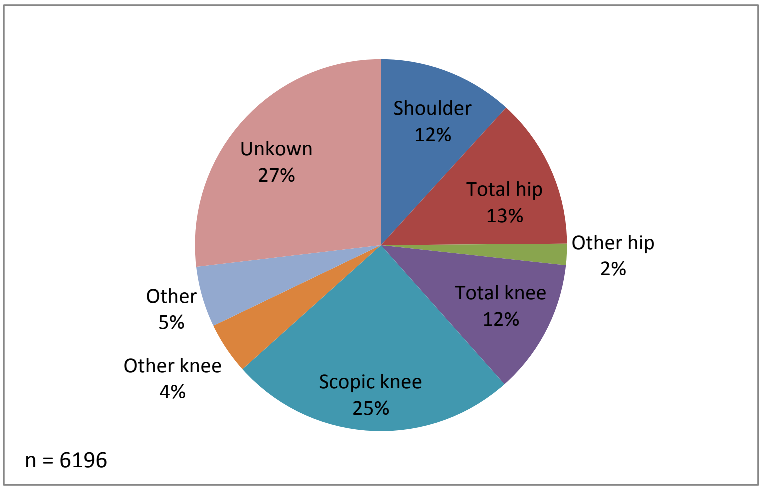

Figure 2.1: Surgery categories based on diagnoses in 2008 and 2009 (source: SAP)

2.1.2 OR department

Since the orthopaedic department is assigned to OR capacity at the 'Sportlaan', we only describe the OR department at this location. The other departments which are assigned to OR capacity at the 'Sportlaan' are gynaecology, ENT (ear-nose-throat), urology, and plastic surgery. Table 2.2 shows key figures of the OR department.

Figure 2.2 displays the layout of the OR department at the 'Sportlaan' location. Before a patient undergoes surgery, he or she is already hospitalised at the ward. Nurses transport the patient to the OR department where the patient arrives in the holding area. In this part of the OR department, the patient is prepared for surgery and sometimes the anaesthesia is already administered. Next, the patient is transported to the operation room itself where he or she undergoes surgery. Afterwards, the patient is transported to the recovery area where the patient is prepared to return to the ward. When patients are discharged from the OR department, they are hospitalised in the wards to recover.

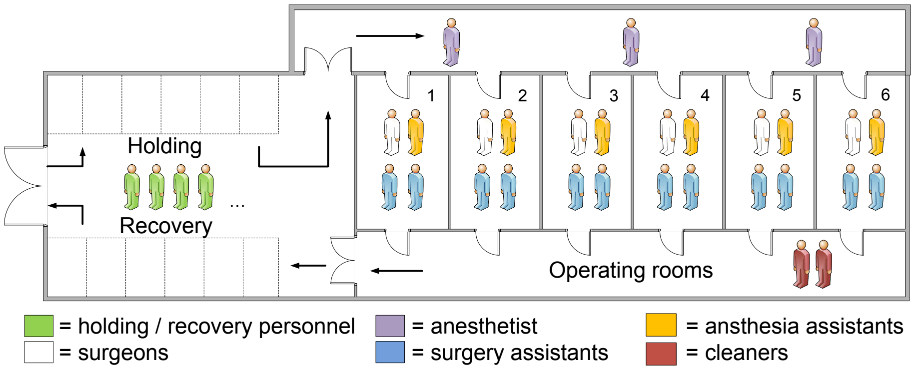

Figure 2.2: Layout of the OR department of the 'Sportlaan' location

We subsequently describe the people that play a role in the OR processes:

Holding and recovery personnel The holding and recovery personnel are nurses who are specialised in preparing the patient for surgery and recovering the patients afterwards. Since the holding and recovery are in the same room, its personnel performs these activities intertwined.

Anaesthetist An anaesthetist is a doctor who is employed by the hospital to administer anaesthetics to patients. An anaesthetist is responsible for the patients of two operating rooms but is not present in the operating room during surgery. Only when patients get unstable, the anaesthetist intervenes during a surgery. Administering the anaesthetics is either done in the holding or just before surgery in the operating room.

Anaesthesia assistant An anaesthesia assistant is someone who takes care of the patient during surgery. He or she will pick up a patient from the holding area and return him to the recovery. During the surgery, the anaesthesia assistant monitors the vital signs of the patient. When a patient gets unstable, the anaesthesia assistant has limited qualification to intervene and has to consult the anaesthetist. Each anaesthesia assistant is dedicated to one OR.

Surgery assistant Surgery assistants assist the surgeon during surgery. Each OR has two dedicated surgery assistants. During the surgery, one surgery assistant performs all the sterile activities while the other performs the non-sterile activities.

Surgeon The surgeon performs the surgery to cure the patient. Unlike the other people at the OR, he is not employed by the hospital but has a partnership with the other doctors in his specialism. The surgeon is sometimes assisted by a resident. Under the supervision of a surgeon, the resident may also carry out a surgery which may cause the surgery to last longer than planned.

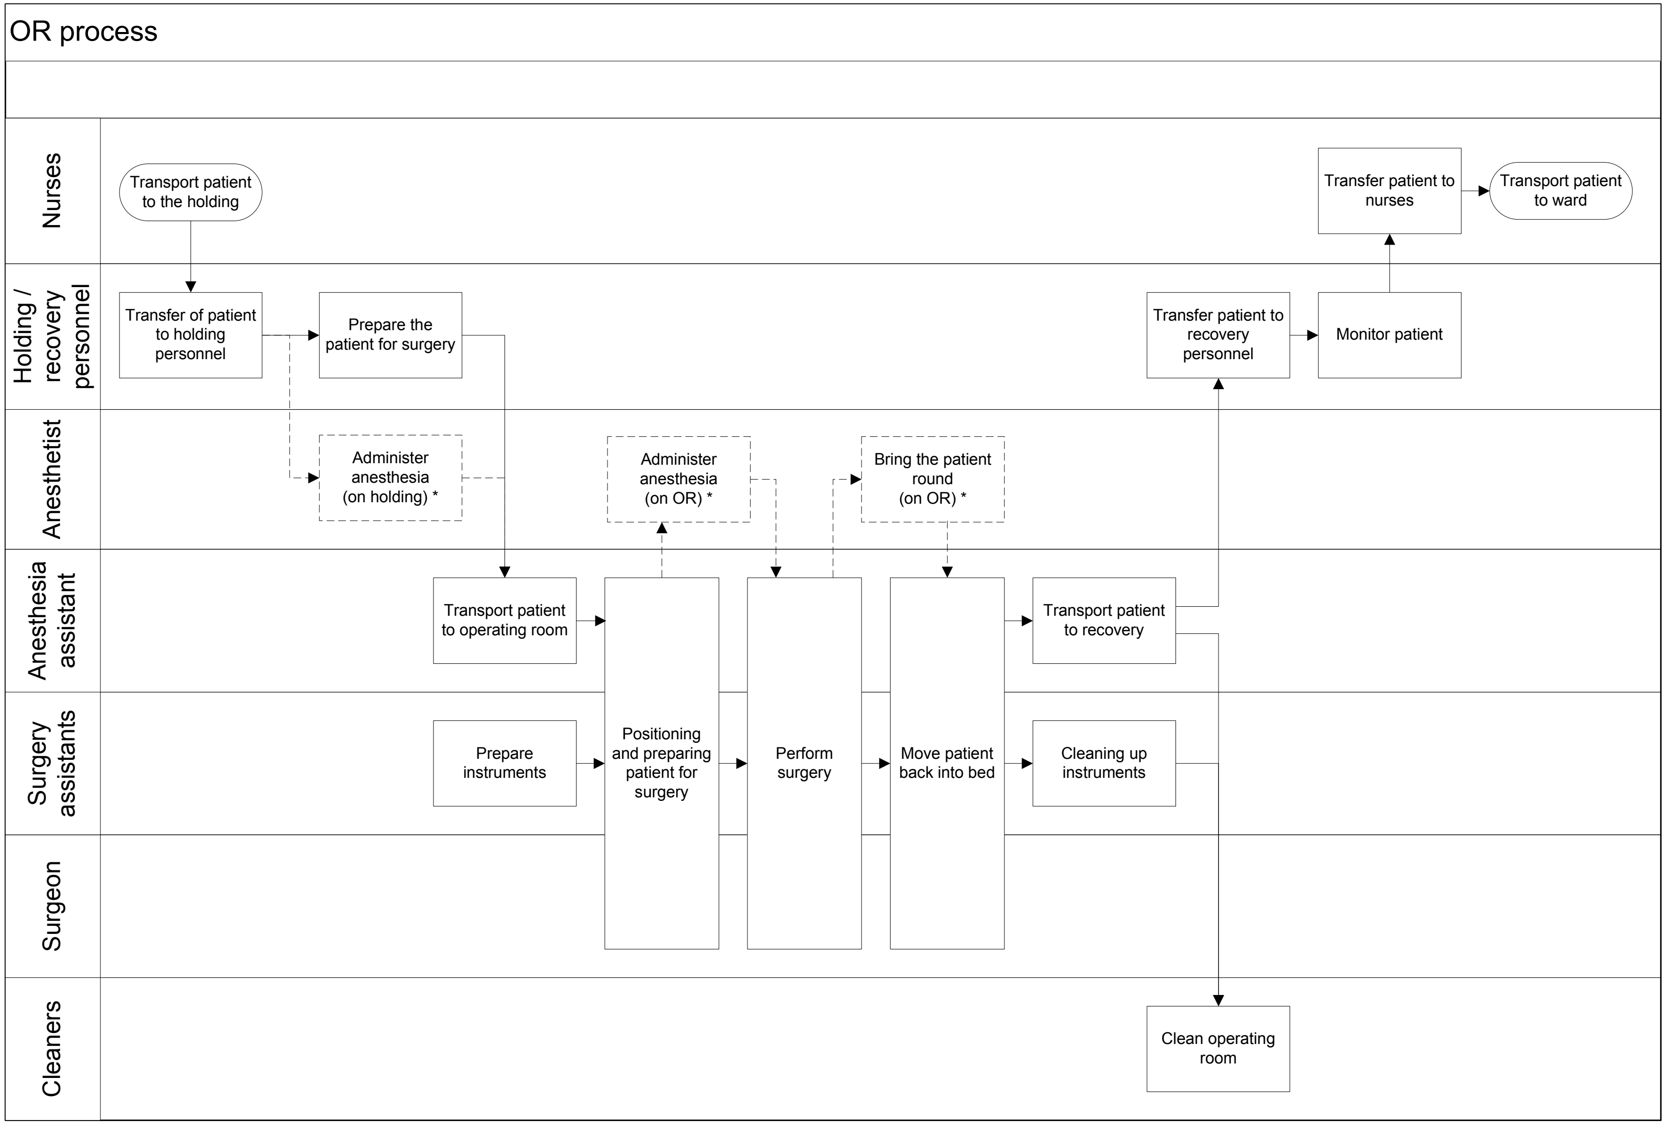

* : a patient is either treated on the holding or on the OR by the anesthetist

Figure 2.3: The general OR process of a single patient

Cleaners Two cleaners clean all the ORs during changeovers of surgeries.

Figure 2.3 represents the general OR process of a single patient. The process is visualised as a critical path of activities. This means that an activity can only start when all its preceding activities are finished. Note that there is a difference in the location where the anaesthesia is administered to the patient. Depending on the type of anaesthetics, the anaesthetist is either present at the holding or at the OR.

Next to the precedence relations for a single surgery, there are precedence relations between surgeries on the same OR. For example, the surgery assistants can only prepare the instruments for a surgery when the instruments from the previous surgery are cleaned up. Furthermore, there are time constraints such as administering anaesthetics not too far ahead from the start of surgery.

The time that is available to perform surgeries depends on the presence of OR personnel. The working shifts of anaesthesia and surgery assistants are displayed in Figure 2.4.

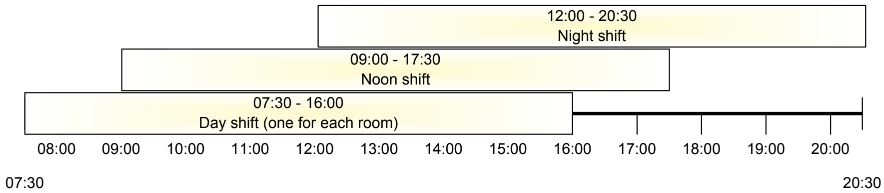

Figure 2.4: Working shift of anaesthesia and surgery assistants

Each shift consists of one anaesthesia and two surgery assistants. For each day, there is a day shift for each operating room and one noon and night shift. During the day, the noon shift replaces the day shift personnel for their coffee and lunch breaks. After 16.00, the noon shift assists with surgeries that are still running at that time. The night shift could also take care of surgeries after 16.00, but this shift is shared with the children's hospital. When multiple surgeries are taking place after 16.00, it could be necessary that the day shift works in overtime. The night shift ends at 20.30, but should be available during the night in case of an emergency surgery.

2.2 Control description

The flow of patients through the OR and the wards is controlled by how patients are scheduled. Figure 2.5 displays an abstraction of the surgery scheduling process. The starting point for the OR planners is the surgeon schedule where surgeons and management specify which surgeon performs surgery on which operating room and day. Based on the surgeon schedule, the OR planners assign the individual surgeries from the waiting list to OR days based on where a surgery fits first. Surgeries are scheduled based on the intervention time that is estimated by the surgeon. We define intervention time as the time between the incision and closing the wound. The OR planners interpret intervention times in terms of units, where one unit is half an hour. Each OR day is loaded with surgeries up to 12 units while the total available time per OR day is 16 units (8 hours). The remaining 4 units account for the changeover times, the variability of surgery durations, and add-on scheduling of urgent surgeries. Based on the intuition of the OR planners, the 4 units of slack is reduced or increased for some OR days.

Surgeons request a surgery by filling in a so-called 'admission form' in which they specify the patient, the type of surgery, the surgery time and the priority (maximum allowed access time). We distinguish four types of surgery requests: elective, semi-urgent, urgent and emergency surgeries. Elective surgery requests can be placed on a waiting list to be scheduled at a later point in time. Semi-urgent surgeries have a restriction on the access time, but the maximum access time is larger than two weeks. Elective and semi-urgent surgeries are scheduled two weeks in advance to inform the patients on time and to give the OR department the opportunity to check the feasibility of the schedule. Urgent surgeries should be performed within two weeks and are added to the elective schedule. Two days in advance, the surgery schedule is submitted to the OR department. Emergency surgeries should be performed within two days and are scheduled by the OR department.

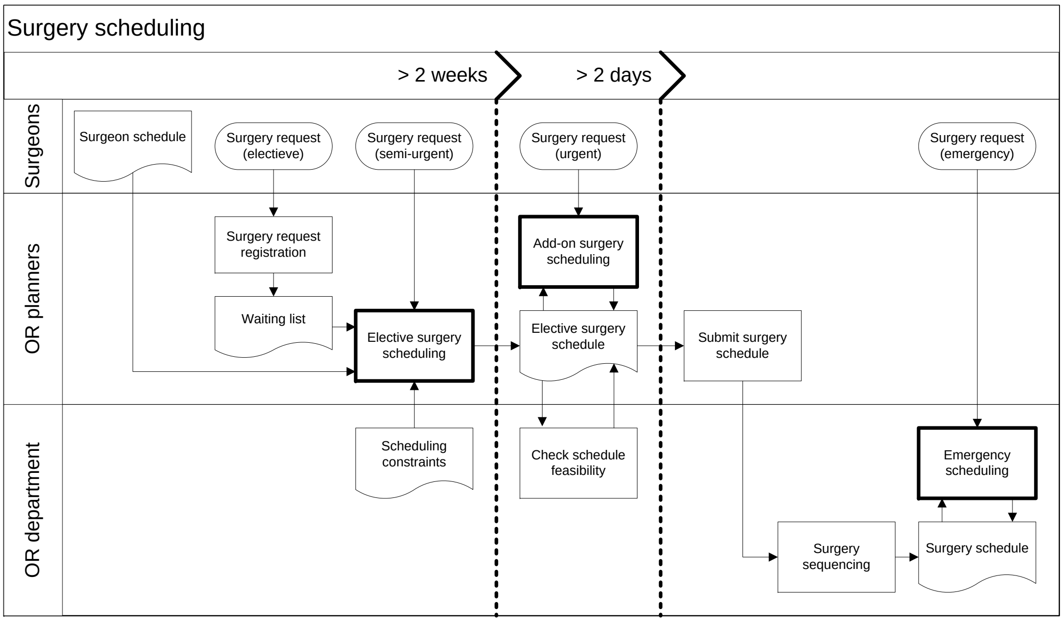

Figure 2.5: The surgery scheduling process

There are three main scheduling activities which are accentuated in Figure 2.5: elective, add-on and emergency scheduling. Note that only the elective and add-on scheduling activities are performed by the OR planners while the emergency scheduling is done by the OR department. In reality, there is no clear-cut distinction between the elective and add-on scheduling activities, but we model them separately to get an understanding of the various scheduling phases.

Figure 2.6 shows the percentages of surgery request types in 2010 until October to indicate the percentage of patients that are scheduled per scheduling phase. The figure shows that a high percentage of elective and semi-urgent patients (88%). This justifies developing a surgery scheduling approach where we only schedule elective and semi-urgent patients.

OR planners have to deal with several scheduling restrictions when scheduling surgeries. The scheduling restrictions are categorised as follows:

Surgeon restrictions For each day a surgeon is dedicated to an operation room as specified in the surgeon schedule. All surgeries that are scheduled on the same OR-day should therefore require the same surgeon. Each surgeon has his or her own speciality and thus performs a subset of all the surgery types. Furthermore, surgeons may have preferences in the surgery sequencing.

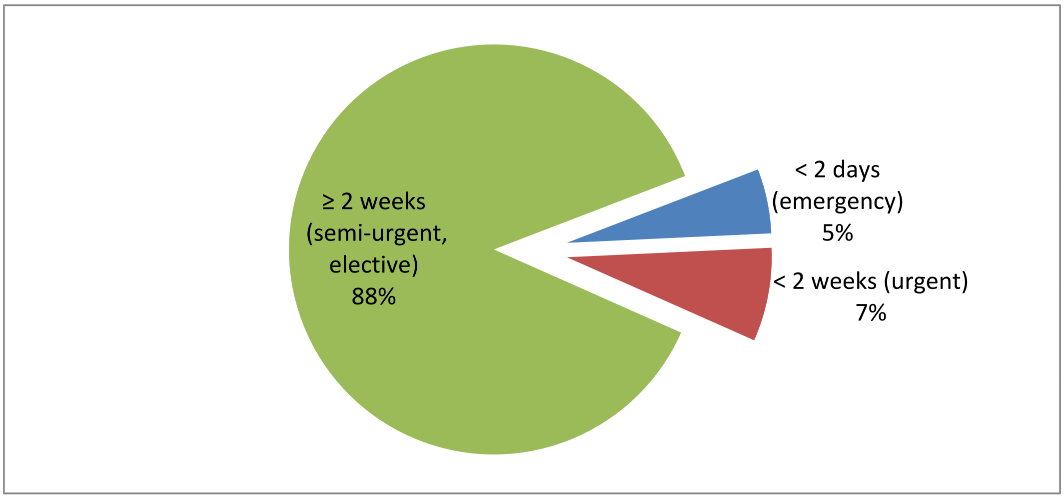

Figure 2.6: Surgery request types in January - October 2010 (source: SAP, $n = 2453$)

Instrument restrictions The number of surgeries of the same type on the same day is limited by the amount of instrument trays that are available for that type of surgery.

Holding and recovery restrictions When scheduling surgeries, the holding and recovery capacity should be taken into account. Too many surgeries with a short duration or many surgeries which end at the same time, could result in an excessive load of the holding and recovery area. This could result into delays of the changeovers between surgeries.

Personnel restrictions The time that is available to schedule surgeries is limited by the working hours of OR personnel. Also, the acceptability of overtime depends on the willingness of the personnel.

Through years of experience, the OR planners know how to avoid resource conflicts of surgeries on the OR department, but these surgery restrictions are not formalised. This causes inefficiencies in the amount of time spend on surgery scheduling because the OR planners and the OR department have to manually check the feasibility of every schedule they come up with.

2.3 Performance analysis

In this section, we introduce indicators to evaluate the performance of the system. We perform this analysis to identify the causes of poor performance. To structure the analysis, we explicitly state each cause we observe. First, we discuss patient-related indicators in Section 2.3.1. Since one of our research objectives is to improve the OR utilisation, we discuss the OR utilisation and identify the core problems of underutilisation in Section 2.3.2. Our second research objective is to reduce the required number of beds. Therefore, Section 2.3.3 presents indicators regarding the wards and determine the causes of a high bed requirement.

2.3.1 Patients

We introduce two indicators regarding patients: access times and the waiting list. Access time is most relevant for patients. The waiting list gives insight in whether there is a structural capacity problem.

Access time The bars in figure 2.7 represent the average access time of surgery categories in 2008, 2009 and 2010. We define the access time as the time between surgery request and the day of surgery. For 80% of the patients of each surgery category, the access time was less or equal than indicated by the line endings on top of the bars. Note that long access time are not always caused by the orthopaedic department but also by patients who postpone their surgery date.

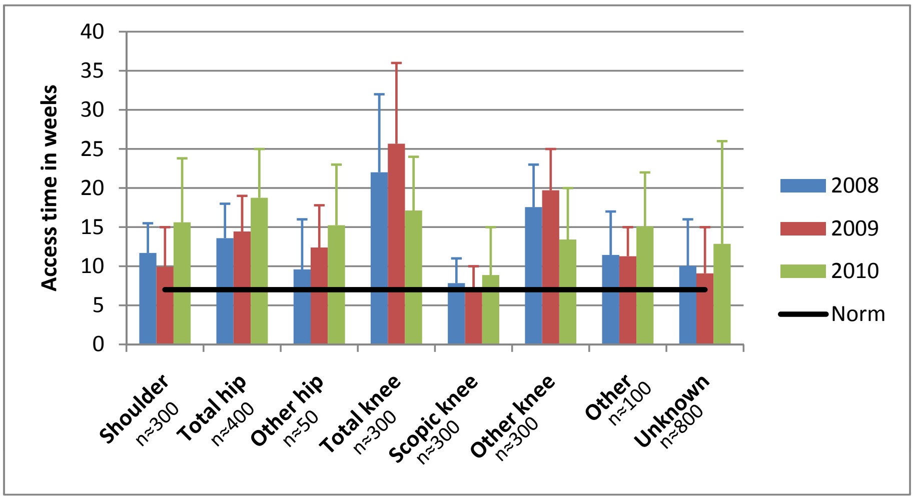

Figure 2.7: Average and 80% quantile of access times in 2008, 2009 and 2010 (source: SAP)

Representatives from several healthcare organisations in the Netherlands established standards about the social acceptability of access times. According to national access time standards (de treeknormen[1]), 80% of the patient population should have a maximum access time of 5 weeks and every patient should have a maximum access time of 7 weeks (displayed by the black line in 2.7). The access times of the orthopaedic department have a considerable deviation compared to the national standard. Figure 2.7 even shows that the surgery categories have at least twice the access time of the national standard. We do not try to reduce access times in our solution approach, since it is not our primary research goal. We assume, however, that access times will reduce indirectly by improving OR utilisation.

Waiting list Figure 2.8 displays the number of people on the waiting list from January 2008 - May 2010 . A patient is on the waiting list on a certain date when a surgery request has been done before or on that date and surgery takes place after or on that date.

The waiting list is stable, but it is likely that it has reached its maximum because access times are too long. Since the orthopaedic department is honest to its patients about the length of the access times, patients might turn to other hospitals although they would prefer HagaZiekenhuis. The hospital could make more money if they could also serve these patients. This requires lowering the access times by increasing the throughput of patients. The easiest way to do this, is to increase the capacity, but this is an ad-hoc solution. Instead, we deal with this problem in this research by making better use the available capacity (improving OR utilisation).

[1] http://www.treeknorm.nl

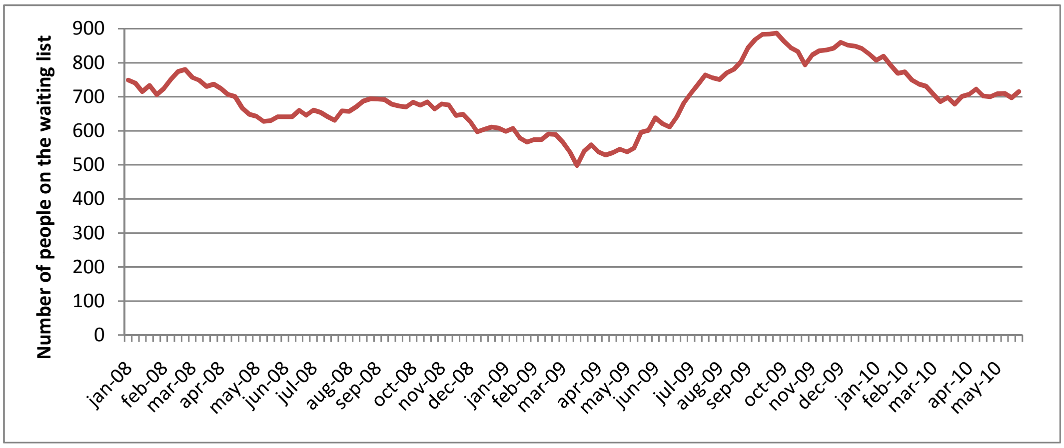

Figure 2.8: Number of people on the waiting list from Jan. 2008 - May 2010 (source: SAP)

Every week, approximately 55 patients undergo orthopaedic surgery. A waiting list of 500 patients (minimum in Figure 2.8) is thus equal to a backlog of 9 weeks. To improve the OR utilisation, it is convenient to have such a large waiting list, because it enables the OR planners to efficiently fill surgery schedules [13]. A large waiting list is thus beneficial for the hospital, but disadvantageous for the patient because it implies long access times. Since hospitals become more competitive and patients more demanding, this could damage the reputation of the HagaZiekenhuis on the long term. Therefore, the waiting lists should not be larger than is strictly necessary to schedule surgeries efficiently.

2.3.2 OR utilisation

The orthopaedic department is assigned to three operating rooms on Mondays, Tuesdays and Wednesdays and two operating rooms on Thursdays and Fridays. In total, the orthopaedic department has 13 OR-days at its disposal per week. During holiday periods, the orthopaedic department is assigned to 8 OR-days per week. Figure 2.9 shows the agreed and actual number of OR days per week for a year. More OR-days were used in some periods due to other specialisms who returned their OR time. Less OR-days were realised when resource problems in the OR department occurred. The adjustments in OR time can occur in as little as a week in advance. Since there is uncertainty about the availability of OR time, the OR planners are reluctant to schedule surgeries too far ahead. In practice, they apply a planning horizon of two or three weeks. When developing a tactical surgery scheduling approach, it is a prerequisite that the availability of OR capacity is known further ahead. Therefore, in this research, we assume that the availability of OR capacity is known when constructing a tactical surgery schedule.

Cause 1. *Uncertainty about the availability of OR capacity.*

Since we positioned our research on a tactical level, we consider the number of available OR-days to be fixed. We are interested in how well the available OR time is utilised. To this end, we introduce a definition that is known in literature as net OR utilisation.

$$\text{Net OR utilisation} = \frac{\sum \text{Session time during office hours (OR out - OR in)}}{\text{Number of OR days * 8 hours}}$$

Session time is defined as the time between a patient entering and leaving the operating

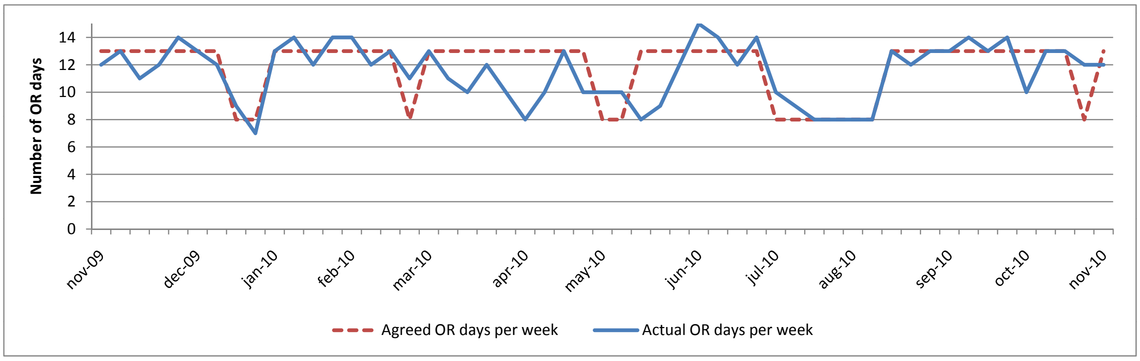

Figure 2.9: Agreed and actual number of OR days per week (source: SAP)

room (OR out - OR in). We only consider the session time during office hours (8:00 - 16:00). We include OR-days in the calculation if they meet the following requirements.

- The OR-day takes place on a working day.
- For all sessions, the time of entering and leaving the operating room is registered.
- At least two surgeries have a positive session time during office hours (8:00 - 16:00).
- All sessions are performed by the same speciality during office hours.
- Session times during office hours have no overlap (overlap: patients enters OR before the previous patient has left).

Figure 2.10 shows an example of an OR-day. In the example, multiple specialisms perform sessions which seems in conflict with the fourth requirement. Nevertheless, the OR day is valid because the gynaecology session is not performed during office hours. From the last session, that takes place (15.00 - 17.00), only the session time during office hours is included in the calculation (15:00 - 16:00 = 1 hour). The net OR utilisation of the example is:

$$\text{Net OR utilisation} = \frac{3 \text{ hours} + 2 \text{ hours} + 1 \text{ hours}}{8 \text{ hours}} = 75\%$$

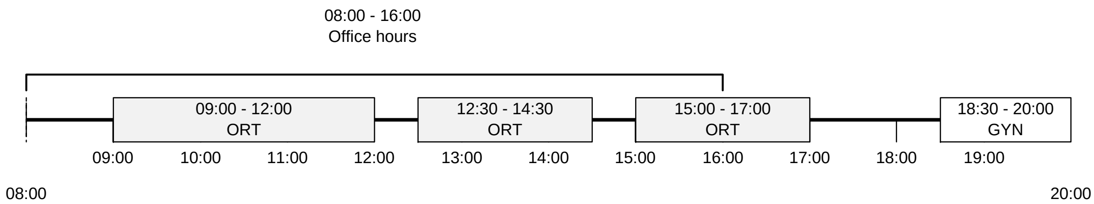

Figure 2.10: Example of an OR day for the net OR utilisation

Figure 2.11 shows the net OR utilisation for the orthopaedic department. On average, there are 50 OR days per month, but only 22 OR days meet the requirements to be included in the calculation. 20% of the OR-days are not used because the time of entering or leaving the operating room is not registered for one or more sessions. 35% of the OR-days are not used due to overlap of sessions. Missing times and overlap are caused by poor registration of session times. To attain a better calculation of the net OR utilisation in the future, session times should be recorded more accurate by OR personnel.

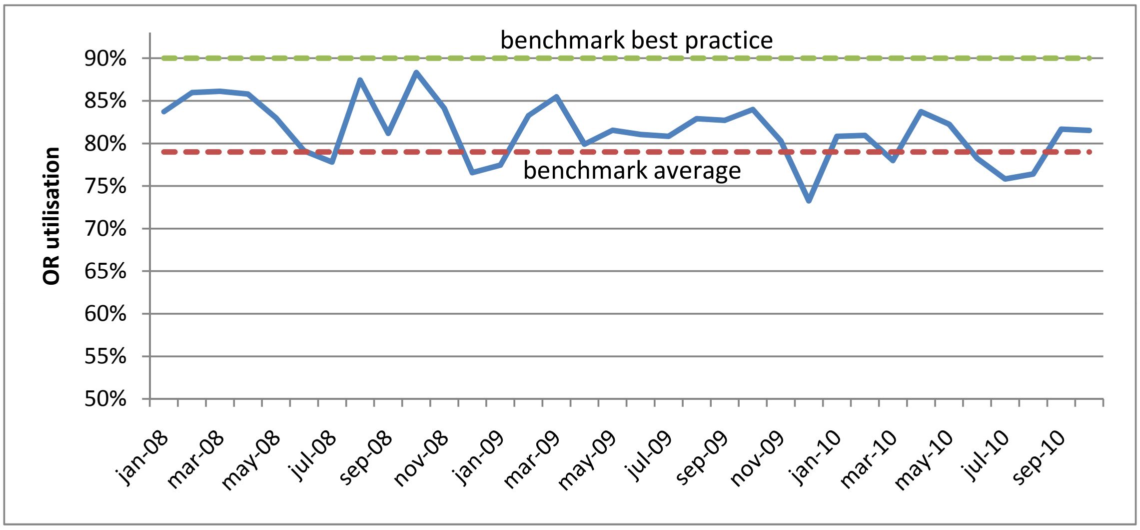

Figure 2.11: Net OR utilisation of the orthopaedic department, n ≈ 22 OR days per month (source: SAP)

Orthopaedic department			Plexus OR benchmark	
2008	2009	2010	average	best practice
82%	80%	78%	79%	90%

Table 2.3: Net OR utilisation (source: SAP, Plexus report)

Plexus, a healthcare consultancy company, performed a study on the OR department of HagaZiekenhuis in 2009. In the report, Plexus presents an OR benchmark based on 45 hospitals in the Netherlands. In Figure 2.11, the dotted lines represent the average and best practice net OR utilisation of the Plexus OR benchmark. Table 2.3 shows the yearly averages of the orthopaedic department. From the data we conclude that the orthopaedic department performs close to the average of the benchmark. On the other hand, the net OR utilisation per month is fluctuating and the yearly net OR utilisation is declining. Compared to the best practice, the orthopaedic department has the potential to improve the net OR utilisation of 2010 by 12%.

In the remainder of this section, we determine the problems that cause underutilisation of OR time due to poor scheduling. From a scheduling point of view, an OR-day consist of four types of periods: start of OR-day, interventions, changeovers and the end of an OR-day. Note that the OR planners schedule surgeries based on intervention times (time between the incision and closing the wound). Figure 2.12 visually represents the periods. We discuss each period in the subsequent paragraphs.

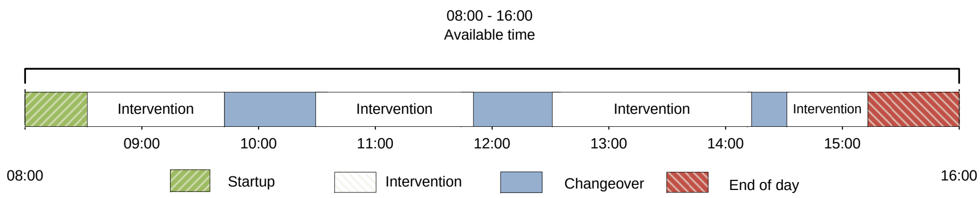

Figure 2.12: Example of an OR day

Start of OR day Figure 2.13 shows the distribution of OR-day starts. We make a distinction between the time the patient arrives at the operating room ('Patient in OR') and the time the incision takes place ('start of intervention'). Note that OR personnel shifts start at 7:30 (Figure 2.4). All operating rooms start at 8:00, but it is not clear whether the patient should be on the operating room, or the surgery should start at that time. 74% of the OR days starts on time, if we assume that a patient should be on the operating room at 8:00. Only 2% of the OR days starts on time if the incision should take place at 8:00.

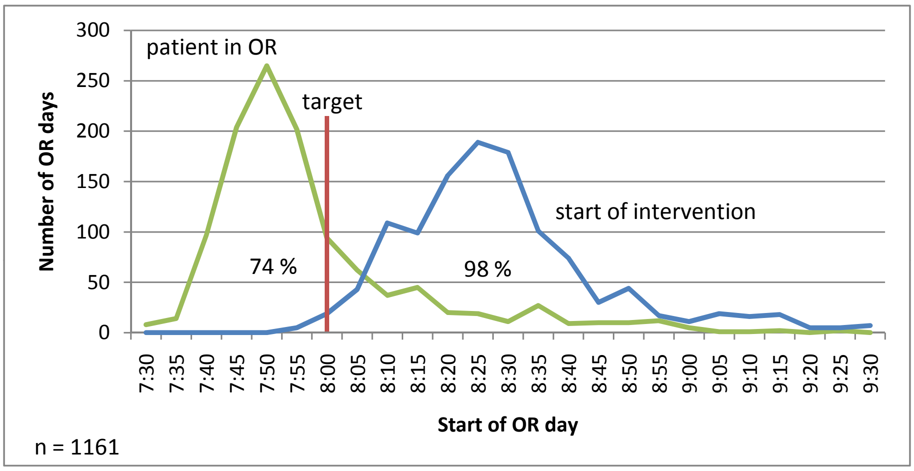

Figure 2.13: Arrival of first patient and start of first surgery in 2008 and 2009 (source: SAP)

Delays are caused by preparation activities (as shown in Figure 2.3) that do not finish on time. The anaesthetist is responsible for two operating rooms and administers the anaesthesia to the patients subsequently. Therefore, it is most likely that one of the two operating room starts late. It may decrease the session time of the first surgery when management would differentiate the starting times of the operating rooms. Apart from that, OR department management should define what they mean by the start of an OR-day. Both causes do not contribute to an improvement of OR utilisation or number of required beds, so we do not look into these causes in this research.

Cause 2. *Equal starting times of operating rooms.*

Cause 3. *Unclear definition of what activity marks the start of an OR-day.*

Intervention durations In Figure 2.14 we show the difference between planned and realised intervention times. Positive values indicate an overestimation of the intervention time (more time was planned than realised). We observe that the interventions times were overestimated for 73% of the surgeries and that the distribution has a high variation. According to Beliën, variability is caused by natural and artificial variability [1]. Natural variability is caused by the variability of the surgery process itself while artificial variability is caused by inaccurate estimations of intervention times. When scheduling surgeries, natural variability is unavoidable because it is caused from stochastic nature of surgery processes. The artificial variability can be prevented by estimating durations more accurately.

We conclude that the estimation of intervention times are inaccurate, because estimations are based on the perception of surgeons. Also, scheduling surgeries in terms of half

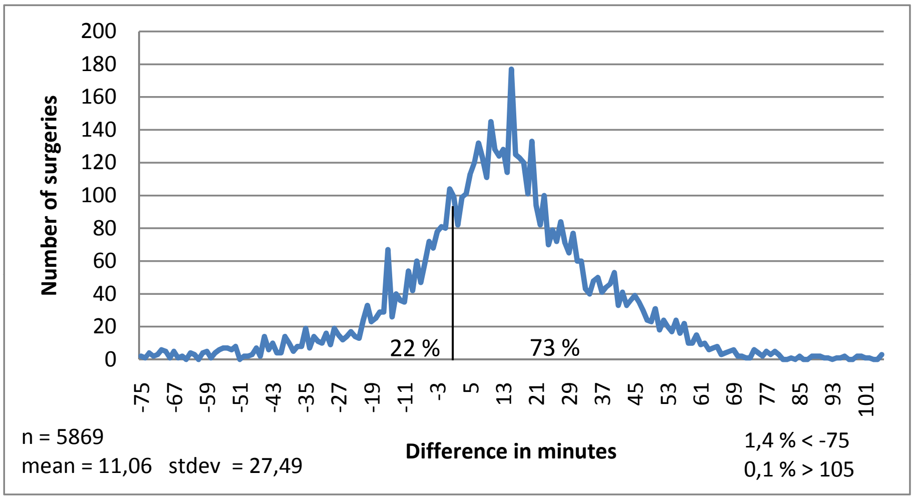

Figure 2.14: Planned - actual intervention times in 2008 and 2009 (source:SAP)

hours provides too little accuracy. Historical data and quantitative tools should be used to improve the estimation of intervention times.

Cause 4. *Inaccurate estimation of intervention times.*

Changeover times The distribution of changeover times is shown in Figure 2.15. The lines are drawn at 15, 30, 45 and 60 minutes to indicate the percentage of surgeries in each set. We define a changeover as the time between closing the wound of a patient and the incision of the next patient on the same operating room. The length of a changeover is defined by the critical path of activities (Figure 2.3) between two subsequent surgeries. It is striking that only 6% of the changeovers is less or equal than the ideal of 15 minutes that OR planners use in the scheduling process. Either the changeovers take too long, or the ideal of 15 minutes of changeover is not realistic. For this research we assume the latter is the case, because it is not our goal to find ways to reduce the changeover times, but how to anticipate on them while scheduling surgeries.

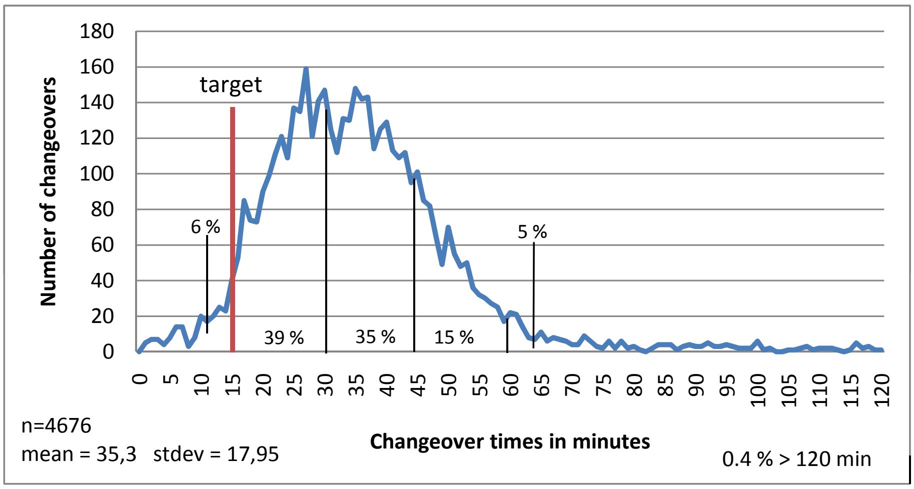

Figure 2.15: Changeover times in 2008 and 2009 (source: SAP)

Changeover times are believed to be 15 minutes between every surgery while on average the changeover time is 35 minutes. Historical data and quantitative tools should be used to improve the estimation of changeover times.

Cause 5. *Inaccurate estimation of changeover times.*

End of OR day Figure 2.16 shows the history of the end times of OR days. The end of an OR day is defined as the moment that the last patient leaves the operating room. As a rule, operating rooms should finish before 16:00 because the day shift ends at that time. The percentage of surgeries that ends after 16.00 is 36%. Since the noon and night shift end at a later time, no actual overtime occurs if no more than 2 operating rooms are still running after 16:00. Operating rooms end at 15.44 on average. This is an acceptable average because a later ending time would increase overtime. Sometimes surgeries are cancelled to prevent that OR personnel have to work in overtime. It is not clear from the data how many OR days ended on time by cancelling one or more surgeries.

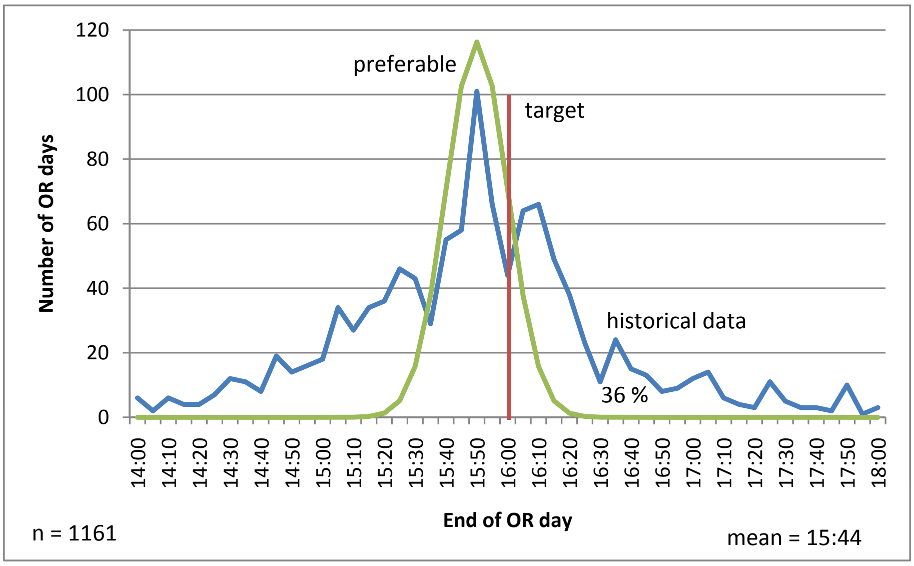

Figure 2.16: End times of OR days in 2008 and 2009 (source: SAP)

It is preferable that a higher percentage of the OR-days ends close to 16:00. Figure 2.16 displays this situation. Apart from inaccurate estimation of intervention and changeover times, an inaccurate estimation of an OR-day ending is caused by reserving slack time for each OR-day based on intuition of the OR planners instead of quantitative criteria. To attain a better estimation of OR-day endings, the variability of surgery durations should be taken into account in the surgery scheduling process.

Cause 6. *Inaccurate estimation of slack.*

2.3.3 Wards

We discuss two indicators concerning the wards: the required number of beds and the length of stay.

Required number of beds Figure 2.17 shows the bed occupancy in 2010. We used a boxplot to represent the data for each day of the week. The line endings indicate the minimum and maximum number of beds. The beginnings and endings of the boxes indicate

Weekdays	*95% quantile*
All days of the week	52
Saturdays	40
Sundays	35

Table 2.4: Required number of beds in 2010 (source: SAP)

the first and third quartile to express the variation in the data (the smaller the boxes, the smaller the variation). The crosses represent the medians. When a patient is discharged in the morning and another is hospitalised in the afternoon, a bed is occupied twice on a day. Therefore, it may seem that more beds are occupied than physically available (45 beds).

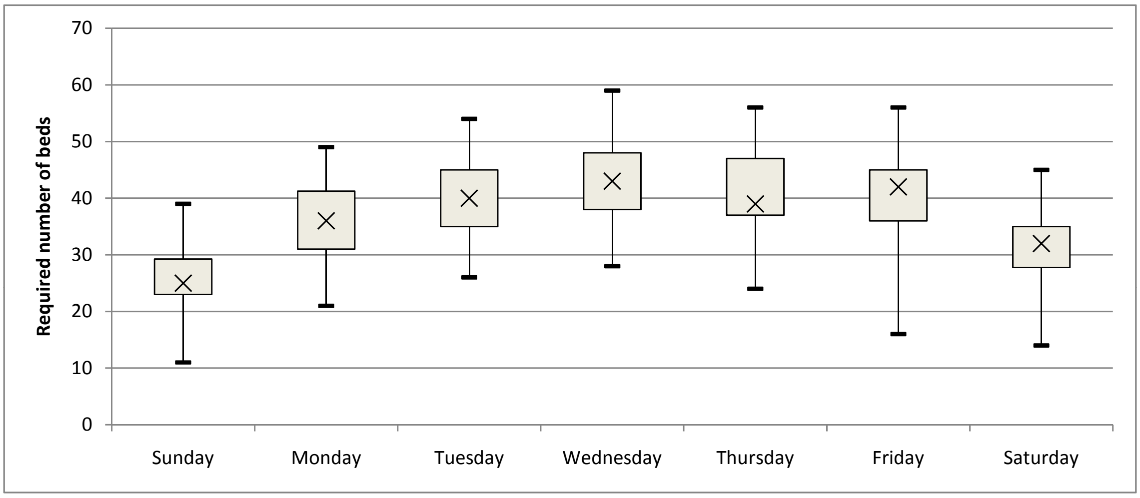

Figure 2.17: Bed occupancy for each day of the week in 2010 (source: SAP)

Based on a t-test (two-sided, $\alpha = 0.025$) we determine that the number of occupied beds on Sunday is significant lower than on other days of the week. On Saturday, there are significant more beds occupied than on Sundays, but significant less than on all other days of the week. On Monday, there are significant more beds occupied than on Sundays and Saturdays, but significant less than on all other days of the week. The number of beds on Tuesdays, Wednesdays, Thursdays and Fridays do not differ significantly from each other.

One of our research goals is to minimise the required number of beds. We make a distinction between the required number of beds for every day of the week, Saturdays and Sundays. Figure 2.18 shows the percentage of days that required at most a given number of beds. For example, 95% of the days required at most 52 beds on all days of the week. Table 2.4 shows the required number of beds at the 95% quantile which corresponds to the horizontal line in Figure 2.18. Variations in the number of occupied beds are caused by not considering the effects on ward capacity when scheduling surgeries. This is understandable because OR planners already have many factors to take into account when scheduling surgeries. Since the surgery restrictions are not explicit and surgeries are scheduled manually, more standardised and/or automated tools should be provided to the OR planners to take the bed occupancy into account and reduce the required number of beds.

Cause 7. *Scheduling without considering ward capacity.*

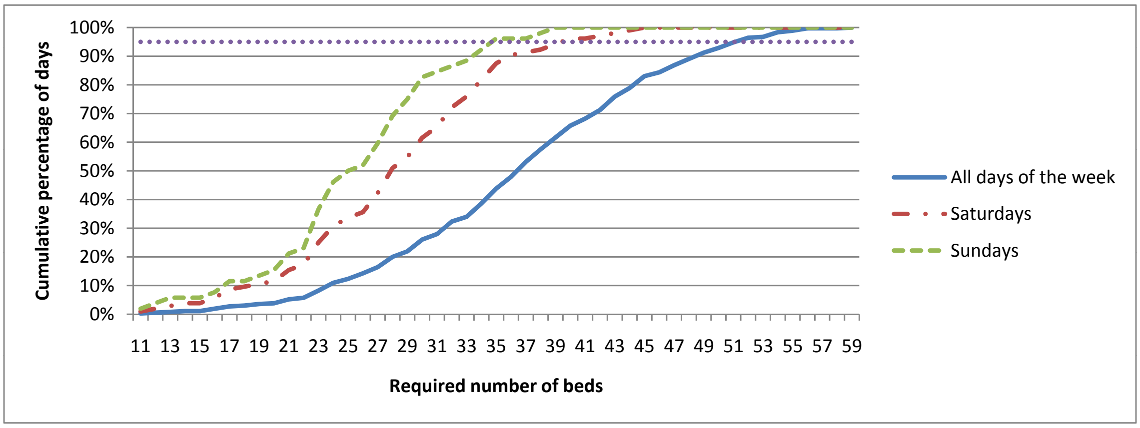

Figure 2.18: Cumulative percentage of required beds in 2010 (source: SAP)

Length of stay The length of stay is defined as the number of days that a patient is hospitalised in the ward. We also include outpatients (length of stay is 1 day). From Figure 2.1 we observe that three surgery types make up for half of the surgeries that are performed: total hip, total knee and knee scopic. The length of stay for these surgery types is shown in Figure 2.19. We grouped the remaining surgery types together in the 'other' category.

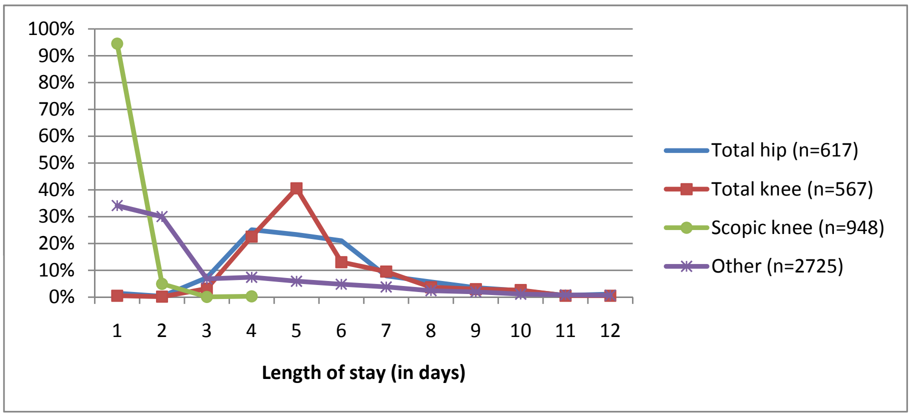

Figure 2.19: Length of stay of three major surgery types in 2009 and 2010 (source: SAP)

The length of stay for the patient population of all surgery types is shown in Figure 2.20. In this graph, the peak at 4 and 5 days is clearly caused by the length of stay of total hip and total knee surgeries. Since the beginning of 2010, about 10% of the day care treatments are performed at an external location. This causes the difference between the percentage of patients with a length of stay of one day in 2009 and 2010.

2.4 Conclusions

In this chapter, we proposed a case study on the orthopaedic department to develop a tactical surgery scheduling approach to maximises OR utilisation and minimise of the number of required beds in the wards. We presented key figures about size and case-mix of the orthopaedic department and described the processes in the OR department. Next,

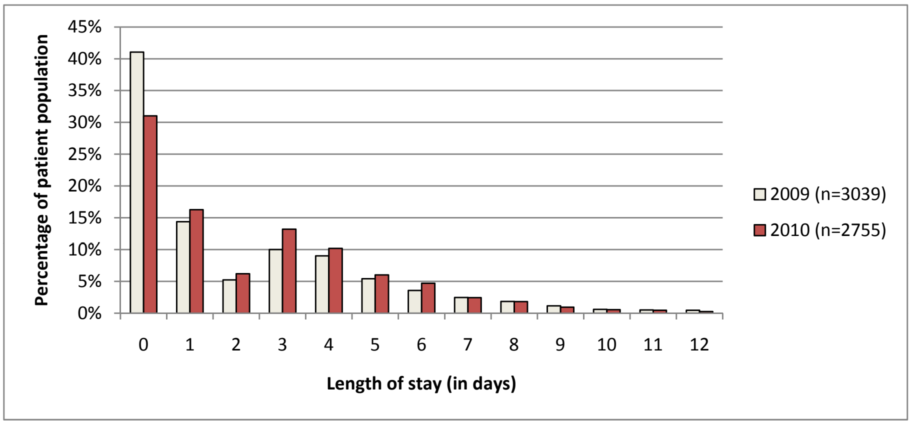

Figure 2.20: Length of stay for the total patient population in 2009 and 2010 (source: SAP)

we described how the system is controlled by explaining the surgery scheduling process. Based on a performance analysis, we identify seven causes of poor performance.

General

1. Uncertainty about the availability of OR capacity
2. Equal starting times of operating rooms
3. Unclear definition of what activity marks the start of an OR-day

OR utilisation

4. Inaccurate estimation of intervention times
5. Inaccurate estimation of changeover times
6. Inaccurate estimation of slack

Required number of beds

7. Scheduling without considering ward capacity

Cause 1 - 3 do not concern the OR utilisation or required number of beds directly. Therefore, we do not look into these causes in this research, but we discuss them in our recommendations (Chapter 6). In Chapter 3, we perform a literature review to find solutions for causes 4 - 6 to maximise OR utilisation. In addition, we search for literature on integrating bed levelling in surgery scheduling (cause 7) to minimise the required number of beds.

Chapter 3

Literature review

Operating room planning and scheduling is gaining more interest in the operation research community. In a recent paper, Cardoen et al. [4] present a literature review on this subject. The reader is referred to this review for a complete overview of the available literature concerning operating room planning and scheduling. We perform a literature review to find solutions for the causes for poor performance as identified in Chapter 2. In this chapter, we look into cause 4 - 7 since these causes are directly related to OR utilisation and the required number of beds.

In Section 3.1, we discuss the concept of master surgery scheduling. This approach is a solution to inaccurate estimation of slack and scheduling without considering ward capacity (cause 6 and 7 in Chapter 2). Section 3.2 discusses how we can obtain a better estimation of intervention durations and changeover times (cause 4 and 5 in Chapter 2). Section 3.3 describes the literature on how bed-levelling can be integrated with surgery scheduling (cause 7 in Chapter 2). We summarise this chapter in Section 3.4.

3.1 Master surgery schedule

A master surgery schedule (MSS) is often mentioned in surgery planning and scheduling literature as a way to optimise OR utilisation, level resource utilisation and create a robust schedule. An MSS is positioned on the framework of Hans et al. [7] (Figure 1.2) as resource capacity planning on a tactical level. There are several definitions of what to schedule in an MSS. Van Oostrum et al. [13] define an MSS as a cyclical schedule of recurrent surgery types, while Beliën and Demeulenmeester [2] define an MSS as a cyclical schedule of blocks of OR-time that are assigned to surgeons or specialisms. For the remainder of this report, we use the MSS definition of Van Oostrum et al. [13]. The cyclic execution of an MSS structures the workload of the OR department and the wards because they can better anticipate on future demand. The paper of Van Houdenhoven et al. [12] shows that the effectiveness of the MSS approach depends on the case-mix of surgery types. Fewer surgery types can be scheduled within an MSS when a high percentage of surgery types occurs infrequently. The percentage of surgeries that can be scheduled within an MSS also depends on how well the surgeries are clustered. In another of paper, Van Oostrum et al. [14] propose an approach to cluster surgeries.

A disadvantage of the MSS approach is that it has little flexibility because it assumes the same resource capacity for each period. The MSS can be updated to account for these changes in resource capacity, but frequent changes contradict with the cyclical nature of an MSS. Another disadvantage of the MSS approach of Van Oostrum is that it defines surgery types on the level of specialisms. This implies that all surgeons of some specialism can perform any kind of surgery type or that multiple surgeons perform surgeries on the same OR-day. In practice, only one surgeon is scheduled per OR-day and the surgeon may perform a subset of surgery types of its specialism due to specialisation.

3.2 Estimating surgery durations

In Chapter 2, we described that OR planners schedule surgeries based on intervention time (the time from incision until closing the wound). In the literature, however, there is a consensus of scheduling surgeries based on the surgery durations which we define as the time between a patient entering and leaving the operating room. From our performance analysis, we found that intervention and changeover times are estimated inaccurately. In accordance with literature, we suggest to schedule surgeries based on their surgery durations, since this reduces the problem of finding a way to accurately estimate the surgery durations.

Stepaniak et al. [11] show that the 3-parameter log-normal distribution is best suitable to model surgery durations. However, Van Oostrum et al. [13] and Hans et al. [8] assume normally distributed surgery durations. One of the advantages of assuming normal distributed surgery durations, is that the sum of surgery durations is also normally distributed. This enables us to calculate the probability of overtime in a relative easy manner. Calculating the probability of overtime is harder when we assume a 3-parameter log-normal distribution, because there is no exact expression for the sum of log-normal distributed random variables. In their paper, Hans et al. [8] purpose to schedule surgeries based on their average durations and reserve additional slack time to account for the uncertainty of surgery durations. The amount of slack on an OR-day depends on maximum probability of overtime and the variances of the surgeries on the OR-day. It appeared that, in total, less slack is required when surgeries with equal variances are performed on the same OR-day. In the financing world, this phenomenon is known as the portfolio effect. The portfolio effect could be exploited to improve the OR utilisation without increasing the probability of overtime.

3.3 Integrate bed-levelling into surgery scheduling

According to Cardoen et al. [4], only half of the papers on operating room planning and scheduling takes integration of upstream and downstream resources into account, such as wards. One of the papers that takes bed leveling into account, is the paper of Beliën and Demeulenmeester [2]. Their goal is to minimise the total expected bed shortage by determining on which days specialisms should be assigned to blocks of OR time. The length of stay is considered to be stochastic and is modelled by a multinominal distribution. Beliën and Demeulenmeester purpose a mixed integer program and meta-heuristic approach to solve their mathematical program. Van Oostrum et al. [15] both tries to maximise the OR utilisation and level the bed occupancy. Their approach comprises two phases. In the first phase, OR utilisation is maximised by generating single OR schedules. This is done by using a column generation technique. In phase two, the bed occupation is levelled by assigning the OR day schedules to specific OR-days. In the model, the surgery durations are considered to be stochastic, but the length of stay is deterministic. An paper by Vanberkel et al. [16] introduces a exact approach to determine the probability distribution of bed occupancy given a master surgical schedule. This model does not optimise the MSS, but serves as an evaluation tool.

3.4 Conclusions

A master surgery scheduling approach can be helpful to maximise OR utilisation and minimise the number of required beds. The approach has some drawbacks like limited flexibility and the assumption that every surgeon can perform every surgery type of a specialism. Instead of scheduling surgeries based on their intervention time, surgeries should be scheduled based on the surgery duration. A 3-parameter log-normal distribution is most appropriate to model surgery durations, but a normal distribution is also used in the literature. Van Oostrum et al. [15] propose a surgery scheduling approach in which they combine the maximisation of OR utilisation and minimisation of the number of required beds. One drawback of the model is that is assumes a deterministic length of stay. The model of Vanberkel et al. [16] can be used to model stochastic length of stay to determine bed occupancy.

Chapter 4

Solution approach

In this chapter, we develop a solution approach to overcome the causes of poor performance that we identified in Chapter 2. We structure this chapter based on the steps to perform a simulation study as proposed by Law [10]. Although we do not perform a simulation study in this research, we believe the framework of Law is also useful in the context of our research. Figure 4.1 visualises the framework. The problem formulation and literature review has already been discussed in Chapter 2 and 3. Based on the causes of poor performance, we construct a conceptual model in Section 4.1. In Section 4.2, we describe how we gather the data that we require for our model. We formulate our technical model in Section 4.3 and validate it in Section 4.4. We summarise this chapter in Section 4.5. In Chapter 5, we perform numerical experiments to determine how the model performs under different circumstances. Tips for implementation are part of Chapter 6.

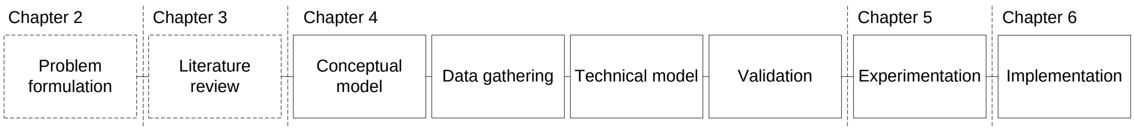

Figure 4.1: Steps to construct a solution approach

4.1 Conceptual model

Based on the framework of Hans et al. [7], presented in Chapter 1 (Figure 1.2), we distinguish a strategic, tactical, and operational level of surgery scheduling. The strategical level comprises of case-mix planning where blocks of OR time are assigned to specialisms. Case-mix planning is based on the production agreements of specialisms with hospital management. On a tactical level, specific surgeons are assigned to the OR blocks that are available for the specialism (surgeon scheduling). Only the availability of the surgeons is a factor in determining the surgeon schedule. On an operational level, case-scheduling is performed: individual patients are assigned to the OR blocks of their surgeon. At this level, the OR planners try to come up with surgery schedules that yield a high OR utilisation, satisfy material constraints and is in accordance with patient availability. The left part of Figure 4.2 visualises the current scheduling approach.

In order to improve the control over OR utilisation and bed occupancy, we propose to apply a master surgery scheduling (MSS) approach based on the research of Van Oostrum et al. [15]. This means that, in addition to surgeon scheduling, surgery types are scheduled on a tactical level. Surgery types are defined as a group of surgical procedures with the same resource demand. On the operational level, individual cases are assigned to the slots of their corresponding surgery type. When constructing an MSS, we take into account surgeon availability, OR utilisation, bed occupancy, material constraints and waiting lists. The right part of Figure 4.2 visualises the proposed surgery scheduling approach.

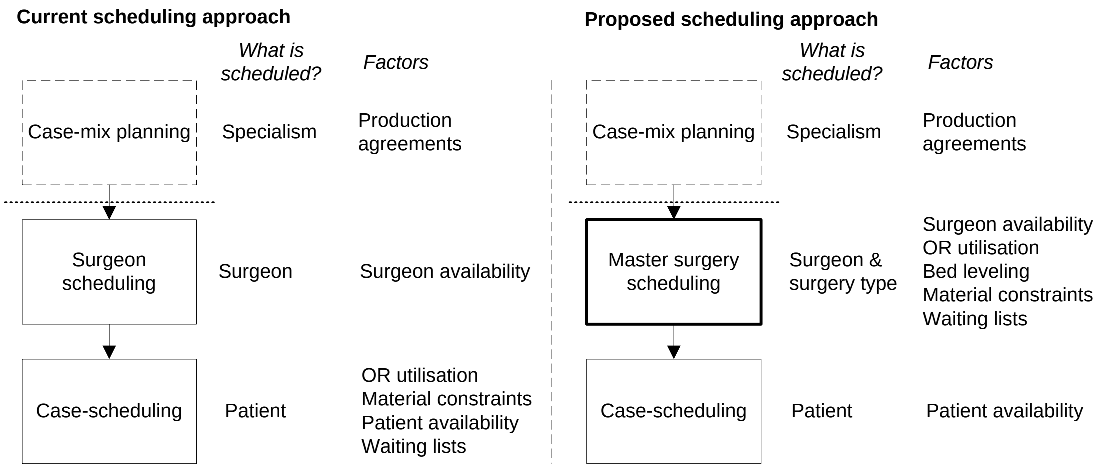

Figure 4.2: Comparison of current and proposed scheduling approach

We assume the surgery durations and length of stay of surgery types to be stochastic. We determine bed occupancy based on stochastic length of stay by using the model of Vanberkel et al. [16]. Management requires that only one surgeon performs all surgeries on an OR-day. In the literature on MSS approaches, the MSS schedule is cyclic. For our research, we propose the MSS to be non-cyclic because of the high level of specialisation of surgeons. The high level of specialisation means that each surgeon performs a small subset of all surgery types. Therefore, the surgery schedule highly depends on the availability of surgeons. For example, suppose a surgeon is absent on a day that he was supposed to perform 6 day treatments according to the MSS. Another surgeon takes over this OR-day and performs 4 surgeries which all have a length of stay of 4 days. This may be necessary, because the surgeon does not perform day treatments due to his specialisation. The adjustment has a strong effect on the bed occupancy: two beds stay empty on the OR-day itself while 4 extra beds are occupied for the next 4 days. We prevent such sub-optimal schedules by constructing schedules based on the availability of surgeons. This requires a non-cyclic MSS approach. An additional advantage of a non-cyclic MSS approach is that management has more flexibility in adjusting their production targets during the year.

In our surgery scheduling approach, we want to optimise the OR utilisation while the model of Van Oostrum minimises the required OR capacity. As a result, we do not schedule a fixed number of surgeries per surgery type (as in the Van Oostrum model). Instead, we let the number of scheduled surgeries per surgery type be variable within a predefined range. This implies that waiting lists are high enough to have the flexibility to schedule a variable number of surgery types in each planning period.

An MSS can be constructed when the availability of OR-days and surgeons is known. Management believes that this information is known 3 months in advance. Therefore, we propose to construct a surgery schedule with a planning horizon of 4 weeks that is determined 3 months in advance. Figure 4.3 shows an example where we construct a MSS for March at the beginning of January.

Our non-cyclic MSS approach addresses the problem of inaccurate estimation of surgery and changeover times (cause 4 and 5, Chapter 2) by scheduling surgery types based on

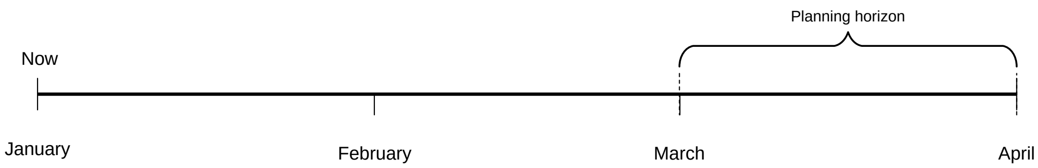

Figure 4.3: Planning horizon for a MSS

their historical surgery duration. As defined before, the surgery duration is the time between a patient entering and leaving the operating room. By determining the amount of surgeries of each surgery type that should be performed on an OR-day, we improve the estimation of slack (cause 6, Chapter 2). Also, cause 7 of Chapter 2 is addressed because we anticipate on the bed occupancy in the MSS.

4.2 Data gathering

We gathered our data from interviews with management, surgeons and OR personnel and from the hospitals information system (SAP). From SAP, we use the OR and ward data of all the patients that were operated in 2010. The definition of most sets is an obvious choice, for example surgeons, operating rooms and days in the planning horizon. Defining surgery types is more troublesome, since SAP contains an excessive amount of surgical procedures types. We define a surgical procedure type as the description of the surgical procedure on the lowest level of registration in SAP. For example, our dataset contains 2073 patients and 555 'surgeon-surgical procedure type' combinations. These surgical procedure types have to be clustered to attain flexibility in the case-scheduling phase. When we do not cluster procedure types, there are slots in the master surgery schedule with a limited amount of patients who can be assigned to them. This makes it difficult for the OR planners to assign patients to an appropriate slot. Suppose that there is only one patient for a certain procedure. This implies that there is only one slot in the MSS to which the patient can be assigned. If the patient is not available for surgery for that slot, the patient has to be assigned to the slot of another procedure type. To prevent this from happening, we cluster the procedure types in order to attain a higher patient population per surgery type slot.

From interviews with surgeons, we clustered 555 surgical procedure types downto 43 surgery types. For one of the scenarios in our numerical experiments (Chapter 5), we apply a clustering procedure based on the research of Van Oostrum et al. [14]. We describe the clustering procedure in Appendix C.

We describe the default data set in Appendix B. In the remainder of this report, we refer to the default data set as the *default scenario*.

4.3 Technical model

We first describe a base model in Subsection 4.3.1. The base model can not be implemented as an ILP model since we can not linearise the model of Vanberkel et al. [16] to determine bed occupancy. We decompose the model based on the approach of Van Oostrum et al. [15]. In accordance with this research, we introduce the concept of *operating room day schedules* (ORDS) which is schedule for a single OR-day. We discuss the first phase of our model decomposition in Subsection 4.3.2. In this phase, we generate a set of ORDSs that maximise OR utilisation. We discuss the second phase in Subsection 4.3.3. This phase comprises the assignment of ORDSs to OR-days to optimise the required number of beds. Figure 4.4 visualises the decomposition approach. In Subsection 4.3.4, we show how we programmed the model into software tools.

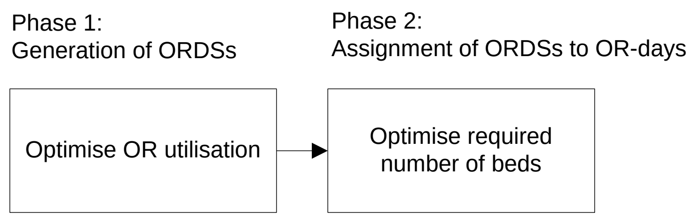

Figure 4.4: Representation of the decomposition approach

4.3.1 Base model

We define T as the set of days in the planning horizon. Set J is the set of identical operating rooms that are available during the planning horizon. An OR-day is a combination of an operating room $j \in J$ and a day $t \in T$ and is represented by the tuple (j, t). The surgery types are denoted by the set I. For modelling purposes, we require that each surgery type is surgeon specific. The purpose of the MSS model is to indicate how many surgeries of each surgery type should be performed on which OR-day. Let V_{ijt} be a decision variable that indicates the number of surgeries of type $i \in I$ on OR-day (j, t). The set of surgeons is given by set S. Let W_{sjt} be a binary decision variable indicating whether surgeon $s \in S$ performs surgery on OR-day (j, t).

In accordance with our research goals, the goal of our model is to maximise the OR utilisation and minimise the number of required beds. OR utilisation is defined as the ratio between utilised and available time. Since the amount of available time is determined in the case-mix planning and therefore is a constant, maximising the OR utilisation is equal to maximising the utilised time. As defined in Section 2.3.2, utilised time is the sum of session times. In accordance with the research of Van Oostrum et al. [15], we assume session times (surgery durations) to be normally distributed with expectation μ_i and variation σ_i^2 (). Since OR capacity is the most expensive resource, our primary objective (4.1) is to maximise the sum of the average surgery durations. Our secondary objective (4.2) is to minimise the number of required beds γ. It is most likely that maximising expected session times will yield solutions that prefers surgeries with a low variation. To impose that surgeries with a high variation have an equal probability of being scheduled, we include the sum of standard deviations in the objective function. We include the weights θ^u and θ^v in the objective function to balance between the sum of averages and the sum of standard deviations of session times.

$$\max \quad \sum_{i \in I} \sum_{j \in J} \sum_{t \in T} (\theta^u \mu_i + \theta^v \sigma_i) V_{ijt} \tag{4.1}$$

$$\min \quad \gamma \tag{4.2}$$

Let A_{jt} be a binary parameter indicating whether OR-day (j, t) is available. As said before, we assume that only one surgeon can be scheduled on an available OR-day:

$$\sum_{s \in S} W_{sjt} \leq A_{jt} \qquad \forall j, t \tag{4.3}$$

Let B_{st} be a binary parameter indicating whether surgeon s is available on day t. A surgeon can only perform surgery if he or she is available that day:

$$\sum_{j \in J} W_{sjt} \leq B_{st} \qquad \forall s, t \tag{4.4}$$

Let E_{is} be a binary parameter indicating whether surgery type i is performed by surgeon s. Since surgery types are surgeon specific, $\sum_s E_{is} = 1 \ \forall i$ by definition. The number of surgeries of type i on OR-day (j,t) can only be positive ($V_{ijt} > 0$) when the surgeon that performs surgery type i is scheduled on that OR-day, i.e. when $W_{sjt} = 1$ and $E_{is} = 1$. In other words: all surgery types on OR-day (j,t) should be performed by the same surgeon. Let M_i be the maximum number of surgeries of type i that can be performed on one OR-day. This results in the following constraint:

$$V_{ijt} \leq \sum_{s \in S} E_{is} W_{sjt} M_i \qquad \forall i, j, t \tag{4.5}$$

Due to financial reasons or restrictions on access times and waiting lists, every surgery type has a minimum ($D_i^{\min}$) and maximum ($D_i^{\max}$) number of surgeries that should be performed in the planning horizon.

$$D_i^{\min} \leq \sum_{j \in J} \sum_{t \in T} V_{ijt} \leq D_i^{\max} \qquad \forall i \tag{4.6}$$

We define R as the set of instruments used during surgeries. Let ω_{ir} be the number of instruments r that are needed for surgery type i. The cleaning of instruments requires one day, so there is a limited capacity of instruments per day. Let $Q_r^{\max}$ be the maximum number of instrument r that can be used on one day.

$$\sum_{i \in I} \sum_{j \in J} V_{ijt} \omega_{ir} \leq Q_r^{\max} \qquad \forall t, r \tag{4.7}$$

For each OR-day (j,t), C is the total available time to perform surgeries. Let $f_{jt}(V)$ denote the probability distribution of the total session time of the surgery types that are scheduled on OR-day (j,t), where V is the vector of all variables V_{ijt}. To assure that the probability that overtime occurs is at most α, we introduce the following constraint:

$$\mathrm{P}(f_{jt}(V) \leq C) \geq 1 - \alpha \qquad \forall j, t \tag{4.8}$$

Similarly, we introduce a function $g_t(V, U, L_i)$ to model the probability distribution of number of occupied beds on day t. The vector U represents the probability distribution of the number of occupied beds on day t resulting from previous planning periods. In addition, we need information about the length of stay of surgery types which we represent by the stochastic variable L. The function is used to make sure that less or equal than γ beds are needed on day t with a probability of at least β.

$$\mathrm{P}(g_t(U, V, L) \leq \gamma) \geq \beta \qquad \forall t \tag{4.9}$$

Appendix A summarises the complete base model.

4.3.2 Phase 1: Generating ORDSs

The primal model selects those ORDSs that maximise OR utilisation. It implicitly assumes that the optimal set of ORDSs is known. To generate an optimal set of ORDS, we use the column generation technique. Solving the primal model yields shadow prices which can be used to solve the pricing model to generate new ORDSs. The new ORDSs are added to the primal model and is solved again. In this manner, new ORDSs keep being added to the primal model until no new ORDSs are generated.

In addition to the notation of the base model, we introduce the set K as the set of ORDSs that are generated in the first modelling phase. Let X_k be a decision variable that represents the number of times that ORDS $k \in K$ is performed in the planning horizon. Let O_{ik} be the number of surgeries of type i in ORDS k. This parameter defines all ORDSs. We formulate the mathematical program as follows:

Primal model

$$\max \quad \sum_{i \in I} \sum_{k \in K} (\theta^u \mu_i + \theta^v \sigma_i) X_k O_{ik} \tag{4.10}$$

subjected to

$$\sum_{k \in K} X_k \leq \sum_{j \in J} \sum_{t \in T} A_{jt} \tag{4.11}$$

$$D_i^{\min} \leq \sum_{k \in K} X_k O_{ik} \leq D_i^{\max} \qquad \forall i \tag{4.12}$$

$$X_k \geq 0 \qquad \forall k$$

The objective function (4.10) is the equivalent of the primary objective (4.1) in the base model. Constraint (4.11) states that the total number of ORDSs should be less or equal to the total number of OR-days that are available. The range constraint (4.12) is a reformulation of constraint (4.6) in the base model.

Since all the surgeries of an ORDS have to be performed by the same surgeon, we generate new ORDSs for each surgeon in each column generation iteration. Let s^* denote the surgeon to generate an ORDS for. We define the set I_s as the subset of surgeries that are performed by surgeon $s \in S$. Let Y_i be a decision variable that indicates how many surgeries of surgery type $i \in I_{s^*}$ are included in the new ORDS. In order to generate an ORDS, we need to linearise the overtime constraint (4.8). In Appendix D, we discuss how to linearise the overtime constraint. We show that the overtime constraint can be simplified to a form where a square root occurs in the overtime constraint. Since a square root function is not linear, we propose to approximate the square root function by means of piecewise linear functions. The beginnings and endings of each interval of the piecewise functions are denoted by a set of $N = \{0, 1, \ldots, m\}$ breakpoints. Let x_n be the breakpoint value and y_n the breakpoint function value. Any point in between two breakpoints is a weighted sum of these two breakpoints. Let ρ_n denote the non-negative weights such that their sum is 1. Figure 4.5 show an example of approximating the square root function for the interval [0, 90]. There are three breakpoints ($N = \{0, 1, 2\}$) with coordinates $(0, 0.8)$, $(10, 4)$ and $(90, 10.8)$. Suppose that we want to know the linear approximation of $\sqrt{80} = 8.9$. This point lies in between breakpoint 1 and 2. Setting $\rho_0 = 0$, $\rho_1 = \frac{1}{8}$ and $\rho_2 = \frac{7}{8}$ results in $\sum_n \rho_n x_n = 80$ and $\sum_n \rho_n y_n = 9.5$. In this example, the linear approximation of $\sqrt{80} \approx 9.5$.

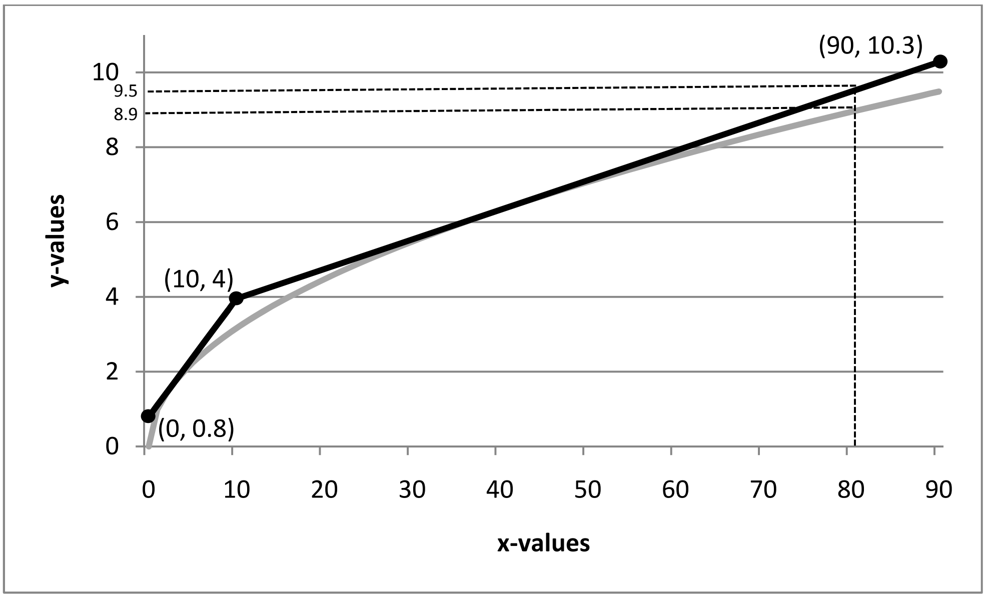

Figure 4.5: Approximation of the square root function by two piecewise linear functions

Let $\pi_i^{\min}$ and $\pi_i^{\max}$ be the shadow prices of the range constraint (4.12). In addition, we introduce constraint (4.14) to make sure that the instrument capacity per day is not violated. This constraint is derived from constraint (4.7) in the base model. We formulate the model to generate an ORDS as follows:

Pricing model

$$\max \quad \sum_{i \in I_{s*}} (\theta^u \mu_i + \theta^v \sigma_i - \pi_i^{\min} - \pi_i^{\max}) Y_i \tag{4.13}$$

subjected to

$$\sum_{i \in I_{s*}} Y_i \omega_{ir} \leq Q_r^{\max} \qquad \forall r \tag{4.14}$$

$$\sum_{i \in I_{s*}} Y_i \mu_i + \Phi^{-1}(\alpha) \sum_{n \in N} \rho_n y_n \leq C \tag{4.15}$$

$$\sum_{n \in N} \rho_n x_n = \sum_{i \in I_{s*}} Y_i \sigma_i^2 \tag{4.16}$$

$$\sum_{n \in N} \rho_n = 1 \tag{4.17}$$

$$Y_i \in \mathbb{N} \qquad \forall i \in I_{s*}$$

The linearisation approach in Appendix D requires us to define a maximum approximation error for the slack in each ORDS. We decide upon an approximation error of at most 1 minute per ORDS. This implies that the approximated slack is at most 1 minute higher than the actual slack of an ORDS.

When the minimum demand for a surgery type is positive for at least one surgery type, i.e. when $\exists i \; D_i^{\min} > 0$, the primal model is infeasible when we run the model for the first time. Therefore, we need to generate an initial set of ORDSs that meet the demand constraint. To this end, we use a LPT-heuristic to generate an initial set of ORDSs that contain the minimum required number of surgeries for each surgery type.

In phase 1, we generate a set of ORDS that maximises OR utilisation. Note, however, that the primal model yields a fractional solution. When we obtain a integer solution in phase 2, it is not likely to be optimal. We investigate the maximal deviation from the optimum in Chapter 5.

4.3.3 Phase 2: Assigning ORDSs to OR-days

Solving the model of phase 1 yields a set of ORDSs that optimise OR utilisation. Phase 2 is concerned with assigning the ORDSs to OR-days to optimise the bed-levelling. It is not possible linearise the bed-levelling constraint (4.9) and therefore, we choose to solve this problem by a simulated annealing approach. Simulated annealing is a local search heuristic. In every iteration of the simulated annealing procedure, the current solution is replaced by a 'neighbour' solution depending on some probability. The probability is based on the objectives of the current and 'neighbour' solution and a temperature parameter. The temperature parameter makes sure that probability of accepting a worse solution decreases as the simulated annealing procedure progresses. The reader is referred to the paper of Kirkpatrick et al. [9] for further explanation of simulated annealing. .

Construct solution In order to apply simulated annealing, we first need to construct a starting solution. Since the primal model in phase 1 yields a fractional solution, we have not yet selected the subset of ORDSs that are assigned to OR-days. We construct an ILP-model that selects the ORDSs that maximise OR utilisation and assigns them to OR-days to come up with an initial solution. Let Z_{kjt} be a binary decision variable that indicates whether ORDS k is performed on OR-day (j, t). We introduce χ_{sk} as a binary parameter that indicates whether surgeon $s \in S$ performs ORDS $k \in K$. Since only one surgeon performs all surgeries on an ORDSs, each ORDS $k \in K$ belongs to a single surgeon, i.e. by definition $\sum_{s \in S} \chi_{sk} = 1 \; \forall k$. Note that we can not guarantee optimality of OR utilisation, since we generate the set of ORDSs by means of the primal model that yields a fractional solution.

$$\max \quad \sum_{i \in I} \sum_{k \in K} \sum_{j \in J} \sum_{t \in T} (\theta^u \mu_i + \theta^v \sigma_i) Z_{kjt} O_{ik} \tag{4.18}$$

subjected to

$$\sum_{k \in K} Z_{kjt} \leq A_{jt} \qquad \forall j, t \tag{4.19}$$

$$\sum_{k \in K} \sum_{j \in J} Z_{kjt} \chi_{sk} \leq B_{st} \qquad \forall s, t \tag{4.20}$$

$$D_i^{\min} \leq \sum_{k \in K} \sum_{j \in J} \sum_{t \in T} Z_{kjt} O_{ik} \leq D_i^{\max} \qquad \forall i \tag{4.21}$$

$$\sum_{i \in I} \sum_{j \in J} \sum_{k \in K} Z_{kjt} O_{ik} \omega_{ir} \leq Q_r^{\max} \qquad \forall r, t \tag{4.22}$$

$$Z_{kjt} \in \{0, 1\} \qquad \forall k, j, t$$

Constraints (4.19), (4.20), (4.21) and (4.22) are the equivalences of constraints (4.3), (4.4), (4.6) and (4.7) in the base model. They respectively represent the OR-day availability, surgeon availability, demand and instrument constraint.

Improve solution Let $h_n^k(x)$ be the discrete distribution of x patients of ORDS k still in recovery on day n within one planning horizon. Since we know how many surgeries of which type are performed on each ORDS (phase 1), we can determine $h_n^k(x)$ by using the model of Vanberkel et al. [16]. In Appendix E, we explain how we use this model to determine $h_n^k(x)$. Since $h_n^k(x)$ is the bed distribution for one planning horizon, $h_n^k(x)$ is only defined for $n \leq |T|$. However, the maximum length of stay of a patient may be longer than the planning horizon. We solve this problem by flipping the probability distribution around for those days of the length of stay that fall into a new planning horizon. Figure 4.6 shows an abstraction of flipping around a probability distribution. This method implies the surgery schedule to be cyclic. Note, however, that we propose a non-cyclic surgery scheduling approach. The assumption of cyclicity thus only applies to determine the bed occupancy since it is likely that the case-mix of subsequent planning horizons will not be significantly different.

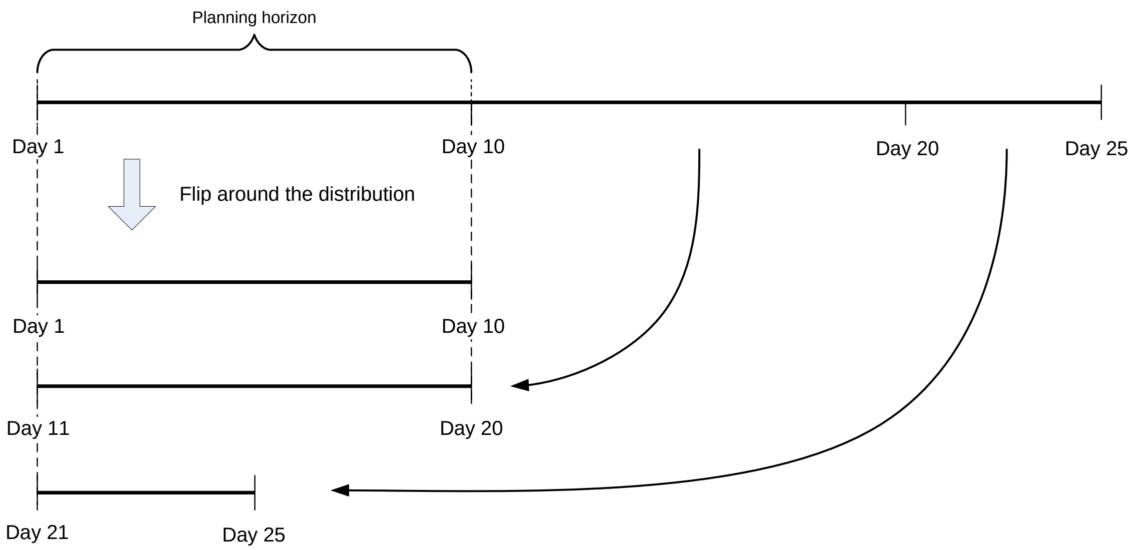

Figure 4.6: An abstraction of flipping around a probability distribution

After constructing an initial solution, we apply simulated annealing (SA) to improve the solution. Our objective is to level the required number of beds per day. Since we assume a stochastic length of stay, also the number of occupied beds on day $t \in T$ is stochastic. Let H_t be the probability distribution of the number of occupied beds on day $t \in T$. To calculate this distribution, we first have to identify the impact of each ORDSs on the number of beds on each day of the planning period. This is done, using the same approach as in the article of Vanberkel et al. [16]. We introduce $\bar{h}_t^k(x)$ as the discrete distribution of x patients of ORDS k still in recovery on day t. Note that $\bar{h}_t^k(x)$ depends on the assignment of ORDSs to OR-days as defined by Z_{kjt}.

$$\bar{h}_t^k(x) = \begin{cases} h^k_{(t-\hat{t}) \bmod |T|}(x) & \text{for } \hat{t} \text{ where } \sum_j Z_{kj\hat{t}} = 1 \text{ and } \hat{t} \leq t \\ \mathbf{0} & \text{otherwise} \end{cases} \tag{4.23}$$

where $\mathbf{0}$ means $\bar{h}_t^k(0) = 1$ and $\bar{h}_t^k(x) = 0 \; \forall \; x > 0$.

In short, we rename h_n^k to relate the ORDSs to the OR-day on which they are scheduled. We can now determine H_t by iteratively adding the distributions of all ORDSs on day t. Adding independent discrete distributions is done by discrete convolutions which we indicate by $*$.

$$H_t = \bar{h}_t^1 * \bar{h}_t^2 * \ldots * \bar{h}_t^{|K|} \tag{4.24}$$

In order to express the bed occupancy on a given day, we determine the number of beds at some quantile β. Let Θ_t^β be the number of occupied beds on day $t \in T$ at quantile β.

$$\Theta_t^\beta = F_t^{-1}(\beta) \qquad \text{where} \qquad F_t(x) = \mathrm{P}(H_t \leq x) \tag{4.25}$$

We define Ψ as the objective function for a given solution as the maximal number of required beds at quantile β in the whole planning horizon.

$$\Psi = \max_{t \in T} Q_t^\beta \tag{4.26}$$

We perform the SA approach as follows:

1. Determine a neighbour solution of the current solution by randomly swapping the assignment of two ORDSs to OR-days. The swaps need to be feasible with respect to the surgeon and resources availability (constraints (4.20) and (4.22)). Let Ψ_{cur} and Ψ_{neigh} be the objective of the current and neighbour solution by using (4.26).

2. Accept the neighbour solution when the objective is better (smaller) or equal to the current solution, i.e. when $\Psi_{\text{neigh}} \leq \Psi_{\text{cur}}$. Otherwise, accept the neighbour solution with some probability φ.

$$\varphi = e^{\frac{\Psi_{\text{cur}} - \Psi_{\text{neigh}}}{\tau}} \tag{4.27}$$

   where τ is the temperature parameter. The parameter τ is decreased by a cooldown factor ν after a set of ω iterations. If the neighbour solution is accepted, this solution becomes the new current solution.

3. Save the new current solution as the best solution if it has a lower objective than the best solution found thus far.

4. The procedure stops when the temperature τ meets a threshold temperature τ^{stop}. Otherwise, continue at step 1.

Determine simulated annealing parameters (cooling scheme) The simulated annealing procedure starts with an initial temperature τ^{start}. We want to choose τ^{start} in such a way that we accept a deterioration of the objective with a high probability at the start of the procedure. We observe that the maximal deterioration of the objective is the maximal amount of surgeries of an ORDS minus the minimal amount of surgeries of an ORDS. We state that we want to accept the maximum deterioration with a probability of 0.95. Using (4.27), we determine the initial temperature to accept a maximum deterioration of the objective with a probability of 0.95.

$$\tau^{\text{start}} = \frac{\min_{k \in K} \sum_{i \in I} O_{ik} - \max_{k \in K} \sum_{i \in I} O_{ik}}{\ln(0.95)} \tag{4.28}$$

Using the same approach, we determine the threshold temperature τ^{stop}. We want to stop the procedure when there is a low probability of accepting the minimal deterioration of the objective. Since the objective is an integer, the minimal deterioration is 1 bed. Using (4.27), we determine what threshold temperature we should choose to accept a deterioration of 1 bed with a probability of 0.001.

$$\tau^{\text{stop}} = \frac{-1}{\ln(0.001)} \tag{4.29}$$

As a rule of thumb, we should perform as many iterations per temperature as there are neighbour solutions. The number of neighbour solutions is equal to the total amount of available OR-days since each ORDS can be swapped with as many ORDSs as there are OR-days. This is an upper bound, since not all swaps are feasible regarding instruments and surgeon constraints.

$$\omega = \sum_{j \in J} \sum_{t \in T} A_{jt} \tag{4.30}$$

The cooldown factor is set to $\nu = 0.95$ by default, but can be decreased to speed up the simulated annealing procedure.

4.3.4 Technical implementation

We use AIMMS[1] to implement phase 1 and the construction of a solution in phase 2 of our technical model. AIMMS is a program to implement (I)LP-models like the ones we formulate in this section. Figure 4.7 displays a screenshot of the model interface. The model requires an Excel-file that contains the input data. Some of the parameters can be adjusted in the interface, such as the weights in the objective function (θ^u and θ^v), the overtime probability (α) and the OR capacity per OR-day (C). The model can be executed step by step (buttons on the left side) or at once by using the *Execute model* button. The model outputs a set of ORDSs and an initial solution of assigning the ORDSs to OR-days. The program writes the output and the original input data to a .csv file.

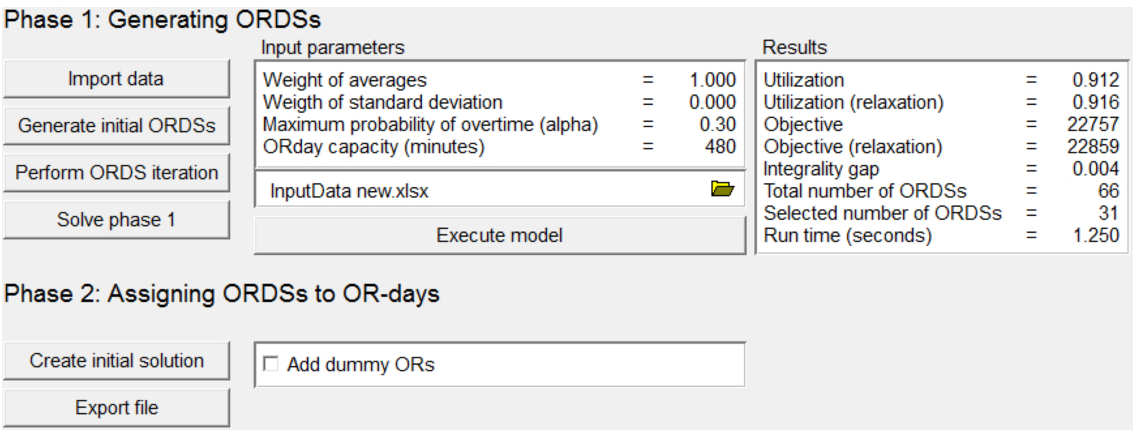

Figure 4.7: Screenshot of the implementation of the model interface in AIMMS

We implement the simulated annealing procedure, from phase 2 in our technical model, in a self-made program that is programmed in Delphi[2]. We use Delphi instead of AIMMS

[1] http://www.aimms.com
[2] http://www.embarcadero.com/products/delphi

since AIMMS requires 700 ms to calculate the bed distribution, while Delphi only needs 4 ms. Since the bed distribution is calculated after every swap, the total run time of simulated annealing procedure was much better in Delphi than in AIMMS. The Delphi program reads the .csv file from the AIMMS model to load the set of ORDSs and the initial assignment of ORDSs to OR-days. Figure 4.8 shows a screenshot of the Delphi program after minimising the number of required beds.

Our major concern, before the implementation of the simulated annealing procedure, was the calculation time of the bed distributions. Since calculating convolutions is computational intensive, determining the bed distribution for many swaps is a major bottleneck in the total running time of the simulated annealing procedure. In the end, we managed to perform 7200 swaps within 32 seconds on a 2,4 GHz dual core laptop based on a planning horizon of 28 days and 52 OR-days.

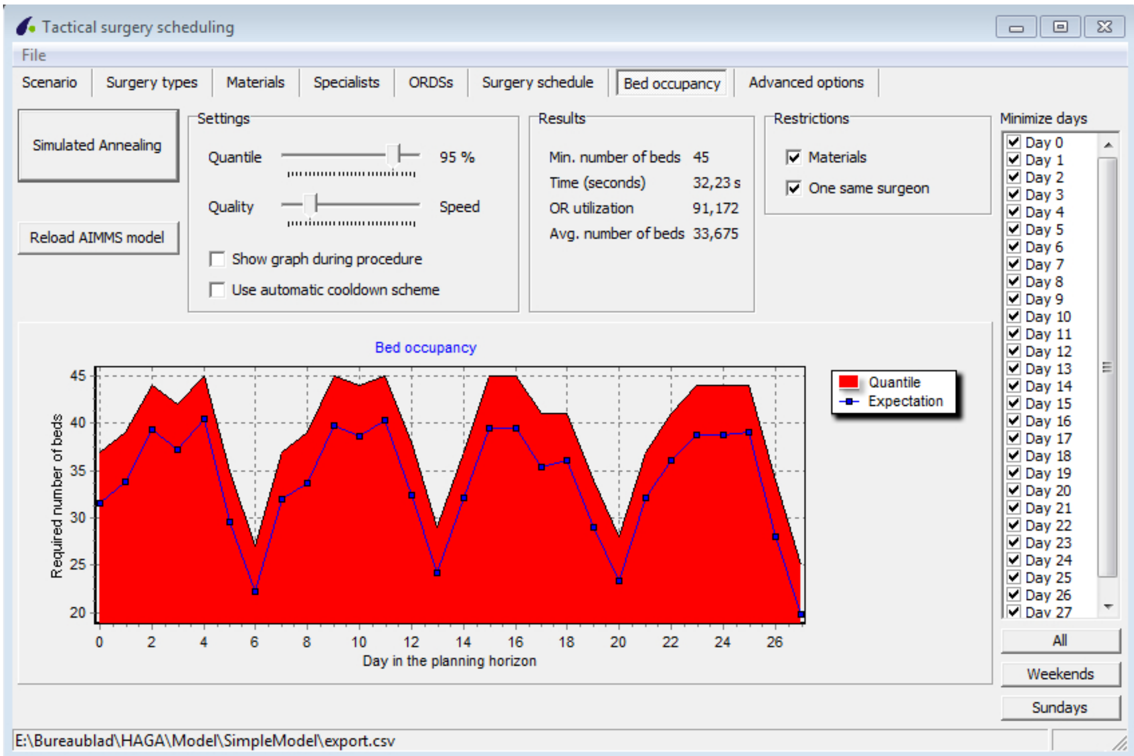

Figure 4.8: Screenshot of the simulated annealing implementation in a Delphi program

4.4 Validation

We validate our model by performing a Monte Carlo simulation on a surgery schedule that results from our surgery scheduling approach. We use a surgery schedule that results from optimising the default scenario (Appendix B). In a single experiment in our simulation, we randomly draw patients from the historical data for each scheduled surgery type in the surgery schedule. Thus, we simulate a realisation of the surgery schedule. If we perform multiple experiments, the OR utilisation and the bed occupancy of the realisation data should match with the expectation of the model.

Table 4.1 shows a comparison of the model expectations and the simulation results of 1000 experiments. The confidence intervals indicate that the expectation of an indicator should lay within that interval with a probability of 95%. The relative error is a measure

	Model	Simulation		
Indicator	*Expectation*	*Average*	*Confidence interval (95%)*	*Relative error*
OR utilisation	91.17%	91.24%	[91.16%, 91.32%]	0.09%
Overtime probability	max 30%	23.86%	[23.61%, 24.11%]	1,04%
Average number of beds	33.68	33.63	[33.53, 33.73]	0.29%
Required number of beds	45	43.37	[43.23, 43.51]	0.32%

Table 4.1: Comparison of model expectations and simulation ($n = 1000$) results

for the precision of the simulation results. We calculate the relative error by dividing the confidence interval half-width by the average.

The expected OR utilisation and average number of beds of the model lay within the confidence interval of the simulation results, so this validates that the expectation of the model is accurate. Also, we require in our model that the probability of overtime has a maximum of 30%. The simulation results show that we meet this requirement, since the average overtime is 23,86%. According to the model, with a probability of at least 95%, we do not require more than 45 beds on each day of the planning horizon. The simulation results show that we require 43.37 beds on average. The average is lower than the 45 beds of the model due to low bed occupancy on weekend days.

We make one more step in validating our calculations of the bed distributions. In Table 4.2, we compare the bed occupancy of the model and the simulation experiments for each day in the planning horizon. For the model, we display the expected number of beds and the required number of beds at a 95% quantile. This data is visualised in the screenshot of the Delphi program in Figure 4.8. For the simulation, we show the averages, confidence intervals, relative errors and the percentage of days that the number of beds is less or equal to the 95% quantile of the model (*sufficient beds*). It appears that for every day of the week, the expected number of beds of the model lays within the confidence interval of the simulation. This validates that the expected number of beds of the model is accurate. As expected, the percentages in the *sufficient beds* column are at least 95%. This validates that, on each day of the planning horizon, having the number of required beds results in having sufficient beds on least 95% of the days.

4.5 Conclusions

In this chapter, we propose a master surgery scheduling (MSS) approach to maximise OR utilisation and minimise the number of required beds. Unlike the models from the literature, our MSS approach is non-cyclic. In addition, we require that only one surgeon performs surgery per OR-day. As a result, we do not only schedule surgery types, but also surgeons. Furthermore, the number of surgeries per surgery type is variable while a fixed number of surgeries is scheduled in most models from literature. We base our MSS model on the research of Van Oostrum et al. [15]. We extend this model by using the paper of Vanberkel et al. [16] to determine the bed occupancy based on stochastic length of stay.

We propose to construct a surgery schedule with a planning horizon of 4 weeks that is determined 3 months in advance. Our MSS approach consists of two phases. In the first phase, we generate a set of operating room day schedules (ORDSs) that optimise OR utilisation. We introduce a new method to linearise the constraint to limit the probability of overtime. The method finds a minimal number of piecewise linear functions to approximate the probability of overtime. In the second phase, we apply a simulated annealing

	Model		**Simulation**			
		Required		*Confidence*	*Relative*	*Sufficient*
Day	*Expectation*	*beds*	*Average*	*interval (95%)*	*error*	*beds*
0	31.6	37	31.4	[31.2, 31.6]	0.6%	97%
1	33.8	39	33.7	[33.5, 33.8]	0.5%	97%
2	39.4	44	39.2	[39.1, 39.4]	0.5%	96%
3	37.2	42	37.2	[37.0, 37.4]	0.5%	97%
4	40.5	45	40.3	[40.2, 40.5]	0.4%	96%
5	29.6	35	29.4	[29.2, 29.6]	0.7%	97%
6	22.2	27	22.1	[21.9, 22.3]	0.9%	96%
7	31.9	37	31.8	[31.6, 32.0]	0.6%	97%
8	33.7	39	33.2	[33.5, 33.9]	0.6%	97%
9	39.8	45	39.8	[39.6, 40.0]	0.5%	97%
10	38.7	44	38.7	[38.5, 38.9]	0.5%	96%
11	40.3	45	40.3	[40.2, 40.5]	0.4%	96%
12	32.4	38	32.5	[32.3, 32.7]	0.6%	98%
13	24.2	29	24.1	[23.9, 24.3]	0.8%	95%
14	32.1	37	32.1	[31.9, 32.3]	0.6%	96%
15	39.5	45	39.5	[39.3, 39.6]	0.5%	97%
16	39.5	45	39.6	[39.4, 39.8]	0.5%	97%
17	35.4	41	35.3	[35.1, 35.5]	0.6%	97%
18	36.1	41	36.0	[35.8, 36.2]	0.5%	96%
19	29.0	34	29.0	[28.8, 29.2]	0.7%	96%
20	23.3	28	23.2	[23.1, 23.4]	0.8%	95%
21	32.1	37	32.1	[31.9, 32.3]	0.6%	96%
22	36.1	41	36.2	[36.0, 36.4]	0.5%	95%
23	38.7	44	38.7	[38.5, 38.9]	0.5%	96%
24	38.7	44	38.7	[38.5, 38.9]	0.5%	97%
25	39.1	44	39.1	[38.9, 39.3]	0.5%	96%
26	28.1	34	28.1	[27.9, 28.3]	0.8%	97%
27	19.8	25	19.9	[19.7, 20.1]	1.0%	96%

Table 4.2: Comparison of model and simulation ($n = 1000$) bed occupancy

approach that swaps the assignment of ORDSs to OR-days around in order to minimise the number of required beds. We define the number of required beds as the amount of beds that yields a probability of at least 95% that sufficient beds are available on each day of the planning horizon.

We validate our model by performing a Monte Carlo simulation on a surgery schedule of the default scenario that results from our model. In the experiments ($n = 1000$), we randomly draw patients from the historical data for each scheduled surgery type in the schedule. Our model is valid, since the expected OR utilisation and bed occupancy from our model matches the simulation results.

Chapter 5

Computational results

In Chapter 4, we propose a surgery scheduling model to maximise OR utilisation and minimise the required number of beds. In this chapter, we perform numerical experiments to determine how the model performs under different circumstances. Our experiments are based on case study data of the orthopaedic department. Some of the data can not be influenced by direct managerial interventions, for example, the surgery durations and length of stay. Changes in the dataset due to these external factors are referred to as scenarios. Data that can be influenced by management, such as the planning horizon and overtime probability, is referred to as experimental factors.

In Section 5.1, we describe the performance indicators of the experiments. Next, we describe the different scenarios in Section 5.2. In Section 5.3, we discuss the experimental factors. Section 5.4 presents the results of the experiments. We conclude this chapter in Section 5.5.

5.1 Performance indicators

The performance indicators are based on the two objectives of this research: maximising OR utilisation and minimising the required number of beds. Regarding the required number of beds, we introduce additional performance indicators.

The required number of beds in the weekends is expected to be lower than on working days, since no surgeries are performed on Saturdays and Sundays. Also, the required numbers of beds on Sundays is lower than on Saturdays, because some patients are discharged while no new patients are hospitalised. Since the orthopaedic department has two nursing wards, it may be worthwhile to minimise the required number of beds on weekend-days. This may enable closing one of the wards on Saturdays and/or Sundays to save personnel costs. Therefore, we make a distinction between minimising the required number of beds in the whole planning period, Saturdays and Sundays.

Another performance indicator is the runtime of the model. Since the model consists of two modelling phases (generating ORDSs and assigning ORDSs to OR-days), we record the runtime for each phase separately. With respect to the generation of ORDSs, we keep track of the amount of generated ORDSs and the number of ORDSs that are assigned to OR-days. Furthermore, we register the objective of the model and its relaxation. The relaxation of the model is solved to optimality, but the integer solution is not likely to be optimal. Therefore, we also record the integrality gap: the deviation between the model relaxation and the integer solution. Table 5.1 summarises the performance indicators.

Generating ORDSs	Assigning ORDSs to OR-days
• Runtime • Total number of ORDSs • Selected number of ORDSs • Integrality gap • Objective • Objective (relaxation) • OR utilisation • OR utilisation (relaxation)	• Runtime • Required number of beds in the planning horizon • Required number of beds on Saturdays • Required number of beds on Sundays

Table 5.1: Performance indicators for numerical experiments

5.2 Scenarios

The following subsections describe the scenarios for our experiment design. Table 5.3 summarises the scenarios.

5.2.1 Default scenario

Most data follows directly from interviews with management, surgeons and OR personnel and from the historical data from SAP (patient data from 2010). For example, the OR-day availability (A_{jt}), surgeon availability (B_{jt}), available time per OR-day (C), surgery durations (μ_i), etc. Some parameters require special explanation. First of all, we set the minimum number of surgeries per surgery type ($D_i^{\min}$) equal to the average number of surgeries per 4 weeks in 2010. We round down when the average is a fractional number. We motivate this choice by considering that we want to develop a model that improves OR utilisation, so we should at least be able to schedule the same amount of surgeries that were performed in 2010. To keep the number of surgeries within bounds, we allow a maximum of 5 additional surgeries per surgery type, i.e $D_i^{\max} = D_i^{\min} + 5 \quad \forall i$. In our default scenario, we set $\theta^v = 0$ and $\theta^u = 1$. We test the effect of an positive standard deviation weight by varying this factor in our experiments. From interviews, we conclude that a overtime probability of $\alpha = 0.3$ is most appropriate. Furthermore, we set the minimum probability of requiring γ beds to 95%, i.e. $\beta = 0.95$. We display the default scenario data in Appendix B.

5.2.2 Longer average surgery duration

Due to technical innovation, surgeries can be performed in a less invasive way. As a result, patients have a reduced length of stay. On the other hand, minimal invasive surgeries may increase surgery durations. Also, when the model is applied to another specialism, the case-mix of surgery types could have higher expected surgery durations (for example: cardiac surgery). We model a longer surgery duration by multiplying the expected duration for every surgery type (parameter μ_i) by a factor 1.5. The minimum demand for each surgery type is divided by a factor 1.5 to prevent having an infeasible solution.

	Days							
Scenarios	0	1	2	3	4	5	6	7
Default scenario	0	0	0.1	0.2	0.3	0.2	0.1	0
Shorter length of stay	0	0	0.1	0.5	0.2	0.1	0	0

Table 5.2: Example of probability distribution of a surgery type

5.2.3 Shorter average surgery duration

Reduction of surgery durations could occur when surgeries are performed more efficiently. For example, when a surgeon gains experience for a new kind of surgery type, it is likely that surgery durations decrease. Other specialisms could have a case-mix of surgery types with shorter surgery durations (for example eye surgery), so decreasing the surgery durations is relevant to see how the model performs for other specialisms. We test this scenario by dividing the expecting duration for every surgery type (parameter μ_i) by a factor 1.5. In order to make a fair comparison with other scenarios, the minimum demand for each surgery type is multiplied by a factor 1.5.

5.2.4 Higher variation of surgery duration

The variation of surgery duration for every surgery type (parameter σ_i) is multiplied by 1.5 to measure the effect of the scenario in which surgery durations have an increased variation. We test this scenario in order to measure the influence of variation of surgery durations on OR utilisation and the required number of beds. Next to that, other specialisms could have a case-mix of surgery types that have a higher variation of surgery durations.

5.2.5 Lower variation of surgery duration

For the same reasons as in the previous subsection, we test a scenario where we model the scenario in which surgery durations have a decreased variation. The variation of surgery durations for every surgery type (parameter σ_i) is divided by 1.5.

5.2.6 More surgery types

The current scenario contains 43 surgery types which are based on the interviews with surgeons. In order to test the scenario with more surgery types, we use the cluster technique from Appendix C to define 117 surgery types. Only procedure types that require the same resources and surgeons are combined.

5.2.7 Shorter length of stay

Due to pressure from health insurers and the government, several projects in HagaZiekenhuis aim at reducing the length of stay of patients. Therefore, it is a likely scenario that the length of stay reduces in the upcoming years. To model this scenario, we reduce the length of stay of surgery types by one day from the third day of hospitalisation. Table 5.2 shows an example of how we adjust the length of stay distribution. We choose to reduce the length of stay by one day after the third day, because a reduction of length of stay is easier to achieve for patients who are hospitalised for a relative long period of time.

Scenarios	*Description*
Current scenario	Default dataset
Longer average surgery duration (μ_i)	Increase by a factor 1.5
Shorter average surgery duration (μ_i)	Decrease by a factor 1.5
Higher variation of surgery duration (σ_i)	Increase by a factor 1.5
Lower variation of surgery duration (σ_i)	Decrease by a factor 1.5
More surgery types (I)	103 surgery types based on clustering heuristic
Shorter length of stay (l_n^i)	Reduce the length of stay by one day after the third day
Increase of hip and knee prostheses patients ($D_i^{\min}$)	Increase by 10%

Table 5.3: Scenarios for numerical experiments

5.2.8 Increase of hip and knee prostheses patients

Due to an ageing society, it is expected that the number of hip and knee prostheses patients increases in the future, since older people generally suffer from worn hips and knees. We model this scenario by increasing the minimum number of hip and knee prostheses surgeries by 10% while maintaining the same amount of OR capacity. A higher percentage results in an infeasible solution.

5.3 Experimental factors

The following subsections describe the experimental factors for our experiment design. Table 5.4 summarises the experimental factors and their values.

5.3.1 Standard deviation weight

In the default dataset, the weight for the OR utilisation and standard deviation in the objective function (θ^u and θ^v in (4.1) of the base model) is 1 and 0, respectively. To increase the preference of scheduling surgery types with a high standard deviation, we set the standard deviation weight θ^v to 1 and 2.

5.3.2 Overtime probability

When an ORDS is generated, the probability that no overtime occurs is at least α. When overtime probability is low, little slack is needed to buffer against overtime. Therefore, a low overtime probability results in a high OR utilisation, and vice versa. In our default setting, the overtime probability is 30%. Next to that, we choose the values 50% and 15% as overtime probabilities in our experiments.

5.3.3 Opening hours

In the current situation, each OR-day is available for 8 hours. OR management considers to increase the opening hours of the OR department for an extra hour in the near future. By adjusting the opening hours of an OR-day in the model (parameter C), we can determine the effects of such a change to the OR utilisation and the required number of beds. Next to that, we increase the opening hours to 12 hours to see the effect of opening hours on OR utilisation. To make a fair comparison, the minimum demand per surgery type (parameter $D_i^{\min}$) is increased by the same proportion as the opening hour increase.

5.3.4 Planning horizon

In developing a solution approach, we decided upon an arbitrary planning horizon of 4 weeks. To investigate what would be an appropriate planning horizon, we vary the planning horizon length to 2 weeks and 8 weeks.

5.3.5 Available days

The number of ORs that are available each day, are given by parameter A_{jt}. By disregarding constraint (4.3) in the base model, we do not restrict the model to schedule a fixed amount of ORDSs per day. Since we relax the problem, we expect to come up with an assignment of ORDSs to OR-days that yields a higher performance. The total number of the available OR-days over the whole planning period remains constant (52 OR-days). Therefore, this experimental factor has no influence on OR utilisation, but on the required number of beds. We allow a maximum of 5 OR-days per day since 5 ORs are physically available at the operating department.

5.3.6 Surgeon availability

As with the previous factor, the surgeon availability corresponds to a constraint in the base model. To determine how the surgeon availability (parameter B_{st}) restricts the performance, we solve the model without constraint (4.4).

5.3.7 One same surgeon per day

This factor corresponds to constraint (4.5) in the base model. This constraint makes sure that a surgeon can only perform surgery on one OR-day each day. If we disregard this constraint, we imply that there is more than one surgeon who performs a specific subset of surgeries. We solve the model without this constraint to determine the effect of this constraint on the performance. For example, it may be worthwhile to hire an additional surgeon who performs foot-related surgery types to improve the OR utilisation and/or the required number of beds.

5.3.8 Minimum and maximum demand

For a planning period, each surgery type has a minimum and maximum amount of surgeries that have to be performed to comply to production agreements or to manage the waiting lists. We solve the model without the demand constraint (4.6) to see what theoretical performance maximum can be achieved. We also perform an experiment where minimum demand reduces by 20% to assess the performance when the demand constraint is less tight.

5.3.9 Instruments

For each instrument r, the parameter $Q_r^{\max}$ denotes how much instruments there are available on a single day. In would be interesting to see how the availability of instruments restricts the performance of a surgery schedule. As an experimental factor, we solve the model without the instrument constraint (equation (4.7) in the base model).

Factors	*Values*		
Standard deviation weight (θ^v)	**0**	1	2
No-overtime probability (α)	50%	**30%**	15%
OR time (C)	**8 hours (0)**	9 hours (+1)	12 hours (+4)
Planning horizon (T)	2 weeks	**4 weeks**	8 weeks
Available days		disregard constraint (4.3)	
Surgeon availability		disregard constraint (4.4)	
One same surgeon per day		disregard constraint (4.5)	
Minimum and maximum demand		disregard constraint (4.6)	
Instruments		disregard constraint (4.7)	

Table 5.4: Factors for numerical experiments

5.4 Results

In our experiments, we test every scenario and experimental factor in isolation, except for the last experiment ('No constraints') in which we disregard the constraints of last five experimental factors as shown in Table 5.4. Figure 5.1 shows the results of the experiments. In the following subsections we discuss the results of the experiments per performance indicator. When we use adjectives like *higher*, *lower*, *more*, *less*, etc. they refer to the default scenario.

5.4.1 Runtime

The runtime of the ORDS generating model is about less then two seconds in most of the experiments. One of the exceptions is the experiment with more surgery types and the experiment where the planning horizon is 8 weeks. In both scenarios, the number of variables is higher compared to the default scenario. The model is NP-hard since the mathematical programs in our model are ILPs. Therefore, the runtime increases exponentially in the number of variables. Three scenarios have a runtime between 3 and 10 seconds. In these scenarios, determining an initial solution of assigning ORDSs to OR-days takes the most time because of tight constraints. The runtime is very low for the last two experiments ('Min. and max. demand' and 'No constraints') because the model is less complicated in absence of the demand constraint.

The simulated annealing procedure needs more time to be executed than the ORDS generation model. Most experiments take about 30 seconds. The planning horizon influences the runtime: a planning horizon of 2, 4 and 8 weeks takes 6, 33 and 188 seconds, respectively. This can be explained by the time it takes to determine the bed distribution after each swap. Calculating convolutions is the most time-consuming process because the simulated annealing procedure calculates the bed distribution for over 7500 times. Every extra OR-day that is added to the bed distribution, increases the runtime exponentially. Another time-consuming factor is the number of scheduled surgeries per ORDS, because it increases the time needed to calculate one convolution. For example, the long duration scenario has a relative short runtime because the number of surgeries per ORDS is smaller than in the default dataset. The opposite holds true for the short duration scenario and the scenario where the OR-day length is 12 hours.

Scenario's	*Generating ORDS's*								*Assigning ORDSs to OR-days*			
	Time (sec)	Total ORDS	Selected ORDS	Integrality gap	Objective	Objective (relaxation)	OR utilisation	OR utilisation (relaxation)	Time (sec)	Min beds total	Min beds Saturdays	Min beds Sundays
Default	1,21	66	31	0,4%	22757	22859	91,2%	91,6%	32,98	45	33	27
Long duration	0,76	52	26	1,0%	22425	22658	89,8%	90,8%	21,11	32	23	18
Short duration	2,82	84	36	0,7%	22515	22682	90,2%	90,9%	67,08	67	49	38
High variation	1,58	65	31	2,0%	22040	22485	88,3%	90,1%	35,27	45	34	28
Low variation	1,86	72	26	0,4%	23089	23189	92,5%	92,9%	34,40	46	35	27
More surgery types	15,28	176	43	1,0%	23231	23462	93,1%	94,0%	17,97	52	37	31
Shorter LOS	1,01	66	30	0,4%	22757	22859	91,2%	91,6%	29,87	43	31	23
Increase hip and knee	7,19	71	33	0,7%	22703	22856	91,0%	91,6%	35,07	46	33	27
Experimental factors												
Std. dev. weight : 1	1,62	69	35	1,0%	30341	30659	90,4%	90,0%	38,95	46	33	27
Std. dev. weight : 2	1,37	73	32	1,8%	38213	38895	90,4%	90,1%	37,64	45	33	27
Overtime probability 50 %	0,82	73	32	0,2%	24723	24767	99,1%	99,2%	37,36	50	36	29
Overtime probability 15 %	9,29	69	30	1,1%	20921	21158	83,8%	84,8%	32,53	43	32	25
Opening hours: 9 hours	1,22	69	30	0,4%	25821	25916	92,0%	92,3%	43,93	51	38	31
Opening hours: 12 hours	3,11	76	33	0,3%	34932	35029	93,3%	93,6%	66,46	69	51	40
Planning horizon: 2 weeks	0,75	53	22	0,7%	11310	11392	90,6%	91,3%	6,26	45	33	26
Planning horizon: 8 weeks	35,92	85	37	0,5%	45402	45637	90,9%	91,4%	188,87	48	37	29
Available days	1,16	66	31	0,4%	22757	22859	91,2%	91,6%	93,77	46	35	27
Surgeon availability	1,14	66	31	0,4%	22757	22859	91,2%	91,6%	32,26	45	33	26
One same surgeon per day	1,16	66	31	0,4%	22757	22859	91,2%	91,6%	32,34	45	33	25
Min. and max. demand	0,37	9	6	0,4%	23502	23587	94,1%	94,5%	22,73	50	36	28
Instruments	1,13	66	29	0,5%	22712	22829	91,0%	91,5%	33,38	45	33	26
20% demand reduction	0,96	58	27	0,6%	22896	23039	91,7%	92,3%	31,74	46	34	28
No constraints	0,36	9	1	0,0%	23833	23833	95,5%	95,5%	42,19	39	10	6

Figure 5.1: Results of experiments

5.4.2 Number of ORDSs

In most experiments, the ORDS generating model generates about 70 ORDSs and selects 30 ORDSs to be scheduled. In the 'More surgery types' scenario, more ORDSs are generated because an increase of surgery types leads to more possible combinations of surgery types on an ORDS. To a lesser extend, the same holds for the 'Short duration' scenario. Due to short surgery durations, more surgeries can be scheduled per ORDS which lead to more possible combinations of surgery types on an ORDS. The last two experiments ('Min. and max. demand' and 'No constraints') only have 9 ORDSs: the most optimal one for every surgeon. In the 'No constraints' experiment, the model schedules the ORDS with the highest OR utilisation on each OR-day in the planning horizon.

5.4.3 Integrality gap

For most scenarios, the integrality gap is about 0.4%. A positive standard deviation weight in the objective function seems to result in a relative high gap ($\geq 1\%$). Still, the OR utilisation of the scenarios with a standard deviation weight of 1 and 2, only differs by 0.4% and 0.3%, respectively. We believe that the integrality gap is small enough to conclude that the integrality gap is negligible.

5.4.4 Objective

By definition, the objective of the relaxation is higher than the objective of the integer solution. It follows from the definition of the objective that it is higher for experiments with a positive standard deviation weight. Since the first term of the objective is the sum of the session times, the objective is higher (lower) for the experiments with increased opening hours or a longer (shorter) planning horizon.

5.4.5 OR utilisation

For most experiments, the OR utilisation of the relaxation is greater or equal to the OR utilisation of the integer solution. This does not hold true for experiments with a positive standard deviation weight because the objective also includes a term for the standard deviation. For these experiments there is no direct connection between the objective and OR utilisation.

The OR utilisation is the lowest for the experiments that require more slack in the ORDSs than the default scenario. The 'High variation' experiment requires a high amount of slack to buffer against the uncertainty of the high variation of the surgery durations. In the experiment where the overtime probability is reduced to 15%, more slack is needed to increase the probability that an ORDS ends within office hours. The other way around, the 'Low variation' and the 'overtime probability 50%' experiment have a high OR utilisation because the experiments require less slack than the default scenario.

Both the 'Long duration' and the 'Short duration' scenario have a lower OR utilisation than the default scenario. One may expect that shorter surgery durations should lead to a higher OR utilisation, because a schedule can be filled more easily with short surgeries. On the other hand, slack increases with the number of surgeries. In the 'Short duration' scenario, the advantage of increased plan-ability does not weigh up to the increased amount of slack.

The 'More surgery types' scenario has a higher OR utilisation than the default scenario. More surgery types lead to more possible combinations of surgery types on an ORDS. As a result, the probability of generating ORDSs with a high OR utilisation increases. In turn, this results in a higher OR utilisation for the whole planning horizon.

As expected, the OR utilisation of the 'Shorter LOS' scenario matches the default scenario, since the length of stay data is not used in generating ORDSs. An increase in hip and knee prostheses surgeries results in an OR utilisation that approximately matches the default scenario.

Increased opening hours cause a higher number of scheduled surgeries per ORDS which results in an increased amount of slack. Still, increased opening hours improve OR utilisation, since the increase of expected surgery durations is higher than the increased amount of slack.

It is striking that most experiments where we disregard constraints ('Available days', 'Surgeon availability', 'One surgeon per day' and 'Instruments' experiment) does not lead to an improved OR utilisation. Apparently, those constraint are no bottlenecks in improving the OR utilisation level of the default experiment. Due to linear approximation in the overtime constraint, the 'Instruments' experiment even has a slightly lower OR utilisation (- 0.1%).

When we lower the minimum demand for each surgery type by 20%, the model has more freedom of scheduling surgeries. This results in a slight increase of OR utilisation (+ 0.6%). In the 'Min. and max. demand' experiment, we disregard the demand constraint completely. The model uses 1 ORDS per surgeon, so the model schedules a small subset of surgery types. The OR utilisation of this experiment is higher than the default scenario (94.1%). In the last experiment ('No constraints'), the model obtains an OR utilisation of 95.5% since only one ORDS is scheduled on each OR-day in the planning horizon.

In Subsection 2.3.2, we determined that the orthopaedic department attained an average OR utilisation of 78% in 2010. If we assume that the standard deviation weight is positive, the model yields an OR utilisation of 90% (equal to the Plexus best practice benchmark). Therefore, it is feasible to improve the average OR utilisation by 12%. Figure 5.2 visualises the comparison of the historical OR utilisation and our surgery scheduling approach.

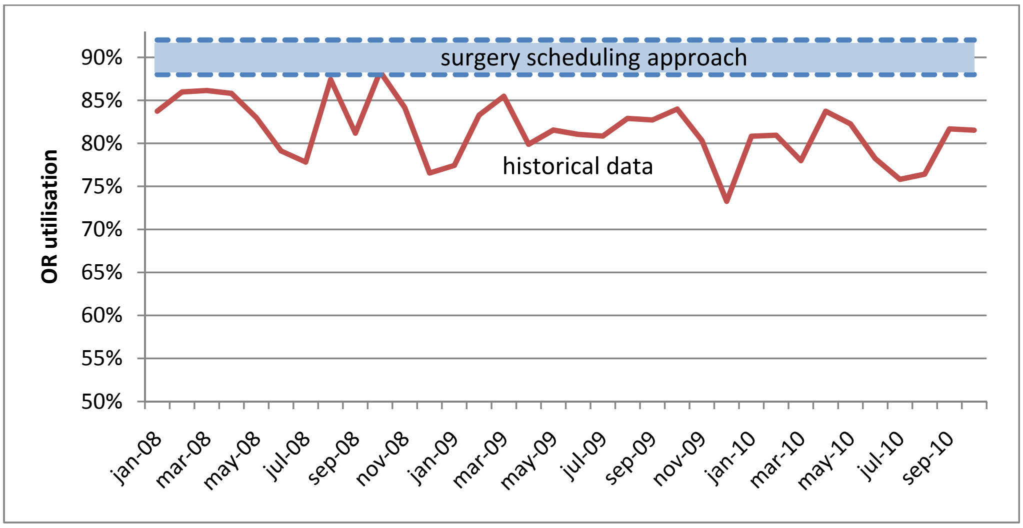

Figure 5.2: Comparison of historical OR utilisation and the surgery scheduling approach

Using our surgery scheduling approach yields a higher OR utilisation than the historical performance because the model generates ORDSs that use the available OR-time more efficiently. Consider, for example, the comparison of scheduling three ORs with or without our surgery scheduling approach in Figure 5.3. The figure displays the expected durations of surgeries (coloured blocks) and the slack (white block) per OR. Randomly filling OR-days with surgeries may result in the left schedule. This schedule has a large amount of unused OR-time. Our surgery scheduling approach results in the right schedule.

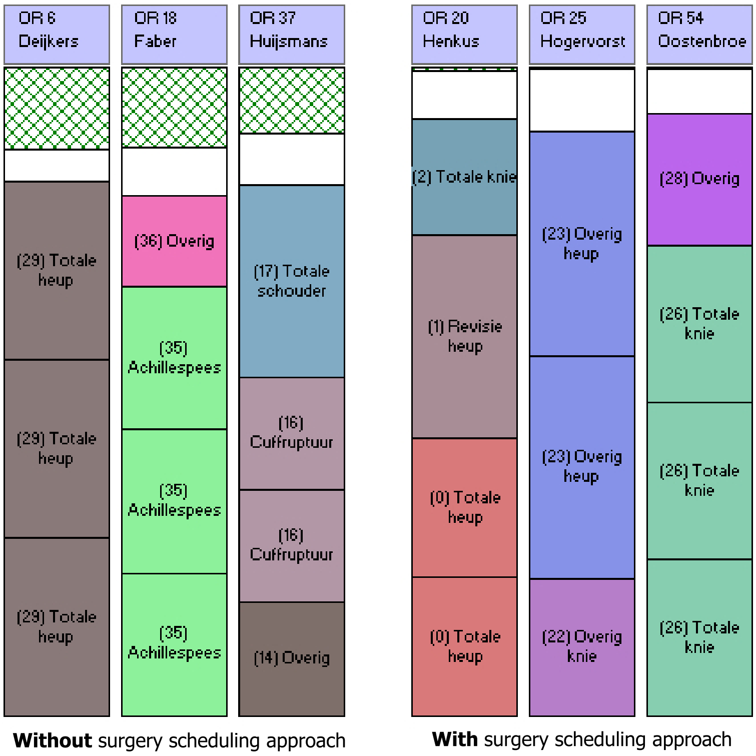

Figure 5.3: Comparison of scheduling with and without our surgery scheduling approach

5.4.6 Required number of beds

Since the simulated annealing procedure is a heuristic, it may get stuck in a local optimum. Therefore, when we compare the experiments with the default scenario, we state that a difference of the required number of beds is significant when the difference is greater or equal than 2 beds. We determine the minimal required number of beds for the whole planning horizon, on Saturdays and on Sundays. For every experiment it holds that either all three indicators are greater or equal to the default scenario or all the three indicators are smaller or equal to the default scenario.

The required number of beds is the highest for experiments with an increased amount of scheduled surgeries per ORDS. This holds for experiments with shorter surgery durations, increased opening hours or increased OR utilisation. In these experiments, more surgeries can be performed per day. As a result, more people are hospitalised and, thus, more beds are required. This line of reasoning explains the increased number of required beds for the 'Short duration', 'More surgery types', 'Overtime probability 50%', 'Opening hours: 9 hours', 'Opening hours: 12 hours' and 'Min. and max. demand' experiment. We require less beds for experiments where less surgeries are scheduled per day, such as the 'Long duration' and the 'Overtime probability 15%' experiment.

As expected, the 'Shorter LOS' experiment requires less beds due to a shorter length of stay per surgery type. A planning horizon of 2 weeks approximately matches the required number of beds in the default scenario. In the experiment with a planning horizon of 8 weeks, the required number of beds slightly increases while we expect that the planning horizon should have no influence on the required number of beds. Due to an increase of the planning horizon, the solution space also increases. Therefore, there is an increased probability that the simulated annealing procedure ends up in a local optimum. The same reasoning applies for the 'Available days' experiment. In this experiment, we do not restrict the model to schedule a fixed amount of ORDSs per day. Since we relax the problem, we expect that the number of required beds is less or equal to the default scenario, but instead, the number of required beds increases. Like in the 'Planning horizon: 8 weeks' experiment, the simulated annealing gets stuck in a local optimum due to a large solution space.

The 'No constraint' experiment has a low required number of beds due to the length of stay of the ORDS that is scheduled on each OR-day in the planning horizon. The ORDS only contains 'scopic knee' surgeries that have a length of stay of one day.

Most experiments match the required number of beds of the default scenario. This is most surprising for experiments where we disregard constraints. We conclude from this that the current set of instruments, surgeon availability and available OR-days have no influence on the required number of beds.

Figure 5.4 shows the bed occupancy when we apply the simulated annealing procedure on the default scenario. The weekends are clearly visible in the graph. Note that there are only 5 days in the planning horizon that require 45 beds.

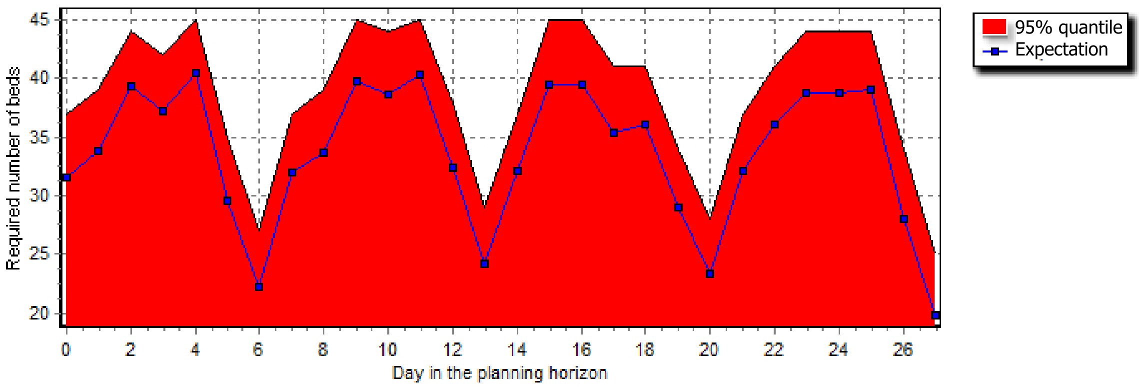

Figure 5.4: Bed occupancy after optimising the default scenario

In Chapter 2 we determined the required number of beds in 2010. Table 5.5 shows a comparison of the required number of beds in 2010 (see subsection 2.3.3) and the default scenario of the model at a 95% quantile. The model reduces the required number of beds by 7 beds and 8 beds on Sundays. Figure 5.5 visualises the number of beds in 2010 and of the surgery scheduling approach. The vertical lines indicate the number of beds that are required in 95% of the days. The figure shows that our surgery scheduling model reduces the variance in bed occupancy.

When a patient is discharged in the morning and another is hospitalised in the afternoon, a bed is occupied twice a day. Therefore, the required number of *physical* beds is lower than the required number of beds we determine in our model. In other words, the model presents a worse-case scenario of bed occupancy. Since the minimal number of required beds on Saturdays (33 beds) and Sundays (27 beds) is higher than the capacity of the biggest ward (24 beds), we conclude that it is not possible to close a ward in the

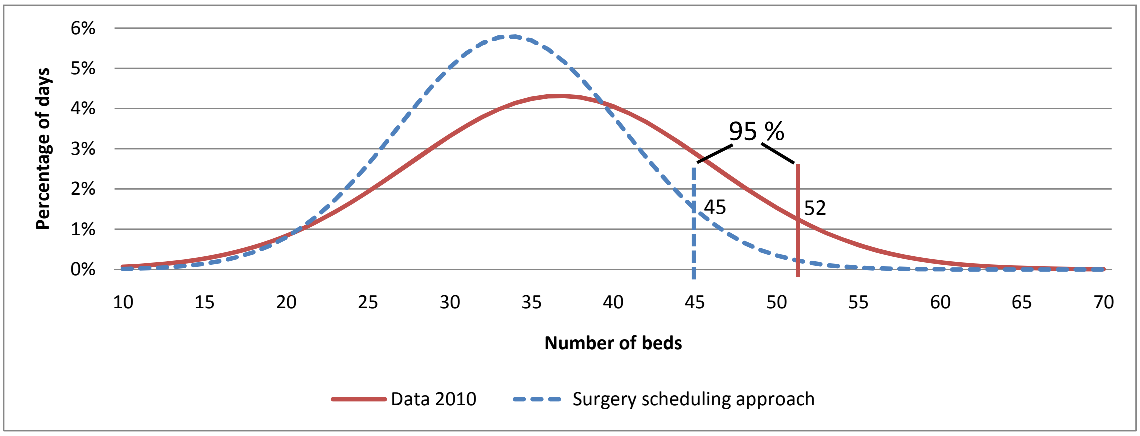

Figure 5.5: Comparison of the number of beds in 2010 and the surgery scheduling approach

Weekdays	*2010*	*Model*
All days of the week	52	45
Saturdays	40	33
Sundays	35	27

Table 5.5: Required number of beds comparison

weekend. Neither can the number of beds be lowered for all days of the week, because the number of physical beds is equal to the required number of beds from our model (45 beds). Since the model presents a worst-case scenario, it is likely that beds can be closed in practice.

5.5 Conclusions

In this chapter, we conduct numerical experiments to determine the performance of the model under different scenarios and experimental factors. Table 5.1, 5.3 and 5.4 respectively describe the performance indicators, scenarios and experimental factors. Figure 5.1 shows the results of the experiments.

The runtime of the ORDS generating model highly depends on the number of variables in the model and takes about 2 seconds for most experiments. The simulated annealing procedure takes about 30 seconds and highly depends on the planning horizon and, to a lesser extend, the number of scheduled surgeries per ORDS. We can not guarantee optimality of the ORDS generating model. The maximal deviation from the optimum varies between 0% and 2%, but it is 0.4% for most experiments. OR utilisation highly depends on the amount of slack in ORDSs. In turn, the amount of slack depends on the overtime probability, the variation of surgeries and the number of surgeries per ORDS. Next to that, we achieve a higher OR utilisation when we introduce more surgery types or increase opening hours. The required number of beds increases with the number of surgeries per ORDS. This holds true for experiments with shorter surgery durations, increased opening hours or increased OR utilisation.

The planning horizon has no influence on OR utilisation, but the required number of beds increases when we apply a planning horizon of 8 weeks. A positive standard deviation weight reduces OR utilisation, but results in a required number of beds that matches the default scenario. The current set of instruments, surgeon availability and available OR-days have no influence on OR utilisation and the required number of beds. Reducing the minimum demand per surgery type, however, improves OR utilisation.

Compared to the historical performance, the model improves the OR utilisation by 12%. No wards can be closed during the weekends. Neither can the number of physical beds be reduced for all days of the week. Since our model presents a worst-case scenario of bed-occupancy, it is likely that beds can be closed in practice.

Chapter 6

Conclusions and recommendations

In Chapter 1, we stated that our research objective is *to develop a tactical surgery scheduling approach that maximises the OR utilisation while minimising the required number of beds in the wards.*

Section 6.1 discusses the conclusions from this research. Section 6.2 gives our recommendations for HagaZiekenhuis.

6.1 Conclusions

As a case study, Chapter 2 describes the system of the orthopaedic department and how the system is controlled. Based on the performance analysis, we identified seven causes of poor performance:

General

1. Uncertainty about the availability of OR capacity
2. Equal starting times of operating rooms
3. Unclear definition of what activity marks the start of an OR-day

OR utilisation

4. Inaccurate estimation of intervention times
5. Inaccurate estimation of changeover times
6. Inaccurate estimation of slack

Required number of beds

7. Scheduling without considering ward capacity

We conducted a literature review in chapter 3 in search for surgery scheduling approaches that could overcome causes 4 - 7. We chose to develop a master surgery scheduling approach based on the article of Van Oostrum et al. [15]. We used the model of Vanberkel et al. [16] to model the stochastic length of stay.

In Chapter 4, we developed a non-cyclic master surgery scheduling approach. Our surgery scheduling approach contributes to literature because we model both the surgery durations and the length of stay as random variables. In addition, our MSS is non-cyclic. Also, we do not only schedule surgery types, but also surgeons. Furthermore, the number of surgeries per surgery type is variable while a fixed number of surgeries is scheduled in most models from literature. We also introduce a new method to linearise the constraint to limit the probability of overtime. The method finds a minimal number of piecewise linear functions to approximate the probability of overtime.

The surgery scheduling approach consists of two phases. In the first phase, we generate a set of operating room day schedules (ORDSs) that optimise OR utilisation. In the second phase, we apply a simulated annealing approach that swaps the assignment of ORDSs to OR-days around to optimise the number of required beds at a 95% quantile.

We conducted numerical experiments in Chapter 5. It appeared that OR utilisation highly depends on the amount of slack that is reserved to buffer against overtime. The amount of slack depends on the maximum probability of overtime, the variance of surgery durations, and the number of surgeries per ORDS. Furthermore, OR utilisation improves when we increase opening hours or when we define more surgery types. The required number of beds increases with the number of surgeries per ORDS. Adjusting the current set of instruments, surgeon availability, and available OR-days has no influence on OR utilisation and the required number of beds.

The orthopaedic department had an average OR utilisation of 78% in 2010. OR utilisation could be improved by 12% by using our surgery scheduling approach. No wards can be closed during the weekends. Neither can the number of physical beds be reduced for all days of the week. Since our model presents a worst-case scenario, it is likely that beds can be closed in practice.

6.2 Recommendations

In this section, we discuss recommendations for HagaZiekenhuis. Our main recommendation is to apply the surgery scheduling approach that we propose in this report. Apart from that, we discuss some general recommendations in Subsection 6.2.1. Subsection 6.2.2 describes recommendations regarding the implementation of the surgery scheduling approach. We conclude with suggestions for further research in Subsection 6.2.3.

6.2.1 General recommendations

In this subsection, we discuss four general recommendation.

Schedule surgeries based on surgery durations Currently, the OR planners schedules surgeries based on intervention time (time from incision until closing the wound) and account for changeover times and slack based intuition and experience. Furthermore, surgeries are scheduled in terms of blocks of half hours. To attain a better estimation of the OR-day endings, surgery durations should be estimated based on historical data and be modelled by a 3-parameter log-normal or a normal distribution.

OR capacity is known before constructing the MSS In the performance analysis of Chapter 2, we found that there is uncertainty about the availability of OR capacity (cause 1). However, our model assumes that the available OR-days are known before constructing a surgery schedule. Unfortunately, the orthopaedic department has no control over the OR capacity, since the OR department is responsible for assigning OR-days to specialisms. We recommend the orthopaedic department to make clear agreements with the OR department on a minimum amount of OR-days they can count on. The surgery scheduling model could then be applied to these minimum amount of OR-days. Each extra OR-day that is assigned to the orthopaedic department could then be scheduled separately.

(Re)define the start of OR-days From our performance analysis of Chapter 2, we found that equal starting times of ORs cause a long time between patients arriving at the OR and the incision. The anaesthetist is responsible for two operating rooms and administers the anaesthesia to the patients subsequently. Therefore, it is most likely that one of the two operating room starts late. It may decrease the surgery duration of the first surgery of the day to differentiate the starting times of the operating rooms. Furthermore, there is no clear definition what activity marks the start of an OR-day. An OR should start at 8.00, but should the patient be on the OR on that time, or should the incision take place at that time? We recommend that the OR department clearly defines what activity marks the start of OR-days. A clear definition also enables the possibility to address employees who frequently start late.

Improve quality of data We recommend Hagaziekenhuis to pay continuous attention on the quality of recorded data. It is of major importance to have reliable data to find bottlenecks in processes, to evaluate performance, and to determine the effect of interventions. In relation to our scheduling approach, we recommend to redefine diagnose codes and relate them to the surgery types of our MSS. Furthermore, time registrations of surgeries should be registered more accurately. The quality of data can be improved by increasing awareness of the importance of reliable data. One way to do this, is to evaluate quality of data regularly and address personnel that often lacks correct registration of data.

6.2.2 Implementation

Ultimately, the tactical surgery schedule should be generated every 4-week period, 3 months in advance. The tactical schedule should be communicated to the OR planners, but also to the OR department and the wards, so that they can anticipate on the resource demand of the surgery schedule. We formulate 6 phases on how to implement the surgery scheduling approach.

Phase 1: Adjust the information system. In this research, we determine the historical surgery duration and length of stay by processing the data manually. To have up-to-date estimations of surgery durations and length of stay per surgery type, the information system should process the historical data automatically. In addition, a surgeon should indicate the surgery type when a surgery request is made. Diagnose codes could be used for this purpose since this feature is already in place in the information system. The diagnose codes should be redefined based on the surgery types. Also, a surgery request should be created in the information system directly instead of filling in a paper 'admission form'. Adjusting the information system in not only useful for our scheduling approach, but also enables the OR planners to already start scheduling surgeries based on historical surgery durations.

Phase 2: Determine the surgery type demand per 4-week period Before applying the surgery scheduling approach, it is important to relate the minimum and maximum demand for each surgery type to the production agreements and the waiting lists. When determining the production agreements per 4-week period, management should also account for holidays, vacation, and conference periods. Integration of the model with existing information systems is useful to get up-to-date information on actual production and current waiting lists when constructing a new surgery schedule. After every planning horizon, the planned and actual production should be evaluated. Based on this evaluation, the minimum and maximum demand for future planning periods should be refined to comply with the production agreements and waiting list policy.

Phase 3: Develop a decision support (DSS) tool Our surgery scheduling approach should be implemented into an operational software tool. Management should decide whether they want to develop this DSS-tool in house or outsource it to a software development company. It is likely that in house development is cheaper, but the technical support of outsource development is better. When management chooses in house development, a multidisciplinary team of software developers, a project leader, SAP experts, medical personnel, and end users should be assembled to develop the DSS-tool. We conservatively estimate that this requires 2 FTE for a year. Based on a yearly salary of €50,000, the development of the DSS-tool requires an investment of €100,000. According to the OR department management, the direct personnel costs for one hour of OR-time is €900. The orthopaedic department is assigned to 616 OR-days per year (40 weeks of 13 OR-days, 12 weeks of 8 OR-days). This means that 1% improvement of the average OR utilisation yields a cost saving of €44,352 per year[1]. Thus, an investment of €100.000 pay back within a year when the DSS-tool improves the average OR utilisation by at least 2.25%. This is feasible, since we have shown that OR utilisation can be improved by 12%. Note that this a conservative calculation, since we did not include overhead costs, increased revenues due to higher OR utilisation, and the positive effects of our surgery scheduling approach on bed occupancy.

Phase 4: Use the DSS-tool for evaluation of the current scheduling practice To gain commitment of the OR planners, the DSS-tool could first be used as an evaluation tool. This evaluation tool does not optimise the surgery schedule, but determines the expected OR utilisation and bed occupancy of a given surgery schedule. After a couple of weeks, the expected performance of the model could be compared with the actual performance to show that the forecasts of the model are accurate.

Phase 5: Perform a pilot in the orthopaedic department The tactical surgery schedule can be implemented by requiring the OR planners to assign patients to their corresponding surgery type slot. In the process, the model may need some fine-tuning of the surgery type definitions and instrument constraints. When, after a while, the actual bed occupancy goes down, management could decide to close beds. To account for urgent surgeries, not all surgery type slots should be filled with elective patients. On the one hand, the OR planners can still try to find an elective patient for a surgery type slot that remains empty a couple of days before surgery. On the other hand, most patients would like to know their surgery date a couple of weeks in advance to prepare themselves mentally and take care of practical issues. OR utilisation will decrease when slots are left empty. Management must therefore decide upon a balance between OR utilisation and patient service by deciding on how to deal with empty slots. Management may determine that only less invasive surgeries or surgeries with a short length of stay may be scheduled a few days before a surgery date. Another option is to ask patients whether they are willing

[1] 1% of (616 OR-days x 8 hours per day x €900 per hour)

to be scheduled for surgery on a short notice. In exchange, a patient has the potential advantage of reduced access time.

Phase 6: Use the DSS-tool for other specialisms When the DSS-tool is in place at the orthopaedic department, the next step could be to use the DSS-tool for other specialisms.

6.2.3 Further research

In this report, we propose a model to construct a tactical surgery schedule. For further research, it is interesting to develop a surgery scheduling approach on an operational level. On an operational level, we schedule individual patients. Since information is known about the patient, we can schedule surgeries based on additional information from surgeons and on procedure types instead of surgery types. With this information, we can reduce the variability in the estimations of surgery duration and length of stay. Reduced variability enables us to further improve the OR utilisation and bed occupancy. In an operational model, we should also incorporate urgent and emergency surgeries. Explicitly modelling urgent and emergency surgeries will further improve the accuracy of the expected OR utilisation and bed occupancy.

We have shown that no wards can be closed in the weekends when maximising OR utilisation and having a fixed number of OR-days per week. In further research, it might be interesting to determine a surgery schedule where one of the wards can be closed during the weekends. This can be attained by allowing a lower OR utilisation or reducing the number of OR-days per week.

It is interesting to investigate the difference between optimising the bed occupancy based on deterministic and stochastic length of stay. The complexity of the model reduces if we would assume deterministic length of stay. By making a comparison between the two approaches, we can determine if it is worthwhile to model stochastic length of stay or that assuming deterministic lengths of stay performs just as good.

In the surgery scheduling approach, we assume that patients who had surgery are the only type of patients that are hospitalised in the ward. This is a valid assumption in the orthopaedic department, but may not be true for specialisms that hospitalise patients without having surgery. Our model could be extended by incorporating non-surgery patients when calculating the number of required beds.

Our surgery scheduling approach maximises OR utilisation by generating a set of ORDSs and selecting a subset of ORDSs to be assigned to OR-days. This implies that minimising the required number of beds does not compromise the OR utilisation. For further research, it would be interesting to see whether the required number of beds can be reduced when we allow a reduction of OR utilisation. This can be implemented by allowing the simulated annealing procedure to not only swap ORDSs within the subset of selected ORDSs, but to use the whole set of ORDSs. A major difficulty for implementation is to make sure that the demand constraints remain satisfied.

In our surgery scheduling approach, we assume that patients occupy a bed for the whole day of hospitalisation and discharge. In practice, however, some patients can occupy the same bed on a day. Therefore, the bed occupancy from our model can not directly be related to the actual bed occupancy in the wards. Further research is required to develop a model that better represents the actual bed occupancy. This can be achieved by determining the bed occupancy for shorter time intervals. For example, we can determine the bed occupancy for parts of the day (morning, afternoon and evening) or in hours.

Bibliography

[1] Beliën, J. *Exact and Heuristic Methodologies for Scheduling in Hospitals: Problems, Formulations and Algorithms.* PhD thesis, Katholieke Universiteit Leuven, Belgium, 2006.

[2] Beliën, J., and Demeulenmeester, E. Building cyclic master surgery schedules with levelled resulting bed occupancy. *European Journal of Operational Research 176* (2007), 1185–1204.

[3] Bisschop, J. Aimms optimizating modelling. Retrieved from: http://www.aimms.com/downloads/manuals/optimization-modeling, March 2010.

[4] Cardoen, B., Demeulemeester, E., and Beliën, J. Operating room planning and scheduling: A literature review. *European Journal of Operational Research 201*, 3 (2010), 921–932.

[5] Carter, M. Diagnosis: mismanagement of resources. *OR/MS Today 19* (2002), 26–32.

[6] Glauberman, S., and Mintzberg, H. Managing the care of health and the cure of disease - part i: Differentiation. *Health Care Management Review 26* (2001), 56–69.

[7] Hans, E., van Houdenhoven, M., and Hulshof, P. A framework for health care planning and control. University of Twente, The Netherlands, 2010.

[8] Hans, E., Wullink, G., van Houdenhoven, M., and Kazemier, G. Robust surgery loading. *European Journal of Operational Research 185*, 3 (2008), 1038–1050.

[9] Kirkpatrick, S., Gelatt, C. J., and Vecchi, M. Optimization by simulated annealing. *Science 220*, 4598 (1983), 671–680.

[10] Law, A. *Simulation Modeling and Analysis*, 4th ed. The McGraw-Hill Companies Inc., 2007.

[11] Stepaniak, P., Heij, C., Mannaerts, G., de Quelerij, M., and de Vries, G. Modeling procedure and surgical times for current procedural terminology-anesthesia-surgeon combinations and evaluation in terms of case-duration prediction and operating room efficiency: A multicenter study. *Anesthesia Analgesia 109*, 4 (2009), 1232–1245.

[12] Van Houdenhoven, M., Van Oostrum, J., Wullink, G., Hans, E., Hurink, J., Bakker, J., and Kazemier, G. Fewer intensive care unit refusals and a higher capacity utilization by using a cyclic surgical case schedule. *Journal of Critical Care 23* (2008), 222–226.

[13] Van Oostrum, J., Bredenhoff, E., and Hans, E. Managerial implications and suitability of a master surgical scheduling approach. *Annals of Operation Research 178*, 1 (2009), 91–104.

[14] Van Oostrum, J., Parlevliet, T., Wagelmans, A., and Kazemier, G. A method for clustering surgical cases to allow master surgical scheduling. Conference Proceedings ORAHS 2008, accepted for publication, 2008.

[15] Van Oostrum, J., Van Houdenhoven, M., Hurink, J., Hans, E., and Kazemier, G. Building cyclic master surgery schedules with levelled resulting bed occupancy. *OR Spectrum 30* (2008), 355–374.

[16] Vanberkel, P., Boucherie, R., Hans, E., Hurink, J., van Lent, W., and van Harten, W. An exact approach for relating recovering surgical patient workload to the master surgical schedule. *Journal of the Operational Research Society. (forthcoming)*.

Appendices

Appendix A

Conceptual model

Indices

$i \in I$	surgery types (clusters of surgical procedure types)
$j \in J$	operating rooms
$t \in T$	days of the planning horizon
$s \in S$	surgeons
$r \in R$	resources

Parameters

A_{jt}	binary parameter indicating OR-day (j,t) is available
B_{st}	binary parameter indicating surgeon s is available on day t
E_{is}	binary parameter indicating surgeon type i is performed by surgeon s $(\sum_s E_{is} = 1\ \forall i)$
M_i	maximum number of surgeries of type i that can be performed on an OR-day
$D_i^{\min}$	minimum number of surgery types i that have to be performed in the planning period
$D_i^{\max}$	maximum number of surgery types i that have to be performed in the planning period
ω_{ir}	the number of instrument r that are needed for surgery type i
$Q_r^{\max}$	maximum number of instrument r that can be used on one day
θ^u	weight of the total average session time in the objective function
θ^v	weight of the total standard deviation of session time in the objective function
μ_i	expected surgery duration of surgery type i
σ_i^2	variance of surgery duration of surgery type i
U_t	random variable representing the number of occupied beds on day t, based on previous planning periods
C	total available time on an OR-day
L_i	stochastic variable that represents the length of stay for surgery type i
α	maximum probability that the total duration of sessions on an OR-day is greater than the total available time (overtime probability)
β	minimum probability that less than γ beds are needed on day t

Functions

$f_{jt}(V)$	probability distribution of the total duration of all surgery types that are scheduled on OR-day (j,t)
$g_t(U,V,L)$	probability distribution of the number of occupied beds on day t

Decision variables

V_{ijt}	number of surgeries of type i on OR-day (j,t)
W_{sjt}	binary variable indicating that surgeon s performs surgery on OR-day (j,t)
γ	maximal number of required beds

$$\max \quad \sum_{i\in I}\sum_{j\in J}\sum_{t\in T}(\theta^u\mu_i + \theta^v\sigma_i)V_{ijt}$$

$$\min \quad \gamma$$

subjected to

$$\sum_{s\in S} W_{sjt} \leq A_{jt} \qquad \forall j,t$$

$$\sum_{j\in J} W_{sjt} \leq B_{st} \qquad \forall s,t$$

$$V_{ijt} \leq \sum_{s\in S} E_{is}W_{sjt}M_i \qquad \forall i,j,t$$

$$D_i^{\min} \leq \sum_{j\in J}\sum_{t\in T} V_{ijt} \leq D_i^{\max} \qquad \forall i$$

$$\sum_{i\in I}\sum_{j\in J} V_{ijt}\omega_{ir} \leq Q_r^{\max} \qquad \forall t,r$$

$$\mathrm{P}(f_{jt}(V) \leq C) \geq \alpha \qquad \forall j,t$$

$$\mathrm{P}(g_t(U,V,L) \leq \gamma) \geq \beta \qquad \forall t$$

$$V_{ijt} \in \mathbb{N} \qquad W_{sjt} \in \{0,1\} \qquad \gamma \in \mathbb{N}$$

Appendix B

Default scenario data

Due to confidentiality, this appendix is not available in the public version of this report.

Appendix C

Clustering surgery types

This appendix discusses a procedure to cluster surgical procedure types. First, we describe the procedure in Section C.1. Next, we discuss our results when applying the procedure on patient data of the orthopaedic department for all patients that were operated in 2010 in Section C.2.

C.1 Clustering procedure

We cluster surgical procedures types based on historical data. Let N be the set of all the patients that are operated in the past. The surgery duration and length of stay of each patient $n \in N$ is denoted by t_n and l_n, respectively. Let Z denote the set of procedure types that have been performed. Procedure type $z \in Z$ is the lowest level of registration in the hospital database. We introduce subset N_z to denote the patients that were operated for procedure type z. Set I is the set of all surgery types (clusters). Subset Z_i includes the procedure types that are clustered to surgery type $i \in I$. Each procedure type z should be assigned to exactly one surgery type i, therefore $Z_i \bigcap Z_{\hat{\imath}} = \emptyset \; \forall i \neq \hat{\imath}$ and $\bigcup_{i \in I} Z_i = Z$.

We present an approach to cluster surgical procedures as proposed by Van Oostrum et al. [14]. The goal of the clustering approach is to create surgery types where the variability of the surgery duration and length of stay within the surgery type is minimised, while the variability between surgery types is maximised. Joining two different surgical procedures in a cluster leads to loss of information, compared to the situation where we assign both surgical procedures to individual surgery types. Van Oostrum et. al. propose the error sum of squares (ESS) as a measure for the loss of information. We denote the error sum of squares of surgery type i by ESS_i. Before formulating the definition of ESS_i, we observe that the average surgery duration $\bar{t_i}$ and length of stay $\bar{l_i}$ for each surgery type i is the average over all the patients that were operated for the procedure types z that are assigned to surgery type i:

$$\bar{t_i} = \frac{\sum\limits_{z \in Z_i} \sum\limits_{n \in N_z} t_n}{\sum\limits_{z \in Z_i} |N_z|} \qquad \bar{l_i} = \frac{\sum\limits_{z \in Z_i} \sum\limits_{n \in N_z} l_n}{\sum\limits_{z \in Z_i} |N_z|} \tag{C.1}$$

We define ESS_i as the maximum of the weighted ESS of surgery duration and length of stay. The weighting factors of the surgery duration and the length of stay are denoted by θ^{sur} and θ^{los} respectively.

$$ESS_i = \max \left\{ \sum_{z \in Z_i} \sum_{n \in N_z} \theta^{sur} (t_n - \bar{t_i})^2, \sum_{z \in Z_i} \sum_{n \in N_z} \theta^{los} (l_n - \bar{l_i})^2 \right\} \tag{C.2}$$

Let $D_{i\hat{\imath}}$ denote the cost (loss of information) of clustering surgery type i and $\hat{\imath}$. $D_{i\hat{\imath}}$ is the difference between the ESS when i and $\hat{\imath}$ are merged, and the ESS when i and $\hat{\imath}$ are left separated:

$$D_{i\hat{\imath}} = ESS_{i\hat{\imath}} - (ESS_i + ESS_{\hat{\imath}}) \qquad \forall i \neq \hat{\imath} \tag{C.3}$$

Procedure types can be clustered by performing the following steps:

1. Start with $|Z|$ surgery types, each containing a single procedure type. Calculate $D_{i\hat{\imath}}$ for every combination of surgery types.

2. Search for the combination of surgery types with the lowest value of $D_{i\hat{\imath}}$. Let this combination be surgery types c and $\hat{c}$.

3. Merge surgery types c and $\hat{c}$ in a new surgery type and delete surgery type c and $\hat{c}$, i.e. $Z \backslash \{c\}$ and $Z \backslash \{\hat{c}\}$. Update $D_{i\hat{\imath}}$ for every surgery type combination.

4. Repeat step 2 and 3 until the cost exceed some threshold value t.

C.2 Clustering results

Before we perform the clustering procedure, we need to define the weights for surgery duration and length of stay (θ^{sur} and θ^{los} respectively) and the threshold value t. We decide that 15 minutes of OR time should be equally important as one day of hospitalisation. Therefore, we set $\theta^{sur} = 1$ and $\theta^{los} = 15$.

Figure C.1 shows an example of the cost ($D_{i\hat{\imath}}$) of each iteration of the clustering procedure. The figure shows that the cost function increase exponentially. Intuitively, the graph can be explained by considering that it is more costly to combine two cluster from a group of 5 big clusters, than to combine two clusters from a group of 20 small clusters. We assumed that it would be best to stop the clustering procedure when the graph reaches an asymptotic value. After trial and error, we decided to set the threshold value $t = 3$.

We apply our clustering procedure on patient data from the hospital database for all patients that were operated in 2010. According to Stepaniak et al. [11], a lognormal distribution is best suitable to model surgery durations and length of stay. Therefore, we take the logarithm of the surgery duration and length of stay of each patient as input data for the clustering procedure.

Table C.1 shows the number of procedure types and the number of surgery types that resulted from the clustering procedure.

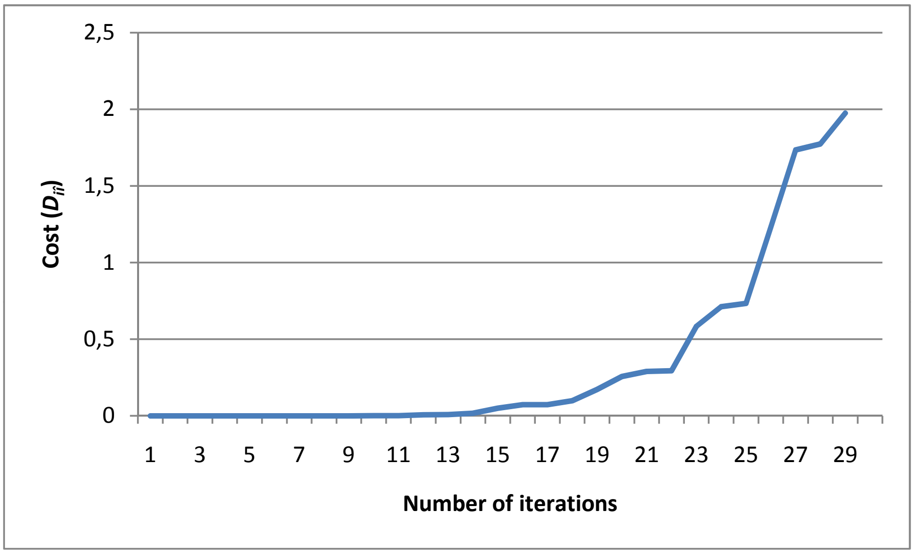

Figure C.1: Example of the cost ($D_{i\hat{i}}$) function of the clustering procedure

Surgeon	*Procedure types*	*Surgery types*
Baas	55	17
Deijkers	73	14
Erp Taalman Kip	23	6
Faber	139	14
Henkus	46	14
Hogervorst	58	12
Lugt	44	14
Oostenbroek	51	13
Huijsmans	66	13
Total	555	117

Table C.1: Result of clustering procedure types into surgery types

Appendix D

Linearisation of the overtime constraint

In this appendix, we show how we linearise the overtime constraint from Section 4.1. Section D.1 discusses how we formulate the overtime constraint in terms of a piecewise linear functions. Section D.2 explains how we determine the linear functions.

D.1 Linearisation

In the problem formulation, the overtime constraint needs to be linearised in order to be solved by a linear solver. As displayed in Subsection 4.3.1 and Appendix A, the overtime constraint is

$$\mathrm{P}(f_{jt}(V) \leq C) \geq 1 - \alpha \qquad \forall j, t \tag{D.1}$$

To linearise the constraints, we assume that the total duration of an OR-day (j, t) is normally distributed. Let μ_{jt} and σ_{jt}^2 be the mean and variance of the random variable $f_{jt}(V)$, i.e $f_{jt}(V) \sim \mathcal{N}(\mu_{jt}, \sigma_{jt}^2)$. In order to express constraint (D.1) linearly in μ_{jt} and σ_{jt}, we rewrite the equation in the standard normal form:

$$\Phi\left(\frac{C - \mu_{jt}}{\sigma_{jt}}\right) \geq \alpha \qquad \Rightarrow \qquad \mu_{jt} + \Phi^{-1}(1-\alpha)\sigma_{jt} \leq C \tag{D.2}$$

In Section 4.1, we define μ_i and σ_i^2 as the mean and variance of the surgery duration of surgery type i. Also, V_{ijt} is the number of surgeries of surgery type i on OR-day (j, t). Using these parameters, we define μ_{jt} and σ_{jt}^2 as follows

$$\mu_{jt} = \sum_{i \in I} V_{ijt}\mu_i \qquad \sigma_{jt}^2 = \sum_{i \in I} V_{ijt}\sigma_i^2 \tag{D.3}$$

Filling in the definition of μ_{jt} and σ_{jt}^2 (equation (D.3)) in the overtime constraint (equation (D.2)) yields

$$\sum_{i \in I} V_{ijt}\mu_i + \Phi^{-1}(1-\alpha)\sqrt{\sum_{i \in I} V_{ijt}\sigma_i^2} \leq C \qquad \forall j, t \tag{D.4}$$

Observe that the constraint is non-linear in the decision variable V_{ijt} due to the square root. In order to make the constraint linear, we propose to approximate the square root function $\sqrt{x}$ by a set of linear functions. Each linear function approximates the square root function for some interval. The beginnings and endings of each interval is denoted by a set of $N = \{0, 1, \ldots, m\}$ breakpoints. Let x_n be the value on the x-axis of breakpoint $n \in N$. By definition, x_0 and x_m are the minimum and maximum x-values for which we approximate $\sqrt{x}$. The remaining x_n values are the intersection points of the linear functions. Each linear function is a tangent line of the square root function in the point t_n for $n \in N \backslash \{0\}$. Let the linear approximation functions be described by $h_n(x) = a_n + b_n x$ for $n \in N \backslash \{0\}$. The constant b_n is the derivative of the square root function in the point t_n. We determine a_n by setting the square root function equal to its linear approximation function.

$$b_n = (\sqrt{t_n})' = \frac{1}{2}\sqrt{\frac{1}{t_n}} \qquad \sqrt{t_n} = a_n + b_n t_n \quad \Rightarrow \quad a_n = \frac{1}{2}\sqrt{t_n} \tag{D.5}$$

Filling in the definitions of a_n and b_n in the linear approximation function $h_n(x)$ yields

$$h_n(x) = a_n + b_n x = \frac{1}{2}\sqrt{t_n} + \frac{1}{2}\sqrt{\frac{1}{t_n}}x \tag{D.6}$$

Figure D.1 shows an example where two linear functions (grey lines) approximate the square root function (black line) in the interval [0, 90]. The linear functions are the tangent lines in the point $t_1 = 7$ and $t_2 = 50$. The points $x_0 = 0$ and $x_2 = 90$ indicate the minimum and maximum of the approximated interval. The intersection point of the two tangent lines is $x_1 = 18.7$.

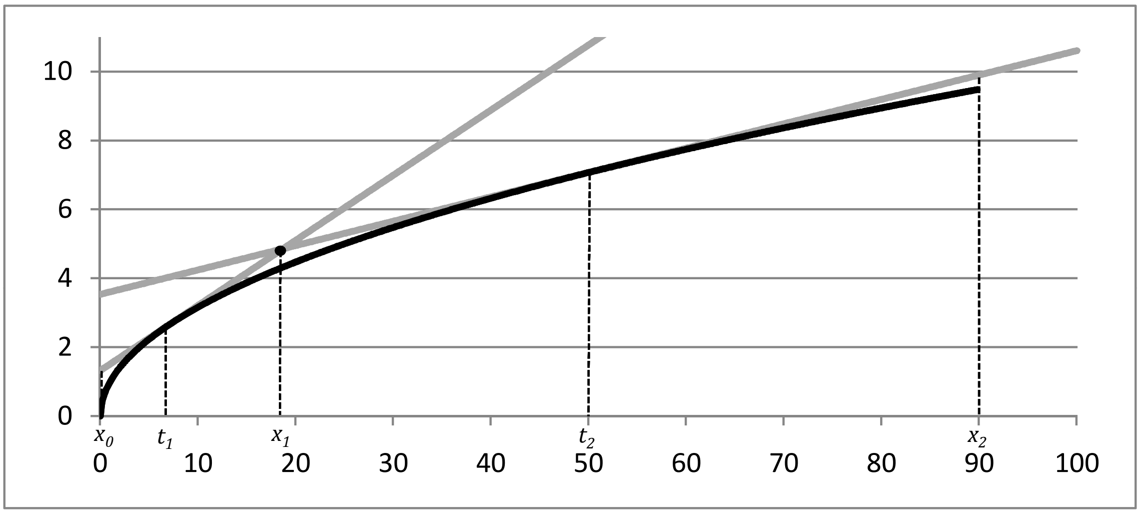

Figure D.1: Approximation of the square root function by two linear functions

Let y_n be the function value of the linear approximation function at breakpoint n, i.e. $y_n = h_n(x_n)$. When we known the breakpoints, we can apply a common method of modelling piecewise linear functions as described in the AIMMS optimisation modelling manual [3]. Observe that any point in between two breakpoints is a weighted sum of these two breakpoints. Let ρ_n denote the non-negative weights such that their sum is 1. The overtime constraint (D.4) can now be written as

$$\sum_{i \in I} V_{ijt}\mu_i + \Phi^{-1}(1-\alpha) \sum_{n \in N} \rho_n y_n \leq C_{jt} \qquad \forall j, t \tag{D.7}$$

with the additional constraints

$$\sum_{n\in N} \rho_n x_n = \sum_{i\in I} V_{ijt}\sigma_i^2 \tag{D.8}$$

$$\sum_{n\in N} \rho_n = 1 \tag{D.9}$$

Furthermore, we require that at most two adjacent ρ_n's are greater than zero. We do not model this as a formal constraint, because most commercial solvers have the ability to specify the last constraint (equation (D.9)) as a special ordered set (SOS) constraint. A solver that is able to use SOS type constraints can prune the branch and bound tree more efficiently compared to a model formulation in which the adjacent requirement is stated explicitly.

D.2 Determining the breakpoints

In this section, we discuss how to determine the number of breakpoints and their values. A high approximation accuracy of the square root function requires a high amount of breakpoints. On the other hand, more breakpoints enlarge the complexity of the ILP-model due to an increasing number of variables. Therefore, we determine the minimal number of breakpoints that is needed to attain a maximum approximation error $\Delta^{\max}$.

Observe from Figure D.1 that for every linear approximation function, the approximation error is the highest at the breakpoints. Let δ_n be the approximation error for breakpoint $n \in N$. Figure D.2 visualises δ_n for an approximation of the square root function by one linear function.

$$\delta_n = h_n(x_n) - \sqrt{x_n} \qquad \forall n > 0 \tag{D.10}$$

$$\delta_0 = h_1(x_0) - \sqrt{x_0} \tag{D.11}$$

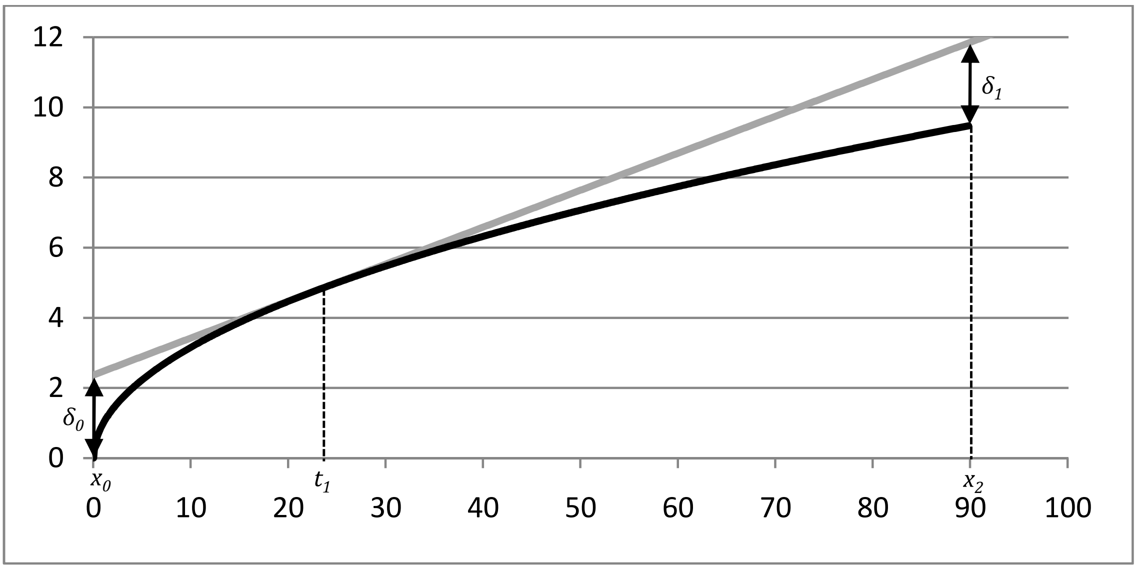

Figure D.2: Approximation of $\sqrt{x}$ in the interval [0,90] by one linear function

Note that $\delta^{\max} = \max_{n\in N} \delta_n$ is minimised when the approximation error at every breakpoint is equal, i.e. when $\delta_0 = \delta_1 = \ldots = \delta_m$. We can find this solution by solving the following set of equations.

$$h_n(x_n) = h_{n+1}(x_n) \qquad \forall n = 1, \ldots, m-1 \qquad \text{(D.12)}$$
$$\delta_n = \delta_{n+1} \qquad \forall n = 0, \ldots, m-1 \qquad \text{(D.13)}$$
$$x_n < t_{n+1} < x_{n+1} \qquad \forall n = 0, \ldots, m-1 \qquad \text{(D.14)}$$

Equation (D.12) follows directly from the definition of x_n as it is the intersectionpoint of two linear functions $h_n(x_n)$ and $h_{n+1}(x_n)$. Equation (D.13) makes sure that the approximation errors at every breakpoint are equal. Equation (D.14) makes sure that an interval of a linear function is positive (¿ 0) and that a tangent point is in between two subsequent breakpoints.

Let us first elaborate on (D.12) by expressing it in terms of t and x values.

$$h_n(x_n) = h_{n+1}(x_n) \quad \Rightarrow \quad a_n + b_n x_n = a_{n+1} + b_{n+1} x_n \quad \Rightarrow \quad x_n = \frac{a_{n+1} - a_n}{b_n - b_{n+1}}$$

Filling in the definition of a_n and b_n from (D.5), we can express x_n as a function of t_n and t_{n+1}.

$$x_n = \frac{\sqrt{t_{n+1}} - \sqrt{t_n}}{\sqrt{\frac{1}{t_n}} - \sqrt{\frac{1}{t_{n+1}}}} = \sqrt{t_{n+1} t_n} \qquad \text{(D.15)}$$

Next, we express equation (D.13) in terms of t and x values.

$$\delta_n = \delta_{n+1} \quad \Rightarrow \quad h_{n+1}(x_n) - \sqrt{x_n} = h_{n+1}(x_{n+1}) - \sqrt{x_{n+1}} \quad \Rightarrow$$
$$x_{n+1} = (2\sqrt{t_{n+1}} - \sqrt{x_n})^2 \qquad \text{(D.16)}$$

We assume that $x_0 = 0$, because the minimum variance of a surgery schedule is zero, e.g. when no surgeries are scheduled on an OR-day. Using (D.15) and (D.16) thus results in the following:

$$\begin{aligned} x_1 &= 4t_1 \\ t_2 &= 4^2 t_1 \\ x_2 &= (2 \cdot 4\sqrt{t_1} - \sqrt{4}\sqrt{t_1})^2 \\ &= (2 \cdot 4 - \sqrt{4})^2 t_1 \\ &= \frac{(2 \cdot 4 - \sqrt{4})^2}{4^2} t_2 \end{aligned}$$

We introduce a variable a_n as the ratio between x_n and t_n. This ratio is a constant when $x_0 = 0$. This gives us the following:

$$a_1 = 4 \qquad \text{(D.17)}$$
$$x_n = a_n t_n \qquad \forall n = 1, \ldots, m \qquad \text{(D.18)}$$
$$t_{n+1} = a_n^2 t_n \qquad \forall n = 1, \ldots, m \qquad \text{(D.19)}$$
$$a_{n+1} = \frac{(2a_n - \sqrt{a_n})^2}{a_n^2} \qquad \forall n = 1, \ldots, m \qquad \text{(D.20)}$$

Since $a_1 = 4$ is known, we can compute a_n for all n. As an example, we display the a-values in Table D.1. The approximation error for the last breakpoint (δ_m) can now be

n	1	2	3	4	5	6	7	8	9	10
a_n	4	2.25	1.78	1.56	1.44	1.36	1.31	1.27	1.23	1.21

Table D.1: Table of a-values

m	1	2	3	4	5	6	7	8	9	10
$\delta^{\max}$	2.37	0.79	0.40	0.24	0.16	0.11	0.08	0.07	0.05	0.04

Table D.2: Approximation errors for $x_m = 90$

expressed in terms of x_m and a_m. Since the approximation errors on all breakpoints are equal, it follows that $\delta^{\max} = \delta_m$.

$$\delta^{\max} = \delta_m = h_m(x_m) - \sqrt{x_m} = \sqrt{x_m}\left(\frac{1}{2\sqrt{a_m}} + \frac{\sqrt{a_m}}{2} - 1\right) \tag{D.21}$$

From Equation (D.21), we observe that the approximation error depends on the maximum value of the interval (x_m) and the number of linear functions (m). Given a value of x_m, we can calculate the approximation error for every number of linear functions. With this, we determine the minimal number of linear functions to achieve a maximum approximation error of $\Delta^{\max}$.

Let us illustrate this with an example. Suppose, we want to approximate the square root function for the interval [0,90]. By definition, if follows that $x_0 = 0$ and $x_m = 90$. Using (D.21) and Table D.1, we determine the maximal approximation error for each number of breakpoints as shown in Table D.2. Suppose that we accept an approximation error of at most 1, i.e. $\Delta^{\max} = 1$. Looking at Table D.2, we conclude that we need 2 linear functions ($m = 2$), since that is the smallest value for m where $\delta^{\max} \leq \Delta^{\max}$. Using (D.15), (D.18) and (D.19), we determine that $t_1 = 2.5$, $t_2 = 40$ and $x_1 = 10$. Figure D.3 shows the result. The y-values that we use in constraint (D.7) are the function values at the breakpoints of the linear approximation function, i.e. $y_n = h_n(x_n)$. This yields $y_0 = 0,791$, $y_1 = 3,95$ and $y_2 = 10,277$.

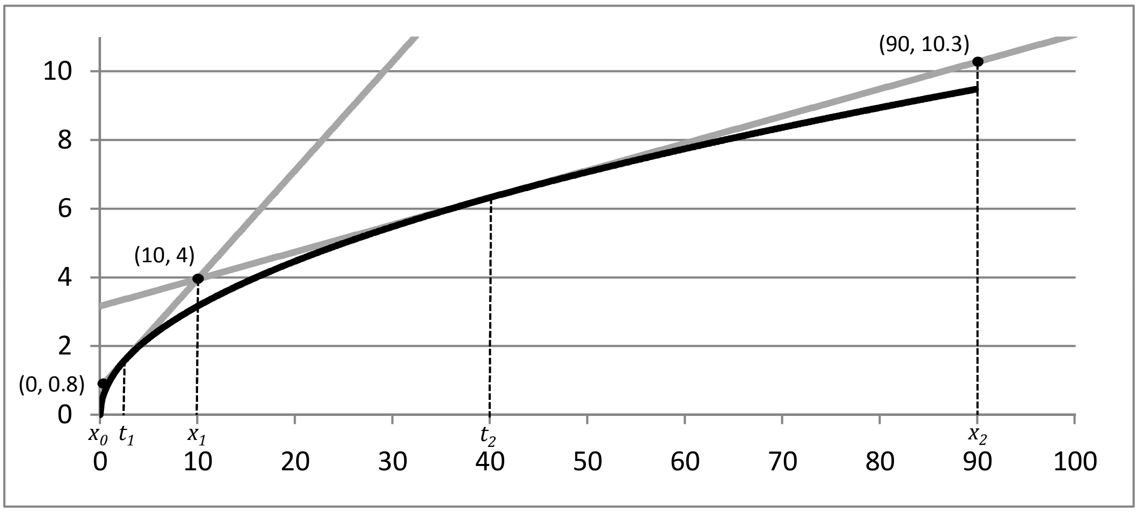

Figure D.3: Approximation of $\sqrt{x}$ in the interval [0,90] by two linear functions

We summarise our approach as follows:

Define a maximum approximation error $\Delta^{\max}$ and a maximum value of the interval x_m for which to approximate the square root function.

1. Compute the a_n-values using (D.17) and (D.20).
2. Calculate the maximal approximation errors $\delta^{\max}$ for each number of linear functions m using (D.21).
3. Choose the lowest m-value (number of linear functions) for which $\delta^{\max} \leq \Delta^{\max}$.
4. Use (D.18) to determine t_m.
5. Determine the rest of the tangent points $t_n \; \forall \; n < m$ by using (D.19).
6. Determine the breakpoints $x_n \; \forall \; n < m$ by using (D.15).
7. Determine the breakpoint function values $y_n = h_n(x_n)$ by using (D.6).

Appendix E

Bed distribution of ORDSs

This appendix discusses how to determine parameter $h_n^k(x)$: the discrete distribution of x patients of ORDS k still in recovery on n day within one planning horizon. Since $h_n^k(x)$ is the bed distribution for one planning horizon, $h_n^k(x)$ is only defined for $n \leq |T|$. The model is based on the article from Vanberkel et al [16]. For this model we need two parameters. First, we need to know how many of which type of surgeries are scheduled per ORDS. Second, we need to know the length of stay for each surgery type.

O_{ik} : number of surgeries of surgery type i in ORDS k.
l_n^i : probability that the length of stay of patient of surgery type i is exactly n days long.

Note that l_n^i is a multinomial distribution. From l_n^i, we derive d_n^i: the probability that a patient, who is still in the ward at day n, is to be discharged that day ($n \in \{0, 1, \ldots, L^i\}$, where L^i is the maximum length of stay for surgery type i). The parameter d_n^i is calculated as follows.

$$d_n^i = \frac{l_n^i}{\prod_{t=0}^{n-1}(1 - d_n^i)}$$

Let $\tilde{h}_n^{ik}(x)$ be the probability that n days after carrying out ORDS k, x patients of surgery type i are still in recovery.

For $n = 0$:

$$\tilde{h}_0^{ik}(x) = \begin{cases} 1 & \text{when } x = O_{ik} \\ 0 & \text{otherwise} \end{cases}$$

For $n > 0$:

$$\tilde{h}_n^{ik}(x) = \sum_{y=x}^{O_{ik}} \binom{y}{x} (d_n^i)^{y-x} (1 - d_n^i)^x h_{n-1}^{ik}(y)$$

Now, h_n^k can be determined by adding the distributions of all the surgery types i on ORDS k. This is done by the convolution of $\tilde{h}_n^{ik}(x)$ over all i. The maximum length of stay of a patient may be longer than the planning horizon. Therefore, we also add the distributions of $\tilde{h}_n^{ik}(x)$ to h_n^k for those days of the length of stay that fall into a new planning horizon.

$$h_n^k = \tilde{h}_n^{0k} * \tilde{h}_n^{1k} * \ldots * \tilde{h}_n^{|I|k} * \tilde{h}_{n+|T|}^{0k} * \ldots * \tilde{h}_{n+|T|}^{|I|k} * \tilde{h}_{n+2|T|}^{0k} * \ldots * \tilde{h}_{n+2|T|}^{|I|k} * \ldots \qquad \forall n \leq |T|$$